



Full-sky, High-resolution Maps of Interstellar Dust

Citation

Meisner, Aaron Michael. 2015. Full-sky, High-resolution Maps of Interstellar Dust. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:17464616>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Full-sky, High-resolution Maps of Interstellar Dust

A dissertation presented

by

Aaron Michael Meisner

to

The Department of Physics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

April 2015

©2015 – Aaron Michael Meisner

All rights reserved.

Full-sky, High-resolution Maps of Interstellar Dust

Abstract

We present full-sky, high-resolution maps of interstellar dust based on data from the Wide-field Infrared Survey Explorer (*WISE*) and *Planck* missions. We describe our custom processing of the entire *WISE* $12\mu\text{m}$ All-Sky imaging data set, and present the resulting $15''$ resolution, full-sky map of diffuse Galactic dust emission, free of compact sources and other contaminating artifacts. Our derived $12\mu\text{m}$ dust map offers angular resolution far superior to that of all other existing full-sky, infrared dust emission maps, revealing a wealth of small-scale filamentary structure.

We also apply the Finkbeiner et al. (1999) two-component thermal dust emission model to the *Planck* HFI maps. We derive full-sky $6.1'$ resolution maps of dust optical depth and temperature by fitting this two-component model to *Planck* 217-857 GHz along with DIRBE/*IRAS* $100\mu\text{m}$ data. In doing so, we obtain the first ever full-sky 100-3000 GHz *Planck*-based thermal dust emission model, as well as a dust temperature correction with $\sim 10\times$ enhanced angular resolution relative to DIRBE-based temperature maps. Analyzing the joint *Planck*/DIRBE dust spectrum, we show that two-component models provide a better fit to the 100-3000 GHz emission than do single-MBB models, though by a lesser margin than found by Finkbeiner et al. (1999) based on FIRAS and DIRBE. We find that, in diffuse sky regions, our two-component 100-217 GHz predictions are on average accurate to within 2.2%, while extrapolating the Planck Collaboration et al. (2013k) single-MBB model systematically underpredicts emission by 18.8% at 100 GHz, 12.6% at 143 GHz and 7.9% at 217 GHz. We calibrate our two-component optical depth to reddening,

and compare with reddening estimates based on stellar spectra. We find the dominant systematic problems in our temperature/reddening maps to be zodiacal light on large angular scales and the cosmic infrared background anisotropy on small angular scales.

Future work will focus on combining our *WISE* $12\mu\text{m}$ dust map and *Planck* dust model to create a next-generation, full-sky dust extinction map with angular resolution several times better than Schlegel et al. (1998).

Contents

1	Introduction	1
1.1	Broad Overview of Interstellar Dust	1
1.2	Observational History: Dust in the Optical	2
1.3	Observational History: Dust in the Infrared	4
1.4	History of Dust as a Cosmological Foreground	5
1.5	Recent Developments	7
1.6	Outline	8
2	WISE 12 micron Dust Emission Map	9
2.1	Introduction	10
2.2	The WISE Data	12
2.3	Point Sources	14
2.4	Point Spread Function	18
2.5	Single Exposure Pipeline	27
2.6	Stacking	35
2.7	Final WSSA Tiles	48
2.8	Discussion and Conclusion	52
3	Planck Thermal Dust Emission Model	56
3.1	Introduction	57

3.2	Data	60
3.3	Preprocessing	60
3.4	Dust Emission Model	72
3.5	Predicting the Observed SED	75
3.6	Global Model Parameters	77
3.7	MCMC Fitting Procedure	80
3.8	Optical Reddening	95
3.9	Comparison of Emission Predictions	103
3.10	Data Release	109
3.11	Conclusions	109
4	Conclusion	114
4.1	Funding Acknowledgments	119
	References	121

Citations to Previously Published Work

The content of Chapter 2, apart from minor modifications, appeared in:

Meisner, A. M., & Finkbeiner, D. P. 2014, *ApJ*, 781, 5

The content of Chapter 3, apart from minor modifications, appeared in

—. 2015, *ApJ*, 798, 88

Chapters 2 and 3 are reproduced by permission of the AAS.

Listing of figures

2.1	WISE nebula fragment sources	15
2.2	WISE catalog curve of growth filtering	16
2.3	Bright source PSF subtraction residuals	19
2.4	Multi-exposure imprint of bright WISE point source	20
2.5	WISE PSF model core and wings	23
2.6	Solar system object masking	34
2.7	Per-quadrant L1b exposure warping	40
2.8	L1b destriping	42
2.9	Off-axis Moon contamination	45
2.10	Example of 12 micron processing results	49
2.11	Full-sky 12 micron dust image	53
3.1	Comparison of Planck 857 GHz and DIRBE/FIRAS prediction	65
3.2	Planck residual Solar dipole corrections	66
3.3	Planck low-order corrections at 100 GHz	67
3.4	DIRBE correlation slopes relative to Planck 857 GHz	77
3.5	Illustration of two-component SED	82
3.6	Example per-pixel posterior distributions for dust model parameters	85
3.7	Comparison of Planck-based and DIRBE-based temperature resolution	87
3.8	Spatial variation of dust grain species' mass fractions	90

3.9	Goodness of fit for competing dust SED models	92
3.10	Calibration of optical depth to reddening	96
3.11	Reddening prediction residuals	97
3.12	Reddening residuals at high ecliptic latitude	99
3.13	Effect of FIR zero levels on reddening predictions	103
3.14	350-3000 GHz thermal dust emission predictions	105
3.15	100-350 GHz thermal dust emission predictions	107
3.16	Full-sky results	110
4.1	Comparison of Planck and Herschel Ophiuchus SED Fits	117

Acknowledgments

First and foremost, I would like to thank Doug for serving as my thesis advisor and for teaching me many invaluable lessons, especially those related to debugging complicated problems, computational best practices and maintaining a big-picture perspective while working on a project's low-level details. Needless to say, my thesis research would not have been possible without Doug's vision and guidance. I would also like to thank the other senior researchers I've worked closely with during my time at Harvard, especially Mario Jurić and Anna Frebel. Thanks also to Chris Stubbs and John Kovac for serving on my committee. I would also like to thank Alyssa Goodman for allowing me to interact with her research group over the past few years, and for being a part of my research exam committee. I would also like to thank all of my fellow graduate students, especially those within the Finkbeiner research group. Last but not least, I would like to thank the *WISE* and *Planck* teams for publicly releasing their data sets to the entire astronomical community.

Chapter 1

Introduction

1.1 Broad Overview of Interstellar Dust

The Galactic interstellar medium (ISM) consists of a variety of physical components. Chief among these are neutral atomic gas (e.g. H I), molecular gas (e.g. H₂, CO), ionized gas (e.g. H II gas) and dust. Astronomers consider dust to be any particulate matter ranging from ~ 10 angstroms to ~ 0.1 microns in size, spanning from ~ 50 atom molecules to mineral grains with millions of atoms. The very smallest dust particles are polycyclic aromatic hydrocarbons (PAHs), organic molecules made up of a small number of carbon rings to which hydrogen atoms are bonded. PAHs are the byproduct of combustion processes, and can be excited by a single UV or optical photon, leading them to emit in the mid-infrared ($\sim 3\mu\text{m}$ - $20\mu\text{m}$) via a series of bands attributable to vibrational stretching and bending modes. Larger grains, typically consisting of amorphous silicate or carbonaceous material, remain in thermal equilibrium with the interstellar radiation field (ISRF) by absorbing photons in the optical and UV, and then re-emitting this energy as a continuum in the far-infrared ($>100\mu\text{m}$).

Interstellar dust is largely confined to the plane of our Galaxy, inhabiting a disk a

few hundred parsecs in scale height. While dust only accounts for a miniscule 0.001% of the Milky Way's total mass, it plays a crucial role in the process of star formation. Absorption of UV photons by dust in dense region allows for the formation of molecular gas and subsequent protostellar collapse, thereby catalyzing the process of star formation. Later in the stellar lifecycle, evolved stars form new dust grains which are eventually dispersed back into the ISM. Interstellar dust is therefore a key element of the 'recycling' process by which stars are born and die. Despite its small fractional mass relative to stars, dark matter, and even interstellar gas, dust converts approximately $\sim 30\%$ of stellar radiation in a typical spiral galaxy from optical to infrared emission via absorption (Cox & Mezger, 1989). Thus it is clear that dust plays a key role in galaxy radiative transfer and evolution.

From the perspective of observational astronomy on Earth, dust also plays a critical, if often unwelcome, role in research across an extremely wide range of wavelengths and astrophysical phenomena. This stems from the fact that the Sun is embedded within the Galaxy's disk, and therefore all lines of sight to distant sources (even those at high Galactic latitude) suffer contamination from dust extinction in the UV and optical, as well as dust emission in the mid-infrared, far-infrared, and microwave regimes.

1.2 Observational History: Dust in the Optical

Arriving at this present understanding of interstellar dust and its importance to observational astronomy has taken centuries of scientific investigation, and our knowledge about the ISM continues to evolve as modern infrared space telescope surveys become ever more sensitive. The earliest observational evidence for interstellar dust came from optical observations of stars. Centuries ago, Sir William

Herschel first noticed empty patches in the Milky Way’s distribution of stars (Herschel, 1785). It remained contested long thereafter whether these dark patches indicated regions of low stellar density or rather absorption by intervening material. The controversy was definitively resolved in the early portion of the twentieth century, thanks to a number of detailed studies of the spatial distribution and colors of Milky Way stars observed in the optical. Starting with Friedrich Georg Wilhelm von Struve in 1847, several studies of the Milky Way’s stellar density suggested increasing extinction with heliocentric distance, at the level of ~ 1 magnitude per kiloparsec (Kapteyn, 1909). In 1919, Barnard weighed in, favoring the hypothesis of intervening obscuration of starlight, citing the cloudy appearance of dark regions on his photographic plates (Barnard, 1919). Trumpler (1930) first detected reddening – a higher amplitude of extinction at bluer wavelengths – based on observations of globular clusters. Finally, in the 1930’s a number of studies further characterized the starlight extinction curve, inferring an approximate scaling of extinction as $1/\lambda$, consistent with absorption by dust particles smaller than the wavelength of visible light (Rudnick, 1936; Hall, 1937; Stebbins et al., 1939). Herschel’s “holes in the heavens” are thus understood to be extreme cases of dust clouds with column density so high as to prevent detection of background stars. The idea of inferring the distribution of interstellar dust from optical observations of stars is currently being pursued with renewed enthusiasm, spurred on by the advent of billion star surveys like Pan-STARRS1 (Kaiser et al., 2002), supercomputers and Markov chains (see e.g. Sale et al., 2009; Green et al., 2014; Schlafly et al., 2014).

1.3 Observational History: Dust in the Infrared

Since these early observations, infrared astronomy has opened up a new window into studies of interstellar dust. Although Herschel first detected infrared radiation in 1800, technical challenges associated with observing astrophysical sources in the infrared delayed the launch of infrared astronomy until the 1960's. In particular these challenges were (1) the need for sensitive infrared bolometers (2) the need to cool these bolometers and other instrumentation sufficiently to reduce thermal backgrounds and (3) the lack of infrared atmospheric transparency. Low et al. (2007) provides a detailed review of the progress in detectors and instrumentation which has enabled modern infrared astronomy. In the 1960's and 1970's, a mixture of ground-based, balloon-based, plane-based, and rocket-based observations from NIR to FIR wavelengths first targeted and detected some of the brightest astrophysical objects in the infrared – the Moon, the outer solar system planets, NGC objects including bright external galaxies, the Galactic center and H II regions. Observations of H II regions and more diffuse areas of the Galactic ISM during the late 1970's and early 1980's were sufficient to hint at the excess of mid-IR emission relative to far-IR emission now known to be attributable to PAH bands (e.g. Andriesse, 1978; Price, 1981).

Launched in 1983, the Infrared Astronomical Satellite (*IRAS*) represented a critical turning point in the field of infrared astronomy and revolutionized the study of interstellar dust. *IRAS* imaged 96% of the sky in four broad infrared bands centered at $12\mu\text{m}$, $25\mu\text{m}$, $60\mu\text{m}$ and $100\mu\text{m}$, with angular resolution ranging from $30''$ to a few arcminutes. *IRAS* first revealed the ubiquitous infrared “cirrus”, diffuse cloud-like structure of interstellar dust emission, permeating even tenuous, high-latitude regions of the ISM (Low et al., 1984). *IRAS* also revealed diffuse Galactic emission at 12

microns, at a level much higher than would be consistent with dust populations responsible for the thermal continuum at longer wavelengths. The ultimate confirmation of interstellar PAHs came from the Infrared Space Observatory (ISO, Kessler et al., 1996) and later the Spitzer Space Telescope (Werner et al., 2004), which were able to perform mid-infrared spectroscopy of interstellar emission, finding bands consistent with laboratory measurements of PAHs.

1.4 History of Dust as a Cosmological Foreground

Returning more specifically to the problem of dust as a cosmological foreground, extinction of extragalactic sources has been a topic of research ever since Hubble noted the Galactic “zone of avoidance” imprinted on his Mount Wilson galaxy counts (Hubble, 1934). Starting from the approximation that Galactic dust is distributed uniformly within a thin cylindrical slab, early efforts focused on modeling Milky Way extinction as a function of Galactic latitude using scalings of the cosecant function, subject to various modifications (e.g. de Vaucouleurs & Malik, 1969). An important step forward occurred when Burstein & Heiles (1978) insisted on accounting for the “patchy” morphology of Galactic dust extinction. Specifically, Burstein & Heiles (1978) proposed to predict reddening by scaling velocity-integrated HI emission under the assumption of a constant dust-to-gas ratio (up to corrections based on the surface density of galaxies). The Burstein & Heiles (1982) map of HI-based reddening predictions became a widely adopted data product throughout the 1980’s and 1990’s.

The Burstein & Heiles (1982) reddening map remained the industry standard until Schlegel et al. (1998, hereafter SFD) proposed and implemented an alternative method of estimating extinction based on far-infrared thermal dust emission. SFD modeled thermal dust emission from $100\mu\text{m}$ - $240\mu\text{m}$ as a modified blackbody, with

extinction due to the emitting grains proportional to the optical depth. SFD employed a custom reprocessing of *IRAS* $100\mu\text{m}$ as a $6.1'$ FWHM template for dust optical depth, and fit the spatially varying dust temperature based on the ratio of DIRBE $100\mu\text{m}$ and DIRBE $240\mu\text{m}$, corrected for zodiacal light. Finally, SFD calibrated the resulting full-sky map of dust optical depth to reddening based on a sample of 384 elliptical galaxies with precisely predicted intrinsic optical colors. SFD further demonstrated that their reddening predictions were twice as accurate as those of Burstein & Heiles (1982). The SFD extinction map also represented a $\sim 6\times$ enhancement in angular resolution relative to Burstein & Heiles (1982).

Schlafly & Finkbeiner (2011) used roughly a quarter million stars with both SDSS photometry and spectroscopy to thoroughly validate the SFD reddening predictions, and to recalibrate SFD reddening estimates at the level of $\sim 15\%$. To this day, SFD remains the industry standard full-sky dust extinction data product.

The full-sky *IRAS*/DIRBE infrared maps of SFD have also been applied to predict thermal dust emission as a foreground for CMB observations (Finkbeiner et al., 1999, henceforth FDS99). The *IRAS*/DIRBE maps from SFD are limited to the wavelength range of $100\mu\text{m}$ - $240\mu\text{m}$ (3000 GHz to 1250 GHz), whereas the CMB peaks in the microwave at ~ 160 GHz. In order to extrapolate from *IRAS*/DIRBE wavelengths to the much lower frequencies relevant for CMB studies, FDS99 used the *COBE*/FIRAS spectra (Mather, 1982) to identify a globally best-fit dust spectrum over the entire 100-3000 GHz frequency range. FDS99 showed that no single MBB power law emissivity could adequately fit the FIRAS spectra; the peak of the dust SED near $160\mu\text{m}$ is too sharp relative to the spectrum's flattening from 100-350 GHz. FDS99 obtained a satisfactory fit by modeling the dust spectrum as the sum of two MBBs, each representing a species of dust grains with distinct absorption/emission properties, and with the two dust temperatures linked by the assumption of coupling

to the same ISRF.

1.5 Recent Developments

In recent years, a new wave of full-sky, high-resolution infrared satellite data has been acquired and publicly released. In the mid-infrared, the Wide-field Infrared Survey Explorer (*WISE*, Wright et al., 2010) has imaged the entire sky in four broad channels centered at 3.4, 4.6, 12 and 22 microns. For a full-sky mid-infrared survey, *WISE* boasts extraordinarily high angular resolution: $6.1''$ at $3.4\mu\text{m}$, $6.4''$ at $4.6\mu\text{m}$, $6.5''$ at $12\mu\text{m}$, and $12''$ at $22\mu\text{m}$. With regard to mapping interstellar dust, the *WISE* 12 micron channel is of most interest, owing to the prominent mid-IR PAH emission bands between $8\mu\text{m}$ and $18\mu\text{m}$.

In the far-infrared, the *Planck* satellite has provided superb maps of the entire sky at nine frequencies from 30 GHz to 857 GHz, with high signal-to-noise and high spatial resolution (Planck Collaboration et al., 2013a, 2015). *Planck* offers angular resolution of $33'$, $24'$, $14'$, $9.66'$, $7.27'$, $5.01'$, $4.86'$, $4.84'$ and $4.63'$ at 30, 44, 70, 100, 143, 217, 353, 545 and 857 GHz, respectively. In addition, *Planck* has measured polarized emission at frequencies of 353 GHz and lower. Although *Planck* was designed primarily as a CMB cosmology experiment, its maps are also dramatically enhancing our knowledge of the Galactic interstellar medium. In particular, *Planck*'s combination of high angular resolution and high sensitivity over a wide range of frequencies allows for the creation of a dust temperature map with $\sim 10\times$ enhanced angular resolution relative to DIRBE-based dust temperature maps.

Akari represents a third recently released full-sky, high-resolution infrared data set (Doi et al., 2015). *Akari* was a Japanese satellite which imaged the entire sky in two relatively narrow bands centered at $60\mu\text{m}$ and $160\mu\text{m}$, as well as two relatively broad

bands centered at $90\mu\text{m}$ and $140\mu\text{m}$. The *Akari* maps have angular resolution of approximately 1.5 arcminutes, but suffer far more contamination from artifacts than those of *Planck* in the FIR.

The recent releases of these high-resolution mid-IR and FIR data sets from *WISE*, *Planck* and *Akari* pose tremendously exciting opportunities to advance our understanding of the interstellar medium and build improved cosmological foreground models. In principle, the full-sky maps provided by these missions should allow for the creation of a next-generation dust data product which predicts extinction with up to $>50\times$ better angular resolution than SFD and includes a dust temperature correction with $>10\times$ enhanced resolution. This dissertation describes my initial efforts towards attaining this goal and details my creation of novel full-sky dust data products based on *WISE* and *Planck*, with an emphasis on enhanced angular resolution.

1.6 Outline

Chapter 2 presents our custom reprocessing of the entire WISE $12\mu\text{m}$ All-Sky imaging data set, resulting in a full-sky map of diffuse Galactic dust emission at $15''$ resolution. Chapter 3 presents our application of the Finkbeiner et al. (1999) thermal dust emission model to the *Planck* HFI maps, resulting in the first ever 100-3000 GHz *Planck*-based thermal dust emission model, as well as a dust temperature correction with $\sim 10\times$ enhanced angular resolution relative to Schlegel et al. (1998). We conclude in Chapter 4.

Chapter 2

WISE 12 micron Dust Emission Map

We describe our custom processing of the entire Wide-field Infrared Survey Explorer (WISE) $12\mu\text{m}$ imaging data set, and present a high-resolution, full-sky map of diffuse Galactic dust emission that is free of compact sources and other contaminating artifacts. The principal distinctions between our resulting co-added images and the WISE Atlas stacks are our removal of compact sources, including their associated electronic and optical artifacts, and our preservation of spatial modes larger than 1.5° . We provide access to the resulting full-sky map via a set of 430 $12.5^\circ \times 12.5^\circ$ mosaics. These stacks have been smoothed to $15''$ resolution and are accompanied by corresponding coverage maps, artifact images, and bit-masks for point sources, resolved compact sources, and other defects. When combined appropriately with other mid-infrared and far-infrared data sets, we expect our WISE $12\mu\text{m}$ co-adds to form the basis for a full-sky dust extinction map with angular resolution several times better than Schlegel et al. (1998).

2.1 Introduction

High-latitude complexes of resolved infrared sources were first identified by Low et al. (1984) using *IRAS* $60\mu\text{m}$ and $100\mu\text{m}$ observations (Wheelock et al., 1994). This so-called “infrared cirrus,” attributable to thermal emission from large dust grains in the interstellar medium, has since been probed in greater detail by a variety of instruments and detected over virtually the entire sky. Beginning in 1989, DIRBE mapped the full sky at ten infrared wavelengths from $1.25\mu\text{m}$ to $240\mu\text{m}$ with a reliable zero point, but inferior 0.7° angular resolution (Boggess et al., 1992). Later, the ISO and IRTS missions (Kessler et al., 1996; Onaka et al., 1996; Tanaka et al., 1996) established that polycyclic aromatic hydrocarbons (PAHs) are responsible for the presence of cirrus at mid-IR wavelengths, displaying strong emission bands at $8.6\mu\text{m}$, $11.3\mu\text{m}$, and $12.6\mu\text{m}$. Present day observatories are now capable of mapping the mid-IR and far-IR cirrus with excellent resolution and sensitivity. For example, *Herschel* has conducted pointed imaging observations from $70\mu\text{m}$ to $500\mu\text{m}$ with corresponding resolution ranging from $5.2''$ to $37''$ (Pilbratt et al., 2010). However, *Herschel* and *Spitzer* (Werner et al., 2004), each with field-of view $\lesssim 5'$, only observed a very small fraction of the sky over their lifetimes.

On the other hand, the Wide-field Infrared Survey Explorer (WISE, Wright et al., 2010) has imaged the full sky in four broad infrared bandpasses at high angular resolution (Cutri et al., 2012). In the two bluest channels ($W1 = 3.4\mu\text{m}$, $W2 = 4.6\mu\text{m}$), the Rayleigh-Jeans tail of stellar emission dominates, with images displaying vast numbers of compact sources, and diffuse dust structure that is typically discernible only at very low Galactic latitudes. At the other end, the reddest band ($W4 = 22\mu\text{m}$) is strongly dominated by our Solar System’s interplanetary dust, its thermal emission peaking between $15\mu\text{m}$ and $25\mu\text{m}$. While the zodiacal foreground

is still bright at $W3 = 12\mu\text{m}$, this bandpass was chosen to coincide with the strongest PAH emission features. As a result, Galactic cirrus is often prominent in $W3$, and has been imaged by WISE with high signal-to-noise over most of the sky in spite of bright zodiacal contamination that varies on large angular scales. Further, with a native resolution of $6.5''$, WISE $W3$ represents a dramatic improvement in angular resolution relative to previous mid-IR and far-IR full-sky data sets. This has motivated us to investigate WISE $12\mu\text{m}$ images as a way of mapping the diffuse mid-IR emission from Galactic interstellar dust and probing small-scale structure in the Milky Way ISM.

However, official WISE mission imaging data products do not include a clean map of the $12\mu\text{m}$ Galactic cirrus. To isolate diffuse $W3$ nebulosity, we have undertaken a custom per-exposure processing of the $W3$ data set and constructed custom stacks from these processed frames. In particular, we have:

- subtracted point sources and artifacts of bright point sources
- masked moving Solar System objects
- masked resolved compact sources
- eliminated single-epoch anomalies (e.g. satellite streaks and cosmic rays)
- corrected single-frame striping artifacts and per-quadrant offsets
- addressed time-dependent foregrounds due to the Moon and zodiacal light
- preserved large angular scale modes

Because PAHs trace the far-IR thermal emission from large grains, a particularly interesting application of diffuse $W3$ maps is the creation of high-resolution foreground templates and extinction estimates. In terms of predicting dust extinction, the industry standard is Schlegel et al. (1998, hereafter SFD), with $6.1'$

angular resolution limited by the asymmetric $4'$ *IRAS* beam. Even if we must smooth *W3* considerably relative to its native resolution, we can still characterize the distribution of interstellar dust over the full sky on angular scales many times smaller than the SFD/*IRAS* beam.

In §3.2 we summarize the relevant details of the WISE survey and its data products. In §3.3.2, we describe the point source catalog which enables us to model the PSF and perform PSF subtractions. In §2.4 we discuss our model of the PSF and related bright source artifacts. In §2.5 we describe our single-exposure processing pipeline that eliminates many contaminants of the diffuse *W3* emission. In §2.6 we discuss our custom stacks built from these processed single exposures. In §2.7 we discuss the full-sky map we are releasing¹ and conclude in §2.8.

2.2 The WISE Data

2.2.1 WISE Survey Strategy

The WISE satellite followed a 95 minute, Sun-synchronous, polar orbit about Earth, always pointing close to zenith and imaging at nearly 90° Solar elongation. Over a single orbit, WISE scanned all ecliptic latitudes β at a particular ecliptic longitude λ , and then all β at longitude $(\lambda + 180^\circ)$. Successive 8.8 s, 0.8° frames overlap by $\sim 10\%$ in the scan direction, while corresponding frames in successive orbits have $\sim 90\%$ overlap.

Because the All-Sky Release spans just under seven months, the vast majority of the sky at low ecliptic latitude has a resulting integer coverage of 10-12. For our purpose of mapping diffuse emission, this substantial redundancy is useful in compensating for single-frame artifacts such as cosmic rays and satellite streaks.

¹See <http://wise.skymaps.info>.

However, for typical locations on the sky, all 10-12 WISE epochs occur within a single ~ 20 hour time period. Unfortunately, the Moon and bright Solar System objects do not move far enough on this timescale to be rejected.

