DEVELOPMENT OF THE TOLTEC DATA REDUCTION PIPELINE AND THE APPLICATION OF HIERARCHICAL BAYESIAN INFERENCE TO TOLTEC DATA

A Dissertation Presented

by

MICHAEL. J. MCCRACKAN

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

May 2024

Department of Astronomy

DEVELOPMENT OF THE TOLTEC DATA REDUCTION PIPELINE AND THE APPLICATION OF HIERARCHICAL BAYESIAN INFERENCE TO TOLTEC DATA

A Dissertation Presented

by

MICHAEL. J. MCCRACKAN

Approved as to style and content by:

Grant. W. Wilson, Chair

Daniela Calzetti, Member

Mark Heyer, Member

Stephane Willocq, Member

Grant W. Wilson, Graduate Program Director
Department of Astronomy

ACKNOWLEDGMENTS

This work would not have been possible without the contributions of many individuals.

Firstly, I wish to express my deepest gratitude to my advisor, Grant Wilson, for his steadfast support, patience, and invaluable mentorship throughout my dissertation research. Additionally, I extend my thanks to my thesis committee members, Daniela Calzetti, Mark Heyer, and Stephane Willocq, for their insightful feedback and guidance. My appreciation also goes to my first and second-year advisors, Min S. Yun, Martin Weinberg, and Neal Katz who introduced me to new subfields and aided me in developing a wide range of skills. Special thanks to Peter Schloerb for offering me the chance to apply my body of work in a completely different area of astrophysics as well as for helpful discussions related to mapmaking.

I am indebted to the TolTEC instrument and software team at UMass, including Zhiyuan Ma, Kamal Souccar, Nat S. DeNigris, and Stephen Kuczarski, for their dedication to building, installing, and testing the TolTEC instrument and its associated software. A special mention to Zhiyuan Ma for his considerable help with the TolTEC data reduction pipeline and for fielding numerous software-related questions.

My thanks to Caleigh Ryan and Amanda Lee for their assistance in identifying commissioning targets of nearby galaxies, refining observing strategies, and testing software. Also, my appreciation to Yuping Tang for sharing his fitting code and providing insight into topics related to hierarchical Bayesian fitting.

Further, I thank the entire TolTEC collaboration for their analysis of TolTEC data, as well as their testing of and recommendations for the TolTEC software I wrote, both of which have been crucial to the development and refinement of my work.

I am grateful to Joseph E. Golec from the Department of Physics at the University of Chicago for creating the Minkasi maps of the Crab Nebula that were instrumental in my research. Appreciation is also due to Artyom Tanashkin and Yuri Shibanov from the Ioffe Physical-Technical Institute of the Russian Academy of Sciences and Aida Kirichenko from the Universidad Nacional Autónoma de México for their aperture photometry analysis on the Crab Nebula maps. Thank you to Robert Gutermuth at UMass for analysis and ancillary data products related to the Monoceros R2 molecular cloud. Additional thanks to Dennis Lee for figures related to the TolTEC half-wave plate.

The kindness and assistance of my fellow students and the faculty at the Department of Astronomy at UMass have been invaluable throughout my time here. I would like to give special acknowledgment to my cohort members Zhiyuan Ji, Patrick Kamienski, and Alyssa Sokol for their camaraderie and support.

Lastly, I acknowledge the financial support granted by the Massachusetts Space Grant Consortium, which was essential to my research. The TolTEC project was funded through the NSF Grant 1636621.

ABSTRACT

DEVELOPMENT OF THE TOLTEC DATA REDUCTION PIPELINE AND THE APPLICATION OF HIERARCHICAL BAYESIAN INFERENCE TO TOLTEC DATA

MAY 2024

MICHAEL. J. MCCRACKAN

B.Sc., UNIVERSITY OF NEW MEXICO

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Grant. W. Wilson

TolTEC is a millimeter-wavelength imaging polarimeter now installed on the 50-meter Large Millimeter Telescope that simultaneously maps the sky using 7718 dual-polarization Lumped Element Kinetic Inductance Detectors distributed among three monochromatic arrays, centered at 1.1, 1.4, and 2.0 mm (273, 214, and 150 GHz). The camera is currently in the commissioning phase and has completed two observational runs, in June and December of 2022.

This work offers a comprehensive review of the TolTEC data reduction and map-making pipeline `Citlali` (v4.0), an open-source, high-performance, and parallelized software package written in C++. `Citlali` rapidly transforms the raw time-ordered data from all categories of TolTEC data into two-dimensional maps of the sky, in addition to performing map coaddition and post-mapmaking point source filtering. The

pipeline's design philosophy, data streaming and parallelization model, timestream reduction stages, mapmaking algorithms, and iterative mapmaking routine are detailed.

Maps of sources observed during TolTEC's 2022 commissioning, including the radio quasar J1159+292, the Crab Nebula, and the Monoceros R2 Giant Molecular Cloud, which were produced using `Citlali`, are presented. The analysis investigates the flux recovery from extended sources by `Citlali`'s iterative mapmaker and compares results from the built-in mapmakers to maps created with the maximum likelihood mapmaker Minkasi.

This work also details a C++ hierarchical Bayesian MCMC software package developed for fitting dust emission SEDs in each pixel of TolTEC maps. This code integrates instrumental PSF data into a forward-fitting model to maintain contributions from higher-resolution observations within the dataset. Both modified blackbody and physically motivated dust models using the Astrodust+PAH model of Hensley and Draine 2023 have been implemented. Results from applying this software to simulated dust SEDs, as well as to WISE, *Spitzer*, and *Herschel* observations of the face-on spiral galaxy NGC 3938, are presented.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

INTRODUCTION

The sky at millimeter wavelengths is host to a wide diversity of astrophysical sources that exist across much of the lifetime of the Universe and are characterized by an enormous range in physical scales and luminosities. The interplay between the photons of the Cosmic Microwave Background (CMB) and the hot gas between galaxies within massive clusters through the Sunyaev–Zeldovich (SZ; Sunyaev and Zeldovich 1970; Sunyaev and Zeldovich 1972) effect, spectral emission lines from molecules deep within cold, dense star-forming clouds and protoplanetary disks, synchrotron radiation from supernovae remnants, and the thermal emission from interstellar dust all fill the sky with millimeter wavelength photons that offer a unique, yet complementary window with which to explore fundamental questions about physics of the Cosmos.

Fueled by the parallel developments of new detector technologies that have enabled the assembly of larger, multi-color detector arrays and the construction of larger single-dish telescopes, interferometer arrays, and airborne experiments, the field of both ground-based and near-Earth orbit submillimeter and millimeter astronomy has been at the center of an explosion in new studies and scientific results in previously poorly explored parameter spaces over the decades. First generation cameras like SCUBA (Holland et al. 1999), LABOCA (Siringo et al. 2009), AzTEC (Wilson et al. 2008) and NIKA (Monfardini et al. 2010) mounted on 15–50-meter telescopes such as the JCMT, IRAM, and the LMT laid much of the groundwork for recent second-generation instruments like SCUBA-2 (Holland et al. 2013) and NIKA2 (Adam et al. 2018) which can now observe the sky at resolutions near $10''$ with much higher sensitivities. This has opened new windows into the small-scale physics within the Milky Way and in nearby galaxies, while simultaneously allowing for studies of faint sources at high redshifts in a fraction of the time of what was previously possible. These cameras are complemented by large field-of-view (FOV) experiments like the Atacama Cosmology Telescope (ACT; Thornton et al. 2016), the South Pole Telescope (SPT; Carlstrom et al. 2011), and CCAT-Prime (Chapman et al. 2022) for investigations into the Large-Scale structure (LSS) of the Universe and the polarization features of the CMB. Upcoming and in-development projects like the Advanced Simons Observatory (SO; Galitzki et al. 2018) and CMB-S4 (Carlstrom et al. 2019) use detector arrays an order of magnitude larger than current instruments and provide sufficient sensitivity to place strong constraints on CMB and fundamental cosmological physics. Airborne instruments like BLASTPol (Pascale and Pascale 2013), SPIDER (Crill et al. 2008), and SOFIA (Helton and Team 2013) observe the signatures of star-formation, magnetic fields, and cosmological B-mode polarizations from high altitudes above much of the Earth's intervening atmosphere which significantly impacts ground-based millimeter observing. Finally, interferometers, particularly the Atacama Large Millimeter/submillimeter Array (ALMA), offer the opportunity to investigate such targets as dust and molecular emission lines within protoplanetary disks and distant galaxies with angular scales of less than 1 arcsecond.

The new ground gained at millimeter wavelengths joins with earlier and ongoing transformational advancements in the fields of infrared and submillimeter astronomy that was propelled forwards by the launch of space-based telescopes, namely *Spitzer* Space Telescope (Werner et al. 2004), WISE (Wright et al. 2010), AKARI (Murakami et al. 2007), the *Herschel* Space Observatory (Pilbratt et al. 2010), and *Planck* (Clements 2017). The James Webb Space Telescope (JWST; Clements 2017) joins this pantheon of space telescopes and offers superior sensitivity and resolution to near-infrared (NIR) and mid-infrared (MIR) wavelengths compared to *Spitzer*. These instruments have conducted large surveys of galactic and extragalactic sources over their lifetimes thus building homogeneous photometric and spectroscopic multiwavelength datasets that can be combined with millimeter observations to place strong, statistically significant constraints on astrophysical phenomenon from just a few microns up to 10 millimeters.

Out of the vast selection of sources that populate the millimeter sky, dust represents a particularly important component and is a primary target of submillimeter and millimeter instruments owing to the ubiquity of its presence in galactic and extragalactic environments and due to its profound effect on galaxies from the microscopic to the macroscopic scales. Dust is a complex building block of the interstellar medium (ISM), being characterized by a range of dust grain sizes, shapes, and compositions, each of which perturbs the optical properties of the global dust population. It is a key player in the star-formation cycle alongside the gas content and stellar populations and is therefore a fundamental driver in galaxy evolution over cosmic time (Galliano et al. 2018). The presence of dust does not only affect the physical nature of galaxies and their interstellar media, but it also significantly impacts our ability to probe their properties through observations as it reshapes the spectral energy distributions (SEDs) of galaxies (Draine 2003; Galliano 2022). It simultaneously obscures processes from direct investigation while also providing new, indirect tracers of their occurrence. Therefore, in order to measure key parameters that describe the structure and evolution of galaxies like star-formation rates (SFR), stellar masses, metallicities, and morphologies, in addition to enable the exploration of the relationships among them, a comprehensive understanding of dust content in galaxies is required. While shorter wavelengths are most affected by dust extinction, the emission features of dust lie squarely within the NIR to millimeter range. The Milky Way dust content has been thoroughly explored and is used as the basis for the construction of models for extragalactic studies. The mass, distribution, and spectral shape of dust within a diverse array of galaxy types, including nearby star-forming spirals, local dwarfs, as well as in Luminous Infrared Galaxies (LIRGs) and Ultraluminous Infrared Galaxies

(ULIRGs) have been extensively explored through large surveys conducted with *Herschel* and *Spitzer* and smaller scale studies with ground-based telescopes (Galliano et al. 2018; Galliano 2022).

These studies have revealed the considerable diversity in dust physics and the environments it is found in, even when only considering local galaxies. In particular, the dust content within nearby dwarf galaxies presents an array of unique and perplexing traits not found within the Milky Way or similar massive, nearby star-forming spirals. The dwarf galaxy classification is a moniker belonging to a varied assortment of galaxy sub-types, though they share the common features of having low stellar masses and luminosities, small physical sizes, and lacking much in the way of heavier elements or metals in their ISMs. Sub-types include dwarf spheroids, ellipticals, irregulars, Blue Compact Dwarfs (BCD) (Sargent and Searle 1970), ultrafaint dwarfs (Simon 2019a), and ultra-compact dwarfs (Hilker et al. 1999; Drinkwater et al. 2000). Dwarf spheroids are the most common type, as well as the smallest and faintest of the dwarfs, with stellar masses less than 10^8 $M_\odot$ and diameters of less than 0.5 kpc (Mateo 1998). They exhibit old stellar populations of 10-12 Gyr and have no ongoing star-formation activity within the last 1-2 Gyrs (Tolstoy et al. 2009). Ultra-faint dwarfs represent an extension of the dwarf spheroid classification to lower luminosities of less than $\mathrm{M}_V = -7.7$ or 10^5 $\mathrm{L}_\odot$ (Simon 2019b). Dwarf ellipticals are the larger and more massive counterparts of the dwarf spheroids, with diameters less than 10 kpc while similarly showing no evidence of ongoing star formation. Dwarf irregulars, like the more general irregular galaxy classification, possess no well-defined shape, contain substantial amounts of gas, especially in the neutral HI phase, and have ongoing star-formation. Many dwarf irregulars and BCDs exhibit signs of undergoing intense periods of starburst activity (Tolstoy et al. 2009; Henkel et al. 2022).

The amount of dust relative to gas in any of the dwarf sub-categories can be more than an order of magnitude lower than what one would naively expect based on the linear trend with metallicity for massive galaxies (Rémy-Ruyer et al. 2014; Galliano et al. 2021). Their dust is typically hotter and has little in the way of the emission features that are more commonly found at NIR to MIR wavelengths in star-forming spirals (Galliano et al. 2005; Rémy-Ruyer et al. 2013). Finally, a significant fraction of SEDs of dwarf galaxies show an enhancement in emission at submillimeter and millimeter wavelengths that is not commonly mirrored by massive galaxies (Rémy-Ruyer et al. 2013; Gordon et al. 2014; Galliano et al. 2011; Galliano 2022). These distinctions point to unique dust grain production, growth, and destruction mechanisms within dwarf galaxy ISMs.

The conditions in the ISMs of dwarf galaxies are thought to be like those of the first galaxies in the Universe (Henkel et al. 2022). Under the theory of hierarchical structure formation within the ΛCDM (White and Rees 1978; Frenk et al. 1988) cosmological framework, primordial galaxies should form within the smallest dark matter halos and build mass to become the massive galaxies we see today through mergers over their lifetimes. Processes such as galaxy downsizing complicate our interpretation of galaxy evolution, with stellar mass measurements out to z=4 (Pérez-González et al. 2008) and star-formation time-scale estimates finding that massive early type galaxies can form earlier and have older stellar populations (Thomas et al. 2010) relative to their smaller counterparts (Silk et al. 2014). However, downsizing does not necessarily contradict a bottom-up formation scenario. It is likely a consequence of massive halos' inability to retain cold gas (Cattaneo et al. 2008), which can be attributed to feedback mechanisms and interactions within dense environments.

The earliest galaxies would not have had sufficient generations of stars to enrich their ISMs with heavier elements. The chemical abundances of elements such as deuterium, helium, and lithium will have undergone little evolution from the initial abundances set by Big Bang nucleosynthesis (Henkel et al. 2022). Nearby dwarf galaxies can be utilized as analogs to these faint, and difficult-to-observe high redshift systems, where the interactions between the dust, gas, and stellar content inside of an extreme environment can be probed down to the scale of individual star-forming complexes. Indeed, blue compact dwarfs, despite undergoing repeated strong starburst phases, are old galaxies with metal abundances preserved from the early universe and can be used to explore how star formation may have occurred in early systems (Cairós et al. 2010).

Acquiring accurate estimates of the total dust masses and improving constraints on the nature of the anomalous excess emission from dust constitute strong motivations for carrying out observations of dwarf galaxies with high-resolution ground-based millimeter cameras, which can map the distributions of excess emission and connect them to the corresponding distributions of key dust parameters within dwarf galaxy ISMs. With strong constraints on the dust SED shape from combinations of space-based submillimeter and ground-based millimeter telescopes, differing dust populations and characteristics can be introduced into SED models to derive a physical interpretation for the unique features of dwarf dust emission. In practice, this proves to be challenging owing to the faint nature of dwarfs across FIR and millimeter bands which makes detections at high significance with earlier millimeter instruments difficult. This necessitated observations of only the most nearby sources, such as the Magellanic Clouds and Local Group dwarfs, or very bright, starburst dwarfs in order to avoid prohibitively long integration times (Simon 2019a). This selection restricts the range of conditions that can be explored, with the dust properties of very metal poor dwarfs with little active star-formation being underexplored. As a result, dwarf galaxy dust content remains a poorly sampled parameter space compared to that of star-forming galaxies and represents a unique opportunity for novel science results with current generation, high resolution, and high sensitivity ground-based millimeter instruments.

The TolTEC camera (Bryan et al. 2018; Wilson et al. 2020) is a millimeter imaging camera that maps the sky at 1.1, 1.4, and 2.0 mm (273, 214, 150 GHz) simultaneously using 7718 polarization sensitive Kinetic Inductance Detectors (KIDs; Day et al. 2003; Doyle et al. 2008). TolTEC is a facility instrument currently undergoing commissioning on the 50-m Large Millimeter Telescope (LMT; Hughes et al. 2020) which provides the camera with a 4' diameter field-of-view and resolutions of 5-10$''$ across its 3 bands. Observations with TolTEC aim to tackle many of the key areas of interest in the field of millimeter astronomy, including making measurements the SZ effect, estimating the polarized emission from star-forming filaments, detecting protostellar cores in molecular clouds, and carrying out spatially resolved studies of the dust content in local galaxies. Furthermore, TolTEC, in conjunction with the LMT, represents an optimal system with which to address the difficulties in investigating nearby dwarf dust content. Observations of bright, nearby, well-studied starburst dwarfs such as NGC 4449 and IC 10 are currently planned and TolTEC is expected to carry out observations of tens of local dwarfs over its lifetime.

THIS THESIS

This thesis details my contributions to the construction, installation, and commissioning of the TolTEC millimeter camera and to the TolTEC project as a whole. My primary role within the TolTEC collaboration is as a member of the software team, where I am the lead and primary author of the TolTEC data reduction and map-making pipeline, `Citlali`. The design, development, deployment, and integration of the pipeline into the broader TolTEC software architecture, and the verification of `Citlali` with raw and simulated TolTEC data constitute the central pillar of my dissertation work. I have also developed a hierarchical Bayesian SED fitting package intended primarily to be used for fitting dust emission within planned observations of nearby star-forming and dwarf galaxies with TolTEC.

Chapter 1 describes the scientific motivation for my work by providing an overview of the properties of the dust content in nearby galaxies, two frequently used dust modeling prescriptions, and a description of the unique qualities of the dust within local low-metallicity dwarf galaxies. An introduction to the TolTEC camera, including its design, setup, and observing strategies, as well as the Large Millimeter Telescope and supporting data reduction hardware is the subject of Chapter 2 to provide context for the requirements of the camera's data reduction pipeline. `Citlali` is described in Chapter 3, where the code structure, core algorithms, integration with other mapmakers, and performance are outlined. Chapter 4 presents preliminary results from the reduction of TolTEC commissioning observation data with `Citlali` and other map-making pipelines. The hierarchical Bayesian SED fitting software code I developed is then detailed in Chapter 5. Results from the fitting of simulated dust SEDs to modified blackbody and physically motivated dust models are presented. Furthermore, fits to integrated aperture photometry and individual pixels from the *Spitzer* and *Herschel* data of the nearby face-on spiral galaxy NGC 3938 are presented. Conclusions are then given in Chapter 6.

CHAPTER 1

DUST CONTENT IN THE LOCAL UNIVERSE

Dust is a key building block of galaxies and plays an important role in their evolution over cosmic time. The interstellar medium exists as a complex mixture of both dust and gas in its molecular, atomic, and ionized phases and its interactions with the stellar populations within galaxies results in the star-formation cycle that drives the dynamical, morphological, and chemical transformation of galaxies in the absence of major disruptive events like interactions and mergers. Within nearby galaxies, the mass abundance of dust is dwarfed by that of the gas content, with only about 1% of the ISM mass being accounted for by dust (Bohlin et al. 1978; Draine et al. 2007; Sandstrom et al. 2012; De Vis et al. 2019). For local dwarf galaxies, the mass fraction can be several orders of magnitude lower (Rémy-Ruyer et al. 2014; De Vis et al. 2017; Galliano et al. 2021). Despite its relative scarcity, the presence of dust cannot be overlooked due to its effects on the physical makeup of galaxies and their radiative properties. It is not an overstatement to say that without dust, the Universe and the galaxies that populate it would be nigh unrecognizable to their present forms.

The dust content in galaxies is an assembly of individual dust grains that form and grow from the heavier elements that flood the ISM during the late-stage processes in the lifecycles of stars (Draine 2003; Galliano 2022). Dust grains are characterized by a considerable complexity in their properties which translates to a similar complexity in their emission features, as well as in their absorption and scattering of radiation from other sources. Dust absorbs and scatters stellar radiation emitted at UV and optical parts of the electromagnetic spectrum and re-radiates it as thermal emission and spectral lines between the NIR to millimeter windows, thus significantly reshaping the SEDs of galaxies (Driver et al. 2016). Figure 1.1 shows several examples of SEDs for different galaxy types, highlighting the impact of the dust thermal emission as well as the dust spectral features at NIR and MIR wavelengths. The relative contribution of the dust thermal emission to the bolometric luminosity of the galaxy increases from bottom to top. Dust is therefore a contributing factor to nearly every measurement that we can make of galaxies. Nearly 30% of the stellar luminosity from typical star-forming galaxies is re-radiated as dust thermal emission (Draine 2003; Galliano 2022), with this number climbing to 99% for ULIRGs at higher redshifts (Bianchi et al. 2018). As dust is co-spatial with the gas content in the ISM, dense, cold molecular clouds can be deeply embedded within dense, dusty regions that obscure direct emission from the protostellar objects and newborn stars within them.

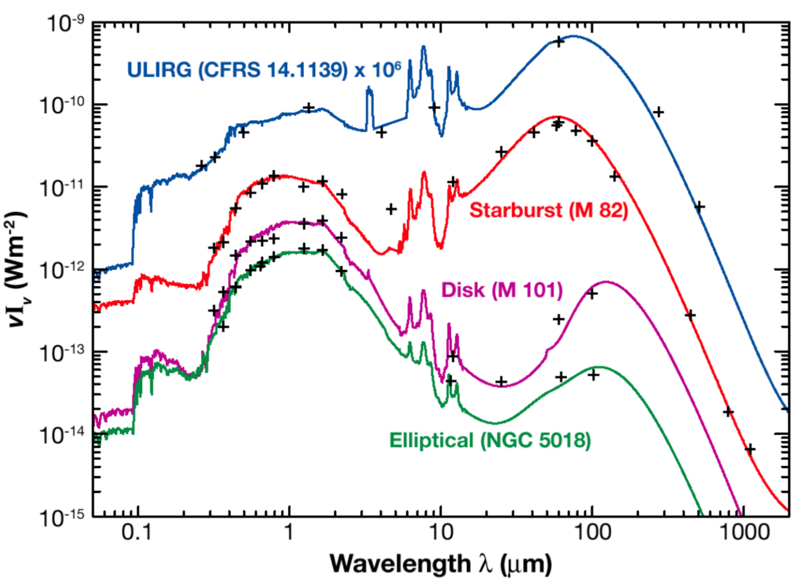

Figure 1.1. Figure reproduced from Lagache et al. 2005. Galaxy SEDs from several different galaxy types, including a local star-forming spiral, an early-type galaxy, a starburst system, and a ULIRG, are plotted.

While dust obscuration and reprocessing may be a considerable nuisance at times, it does offer a direct window through which to investigate the physics of dust grains themselves and the physical conditions of the environment they populate. The wavelength dependence of a dust grain's re-radiation of absorbed photons as thermal radiation is influenced by both the physical environment of the ISM that the dust grain inhabits and the structure of the grain itself. Grains are heated by the photons of the interstellar radiation field (ISRF) and will emit at different wavelengths depending on the strength of the local and diffuse radiation field components. Each grain can be described in terms of an absorption cross-section which is a function of the grain size and its dielectric constant. As a result, the peak and spectral shape of dust thermal emission SEDs are signatures of local ISM temperature and the types of grains within the emitting population (Hensley and Draine 2023). Dust spectroscopic features at MIR wavelengths arise from electronic, vibrational, and rotational transitions that depend on the chemical makeup and structure of the grain (Draine 2003). Using observational constraints on the wavelength-dependent absorption and scattering by dust, measurements of dust analogs in laboratory environments, and modeling of grain size distributions, shapes, and assembly, a model of the interstellar dust population in differing environments can be constructed and used to extract information about the dust content from galaxy SEDs.

This of course is no simple nor straightforward feat. Complications arise from the fact that the actual three-dimensional structure of the ISM is difficult to ascertain from the two-dimensional projections provided by observations. With the exception of the very diffuse ISM, different lines of sight are likely to be characterized by a mixture of temperatures and dust grain species that contribute to an average SED described by effective dust and ISM properties rather than those of any single actual component. Galaxies are also viewed at different inclination angles which can impact the recovery of key parameters like metallicity due to differences in dust attenuation.

Three-dimensional radiative transfer (RT) methods provide an avenue to explore the dust emission and extinction and the effects of projection and inclination for simulated 3D distributions of mixtures of dust grain types, though accurately modeling complex structures is difficult (Popescu et al. 2000; Galliano 2022). In addition, while laboratory-based dust analogs are used to infer the compositions of interstellar grains, exact recreations of the conditions in which the latter forms and grows is challenging, if not impossible. Therefore, while general characteristics may be deduced, their chemical and structural makeup cannot be fully determined from analogs alone and they must be used in conjunction with independent observational constraints.

1.0.1 What are Dust Grains Made Of?

Dust grains are solid structures that are assembled from many of the most common metals found within the ISM, namely carbon (C), oxygen (O), magnesium (Mg), silicon (Si), and iron (Fe) (Siebenmorgen et al. 2014). The presence of these elements within dust grains in the Milky Way (MW) can be inferred through comparisons of solar photosphere elemental abundances with those of the gas in the ISM (Asplund et al. 2009), through absorption spectroscopy (Jenkins 2009), and dust polarization spectra (Siebenmorgen et al. 2014). These studies find, with exceptions, that many dust grain candidate elements are less abundant or depleted in the ISM, particularly in dense ISM regions where grain growth from raw materials in the gas is expected to be more efficient. Our understanding of the chemical composition of dust grains is further informed by the spectroscopic features exhibited by different dust grain species. The 2175 Å bump, prominent emission lines at 9.4 µm and 18 µm (Woolf and Ney 1969; Breemen et al. 2011), the 3.4 µm and 6.85 µm absorption features (Tielens et al. 1996; Chiar et al. 2000; Hensley and Draine 2021), and diffuse interstellar bands (Herbig 1995) have each been used to constrain dust candidates. Presolar grains from meteorites also provide a way to directly probe the dust content in the local ISM (Draine 2003; Nittler and Ciesla 2016; Hensley and Draine 2023).

Of the many candidates, the most abundant grain type is thought to be in the form of silicates – anions containing Si and O – bonded with either Mg or Fe into tetrahedrons. They occupy about 2/3 of the dust mass in local galaxy ISMs (Draine 2003; Galliano 2022) and assemble into amorphous structures rather than crystalline arrangements (Do-Duy et al. 2020). The remaining mass in dust comes from carbonaceous grains, where carbon atoms bond with other carbons and hydrogen into a variety of structures including graphite sheets, hydrogenated amorphous carbons, and polycyclic aromatic hydrocarbons (PAH) aromatic rings. The size distributions of silicate and carbonaceous grains can be constrained by dust extinction curves (Mathis et al. 1977; Siebenmorgen et al. 2014) and by fitting the polarization spectrum (Das et al. 2010) and are found to vary between 0.3 nm and 0.3 µm in size. Within the densest regions of the ISM, grains can be enshrouded in molecular mantles of frozen CO, CO_2, and H_2O (Boogert et al. 2015). These contribute additional absorption bands at MIR wavelengths within these regions.

1.0.2 Dust Grain Formation, Growth, and Destruction

The raw metals for dust are ejected into the interstellar medium through the stellar outflows of Asymptotic Giant Branch (AGB) stars (Sargent et al. 2010; Michałowski 2015) and from core-collapse supernovae (SNII; Bevan et al. 2017; De Looze et al. 2017). Main sequence stars with $M_{\odot} < 8\ M_{\odot}$ will undergo the AGB phase which is characterized by thermal pulses and mass loss that can eject heavier core elements through stellar winds to eventually join the ISM (Dell'Agli et al. 2015). Models predict that approximately 10 - 40% of heavier elements ejected will condense into dust (Morgan and Edmunds 2003; Ventura et al. 2012; Schneider et al. 2014). The additional contribution from cooling gas emitted from SNIIs is relatively unconstrained, with predictions ranging between 10^{-3} - 1 $M_{\odot}$ (Todini and Ferrara 2001; Ercolano et al. 2007; Bianchi and Schneider 2007; Bocchio et al. 2016; Marassi et al. 2019) and actual measurements from *Herschel* finding values between $0.03 - 1.1 M_{\odot}$ (Barlow et al. 2010; Arendt et al. 2014; De Looze et al. 2017; Bevan et al. 2017; Priestley et al. 2019; Gomez et al. 2012; Temim and Dwek 2013; De Looze et al. 2019). Even accounting for the uncertainties in the SNII dust production rate, dust grains must continue to grow throughout the ISM to account for the presence of large quantities of dust at high redshifts (Bertoldi et al. 2003; Priddey et al. 2003; Rowlands et al. 2014; Watson et al. 2015; Michałowski 2015). Furthermore, within dense Milky Way ISM regions, FIR dust cross sections are larger, and MIR emission features weaken, indicating a bias favoring larger grain sizes for such environments (Stepnik et al. 2003; Köhler et al. 2015; Galliano et al. 2018). Dust grains can amalgamate together and accrete further mass from the metals and molecules in the gas, thus depleting it.

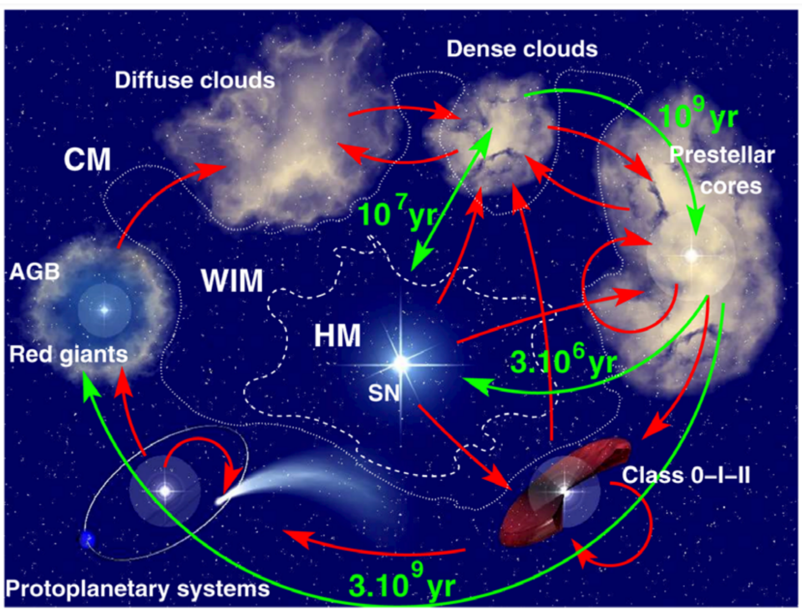

Figure 1.2. Figure reproduced from Demyk 2011 demonstrating the stages and timescales of the star formation cycle and dust evolution in the ISM.

The lifetime of a dust grain within the various phases of the ISM is uncertain, with estimates between $1 - 20 \times 10^8$ years (Barlow 1978; Watson et al. 2015; Dwek and Scalo 1980; Bocchio et al. 2014; Zhukovska et al. 2016). Dust grains can be broken apart or destroyed through the reverse shocks of the same SNII events that produced them (Bocchio et al. 2014), as well as by thermal sputtering or collisions with gas, other grains, and cosmic rays, and photodesorption from high energy photons (Galliano et al. 2018). Theoretical predictions expect silicate grains to have the highest survivability and carbonaceous grains, particularly amorphous hydrogenate carbons, to be more easily destroyed within SNII shocks (Bocchio et al. 2014; Slavin et al. 2015).

1.0.3 Dust and the Star Formation Cycle

Dust plays a synergistic role with the gas content of the ISM that directly contributes to the formation of new stars within molecular clouds (Kennicutt 1998; Kennicutt et al. 2009) through the star formation cycle (Figure 1.2). Dust extinction of high energy UV and X-ray photons as well as cosmic rays prevents them from dissociating molecular hydrogen (H_2), the fuel for star-formation (Li and Greenberg 2003). The surfaces of dust grains also operate as an energetically favorable catalyzing site where atomic hydrogen atoms collide, lose kinetic energy, and more readily coalesce into molecular hydrogen (Gould and Salpeter 1963; Wolfire et al. 1995; Le Bourlot et al. 2012; Bron et al. 2014; Wakelam et al. 2017) before being liberated back into the gas, thus increasing the efficiency of H_2 formation. Furthermore, the re-emission of absorbed radiation by dust grains as thermal radiation offers a cooling channel for molecular clouds to radiate away gravitational energy as they collapse under their own gravitational pull to form a protostellar object (Klessen and Glover 2016).

1.0.4 Dust SED Modeling

Numerous models and fitting approaches have been developed to fit the dust continuum and spectral emission across NIR to millimeter wavelengths. Each incorporates its own assumptions and simplifications, and they vary considerably in terms of model complexity and computational requirements. The simplest strategy is to fit only the SED from dust thermal and spectral emission itself either directly to a model parametrized in terms of the dust mass, spectral shape, and the ISM conditions or to scale an existing SED template derived from separate observations to fit the observations. These methods therefore only require observations of the dust emission and can be used when few observations are available. Other approaches use the principle of energy balance between the stellar and dust components to fit the entire panchromatic SED across UV and millimeter wavelengths. Fitting codes such as CIGALE (Boquien et al. 2019), MAGPHYS (Cunha et al. 2008), and Prospector (Leja et al. 2017) have become popular when fitting with constraints across UV to IR wavelengths and employ the Bayesian statistical framework to derive full posterior distributions for key galaxy parameters like stellar masses, SFRs, dust masses and can incorporate realistic grain models.

A full model of the propagation of radiation through a simulated medium of dust can be carried out with radiative transfer fitting and compared to observations (Popescu et al. 2011; Nersesian et al. 2020; Galliano 2022). Heterogeneous mixtures of dust grain compositions and non-local heating from the diffuse ISRF can be fully incorporated. However, while self-consistent solutions of complex media can be derived through radiative transfer modeling, it is a highly computationally expensive approach, particularly for large samples or spatially resolved studies seeking to fit on a per-pixel basis. Furthermore, it is difficult to constrain the actual three-dimensional structure of the ISM through observations owing to the two-dimensional projections of observations and inclination effects.

Below, I detail the two strategies implemented in the SED fitting code described in Chapter 4, which are the modified blackbody model and the physically motivated dust models.

1.0.4.1 Modified Blackbody

The modified blackbody (MBB) or graybody model (Hildebrand 1983) is the most utilized model for fitting the dust emission at FIR wavelengths. It assumes that the dust grains are an isothermal population in equilibrium with the ISM environment and is characterized by a single temperature such that the flux density can be written as

$$I_\nu = (1 - e^{-\tau_\nu}) \times B_\nu(T). \tag{1.1}$$

Here, τ_ν is the optical depth at frequency ν and $B_\nu(T)$ is the Planck function for the dust temperature T. The dust is assumed to be optically thin, such that Equation 1.1 can be reduced to

$$I_\nu = \tau_\nu \times B_\nu(T), \tag{1.2}$$

owing to the fact that $\tau_\nu \ll 1$. The optical depth τ_ν can be expanded in terms of the dust surface mass density Σ_d, and the opacity κ_ν, giving

$$\tau_\nu = \kappa_\nu \times \Sigma_{dust}. \tag{1.3}$$

The MBB model approximates the opacity as the power-law

$$\kappa_\nu = \kappa_0 \times \left(\frac{\nu}{\nu_0}\right)^\beta, \tag{1.4}$$

where κ_0 is the reference opacity at frequency ν_0 and β is the power law spectral index. The resulting expression for flux density becomes

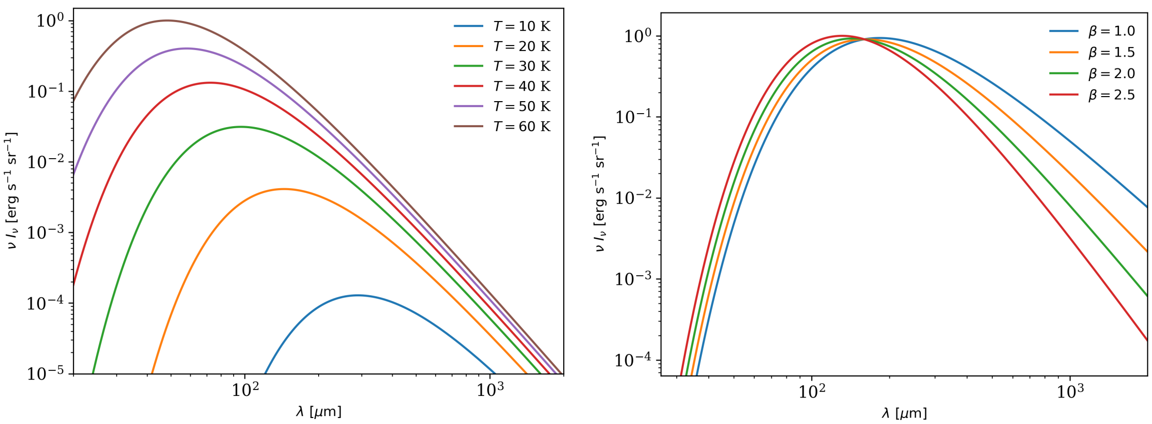

Figure 1.3. Example SEDs for modified blackbodies. Left: SEDs derived from varying the dust temperature between 10 K and 60 K, while keeping Σ_{dust} and β constant. Right: The same, but varying β between 1.0 and 2.5 while keeping the other parameters constant.

$$I_\nu = \kappa_0 \left(\frac{\nu}{\nu_0}\right)^\beta \Sigma_{dust} B_\nu(T). \tag{1.5}$$

The free parameters for the MBB model are therefore given by Σ_{dust}, T, and β. These effect the overall scale of the SED, the location of the peak, and the slope at FIR wavelengths and beyond respectively. Figure 1.3 illustrates the effect of varying the dust temperature and emissivity index β for a range of values.

The simplicity of this approach introduces a few drawbacks. Dust grains at MIR wavelengths are stochastically heated and not in thermal equilibrium with the ISM. They therefore cannot be fit effectively with MBB models. The assumption of a simple power law for the dependence on wavelength for FIR opacity is not entirely supported by laboratory analogs, which have been found to demonstrate more complex behaviors, and the choice of the reference opacity κ_0 can affect the recovered dust mass as they are correlated (Galliano 2022). Furthermore, the ISM along a line-of-sight and within a telescope beam is likely a superposition of multiple temperature components and dust grain populations, such that the fitted T and β will represent an average value only and will not map to the actual physical values for any single

ISM component. Consequently, it is challenging to constrain grain compositions with MBB models alone.

Another complication is that T and β are mathematically anti-correlated in the MBB models (Shetty et al. 2009b; Shetty et al. 2009a; Kelly et al. 2012; Galliano 2018; Tang et al. 2021; Galliano 2022) such that an unphysical inverse relationship can be introduced between them when performing least-squares or non-hierarchical Bayesian fitting. This anti-correlation can compound on spurious correlations from statistical uncertainties and correlated calibration errors across instruments to obscure or exaggerate any potential intrinsic physically meaningful correlations between T and β.

One method to avoid the impact of the T-β anti-correlation is to fix one parameter and only fit the other and Σ_{dust}. The emissivity is usually fixed owing to the constraints on silicate and graphite grains from laboratory analogs and the greater variation in dust temperatures across the different ISM phases. An estimate of β can be determined from observational data by taking the flux ratio between two flux bands in the Rayleigh-Jeans limit of the SED which cancels out the dust mass and opacity leaving

$$\frac{I(\nu_1)}{I(\nu_2)} = \left(\frac{\nu_1}{\nu_2}\right)^{\beta} \times \frac{B_{\nu,1}(T)}{B_{\nu,2}(T)}. \tag{1.6}$$

In this way, β variations within different pixels can still be explored while fixing β. Fitting with a free β can produce underestimates of the dust mass relative to fixed MBB (Bianchi 2013; Rémy-Ruyer et al. 2013) and morphological dependent differences from other modeling approaches (Galliano et al. 2021). Fixing β also introduces dust mass underestimates and may not be realistic if the grain composition varies (Galliano 2022).

A more statistically rigorous approach to improving MBB fit results is to incorporate the parameter degeneracies into the model itself, which can be accomplished through hierarchical or multi-level Bayesian fitting (Kelly et al. 2012; Galliano 2018; Lamperti et al. 2019; Galliano 2022). A description and a particular implementation of hierarchical Bayesian modeling (Gelman and Hill 2007) to MBB fits is given in more detail in Chapter 5. In brief, information about the distributions and relationships between model parameters is incorporated into the model through Bayesian priors and is determined from the data itself instead of being built from theoretical predictions or separate observational constraints. This is a powerful technique for handling errors and correlations introduced throughout all stages of data acquisition and analysis. It has been shown to accurately recover intrinsic correlations between T and β in MBB fits to simulated SEDs (Kelly et al. 2012; Galliano 2018; Lamperti et al. 2019; Galliano 2022) and reduce uncertainties in the recovered probability distributions.

1.0.4.2 Physically Motivated Dust Models

Fitting with physically motivated dust models has two major advantages over MBB model fits in that the physical models can incorporate both realistic dust grain emissivities and the mixing of different physical conditions or temperature components within the ISM. In this approach, the dust is modeled as being heated by a distribution describing the contributions from the local and diffuse ISRF. The differential dust mass dM_{dust} is heated by a range of ISRF values dU and can be well-described in terms of a power law (Dale et al. 2001) with a power-law index alpha in the form

$$\frac{1}{M_{\rm dust}} \frac{dM}{dU} \propto U^{-\alpha}, \tag{1.7}$$

where M_{dust} is the total dust mass, U is the value of the ISRF normalized by the mean value of the local ISRF in the Milky Way of $2.2 \times 10^{-5}\,\mathrm{W\ m^{-2}}$ (Mathis et al. 1983). The proportionality constant is expressed as

$$N = \begin{cases} \frac{(1-\alpha)}{(U_- + U_\Delta)^{1-\alpha} - (U_-)^{1-\alpha}} & \text{if } \alpha > 1 \\ \\ \frac{1}{\ln(U_- + U_\Delta) - \ln(U_-)} & \text{if } \alpha = 1, \end{cases} \tag{1.8}$$

with U_- being the lower limit on the ISRF and U_Δ the range of ISRF values that are heating the ISM such that $U_+ = U_- + U_\Delta$. The simple power law formalism has been demonstrated to hold even for complex, clumpy ISM distributions in 3D radiative transfer simulations (Galliano 2022), though it is not the only parametrization with others being developed given different observational constraints (Draine and Li 2007). The mean value of the ISRF heating the dust can be written as

$$\langle U \rangle = \begin{cases} \frac{1-\alpha\left((U_- + U_\Delta)^{2-\alpha} - U_-^{2-\alpha}\right)}{2-\alpha\left((U_- + U_\Delta)^{1-\alpha} - U_-^{1-\alpha}\right)} & \text{if } \alpha \neq 1 \text{ and } \alpha \neq 2 \\ \\ \frac{U_\Delta}{\ln(U_- + U_\Delta) - \ln(U_-)} & \text{if } \alpha = 1 \\ \\ \frac{\ln(U_- + U_\Delta) - \ln(U_-)}{U_-^{-1} - (U_- + U_\Delta)^{-1}} & \text{if } \alpha = 2, \end{cases} \tag{1.9}$$

and is a useful, physically meaningful parameter that can play a similar role to the dust temperature of MBB models (Galliano 2018). A dust temperature can be inferred in this formalism given U_- and the emissivity β of the dust being heated through

$$T = T_0 \times U_-^{\left(\frac{1}{4+\beta}\right)}, \tag{1.10}$$

where T_0 is the dust temperature in the solar neighborhood with a value of 18.3 K (Aniano et al. 2012; Nersesian et al. 2019). This relies on the assumption that the dust is heated by a Milky Way-like ISRF distribution (Mathis et al. 1983).

With Equation 1.7, the luminosity of a dust mass heated by the ISRF at a given frequency is calculated by incorporating the monochromatic emissivities of a modeled dust population and integrating between U_- and U_+ to give

$$L_\nu = \frac{1}{M_{\rm dust}} \times N \times \int_{U_-}^{U_+} j_\nu(U) \times U^{-\alpha}\, dU. \tag{1.11}$$

The monochromatic emissivity in Equation 1.11 can be for a single dust grain population or a linear combination of several different types, which allows for the inclusion of large and small silicates and graphite grains, ionized and neutral PAH molecules, and spinning magnetized grains into a single composite population. The relative abundance of each type can be adjusted by setting that type's mass fraction; for example, the PAH grain fraction, q_{PAH} can be fit along with the other parameters. Many dust models have been developed over the course of many decades and include but are not limited to the silicate, graphite, PAH models of Zubko et al. 2004, Draine and Li 2007, and Draine et al. 2014, the THEMIS modeling framework which incorporates grains consisting of both cores and mantles (Jones et al. 2013; Köhler et al. 2014; Jones et al. 2017), and the Astrodust model that employs a single unified grain type to account for recent FIR polarization constraints on grain compositions (Draine and Hensley 2021; Hensley and Draine 2023). These models have different observational features, including variations in their MIR and FIR slopes (Galliano 2022; Hensley and Draine 2023) and comparisons between fitted results with multiple models can help place bounds on fundamental dust physics. The total extinction as a function of wavelength for several dust models is plotted in Figure 1.4.

The emissivity of a population of N different dust grain compositions being heated by a single ISRF value takes the form of

$$j_\nu = \sum_{i=1}^{N} \int da \left(\frac{1}{N_H}\frac{dn_i}{da}\right) \times \int dT \left(\frac{dP}{dT}\right)_{i,a} \times C_{\rm abs,i}(\nu, a) \times B_\nu(T). \tag{1.12}$$

Each grain type i is characterized by a size distribution dn_i/da and an absorption cross-section $C_{\rm abs,i}(\nu, a)$ that varies with the grain radius a and composition, both of

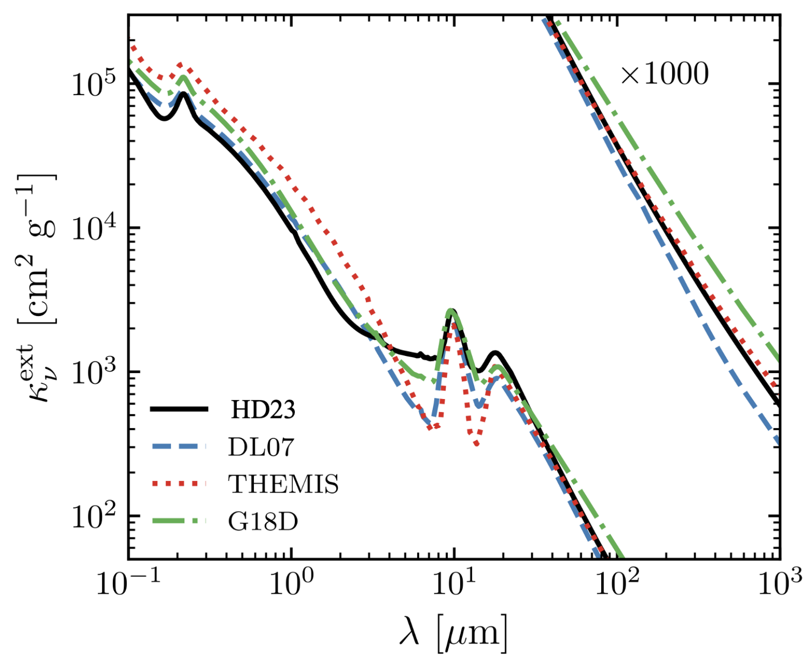

Figure 1.4. Figure reproduced from Hensley and Draine 2023, plotting the total extinction as a function of wavelength for a range of physical dust models. Models include Hensley and Draine 2023 (HD23), Draine and Li 2007 (DL07), Jones et al. 2017 (THEMIS), and Guillet et al. 2018 (G18D).

which can be estimated from laboratory analogs. Separate cross-sections are required for unpolarized and polarized emission from grains (Hensley and Draine 2023). The dP/dT term describes the distribution of grain temperatures and will take different forms depending on whether the grains are in thermal equilibrium with the ISM or not. The absorption cross-section and temperature distribution modify the idealized blackbody thermal emission and are integrated over the range of dust temperatures and grain sizes considered in the model.

The parameters for the physically motivated models are then M_{dust}, α, U_{-}, and U_{Δ}. Example dust SEDs that are irradiated by different values of the ISRF are shown in Figure 1.5 with stronger ISRFs serving to shift the peak of the SED to towards shorter wavelengths. The PAH mass fraction q_{PAH} or the equivalent mass in small hydrocarbons can be included if MIR bands are being fit. Contrary to MBB models, the complexity of the physically motivated dust models makes them more computationally expensive to re-compute the dust emissivities, particularly when using Monte Carlo Markov Chain (MCMC) fitting techniques where thousands of models may need to be calculated. Emissivities are therefore usually pre-computed for a grid of individual or integrated values of U which can be used or interpolated across during fitting. The physical grounding of these models makes them a powerful tool for extracting physically meaningful information about the dust and the ISM in a way that MBB models cannot. They can also incorporate MIR wavelength constraints that do not meet the thermal equilibrium requirements of MBB fitting. The physically motivated fitting approach is still susceptible to resolution effects, where different dust mixtures cannot be distinguished if considering integrated measurements or when incorporating lower-resolution data into fits. The model parameters, as well as underlying choices related to the grain size distributions and the ISRF prescription, also show spurious model and noise-induced correlations mirroring the T-β anti-correlation of MBB models (Galliano 2018; Galliano 2022). A hierarchical Bayesian model can similarly be used here to mitigate this effect and improve parameter recovery.

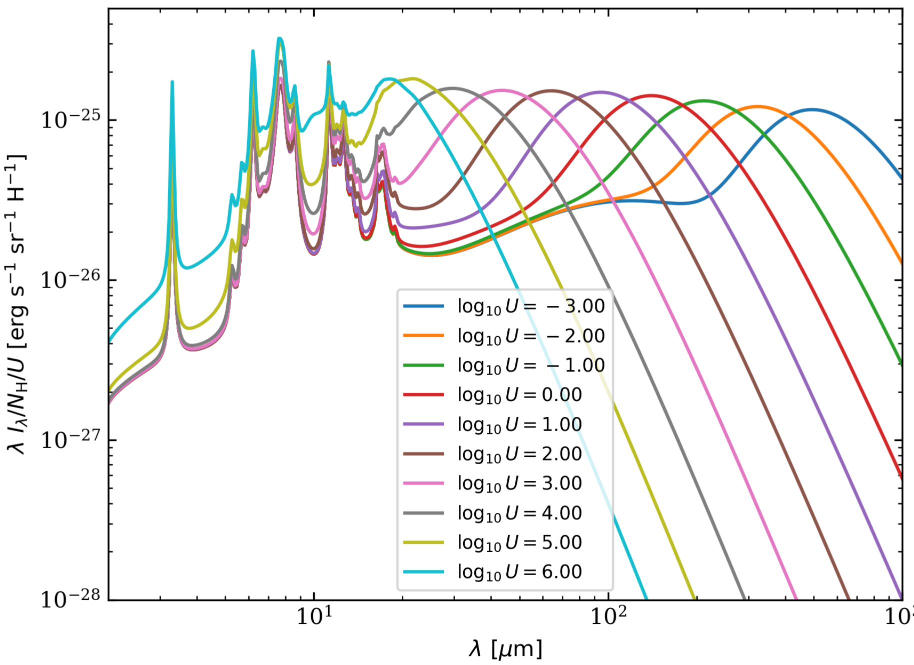

Figure 1.5. Dust SEDs showing the effect of the interstellar radiation field strength U on the dust emission. The dust model is the Astrodust+PAH model of Hensley and Draine 2023 and is being heated by a single value of $\log(U)$ ranging between -3 and 6.

1.0.5 Surveys of Dust Content in Nearby Galaxies

Dust is most easily studied within the Milky Way and in local galaxies where the faintest and smallest scale structures can be explored. Galaxies in the local universe exhibit a considerable range of dust and star-formation properties and are ideal laboratories for carrying out observations to constrain and calibrate dust models and dust scaling relationships. Space-based instruments have thus far provided the largest datasets of nearby galaxy dust content observations. In particular, *Herschel* (Pilbratt et al. 2010) has been a fundamental driver in furthering our knowledge of the ISM of galaxies. The broad wavelength coverage of *Herschel* photometry from 70-500 µm, when combined with MIR data from *Spitzer* (Werner et al. 2004), provides flux measurements across a significant portion of the dust SED, thereby allowing for more accurate dust mass estimates and for improved constraints on the dust composition, such as the dust grain opacity at FIR wavelengths. Along with *Planck*, *Herschel* is sensitive to the cold dust components that are not traced by telescopes like *Spitzer* and JWST, and consequently recovers a greater fraction of the total dust mass. *Herschel*'s superior, sub-arcminute resolution compared to Infrared Astronomical Satellite (IRAS; Neugebauer et al. 1984) and the Infrared Space Observatory (ISO; Kessler et al. 1996) has enabled the investigation of the spatial variation of ISM properties within galaxies, which can trace physical differences in the nature of the dust, gas, and star-formation within different regions (Madden and Cormier 2019).

Surveys with *Herschel* include the *Herschel* Reference Survey, Key Insights on Nearby Galaxies: a Far-Infrared Survey with *Herschel* (KINGFISH, Dale et al. 2012), the *Herschel* Dwarf Galaxy Survey (DGS; Madden et al. 2013), the *Herschel* Inventory of The Agents of Galaxy Evolution survey (HERITAGE, Meixner et al. 2010), and DustPedia (Davies et al. 2017). Together these provide observations of the warm and cold dust content of nearly 1000 nearby local galaxies out to a distance of 200 Mpc over several orders of magnitude in stellar masses $3 \times 10^6 - 3 \times 10^{11}\, M_\odot$ and metallicities in ($Z/Z_\odot$ = 0.03 – 1.20). Figure 1.6 plots the distribution of stellar masses and metallicities for the KINGFISH, DGS, and DustPedia surveys.

These datasets are well complemented by surveys and individual observations of local galaxies with ground-based submillimeter and millimeter instruments which extend beyond *Herschel*'s longest wavelength band and provide higher resolution measurements compared to the 36.3″ beamsize of the SPIRE 500 µm band. Surveys at wavelengths of 450 µm and 850 µm include the JCMT Nearby Galaxies Survey (Wilson et al. 2009; Pattle et al. 2023), which observed more than 100 massive galaxies within 25 Mpc, and the JCMT dust and gas In Nearby Galaxies Legacy Exploration (JINGLE; Saintonge et al. 2018) survey that targeted 193 galaxies with *Herschel*

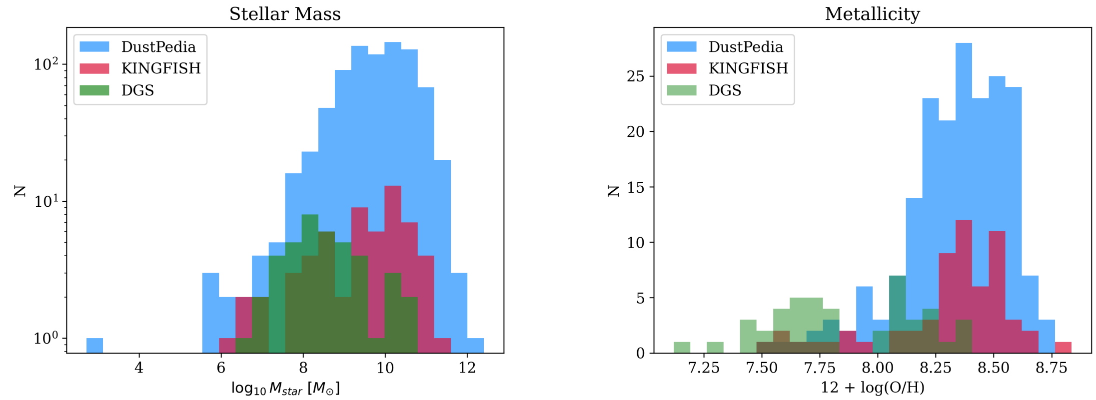

Figure 1.6. Left: Histograms of the stellar mass distributions for KINGFISH, the DGS, and DustPedia galaxies. Stellar mass measurements were acquired from Kennicutt et al. 2011 for KINGFISH, Rémy-Ruyer et al. 2013 for the DGS, and Clark et al. 2018 for DustPedia. Right: Histograms of galaxies with metallicity measurements from the same samples. Metallicities are from Kennicutt et al. 2011, Madden et al. 2013, and De Vis et al. 2019 for KINGFISH, the DGS, and DustPedia respectively.

observations at $0.01 < z < 0.02$. NIKA2 (Adam et al. 2018) on the IRAM 30m telescope is conducting the Interpreting the Millimetre EMission of Galaxies survey (IMEGIN) to observe ∼40 local galaxies within 30 Mpc at 1.2 and 2.0 mm as well as the Submillimeter Excess In Nearby Fairly-Extended Low-metallicity Dwarfs (SEINFELD) survey to explore dust content in local dwarf galaxies. Targeted observations of nearby dwarf galaxies have been carried out with SCUBA at 450 µm and 850 µm (Galliano et al. 2003), LABOCA at 870 µm (Galametz et al. 2009), AzTEC at 1.1 mm (Calzetti et al. 2018), and NIKA2 (Ejlali et al. 2022; Ejlali et al. 2023). These provide invaluable additional wavelength constraints for galaxies in the DGS.

1.0.6 Dust in Dwarf Galaxies

A notable finding from investigations into data from nearby galaxy surveys is the discrepancy between the dust properties of dwarf galaxies and those of massive galaxies. Galliano et al. 2021 explored the dust content within the combination of the DustPedia sample of 800 local galaxies with the DGS survey of 48 nearby dwarf galaxies using a hierarchical Bayesian fitting method with the THEMIS physically motivated dust model. As shown in Figure 1.7, they identified a strong relationship between the dust-to-gas mass ratios and the metallicity with a sharp drop in M_{dust}/M_{gas} near $Z/Z_{\odot} = 0.4$ by several orders of magnitude. This agrees with earlier findings of Rémy-Ruyer et al. 2014 using KINGFISH, which is dominated by star-forming spirals, and dwarf galaxies in the DGS. Chemical models can be used to interpret this observation as an evolution in the relative importance of the primary grain formation and growth mechanisms as a function of metallicity (Asano et al. 2013; Zhukovska 2014; De Vis et al. 2017). For the lowest metallicity systems below $Z/Z_{\odot} = 0.2$, the timescale of grain condensation from supernovae remnants relative to the grain growth is very small owing to the low density of the ISM making dust grain interactions much rarer. The M_{dust}/M_{gas} ratio increases rapidly with metallicity as the ISM grain growth efficiency grows outpacing the SNII condensation rate. Eventually, the destruction of grains due to SNII becomes important within the highest metallicity region.

Using modified blackbody fits to KINGFISH and the DGS, Rémy-Ruyer et al. 2013 found the effective dust temperatures in the DGS to be systematically higher relative to those of KINGFISH with a median temperature of 32 K compared to 23 K and are characterized by much greater scatter. This is in agreement with a host of earlier studies (Galliano et al. 2005; Madden et al. 2013) that have demonstrated that dwarf dust emission peaks at wavelengths shorter than the $>$ 100-µm peak of both regular and starburst galaxies, with some starburst and blue compact dwarfs having peaks between 35 and 70 µm (Madden and Cormier 2019). While temperatures derived through MBB fits are typically overestimated to that inferred from physically motivated dust models, thereby leading to an underestimate in the total dust mass, there is a physical expectation of higher temperatures owing to the lack of dust shielding contributing to a harder ISRF. Rémy-Ruyer et al. 2015 subsequently fit the

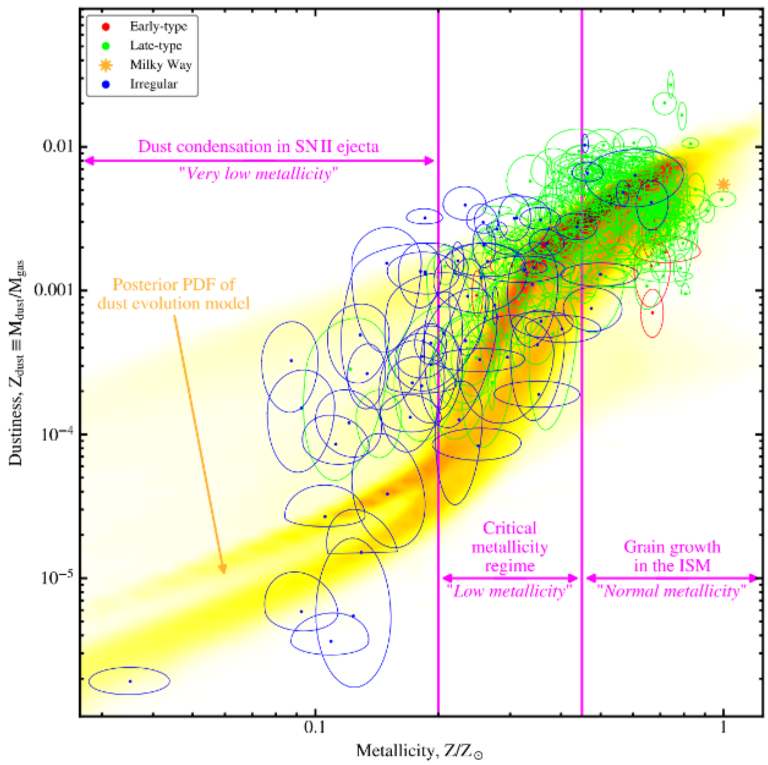

Figure 1.7. Figure reproduced from Galliano 2022. The dust-to-gas mass ratio (dustiness) versus metallicity for a subset of galaxies in the DustPedia sample. Ellipses represent uncertainties from the model fits. Contours are from the dust evolution model in Galliano et al. 2021.

dust SEDs of KINGFISH and DGS galaxies to the dust model of Galliano et al. 2011 and identified a similar difference in temperatures (Figure 1.8). Dwarf dust peaks are broader than star-forming spirals alluding to the presence of a clumpy ISM with multiple different temperature components (Boselli et al. 2010; Boselli et al. 2012; Ciesla et al. 2014). The far-IR slope inferred from fits to DGS galaxies and the LMC is therefore shallower, with an average metallicity independent spectral index of $\beta = 1.7$, compared to the Milky Way slope of $\beta = 2.0$ (Galliano et al. 2011; Rémy-Ruyer et al. 2014).

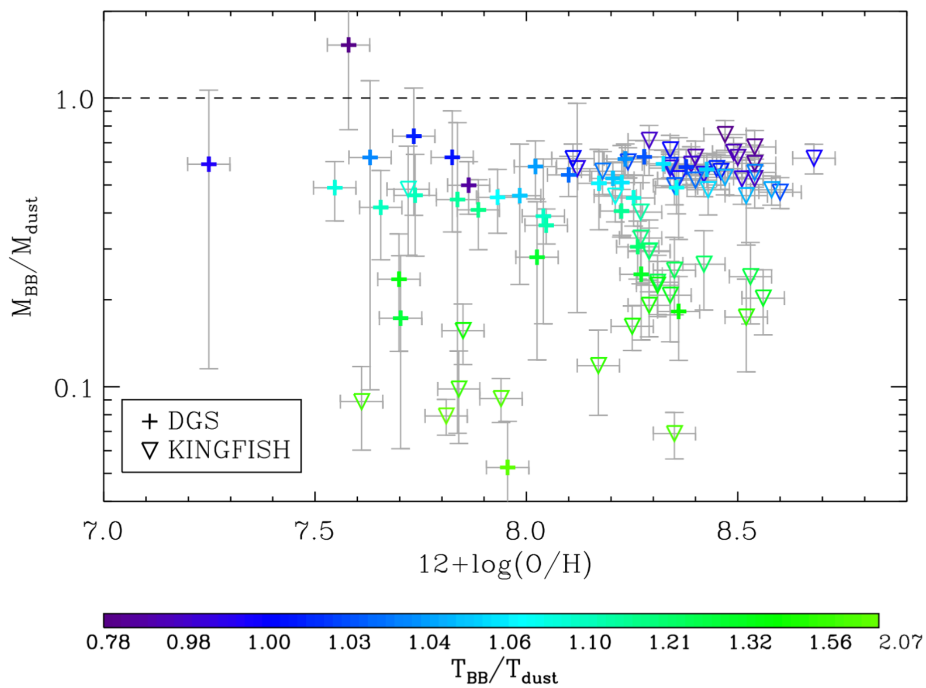

Figure 1.8. Figure reproduced from Rémy-Ruyer et al. 2015. The ratio of dust masses derived from MBB fits to those from physical dust model fits using the dust model of Galliano et al. 2011. Galaxies from KINGFISH and the DGS are included. The color represents the ratio of derived temperatures, with T_{dust} being derived from Equation 1.10 with $T_0 = 19.7$ K.

A consequence of the more intense ISRF is a reduction in the emission from PAH grains in dwarf galaxy SEDs, with the strength of the spectral features and the inferred PAH mass fraction varying inversely with the metallicity (Rémy-Ruyer et al.

2013; Madden and Cormier 2019; Hunt et al. 2010; Henkel et al. 2022). In some low metallicity systems PAH emission is completely absent (Engelbracht et al. 2005). These small dust grains are likely destroyed by the more intense radiation, which has also been observed within HII regions (Madden et al. 2006; Galliano et al. 2011). If PAH grains preferentially form in molecular clouds, the ISM porosity of dwarfs would result in a lower production efficiency. Finally, slower production of carbon from AGB stars is another potential mechanism for inhibiting PAH formation (Galliano et al. 2008).

Similar to the PAH content, the CO emission in dwarfs is also underrepresented, even for dwarfs undergoing intense star formation (Leroy et al. 2005; Schruba et al. 2012; Cormier et al. 2014; Hunter et al. 2024). Shielding of molecular clouds by CO molecules becomes inefficient at very low metallicities where the elemental abundances mean there are fewer atoms to form CO from. Consequently, compared to more massive galaxies there will be a greater fraction of CO-dark H_2 gas, whose presence has been inferred through comparisons of the excess IR emission relative to the measured gas mass in the LMC (Bernard et al. 2008). The difference in structure between molecular clouds for a low metallicity ISM is shown in Figure 1.9 This necessitates a higher X_{CO} conversion factor for low metallicity environments though its value remains poorly constrained at $Z/Z_{\odot} < 1/5$ owing to the difficulty in detecting CO (Leroy et al. 2007; Leroy et al. 2011; Galliano et al. 2011; Rémy-Ruyer et al. 2013).

Another key trait of dwarf galaxy SEDs is the observation of excess emission, primarily identified in submillimeter bands (Reach et al. 1995; Lisenfeld et al. 2001; Galametz et al. 2009; Bot et al. 2010; Galametz et al. 2011; Galliano et al. 2011; Grossi et al. 2015; Chang et al. 2021), but also at millimeter wavelengths (Galliano et al. 2003; Galliano et al. 2005), relative to fits to standard dust models. While definitions vary slightly, the excess can be described in terms of the higher-than-expected fit

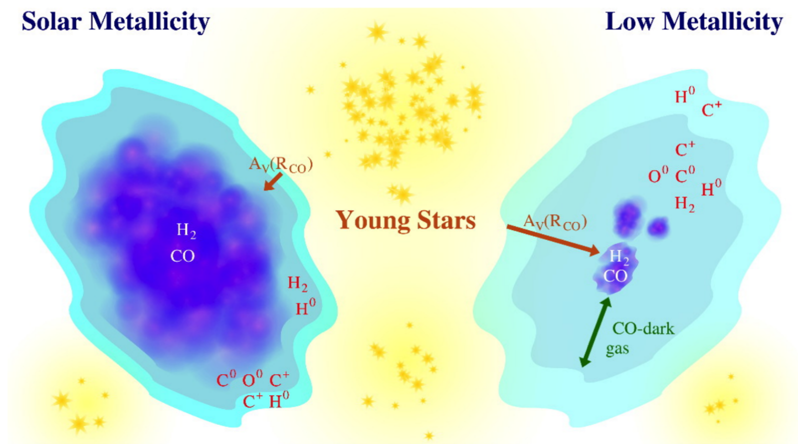

Figure 1.9. Figure reproduced from Hunter et al. 2024. Comparison of the extent of H_2 gas not traced by CO for a molecular cloud within a solar metallicity ISM environment and one in a low metallicity ISM.

residuals for submillimeter and millimeter observations relative to a single isothermal dust population in the case of the STMBB models or relative to the expected slope from a single dust grain mixture in the physically motivated dust models. The excess is often defined directly in terms of the residual in the relative flux in some band as

$$R_\nu = \frac{f_{\nu,\mathrm{model}} - f_{\nu,\mathrm{fit}}}{f_{\nu,\mathrm{model}}}, \tag{1.13}$$

where ν is most often the 500-µm *Herschel* SPIRE band but can be longer owing to the appearance of excess emission beyond Herschel's coverage (Rémy-Ruyer et al. 2013; Gordon et al. 2014). In some dwarfs for instance, such as NGC 1569, the excess has only been observed beyond 850 µm and at 1.3 mm and is not detected in the *Herschel*/SPIRE 500-µm band (Lisenfeld et al. 2001; Galliano et al. 2003). Both Magallenic clouds have been found to exhibit an excess (Israel et al. 2010; Bot et al. 2010) except for some star-forming complexes within them (Galametz et al. 2013). The LMC and SMC allow for very high-resolution, spatially resolved studies of the distribution of the excess. While present in massive star-forming galaxies (Dumke et al. 2004; Galametz et al. 2009; Galametz et al. 2011) including the Milky Way (Reach et al. 1995; Paradis et al. 2012) and potentially M33 (Hermelo et al. 2016), it is more commonly observed in low metallicity systems and with greater residuals relative to model fits. Rémy-Ruyer et al. 2013 found that 41% of the DGS with *Herschel* SPIRE 500 μm detections exhibited excess emissions, with the lowest metallicity systems having emission 150% higher than predictions.

Many hypotheses have been advanced to explain the excess emission. A secondary population of very cold dust (VCD) grains with an effective temperature of less than 10 K would enhance the SED at submillimeter and millimeter wavelengths in a way consistent with the observed excess. The resulting dust mass would, however, also be significantly increased by as much as 80% (Galliano et al. 2003; Galliano et al. 2005; Galliano 2022) which would be unrealistic in many scenarios given the molecular gas masses which can be independently measured. Such VCD components would need to be distributed in dense clumps. However, for the LMC, the excess emission is most prominent in less dense regions, which rules out this scenario in this instance (Galliano et al. 2011). This has similarly been found to be the case for non-barred spiral galaxies (Hunt et al. 2015).

Instead of varying the dust temperature, the spectral properties of the grains can be varied by introducing a temperature dependence to the grain emissivity (Meny et al. 2007), which is observed by some laboratory measurements (Demyk et al. 2017b; Demyk et al. 2017a; Galliano 2022). The variations in the opacity of some laboratory analogs with temperature are shown in Figure 1.10. In this case, the emissivity index β will decrease with temperature thus producing a shallower slope for colder dust emitting at submillimeter wavelengths and beyond. Such a dust grain would likely be comprised of amorphous carbon grains instead of the more structured graphite carbon grains (Meny et al. 2007; Nashimoto et al. 2020). This approach has had some success in explaining the LMC excess (Paradis et al. 2012), but not the SMC (Bot et al. 2010). Gordon et al. 2014 similarly found that the LMC and SMC submillimeter excesses as measured in the HERITAGE survey are best fit through a MBB characterized by two spectral emissivity indices rather than a TTMBB model, pointing to a change in the emissivity as the culprit rather than additional dust populations.