2.2.2 WISE Data Products

All WISE data products utilized throughout this work, including single exposures and catalogs, are drawn from the All-Sky Data Release (Cutri et al., 2012). The All-Sky Release features two distinct options for accessing WISE images. The “Atlas images” are a set of 18,240 co-adds, 1.5° on a side. While it might seem convenient for these co-adds to form the basis of our $12\mu\text{m}$ dust map, various considerations led us to conclude that we were best served by starting our analysis at the single-exposure level, then assembling our own custom stacks. First, each Atlas co-add is designed only to minimize the appearance of exposure boundaries within its footprint, but will not necessarily agree with its neighboring Atlas tiles on their overlaps. Additionally, the Atlas images contain all compact sources observed by WISE, which require extremely careful PSF subtraction in order to isolate the diffuse cirrus. The Atlas PSF must be defined in a complex way and would be difficult both to model and apply, given the multiple epochs combined at each Atlas pixel and the fact that the Atlas images have been smoothed.

Because of these Atlas co-add considerations, we have instead obtained WISE imaging data by downloading all ~ 1.5 million *W3* exposure sets, including the “intensity”, “mask”, and “uncertainty” files. The single exposures are referred to as “Level 1b” (L1b) images, are archived by the NEOWISE project (Mainzer et al., 2011), and are the lowest-level publicly available WISE imaging data. L1b images are square, 1016 pixels on a side, and have a platescale of $2.75''/\text{pixel}$. The x (y) axis of each L1b exposure is oriented perpendicular (antiparallel) to the scan direction.

To remove compact sources, we also make use of the WISE All-sky Source Catalog (§3.3.2) and WISE All-Sky Known Solar System Object Possible Association List (§2.5.3).

2.2.3 L1b Metadata Cuts

We discard a small subset of L1b exposures which, based on metadata alone, we know in advance will not be worthy of further processing. Our cuts generally follow the recommendations and cautions of Cutri et al. (2012). Specifically, we reject frames with `saa_sep` $\leq -5^\circ$, or `moon_sep` $\leq 12^\circ$, or `w3msknumsat` $> 500,000$ pixels, or with `dtanneal` $\leq 1,000$ s at high ecliptic latitude. These cuts remove 4.9% of the L1b images.

2.3 Point Sources

A major goal of our L1b processing is the PSF subtraction of each pointlike source from every exposure in which it appears, along with correction of the artifacts associated with bright sources. In order to model the PSF and artifacts, and eventually to subtract appropriately centered/rescaled PSFs, we require a point source catalog. We opted not to build a custom point source catalog from scratch. Rather, we make use of the source positions and fluxes from the WISE All-Sky Source Catalog, only considering those sources with `w3snr` > 2 (Cutri et al., 2012).

However, cutting only on `w3snr` > 2 did not yield a satisfactory catalog. In the following subsections, we describe the additional cuts and custom augmentations we have made in order to obtain a catalog that best subtracts pointlike sources while preserving diffuse nebulosity. In particular, we were forced to address issues with very faint and very bright sources; at intermediate fluxes, we found the WISE All-Sky




Figure 2.1 $0.15^\circ \times 0.28^\circ$ sub-region of a *W3* L1b image, centered at $(l, b) = (212.5^\circ, -0.45^\circ)$. Left: L1b data. Center: All-Sky Catalog, all nominal *W3* detections. Right: list of pointlike sources based on the COG cuts of §2.3.1.

Source Catalog extremely reliable.

2.3.1 Nebula Fragment Rejection

The WISE All-sky Source Catalog was engineered to emphasize completeness over purity. Consequently, a striking feature of the *W3*-detected sample ($w3snr > 2$) is the tendency for multitudes of dubious, faint sources to trace any sufficiently bright, filamentary dust structure (see Figure 2.1). Leaving this issue unaddressed, the filaments in PSF-subtracted images take on a distinctly “chewed up” appearance, compromising our estimate of the dust cirrus.

We were unable to identify any trivial cuts on e.g. signal-to-noise or source color that rectified this problem. However, we observed that the sources clustered on filaments tend to have a curve of growth (COG) more similar to that of a constant surface brightness profile than that of a PSF (see Figure 2.2). We measured the COG via a series of aperture magnitudes, $w3mag_n$, for $n = 1, \dots, 6$, available in the

All-Sky Catalog. The PSF COG was computed by integrating our model from §2.4 in the relevant apertures. We were able to satisfactorily reject nebula fragments by cutting those sources with COG more consistent with constant surface brightness than with the $W3$ PSF model. We also discarded sources with an undefined error $w3sigm_1$ on their innermost aperture magnitude, $w3mag_1$. The PSF versus constant surface brightness comparison was calculated in the sense of minimum absolute deviation, as the majority of All-Sky Catalog errorbars ($w3sigm_n$) were undefined for the faint sources of interest, precluding a formal likelihood analysis.

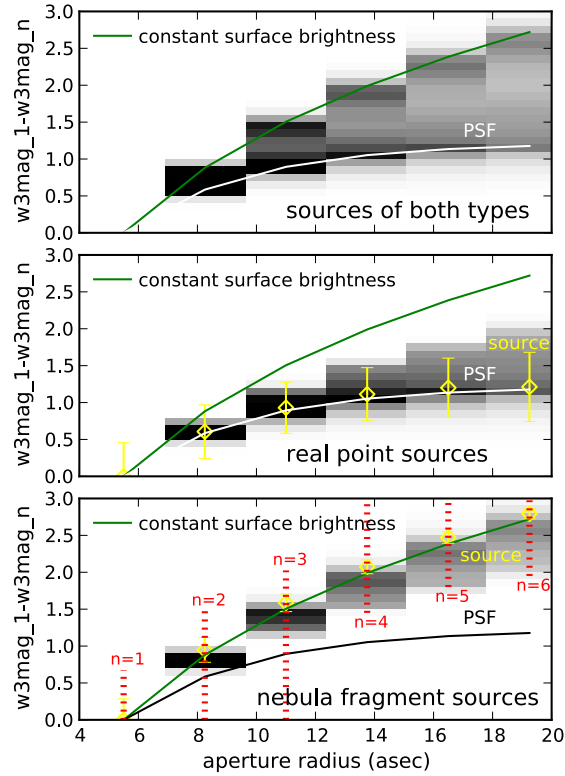


Figure 2.2 Top: $W3$ aperture 1 minus $W3$ aperture n vs. radius for a set of $\sim 250,000$ All-sky Catalog sources (grayscale), along with the expected curve for a PSF (white line) and constant surface brightness (green line). Middle: same for sources classified as true point sources by our COG cut, with an example such source shown (yellow). Bottom: same, but for sources classified as nebula fragments by our COG cut, along with an example such source (yellow).

The objects thus removed constituted a non-negligible fraction of All-Sky Catalog sources with $w3snr > 2$. Prior to the COG cut, the All-Sky Catalog contained ~ 167 million sources detected in *W3*; applying the COG filter reduced the number of remaining *W3* sources to ~ 41 million.

2.3.2 Custom Bright Source Fluxes

By employing our detailed model of the PSF wings (see §2.4), we were able to derive custom bright source parameters that yielded cosmetically optimal PSF subtractions. Our procedure ignores the saturated core of each bright source, fitting only an annulus containing the wings.

Specifically, we fit custom fluxes for those $\sim 15k$ sources in the All-Sky Source Catalog brighter than $w3mpro = 2$ (>5 Jy). This threshold has no special significance with regard to the WISE instrumentation, and its choice is driven by pragmatic concerns relating to e.g. our available computational resources. To fit each bright source, we consider an annulus about its centroid with outer radius $2.1'$. The inner boundary of the annulus was determined via a flood-fill of pixels near saturation ($>20,000$ DN), starting from the centermost pixel. After dilating this saturation mask by 3 L1b pixels, we excluded the resulting interior region from the flux fit. We then perform a linear least squares fit to the annulus pixel values. The model is the sum of a rescaled PSF and a constant offset representing the background level. Thus, there are two parameters: the background offset and PSF rescaling factor, from which we infer the source flux. For each source, such a fit is performed for each L1b image in which the entire fitting annulus falls within the image boundaries. The L1b pixels of each fit are inverse variance weighted based on the corresponding L1b uncertainty image. The flux assigned is taken to be the median of the individual L1b fluxes. Typically 9 L1b fits per source are performed.

We also computed custom centroids for bright sources in order to achieve cosmetically optimal PSF subtractions. After computing our custom fluxes, we recompute the centroid by fixing the background and flux, but fitting x , y offsets via Levenberg-Marquardt chi-squared minimization.

Figure 3.11 shows a representative sample of 25 bright sources drawn from ~ 300 square degrees of the sky which we processed in turn with the raw All-sky Catalog parameters and our custom parameters. On the whole, our custom parameters clearly provide substantial improvement in the monopole and dipole components of the residuals. For the purposes of this work, we only care about the cosmetic quality of our point source subtracted maps, and have therefore chosen bright source subtractions based on our custom parameters. We caution, however, that our custom bright source parameters have not been thoroughly vetted against other mid-infrared calibrators, whereas those of the official All-sky Source Catalog have been (Cutri et al., 2012).

2.4 Point Spread Function

Because the space-based WISE mission did not have to contend with atmospheric effects, we expect its PSF to be rather stable, and that a detailed characterization will lead to excellent PSF subtractions. Figure 2.4 shows a schematic of our model for the imprint² of a single bright point source in the $W3$ data. The model includes some features common to astronomical imaging: a bright PSF core, faint but extensive wings, and diffraction spikes. Additionally, appearing within the same exposure as the bright source itself, a doughnut-like optical “ghost” results from the WISE optics. This ghost is offset $9.5'$ in the scan direction relative to the source centroid. When a

²We use this term to refer to the PSF and additional non-linear after-effects in subsequent exposures.




Figure 2.3 Comparison of bright source subtraction residuals using All-sky Source Catalog parameters versus our custom-fitted parameters. We have optimized the fit to subtract the PSF wings. Blue outlines indicate pixels flagged with our “bright region of point source profile” mask bit (see Table 2.2). The diffraction spikes and ghost also appear explicitly in the bit-mask.

source is bright enough to saturate in its core, the following exposure one frame cycle later displays a saucer-like electronic defect termed the “first latent”. Another frame cycle after that (two exposures following the bright source detection) the electronics are still recovering and a fainter “second latent” is observed. For further details about these effects, see the Explanatory Supplement (Cutri et al., 2012). In the following subsections we describe our models for each of these features, which we later use to PSF subtract pointlike objects and correct bright source artifacts.

Our PSF modeling was iterative; after creating an initial model for the PSF wings based on out-of-the-box fluxes from the All-Sky Catalog, we used this model to fit custom bright source fluxes (§2.3.2). We then fed our refined fluxes back into the PSF/artifact modeling procedure. The models eventually used to create the stacks of §2.7 were all based on our custom bright source fluxes. Henceforth, the name `w3mag` will refer to our custom bright source fluxes, whereas `w3mpro` will refer specifically to values quoted directly from the All-Sky Catalog.

2.4.1 PSF Core

We cannot rely upon one single sample of WISE sources to model all portions of the W3 PSF. For instance, the extremely bright sources that have high signal-to-noise in their outer wings will be saturated in their cores, and thus useless in modeling the innermost several FWHM.




Figure 2.4 Our model for the imprint of a single bright point source, including PSF wings, core, ghost and first, second latents.

$W3$ sources begin to saturate in their very central pixels at $3.5 < w3mpro < 4.0$. In order to leverage the brightest non-saturated sources, we modeled the PSF’s inner $1.15'$ (10.6 FWHM, the “core”) using ~ 4600 All-Sky Catalog sources with $4 < w3mpro < 5$ and $|b| > 40^\circ$.

To fully reap the benefits of PSF subtraction at the single exposure level, our model must include spatial PSF variation across the chip, as a function of source centroid (x_c, y_c) . We begin by gathering cutouts from individual L1b images in which each source in our sample appears. These cutouts include those in which the PSF core is only partially contained within the image. This permits us to model the PSF variation all the way to the chip’s edges. Typically 10 such cutouts per source were collected. We sinc shift each cutout in both x and y directions by a fraction of a pixel such that all cutouts are precisely centered on the source centroid. Additionally, each L1b cutout is background subtracted and rescaled to an arbitrary common source flux. The (x_c, y_c) location of each cutout’s centroid within its original L1b image is also stored.

For each pixel location with respect to the PSF model centroid, the set of rescaled pixel values available from these cutouts is fit with a third order polynomial surface in (x_c, y_c) . We chose third order because our measurements of the PSF FWHM as a function of position in the chip revealed a dependence that required third order terms. The per-pixel fit is performed independent of other neighboring pixels and carried out via weighted linear least squares, including iterative outlier rejection. The weights are inverse variances based on propagation of errors from the L1b uncertainty masks. The result is a polynomial model of the PSF core wherein the zeroth order term represents the PSF at the center of the chip, and higher order terms correct for PSF variation across the chip.

2.4.2 PSF Wings

Our procedure for modeling the PSF wings is quite similar to that employed in modeling the PSF core (§2.4.1). Our model for the PSF wings consists of a rectangular region $14.9'$ on a side, and thus characterizes the *W3* PSF behavior over a diameter of ~ 140 FWHM.

In assembling a sample of sources with which to model the PSF wings, we aimed to use approximately the $\sim 100^{th}$ - 1100^{th} brightest high-latitude sources on the sky. Specifically, our sample was defined to be those sources with $-2 < \mathbf{w3mag} < 1$, and with $|b| > 15^\circ$. This sample contains 1046 unique sources. We found that the handful of $\mathbf{w3mag} < -2$ sources were predisposed towards anomalous defects such as bleeding, and thus would contribute systematic problems to the PSF model despite having nominally high signal-to-noise in the wings.

We again collect cutouts as in §2.4.1, but model the per-pixel spatial variation with only first order terms in x_c, y_c . Higher order terms did not appear necessary and tended to destabilize the per-pixel polynomial fit in outer regions of the wings where signal-to-noise is low. The resulting model of the PSF core and wings, tapered according to §2.4.5, is shown on a logarithmic stretch in Figure 2.5. The diffraction spikes are not treated specially relative to other pixels in the model of the PSF wings.

2.4.3 Ghost

The optical ghost location with respect to the PSF core depends strongly on the centroid coordinates of its parent source, particularly x_c . As x_c ranges from one edge of the image to the other, $x_c \approx 0$ to $x_c \approx 1015$, the ghost centroid shifts in position relative to x_c by $0.5'$ (~ 11 L1b pixels). We initially attempted to model this translation of the ghost relative to x_c with an approach similar to that in §2.4.1,

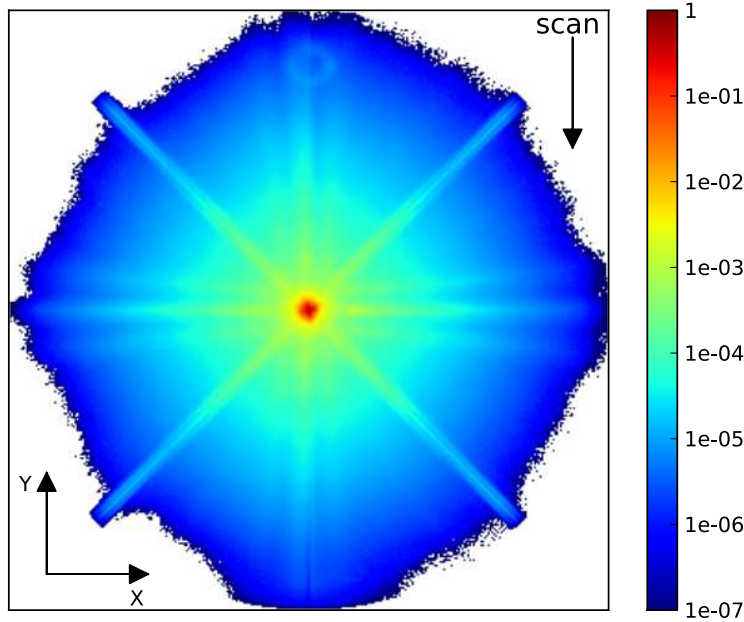


Figure 2.5 Our model of the *W3* PSF core and wings at chip center, on a logarithmic stretch, tapered according to §2.4.5. The image is 14.9' (325 L1b pixels) on a side. L1b *x*-axis, *y*-axis orientations, as well as the scan direction, are labeled with arrows.

§2.4.2, employing high-order per-pixel polynomials. However, we found this model did not satisfactorily capture the ghost's translation and became unstable for x_c near the L1b edges.

We thus resorted to a ghost model in which the ghost shape remains constant, but can translate with respect to x_c and can be amplified/suppressed by an overall rescaling factor f_g that varies with (x_c, y_c) . To create such a model, we again started by extracting many L1b cutouts. The same 1046 sources used to model the PSF wings were also used to model the ghost. We began by computing the centroid of the ghost in each cutout, and fit a second-order polynomial model to the ghost's offset relative to (x_c, y_c) . We then used this ghost translation model to rectify all the cutouts to a common fiducial ghost centroid with respect to (x_c, y_c) . The rectified cutouts allowed us to compute the ghost profile via a weighted average with iterative

outlier rejection. To determine f_g , we modeled each rectified cutout as a sum of an offset and f_g times the ghost profile. Finally, we fit a second-order polynomial in (x_c, y_c) to the set of f_g values we obtained. With $f_g = 1$ for (x_c, y_c) at the chip center, our resulting model always has $0.85 < f_g < 1.06$.

2.4.4 Latents

First Latent

The first latent always appears with the same centroid coordinates (x_c, y_c) as its parent bright source, but one exposure later. By “one exposure later” we do not necessarily mean an increment of one in the L1b `frame_num` parameter. Rather, we mean an exposure which follows the bright source detection by exactly one frame cycle, 11.1 s. There is not always such an exposure for every bright source detection.

We collect the relevant cutouts by identifying those frames that follow detection of sources with $-3 < \text{w3mag} < 1$, $|b| > 15^\circ$ by 11.1 s. There are 1088 such sources. Our latent model is built on a rectangular footprint, 14.9' on a side. To characterize the first latent shape, we assume the latent scales linearly with source flux, as do other PSF components. We then rescale our cutouts, and determine the latent profile via a weighted average with iterative outlier rejection.

In reality, the first latent does not scale linearly with parent source flux. As parent flux decreases, the first latent amplitude increases relative to the parent source flux (until the parent source no longer saturates, in which case no latent will arise). Our model accounts for this fact by assuming the latent shape to be independent of parent flux, but allowing for a scalar nonlinearity factor f_l that is a function of parent flux. f_l is defined to be a correction relative to perfect linear scaling, and would be identically unity if the latent behaved as a true component of the PSF, proportional in amplitude to parent source flux. We compute f_l by creating a first latent model in

each bin of 0.2 mag and computing their relative amplitudes. A smooth model for f_l results from fitting a polynomial to these relative amplitudes. f_l varies between 0.38 and 1.25.

Second Latent

We similarly take the second latent to be an artifact appearing 22.2s after a bright source detection, with (x_c, y_c) identical to that of the parent source. The procedure for fitting the second latent profile is analogous to that for the first latent. To model the second latent we use those sources with $-3 < \text{w3mag} < -0.5$ and $|b| > 15^\circ$. Our second latent model is rectangular, $12.4'$ on a side. Because the second latent appeared less strongly nonlinear and was more difficult to characterize in narrow bins of w3mag than the first latent, no nonlinearity correction was computed. We also found that higher-order latents were too faint to reliably model.

2.4.5 Tapering

Our models for the PSF wings and ghost, as well as the latents, are defined on rectangular footprints, and may not necessarily asymptote identically to zero precisely at their very edges. Thus, subtracting our rectangular models would introduce extra power in modes oriented parallel to the x, y axes of the chip. For this reason we taper the aforementioned models to zero before performing subtractions. This is achieved by premultiplying the models by weights that decline smoothly from unity in each model's center to zero at each model's boundary. We denote the angular extent of a particular model in the x (y) direction θ_x (θ_y). In order to define our tapering weight we first compute a radial coordinate r relative to the model's central pixel:

$$r = \sqrt{(\theta_y/\theta_x)^2 \Delta x^2 + \Delta y^2} \quad (2.1)$$

Table 2.1. Model Sizes & Tapering Parameters

Model	θ_x	θ_y	r_{in}	r_{out}
PSF Wings (bright)	14.9'	14.9'	6.7'	7.4'
PSF Wings (faint)	5.3'	5.3'	1.9'	2.6'
Ghost	6.2'	4.0'	1.7'	2.0'
First Latent	14.9'	14.9'	6.7'	7.4'
Second Latent	12.4'	12.4'	5.5'	6.2'

Where Δx and Δy are displacements relative to the model centroid. For $r \leq r_{in}$, the tapering weight is unity, and for $r \geq r_{out}$ the tapering weight is zero. For $r_{in} < r < r_{out}$, the weight transitions from 1 to 0 linearly with respect to r . Table 2.1 lists θ_x , θ_y , r_{in} , r_{out} for the models that we taper prior to subtraction.

2.4.6 Relating Counts and Magnitudes

In order to relate magnitudes (`w3mag`) to L1b counts (DN), we must accurately determine the “zero-point magnitude” m_0 , such that a source with integrated flux of 1 count corresponds to `w3mpro` = m_0 .

We chose to define our m_0 with respect to an aperture of diameter 1.8'. We then conducted aperture photometry on the same set of high-latitude, unsaturated, moderately bright sources used to model the PSF core. In order to determine m_0 , we performed a least-squares fit with iterative outlier rejection to the trend of `w3mpro` vs. $2.5 \log_{10}(F_{DN})$, F_{DN} signifying the total counts within the 1.8' aperture. The result determined $m_0 = 17.645$. Reassuringly, the best-fit slope was indeed very nearly unity. More generally, we found no evidence for significant nonlinearity of `w3mpro` for unsaturated sources at any point throughout this work.

We found that m_0 was sufficiently stable with time throughout the All-Sky Release

portion of the WISE mission to be considered a constant. Similarly, we did not find any evidence for temporal variation of the PSF shape.

2.5 Single Exposure Pipeline

We have built a single-exposure pipeline meant to remove and/or flag several contaminants affecting the diffuse $12\mu\text{m}$ emission we wish to isolate. Each of the ~ 1.5 million L1b exposures is processed independently. The pipeline attempts to remove all traces of pointlike sources by combining the All-Sky Catalog (including our custom bright fluxes and COG cuts) with our models of the PSF and related defects. SSOs are also interpolated over and flagged. Additional single-exposure mask bits are set, and propagate into extensions of the tile co-adds, as discussed in §2.7. The following subsections detail our L1b processing steps.

2.5.1 L1b Mask

Each raw L1b intensity image is accompanied by a corresponding `*-w3-msk-1b.fits` bit-mask provided by the WISE team. We will refer to such files as msk files, so as not to confuse them with our own custom L1b masks. The first step in our L1b processing is to create a mask of bad pixels in the raw L1b image which we must interpolate over, based on the msk file. We dilate the cosmic ray mask by a kernel of sidelength 3 pixels. Bad pixels are then taken to be those flagged in the dilated CR mask, or with any of the static bad pixel msk bits set, or with a NaN value. We immediately interpolate over this mask before proceeding.

2.5.2 Static Point Sources

Next we move on to correcting for static pointlike sources. From our custom catalog, we retrieve a list of all those sources whose PSF wings or ghost image, as represented by our models, have non-zero overlap with the L1b image at hand. This includes objects with centroids that fall outside the boundaries of the L1b image.

For each point source in the list, we compute its centroid coordinates (x_c, y_c) based on the L1b WCS and create a PSF model including the core, wings and ghost. The core and wings are computed from the per-pixel polynomials in (x_c, y_c) described in §2.4.1, §2.4.2. The ghost image is derived based on the model of §2.4.3, taking (x_c, y_c) into account through the ghost centroid translation and the amplification factor f_g . When computing these models, we always bound x_c, y_c to lie within the image, so as to not extrapolate beyond the limitations of our various polynomial corrections in x_c, y_c . We renormalize the PSF according to each source’s **w3mag** value and the prescription of §2.4.6.

Before finally subtracting this PSF model from the L1b image, two further steps are taken: (1) we sinc shift appropriately such that the observed and model source centroids match at the sub-pixel level (2) we apply the tapering prescription of §2.4.5 so as not to introduce rectangular boundaries into the model-subtracted images. These two steps are also applied before subtracting the first and second latent models as described in §2.5.2, §2.5.2 below.

It was not necessary to apply the full 14.9’ model of the PSF wings to arbitrarily faint sources. Rather, for sources deemed faint, we subtracted a 5.3’ cutout of the PSF core/wings. This smaller model does not include the ghost. We set the threshold distinguishing bright versus faint at **w3mag** = 4.2 for typical coverage, though the exact cutoff was scaled on a source-by-source basis to account for coverage **w3m**.

For sufficiently bright sources, the innermost pixels of the PSF subtraction tended to show strong residuals. Consequently, we interpolate over such regions, using a circular mask of size that varies with `w3mag`. Such interpolation only occurs for sources with `w3mpro` ≤ 9 . We flagged pixels interpolated over in this manner with bit 4 of our mask. These interpolations affect a very small fraction, $<0.25\%$, of the sky.

Based on the appropriately shaped and scaled PSF model derived for each source, a corresponding bit-mask is created. Saturated pixels within the PSF core ($>30,000$ DN) are flagged with bit 0. Bright regions of the PSF profile (>250 DN) are flagged with bit 2. Pixels within the ghost model that have values greater than 10^{-4} times the peak value of the PSF core are flagged with bit 1. This mask bit is meant to trace the general outline of the entire ghost model. Especially bright pixels within the ghost model (>20 DN) are flagged with bit 5.

First Latent

The presence of a first latent in a given L1b frame indicates that a bright source at the latent's (x_c, y_c) location was observed in a frame 11.1 s prior. For each L1b frame, we attempt to identify the frame 11.1 s prior to it. If no such frame exists, then no first latent corrections are made. If such a frame exists, we obtain a list of point sources bright enough to induce a first latent that were observed within the prior frame. The threshold for inducing a latent is `w3mag` < 4 . An appropriately rescaled latent model, taking into account the nonlinearity correction f_l , is then subtracted at the location of each bright source in the prior exposure. The first latent flag (bit 3, see Table 2.2) is set for pixels within a diameter of 9 pixels relative to each first latent centroid.

Table 2.2. Bit-mask summary

Bit	Description
0	saturated point source core
1	point source ghost
2	bright region of point source profile
3	first latent of point source
4	PSF subtraction residual interpolated over
5	bright region of point source ghost
6	SSO interpolation
7	resolved compact source
8	second latent of point source
9	third latent of point source
10	fourth latent of point source
11	bright SSO ghost
12	bright SSO latent
13	point source diffraction spike
14	saturated pixel not in static mask
15	Moon contamination
16	RC3 optical galaxy
17	big object (M31, LMC, SMC)
18	Solar System planet
19	reference comparison failure

Table 2.2—Continued

Bit	Description
20	line-like defect
21	low integer frame coverage
22	ecliptic plane

Second Latent

The procedure for correcting and flagging the second latent is virtually identical to that for correcting the first latent. The important distinction is that the relevant list of bright sources will correspond to the frame 22.2 s prior to the L1b image being processed, if such a prior frame exists. We set the threshold for second latent correction at $w3mag < -0.5$.

Third, Fourth Latents

As mentioned in §2.4.4, we only modeled the first and second latents. While we did not model the third and fourth latents, we do include mask bits to flag their locations (see Table 2.2). We carefully inspected the imprints of very bright sources in early versions of our stacks to determine which sources required mask bits to be set for third and fourth latents. We found that sources with $w3mag < 0$ required the third latent to be flagged, and those with $w3mag < -1.5$ required the fourth latent be flagged as well. Flagging the third (fourth) latent simply involved acquiring the source list for the frame 33.3 s (44.4 s) prior to the frame under consideration.

Diffraction Spike Mask

While the diffraction spikes were modeled simply as a part of the PSF wings, we judged that a special mask bit labeling potential diffraction spike residuals might be

of value. Using our PSF model, we calibrated the angular extent over which diffraction spikes would be detected as a function of source flux and integer coverage, conservatively assuming all epochs to have identical scan direction and negligible background structure. The threshold for assigning a diffraction spike mask is `w3mag` < 4.25 for typical coverage.

We then computed two corrections to this simplistic model to handle relevant edge cases. The first problematic scenario arises near the ecliptic poles, where scans converge from many different directions and integer coverage is very high. Applying the naive model would result in excessively large, nearly circular diffraction spike masks in our tile stacks at very high $|\beta|$. Consequently, we create an effective magnitude by adding to `w3mag` a term that accounts for diffraction spike attenuation due to varying L1b exposure orientation. We then use this effective magnitude instead of `w3mag` to compute the appropriate diffraction spike extent.

The second case is that of bright, highly structured background emission, e.g. in the Galactic plane. In such regions, the systematic variation and elevated Poisson noise of the background will diminish the angular extent over which diffraction spikes and their residuals are readily apparent. To quantify such background structure, we rely upon SFD *IRAS* $100\mu\text{m}$ (i100). We therefore computed \mathcal{R}_{100} , an $N_{side} = 64$ HEALPix map (Górski et al., 2005a), wherein each $\sim 1^\circ$ pixel stores the RMS of the i100 values sampled at the $N_{side} = 512$ pixel centers nested within its boundaries. \mathcal{R}_{100} is typically of order 10^{-1} MJy/sr at high $|b|$, but can be of order 10^3 MJy/sr in bright regions like the Galactic center. A term logarithmic in \mathcal{R}_{100} was added to the effective magnitude we used to compute the appropriate diffraction spike mask extent for $\mathcal{R}_{100} > 3$ MJy/sr. This correction adequately reduced the diffraction spike extent in regions with substantial diffuse structure.

2.5.3 Solar System Objects

To some extent, the outlier rejection steps in our mosaicking process (§2.7) eliminate moving objects, without any special treatment of SSOs in particular. However, there are many cases in which moving objects proceed slowly enough so as to contaminate the same location on the celestial sphere at multiple epochs. Furthermore, the WISE team provides a convenient catalog of potential SSO detections – the WISE All-Sky Known Solar System Object Possible Association List (Cutri et al., 2012, henceforth PAL). For these reasons, we chose to incorporate steps for masking and flagging SSOs into our L1b processing pipeline.

The PAL uses the orbits of known SSOs to calculate the expected locations of these objects within each L1b exposure. If there exists a WISE detection of some source nearby to one of these predicted locations, the measured properties of that source will be recorded in the PAL. Importantly, each PAL entry contains an L1b frame identification, coordinates of the source within that L1b image, and the source’s profile-fit magnitude, again designated `w3mpro`. In some cases, the associations may be inertial sources, or spurious detections, as PAL entries are based on single-epoch photometry. For our purposes, sample purity is not of much concern. Even if we interpolate unnecessarily in some cases, the redundant WISE coverage will tend to render such a mistake imperceptible.

For each SSO in a given L1b exposure, we produce a mask centered about the SSO centroid that scales in size with the detected object’s brightness. These masks each contribute to bit 6 of our L1b bit-mask (see Table 2.2), which denotes the pixels we have interpolated over within the L1b image to remove the SSO core/wings. The mask about an SSO’s centroid consists predominantly of a circular region, with diameter varying from 0.3′ (7 L1b pixels) for `w3mpro` ≥ 10 to 3.5′ (83 L1b pixels) for

$w3mpro \leq 1.4$. In addition, we apply our diffraction spike model of §2.5.2 to SSOs with $w3mpro < 3$. We interpolate over these SSO diffraction spikes at the L1b level, and the affected pixels are also flagged with bit 6. We considered PSF subtracting SSOs rather than simply interpolating over them, but found that the single-epoch PAL photometry was too unreliable and yielded a cosmetically inferior result.

In addition to addressing each SSO’s core, wings and diffraction spikes, we take further actions with regard to artifacts of very bright SSOs. For SSOs with $w3mpro < 2.5$, we mask and interpolate over an elliptical region surrounding the object’s ghost image. SSO ghost interpolation is registered in our bit-mask with bit 11. Similarly, for SSOs brighter than $w3mpro = 1$, we interpolate over the first latent and correspondingly set mask bit 12. Figure 2.6 provides an example of our SSO masking and interpolation results.




Figure 2.6 $0.4^\circ \times 0.4^\circ$ co-add cutout at $(l, b) = (355.68^\circ, 43.84^\circ)$, $(\lambda, \beta) = (226.74^\circ, 13.19^\circ)$, illustrating our SSO mask bits. Green, orange, and red represent bits 6, 11, and 14 respectively (see Table 2.2). The bright SSO (381 Myrrha) reveals the detailed nature of bit 6, which incorporates the diffraction spike model of §2.5.2.

2.5.4 Bit-mask Summary

Table 2.2 lists all mask bits included in the WSSA tiles of our data release.

Throughout §2.5, bits 0-6, 8-13 have been explained. Bit 14 flags pixels that are

saturated ($>30,000$ DN), but not marked as suspect by any of the msk file static mask bits. This can be useful downstream in identifying regions of extremely bright nebulosity that have reached saturation. Bits 7, 15-18 are only defined at the co-add level and are explained in §2.6.5.

2.5.5 Smoothing and Rebinning

Before writing our processed L1b files, we smooth to $15''$ FWHM, trim the smoothed image to 1000×1000 , and bin down by a factor of two in each dimension. We write the outputs of our L1b processing pipeline as multi-extension FITS files. For each raw L1b file in the All-Sky Release we write a corresponding processed file. The primary extension is the fully “cleaned” image, where we have removed point sources and their artifacts and SSOs. Extension 1 is the “dirty” image, which is merely a smoothed and rebinned version of the raw L1b image after interpolating as described in §2.5.1. Extension 2 is our bit-mask of §2.5.4, after dilation and rebinning that mimics the smoothing and rebinning of the first two extensions. These processed L1b images are not included in our data release, but will be made available upon request.

2.6 Stacking

We chose to co-add our processed L1b images on astrometric footprints analogous to those of the *IRAS* Sky Survey Atlas (ISSA) tiles (Wheelock et al., 1994). We define a set of “WSSA” astrometric footprints with the same central (α, δ) values and same orientations relative to north as the original ISSA tiles, but precessed to J2000. As a result, each of our WSSA tiles does not identically cover the same region of the sky as its ISSA counterpart. Also, to accommodate our $15''$ FWHM, WSSA tiles have a pixel scale of $5.625''/\text{pixel}$, as opposed to $1.5''/\text{pixel}$ for ISSA tiles. There are 430 such

$12.5^\circ \times 12.5^\circ$ WSSA tiles, with centers tracing out a series of isolatitude rings in J2000 equatorial coordinates. These tiles are not mutually exclusive; over the full sky, the median tile coverage is 1, but the mean coverage is 1.6, with a maximum of 5.

Our basic strategy in producing each WSSA tile is to start from the tile’s center and build outward, matching each image with the existing stack via a single offset before adding it in. Admittedly, this is a simple-minded approximate solution to the matrix inversion problem posed by best matching overlaps between neighboring exposures. In order to ensure that our simplistic procedure is feasible, we take two additional steps before beginning the stacking process, described in the following two subsections.

2.6.1 Gradient Correction

Overlapping L1b images may differ by more than a simple relative offset. This is especially true in regions that were imaged six months apart with opposite scan directions or near the ecliptic poles. In such cases, the zodiacal light can differ substantially between epochs. Even though the zodiacal light at a given epoch generally varies on spatial scales larger than a degree, this foreground is so strong that its gradient is readily apparent in single L1b exposures. As such, to ensure that our stacking can match neighboring exposures with a single offset, it is important to compute a correction for each exposure that removes time-dependent first order structure.

To arrive at such a gradient correction, we compare each L1b exposure to SFD $100\mu\text{m}$ on the same footprint. For each cleaned L1b exposure, we interpolate generously over bright source artifacts using our masks, then smooth to $6.1'$ FWHM to match i100 resolution and bin down $5\times$ in each dimension. We denote each pixel with coordinate x_i, y_i in this smoothed, rebinned $W3$ image as w_i . We then fit the

following model to w_i :

$$w_i = f_{wi}i_i + \nabla_x x_i + \nabla_y y_i + C \quad (2.2)$$

Where f_{wi} , ∇_x , ∇_y , and C are the scalar parameters to be fit and i_i are the i100 values corresponding to w_i . The fit is performed via unweighted linear least squares, with iterative outlier rejection. ∇_x , ∇_y then specify x , y gradients that make the processed L1b exposure most consistent with Galactic emission, free of slowly varying time-dependent foregrounds. f_{wi} is a factor that allows small-scale dust structure to be nulled, accounting for the fact that conversion factor between $W3$ and i100 will be different for different L1b pointings. Similarly, C absorbs any constant offset arising from the WISE and i100 zero points. Our model for the time-dependent gradient in each exposure is computed from ∇_x , ∇_y , and subtracted before the mosaicking process begins. We only desire the gradient terms, and make no use of f_{wi} , C downstream.

2.6.2 Pairwise Rejection

Despite our metadata cuts (§2.2.3), some L1b images with severe systematic problems remain and could potentially corrupt the stacking process if allowed to contribute to our co-adds. To remove such images, we employ a pairwise comparison between overlapping L1b exposures.

Specifically, each gradient-corrected, processed L1b image is compared to up to 20 of its nearest neighbors at the locations of $N_{side} = 2048$ HEALPix pixel centers on which the frames overlap. For each pair, we compute the RMS of the differences between the sampled values. Then, for each image, we compute the median of these RMS values, \mathcal{R}_{med} . We then reject the highest \mathcal{R}_{med} frame with $\mathcal{R}_{med} > 500$ DN,

updating \mathcal{R}_{med} for all other frames after each rejection, until no frames remain with $\mathcal{R}_{med} > 500$ DN. $\sim 0.4\%$ of images are rejected by the pairwise comparison, but these frames would have inflicted a disproportionate, negative influence on our stacks given the severity of their systematic problems.

2.6.3 Initial Stack

For each WSSA tile, we co-add the constituent L1b exposures onto the relevant astrometric footprint, with pixel scale of $5.625''/\text{pixel}$. $12.5^\circ \times 12.5^\circ$ co-adds are thus 8,000 pixels on a side. Such tiles are well-sampled because we have smoothed our processed L1b images to $15''$ FWHM.

We obtain a list of L1b exposures that will overlap the tile’s astrometric footprint, sorted by proximity to the tile center. We apply the gradients of section §2.6.1 to each cleaned L1b image, and then use the pairwise rejection step to reject the typically small number of problematic exposures. Starting with the exposure closest to the center, we add the cleaned L1b exposures into the stack one at a time, applying a single offset to make each exposure best agree with the existing stack. This offset is simply the median value of the difference between the exposure at hand and the existing stack on their overlap. As we stack, we also keep track of the minimum and maximum single-exposure pixel values at each pixel in the co-add footprint. We also store an integer coverage map, incrementing the relevant pixels each time an L1b image is added to the stack. For locations in the completed co-add that have coverage > 2 , we subtract out the average of the minimum and maximum images at those locations, and adjust the coverage map appropriately. We will refer to the average of the minimum and maximum images as the “transient artifact image”. This outlier rejection step eliminates satellite streaks and any other unusual single-epoch features not specifically addressed by our L1b processing pipeline. We also apply the

per-frame offsets to the dirty extension of our L1b images to create a corresponding stack with all point sources, bright source defects, and SSOs present.

2.6.4 Reference Comparison

With our initial WSSA tile stacks in hand, we can perform some further refinements, using the existing stack as a reference. There are two typically low-level problems we address by making use of our reference stack: (1) per-quadrant polynomial offsets in individual L1b exposures and (2) striping artifacts in L1b images.

Per-quadrant Polynomial Warps

We have already addressed scalar intensity offsets and gradient corrections on a per-exposure basis. But within a given exposure, the four quadrants can be offset relative to each other due to four-amplifier detector readout. These offsets can be simple scalars, but in some cases involve higher order polynomials (e.g. a gradient across one quadrant that is absent in the other quadrants). We refer to these polynomial offsets as per-quadrant “warps”. For fixed position on the celestial sphere within our co-adds, such warps in the various contributing L1b quadrants are not correlated. As such, any warping will be suppressed by $\gtrsim 90\%$ in our initial stacks.

Thus, we may reasonably consider our initial stacks to be nearly free from the effect of such per-quadrant warping, and as a result use them as reference templates to correct individual L1b quadrants. Doing so can help eliminate any jagged L1b quadrant edges that persist at low levels in the initial stacks. For each tile, we gather the same sample of L1b exposures as in §2.6.3. We then rebuild the mosaic from scratch, one L1b quadrant at a time, first fitting a per-quadrant warp relative to the reference, and then subtracting the warp model before adding the corrected quadrant into the revised stack.




Figure 2.7 $1.2^\circ \times 0.6^\circ$ region within WSSA tile 221, at $(l, b) \approx (5^\circ, 48^\circ)$. Top: before reference warping. Bottom: after reference warping, with jagged L1b exposure boundaries eliminated.

Our warp model is simply a polynomial in the the relevant cleaned L1b quadrant's x, y coordinates. For each such quadrant, we fit a fourth order polynomial to the difference between the cleaned L1b image and the reference stack. The fit is performed via unweighted linear least squares. With the best-fit polynomial computed, we evaluate the χ^2 goodness of fit between the reference stack and the corrected L1b quadrant. We take the expected pixel variance in each pixel of the cleaned L1b image to be:

$$\sigma_i^2 = f_s^2 \left(R^2 + \frac{N_i}{g} \right) + (\delta g N_i)^2 + (\nabla N|_i \delta \theta)^2 \quad (2.3)$$

Where N_i is the cleaned L1b pixel value in DN, $R = 16.94$ DN is the readnoise, $g = 6.83 e^-/\text{DN}$ is the gain, f_s is a constant that accounts for our smoothing/rebinning, δg is an estimate of the fractional gain variation, $\nabla N|_i$ is the magnitude of the

reference image gradient evaluated at pixel i , and δ_θ is an estimate of the L1b astrometric error. N_i is floored at zero in the second term, since it is possible that e.g. a bright source residual could have negative values. $\nabla N|_i$ is calculated based on the reference stack, not the L1b image being compared to the stack. Based on studies of point sources presented in Cutri et al. (2012), we chose $\delta g = 0.025$ and $\delta_\theta = 0.4''$. The first two terms account for statistical errors in the cleaned L1b image. The second term allows for the possibility that the per-quadrant L1b gain may vary slightly. The final term accounts for astrometric imperfections that could cause the L1b exposure to be out of alignment with the reference stack. We assume that statistical errors in the reference stack are negligible relative to those in the L1b quadrant of interest.

If the $\chi^2/\text{D.O.F.}$ is greater than 7 for a particular L1b quadrant after applying the best-fit polynomial warp, that quadrant is excluded from the stack. 2% of L1b quadrants were thus rejected. Figure 2.7 shows a region in which L1b boundaries in the initial stack of §2.6.3 are corrected by our quadrant warping procedure.

Destriping

L1b exposures display low-level striping. Complicating matters, this striping occurs with different amplitudes and on different wave numbers from exposure to exposure, and even from quadrant to quadrant within single exposures. Still, a particular set of relatively few modes tend to be enhanced, as the stripes within a given quadrant are typically oriented parallel to the scan direction and have wavelengths of $\sim 2.5'$ - $4'$. To avoid high-pass filtering diffuse structure, our destriping never modifies modes larger than $4.6'$. We assume the striping to be an additive rather than multiplicative defect.

As such, our first step towards correcting L1b striping is to make a mask in Fourier space of the typically contaminated modes, using $\sim 20,000$ cleaned L1b images at very high $|b|$, where there are few compact sources and diffuse structure is minimal. For




Figure 2.8 $1.6^\circ \times 1.3^\circ$ sub-region within WSSA tile 339, at $(l, b) \approx (188.7^\circ, 56.2^\circ)$. Top: before destriping. Middle: after destriping. Bottom: stripe model. The stripes are parallel to the scan direction.

each of these L1b exposures, we compute the isotropic Fourier power spectrum in each quadrant. We label any mode “contaminated” which has a modulus-squared that is a $>5\sigma$ outlier relative to the appropriate isotropic value. We construct a mask of frequently contaminated modes by counting the number of occurrences of such outliers at each pixel within the quadrant FFT.

For each quadrant that successfully passes the quadrant warping step of §2.6.4, we attempt a destriping correction before finally adding the quadrant into the reference-corrected stack. We take the FFT of the quadrant under consideration and identify any $>5\sigma$ outliers relative to the isotropic power spectrum. We then modify the FFT values of those outliers which coincide with modes we previously flagged as frequently contaminated. We rescale such outliers to have a modulus that matches the appropriate isotropic value, leaving the phase unchanged. Transforming back to the spatial domain yields a destriped image.

For most L1b quadrants, this correction removes all apparent striping artifacts. However, there are cases in which the correction can go awry. For example, the residuals of bright source artifacts can be aligned with the scan direction, wreaking havoc with our analysis meant to deal with low-level striping of the same orientation. We therefore incorporated destriping into our reference comparison process, in order to robustly identify such problematic instances. Specifically, we perform a chi-squared comparison between each L1b quadrant and the reference stack, both before and after destriping. We use the pixel variances of §2.6.4. We then apply the destriping correction only if (1) the destriping correction improves the $\chi^2/\text{D.O.F}$ and (2) more individual pixels within the quadrant have their χ^2 improved than worsened by destriping. Since the striping is most apparent in regions of very low signal/background, and the destriping correction is very successful in such regions, we find this “do no harm” approach to be satisfactory. Figure 2.8 shows an example of

our destriping correction.

2.6.5 Mosaic-Level Mask Bits

The following subsections describe mask bits defined at the mosaic level, rather than the L1b level.

Moon Mask

While our metadata cuts remove L1b frames pointed near the Moon (`moon_sep` $\leq 12^\circ$), off-axis Moon-glow can adversely affect exposures with much larger `moon_sep` values. In fact, prominent off-axis Moon illumination features extend out to many tens of degrees away from the Moon’s location. Accounting for the WISE survey strategy, this illumination detrimentally contaminates a non-negligible fraction, $\sim 5\text{-}10\%$, of the sky. Adverse effects on our stacks arise because the extended Moon-glow violates our assumption that the differences between exposures can be accounted for by an offset and linear gradients. For example, Moon-glow diffraction spikes can imprint strong, higher-order structure within a single L1b frame (see Cutri et al. 2012, §2.4b).

We derived an empirical model of *W3* off-axis Moon illumination in order to flag portions of our stacks which suffer serious Moon-glow corruption. In ecliptic coordinates, Moon-contaminated regions trace out a series of narrow vertical bands, a few degrees wide and generally separated by $\sim 30^\circ$ in λ . We did not attempt to characterize sub-exposure Moon-glow structure, but instead created a model of the additive offset to the WISE zero point due to Moon illumination. Our model is not a function of (λ, β) , but rather Moon angular separation (`moon_sep`) and Moon position angle relative to the WISE optics (termed $\phi_m = \text{moonpa} - \text{w3pa}$). To arrive at such a model, we generated a full-sky *W3* map, taking the L1b zero point at face

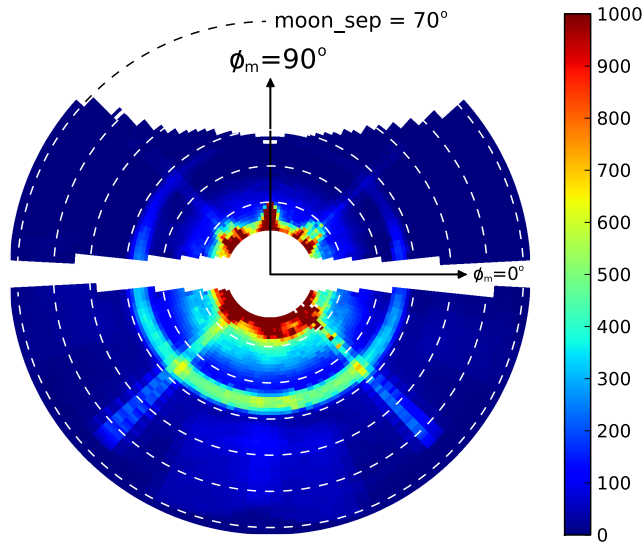


Figure 2.9 Our model for *W3* off-axis Moon-glow as a function of angular separation `moon_sep` (radial coordinate) and position angle ϕ_m (azimuthal coordinate). The units are *W3* DN, and dashed white circles denote intervals of 10° in `moon_sep`. Diffraction spikes are apparent out to `moon_sep` $> 60^\circ$. $0^\circ < \phi_m < 180^\circ$ tends to correspond to lower Moon phase than $-180^\circ < \phi_m < 0^\circ$, and is thus fainter.

value. We then subtracted $1.2\times$ the Kelsall model to roughly cancel the zodiacal light. Next, we smoothed to $6.1'$, and subtracted scaled SFD i100 assuming $i100/W3 = 17.5$, to approximately cancel diffuse Galactic emission. We designated particular, narrow ranges of λ as Moon-affected, and created a Moon-free prediction by interpolating over such regions at constant β . Restricting to $|b| > 15^\circ$, the difference between the Moon-free prediction and the corresponding L1b data constitutes our Moon-glow estimate, which yields the profile shown in Figure 2.9 when binned in $(\text{moon_sep}, \phi_m)$. We also created analogous profiles in Moon phase and Moon distance to enable corrections for these factors when predicting Moon contamination for individual L1b pointings.

As we stack, at each pixel in the co-add, we keep an integer count of the number of Moon-affected frames contributing to that pixel. We define Moon-affected frames to

be those for which our model predicts moon contamination of $S_m > 100$ DN, and $S_m/\mathcal{R}_{100} > 450$. Pixels with Moon-contamination count > 3 at the end of the stacking process are flagged with bit 15, as listed in Table 2.2.

In rare cases typically associated with moon contamination, all L1b quadrants at a particular sky location are rejected during the reference comparison step. Mask bit 19 labels such locations, at which we have tapered in the reference stack to avoid leaving regions of missing data in the final mosaic. 0.15% of the sky is affected.

Resolved Compact Sources

Though we have carefully removed pointlike sources, there are many resolved compact sources in the WISE All-Sky Catalog. These objects, which include galaxies and multiple sources, do not belong in our map of diffuse Galactic cirrus. Instead of attempting to model and subtract such sources, we have simply added mask bit 7 to flag their locations, allowing users to ignore or interpolate over the affected regions.

To obtain a catalog of resolved, compact $W3$ sources, we cross-matched WISE All-Sky Catalog objects with `w3snr` > 2 , `ext_flg` = 5 to 2MASS XSC sources, with a matching radius of $5''$. We excluded XSC objects with `cc_flag` = a or `cc_flag` = z, and additionally discarded the 3712 XSC sources believed to be fragments of diffuse Milky Way dust (Skrutskie et al., 2006, §2.3b). This resulted in a list of $\sim 150,000$ resolved $W3$ sources. In order to obtain the best possible set of attributes for each such object, including the semi-major axis to semi-minor axis ratio r_{ab} and position angle θ_{PA} , we cross-matched our resolved $W3$ source list to the PGC catalog (Paturel et al., 2003), with a matching radius of $1'$.

We proceed to merge attributes from the XSC and PGC catalogs. We assign PGC values for r_{ab} and θ_{PA} to all PGC-matched sources with both values available in the PGC catalog. For objects with PGC major axis $> 1'$, but lacking either r_{ab} or θ_{PA} in

the PGC catalog, r_{ab} and θ_{PA} are assigned to their XSC values. All other sources are assigned a default value of $r_{ab} = 1$. Relatively large sources with PGC major axis $>1'$ are assigned their respective PGC major axis value. All other sources are assigned a major axis of $1'$. To be conservative and avoid overly line-like masked regions, we cap r_{ab} at 2. The LMC, SMC, and M31 are excluded from the bit 7 mask and flagged separately (see §2.6.6).

For each WSSA tile, we then create a resolved object mask by flagging elliptical regions corresponding to the relevant entries from our WISE/XSC/PGC catalog. Over the entire sky, only 0.07% of pixels are thus flagged.

Our cross-matching procedure may occasionally discard true $W3$ sources in cases of exotic morphology (e.g. Barnard’s Galaxy), or in the event of a WISE All-Sky Catalog or XSC failure. To be conservative, we created mask bit 16, which uses the available PGC parameters of any prominent optical galaxies missing from our cross-match to again flag appropriate elliptical regions. To isolate such a sample of ‘prominent’ optical sources, we restricted to the subset of PGC sources with an RC3 counterpart (Corwin et al., 1994). Before generating this mask, we carefully excluded 18 large ($>5'$ diameter), dIrr or dSph galaxies which are completely invisible in $W3$ (e.g. the Fornax Dwarf). In all, bit 16 masks only 0.005% of the sky.