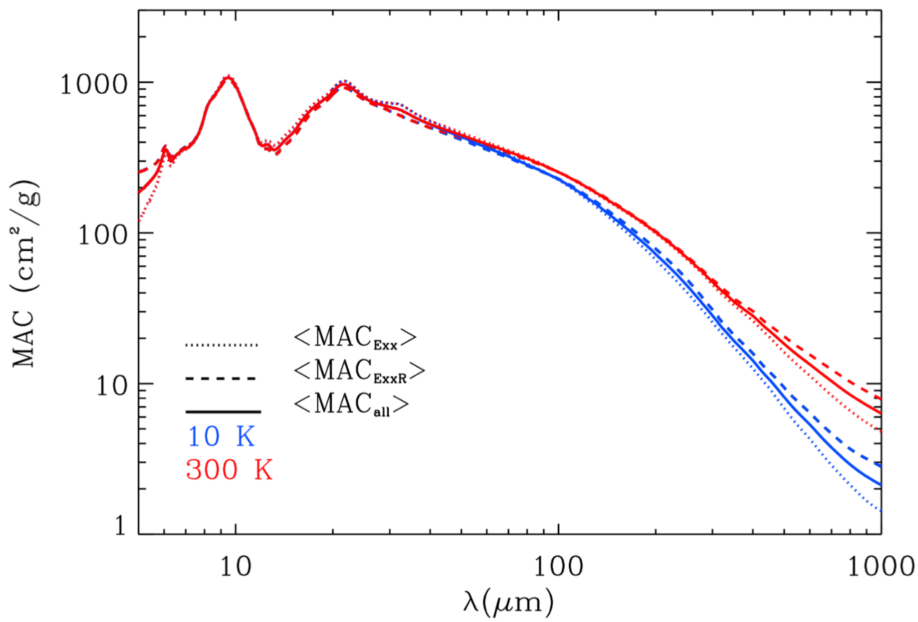

Figure 1.10. Figure reproduced from Demyk et al. 2017a. The temperature dependence of the average opacity (mass absoprtion coefficient) from several laboratory dust analogs.

Differing emission mechanisms represent a third hypothesis for the source of excess emission. These may include dipole radiation from magnetic nanograins comprised of iron and iron oxides, as well as emission from small, spinning PAH-like grains that contribute to the anomalous microwave emission at centimeter wavelengths (Draine and Hensley 2012; Hensley et al. 2022). A fit to the SMC including both spinning and magnetic grains is plotted in Figure 1.11. Draine and Hensley 2012 found that for the SMC, spinning dust grains were insufficient to explain the observed excess, but could be modeled with the inclusion of magnetic nanograins.

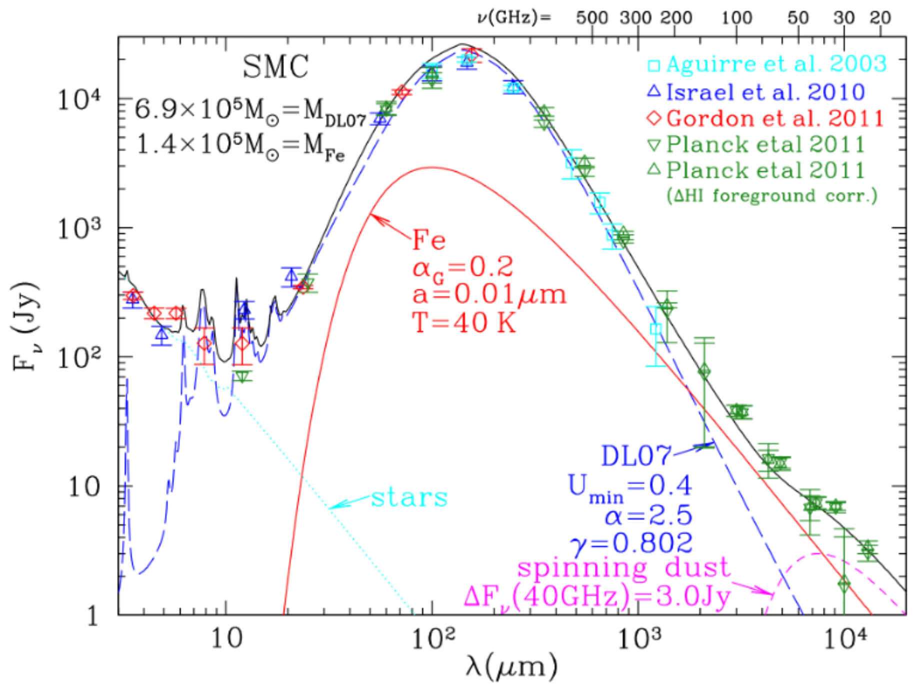

Figure 1.11. Figure reproduced from Draine and Hensley 2012. SED fit to the Small Magellanic Cloud dust emission that incorporates both spinning dust grains and magnetic ferrous nanograins at a temperature of 40 K.

Inclination effects may also play a role in the interpretation of emission as an excess or not in some cases. Using radiative transfer modeling in combination with UV and submillimeter observations to constrain the distributions of stars and dust and incorporate inclination information, Thirlwall et al. 2020 found that no adjustment of dust properties was required to fit the SED in Messier 33, whereas previous studies had identified the galaxy as showing an excess that required a different dust prescription (Hermelo et al. 2016; Williams et al. 2019).

CHAPTER 2

THE TOLTEC CAMERA

TolTEC (Bryan et al. 2018; Wilson et al. 2020) is an NSF-funded millimeter imaging polarimeter currently installed on the 50-m Large Millimeter Telescope (LMT/ Gran Telescopio Milimétrico Alfonso Serrano; Hughes et al. 2020), which sits atop the dormant volcano Sierra Negra in Puebla, Mexico at an altitude of 4600 m. Figure 2.1 shows the instrument in the receiver cabin at the LMT and Figure 2.2 shows the warm optics components within the receiver cabin. The camera employs 7718 cryogenically cooled dual polarization superconducting Lumped Element Kinetic Inductance Detectors (LEKIDs; Austermann et al. 2018) divided among 3 monochromatic focal plane detector arrays installed along a single optical path to image the sky at 1.1, 1.4, and 2.0 mm (273, 214, and 150 GHz) simultaneously. As a result, each observation with TolTEC can produce a maximum of 9 maps of the same fields, with the additional factor of 3 being accounted for by the 3 different Stokes polarization parameters I, Q, and U. Owing to the large collecting area of the LMT, TolTEC's theoretical diffraction limited beam FWHMs are $5.0''$, $6.3''$, and $9.5''$ at 273, 214, and 150 GHz respectively. The full field-of-view is $4'$ in diameter, nearly matching the full field provided by the LMT. Table 2.1 provides a comparison of the specifications of the TolTEC camera with other ground-based submillimeter and millimeter instruments.

Construction of the instrument began at the University of Massachusetts Amherst Department of Astronomy's Cryogenics Detector Laboratory (CDL) in 2016 and the first in-lab testing of the fully constructed and cooled camera occurred in late 2019. TolTEC was transported and installed on the LMT in December 2021 and two com-

Figure 2.1. TolTEC (far left) and its associated cryogenics and and readout hardware in the receiver cabin at the LMT.

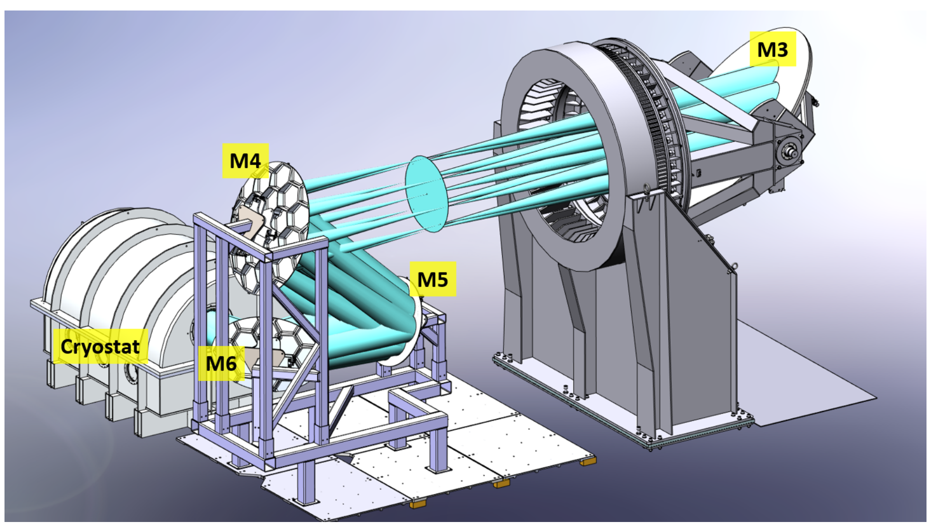

Figure 2.2. Figure reproduced from DeNigris 2024 showing the warm optics of TolTEC in the receiver cabin. Photons are reflected from a series of mirrors (M3 through M6) before entering into the cryostat window. The primary (M1) and secondary mirrors (M2) are not pictured.

Table 2.1. Instrument Specifications and Comparison

Instrument	TolTEC[1]	NIKA2[2]	SCUBA-2[3]
Frequency (GHz)	273/214/150	260/150	666/353
Wavelength (mm)	1.1/1.4/2.0	1.2/2.1	0.45/0.85
FWHM (arcsec)	5/6.3/9.5	11.0/17.5	7.9/13.0
FOV (arcmin)	4	6.5	3.2
Detectors	7718 KIDs	3000 KIDs	10000 Bolometers
Telescope Diameter (m)	50	30	15

[1] Wilson et al. 2020
[2] Adam et al. 2018
[3] Dempsey et al. 2013

missioning observation runs were completed in June-July and December 2022. Wilson et al. 2020 provide an in-depth discussion of the instrument design and detector and cold optics characterization prior to on-sky commissioning. A full overview of TolTEC, in-lab testing, and commissioning is given in DeNigris 2024. Details on the cryogenic systems of TolTEC are given in DeNigris et al. 2020. Information on the warm optics components is provided in Lunde et al. 2020 and Lunde et al. 2022. In-lab characterization is also described in Wilson et al. 2020.

A key deliverable of the TolTEC project is the completion of 10 legacy surveys aimed at answering fundamental questions in millimeter-wavelength astronomy. The surveys will each total 100 hours in integration time and all data will be released publicly following scientific verification of the data products. Of the 10 planned surveys, the first 4 have already been defined through a public workshop of the TolTEC science team members that was carried out at the University of Massachusetts Amherst in late 2018 and will begin observations following the completion of the instrument's commissioning phase. These 4 surveys include the (1) Large-Scale Structure Survey, (2) the Ultra-deep Survey, (3) the Clouds-to-Cores Survey, and (4) the Fields-in-Filaments Survey. In addition to the legacy surveys, TolTEC accepts the submission of PI-led science proposals as part of the LMT's Call for Proposals from the US and Mexican astronomy communities.

2.1 TolTEC Detector Arrays

The use of Kinetic Inductance Detectors (KIDs) (Day et al. 2003; Doyle et al. 2008) in the field of astronomy has increased significantly within the last decade owing to advancements in their design resulting in improvements to their noise characteristics, as well due to the ease and lower cost with which large detector arrays of 10^3 or 10^4 individual KIDs can be assembled. Ground-based cameras like NIKA (Monfardini et al. 2010), NIKA2 (Adam et al. 2018), MUSCAT (Tapia et al. 2022), and CCAT (Chapman et al. 2022) all employ KIDs for their detector arrays.

TolTEC utilizes LeKIDs (Doyle et al. 2008) operating at microwave wavelengths or MKIDs. These are superconducting detectors built by combining an inductor and capacitor into an LC resonator circuit. When a photon is absorbed by the detector, it breaks one of the electron Cooper pair which form in superconducting media. The Cooper pair becomes a pair of quasiparticle whose presence alters the kinetic inductance and subsequently the resonance frequency of the circuit. This temporary change in resonance frequency, f_r, of the circuit can be measured and used to directly measure the strength of the incoming radiation. When designing large arrays of MKIDs, the resonator frequency of each detector is set by editing the capacitance of the capacitor in each MKID to produce a set of unique resonance frequencies that are separated from one another in frequency space thereby allowing for detectors to be distinguished during readout. The MKIDs are driven at their measured resonance frequency when not on source and coupled to a microwave transmission feedline to the readout hardware via capacitive coupling. This simple approach allows for nearly 10^3 individual MKIDs to be read out on a single feedline with microwave multiplexing, thus dramatically reducing the amount of required hardware to support data acquisition. This contrasts with Transition Edge Sensors (TES), which are not only more complicated to manufacture but also require a Superconducting QUantum Interference Device (SQUID) in order to enable microwave multiplexing. This is particularly advantageous for astronomical applications where large-format ground-based cameras operating at optical, FIR, and mm wavelengths can be assembled at lower costs, within smaller configurations, and without extensive supporting hardware for readouts that can cause undesirable system heating (Sueno et al. 2022). Some scatter is introduced into the intrinsic resonance frequencies of some fraction of MKIDs owing to fabrication errors, which can result in frequency collision between two detectors. Such cases can be ignored during data reduction, modeled if they do not completely overlap in frequency space, or re-edited to adjust their capacitance further and shift their resonance frequencies.

The noise spectra of KIDs are characterized by a noise component that is well-modeled by a 1/f power law form. The uncorrelated component of the low-frequency noise features arises due to two-level system (TLS) effects where the movement of electrons between ground and excited states because of the electric field from the microwave readout feedline (Austermann et al. 2018; Sueno et al. 2022). The design of the MKID and the choice of materials used in their fabrication can mitigate TLS noise. Other sources of noise include the number and recombination of quasiparticles, whose mean lifetime sets the upper limit for the frequency response of MKID and the heat from low noise amplifiers (LNAs) used for signal amplification in readout hardware.

When observing from the ground, atmospheric turbulence adds another correlated low-frequency noise component to the raw data which also varies over time (Choi et al. 2020). In order to filter this contribution, data cleaning algorithms during reduction, modulation of the optical signal, and careful selection of on-sky mapping patterns can be employed.

The calculation of the optical signal from the raw KID readout is fully detailed in Ma et al. 2020 and Wilson et al. 2020. The actual quantity measured by KID readouts is the scattering parameter or transmission coefficient S_{21}. This is a complex number defined as $S_{21} = I + iQ$, where I and Q are the in-phase and quadrature components. The scattering parameter is not proportional to the astrophysical signal and therefore requires a transformation before the raw timestreams can be used to make flux-calibrated maps. The fractional shift between the frequency that the tone is being probed, f_p, at and the actual resonance frequency f_r for a given optical loading is the desired quantity. This is known as the detuning parameter x and can be determined by fitting S_{21} to a function parameterized as

$$S_{21}(f_p) = G\left(\tilde{S}_{21}(f_p, f_r, Q_r)\right), \tag{2.1}$$

where G is a complex gain function, $\tilde{S}_{21}$ is the canonical form of S_{21}, and Q_r is the resonator's quality factor. The gain function can be written as

$$G = \tilde{S}_{21} + Kf + M \tag{2.2}$$

and $\tilde{S}_{21}$ as

$$\tilde{S}_{21} = \frac{Q_r^2}{1 + 2Q_r xi}. \tag{2.3}$$

Consequently, x can be determined by fitting S_{21} to Equation 2.2 in the I-Q plane to give derived values for G, K, M, f_r, and Q_r. With the fit in hand, it can then be solved for the raw TODs during data reduction. The $\tilde{S}_{21}$ parameter can be re-expressed in terms of x and its corresponding quadrature channel r as

$$\tilde{S}_{21} = \tilde{I} + i\tilde{Q} = \frac{1}{x + ir}, \tag{2.4}$$

where comparisons with Equation 2.3 will reveal that $r = 1/2Q_r$. For a perfect fit to the resonator model, x and r will each be exclusively sensitive to the incoming signal and readout noise respectively, thus making r a useful metric for examining resonance model fit quality. An example of a single resonator's response in the I-Q plane and in the S_{21} and frequency space is plotted in Figure 2.3.

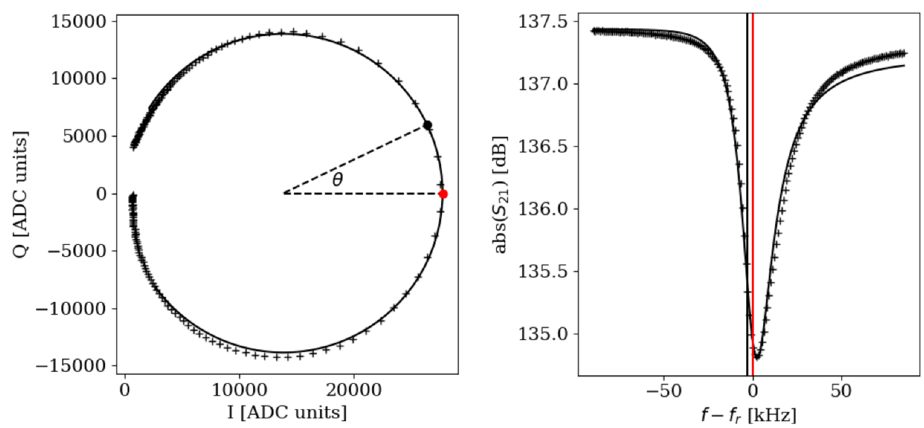

Figure 2.3. Figure from Wilson et al. 2020. Left: Response of a single KID in the I-Q plane. The black and red dots are the real resonator frequency and the probe tone frequency respectively. Right: The resonator scattering parameter S_{21} as a function of frequency for the same KID as the left panel. The red and black lines map to the red and black points.

TolTEC's MKIDs are built using titanium (Ti) and titanium-nitride (TiN) in a TiN/Ti/TiN substrate trilayer, which has been shown to reduce TLS noise (Austermann et al. 2018). The detectors are operated at a temperature of 100 mK. The pixels on the detector arrays consist of two MKIDs, each sensitive to a single linear polarization, oriented orthogonal to one another in pairs at angles of 0 and 90 for one configuration and 45 and 135 degrees for the other. This means each pixel measures two orthogonal linear polarizations for the same on-sky position from which the Stokes parameters can be derived. TolTEC's arrays consist of single-color MKIDs with detector counts chosen to optimize detector spacing for each band. With microwave multiplexing, the arrays are divided among 13 different readout networks with 7, 4, and 2 networks dedicated to the 273, 214, and 150 GHz arrays respectively. Figure 2.4 illustrates the physical design positions of the detectors on the arrays.

The TolTEC readout hardware computing system consists of 13 independent ROACH-2 Field Programmable Gate Arrays (FPGA) boards matched to each detector array. These are connected to two data recording and analysis computers that write out the raw time-ordered data (TODs) at a maximum sampling rate of 488 Hz.

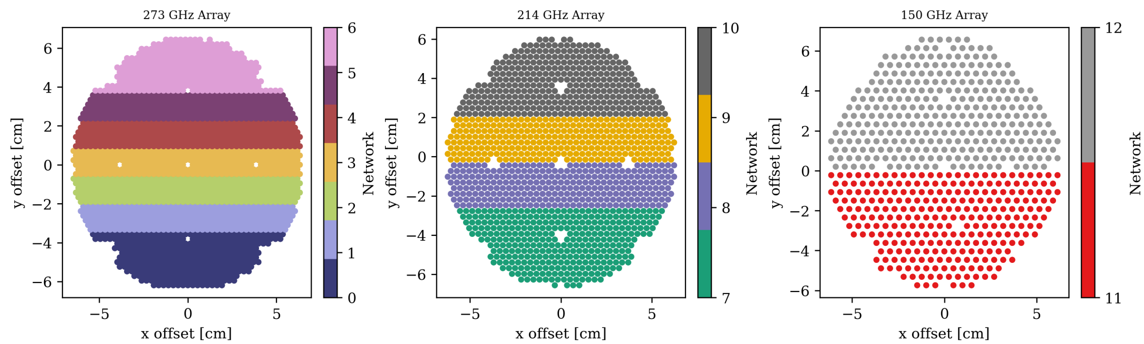

Figure 2.4. Design positions of the KIDs for the 3 TolTEC arrays, with 4012, 2534, and 1172 KIDS on each of the arrays. The colors represent the 13 different readout networks. Each point contains two orthogonal KIDs. The holes are the locations of the dark detectors that are not optically sensitive.

Each sample is time-tagged using a Global Positioning System (GPS) synchronous Packets Per Second (PPS) signal that is common to all networks to ensure each can be fully aligned temporally with one another.

2.2 The Large Millimeter Telescope

At 50 meters in diameter, the LMT (Hughes et al. 2010; Figure 2.5) is the largest steerable single-dish millimeter telescope operating at 1.1 mm in the world, which contributes greatly to TolTEC's high resolving power. The telescope is designed to support observations between 0.85 mm and 4.0 mm (353 GHz and 75 GHz). The primary mirror is an active surface consisting of 180 different segments whose positions can be adjusted independently in real-time to adjust for deviations in the primary mirror shape introduced by gravitational effects as well as stochastic variations from temperature and wind effects over the course of an observing night. The system is designed to operate with a surface RMS tolerance of 75 μm which is well below the requirements for TolTEC. The LMT's 2.5 m secondary mirror's position in the parallel and orthogonal directions relative to the primary mirror can be adjusted along with its angle to allow for telescope focusing and pointing alignment, as well as beam-switching observing modes (Mangum et al. 2007). Measurements of the telescope's pointing accuracy place the pointing and absolute pointing errors at an RMS of less than 2.5″.

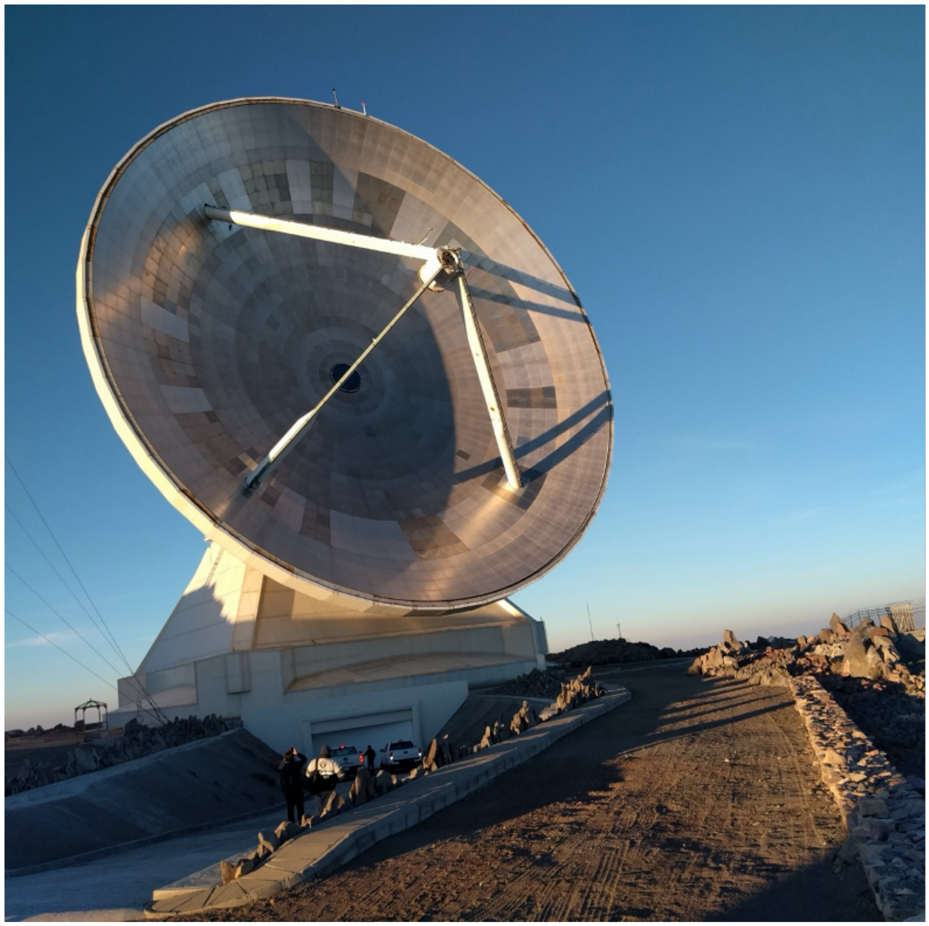

Figure 2.5. The Large Millimeter Telescope. Image credit: Gopal Narayanan.

All information related to the state of the telescope and its on-sky position and pointing corrections to its boresight in the relevant coordinate reference frames are tracked by the LMT's Telescope Control System (TCS). These are written to disk at a sampling rate of 20 Hz for each observation. Each sample within the telescope pointing vectors is time-tagged with the same GPS synchronous PPS signal as the TolTEC raw data samples such that they can be properly aligned during data reduction. Owing to the different sampling rates, however, an interpolation onto TolTEC's higher sampling frequency is required.

The LMT site includes a meteorological station and a radiometer system that measures the ambient temperature, pressure, humidity, wind speed, and direction, in addition to the zenith atmospheric opacity at 225 GHz (Ferrusca and Contreras R. 2014). The average 225 GHz opacity derived from radiometer measurements over the course of a full year is shown in Figure 2.6. The ideal observing season at the LMT site occurs between the months of October to May with 25%, 50%, and 80% of the time exhibiting opacities below 0.06, 0.10, and 0.28 respectively (Ferrusca and Contreras R. 2014; Zeballos et al. 2016; Bryan et al. 2018). By contrast, the summer months of June to September offer the poorest observing conditions at the site, with values of $\tau_{225GHz} > 0.15$ being more common. Under such conditions, observations above 150 GHz are challenging, though occasional nights of better weather conditions do occur.

2.3 Continuously Rotating Half-Wave Plate

The low-frequency noise features of KIDs degrade the ability to recover faint astrophysical signals, which makes measurements of the polarized emission from sources particularly challenging, as the fraction of polarized light in nearby star-forming regions is often on the order of a few percent of the total intensity signal. One technique to overcome this limitation is to modulate the incoming light in a way that shifts the desired signal to higher frequencies where the noise is lower. This modulation can be achieved by introducing a waveplate into the optical chain to shift the phase between two orthogonal linear or circular polarization vectors thus rotating the polarization angle. When the waveplate is rotated at a constant angular velocity, Ω_w, it boosts the frequency of the polarized signal to $4 \times \Omega_w$.

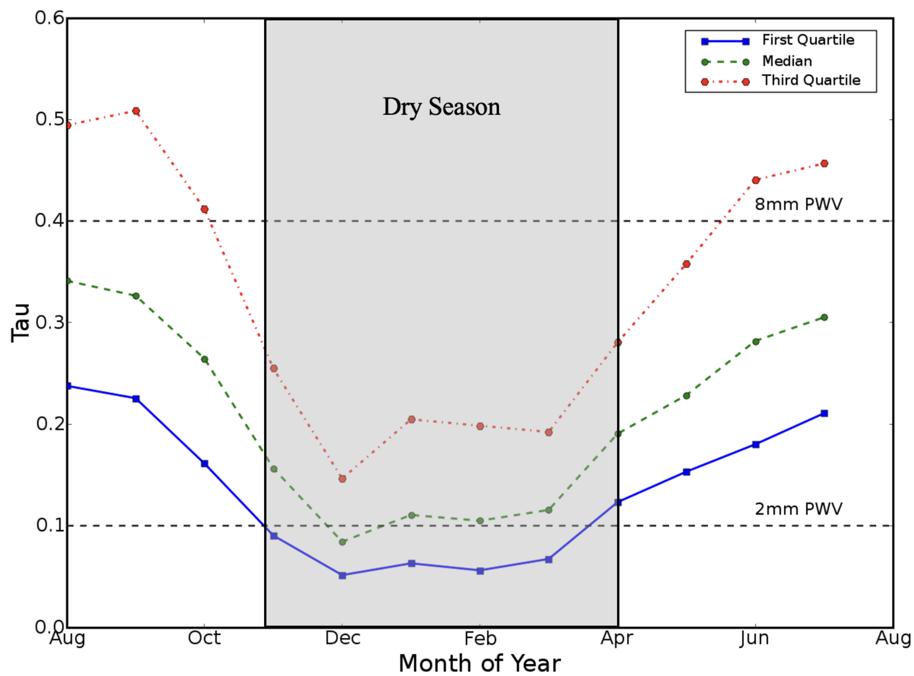

Figure 2.6. Measured average atmospheric opacity at 225 GHz at the LMT site for each month given different precipitable water vapor (PWV) levels. Figure adapted from Hughes et al. 2020 and http://lmtserver.astro.umass.edu/site.html

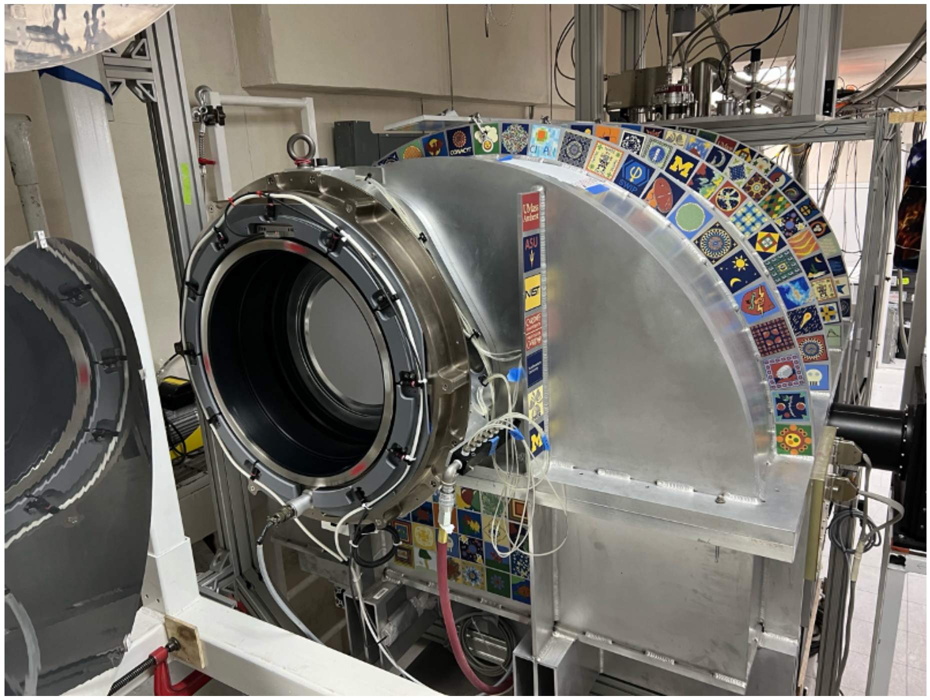

Figure 2.7. TolTEC's Continuously Rotating Half-Wave Plate mounted around the optical window on the front of the cryostat. Image credit: Dennis Lee

Two achromatic half-wave (HWP) plates, which modulate linear polarizations have been constructed for use with TolTEC (Lee et al. 2022). Their bandpasses are centered between TolTEC's own bands at 273 and 214 GHz and 214 and 150 GHz respectively and they are rotated at a frequency of 2 Hz by a Continuously Rotating HWP Rotator (CRHWPR). As shown in Figure 2.7, the CRHWPR is mounted outside the cryostat directly in front of TolTEC's observing window instead of being a cold optics component, thereby allowing the HWP to be removed to prevent intensity-only observations from being affected by their bandpass coverage. An example of the modulation induced in the data by the spinning HWP is shown in Figure 2.8.

The HWP orientation angle is tracked using an optical encoder, which measures the time of each complete revolution of the HWP. The angle relative to the initial encoder position can then be readily determined. The encoder values are written to disk at a sampling rate of 1 kHz by dedicated hardware for the CRHWPR system. As is the case with the TolTEC data and LMT metadata, the samples are time-tagged with a GPS synchronous PPS signal for eventual interpolation onto the TolTEC sampling time grid.

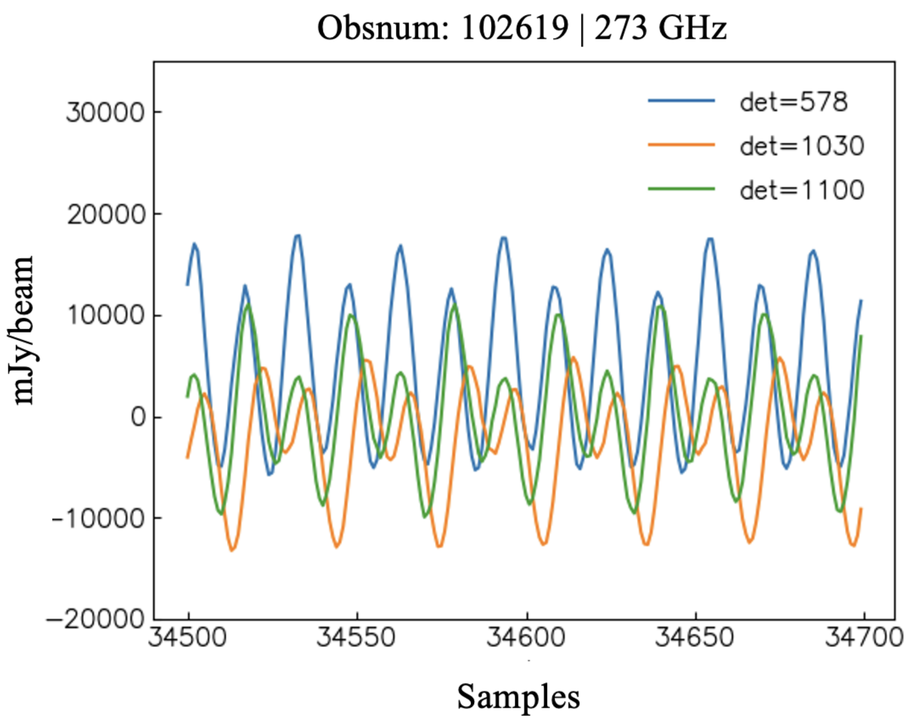

Figure 2.8. Timestreams from 3 detectors on the 273 GHz array showing amplitude modulation due to the rotating HWP. Image credit: Dennis Lee

2.4 TolTEC Data Handling

Network access to the LMT is severely limited, which has a significant impact on how the camera's status is monitored remotely and how its raw data products are managed and reduced. TolTEC's maximum sampling frequency of 488 Hz translates to a data rate of approximately 30 MB/s or nearly 2 TB per observing night. A 30-minute observation utilizing all three detector arrays outputs a total of 54 GB of data, which cannot be readily copied off-site during an observing night. For this reason, data reductions and the generation of necessary calibration data products and quick-look maps for data validation must occur on-site. Two dedicated workstations have been installed to carry out full data reductions, perform detector setup and characterization, and host data visualization servers. Each has a single socket 32-core Intel Xeon Gold 5218 processor and 93 GB of DDR4-2933 RAM running the Ubuntu 20.04 operating system and are referred to as TACO and TACA. These therefore represent the minimum required specifications that the TolTEC data reduction pipeline must adhere to. Data is copied simultaneously from the readout hardware to both TACO and TACA to allow for parallel data reductions and analysis to take place.

For off-site data reductions, in-depth analysis, scientific verification efforts, and software development, members of the TolTEC collaboration have access to 3 high-performance computing clusters (HPC). The primary HPC used for commissioning data exploration is the Unity[1] high-performance computing cluster located at the Massachusetts Green High Performance Computing Center (MGHPCC), on which there are 7 dedicated compute nodes for TolTEC data processing. Each node consists of two 32-core Intel Xeon Gold 5218 processors with 384 GB of DDR4-2666 RAM and are also running on Ubuntu 20.04. Unity is connected to the Northeast Storage Exchange (NESE) data lake where there are 400 TB of storage available for TolTEC data. TolTEC software has also been installed and utilized on the Quest HPC at Northwestern University and the National Institute of Astrophysics, Optics and Electronics's (INAOE) Mextli HPC. The ability to take advantage of these HPC resources through parallelized data reductions is a central consideration in the design of TolTEC's reduction pipeline.

2.5 Instrument Operations

A single observing night with TolTEC consists of not only science observations but also on-sky detector setup and characterization processes, as well as science readiness

[1] https://unity.rc.umass.edu/

observations. These are usually carried out in the order of (1) detector identification and tuning, (2) flux calibration beammaps, (3) focusing, (4) first-order astigmatism deformation corrections of the LMT's primary mirror, and (5) pointing offset determination.

2.5.1 Detector Setup

Before observations of astrophysical sources can start, the resonance frequencies of each KID must be measured for the current level of background loading. With the telescope pointing at zenith, a sweep of the entire frequency range occupied by the KIDs is performed. Each detector network is driven with 1000 separate probe frequencies over a range of 500 MHz which are then shifted in 2 kHz steps to sample the full range. After a peak finding algorithm is applied to acquire a rough estimate of the resonance frequency location, a denser grid of tones spanning 175 kHz around each peak is used to sample the resonator shape. They are then fit to Equation 2.3 and the values are written to disk to solve the models for the detuning parameter x in observations of sources. An example sweep of the first readout network of the 273 GHz array is plotted in Figure 2.9. The initial coarse sweep can be carried out once per night or shared across nights if the observing conditions have not changed significantly. The denser sweep or detector tune is performed prior to each observation as large shifts in loading that occur when moving between sources at radically different elevations can result in poor re-recovery of the tone frequencies.

2.5.2 Observing Strategies

One consequence of the sweep and tune procedure is that TolTEC and its KIDs are not sensitive to temporally and spatially invariant signals. As a result, the LMT cannot remain stationary and integrate on a particular region of the sky nor can it alternate between an on-source and blank sky field when observing with TolTEC. Instead, the on-the-fly mapping strategy (OTF) is used in which the LMT is continu-

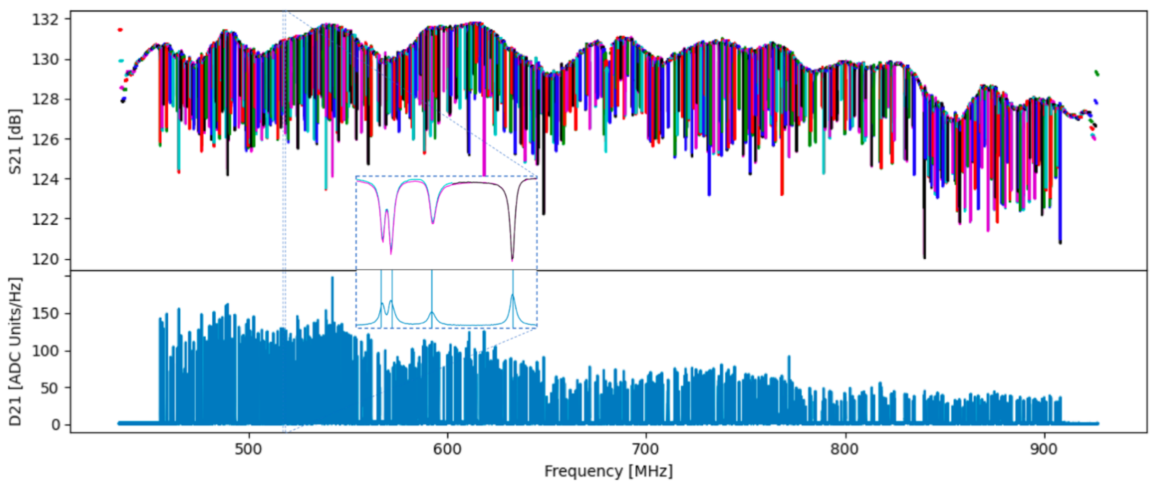

Figure 2.9. Figure from Wilson et al. 2020. Top row: A network frequency sweep for detector network 0 on the 273 GHz array intended to measure the frequency locations of each KID. Bottom row: The derivative of S_{21} with respect to frequency for the same sweep.

ously slewed over the source field along a pre-computed trajectory or mapping pattern (Mangum et al. 2007). The choice of what map pattern to use for a desired region size, integration time, source type, and flux depth is a complex parameter space and can significantly impact observing efficiency, coverage of the field, and the quality of the final maps produced. The sky signal should be sampled above the Nyquist sampling interval of $\lambda/2\mathrm{D}$ as under-sampling can introduce aliasing artifacts and enlarge the beam along the scan direction which adds noise and limits the recovery of small-scale features. Improper or insufficient removal of the low-frequency correlated noise from the atmosphere results in non-uniformity or striping in the maps along the telescope's direction of motion (Poletti et al. 2017; Choi et al. 2020). This noise striping can be mitigated by using OTF strategies that pass over the source from multiple different directions to improve the "cross-linking" between scans (Tegmark 1997). For mapping strategies without good cross-linking, observations of the same source can be taken at different times of the night to take advantage of sky rotation of the mapping pattern or coadded observations with rotated versions of the same pattern.

TolTEC uses two primary mapping patterns for most of its observations which were informed based on previous experience with the AzTEC camera (Wilson et al. 2008) and are illustrated in Figure 2.10. The first is the Raster map, which consists of long, straight scans in one direction interspersed with short steps or "turnarounds" in the orthogonal axis. These are ideal for mapping regions larger than the instrument's field of view. For a constant mapping speed, the overhead of a Raster map is inversely related to the length of the scan owing to the relatively constant time spent at the turnarounds. Unless Raster maps are coadded, they do not have good cross-linking and are more susceptible to scan synchronous noise effects than other patterns.

The second pattern uses the parametric Lissajous curve to generate a trajectory. The position as a function of time t of the Lissajous pattern in the arbitrary coordinate reference frame (x,y) is given by

$$\begin{aligned} x(t) &= x_{\text{len}} \sin(\omega_x t + \delta) \\ y(t) &= y_{\text{len}} \sin(\omega_y t), \end{aligned} \tag{2.5}$$

with x_{len} and y_{len} determining the overall size of the mapped region, ω_x and ω_y controlling the speed, and delta being a phase offset that tunes the angle between the oscillations. Unlike the Raster, the Lissajous pattern is more well suited for fields at the scale or smaller than the field-of-view. The pattern demonstrates a much higher degree of cross-linking and incurs no overhead. However, depending on the parameter choices of Equation 2.5, the Lissajous pattern can result in uneven coverage with a relatively larger fraction of the time spent near the edges of the map.

Combinations of two Lissajous as well as a pattern consisting of a Raster modulated by a Lissajous, referred to as a Rastajous map, are also used to improve coverage and to take advantage of the benefits provided by both patterns.

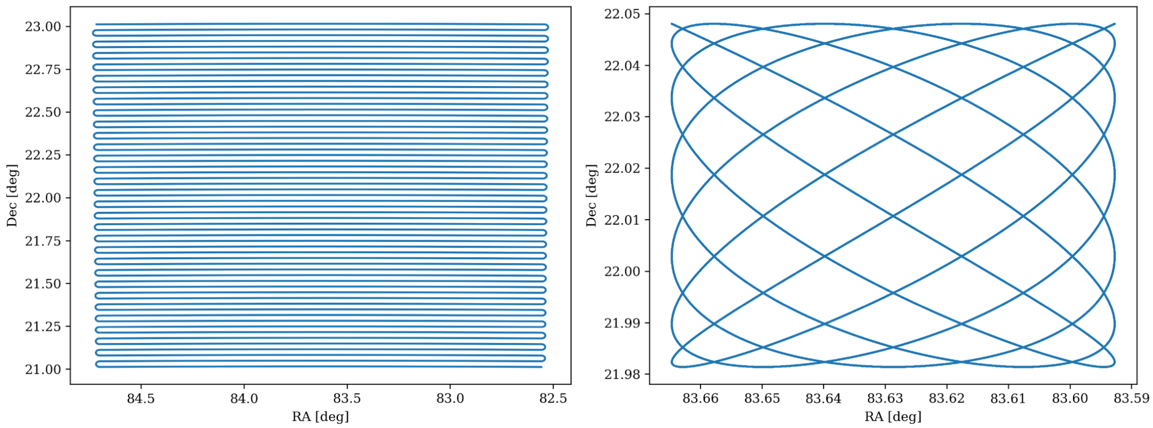

Figure 2.10. Left panel: Telescope boresight trajectory for a simulated Raster map. Right panel: An example of a Lissajous mapping pattern over a smaller field.

2.5.3 Beammaps

Observations that are intended to measure flux calibration factors, fit the on-sky detector centroids and beam sizes, and flag false or poorly behaving detectors are referred to as beammaps. The deviations in the actual or estimated resonance frequencies of TolTEC's KIDs from their design specifications that arise from fabrication errors and measurement noise make mapping a resonator to an actual KID on the detector arrays a non-trivial task. This mapping is a required quantity for coadding individual detectors into a single map. For this reason, it is necessary to empirically measure the detector positional offsets relative to the telescope boresight through an observation of a bright (typically >1 Jy/beam at 273 GHz) point source. Solar System bodies, including the planets Neptune and Uranus, the moons Titan and Callisto, as well as asteroids/dwarf planets such as Ceres, 4 Vesta, and 2 Pallas represent idea flux calibration sources at millimeter wavelengths owing to their well-constrained emission properties and brightness. The Submillimeter Array (SMA; Ho et al. 2004) maintains a public database of submillimeter/millimeter calibrators that are routinely re-observed at 870 μ m, 1.1 mm, and 1.3 mm. These sources consist of quasars, blazars, and BL Lacertae objects which were used as TolTEC beammapping targets during commissioning. Such sources can be highly variable due to AGN activity, however, thus making them less ideal than Solar System bodies if they were not observed with the SMA near in time to the TolTEC observation.

A beammap consists of a dense Raster of 300 azimuth scans at $5'$ in length, with elevation steps of $1.5''$. The telescope is slewed at a rate of $50''$/s and takes approximately 30 minutes to complete. This ensures that the source is more than Nyquist sampled in the elevation direction with approximately 3-7 scans crossing the source in each of the TolTEC bands. When regular observing begins with TolTEC, ideally at least two beammaps will be taken each night with one near the beginning and one close the end to bracket any other observations taken during that night and allow for the interpolation between beammapping derived data products as a function of observing time.

Since reductions of beammaps must be carried out without the metadata provided by beammaps themselves and because they are used to make per-detector maps and derive per-detector characteristics for use in the other observation types, the data reduction process differs considerably from that of the science and other calibration observations, making use of an iterative mapmaking technique by default. The unique reduction engine used for beammaps is described in Section 3.12.1.3.

2.5.4 Focus Observations

Variations in the ambient temperature and windspeed over the course of an observing night have a strong effect on the quality of observations and introduce both systematic and random deviations into the secondary mirror position. Consequently, the telescope will drift out of the optimal focus configuration and the beam shapes will become distorted. Focus observations are carried out at the start of the night and are repeated when pointing maps are determined to be out-of-focus by eye. Approximately 9 separate maps of the same source are made in succession with the vertical position of the secondary mirror shifted between -4.0 and 4.0 mm. Each map is a $4' \times 4'$ double Lissajous pattern with an integration time between 30-60 seconds depending on the source flux. In the data reduction process, the source at 273 GHz in each of the maps is fit to a 2D Gaussian and the source amplitudes across all the maps are subsequently fit to a parabola as a function of the secondary mirror offset. The offset where the parabola is at its peak is taken to conform to the optimal focus location. A set of fitted focus maps is shown in Figure 2.11.

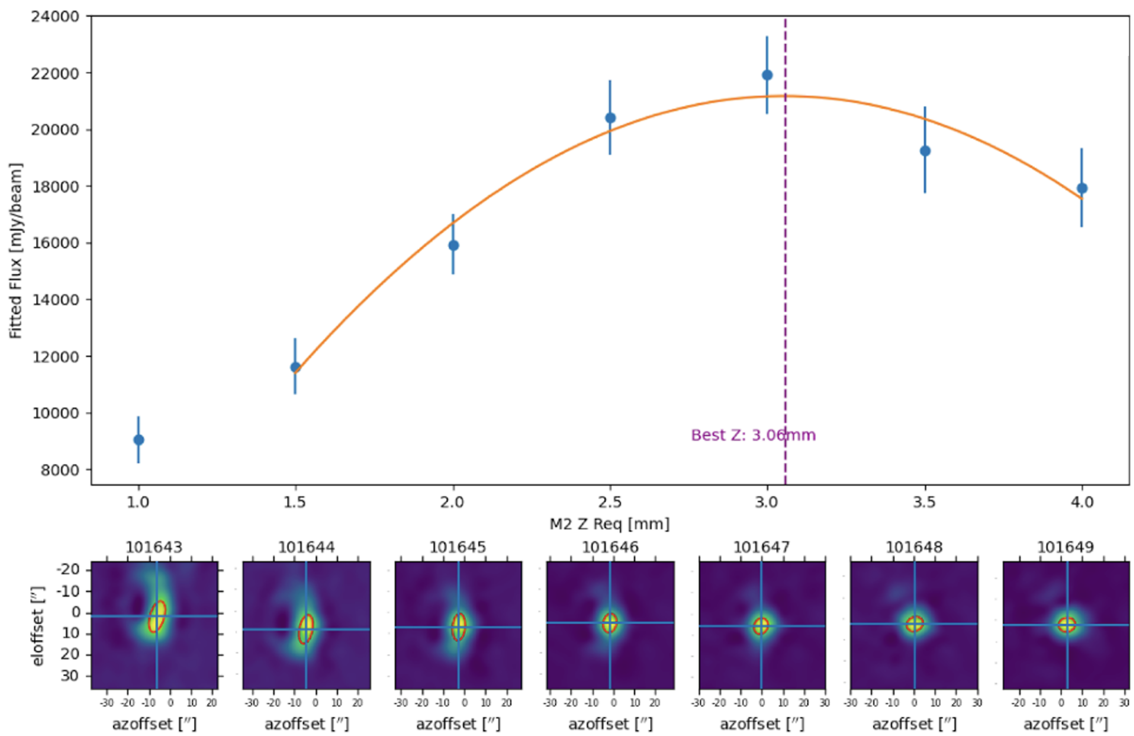

Figure 2.11. Top panel: Parabolic fit to source amplitudes derived from 2D Gaussian fits to observations of point sources. The position of the LMT's secondary mirror was varied for each point. The best focus position is marked by the dashed line. Bottom Panel: Maps of the source at M2 position.

2.5.5 Astigmatism Observations

The setup of the astigmatism correction observations is nearly identical to that of the focusing observations, with the same mapping pattern and data reduction process being used. Instead of shifting the secondary mirror position, the shape of the primary mirror is deliberately deformed using the LMT's active surface to adjust its level of astigmatism. A series of 4 maps is then taken at deformations of -300 and 300 mm and fit in the same way as the pointing maps. To correct higher-order deviations in the antenna shape, out-of-focus holography techniques are required, which fit the beam pattern in a series of focus maps to Zernike polynomials to derive a more precise deviation model. Figure 2.12 shows the fitted amplitudes from a set of astigmatism observations.

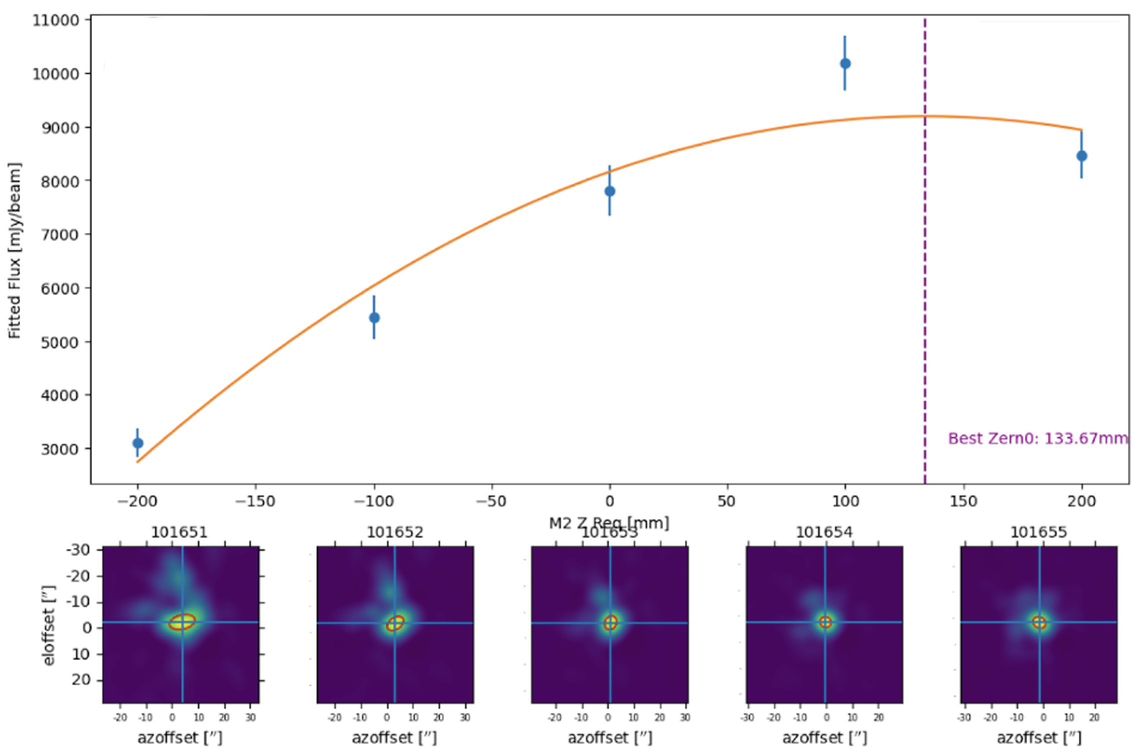

Figure 2.12. Same as Figure 2.11, but for an astigmatism correction observation where the primary mirror shape is modified instead. The best astigmatism value is marked by the dashed line.

2.5.6 Pointing Correction Observations

Pointing observations are carried out immediately before and after science observations so that measured offsets can be interpolated between and to confirm the focus of the telescope. The mapping pattern is identical to the focus and astigmatism maps. Focus, astigmatism, and pointing observations all utilize the same reduction framework which differs from both the beammapping and science processing. The source is similarly fit to a 2D Gaussian model in the Azimuth and Elevation reference frame, with the centroid position providing the pointing offset estimate. An example of a quick-look map of a pointing observation is illustrated in Figure 2.13.

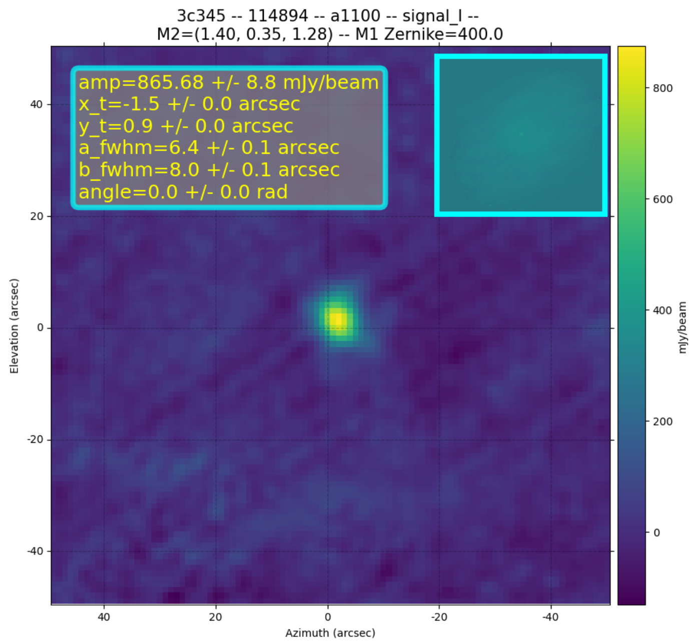

Figure 2.13. Quick-look 273 GHz map for a pointing observation of 3c345. The map has been smoothed from its native resolution with a $\sigma = 2.0''$ Gaussian kernel. Parameters are derived from a 2D Gaussian fit to the smoothed map. The cutout is a sum of the signal and uncertainty maps.

CHAPTER 3

THE TOLTEC DATA REDUCTION PIPELINE

The extraction of the astrophysical signal from raw data in the time domain requires the development of specialized software tools tailored to handle the needs and idiosyncrasies of the astronomical instrument or instruments in question, their data products and associated metadata, the observation setup, and the sky and source properties. This process usually involves a transformation of the raw data streams into a different coordinate reference frame that results in a large compression of its size and dimensionality and is thus representative of a "reduction" in the complexity of the data sets. These data reduction software packages or pipelines are usually independent of software developed with the goal of performing scientific analysis and instead provide the inputs to those tools. As a result, reduction pipelines are a necessary component of any instrument and require extensive and careful planning in conjunction with knowledge of the instrument design and expected performance prior to the acquisition of data.

Ideally, reduction pipelines should recover an unbiased estimate of the source signal and the uncertainty in its measurement and remove or filter all unwanted instrumental effects, noise contributions, and artifacts to the data as part of the process. For ground-based observing, this is particularly challenging as a consequence of the varying foreground atmosphere and background loading environment at the telescope, in addition to the potential for external sources of interference or signal pick-up from the ground. The optimal pipeline would also not introduce artifacts of its own and have an identity transfer function for all astrophysical source types. This is similarly difficult as the algorithms and techniques used to filter undesired contaminants to the data will also attenuate the signal in complex, often nonlinear ways that vary with the source brightness, morphology, and scale. As such, the pipeline's transfer function must be estimated and be well-understood and characterized such that its impact can be corrected for in post-processing stages.

For ground-based imaging cameras operating at submillimeter and millimeter wavelengths, the primary data products of their data reduction pipelines are usually the two-dimensional maps of the source as a function of sky position that are derived from arrays of incoherent detector timestreams. Pipeline development has a long history in this field and includes such software packages as the SCUBA User Reduction Facility (Jenness and Lightfoot 1998; Jenness and Lightfoot 2014), Macana for AzTEC (Scott et al. 2008), the Sub-Millimetre User Reduction Facility (SMURF) for SCUBA-2 (Chapin et al. 2013; Jenness et al. 2013), the NIKA processing pipeline (Catalano et al. 2014; Adam et al. 2014). The advent of second-generation cameras like SCUBA-2, NIKA2, and now TolTEC has increased the requirements for data reduction software by a wide margin in terms of having to accommodate newer detector technologies like KIDs, simultaneous multi-color imaging, polarization sensitivity, and a much larger data volume.

As a point of comparison, TolTEC's predecessor, AzTEC, imaged the sky at a single band using 144 bolometer elements with a sampling frequency of 64 Hz; TolTEC's detector count is a factor of 53 times larger, divided across 3 arrays for which data is collected in parallel at a sampling rate nearly 8 times faster. The resulting increase in the data rate and output file sizes is a factor of 400. The performance of computing hardware has not seen a similar growth in the same time frame. TolTEC also supports additional hardware like the HWP, and its raw data products use inputs from KID model fitting for every observation. This evolution necessitates a rethinking of fundamental design choices in the pipeline design decisions from the ground up for current and future instrumentation. Access to high-performance compute clusters and the parallelization of reduction algorithms across compute threads is now a critical consideration, though these resources are not always usable, such as when reducing data on-site while observing.

A second concern of the increased detector count other than the added computational requirements is the need to perform validation of data at various stages of processing. As millimeter astronomy continues its shift into the “big data” realm, it becomes increasingly difficult to examine observation and instrument characteristics at the level of individual detectors. Data cuts and flagging to remove false or outlier detectors are now most optimally done automatically with summary statistics and metrics of detector populations becoming useful.

With these considerations in mind, we have developed the standalone high-performance C++ data reduction software package `Citlali`[1] to reduce TolTEC raw data timestreams into science-ready maps for all categories of TolTEC observations. This chapter will give an overview of the pipeline's current version (v4.0) describing its architecture and detailing its reduction algorithms and their arrangement into separate sub-pipelines to handle the unique needs of beammmaping, focus, astigmatism, pointing, and science observations. The discussion here supersedes the overview of the pipeline given in McCrackan et al. 2022 which covers the depreciated v2.0 specifically.

3.1 TolTEC Software Architecture

`Citlali` is one component of the overall TolTEC software architecture that has been constructed to meet TolTEC's diverse data handling, analysis, simulation, and visualization needs. The software stack can be broadly divided between the high-

[1]https://github.com/toltec-astro/citlali

level Python data management framework `TolTECA`[2] (Ma et al. 2020) and the low-level C++ data reduction software, which includes the data reduction pipeline `Citlali` and the KIDs processing library `kidscpp`, which carries out the model fits for the KIDs. The low-level C++ data reduction software share between them a set of libraries for managing CMake modules for compilation (`tula_cmake`) and assembling frequently used algorithms and functions (`tula`). All TolTEC data processing and reduction software is fully open-source and hosted on its own repository on the TolTEC Project's GitHub.

The `TolTECA` Python framework is described in detail in Ma et al. 2020. Briefly, it is a pip installable package that serves as the higher-level user interface and data manager for most of the TolTEC software packages. Database management, automatic collection of raw and metadata products for the execution of `Citlali` and the hosting of web-based visualization tools developed using the Plotly Dash[3] library are handled by `TolTECA`. It also can generate simulated raw data using real or synthetic KIDs model parameters of Equation 2.4, an on-the-fly mapping pattern, 1/f and readout noise contributions, and a realistic atmosphere model derived from the Time Ordered Astrophysics Scalable Tools 3 (TOAST3) software framework (Kisner et al. 2021).

3.2 `Citlali` Codebase

The `Citlali` codebase is written entirely in the C++ programming language to take advantage of the superior performance and memory management options relative to a scripted language like Python. This choice is heavily motivated by the need to reduce calibration and science readiness observations several times each observing

[2] https://github.com/toltec-astro/tolteca

[3] https://plotly.com/dash/

night, where reduction performance on a single workstation is essential for maximizing observation efficiency with the LMT and TolTEC. C++ allows for more explicit control over memory allocation and deallocation with pointers and pass-by-reference and many C++ libraries are built from the ground up to minimize the overall memory usage and avoid the creation of temporary copies. This lowers the threshold for system requirements on which `Citlali` can be run while further optimizing performance.

Parallelization of code occurs at a very low level in C++ and can be much more efficient for CPU-bound tasks than Python's own implementation. `Citlali` is a multi-threaded software package and can therefore make efficient use of the resources available on HPCs while still maintaining high performance on CPU and memory-limited hardware.

Potential downsides to the use of C++ include the lower degree of portability relative to Python, its more complex syntax, and the smaller pool of developers who are familiar with the language. For these reasons, `Citlali` employs modern C++ programming practices to make extensions to the code and its readability easier. It uses automatic build systems to simplify cross-platform compilation, expression template metaprogramming to create type-independent functions, and external well-documented libraries for critical reduction steps. The codebase is written with an object-oriented focus and in a hierarchical and modular sequential pipeline structure such that new reduction steps can added with little adjustment to other components.

Due to the emphasis on modular code design, much of the codebase is shared across the different sub-pipelines or reduction engines that are written to reduce specific categories of TolTEC observations and produce their required outputs. These reduction engines consist of the science (`citlali::lali`), pointing (`citlali::pointing`), and beammap (`citlali::beammap`) classes which compose the timestream and mapmaking algorithms together into different configurations. Other than the reduced level of code duplication, a key advantage of this is that any observation can be reduced in any of the different modes and that there is a single unified user interface – the `Citlali` configuration file – that is the same across every observation.

`Citlali` makes extensive use of open-source software libraries which provide further optimizations for computationally expensive reduction steps, access to special functions, and high-level abstractions of lower-level routines to streamline development and improve code legibility. In selecting libraries, a preference was given to those that are currently in active development, are well-documented, and are used by other projects both in astronomy and industry. Libraries related to data storage, performance, and parallelization that are used include:

- **Eigen v3.4.0+ (Guennebaud, Jacob, et al. 2022):** A critical design decision early in the development of `Citlali` was the choice of storage structures for vectors and arrays of timestream and map data. C++ includes native array and vector support, but these are limited in their ability to handle large multi-dimensional data structures, have poor or no dynamic memory allocation support, and possess insufficient built-in support for matrix arithmetic and linear algebra. For this reason, `Citlali` uses the Eigen library as its default method for storing non-scalar numeric variables. Eigen is a high-level C++ header-only template library built for vector and matrix storage, manipulation, and linear algebra. Its vector, matrix, and array classes support containers of arbitrary sizes for most standard data types. Both dense and sparse structures are implemented. Higher dimensional tensors are available and are used occasionally in `Citlali`. Eigen uses expression template metaprogramming and under-the-hood design optimizations, allowing for functions and algorithms that use it to take advantage of lazy evaluation, loop unraveling, and automatic vectorization of Single Instruction/Multiple Data (SIMD) instruction sets. Modules for Fast Fourier Transforms (FFTs), matrix decompositions, polynomial solvers, and multidimensional spline fitting, and interpolation are available.

- **Ceres Solver v2.1+ (Agarwal et al. 2022):** Reductions of all calibration observations performed by TolTEC require a fit of the observed point source to determine its amplitude, offset, and FWHM. For beammaps, a fit is performed for every detector. Citlali uses Ceres Solver for its source fitting, which is a non-linear least squares library developed by Google that provides support for bounded or unconstrained problems, automatic and numeric differentiation, loss functions for handling outliers, and covariance matrix evaluation for uncertainty estimation.

- **Sparse Eigenvalue Computation Toolkit as a Redesigned ARPACK (SPECTRA[4]) v1.0.1:** The atmospheric removal algorithm implemented in `Citlali` uses a principal component analysis-based approach and is its most computationally expensive timestream reduction stage. It requires the calculation of many of the largest eigenvalues in the detector timestream correlation matrices. Although Eigen natively supports eigenvalue decomposition of its matrices, the SPECTRA library can compute a subset of the smallest or largest eigenvalues in a matrix without performing a full decomposition which is more efficient.

- **Generic Reusable Parallel Pattern Interface (GrPPI) v.0.4 (Rio Astorga et al. 2017):** `Citlali` uses a range of parallel patterns at different parts of the pipeline when working with either timestream or map data. GrPPI is a header library that allows for complex parallelization patterns to be composed together in pipeline-like structures, including simple parallel loops as well as parallel farms that spawn new workers based on inputs from a sequential data stream. It supports most major parallel programming frameworks including ISO C++ Threads, OpenMP, Intel Thread Building Blocks (TBB), and FastFlow.

[4]https://spectralib.org/

- **Boost[5] v1.77+:** The Boost Random Number Library and Special Functions Library are employed for generating random deviates for jackknifed noise maps and for computing Bessel functions for timestream filtering and mapmaking respectively.

- **FFTW v3.3.9+ (Frigo and Johnson 2005):** Fast Fourier Transforms are required for computing the power spectral densities of timestreams and maps and are also used in the convolutions carried out in map filtering. After benchmarking, it was found that FFTW outperformed Eigen's own FFT implementation while also enabling under-the-hood parallelization for large matrices.

- **Mlinterp (Azimzadeh 2017):** This provides support for efficient linear interpolation in arbitrary dimensions and is used for timestream alignment between TolTEC raw data, telescope boresight vectors, HWP orientation angles, and pointing offset measurements.

Other external libraries are used in `Citlali` for the purpose of data I/O and command line output and logging. These include the network Common Data Format-4 C++ (netCDF4; Rew and Davis 1990) API for reading in raw data files, the CCfits[6] library for writing output maps into the Flexible Image Transport System (FITS; Wells et al. 1981) format, and the yaml-cpp[7] parser library for configuration file input. Command line output is performed by the combination of the `spdlog`[8], `fmt`[9],

---

[5]https://www.boost.org/

[6]https://heasarc.gsfc.nasa.gov/fitsio/CCfits/

[7]https://github.com/jbeder/yaml-cpp

[8]https://github.com/gabime/spdlog

[9]https://github.com/fmtlib/fmt

and `re2`[10] libraries to provide formatted info, warning, debugging, and error output, as well as regular expression handling.

3.3 Data Streaming and Parallelization Strategy

With the exception of observations with very short integration times like pointing maps, the size of most TolTEC raw data products is a significant fraction of the available RAM on the timely analysis workstations at the LMT. Many of the reduction algorithms, despite optimizations, still incur additional memory overhead. Running on multiple threads and enabling map coaddition and the generation of jackknifed noise maps require additional memory, though this additional load is typically much lower than the size of the raw data itself. To ensure future scalability of the code, `Citlali` uses a data streaming model that does not necessitate the entire raw timestreams to be stored in memory for reductions other than the beammapping engine. In this model, the raw timestreams are subdivided into smaller chunks that are passed sequentially into a parallel farm pattern over different compute threads as they are read in from disk. This has the benefit of minimizing I/O bottlenecks as one or more chunk is being reduced while another is being populated. Consequently, the parallelization efficiency of `Citlali` is influenced by the relative difference between the disk read speed to the amount of time required to reduce an individual chunk.

An alternative scheme would be reducing each chunk sequentially but instead, parallelizing over individual detectors when possible. Most timestream reduction algorithms involve a loop over detectors and this strategy would have potential benefits if the number of chunks in an observation is less than the number of available physical compute nodes. Except in a few instances, however, this is significantly less efficient due to the more fine-grained implementation required for such parallelization. Par-

[10]https://github.com/google/re2

allelized code introduces a small amount of thread startup and communication overhead, which can become a significant limiting factor if they are used too frequently. Per-detector parallelization is used for one of the mapmaking algorithms.

Per-map parallelization is also used when working in the map domain where chunk-based parallelization is no longer possible. This reduces to a per-detector parallelization scheme in the beammap reduction mode.

3.4 Software Compilation and Installation

`Citlali` uses the CMake[11] v3.20+ build system to automate the generation of the required C/C++ build files across different platforms. As the kidscpp library is a required component of `Citlali`'s timestream reduction, it is fetched remotely during setup, which in turn fetches the `tula` and `tula_cmake` helper libraries. The `tula_cmake` library includes all CMake modules for compilation of the external libraries used in `Citlali`. When possible, the CMake modules provide 3 options for finding libraries, including finding an existing installation, downloading the library's source files from its repository or host server, or using the Conan[12] C/C++ software package manager v1.55.

`Citlali` has been successfully compiled on Ubuntu 20.04, CentOS Linux 7.5, Red Hat Enterprise Linux 7.9, and Intel-based MacOS 12, with either the GCC 11+ and LLVM 13+ compilers.

3.5 Runtime Configuration

`Citlali` accepts a standardized configuration file written in the YAML data serialization language as its sole input to modify reduction parameters at runtime.

[11] https://cmake.org/

[12] https://conan.io/

YAML[13] files are a minimal syntax, human-readable file format that support inputs in most standard numerical and character data types in the form of scalars, lists, or dictionaries within a hierarchical key-value node structure. All `Citlali` reduction use the same configuration file regardless of the observation goal and the reduction engine can be selected through a single configuration option. Mirroring the modular nature of `Citlali` most reduction steps can be toggled on or off at runtime. Part of the configuration file is shown in Figure 3.1.

```yaml
runtime: # parameters to control how citlali runs
  verbose: true # extra debug output in terminal (same as -l debug).  adds summary files to obsnum/logs/ directory.
  meta:
    version: tolteca... # extra info from tolteca
  n_threads: 1 # number of threads to use
  output_dir: /path/to/redu/directory/ # path to output directory
  parallel_policy: seq # parallelization policy (seq=sequential, omp=OpenMP parallelization)
  reduction_type: science # reduction engine (science, pointing, beammap)
  use_subdir: true # create "redu00, redu01,..." sub-directories.  increments automatically.
timestream: # timestream reduction parameters
  type: xs # type of TOD (xs, rs, is, qs)
  enabled: true # enable or disable TOD processing (to be implemented)
  precompute_pointing: false # compute detector pointing once (ignored for now)
  chunking: # how to chunk timestreams
    force_chunking: false # ignore map pattern chunking and use config file chunk length
    length_sec: 10.0 # chunk length for lissajous/rastajous maps (ignored if raster map)
  polarimetry: # polarized reduction
    enabled: false # enable or disable polarized reduction
    ignore_hwpr: auto # override automatic hwpr finding (true, false, auto)
    grouping: fg # use all matched detectors (fg) or only matched detectors with a pair (loc)
  output: # controls options for writing out TOD as netCDF file
    subdir_name: null # optional subdirectory name for tod output files (ignored if null)
    stats: # controls for what variables to include in stats file
      eigenvalues: false # include first 64 eigenvalues in stats file
  raw_time_chunk:
    despike: # remove cosmic ray spikes
      enabled: false # enable or disable despiking
      min_spike_sigma: 8.0 # minimum spike sigma value (n x sigma of TOD)
      time_constant_sec: 0.015 # spike decay time
      window_size: 32 # number of samples to flag on either side of spike (ignored if TOD filtering)
    downsample: # downsample the data
      enabled: false # enable or disable downsampling
      factor: 1 # downsample factor (integer only, ignored if <=0)
      downsampled_freq_Hz: 0 # specify downsample frequency (used if factor<=0, must be <= raw sampling factor)
    filter: # tod filtering
      enabled: false # enable or disable tod filtering
      a_gibbs: 50.0 # Kaiser-Bessel filter shape param
      freq_high_Hz: 16.0 # upper freq limit
      freq_low_Hz: 0.0 # lower freq limit
      n_terms: 32 # size of filter in samples
```

Figure 3.1. Example of the `runtime` and `timestream` configuration YAML nodes from the `Citlali` configuration file.

No default values are assumed for configuration parameters in `Citlali`. This was a specific choice to prevent reductions using ill-fitting parameter choices outside of the

[13] https://yaml.org/

user's knowledge, particularly during TolTEC's commissioning phase when optimal reduction parameters have not yet been determined.

3.6 Pipeline Inputs

`Citlali` requires an assortment of file inputs to reduce data and create maps, all of which are provided to `Citlali` within its configuration file. Some are optional depending on the observation goal and the reduction type selected. These inputs include:

- TolTEC raw data files from each of the 13 readout networks
- LMT telescope file
- HWP data file
- KIDs model parameter table
- Array Property Table
- Reduced `Citlali` maps
- Synthetic source maps

The TolTEC raw data files are stored in the netCDF4 file format and include the full timestreams for every detector on the corresponding readout network. These are the files output directly from the readouts without any pre-processing applied. Files from an arbitrary number of different observations can be added which Citlali will then reduce sequentially. All networks do not need to be provided even if included in the KIDs model file and APT, as `Citlali` will filter out those networks at runtime.

The LMT telescope file is in the netCDF4 file that includes the information required to reconstruct the boresight pointing of the telescope in the Az/El and Ra/Dec coordinate frames. Each observation or set of TolTEC raw data files has a single matching telescope file. The measured optical depth at 225 GHz from the LMT site's radiometer is stored here.

The HWP data file includes the encoder angle timestream and is an optional input to `Citlali` as the HWP can be removed. It is also stored in the netCDF4 format. The HWP, the telescope, and the raw data files each include the GPS PPS signal measurements for every data sample.

The KID model parameter table is an output of the standalone `kidscpp` library and includes fitted model parameters from equations 2.2 and 2.3 for each detector as derived from the tune observations. It is written in the Enhanced Character-Separated Values (ECSV) format which merges a CSV data table with a YAML metadata header. These model parameters are kept separate from the raw data files to allow for updates to the model in the future without requiring potentially risky editing of the raw data files.

The final file input file to `Citlali` for the standard reduction modes is the Array Property Table (APT) which is unique in that it is both an input and an output of `Citlali` itself. It is an ECSV table that includes per-detector characteristics and is a primary output of the beammap reduction engine of `Citlali`. As such it is not a required input when reducing in the beammap mode but is needed for the other modes. Important columns for other reduction modes include the parameters and errors from fits of the beammap source in maps made for each detector, bad detector flags derived from those fits, sensitivities, and the flux calibration factors. A few metadata items extracted from the APT are shown in Figure 3.2.

For its iterative mapmaker, `Citlali` reads in the maps it wrote in either the previous iteration or from a user-specified path. The pipeline can also optionally accept an image in the FITS format to generate synthetic timestreams for exploring the `Citlali`'s transfer function properties.

```yaml
meta:
  obsnum: 102518
  source: 1159+292
  project_id: 2022_TolTEC_Commissioning
  a1100_flux:
    - 2000
    - "units: mJy/beam"
    - a1100 flux density
  a1400_flux:
    - 2000
    - "units: mJy/beam"
    - a1400 flux density
  a2000_flux:
    - 2000
    - "units: mJy/beam"
    - a2000 flux density
  date: 2024-03-17.20:39:45
  mjd: 59932
  Radesys: altaz
  a_fwhm:
    - "units: arcsec"
    - fitted azimuthal FWHM
  a_fwhm_err:
    - "units: arcsec"
    - fitted azimuthal FWHM error
  amp:
    - "units: xs"
    - fitted amplitude
  amp_err:
    - "units: xs"
    - fitted amplitude error
  angle:
    - "units: rad"
    - fitted rotation angle
  angle_err:
    - "units: rad"
    - fitted rotation angle error
  array:
    - "units: N/A"
    - array index
  b_fwhm:
    - "units: arcsec"
    - fitted altitude FWHM
  b_fwhm_err:
    - "units: arcsec"
    - fitted altitude FWHM error
```

Figure 3.2. Subset of the meta information contained in the `Citlali` APT.

Other inputs that are neither files or reduction algorithm configuration parameters include pointing offset measurements and the flux density for the beammap source in each of the TolTEC bands. The pointing offsets are derived from pointing observations and can be provided as either a single value to be used as a constant offset or as a pair which are interpolated between. The Modified Julian Date (MJD) of the pointing observations can be provided to improve the accuracy of the correction when interpolation is performed. The beammap source fluxes are used to determine the flux calibration factors.

3.7 Pipeline Outputs

`Citlali` generates a new directory for every reduction which is populated with the reduction output files and a copy of the configuration file provided to it at runtime. Subdirectories for each individual observation from input file list and for the coadded data products are created, which in turn contain sub-directories for those products associated with raw or filtered maps. The structure for the `Citlali` output directory is illustrated in Figure 3.3.

The output files consist of:

- Raw and filtered maps
- Raw and filtered noise maps
- Map PSD files and histogram files
- Statistics file
- Array Property Table
- Pointing Property Table
- Timestream files

```
.
└── redu00
    ├── 102518
    │   ├── index.yaml
    │   └── raw
    │       ├── index.yaml
    │       ├── toltec_commissioning_a1100_science_102518_citlali.fits
    │       ├── toltec_commissioning_a1400_science_102518_citlali.fits
    │       ├── toltec_commissioning_a2000_science_102518_citlali.fits
    │       ├── toltec_commissioning_science_102518_hist.nc
    │       ├── toltec_commissioning_science_102518_psd.nc
    │       └── toltec_commissioning_science_102518_stats.nc
    ├── citlali_o102518_0_2_c1.yaml
    ├── coadded
    │   ├── filtered
    │   │   ├── index.yaml
    │   │   ├── toltec_commissioning_a1100_filtered_citlali.fits
    │   │   ├── toltec_commissioning_a1100_noise_filtered_citlali.fits
    │   │   ├── toltec_commissioning_a1400_filtered_citlali.fits
    │   │   ├── toltec_commissioning_a1400_noise_filtered_citlali.fits
    │   │   ├── toltec_commissioning_a2000_filtered_citlali.fits
    │   │   ├── toltec_commissioning_a2000_noise_filtered_citlali.fits
    │   │   ├── toltec_commissioning_hist_filtered.nc
    │   │   └── toltec_commissioning_psd_filtered.nc
    │   ├── index.yaml
    │   └── raw
    │       ├── index.yaml
    │       ├── toltec_commissioning_a1100_citlali.fits
    │       ├── toltec_commissioning_a1100_noise_citlali.fits
    │       ├── toltec_commissioning_a1400_citlali.fits
    │       ├── toltec_commissioning_a1400_noise_citlali.fits
    │       ├── toltec_commissioning_a2000_citlali.fits
    │       ├── toltec_commissioning_a2000_noise_citlali.fits
    │       ├── toltec_commissioning_hist.nc
    │       └── toltec_commissioning_psd.nc
    └── index.yaml
```

Figure 3.3. Structure of the standard `Citlali` output directory, with individual observation and coadded raw and filtered files divided into various subdirectories. `Citlali` will create a new directory of this type for every reduction.

`Citlali`'s maps are written in the FITS file format, which was chosen due to its widespread use across astronomy and the availability of existing software packages, particularly the Python library `astropy` (Collaboration et al. 2013; Collaboration et al. 2018; Collaboration et al. 2022) and the SAOImageDS9[14] application, to read in and manipulate FITS image data directly. Each detector array receives its own file, with the maps being stored as separate FITS extension Header Data Unit (HDU) layers. Raw and filtered maps are also maintained as separate files. The full list of

[14]https://sites.google.com/cfa.harvard.edu/saoimageds9

HDUs for single file is given in Figure 3.4. The maps output include the (1) signal, (2) weight, (3) synthetic source, (4) coverage, (5) coverage Boolean, and (6) signal-to-noise (S/N) maps. The coverage Boolean map identifies all pixels above a pixel weight threshold given by the user at runtime and is intended mostly for visualization purposes. `Citlali` can make maps for each array, network, or detector, so the total number of extension layers within a file varies between files. There will be 3 copies of each map type if a polarization reduction is enabled, corresponding to the 3 Stokes parameters. In beammap mode, owing to the fact that a map is made for each detector, only the signal and weight maps are included to reduce the file sizes and increase the file write speed.

```
No.    Name      Ver    Type      Cards   Dimensions   Format
  0  PRIMARY       1 PrimaryHDU      89   ()
  1  signal_I      1 ImageHDU        33   (1199, 1217, 1, 1)   float64
  2  weight_I      1 ImageHDU        34   (1199, 1217, 1, 1)   float64
  3  kernel_I      1 ImageHDU        35   (1199, 1217, 1, 1)   float64
  4  coverage_I    1 ImageHDU        33   (1199, 1217, 1, 1)   float64
  5  coverage_bool_I    1 ImageHDU        34   (1199, 1217, 1, 1)   float64
  6  sig2noise_I    1 ImageHDU        33   (1199, 1217, 1, 1)   float64
  7  signal_Q      1 ImageHDU        33   (1199, 1217, 1, 1)   float64
  8  weight_Q      1 ImageHDU        34   (1199, 1217, 1, 1)   float64
  9  kernel_Q      1 ImageHDU        35   (1199, 1217, 1, 1)   float64
 10  coverage_Q    1 ImageHDU        33   (1199, 1217, 1, 1)   float64
 11  coverage_bool_Q    1 ImageHDU        34   (1199, 1217, 1, 1)   float64
 12  sig2noise_Q    1 ImageHDU        33   (1199, 1217, 1, 1)   float64
 13  signal_U      1 ImageHDU        33   (1199, 1217, 1, 1)   float64
 14  weight_U      1 ImageHDU        34   (1199, 1217, 1, 1)   float64
 15  kernel_U      1 ImageHDU        35   (1199, 1217, 1, 1)   float64
 16  coverage_U    1 ImageHDU        33   (1199, 1217, 1, 1)   float64
 17  coverage_bool_U    1 ImageHDU        34   (1199, 1217, 1, 1)   float64
 18  sig2noise_U    1 ImageHDU        33   (1199, 1217, 1, 1)   float64
```

Figure 3.4. A list of the Header Data Units included in each FITS extension layer in a single `Citlali` output FITS file. With the exception of the `PRIMARY` HDU, which only contains a header, each includes both a header and a data array.

The FITS World Coordinate System (WCS) standard conventions are adopted, such that the files can be used with the `astropy.wcs` module. All maps for an individual observation are made on the same pixel grid and thus share the same WCS information which is stored in every map header. An example header with its WCS information from a signal map is shown in Figure 3.5. The coordinate reference frame will be either the Right Ascension and Declination celestial tangent plane projection for science map reductions or an offset coordinate frame in Azimuth and Elevation with the maps centered at (0,0) for calibration maps. The WCS units will be in units of degrees for the former and arcseconds for the latter. Metadata related to reduction parameters and derived quantities, observation identification information, and telescope configuration is stored in the primary header.

```
XTENSION= 'IMAGE   '           / IMAGE extension
BITPIX  =                  -64 / number of bits per data pixel
NAXIS   =                    4 / number of data axes
NAXIS1  =                  631 / length of data axis 1
NAXIS2  =                  595 / length of data axis 2
NAXIS3  =                    1 / length of data axis 3
NAXIS4  =                    1 / length of data axis 4
PCOUNT  =                    0 / required keyword; must = 0
GCOUNT  =                    1 / required keyword; must = 1
EXTNAME = 'signal_I'
HDUVERS =                    1
EQUINOX =                2000. / WCS: Equinox
CTYPE1  = 'RA---TAN'           / WCS: Projection Type 1
CUNIT1  = 'deg     '           / WCS: Axis Unit 1
CRVAL1  =             91.94292 / WCS: Ref Pixel Value 1
CDELT1  =        -0.0002777778 / WCS: Pixel Scale 1
CRPIX1  =                 316. / WCS: Ref Pixel 1
CTYPE2  = 'DEC--TAN'           / WCS: Projection Type 2
CUNIT2  = 'deg     '           / WCS: Axis Unit 2
CRVAL2  =            -6.385833 / WCS: Ref Pixel Value 2
CDELT2  =         0.0002777778 / WCS: Pixel Scale 2
CRPIX2  =                 298. / WCS: Ref Pixel 2
CTYPE3  = 'FREQ    '           / WCS: Projection Type 3
CUNIT3  = 'Hz      '           / WCS: Axis Unit 3
CRVAL3  =         2.725386E+11 / WCS: Ref Pixel Value 3
CDELT3  =                   1. / WCS: Pixel Scale 3
CRPIX3  =                   1. / WCS: Ref Pixel 3
CTYPE4  = 'STOKES  '           / WCS: Projection Type 4
CUNIT4  = '        '           / WCS: Axis Unit 4
CRVAL4  =                   0. / WCS: Ref Pixel Value 4
CDELT4  =                   1. / WCS: Pixel Scale 4
CRPIX4  =                   1. / WCS: Ref Pixel 4
UNIT    = 'mJy/beam'           / Unit of map
END
```

Figure 3.5. Example of the header for a signal map from a science observation showing the WCS setup.

The FITS files for the jackknifed noise maps are very similar to those of the data maps using a separate extension layer for each noise realization. Figure 3.6 shows the different extension layers of a noise FITS map. They also make use of the same WCS as the observation files.

No.	Name	Ver	Type	Cards	Dimensions	Format
0	PRIMARY	1	PrimaryHDU	89	()	
1	signal_0_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
2	signal_1_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
3	signal_2_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
4	signal_3_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
5	signal_4_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
6	signal_5_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
7	signal_6_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
8	signal_7_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
9	signal_8_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
10	signal_9_I	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
11	signal_0_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
12	signal_1_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
13	signal_2_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
14	signal_3_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
15	signal_4_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
16	signal_5_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
17	signal_6_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
18	signal_7_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
19	signal_8_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
20	signal_9_Q	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
21	signal_0_U	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
22	signal_1_U	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
23	signal_2_U	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
24	signal_3_U	1	ImageHDU	34	(1199, 1217, 1, 1)	float64
25	signal_4_U	1	ImageHDU	34	(1199, 1217, 1, 1)	float64

Figure 3.6. The same as Figure 3.4 except for the jackknifed noise maps FITS file. Each layer is an individual noise realization.

The map units will be determined based on the configuration file parameters `mapmaking->cunit` and `timestream->type`. The former accepts values of mJy/beam, MJy/Sr, Jy/pixel, and μK_{CMB}. Internally, mJy/beam is used as the default unit and the flux calibration factors within the APT are always in this unit. The beam size refers to the average beam area for each wavelength within the beammap observation that made the APT and can vary between APTs. As a result, all detectors within an array receive the same conversion factor. For this reason, reducing science and non-beammap calibration observations in mJy/beam requires the fewest number of unit conversions. The standard FITS beam shape parameters keys `BMAJ`, `BMIN`, and `BPA` are included in the primary HDUs. The conversion to uK is given by

$$\frac{dS_\nu}{dT_0} = \frac{2\pi\theta^2}{2\sqrt{2\ln(2)}} \times \frac{dB_\nu(T_0)}{dT_0}, \tag{3.1}$$

where T_0 is the CMB temperature, θ is the FWHM of the beam for frequency ν, B_ν is the Planck function at T_0.

The `timestream->type` parameter controls what type of data is used from the output of the KIDs model solving to create the time chunk data, with allowed values being the detuning parameter, x, its quadrature channel, r, and the corresponding S_{21} in-phase and quadrature channels I and Q. The flux calibrations are only physically meaningful for x and r and only make sense if the beammap was reduced with the same unit.

Two-dimensional PSDs and histograms for each signal map and the average of the noise maps are output as netCDF4 files. Pixels with weights below the weight threshold that controls the coverage Boolean map are not included in these calculations. The statistics or "stats" file is similarly written in the netCDF4 format and includes summary information of the timestreams on a per-time chunk and per-detector level. Statistics like the detector means, RMS values, and the fraction that were flagged by weight cuts are helpful in validating and characterizing the behavior of the individual detector networks. The eigenvalues and eigenvectors from the atmospheric cleaning can be optionally included.

As detailed in Section 3.6, the APT is an output for the beammaping engine only and incorporates per detector information determined from that mode. The Pointing Property Table (PPT) is an analogous ECSV table output by the `citlali::pointing` reduction engine and is a key output for focus, astigmatism, and pointing reductions. It includes fitted parameters and errors for the source in each array's signal map instead of for each detector. Unlike the APT, if map filtering or coaddition is enabled, a PPT is generated for those as well. The source coordinates from the PPT in the azimuth and elevation coordinate frame can be used as the inputs for the pointing offsets in the configuration file.

`Citlali` can output entire partially reduced timestreams following flux calibration or after atmospheric cleaning as netCDF4 files. As all timestream reduction algorithms can be disabled, any combination of reduction steps can be applied prior to their output. These files also include the flagging and synthetic source timestreams, as well as processed telescope and per-detector pointing arrays. The primary purpose of these files is to enable the reduction of TolTEC data with external data reduction pipelines and mapmakers without needing to develop additional software tools for translation. The flux-calibrated timestreams do not require information related to the KIDs models, TolTEC, or the LMT to be processed making them easier to incorporate into standalone software pipelines. These are also convenient for testing and verifying reduction algorithms and for data characterization and verification during commissioning.

3.8 Verification of Algorithms

Many of `Citlali`'s fundamental algorithms, particularly those related to timestream reduction, mapmaking, and map filtering are derived from the AzTEC data reduction pipeline (Scott et al. 2008). Early in the `Citlali` development cycle, these algorithms were translated to the new codebase and rewritten to add optimizations, improve modularity, and deal with TolTEC-specific data features. During this period, the newly implemented algorithms underwent a testing phase where they were directly applied to AzTEC raw timestreams and compared to results produced from the AzTEC pipeline, which has undergone extensive verification through published scientific results. Later, the TolTEC observation simulator was used to create simulated datasets with realistic detector and atmospheric noise which allowed for detailed characterization and benchmarking of each phase of the pipeline.

Commissioning observations from 2022 allowed for the first full integration of `Citlali` with the rest of the TolTEC software stack as well as providing critical tests of every pipeline stage through analyses with real data. In particular, the observations informed developments of better detector flagging routines, the iterative mapmaking algorithm, and improvements to the polarization reduction algorithm.

3.9 Pipeline Structure

A high-level view of the structure of `Citlali` is illustrated in Figure 3.7. The pipeline can be logically divided into four components which are executed in order, though some may be repeated when reducing multiple observations or when the iterative mapmaker is enabled. These phases are the (1) initial configuration, (2) observation setup, (3) timestream reduction, and (4) mapmaking, and post-mapmaking stages. The following sections will describe each phase.

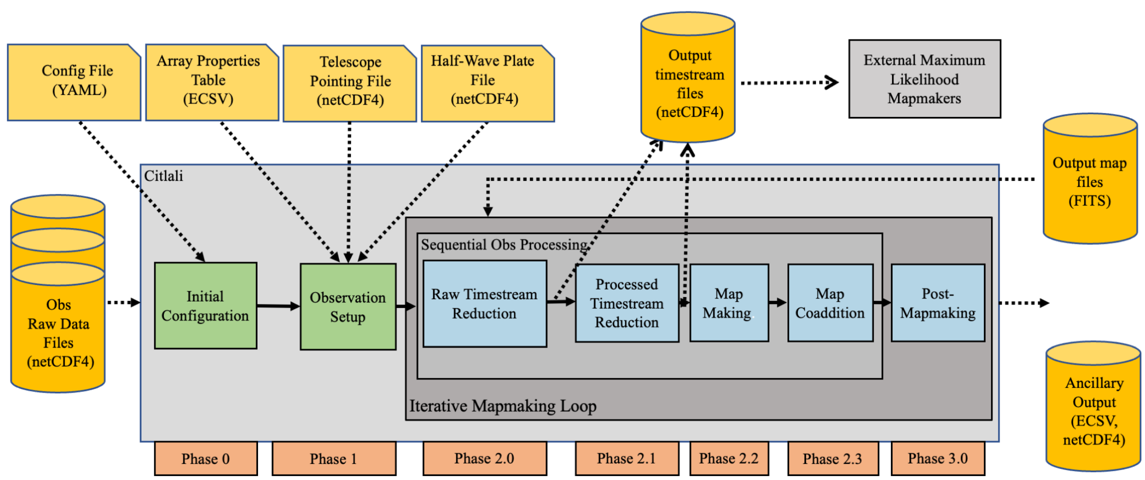

Figure 3.7. High-level view of the `Citlali` data reduction pipeline.

3.10 Initial Configuration

The first phase is the simplest, where the `Citlali` YAML configuration file is read in, and its values are validated against allowed values and ranges to ensure all required keys are present, that the values are reasonable, and that the datatypes are correct. If erroneous values are found, `Citlali` will cleanly exit and print a list of the configuration keys that require adjustment. Some preliminary setup is performed here, which mainly involves populating the timestream processing and mapmaking classes with values derived from configuration inputs. The pipeline then selects the corresponding reduction engine class for use depending on the reduction type from the config file.

The file list is also parsed and all files and observation-specific parameters like pointing offsets and source fluxes are assembled into a vector of data classes for each observation. No actual data other than file header information for input validation is read in at this stage.

3.11 Pipeline Setup

After parsing the configuration file, the number and dimensions of the maps are calculated in a loop across all included observations. The number of maps simply depends on the map grouping provided by the user (i.e. array, network, or detector), which files are included in the list of raw data inputs, and if polarized maps are requested. This is allowed to vary between observations if necessary. Before the calculation of the map dimensions can occur, three other quantities must be determined. First, the telescope pointing vectors must be interpolated from their native sampling frequency onto a time grid common to all the TolTEC networks. The boresight tangent plane projection coordinates and the start and end vector indices for the data time chunk are also required. The current observation's telescope file and APT are therefore read in and parsed at this point.

3.11.1 Timing Interpolation

The absolute Coordinated Universal Time (UTC) of each sample in the telescope vectors and raw timestream data can be calculated from the GPS PPS synchronous signal with

$$t = t_0 + \mathrm{PPS} + \frac{\mathrm{Clock} - \mathrm{PPS}}{f_{\mathrm{fpga}}}, \tag{3.2}$$

where T_0 is the UTC of the first sample, PPS is the PPS count, Clock is the clock count, and f_{fpga} is the frequency of the network FPGA readout. The value for T_0 is calculated by

$$t_0 = \mathrm{int}(\mathrm{sec}_0 + \mathrm{nanosec}_0 \times 10^{-9} - 0.5), \tag{3.3}$$

with sec_0 and $\mathrm{nanosec}_0$ being the second and nanosecond components of the PPS signal respectively. Each of the TolTEC networks may begin recording data at slightly different times, so the common time grid used consists of all samples after the time that the last network begins recording and ends when the first network stops recording. A start and end index are stored for each network and used to offset the raw data when it is read in during later phases.

Once the grid is calculated, all telescope vectors are interpolated over. The pointing offsets from the configuration file are also interpolated across to determine the necessary correction for each sample.

3.11.2 Coordinate Frame Calculations

The tangent plane or local geodetic reference frame is used for all `Citlali` maps. In the Azimuth and Elevation coordinate system, the tangent plane coordinates are calculated with

$$Az_{tan} = cos(El - El_{cor}) \times (Az - Az_{source}) - Az_{cor} \tag{3.4}$$

and

$$El_{tan} = El - El_{source} - El_{cor}. \tag{3.5}$$

In these equations, $(Az,\ El)$ are the telescope boresight coordinates in absolute coordinates, $(Az_{cor},\ El_{cor})$ are corrections from the LMT pointing model and observer offsets, and $(Az_{source},\ El_{source})$ are the source coordinates.

In the Ra and Dec frame, the tangent plane calculations are given by

$$\alpha_{tan} = \cos(\delta) \times \frac{\sin(\alpha - \alpha_{center})}{\cos(c)} \tag{3.6}$$

and

$$\delta_{tan} = \frac{\cos(\delta_{center})\sin(\delta) - \sin(\delta_{center})\cos(\delta)\cos(\alpha - \alpha_{center})}{\cos(c)}. \tag{3.7}$$

Here, (δ,α) are the boresight coordinates and $(\delta_{center},\ \alpha_{center})$ is the map tangent point which is typically taken to be the center of the map. The denominator term is

$$\cos(c) = \sin(\delta_{center})\sin(\delta) + \cos(\delta_{center})\cos(\delta)\cos(\alpha - \alpha_{center}) \tag{3.8}$$

.

3.11.3 Time Chunking

For Raster maps, the default time chunk size is equivalent to the length of an individual scan. The LMT telescope file includes a `HOLD` signal that records when the telescope is on and off trajectory, which is used during reduction to identify and exclude the samples corresponding to the turnarounds. Figure 3.8 plots the telescope trajectory for a beammap observation, highlighting the identification of the samples where the telescope is turning around.

Observations using the Lissajous, double Lissajous, and the Rastajous pattern are instead divided into a series of equal-length time chunks whose duration in seconds can be adjusted in the configuration file. Raster map chunking can also optionally be overwritten to use this mode.

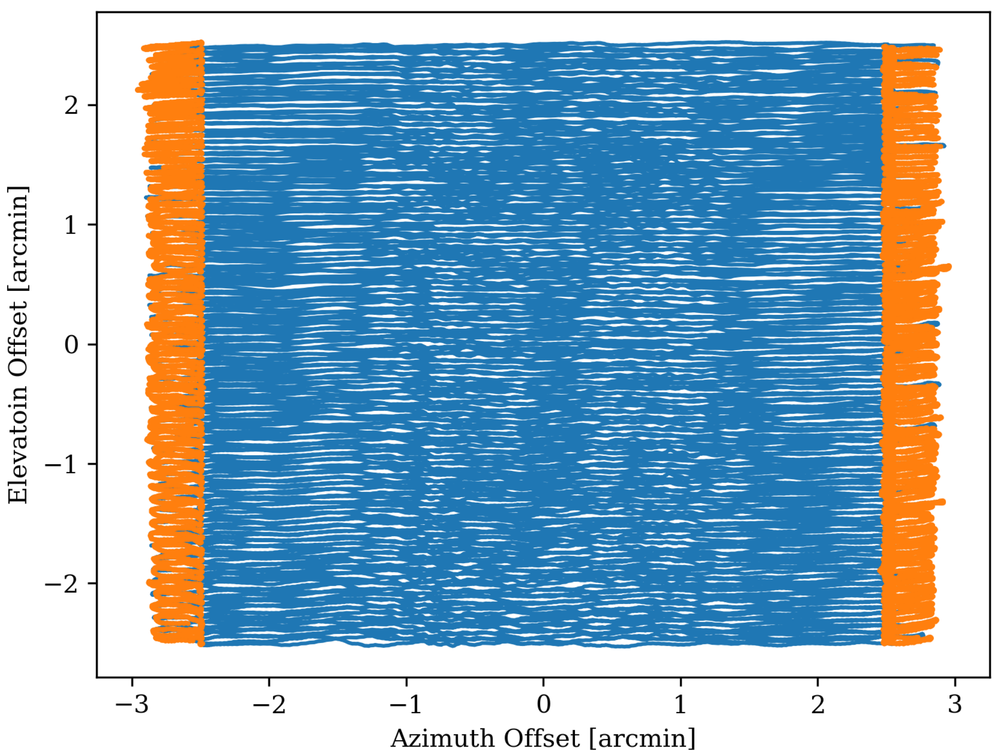

Figure 3.8. Raster map pattern from a beammapping observation with TolTEC. The orange curves represent the turnaround points as measured by the LMT's `HOLD` signal and will not be included in the final maps.

A bandpass filter may be applied during the later timestream reduction stages, which has the effect of reducing the size of the time chunk owing to the filter convolution. In order to prevent the loss of a large fraction of data, each time chunk, except for the first and last, are extended by half of the filter size on either end, such that they include overlapping data. The setup is illustrated in Figure 3.9. After the filter has been applied, this extra data is discarded.

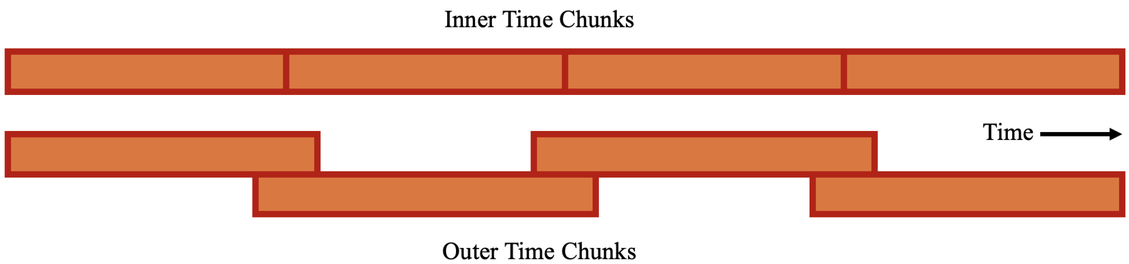

Figure 3.9. Diagram of the chunking strategy used in `Citlali`. Each orange block represents an individual time chunk in the timestream. The outer chunks include overlapping data and are read from disk. After timestream filtering, only the inner chunks are used.

3.11.4 On-sky Detector Pointing

With the interpolated telescope vectors, boresight tangent plane coordinates, and time-chunk indices, the global maximum and minimum coordinates of all detectors in the tangent plane coordinates are then calculated. The APT provides the measured offset coordinates ($Az_{off,0}$, $El_{off,0}$) for each detector at an elevation of 0 degrees. These values will be 0 when beammapping, such that the maps will be centered on the boresight as opposed to each detector's offset position for that mode. For each sample, these are rotated to the current elevation (El):

$$\begin{pmatrix} Az_{\text{off}} \\ El_{\text{off}} \end{pmatrix} = \begin{pmatrix} \cos(El) & -\sin(El) \\ \sin(El) & \cos(El) \end{pmatrix} \begin{pmatrix} Az_{\text{off},0} \\ El_{\text{off},0} \end{pmatrix}. \tag{3.9}$$

For maps made in the Az/El frame, a detector's position is just the sum of its coordinates (Az_{off},El_{off}) and the boresight coordinates or

$$\begin{aligned} Az_{det} &= Az_{off} + Az_{tan} \\ El_{det} &= El_{off} + El_{tan}. \end{aligned} \tag{3.10}$$

To determine the detector pointing in the Ra/Dec frame an additional rotation of (-Az_{off}, El_{off}) by the parallactic angle PA is required in addition to adding the boresight offset (α_{tan}, δ_{tan}):

$$\begin{pmatrix} \alpha_{\mathrm{det}} \\ \delta_{\mathrm{det}} \end{pmatrix} = \begin{pmatrix} \cos(PA) & -\sin(PA) \\ \cos(PA) & \sin(PA) \end{pmatrix} \begin{pmatrix} -Az_{\mathrm{off}} \\ El_{\mathrm{off}} \end{pmatrix} + \begin{pmatrix} \alpha_{\mathrm{tan}} \\ \delta_{\mathrm{tan}} \end{pmatrix}. \tag{3.11}$$

A detector's pointing is not permanently stored in `Citlali` but is instead calculated on-the-fly whenever it is needed to save memory.

3.11.5 Map Setup

The required number of map rows and columns is simply defined by the smallest region that includes the maximum and minimum tangent plane coordinates from all detectors that were not flagged as bad in the APT. The pixel size is provided as a configuration file parameter, with the standard value used during commissioning being $1''$. All three detector arrays are co-aligned and have the same sized footprint on the sky so the same map and pixel size can be used for all of them without unnecessarily adding extra pixels.

If map coaddition is requested, its dimensions are calculated from the union of the individual map sizes. `Citlali`'s coaddition routine uses the same pixel size for the coadded maps as the individual observation maps and assumes all observations to be co-centered at the same on-sky coordinates. Memory for the coadded maps is allocated immediately after this calculation.

3.12 Timestream Reduction and Mapmaking

Phase 2 timestream and mapmaking reduction begin after the map dimensions are calculated and stored and the memory for the coadded maps is allocated. Here, `Citlali` loops through each observation, recomputes timing grids for the raw data, pointing information, and HWP angle, and then begins the reduction. The reduction loop through observations may be performed iteratively using information from the previous iteration.

`Citlali` includes many timestream processing algorithms and mapmaking strategies which are selectable and customizable in the runtime configuration file. The timestream reduction algorithms serve to remove instrumental effects or badly behaving detectors and to filter the atmospheric and other noise contributions. The mapmaking routines specifically refer to the transformations carried out to translate the data from the time domain to the map-space. Each algorithm consists of a self-contained class that accepts the minimal amount of contextual information required to carry out its operations, thereby making the development and addition of new components more manageable.

Within `Citlali` a time chunk represents the minimum-sized block of timestream data that the pipeline will process and can be thought of as `Citlali`'s fundamental data object. Two separate time chunk classes are defined to hold all data that is unique and relevant to a single chunk. This includes the data, synthetic source, and flag timestream matrices, in addition to detector weights and the sub-set of the telescope vectors, pointing corrections, and HWP orientation angles corresponding to that chunk. The APT is not included as it is common to all time chunks and does not incorporate per-sample information. Also, as previously discussed, no per-detector pointing information is stored as these are calculated when needed. The two classes are the raw time chunk data (`citlali::RTCData`) and processed time chunk data (`citlali::PTCData`) classes. Both inherit from an underlying time chunk class (`citlali::TCData`) that incorporates variables common to both chunk variants. Each retains a history of which reduction stages have been applied to it as well as some basic statistics.

In the pipeline, a `citlali::RTCData` instance is declared and populated from the raw data, converted into a `citlali::PTCData` through successive time chunk processing, further processed, and then passed to the mapmaking classes after which it is deleted. The transformation of the raw chunk into a processed chunk occurs immediately following timestream decimation and prior to atmospheric cleaning. The division is motivated by the fact that decimation incurs a reduction in the size of the timestream arrays along the time axis and that the `citlali::beammap` engine requires the storage and re-use of every chunk before atmospheric cleaning. The incorporation of external mapmaking pipelines also prompts the timestreams to be optionally written to disk before atmospheric cleaning is carried out thus making it a logical dividing point.

The time chunk classes are paired with two corresponding processing classes for raw (`citlali::RTCProc`) and processed (`citlali::PTCProc`) time chunk processing that perform the actual timestream reduction operations. Each includes a `run()` function that accepts either a `citlali::RTCData` or `citlali::PTCData` as input and applies all enabled timestream reduction steps sequentially. The various reduction algorithm classes like `timestream::Despiker` and `timestream::Cleaner` are declared in their respective processing class. The member classes of the `citlali::RTCProc` class are illustrated in Figure 3.10. They also inherit from a base processing class, `citlali::TCProc`, and share functions to carry out bad detector flagging, time-domain source addition and subtraction, region masking, inverse mapmaking routines for iterative mapmaking, and the ability to append the data they hold to `Citlali`'s outputs timestream netCDF files.

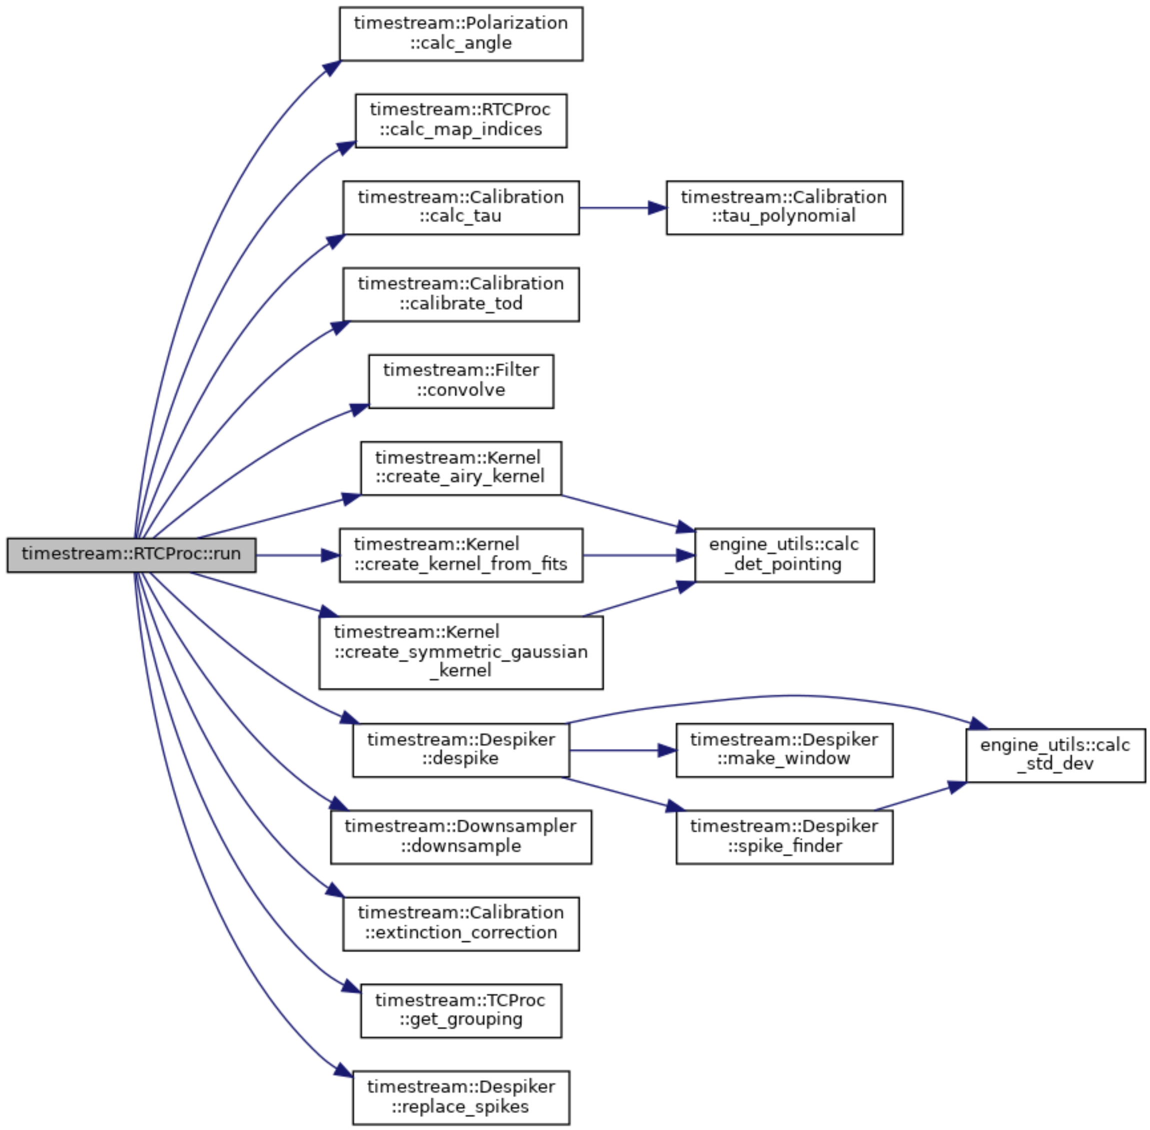

Figure 3.10. Timestream reduction steps called by the `citlali::RTCProc` class.

The objects related to the storage and manipulation of data in the map domain consist of a class for holding all data and noise maps and related information for either a single observation or coaddition (`mapmaking::MapBuffer`) and individual classes for the two different mapmaking routines currently implemented (`mapmaking::naive_mm` and `mapmaking::jinc_mm`). The mapmaker classes accept the `citlali::PTCProc` time chunk classes as inputs and add their samples into the various map matrices of the `mapmaking::MapBuffer` class.

3.12.1 Reduction Engines

The timestream processing and map classes thus far discussed are collected within the 3 high-level reduction engine classes, `citlali::lali`, `citlali::pointing`, and `citlali::beammap`, whose operations are the primary drivers of the stage 2 pipeline phase. These engines are written to provide the unique assembly of lower-level reduction steps required by science, pointing, focus, astigmatism, and beammap observations. `Citlali` uses the C++ `std::variant` class template functionality to instantiate these classes and enable easy switching between them in a single `main.cpp` file with the single configuration file option `runtime->reduction_type`. These inherit from a `citlali::Engine` class that manages per-observation routines like output directory setup, file writing, and map filtering. All timestream processing, map and mapmaking, and post-mapmaking filtering classes are declared in `citlali::Engine`.

All of the reduction engines declare a `setup()`, `pipeline()`, and `output()` function that are called in that order in `main.cpp`. The `setup()` function re-initializes all variables and containers necessary for the reduction of each observation while also creating per-observation output directories. The actual timestream reduction and mapmaking is performed within `pipeline()` which accepts a KIDs processing class, `citlali::KidsProc`, as its only input. This class acts as the interface between `kidscpp` library that is also compiled with `Citlali` and the rest of the pipeline. It is used to extract data and header information from the raw TolTEC data files and solve the KIDs model with the tune ECSV file. As its name suggests, `output()` writes the map FITS, APT, PPT, PSD, histogram, and stats files for each observation. It is re-used for coadded FITS files as well. The netCDF4 timestream files are not written here as they are appended while the time-chunks are reduced so that the entire reduced timestream does not need to be stored in memory. Some final header information that requires the final maps to be generated is added here.

3.12.1.1 Science Observation Reductions

Science observation reductions are reduced with `citlali::lali` which is the simplest of the reduction engines. A schematic is illustrated in Figure 3.11. In the `pipeline()` function a `grppi::pipeline` object is used to produce a single-threaded data stream of populated `citlali::RTCData` classes from the raw TolTEC data files in conjunction with the pre-computed time chunk indices. The `citlali::RTCData` classes are then fed directly into a `grppi::farm` parallel pattern which will create a new thread to carry out all timestream and mapmaking reduction stages on that chunk. Once the contributions from all time chunks have been added to into the observation map, the maps are normalized and the PSDs and histograms are calculated, which occurs outside of the farm. These are however still parallelized over the number of maps but at a much lower level. File output for the current observation occurs directly after this.

3.12.1.2 Focus, Astigmatism, and Pointing Observation Reductions

The primary unique data product for all the non-beammap calibration observations are the fitted parameters for the observed point source. Consequently, their reductions are identical and can use the exact same code regardless of whether the fits are used in different ways post-reduction. The `citlali::pointing` reduction engine is very similar to the `citlali::lali` engine as shown in Figure 3.11, with the

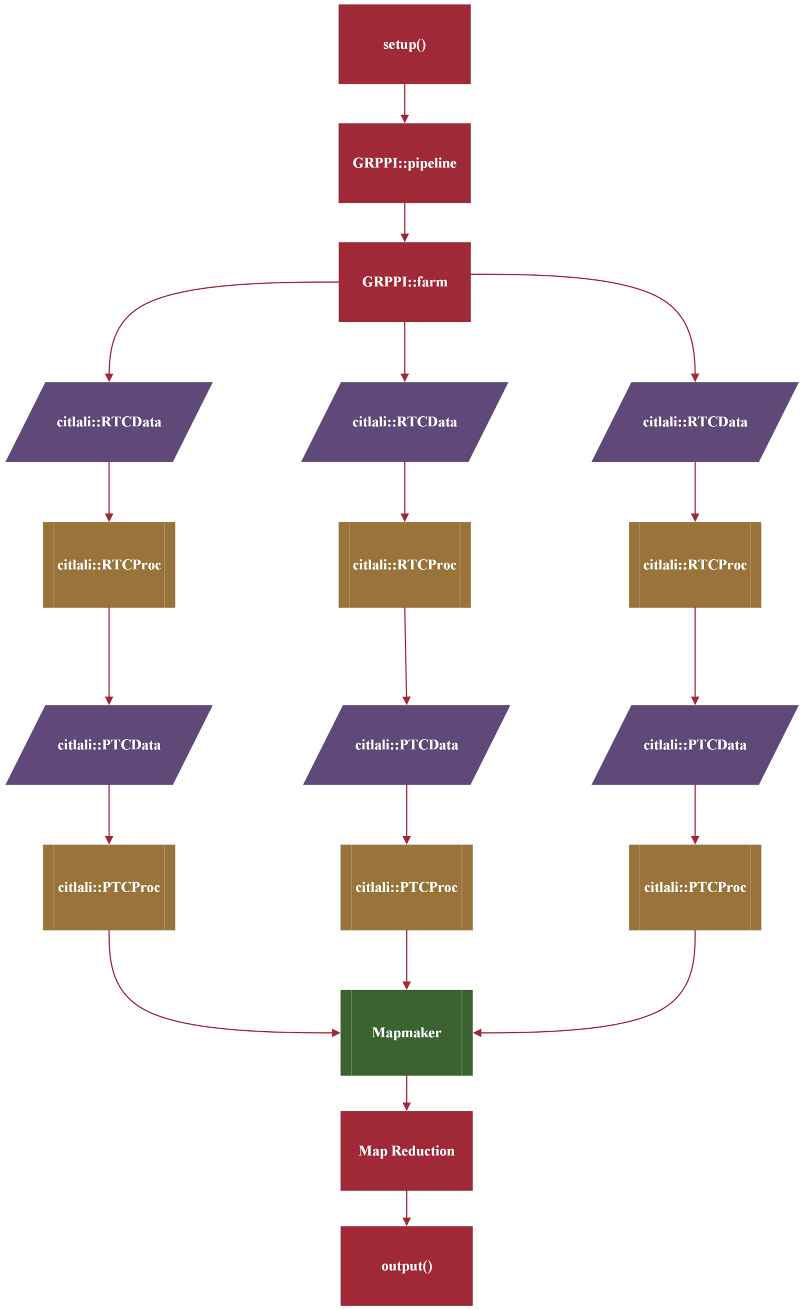

Figure 3.11. Diagram of the `citlali::lali` and `citlali::pointing` reduction engines. The two engines primarily differ in the "Map Reduction" step. Square blocks represent reduction functions or classes and the parallelogram are data classes. Diverging arrows represent processes being spawned to individual threads by the `GRPPI::farm` pattern, which then are collected during the mapmaking step.

primary difference being the addition of the map fitting step that fits point sources in the maps to a two-dimensional elliptical Gaussian with the Ceres-Solver library. Despite the similarities, it is kept as a different reduction engine from `citlali::lali` to maintain a distinction between the specific needs of each of the observation types and to allow for future expansion that may cause the reduction approaches to differ more substantially. The routine uses the S/N map to find an initial guess position within a sub-region of the map selected by the user. The `setup()` function initializes the PPT which is populated after the fits are performed following map normalization in `run()`, and written to the ECSV table in `output()`. While non-beammap calibration observations are not expected to be coadded, a fit is also performed for these maps in addition to raw and filtered maps as well. When iterative mapmaking is enabled, fits will also be performed for the maps of each iteration as the iteration loop wraps around the `setup()`, `run()`, and `output()` functions in `main.cpp`.

3.12.1.3 Beammap Observation Reductions

The `citlali::beammap` engine differs substantially from the other two classes. With the goal of beammaps being to determine individual detector offset positions, beam shapes, and flux calibration factors, individual detector maps must be produced and fit to a two-dimensional Gaussian model. The atmospheric removal step that is necessary before the maps can be produced and fit significantly impacts the estimate of the flux calibration factors in a negative way. The point source flux will be attenuated by the cleaning algorithm which will then bias the flux calibration factor to predict a higher value. This will likely be a wavelength-dependent effect as the same cleaning parameters when applied to the longer wavelength bands will produce a greater reduction in source flux on the source owing to lower opacity and larger beamsizes. `Citlali` does include an optimal point source filter that can use the synthetic point source maps to correct for flux lost during atmospheric removal, but this is too computationally expensive for beammaps when reducing on-site as each detector's map must be filtered and will require jackknifed noise realizations. The filter also has the effect of smoothing the maps and increasing the beam width.

To correct for this undesired effect, an iterative timestream reduction and mapmaking routine is used in `citlali::beammap` which is illustrated in Figure 3.12. In this approach, the timestream reduction is broken into two components that mirror the separation in the raw and processed time chunk classes. The `grppi::pipeline` and `grppi::farm` from the `citlali::lali` and `citlali::pointing` reduction engines is the first component, but instead of carrying out all timestream and mapmaking routines it now ends before atmospheric cleaning is applied to the time chunks. The `citlali::PTCData` classes output from `citlali::RTCProc` are collected into a `std::vector` to be used after the parallel farm has been completed.

Following the completion of the first pipeline, an iteration while loop is started. A copy of the vector of `citlali::PTCData` classes is made as the original timestreams before cleaning are required for each iteration. As a result, `citlali::beammap` must hold in memory two copies of the entire timestream data within all the time chunks, though this is usually much smaller than the raw data when timestream decimation is used. On the first iteration, the `citlali::PTCProc` reduction algorithms, namely cleaning, flagging, and detector weight calculation, are applied to this copy. Individual detector maps from the cleaned timestreams are made and fit to a 2D symmetric Gaussian model. Symmetry is enforced since the cleaning algorithms and Raster map pattern introduce a compression of the beam width in the scan direction. On subsequent iterations, the source as determined by the two-dimensional fit is subtracted from the uncleaned timestreams and re-added after it. This has the effect of removing a large fraction of the source contribution from the detector correlation matrix eigenvalues used in the principal component analysis cleaning algorithm. This contribution is therefore not attenuated and is carried through to the final maps. This

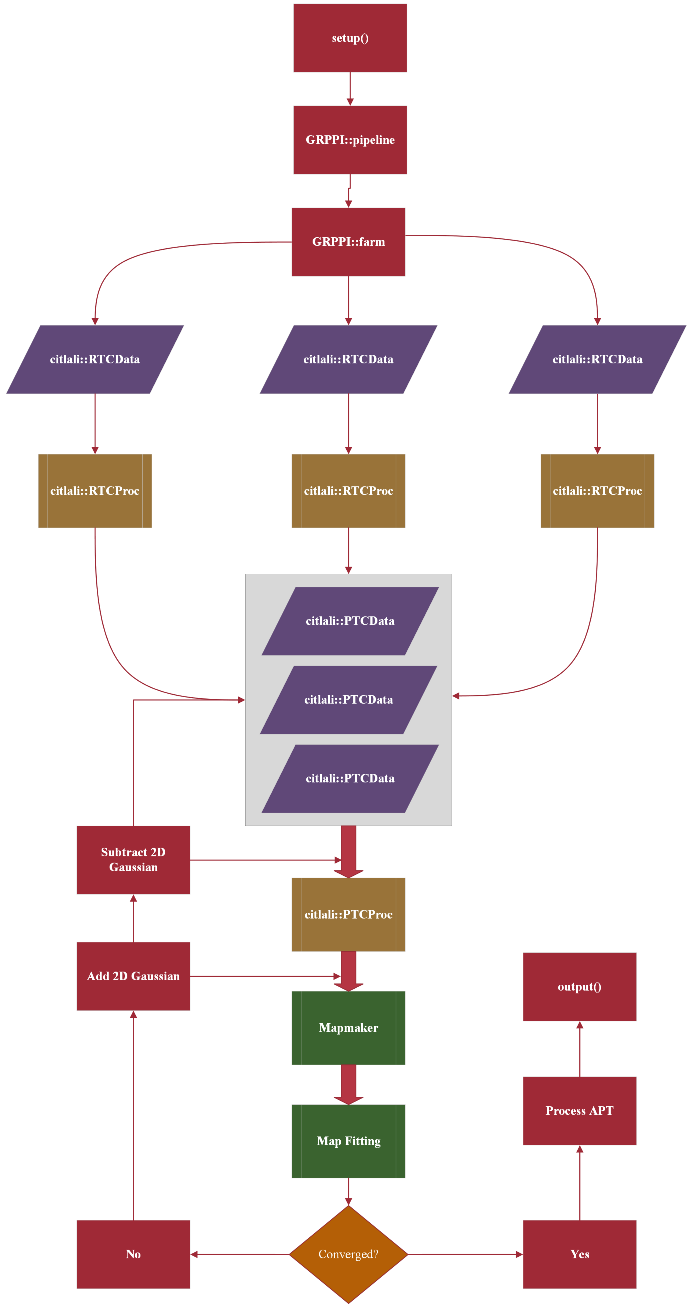

Figure 3.12. Similar to Figure 3.11 but for the `citlali::beammap` reduction engine. Thick lines represent parallelized for loops over time chunks or detectors. The first farm pattern ends with the `citlali::PTCData` classes being collected into a vector. A copy of these are processed in the later stages. If the map fits converge, the APT is generated and written to disk. Otherwise, the process is repeated using the `citlali::beammap` iterative mapmaking algorithm.

is repeated until the percent difference between fits of two iterations is below a user-specified tolerance indicating convergence or until the maximum number of iterations is reached. The convergence of each detector is checked at the end of each iteration and converged detectors are not re-fit on later iterations.

Parallelization of the iterative `citlali::beammap` pipeline is also slightly more complex than in the other reduction engines to maintain efficiency when timestream output is enabled. The `citlali::PTCProc` reductions of each `citlali::PTCData` copy and mapmaking are parallelized using two grppi::map patterns over time chunks, which are the exact equivalent of a parallel for loop. Timestream output occurs between these two loops as this must be performed sequentially. Map fitting is parallelized over the number of maps which in this case is the number of detectors.

After all maps have been fit, the APT is constructed in `citlali::beammap`. The fitted parameters are checked against limits provided in the `Citlali` configuration file and are flagged as bad if they fall outside of the allowed range. Spurious features in frequency sweeps that were erroneously identified as real KIDs will generally produce very noisy maps with no source present in them. The limits include upper and lower bounds on the beam FWHM, array footprint size, and S/N. The median sensitivity across all time chunks for each detector is also calculated and can be used for flagging. Optimal ranges for these parameters are informed based on theoretical expectations of the instrument and detector characteristics. Resonators colliding in frequency space can result in two sources appearing in a single map and are difficult to identify automatically at this stage. They are instead flagged during timestream reduction by placing a lower limit on the frequency separation on the tones.

After flagging, the fitted detector positions in the Azimuth and Elevation coordinate frame are offset with respect to an unflagged reference detector which is chosen to be the one closest to (0,0). The centroids are then de-rotated by the mean elevation of the beammap source to the horizon as they will be re-rotated to the elevation of each sample in the other reduction engines according to the calculations in Section 3.11.4.

Finally, the flux calibration factors for unflagged detectors are calculated by

$$FCF_\nu = \left(\frac{F_{calib}}{x_{fit}}\right) \times e^{-\tau_\nu}, \tag{3.12}$$

where F_{calib} is the calibration source flux in units of mJy/beam, x_{fit} is the amplitude from the beammap source fit in units of the detuning parameter, and τ_ν is the mean opacity for each of the TolTEC bands. The opacity calculation is described in Section 3.12.2.2.

3.12.2 Raw Timestream Reduction Algorithms

In this section, I describe the main timestream reduction algorithms that are handled by the `citlali::RTCProc` class and are applied to the timestream data before atmospheric filtering in the order that they are carried out.

3.12.2.1 KIDs Model Solving and Flux Calibration

The first timestream reduction step that occurs is the transformation of the timestream data from the signal in the raw (I,Q) plane to the (x,r) frame by using the fitted solution to 2.3. The FCF from the APT given in Equation 3.12 is then multiplied against all samples for each detector, as are any unit conversions from mJy/beam to the other supported units that were calculated in phase 1.

3.12.2.2 Extinction Correction

Even at high altitudes in dry locations, water and oxygen molecules in the troposphere introduce broadband absorption at millimeter wavelengths which must be corrected for to derive accurate flux density estimates of astrophysical sources (Ferrusca and Contreras R. 2014). The LMT radiometer records the zenith atmospheric opacity at 225 GHz in 15-minute intervals. `Citlali` uses the `am` atmosphere model

(Paine 2022) to derive an estimate of the opacity in each of its bands based on the measured zenith τ_{225GHz} value. The `am` model is a radiative transfer code that calculates the transmission as a function of frequency through simulated atmospheric layers whose properties are constrained from observations with the NASA Modern-Era Retrospective analysis for Research and Applications (MERRA) mission.

Simulated atmospheric transmission for the 25, 50, 75, and 95 percentile weather conditions at the LMT site have been computed for elevations between 10 degrees and 80 degrees in steps of 10 degrees. The 25 percentile transmission as a function of frequency and elevation is plotted in Figure 3.13. The transmission T_x at elevation angle θ_{El} is related to the opacity by

$$T_x(\theta_{El}) = \exp(-\tau_\nu(\theta_{El})). \tag{3.13}$$

The optical depth at elevation angle θ_{El} can be derived from the zenith opacity with

$$\tau_\nu(\theta_{El}) = \tau_\nu(90^\circ) \times A, \tag{3.14}$$

where A is the airmass and is estimated as

$$A = \sec(90^\circ - \theta_{El}) \left(1 - 0.0012 \left(\sec(90^\circ - \theta_{El})^2 - 1\right)\right), \tag{3.15}$$

as given in Young and Irvine 1967. The ratio of the transmission at 225 GHz to the transmission in each of the TolTEC bands is fitted to an order 6 polynomial as a function of elevation for each percentile. Figure 3.14 shows the fit results for the 25 percentile atmosphere case. During reduction, the transmission at τ_{225GHz} at the elevation of every sample is calculated using equations 3.13, 3.14, and 3.15 and the polynomial fit is then used to find the TolTEC band transmissions for each sample. The same equations are then inverted to give τ_{273GHz}, τ_{214GHz}, and τ_{150GHz}. `Citlali` compares the measured zenith τ_{225GHz} to each of the percentile model's τ_{225GHz} at 80 degrees and automatically selects the percentile model that has the nearest value.

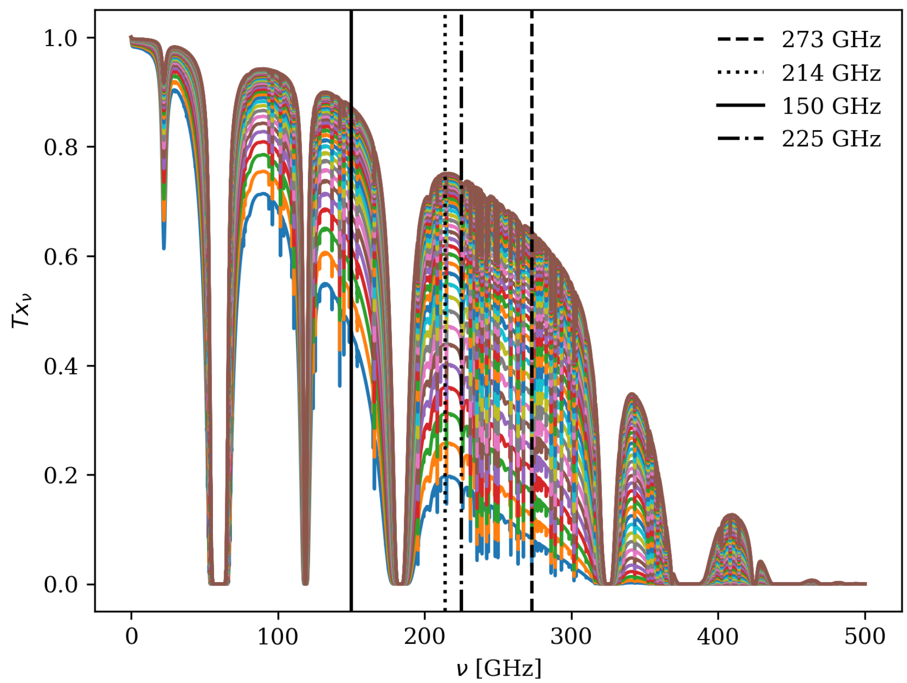

Figure 3.13. Simulated atmospheric transmission as a function of frequency as calculated by the `am*` atmosphere model. Each curve is the transmission for a different elevation with values 10 and 80 degrees. The vertical lines mark the frequencies of the radiometer and the TolTEC bands.

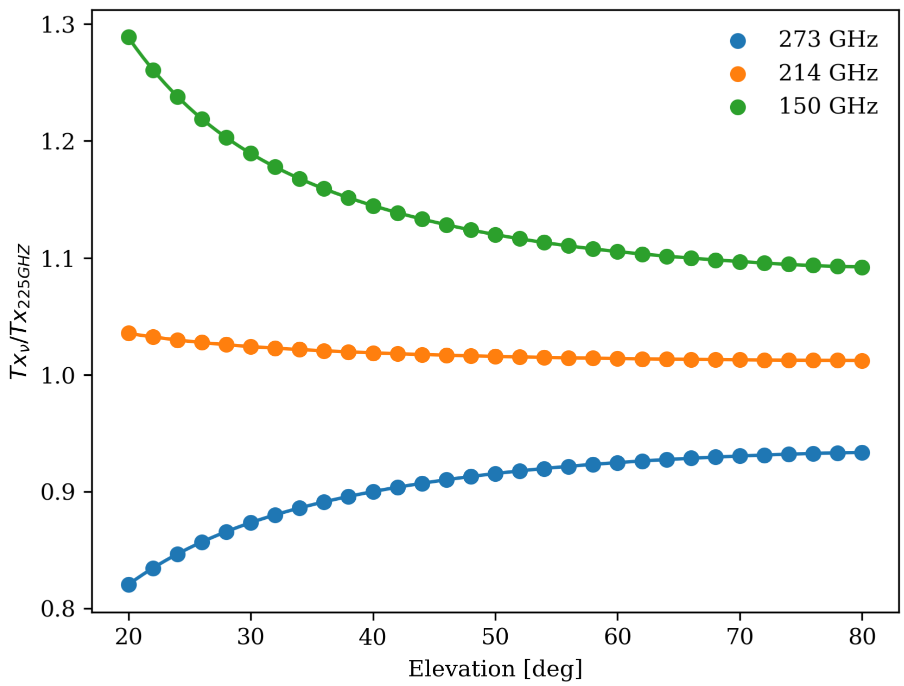

Figure 3.14. The data points are the simulated atmospheric transmission as a function of elevation for the 3 TolTEC bands. The lines are order 6 polynomial fits to the points. The atmosphere model is assuming 25 percentile observing conditions at the LMT site.

The extinction correction is then applied to the FCF as in Equation 3.12 on a per-sample basis. A mean value of the extinction-corrected FCF is stored in the `citlali::PTCData` class to calibrate the detector sensitivities if these are used during detector weighting.

3.12.2.3 Synthetic Timestreams

The timestream and mapmaking algorithms implemented in `Citlali` impose a non-unity transfer function on the data in the time and map domains that will filter the incoming signal. In particular, the PCA-based atmospheric removal process can act similarly to a high pass filter and remove flux from extended emission near the size of the instrument field-of-view. It also subtracts a component of the flux from bright point sources. To enable measurements of the transfer function without needing to create full simulated datasets, `Citlali` creates a parallel set of noiseless timestreams with an injected synthetic source that is reduced in the exact same way as the real timestreams. All subsequent timestream reduction algorithms other than despiking are applied and mapmaking of the source is also performed.

The dimensions of the synthetic timestreams are the same as that of the data timestreams and they use the same pointing solution. The built-in options for the synthetic source type are either a two-dimensional Gaussian or Airy disk with unit amplitude. The source is placed at the fitted detector centroids from the APT such that it will be at the center of its map. The average FWHM for each detector array is used for the source width but can also be manually entered into the configuration file to allow for the transfer function effects on extended emission to be determined. In the standard configuration, the source approximates the point spread function (PSF) of each band. Synthetic source timestreams with the standard Gaussian source are plotted in Figure 3.15. A FITS image can also be input to the pipeline which

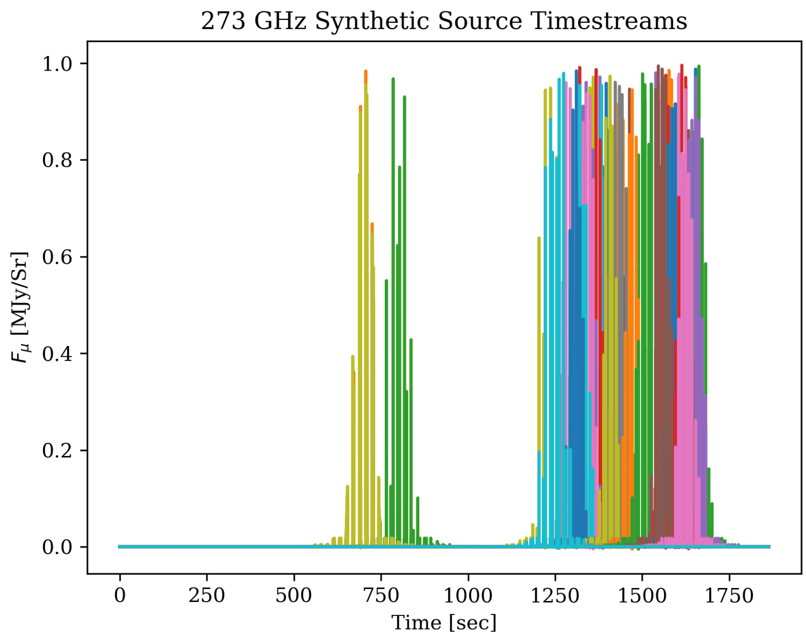

Figure 3.15. Synthetic source timestreams generated by `Citlali` for a number of detectors. The timestreams include only a 2D Gaussian at the coordinates corresponding to the center of the map and have no added noise.

will be transformed into a timestream via inverse mapmaking and allows for the characterization of the pipeline transfer function on more complex source topologies.

3.12.2.4 Despiking

Cosmic rays are a concern at millimeter wavelengths and can introduce large spikes into timestream data that bias estimates of the timestream variance which are used in determining map weights. `Citlali` searches for cosmic ray spikes by differencing adjacent timestream samples and identifying those that deviate from the mean value by a user-determined multiplicative factor of the standard deviation of the signal over the time chunk. This is carried out for each detector iteratively until no further spikes are found. The number of samples corresponding to the estimated decay time for the spike is then flagged and not used during atmospheric cleaning or mapmaking. If timestream filtering is enabled, the size of this region is instead matched to the size of the filter to prevent other samples from being corrupted.

`Citlali` can replace the flagged regions with a linear interpolation between the spiked region endpoints when continuous data is required. Gaussian noise equivalent to the estimated noise of the unflagged data is added.

3.12.2.5 Timestream Filtering

A bandpass filter is included to filter out noise contributions to the data. These range from variations at frequencies greater than D/lambda, where no signal of astrophysical signal is present, and the low-frequency 1/f noise from the detectors and atmosphere. `Citlali` uses a finite impulse response (FIR) filter with a Kaiser-Bessel window for its bandpass filter (Kaiser 1974). The Kaiser-Bessel coefficients are given by

$$w_i = \frac{I_0(\alpha\sqrt{1-\left(\frac{i-N}{N}\right)^2})}{I_0(\alpha)}, \tag{3.16}$$

where I_0 is the zeroth order Bessel function of the first kind, N is the length of the filter, and α is the window shape parameter. For a desired attenuation factor, a_{gibbs}, the shape parameter can be written as

$$\alpha = \begin{cases} 0 & \text{if } a_{gibbs} < 21 \\ 0.1102(a_{gibbs} - 8.7) & \text{if } a_{gibbs} > 50 \\ 0.5842(a_{gibbs} - 21.0)^{0.4} + 0.07886(a_{gibbs} - 21.0) & \text{if } 21 \leq a_{gibbs} \leq 50. \end{cases} \tag{3.17}$$

The coefficients of the impulse response of the ideal filter are

$$A_i = \frac{\sin(2\pi i f_{\text{upper}}) - \sin(2\pi i f_{\text{lower}})}{\pi i}, \tag{3.18}$$

with f_{upper} and f_{upper} being the bandpass edge frequencies normalized by the Nyquist frequency. The final windowed function coefficients are then simply the product of equations 3.16 and 3.18.

Setting f_{lower} or f_{upper} to 0 results in either a lowpass or highpass filter respectively. `Citlali` uses the `Eigen::Tensor` convolution routine to convolve the filter with the timestream data on a per-time chunk basis. A lowpass filter with a similar setup to the ones used during TolTEC commissioning is shown in Figure 3.16.

3.12.2.6 Decimation

`Citlali` can decimate the timestream data by an integer factor to improve the performance of later reduction algorithms and to reduce the overall memory bandwidth of the pipeline. This is useful for generating quick-look data products on the memory-limited timely analysis workstations. Decimation must be used in conjunction with a lowpassing filter to prevent aliasing of high-frequency artifacts; together decimation and lowpassing operate as an anti-aliasing filter. All per-sample vectors

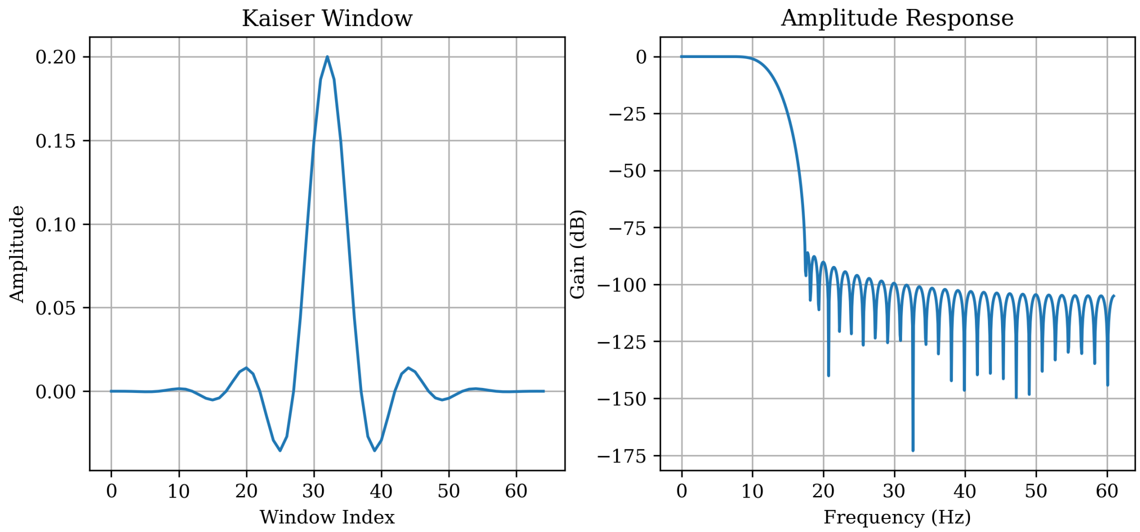

Figure 3.16. Left: Kaiser Bessel window assuming a windows size of 32 points, an upper cutoff frequency of 20 Hz, and $\alpha = 8.6$. Right: Frequency response of the lowpass filter from the left panel.

and matrices in the time chunk class are decimated in the same way as the timestream data.