2.6.6 Low Integer Frame Coverage

In regions with low integer frame coverage <5 , compact sources are listed only in the WISE All-sky Reject Table, and are excluded from the All-sky Source Catalog (Cutri et al., 2012). Rather than attempt to identify valid sources among Reject Table entries, we opted instead to flag such regions with mask bit 21. 1.2% of the sky is affected.

Big Objects

The SFD mask contains a so-called “big object” bit, which labels regions affected by extragalactic sources with exceptionally large apparent sizes: the LMC, SMC, and M31. Bit 17 of our masks is meant to flag these same three objects. Our big object mask is identical to that of SFD for the irregularly shaped LMC and SMC. For M31, we use an ellipse based on PGC catalog parameters.

Low Ecliptic Latitude

Unfortunately, certain low ecliptic latitude regions were only observed by WISE while obstructed from view by a planet within our Solar System. In such cases, the planet appears as an extraordinarily bright, compact source and renders measurement of diffuse Galactic emission in its vicinity completely hopeless. As a result, we have flagged circular regions about the centroid of each planet imprint, typically several degrees in diameter. We have also added mask bit 22 to flag all pixels immediately within the ecliptic plane, as these should generally be handled with extra caution. The approximate range of ecliptic latitudes flagged is $-2^\circ < \beta < 2^\circ$. Additionally, at low ecliptic latitude, we have noticed unexpected line-like features $\sim 5.5'$ wide and extending many degrees in length. The origin of these lines is not known, and they do not appear to follow trajectories obviously indicative of an instrumental or processing malfunction, e.g. constant δ or constant β . Mask bit 20 marks the locations of such lines, which affect less than 0.15% of the sky.

2.7 Final WSSA Tiles

In the following subsections we outline the final steps taken to convert our WSSA co-adds of §2.6 into the data products we have released.




Figure 2.10 $4.3^\circ \times 3.8^\circ$ sub-region of WSSA tile 51, at $(l, b) \approx (266^\circ, -6^\circ)$, illustrating three of the extensions included in each of our tiles. Top: primary extension, diffuse Galactic *W3* emission. Middle: ‘dirty’ extension, including point sources, satellite streaks, and SSOs. Bottom: colorscale rendering of ‘OR mask’ extension. Red crosses denote point source diffraction spikes.

2.7.1 Zero Point

Our mosaicking process only attempts to create a smooth image over each tile footprint by best matching each L1b exposure with its neighbors. However, this process does not ensure that neighboring tiles agree on their overlaps. Further, relics of the zodiacal light and Moon contamination that aren't perfectly removed by our SFD-based gradient correction can persist, leading to wandering of the tile zero points on relatively large scales of several degrees.

We investigated the possibility of obtaining a zero point by taking the L1b zero point at face value, and then subtracting out models for the time-dependent foregrounds. Despite our best efforts to subtract the Moon-glow model of §2.6.5 and various customizations of the Kelsall zodiacal light model (Kelsall et al., 1998), such an approach ultimately yielded a cosmetically unappealing result. Based on the Planck comparison we describe below, we estimate that, at high $|b|$, the true $W3$ Galactic emission is $\sim 0.5\%$ as bright as the $|\beta| = 90^\circ$ zodiacal light. While the Kelsall model appeared reasonable at the several percent level, this remained inadequate due to the overwhelming brightness of the zodiacal light at $12\mu\text{m}$.

We therefore resorted a zero point based on Planck (Planck Collaboration et al., 2013a). Henceforth, the term Planck will refer to Planck 857 GHz ($350\mu\text{m}$), corrected for zodiacal emission according to Planck Collaboration et al. (2013h). To ensure this Planck map contained only diffuse Galactic emission, we also interpolated over the SFD compact source mask. The great advantage of Planck relative to other full-sky, mid-IR and far-IR datasets is the very faint zodiacal foreground. The basic strategy will be to replace modes on scales of several degrees in our WISE tiles with appropriately rescaled and smoothed Planck.

In order to accomplish this goal, we must derive a map of the spatially varying

coefficient f_{pw} that converts $350\mu\text{m}$ emission to $12\mu\text{m}$ emission. To derive f_{pw} we correlate cleaned WISE with Planck on a bandpass between $4'$ and $15'$. We break each tile into many $0.5^\circ \times 0.5^\circ$ rectangular patches and perform a linear least squares fit to the bandpass filtered WISE sub-image, where the model is bandpass Planck multiplied by f_{pw} , plus a constant offset. For each such fit, we also compute the Pearson correlation coefficient c of bandpass WISE versus bandpass Planck. We perform such fits over the full sky. For a variety of reasons, including poor signal to noise in Planck or WISE, or systematic issues in in Planck or WISE, c is often closer to zero than to unity. In such cases, the fit is driven to $f_{pw} \rightarrow 0$ for reasons that are instrumental rather than astrophysical. In order to obtain a map of f_{pw} on large scales, we make a cut on $c > 0.7$ and smooth the map of f_{pw} at those locations to 15° . With WISE in $W3$ DN and Planck in MJy/sr f_{pw} , varies between 2.2 and 5.8. Converting WISE to MJy/sr according to Cutri et al. (2012) so that both data sets have the same units, this translates to $0.036 < f_{pw} < 0.095$.

To replace the large scale modes of WISE with those of Planck, we apply a $12.5'$ median filter to both data sets, and then smooth both to 2° . We inpainted over regions with strong molecular emission, defined as Planck Type 3 CO >15 K km/s (Planck Collaboration et al., 2013g). We subtract 2° WISE from each tile, and then add in smoothed Planck, multiplied by f_{pw} . In order to derive tiles free of edge effects from the smoothing process, we initially built our WSSA stacks on $14^\circ \times 14^\circ$ padded astrometric footprints.

Lastly, we convert the WISE tiles with Planck-based zero point from $W3$ DN to MJy/sr according to the prescription of Cutri et al. (2012), §4.4h.

Table 2.3. WSSA Tile Extensions

Extension	Description
0	cleaned co-add
1	dirty co-add
2	integer coverage
3	minimum value image
4	maximum value image
5	AND bit-mask
6	OR bit-mask
7	transient artifact image

2.7.2 Tile Format

We release the 430 *W3* WSSA tiles as multi-extension FITS files. Table 3 lists the extensions included. Refer to the tile headers for further details.

2.8 Discussion and Conclusion

We have described and released³ a full-sky, 15'' resolution map of diffuse Galactic 12 μ m emission based on WISE images. Indeed, WISE *W3* provides valuable and striking small angular scale detail regarding the spatial structure of the ISM over much of the sky. We strongly encourage all interested astronomers to sample values from our mosaics, and also visually inspect the co-adds. We anticipate this to be the first in a series of papers, throughout which we demonstrate that an enhanced-resolution dust map based on WISE *W3* better predicts extinction than does SFD, and ultimately release such a full-sky data product for accessing improved reddening estimates.

Figure 2.11 shows the 12 μ m map we have derived, binned to 7' resolution. On large

³See <http://wise.skymaps.info>.




Figure 2.11 Our full-sky, $12\mu\text{m}$ map, binned to $7'$ resolution, in Hammer-Aitoff projection, on a logarithmic stretch.

scales, the map looks very similar to the Planck $350\mu\text{m}$ data from which we derived our $12\mu\text{m}$ zero-point. The notable exceptions are stripes of moon contamination, which dominate our map over $\sim 5\%$ of the sky.

We conclude by briefly pointing out aspects of the WISE All-sky release processing that turned out to be suboptimal for the purpose of mapping diffuse emission, and finally compare our $12\mu\text{m}$ map to those of DIRBE and *IRAS*.

2.8.1 All-sky Release Processing

The “dynamic calibration sky offset correction” (Cutri et al., 2012, §4.4a) introduced several problematic features into the L1b frames which we did not address for this initial $12\mu\text{m}$ tile release. In particular, long-term latents due to bright sources are mirrored in acausal fashion to affect frames temporally *preceding* bright source detection. Examples of such features can be seen in Figure 2.8 as low-level, dark ‘divots’ $\sim 1\text{-}2'$ in size. These acausal latents are the dominant remaining imperfection in many regions of our co-adds with low dust signal. In a related problem, the dynamic sky offset procedure can often assign up to $\sim 4\%$ of pixels in a given L1b image to exactly the median pixel value within the frame. Because these issues would be better addressed with lower-level data products than the L1b images, we have so far opted not to address them.

2.8.2 DIRBE & IRAS Comparison

The relevant DIRBE data product to compare with our *W3* map is the band 5 “Zodi-Subtracted Mission Average” (ZSMA5). Because DIRBE pointed at a variety of Solar elongations, including some $< 90^\circ$, the DIRBE-observed zodiacal light is even more pronounced and complex than that in *W3*. Arguably, the most prominent feature in ZSMA5 is the *residual* zodiacal light left over after subtracting the DSZA

model. The ecliptic plane is obvious in such residuals, and at high $|b|$ there remains a clear gradient of average flux with β , with magnitude of $\sim 6 \times 10^{-3}$ (MJy/sr)/deg, or ~ 0.2 MJy/sr between $|\beta| \sim 60^\circ$ and $|\beta| \sim 90^\circ$. By comparison, our Planck-based far-IR zero point suggests that at $|b| \gtrsim 80^\circ$, the *entire* non-IPD $12\mu\text{m}$ signal is ~ 0.03 MJy/sr. There is also a significant offset between ZSMA5 and our map in regions near the ecliptic poles, with amplitude ~ 0.6 MJy/sr. While it might be interesting to investigate these differences, they are immaterial to our long-term goal of predicting dust extinction. Thanks to our far-IR zero point, our *W3* map only shows traces of zodiacal light in rare circumstances, very near the ecliptic plane, where the zodi can vary strongly on scales smaller than 2° . Our *W3* map is also superior to ZSMA5 in that point sources remain unsubtracted in the latter, impacting a substantial fraction of the sky. On the other hand, DIRBE band 5 is truly a full-sky map, whereas $\sim 5\%$ of diffuse *W3* emission is badly contaminated by Moon-glow and/or Solar System planets.

IRAS $12\mu\text{m}$ has a separate “reject” data set at low $|\beta|$ that is not consistent with the remainder of their map. Thus, our *W3* map improves upon *IRAS* $12\mu\text{m}$ by virtue of (1) its consistent, zodi-free zero point (2) its factor of 16 enhancement in angular resolution and (3) our removal and masking of compact sources.

We thank Roc Cutri for his guidance in making use of the WISE data products. We also thank Bruce Draine for his careful reading of an early version of this manuscript.

Chapter 3

Planck Thermal Dust Emission Model

We apply the Finkbeiner et al. (1999) two-component thermal dust emission model to the *Planck* HFI maps. This parametrization of the far-infrared dust spectrum as the sum of two modified blackbodies serves as an important alternative to the commonly adopted single modified blackbody (MBB) dust emission model. Analyzing the joint *Planck*/DIRBE dust spectrum, we show that two-component models provide a better fit to the 100-3000 GHz emission than do single-MBB models, though by a lesser margin than found by Finkbeiner et al. (1999) based on FIRAS and DIRBE. We also derive full-sky 6.1' resolution maps of dust optical depth and temperature by fitting the two-component model to *Planck* 217-857 GHz along with DIRBE/*IRAS* 100 μ m data. Because our two-component model matches the dust spectrum near its peak, accounts for the spectrum's flattening at millimeter wavelengths, and specifies dust temperature at 6.1' FWHM, our model provides reliable, high-resolution thermal dust emission foreground predictions from 100 to 3000 GHz. We find that, in diffuse sky regions, our two-component 100-217 GHz predictions are on average accurate to within 2.2%, while extrapolating the Planck Collaboration et al. (2013k) single-MBB model systematically underpredicts emission by 18.8% at 100 GHz, 12.6% at 143 GHz and 7.9% at 217 GHz. We calibrate our two-component optical depth to reddening,

and compare with reddening estimates based on stellar spectra. We find the dominant systematic problems in our temperature/reddening maps to be zodiacal light on large angular scales and the cosmic infrared background anisotropy on small angular scales.

3.1 Introduction

The presence of Galactic interstellar dust affects astronomical observations over a wide range of wavelengths. In the mid-infrared and far-infrared, Galactic dust emission contributes significantly to the total observed sky intensity. At optical and ultraviolet (UV) wavelengths, dust grains absorb and scatter starlight. Observations of interstellar dust emission/absorption can improve our understanding of the physical conditions and composition of the interstellar medium (ISM), an environment which plays a crucial role in Galactic evolution and star formation. Equally, or perhaps even more important to the practice of astronomy, however, is accurately accounting for dust as a foreground which reddens optical/UV observations of stars/galaxies and superimposes Galactic emission on low-frequency observations of the cosmic microwave background (CMB).

Over the past decades, satellite observations have dramatically enhanced our knowledge about infrared emission from the ISM. The *Infrared Astronomy Satellite* (*IRAS*), with its $\sim 4'$ resolution, revolutionized the study of Galactic dust emission, first revealing the high-latitude “infrared cirrus” using $60\mu\text{m}$ and $100\mu\text{m}$ observations (Low et al., 1984; Wheelock et al., 1994) and highlighting the importance of detailed dust mapping in the far-infrared/submillimeter as a key foreground for cosmology. Later, the Diffuse Infrared Background Experiment (DIRBE) aboard the *COBE* satellite provided complementary full-sky measurements at ten infrared wavelengths from $1.25\mu\text{m}$ to $240\mu\text{m}$, boasting a reliable zero point despite inferior

$\sim 0.7^\circ$ angular resolution (Boggess et al., 1992). *COBE*/FIRAS (Mather, 1982) also provided full-sky infrared dust spectra at 7° resolution in 213 narrow frequency bins between 30 GHz and 2850 GHz.

Finkbeiner et al. (1999, hereafter FDS99) used these FIRAS data to derive a globally best-fit model of dust emission applicable over a very broad range of frequencies. FDS99 showed that no model consisting of a single modified blackbody (MBB) could accurately match the FIRAS/DIRBE spectrum at both the Wien and Rayleigh-Jeans extremes. To fit the thermal dust spectrum between 100 and 3000 GHz, FDS99 therefore proposed an emission model consisting of two MBBs, each with a different temperature and emissivity power law index. Physically, these two components might represent distinct dust grain species within the ISM, or they might simply provide a convenient fitting function. By combining this best-fit two-component model with a custom reprocessing of DIRBE and *IRAS* $100\mu\text{m}$ data, FDS99 provided widely used foreground predictions with $6.1'$ FWHM, limited largely by their 1.3° resolution DIRBE-based temperature correction.

The *Planck* 2013 data release (Planck Collaboration et al., 2013b) represents an important opportunity to revisit foreground predictions in light of *Planck*'s superb, relatively artifact-free broadband data covering the entire sky and a wide range of frequencies. Towards this end, Planck Collaboration et al. (2013k) has conducted a study modeling *Planck* 353 GHz, 545 GHz, 857 GHz and DIRBE/*IRAS* $100\mu\text{m}$ emission with a single-MBB spectrum. More recently, Planck Collaboration et al. (2014b) has applied the Draine & Li (2007) dust grain model to *Planck*, *IRAS*, and *WISE* emission between 353 GHz and $12\mu\text{m}$. Here we investigate the FDS99 two-component dust emission model as an alternative parametrization for the 100-3000 GHz dust spectral energy distribution (SED) composed of *Planck* High Frequency Instrument (HFI), DIRBE and *IRAS* data. In doing so, we obtain

Planck-based maps of dust temperature and optical depth, both at $6.1'$ resolution. Because we employ a model that has been validated with FIRAS down to millimeter wavelengths and optimized for *Planck*, our derived parameters are useful in constructing high-resolution predictions of dust emission over a very broad range of wavelengths. This includes low frequencies (100-350 GHz), which Planck Collaboration et al. (2013k) caution their model may not adequately fit, and also wavelengths near the peak of the dust SED, relevant to e.g. *AKARI* 140-160 μm (Doi et al., 2012). We also anticipate our derived optical depth map will serve as a valuable cross-check for extinction estimates based directly upon optical observations of stars (e.g. Schlafly et al., 2014) and as a baseline for next-generation dust extinction maps incorporating high-resolution, full-sky infrared data sets such as *WISE* (Wright et al., 2010; Meisner & Finkbeiner, 2014) and *AKARI*.

In §3.2 we introduce the data used throughout this study. In §3.3 we describe our preprocessing of the *Planck* maps to isolate thermal emission from Galactic dust. In §3.4 we explain the two-component emission model we apply to the *Planck*-based dust SED. In §3.5, we discuss the details of predicting *Planck* observations based on this dust model. In §3.6 we derive constraints on our model's global parameters in light of the *Planck* HFI maps. In §3.7 we detail the Markov chain Monte Carlo (MCMC) method with which we have estimated the spatially varying parameters of our model. In §3.8 we calibrate our derived optical depth to reddening at optical wavelengths. In §3.9 we compare our two-component thermal dust emission predictions to those of Planck Collaboration et al. (2013k). In §3.10 we present the full-sky maps of dust temperature and optical depth we have obtained, and conclude in §3.11.

3.2 Data

All *Planck* data products utilized throughout this work are drawn from the *Planck* 2013 release (Planck Collaboration et al., 2013b). Specifically, we have made use of all six of the zodiacal light corrected HFI intensity maps (`R1.10_nominal_ZodiCorrected`, Planck Collaboration et al., 2013i). Our full-resolution (6.1' FWHM) SED fits neglect the two lowest HFI frequencies, 100 and 143 GHz, as these have FWHM of 9.66' and 7.27' respectively.

To incorporate measurements on the Wien side of the dust emission spectrum, we include 100 μm data in our SED fits. In particular, we use the Schlegel et al. (1998, henceforth SFD) reprocessing of DIRBE/*IRAS* 100 μm , which we will refer to as `i100`, and at times by frequency as 3000 GHz. The `i100` map has angular resolution of 6.1', and was constructed so as to contain only thermal emission from Galactic dust, with compact sources and zodiacal light removed, and its zero level tied to HI. We use the `i100` map as is, without any custom modifications.

In some of our FIR dust SED analyses which do not require high angular resolution, specifically those of §3.6, §3.7.4, and §3.7.5, we also make use of the SFD reprocessings of DIRBE 140 μm (2141 GHz) and 240 μm (1250 GHz).

3.3 Preprocessing

The following subsections detail the processing steps we have applied to isolate Galactic dust emission in the *Planck* maps in preparation for SED fitting.

3.3.1 CMB Anisotropy Removal

We first addressed the CMB anisotropies before performing any of the interpolation/smoothing described in §3.3.2/§3.3.3. The CMB anisotropies are

effectively imperceptible upon visual inspection of *Planck* 857 GHz, but can be perceived at a low level in *Planck* 545 GHz, and are prominent at 100-353 GHz relative to the Galactic emission we wish to characterize, especially at high latitudes. To remove the CMB anisotropies, we have subtracted the Spectral Matching Independent Component Analysis (SMICA, Planck Collaboration et al., 2013f) model from each of the *Planck* maps, applying appropriate unit conversions for the 545 and 857 GHz maps with native units of MJy/sr. Low-order corrections, particularly our removal of Solar dipole residuals, are discussed in §3.3.5.

3.3.2 Compact Sources

After subtracting the SMICA CMB model, we interpolate over compact sources, including both point sources and resolved galaxies. Removing compact sources at this stage is important as it prevents contamination of compact-source-free pixels in our downstream analyses which require smoothing of the *Planck* maps. SFD carefully removed point sources and galaxies from the i100 map everywhere outside of $|b| < 5^\circ$. We do not perform any further modifications of the i100 map to account for compact sources. To mask compact sources in the *Planck* 217-857 GHz maps, we use the SFD compact source mask. At 100, 143 GHz we use the compact source masks provided by the *Planck* collaboration in the file `HFI_Mask_PointSrc_2048_R1.10.fits`. Given our pixelization (see §3.7.1), 1.56% of pixels are masked at 217-857 GHz (1.05%, 1.02% at 100, 143 GHz).

3.3.3 Smoothing

For our full-resolution model, we wish to simultaneously fit i100 along with the four highest-frequency *Planck* bands. To properly combine these maps, they must have the same point spread function (PSF). i100, with its 6.1' symmetric Gaussian beam,

has the lowest angular resolution of the relevant maps. To match PSFs, we have therefore smoothed each of the *Planck* maps under consideration to $l=100$ resolution by considering each native *Planck* map to have a symmetric Gaussian beam and smoothing by the appropriate symmetric Gaussian such that the resulting map has a $6.1'$ FWHM. The FWHM values we assign to the native *Planck* maps are taken from Planck Collaboration et al. (2013d), and are listed in Table 3.1.

3.3.4 Molecular Emission

Because the FIRAS spectra consist of many narrow frequency bins, FDS99 were able to discard the relatively small number of frequency intervals contaminated by strong molecular line emission. Unfortunately, while the *Planck* data considered in this study are of high angular resolution, the broad *Planck* bandpasses do not allow us to adopt the same approach as FDS99 in dealing with line emission. Instead, we must subtract estimates of the molecular line contamination from each *Planck* band in order to best isolate the thermal continuum we wish to characterize. The most prominent molecular line emission in the *Planck* bands of interest arises from the three lowest CO rotational transitions: $J=1\rightarrow 0$ at 115 GHz, $J=2\rightarrow 1$ at 230 GHz and $J=3\rightarrow 2$ at 345 GHz, respectively affecting the *Planck* 100, 217 and 353 GHz bands. The $J=1\rightarrow 0$ line also imparts a signal upon *Planck* 143 GHz, but at a negligible level, $\sim 1000\times$ fainter relative to the dust continuum than $J=1\rightarrow 0$ at 100 GHz. More specifically, the ratio of $J=1\rightarrow 0$ intensity to thermal dust emission in *Planck* 143 GHz is ≥ 0.001 for only $< 2\%$ of the sky.

To correct for molecular emission, we employ the *Planck* Type 3 CO data product, which boasts the highest S/N among the available full-sky CO maps based on the *Planck* HFI and Low Frequency Instrument (LFI) data (Planck Collaboration et al., 2013g). The native angular resolution of the Type 3 CO map is $5.5'$. We therefore

begin by smoothing the raw Type 3 CO map to match the PSF of the smoothed *Planck* intensity maps we wish to correct for molecular emission.

We must apply the appropriate unit conversions to the Type 3 CO map before subtracting it from the *Planck* intensity maps, which have native units of K_{CMB} at the frequencies of interest. The Type 3 CO map is provided in units of K_{RJ} km/s of $J=1\rightarrow 0$ emission. To convert this quantity to K_{CMB} , we assume that all of the CO emission arises from the ^{12}CO isotope, and derive the *Planck*-observed CO intensity in units of K_{CMB} as follows:

$$I_{CO,\nu_i,N,N-1} = I_3 F_{12CO,\nu_i,N,N-1} R_{N,N-1} \quad (3.1)$$

Where $I_{CO,\nu_i,N,N-1}$ is the intensity in K_{CMB} in *Planck* band ν_i due to the CO transition from $J=N$ to $J=(N-1)$. I_3 represents the appropriately smoothed Type 3 CO amplitude in K_{RJ} km/s of $J=1\rightarrow 0$ emission. The $F_{12CO,\nu_i,N,N-1}$ are conversion factors between K_{RJ} km/s and K_{CMB} for particular band/transition pairs. The relevant values, calculated with the *Unit Conversion and Colour Correction* software utilities (v1.2), are:

$$F_{12CO,100,1,0} = 1.478 \times 10^{-5} K_{CMB} / (K_{RJ} \text{ km/s}),$$

$$F_{12CO,217,2,1} = 4.585 \times 10^{-5} K_{CMB} / (K_{RJ} \text{ km/s}), \text{ and}$$

$$F_{12CO,353,3,2} = 1.751 \times 10^{-4} K_{CMB} / (K_{RJ} \text{ km/s}).$$

$R_{N,N-1}$ represents the line ratio of the transition from $J=N$ to $J=(N-1)$ relative to the $J=1\rightarrow 0$. Thus, $R_{1,0}=1$, and we further adopt $R_{2,1}=0.595$ and $R_{3,2}=0.297$ based on Planck Collaboration et al. (2013g). These line ratios are assumed to be constant over the entire sky.

Formally, then, the CO contamination in band ν_i is given by:

$$I_{CO,\nu_i} = \sum_N I_{CO,\nu_i,N,N-1} \quad (3.2)$$

It happens that, for each of the *Planck* bands in which CO emission is non-negligible (100, 217 and 353 GHz), only a single N contributes ($N=1$, $N=2$ and $N=3$, respectively).

Unfortunately, the Type 3 CO map at 6.1' FWHM is rather noisy, and the vast majority of the sky has completely negligible CO emission. Thus, in order to avoid adding unnecessary noise outside of molecular cloud complexes and at high latitudes, we have zeroed out low-signal regions of the Type 3 CO map. We identify low-signal regions as those with $\mathcal{I}_3 < 1 \text{ K}_{RJ} \text{ km/s}$, where \mathcal{I}_3 is the Type 3 CO map smoothed to 0.25° FWHM. As a result of this cut, 90% of the sky remains unaffected by our CO correction, particularly the vast majority of the high Galactic latitude sky.

3.3.5 Zero Level

Although we wish to isolate and model thermal emission from Galactic dust, the *Planck* maps contain additional components on large angular scales. At each frequency, there can exist an overall, constant offset that must be subtracted to set the zero level of Galactic dust by removing the mean cosmic IR background (CIB, Hauser & Dwek, 2001), as well as any instrumental offset. Additionally, faint residuals of the Solar dipole remain at low frequencies. We will address these issues by separately solving two sub-problems: first, we set the absolute zero level of *Planck* 857 GHz relative to external data, and second we fit the 100-545 GHz offsets and low order corrections by correlating these *Planck* bands against *Planck* 857 GHz.