3.12.3 Processed Timestream Reduction Algorithms

As with the previous section, here I outline the main reduction algorithms that occur after and including the atmospheric filtering in order of operation. These begin after the transformation of the `citlali::RTCData` into the `citlali::PTCData` time chunk classes. These are managed by the `citlali::PTCProc` class.

3.12.3.1 Atmospheric Filtering

The contribution from atmospheric thermal emission at millimeter wavelengths can be many orders of magnitude brighter than the observed astrophysical source flux, making it the primary source of noise in the raw data, while also having the added complexity of varying in both time and as a function of position on the sky. It constitutes the strongest mitigating factor to recovering an unbiased estimate of astrophysical signals with ground-based millimeter cameras. It has two key characteristics which are that it is not stationary like signals of cosmic origin and is correlated across many detectors. To illustrate this, Figure 3.17 shows detector timestreams before atmospheric removal, each of which has similar, correlated long wavelength modes in addition to large amplitude variations. A plethora of techniques have been developed for atmospheric removal including template-based subtraction, principal component analysis, most correlated pixel methods, and maximum likelihood mapmaking.

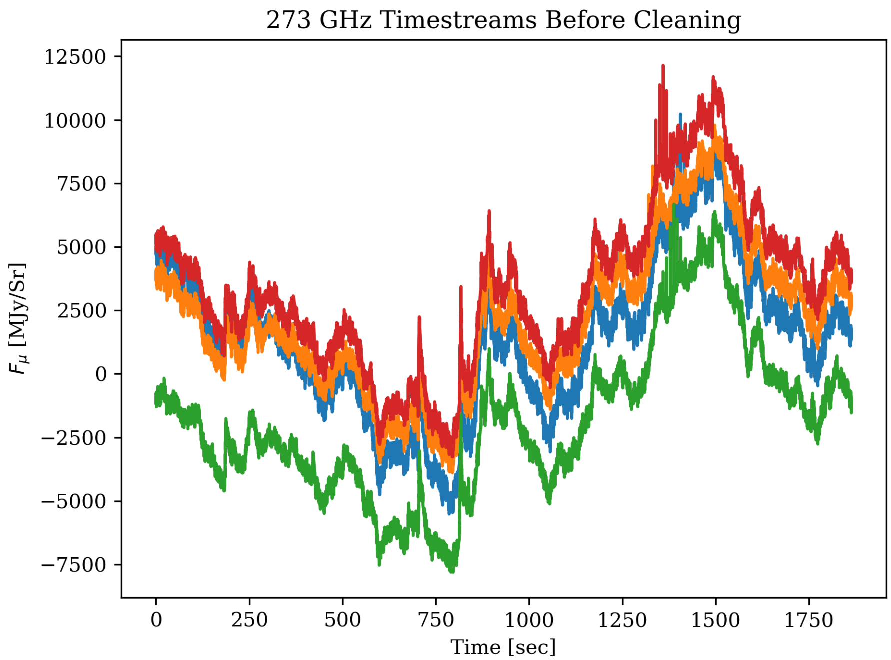

Figure 3.17. Flux calibrated and extinction corrected timestreams from a beammapping observation of J1159+292 for several detectors on the 273 GHz array prior to atmospheric removal. The negative fluxes occur due to the KID tune procedure which sets a relative zero point for the fluxes.

`Citlali` employs the principal component analysis-based approach and is based on the algorithm developed for the AzTEC data reduction pipeline (Scott et al. 2008). PCA is a statistical technique that transforms data into a set of orthogonal components representing the directions of maximum variance, often used for dimensionality reduction. First, the timestreams are mean subtracted and the detector covariance matrix $d^T d$ is calculated from the $N_{samples} \times N_{detectors}$ time chunk matrices d. The matrix $d^T d$ has dimensions $N_{detectors} \times N_{detectors}$ and encodes the correlations among the detectors. Samples and detectors flagged in the APT or by the despiker and timestream variance cuts are zeroed out prior to computation and therefore have no impact on the elements of the correlation matrix. The covariance matrix is then eigenvalue decomposed using the SPECTRA library which will calculate a user-specified number of the largest eigenmodes. The timestream matrices are then reconstructed with these eigenvalues then zeroed out thus removing their contribution from the data. The largest eigenvalues map to the strongest correlations amongst detectors within the data, so this method works under the assumption that the atmosphere is the largest correlated signal across most detectors and that the correlations it introduces do not change on timescales shorter than the length of the time chunk. This assumption holds well for point sources, where only a handful of detectors observe the source at a time, though they are still diminished. For extended emission, the coupling with the atmosphere will be greater and introduce regions of negative or over-subtracted signal into the maps. An iterative mapmaking technique is required to recover the subtracted signal in this case when using PCA. The result of applying this cleaning approach to the detector timestreams of Figure 3.17 is shown in Figure 3.18, where most of the correlated low-frequency signal is removed, thus producing flat, zero-mean source-dominated timestreams. The removal of low-frequency correlated noise can be seen in the detector timestream PSDs of Figure 3.19.

An important consideration is which subset of detectors to calculate and remove the eigenvalues from for each chunk. TolTEC presents 3 obvious choices for subdivision with these being across all detectors and arrays, for each array, or each network. Large correlations across each network's detectors may either be averaged out over all networks on that array or dominate those of other networks which generally necessitates a higher number of eigenvalues to be removed when computed over each array.

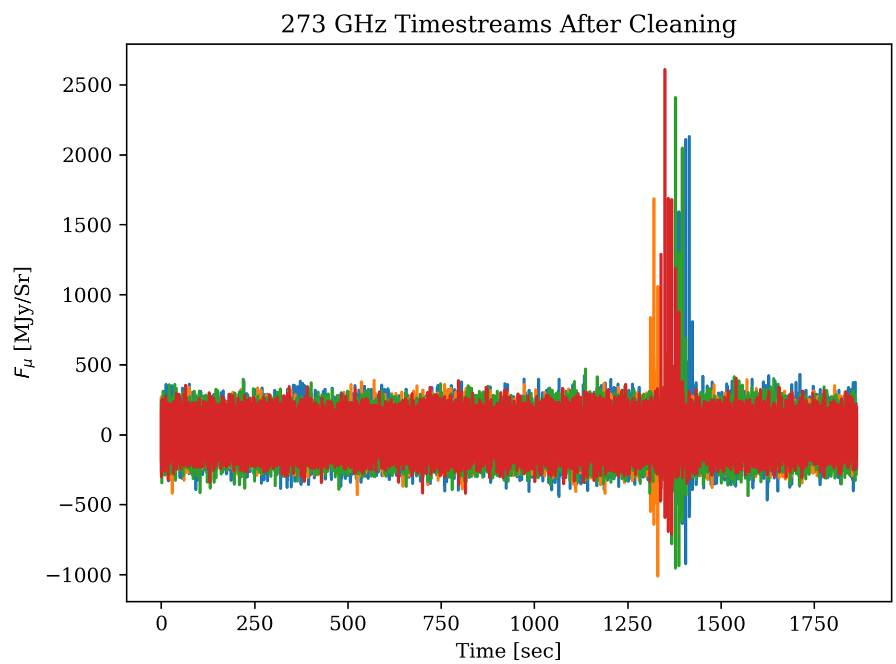

Figure 3.18. Timestreams from the same detectors as those in Figure 3.17, but after the PCA atmospheric cleaning has been applied on a detector per-network level. The peaks are the bright point source J1159+292. The negatives occur due to the mean subtraction before PCA.

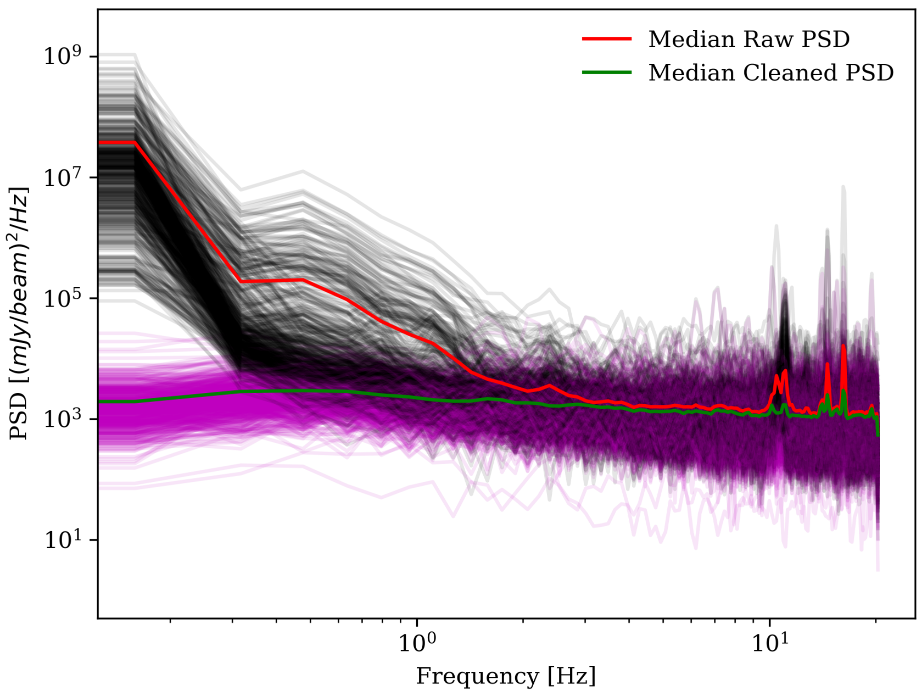

Figure 3.19. Timestream PSDs from the same observation as Figures 3.17 and 3.18. The black and magenta curves are PSDs from the same individual detectors before and after cleaning. The red and green curves are the median PSDs of the plotted detectors for both cases.

Calculating and removing eigenmodes across all arrays was empirically found to be less effective on commissioning data than either array or network level cleaning due to this effect while also being slower due to the increased covariance matrix dimensions.

The eigenmodes calculated from the timestream data are also subtracted from the synthetic source timestreams to enable corrections for lost flux in the post-mapmaking filtering stage and for transfer function calculations.

3.12.3.2 Detector Weighting

`Citlali` calculates a weighting for each detector to be used in the mapmaking algorithms. For most observations, the weight can be defined as the inverse of the variance of the unflagged samples in the time chunk for that detector such that noisy detectors will be given lower weights in the final maps. However, in observations of bright sources, the variance can be skewed for time chunks that include the source, resulting in the corresponding pixels in the map being heavily downweighted. This effect is illustrated in maps constructed from the timestream weights shown in Figure 3.20. For these cases, an estimate for the weight of each detector can be derived from the beammap observations's detector sensitivities, S, which can be calculated from $1/(f_S \times S^2)$, where f_S is the timestream sampling frequency after decimation.

3.12.3.3 Detector Flagging

The timestream variances calculated across each detector and time chunk are used to identify outlier detectors which are then either completely flagged within the current time chunk or downweighted to the median weight of all detectors not flagged in the APT. On commissioning data, when using a 3σ upper limit to define outliers, this usually impacts around 100-200 detectors for each of the arrays. Some examples of outlier detectors that were flagged during commissioning observations are plotted in Figure 3.21.

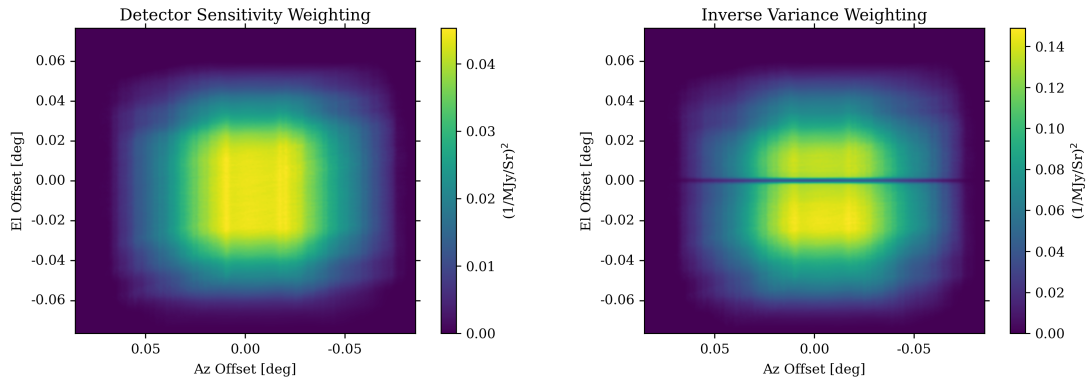

Figure 3.20. Pixel weight maps from a beammapping observation reduced using weights calculated from the detector sensitivities (left) and timestream variances (right). The samples along the scans that see the bright point source are down-weighted due to the mean subtraction that occurs during timestream reduction.

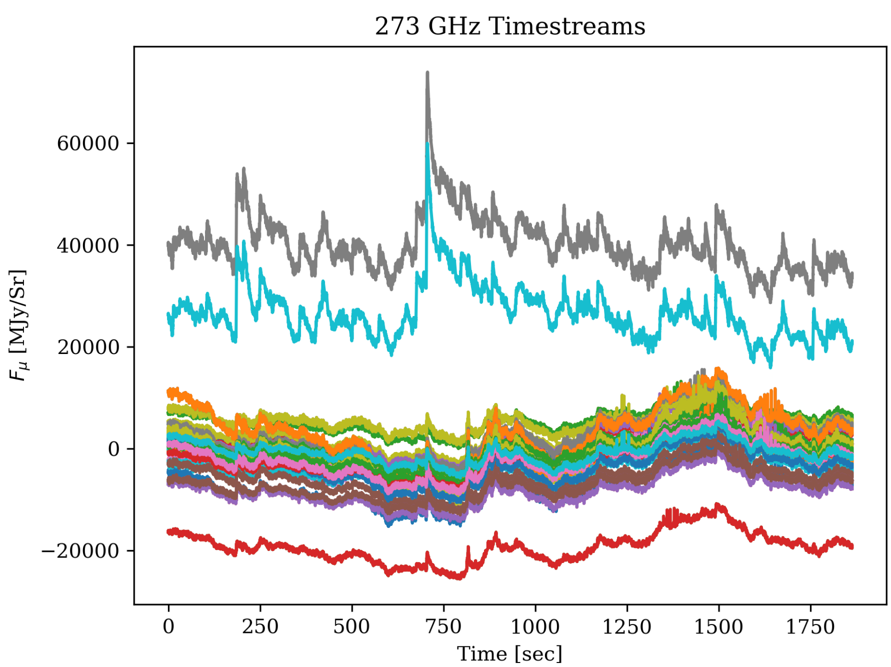

Figure 3.21. Similar to Figure 3.17, but also including outlier detector timestreams (gray, light blue, and red) that would normally be flagged during beammapping or by `Citlali`'s weight cuts.

During commissioning, a software issue in the KID resonator identification algorithm was found which resulted in tones with nearby resonator frequencies being merged and counted as a single resonator. A secondary flagging stage was therefore added to `Citlali` which flags resonators whose frequencies are within a user-specified limit (typically 60 kHz). Examples of nearby tones are shown in Figure 3.22. Approximately 40% of all detectors were merged with most of these cases (98%) being flagged during the beammapping reduction. An additional 2% of the total detector number are typically added on top of the APT flags during science and non-beammap calibration maps due to the merging. This issue was corrected before the 2024 commissioning commenced.

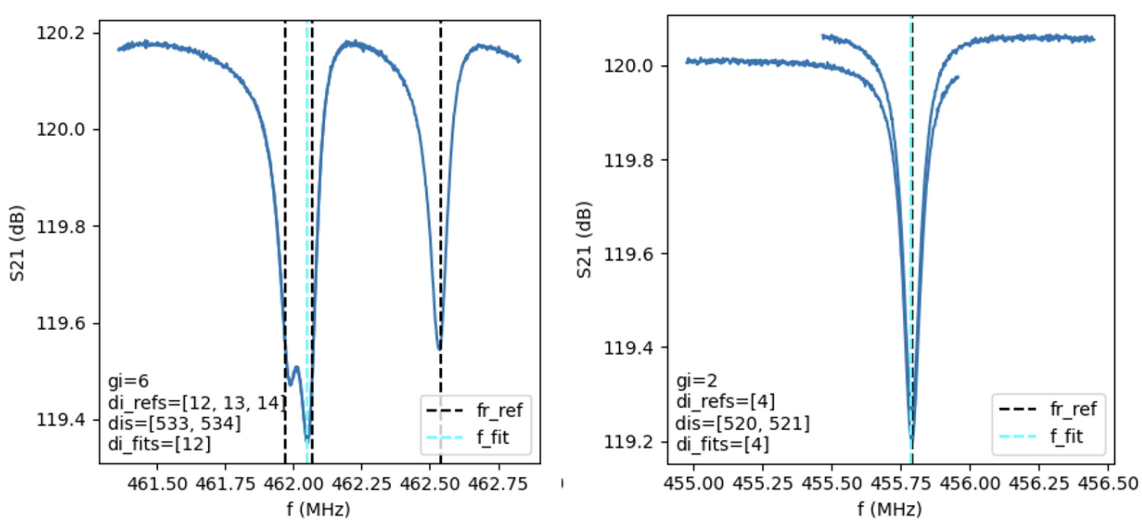

Figure 3.22. Examples of colliding KIDs resonances that overlap one another in frequency space. Image Credit: Zhiyuan Ma

3.12.4 Mapmaking

Following the mapmaking formalism outlined in (Benton 2015), the relationship between the timestream data recorded by the instrument and the signal on the sky can be succinctly expressed in the form

$$d = Pm + n, \tag{3.19}$$

where d is a vector of all timestream data of length $N_{samples} \times N_{detectors}$, P is the pointing matrix, m is a vector of length $N_{pixels} \times N_{pixels}$ corresponding to the map of the sky, and n is the noise contribution to each element of d. The pointing matrix maps each sample of d to an element or pixel of m and is therefore a large matrix of $N_{samples} \times N_{detectors} \times N_{pixels}$ elements, but is a sparse matrix, with only a few nonzero elements across each row. The χ^2 statistic for 3.19 is

$$\chi^2 = (d - P\hat{m})^T \Sigma^{-1} (d - P\hat{m}). \tag{3.20}$$

In this case, $\hat{m}$ is an estimate for the map and Σ is the covariance matrix of the noise vector n. Minimizing this gives the Generalized Least Squares (GLS) solution as

$$\hat{m} = (P^T \Sigma^{-1} P)^{-1} P^T \Sigma^{-1} d. \tag{3.21}$$

If the noise is Gaussian, this is equivalent to the maximum-likelihood solution. Two difficulties for evaluating this in practice are the challenges of storing the matrices in memory and inverting them, as well as acquiring a statistically rigorous estimate of the noise covariance. As an example, for a 30-minute TolTEC beammap observation measuring intensity only at the full sampling frequency of 488 Hz and with 1″ pixels (300″ × 300″ maps), P and Σ both contain 6×10^{14} elements, requiring 5000 TB of available memory to hold the full observation.

Mapmakers usually rely on iterative numerical techniques like the Preconditioned Conjugate Gradient (PCG) method to when attempting to solve Equation 3.21 without storing the zero elements. A simple estimate of the noise can be derived by assuming it is white, making all elements of Σ other than the main diagonal zero. The main diagonal can then be estimated from the variance in the timestreams or from their power spectra. This assumption is not strictly valid owing to the 1/f noise contribution from both the atmosphere and the detector themselves which must first be filtered through highpassing and atmospheric subtraction. This filtering will result in $\hat{m}$ being a biased estimator of the source flux.

The noise can also be estimated assuming that it is stationary and that samples far from one another in time are independent. This makes the covariance matrix Σ a band-diagonal symmetric Toeplitz matrix with nonzero elements along the main diagonal and those directly adjacent to it (Benton 2015).

The mapmaking strategy implemented for `Citlali` is an attempt to strike a balance between "correctness" and reduction efficiency. A fast mapmaker that can be run with little computational resources is required for on-site reductions, while a mapmaker that can make use of HPC resources to produce maps sufficient for scientific analysis is also needed. `Citlali` therefore includes two mapmaking implementations that partially reflect this trade-off, an iterative algorithm for timestream and mapmaking reductions, and the ability to output partially reduced timestreams to be used as inputs to standalone maximum-likelihood mapmakers like Minaksi or TOAST3.

3.12.4.1 Naive Mapmaker

The naive mapmaker is the approach derived under the assumption of white noise. The mapmaking equation can then be solved implicitly by looping through the timestream data and coadding each detector sample into the map pixel that contains its pointing sample (Figure 3.23). This results in a nearest-neighbor gridding of the data onto the map array which is equivalent to convolving the timestream data with the rectangular function $\Pi(t)$, which in the Fourier (u,v) plane is given by

$$\mathcal{F}\left(\Pi(t)\right)(u,v) = \frac{\sin(\pi t u)}{\pi t u} \times \frac{\sin(\pi t v)}{\pi t v}, \tag{3.22}$$

where t is the pixel size assuming square pixels. A signal map pixel value then becomes the weighted average of all samples that fall within its bounds:

$$S_p(i,j) = \frac{\sum_{d=0}^{N_{\mathrm{dets}}} (s_d \cdot w_d \cdot \delta(x_p - x_d))}{\sum_{d=0}^{N_{\mathrm{dets}}} w_d \cdot \delta(x_p - x_d)}. \tag{3.23}$$

In this expression, p is the current pixel, d is the current detector, and $\delta(x_p - x_d)$ determines if the detector pointing is within the current pixel. The actual timestream flux is s_d. The weights w_d are derived as described in Section 3.12.3.2. Maps of the synthetic source are made in the same way only replacing the timestream data sample with the source sample for that detector. Weight (inverse error) and coverage (integration time per pixel) maps are similarly produced.

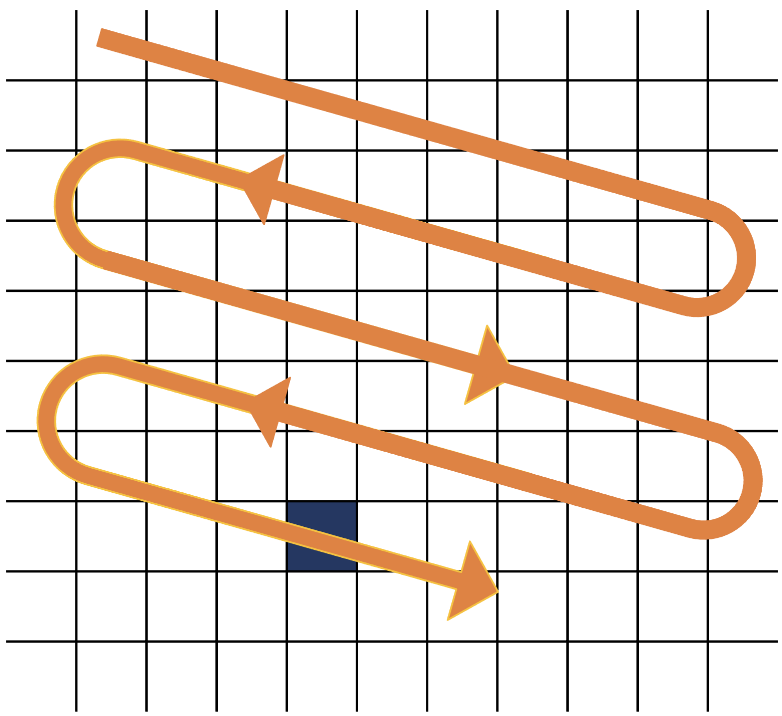

Figure 3.23. Diagram of the naive mapmaking algorithm. The grid represents individual pixels in a map and the orange curve is the trajectory of a detector given the mapping pattern and its offset position. The shaded cell is a pixel that a detector sample is added into. For the naive mapmaker, only pixels whose boundaries the sample falls into receive a contribution from that sample.

The naive mapmaker is the fastest and simplest mapmaking algorithm possible and is used for the on-site reductions carried out with `Citlali` and can produce maps many times faster than the integration time for most observation categories.

3.12.4.2 Jinc Filtered Mapmaker

An aspect of the data not accounted for in the naive mapmaking method is the presence of receiver or background noise at frequencies higher than the fastest sampling frequency of the sky which for a telescope of diameter of diameter D and detectors sensitive to photon wavelength Λ is λ/D. In conjunction with over-sampled maps, this produces additional pixel-pixel noise variations that are not representative of the true noise level and reduce the map S/N. This noise is strictly not due to source variations and can safely be filtered without removing source flux. This can be accomplished during mapmaking by convolving each data sample with a function that has a unity value below λ/D and is truncated at higher spatial frequencies (Mangum et al. 2007). The function used in `Citlali` is

$$C(r') = 4 \times \mathrm{jinc}\left(\frac{2\pi r'}{a}\right) \times \exp\left(-\left(\frac{2r'}{b}\right)^c\right) \times \mathrm{jinc}\left(\frac{3.81706r'}{r_{\max}}\right). \tag{3.24}$$

The jinc function is $J_1(x)/x$, where $J_1(x)$ is the Bessel function of the first kind. The first jinc function is Fourier transform of the indicator function $\mathbf{1}_A(x)$, which is zero for values larger than $2\pi r'/a$, with r' being $r' = rD/\lambda$. The parameter a is a scaling factor. The exponential function with shape parameters b and c is added to dampen the jinc sidelobes. The final jinc function truncates the response at r_{max} which is some multiple of λ/D and is the maximum radius to which the function is calculated. The parameters a, b, and c control the balance between the S/N across the map and the response at higher spatial frequencies. The values used in `Citlali` are a=1.1, b=FWHM/3, and c=2 where FWHM is the theoretical beam width for each band, which is the midpoint of the S/N and frequency response trade-off. The response of the filter as a function of λ/D for the default TolTEC parameters is illustrated in Figure 3.24.

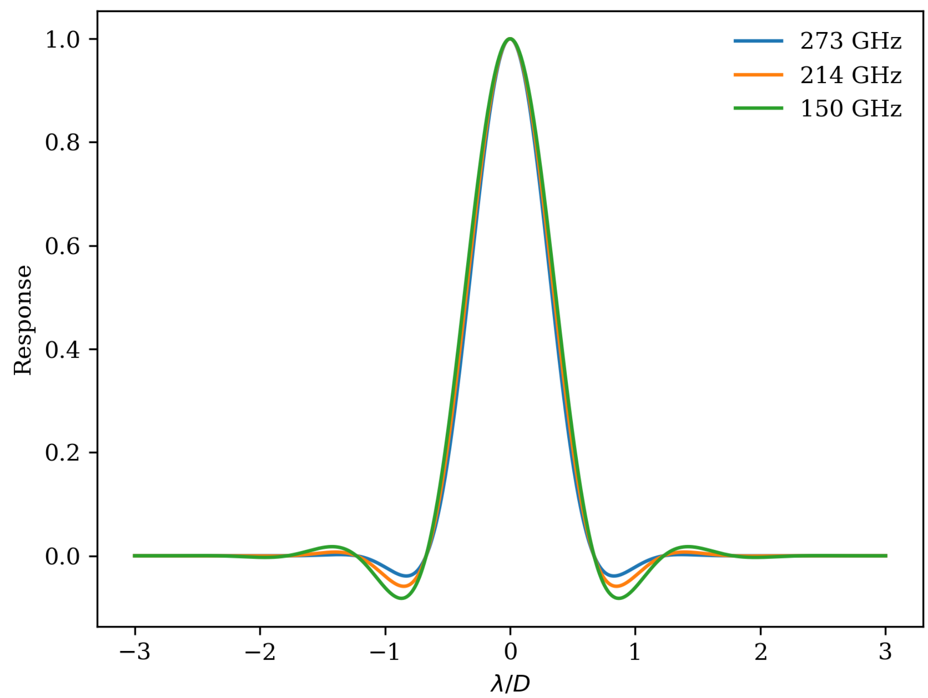

Figure 3.24. Response of the jinc filter given by Equation 3.24 for each of the TolTEC arrays assuming the nominal values of the shape parameters a, b, and c and $r_{max} = 1.5$.

The mapmaker using this convolution is referred to as the Jinc filtered mapmaker and is implemented as `Citlali`'s second mapmaker. Each row of the pointing matrix P, in Equation 3.19 will now be populated by additional elements and may no longer be sparse depending on the choice of r_{max}. Similar to the naive mapmaker, Equation 3.21 is solved implicitly in a loop over detectors and samples. To improve performance, the value of Equation 3.24 is not calculated for all pixels within r_{max} for each sample. Instead, Equation 3.24 is precomputed in the pipeline phase 1 as a matrix that is then multiplied by the data sample value, and added to the same-sized block of pixels around the nearest pixel, which can make use of Eigen's vectorization capabilities. This approach is illustrated in Figure 3.25. Regardless, the Jinc mapmaker remains much more computationally expensive than the naive mapmaker with most reductions taking 3-4 times longer. The Jinc mapmaker therefore enforces per-detector parallelization rather than per-time chunk to maximize the number of compute threads being used and will be most efficient when used for short observations.

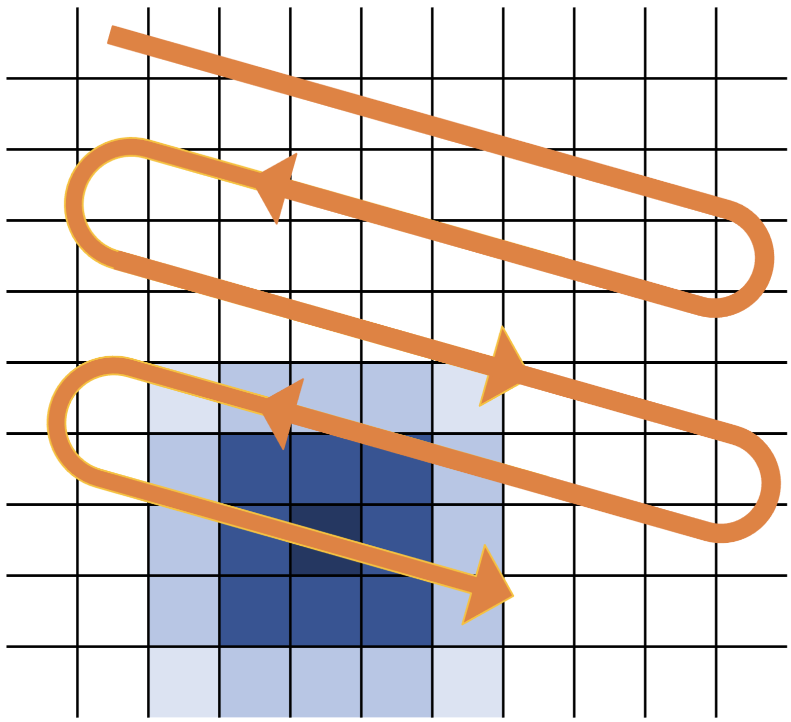

Figure 3.25. Same as Figure 3.23, but for the jinc mapmaker. The sample falls within the bounds of the darkest shaded cell, but all shaded pixels are added to due to the convolution of the sample with Equation 3.24. The block of shaded pixels is precomputed and multiplied by each sample's weighted signal value during mapmaking.

Detector weighting and normalization are applied in the exact same manner as the naive mapmaker. The synthetic source and weight maps are also convolved with Equation 3.24. A comparison of the differences between the naive and jinc mapmaker from a commissioning observation of a point source is illustrated in Figure 3.26, where the unfiltered naive mapmaker result is characterized by a greater pixel-pixel scatter relative to the unfiltered jinc mapmaker curve.

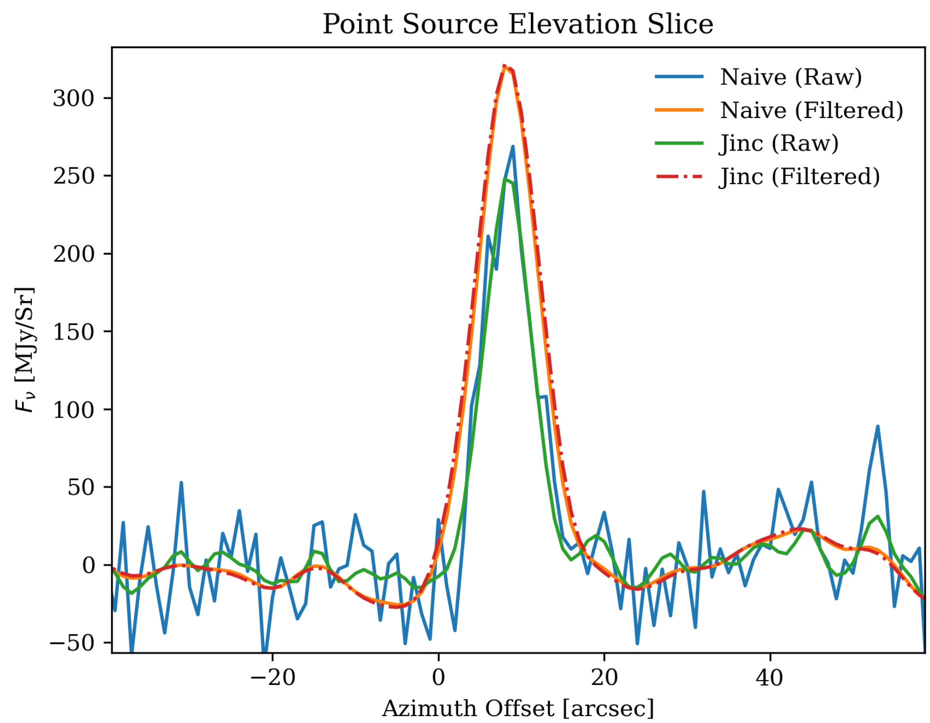

Figure 3.26. Elevation slice through maps of a bright point source (J1041+061) made with the naive mapmaker and the Jinc mapmaker. The filtered maps refer to those passed through the optimal point source filter (Section 3.13.1).

3.12.4.3 Maximum Likelihood Mapmakers

As discussed in Section 3.7, `Citlali` can output the data timestreams and their associated metadata before and after atmospheric filtering. These partially reduced and flux-calibrated output timestream files can be used as inputs to existing standalone maximum-likelihood mapmakers that solve Equation 3.21 using iterative PCG techniques. This allows these external mapmakers to take advantage of `Citlali`'s KIDs processing, timestream reduction, and flux calibration which have already been implemented and optimized while leveraging their own mapmaking capabilities. Two maximum-likelihood mapmakers have been integrated into the TolTEC software stack, with these being the Minkasi maximum-likelihood mapmaker (Sievers 2023; Romero et al. 2020) and the TOAST3 software framework (Kisner et al. 2021) which is also used for atmospheric simulations. These are not replacements for the naive or Jinc mapmakers, however, as the memory requirements and reduction time for maximum likelihood mapmakers are typically more than an order of magnitude higher.

Most of the integration testing and commissioning analysis has been carried out with the Minkasi mapmaker, with TOAST3 efforts being primarily focused on recovering polarized emission. The Minkasi mapmaker is the data reduction pipeline and mapmaker for the MUSTANG2 project. It uses a singular value decomposition (SVD) of the data timestreams to determine a smoothed measure of the noise power spectra.

3.12.4.4 Noise Maps

`Citlali` can generate noise map realizations using the jackknifing method for estimating map noise properties and for use in its optimal point source filter. The weighted samples in a time chunk are multiplied by a vector of randomly generated $\pm$ 1 and a map is made for each realization using the same mapmaking algorithm as the signal map. The vector of random deviates can be held constant or vary with each detector, with the latter serving to further randomize out most residual map features to produce lower noise estimates. Example histograms of the distributions of noise pixels for maps made with the Naive mapmaker and the Jinc mapmaker are shown in Figure 3.27.

Since many noise maps may be required (possibly >100 for each array and Stokes parameter), individual maps of every realization are not maintained for both observations and coadded maps if coaddition is enabled. Instead, the noise samples are directly added into the coadded map during mapmaking. This approach is not used

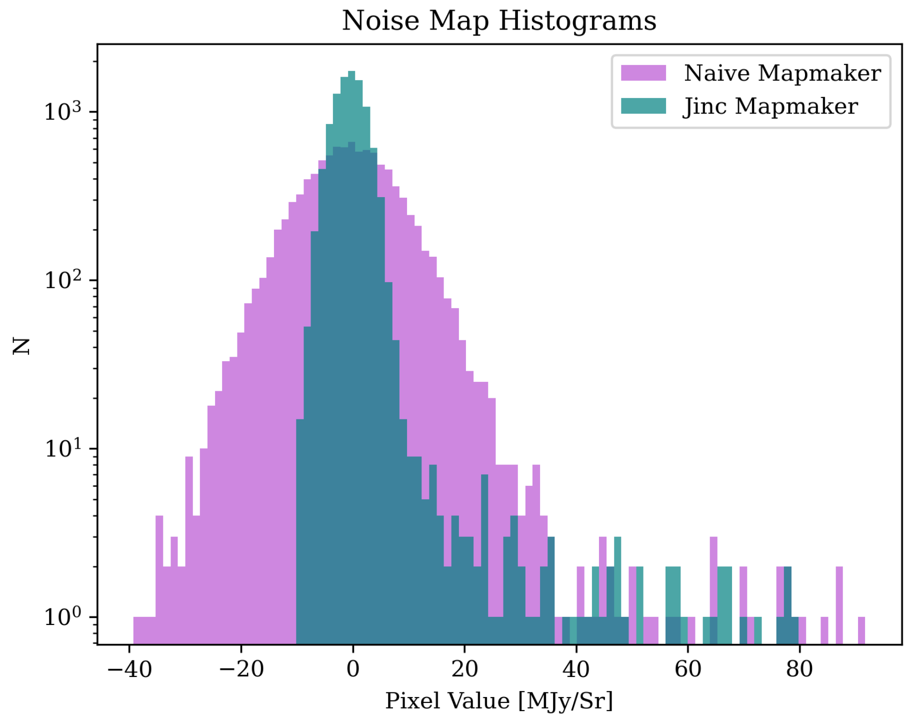

Figure 3.27. Histograms derived from the averages of 10 jackknifed noise map realizations made with the naive and Jinc mapmakers. The original observation was a pointing map of J1041+061. The latter shows less variance due to the removal of high-frequency noise by the mapmaking algorithm. The counts at high flux densities originate from source leakage into the noise maps.

for the other map types owing to the significant data compression that occurs during mapmaking, thereby making it more efficient to coadd in the map domain when possible.

3.12.4.5 Polarimetry Mapmaking

Note: `Citlali`'s polarization reduction remains under active development. As of the time of this writing, the instrumental polarization has also not yet been measured.

As discussed in 2.1, each pixel on the TolTEC detector arrays consist of two orthogonal KIDs that are sensitive to a single linear polarization. As a result of sky rotation during observations, the angle of each KID with respect to the source will rotate allowing any arbitrary source polarization angle to be measured by the detector arrays. Following the notation of Benton 2015, the signal measured by a detector t for a particular sample i can be most generally written as the sum of the signal from each of the linear Stokes parameters I, Q, and U for the current pixel p of map m as

$$d_{i,t} = I_p + Q_p \cdot \cos(2\theta_{total}) + U_p \cdot \sin(2\theta_{total}). \tag{3.25}$$

The angle θ_{total} is the relative angle between the KID and the source polarization angle and is the combination of all angles that produce a full rotation from the detector's reference frame to the sky frame. For TolTEC is given by

$$\theta_{total} = \theta_{El} + \theta_{PA} + 2\beta + \theta_{detector} + \theta_{array} + \epsilon, \tag{3.26}$$

where θ_{El} is the source elevation, θ_{PA} is the parallactic angle, β is the HWP rotator angle if installed, $\theta_{detector}$ is the KID orientation angle (0, 45, 90, 135 degrees), and θ_{array} is the installation angle for each array, which is 90 degrees at 273 GHz and -90 degrees at 214 and 150 GHz. The final angle, ϵ, accounts for any offsets due to misalignments in the optics or array orientations and must be determined empirically.

With the introduction of Q and U, the pointing matrix P from Equation 3.21 now contains 3 elements per row. For the implicit solution of the naive and Jinc mapmakers in `Citlali`, the elements of the 3 x 3 matrix $P^T P$ and the 3 x 1 vector $P^T d$ must now be stored for each pixel. To save memory, `Citlali` re-uses the map matrices it creates for the Stokes Q and U maps for $P^T d$ and overwrites them following the solution of Equation 3.21 for each pixel. Full storage of the pointing matrix or its nonzero elements is not required.

As a result, for polarized reductions, unlike maps only measuring total intensity, the full association of the resonance frequency measured in a tune observation combined with the fitted source centroid found during beammap to the actual KID on the physical array must be determined. The two co-centered orthogonal detectors are deliberately fabricated with resonance frequencies separated from one another to aid in this identification, with all detectors belonging to 4 different frequency ranges or “frequency groups” matched to each orientation angle, which are shown in Figure 3.28. Fabrication errors, effects introduced by the optics, and uncertainties in the beammap fitting will result in a fraction of detectors being unmatched. These are flagged and not included in the Stokes Q or U maps that `Citlali` makes.

A simpler, yet very useful, measurement of the Stokes Q and U signals can be made by producing maps for each of the detector orientations or frequency groups and differencing those of the orthogonal orientations. `Citlali` includes the functionality to output maps of each frequency group for this purpose. This approach, however, cannot take the per-sample HWP rotation angle into account.

3.12.4.6 Map Coaddition

`Citlali` assumes that each observation map in the input list is co-centered on the same sky coordinates when coadding. In this scenario, the zeroth element of the observation maps is offset from the coadded map zeroth element with

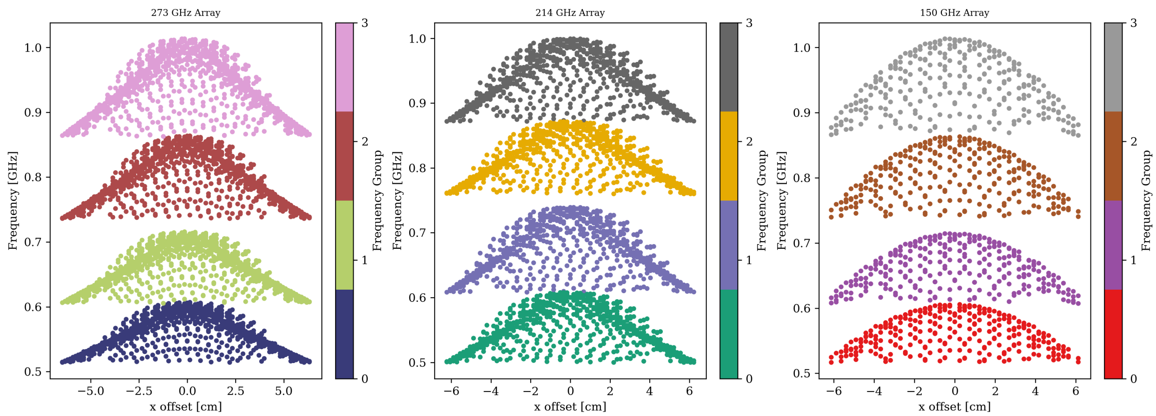

Figure 3.28. Design KID resonance frequencies plotted against the physical location in one direction. The frequencies are intentionally separated into different frequency groups (plotted as colors) offset from one another to allow for easier matching of measured resonance frequencies to the physical detector location.

$$\begin{aligned}\delta_{row} &= 0.5 \times (rows_{coadd} - rows_{obs}) \\ \delta_{col} &= 0.5 \times (cols_{coadd} - cols_{obs}).\end{aligned} \tag{3.27}$$

Here, the subscripts coadd and obs refer to the coadded map and observation map matrices respectively. For an observation signal map S_i with a corresponding weight map W_i, the coadded signal map is

$$S = \frac{\sum_{i=0}^{N_{\text{maps}}} (S_i \times W_i)}{\sum_{i=0}^{N_{\text{maps}}} W_i}. \tag{3.28}$$

The coaddition of the jackknifed noise maps is handled differently compared to the map types. Since no noise maps are allocated for individual observation maps when co-addition is enabled, the calculated noise timestream values are added directly into the coadded noise map matrices and no map-space coaddition is required. They are still normalized by the coadded weight maps in the same manner as the others. This approach is not used for the other map types owing to the significant data compression that occurs during mapmaking, thereby making it much more efficient to co-add in the map domain when possible, particularly when using the slower Jinc mapmaker.

3.12.4.7 Iterative Mapmaker

The effect of the PCA atmospheric cleaning on extended emission is difficult to qualify or correct for in a quantitative manner and will vary greatly from one source to another and between multiple observations under different conditions and instrumental setups of the same source. For this reason, `Citlali` includes an iterative timestream reduction and mapmaker that is similar, yet distinct to the iterative procedure in the beammap reduction engine to mitigate the effects of its PCA reduction. Iterative mapmakers have been developed for use with previous instruments like SCUBA-2 (Chapin et al. 2013), and AzTEC (Liu et al. 2010). The approach implemented in `Citlali` is based on the AzTEC algorithm and is named Flux Recovery Using Iterative Techniques (FRUITloops).

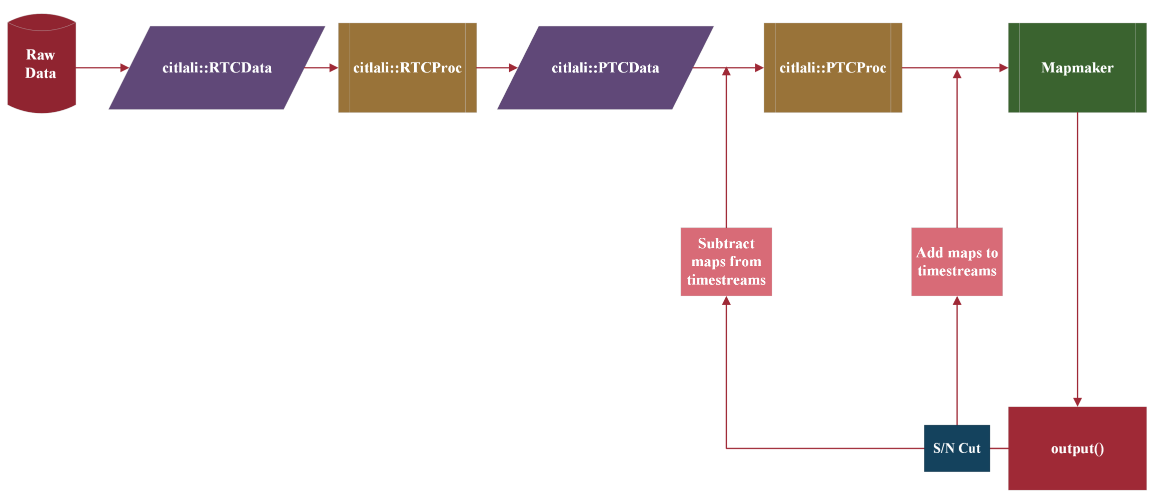

Figure 3.29. Diagram of the `Citlali` FRUITLoops algorithm which wraps the reduction engine functions in a iterative mapmaking loop. For each iteration, high significance pixels from the previous iteration's maps are subtracted and added before and after atmospheric cleaning.

Figure 3.29 shows the flow the FRUITloops algorithm. It consists of a loop around the phase 2 timestream and mapmaking reduction and phase 3 post-mapmaking stages of the pipeline. Phases 0 and 1 are only performed once. During the first iteration, each observation is reduced, data and noise maps are generated and normalized, and coaddition is performed if requested. The map files are then written to disk. On the next iteration, either a S/N, a flux cut, or both are applied to the previous iteration's maps (or the maps in a user-specified path) such that all pixels below those limits are zeroed out. The maps are re-cast into timestream space using an inverse of the naive mapmaker and subtracted from the raw data timestreams after flux calibration. Each observation may have its own raw or filtered map or the coadded versions subtracted to take advantage of the reduced noise level in the latter. Atmospheric cleaning is performed, and the current iteration noise maps are calculated before the same pixel values are re-added to the cleaned timestreams. The remaining processes are then applied as in the first iteration. This is then repeated until a maximum number of iterations is reached.

Provided reasonable S/N and flux limits are used the source flux in the maps will converge with a much higher level of recovery than with a single iteration and the over-subtracted negative regions PCA typically introduces will be minimized. The optimal selection of the free parameters will depend on the source morphology and size, the noise floor, and the relative difference in magnitude between the source and atmospheric amplitudes. Bright extended emission from sources like the Crab Nebula (Messier 1) will require more iterations to converge.

3.13 Post-Mapmaking

The processes that occur after maps have been made and normalized include source fitting for the `citlali::pointing` reduction engine, map PSD and histogram calculation, optional map filtering, and file output. These steps are all enclosed within the FRUITloops mapmaking iteration and will be carried out and their outputs generated for every observation and iteration.

3.13.1 Optimal Point Source Filter

An alternative, more mathematically rigorous scheme to iterative mapmaking for recovering point source flux lost during atmospheric cleaning is to filter the maps using the synthetic source maps as the convolution kernel or filter template. The default synthetic sources are derived from the average of 2D Gaussian fits to each detector's map of the beammapping point source and are therefore representative of the PSF of each band. The synthetic source timestreams are subject to the same atmospheric cleaning as the data and can be scaled to their original unit amplitude to correct for the flux loss incurred by the PCA approach. Under the conditions that the pixel noise is independent and white, this is equivalent to fitting the signal map to the synthetic source and deriving the best-fit amplitudes (Scott et al. 2008).

`Citlali` uses the GLS algorithm developed out in (Perera et al. 2013), whose notation I repeat below. It works on two assumptions. First, the PSF is assumed to be constant across the field. Second, the pixel-pixel noise correlations are a function of the distance between two pixels only. In this case, for pixel p, the χ^2 statistic to be minimized in the optimal filter is

$$\chi_p^2 = \sum_{k,l=1}^{N_{\mathrm{Pix}}} \left(d(x_k) - s_p f(x_k - x_p) W_{kl} \left[d(x_l) - s_p f(x_l - x_p)\right]\right), \tag{3.29}$$

where (k,l) is the row and column of the current pixel in the summation, d is the unfiltered map at (k,l), s_p is the amplitude of the convolution kernel, and $f(x_l - x_p)$ is the value of the convolution kernel at (k,l) when centered at p. The value W_kl is the inverse covariance of the pixel (k,l) and is given by

$$W_{kl} = \frac{1}{N_{\mathrm{Pix}}^2} \frac{1}{\epsilon(x_k)} \left(\sum_{x_k, x_l} \frac{\epsilon^{(2\pi j k_a (x_k - x_l))}}{V^2(k_a)} \right) \frac{1}{e(x_l)}, \tag{3.30}$$

with $\epsilon(x_k)$ and $\epsilon(x_l)$ being the noise standard deviations at (k,l) and $V^2(k_a)$ is the normalized PSD of the data. The inclusion of W_{kl} allows for maps with uneven coverage to be filtered. Owing to the presence of 1/f noise, the PSD will not be flat, and the covariance matrix then becomes band diagonal. Due to the division by $V^2(x_k)$ the noise is whitened, and pixel-pixel noise will be filtered. The maximum likelihood solution for s_p that minimizes Equation 3.29 then becomes

$$S_p = \sum_{k,l} \left(\frac{W_{kl} d(x_k) f(x_l - x_p)}{\sum_{k,l} (W_{kl} f(x_p - x_k) f(x_p - x_l))} \right) = \frac{N_p}{D_p}. \tag{3.31}$$

The convolution kernel used when filtering maps with point sources is usually chosen to be a rotationally symmetric version of the synthetic source map to account for potential variations in the PSF across the field (Scott et al. 2008). The PSD is derived from the average power spectra of the jackknifed noise realization maps. The numerator N_p and denominator D_p are solved separately in `Citlali` and combined, with the latter being solved iteratively. The denominator requires a large number of FFTs to be calculated and can therefore become computationally expensive for large (1 deg^2) maps. `Citlali` employs the FFTW library for the 2D FFT calculations in conjunction with the library's multi-threaded OpenMP-based options.

A consequence of the map filtering is that the beam in the filtered maps is broadened relative to that of the raw maps due to the convolution with the synthetic source. `Citlali` filters the signal, weight, synthetic source, and all noise maps. The filtering kernel is not restricted to only the simulated PSF map but can also be used with more complex kernels. `Citlali` permits 2D Gaussian and Airy patterns of arbitrary widths to be used for the recovery of more extended emission. If a FITS image is used to create the synthetic source maps those will also be used here.

Figure 3.30 shows maps produced by the unfiltered and filtered naive and jinc mapmakers. Pixel-pixel variations are evident in the filtered naive result but largely reduced when using the other approaches. Filtering in conjunction with either the

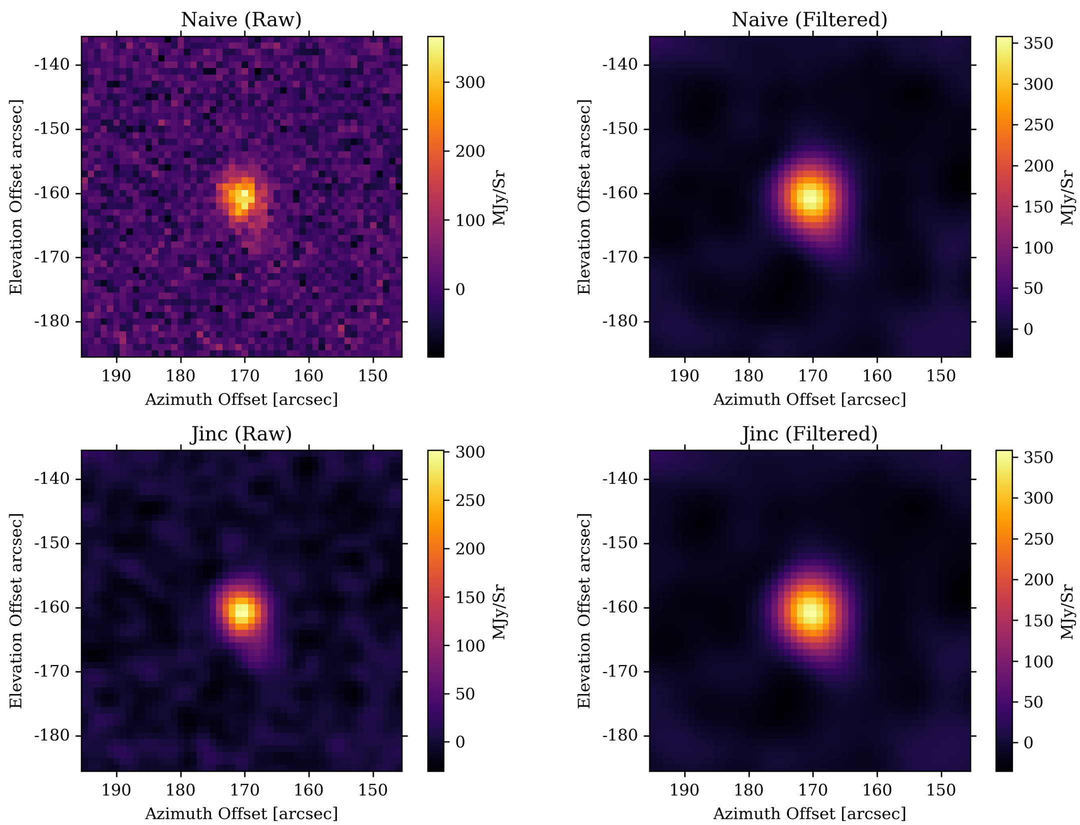

Figure 3.30. Maps of the source J1041+061 made with the naive (top row) and Jinc (bottom row) mapmakers as well as with (right column) and without (left column) the optimal point source filter.

naive or Jinc mapmakers are nearly identical. The recovery of lost subtracted flux and the broadening of the beam can be observed in an elevation slice through a synthetic source map in Figure 3.31.

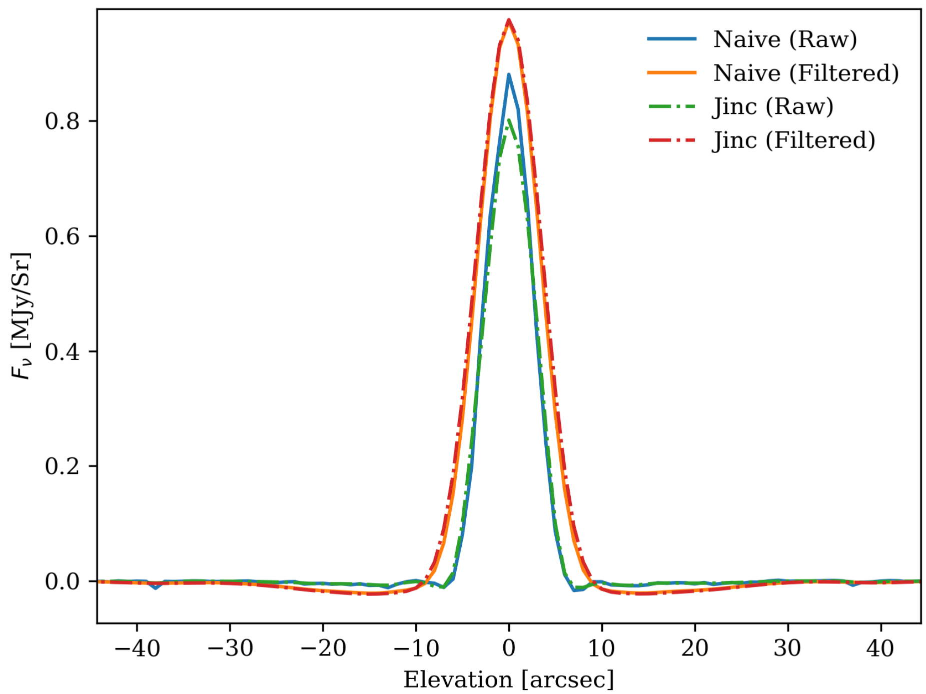

Figure 3.31. Elevation slice through raw and Wiener filtered maps made from the synthetic source timestreams with the naive and Jinc mapmakers.

3.14 Benchmarking and Performance

The memory footprint of `Citlali` is determined primarily by the observation and coadded map arrays, with the total memory allocation depending upon the pixel size, map dimensions, and the number of maps. The latter varies based on the type of reduction, the number of detector arrays enabled, whether noise maps are requested, and whether Stokes Q and U maps are being generated. Temporary memory allocations from the reduction and mapmaking stages will incur additional overhead, with the largest contribution to temporary memory coming from the atmospheric cleaning algorithm. This is due to its computation of detector correlation matrices and eigenvalue decompositions.

For an equivalent reduction setup and observation parameters, reductions using the `citlali::beammap` mode will require the most memory compared to science or pointing modes. This higher demand stems not only from the fact that each detector's signal and weight maps are stored in memory but also because a full copy of the partially reduced timestreams is needed for the iterative reduction algorithm. Observations are typically decimated where possible, so the overhead from the timestream copy is usually less than that of the raw data. Memory usage was measured for pointing, beammap, and two science observations—Messier 1 and Monoceros R2—across different thread counts. Details of the observation configuration for the science reductions are provided in Section 4.2 for Messier 1 and in Section 4.3 for Monoceros R2. The observations were processed using the currently adopted default reduction parameters, as determined from commissioning observations. The results for a single thread are presented in Table 3.1. There is little variation with an increased number of threads, indicating that the memory footprint is dominated by the memory allocated for maps and timestreams outside of parallel regions and that the I/O speed is roughly equal to or less than the reduction time of individual time chunks.

`Citlali` has two potential computational bottlenecks: file I/O and memory bandwidth. Due to `Citlali`'s data streaming model, while one time chunk is being read from the disk, one or more are being reduced. Therefore, if the disk read speed is slower than the reduction time for a single time chunk within the initial `grppi::farm` instance, the pipeline will be I/O bound. An I/O bottleneck is more likely in `citlali::beammap` reduction mode, due only to the raw time chunk processing being performed while the raw data is read from disk.

Most stages in `Citlali` are, in principle, embarrassingly parallel. However, there is a parallelization overhead as more processors are utilized, meaning that reduction

Table 3.1. Single Threaded Memory Overhead for `Citlali`.

Observation Type	Integration Time	Map Maker	Memory Used
Pointing	60 sec	Naive	3.25 MB
Pointing	60 sec	Jinc	3.32 MB
Beammap	30 min	Naive	26.68 GB
Messier 1	12 min	Naive	1.17 GB
Messier 1	12 min	Jinc	1.15 GB
Monr2	10 min	Naive	546.53 MB
Monr2	10 min	Jinc	368.12 MB

time will not scale linearly with processor count. Particularly for `citlali::beammap` reductions, parallelized regions must necessarily be divided during the transformation from the time domain to map space, which limits the code's scalability. Additionally, since parallelization primarily occurs across time chunks, the number of time chunks sets an upper limit on the maximum number of cores that can be employed.

Figure 3.32 illustrates the speed-up factors, defined as the ratio of the reduction time for one thread to that of other thread counts, for different reduction modes. Figure 3.33 shows the actual reduction times for different numbers of cores. Only physical compute threads were used. The Monoceros R2 observation's naive reduction shows the greatest speed-up factor due to its 22 time chunks, which facilitates easy parallelization across many threads. The Messier 1 naive reduction shows lower gains because it is a coaddition of multiple observations that are reduced sequentially. For pointing observations with 10-second time chunks, each has at most 6 chunks, meaning that no efficiency gains are achieved beyond approximately 6 threads. The speed-up factor for `citlali::beammap` reductions is somewhat deceptive; its runtime is extended by sequential processing steps, including preliminary observation setup tasks like timestream alignment and map allocation, as well as file output time. For the parallelized stages, the speed factor is approximately 2.5 times for 28 threads. Jinc mapmaker reductions show poorer speed-up factors compared to the naive mapmaker, owing to a lower level of parallelization. Only the mapmaking stage itself is multi-threaded with Jinc mapmaking enabled. However, the current implementation of per-detector parallelization for the Jinc mapmaker still yields shorter reduction times compared to per-chunk parallelization, especially for observations with significantly fewer chunks relative to the number of available threads.

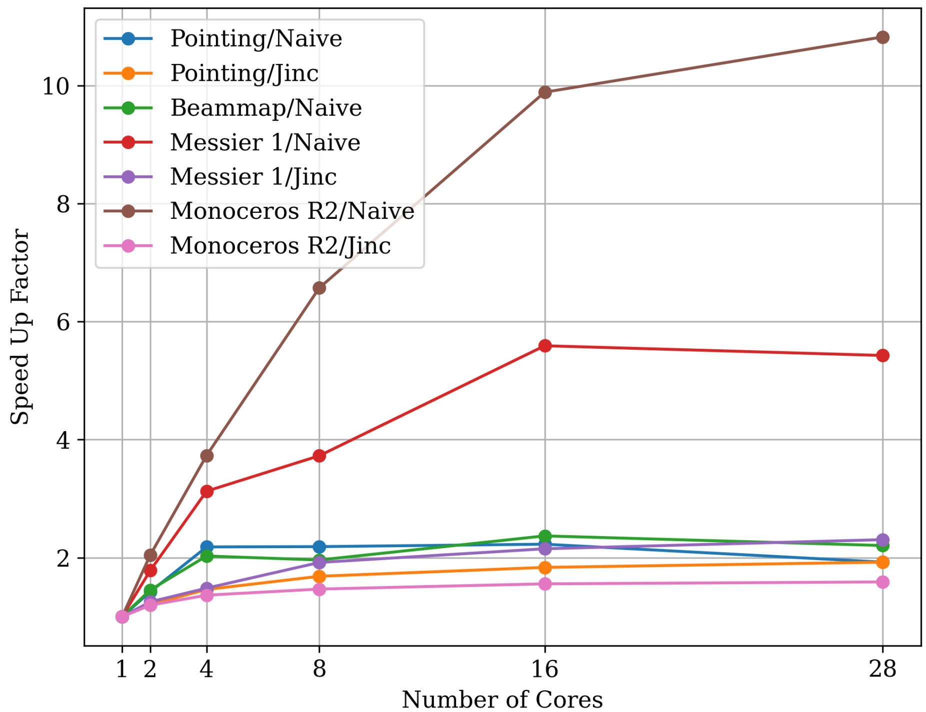

Figure 3.32. Speed up factor for reductions with `Citlali` of different observation types and mapmakers using different numbers of compute cores. Results are for a single mapmaking iteration.

The reduction time for the iterative mapmaker scales from the single-iteration reduction times in an approximately linear fashion with the number of iterations, although some initial setup processes are not repeated.

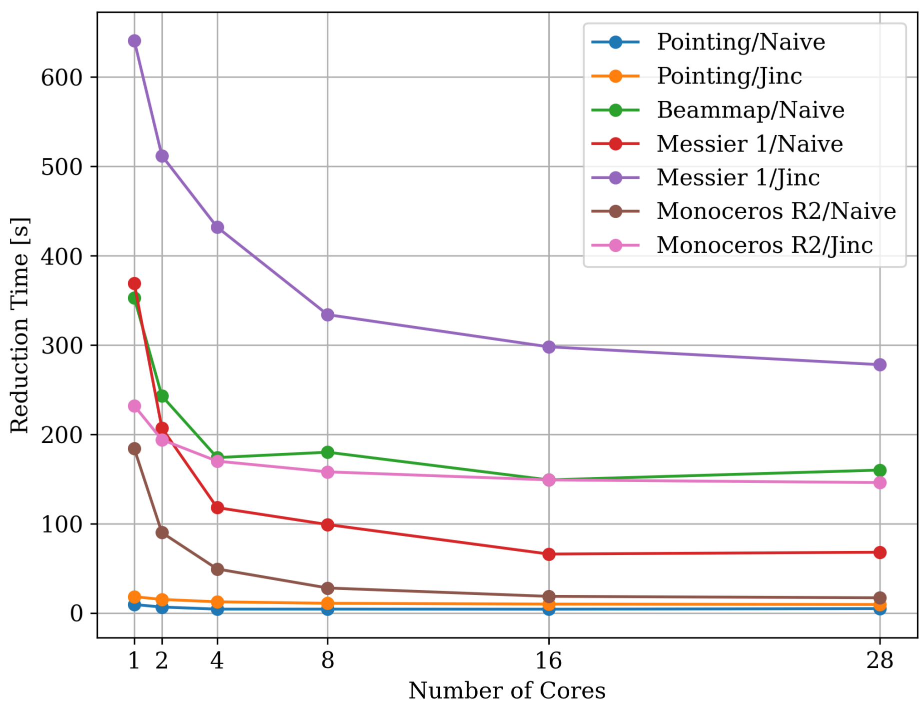

Figure 3.33. Similar to Figure 3.32, but plotting the actual reduction time as a function of number of cores used. Results are for a single mapmaking iteration.

CHAPTER 4

TOLTEC COMMISSIONING MAPS

This chapter presents results obtained from `Citlali` data reductions and subsequent analyses of TolTEC commissioning observation data acquired during the instrument's commissioning runs. I focus primarily on maps of two extended Galactic sources and the beammap used between them. All maps and derived quantities shown here are preliminary and are still undergoing scientific verification. They do not represent final data products and should not be used for scientific analyses.

As of the time of this writing, two full commissioning runs have been completed with TolTEC, consisting of an initial run of 8 nights between June 16, 2022 and July 7, 2022 and a second phase of 13 nights between December 9, 2022 and December 23, 2022. An additional 4 observing nights were carried out in the spring of 2023. No observations were performed for the remainder of 2023 due to poor weather conditions, a wildfire at the base of Sierra Negra, and an extended 6-month power outage to the telescope as a result of electrical storm damage. Following the restoration of power, commissioning observations resumed in late March 2024.

As it was the first time observing the sky with TolTEC on the LMT, the June-July 2022 commissioning run focused primarily on instrument setup and characterization as well as the integration of the data reduction and visualization software packages into the telescope's computing infrastructure. These tests included procedures to confirm the alignment of the various components of the warm and cold optics chain, verifying the resonator identification algorithms given the background loading conditions in the LMT receiver cabin, and exercising the beammapping, focus, astigmatism, and pointing correction pipeline reduction modes. The summer months are characterized by much higher humidity and increased precipitation than the winter months and are suboptimal for observations below 3 mm. For this reason, combined with the intention to improve the noise characteristics of the arrays, only bright point sources were observed in this run.

Following extensive software upgrades and hardware modifications informed by the findings of the June-July commissioning run, the December 2022 observing run drew targets from a more diverse list of sources, including extended objects and point source fields. A list of the observed non-point source objects is given in table 4.1 which were chosen primarily based on their availability at different times of the night and the likelihood of detection. Of the 5 observed, 2 were detected, with these being the supernovae remnant Messier 1 and the Monoceros R2 GMC. Both sources are very bright at millimeter wavelengths, making detections relatively trivial even under poor observing conditions and noise environments. The face-on spiral galaxy Messier 74 and the galaxy cluster MACS J0717.5+3745 were not detected due to detector network settings that caused higher-than-expected timestream noise. The non-detection of sources in the COSMOS field was the result of a coordinate error.

In addition to extended emission, commissioning tests with the HWP installed were also undertaken during the December run. Polarization calibrator sources 3c147 and 3c286 were observed in three different modes: without the HWP, with the HWP installed but not spinning, and with a spinning HWP.

Analysis and scientific verification of timestream data and maps from the December commissioning run was carried out in 2023. Primary goals were to explore the intrinsic data quality to ascertain if further hardware adjustments were necessary, determine the optimal data reduction parameters, and to compare source flux recovery and calibration in maps produced by `Citlali`'s internal mapmakers and the Minkasi maximum likelihood mapmaker for the different source categories.

Table 4.1. Commissioning Observation Parameters

Source	RA	Dec	Num Obs	Detected?
Messier 1	5h34m31.95s	22d00m52.15s	7	yes
Monoceros R2	6h07m46.3s	6d23m09s	5	yes
Messier 74	1h36m41.7s	15d47m01.1s	2	no
COSMOS	10h00m30s	2d12m20.00016s	4	no
MACS J0717.5+3745	7h17m30s	37d45m00s	2	no

4.1 Beammaps

The beammapping observing strategy is outlined in Section 2.5.3 and the data reduction approach is described in Section 3.12.1.3. Beammapping observations were taken of the Solar System planets Saturn, Neptune, and Uranus, as well as Submillimeter Array Calibrator List targets such as 3c279, 3c84, BL LAC, and OJ287. No other Solar System bodies such as asteroids or dwarf planets were observed due to lack of observability during the commissioning runs and because of constraints from the observing and instrument testing schedule.

The quality of the beammaps varied considerably between commissioning nights due to ongoing instrument and optics adjustments. During regular operations, at least two beammaps will be acquired each night. For commissioning, however, an effort was made to obtain the best possible Array Property Table and use it across different observing nights and for the reductions of the extended sources. This was done to streamline observations and instrument testing. A beammap taken on December 19, 2022 (MJD=59932) of the source J1159+292 (Ton 599 or 4C+29.45) was selected owing to its proximity to the observations of the COSMOS field and the quality of the raw data and weather conditions. The properties of J1159+292 are given in Table 4.2. The source is a flat spectrum radio quasar and is highly variable (Rajput and Pandey 2021) in both total intensity and polarization fraction and is frequently monitored by the SMA at 870 µm and 1.1-1.4 millimeters. It undergoes frequent flaring on short timescales with flux variations of 200-300% (Figure 4.1), making absolute flux calibration difficult. The flux in each of the TolTEC bands was estimated by measuring the spectral index from the most recent SMA observations at 1.1 mm and 850 µm and extrapolating to longer wavelengths. Uncertainties are expected to be at the 15% level.

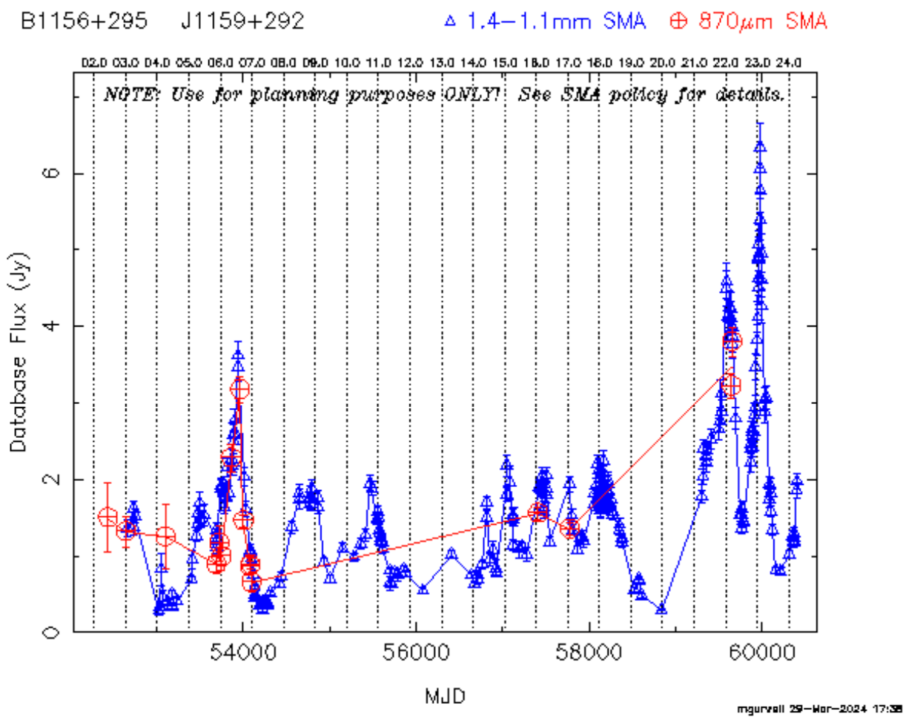

Figure 4.1. Plot of the flux at 870 µm and 1.1 mm made by the SMA over time of the radio galaxy J1159+292 observed in beammap observation 102518. Figure reproduced from: http://sma1.sma.hawaii.edu/callist/callist.html?plot=1159%2B292

Figure 4.2 shows the on-sky detector offsets from the APT produced by the beammapping reduction after flagging and rotation to the horizon. Timestream reduction steps include lowpassing at 20 Hz, decimation to 40 Hz, and atmospheric cleaning with 5 eigenvalues. The naive mapmaker was used to generate the maps used in the final fitting. The beammapping reduction was run for a total of 10 iterations. The network alignments are rotated with respect to the design APT (Figure 2.4) owing to the installation angles of the arrays. Networks 6 and 10 were disabled for this observation. Detector flagging was performed during reduction with limits on the beam centroids, FWHMs, S/N, and sensitivity, followed by a small level of manual flagging. The statistics for the detectors across each network are given in Table 4.2. The median percentage of flagged detectors for 102518 for all networks is 42.3% and is primarily due to the now-corrected software bug that resulted in nearby resonators being merged.

Table 4.2. J1159+292 Parameters

Parameter	**Value**
Observation ID	102518
RA	11h59m31.8339s
Dec	+29d14m43.826s
z[1]	0.725
F_{273GHZ}	3.426 Jy/beam
F_{214GHZ}	3.475 Jy/beam
F_{150GHZ}	3.676 Jy/beam

[1] Rajput and Pandey 2021

Maps of the beams derived from coadding all detectors are plotted in Figure 4.4. Sidelobe feature are visible across all 3 arrays and have similar structure within each indicating that they are the result of residual deformations in the primary mirror surface as opposed to being introduced by the reduction proceedure or detector array characteristics. Two important results from the beammapping analysis include the confirmation that the arrays have the same footprint on the sky and that they are co-aligned. The array extents closely match the predicted FOV of $220''$ given the detector spacing and array dimensions. This was previously confirmed during in-lab testing following a fix to the dichroic filters in the cold optics (Wilson et al. 2020).

The distribution of the flux calibration factors across the arrays and histograms of the fitted beam FWHMs in both Azimuth and Elevation are illustrated in Figure 4.3. The median log(FCF) values for each array are 9.04 ± 0.08, 8.53 ± 0.11, and 8.41 ± 0.09 in units of mJy/beam/x with x being the KIDs scattering parameter. Variations of the FCF across different networks are visible due to varying readout settings and detector quality factors. There is no significant dependence in either axis or with

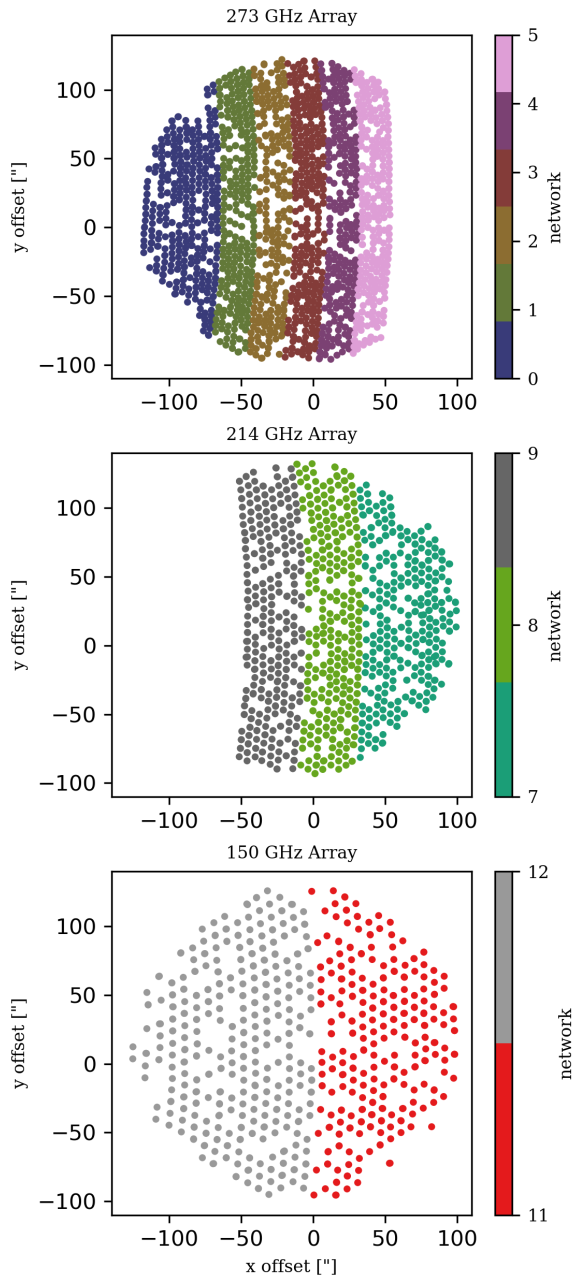

Figure 4.2. Fitted azimuth and elevation offsets after rotation to the horizon from the APT of J1159+292 for the 3 TolTEC arrays. The colors represents the detector readout networks. Networks 6 and 10 were disabled. Detectors flagged by `Citlali` during the beammapping reduction are not plotted.

distance from the array center which would be indicative of optics misalignment. An asymmetry is observed, however, in the beam widths with a slight broadening in the scan direction. The average FWHMs and fraction of unflagged detectors are listed in Table 4.1. A narrowing of the beam in the scan direction can occur from the atmospheric filtering stage of the data reduction pipeline, but this is mitigated by the iterative beammapping procedure. Therefore, this effect is likely due to a residual focusing offset. The median beamwidths for each array are $6.07 \pm 0.37''$ $7.51 \pm 0.36''$, and $10.74 \pm 0.31''$ which are $1''$ larger than the theoretical instrument beam FWHMs. Example cutouts of random detector maps from each array are plotted in Figure 4.5 and show a range of beam morphologies. Each has been smoothed with a Gaussian beam with $\sigma = 3''$.

Table 4.3. Beammap Network Statistics

NW	**Found**	**Good**	**Bad**	**Good %**	**Az FWHM** [arcsec]	**El FWHM** [arcsec]
0	682	391	291	57.33	6.56	5.69
1	518	328	190	63.32	6.42	5.73
2	557	256	301	45.96	6.40	5.38
3	565	328	237	58.05	6.53	5.53
4	568	285	283	50.18	6.82	5.57
5	515	303	212	58.83	6.62	5.53
6	N/A	N/A	N/A	N/A	N/A	N/A
7	579	329	250	56.82	8.15	7.04
8	489	352	137	71.98	7.96	6.89
9	573	341	232	59.51	8.02	7.00
10	N/A	N/A	N/A	N/A	N/A	N/A
11	524	276	250	52.67	11.27	9.82
12	634	366	268	57.73	11.30	10.39
total	6204	3555	2651	57.30	N/A	N/A

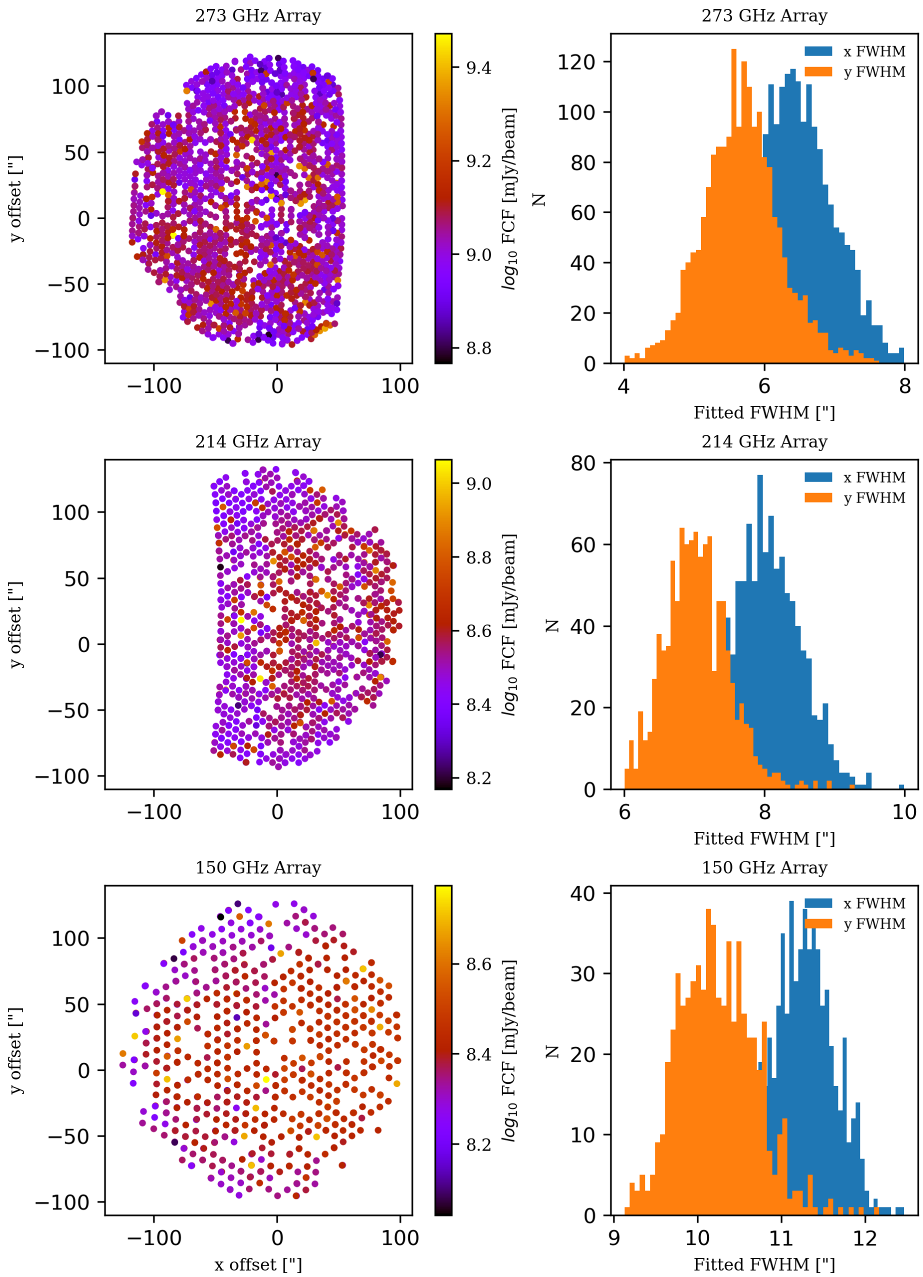

Figure 4.3. Left column: Fitted detector offsets as in Figure 4.2, but with the flux calibration factor plotted as color. Right column: Fitted azimuth (x) and elevation (y) beam FWHMs for each unflagged detector.

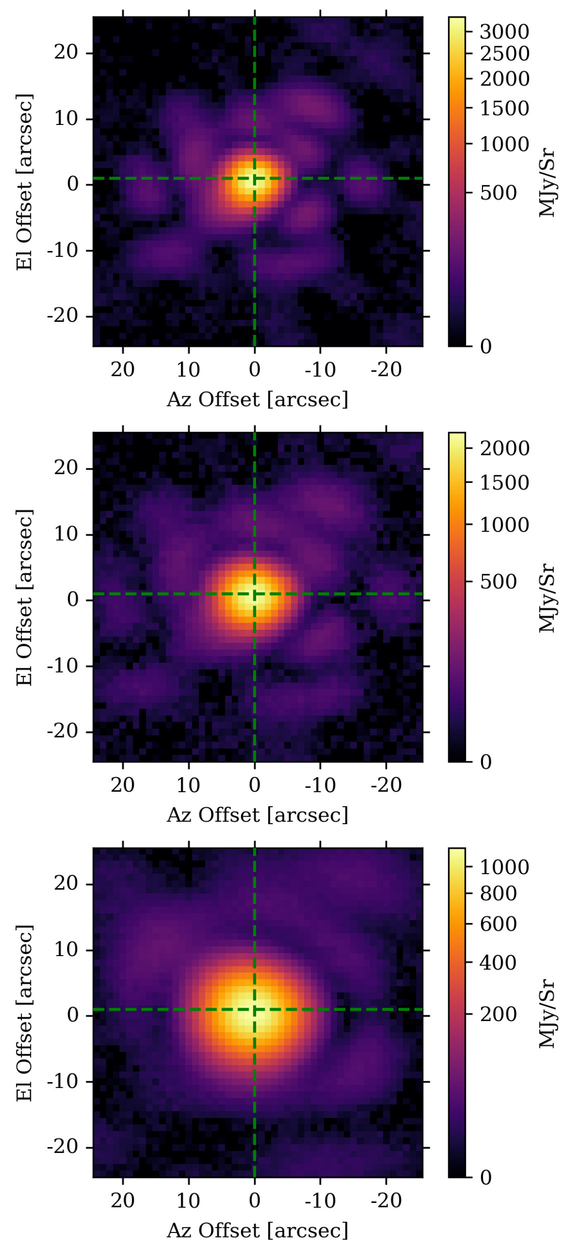

Figure 4.4. Beam profiles for each array made from coadding all detector maps using the `citlali::pointing` reduction engine, using 10 iterations of the naive mapmaker. The color scale is nonlinear due to the large dynamic range between the flux of the peak and sidelobes.

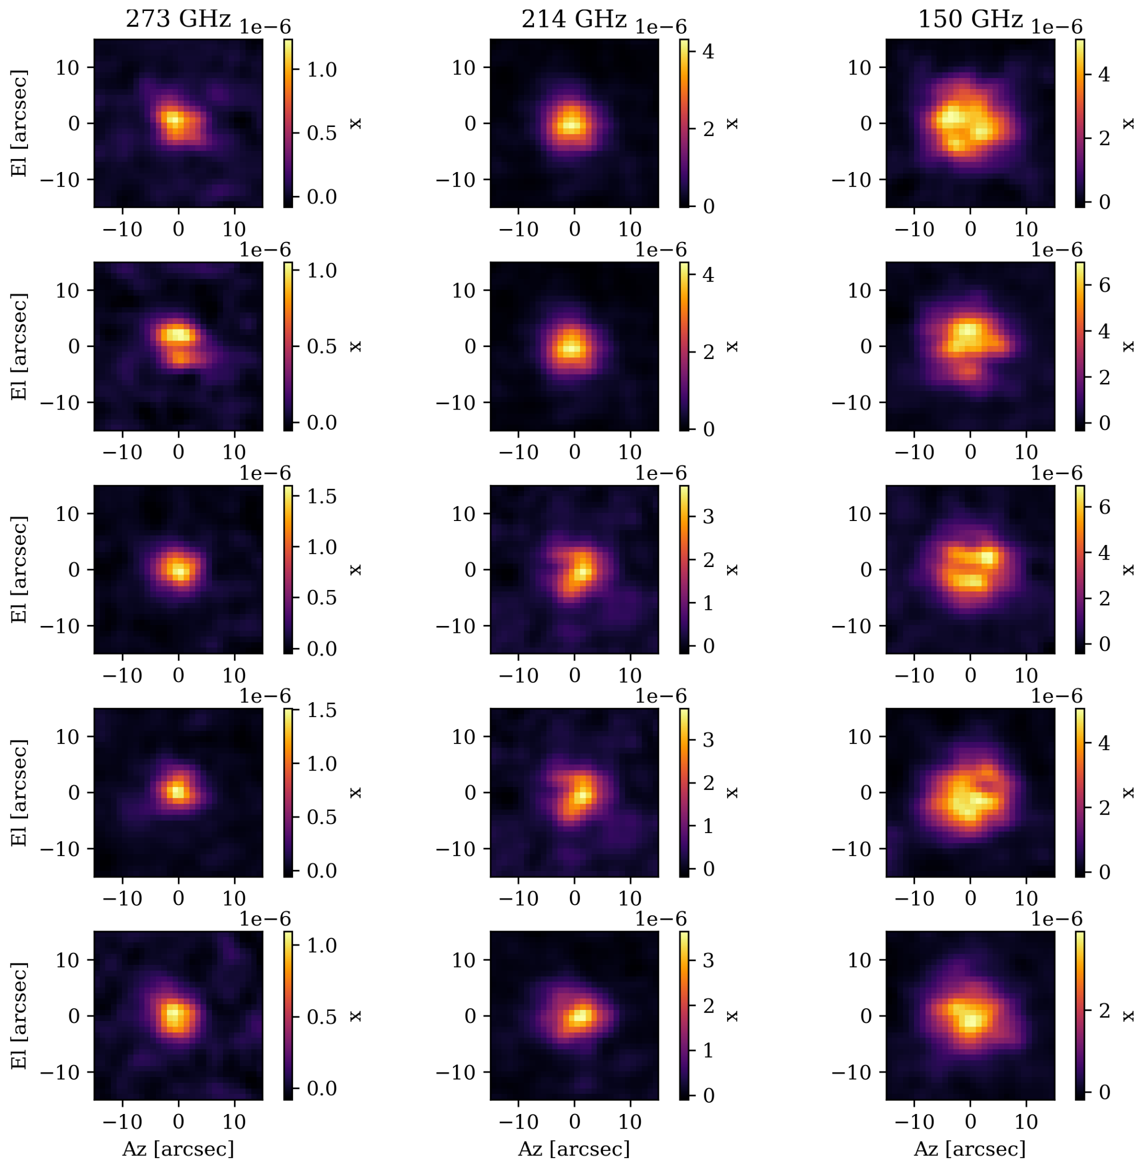

Figure 4.5. Beam profiles from random detectors selected from each TolTEC array (columns). Map units are the raw timestream KID detuning parameter x. Maps have been smoothed with a Gaussian with $\sigma = 3''$.

4.2 Messier 1

The Crab Nebula (Messier 1 or Tau A) is an extremely bright (∼200 Jy within TolTEC's bands) supernovae remnant that is dominated by nonthermal emission in the form of synchrotron radiation (see Bühler and Blandford 2014 for a review). It has been extensively observed from gamma ray to radio wavelengths and is ideal for commissioning observations of millimeter wavelengths owing to its well-constrained spectral index, high polarization fraction, and nearly constant polarization angle across 23-353 GHz. Furthermore, it has an angular scale of $5' \times 7'$ making it an optimal target for testing the flux recovery of data reduction and mapmaking pipelines for extended emission. The nebula has been recently observed in total intensity and Stokes Q and U at 150 GHz with both NIKA (Ritacco et al. 2018) and at 260 GHz with NIKA2 (Ritacco et al. 2022).

A total of 7 observations each centered at RA=5h34m31.95s and Dec=22d00m52.15s were taken of the Crab Nebula with the parameters of each being listed in Table 4.2. While measurements of the polarized signal were intended, the HWP was not installed for these observations due to time and observing constraints. The first 6 observations were taken on December 16, 2022, and consist of a set of Raster maps with short integration times of 2 minutes with every second map being rotated by 45 degrees to enhance cross-linking between scans in the final coadded maps. The Raster map dimensions are selected to be much larger than the size of the source to allow for a better estimate of the background loading and to improve the removal of the atmosphere in the PCA algorithm. The average τ_{225GHz} across all 6 maps was 0.048. The final observation is a 23-minute Lissajous map over a smaller field around the source taken on December 19, 2022. The measured opacity was slightly higher than for the Raster maps at 0.053 at 225 GHz. Analysis was focused on the Raster maps for data reduction and analysis in this work due to the larger field.

Figure 4.6 shows total intensity maps of the Crab Nebula in each of TolTEC's 3 bands made from the coaddition of the 6 Raster maps. The data was reduced with `Citlali`'s science reduction engine and using the pipeline's Jinc filtered iterative mapmaker and $1''$ pixels. The APT from the beammap of J1159+292 was used for the final map analysis. Pointing offsets were determined from a pointing observation (LMT Observation ID=102284) of J0510+180 carried out immediately prior

Table 4.4. Messier 1 Observation Parameters

Observation ID	Mapping Pattern	Map Size [arcmin]	Angle [deg]	Integration Time [min]
102285	Raster	17 x 15	0	2.2
102286	Raster	17 x 15	0	2.2
102287	Raster	17 x 15	45	2.2
102288	Raster	17 x 15	45	2.2
102289	Raster	17 x 15	90	2.2
102290	Raster	17 x 15	90	2.2
102487	Lissajous	10 x 10	N/A	23

to the Raster maps with the measured offsets being (Az,El) = (-4.62", 3.61"). The timestreams were lowpassed at 50 Hz and decimated by a factor of 3, followed by atmospheric cleaning with the largest 10 eigenvalues being removed. Detectors were weighted by the APT detector sensitivities instead of the timestream variances due to the high source flux. Detectors with variances less than a factor of 3 times the median were downweighed to the median value. The mapmaker was run for 150 iterations with a S/N cut of 2. The resulting maps were tested for convergence with an average difference in flux of $< 1\%$ between 100 iterations and 150 iterations being measured. No post-mapmaking Wiener filtering was performed as the iterative mapmaker serves the role of correcting for flux lost during atmospheric filtering.

For each array, 10 jackknifed noise map realizations were also made using the Jinc filtered mapmaker and give 7.8 mJy/beam, 8.23 mJy/beam, and 8.04 mJy/beam for the average map RMS at 273, 214, and 150 GHz respectively. The reduction was repeated with the same reduction parameters except for replacing the Jinc mapmaker with the naive mapmaker, which gives map RMS values of 23.97, 29.67, and 20.52

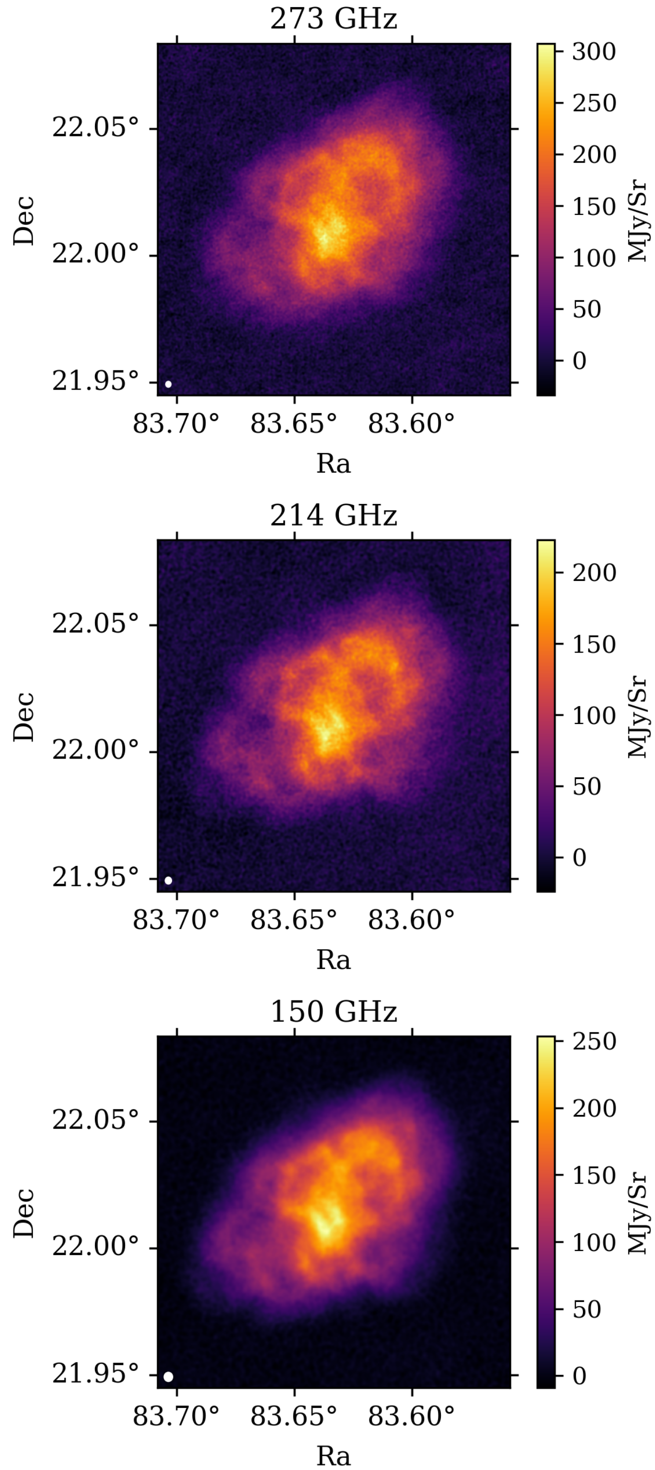

Figure 4.6. Maps of the Crab Nebula made with 150 iterations of `Citlali`'s iterative Jinc mapmaker.

mJy/beam. The difference is attributable to the removal of high-frequency pixel-pixel noise by the Jinc mapmaker.

Preliminary aperture photometry was performed on the `Citlali` iterative Jinc maps to compare flux measurements from existing literature sources. The analysis was carried out by Artyom Tanashkin with additional support from Yuri Shibanov and Aida Kirichenko. Figure 4.7 shows the different-sized apertures used for estimating the flux, with the background remaining nearly zero for all aperture radii considered. The SED of the Crab Nebula between 10-10^6 GHz was obtained from Arendt et al. 2011 and is shown in Figure 4.8 with the TolTEC fluxes overplotted. The 214 GHz and 150 GHz bands are in good agreement with the measured SED while the 273 GHz is systematically higher. The deviation may result from both calibration uncertainties and suboptimal reduction parameter settings. The 273 GHz array includes the largest contribution from the foreground atmosphere of the 3 bands and may not be entirely subtracted. Flux calibration uncertainties carry over from the beammap source dominating the measurement errors at the 15% level.

The recovery of the source flux in the Crab Nebula maps is significantly affected by the atmospheric filtering stages in `Citlali` due to its bright and extended nature. The right column of Figure 4.9 shows the maps produced by the initial iteration of the `Citlali` Jinc filtered mapmaker, highlighting significant flux lost and regions of negative flux introduced by the PCA algorithm. To test how effectively the iterative mapmaker recovers the source flux, the maps were also reduced with the Minkasi maximum likelihood mapmaker to compare results from the radically different mapmaking approaches. This analysis was carried out by Joseph E. Golec. The timestreams from `Citlali` were output after only having the KIDs model solving, flux calibration, and extinction correction applied to them before being passed into Minkasi. The Preconditioned Conjugate Gradient solver was run for 50 iterations. After the PCG solution was reached, the resulting maps were subtracted from the timestreams to generate

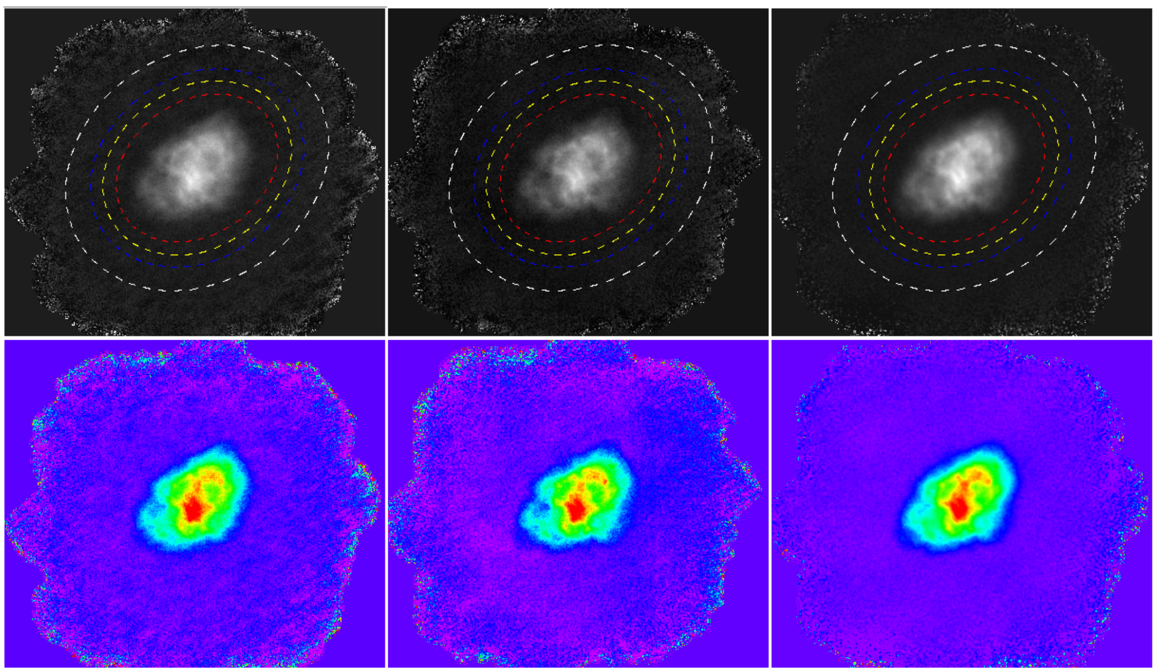

Figure 4.7. Top row: The same maps as in Figure 4.6, but with the apertures used to calculate the integrated flux (red) and background overplotted. Multiple apertures were tested to estimate the background, but each gave similar results. Bottom Row: Maps plotted with a different colormap to highlight the flat background. Figure and analysis carried out by Artyom Tanashkin.

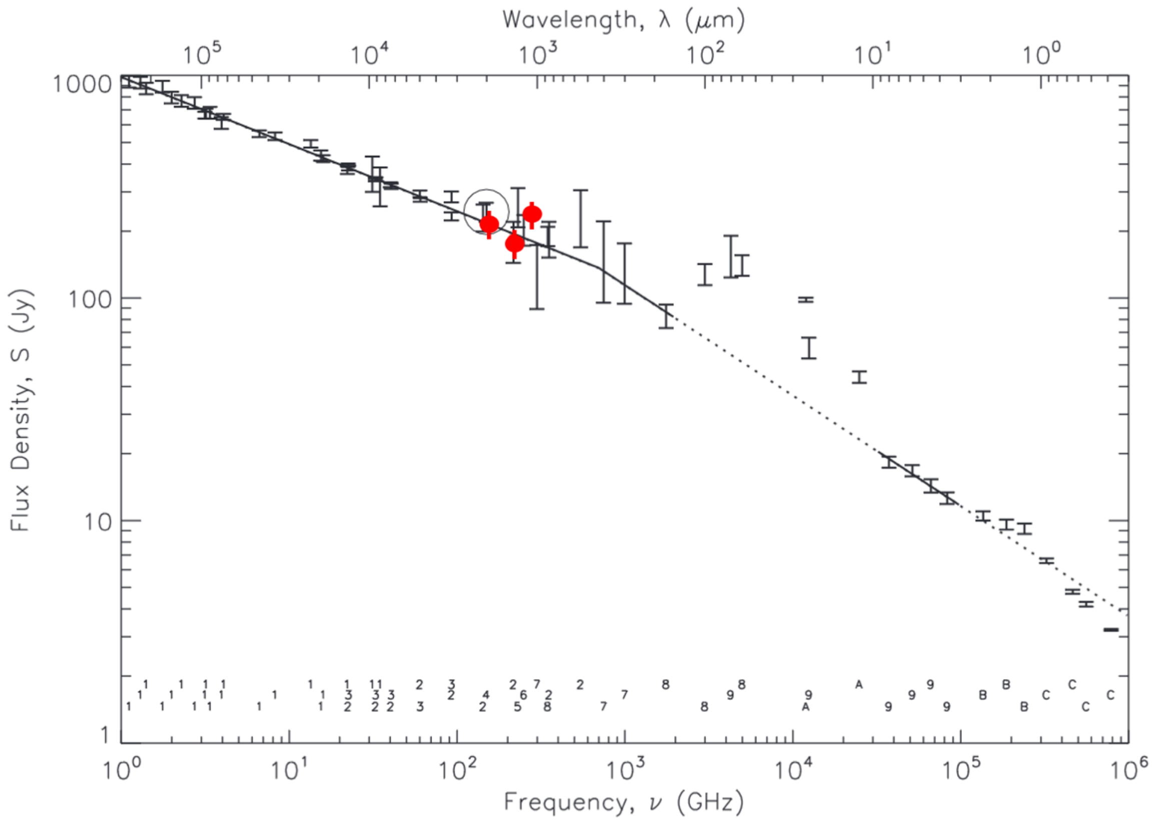

Figure 4.8. Figure reproduced from Arendt et al. 2011. Crab Nebula SED with the integrated fluxes from the `Citlali` iterative Jinc maps overplotted (red points). Error bars represent the 15 % calibration uncertainties. Figure editing and analysis carried out by Artyom Tanashkin.

a new noise estimate, and the mapmaker was re-run. This process was repeated 30 times, with convergence being reached after around 10-15 iterations. The Minkasi maps from the final iteration are shown in the left column of Figure 4.9. Overall, Minkasi outperforms earlier iterations of both `Citlali`'s naive and Jinc mapmaker but underpredicts the flux relative to the final converged `Citlali` maps. There are also extended regions of negative pixels in the Minkasi maps that surround the source, with the 150 GHz band being the most affected. Figure 4.10 plots the flux difference between the 150 iteration `Citlali` Jinc map and the Minkasi maps for each pixel versus the `Citlali` flux in the 3 maps after reprojecting the latter onto the same WCS using the Python package `reproject`[1]. The 273 GHz and 150 GHz demonstrate similar behavior with the highest flux pixels showing the greatest discrepancy between the two mapmakers. There is little variation of the flux difference with the `Citlali` flux in the 214 GHz band though there remains a positive offset. The distinction between the 214 GHz band and the others likely arises at least in part due to the higher noise properties of that array during observing.

The origin of the negative flux in the Minkasi maps is not yet well understood. As a maximum-likelihood mapmaker, it does not utilize the same atmospheric removal approach as `Citlali` such that a bias should not be introduced through filtering. It is plausible that the bright nature of the source affects the noise estimate or that pixel-pixel noise and timestream and instrument gain variations result in a poor flux model estimate (Næss 2019).

Preliminary Stokes Q and U maps of the Crab Nebula are shown in Figure 4.11. As the HWP was not installed while observing and since the polarized reduction mode of `Citlali` is still under development, each Stokes map is made by differencing the maps generated from the different detector orientations. In this approach, the

[1]https://reproject.readthedocs.io/en/stable/

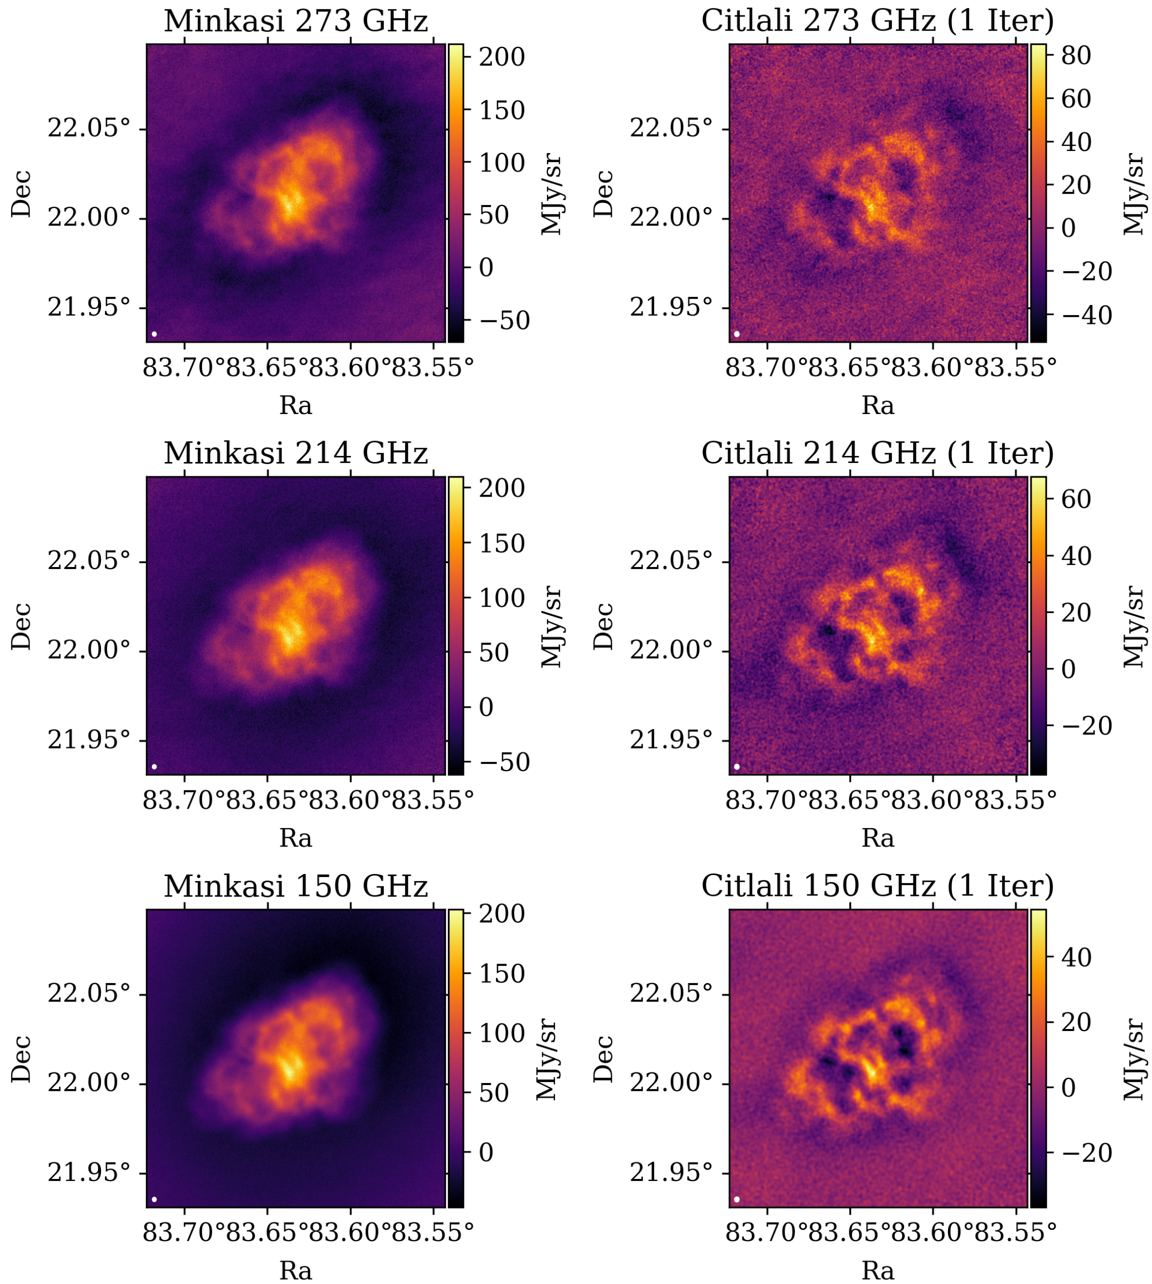

Figure 4.9. Left column: Crab Nebula maps made with the Minkasi maximum likelihood mapmaker. Minkasi maps were generated by Joseph E. Golec. Right column: Maps made from a single iteration of `Citlali`'s Jinc mapmaker.

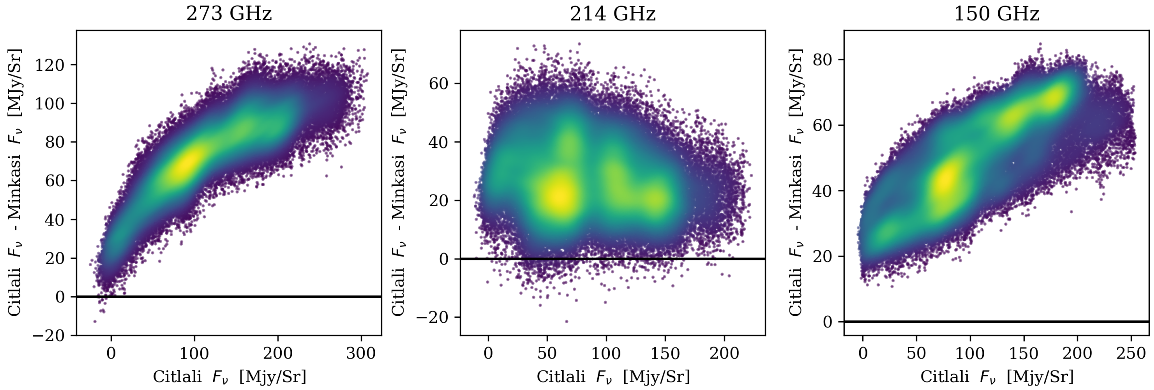

Figure 4.10. Differences between the Crab Nebula fluxes from the `Citlali` 150 iteration Jinc maps and the Minkasi maps plotted against the `Citlali` map's fluxes for each pixel in the maps. The colors represent the density of the points from a Kernel Density Estimate. The black line is the zero-line where the `Citlali` and Minkasi fluxes would be identical.

Stokes Q map is derived from the subtraction of the detectors at 45 degrees (frequency group=1) from those at 135 degrees (frequency group=3) and the Stokes U map is the difference between the 90 degrees (frequency group=2) detectors and those at 0 degrees (frequency group=1). This convention was chosen to match the sign of the emission in the Stokes Q and U maps from observations with NIKA at 150 GHz (Ritacco et al. 2018). The data reduction was performed in the same way as the total intensity maps, using the `Citlali` iterative Jinc mapmaker for the same number of iterations. The resonance frequencies and APT positions were matched to the physical detector locations by hand with 25, 17, and 5 unflagged detectors remaining unmatched at 273, 214, and 150 GHz. Overall, the morphology of the Stokes Q and U map at 150 GHz is in good qualitative agreement with the corresponding NIKA maps. Owing to the first-order nature of the polarization reduction method and the uncertainty regarding polarization angle systematics, I do not compare the polarization fraction or orientation.

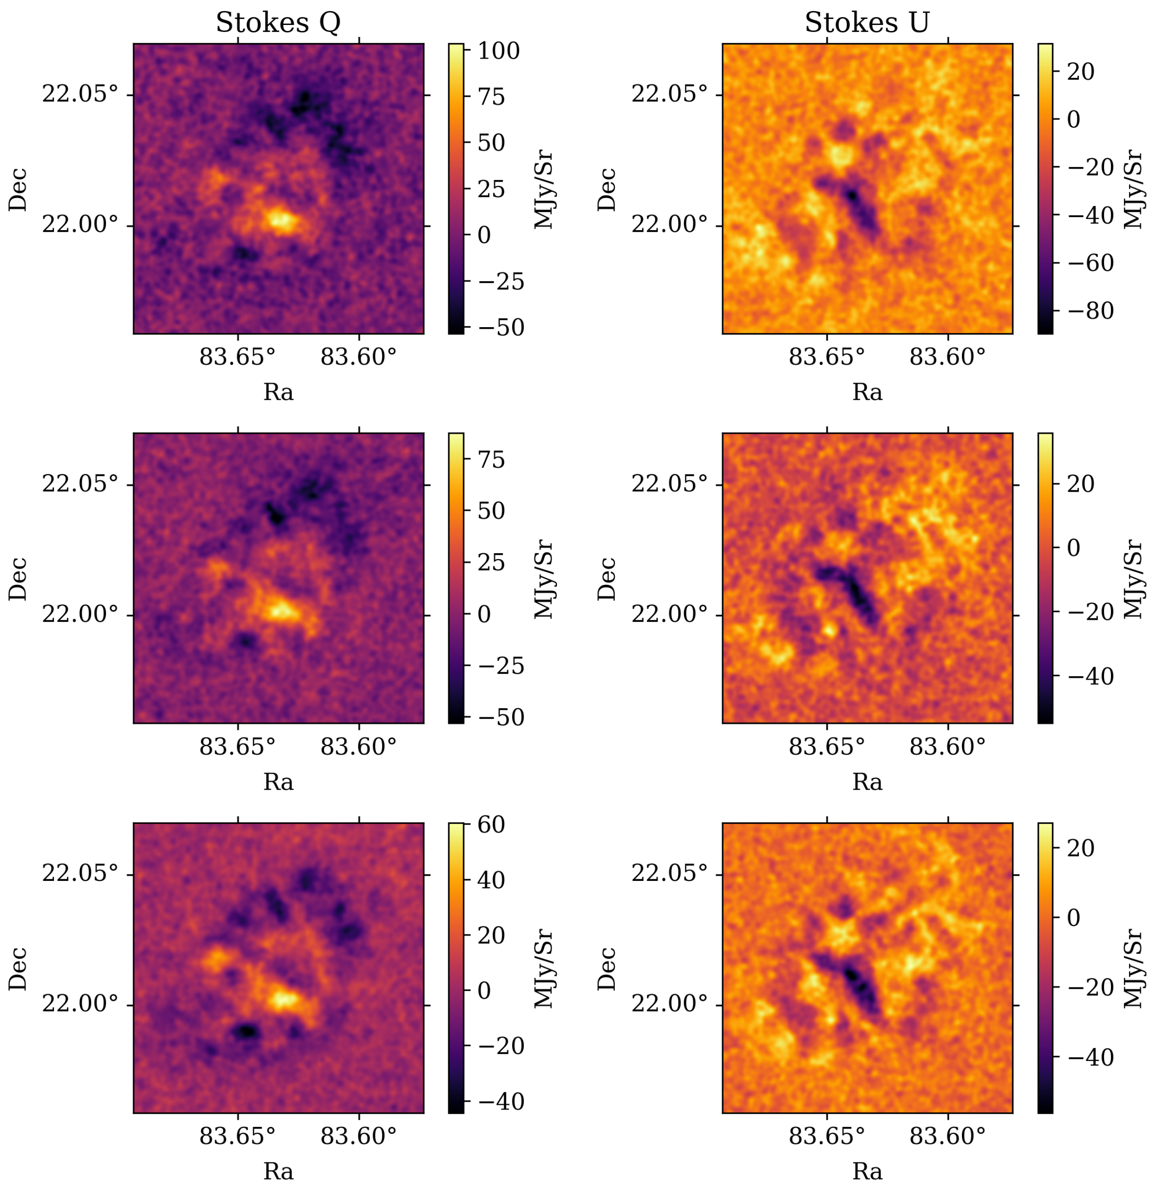

Figure 4.11. Preliminary Stokes Q (left column) and U (right column) maps made with `Citlali`'s iterative Jinc mapmaker by differencing detectors from the different frequency groups. The HWP was not installed for these observations.

4.3 Monoceros R2

The Monoceros R2 (hereafter MonR2) Giant Molecular Cloud is a large, nearby star-forming region at a distance of 893 ± 42 parsecs (Dzib et al. 2016; Jiang and Hillenbrand 2024). The entire cloud structure is extended over a field of 3 x 6 degrees and is characterized by bright reflection nebulae connected through filamentary structures (Pokhrel et al. 2016). Unlike the Crab Nebula, dust reprocessing and nonthermal free-free emission are the primary source of photons from MonR2. The region is similarly useful for testing the recovery of extended emission at scales near or larger than the instrument FOV, while also offering small-scale structures that are advantageous for confirming pointing alignment. MonR2 was observed at 273 GHz with AzTEC when it was mounted on the LMT when the telescope was in its 32 m configuration in 2014 and 2015 (Sokol et al. 2019).

MonR2 was observed in 5 separate observations with TolTEC which were taken in succession to one another on December 21. 2022. A summary of the observation properties is given in Table 4.3. They consist of two Lissajous maps, two Raster maps, and a single Rastajous pattern. The instrument sampling rate was varied between the observations due to the range in mapping speeds used. The Lissajous maps observed a smaller field around the brightest central region of MonR2, whereas the Raster maps covered a 1 square-degree field. The Rastajous map was designed to be in between the two sizes. Of the 5 observations, the first Lissajous map had the highest data quality and was used for subsequent analyses. The mapping speed of the Raster maps was erroneously set too high resulting in the maps being under sampled thus filtering out much of the extended structure. The second Lissajous and Rastajous maps show signs of being out-of-focus with the sources being stretched in the scan direction. Two pointing observations (102580 and 102586) were acquired before the first Lissajous and after the final Rastajous observation. The beam in the second pointing observation is very extended in Elevation with an axial ratio greater than approximately 1.4. The mean pointing offset from both observations was (Az,El) = (1.02″, 3.82″). The optical depth at 225 GHz for all observations varied between 0.058-0.06.

Table 4.5. MonR2 Observation Parameters

Obs ID	**RA**	**Dec**	**Map Pattern**	**Map Size** [arcmin]	**Map Speed** [arcsec/s]	**Integration Time** [min]
102581	6h07m46s	6d23m09s	Lissajous	8 x 8	50	10.0
102582	6h08m00s	6d20m00s	Raster	60 x 60	450	10.8
102583	6h08m00s	6d20m00s	Raster	60 x 60	450	10.8
102584	6h07m46s	6d23m09s	Lissajous	8 x 8	50	10.0
102585	6h07m60s	8d34m50s	Rastajous	15 x 15	50	8.4

Data reduction with `Citlali` for the first Lissajous map was carried out using the same manner as for the Crab Nebula reduction, with the same parameters chosen for lowpassing, decimation, cleaning, and detector weighting being used. The iterative Jinc mapmaker was run for 100 iterations with an S/N cut of 2.0. The APT from beammap of J1159+292 was again used for detector offsets and flux calibration. After comparing with the AzTEC 273 GHz, a residual pointing offset remained after applying the pointing corrections from the nearby pointing observations. An offset of (Az,El) = (-4.52, 4.28) was instead applied to line up the bright central source with the AzTEC map. The final maps are shown in Figure 4.12 which have average RMS values of 4.8, 4.7, and 2.3 mJy/beam. As with the Crab Nebula, the naive mapmaker reduction has higher RMS values at 13.5, 16.9, and 11.6 mJy/beam.

A shift of approximately 4″ in the position of the central bright source can be observed between the three TolTEC bands (Figure 4.13). This effect cannot be due to misalignment between the three arrays as beammapping and point source calibration observations confirm that the optics are aligned at the sub-arcsecond level. This is possibly the result of free-free emission constituting a greater fraction of the source

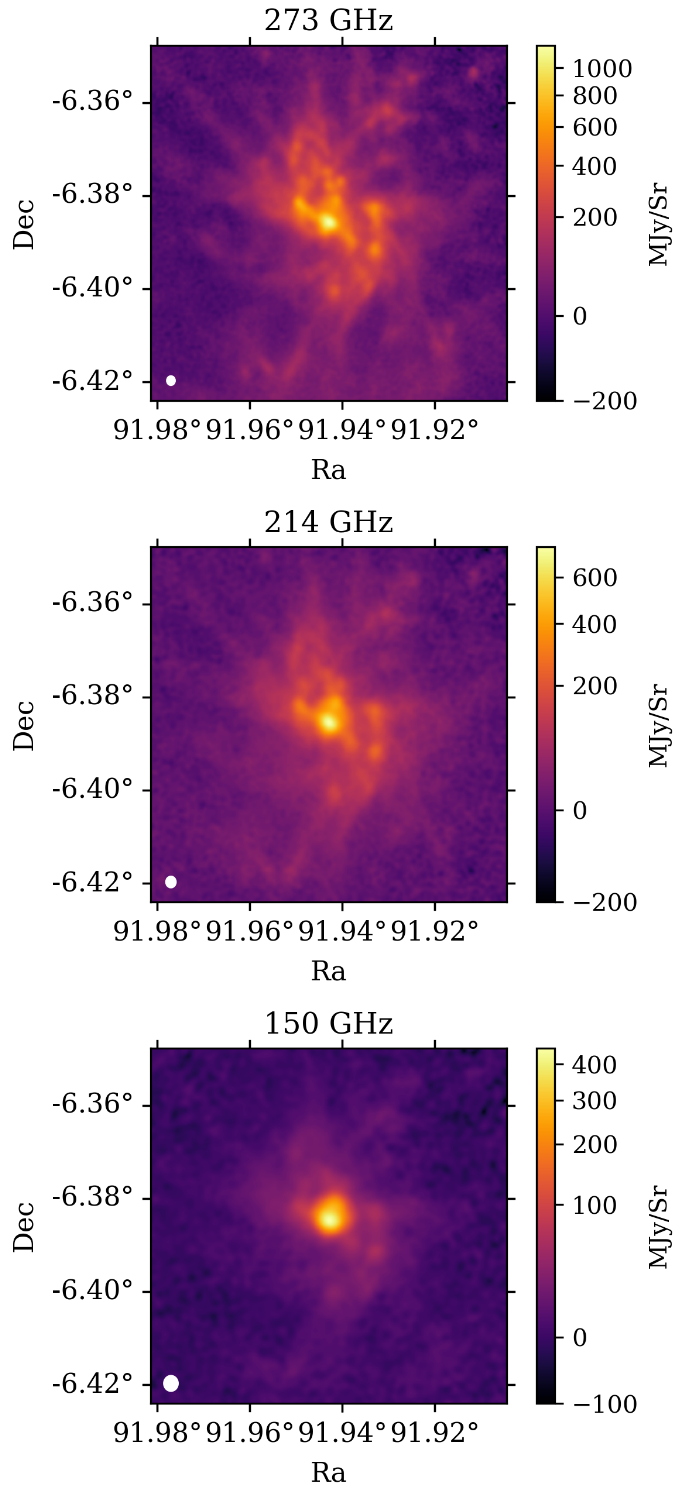

Figure 4.12. Maps of MonR2 made with 100 iterations of `Citlali`'s iterative Jinc mapmaker. The color scale is nonlinear to highlight the fainter, filamentary features in the maps.

flux at 150 GHz than at 273 GHz where dust thermal emission will be the dominant emission mechanism.

I compare the `Citlali` 273 GHz map to the AzTEC 273 GHz map from Sokol et al. 2019 in Figure 4.14. The data reduction for the AzTEC map was similar to that of the `Citlali` maps with PCA also being used for atmospheric filtering. Histograms of the pixel flux distribution are presented in Figure 4.15. Two important distinctions between the reduction, however, are that the AzTEC maps do not use an iterative mapmaker nor are they filtered to recover flux lost from atmospheric cleaning. Instead, a naive mapmaking algorithm was used, and the maps were smoothed with a Gaussian filter from their native resolution of $8.5''$ to $12''$. Consequently, the recovered flux is much less than that of the `Citlali` map, and regions of negative flux are present. The TolTEC map recovers a peak flux a factor of 1.62 higher than the AzTEC map. The AzTEC map has a depth of 5.96 mJy/beam after converting to the TolTEC 273 GHz beamsize.

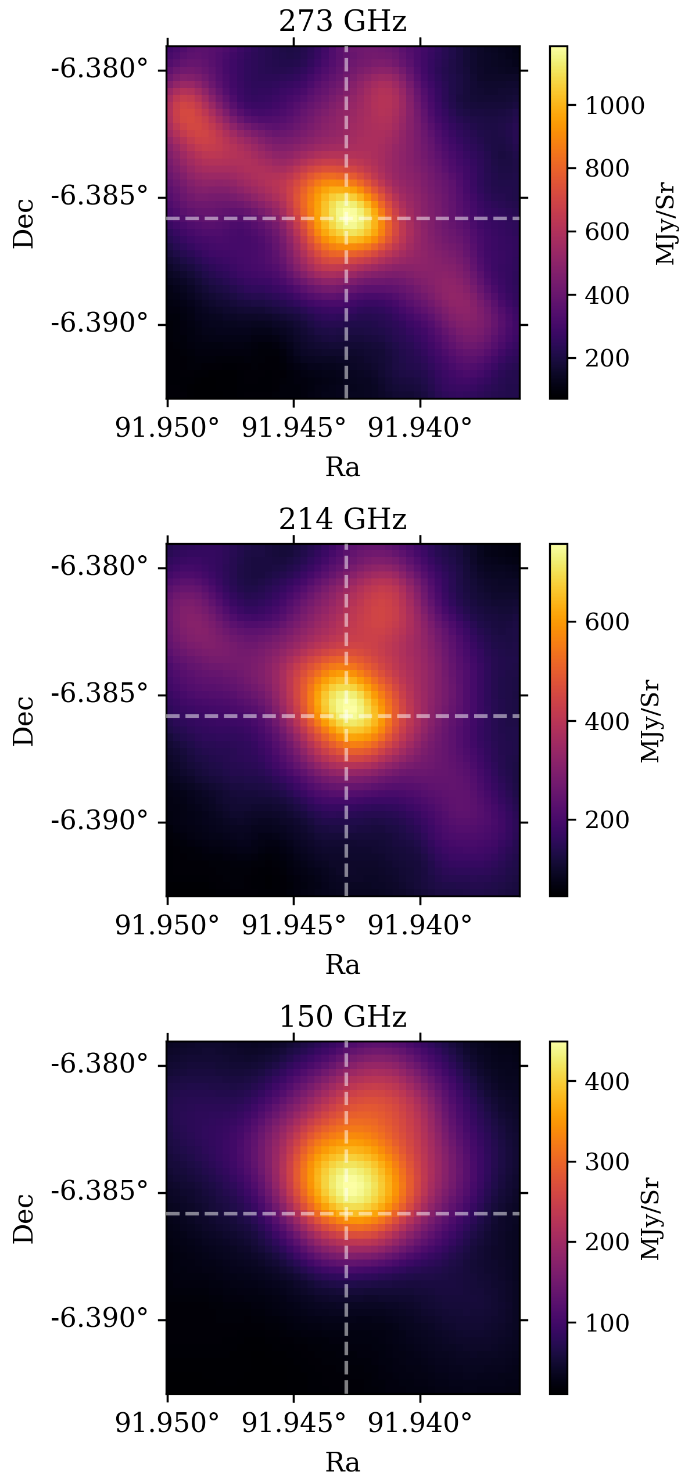

Figure 4.13. Zoom in on the bright central region of the MonR2 maps from Figure 4.12. The white dashed lines are at the peak of the source at 273 GHz.

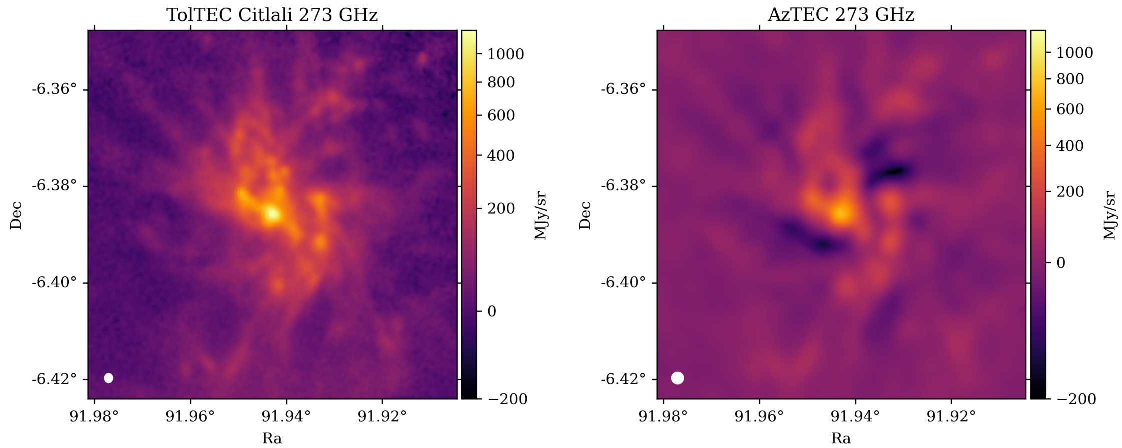

Figure 4.14. Left panel: Same 273 GHz map as in the top panel of Figure 4.12, but plotted on the same color scale as the AzTEC map on the right panel. Right Panel: AzTEC map of MonR2 at 273 GHz from Sokol et al. 2019.

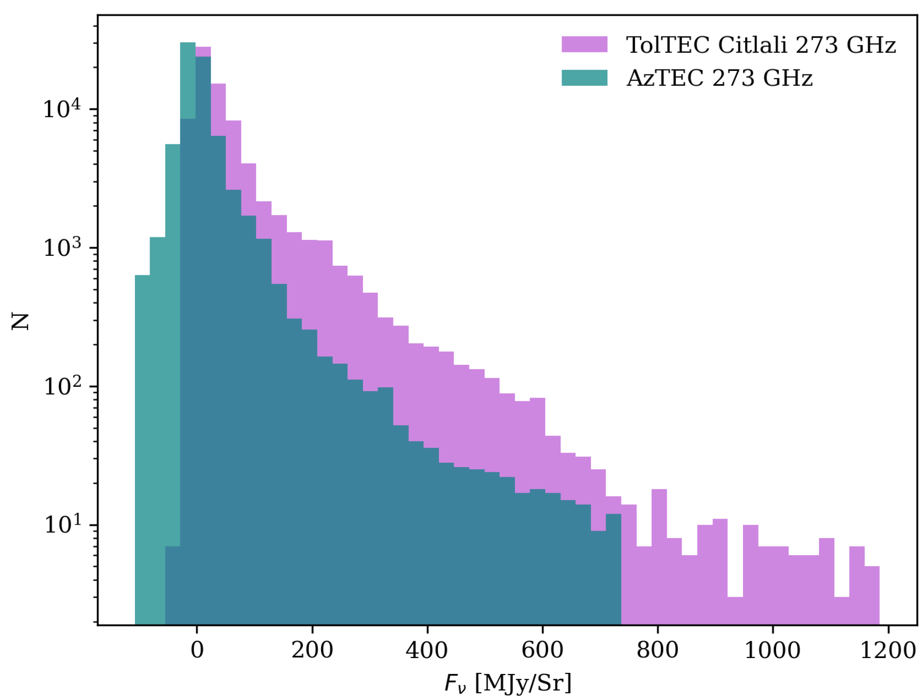

Figure 4.15. Histograms comparing the flux distributions in the `Citlali` 100 iteration Jinc map of MonR2 with that of the AzTEC map plotted in Figure 4.14.

Finally, I also created maps with the Minkasi maximum likelihood mapmaker for comparison. The data reduction steps with Minkasi were identical to those used in the reduction of the Crab Nebula. The Minkasi maps are plotted in Figure 4.16. Similar negative features to the Crab maps are also seen, with regions between the fainter filamentary structures having negative flux. The difference between the flux in the `Citlali` iterative Jinc maps and the Minkasi maps is shown in Figure 4.17. Overall, there is greater agreement compared to the Crab Nebula, though there remains scatter. The Minkasi map recovers more flux from the brightest pixels in the 273 GHz map relative to `Citlali` which is in contrast to the case with the Crab Nebula. The brightest regions are near the size of the beam in MonR2 as opposed to being extended over several arcminutes as in the Crab Nebula and therefore likely have a smaller impact on the nose estimate in the Minkasi mapmaker.

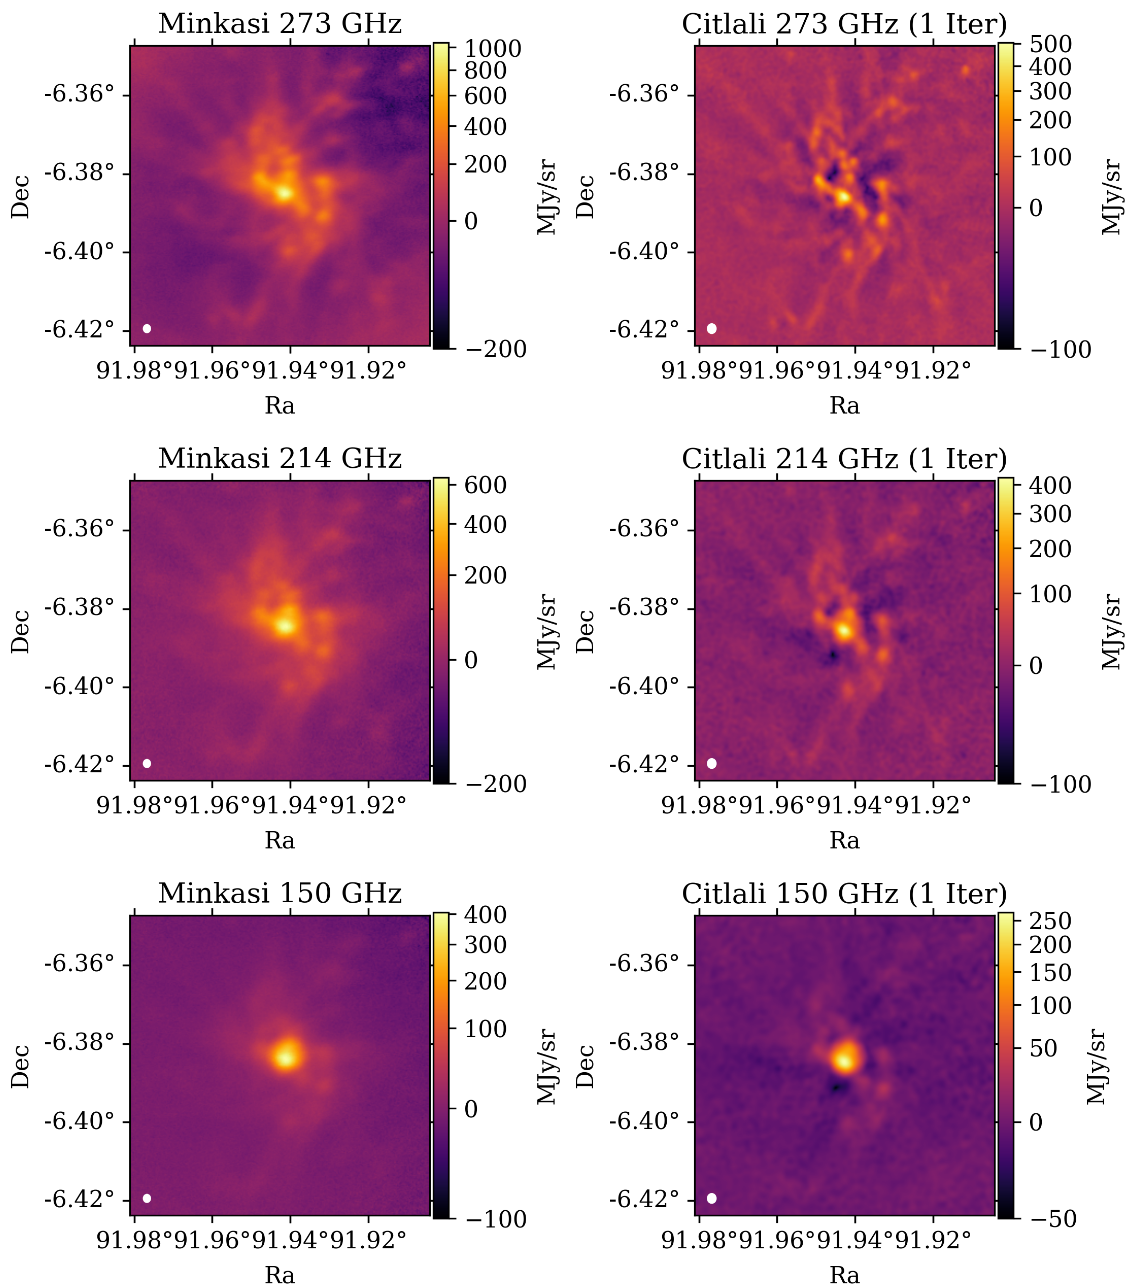

Figure 4.16. Left column: Maps of MonR2 made with the Minkasi mapmaker for the 3 TolTEC bands. Right column: Initial iteration `Citlali` Jinc maps of MonR2.

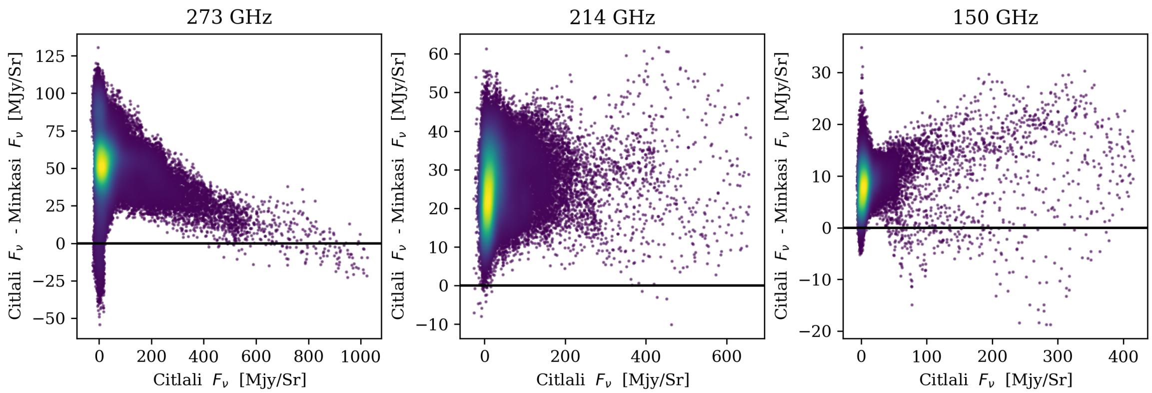

Figure 4.17. Difference between the `Citlali` 100 iteration Jinc map fluxes and the Minkasi fluxes for each pixel plotted against the `Citlali` fluxes. The color represents the density of the points.

CHAPTER 5

HIERARCHICAL BAYESIAN SED MODELING

In the analysis and modeling of dust SEDs, it is critical to consider the range of factors that influence the estimation of key dust parameters before they can be interpreted scientifically. The choice of the dust model for both the dust and other SED components, the statistical framework employed for carrying out the fit, data quality and wavelength coverage, and the inherent nature of the astrophysical source all can contribute to systematic variations and stochastic noise introducing additional uncertainties and biases into the derived dust parameters.

As described in Section 1.0.4, existing studies seeking to fit dust SEDs commonly use either a single or multiple component modified blackbody or graybody models between FIR and millimeter wavelengths and use the physically motivated dust models when NIR and MIR constraints are also available. The graybody models are mathematically straightforward with few parameters and therefore require minimal wavelength constraints for fitting but can produce parameter estimates that are less representative of actual dust grain properties due to their inability to account for the mixed physical conditions within galaxy ISMs. In contrast, the physical dust models leverage the constraints from independent observations and dust grain laboratory analogs to construct a detailed representation of the dust, factoring in grain shapes, compositions, and size distributions, and heat it with a stellar radiation field distribution to generate dust thermal and spectral emission features.