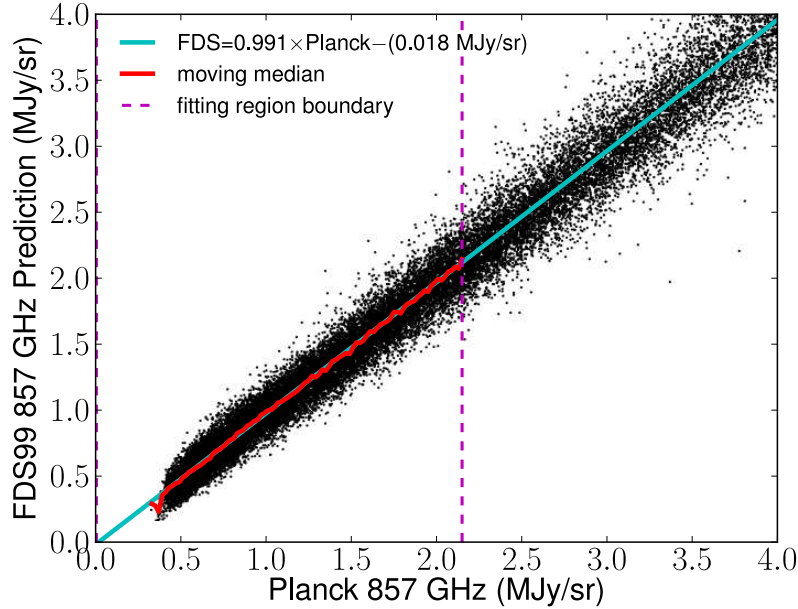


Figure 3.1 Scatter plot of FDS99-predicted 857 GHz thermal dust emission versus *Planck* 857 GHz observations, illustrating our absolute zero level determination described in §3.3.5.

Absolute Zero Level

In Planck Collaboration et al. (2013k), the absolute zero level of thermal dust emission was set by requiring that *Planck* infrared emission tends to zero when HI is zero, assuming a linear correlation between these two measurements at low column density. However, this approach is less than completely satisfying in that there appear to be different slopes of *Planck* 857 GHz versus HI for different ranges of HI intensity. In particular, *Planck* 857 GHz appears to “flatten out” at very low HI, as shown in Figure 5 of Planck Collaboration et al. (2013k). More quantitatively, we have found using the LAB HI data (Kalberla et al., 2005) for $-72 < v_{LSR} < +25$ km/s that the best-fit slope for $HI < 70$ K km/s is a factor of ~ 1.9 lower than the best fit slope for $110 \text{ K km/s} < HI < 200 \text{ K km/s}$, and as a result the implied zero level offsets for *Planck* 857 GHz differ by ~ 0.37 MJy/sr.




Figure 3.2 Scatter plots of *Planck* 100, 143, 217, 353, and 545 GHz versus *Planck* 857 GHz. Left: before applying our best-fit zero level offsets and additional low-order corrections. Right, top four panels: *Planck* 143-545 GHz after correcting for each band's best-fit offset and residual Solar dipole. Bottom right: *Planck* 100 GHz after applying the spherical harmonic corrections of Equation 3.4. The dashed red line shows the best-fit linear relationship in all cases.

Because of this ambiguity in the relationship between 857 GHz and HI emission, we decided to instead constrain the *Planck* 857 GHz zero level by comparison to the FDS99-predicted 857 GHz thermal dust emission. This renders our *Planck* 857 GHz absolute zero level tied indirectly to HI through the FDS99 100 μ m and 240 μ m zero levels.

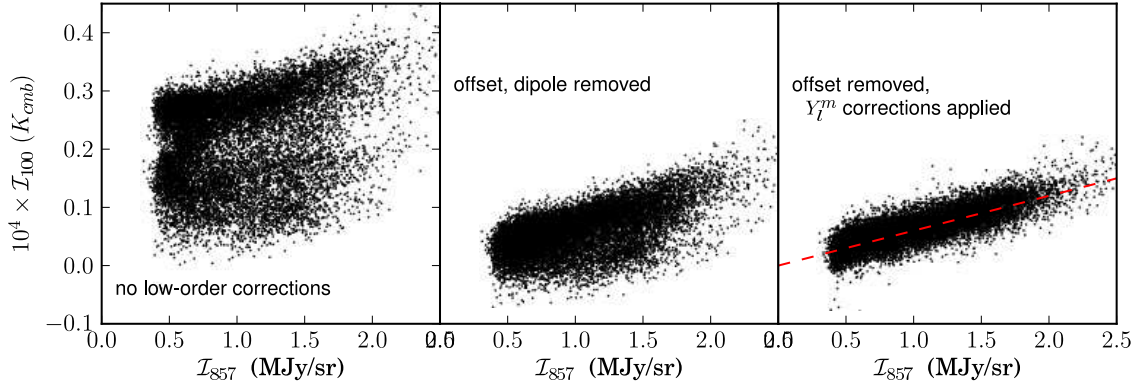


Figure 3.3 Summary of low-order corrections at 100 GHz. Left: prior to our low-order corrections, a $\sim 17\mu\text{K}$ zero level offset is present and strong low-order problems reduce the linearity of the 100 GHz trend versus 857 GHz. Center: scatter plot of *Planck* 100 GHz versus 857 GHz after applying the best-fit offset and residual Solar dipole corrections derived with Equation 3.3 to *Planck* 100 GHz. The correlation is strengthened, but remains far less tight than for 143-545 GHz (see right column of Figure 3.2, top four rows). Right: after applying the spherical harmonic corrections of Equation 3.4 to *Planck* 100 GHz, the correlation versus 857 GHz is far more tightly linear than following the dipole correction.

We perform a linear fit to the FDS99-predicted 857 GHz values as a function of *Planck* 857 GHz. For this purpose, we employ a version of the *Planck* 857 GHz map with zodiacal light and point sources removed and smoothed to 1° FWHM, which we will refer to as \mathcal{I}_{857} . We consider \mathcal{I}_{857} to be the independent variable, as it has much higher S/N than the FDS99 prediction, henceforward referred to as \mathcal{F}_{857} . Note that \mathcal{F}_{857} is not simply the FDS99 model evaluated at 857 GHz, but also incorporates the color correction factor of §3.5, using the FDS99 temperature map to determine the

dust spectrum shape. We rebin to $N_{side}=64$ and restrict to pixels with $\mathcal{I}_{857}<2.15$ MJy/sr. Since *Planck* 857 GHz smoothed to degree resolution has very high S/N, we can safely perform such a cut on \mathcal{I}_{857} . Figure 3.1 shows a scatter plot of \mathcal{I}_{857} versus \mathcal{F}_{857} , with a moving median and linear fit overplotted. The linear fit was performed with uniform weights and iterative outlier rejection. The best-fit linear model is given by $\mathcal{F}_{857}=0.991\mathcal{I}_{857}-0.018$ MJy/sr. It is encouraging that the slope is quite close to unity. It is also encouraging that our choice of *Planck* 857 GHz threshold at 2.15 MJy/sr is unimportant; any threshold value between 1.3 MJy/sr (28th percentile in \mathcal{I}_{857}) and 3.9 MJy/sr (61st percentile in \mathcal{I}_{857}) yields a zero level offset within 0.01 MJy/sr of our adopted value.

The formal statistical error on the best-fit 857 GHz offset is quite small, ~ 0.002 MJy/sr. The systematics likely to dominate the actual uncertainty on our FDS-based zero level are imperfections in the *Planck*/i100 zodiacal light models and the FDS99 temperature map. To quantify these systematic uncertainties, we split the sky into four quadrants, with boundaries at $b=0^\circ$ and $l=0^\circ, l=180^\circ$. We again restricted to $\mathcal{I}_{857}<2.15$ MJy/sr, and repeated the regression in each quadrant. The rms of the per-quadrant slopes was found to be 0.0188, while the rms of the per-quadrant offsets was 0.0586 MJy/sr. Our adopted ~ 0.06 MJy/sr zero level uncertainty is sufficiently large to be consistent with the possible error introduced by assuming no appreciable Solar dipole signal in the *Planck* 857 GHz map. If we allow for a dipole template in our FDS99 versus *Planck* linear regression at 857 GHz, the best-fit dipole amplitude is only 0.02 MJy/sr.

Relative Zero Level

In the course of this study we use not only *Planck* 857 GHz, but also all of the remaining *Planck* HFI bands, as well as i100. To derive the zero level offsets that

Table 3.1. Input Map Properties & Pre-processing

ν (GHz)	Instrument(s)	Offset (K_{CMB})	Dipole (K_{CMB})	$ss57,\nu \times u_\nu$	$\sigma_{ss57,\nu} \times u_\nu$	n_ν (K_{CMB})	c_ν	FWHM (')
100	<i>Planck</i> HFI	$1.69 \times 10^{-5} \pm 3.61 \times 10^{-7}$	-1.08×10^{-5}	1.46×10^{-3}	2.92×10^{-5}	7.77×10^{-5}	0.0054	9.66
143	<i>Planck</i> HFI	$3.58 \times 10^{-5} \pm 7.58 \times 10^{-7}$	-1.08×10^{-5}	4.68×10^{-3}	9.37×10^{-5}	3.25×10^{-5}	0.0054	7.27
217	<i>Planck</i> HFI	$7.79 \times 10^{-5} \pm 2.60 \times 10^{-6}$	-1.40×10^{-5}	2.09×10^{-2}	4.19×10^{-4}	4.51×10^{-5}	0.0054	5.01
353	<i>Planck</i> HFI	$2.76 \times 10^{-4} \pm 1.95 \times 10^{-5}$	-3.08×10^{-5}	9.32×10^{-2}	1.86×10^{-3}	1.51×10^{-4}	0.012	4.86
		Offset (MJy/sr)	Dipole (MJy/sr)	$ss57,\nu$	$\sigma_{ss57,\nu}$	n_ν (MJy/sr)		
545	<i>Planck</i> HFI	$7.27 \times 10^{-2} \pm 1.99 \times 10^{-2}$	1.63×10^{-2}	3.31×10^{-1}	6.62×10^{-3}	0.046	0.10	4.84
857	<i>Planck</i> HFI	$1.82 \times 10^{-2} \pm 6.02 \times 10^{-2}$	-	1.0	2.0×10^{-2}	0.046	0.10	4.63
1250	DIRBE	$7.06 \times 10^{-2} \pm 1.19 \times 10^{-1}$	-	1.98	3.97×10^{-2}	0.42	0.10	42
2141	DIRBE	$1.04 \times 10^{-1} \pm 1.54 \times 10^{-1}$	-	2.56	5.12×10^{-2}	0.79	0.10	42
3000	DIRBE/ <i>IRAS</i>	$0.0 \pm 4.3 \times 10^{-2}$	-	1.27	2.53×10^{-2}	0.06	0.10	6.1

Note. — Column 1: Approximate band center frequency of each input map. Note that 1250 GHz and 2141 GHz refer to the SFD98 reprocessings of DIRBE 240 μ m and 140 μ m respectively. Column 2: Instrument(s) from which the input map at each frequency has been obtained. Column 3: Zero level offset subtracted from each raw input map. Column 4: Best-fit residual Solar dipole amplitude according to Equation 3.3. Column 5: Dimensionless correlation slope of each map relative to *Planck* 857 GHz. These are the correlation slopes used in the analysis of §3.6, specifically Equation 3.10. Column 6: Adopted uncertainty on the dimensionless correlation slopes relative to *Planck* 857 GHz, for use in Equation 3.10. Column 7: n_ν represents the adopted per-pixel statistical noise level at full resolution, which contributes to the error budget of Equation 3.14. Column 8: Multiplicative fractional uncertainty on each input map, for use in the error budget of Equation 3.14. Column 9: Native angular resolution of each input map.

must be applied to each of the five lowest-frequency *Planck* bands, we perform a regression versus the *Planck* 857 GHz map corrected for the best-fit absolute zero level offset from §3.3.5. We assume no offset need be applied to i100, which already has its zero level tied to HI by SFD.

The need for additional low-order corrections beyond simple scalar offsets became evident upon inspecting the HFI maps at 100-545 GHz. In particular, we noticed the presence of a low-level dipole pattern, with an orientation consistent with that of the Solar dipole. Our strategy will be to simultaneously fit both this residual dipole and the zero-level offset amplitude for each band. To most precisely recover these amplitudes, it is necessary to have the highest available S/N in the independent variable of our regression. For this reason we have used *Planck* 857 GHz as a reference for the 100-545 GHz bands, as opposed to the FDS99 predictions or HI data. In doing so, we assume *Planck* 857 GHz contains no appreciable Solar dipole residual.

We perform one regression per HFI band (other than 857 GHz) to simultaneously fit for the zero level offset, the slope relative to 857 GHz, and the residual dipole amplitude. For each 100-545 GHz HFI band, we restrict to regions of low column density ($\text{HI} < 200 \text{ K km s}^{-1}$ for $-72 < v_{LSR} < +25 \text{ km s}^{-1}$) and fit the following model:

$$\mathcal{I}_{\nu_i,p} = m\mathcal{I}_{857,p} + b + d\mathcal{D}_p \quad (3.3)$$

With p denoting a single $N_{side}=64$ HEALPix pixel (Górski et al., 2005b) in the maps \mathcal{I}_{857} , \mathcal{I}_{ν_i} , and \mathcal{D} . Here \mathcal{I}_{857} is the *Planck* 857 GHz map with zodiacal emission, compact sources, and the constant offset of §3.3.5 removed, smoothed to 1° resolution. \mathcal{I}_{ν_i} is the corresponding 1° resolution *Planck* HFI map with zodiacal emission, CMB anisotropies, and compact sources removed. In the context of Equation 3.3, $\nu_i \in \{100, 143, 217, 353, 545\}$ GHz. Note that \mathcal{I}_{ν_i} is always in the native units of the

relevant *Planck* band. \mathcal{D} is a scaling of the Solar dipole pattern oriented toward $(l, b) = (263.99^\circ, 48.26^\circ)$, with unit amplitude. Because $\sim 18,000$ pixels satisfy the low HI cut, we have an overconstrained linear model with three parameters: m , d , and b . m represents the best-fit slope of *Planck* band ν_i versus *Planck* 857 GHz assuming they are linearly related. d is the residual Solar dipole amplitude, and its best-fit value represents the scaling of the Solar dipole that makes the *Planck* band ν_i versus 857 GHz correlation most tightly linear. b represents the constant offset that must be subtracted from the band ν_i map to make its zero level consistent with that of the 857 GHz map. For each band ν_i , we obtain estimates of m , d , and b by performing a linear least squares fit with uniform weights and iterative outlier rejection. Figure 3.2 shows scatter plots of the band ν_i versus 857 GHz correlation before (left) and after (right) correcting for the best-fit offset and residual dipole, for each $\nu_i \in \{143, 217, 353, 545\}$ GHz. Not only are the tightened correlations striking in these scatter plots, but the residual dipole subtractions appear very successful in the two-dimensional band ν_i maps themselves. Before performing thermal dust fits, we therefore subtract the best-fit b and $d\mathcal{D}$ from each 143-545 GHz map. The best-fit offsets and residual dipole amplitudes are listed in Table 3.1, along with other important per-band parameters, such as the fractional multiplicative calibration uncertainty c_ν .

We found that a dipole correction alone could not sufficiently rectify the *Planck* 100 GHz map (see Figure 3.3). Therefore, for 100 GHz, we performed a modified version of the Equation 3.3 fit, using the following model:

$$\mathcal{I}_{100,p} = m\mathcal{I}_{857,p} + b + \sum_{l=1}^4 \sum_{m=-l}^l a_{lm} Y_l^m(\theta_p, \phi_p) \quad (3.4)$$

Where Y_l^m are the real spherical harmonics, and the a_{lm} are their corresponding real coefficients. The angle ϕ_p is taken simply to be $l_{gal,p}$ and $\theta_p = (90^\circ - b_{gal,p})$. Thus,

we have replaced the Solar dipole term with a sum of 24 spherical harmonic templates, which, when multiplied by the best-fit a_{lm} coefficients and subtracted from *Planck* 100 GHz make the relation between 100 GHz and 857 GHz most tightly linear. Figure 3.3 illustrates the improved correlation of 100 GHz vs. 857 GHz when including the spherical harmonic corrections relative to the dipole-only correction. The spherical harmonic decomposition of Equation 3.4 did not improve the correlations at higher frequencies enough to warrant replacing the dipole-only correction in those cases.

3.4 Dust Emission Model

At sufficiently high frequencies, Galactic thermal dust emission can be adequately modeled as a single MBB with power-law emissivity (e.g. SFD; Planck Collaboration et al., 2013k). However, it has long been recognized, particularly in view of the FIRAS spectra, that the dust SED flattens toward the millimeter in a manner which is not consistent with a simple extrapolation of single-MBB models to low frequencies. In the diffuse ISM, Reach et al. (1995) found an improved fit to the FIRAS data using an empirically motivated superposition of two $\beta=2$ MBBs, one representing a ‘hot’ grain population ($T \approx 16-21$ K), the other a ‘cold’ grain population ($T \approx 4-7$ K). FDS99 built a more physically motivated two-MBB model, in which different grain emission/absorption properties account for the differing temperatures of each population, and these temperatures are coupled by assuming thermal equilibrium with the same interstellar radiation field (ISRF).

The primary FDS99 analysis considered the intrinsic grain properties of each species, for example the emissivity power law indices, to be constant over the sky, and performed a correlation slope analysis to constrain these parameters with FIRAS and

DIRBE observations. FDS99 also constructed a DIRBE $240\mu\text{m}/100\mu\text{m}$ ratio to account for temperature variation at $\sim 1.3^\circ$ resolution. In this work we seek to apply the FDS99 emission model to the *Planck* data set, which offers a dramatic enhancement in angular resolution relative to the FIRAS spectra. The *Planck* data thereby allow us to derive an improved temperature correction at near-*IRAS* resolution (§3.7.3), re-evaluate the best-fit global dust properties (§3.6, §3.7.5), and fit additional two-component model parameters as a function of position on the sky (§3.7.4).

The shape of the two-component model spectrum we will consider is given by:

$$M_\nu \propto \left[f_1 q_1 \left(\frac{\nu}{\nu_0} \right)^{\beta_1} B_\nu(T_1) + f_2 q_2 \left(\frac{\nu}{\nu_0} \right)^{\beta_2} B_\nu(T_2) \right] \quad (3.5)$$

Where B_ν is the Planck function, T_1 is the ‘cold’ dust temperature, T_2 is the ‘hot’ dust temperature, and β_1 and β_2 are the emissivity power-law indices of the cold and hot dust components respectively. q_1 represents the ratio of FIR emission cross section to optical absorption cross section for species 1, and similarly q_2 for species 2. f_1 and f_2 dictate the relative contributions of the two MBB components to the combined SED. Thus, f_1 and f_2 can be thought of as encoding the mass fraction of each species, although technically f_1 (f_2) is the optical absorption cross-section weighted mass fraction for species 1 (2). Following the convention of FDS99, we choose $\nu_0=3000$ GHz and take $f_2=(1-f_1)$.

Mathematically, this two-MBB model requires specification of seven parameters for every line of sight: T_1 , T_2 , β_1 , β_2 , f_1 , q_1/q_2 and the normalization of M_ν . However, under the assumption that the temperature of each species is determined by maintaining thermal equilibrium with the same ISRF, $T_1=T_1(T_2, \beta_1, \beta_2, q_1/q_2)$ is fully determined by these other parameters. T_1 is always related to T_2 via a simple

Table 3.2. Global Model Parameters

Number	Model	f_1	q_1/q_2	β_1	β_2	T_2	T_1	n	D.O.F.	χ^2	χ^2_ν
1	FDS99 best-fit	0.0363	13.0	1.67	2.70	15.72	9.15	1.018	7	23.9	3.41
2	FDS99 general	0.0485	8.22	1.63	2.82	15.70	9.75	0.980	3	3.99	1.33
3	single MBB	0.0	1.59	19.63	...	0.999	6	33.9	5.65

power law, although the prefactor and exponent depend on the parameters q_1/q_2 , β_1 and β_2 (see FDS99 Equation 14).

These considerations still leave us with six potentially free parameters per line of sight. Unfortunately, fitting this many parameters per spatial pixel is not feasible for our full-resolution 6.1' fits, as these are constrained by only five broadband intensity measurements. Hence, as in FDS99, we deem certain parameters to be “global”, i.e. spatially constant over the entire sky. In our full-resolution five-band fits, we designate β_1 , β_2 , f_1 and q_1/q_2 to be spatially constant. This same approach was employed by FDS99, and the globally best-fit values obtained by FDS99 for these parameters are listed in the first row of Table 3.2. With these global parameters, FDS99 found $T_2 \approx 16.2\text{K}$, $T_1 \approx 9.4\text{K}$ to be typical at high-latitude. In §3.6, we discuss the best-fit global parameters favored by the *Planck* HFI data; these are listed in the second row of Table 3.2.

Fixing the aforementioned four global parameters, our full-resolution, five-band fits have two remaining free parameters per line of sight: the hot dust temperature T_2 determines the SED shape and the normalization of M_ν determines the SED amplitude. In the lower-resolution fits of §3.7.4 which include all HFI bands, we will allow f_1 to be a third free parameter, still holding β_1 , β_2 , and q_1/q_2 fixed.

To calculate the optical depth in the context of this model, we assume optically thin conditions, meaning that $\tau_\nu = M_\nu/S_\nu$, where M_ν is the appropriately scaled

two-component model intensity and the source function is given by:

$$S_\nu = \frac{f_1 q_1 (\nu/\nu_0)^{\beta_1} B_\nu(T_1) + f_2 q_2 (\nu/\nu_0)^{\beta_2} B_\nu(T_2)}{f_1 q_1 (\nu/\nu_0)^{\beta_1} + f_2 q_2 (\nu/\nu_0)^{\beta_2}} \quad (3.6)$$

3.5 Predicting the Observed SED

The thermal dust emission model of §3.4 predicts the flux density per solid angle M_ν in e.g. MJy/sr for any single frequency ν . In practice, however, we wish to constrain our model using measurements in the broad *Planck*/DIRBE bandpasses, each with $\Delta\nu/\nu \sim 0.3$. Both the *Planck* and DIRBE data products quote flux density per solid angle in MJy/sr under the ‘IRAS convention’. More precisely, each value reported in the *Planck* maps gives the amplitude of a power-law spectrum with $\alpha=-1$, evaluated at the nominal band center frequency, such that this spectrum integrated against the transmission reproduces the bolometer-measured power. Because our model spectra do not conform to the $\alpha=-1$ convention, we have computed color correction factors to account for the $\text{MBB}(T, \beta)$ spectral shape and the transmission as a function of frequency:

$$b_{\nu_i}(T, \beta) = \frac{\int \nu^\beta B_\nu(T) \mathcal{T}_{\nu_i}(\nu) d\nu \left[\int (\nu_{i,c}/\nu) \mathcal{T}_{\nu_i}(\nu) d\nu \right]^{-1}}{\nu_{i,c}^\beta B_{\nu_{i,c}}(T)} \quad (3.7)$$

Here $\nu_{i,c}$ is the nominal band center frequency of band ν_i , with $\nu_{i,c} \in \{100, 143, 217, 353, 545, 857, 1249.1352, 2141.3747, 2997.92458\}$ GHz. $\mathcal{T}_{\nu_i}(\nu)$ represents the relative transmission as a function of frequency for band ν_i . For the HFI maps, $\mathcal{T}_{\nu_i}(\nu)$ is given by the *Planck* transmission curves provided in the file `HFI_RIMO_R1.10.fits` (Planck Collaboration et al., 2013c). For i100 and DIRBE $140\mu\text{m}$, $240\mu\text{m}$, we have adopted the corresponding DIRBE transmission curves.

The two-component model prediction in band ν_i under the IRAS convention, termed \tilde{I}_{ν_i} , is then constructed as a linear combination of color-corrected MBB terms:

$$\tilde{I}_{\nu_i} \propto \sum_{k=1}^2 b_{\nu_i}(T_k, \beta_k) f_k q_k (\nu_{i,c}/\nu_0)^{\beta_k} B_{\nu_{i,c}}(T_k) \quad (3.8)$$

The color correction of Equation 3.7 therefore allows us to predict \tilde{I}_{ν_i} by computing monochromatic flux densities at the central frequency $\nu_{i,c}$ and then multiplying by factors $b_{\nu_i}(T, \beta)$. In practice, we interpolated the color corrections off of a set of precomputed, one-dimensional lookup tables each listing $b_{\nu_i}(T, \beta)$ for a single β value as a function of T . We thus avoided the need to interpolate in both β and T by computing only a small set of one dimensional correction factors for the particular set of β values of interest (e.g. $\beta=1.67, 2.7, 1.63, 2.82 \dots$, see Table 3.2). This color correction approach makes the MCMC sampling described in §3.7.3 much more computationally efficient by circumventing the need to perform the integral in the numerator of Equation 3.7 on-the-fly for each proposed dust temperature. We have chosen to compute the color corrections on a per-MBB basis because this approach is very versatile; all possible two-component (and single-MBB) models are linear combinations of MBBs, so we can apply all of our color correction machinery even when we allow parameters other than temperature (e.g. f_1) to vary and thereby modify the dust spectrum shape.

With these color corrections and the formalism established in §3.4 in hand, we can mathematically state the model we will use e.g. during MCMC sampling to predict the observed SED. The predicted observation in band ν_i is given by:

$$\tilde{I}_{\nu_i} = \frac{\sum_{k=1}^2 b_{\nu_i}(T_k, \beta_k) f_k q_k (\nu_{i,c}/\nu_0)^{\beta_k} B_{\nu_{i,c}}(T_k) u_{\nu_i}^{-1}}{\sum_{k=1}^2 b_{545}(T_k, \beta_k) f_k q_k (545\text{GHz}/\nu_0)^{\beta_k} B_{545}(T_k)} \tilde{I}_{545} \quad (3.9)$$




Figure 3.4 Linear fits of SFD-reprocessed DIRBE 240 μm (left), 140 μm (center), and 100 μm (right) as a function of *Planck* 857 GHz. The red lines illustrate the DIRBE correlation slopes used in our dust emission model optimization of §3.6.

This equation is quite similar to Equation 3.8, but with two important differences. First, the normalization of \tilde{I}_{ν_i} is now specified by \tilde{I}_{545} , which represents the IRAS convention *Planck* 545 GHz intensity. The denominator serves to ensure that, for the case of $\nu_i=545$ GHz, \tilde{I}_{545} is self-consistent. Second, each term in the numerator is multiplied by a unit conversion factor $u_{\nu_i}^{-1}$. This factor is necessary because some of the *Planck* maps of interest have units of K_{CMB} (100-353 GHz), while the remaining maps (545-3000 GHz) have units of MJy/sr. We have adopted the strategy of predicting each band in its native units, whether MJy/sr or K_{CMB} . For this reason, we always evaluate $B_{\nu_i,c}$ in Equation 3.9 in MJy/sr and let $u_{\nu_i}=1$ (dimensionless) for $\nu_i \geq 545$ GHz. For $\nu_i \leq 353$ GHz, u_{ν_i} represents the conversion factor from K_{CMB} to MJy/sr, given by Planck Collaboration et al. (2013c) Equation 32.