Both models are used in conjunction with the χ^2 minimization, Maximum Likelihood Estimator (MLE), or Bayesian fitting frameworks to estimate parameters from single sources like integrated galaxy flux measurements or pixels, as well as simultaneously across populations of independent observations, namely individual galaxies or all the pixels in image data. With measurements of populations, information regarding the distributions of parameters and the correlations among them can be gleaned. The least-squares and MLE approaches are frequentist in nature and, while more computationally efficient at deriving best-fit values, come with a number of limitations including the inability to easily incorporate existing knowledge of the parameter distributions, assumptions about the form of the likelihood, sample size biases, and requirements that the noise be Gaussian and its variance constant. Comparison between models also presents some challenges. For Bayesian model inference, on the other hand, the prior distributions are an integral component during fitting, and a sampling from the full parameter probability distributions referred to as the posteriors are the primary result. Sampling of the posteriors is performed using Monte Carlo Markov Chain (MCMC) random walks which raise the computational requirements for the Bayesian approach.

Regardless of the fitting method, interpreting the results from MBB and the physical dust models must be handled with care, as not only are there systematic biases between the recovered dust masses and temperatures from fits to the same source with each model, but also due to the fact that many of the dust parameters are correlated in a mathematical and, potentially, in a physical sense. These include the anti-correlations between the dust temperature T and spectral emissivity index β for the MBB model (Sajina et al. 2006) and an analogous anti-correlation between the grain size distribution and mean starlight efficiency $\langle U \rangle$ in the physically motivated dust models (Galliano 2022). The former arises because an increased dust temperature shifts the peak to shorter wavelengths and raises the flux into the Rayleigh-Jeans limit, which can compensate for the shallower slope produced by a smaller β. The grain size and ISRF distribution correlation occurs as a result of the SED at MIR wavelengths being the result of emission from stochastically heated small grains as well as larger, hot grains in thermal equilibrium with the ISM. A larger shift in the grain size distribution that favors small grains or a move towards more intense ISRF values can result in a similar SED signature. Unlike the T-beta degeneracy, the grain size distribution is usually constrained through independent means and is usually fixed while fitting. However, various correlations also exist between the physical model dust parameters M_{dust}, α, U_{-}, U_{Δ}, and q_{pah} (Galliano 2022).

The consequences of these effects are most apparent when attempting to fit for multiple independent observations and explore parameter scaling relationships. The mathematically induced correlations act together with measurement uncertainties and intrinsic source variations to bias parameter estimates and introduce further spurious, unphysical correlations between them. Owing to the assumption of isothermal dust in MBB models, multiple temperature components along the line of sight will produce incorrect estimates of the dust temperature which will then bias β in the opposite direction. Adding random noise onto each measurement scatters the estimated temperatures further and results in a characteristic "banana" shape of the T-β parameter space that does not represent the actual connection between the parameters. Uncertainties at the 5% level have been shown to produce an anti-correlation (Shetty et al. 2009b). Galliano 2018 demonstrated that both random noise and calibration uncertainties correlated across observation bands in simulated SEDs perturb the shapes of parameter covariance distributions for MBBs, BEMBBs, and physical dust models and investigated the extent of the effect with S/N and wavelength coverage variations. A strong negative correlation was observed between the mean ISRF value $\langle U \rangle$ and M_{dust} when fitting with least-squares and single-level Bayesian methods.

The existence of the T-β anti-correlation has long been known in observational data (Keene et al. 1980; Blain et al. 2003; Sajina et al. 2006; Kelly et al. 2012). It has been observed in MBB fits to starless cores (Schnee et al. 2010), the Central Molecular

Zone of the Milky Way (Tang et al. 2021), and cold dense clouds (Paradis et al. 2014; Juvela et al. 2015). The nature of the actual underlying physical relationship remains unclear. Under a simple scenario of enhanced dust grain growth within dense ISM environments due to the increased frequency of dust-dust interactions and collisions, a decrement in the emissivity index would be expected with increasing dust and gas column density. With the dust temperature decreasing in dense environments as a result of dust shielding by the outermost layers, a positive correlation in the T-β space would be favored. Lower β values relative that of the diffuse ISM have been observed in some proto-planetary disks and dense molecular clouds which support this idea. However, other observations, simulations, and laboratory analogs demonstrate that an anti-correlation may have a physical basis. Some amorphous carbon laboratory analogs do indeed exhibit intrinsic anti-correlations (Agladze et al. 1996; Boudet et al. 2005). Ysard et al. 2015 employed the dust models of Jones et al. 2013 in fits of the *Planck* High Frequency Instrument (HFI) survey and found that the inclusion of dust mantles onto grains can produce an anti-correlation whose scatter is completely in agreement with observations. Galliano et al. 2021 reproduced the T-β anti-correlation in MBB fits (Figure 5.1 to the combined DustPedia and *Herschel* Dwarf Galaxy surveys.

These spurious and unphysical parameter correlations can be minimized through the application of a hierarchical or multilevel Bayesian model, which naturally incorporates and accounts for measurement and calibration uncertainties and their introduced correlations in the fitting through Bayesian priors. In hierarchical Bayesian fitting, the prior distributions which usually are used to encode external knowledge about the parameters are constructed from the fitted parameter distributions and described with a new set or level of parameters and priors known as the hyperparameters and hyperpriors. Kelly et al. 2012 developed an approach using a hierarchical Bayesian model to add random and correlated noise components as hyperparameters

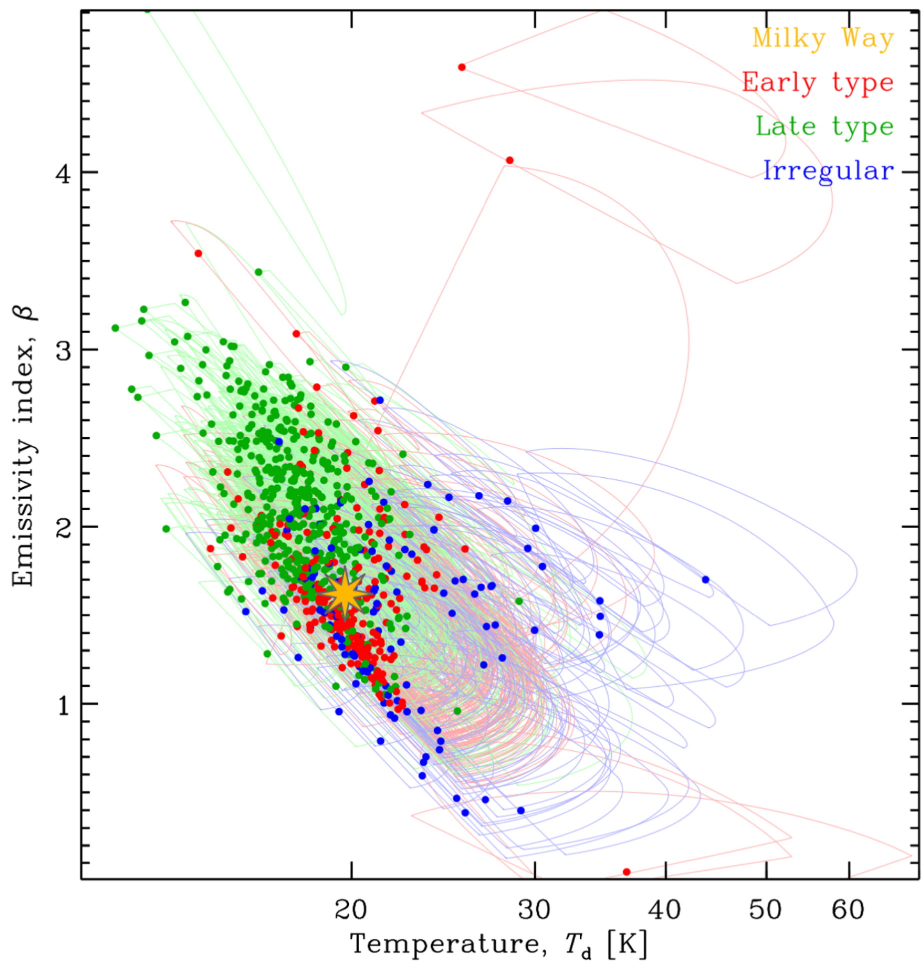

Figure 5.1. Figure reproduced from Galliano et al. 2021 showing the T-β anti-correlation recovered from a hierarchical Bayesian STMBB fit to galaxies of the Dust-Pedia and DGS samples. Colors represent the morphological type. The ellipses are the 1σ Skewed Uncertainty Ellipses derived from the posterior distribution.

to MBB model fits of simulated dust SEDs that correctly recovered the true correlations between T and β. This approach has since been utilized to develop a suite of MCMC fitting codes for MBB and TTMBB models (Veneziani et al. 2013; Lamperti et al. 2019; Tang et al. 2021), as well as physical dust models (Galliano 2018; Galliano et al. 2021).

This chapter details the C++ hierarchical Bayesian MCMC SED fitting software tool that I have developed primarily for the purpose of fitting the dust SEDs within each pixel of TolTEC observations of nearby massive and dwarf galaxies, though its modular nature allows for the incorporation of a range of SED models. The approach implemented here closely follows the work described in Tang et al. 2021 who wrote a hierarchical Bayesian fitting software package also in C++ to perform MBB fits to each pixel in observations of the Central Molecular Zone (CMZ) carried out with *Herschel*, AzTEC, Bolocam, and *Planck*. A forward-fitting strategy where the model in each wavelength band is effectively convolved with the corresponding instrumental PSF for every step in the MCMC chains was utilized to ensure that the results could fully leverage the higher resolution bands without needing to degrade all observations to the lowest resolution of the incorporated data. In Tang et al. 2021, for the single temperature MBB model, only the hyperparameters related to T and β were used in the hierarchical model. A multi-temperature MBB model that assumes a polytropic equation of state was also implemented, which included hyperparameters for T and the polytropic index. The code described here implements the hyperparameters and hyperpriors in a model-independent framework allowing for all parameters of any included model and parameters from linear combinations of separate models to be folded into the hierarchical model.

5.1 Bayesian Formalism

Following both the notation and formalism outlined in Kelly et al. 2012 and Galliano 2018, the observed flux density F_{obs} for source i and band j can be written as

$$F_{\mathrm{obs}_j} = \delta_j \cdot F_{\mathrm{model}_j}(x_i) + \epsilon_{ij}, \tag{5.1}$$

with $F_{\mathrm{model}_j}(x_i)$ being the modeled flux for the same source and band given the fitted parameter vector x_i, δ_j a calibration uncertainty for band j, and ϵ_{ij} the measurement noise for each pixel. At the core of any Bayesian inference is the Bayes theorem which is given by

$$p(x_i|F_{\mathrm{obs}_j}, F_{\mathrm{model}_j}) = \frac{p(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j}) \cdot p(x_i|F_{\mathrm{model}_j})}{p(F_{\mathrm{obs}_j}|F_{\mathrm{model}_j})}. \tag{5.2}$$

The Bayes theorem relates the posterior distribution $p(x_i|F_{\mathrm{obs}_j}, F_{\mathrm{model}_j})$ which is the probability of acquiring the parameter vector of the particular model given the data, to the product of the prior and the likelihood probability. The prior, $p(x_i|F_{\mathrm{model}_j})$, is a probability distribution that describes what is known about the parameters before the inference and represents the initial belief or knowledge about the parameters of the model. The choice of prior can range from very informative, specifying expected values and variances based on prior knowledge, to non-informative or weakly informative, which have minimal impact on the posterior distribution. The likelihood, $p(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j})$, encodes the probability of obtaining the data given the model, the current choice of parameters, and the measurement uncertainty. It is often assumed to be a normal distribution and can be written in the form

$$L(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j}) = \exp\left(-0.5\left(\frac{F_{\mathrm{obs}_j} - F_{\mathrm{model}_j}(x_i) \cdot \delta_j}{\sigma_{ij}}\right)^2\right). \tag{5.3}$$

The normalization constant $p(F_{\mathrm{obs}_j}|F_{\mathrm{model}_j})$ is the marginal likelihood which is the evidence or probability of obtaining the data given the model after integrating or marginalizing the likelihood function and prior over all the parameters in the model. When model fitting under the Bayesian framework, the normalization of the posterior is usually irrelevant unless comparisons between two models are required. We can therefore ignore the marginal likelihood and write the Bayes theorem as a proportionality or

$$p(x_i|F_{\mathrm{obs}_j}, F_{\mathrm{model}_j}) \propto p(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j}) \cdot p(x_i|F_{\mathrm{model}_j}). \tag{5.4}$$

The idea of Bayesian MCMC fitting is to sample from $p(x_i|F_{\mathrm{obs}_j}, F_{\mathrm{model}_j})$ by randomly sampling the parameters and evaluating Equation 5.4 for the corresponding parameter vector x_i.

5.2 Hierarchical Bayesian Formalism

In the hierarchical Bayesian extension to the single-level Bayesian modeling framework, we introduce a new set of variables, known as the hyperparameters, that describe the distribution of the parameters x_i across all sources. The priors for the parameters are then inferred based on the choice of probability distribution for the hyperparameters and the values of the parameters themselves. If x_h represents the vector of hyperparameters, Equation 5.4 can then be rewritten as

$$p(x_i|F_{\mathrm{obs}_j}, F_{\mathrm{model}_j}, x_h) \propto p(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j}) \cdot p(x_i|x_h) \cdot p(x_h). \tag{5.5}$$

Here, $p(x_h)$ is the prior on the hyperparameters or their hyperpriors. The parameter probability distribution $p(x_i|x_h)$ describes the probability of the model parameters given the values of the hyperparameters. It is often assumed to follow a multivariate Student's t-distribution to ensure a robust estimate in the case of outliers:

$$p(x_i|x_h) = \frac{1}{\sqrt{|\Sigma|}}\left(1 + \frac{1}{(x_i-\mu)^T\Sigma^{-1}(x_i-\mu)}\right)^{-(d+q)/2}. \tag{5.6}$$

In this case, the hyperparameters are the vector of parameter means, μ, and the parameter covariance matrix, Σ which incorporates the variances and correlation coefficients that describe the model parameter distribution. The degrees of freedom, d, controls the shape of the Student-t distribution and the extent of outlier tails. I adopt $d = 8$, as used in Kelly et al. 2012 and Galliano 2018. The value q is the number of parameters in the model which may include all of the parameters that make up the model or only a subset of the most correlated ones to reduce dimensionality.

5.2.1 Hyperpriors

With the hierarchical model parameter distribution and hyperparameters chosen, we can now define additional priors, $p(\mu)$ and $p(\Sigma)$, for the hyperparameters themselves, such that Equation 5.4 becomes

$$p(x_i|F_{\mathrm{obs}_j}, F_{\mathrm{model}_j}, \mu, \Sigma) \propto p(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j}) \cdot p(x_i|\mu, \Sigma) \cdot p(\mu) \cdot p(\Sigma). \tag{5.7}$$

There are few requirements for the hyperprior on μ, so a uniform distribution is commonly assumed. The hyperprior on Σ is more involved owing to the fact that the parameter covariance matrix must be both symmetric and positive definite such that $|\Sigma| > 0$. It is therefore convenient to sample each of the distribution variances and correlation coefficients that make up Σ sequentially and place independent priors on them. Barnard et al. 2000 devised a separation strategy for Σ where the matrix can be decomposed as

$$\Sigma = SRS. \tag{5.8}$$

The matrix S is a diagonal matrix consisting of the distribution variances and R is the correlation matrix. The hyperparameter priors can then be re-expressed as

$$p(\mu) \cdot p(\Sigma) \sim p(\mu) \cdot p(S) \cdot p(R). \tag{5.9}$$

For $p(S)$, previous studies have used a normal distribution on the logarithm of each variance (Kelly et al. 2012; Galliano 2018) or a half-Cauchy distribution (Lamperti et al. 2019). I employ the former which was found to be slightly more robust for simulated distributions with low S/N fluxes. As given in Galliano 2018, it takes the form

$$p(S) = \prod_{k=1}^{n} \frac{1}{\sqrt{2\pi}\sigma_{\chi^2}} \exp\left(-\frac{(\ln(S_{kk}) - \ln(S_{kk}^{\chi_k^2}))^2}{2\sigma_{\chi_k^2}^2}\right), \tag{5.10}$$

where S_{kk} is the natural log of the variance for parameter k, $S_{kk}^{\chi^2}$ is the variance determined from a χ^2 fit, and $\sigma_{\chi_i^2}$ is the variance of the natural log of $S_{kk}^{\chi^2}$. A value of 10 is used for $\sigma_{\chi_i^2}$ for all k resulting in $p(S)$ being a relatively weakly informative prior. For sampling the correlation coefficients, an inverse Wishart distribution is used, where $p(R)$ becomes

$$p(R) = |R|^{\frac{q(q-1)}{2}-1} \times \left(\prod_{k=1}^{n} |R_{kk}| \right)^{-\frac{q+1}{2}}, \tag{5.11}$$

with q being the number of model parameters included in Σ, and R_{kk} is the principal submatrix made by removing row and column k from R. The inverse Wishart distribution produces a jointly uniform prior on the interval [-1,1] for all correlation matrices and can be used to acquire a probability and interval such that a randomly sampled correlation coefficient will result in a positive definite matrix given that all other coefficients are held constant. Each marginal probability, however, will not be uniform as the values of other correlation coefficients restrict the value for what the value of a newly sampled coefficient can be. This effect is illustrated in Figure 5.2. The interval for sampling each new correlation coefficient from to ensure $|R| > 0$ can be determined from the roots of $|R|$ after recognizing that it is a quadratic equation in R.

5.3 Calibration Uncertainties

The probability distribution for the calibration uncertainties can be assumed to take the form of a multivariate normal or multivariate Student t-distribution with a covariance matrix describing the relationships between each correlation coefficient δ_j for band j. It has been shown that incorporating calibration uncertainties as additional parameters can result in significantly longer MCMC convergence times since they can be highly correlated between bands (Lamperti et al. 2019). More involved sampling techniques like the ancillarity–sufficiency interweaving strategy (ASIS; Yu and Meng 2011) are required for convergence (Kelly et al. 2012; Galliano 2018). I do not account for the correlation uncertainties in the hierarchical model in the current implementation of the code described here owing to the additional computational

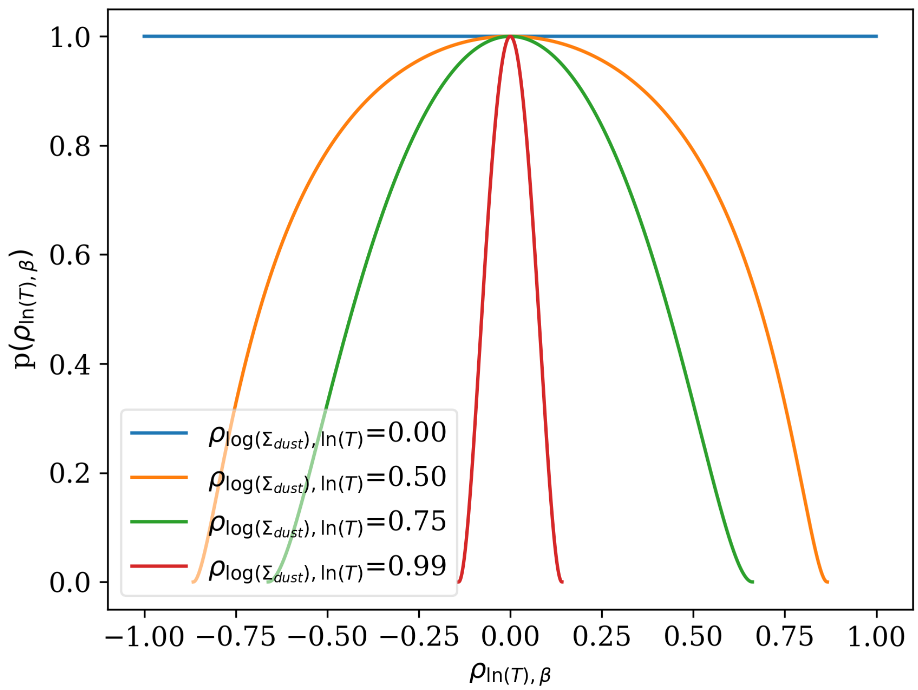

Figure 5.2. Example distributions of the MBB hyperprior $p(R)$ for the $\ln(T) - \beta$ correlation coefficient, $\rho_{\ln(T),\beta}$ given different values of $\rho_{\log(\Sigma_{dust}),\ln(T)}$ and with $\rho_{\log(\Sigma_{dust}),\beta} = 0$. The allowed probabilities for $\rho_{\ln(T),\beta}$ will be restricted to ensure the correlation matrix R is positive definite.

requirements of the PSF convolutions performed each MCMC step carried out and the preliminary nature of the results being presented. I do, however, plan to fully integrate calibration uncertainties in conjunction with additional samplers, namely the ASIS and the No-U-Turn Sampler (NUTS; Hoffman and Gelman 2011) in a future version.

5.4 Slice-within-Gibbs Sampling

I use the slice-within-Gibbs (Neal 2000) sampling strategy as implemented in Tang et al. 2021. Gibbs sampling is an MCMC algorithm often used for obtaining parameters from a multivariate probability distribution when direct sampling is difficult. The essence of Gibbs sampling lies in its iterative approach, where each parameter is updated sequentially, sampling from the conditional distribution of that parameter given the value of all the other parameters from the previous MCMC step. This process exploits the fact that sampling from the conditional distribution of a subset of variables is often simpler than sampling from the joint distribution of all variables. Over many iterations, the sequence of samples thus generated converges to the target distribution. The sampling of each parameter in turn is critical for updating the hyperparameters where each new correlation coefficient must be tested against the current correlation coefficients to ensure that the covariance matrix Σ is positive definite.

The slice-within-Gibbs algorithm is outlined in Figure 5.3. Slice sampling operates by defining a "slice" or a region under the probability density function at a certain level and then uniformly samples from this slice. It starts by using the current parameter guess and a corresponding slice by evaluating the probability density at that point. A new sample is then drawn from within this slice, ensuring it adheres to the target distribution. The key advantage of slice sampling over other MCMC methods is its self-tuning property; it automatically adjusts the size of the slice according to

Input:

$\ln(p)$ = natural logarithm of the posterior distribution

x_0 = current parameter vector

y = the slice height, defined as $\ln(p(x_0)) - e$, where $e \sim$ Exponential(1)

(L, R) = the interval to sample from

ϵ = the minimum tolerance for the separation between $\hat{L}$ and $\hat{R}$

Output:

x_1 = new parameter vector

$\bar{L} \leftarrow L,\ \bar{R} \leftarrow R$

repeat

    $U \sim$ Uniform$(0, 1)$

    $x_1 \leftarrow \bar{L} + U \cdot (\bar{R} - \bar{L})$

    if $y < \ln(p(x_1))$ **then**

        Accept(x_1) and **exit loop**

    end if

    if $x_1 < x_0$ **then**

        $\bar{L} \leftarrow x_1$

    else

        $\bar{R} \leftarrow x_1$

    end if

until $\bar{L} - \bar{R} < \epsilon$

Figure 5.3. Figure adapted from Neal 2000 and Tang 2019. The Gibbs slice shrinkage algorithm used to sample new parameter and hyperparameter values.

the shape of the distribution, which helps in efficiently exploring the sample space without the need for manual tuning of parameters. This attribute makes slice sampling particularly useful in situations where the probability distribution has complex characteristics. Furthermore, unlike Metropolis-Hastings and other similar MCMC sampling strategies, there is no acceptance or rejection fraction of proposed parameters for each step of the MCMC chain in slice sampling with every step accepting a new parameter. One disadvantage of Gibbs sampling is that it can be inefficient if the variables are highly correlated.

5.5 SED Models

A core design choice of our code is the separation between the model definitions and both the parameter sampling and hierarchical model. This allows a range of models to be added and used in conjunction with one another to model different astrophysical emission mechanisms while employing the same hierarchical Bayesian formalism. The code incorporates models to describe dust, stellar, and nonthermal emission in order to fit galaxy SEDs across NIR and millimeter wavelengths, as well as a model for the thermal Sunyaev-Zeldovich effect to support future TolTEC observations of galaxy clusters.

For most models, the code explicitly calculates the models for each step of the MCMC within every pixel and does not use a precomputed model grid. The sole exception is the physically motivated dust model, as directly computing the dust emissivities from their dielectric functions is complex and computationally intensive.

5.5.1 Single Temperature Modified Blackbody

I implement the single temperature MBB (STMBB) model for fitting the dust emission between FIR and millimeter wavelengths. An in-depth discussion of the

STMBB model and its variants is given in Section 1.0.4.1. I adopt the most general form for the model without assuming that the dust is optically thin (i.e. $\tau_\nu << 1$) as

$$I_\nu = \left(1 - e^{-\kappa_0\left(\frac{\nu}{\nu_0}\right)^\beta \Sigma_{dust}}\right) \times B_\nu(T). \tag{5.12}$$

I use a value of $\kappa_0 = 9.93\ cm^2/g$ at $\nu_0 = 160$ µm following the recommendation of Galliano 2022 to not use submillimeter wavelengths owing to the higher amount of scatter in laboratory dust analogs (Coupeaud et al. 2011). The value is taken from the reported opacity of the Astrodust model from Chastenet et al. 2021. Furthermore, as shown in Figure 1.4, the total extinction from the Astrodust model (Hensley and Draine 2023), which I use for the physically motivated dust model fitting in this work, and the THEMIS model are in good agreement at FIR wavelengths. Following Kelly et al. 2012, the dust temperature is sampled as $\ln(T)$ to normalize its range relative to the other parameters.

5.5.2 Physically Motivated Dust Model

The code implements the Astrodust + PAH dust model from Hensley and Draine 2023 for fitting dust emission with NIR and MIR constraints. It differs from contemporaneous dust models in that it approximates larger dust grains ($a \geq 0.02$ µm) as a single amalgamation of different materials as opposed to two separate populations of silicate and carbonaceous grains. This choice is motivated by measurements of the dust polarization fraction with *Planck* and BLASTPol between 250 µm and 3 millimeters which have variations of less than 10% as a function of wavelength. The polarized emission properties of individual silicate and carbonaceous grains are expected to be different such that different mixtures of grain types in the ISM will exhibit varying polarization fractions. If the dust grains are indeed divided into two categories, the emission from larger silicate grains would become increasingly important to the dust SEDs at longer wavelengths, which would lead to an evolution in the polarization fraction.

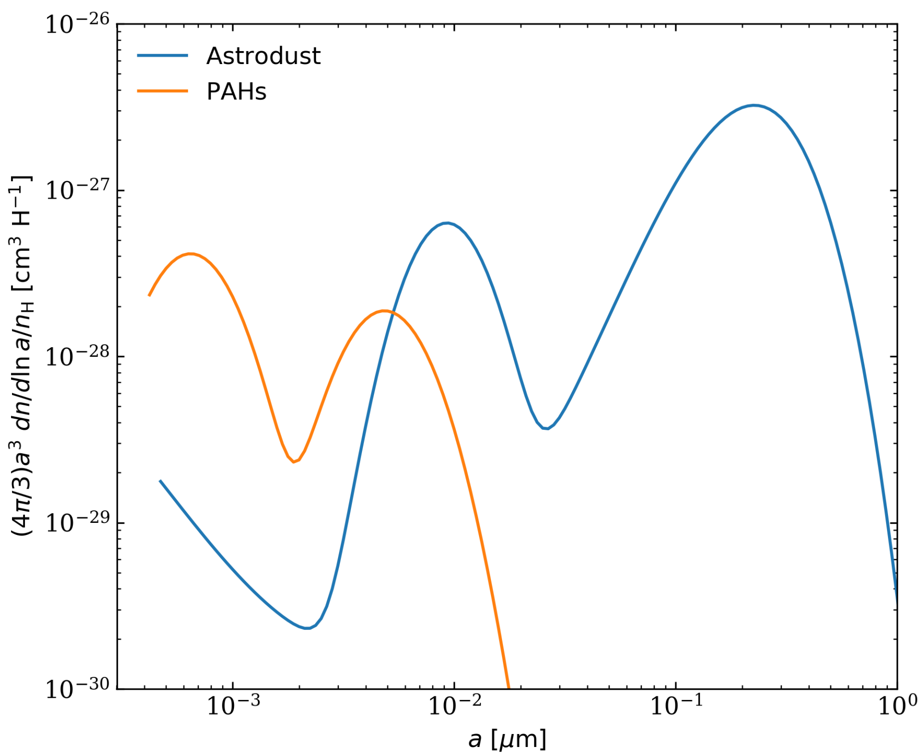

Figure 5.4. Size distributions for the Astrodust (blue) grains and PAH (orange) grains.

The Astrodust model assumes a spheroidal shape for the larger grains as asymmetry is a requirement for dust grains to emit and absorb polarized light (Draine and Hensley 2021). The grains are comprised primarily of silicates bound to Mg, Fe, and Ni with about 50% of the mass of the grains being accounted for by silicates. Iron and additional ferrous compounds such as Fe_3O_4 and FeS are added to match the measured iron depletions of the ISM and make up 30% of the grain mass. The remaining 20% of the mass is incorporated as carbon compounds like hydrocarbons and $CaCO_3$, as well as Al_2O_3 and SiO_2. The grain size distribution is estimated assuming a parametric form using MIR extinction and FIR to millimeter emission constraints from observations of the diffuse Milky Way ISM. The size distribution of Astrodust is shown in Figure 5.4 and is distinctly bimodal. Similar multi-modal distributions have been derived before when using polarization constraints (Kim and Martin 1995; Draine and Fraisse 2009) and are also a consequence of grain growth and evolution models (Li et al. 2021; Hirashita and Il'in 2022; Hensley and Draine 2023).

The model also includes PAH grains ($a < 0.01$ µm) that account for the remaining carbon mass of the ISM in order to explain MIR emission and extinction spectra. The derived PAH mass fraction for the Milky Way, q_{PAH}, is 5.8%. The PAH grains are modeled as being non-aligned such that they do not contribute to the observed dust polarization features. The derived PAH size distribution is similarly bimodal to that of the Astrodust grains and are also illustrated in Figure 5.4.

I use the publicly available model emissivities[1] for both the Astrodust and PAH grains to fit the dust SEDs. The model corresponds to the Milky Way extinction value of $R_V = 3.1$. I only include thermal emission from non-aligned grains and do not consider the total or polarized emission from aligned or spinning grains in this work. The emission quantities are tabulated for values of the ISRF field strength of $-3 < \log(U) < 6$ and for wavelengths of between $0.1 < \lambda < 3 \times 10^4$ µm in steps of 1.3×10^{-3} µm. I use the dust-to-gas ratios of $M_{astrodust}/M_H = 0.0064\Sigma_H$ and $M_{PAH}/M_H = 0.0007\Sigma_H$ to convert the tabulated emissivities from $\nu I_\nu/N_H$ to $\nu I_\nu/g$ to derive dust masses as given in Hensley and Draine 2023. The Astrodust and PAH emissivities are available either individually to allow them to be fit independently or as a total. I use the former to enable the fitting of q_{PAH} and thus model the dust emissivities as

$$j_\nu(U) = (1 - q_{\mathrm{PAH}}) \times j_{\nu,\mathrm{astrodust}}(U) + q_{\mathrm{PAH}} \times j_{\nu,\mathrm{PAH}}(U), \tag{5.13}$$

[1]https://dataverse.harvard.edu/dataverse/astrodust

where $j_{\nu,\mathrm{astrodust}}(U)$ and $j_{\nu,\mathrm{PAH}}(U)$ are the Astrodust and PAH emissivities at frequency ν and ISRF value U. A two-dimensional bicubic interpolation across $\log(U)$ and wavelength is used to derive j_ν for any arbitrary value of each parameter. For the starlight distribution, we use the Dale et al. 2001 parametrization of $dM_{dust}/dU \sim U^{-\alpha}$ (see Section 1.0.4.2 for a more in-depth description) and numerically integrate $j_\nu \times U^{-\alpha}$ over 1000 points between U_- and U_Δ which are sampled in the MCMC fitting algorithm.

5.5.3 Stellar Emission

Starlight emission from old stellar populations of cool, low mass stars contributes to SEDs at NIR wavelengths and can affect estimates of the contribution to the dust mass from small grains. I include a stellar emission component which is modeled as a Rayleigh-Jeans law in the form of

$$F_\nu = \frac{2\nu^2 k_B F_{\mathrm{scale}}}{c^2}, \tag{5.14}$$

where k_b is the Boltzmann constant and F_{scale} is a scaling factor accounting for the temperature and bolometric flux value and is the single fitted parameter. The effects of dust extinction and reddening aren't taken into account in the stellar SED, so the applicability of this simple approximation below NIR wavelengths is limited. The fit is therefore restricted to $\lambda < 4.0$ µm.

5.5.4 Nonthermal Emission

Nonthermal contributions to the SED arising from free-free and synchrotron emission can represent a significant fraction of the total flux for millimeter observations. I adopt the prescription outlined in Galliano 2018 to fit nonthermal emission, which takes the form of

$$F_\nu = \frac{f_1}{\nu_1}\left(f_{\rm ff}\left(\frac{\nu}{\nu_1}\right)^{-0.1} + (1-f_{\rm ff})\left(\frac{\nu}{\nu_1}\right)^{\alpha}\right). \tag{5.15}$$

Here, F_1 is flux at 1 cm with $\nu_1 = c/1\text{cm} = 2930$ GHz, f_{ff} is the fraction of free-free flux, and α is the synchrotron index.

5.5.5 Thermal Sunyaev-Zeldovich Effect

TolTEC will carry out observations of the thermal SZ effect (Sunyaev and Zeldovich 1970; Sunyaev and Zeldovich 1972) in galaxy clusters which will take advantage of the camera's simultaneous measurement of the flux increment, the null, and decrement of CMB photons within its 273, 214, and 150 GHz bands. Following Carlstrom et al. 2002, the spectral distortion to the CMB spectrum from the thermal SZE can be expressed as

$$\Delta I_{\rm SZE} = g(x) I_0 y, \tag{5.16}$$

with $x = \frac{h\nu}{k_B T_{\rm CMB}}$, $I_0 = \frac{2(k_B T_{\rm CMB})^3}{(hc)^2}$, and

$$G(x) = \frac{x^4 e^x}{(e^x-1)^2}\left(\frac{e^x+1}{e^x-1} - 4\right)(1 - \delta_{\rm SZE}(x, T_{\rm CMB})). \tag{5.17}$$

The term $\delta_{SZE}(x, T_{CMB})$ is a relativistic correction, which is ignored in the current implementation. The Compton-y parameter is the single model parameter and is defined as the integral of the electron pressure, $n_e k_B T_e$ where n_e is the electron number density, and T_e is the electron temperature along the line of sight through the cluster:

$$y = \int \frac{k_B T_e}{m_e c^2} n_e \sigma_T \, dl. \tag{5.18}$$

Here, m_e is the electron mass, σ_T is the Thomson scattering cross-section. An important consideration that can bias estimates of the Compton-y parameter when fitting the SZ effect in galaxy clusters is the presence of foreground dusty star-forming galaxies (DSFGs) which will take the form of point sources in the observed field. DSFGs can be simultaneously fit using MBB models and incorporated with the SZ model in the hyperparameters to mitigate their influence.

5.6 Forward Fitting Strategy

The code uses the forward-fitting strategy outlined in Tang 2019 and Tang et al. 2021, with the entire fitting process being illustrated in Figure 5.5. The input data consists of a set of images at different wavelengths that are aligned onto the same pixel grid, in addition to images of the PSFs for each band that are assumed to have the same pixel scale as the data. The images are sampled at the native resolution of the instrument and need not be convolved to the same resolution, as the PSF information is used to generate the model images. A mask image is an optional input that restricts which pixels are used in the calculation of the hyperpriors to limit the potential biasing due to the inclusion of pixels that do not measure source flux. The mask does not exclude pixels from the fit itself to maintain the same number of measurements across all pixels. The outputs of the code are (1) the MCMC chains for every pixel, parameter, and hyperparameter, (2) median and maximum a posteriori maps for every model parameter, and (3) model flux maps derived from the output model parameter images.

For each step in the MCMC routine, the parameters are sampled for each pixel within the common pixel grid, and model flux images are generated for each band given the current choice of SED models. The model images can be generated for a single representative wavelength or can be generated from an integral over the instrument bandpasses, where the pixel flux is then given by

$$F_{\nu,0} = \frac{\int F_\nu \cdot g_\nu \, d\nu}{\int \frac{\nu_0}{\nu} \cdot g_\nu \, d\nu}. \tag{5.19}$$

where g_ν is the normalized bandpass transmission at frequency ν. After model generation, the model fluxes are convolved with the corresponding PSF to dilute it to the

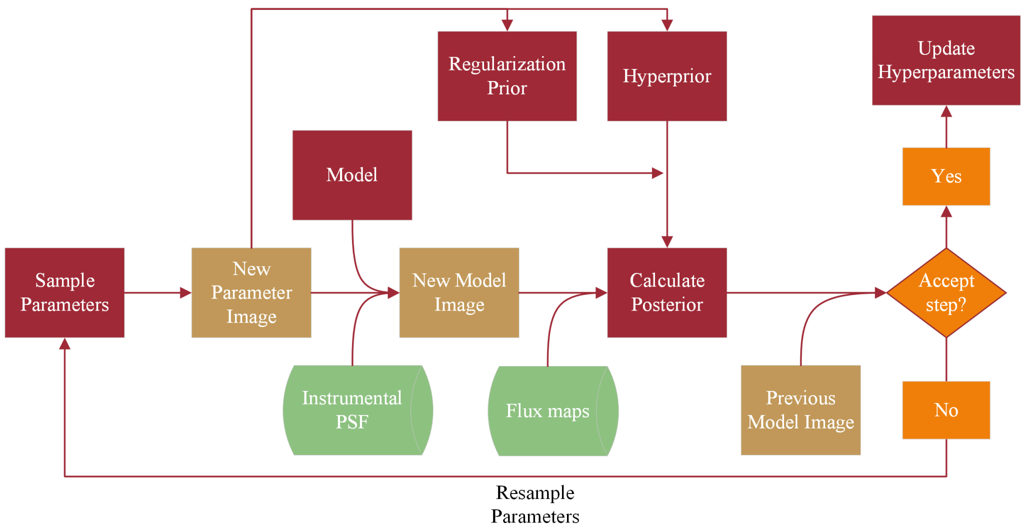

Figure 5.5. The fitting strategy used in the hierarchical Bayesian routine presented here. Red blocks represent functions, tan blocks are values calculated over the common pixel grid, and green are raw data inputs.

resolution of the observation. Our approach differs from Tang et al. 2021 in that the entire convolution is carried out over the full pixel grid at once for every proposed MCMC step as opposed to dividing the input grid into smaller blocks and only convolving within a radius of $8 \times \sigma_{psf}$, where σ_{psf} is the standard deviation of the PSF. The posterior is then evaluated for each pixel taking into account the proposed parameters and current hyperparameters according to Equation 5.7, and the parameters are accepted or re-sampled according to the slice within Gibbs acceptance criteria. The hyperparameters are then sampled and updated after the model parameters for each pixel has been updated.

In addition to its use in Tang et al. 2021 to study the CMZ, similar approaches that incorporate the PSF information into Bayesian inference fits have been employed previously to fit the dust continuum emission from the *Herschel* infrared Galactic Plane (Hi-GAL) survey (Marsh et al. 2015; Marsh et al. 2017). These studies demonstrated an improvement in the accuracy of the recovered dust column density partly as a result of the improved resolution of the incorporated data.

A quantitative determination of the resolution of each parameter map is challenging as different model parameters are more strongly constrained by observations at different wavelengths which can have PSF FWMHs that differ by an order of magnitude. For example, in MBB model fitting, the spectral index β will not depend as strongly on observations with $\lambda < 100$ µm as it would for observations at submillimeter and millimeter wavelengths. If fitting with only *Herschel* observations, the lower resolution SPIRE bands will likely have a greater effect on the recovered resolution for the β image compared to the higher resolution PACS bands. When fitting with physical dust models, the PAH component will be better constrained by higher resolution MIR observations and depend much less on FIR bands.

5.7 Bayesian Regularization

One complication that is introduced by the forward-fitting method is a tendency to overfit each pixel which arises due to intrinsic and noise variations within the higher-resolution data. The effect can be illustrated by considering the fit to a compact source with an angular extent smaller than that of the lowest resolution band. Such a source would appear as a PSF-like object in the lower-resolution bands but may be resolved in the higher-resolution data. The best-fit parameter maps would be a delta function for the observations with coarser resolutions but extended for more finely-grained bands. The extent of overfitting depends on the pixel scale, with larger pixel sizes minimizing variations in the higher-resolution data because of the pixel binning.

Overfitting in Bayesian inference is a well-studied phenomenon, particularly in the context of machine learning where the number of parameters is significantly higher than the number of observations. It can be mitigated by Bayesian likelihood regularization, which adds a penalty term to the likelihood in the form of

$$\ln\left(p(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j})\right) = \ln\left(p(F_{\mathrm{obs}_j}|x_i, F_{\mathrm{model}_j})\right) + \lambda G, \tag{5.20}$$

where G is a function that depends on the model parameters and lambda is a scaling factor. The form of G can be chosen to reduce the probability of certain parameter values. Existing forms include L1 norm and L2 norm regularization, which suppress parameters with large magnitudes, as well as gradient and curvature regularization that minimize variations between adjacent pixels and generally assume a flat and a planar distribution for the image pixels respectively. Tang et al. 2021 adapted the gradient regularization technique of Warren and Dye 2003, with G being expressed as

$$G = \frac{1}{2}\sum_{\mathrm{p}} \left(\lambda_p \left(p_{i,j} - p_{i+1,j}\right)^2 + \left(p_{i,j} - p_{i,j+1}\right)^2 + \left(p_{i,j} - p_{i-1,j}\right)^2 + \left(p_{i,j} - p_{i,j-1}\right)^2\right). \tag{5.21}$$

Here, $p_{i,j}$ is the model parameter at pixel (i,j) and λ_p is a scaling parameter that controls the level of regularization or smoothing. It can be expressed in terms of a parameter standard deviation in the form of $\lambda_p = 1/2\sigma_p^2$. The likelihood is therefore modified depending on the magnitude of the difference between the parameter values at (i,j) and those of nearby pixels. The optimal value of λ_p varies with the pixel scale, image resolutions, and model parameters. Values in the results presented here were determined empirically based on a visual inspection of the best-fit model images and parameter distributions. I also explored the curvature regularization function described in Warren and Dye 2003, but found it produced similar results to the gradient regularization for the values of λ used here.

5.8 Code Overview

The codebase incorporates many of the software lessons related to optimization and parallelization learned during the development of the TolTEC data reduction pipeline, which is described in more detail in Section 3.2. Carrying out the model evaluation and PSF convolution for every step of the MCMC chain is computationally expensive. The images considered in this study typically have $> 10^3$ pixels and may require $> 10^4$ samples to converge. To accommodate the computational requirements, I have written the code in C++, using Eigen, FFTW, and GrPPI to improve efficiency. The GNU Scientific Library[2] (GSL) is used for generating random deviates and for the interpolation of the physical dust model emissivity grids. The decision to carry out the PSF convolutions on the entire images was motivated by the use of FFT-based convolutions which are more efficient for large image sizes and to enable vectorization of the model generation across all pixels. Typical runtimes required to reach convergence on a single chain vary between ∼30 minutes for MBB fits with *Herschel* data only to approximately ∼10 hours for physical dust models using both *Spitzer* and *Herschel* constraints.

Parallelization is handled on a per-chain basis, where multiple independent MCMC chains can be used to sample the parameter space simultaneously. This helps mitigate MCMC chains from becoming stuck due to poor mixing or multi-modal posterior distributions.

5.9 Results

In this section, I present the results of the application of the hierarchical Bayesian fitting package outlined above to simulated SEDs as well as observations of the nearby galaxy NGC 3938. The parameter space to explore is large, consisting of all the parameters and hyperparameters from the various implemented models, their correlations, and the ranges of possible S/N and resolutions of the data. For this reason, I focus primarily on demonstrating the correctness of results and restrict the analysis to the dust and stellar models as these directly relate to the primary intended pur-

[2]http://www.gnu.org/software/gsl/

pose of the fitting code, which is to fit dust SEDs in nearby star-forming and dwarf galaxies with TolTEC constraints.

5.9.1 Simulations

I first test the fitting algorithm and model implementations by applying them to a population of synthetic SEDs generated from the MBB and physical dust model prescriptions. I begin by exploring the fitting without considering the instrumental resolutions or likelihood regularization, which is equivalent to assuming that all bands are at the same resolution and each measurement is independent outside of its distribution as sampled by the hyperparameters. This would be representative of the case where each source is the integrated flux from a galaxy as opposed to pixels in an image of a single source and is useful for characterizing the MCMC sampler and dust models. A total of 625 synthetic SEDs on 25×25 grids were generated by randomly drawing model parameters from a multivariate normal distribution. This allows for explicit control of the intrinsic correlations among the input parameters. I add Gaussian distributed noise to the flux estimates of each source that corresponds to the average S/N across all the sources. The MCMC is run for at least 20000 samples for all cases and tested for convergence, with the median value of the final 5000 samples used as the parameter value.

5.9.1.1 Simulated Single Temperature MBBs

Figures 5.6, 5.7, and 5.8 show the T-β distributions for hierarchical and non-hierarchical fits to simulated single temperature MBB models sampled at the *Herschel*/PACS, *Herschel*/SPIRE, and TolTEC bands for varying noise. The input multivariate normal distribution is given in Figure 5.1. All correlation coefficients including $\rho_{T,\beta}$, are set to zero for the input data. For all S/N values considered, the fitting results using the hierarchical model outperform the non-hierarchical fit, with the former being characterized by less scatter and little to no introduction of a correlation between the parameters. Above S/N ≥ 50, the distributions and uncertainties are nearly identical. The noise-induced anti-correlation is clearly visible in the non-hierarchical fit for S/N ≤ 5. The scatter in the hierarchical fits tends to underpredict the intrinsic scatter of the parameters, particularly for the low S/N cases. This effect was previously observed in Galliano 2018 who found it to be the result of the changing shape of the likelihood relative to the prior distribution as the noise increases. The likelihood function broadens as the noise increases thereby amplifying the contribution from the nearly flat priors. This has the effect of flattening the posterior distribution. Both fits in the lowest S/N bin are biased, though the hierarchical fit remains much closer to the true values, and the individual source uncertainties are markedly larger.

Table 5.1. MBB Input Distribution Parameters

Parameter	Value
$\mu_{\log \sigma_{\rm dust}}, \mu_T, \mu_\beta$	$-7, 20, 2$
$\sigma_{\log \sigma_{\rm dust}}, \sigma_T, \sigma_\beta$	$2, 5, 0.1$

Of particular interest for the study of nearby galaxies with TolTEC is the extent of the improvement that millimeter wavelength constraints can add to the recovery of MBB dust parameters and to the reduction of uncertainties and biases. Fits with and without TolTEC constraints are shown in Figure 5.9 for simulated dust SEDs with an average S/N value of 10. There is a reduction in the scatter and uncertainties across the sources relative to the *Herschel* only case, especially in β. A slight residual anti-correlation remains in the T-β distribution derived from the *Herschel* only result. Figure 5.10 shows the posteriors and parameter covariance distributions for all MBB dust parameters in both cases. While $\log(\Sigma_{dust})$ and T are similar, β is better constrained. The estimate of T depends more heavily on measurements near the peak of the dust SED and will be less affected by the inclusion of millimeter observations. While $\log(\Sigma_{dust})$ does depend on longer wavelength data to constrain the cold dust,

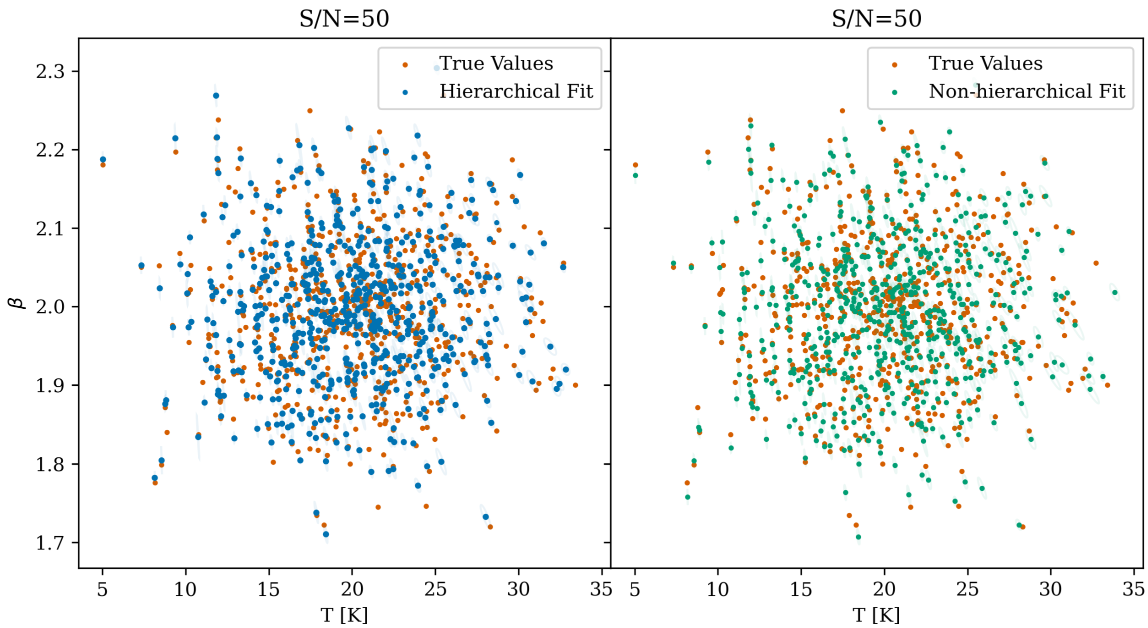

Figure 5.6. T-β distributions from fits to simulated MBB SEDs (orange points) with an average S/N=50. The left panel shows the fitted values using the hierarchical model (blue points) and the right panel is from a non-hierarchical Bayesian MCMC fit. The ellipses around each point are the 1-σ parameter uncertainties derived from the MCMC samples of that source.

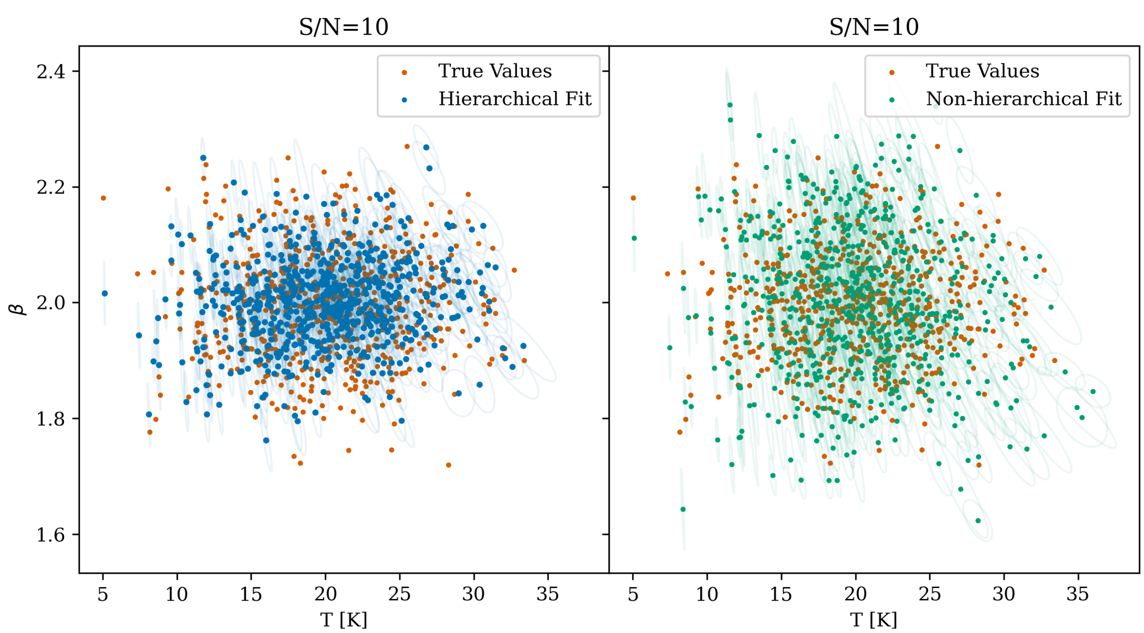

Figure 5.7. Same as Figure 5.6 but for an average S/N=10.

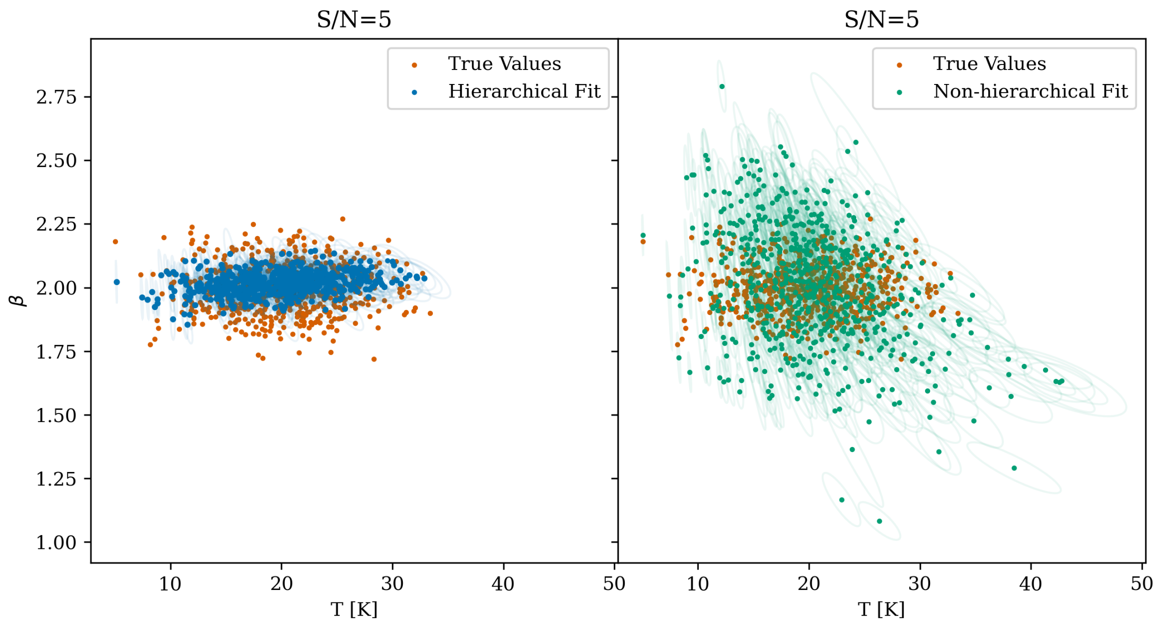

Figure 5.8. Same as Figure 5.6 but for an average S/N=5.

it is not subject to the same noise degeneracies as T and β and therefore doesn't vary significantly with additional millimeter observations.

5.9.1.2 Simulated Physical Dust Models

The previous analysis was repeated with the Astrodust + PAH physical dust model instead of the single temperature MBB. The parameters of the multivariate normal distribution from which the model parameters are drawn are given in Table 5.2. The flux values are sampled at the *Spitzer*/IRAC, *Spitzer*/MIPS 24 µm, *Herschel*/PACS, and *Herschel*/SPIRE, and TolTEC bands with varying levels of noise added to them. The correlations among parameters, particularly U_{-}, U_{Δ}, and α, can be complex so I combine their effects into the mean value of the radiation field $\langle U \rangle$ as given in Equation 1.9. The distribution of $\langle U \rangle$ and $\log(\Sigma_{dust})$ for different S/N values are illustrated in Figures 5.11, 5.12, and 5.13. Compared to the MBB case, the hierarchical fit of the physical dust model performs better even for high noise cases. Parameter chain histograms and joint posterior scatter plots for all model parameters are provided

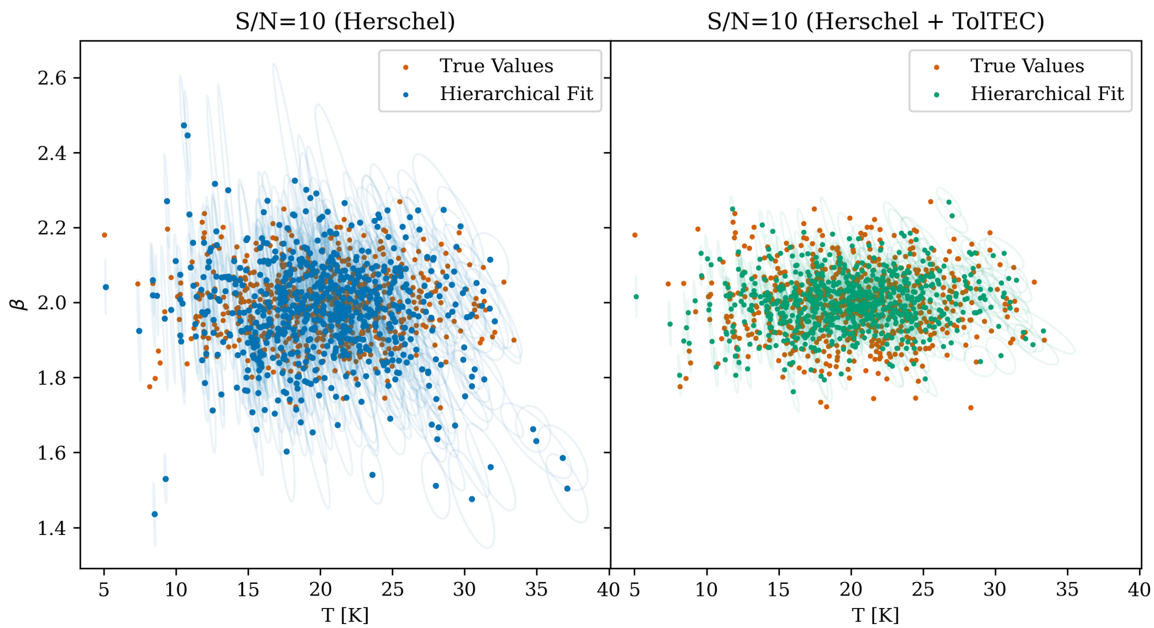

Figure 5.9. T-β distributions derived from hierarchical Bayesian MCMC fits to simulated SEDs with an average S/N=10. The left panel includes only fluxes within the *Herschel* PACS and SPIRE bands, whereas the right panel also adds fluxes at the 3 TolTEC bands.

for the S/N=5 simulation in Figure 5.14. The hierarchical fit results in much better constraints and reduces the strong correlations between in U_{-}, U_{Δ}, α, and $\log(\Sigma_{dust})$ seen in the non-hierarchical fit.

Table 5.2. Physically Motivated Model Input Distribution Parameters

Parameter	Value
$\mu_{\log\Sigma_{\rm dust}}, \mu_{\alpha}, \mu_{U_{-}}, \mu_{U_{\Delta}}, \mu_{q_{\rm pah}}$	$7, 2, 0, 3, 0.048$
$\sigma_{\log\Sigma_{\rm dust}}, \sigma_{\alpha}, \sigma_{U_{-}}, \sigma_{U_{\Delta}}, \sigma_{q_{\rm pah}}$	$0.1, 0.1, 0.1, 0.1, 0.01$

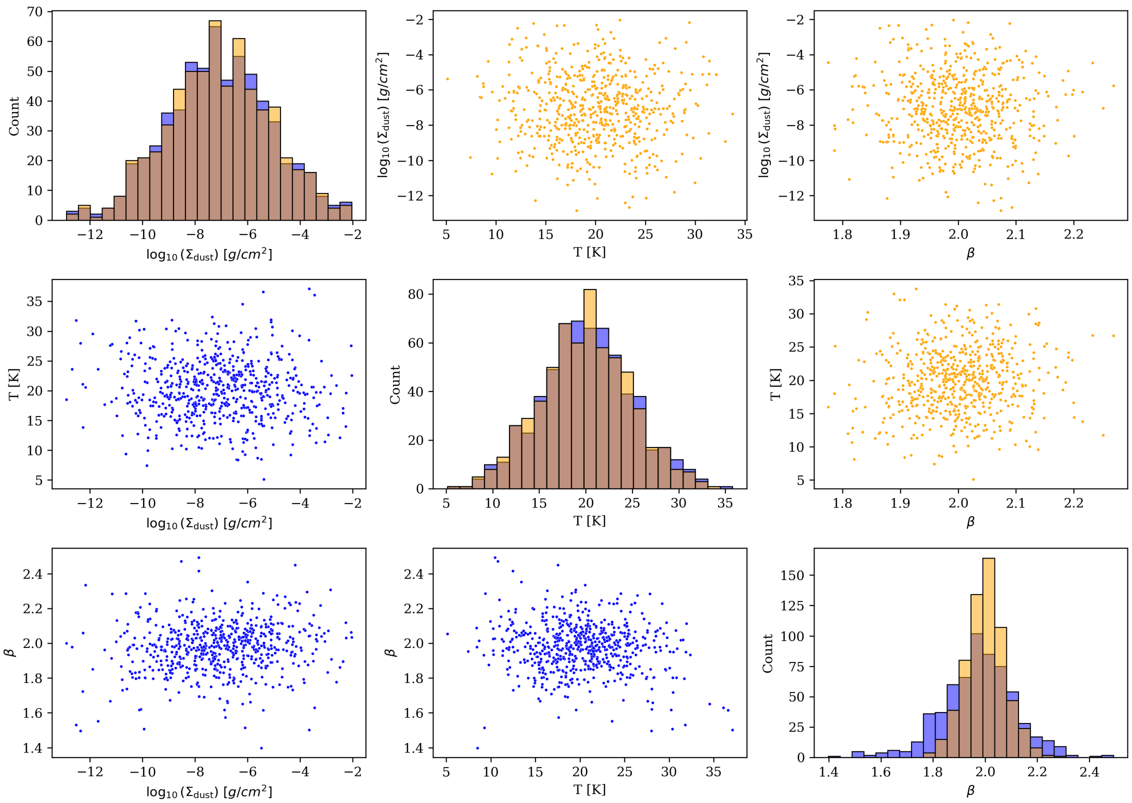

Figure 5.10. Posterior histograms and parameter-parameter distributions for hierarchical fits to the S/N=10 *Herschel* only (blue) and *Herschel* + TolTEC (orange) cases.

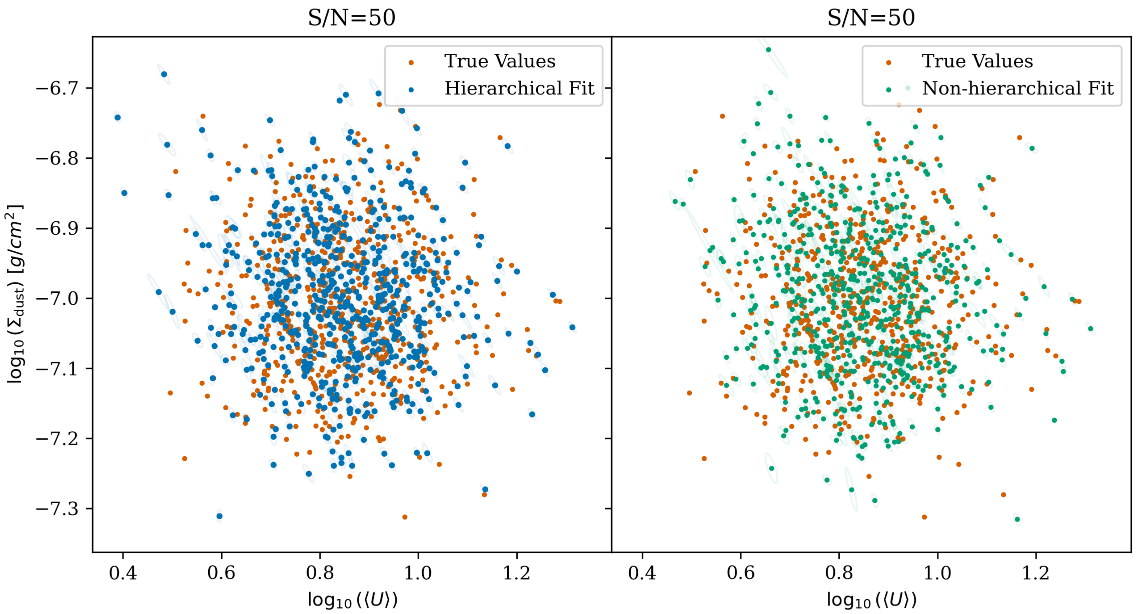

Figure 5.11. $\log(\Sigma_{dust})$-$\langle U \rangle$ distributions from fits to simulated physical dust model SEDs (orange points) with an average S/N=50. The left panel shows the fitted values using the hierarchical model (blue points) and the right panel is from a non-hierarchical Bayesian MCMC fit. The ellipses around each point are the 1-σ parameter uncertainties derived from the MCMC samples of that source.

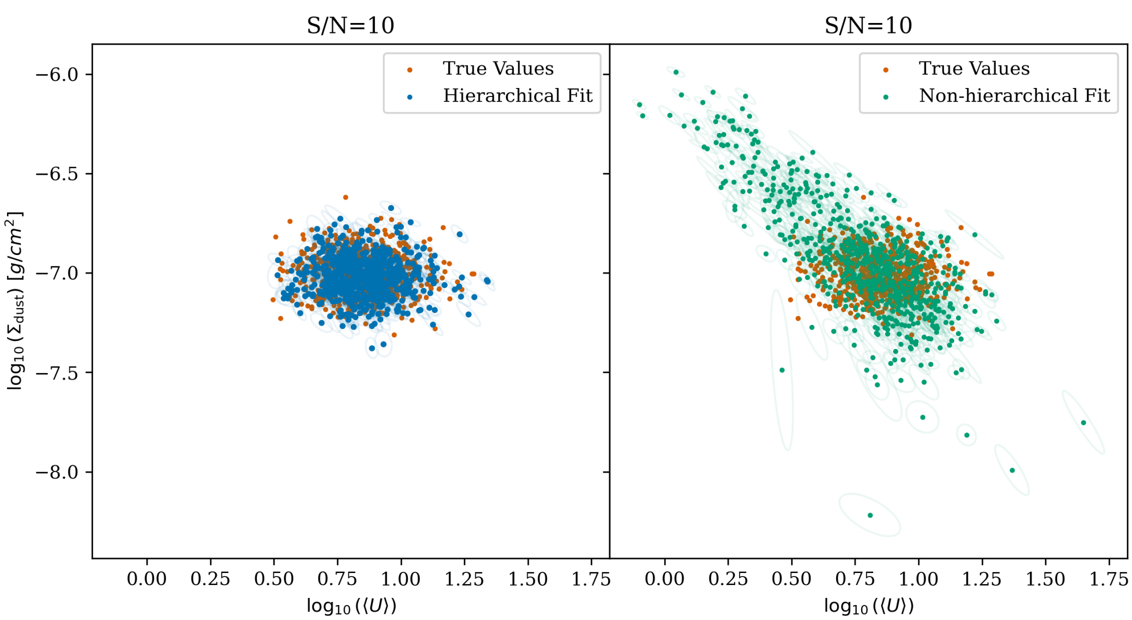

Figure 5.12. Same as Figure 5.11 but for an average S/N=10.

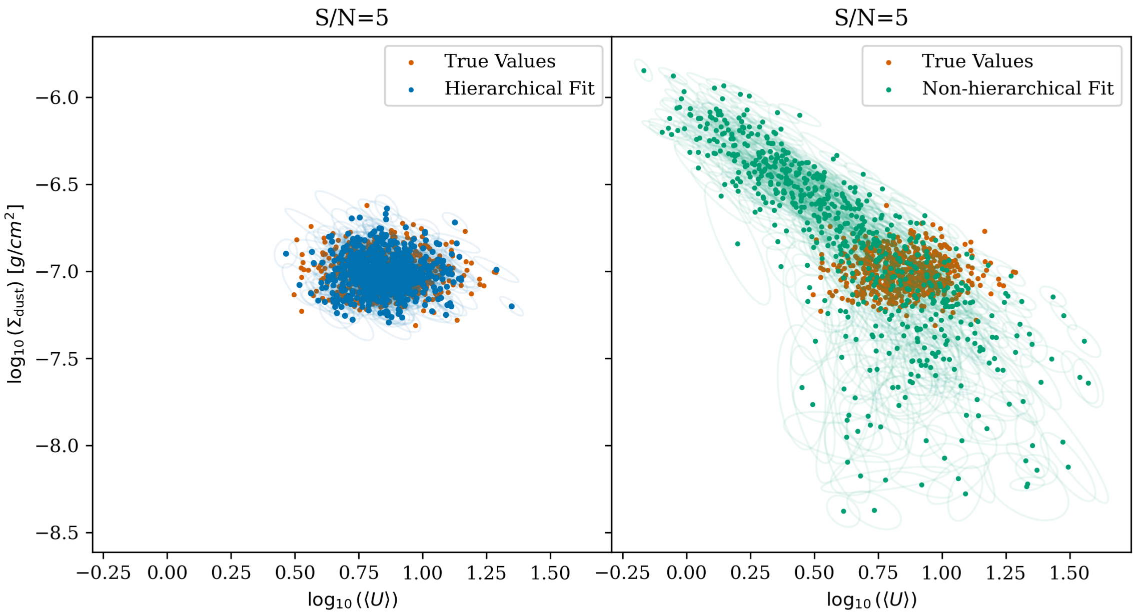

Figure 5.13. Same as Figure 5.11 but for an average S/N=5.

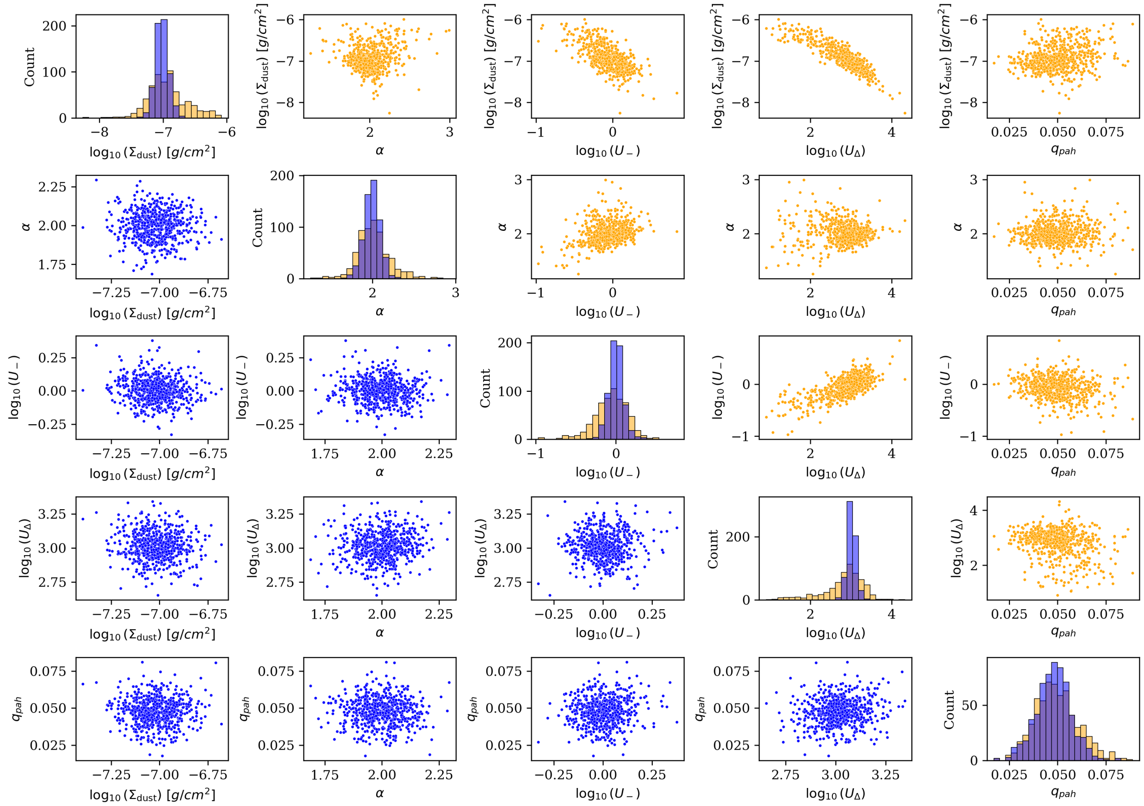

Figure 5.14. Posteriors and parameter-parameter scatter plots for the hierarchical (blue) and non-hierarchical (orange) fits to the simulated S/N=10 physical dust SEDs.

A comparison of the fits with and without TolTEC constraints is shown in Figure 5.15 for the $\langle U \rangle$ and Σ_{dust} relationship for the S/N=5 case. Unlike the MBB model, there is little improvement in the parameter estimations, which is likely due to the inclusion of MIR constraints.

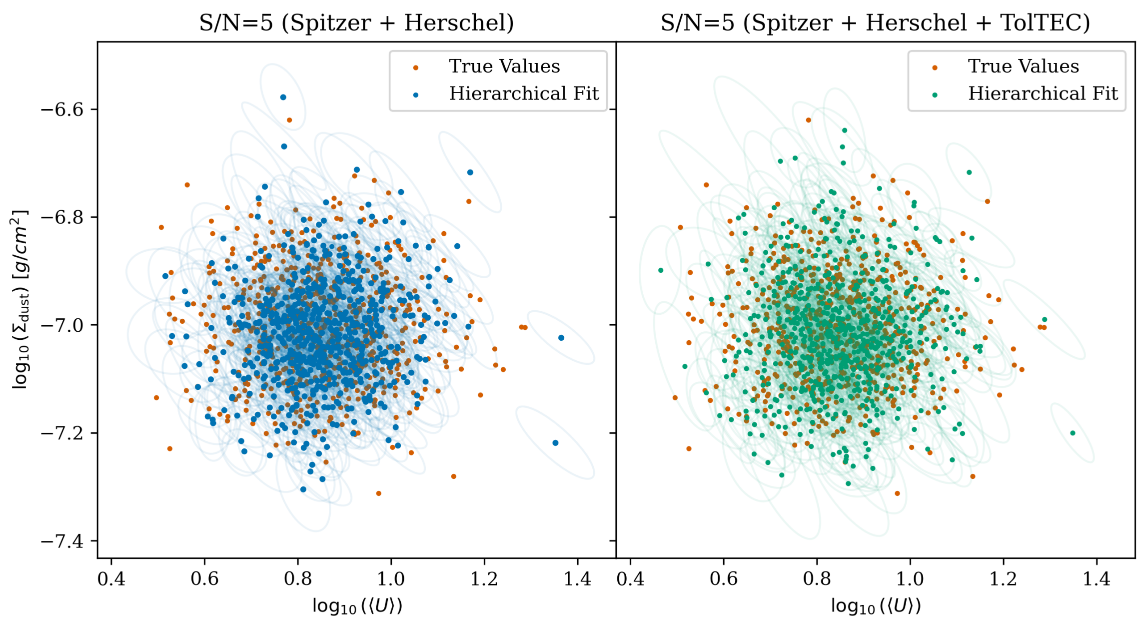

Figure 5.15. $\log(\Sigma_{dust})$-$\langle U \rangle$ distributions derived from hierarchical Bayesian MCMC fits to simulated physical dust model SEDs with an average S/N=5. The left panel includes only fluxes from all *Spitzer* and *Herschel* bands, whereas the right panel also adds fluxes at the 3 TolTEC bands.

5.9.2 NGC 3938 Analysis

I apply the fitting code on real data of the face-on nearby spiral galaxy NGC 3938 using archival WISE, Spitzer, and *Herschel* observations. NGC 3938 was selected for analysis primarily due to its inclusion in the list of TolTEC commissioning targets, its relatively small angular size, and its well-defined spiral structure which may show variations in dust properties. Image data of NGC 3938 was acquired from the DustPedia database which has collected multi-wavelength photometric data from 875 local galaxies and processed them to build a large homogeneous data set (Davies et al. 2017; Clark et al. 2018). Aperture-matched photometric fits to an MBB model as well as from panchromatic models with CIGALE using the THEMIS and Draine et al. 2014 dust models are also tabulated for most galaxies.

The data reduction for each band by the DustPedia team is detailed in Clark et al. 2018. Briefly, the *Herschel* SPIRE and PACS maps were processed with the *Herschel*

Interactive Processing Environment (HIPE) naive mapmaker and the SCANAMORPHOUS software package respectively. *Spitzer* IRAC data of NGC 3938 was acquired from the *Spitzer* Survey of Stellar Structure in Galaxies (S4G; Sheth et al. 2010) survey and MIPS observations used observations from the *Spitzer* Legacy/Exploration Science Programs (SEIP). Finally, the WISE data was obtained from the ALLWISE data release Image Atlas (Cutri et al. 2021).

The PACS 100 μm and 160 μm bands are used as opposed to the *Spitzer* MIPS alternatives to make use of the superior resolution offered by PACS. While maps are provided in units of Jy/pixel, pixel sizes vary between the bands. I use the Python package `reproject` which provides a flux-conserving regridding algorithm to align images onto a common pixel grid using their associated WCS information. A pixel size of $7''$ was chosen as a trade-off between the incorporation of information from the higher resolution bands and the level of regularization needed to prevent overfitting. The maps were then converted into units of MJy/Sr.

Uncertainty images are also provided by the DustPedia database for some bands (Clark et al. 2018), but are not available for the *Spitzer* IRAC 3.6 μm and 4.5 μm or any WISE bands. Therefore, to estimate uncertainties in a consistent way across all images, I fit the pixel histograms to a Gaussian function and use the derived standard deviation from the fit as the map uncertainty. The values are verified against blank-sky aperture estimates which are found to be in good agreement with one another. An example of the fit to the SPIRE 500 μm image is shown in Figure 5.16.

I also obtained aperture photometry available from the DustPedia database. For NGC 3938, an elliptical aperture with semi-major axis and semi-minor axes of $272''$ and $241''$ was used to determine the source flux. Uncertainties for the aperture photometry are also included and are derived from an iterative algorithm using a series of random apertures to estimate the map standard deviation. This procedure and the instrument calibration uncertainties are described in Clark et al. 2018 and included

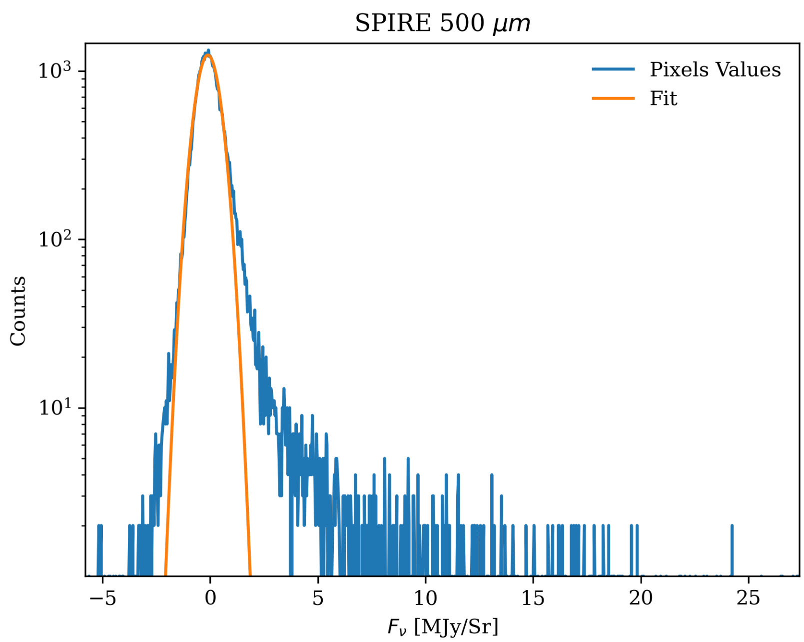

Figure 5.16. Example of the statistical uncertainty estimation method for the NGC3938 *Herschel* SPIRE 500 µm band. The blue curve is the histogram of pixel values in the map. The orange curve is a Gaussian fit to the pixels around the peak. The standard deviation of the fitted Gaussian is used for the map uncertainty.

in the uncertainty estimates on the aperture photometry and are between 2.9% and 7%.

I use the PSFs provided in Aniano et al. 2011 to dilute the model images to the instrumental resolution. These PSFs incorporate realistic characteristics including distortions and asymmetries and are aligned onto a common pixel grid. I re-grid each PSF from its original pixel size of $1''$ to $7''$ in the same manner as the source maps. I normalize the PSFs such that the sum of all pixels is unity.

Bandpasses were acquired from the database provided by the Spanish Virtual Observatory's Filter Profile Service (Rodrigo and Solano 2020). The native filters were downsampled in frequency space by a factor of 10 to improve reduction speed while preserving the bandpass shape.

5.9.2.1 Aperture Photometry

The result of fits to the DustPedia aperture photometry measurements with WISE, Spitzer, and Hershel is illustrated in Figure 5.17. As this is a measurement from a single source, the hierarchical Bayesian model is not used. The fit with the Astrodust + PAH model included all bands from the three instruments. A stellar contribution was included in the physical dust model fit to account for the NIR SED bump. The sampler was run for 10^5 samples, with convergence being reached within $2{\times}10^3$ samples for all models as determined via a visual inspection and autocorrelation tests.

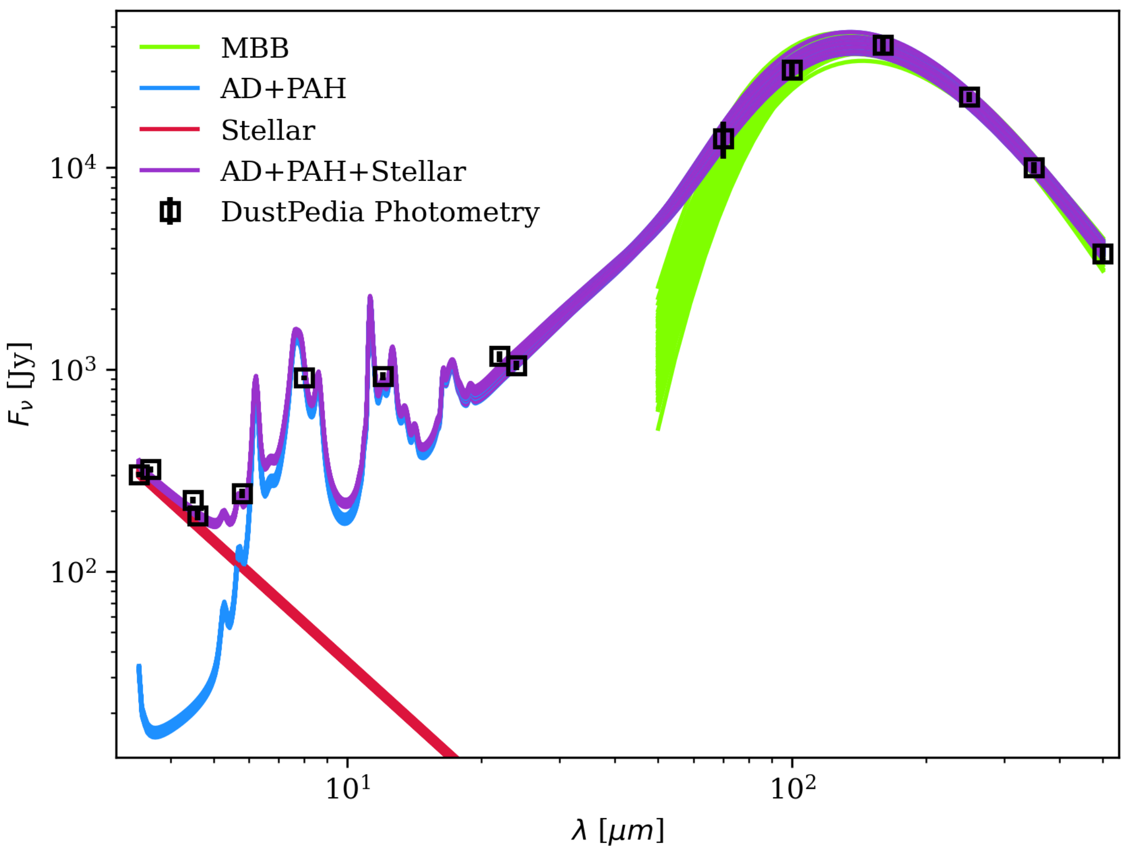

Figure 5.17. SED fits to the DustPedia integrated fluxes values of NGC 3938. The green curve is a fit of the *Herschel* PACS and SPIRE bands to a Modified Blackbody dust model. The blue curve is a fit to all *Herschel* and *Spitzer* bands to the Astrodust + PAH physical dust model. The red line is the stellar contribution and the purple line is the sum of the Astrodust+PAH fit and the stellar fit. The width of the fitted SEDs are derived from generating models from randomly selecting samples in the MCMC parameter chains after they have converged.

The MBB model the best fit values are given in Table 5.3. Using the distance of 19.4 Mpc from the DustPedia database, the derived dust mass is $3.58 \times 10^7 \pm 8.18 \times 10^6\ M_\odot$ which is in agreement with the reported DustPedia value of $3.07 \times 10^7 \pm 2.76 \times 10^6\ M_\odot$. The dust temperature is also consistent with the reported value of $21.991 \pm 0.634\ K$. I also try fitting with a β fixed to 1.79 to test the robustness of the free β solution and derive $M_{dust} = 4.40 \times 10^7 \pm 3.2 \times 10^6\ M_\odot$ and $T = 22.05 \pm 0.45K$.

Table 5.3. Best Fit Values for the MBB Model Parameters.

Parameter	Value
$\log(\Sigma_{\rm dust})$ [g/cm^2]	-10.63 ± 0.11
T [K]	22.29 ± 1.28
β	1.75 ± 0.18

For the Astrodust + PAH model fit, α was fixed to a value of 2.0 to improve convergence of the MCMC chain. The best fit values for the dust parameters are given in Table 5.4. The dust mass derived from $\log(\Sigma_{dust})$ is $M_{dust} = 3.28 \times 10^7 \pm 2.51 \times 10^6\ M_{\odot}$. The DustPedia CIGALE fits (analysis described in Nersesian et al. 2019) finds $M_{dust} = 2.746 \times 10^7 \pm 3.494 \times 10^6\ M_{\odot}$ and $M_{dust} = 7.676 \times 10^7 \pm 7.566 \times 10^6\ M_{\odot}$ when using the THEMIS and the Draine et al. 2014 dust models respectively. The former is in good agreement with the results derived here, while the latter is 2.3 times higher. Integrated flux measurements of NGC 3938 have also been fit to the Draine and Li 2007 dust model by Draine et al. 2007 ($M_{dust} = 4.90 \times 10^7\ M_{\odot}$) and by Aniano et al. 2020 ($M_{dust} = 5.20 \times 10^7 \pm 1.40 \times 10^7\ M_{\odot}$). Variations in the dust masses originate primarily from differences in the FIR opacities and choice of dust-to-gas ratios. The THEMIS opacity at 160 µm is slightly higher with a value of 14.2 cm^2/g (Galliano 2018) and the Draine et al. 2007 model uses 10.2 cm^2/g. A systematic comparison of dust model parameters including the dust mass between several dust models was carried out by Chastenet et al. 2021 in a spatially resolved study of the face-on spiral Messier 101. Chastenet et al. 2021 found that the dust masses obtained by using THEMIS and Astrodust models were very similar, whereas dust masses from the Draine et al. 2007 model were ~ 1.3 times higher than either model, which is consistent with the trend observed here.

The inferred PAH mass fraction is consistent with Draine et al. 2007 (0.046), but higher than the inferred value from the CIGALE DL14 model (0.028).

Table 5.4. Best Fit Values for the Astrodust + PAH Model Parameters.

Parameter	**Value**
$\log(\Sigma_{\rm dust})$ [g/cm^2]	-10.70 ± 0.03
α (fixed)	2.00
$\log(U_{-})$	-0.15 ± 0.05
$\log(U_{\Delta})$	3.16 ± 0.03
$q_{\rm PAH}$	0.035 ± 0.003

5.9.2.2 Spatially Resolved STMBB Fit

I fit each pixel using both *Herschel* PACS and SPIRE images to an STMBB using the full hierarchical model and instrumental PSFs. The PACS 70 µm band was excluded to ensure no contributions from small grains that are not in thermal equilibrium with the ISM. Multiple image sizes were considered, ranging between 21×21 pixels and 61×61 pixels, each with 7″ pixels, to test the effect of non-source background pixels on the values and distribution of the parameters. The 61×61 flux maps are plotted in Figure 5.18. The image pixel sizes and regularization parameter σ_p were determined empirically based on a trade-off between the observed pixel-to-pixel noise and the resolution of the dataset. Values between 0.01 to 1.0 were explored for all parameters. For each case, the sampler was run for 2×10^4 samples, and parameter maps were derived from the median values of the final 10^3 samples. Fits without the hierarchical model were also carried out for comparison using the same configuration.

The parameter maps derived from the fit to a 61×61 image are shown in Figure 5.19. The angular size of the maps, approximately 427″, corresponds to the size of the aperture used in the integrated photometry fits in Section 5.9.2.1 and includes a large fraction of background pixels. Qualitatively, the hierarchical model fit recovers much of the spiral structure and central bulge at a resolution higher than what would be obtained if all images were smoothed to the 36″ resolution of the SPIRE 500 µm band. The ranges for the parameters on the source in the hierarchical model are approximately $\log(\Sigma_{\rm dust}) = [-5, -3.9]$, $T = [12.5, 22.6]$, and $\beta = [1.8, 2.7]$. For the non-hierarchical case, $\log(\Sigma_{\rm dust})$ and T cover a similar range, but β reaches a

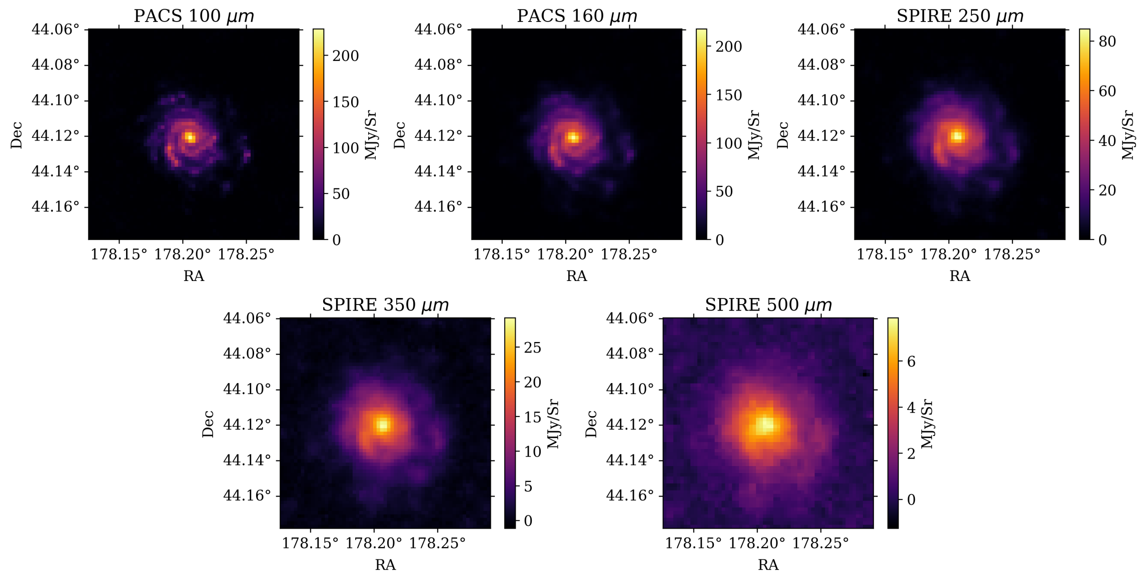

Figure 5.18. *Herschel* PACS and SPIRE flux maps of NGC 3938 from the 61×61 test case. The raw maps have been regridded onto the same pixel grid with a pixel size of 7″.

lower limit of 1.3 in the outer spiral arms. There is clear evidence of a strong anti-correlation between T and β in the hierarchical fit, with a marked enhancement of β and reduction in T within the central bulge. The effect is not as evident in the non-hierarchical fit, with the low β region in the outer arms not showing an increase in T compared to other parts of the galaxy. The value of β is elevated in the background pixels relative to that of the galaxy, which likely derives from the differences in the noise level between the bands.

Residual maps for the 61×61 image fits are shown in Figure 5.20 and Figure 5.21, with average map residuals being between 3-7% for the hierarchical fit and 3-6% for the non-hierarchical fit, respectively. There is little difference in the residuals between the two models, but some spiral structure is visible, particularly in the higher-resolution bands, though most of the spiral structure is smoothed out in the SPIRE 350 µm and 500 µm bands such that the residuals should appear smoother.

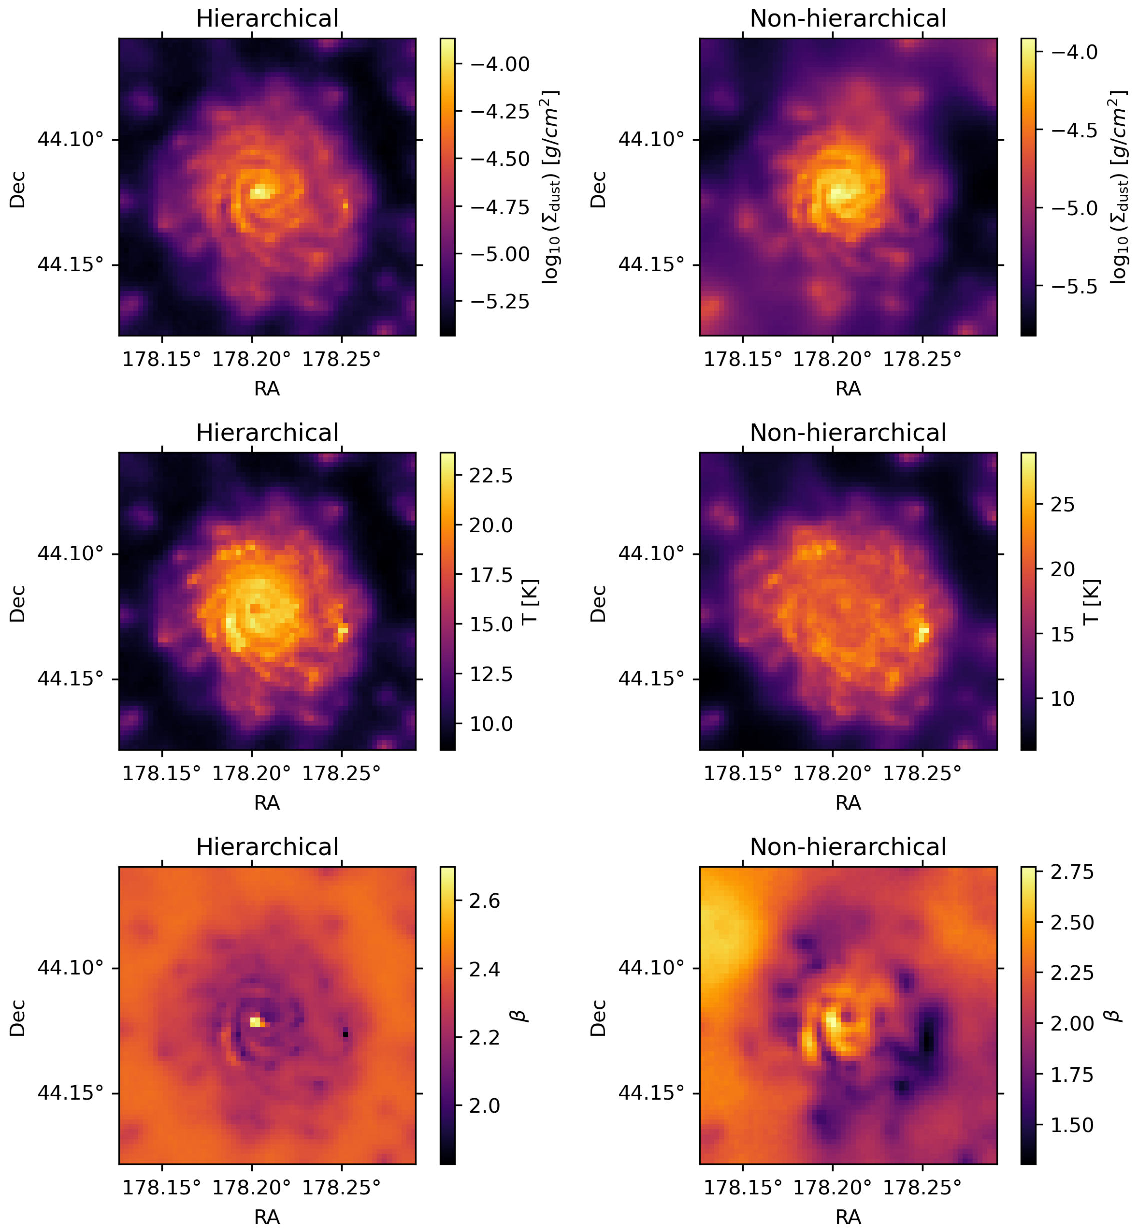

Figure 5.19. $\log(\Sigma_{dust})$, T, and β maps from hierarchical (left column) and non-hierarchical (right column) fits to the unmasked 61×61 maps of NGC 3938. Values used for each pixel are the median of the final 10^3 samples of each parameter.

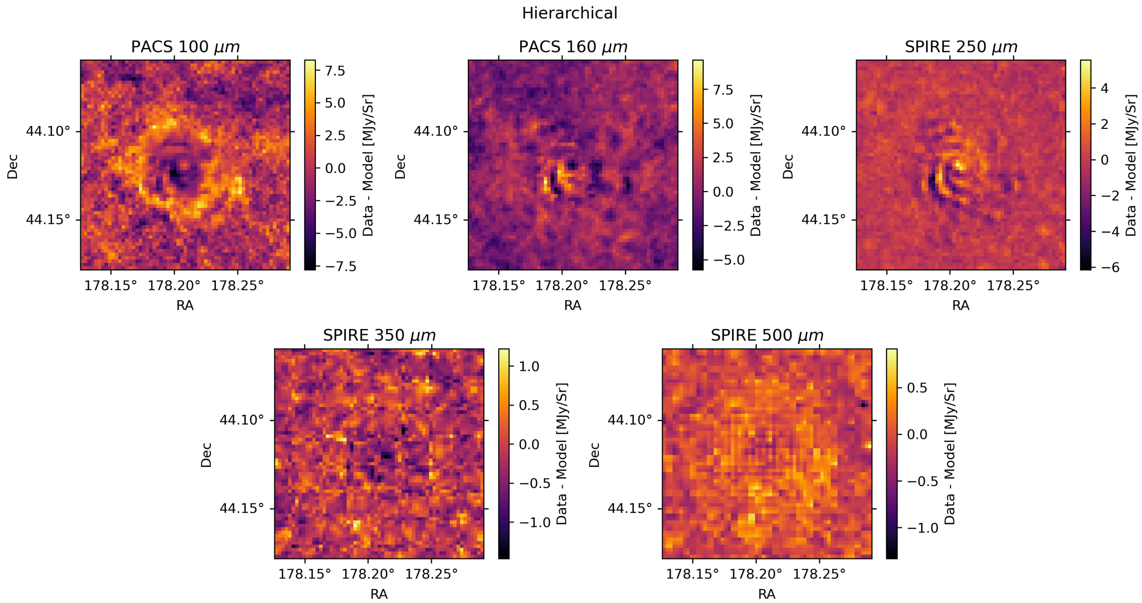

Figure 5.20. Residuals of the hierarchical fit to the unmasked 61×61 map of NGC 3938. Residuals are approximately at the 3-7% level.

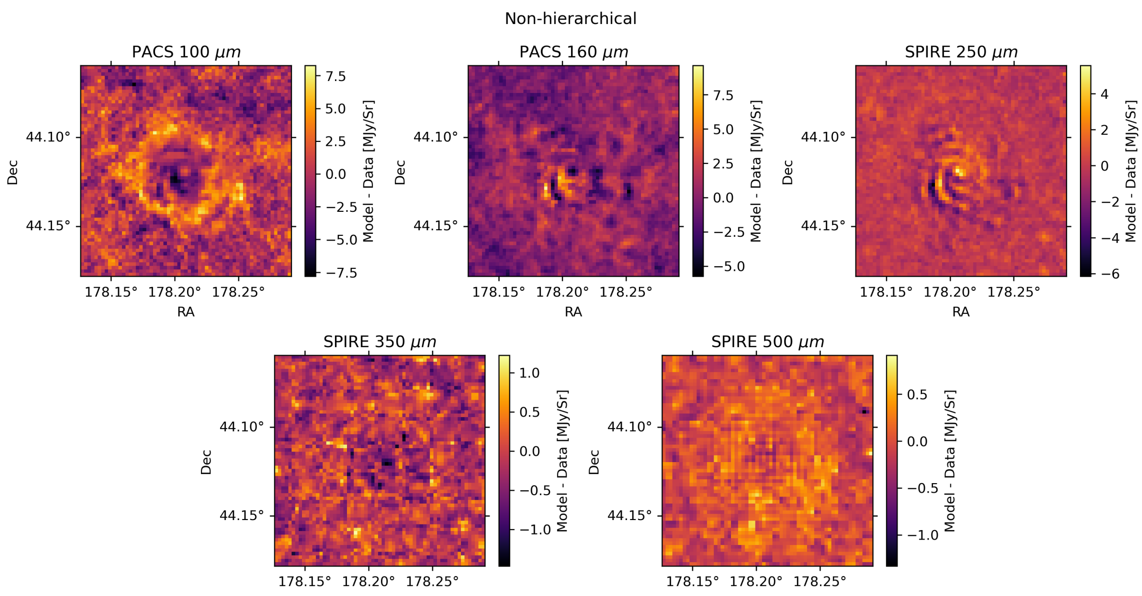

Figure 5.21. Same as Figure 5.20, but for the non-hierarchical case. Residuals are approximately 3-6% of the map fluxes.

Figure 5.22 shows the posterior and joint posterior distributions for every pixel for the hierarchical and non-hierarchical cases. All posteriors show signs of bimodality, which is expected due to the background pixels. The sharp lower limits on $\log(\Sigma_{\rm dust})$ and T and the upper limit on β in the hierarchical posteriors are not due to prior limits but are the result of the narrow distribution of the background. Overall, the hierarchical posteriors are more constrained, recovering a systematically higher $\log(\Sigma_{\rm dust})$ value and narrower ranges for T and β. There are strong correlations among the hierarchical model parameters. The correlation coefficients derived from the hyperparameters are $(\rho_{\log(\Sigma_{\rm dust}),\ln(T)},\ \rho_{\log(\Sigma_{\rm dust}),\beta},\ \rho_{\ln(T),\beta})$ = (0.98, -0.97, -0.95), implying the parameters are almost completely correlated. However, these correlations are not physical and instead are indicative of what is derived when fitting an image with constant $\log(\Sigma_{\rm dust})$, T, and β values. The more numerous background pixels, which show up as the long tails characterized by little scatter in the joint posterior plots, largely override any intrinsic physical correlations that may exist and introduce this unphysical relationship. This effect is amplified by the inclusion of regularization in the fitting routine, which enforces a smoother parameter distribution, with the likelihoods of source pixels being influenced by adjacent off-source pixels.

The non-hierarchical joint posterior scatter plots show considerably more scatter relative to the hierarchical fit, somewhat obscuring any correlations. The positive $\log(\Sigma_{\rm dust})$-T correlation of the hierarchical fit is reproduced, and an anti-correlation in T-β is also visible. Without the hyperparameters, these will not be subject to the same background-induced correlation as the hierarchical result, which implies that the $\log(\Sigma_{\rm dust})$-T correlation may be real or induced by the fitting methodology. The negative T-β relationship is the noise-induced anti-correlation typical of non-hierarchical MBB fits.

The fits are repeated for a smaller 31×31 image centered on the same region to reduce the number of background pixels. The parameter maps are shown in Fig-

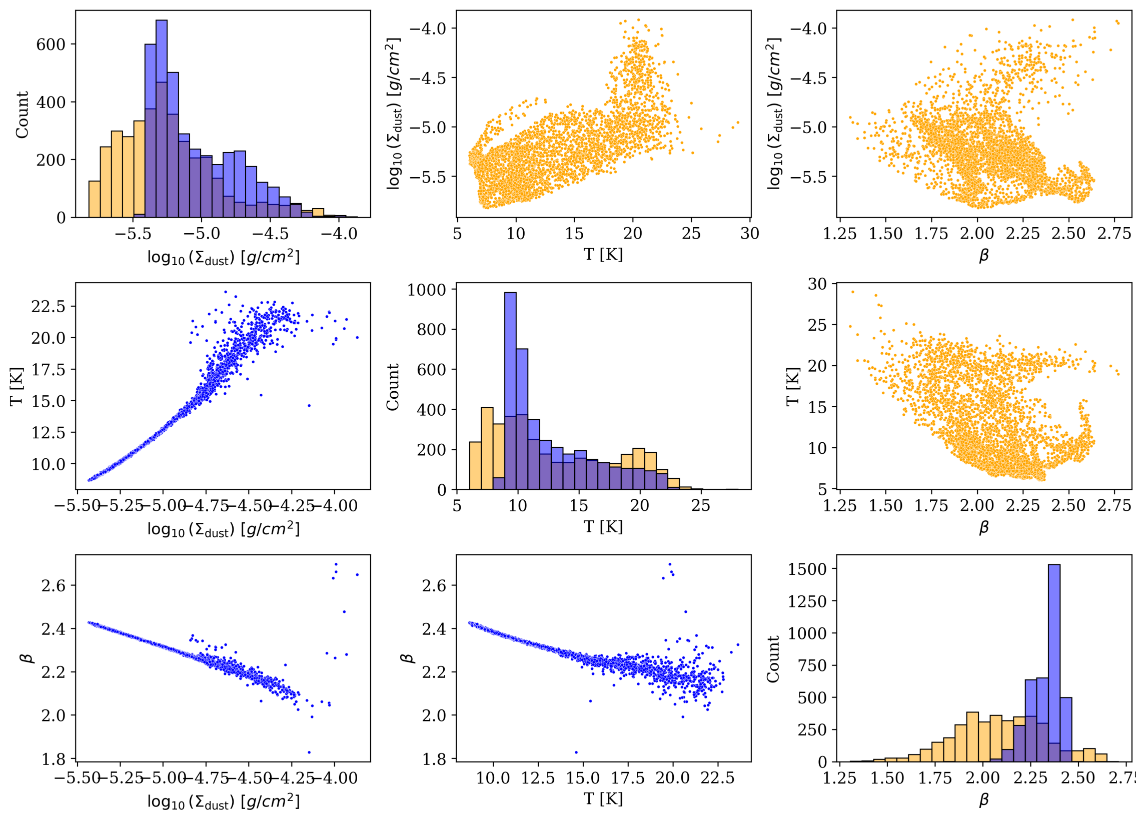

Figure 5.22. Both individual and joint posterior plots from the hierarchical (blue points) and non-hierarchical (orange points) MBB fits of the unmasked 61×61 maps of NGC 3938.

ure 5.23. Residuals are again on the order of 6% and 5% for both fits, respectively. Compared to the 61×61 image fit, the hierarchical and non-hierarchical fits are more similar due to the restriction to higher S/N pixels, with the largest deviations occurring near the edges of the galaxy. Most pixels have S/N > 50 across all bands. The strong anti-correlation near the center is now also reduced.

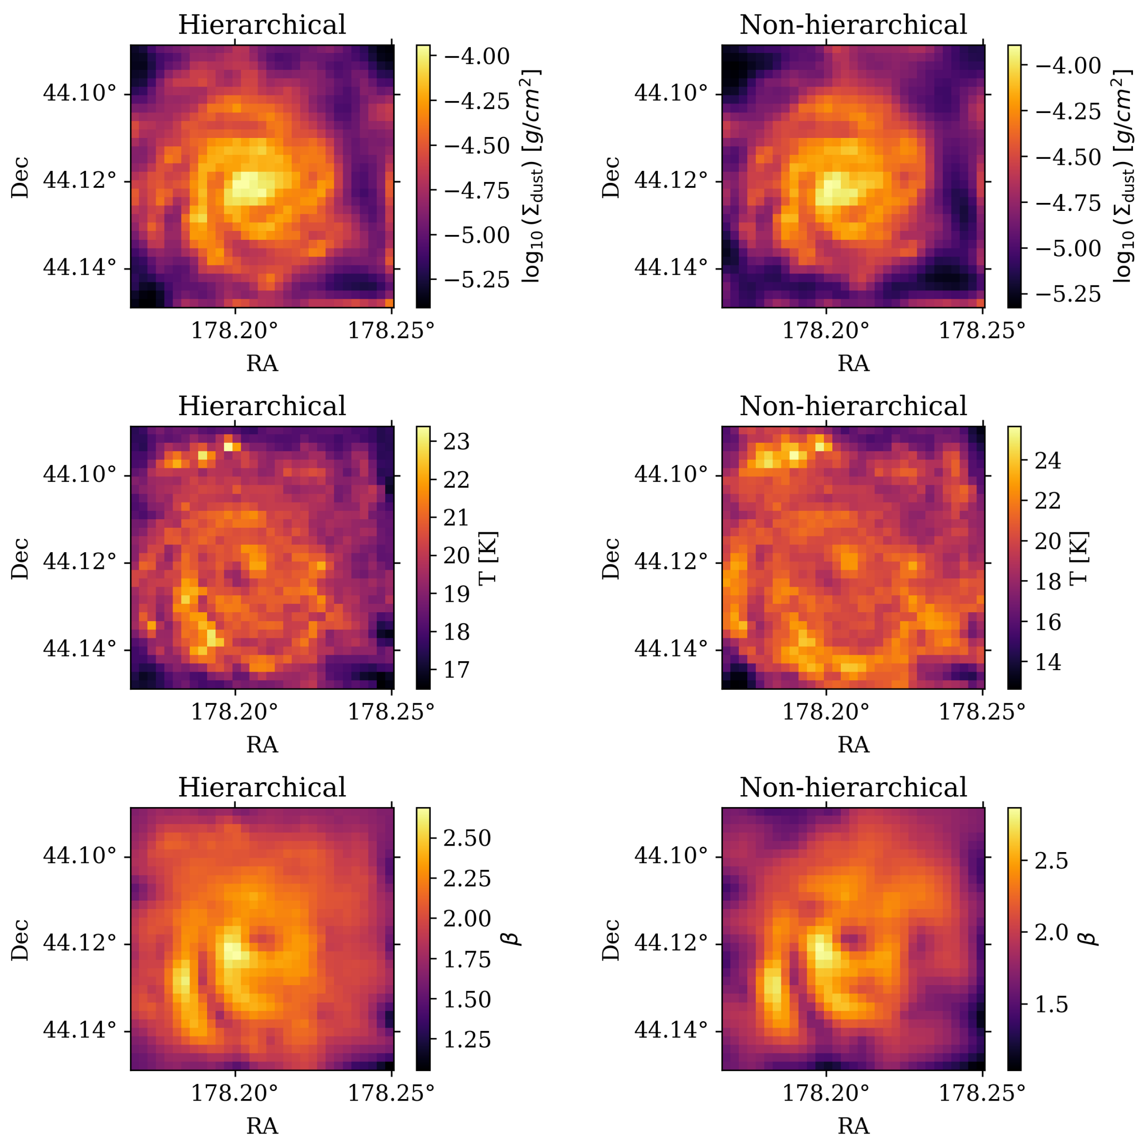

Figure 5.23. $\log(\Sigma_{dust})$, T, and β maps from hierarchical (left column) and non-hierarchical (right column) fits to the unmasked 31×31 maps of NGC 3938. Values used for each pixel are the median of the final 10^3 samples of each parameter.

The increased similarity in the parameter maps is mirrored in the posterior distributions as illustrated in Figure 5.24. There is little to no bimodality observed in the posterior distributions, and the two fitting cases have similar shapes and mean values. The β distribution shows the most significant difference between the hierarchical and non-hierarchical fitting methods, with the latter being more evenly distributed over its range. There is reduced scatter in the hierarchical T and β posteriors, with most of the reduction in the T distribution. The correlations among the parameters are significantly different from those of the 61×61 image for the hierarchical model. The correlation coefficients are $(\rho_{\log(\Sigma_{\rm dust}),\ln(T)},\ \rho_{\log(\Sigma_{\rm dust}),\beta},\ \rho_{\ln(T),\beta}) = (0.67,$ $0.68,\ 0.77)$. There is some reduced scatter towards lower $\log(\Sigma_{\rm dust})$ originating from the remaining low S/N pixels in the maps.

Determining whether the correlations are indeed intrinsic and physically meaningful is challenging. Nearly identical correlations are derived for an even smaller 21×21 image consisting of the brightest pixels only. At the resolutions of the *Herschel* bands, each pixel will contain a range of physical conditions, particularly for the non-diffuse ISM regions, and will not be characterized by a single dust temperature. The positive T-β correlation is not due to the mathematical degeneracy of T and β, which would introduce a negative relationship. Furthermore, it is unlikely to be the result of excessive smoothing from the Bayesian regularization procedure, which has a similar effect to the inclusion of background pixels and will, therefore, also result in an anti-correlation. The non-hierarchical T-β distribution is mostly flat but characterized by a large amount of scatter. The pixels below T=17.5 K and β=2 are near the edge of the map where some background remains and do not imply a positive correlation in conjunction with the rest of the points. As discussed in Section 5, grain growth in dense environments can result in a negative correlation between $\log(\Sigma_{\rm dust})$-β and a positive one for T-β. The hierarchical fits results are in agreement with this model, but a positive $\log(\Sigma_{\rm dust})$-β relation is recovered contrary

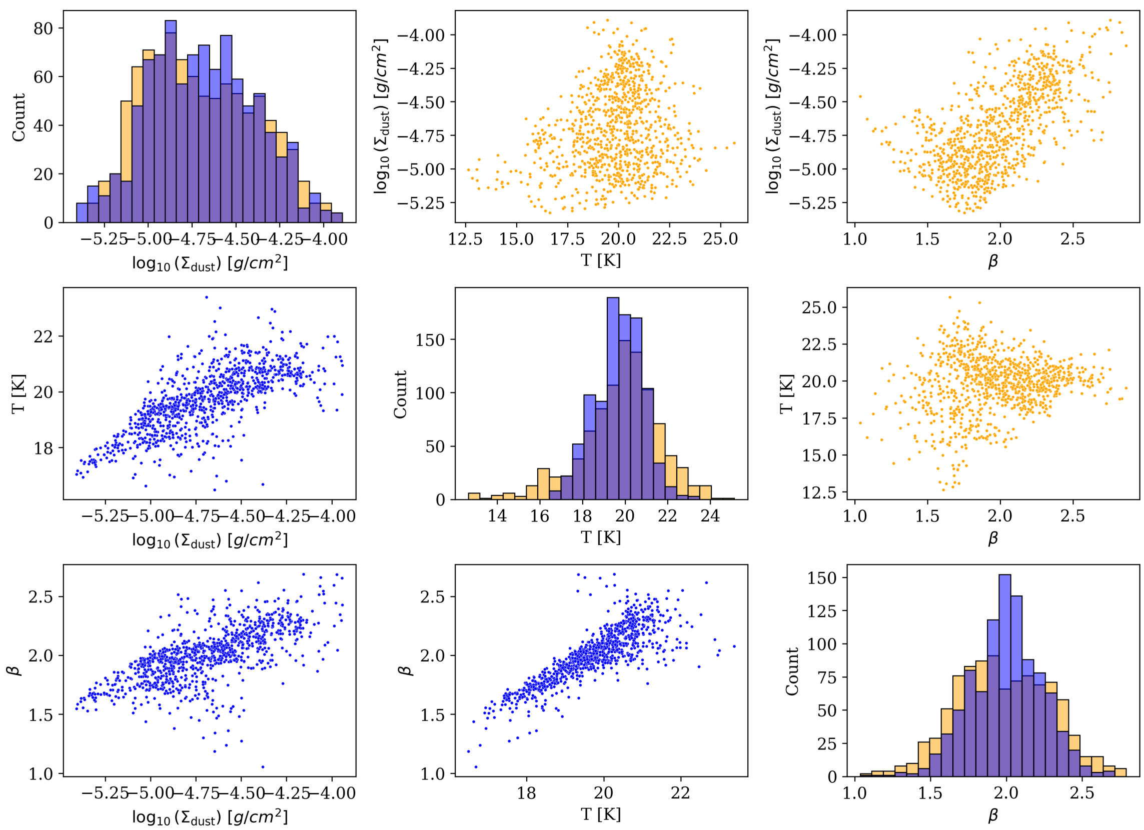

Figure 5.24. One-dimensional and joint posterior plots from the hierarchical (blue points) and non-hierarchical (orange points) MBB fits of the unmasked 61×61 maps of NGC 3938.

to what is expected. However, a negative T-β relation has also been predicted based on laboratory analogs and dust models involving mantle accretion. The negative T-β correlation found from the resolved study of the CMZ in Tang et al. 2021 that uses the same forward fitting and hierarchical Bayesian model implies that the result here may not be strictly model-dependent, and may at least partially reflect the source properties and data characteristics. Tang et al. 2021 also reported an increase of β in dense environments, but an anti-correlation between $\Sigma_{\rm dust}$ and T. However, the CMZ is resolved at a much higher resolution and will not suffer from resolution-dependent effects and the mixing of physical conditions to the same level as a distant galaxy.

In Figure 5.26, the 61×61 map is re-fit but with a pixel mask (Figure 5.25) applied to the pixels, which determines which pixels are used in the hyperparameter and hyperprior calculation. The non-hierarchical fit is not shown as it is identical to that of Figure 5.19 due to its lack of hyperparameters. A modest S/N value of 3 was selected for the lower limit, with only pixels above this threshold being utilized in the hyperparameter MCMC sampling. Two cases were considered, one with the calculated hyperparameter probabilities still applied to the posterior of the masked pixels and one where they were not. The former can lead to discontinuities at the boundary of the mask due to the Bayesian regularization, which will smooth the parameters on either side of the mask. A more robust smoothing algorithm is required to account for this. The resulting parameter maps of either case are similar to those of the 31×31 case, particularly for β, where the anti-correlation in the bulge of the unmasked 61×61 map is not present. There are some deviations between the two masking approaches with the T and β in edge pixels in the case where the hyperparameters are not applied to the masked regions being lower than in the other case. Figure 5.27 shows the posterior and joint posterior plots for the unmasked pixels for each case. The distributions and correlations of the masked pixels are nearly identical for both cases but are slightly different from the 31×31 image with correlation coefficients of $(\rho_{\log(\Sigma_{\rm dust}),\ln(T)},\ \rho_{\log(\Sigma_{\rm dust}),\beta},\ \rho_{\ln(T),\beta}) = (0.51,\ 0.70,\ 0.57)$ and $(\rho_{\log(\Sigma_{\rm dust}),\ln(T)},\ \rho_{\log(\Sigma_{\rm dust}),\beta},\ \rho_{\ln(T),\beta}) = (0.56,\ 0.70,\ 0.60)$ for the included and excluded masked pixel cases, respectively.

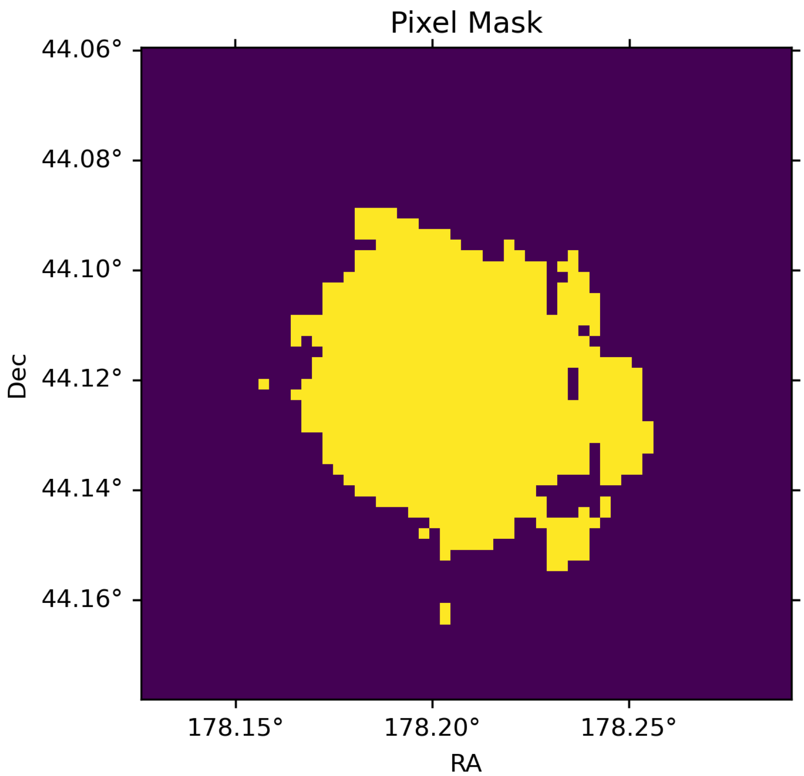

Figure 5.25. Pixel mask for deciding which pixels are entered into the hyperprior calculation. The S/N cutoff value was 3.

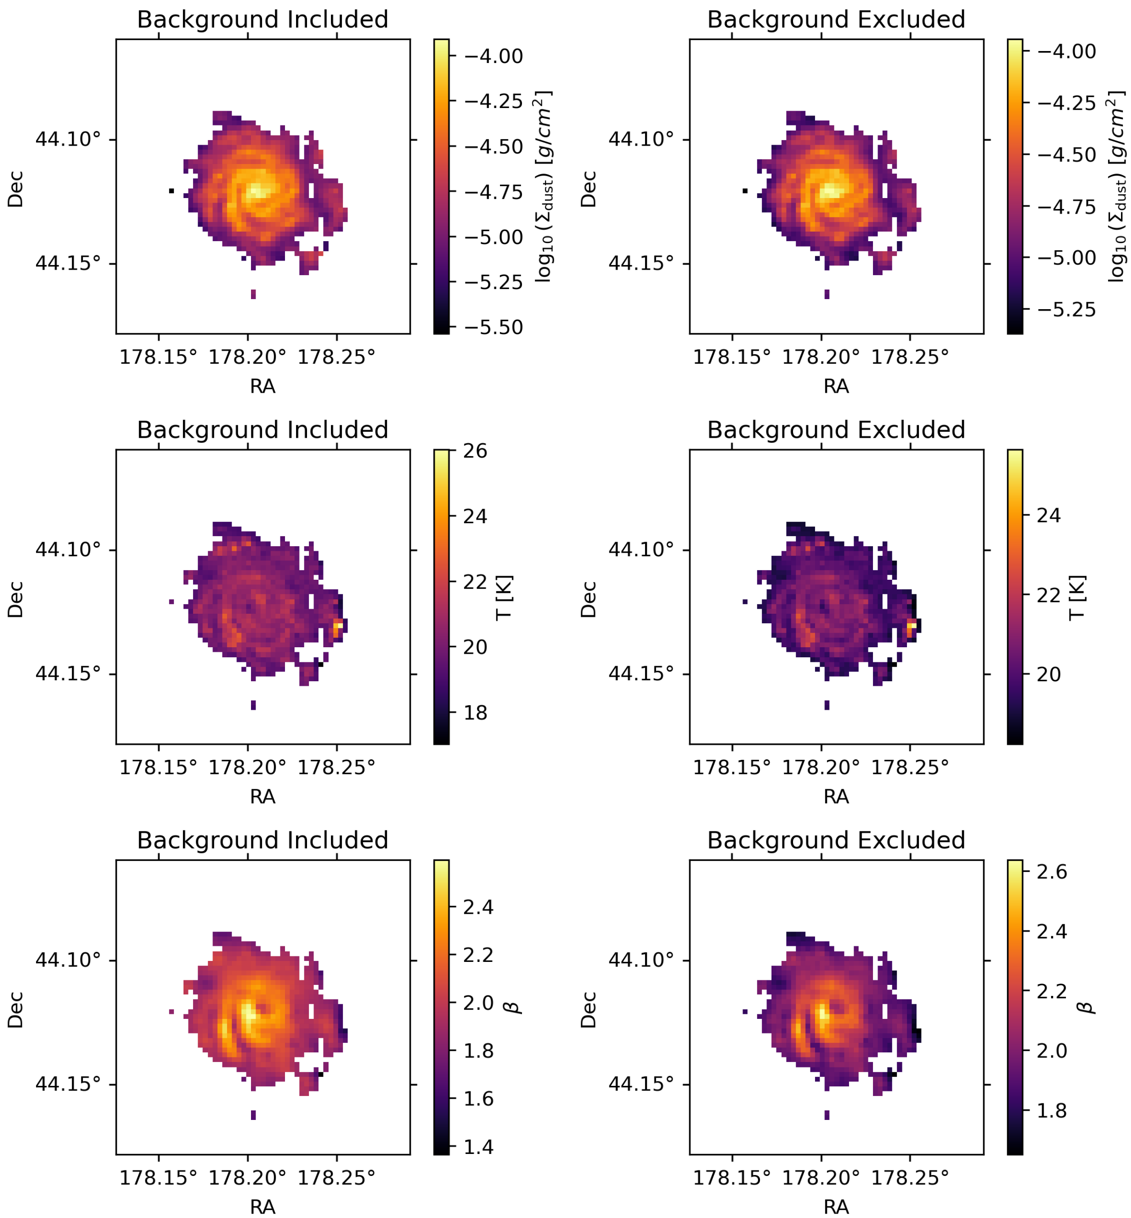

Figure 5.26. $\log(\Sigma_{dust})$, T, and β maps from hierarchical fits to the masked 61×61 maps of NGC 3938, where masked background pixels are not used in the hyperparameter sampling. Values used for each pixel are the median of the final 10^3 samples of each parameter. The left column shows maps where the hyperparameter probabilities were also applied to the background pixel posteriors. The right column is the case where the background pixel posteriors were sampled independently of the hyperparameters.

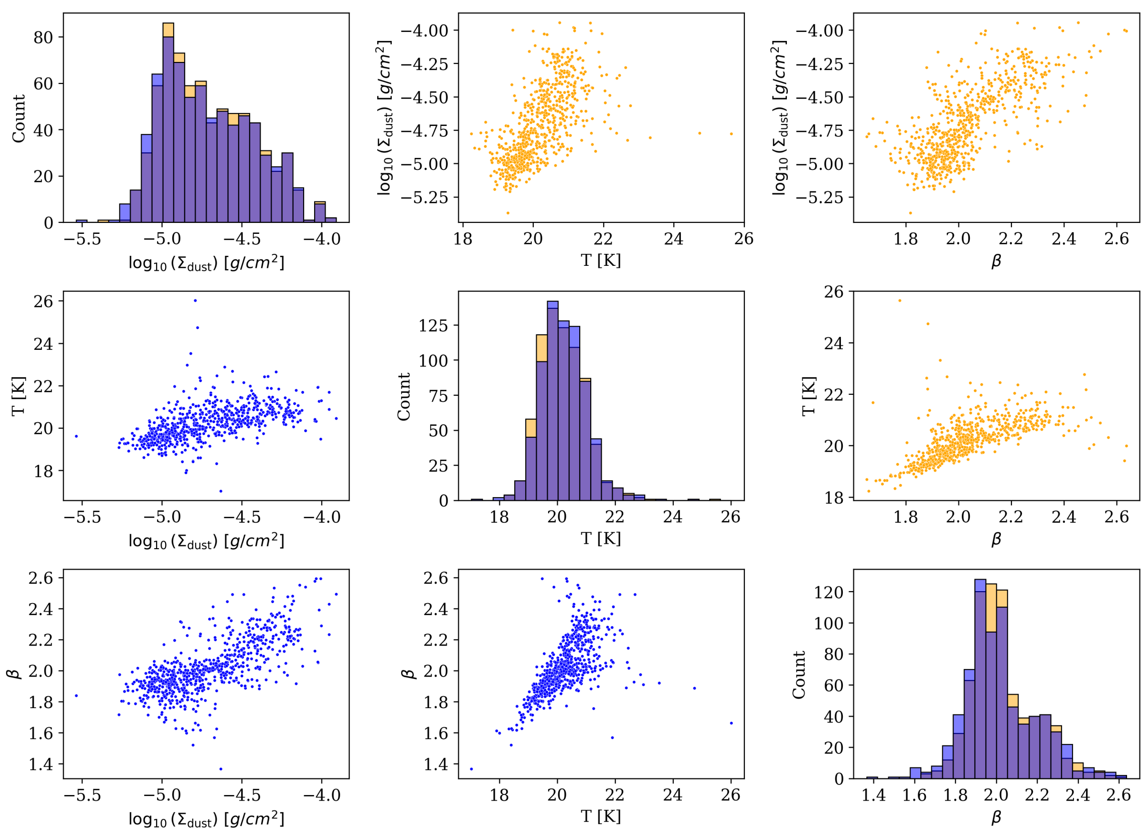

Figure 5.27. Posterior distributions for masked 61×61 fit. The blue (hyperprior applied to background pixels) and orange (hyperprior not applied to background pixels) points correspond to the left and right columns of Figure 5.26 respectively.

CHAPTER 6

CONCLUSION

TolTEC is a millimeter-wavelength imaging polarimeter currently installed on the 50-meter Large Millimeter Telescope that maps the sky simultaneously at 1.1, 1.4, and 2.0 mm (273, 214, and 150 GHz) with 7,718 dual-polarization Kinetic Inductance Detectors. It achieves diffraction-limited beamsizes between $5''$ and $10''$ with a full field of view of $4'$ in diameter. It will explore a wide variety of astrophysical phenomena, including galaxy clusters, high redshift luminous infrared galaxies, galactic star-forming molecular clouds, and interstellar dust (Chapter 1). The camera, telescope, and observing strategies were outlined in Chapter 2.

The primary result of my work related to the TolTEC project is `Citlali` the camera's end-to-end data reduction and mapmaking pipeline, which is described in Chapter 3. The main points of Citlali can be summarized as follows:

- It has been developed as an open-source, high-performance software framework written entirely in C++ to maximize performance on memory- and CPU-limited workstations.

- It uses a parallelized data streaming model to reduce the pipeline's overall memory footprint and to enable near real-time reduction of the data and is therefore scalable to take advantage of resources offered by high-performance computing clusters.

- `Citlali` transforms the raw time-ordered data into maps of the sky for all categories of TolTEC observations (beammapping, focus, astigmatism, pointing, and science) using a single highly configurable workflow.

- A range of timestream reduction algorithms have been implemented, including cosmic ray despiking, timestream filtering, and a principal component analysis-based atmospheric removal routine.

- Two built-in mapmaking algorithms are implemented: a fast naive mapmaker for quick-look data products and a more computationally expensive jinc filtered mapmaker optimized for removing high-frequency pixel-pixel noise within the maps. An iterative mapmaking routine to recover flux lost during atmospheric filtering has also been implemented.

- The pipeline can write partially reduced, flux-calibrated timestreams to disk to be used as inputs into standalone maximum likelihood mapmakers, such as TOAST3 and Minkasi.

Development of `Citlali` continues in parallel with the commissioning of the TolTEC camera. The inclusion of the half-wave plate signal to generate Stokes Q and U maps, improvements and further characterization of the iterative mapmaker, and additional map filtering routines are in active development. A built-in maximum likelihood mapmaker is also being implemented, which will be able to leverage C++ and Eigen to carry out the computationally expensive preconditioned conjugate gradient solution. Distributed parallelization over different computational nodes and GPU-based FFT calculations are being investigated.

As of the date of this writing, TolTEC has undergone two commissioning phases in 2022 and has resumed operations in March 2024. Chapter 4 presented preliminary maps and analyses of the on-sky data acquired in 2022 using both `Citlali` and Minkasi, including beam mapping observations of the radio quasar J1159+292 as well as maps of extended emission from the Crab Nebula and the Monoceros R2 Giant Molecular Cloud. Key takeaways include:

- On-sky detector positions, beam sizes, and flux calibration factors were derived from the beam map observation of J1159+292. Optical alignment of the 3 detector arrays was confirmed, and minimal radial dependence of per-detector characteristics was found.

- The ability to recover bright extended emission from the Crab Nebula and MonR2 with `Citlali`'s iterative Jinc mapmaker was demonstrated, with preliminary integrated flux measurements from the Crab Nebula being in good agreement with values from the literature. Maps from `Citlali` and Minkasi were compared, with maps produced by Minkasi being found to exhibit regions of negative flux surrounding the bright source. `Citlali` recovers more flux than Minkasi for the Crab Nebula but slightly underpredicts the flux of the bright regions of MonR2 relative to Minkasi.

- Preliminary Stokes Q and U maps were presented from observations of the Crab Nebula without the TolTEC half-wave plate installed. While polarization fractions were not derived, the morphology of the maps is consistent with existing NIKA 150 GHz maps.

The dust content of nearby galaxies will be mapped through observations with TolTEC. Commissioning observations of the star-forming spirals NGC 3938, NGC 4736, and M74 and the starburst irregular dwarf galaxies NGC 4449 and IC10 are currently planned. Chapter 5 introduced the hierarchical Bayesian MCMC fitting code I wrote to fit dust SEDs to modified blackbody and physically motivated dust models within each pixel of maps of local galaxies. Results from this chapter encompass:

- Simulated dust SEDs with varying noise levels were fit to single temperature modified blackbodies and the Astrodust+PAH dust model (Hensley and Draine 2023) using both a hierarchical and non-hierarchical model. The hierarchical fit recovers the intrinsic correlations among the input dust parameters and is characterized by reduced scatter and bias. The addition of observations at the TolTEC wavelengths to modified blackbody fits of simulated *Herschel* observations improves the recovery of the dust spectral emissivity index β.

- Integrated WISE, *Spitzer* and *Herschel* flux measurements from the DustPedia database of the face-on spiral galaxy NGC 3938 were fit to both dust models to test the fitting code on real data. The fitted dust mass and temperature are in agreement with values from the literature for modified blackbodies, THEMIS, Draine and Li 2007, and the updated Draine et al. 2014 dust models.

- A preliminary spatially resolved fit of Herschel PACS and SPIRE images of NGC 3938 to a single-temperature MBB model while also incorporating each band's resolution information was performed. Maps of Σ_{dust}, T, and β were presented, and the impact of background pixels on the correlations among dust parameters was investigated. A positive correlation among all dust parameters was identified in the fits using the hierarchical Bayesian model after background subtraction. The physical interpretation of these correlations remains uncertain. Expanded testing with simulated observations and realistic models of galaxies is required to fully explore the consequences of the choice of common pixel size, wavelength coverage, Bayesian regularization, and PSF inclusion on the recovered correlations.

- Additional work involving the hierarchical Bayesian fitting code is underway, including a spatially resolved fit of NGC 3938 to the Astrodust+PAH dust model, as well as a modified blackbody fit to SOFIA HAWC+, *Herschel* and TolTEC observations of the Monoceros R2 Giant Molecular Cloud. Further investigations of the ability to discern the signatures of the Sunyaev-Zeldovich effect and foreground dusty star-forming galaxies are also being conducted.

BIBLIOGRAPHY

Adam, R. et al. (Sept. 2014). "First observation of the thermal Sunyaev-Zel'dovich effect with kinetic inductance detectors". In: *Astronomy and Astrophysics* 569, A66, A66. DOI: 10.1051/0004-6361/201322902. arXiv: 1310.6237 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...569A..66A.

Adam, R. et al. (Jan. 2018). "The NIKA2 large-field-of-view millimetre continuum camera for the 30 m IRAM telescope". In: *Astronomy and Astrophysics* 609, A115, A115. DOI: 10.1051/0004-6361/201731503. arXiv: 1707.00908 [astro-ph.IM].

Agarwal, Sameer, Keir Mierle, and The Ceres Solver Team (Mar. 2022). *Ceres Solver*. http://ceres-solver.org. Version 2.1. A large scale non-linear optimization library. URL: https://github.com/ceres-solver/ceres-solver.

Agladze, N. I. et al. (May 1996). "Laboratory Results on Millimeter-Wave Absorption in Silicate Grain Materials at Cryogenic Temperatures". In: *Astrophysical Journal* 462, p. 1026. DOI: 10.1086/177217. URL: https://ui.adsabs.harvard.edu/abs/1996ApJ...462.1026A.

Aniano, G. et al. (Oct. 2011). "Common-Resolution Convolution Kernels for Space- and Ground-Based Telescopes". In: *PASP* 123.908, p. 1218. DOI: 10.1086/662219. arXiv: 1106.5065 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2011PASP..123.1218A.

Aniano, G. et al. (Sept. 2012). "Modeling Dust and Starlight in Galaxies Observed by Spitzer and Herschel: NGC 628 and NGC 6946". In: *Astrophysical Journal* 756.2, 138, p. 138. DOI: 10.1088/0004-637X/756/2/138. arXiv: 1207.4186 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2012ApJ...756..138A.

Aniano, G. et al. (Feb. 2020). "Modeling Dust and Starlight in Galaxies Observed by Spitzer and Herschel: The KINGFISH Sample". In: *Astrophysical Journal* 889.2, 150, p. 150. DOI: 10.3847/1538-4357/ab5fdb. arXiv: 1912.04914 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2020ApJ...889..150A.

Arendt, R. G. et al. (June 2011). “The Radio-2 mm Spectral Index of the Crab Nebula Measured with Gismo”. In: *Astrophysical Journal* 734.1, 54, p. 54. DOI: 10.1088/0004-637X/734/1/54. arXiv: 1103.6225 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2011ApJ...734...54A.

Arendt, Richard G. et al. (May 2014). “Interstellar and Ejecta Dust in the Cas A Supernova Remnant”. In: *Astrophysical Journal* 786.1, 55, p. 55. DOI: 10.1088/0004-637X/786/1/55. arXiv: 1403.3008 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014ApJ...786...55A.

Asano, Ryosuke S. et al. (Mar. 2013). “Dust formation history of galaxies: A critical role of metallicity for the dust mass growth by accreting materials in the interstellar medium”. In: *Earth, Planets and Space* 65.3, pp. 213–222. DOI: 10.5047/eps.2012.04.014. arXiv: 1206.0817 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2013EP&S...65..213A.

Asplund, Martin et al. (Sept. 2009). “The Chemical Composition of the Sun”. In: *Annual Review of Astronomy and Astrophysics* 47.1, pp. 481–522. DOI: 10.1146/annurev.astro.46.060407.145222. arXiv: 0909.0948 [astro-ph.SR]. URL: https://ui.adsabs.harvard.edu/abs/2009ARA&A..47..481A.

Austermann, J. E. et al. (Nov. 2018). “Millimeter-Wave Polarimeters Using Kinetic Inductance Detectors for TolTEC and Beyond”. In: *Journal of Low Temperature Physics* 193.3-4, pp. 120–127. DOI: 10.1007/s10909-018-1949-5. arXiv: 1803.03280 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2018JLTP..193..120A.

Azimzadeh, Parsiad (June 2017). *mlinterp*. https://github.com/parsiad/mlinterp. Accessed: 2024-04-11.

Barlow, M. J. (May 1978). “The destruction and growth of dust grains in interstellar space - I. Destruction by sputtering.” In: *Monthly Notices of the Royal Astronomical Society* 183, pp. 367–395. DOI: 10.1093/mnras/183.3.367. URL: https://ui.adsabs.harvard.edu/abs/1978MNRAS.183..367B.

Barlow, M. J. et al. (July 2010). “A Herschel PACS and SPIRE study of the dust content of the Cassiopeia A supernova remnant”. In: *Astronomy and Astrophysics* 518, L138, p. L138. DOI: 10.1051/0004-6361/201014585. arXiv: 1005.2688 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...518L.138B.

Barnard, John, Robert McCulloch, and Xiao Li Meng (Oct. 2000). “Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage”. English (US). In: *Statistica Sinica* 10.4, pp. 1281–1311. ISSN: 1017-0405.

Benton, Steven James (Jan. 2015). “Mapping submillimetre polarization with BLAST-Pol”. PhD thesis. University of Toronto, Canada. URL: https://ui.adsabs.harvard.edu/abs/2015PhDT........83B.

Bernard, Jean-Philippe et al. (Sept. 2008). “Spitzer Survey of the Large Magellanic Cloud, Surveying the Agents of a Galaxy’s Evolution (sage). IV. Dust Properties in the Interstellar Medium”. In: *The Astronomical Journal* 136.3, pp. 919–945. DOI: 10.1088/0004-6256/136/3/919. URL: https://ui.adsabs.harvard.edu/abs/2008AJ....136..919B.

Bertoldi, F. et al. (July 2003). “Dust emission from the most distant quasars”. In: *Astronomy and Astrophysics* 406, pp. L55–L58. DOI: 10.1051/0004-6361:20030710. arXiv: astro-ph/0305116 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003A&A...406L..55B.

Bevan, Antonia, M. J. Barlow, and D. Milisavljevic (Mar. 2017). “Dust masses for SN 1980K, SN1993J and Cassiopeia A from red-blue emission line asymmetries”. In: *Monthly Notices of the Royal Astronomical Society* 465.4, pp. 4044–4056. DOI: 10.1093/mnras/stw2985. arXiv: 1611.05006 [astro-ph.SR]. URL: https://ui.adsabs.harvard.edu/abs/2017MNRAS.465.4044B.

Bianchi, S. (Apr. 2013). “Vindicating single-T modified blackbody fits to Herschel SEDs”. In: *Astronomy and Astrophysics* 552, A89, A89. DOI: 10.1051/0004-6361/201220866. arXiv: 1302.5699 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2013A&A...552A..89B.

Bianchi, S. et al. (Dec. 2018). “Fraction of bolometric luminosity absorbed by dust in DustPedia galaxies”. In: *Astronomy and Astrophysics* 620, A112, A112. DOI: 10.1051/0004-6361/201833699. arXiv: 1810.01208 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2018A&A...620A.112B.

Bianchi, Simone and Raffaella Schneider (July 2007). “Dust formation and survival in supernova ejecta”. In: *Monthly Notices of the Royal Astronomical Society* 378.3, pp. 973–982. DOI: 10.1111/j.1365-2966.2007.11829.x. arXiv: 0704.0586 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2007MNRAS.378..973B.

Blain, A. W., V. E. Barnard, and S. C. Chapman (Jan. 2003). “Submillimetre and far-infrared spectral energy distributions of galaxies: the luminosity-temperature relation and consequences for photometric redshifts”. In: *Monthly Notices of the Royal Astronomical Society* 338.3, pp. 733–744. DOI: 10.1046/j.1365-8711.2003.06086.x. arXiv: astro-ph/0209450 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003MNRAS.338..733B.

Bocchio, M. et al. (Mar. 2016). "Dust grains from the heart of supernovae". In: *Astronomy and Astrophysics* 587, A157, A157. DOI: 10.1051/0004-6361/201527432. arXiv: 1601.06770 [astro-ph.HE]. URL: https://ui.adsabs.harvard.edu/abs/2016A&A...587A.157B.