3.6 Global Model Parameters

While we ultimately aim to obtain *Planck*-resolution maps of the spatially varying dust temperature and optical depth, we start by applying the machinery/formalism thus far developed to reassess the best-fit global two-component model parameters in

light of the *Planck* HFI data.

FDS99 determined the best-fit values of the two-component model global parameters β_1 , β_2 , q_1/q_2 and f_1 via a correlation slope analysis incorporating DIRBE and FIRAS data. Here we seek to estimate these same global parameters via an analogous correlation slope analysis in which we swap the *Planck* HFI maps for FIRAS at low frequencies, while still relying on DIRBE at higher frequencies. We also seek to determine via this correlation slope analysis whether or not the combination of *Planck*+DIRBE data favors two-component models over single-MBB models in the same way that the FIRAS+DIRBE data did in the FDS99 analysis.

In the two-component model case, based on a spectrum of *Planck* and DIRBE correlations slopes, we wish to obtain estimates for six free parameters: β_1 , β_2 , q_1/q_2 , f_1 , T_2 and the overall spectrum normalization n . The constraints we employ are the correlation slopes of each of the *Planck* HFI bands, as well as DIRBE $100\mu\text{m}$ (3000 GHz), $140\mu\text{m}$ (2141 GHz) and $240\mu\text{m}$ (1250 GHz) relative to *Planck* 857 GHz, i.e. dI_{ν_i}/dI_{857} . We will refer to the slope for band ν_i relative to *Planck* 857 GHz as s_{857,ν_i} . The slopes for *Planck* 100-545 GHz are taken to be those derived from the relative zero level fits of §3.3.5, and are illustrated by the dashed red lines in the right-hand column plots of Figure 3.2. The 857 GHz slope is unity by definition.

At 1250, 2141 and 3000 GHz, we use the SFD-reprocessed DIRBE maps. For each DIRBE band, we determine s_{857,ν_i} by performing a linear fit to DIRBE as a function of *Planck* 857 GHz, after both have been zodiacal light subtracted and smoothed to 1° FWHM, also restricting to the low HI mask of §3.3.5 (see Figure 3.4).

Counting 857 GHz, we thus have nine correlation slope constraints for six free parameters. Including DIRBE $140\mu\text{m}$ and $240\mu\text{m}$ is critical in making the problem at hand sufficiently overconstrained, and also in providing information near the peak of the dust SED at $\sim 160\mu\text{m}$, which is particularly sensitive to the presence of a single

versus multiple MBB components.

We assume an uncertainty of 2% on each of the s_{857,ν_i} and minimize the chi-squared given by:

$$\chi^2 = \sum_{i=0}^8 \frac{\left[s_{857,\nu_i} - n \frac{\tilde{I}_{\nu_i}(\beta_1, \beta_2, f_1, q_1/q_2, T_2)}{\tilde{I}_{857}(\beta_1, \beta_2, f_1, q_1/q_2, T_2)} \right]^2}{\sigma_{s_{857,\nu_i}}^2} \quad (3.10)$$

Where $\nu_i \in \{100, 143, 217, 353, 545, 857, 1250, 2141, 3000\}$ GHz. Note that this formula encompasses the general two-component case; in the single-MBB case, we take $f_1=0$ and hence q_1/q_2 , β_1 and T_1 are immaterial, but Equation 3.10 still applies. Note also that no ‘priors’ are included to preferentially drag our results towards agreement with those of FDS99. The correlation slopes $s_{857,\nu}$ and their adopted uncertainties are listed in the fifth and sixth columns of Table 3.1.

The results of our chi-squared minimization are listed in Table 3.2. First (model 1), we fix β_1 , β_2 , q_1/q_2 and f_1 to the best-fit values from the FDS99 analysis based on DIRBE+FIRAS. We then allow n and T_2 to vary so as to best match our DIRBE+*Planck* spectrum. This results in a reduced chi-squared of $\chi_\nu^2=3.41$. Reassuringly, n is quite close to unity. It should be noted though that our best-fit T_2 is ~ 0.5 K lower than that found by FDS99 for the same values of β_1 , β_2 , q_1/q_2 and f_1 .

Next (model 2), we consider the fully-general two-component model, allowing all six model parameters to vary. In this case, the reduced chi-squared of the best fit parameters is $\chi_\nu^2=1.33$, signifying that our introduction of four additional free parameters is justified. The best-fit β_1 and β_2 are both consistent with the corresponding FDS99 values to within 5%. $q_1/q_2=8.22$ represents a $\sim 40\%$ lower value than found by FDS99, while $f_1=0.0458$ represents a $\sim 25\%$ increase relative to FDS99. Again, our best-fit high-latitude T_2 is ~ 0.5 K lower than the typical value of $\langle T_2 \rangle=16.2$ K from FDS99.

Lastly, we calculate the optimal single-MBB fit to the *Planck*+DIRBE correlation slope spectrum. The best-fit single MBB has $\beta=1.59$, $T=19.63$, and $\chi^2_\nu=5.65$, indicating a significantly worse fit to the data than our best-fit two-component model (model 2). Thus, our *Planck*+DIRBE correlation slope analysis has confirmed the main conclusion of FDS99 and others e.g. Reach et al. (1995), that the FIR/submm dust SED prefers two MBBs to just one, but, for the first time, independent of FIRAS. Still, it is apparent that the improvement in χ^2_ν for single-MBB versus double MBB models found here is substantially less dramatic ($\Delta\chi^2_\nu=4.32$) than that found in FDS99 ($\Delta\chi^2_\nu=29.2$). This is likely attributable to the exquisite narrow-band frequency coverage of FIRAS, especially near the dust SED peak, which makes FIRAS a better suited data set than *Planck* for a detailed analysis of the globally best-fit dust SED model. In §3.7.5, we confirm the basic conclusions of this section via an approach in which we allow the dust temperature to vary spatially. The analysis of §7.5 also allows us to confirm the conclusions of this section while including a fully detailed uncertainty model; our assumption of 2% per-band uncertainties on the correlation slopes is largely a statement that we seek a model accurate to 2% from 100-3000 GHz, although the fact that our χ^2_ν values are order unity suggests that the assumed uncertainties are not grossly over or underestimated.

3.7 MCMC Fitting Procedure

The following subsections detail our procedure for constraining the two-component dust emission model parameters which are permitted to vary spatially. We use the MCMC procedure described to perform two types of fits: (1) full-resolution 6.1' fits, in which only the SED normalization and dust temperatures vary spatially, and (2) lower-resolution fits in which f_1 is also allowed to vary from one line of sight to

another.

3.7.1 Pixelization

For the purpose of fitting, we divide the sky into ~ 50 million pixels of angular size $\sim 1.72'$, defined by the HEALPix pixelization in Galactic coordinates, with $N_{side}=2048$. This pixelization is convenient because it is the format in which the *Planck* HFI maps were released, and because it adequately samples the $6.1'$ FWHM maps under consideration in our full-resolution fits. Our procedure will fit the intensity measurements in each spatial pixel independently.

3.7.2 Sampling Parameters

As discussed in §3.4, our full-resolution fits consider the “global” parameters f_1 , q_1/q_2 , β_1 , β_2 to be spatially constant. We employ the best-fit *Planck*+DIRBE global parameters of Table 3.2, model 2. For each line of sight, only the dust spectrum normalization and dust temperatures are allowed to vary. In order to predict the dust SED for a given pixel, we are thus left with two remaining degrees of freedom, and must choose an appropriate set of two parameters to sample and thereby constrain via MCMC. To determine the SED normalization in each pixel, we draw samples in \tilde{I}_{545} , the ‘IRAS convention’ intensity in the *Planck* 545 GHz bandpass, as defined in Equation 3.9. With the four aforementioned global parameters fixed, the dust spectrum shape is determined entirely by the two dust temperatures, which are coupled. To constrain the dust temperatures, we sample in T_2 , the hot dust temperature. For each sample in T_2 , we compute the corresponding value of T_1 , thereby fully specifying the SED shape. In principle, we could sample in either T_1 or T_2 , but have chosen to sample in T_2 because emission from this component dominates in the relatively high frequency bands which most strongly constrain the dust




Figure 3.5 Top: Summary of observed SEDs and best-fit thermal dust emission models for $\sim 13,000$ $N_{side}=2048$ pixels with similar best-fit temperatures and optical depths ($15.695\text{ K} < T_2 < 15.705\text{ K}$, $2.3 \times 10^{-5} < \tau_{545} < 2.5 \times 10^{-5}$). This region of parameter space was arbitrarily chosen in order to obtain a large number of pixels within a narrow T_2 interval and small fractional range in τ_{545} . Black points represent the average observed intensities after rescaling each pixel to $\tau_{545}=2.4 \times 10^{-5}$, while red error bars represent the typical per-pixel uncertainties at each frequency. For each pixel, the best-fit two-component model is derived via the MCMC procedure of §3.7.3, based on *Planck* 217-857 GHz and SFD $100\mu\text{m}$ at full $6.1'$ resolution. Note that the two lowest-frequency data points were not used to derive the average two-component fit shown (blue line), while the three lowest-frequency data points were not used to derive the average Planck Collaboration et al. (2013k) single-MBB fit shown (cyan line). Bottom: Comparison of average data, average two-component model and average Planck Collaboration et al. (2013k) single-MBB model after dividing out the average two-component model. Black error bars represent the uncertainty on the mean observed spectrum. The two-component fit is consistent with the average data from 100-3000 GHz, whereas extrapolating the Planck Collaboration et al. (2013k) model to 100-217 GHz yields predictions which are significantly low relative to the observed SED.

temperatures.

For the lower resolution fits described in §3.7.4, we sample in three parameters: \tilde{I}_{545} , T_2 , and f_1 .

3.7.3 Markov Chains

In our full-resolution fits, we use a MCMC approach to constrain the parameters \tilde{I}_{545} and T_2 . For each pixel, we run a Metropolis-Hastings (MH) Markov chain sampling the posterior probability of the observed 217-3000 GHz thermal dust SED as a function of the two parameters \tilde{I}_{545} and T_2 . More specifically, for each pixel, we are sampling the posterior given by:

$$P(\tilde{I}_{545}, T_2 | \mathbf{I}) \propto \mathcal{L}(\mathbf{I} | \tilde{I}_{545}, T_2) P(T_2) P(\tilde{I}_{545}) \quad (3.11)$$

Here \mathbf{I} denotes the vector of observed thermal dust intensities quoted under the ‘IRAS convention’: $\mathbf{I} = (I_{217}, I_{353}, I_{545}, I_{857}, I_{3000})$. The likelihood function is given by:

$$\mathcal{L}(\mathbf{I} | \tilde{I}_{545}, T_2) = \exp \left[-\frac{1}{2} (\mathbf{I} - \tilde{\mathbf{I}})^T \Sigma^{-1} (\mathbf{I} - \tilde{\mathbf{I}}) \right] \quad (3.12)$$

Here $\tilde{\mathbf{I}}$ is the vector of predicted observations based on Equation 3.9 and the proposed values of \tilde{I}_{545} and T_2 : $\tilde{\mathbf{I}} = (\tilde{I}_{217}, \tilde{I}_{353}, \tilde{I}_{545}, \tilde{I}_{857}, \tilde{I}_{3000})$. Σ is the per-pixel covariance matrix constructed based on the uncertainties in the observed intensities:

$$\Sigma = \begin{pmatrix} \sigma_{217}^2 & \dots & \rho_{217,3000} \sigma_{217} \sigma_{3000} \\ \vdots & \ddots & \vdots \\ \rho_{3000,217} \sigma_{3000} \sigma_{217} & \dots & \sigma_{3000}^2 \end{pmatrix} \quad (3.13)$$

For each pixel p in band ν_i , the variance of the measured value $I_{\nu_i}(p)$ is taken to be:

$$\sigma_{\nu_i}^2(p) = c_{\nu_i}^2 I_{\nu_i}^2(p) + c_{\nu_i}^2 \sigma_{CMB,\nu_i}^2 + (\delta O_{\nu_i})^2 + n_{\nu_i}^2 + \sigma_{CO,\nu_i}^2(p) + \sigma_{CIBA,\nu_i}^2 \quad (3.14)$$

This error budget is modeled after Planck Collaboration et al. (2013k) Equation B.1, but with some modifications and additions. The first term accounts for the multiplicative uncertainty on the input maps. Table 3.1 lists the multiplicative calibration uncertainty c_ν for each band. These values are taken from Table 11 of Planck Collaboration et al. (2013e). The second term represents an uncertainty due to our subtraction of the SMICA CMB model. The analogous term in Planck Collaboration et al. (2013k) Equation B.1 is $(c_\nu \times \text{SMICA}(p))^2$, i.e. an uncertainty proportional to the CMB model amplitude in each pixel. Because this term's spatial dependence can imprint the CMB anisotropies on the derived parameters, we have chosen to replace $\text{SMICA}(p)$ with a spatially constant, RMS value for the CMB amplitude, σ_{CMB,ν_i} . δO_{ν_i} represents the uncertainty in the band ν_i zero level offset, and the values of δO_{ν_i} can be read off from the second column of Table 3.1. n_{ν_i} represents the instrumental noise in band ν_i . Because using per-pixel noise estimates based on the *Planck* `ii_cov` parameter can imprint features of the survey pattern onto the derived parameters, we have adopted a conservative, spatially constant value of n_{ν_i} for each band. These values of n_{ν_i} are listed in Table 3.1. The next term accounts for the uncertainty on the CO emission correction, taking $\sigma_{CO,\nu_i}(p) = 0.15 \times I_{CO,\nu_i}(p)$ (see §3.3.4, specifically Equation 3.2).




Figure 3.6 Gridded posterior PDFs for three $N_{side}=2048$ HEALPix pixels, based on *Planck* 217-857 GHz and SFD $100\mu\text{m}$ at full $6.1'$ resolution. The colorscale is linear in $\log(P)$, with black corresponding to the maximum of $\log(P)$ and white representing $\max[\log(P)] - 5$. Light green crosses and ellipses mark the best-fit parameters and 1σ uncertainties based on our MCMC sampling of the posteriors. Our MCMC parameter and uncertainty estimates are in good agreement with those based on gridded posteriors. These three pixels are also representative in that we find the posterior distributions from Equation 3.11 are in general extremely well-behaved, showing no multimodality or other pathological qualities. Left: Low S/N pixel at high latitude in the Galactic north. Center: High S/N pixel in the Polaris flare region. Right: Low S/N pixel at high latitude in the Galactic south.

Finally, we include a term to account for the RMS amplitude of the cosmic infrared background anisotropy (CIBA) in band ν_i , σ_{CIBA,ν_i} . The values for the CIBA RMS amplitudes are obtained by assuming a $T=18.3$ K, $\beta=1.0$ MBB spectrum for the CIB, with 857 GHz normalization from Planck Collaboration et al. (2011b). The CIBA not only contributes to the per-band variance $\sigma_{\nu_i}^2$, but also to the inter-frequency covariances; this is why we have included the off-diagonal terms in the covariance matrix of Equation 3.13. In our noise model, the CIBA is the only source of inter-frequency covariance. Thus, the off-diagonal covariance matrix element between bands ν_i and ν_j is given by:

$$\Sigma_{ij} = \rho_{\nu_i,\nu_j} \sigma_{\nu_i} \sigma_{\nu_j} = \rho_{CIBA,\nu_i,\nu_j} \sigma_{CIBA,\nu_i} \sigma_{CIBA,\nu_j} \quad (3.15)$$

With values for ρ_{CIBA,ν_i,ν_j} from Planck Collaboration et al. (2013j). The approach we have taken in accounting for the CIBA is similar to that of Planck Collaboration et al. (2013k), Appendix C, in that we treat the CIBA amplitude in each pixel as a Gaussian random draw. However, instead of performing a separate analysis to gauge the uncertainty on derived dust parameters due to the CIBA, we allow the CIBA covariance to propagate naturally into our uncertainties via the likelihood function. Still, our treatment of the CIBA is a major oversimplification; a more sophisticated approach that accounts for the detailed CIBA spatial structure, or even removes the CIBA by subtraction would be preferable.

We include the following prior on the hot dust temperature:

$$P(T_2) = \mathcal{N}(T_2|\bar{T}_2, \sigma_{\bar{T}_2}) \quad (3.16)$$

With $\bar{T}_2 = 15.7$ K and $\sigma_{\bar{T}_2} = 1.4$ K. The T_2 prior mean is chosen based on the typical high-latitude T_2 value derived from the correlation slope analysis of §3.6. We find, as desired, that this relatively broad T_2 prior has little influence on the derived temperatures, other than to regularize the rare pixels with one or more defective intensities which might otherwise yield unreasonable parameter estimates. In principle, there can also be an informative prior on \tilde{I}_{545} . However, we have chosen to assume a uniform prior on the SED normalization and, as a matter of notation, will omit $P(\tilde{I}_{545})$ henceforward. In practice we always perform computations using logarithms of the relevant probabilities.

For each pixel, we initialize the Markov chain with parameters $\tilde{I}_{545} = I_{545}$ and T_2 consistent with the FDS99 DIRBE $100\mu\text{m}/240\mu\text{m}$ ratio map \mathcal{R} . The initial proposal distribution is a two-dimensional normal distribution, with $\sigma_{T_2}=0.25$ K, $\sigma_{\tilde{I}_{545}}=\max(0.01 \times I_{545}, 0.05 \text{ MJy/sr})$ and $\rho_{T_2,\tilde{I}_{545}}=0$. We run 5 iterations of burn-in,




Figure 3.7 Comparison of temperature maps based on FIR dust emission over a $10.5^\circ \times 8.3^\circ$ region centered about $(l, b) = (111.6^\circ, 18.3^\circ)$. Top: SFD temperature map based on DIRBE $100\mu\text{m}$ and $240\mu\text{m}$, with $\sim 1.3^\circ$ resolution. Center: $6.1'$ resolution two-component temperature based on *Planck* 217-857 GHz and SFD $100\mu\text{m}$. Bottom: Planck Collaboration et al. (2013k) single-MBB temperature map based on *Planck* 353-857 GHz and $100\mu\text{m}$ data, with $5.1'$ FWHM. Both temperature maps incorporating *Planck* observations clearly show a major improvement in angular resolution relative to SFD.

each consisting of 500 MH steps. After each burn-in iteration, we rescale the proposal distribution so as to ultimately attain an acceptance fraction f_{acc} as close as possible to the optimal value $f_{opt} = 0.234$. This is accomplished by multiplying the proposal distribution standard deviations by f_{acc}/f_{opt} .

After burn-in, we estimate the parameters and their uncertainties by performing 10,000 sampling steps, with $T_{2,j}$ and $\tilde{I}_{545,j}$ denoting the proposed parameter values at the j^{th} step since the end of burn-in. From these 10,000 samples, we compute estimates of each parameter's mean, $\langle T_2 \rangle = \langle T_{2,j} \rangle$, $\langle \tilde{I}_{545} \rangle = \langle \tilde{I}_{545,j} \rangle$, of each parameter's variance, $\sigma_{T_2}^2 = \langle T_{2,j}^2 \rangle - \langle T_{2,j} \rangle^2$, $\sigma_{\tilde{I}_{545}}^2 = \langle \tilde{I}_{545,j}^2 \rangle - \langle \tilde{I}_{545,j} \rangle^2$ and of the covariance $\sigma_{T_2} \sigma_{\tilde{I}_{545}} = \langle T_{2,j} - \langle T_2 \rangle \rangle \langle \tilde{I}_{545,j} - \langle \tilde{I}_{545} \rangle \rangle$.

After obtaining this initial estimate of the covariance matrix for each pixel, we re-run a second iteration of the entire MCMC procedure, starting from the first burn-in period. On this iteration, for each pixel, we begin with a proposal distribution that is a two-dimensional Gaussian with covariance equal to the first-pass covariance estimate. This gives the each pixel's proposal distribution approximately the 'right shape', whereas on the first pass we started by simply guessing the relative widths of the proposal distribution in \tilde{I}_{545} , T_2 , and also assumed that the $\rho_{T_2, \tilde{I}_{545}} = 0$.

Lastly, during post burn-in sampling, we also estimate the monochromatic two-component intensity at 545 GHz, $M_{545} = \langle M_{545,j} \rangle = \langle \tilde{I}_{545,j} / b_{545}(T_{2,j}, \beta_2) \rangle$, its variance, and the 545 GHz optical depth $\tau_{545} = \langle \tau_{545,j} \rangle = \langle M_{545,j} / S_{545,j} \rangle$ and its variance. τ_{545} and M_{545} are more readily useful than the sampling parameters themselves for translating our fit results into predictions of reddening (§3.8) and thermal dust emission (§3.9.2), respectively. At high Galactic latitude, we find a typical T_2 uncertainty of 0.45 K, and typical \tilde{I}_{545} fractional uncertainty of 13%. Figure 3.5 illustrates the two-component model SED and the intensity measurements which constrain our fits, while Figure 3.6 shows example posterior PDFs for three

pixels. Figure 3.7 shows a map of our derived hot dust temperature at full-resolution, for a patch of sky in the Polaris flare region.

We validated the parameters and uncertainties recovered from our MCMC procedure by comparing with results based on finely gridded posterior calculations performed on a random subset of pixels. These comparisons verified that the proposal distribution rescaling and reshaping steps that we employ do improve the accuracy of the recovered parameters/uncertainties, and that the parameters/uncertainties ultimately derived are highly reliable. We can quantify the fidelity of our MCMC parameter estimates by noting that the RMS fractional discrepancy between MCMC and gridded posterior means is 0.25% for \tilde{I}_{545} and 0.07% (~ 0.01 K) for T_2 . Regarding the accuracy of our uncertainty estimates, we find RMS fractional discrepancies of 2.2% for $\sigma_{\tilde{I}_{545}}$ and 2.4% for σ_{T_2} . Aside from these small statistical scatters, we find no biases in our MCMC estimates of the parameters and their uncertainties.

3.7.4 Low-resolution Fits

As mentioned in §3.4, the combination of high S/N and high angular resolution afforded by the *Planck* HFI maps provides us with the opportunity to allow additional parameters of the two-component model, previously fixed by FDS99, to vary spatially. Specifically, we consider allowing f_1 to vary, while maintaining β_1 , β_2 , and q_1/q_2 spatially constant. In principle, we could alternatively introduce a third free parameter by permitting β_1 , β_2 or q_1/q_2 to vary while holding f_1 fixed. However, a model in which f_1 varies continuously from one line of sight to another is the most natural three-parameter scenario, in that f_1 variation can be attributed to continuous changes in the dust species' mass fractions, whereas continuous variations in the other global parameters, which represent grain emission/absorption properties, seem less plausible.

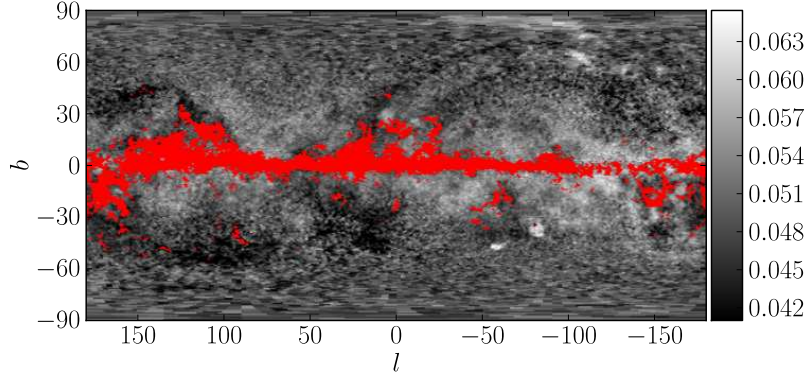


Figure 3.8 1° FWHM full-sky map of f_1 derived from our low-resolution fits described in §3.7.4. Red coloring masks pixels with appreciable molecular emission, as defined in §3.3.4. Such pixels should not be trusted in this analysis, which is sensitive to the SED shape at low frequencies affected by CO line emission. Variations in f_1 along the ecliptic plane are spurious results of imperfect zodiacal light subtractions. However, interesting astrophysical variations of f_1 are evident, particularly the trend of increasing f_1 with decreasing absolute Galactic latitude, the relatively low f_1 values in the Polaris flare and R Coronae Australis regions, and the clouds with relatively high f_1 values near the north Galactic pole.

In order for our variable f_1 fits to remain sufficiently constrained following the introduction of a third free parameter, we enhance per-pixel S/N by smoothing the input maps to 1° FWHM, and pixelize at $N_{side} = 64$. To best constrain the model parameters in each pixel, we also include *Planck* 100 GHz and 143 GHz, and DIRBE $140\mu\text{m}$ and $240\mu\text{m}$, all at 1° resolution.

We now run Markov chains sampling in all three of f_1 , \tilde{I}_{545} and T_2 , with the posterior given by:

$$P(\tilde{I}_{545}, T_2, f_1 | \mathbf{I}) \propto \mathcal{L}(\mathbf{I} | \tilde{I}_{545}, T_2, f_1) P(T_2) P(f_1) \quad (3.17)$$

The likelihood here is conceptually the same as that of Equation 3.12, but now depends on f_1 , which can vary from proposal to proposal within each individual pixel. The other difference is that \mathbf{I} and $\tilde{\mathbf{I}}$ now include 100 GHz, 143 GHz, $140\mu\text{m}$ and

240 μm , in addition to the five bands used for the full-resolution fits.

The prior $P(T_2)$ from Equation 3.16 remains unchanged. We adopt the following prior on f_1 :

$$P(f_1) = \mathcal{N}(f_1 | \bar{f}_1, \sigma_{\bar{f}_1}) \quad (3.18)$$

With $\bar{f}_1=0.0485$ (from Table 3.2, model 2) and $\sigma_{\bar{f}_1}=0.005$. This is a fairly stringent prior, but we must restrict the fit from wandering with too much freedom, as we are attempting to constrain three parameters using an SED with only nine intensity measurements, several of which are quite noisy. Again, we have adopted a uniform prior on \tilde{I}_{545} , and, as mentioned previously, we have omitted it from Equation 3.17 as a matter of notation.

The resulting full-sky map of f_1 is shown in Figure 3.8. A general trend of increasing f_1 towards lower absolute Galactic latitudes is apparent. The other most salient features are the relatively low values of f_1 in the Polaris flare and R Coronae Australis regions, and the relatively high f_1 clouds near the north Galactic pole.

3.7.5 Global Parameters Revisited

The posterior sampling framework thus far described also affords us an opportunity to evaluate the goodness-of-fit for competing dust SED models, and thereby cross-check the conclusions of our correlation slope analysis in §3.6. The basic idea will be to continue evaluating the posterior of Equation 3.11, but at low resolution ($N_{side}=64$), including all HFI bands as well as DIRBE 100 μm , 140 μm and 240 μm , and switching to a uniform prior on T_2 . Under these circumstances, the chi-squared corresponding to the best-fit parameters for pixel p , termed χ_p^2 , is simply $-2 \times \log(P_{max})$. We will refer to the per-pixel chi-squared per degree of freedom as $\chi_{p,\nu}^2$.

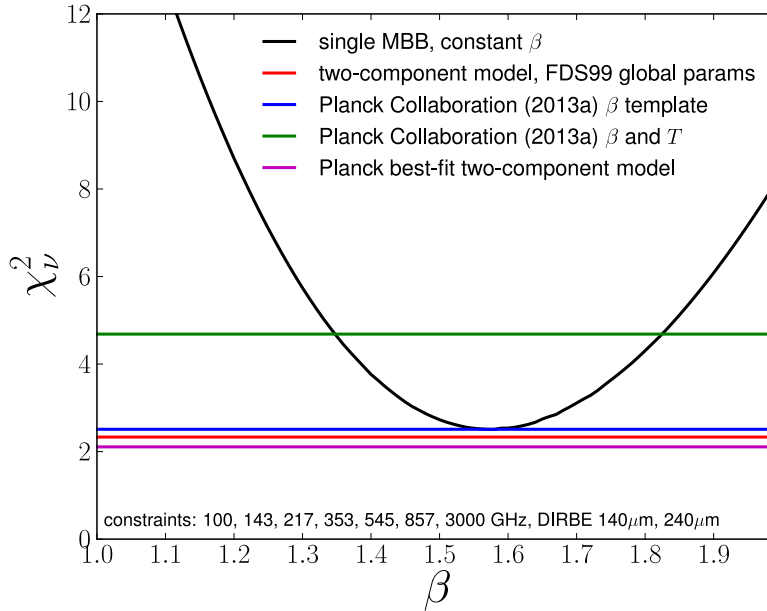


Figure 3.9 Comparison of goodness-of-fit, $\chi_\nu^2 = \langle \chi_{p,\nu}^2 \rangle$, for various dust SED models, as described in §3.7.5. For single-MBB models with spatially constant β , we varied β between 1 and 2 (horizontal axis), achieving reduced chi-squared χ_ν^2 shown by the black line, with $\beta=1.57$ providing the best single-MBB fit. Horizontal lines indicate χ_ν^2 for other dust emission models considered, including the FDS99 best-fit two-component model (Table 3.2, model 1, red) and the Planck Collaboration et al. (2013k) single-MBB model (green). The minimum χ_ν^2 is achieved with two-component ‘model 2’ from Table 3.2 (magenta).

Because we seek to compare the goodness-of-fit for various dust SED models in the diffuse ISM, we restrict to a set of $\sim 10,800$ pixels ($\sim 22\%$ of the sky), with $|b| > 30^\circ$ and $|\beta| > 10^\circ$. We also avoid the SMICA inpainting mask, pixels with appreciable CO contamination, and compact sources. The goodness-of-fit ‘objective function’ we employ to judge the quality of a particular dust SED model is $\langle \chi_{p,\nu}^2 \rangle$, where the average is taken over the aforementioned set of $\sim 10,800$ pixels. $\langle \chi_{p,\nu}^2 \rangle$ is also equivalent to the reduced chi-squared, χ_ν^2 , when considering the total number of free parameters to be the number of pixels multiplied by the number of free parameters per pixel (and similarly for the total number of constraints), and taking $\chi^2 = \sum \chi_p^2$.

We calculate χ_ν^2 for various dust SED models, independently minimizing each χ_p^2 by finding pixel p ’s best-fitting dust temperature and normalization, then evaluating $\langle \chi_{p,\nu}^2 \rangle$. First, we consider single-MBB models with β spatially constant (see the black line in Figure 3.9). $\beta=1.57$ yields the best fit, with $\chi_\nu^2=2.51$. This result is in excellent agreement with that of §3.6, where we found the best-fit single-MBB model to have $\beta=1.59$.

We also evaluated χ_ν^2 for single-MBB models in which β varies spatially. In these cases, we adopted the 0.5° resolution β map from Planck Collaboration et al. (2013k). We started by calculating χ_ν^2 using the Planck Collaboration et al. (2013k) temperature map, finding $\chi_\nu^2=4.68$. Note that in this case no per-pixel chi-squared minimization was involved, as we simply evaluated χ_p^2 for each pixel based on the fully-specified Planck Collaboration et al. (2013k) emission model. Next, we tested a single-MBB model for which we adopted the Planck Collaboration et al. (2013k) β map, but allowed the per-pixel temperature and normalization to vary so as to minimize χ_p^2 . In this case, we found $\chi_\nu^2=2.51$, effectively identical to the value found for the spatially constant $\beta=1.57$ single-MBB model. This is perhaps unsurprising, as the average β value from Planck Collaboration et al. (2013k) over the mask in

question is $\langle\beta\rangle=1.58$. This result does suggest, however, that in diffuse regions the half-degree variations in β are not materially improving the goodness-of-fit over the full frequency range 100-3000 GHz relative to a model with appropriately chosen spatially constant β .

We move on to evaluate two-component models, first calculating χ_ν^2 with the FDS99 global parameters (Table 3.2, model 1). We find $\chi_\nu^2=2.33$, a slight improvement relative to the best-fitting single-MBB models. Finally, we calculate χ_ν^2 for Table 3.2 model 2, the two-component model favored by our *Planck*+DIRBE correlation slopes. In this case, we achieve the best goodness-of-fit out of all the models we have tested, with $\chi_\nu^2=2.11$.

Thus, our degree-resolution goodness-of-fit analysis has generally confirmed the conclusions of §3.6. We find the single-MBB β value favored by the combination of *Planck* and DIRBE to be nearly identical here ($\beta=1.57$) versus in §3.6 ($\beta=1.59$). As in §3.6, we also find that the *Planck*+FIRAS and *Planck*+DIRBE best-fit two-component models from Table 3.2 outperform single-MBB alternatives, though only by a relatively small margin in χ_ν^2 . Still, because our present analysis has $\sim 75,500$ degrees-of-freedom, $\Delta\chi_\nu^2=0.4$ formally corresponds to an enormously significant improvement in χ^2 . The agreement between our correlation slope analysis and the present goodness-of-fit analysis is especially encouraging for three main reasons: (1) in the present analysis, dust temperature has been allowed to vary on degree scales, whereas in §3.6 we assumed a single global dust temperature (2) the present analysis employs a fully detailed, per-pixel uncertainty model and (3) in the present analysis, our zero-level offsets factor into the dust temperature, whereas in §6 this was not the case, meaning the former and latter analyses agree in spite of their potential to be affected by rather different systematics.

3.8 Optical Reddening

While the temperature and optical depth maps thus far derived are useful for making thermal dust emission foreground predictions, estimating optical reddening/extinction is another important application of the τ_{545} map. Translating our two-component optical depth to reddening is especially valuable because our T_2 map has $\sim 13\times$ better angular resolution than the SFD temperature correction, and thus there is reason to believe our two-component reddening estimates may be superior to those of SFD. However, as discussed in §3.11.2, we do not yet advocate for the wholesale replacement of SFD, and more detailed work is still necessary to determine/quantify the extent to which *Planck*-based dust maps might improve reddening estimates relative to SFD.

3.8.1 Reddening Calibration Procedure

We calibrate optical depth to reddening empirically rather than derive a relationship between τ_{545} and reddening by introducing additional assumptions about the dust grain physics and size distribution. To achieve this empirical calibration, we must adopt a set of calibrator objects for which true optical reddening is known. There are various possibilities at our disposal. Planck Collaboration et al. (2013k) calibrated their radiance and τ_{353} maps to $E(B - V)$ using broadband Sloan Digital Sky Survey (SDSS; York et al., 2000) photometry for a set of $\sim 10^5$ quasars. The SFD calibration was originally tied to a sample of 384 elliptical galaxies, but was later revised by Schlafly & Finkbeiner (2011, hereafter SF11) based on $\sim 260,000$ stars with both spectroscopy and broadband photometry available from the SEGUE Stellar Parameter Pipeline (SSPP, Lee et al., 2008).

To calibrate our two-component optical depth to reddening, we make use of the




Figure 3.10 Linear fit of $E(B - V)_{SSPP}$ as a function of two-component 545 GHz optical depth, illustrating our procedure for calibrating optical depth to reddening, as described in §3.8.1.

stellar sample from SF11. Given a library of model stellar atmospheres, the spectral lines of these stars can be used to predict their intrinsic optical broadband colors. The ‘true’ reddening is then simply the difference between the observed $g - r$ color and the $g - r$ color predicted from the spectral lines. Applying a color transformation then yields ‘true’ $E(B - V)$ values for $\sim 260,000$ lines of sight. Throughout our SSPP calibration analysis, we restrict to the $\sim 230,000$ lines of sight with $|b| > 20^\circ$ in order to avoid stars which may not lie behind the full dust column. In this section and §8.2, we make absolute latitude cuts (in both b and β) at 20° , to match the footprint of SF11 and adapt to the non-uniform distribution of SSPP stars on the sky. The calibration of two-component optical depth to $E(B - V)$ is performed as a linear regression of $E(B - V)_{SSPP}$ versus τ_{545} . τ_{545} is considered to be the independent variable in this regression, as we ultimately wish to predict $E(B - V)$ as a function of




Figure 3.11 (top left) Residuals of $E(B - V)_{2comp}$ relative to $E(B - V)_{SSPP}$ as a function of $E(B - V)_{SFD}$. The grayscale represents the conditional probability within each $E(B - V)_{SFD}$ bin. The central black line shows the moving median. The upper and lower black lines represent the moving 75th and 25th percentiles respectively. (bottom left) Residuals of $E(B - V)_{2comp}$ relative to $E(B - V)_{SSPP}$ as a function of hot dust temperature T_2 . (top right) Same as top left, but illustrating the residuals of $E(B - V)_{mbb}$, our calibration of the Planck Collaboration et al. (2013k) τ_{353} to $E(B - V)_{SSPP}$. (bottom right) Same as bottom left, but showing the $E(B - V)_{mbb}$ residuals as a function of the single-MBB dust temperature from Planck Collaboration et al. (2013k). The temperature axes always range from the 0.4th percentile temperature value to the 99.6th percentile temperature value.

optical depth, and τ_{545} has much higher S/N than the SSPP $E(B - V)$ estimates.

This regression is illustrated in Figure 3.10. As expected, there is a strong linear correlation between $E(B - V)_{SSPP}$ and τ_{545} . The conversion factor from τ_{545} to $E(B - V)$ is 2.62×10^3 . Reassuringly, the best-fit offset is close to zero, ~ 2.6 mmag.

Figure 3.11 shows the residuals of our τ_{545} -based reddening predictions, $E(B - V)_{2comp}$, relative to the corresponding SF11 reddening measurements, $E(B - V)_{SSPP}$, as a function of SFD reddening, $E(B - V)_{SFD}$, (top left panel) and as a function of hot dust temperature (bottom left panel). For comparison, the right panels show analogous residual plots, but with respect to reddening predictions based on our calibration of the Planck Collaboration et al. (2013k) 353 GHz optical depth to $E(B - V)_{SSPP}$, using the same regression procedure employed to calibrate $E(B - V)_{2comp}$. We refer to these reddening predictions based on the Planck Collaboration et al. (2013k) single-MBB model and calibrated to the SF11 measurements as $E(B - V)_{mbb}$.

All four residual plots in Figure 12 show systematic problems at some level. The most striking systematic trend is the ‘bending’ behavior of the reddening residuals versus $E(B - V)_{SFD}$ (top panels), with the median residual bottoming out near -10 mmag at $E(B - V)_{SFD} \approx 0.15$ mag. This behavior is common to both $E(B - V)_{2comp}$ and $E(B - V)_{mbb}$, and in fact was first noted in the residuals of $E(B - V)_{SFD}$ itself relative to $E(B - V)_{SSPP}$ by SF11 (see their Figure 6). Such a bending behavior is troubling because it could indicate a nonlinearity common to many FIR reddening predictions based on column densities inferred from dust emission. Alternatively, because the SF11 stars are distributed over the sky in a highly non-uniform manner, the bend could arise from aliasing of discrepancies particular to certain sky regions (e.g. inner vs. outer Galaxy) on to the $E(B - V)_{SFD}$ axis.

The obvious culprit for any potential nonlinearity in FIR-based reddening estimates




Figure 3.12 Same as Figure 3.11, but restricting to high ecliptic latitude, $|\beta| > 20^\circ$. In both the top left and top right plots, the bending of the reddening residuals as a function of $E(B-V)_{SFD}$ seen in Figure 3.11 has been eliminated. Further, the two-component reddening residual temperature dependence (bottom left) has been significantly reduced relative to the corresponding trend shown in Figure 3.11. For $E(B-V)_{SFD} \gtrsim 0.3$ mag, the top row plots appear noisy because there are an insufficient number of remaining SSPP points of comparison.

is a faulty temperature correction. For this reason, we have included the bottom panels of Figure 3.11, to check for the presence of a temperature dependence of the reddening residuals. Indeed, in both the two-component and single-MBB cases there exists some systematic dependence of the reddening residuals on temperature. For $T_{mbb} \gtrsim 19$ K, the median residual is reasonably flat, but at lower temperatures (the lowest temperature $\sim 20\%$ of SSPP sight lines), the median shows trends at the ~ 10 mmag level. On the other hand, the median residual in the two-component case trends downward with increasing T_2 over the entire T_2 range shown, with a peak-to-peak amplitude of ~ 20 mmag.

3.8.2 Rectifying the Reddening Residuals

In this section we describe our attempts to eliminate the systematic problems in the two-component reddening residuals shown in the left column of Figure 3.11. We employed two main strategies: (1) recomputing the two-component τ_{545} by re-running our Markov chains after modifying the input maps and/or changing the particular two-component model parameters adopted and (2) making spatial cuts to isolate sky regions in which the residuals are especially pristine (or especially problematic).

The following is a list of dust model modifications we tested, but which proved to have little impact on the reddening residual trends as a function of either

$E(B - V)_{SFD}$ or T_2 :

- Varying each of the global two-component model parameters β_1 , β_2 , q_1/q_2 and f_1 individually while holding the others fixed.
- Allowing f_1 to vary spatially as in the fits of §3.7.4.
- Changing the mean and/or variance of the T_2 prior.
- Varying multiple global parameters at a time e.g. both f_1 and q_1/q_2 , restricting to regions of parameter space favored by our goodness-of-fit analyses described

in §3.6 and §3.7.5.

We additionally investigated the following spatial cuts which did not resolve the dominant problems noted in the reddening residuals:

- Separating Celestial north and south.
- Separating Galactic north and south.
- Separating inner and outer Galaxy.
- Combining the above two sets of cuts i.e. separating the Galaxy into quadrants.

- Combining these spatial cuts with the dust model changes of the previous list.

However, we found that changing the zero level offsets of the input maps had a significant effect on the strength of the anticorrelation between median reddening residual and T_2 . In particular, we experimented with perturbing the zero level offset of *Planck* 857 GHz while correspondingly changing the zero levels of the remaining *Planck* maps based on the prescription of §3.3.5. We also experimented with changing the zero level of SFD i100, independent of the other zero levels. Unfortunately, completely flattening the reddening residual dependence on T_2 required unreasonably large zero level modifications. For example, flattening the T_2 residual required adding $\gtrsim 0.6$ MJy/sr to the i100 map. Such an offset is implausible, being an order of magnitude larger than the nominal i100 zero level uncertainty quoted by SFD, and comparable to the entire 3000 GHz CIB monopole signal. Furthermore, we note that even these large zero level modifications had virtually no effect in eliminating the reddening residual ‘bend’ versus $E(B - V)_{SFD}$. Thus, changing the zero level offsets showed hints of promise in rectifying the reddening residual temperature dependence, but could not by itself completely resolve the systematic trends in reddening residuals.

The only solution we have been able to identify that both removes the ‘bend’ vs. $E(B - V)_{SFD}$ and simultaneously reduces the temperature dependence of the reddening residuals is cutting out the ecliptic plane by restricting to $|\beta| > 20^\circ$. In this case, we completely eliminated the bending behavior of the residual versus $E(B - V)_{SFD}$, and significantly reduced the T_2 dependence to a peak-to-peak amplitude of only ~ 10 mmag (see Figure 3.12). Figure 3.12 still includes the single-MBB plots (right column), to show that the bend versus $E(B - V)_{SFD}$ is eliminated by the $|\beta|$ cut, even for the single-MBB model. However, the single-MBB residuals still differ systematically from zero for $T \lesssim 19$ K. Perhaps the improvements in the two-component reddening residuals after restricting to high ecliptic latitude should come as no surprise, given that the ecliptic plane is the most obvious systematic problem with our temperature map (see the full-sky results shown in Figure 3.16).

After cutting the ecliptic plane, we found that only small zero level perturbations were required to fully flatten the temperature residuals, while still maintaining flat residuals versus $E(B - V)_{SFD}$. The optimal offsets we found were ± 0.08 MJy/sr to ± 100 and ± 857 GHz respectively (see Figure 3.13). These offsets are well within reason, given the nominal zero level uncertainties quoted in Table 3.1.

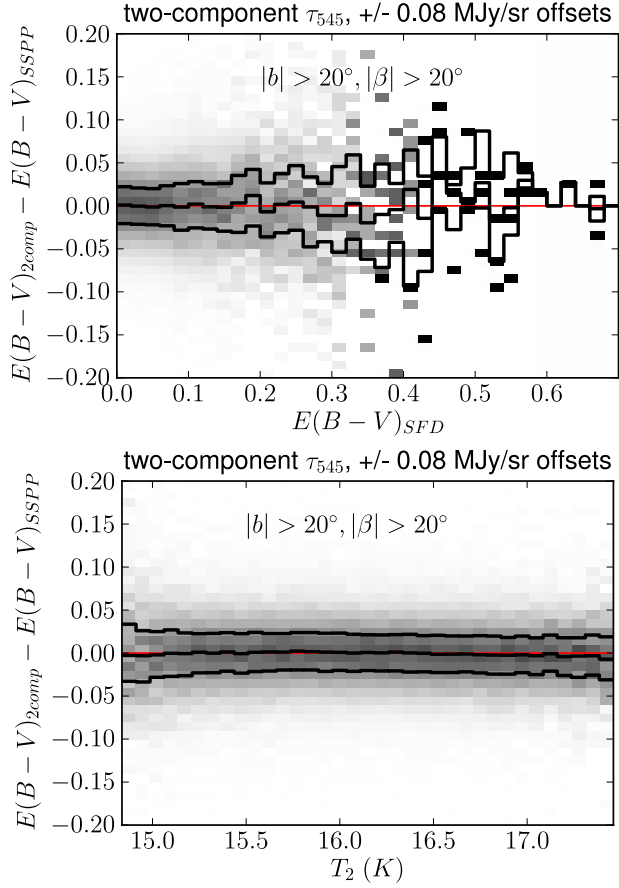


Figure 3.13 Two-component reddening residuals after restricting to high ecliptic latitude ($|\beta| > 20^\circ$) and perturbing the $i100$ and 857 GHz zero levels by $+0.08$ MJy/sr and -0.08 MJy/sr respectively. The bending behavior as a function of $E(B-V)_{SFD}$ has been eliminated, and virtually no temperature dependence remains. For $E(B-V)_{SFD} \gtrsim 0.3$ mag, the top plot appears noisy because there are an insufficient number of remaining SSPP points of comparison following our cut on ecliptic latitude.

3.9 Comparison of Emission Predictions

3.9.1 The 353-3000 GHz Frequency Range

Here we compare our two-component emission predictions to those of the Planck Collaboration et al. (2013k) single-MBB model in the 353-3000 GHz range.

This frequency range represents the overlap between the recommended range of applicability for the Planck Collaboration et al. (2013k) model and the 100-3000 GHz frequency range of our two-component model. Since we have used input maps that are very similar to those of Planck Collaboration et al. (2013k), and since our model and the Planck Collaboration et al. (2013k) model both fit the data well in this frequency range, good agreement between our two-component predictions and those of the Planck Collaboration et al. (2013k) single-MBB model is to be expected.

We compare the emission models in this frequency range by using each model in turn to predict the observed *Planck* 353, 545, and 857 GHz maps, as well as the 3000 GHz DIRBE/*IRAS* map. We rebin to $N_{side}=64$ and restrict to the diffuse sky regions of our mask from §3.7.5. We summarize this comparison by producing a per-band scatter plot of the Planck Collaboration et al. (2013k) prediction versus the two-component prediction, and performing a linear regression between these two quantities. Before plotting and performing these regressions, we adjusted the Planck Collaboration et al. (2013k) predictions to account for the differing zero level offsets used in this work and in Planck Collaboration et al. (2013k). For instance, at 3000 GHz, Planck Collaboration et al. (2013k) added 0.17 MJy/sr to the SFD98 zero level, whereas we made no such modification; therefore, for the sake of comparison, we subtracted 0.17 MJy/sr from the Planck Collaboration et al. (2013k) predictions before plotting and performing the 3000 GHz regression.

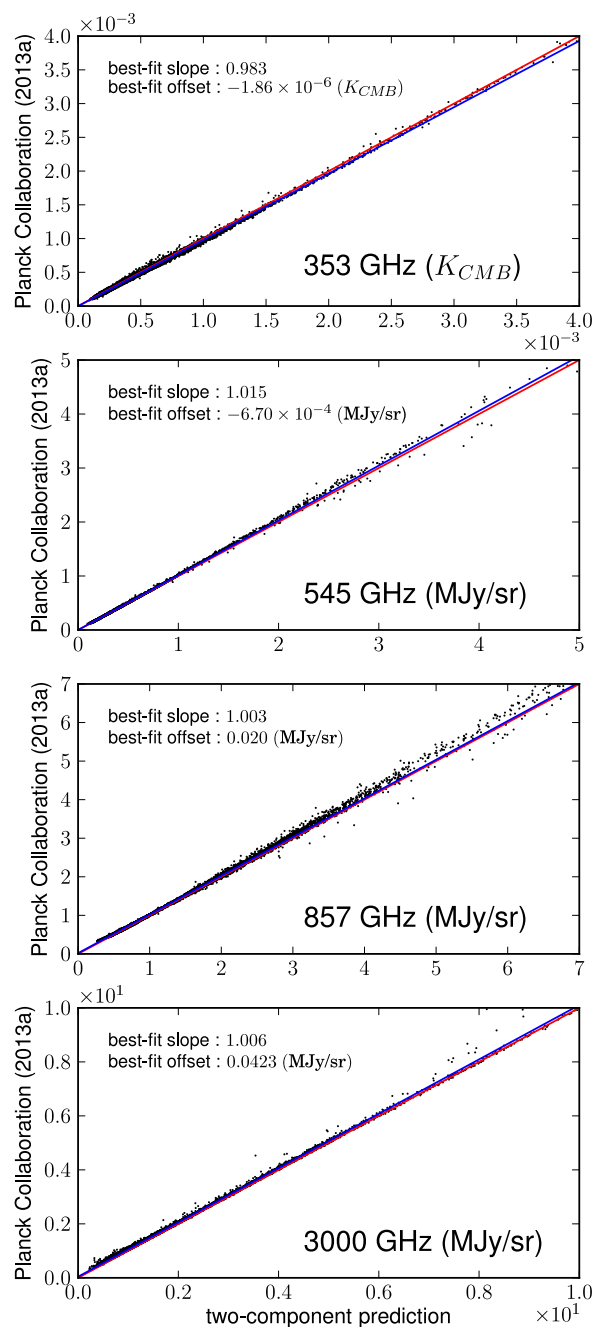


Figure 3.14 Scatter plots of Planck Collaboration et al. (2013k) single-MBB predictions (vertical axes) versus our two-component predictions (horizontal axes), rebinning to $N_{side}=64$ and restricting to the diffuse regions of §3.7.5. The lines of best fit are shown in blue, and red lines represent perfect agreement between the two predictions. Note that a per-band offset has been applied to the Planck Collaboration et al. (2013k) predictions to account for the differing zero level offsets used in building the two models. After accounting for the different zero levels, the best fit offsets between predictions are consistent with zero to within the uncertainties quoted in Table 3.1. The slopes are also within $\leq 1.7\%$ of unity.

The slopes obtained from these linear fits indicate very good agreement between the single-MBB and two-component models, with values between 0.983-1.015 (agreement at the $\leq 1.7\%$ level). The offsets are also consistent with zero to within the uncertainties quoted in Table 3.1. We do not find evidence that our two-component model provides emission predictions in the 353-3000 GHz range which are superior to those of Planck Collaboration et al. (2013k). From 353-3000 GHz and in diffuse sky regions, the main difference between emission predictions from these two models will be overall offsets due to differing input map zero levels.

3.9.2 The 100-217 GHz Frequency Range

FDS99 originally performed their FIRAS+DIRBE dust SED analysis for the sake of accurately forecasting low-frequency CMB foregrounds. Recently, Galactic CMB foregrounds, especially in the 100-150 GHz frequency range, have become a focal point of cosmology owing to the Ade et al. (2014) B -mode polarization results. Here we show that our two-component foreground predictions remain accurate on average to within 2.2% from 100-217 GHz, and we quantify the benefit of using our two-component emission predictions in this frequency range relative to extrapolating the Planck Collaboration et al. (2013k) single-MBB model.

To assess the accuracy of low-frequency emission predictions, we compare the observed *Planck* HFI map at each of 100, 143, 217 GHz to the corresponding single-MBB and two-component predictions, with all maps smoothed to 1° FWHM and binned down to $N_{side}=64$. We restrict to the same set of pixels used for the goodness-of-fit analysis of §3.7.5, with $|b| > 30^\circ$ and $|\beta| > 10^\circ$, also avoiding molecular emission, the SMICA inpainting mask, and compact sources. We then perform a linear fit between the *Planck* observed emission and the predicted emission at each frequency and for each emission model. For these fits, we consider the

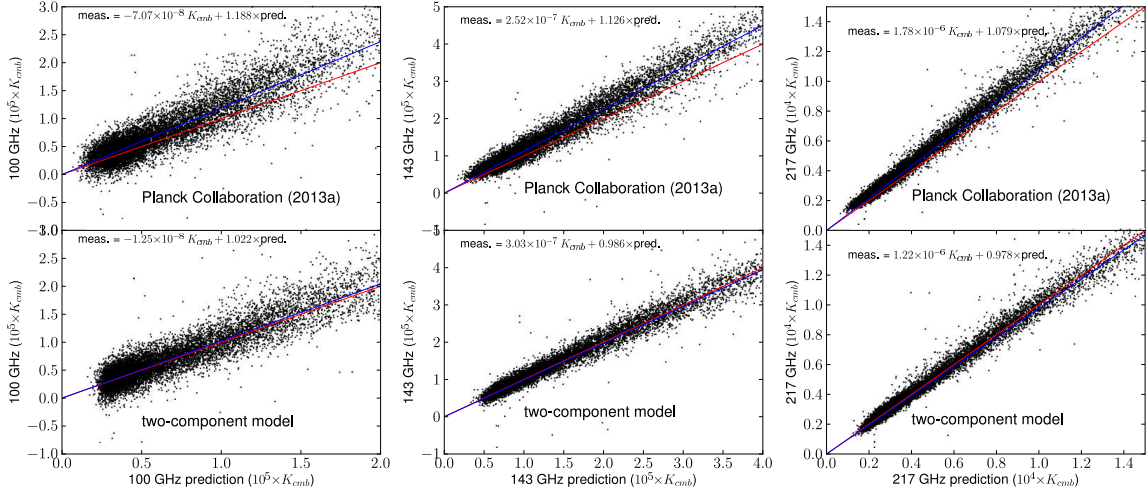


Figure 3.15 Comparison between low-frequency thermal dust emission predictions from our best-fit two-component model (Table 3.2, model 2) and those based on extrapolation of the Planck Collaboration et al. (2013k) model. The top row shows scatter plots of the Planck Collaboration et al. (2013k) predictions versus observed *Planck* 100 GHz (left), *Planck* 143 GHz (center) and *Planck* 217 GHz (right). The bottom row shows scatter plots of the corresponding two-component predictions versus *Planck* observations. In all cases, the blue line indicates the best-fit linear relationship, while the red line represents a perfect match between predictions and observations. The lines of best-fit illustrate that the single-MBB model systematically underpredicts emission (in the multiplicative sense) by 18.8%, 12.6% and 7.9% at 100, 143 and 217 GHz respectively. On the other hand, by the same metric, the two-component model predictions at 100-217 GHz are always accurate to within $\leq 2.2\%$. The two-component fit results shown are based on 217-3000 GHz observations, meaning that the 100 GHz and 143 GHz predictions are truly extrapolations, while the 217 GHz agreement is enforced by the fitting process itself to some extent.

predicted emission to be the independent variable, since it has higher S/N than the observations, especially at 100 and 143 GHz. We also assign pixel weights proportional to the predicted emission, so that the best-fit lines faithfully capture the linear trend exhibited without being biased by the large number of very low S/N pixels with minimal emission. Scatter plots between the predicted and observed emission are shown in Figure 3.15. The best-fit lines are overplotted and their equations are given in the top left corner of each subplot.

In both the single-MBB and two-component cases, all of the best fit offsets are within the uncertainties quoted in Table 3.1. On the other hand, the top row of Figure 3.15 shows that the Planck Collaboration et al. (2013k) single-MBB extrapolations yield slopes substantially different from unity: 1.079 at 217 GHz, 1.126 at 143 GHz, and 1.188 at 100 GHz. The fact that the slopes are larger than unity indicates that the Planck Collaboration et al. (2013k) extrapolations are systematically low. The systematic underprediction evidently becomes gradually more pronounced as lower frequencies are considered, with a 7.9% underprediction at 217 GHz, a 12.6% underprediction at 143 GHz and an 18.8% underprediction at 100 GHz. A deficit in single-MBB predictions relative to the observed *Planck* 100-217 GHz emission was also noted in Planck Collaboration et al. (2011a), e.g. their Figure 7.

For the case of the two-component model, we perform full-resolution 217-3000 GHz fits using the *Planck*+DIRBE favored global parameters (Table 3.2, model 2), then smooth to 1° FWHM and bin down to $N_{side} = 64$ before predicting the 100-217 GHz emission. The bottom row of Figure 14 shows that each of the best-fit lines is very similar to the corresponding red line which represents a perfect match between predicted and observed emission. More quantitatively, the two-component slopes are all within 2.2% of unity: 0.978 at 217 GHz, 0.986 at 143 GHz and 1.022 at 100 GHz. We note that at 217 GHz, the good agreement is in some sense predetermined by the

fact that *Planck* 217 GHz has been included in our two-component MCMC fits. On the other hand, the 143 and 100 GHz predictions are based on extrapolation.

We conclude from these predicted versus observed emission comparisons that our two-component model outperforms extrapolation of the Planck Collaboration et al. (2013k) single-MBB model at predicting Galactic thermal dust emission in diffuse regions from 100-217 GHz. It should be reiterated, once again, that Planck Collaboration et al. (2013k) did not intend for their single-MBB model to be extrapolated to frequencies below 350 GHz (see their §7.2.1), whereas we optimized our two-component model to be valid over the entire 100-3000 GHz frequency range. Our two-component model thus represents the first *Planck* based thermal dust emission model valid over the entire 100-3000 GHz frequency range.

3.10 Data Release

We are releasing a set of $N_{side}=2048$ HEALPix maps in Galactic coordinates which summarize the results of our full-resolution two-component dust fits. Low-resolution renderings of our full-sky dust temperature and optical depth maps are shown in Figure 3.16. Our data release also includes software utilities for obtaining emission and reddening predictions from our *Planck*-based two-component fits. Refer to the data release documentation and FITS file headers for further details.¹

3.11 Conclusions

3.11.1 Single-MBB versus Two-component emission

A major aim of this work has been to determine whether the FDS99 two-component dust emission model remains favored over single-MBB models when swapping the

¹<http://faun.rc.fas.harvard.edu/ameisner/planckdust>




Figure 3.16 (top) Hot dust temperature derived from our full-resolution two-component model fits of *Planck* 217-857 GHz and SFD $100\mu\text{m}$, downbinned to $27.5'$ resolution. (bottom) Corresponding full-sky map of best-fit two-component 545 GHz optical depth.

Planck HFI maps for FIRAS at frequencies below 1250 GHz. We compared dust SED models in two ways (1) by fitting a 100-3000 GHz spectrum composed of per-band correlation slopes versus *Planck* 857 GHz (2) by finding the best-fit dust temperature and optical depth per line-of-sight, with each pixel's SED comprised of 100-3000 GHz *Planck*+DIRBE data, and comparing the average goodness-of-fit under various emission models.

In both the correlation slope analysis of §3.6 and the goodness-of-fit analysis of §3.7.5 we found that the best-fit *Planck*+DIRBE two-component model (Table 3.2, model 2) outperformed the best-fit single-MBB model, but by a lesser margin in χ^2_ν than found by FDS99 using FIRAS+DIRBE. Specifically, our best-fit *Planck*+DIRBE two-component model yielded an improvement of $\Delta\chi^2_\nu=3.41$ (§3.6) and $\Delta\chi^2_\nu=0.4$ (§3.7.5). This represents a far less dramatic contrast in χ^2_ν than found by the FDS99 correlation slope analysis, $\Delta\chi^2_\nu=29.2$. Perhaps a relative lack of discrimination amongst competing dust SED models when relying on *Planck*+DIRBE is to be expected, given that our constraints include only nine broad frequency channels, whereas FDS99 employed >200 narrow bands. Still, $\Delta\chi^2_\nu=0.4$ from §3.7.5 is formally of enormous significance, given the $\sim 75,000$ degrees of freedom in that analysis.

Nevertheless, we have established that the two-component emission model remains viable in light of the *Planck* HFI data, and that the FIR/submm dust SED's preference for two MBB components rather than just one is not simply an idiosyncrasy of the FIRAS spectra. Furthermore, we showed in §3.9.2 that our 100-217 GHz two-component emission predictions are on average accurate to within 2.2%, whereas extrapolating the Planck Collaboration et al. (2013k) single-MBB model systematically underestimates low-frequency dust emission by 18.8% at 100 GHz, 12.6% at 143 GHz and 7.9% at 217 GHz. We therefore recommend that those interested in thermal dust foregrounds in the 100-3000 GHz frequency range use our

data release to predict unpolarized dust emission, at the very least in order to help determine the level at which the choice of dust emission model may influence their conclusions.

3.11.2 Towards a Replacement for SFD

Because of the broad frequency coverage and high angular resolution afforded by the *Planck* HFI full-sky maps, we initially speculated that a *Planck* based extinction map might easily outperform SFD, the most commonly used optical reddening map.

However, at this point in time, we do not yet recommend that the results presented in this work be considered a replacement for SFD in terms of optical extinction/reddening estimates.

The CIBA remains a major imperfection that still requires further investigation. The CIB anisotropies are very evident in low-dust regions of our maps of optical depth and predicted dust emission. As described in §3.7.3, we have propagated the CIBA RMS amplitudes and inter-frequency covariances into our uncertainty estimates through the likelihood function in our MCMC procedure. However, this treatment falls far short of actually removing the spatial imprint of the CIBA on our derived parameters. The CIB anisotropies are more prominent in our optical depth map relative to that of SFD because of the lower-frequency *Planck* maps we rely upon to achieve a high-resolution temperature correction.

Imperfect zodiacal light (zodi) corrections represent a second major limitation of our results. The ecliptic plane’s prominence in our full-sky temperature map (Figure 3.16) suggests that the zodiacal light subtractions performed on the input maps are not ideal. Our comparisons of the FIR maps used in this study against HI emission bear out this notion, further revealing that the imperfect zodi corrections are not limited to $\ell > 100$, but in fact are noticeable in all of the HFI

R1.10_nominal_ZodiCorrected maps as well. We deemed it infeasible to reconsider all of the *Planck* zodi corrections in addition to the 3000 GHz zodi correction as a part of this study, especially considering that the forthcoming *Planck* 2014 release is expected to include a revised/improved zodi subtraction.

Irrespective of the notable imperfections in our results, more detailed comparisons between our reddening estimates here and those of SFD are required to determine/quantify which map is superior in particular applications. One definitive improvement of our reddening estimates relative to those of SFD is our ability to quote reddening uncertainties, which results from the probabilistic framework of §3.7.3. The extinction estimates from this work can also be employed as an alternative to those of SFD, to gauge the impact of dust map choice in a specific end user's application.

Chapter 4

Conclusion

The work presented within this dissertation lays the foundation for a definitive next-generation, full-sky dust data product based on *WISE* and *Planck*. Toward this end, a number of major challenges still remain to be addressed. Future surveys will both benefit from and help to achieve advancements in dust-related foreground models.

On small angular scales, the dominant limitation of *Planck*-based dust models at high Galactic latitude arises from the cosmic infrared background anisotropies (CIBA). Despite the nominally high temperature resolution provided by *Planck*, much of the small-scale temperature and optical depth variation at high Galactic latitude seen in Meisner & Finkbeiner (2015) and Planck Collaboration et al. (2013k) is attributable to the CIBA. While the typical amplitude of the CIBA imprint on a given line of sight's reddening is small, the CIBA is correlated with the large-scale structure in such a way as to potentially introduce systematic biases in cosmological inferences when dereddening large optical surveys. To begin quantifying the level at which *Planck*-based dust models limit precision cosmology in the optical, it would be informative to propagate the effect of CIBA-contaminated *Planck* extinction maps through standard galaxy clustering analyses.

Ideally, we would like not only to quantify the effect of the CIBA on cosmological measurements, but rather remove the CIBA imprint from our dust map entirely, on a per line-of-sight basis. Unfortunately, it is difficult to predict the detailed CIBA spatial structure based on presently cataloged three dimensional positions of galaxies, as typical redshifts in existing spectroscopic galaxy surveys are significantly lower than those of the dominant CIBA emitters (Fernandez-Conde et al., 2008). The forthcoming Dark Energy Spectroscopic Instrument project (DESI, Levi et al., 2013) may enable more detailed modeling of the CIBA based on the three dimensional distribution of galaxies by extending to higher redshifts than existing wide-field spectroscopic surveys such as BOSS (Dawson et al., 2013). In the coming years, I plan to work on the Dark Energy Camera Legacy Survey (DECaLS), a deep optical imaging survey which will inform DESI target selection.

On large angular scales, the zodiacal light has long been the bane of dust maps based on infrared emission. Most recently, we have shown that despite *Planck*'s broad spectral coverage down to millimeter wavelengths, the zodiacal light still propagates into *Planck*-based extinction estimates and represents the dominant large angular scale problem in the Meisner & Finkbeiner (2015) reddening predictions at high Galactic latitude. It remains to be seen which if any future infrared data set (or re-analysis of existing data sets) might unlock a breakthrough in eliminating zodiacal light residuals from Galactic dust emission maps. However, it may be possible to circumvent the zodiacal light by tying the large angular scale modes of infrared-based extinction maps to maps of reddening based directly on optical observations of stars (e.g. Schlafly et al., 2014), which are wholly unaffected by zodiacal emission.


Notwithstanding current imperfections in our *Planck*-based thermal dust model, we are already able to pursue a number of promising applications. In one highlight, we are using the Meisner & Finkbeiner (2015) emission predictions to calibrate the zero

levels of Herschel mosaics of Gould Belt molecular clouds including Ophiuchus, Perseus and Serpens (André et al., 2010; Ridge et al., 2006). Our two-component predictions allow us to replace the large angular scale modes in Herschel with those from *Planck*. This technique eliminates the need for sub-optimal procedures which have historically involved assuming that a constant additive offset can rectify the Herschel zero level over large regions tens of square degrees in size. Using our *Planck*-based zero level technique, we have already obtained stunning 36'' resolution maps of dust temperature and optical depth using Herschel PACS/SPIRE 100 μ m-500 μ m data (see Figure 4.1).


In order to improve our understanding of the interstellar medium, it will likely prove worthwhile to consider more physically motivated models of dust emission and extinction than those adopted in e.g. Meisner & Finkbeiner (2015) and Planck Collaboration et al. (2013k). The fact that the Finkbeiner et al. (1999) two-MBB model fits the unpolarized dust SED at the 2% level or better from 100-3000 GHz (Meisner & Finkbeiner, 2015) in the diffuse ISM is quite remarkable, given the model's simplicity. Even so, it is already clear that the simple two-MBB model cannot fully explain the relatively complex frequency dependence of dust polarization observed by *Planck* (Planck Collaboration et al., 2014a). This points to the need for a more ambitious effort toward simultaneously fitting physical models of dust grain emission and extinction to all available data, including optical reddening (e.g. from Pan-STARRS1) and *polarized* dust emission from *Planck*. Within this rephrasing of the problem at hand, it will likely become preferable to fit a spatially varying ISRF intensity and spectrum rather than a dust temperature.

Additionally, no evidence yet exists to support the hypothesis that our two MBB components in fact correspond to two physically distinct grain species inhabiting the ISM. Identifying spatial variations of dust grain composition using an ISM tracer other


Planck-based temperature, beta=1.71 MBB



Herschel-based temperature, beta=1.71 MBB



Planck-based 545 GHz optical depth



Herschel-based 545 GHz optical depth




Figure 4.1 A comparison of *Planck*-based and *Herschel*-based dust SED fit results for the Ophiuchus region. Prior to SED fitting, we have rectified the *Herschel* emission mosaic zero levels by layering the small-scale modes from *Herschel* on top of 6.1' FWHM *Planck*-based predictions of the *Herschel*-observed emission. The *Herschel*-based SED parameters have 36'' FWHM, limited by *Herschel* 500 μ m native angular resolution.

than FIR emission could provide a first test of this hypothesis. For example, with Pan-STARRS1 (Kaiser et al., 2002), it is possible to map the extinction curve shape (parametrized by R_V) and thus check whether inferences about the spatial variation of dust composition from optical extinction and FIR emission are in agreement.

Of course, any discussion of pushing towards a more physical understanding of the ISM must culminate in acknowledging that such an effort, in principle, is much more appropriately undertaken in three spatial dimensions rather than two projected dimensions. Indeed, the interplay between high-resolution dust maps in two dimensions and three-dimensional reddening maps (e.g. Green et al., 2014) presents an exciting array of future research directions. As already mentioned, star-based dust maps may provide a natural way to circumvent the zodiacal light contamination which is so difficult to remove from infrared-based dust maps. Further, optical star-based reddening maps, which typically have angular resolution comparable to that of *IRAS*, may represent the most natural baseline model onto which we can layer the small-angular scale detail available from *WISE*. In the long term, the goal of merging infrared dust mapping efforts with optical dust mapping in three-dimensions will be to simultaneously constrain the three dimensional positions of stars, the six dimensional ISRF, the three dimensional distribution of both large dust grains and PAHs, the Galactic magnetic field and the spatial variation of the extinction law.

Returning to the more concrete goal of a next-generation dust extinction map in two dimensions, it seems clear that we can construct an aesthetically pleasing, high-resolution dust model that looks visually plausible based on *WISE* and *Planck*. But in order for the resulting dust data product to become widely adopted, we must convincingly *quantify* the advantages of estimating extinction with enhanced angular resolution. The subtlety of this task stems from the fact that, for typical photometric measurements, the observational statistical uncertainties are much larger than the

systematic errors due to reddening corrections with sub-optimal angular resolution. Naturally, we will circumvent this issue at the computational cost of combining information from hundreds of millions of optical sources. This will undoubtedly involve the use of large databases of optical photometry and spectroscopy, drawn from Pan-STARRS1, SDSS, BOSS, DES, DESI, GAIA, and LSST among other existing and future surveys.

The demand for accurate cosmological foreground models seems poised to remain a critical and central theme of astronomy for the foreseeable future, especially in light of recent developments in B-mode cosmology (e.g Ade et al., 2014; BICEP2/Keck et al., 2015). The pursuit of ever-refined interstellar dust models is among a number of astrophysics research topics pushing astronomers to devote their careers to confronting and embracing data-intensive computational challenges.

4.1 Funding Acknowledgments

D.F. and A.M. acknowledge support of NASA grant NNX12AE08G for this research. This research was conducted with Government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a. We gratefully acknowledge support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE1144152.

This research made use of the NASA Astrophysics Data System (ADS) and the IDL Astronomy User’s Library at Goddard.

This research makes use of data products from the Wide-field Infrared Survey Explorer, which is a joint project of the University of California, Los Angeles, and the Jet Propulsion Laboratory/California Institute of Technology, funded by the National

Aeronautics and Space Administration. This publication also makes use of data products from NEOWISE, which is a project of the Jet Propulsion Laboratory/California Institute of Technology, funded by the Planetary Science Division of the National Aeronautics and Space Administration. This publication makes use of data products from the Two Micron All Sky Survey, which is a joint project of the University of Massachusetts and the Infrared Processing and Analysis Center/California Institute of Technology, funded by the National Aeronautics and Space Administration and the National Science Foundation.

Based on observations obtained with Planck (<http://www.esa.int/Planck>), an ESA science mission with instruments and contributions directly funded by ESA Member States, NASA, and Canada.

References

- Ade, P., et al. 2014, *Phys.Rev.Lett.*, 112, 241101
- André, P., Men'shchikov, A., Bontemps, S., et al. 2010, *A&A*, 518, L102
- Andriesse, C. D. 1978, *A&A*, 66, 169
- Barnard, E. E. 1919, *ApJ*, 49, 1
- BICEP2/Keck, Planck Collaborations, :, et al. 2015, *ArXiv e-prints*
- Boggess, N. W., Mather, J. C., Weiss, R., et al. 1992, *ApJ*, 397, 420
- Burstein, D., & Heiles, C. 1978, *ApJ*, 225, 40
- . 1982, *AJ*, 87, 1165
- Corwin, Jr., H. G., Buta, R. J., & de Vaucouleurs, G. 1994, *AJ*, 108, 2128
- Cox, P., & Mezger, P. G. 1989, *A&A Rev.*, 1, 49
- Cutri, R. M., Wright, E. L., Conrow, T., et al. 2012, Explanatory Supplement to the WISE All-Sky Data Release Products,
<http://wise2.ipac.caltech.edu/docs/release/allsky/expsup/>
- Dawson, K. S., Schlegel, D. J., Ahn, C. P., et al. 2013, *AJ*, 145, 10
- de Vaucouleurs, G., & Malik, G. M. 1969, *MNRAS*, 142, 387
- Doi, Y., Komugi, S., Kawada, M., et al. 2012, *Publication of Korean Astronomical Society*, 27, 111
- Doi, Y., Takita, S., Ootsubo, T., et al. 2015, *ArXiv e-prints*
- Draine, B. T., & Li, A. 2007, *ApJ*, 657, 810
- Fernandez-Conde, N., Lagache, G., Puget, J.-L., & Dole, H. 2008, *A&A*, 481, 885
- Finkbeiner, D. P., Davis, M., & Schlegel, D. J. 1999, *ApJ*, 524, 867
- Górski, K. M., Hivon, E., Banday, A. J., et al. 2005a, *ApJ*, 622, 759

- . 2005b, *ApJ*, 622, 759
- Green, G. M., Schlafly, E. F., Finkbeiner, D. P., et al. 2014, *ApJ*, 783, 114
- Hall, J. S. 1937, *ApJ*, 85, 145
- Hauser, M. G., & Dwek, E. 2001, *ARA&A*, 39, 249
- Herschel, W. 1785, *Royal Society of London Philosophical Transactions Series I*, 75, 213
- Hubble, E. 1934, *ApJ*, 79, 8
- Kaiser, N., Aussel, H., Burke, B. E., et al. 2002, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 4836, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, ed. J. A. Tyson & S. Wolff, 154–164
- Kalberla, P. M. W., Burton, W. B., Hartmann, D., et al. 2005, *A&A*, 440, 775
- Kapteyn, J. C. 1909, *ApJ*, 29, 46
- Kelsall, T., Weiland, J. L., Franz, B. A., et al. 1998, *ApJ*, 508, 44
- Kessler, M. F., Steinz, J. A., Anderegg, M. E., et al. 1996, *A&A*, 315, L27
- Lee, Y. S., Beers, T. C., Sivarani, T., et al. 2008, *AJ*, 136, 2022
- Levi, M., Bebek, C., Beers, T., et al. 2013, *ArXiv e-prints*
- Low, F. J., Rieke, G. H., & Gehrz, R. D. 2007, *ARA&A*, 45, 43
- Low, F. J., Young, E., Beintema, D. A., et al. 1984, *ApJ*, 278, L19
- Mainzer, A., Bauer, J., Grav, T., et al. 2011, *ApJ*, 731, 53
- Mather, J. C. 1982, *Optical Engineering*, 21, 769
- Meisner, A. M., & Finkbeiner, D. P. 2014, *ApJ*, 781, 5
- . 2015, *ApJ*, 798, 88
- Onaka, T., Yamamura, I., Tanabe, T., Roellig, T. L., & Yuen, L. 1996, *PASJ*, 48, L59
- Paturel, G., Petit, C., Prugniel, P., et al. 2003, *A&A*, 412, 45
- Pilbratt, G. L., Riedinger, J. R., Passvogel, T., et al. 2010, *A&A*, 518, L1
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2011a, *A&A*, 536, A19

- . 2011b, *A&A*, 536, A18
- . 2013a, ArXiv e-prints
- . 2013b, ArXiv e-prints
- . 2013c, ArXiv e-prints
- . 2013d, ArXiv e-prints
- . 2013e, ArXiv e-prints
- . 2013f, ArXiv e-prints
- . 2013g, ArXiv e-prints
- . 2013h, ArXiv e-prints
- . 2013i, ArXiv e-prints
- . 2013j, ArXiv e-prints
- Planck Collaboration, Abergel, A., Ade, P. A. R., et al. 2013k, ArXiv e-prints
- Planck Collaboration, Ade, P. A. R., Alves, M. I. R., et al. 2014a, ArXiv e-prints
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014b, ArXiv e-prints
- Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2015, ArXiv e-prints
- Price, S. D. 1981, *AJ*, 86, 193
- Reach, W. T., Dwek, E., Fixsen, D. J., et al. 1995, *ApJ*, 451, 188
- Ridge, N. A., Di Francesco, J., Kirk, H., et al. 2006, *AJ*, 131, 2921
- Rudnick, J. 1936, *ApJ*, 83, 394
- Sale, S. E., Drew, J. E., Unruh, Y. C., et al. 2009, *MNRAS*, 392, 497
- Schlaflly, E., Green, G., Finkbeiner, D., et al. 2014, *ApJ*, 789, 15
- Schlaflly, E. F., & Finkbeiner, D. P. 2011, *ApJ*, 737, 103
- Schlegel, D. J., Finkbeiner, D. P., & Davis, M. 1998, *ApJ*, 500, 525
- Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, *AJ*, 131, 1163
- Stebbins, J., Huffer, C. M., & Whitford, A. E. 1939, *ApJ*, 90, 209
- Tanaka, M., Matsumoto, T., Murakami, H., et al. 1996, *PASJ*, 48, L53

Trumpler, R. J. 1930, PASP, 42, 214

Werner, M. W., Roellig, T. L., Low, F. J., et al. 2004, ApJS, 154, 1

Wheelock, S. L., Gautier, T. N., Chillemi, J., et al. 1994, NASA STI/Recon Technical Report N, 95, 22539

Wright, E. L., Eisenhardt, P. R. M., Mainzer, A. K., et al. 2010, AJ, 140, 1868

York, D. G., Adelman, J., Anderson, Jr., J. E., et al. 2000, AJ, 120, 1579