Bocchio, Marco, Anthony P. Jones, and Jonathan D. Slavin (Oct. 2014). "A re-evaluation of dust processing in supernova shock waves". In: *Astronomy and Astrophysics* 570, A32, A32. DOI: 10.1051/0004-6361/201424368. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...570A..32B.

Bohlin, R. C., B. D. Savage, and J. F. Drake (Aug. 1978). "A survey of interstellar H I from Lalpha absorption measurements. II." In: *Astrophysical Journal* 224, pp. 132–142. DOI: 10.1086/156357. URL: https://ui.adsabs.harvard.edu/abs/1978ApJ...224..132B.

Boogert, A. C. Adwin, Perry A. Gerakines, and Douglas C. B. Whittet (Aug. 2015). "Observations of the icy universe." In: *Annual Review of Astronomy and Astrophysics* 53, pp. 541–581. DOI: 10.1146/annurev-astro-082214-122348. arXiv: 1501.05317 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015ARA&A..53..541B.

Boquien, M. et al. (Feb. 2019). "CIGALE: a python Code Investigating GALaxy Emission". In: *Astronomy and Astrophysics* 622, A103, A103. DOI: 10.1051/0004-6361/201834156. arXiv: 1811.03094 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019A&A...622A.103B.

Boselli, A. et al. (July 2010). "FIR colours and SEDs of nearby galaxies observed with Herschel". In: *Astronomy and Astrophysics* 518, L61, p. L61. DOI: 10.1051/0004-6361/201014534. arXiv: 1005.1537 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...518L..61B.

Boselli, A. et al. (Apr. 2012). "Far-infrared colours of nearby late-type galaxies in the Herschel Reference Survey". In: *Astronomy and Astrophysics* 540, A54, A54. DOI: 10.1051/0004-6361/201118602. arXiv: 1201.2305 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2012A&A...540A..54B.

Bot, C. et al. (Nov. 2010). "Submillimeter to centimeter excess emission from the Magellanic Clouds. II. On the nature of the excess". In: *Astronomy and Astrophysics* 523, A20, A20. DOI: 10.1051/0004-6361/201014986. arXiv: 1008.2875 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...523A..20B.

Boudet, N. et al. (Nov. 2005). "Temperature Dependence of the Submillimeter Absorption Coefficient of Amorphous Silicate Grains". In: *Astrophysical Journal* 633.1, pp. 272–281. DOI: 10.1086/432966.

Breemen, J. M. van et al. (Feb. 2011). "The 9.7 and 18 μm silicate absorption profiles towards diffuse and molecular cloud lines-of-sight". In: *Astronomy and Astrophysics* 526, A152, A152. DOI: 10.1051/0004-6361/200811142. arXiv: 1012.1698 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2011A&A...526A.152V.

Bron, Emeric, Jacques Le Bourlot, and Franck Le Petit (Sept. 2014). "Surface chemistry in the interstellar medium. II. H_2 formation on dust with random temperature fluctuations". In: *Astronomy and Astrophysics* 569, A100, A100. DOI: 10.1051/0004-6361/201322101. arXiv: 1407.4473 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...569A.100B.

Bryan, Sean et al. (July 2018). "Optical design of the TolTEC millimeter-wave camera". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 10708. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 107080J, 107080J. DOI: 10.1117/12.2314130. arXiv: 1807.00097 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2018SPIE10708E..0JB.

Bühler, R. and R. Blandford (June 2014). "The surprising Crab pulsar and its nebula: a review". In: *Reports on Progress in Physics* 77.6, 066901, p. 066901. DOI: 10.1088/0034-4885/77/6/066901. arXiv: 1309.7046 [astro-ph.HE]. URL: https://ui.adsabs.harvard.edu/abs/2014RPPh...77f6901B.

Cairós, L. M. et al. (Sept. 2010). "Mapping the properties of blue compact dwarf galaxies: integral field spectroscopy with PMAS". In: *Astronomy and Astrophysics* 520, A90, A90. DOI: 10.1051/0004-6361/201014004. arXiv: 1004.2858 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...520A..90C.

Calzetti, D. et al. (Jan. 2018). "Spatially Resolved Dust, Gas, and Star Formation in the Dwarf Magellanic Irregular NGC 4449". In: *Astrophysical Journal* 852.2, 106, p. 106. DOI: 10.3847/1538-4357/aaa1e2. arXiv: 1712.06233 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2018ApJ...852..106C.

Carlstrom, J. E. et al. (May 2011). "The 10 Meter South Pole Telescope". In: *PASP* 123.903, p. 568. DOI: 10.1086/659879. arXiv: 0907.4445 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2011PASP..123..568C.

Carlstrom, John et al. (Sept. 2019). "CMB-S4". In: *Bulletin of the American Astronomical Society*. Vol. 51, 209, p. 209. DOI: 10.48550/arXiv.1908.01062. arXiv: 1908.01062 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2019BAAS...51g.209C.

Carlstrom, John E., Gilbert P. Holder, and Erik D. Reese (Jan. 2002). "Cosmology with the Sunyaev-Zel'dovich Effect". In: *Annual Review of Astronomy and As-*

trophysics 40, pp. 643–680. DOI: 10.1146/annurev.astro.40.060401.093803. arXiv: astro-ph/0208192 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2002ARA&A..40..643C.

Catalano, A. et al. (Sept. 2014). "Performance and calibration of the NIKA camera at the IRAM 30 m telescope". In: *Astronomy and Astrophysics* 569, A9, A9. DOI: 10.1051/0004-6361/201423557. arXiv: 1402.0260 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...569A...9C.

Cattaneo, A. et al. (Sept. 2008). "Downsizing by shutdown in red galaxies". In: *Monthly Notices of the Royal Astronomical Society* 389.2, pp. 567–584. DOI: 10.1111/j.1365-2966.2008.13562.x. arXiv: 0801.1673 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2008MNRAS.389..567C.

Chang, Zhengxue et al. (July 2021). "Investigating Cold Dust Properties of 12 Nearby Dwarf Irregular Galaxies by Hierarchical Bayesian Spectral Energy Distribution Fitting". In: *Astrophysical Journal* 915.1, 51, p. 51. DOI: 10.3847/1538-4357/abfe67. URL: https://ui.adsabs.harvard.edu/abs/2021ApJ...915...51C.

Chapin, Edward L. et al. (Apr. 2013). "SCUBA-2: iterative map-making with the Sub-Millimetre User Reduction Facility". In: *Monthly Notices of the Royal Astronomical Society* 430.4, pp. 2545–2573. DOI: 10.1093/mnras/stt052. arXiv: 1301.3652 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2013MNRAS.430.2545C.

Chapman, Scott C. et al. (Aug. 2022). "CCAT-prime: the 850 GHz camera for prime-cam on FYST". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 12190. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1219005, p. 1219005. DOI: 10.1117/12.2630628. arXiv: 2208.10634 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2022SPIE12190E..05C.

Chastenet, Jérémy et al. (May 2021). "Benchmarking Dust Emission Models in M101". In: *Astrophysical Journal* 912.2, 103, p. 103. DOI: 10.3847/1538-4357/abe942. arXiv: 2101.09236 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2021ApJ...912..103C.

Chiar, J. E. et al. (July 2000). "The Composition and Distribution of Dust along the Line of Sight toward the Galactic Center". In: *Astrophysical Journal* 537.2, pp. 749–762. DOI: 10.1086/309047. arXiv: astro-ph/0002421 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2000ApJ...537..749C.

Choi, S. K. et al. (Mar. 2020). "Sensitivity of the Prime-Cam Instrument on the CCAT-Prime Telescope". In: *Journal of Low Temperature Physics* 199.3-4, pp. 1089–

1097. DOI: 10.1007/s10909-020-02428-z. arXiv: 1908.10451 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2020JLTP..199.1089C.

Ciesla, L. et al. (May 2014). “Dust spectral energy distributions of nearby galaxies: an insight from the Herschel Reference Survey”. In: *Astronomy and Astrophysics* 565, A128, A128. DOI: 10.1051/0004-6361/201323248. arXiv: 1402.3597 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...565A.128C.

Clark, C. J. R. et al. (Jan. 2018). “DustPedia: Multiwavelength photometry and imagery of 875 nearby galaxies in 42 ultraviolet-microwave bands”. en. In: *Astronomy & Astrophysics, Volume 609, id.A37, ¡NUMPAGES¿30¡/NUMPAGES¿ pp.* 609, A37. ISSN: 0004-6361. DOI: 10.1051/0004-6361/201731419. URL: https://ui.adsabs.harvard.edu/abs/2018A%5C%26A...609A..37C/abstract (visited on 11/11/2022).

Clements, David L. (Oct. 2017). “An introduction to the Planck mission”. In: *Contemporary Physics* 58.4, pp. 331–348. DOI: 10.1080/00107514.2017.1362139. arXiv: 1707.09220 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2017ConPh..58..331C.

Collaboration, Astropy et al. (Oct. 2013). “Astropy: A community Python package for astronomy”. In: *Astronomy and Astrophysics* 558, A33, A33. DOI: 10.1051/0004-6361/201322068. arXiv: 1307.6212 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2013A&A...558A..33A.

Collaboration, Astropy et al. (Sept. 2018). “The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package”. In: *The Astronomical Journal* 156.3, 123, p. 123. DOI: 10.3847/1538-3881/aabc4f. arXiv: 1801.02634 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2018AJ....156..123A.

Collaboration, Astropy et al. (Aug. 2022). “The Astropy Project: Sustaining and Growing a Community-oriented Open-source Project and the Latest Major Release (v5.0) of the Core Package”. In: *Astrophysical Journal* 935.2, 167, p. 167. DOI: 10.3847/1538-4357/ac7c74. arXiv: 2206.14220 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2022ApJ...935..167A.

Cormier, D. et al. (Apr. 2014). “The molecular gas reservoir of 6 low-metallicity galaxies from the Herschel Dwarf Galaxy Survey. A ground-based follow-up survey of CO(1-0), CO(2-1), and CO(3-2)”. In: *Astronomy and Astrophysics* 564, A121, A121. DOI: 10.1051/0004-6361/201322096. arXiv: 1401.0563 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...564A.121C.

Coupeaud, A. et al. (Nov. 2011). “Low-temperature FIR and submillimetre mass absorption coefficient of interstellar silicate dust analogues”. In: *Astronomy and Astrophysics* 535, A124, A124. DOI: 10.1051/0004-6361/201116945. arXiv: 1109.2758 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2011A&A...535A.124C.

Crill, B. P. et al. (July 2008). “SPIDER: a balloon-borne large-scale CMB polarimeter”. In: *Space Telescopes and Instrumentation 2008: Optical, Infrared, and Millimeter*. Ed. by Jr. Oschmann Jacobus M., Mattheus W. M. de Graauw, and Howard A. MacEwen. Vol. 7010. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 70102P, 70102P. DOI: 10.1117/12.787446. arXiv: 0807.1548 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2008SPIE.7010E..2PC.

Cunha, Elisabete da, Stéphane Charlot, and David Elbaz (Aug. 2008). “A simple model to interpret the ultraviolet, optical and infrared emission from galaxies”. In: *Monthly Notices of the Royal Astronomical Society* 388.4, pp. 1595–1617. DOI: 10.1111/j.1365-2966.2008.13535.x. arXiv: 0806.1020 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2008MNRAS.388.1595D.

Cutri, R. M. et al. (Feb. 2021). *VizieR Online Data Catalog: AllWISE Data Release (Cutri+ 2013)*. URL: https://ui.adsabs.harvard.edu/abs/2014yCat.2328....0C.

Dale, D. A. et al. (Jan. 2012). “Herschel Far-infrared and Submillimeter Photometry for the KINGFISH Sample of nearby Galaxies”. In: *Astrophysical Journal* 745.1, 95, p. 95. DOI: 10.1088/0004-637X/745/1/95. arXiv: 1112.1093 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2012ApJ...745...95D.

Dale, Daniel A. et al. (Mar. 2001). “The Infrared Spectral Energy Distribution of Normal Star-forming Galaxies”. In: *Astrophysical Journal* 549.1, pp. 215–227. DOI: 10.1086/319077.

Das, H. K., N. V. Voshchinnikov, and V. B. Il’in (May 2010). “Interstellar extinction and polarization - a spheroidal dust grain approach perspective”. In: *Monthly Notices of the Royal Astronomical Society* 404.1, pp. 265–274. DOI: 10.1111/j.1365-2966.2010.16281.x. arXiv: 1001.0655 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2010MNRAS.404..265D.

Davies, J. I. et al. (Apr. 2017). “DustPedia: A Definitive Study of Cosmic Dust in the Local Universe”. In: *PASP* 129.974, p. 044102. DOI: 10.1088/1538-3873/129/974/044102. arXiv: 1609.06138 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017PASP..129d4102D.

Day, Peter K. et al. (Oct. 2003). “A broadband superconducting detector suitable for use in large arrays”. In: *Nature* 425.6960, pp. 817–821. DOI: 10.1038/nature02037. URL: https://ui.adsabs.harvard.edu/abs/2003Natur.425..817D.

De Looze, I. et al. (Mar. 2017). “The dust mass in Cassiopeia A from a spatially resolved Herschel analysis”. In: *Monthly Notices of the Royal Astronomical Society* 465.3, pp. 3309–3342. DOI: 10.1093/mnras/stw2837. arXiv: 1611.00774 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017MNRAS.465.3309D.

De Looze, I. et al. (Sept. 2019). “The dust content of the Crab Nebula”. In: *Monthly Notices of the Royal Astronomical Society* 488.1, pp. 164–182. DOI: 10.1093/mnras/stz1533. arXiv: 1906.02203 [astro-ph.HE]. URL: https://ui.adsabs.harvard.edu/abs/2019MNRAS.488..164D.

De Vis, P. et al. (Oct. 2017). “Using dust, gas and stellar mass-selected samples to probe dust sources and sinks in low-metallicity galaxies”. In: *Monthly Notices of the Royal Astronomical Society* 471.2, pp. 1743–1765. DOI: 10.1093/mnras/stx981. arXiv: 1705.02340 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017MNRAS.471.1743D.

De Vis, P. et al. (Mar. 2019). “A systematic metallicity study of DustPedia galaxies reveals evolution in the dust-to-metal ratios”. In: *Astronomy and Astrophysics* 623, A5, A5. DOI: 10.1051/0004-6361/201834444. arXiv: 1901.09040 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019A&A...623A...5D.

Dell’Agli, F. et al. (Dec. 2015). “AGB stars in the SMC: evolution and dust properties based on Spitzer observations”. In: *Monthly Notices of the Royal Astronomical Society* 454.4, pp. 4235–4249. DOI: 10.1093/mnras/stv2298. arXiv: 1510.01230 [astro-ph.SR]. URL: https://ui.adsabs.harvard.edu/abs/2015MNRAS.454.4235D.

Dempsey, J. T. et al. (Apr. 2013). “SCUBA-2: on-sky calibration using submillimetre standard sources”. In: *Monthly Notices of the Royal Astronomical Society* 430.4, pp. 2534–2544. DOI: 10.1093/mnras/stt090. arXiv: 1301.3773 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2013MNRAS.430.2534D.

Demyk, K. (Jan. 2011). “Interstellar dust within the life cycle of the interstellar medium”. In: *European Physical Journal Web of Conferences*. Vol. 18. European Physical Journal Web of Conferences, 03001, p. 03001. DOI: 10.1051/epjconf/20111803001. URL: https://ui.adsabs.harvard.edu/abs/2011EPJWC..1803001D.

Demyk, K. et al. (Apr. 2017a). “Low temperature MIR to submillimeter mass absorption coefficient of interstellar dust analogues. I. Mg-rich glassy silicates”.

In: *Astronomy and Astrophysics* 600, A123, A123. DOI: 10.1051/0004-6361/201629711. arXiv: 1701.07225 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017A&A...600A.123D.

Demyk, K. et al. (Oct. 2017b). "Low-temperature MIR to submillimeter mass absorption coefficient of interstellar dust analogues. II. Mg and Fe-rich amorphous silicates". In: *Astronomy and Astrophysics* 606, A50, A50. DOI: 10.1051/0004-6361/201730944. arXiv: 1706.09801 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017A&A...606A..50D.

DeNigris, N. S. et al. (Jan. 2020). "Developing a Large-Scale Cryogenic System for the Simultaneous Operation of Three Detector Focal Planes in TolTEC, A New Multichroic Imaging Polarimeter". In: *Journal of Low Temperature Physics* 199.3-4, pp. 789–797. DOI: 10.1007/s10909-019-02319-y. URL: https://ui.adsabs.harvard.edu/abs/2020JLTP..199..789D.

DeNigris, Nat S (2024). "TolTEC: A New Multichroic Imaging Polarimeter for the Large Millimeter Telescope". PhD thesis. DOI: 10.7275/36464651.

Doyle, S. et al. (Apr. 2008). "Lumped Element Kinetic Inductance Detectors". In: *Journal of Low Temperature Physics* 151.1-2, pp. 530–536. DOI: 10.1007/s10909-007-9685-2. URL: https://ui.adsabs.harvard.edu/abs/2008JLTP..151..530D.

Draine, B. T. (Jan. 2003). "Interstellar Dust Grains". In: *Annual Review of Astronomy and Astrophysics* 41, pp. 241–289. DOI: 10.1146/annurev.astro.41.011802.094840. arXiv: astro-ph/0304489 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003ARA%5C&A..41..241D.

Draine, B. T. and Brandon Hensley (Sept. 2012). "The Submillimeter and Millimeter Excess of the Small Magellanic Cloud: Magnetic Dipole Emission from Magnetic Nanoparticles?" In: *Astrophysical Journal* 757.1, 103, p. 103. DOI: 10.1088/0004-637X/757/1/103. arXiv: 1205.6810 [astro-ph.GA].

Draine, B. T. and Brandon S. Hensley (Mar. 2021). "The Dielectric Function of "Astrodust" and Predictions for Polarization in the 3.4 and 10 μm Features". In: *Astrophysical Journal* 909.1, 94, p. 94. DOI: 10.3847/1538-4357/abd6c6. arXiv: 2009.11314 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2021ApJ...909...94D.

Draine, B. T. and Aigen Li (Mar. 2007). "Infrared Emission from Interstellar Dust. IV. The Silicate-Graphite-PAH Model in the Post-Spitzer Era". In: *Astrophysical Journal* 657.2, pp. 810–837. DOI: 10.1086/511055. arXiv: astro-ph/0608003 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2007ApJ...657..810D.

Draine, B. T. et al. (July 2007). "Dust Masses, PAH Abundances, and Starlight Intensities in the SINGS Galaxy Sample". In: *Astrophysical Journal* 663.2, pp. 866–894. DOI: 10.1086/518306. arXiv: astro-ph/0703213 [astro-ph].

Draine, B. T. et al. (Jan. 2014). "Andromeda's Dust". In: *Astrophysical Journal* 780.2, 172, p. 172. DOI: 10.1088/0004-637X/780/2/172. arXiv: 1306.2304 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2014ApJ...780..172D.

Draine, Bruce T. and Aurélien A. Fraisse (May 2009). "Polarized Far-Infrared and Submillimeter Emission from Interstellar Dust". In: *Astrophysical Journal* 696.1, pp. 1–11. DOI: 10.1088/0004-637X/696/1/1. arXiv: 0809.2094 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2009ApJ...696....1D.

Drinkwater, M. J. et al. (Dec. 2000). "Compact Stellar Systems in the Fornax Cluster: Super-massive Star Clusters or Extremely Compact Dwarf Galaxies?" In: *Publications of the Astronomical Society of Australia* 17.3, pp. 227–233. DOI: 10.1071/AS00034. arXiv: astro-ph/0002003 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2000PASA...17..227D.

Driver, Simon P. et al. (Feb. 2016). "Galaxy And Mass Assembly (GAMA): Panchromatic Data Release (far-UV-far-IR) and the low-z energy budget". In: *Monthly Notices of the Royal Astronomical Society* 455.4, pp. 3911–3942. DOI: 10.1093/mnras/stv2505. arXiv: 1508.02076 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2016MNRAS.455.3911D.

Dumke, M., M. Krause, and R. Wielebinski (Feb. 2004). "Cold dust in a selected sample of nearby galaxies. I. The interacting galaxy NGC 4631". In: *Astronomy and Astrophysics* 414, pp. 475–486. DOI: 10.1051/0004-6361:20031636. arXiv: astro-ph/0311591 [astro-ph].

Do-Duy, Tho et al. (Apr. 2020). "Crystalline silicate absorption at 11.1 μm: ubiquitous and abundant in embedded YSOs and the interstellar medium". In: *Monthly Notices of the Royal Astronomical Society* 493.3, pp. 4463–4517. DOI: 10.1093/mnras/staa396. URL: https://ui.adsabs.harvard.edu/abs/2020MNRAS.493.4463D.

Dwek, E. and J. M. Scalo (July 1980). "The evolution of refractory interstellar grains in the solar neighborhood". In: *Astrophysical Journal* 239, pp. 193–211. DOI: 10.1086/158100. URL: https://ui.adsabs.harvard.edu/abs/1980ApJ...239..193D.

Dzib, Sergio A. et al. (Aug. 2016). "VLBA Determination of the Distance to Nearby Star-forming Regions. VII. Monoceros R2". In: *Astrophysical Journal* 826.2, 201,

p. 201. DOI: 10.3847/0004-637X/826/2/201. arXiv: 1606.01757 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2016ApJ...826..201D.

Ejlali, G. et al. (July 2022). "Dust Emission in Galaxies at Millimeter Wavelengths: Cooling of star forming regions in NGC6946". In: *mm Universe @ NIKA2 - Observing the mm Universe with the NIKA2 Camera*. Vol. 257. European Physical Journal Web of Conferences, 00016, p. 00016. DOI: 10.1051/epjconf/202225700016. arXiv: 2111.03844 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2022EPJWC.25700016E.

Ejlali, G. et al. (Oct. 2023). "Constraining Millimeter Dust Emission in Nearby Galaxies with NIKA2: the case of NGC2146 and NGC2976". In: *arXiv e-prints*, arXiv:2310.03428, arXiv:2310.03428. DOI: 10.48550/arXiv.2310.03428. arXiv: 2310.03428 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2023arXiv231003428E.

Engelbracht, C. W. et al. (July 2005). "Metallicity Effects on Mid-Infrared Colors and the 8 μm PAH Emission in Galaxies". In: *Astrophysical Journal Letters* 628.1, pp. L29–L32. DOI: 10.1086/432613. arXiv: astro-ph/0506214 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2005ApJ...628L..29E.

Ercolano, B., M. J. Barlow, and B. E. K. Sugerman (Mar. 2007). "Dust yields in clumpy supernova shells: SN 1987A revisited". In: *Monthly Notices of the Royal Astronomical Society* 375.3, pp. 753–763. DOI: 10.1111/j.1365-2966.2006.11336.x. arXiv: astro-ph/0611719 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2007MNRAS.375..753E.

Ferrusca, D. and J. Contreras R. (July 2014). "Weather monitor station and 225 GHz radiometer system installed at Sierra Negra: the Large Millimeter Telescope site". In: *Ground-based and Airborne Instrumentation for Astronomy V*. Ed. by Suzanne K. Ramsay, Ian S. McLean, and Hideki Takami. Vol. 9147. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 914730, p. 914730. DOI: 10.1117/12.2055005. URL: https://ui.adsabs.harvard.edu/abs/2014SPIE.9147E..30F.

Frenk, Carlos S. et al. (Apr. 1988). "The Formation of Dark Halos in a Universe Dominated by Cold Dark Matter". In: *Astrophysical Journal* 327, p. 507. DOI: 10.1086/166213. URL: https://ui.adsabs.harvard.edu/abs/1988ApJ...327..507F.

Frigo, Matteo and Steven G. Johnson (2005). "The Design and Implementation of FFTW3". In: *Proceedings of the IEEE* 93.2. Special issue on "Program Generation, Optimization, and Platform Adaptation", pp. 216–231.

Galametz, M. et al. (Dec. 2009). "Probing the dust properties of galaxies up to submillimetre wavelengths. I. The spectral energy distribution of dwarf galaxies using LABOCA". In: *Astronomy and Astrophysics* 508.2, pp. 645–664. DOI: 10.1051/0004-6361/200912963. arXiv: 0910.0043 [astro-ph.CO].

Galametz, M. et al. (Aug. 2011). "Probing the dust properties of galaxies up to submillimetre wavelengths. II. Dust-to-gas mass ratio trends with metallicity and the submm excess in dwarf galaxies". In: *Astronomy and Astrophysics* 532, A56, A56. DOI: 10.1051/0004-6361/201014904. arXiv: 1104.0827 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2011A%5C&A...532A..56G.

Galametz, M. et al. (May 2013). "The thermal dust emission in N158-N159-N160 (LMC) star-forming complex mapped by Spitzer, Herschel and LABOCA". In: *Monthly Notices of the Royal Astronomical Society* 431.2, pp. 1596–1617. DOI: 10.1093/mnras/stt280. arXiv: 1302.2825 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2013MNRAS.431.1596G.

Galitzki, Nicholas et al. (July 2018). "The Simons Observatory: instrument overview". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 10708. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1070804, p. 1070804. DOI: 10.1117/12.2312985. arXiv: 1808.04493 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2018SPIE10708E..04G.

Galliano, F. et al. (Aug. 2003). "ISM properties in low-metallicity environments. II. The dust spectral energy distribution of NGC 1569". In: *Astronomy and Astrophysics* 407, pp. 159–176. DOI: 10.1051/0004-6361:20030814. arXiv: astro-ph/0306192 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003A%5C&A...407..159G.

Galliano, F. et al. (May 2005). "ISM properties in low-metallicity environments. III. The spectral energy distributions of II Zw 40, He 2-10 and NGC 1140". In: *Astronomy and Astrophysics* 434.3, pp. 867–885. DOI: 10.1051/0004-6361:20042369. arXiv: astro-ph/0501632 [astro-ph].

Galliano, F. et al. (Dec. 2011). "Non-standard grain properties, dark gas reservoir, and extended submillimeter excess, probed by Herschel in the Large Magellanic Cloud". In: *Astronomy and Astrophysics* 536, A88, A88. DOI: 10.1051/0004-6361/201117952. arXiv: 1110.1260 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2011A%5C&A...536A..88G.

Galliano, Frédéric (May 2018). "A dust spectral energy distribution model with hierarchical Bayesian inference - I. Formalism and benchmarking". In: *Monthly Notices of the Royal Astronomical Society* 476.2, pp. 1445–1469. DOI: 10.1093/mnras/

sty189. arXiv: 1801.06660 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2018MNRAS.476.1445G.

Galliano, Frédéric (Feb. 2022). “A Nearby Galaxy Perspective on Interstellar Dust Properties and their Evolution”. In: *Habilitation Thesis*, p. 1. arXiv: 2202.01868 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2022HabT.........1G.

Galliano, Frédéric, Eli Dwek, and Pierre Chanial (Jan. 2008). “Stellar Evolutionary Effects on the Abundances of Polycyclic Aromatic Hydrocarbons and Supernova-Condensed Dust in Galaxies”. In: *Astrophysical Journal* 672.1, pp. 214–243. DOI: 10.1086/523621. arXiv: 0708.0790 [astro-ph].

Galliano, Frédéric, Maud Galametz, and Anthony P. Jones (Sept. 2018). “The Interstellar Dust Properties of Nearby Galaxies”. In: *Annual Review of Astronomy and Astrophysics* 56, pp. 673–713. DOI: 10.1146/annurev-astro-081817-051900. arXiv: 1711.07434 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2018ARA%5C&A..56..673G.

Galliano, Frédéric et al. (May 2021). “A nearby galaxy perspective on dust evolution. Scaling relations and constraints on the dust build-up in galaxies with the DustPedia and DGS samples”. In: *Astronomy and Astrophysics* 649, A18, A18. DOI: 10.1051/0004-6361/202039701. arXiv: 2101.00456 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2021A%5C&A...649A..18G.

Gomez, H. L. et al. (Nov. 2012). “A Cool Dust Factory in the Crab Nebula: A Herschel Study of the Filaments”. In: *Astrophysical Journal* 760.1, 96, p. 96. DOI: 10.1088/0004-637X/760/1/96. arXiv: 1209.5677 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2012ApJ...760...96G.

Gordon, Karl D. et al. (Dec. 2014). “Dust and Gas in the Magellanic Clouds from the HERITAGE Herschel Key Project. I. Dust Properties and Insights into the Origin of the Submillimeter Excess Emission”. In: *Astrophysical Journal* 797.2, 85, p. 85. DOI: 10.1088/0004-637X/797/2/85. arXiv: 1406.6066 [astro-ph.GA].

Gould, Robert J. and Edwin E. Salpeter (Aug. 1963). “The Interstellar Abundance of the Hydrogen Molecule. I. Basic Processes.” In: *Astrophysical Journal* 138, p. 393. DOI: 10.1086/147654. URL: https://ui.adsabs.harvard.edu/abs/1963ApJ...138..393G.

Grossi, M. et al. (Feb. 2015). “The Herschel Virgo Cluster Survey. XVIII. Star-forming dwarf galaxies in a cluster environment”. In: *Astronomy and Astrophysics* 574, A126, A126. DOI: 10.1051/0004-6361/201424866. arXiv: 1411.3960 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015A&A...574A.126G.

Guennebaud, Gaël, Benoît Jacob, et al. (Oct. 2022). *Eigen.* http://eigen.tuxfamily.org. Version 3.4.0. Eigen is a high-level C++ library for linear algebra, matrix and vector operations, numerical solvers, and related algorithms. URL: http://eigen.tuxfamily.org/index.php?title=Main_Page.

Guillet, V. et al. (Feb. 2018). "Dust models compatible with Planck intensity and polarization data in translucent lines of sight". In: *Astronomy and Astrophysics* 610, A16, A16. DOI: 10.1051/0004-6361/201630271. arXiv: 1710.04598 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2018A&A...610A..16G.

Helton, A. L. and SOFIA Science Team (Jan. 2013). "SOFIA - Stratospheric Observatory for Infrared Astronomy". In: *Proceedings of The Life Cycle of Dust in the Universe: Observations*, 59, p. 59. DOI: 10.22323/1.207.0059. URL: https://ui.adsabs.harvard.edu/abs/2013lcdu.confE..59H.

Henkel, Christian, Leslie K. Hunt, and Yuri I. Izotov (Jan. 2022). "The Interstellar Medium of Dwarf Galaxies". In: *Galaxies* 10.1, 11, p. 11. DOI: 10.3390/galaxies10010011. arXiv: 2202.08231 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2022Galax..10...11H.

Hensley, Brandon S. and B. T. Draine (Jan. 2021). "Observational Constraints on the Physical Properties of Interstellar Dust in the Post-Planck Era". In: *Astrophysical Journal* 906.2, 73, p. 73. DOI: 10.3847/1538-4357/abc8f1. arXiv: 2009.00018 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2021ApJ...906...73H.

Hensley, Brandon S. and B. T. Draine (May 2023). "The Astrodust+PAH Model: A Unified Description of the Extinction, Emission, and Polarization from Dust in the Diffuse Interstellar Medium". In: *Astrophysical Journal* 948.1, 55, p. 55. DOI: 10.3847/1538-4357/acc4c2. arXiv: 2208.12365 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2023ApJ...948...55H.

Hensley, Brandon S., Claire E. Murray, and Mark Dodici (Apr. 2022). "Polycyclic Aromatic Hydrocarbons, Anomalous Microwave Emission, and their Connection to the Cold Neutral Medium". In: *Astrophysical Journal* 929.1, 23, p. 23. DOI: 10.3847/1538-4357/ac5cbd. arXiv: 2111.03067 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2022ApJ...929...23H.

Herbig, G. H. (Jan. 1995). "The Diffuse Interstellar Bands". In: *Annual Review of Astronomy and Astrophysics* 33, pp. 19–74. DOI: 10.1146/annurev.aa.33.090195.000315. URL: https://ui.adsabs.harvard.edu/abs/1995ARA&A..33...19H.

Hermelo, I. et al. (May 2016). "Millimeter and submillimeter excess emission in M 33 revealed by Planck and LABOCA". In: *Astronomy and Astrophysics* 590, A56, A56. DOI: 10.1051/0004-6361/201525816. arXiv: 1603.02125 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2016A&A...590A..56H.

Hildebrand, R. H. (Sept. 1983). "The determination of cloud masses and dust characteristics from submillimetre thermal emission." In: *Quarterly Journal of the Royal Astronomical Society* 24, pp. 267–282. URL: https://ui.adsabs.harvard.edu/abs/1983QJRAS..24..267H.

Hilker, M., L. Infante, and T. Richtler (July 1999). "The central region of the Fornax cluster. III. Dwarf galaxies, globular clusters, and cD halo - are there interrelations?" In: *Astronomy and Astrophysics Supplement* 138, pp. 55–70. DOI: 10.1051/aas:1999495. arXiv: astro-ph/9905112 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/1999A&AS..138...55H.

Hirashita, Hiroyuki and Vladimir B. Il'in (Feb. 2022). "Evolution of dust grain size distribution and grain porosity in galaxies". In: *Monthly Notices of the Royal Astronomical Society* 509.4, pp. 5771–5789. DOI: 10.1093/mnras/stab3455. arXiv: 2111.12868 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2022MNRAS.509.5771H.

Ho, Paul T. P., James M. Moran, and Kwok Yung Lo (Nov. 2004). "The Submillimeter Array". In: *Astrophysical Journal Letters* 616.1, pp. L1–L6. DOI: 10.1086/423245. arXiv: astro-ph/0406352 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2004ApJ...616L...1H.

Hoffman, Matthew D. and Andrew Gelman (Nov. 2011). "The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo". In: *arXiv e-prints*, arXiv:1111.4246, arXiv:1111.4246. DOI: 10.48550/arXiv.1111.4246. arXiv: 1111.4246 [stat.CO]. URL: https://ui.adsabs.harvard.edu/abs/2011arXiv1111.4246H.

Holland, W. S. et al. (Mar. 1999). "SCUBA: a common-user submillimetre camera operating on the James Clerk Maxwell Telescope". In: *Monthly Notices of the Royal Astronomical Society* 303.4, pp. 659–672. DOI: 10.1046/j.1365-8711.1999.02111.x. arXiv: astro-ph/9809122 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/1999MNRAS.303..659H.

Holland, W. S. et al. (Apr. 2013). "SCUBA-2: the 10 000 pixel bolometer camera on the James Clerk Maxwell Telescope". In: *Monthly Notices of the Royal Astronomical Society* 430.4, pp. 2513–2533. DOI: 10.1093/mnras/sts612. arXiv: 1301.3650 [astro-ph.IM].

Hughes, David H. et al. (July 2010). “The Large Millimeter Telescope”. In: *Ground-based and Airborne Telescopes III*. Ed. by Larry M. Stepp, Roberto Gilmozzi, and Helen J. Hall. Vol. 7733. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 773312, p. 773312. DOI: 10.1117/12.857974. URL: https://ui.adsabs.harvard.edu/abs/2010SPIE.7733E..12H.

Hughes, David H. et al. (Dec. 2020). “The Large Millimeter Telescope (LMT) Alfonso Serrano: current status and telescope performance”. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Vol. 11445. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1144522, p. 1144522. DOI: 10.1117/12.2561893.

Hunt, L. K. et al. (Apr. 2015). “Cool dust heating and temperature mixing in nearby star-forming galaxies”. In: *Astronomy and Astrophysics* 576, A33, A33. DOI: 10.1051/0004-6361/201424734. arXiv: 1409.5916 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015A%5C&A...576A..33H.

Hunt, Leslie K. et al. (Mar. 2010). “The Spitzer View of Low-Metallicity Star Formation. III. Fine-Structure Lines, Aromatic Features, and Molecules”. In: *Astrophysical Journal* 712.1, pp. 164–187. DOI: 10.1088/0004-637X/712/1/164. arXiv: 1002.0991 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010ApJ...712..164H.

Hunter, Deidre A., Bruce G. Elmegreen, and Suzanne C. Madden (Feb. 2024). “The Interstellar Medium in Dwarf Irregular Galaxies”. In: *arXiv e-prints*, arXiv:2402.17004, arXiv:2402.17004. DOI: 10.48550/arXiv.2402.17004. arXiv: 2402.17004 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2024arXiv240217004H.

Israel, F. P. et al. (Sept. 2010). “submillimeter to centimeter excess emission from the Magellanic Clouds. I. Global spectral energy distribution”. In: *Astronomy and Astrophysics* 519, A67, A67. DOI: 10.1051/0004-6361/201014073. arXiv: 1006.2232 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...519A..67I.

Jenkins, Edward B. (Aug. 2009). “A Unified Representation of Gas-Phase Element Depletions in the Interstellar Medium”. In: *Astrophysical Journal* 700.2, pp. 1299–1348. DOI: 10.1088/0004-637X/700/2/1299. arXiv: 0905.3173 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2009ApJ...700.1299J.

Jenness, T. and J. F. Lightfoot (Jan. 1998). “Reducing SCUBA Data at the James Clerk Maxwell Telescope”. In: *Astronomical Data Analysis Software and Systems VII*. Ed. by Rudolf Albrecht, Richard N. Hook, and Howard A. Bushouse. Vol. 145. Astronomical Society of the Pacific Conference Series, p. 216. URL: https://ui.adsabs.harvard.edu/abs/1998ASPC..145..216J.

Jenness, Tim and John Lightfoot (Mar. 2014). *SURF: Submm User Reduction Facility*. Astrophysics Source Code Library, record ascl:1403.008. URL: https://ui.adsabs.harvard.edu/abs/2014ascl.soft03008J.

Jenness, Tim et al. (Oct. 2013). *SMURF: SubMillimeter User Reduction Facility*. Astrophysics Source Code Library, record ascl:1310.007. URL: https://ui.adsabs.harvard.edu/abs/2013ascl.soft10007J.

Jiang, Sally D. and Lynne A. Hillenbrand (Mar. 2024). "The Emerging Stellar Complex in Mon R2: Membership and Optical Variability Classification". In: *arXiv e-prints*, arXiv:2403.03843, arXiv:2403.03843. DOI: 10.48550/arXiv.2403.03843. arXiv: 2403.03843 [astro-ph.SR]. URL: https://ui.adsabs.harvard.edu/abs/2024arXiv240303843J.

Jones, A. P. et al. (Oct. 2013). "The evolution of amorphous hydrocarbons in the ISM: dust modelling from a new vantage point". In: *Astronomy and Astrophysics* 558, A62, A62. DOI: 10.1051/0004-6361/201321686. arXiv: 1411.6293 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2013A&A...558A..62J.

Jones, A. P. et al. (June 2017). "The global dust modelling framework THEMIS". In: *Astronomy and Astrophysics* 602, A46, A46. DOI: 10.1051/0004-6361/201630225. arXiv: 1703.00775 [astro-ph.GA].

Juvela, M. et al. (Dec. 2015). "Galactic cold cores. VI. Dust opacity spectral index". In: *Astronomy and Astrophysics* 584, A94, A94. DOI: 10.1051/0004-6361/201425269. arXiv: 1509.08023 [astro-ph.GA].

Kaiser, James F. (Apr. 1974). "Nonrecursive digital filter design using the I_0-sinh window function". In: *Proceedings of the 1974 IEEE International Symposium on Circuits and Systems*. San Francisco, CA, USA.

Keene, J. et al. (Aug. 1980). "Far-infrared observations of the globule B335". In: *Astrophysical Journal Letters* 240, pp. L43–L46. DOI: 10.1086/183320. URL: https://ui.adsabs.harvard.edu/abs/1980ApJ...240L..43K.

Kelly, Brandon C. et al. (June 2012). "Dust Spectral Energy Distributions in the Era of Herschel and Planck: A Hierarchical Bayesian-fitting Technique". In: *Astrophysical Journal* 752.1, 55, p. 55. DOI: 10.1088/0004-637X/752/1/55. arXiv: 1203.0025 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2012ApJ...752...55K.

Kennicutt Robert C., Jr. (May 1998). "The Global Schmidt Law in Star-forming Galaxies". In: *Astrophysical Journal* 498.2, pp. 541–552. DOI: 10.1086/305588. arXiv: astro-ph/9712213 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/1998ApJ...498..541K.

Kennicutt Robert C., Jr. et al. (Oct. 2009). “Dust-corrected Star Formation Rates of Galaxies. I. Combinations of Hα and Infrared Tracers”. In: *Astrophysical Journal* 703.2, pp. 1672–1695. DOI: 10.1088/0004-637X/703/2/1672. arXiv: 0908.0203 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2009ApJ...703.1672K.

Kennicutt, R. C. et al. (Dec. 2011). “KINGFISH—Key Insights on Nearby Galaxies: A Far-Infrared Survey with Herschel: Survey Description and Image Atlas”. In: *PASP* 123.910, p. 1347. DOI: 10.1086/663818. arXiv: 1111.4438 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2011PASP..123.1347K.

Kessler, M. F. et al. (Nov. 1996). “The Infrared Space Observatory (ISO) mission.” In: *Astronomy and Astrophysics* 315.2, pp. L27–L31. URL: https://ui.adsabs.harvard.edu/abs/1996A&A...315L..27K.

Kim, Sang-Hee and P. G. Martin (May 1995). “The Size Distribution of Interstellar Dust Particles as Determined from Polarization: Spheroids”. In: *Astrophysical Journal* 444, p. 293. DOI: 10.1086/175604. URL: https://ui.adsabs.harvard.edu/abs/1995ApJ...444..293K.

Kisner, Theodore et al. (Oct. 2021). *hpc4cmb/toast: Update Pybind11*. URL: https://zenodo.org/record/5559597 (visited on 08/10/2023).

Klessen, Ralf S. and Simon C. O. Glover (Jan. 2016). “Physical Processes in the Interstellar Medium”. In: *Saas-Fee Advanced Course*. Ed. by Yves Revaz et al. Vol. 43. Saas-Fee Advanced Course, p. 85. DOI: 10.1007/978-3-662-47890-5_2. arXiv: 1412.5182 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2016SAAS...43...85K.

Köhler, M., A. Jones, and N. Ysard (May 2014). “A hidden reservoir of Fe/FeS in interstellar silicates?” In: *Astronomy and Astrophysics* 565, L9, p. L9. DOI: 10.1051/0004-6361/201423985. arXiv: 1405.4208 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...565L...9K.

Köhler, M., N. Ysard, and A. P. Jones (July 2015). “Dust evolution in the transition towards the denser ISM: impact on dust temperature, opacity, and spectral index”. In: *Astronomy and Astrophysics* 579, A15, A15. DOI: 10.1051/0004-6361/201525646. arXiv: 1506.01533 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015A&A...579A..15K.

Lagache, Guilaine, Jean-Loup Puget, and Hervé Dole (Sept. 2005). “Dusty Infrared Galaxies: Sources of the Cosmic Infrared Background”. In: *Annual Review of Astronomy and Astrophysics* 43.1, pp. 727–768. DOI: 10.1146/annurev.astro.43.072103.150606. arXiv: astro-ph/0507298 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2005ARA&A..43..727L.

Lamperti, Isabella et al. (Nov. 2019). “JINGLE - V. Dust properties of nearby galaxies derived from hierarchical Bayesian SED fitting”. In: *Monthly Notices of the Royal Astronomical Society* 489.3, pp. 4389–4417. DOI: 10.1093/mnras/stz2311. arXiv: 1909.05266 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019MNRAS.489.4389L.

Le Bourlot, J. et al. (May 2012). “Surface chemistry in the interstellar medium. I. H_2 formation by Langmuir-Hinshelwood and Eley-Rideal mechanisms”. In: *Astronomy and Astrophysics* 541, A76, A76. DOI: 10.1051/0004-6361/201118126. arXiv: 1202.0374 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2012A&A...541A..76L.

Lee, Dennis et al. (Aug. 2022). “The TolTEC camera: polarimetric commissioning and performance of the continuously rotating half-wave plate”. In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 12190. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 121901O, 121901O. DOI: 10.1117/12.2627725. URL: https://ui.adsabs.harvard.edu/abs/2022SPIE12190E..1OL.

Leja, Joel et al. (Mar. 2017). “Deriving Physical Properties from Broadband Photometry with Prospector: Description of the Model and a Demonstration of its Accuracy Using 129 Galaxies in the Local Universe”. In: *Astrophysical Journal* 837.2, 170, p. 170. DOI: 10.3847/1538-4357/aa5ffe. arXiv: 1609.09073 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017ApJ...837..170L.

Leroy, A. et al. (June 2005). “The Molecular Interstellar Medium of Dwarf Galaxies on Kiloparsec Scales: A New Survey for CO in Northern, IRAS-detected Dwarf Galaxies”. In: *Astrophysical Journal* 625.2, pp. 763–784. DOI: 10.1086/429578. arXiv: astro-ph/0502302 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2005ApJ...625..763L.

Leroy, Adam et al. (Apr. 2007). “The Spitzer Survey of the Small Magellanic Cloud: Far-Infrared Emission and Cold Gas in the Small Magellanic Cloud”. In: *Astrophysical Journal* 658.2, pp. 1027–1046. DOI: 10.1086/511150. arXiv: astro-ph/0611687 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2007ApJ...658.1027L.

Leroy, Adam K. et al. (Aug. 2011). “The CO-to-H_2 Conversion Factor from Infrared Dust Emission across the Local Group”. In: *Astrophysical Journal* 737.1, 12, p. 12. DOI: 10.1088/0004-637X/737/1/12. arXiv: 1102.4618 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2011ApJ...737...12L.

Li, Aigen and J. Mayo Greenberg (Jan. 2003). “In dust we trust: an overview of observations and theories of interstellar dust”. In: *Solid State Astrochemistry*. Ed. by

Valerio Pirronello, Jacek Krelowski, and Giulio Manicò. Vol. 120, pp. 37–84. DOI: 10.48550/arXiv.astro-ph/0204392. arXiv: astro-ph/0204392 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003ssac.proc...37L.

Li, Qi et al. (Oct. 2021). "The origin of the dust extinction curve in milky way-like galaxies". In: *Monthly Notices of the Royal Astronomical Society* 507.1, pp. 548–559. DOI: 10.1093/mnras/stab2196. arXiv: 2012.03978 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2021MNRAS.507..548L.

Lisenfeld, U. et al. (Dec. 10, 2001). "(Sub)millimetre emission from NGC 1569: an abundance of very small grains". In: *Astronomy &; Astrophysics* 382.3, pp. 860–871. ISSN: 1432-0746. DOI: 10.1051/0004-6361:20011782. arXiv: astro-ph/0112212 [astro-ph].

Liu, Guilin et al. (Mar. 2010). "An Investigation of the Dust Content in the Galaxy Pair NGC 1512/1510 from Near-Infrared to Millimeter Wavelengths". In: *The Astronomical Journal* 139.3, pp. 1190–1198. DOI: 10.1088/0004-6256/139/3/1190. arXiv: 1001.1764 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010AJ....139.1190L.

Lunde, Emily et al. (Dec. 2020). "The optical design and performance of TolTEC: a millimeter-wave imaging polarimeter". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Vol. 11453. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 114534A, 114534A. DOI: 10.1117/12.2562798.

Lunde, Emily et al. (Aug. 2022). "The TolTEC camera: optical alignment and characterization". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 12190. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1219016, p. 1219016. DOI: 10.1117/12.2630340. URL: https://ui.adsabs.harvard.edu/abs/2022SPIE12190E..16L.

Ma, Zhiyuan et al. (Dec. 2020). "The TolTEC data analysis pipeline and software stack". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Vol. 11452. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 114522O, 114522O. DOI: 10.1117/12.2560735.

Madden, S. C. et al. (Feb. 2006). "ISM properties in low-metallicity environments". In: *Astronomy and Astrophysics* 446.3, pp. 877–896. DOI: 10.1051/0004-6361:20053890. arXiv: astro-ph/0510086 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2006A&A...446..877M.

Madden, S. C. et al. (June 2013). "An Overview of the Dwarf Galaxy Survey". In: *PASP* 125.928, p. 600. DOI: 10.1086/671138. arXiv: 1305.2628 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2013PASP..125..600M.

Madden, Suzanne C. and Diane Cormier (Oct. 2019). "Dwarf Galaxies: Their Low Metallicity Interstellar Medium". In: *Dwarf Galaxies: From the Deep Universe to the Present*. Ed. by Kristen B. W. McQuinn and Sabrina Stierwalt. Vol. 344, pp. 240–254. DOI: 10.1017/S1743921318007147. arXiv: 1810.09953 [astro-ph.GA].

Mangum, J. G., D. T. Emerson, and E. W. Greisen (Nov. 2007). "The On The Fly imaging technique". In: *Astronomy and Astrophysics* 474.2, pp. 679–687. DOI: 10.1051/0004-6361:20077811. arXiv: 0709.0553 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2007A&A...474..679M.

Marassi, S. et al. (Apr. 2019). "Supernova dust yields: the role of metallicity, rotation, and fallback". In: *Monthly Notices of the Royal Astronomical Society* 484.2, pp. 2587–2604. DOI: 10.1093/mnras/sty3323. arXiv: 1812.00009 [astro-ph.SR]. URL: https://ui.adsabs.harvard.edu/abs/2019MNRAS.484.2587M.

Marsh, K. A., A. P. Whitworth, and O. Lomax (Dec. 2015). "Temperature as a third dimension in column-density mapping of dusty astrophysical structures associated with star formation". In: *Monthly Notices of the Royal Astronomical Society* 454.4, pp. 4282–4292. DOI: 10.1093/mnras/stv2248. arXiv: 1509.08699 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2015MNRAS.454.4282M.

Marsh, K. A. et al. (Nov. 2017). "Multitemperature mapping of dust structures throughout the Galactic Plane using the PPMAP tool with Herschel Hi-GAL data". In: *Monthly Notices of the Royal Astronomical Society* 471.3, pp. 2730–2742. DOI: 10.1093/mnras/stx1723. arXiv: 1707.03808 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017MNRAS.471.2730M.

Mathis, J. S., P. G. Mezger, and N. Panagia (Nov. 1983). "Interstellar radiation field and dust temperatures in the diffuse interstellar medium and in giant molecular clouds". In: *Astronomy and Astrophysics* 128, pp. 212–229. URL: https://ui.adsabs.harvard.edu/abs/1983A&A...128..212M.

Mathis, J. S., W. Rumpl, and K. H. Nordsieck (Oct. 1977). "The size distribution of interstellar grains." In: *Astrophysical Journal* 217, pp. 425–433. DOI: 10.1086/155591. URL: https://ui.adsabs.harvard.edu/abs/1977ApJ...217..425M.

McCrackan, Michael et al. (Aug. 2022). "The TolTEC camera: the citlali data reduction pipeline engine". In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Vol. 12189. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 121891H, 121891H. DOI: 10.1117/12.2629095. URL: https://ui.adsabs.harvard.edu/abs/2022SPIE12189E..1HM.

Meixner, M. et al. (July 2010). "HERschel Inventory of The Agents of Galaxy Evolution (HERITAGE): The Large Magellanic Cloud dust". In: *Astronomy and Astrophysics* 518, L71, p. L71. DOI: 10.1051/0004-6361/201014662. arXiv: 1006.0985 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...518L..71M.

Meny, C. et al. (June 2007). "Far-infrared to millimeter astrophysical dust emission. I. A model based on physical properties of amorphous solids". In: *Astronomy and Astrophysics* 468.1, pp. 171–188. DOI: 10.1051/0004-6361:20065771. arXiv: astro-ph/0701226 [astro-ph].

Michałowski, Michał J. (May 2015). "Dust production 680-850 million years after the Big Bang". In: *Astronomy and Astrophysics* 577, A80, A80. DOI: 10.1051/0004-6361/201525644. arXiv: 1503.08210 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015A&A...577A..80M.

Monfardini, A. et al. (Oct. 2010). "NIKA: A millimeter-wave kinetic inductance camera". In: *Astronomy and Astrophysics* 521, A29, A29. DOI: 10.1051/0004-6361/201014727. arXiv: 1004.2209 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...521A..29M.

Morgan, H. L. and M. G. Edmunds (Aug. 2003). "Dust formation in early galaxies". In: *Monthly Notices of the Royal Astronomical Society* 343.2, pp. 427–442. DOI: 10.1046/j.1365-8711.2003.06681.x. arXiv: astro-ph/0302566 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003MNRAS.343..427M.

Murakami, Hiroshi et al. (Oct. 2007). "The Infrared Astronomical Mission AKARI*". In: *Publications of the Astronomical Society of Japan* 59, S369–S376. DOI: 10.1093/pasj/59.sp2.S369. arXiv: 0708.1796 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2007PASJ...59S.369M.

Næss, Sigurd K. (Dec. 2019). "How to avoid X'es around point sources in maximum likelihood CMB maps". In: *Journal of Cosmology and Astroparticle Physics* 2019.12, 060, p. 060. DOI: 10.1088/1475-7516/2019/12/060. arXiv: 1906.08030 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2019JCAP...12..060N.

Nashimoto, Masashi et al. (Sept. 2020). "Cosmic Amorphous Dust Model as the Origin of Anomalous Microwave Emission". In: *Astrophysical Journal Letters* 900.2, L40, p. L40. DOI: 10.3847/2041-8213/abb29d. arXiv: 2009.00137 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2020ApJ...900L..40N.

Neal, Radford M. (Sept. 2000). "Slice Sampling". In: *arXiv e-prints*, physics/0009028, physics/0009028. DOI: 10.48550/arXiv.physics/0009028. arXiv: physics/0009028 [physics.data-an]. URL: https://ui.adsabs.harvard.edu/abs/2000physics...9028N.

Nersesian, A. et al. (Apr. 2019). "Old and young stellar populations in DustPedia galaxies and their role in dust heating". In: *Astronomy and Astrophysics* 624, A80, A80. DOI: 10.1051/0004-6361/201935118. arXiv: 1903.05933 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019A&A...624A..80N.

Nersesian, Angelos et al. (May 2020). "High-resolution, 3D radiative transfer modelling. III. The DustPedia barred galaxies". In: *Astronomy and Astrophysics* 637, A25, A25. DOI: 10.1051/0004-6361/201936176. arXiv: 2004.03616 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2020A&A...637A..25N.

Neugebauer, G. et al. (Mar. 1984). "The Infrared Astronomical Satellite (IRAS) mission." In: *Astrophysical Journal Letters* 278, pp. L1–L6. DOI: 10.1086/184209. URL: https://ui.adsabs.harvard.edu/abs/1984ApJ...278L...1N.

Nittler, Larry R. and Fred Ciesla (Sept. 2016). "Astrophysics with Extraterrestrial Materials". In: *Annual Review of Astronomy and Astrophysics* 54, pp. 53–93. DOI: 10.1146/annurev-astro-082214-122505. URL: https://ui.adsabs.harvard.edu/abs/2016ARA&A..54...53N.

Paine, Scott (July 2022). *The am atmospheric model*. eng. DOI: 10.5281/zenodo.6774376. URL: https://zenodo.org/record/6774376 (visited on 08/10/2023).

Paradis, D. et al. (Jan. 2012). "Detection and characterization of a 500 μm dust emissivity excess in the Galactic plane using Herschel/Hi-GAL observations". In: *Astronomy and Astrophysics* 537, A113, A113. DOI: 10.1051/0004-6361/201117956. arXiv: 1111.1852 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2012A&A...537A.113P.

Paradis, D. et al. (Dec. 2014). "Modeling and predicting the shape of the far-infrared to submillimeter emission in ultra-compact HII regions and cold clumps". In: *Astronomy and Astrophysics* 572, A37, A37. DOI: 10.1051/0004-6361/201322566. arXiv: 1409.6892 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...572A..37P.

Pascale, Enzo and Pascale (Jan. 2013). "The Balloon-borne Large Aperture Submillimetre Telescope (BLAST) and BLASTPol". In: *Astrophysics from Antarctica*. Ed. by Michael G. Burton, Xiangqun Cui, and Nicholas F. H. Tothill. Vol. 288, pp. 154–160. DOI: 10.1017/S174392131201681X. URL: https://ui.adsabs.harvard.edu/abs/2013IAUS..288..154P.

Pattle, Kate, Walter Gear, and Christine D. Wilson (June 2023). "The JCMT nearby galaxies legacy survey: SCUBA-2 observations of nearby galaxies". In: *Monthly Notices of the Royal Astronomical Society* 522.2, pp. 2339–2368. DOI: 10.1093/mnras/stad652. arXiv: 2302.14800 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2023MNRAS.522.2339P.

Perera, T. A. et al. (July 2013). "An Efficient and Optimal Filter for Identifying Point Sources in Millimeter/Submillimeter Wavelength Sky Maps". In: *PASP* 125.929, p. 838. DOI: 10.1086/671756. arXiv: 1304.0413 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2013PASP..125..838P.

Pérez-González, Pablo G. et al. (Mar. 2008). "The Stellar Mass Assembly of Galaxies from z = 0 to z = 4: Analysis of a Sample Selected in the Rest-Frame Near-Infrared with Spitzer". In: *Astrophysical Journal* 675.1, pp. 234–261. DOI: 10.1086/523690. arXiv: 0709.1354 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2008ApJ...675..234P.

Pilbratt, G. L. et al. (July 2010). "Herschel Space Observatory. An ESA facility for far-infrared and submillimetre astronomy". In: *Astronomy and Astrophysics* 518, L1, p. L1. DOI: 10.1051/0004-6361/201014759. arXiv: 1005.5331 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2010A&A...518L...1P.

Pokhrel, R. et al. (Sept. 2016). "A Herschel-SPIRE survey of the Mon R2 giant molecular cloud: analysis of the gas column density probability density function". In: *Monthly Notices of the Royal Astronomical Society* 461.1, pp. 22–35. DOI: 10.1093/mnras/stw1303. arXiv: 1606.01752 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2016MNRAS.461...22P.

Poletti, Davide et al. (Apr. 2017). "Making maps of cosmic microwave background polarization for B-mode studies: the POLARBEAR example". In: *Astronomy and Astrophysics* 600, A60, A60. DOI: 10.1051/0004-6361/201629467. arXiv: 1608.01624 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2017A&A...600A..60P.

Popescu, C. C. et al. (Oct. 2000). "Modelling the spectral energy distribution of galaxies. I. Radiation fields and grain heating in the edge-on spiral NGC 891". In: *Astronomy and Astrophysics* 362, pp. 138–150. DOI: 10.48550/arXiv.astro-ph/0008098. arXiv: astro-ph/0008098 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2000A&A...362..138P.

Popescu, C. C. et al. (Mar. 2011). "Modelling the spectral energy distribution of galaxies. V. The dust and PAH emission SEDs of disk galaxies". In: *Astronomy and Astrophysics* 527, A109, A109. DOI: 10.1051/0004-6361/201015217. arXiv: 1011.2942 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2011A&A...527A.109P.

Priddey, Robert S. et al. (Oct. 2003). "Quasars as probes of the submillimetre cosmos at z ¿ 5 - I. Preliminary SCUBA photometry". In: *Monthly Notices of the Royal Astronomical Society* 344.4, pp. L74–L78. DOI: 10.1046/j.1365-8711.2003.07076.x. arXiv: astro-ph/0308132 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003MNRAS.344L..74P.

Priestley, F. D., M. J. Barlow, and I. De Looze (May 2019). "The mass, location, and heating of the dust in the Cassiopeia A supernova remnant". In: *Monthly Notices of the Royal Astronomical Society* 485.1, pp. 440–451. DOI: 10.1093/mnras/stz414. arXiv: 1902.01675 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019MNRAS.485..440P.

Rajput, Bhoomika and Ashwani Pandey (Dec. 2021). "γ-ray Flux and Spectral Variability of Blazar Ton 599 during Its 2021 Flare". In: *Galaxies* 9.4, 118, p. 118. DOI: 10.3390/galaxies9040118. arXiv: 2112.06681 [astro-ph.HE]. URL: https://ui.adsabs.harvard.edu/abs/2021Galax...9..118R.

Reach, W. T. et al. (Sept. 1995). "Far-Infrared Spectral Observations of the Galaxy by COBE". In: *Astrophysical Journal* 451, p. 188. DOI: 10.1086/176210. arXiv: astro-ph/9504056 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/1995ApJ...451..188R.

Rémy-Ruyer, A. et al. (Sept. 2013). "Revealing the cold dust in low-metallicity environments. I. Photometry analysis of the Dwarf Galaxy Survey with Herschel". In: *Astronomy and Astrophysics* 557, A95, A95. DOI: 10.1051/0004-6361/201321602. arXiv: 1309.1371 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2013A%5C&A...557A..95R.

Rémy-Ruyer, A. et al. (Mar. 2014). "Gas-to-dust mass ratios in local galaxies over a 2 dex metallicity range". In: *Astronomy and Astrophysics* 563, A31, A31. DOI: 10.1051/0004-6361/201322803. arXiv: 1312.3442 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014A%5C&A...563A..31R.

Rémy-Ruyer, A. et al. (Oct. 2015). "Linking dust emission to fundamental properties in galaxies: the low-metallicity picture". In: *Astronomy and Astrophysics* 582, A121, A121. DOI: 10.1051/0004-6361/201526067. arXiv: 1507.05432 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015A%5C&A...582A.121R.

Rew, R. and G. Davis (1990). "NetCDF: An Interface for Scientific Data Access". In: *IEEE Computer Graphics and Applications* 10.4, pp. 76–82. DOI: 10.xxxxx/yyyyy.

Rio Astorga, David del et al. (May 2017). "A generic parallel pattern interface for stream and data processing". In: *Concurrency and Computation: Practice and Experience* 29.24. ISSN: 1532-0634. DOI: 10.1002/cpe.4175.

Ritacco, A. et al. (Aug. 2018). "NIKA 150 GHz polarization observations of the Crab nebula and its spectral energy distribution". In: *Astronomy and Astrophysics* 616, A35, A35. DOI: 10.1051/0004-6361/201731551. arXiv: 1804.09581 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2018A&A...616A..35R.

Ritacco, A. et al. (July 2022). "Crab nebula at 260 GHz with the NIKA2 polarimeter: Implications for the polarization angle calibration of future CMB experiments". In: *mm Universe @ NIKA2 - Observing the mm Universe with the NIKA2 Camera*. Vol. 257. European Physical Journal Web of Conferences, 00042, p. 00042. DOI: 10.1051/epjconf/202225700042. arXiv: 2111.02143 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2022EPJWC.25700042R.

Rodrigo, C. and E. Solano (July 2020). "The SVO Filter Profile Service". In: *XIV.0 Scientific Meeting (virtual) of the Spanish Astronomical Society*, 182, p. 182. URL: https://ui.adsabs.harvard.edu/abs/2020sea..confE.182R.

Romero, Charles E. et al. (Mar. 2020). "Pressure Profiles and Mass Estimates Using High-resolution Sunyaev-Zel'dovich Effect Observations of Zwicky 3146 with MUSTANG-2". In: *Astrophysical Journal* 891.1, 90, p. 90. DOI: 10.3847/1538-4357/ab6d70. arXiv: 1908.09200 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2020ApJ...891...90R.

Rowlands, K. et al. (June 2014). "The dust budget crisis in high-redshift submillimetre galaxies". In: *Monthly Notices of the Royal Astronomical Society* 441.2, pp. 1040–1058. DOI: 10.1093/mnras/stu605. arXiv: 1403.2995 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014MNRAS.441.1040R.

Saintonge, Amélie et al. (Dec. 2018). "JINGLE, a JCMT legacy survey of dust and gas for galaxy evolution studies - I. Survey overview and first results". In: *Monthly Notices of the Royal Astronomical Society* 481.3, pp. 3497–3519. DOI: 10.1093/mnras/sty2499. arXiv: 1809.07336 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2018MNRAS.481.3497S.

Sajina, A. et al. (June 2006). "The 1-1000μm spectral energy distributions of far-infrared galaxies". In: *Monthly Notices of the Royal Astronomical Society* 369.2, pp. 939–957. DOI: 10.1111/j.1365-2966.2006.10361.x. arXiv: astro-ph/0603614 [astro-ph].

Sandstrom, Karin M. et al. (Jan. 2012). "The Spitzer Spectroscopic Survey of the Small Magellanic Cloud (S^4MC): Probing the Physical State of Polycyclic Aromatic Hydrocarbons in a Low-metallicity Environment". In: *Astrophysical Journal* 744.1, 20, p. 20. DOI: 10.1088/0004-637X/744/1/20. arXiv: 1109.0999 [astro-ph.CO].

Sargent, Benjamin A. et al. (June 2010). "The Mass-loss Return from Evolved Stars to the Large Magellanic Cloud. II. Dust Properties for Oxygen-rich Asymptotic Giant Branch Stars". In: *Astrophysical Journal* 716.1, pp. 878–890. DOI: 10.1088/0004-637X/716/1/878. arXiv: 1407.6996 [astro-ph.SR]. URL: https://ui.adsabs.harvard.edu/abs/2010ApJ...716..878S.

Sargent, Wallace L. W. and Leonard Searle (Dec. 1970). "Isolated Extragalactic H II Regions". In: *Astrophysical Journal Letters* 162, p. L155. DOI: 10.1086/180644. URL: https://ui.adsabs.harvard.edu/abs/1970ApJ...162L.155S.

Schnee, Scott et al. (Jan. 2010). "The Dust Emissivity Spectral Index in the Starless Core TMC-1C". In: *Astrophysical Journal* 708.1, pp. 127–136. DOI: 10.1088/0004-637X/708/1/127. arXiv: 0911.0892 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2010ApJ...708..127S.

Schneider, Raffaella et al. (Aug. 2014). "Dust production rate of asymptotic giant branch stars in the Magellanic Clouds". In: *Monthly Notices of the Royal Astronomical Society* 442.2, pp. 1440–1450. DOI: 10.1093/mnras/stu861. arXiv: 1404.7132 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014MNRAS.442.1440S.

Schruba, Andreas et al. (June 2012). "Low CO Luminosities in Dwarf Galaxies". In: *The Astronomical Journal* 143.6, 138, p. 138. DOI: 10.1088/0004-6256/143/6/138. arXiv: 1203.4231 [astro-ph.CO].

Scott, K. S. et al. (Apr. 2008). "AzTEC millimetre survey of the COSMOS field - I. Data reduction and source catalogue". In: *Monthly Notices of the Royal Astronomical Society* 385.4, pp. 2225–2238. DOI: 10.1111/j.1365-2966.2008.12989.x. arXiv: 0801.2779 [astro-ph].

Sheth, Kartik et al. (Dec. 2010). "The Spitzer Survey of Stellar Structure in Galaxies (S4G)". In: *PASP* 122.898, p. 1397. DOI: 10.1086/657638. arXiv: 1010.1592 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010PASP..122.1397S.

Shetty, Rahul et al. (May 2009a). "The Effect of Line-of-Sight Temperature Variation and Noise on Dust Continuum Observations". In: *Astrophysical Journal* 696.2, pp. 2234–2251. DOI: 10.1088/0004-637X/696/2/2234. arXiv: 0902.3477 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2009ApJ...696.2234S.

Shetty, Rahul et al. (May 2009b). “The Effect of Noise on the Dust Temperature-Spectral Index Correlation”. In: *Astrophysical Journal* 696.1, pp. 676–680. DOI: 10.1088/0004-637X/696/1/676. arXiv: 0902.0636 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2009ApJ...696..676S.

Siebenmorgen, R., N. V. Voshchinnikov, and S. Bagnulo (Jan. 2014). “Dust in the diffuse interstellar medium. Extinction, emission, linear and circular polarisation”. In: *Astronomy and Astrophysics* 561, A82, A82. DOI: 10.1051/0004-6361/201321716. arXiv: 1308.3148 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2014A&A...561A..82S.

Sievers, Jonathan (2023). *minkasi*. URL: https://github.com/sievers/minkasi.

Silk, J., A. Di Cintio, and I. Dvorkin (Dec. 2014). “Galaxy Formation”. In: *Proceedings of the International School of Physics 'Enrico Fermi' Course 186 'New Horizons for Observational Cosmology' Vol. 186*. Vol. 186, pp. 137–187. DOI: 10.3254/978-1-61499-476-3-137. arXiv: 1312.0107 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2014nhoc.conf..137S.

Simon, Joshua D. (Aug. 2019a). “The Faintest Dwarf Galaxies”. In: *Annual Review of Astronomy and Astrophysics* 57, pp. 375–415. DOI: 10.1146/annurev-astro-091918-104453. arXiv: 1901.05465 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019ARA%5C&A..57..375S.

Simon, Joshua D. (Aug. 2019b). “The Faintest Dwarf Galaxies”. In: *Annual Review of Astronomy and Astrophysics* 57, pp. 375–415. DOI: 10.1146/annurev-astro-091918-104453. arXiv: 1901.05465 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019ARA&A..57..375S.

Siringo, G. et al. (Apr. 2009). “The Large APEX BOlometer CAmera LABOCA”. In: *Astronomy and Astrophysics* 497.3, pp. 945–962. DOI: 10.1051/0004-6361/200811454. arXiv: 0903.1354 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2009A&A...497..945S.

Slavin, Jonathan D., Eli Dwek, and Anthony P. Jones (Apr. 2015). “Destruction of Interstellar Dust in Evolving Supernova Remnant Shock Waves”. In: *Astrophysical Journal* 803.1, 7, p. 7. DOI: 10.1088/0004-637X/803/1/7. arXiv: 1502.00929 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015ApJ...803....7S.

Sokol, Alyssa D. et al. (Feb. 2019). “Early science with the Large Millimetre Telescope: An LMT/AzTEC 1.1 mm Survey of dense cores in the Monoceros R2 giant molecular cloud”. In: *Monthly Notices of the Royal Astronomical Society* 483.1, pp. 407–424. DOI: 10.1093/mnras/sty3107. arXiv: 1811.07024 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019MNRAS.483..407S.

Stepnik, B. et al. (Feb. 2003). "Evolution of dust properties in an interstellar filament". In: *Astronomy and Astrophysics* 398, pp. 551–563. DOI: 10.1051/0004-6361:20021309. URL: https://ui.adsabs.harvard.edu/abs/2003A&A...398..551S.

Sueno, Y. et al. (Mar. 2022). "Characterization of two-level system noise for a microwave kinetic inductance detector comprising niobium film on a silicon substrate". In: *Progress of Theoretical and Experimental Physics* 2022.3, 033H01, 033H01. DOI: 10.1093/ptep/ptac023. arXiv: 2110.00127 [physics.ins-det]. URL: https://ui.adsabs.harvard.edu/abs/2022PTEP.2022c3H01S.

Sunyaev, R. A. and Ya. B. Zeldovich (Apr. 1970). "Small-Scale Fluctuations of Relic Radiation". In: *Astrophysics and Space Science* 7.1, pp. 3–19. DOI: 10.1007/BF00653471. URL: https://ui.adsabs.harvard.edu/abs/1970Ap&SS...7....3S.

Sunyaev, R. A. and Ya. B. Zeldovich (Nov. 1972). "The Observations of Relic Radiation as a Test of the Nature of X-Ray Radiation from the Clusters of Galaxies". In: *Comments on Astrophysics and Space Physics* 4, p. 173. URL: https://ui.adsabs.harvard.edu/abs/1972CoASP...4..173S.

Tang, Yuping (Jan. 2019). "AzTEC survey of the central molecular zone: Modeling dust SEDs with hierarchical bayesian analysis". PhD thesis. UMass Amherst. URL: https://ui.adsabs.harvard.edu/abs/2019PhDT........99T.

Tang, Yuping, Q. Daniel Wang, and Grant W. Wilson (Aug. 2021). "AzTEC survey of the central molecular zone: increasing spectral index of dust with density". In: *Monthly Notices of the Royal Astronomical Society* 505.2, pp. 2377–2391. DOI: 10.1093/mnras/staa3230. arXiv: 2008.12361 [astro-ph.GA].

Tapia, M. et al. (Aug. 2022). "The Mexico UK Sub-mm Camera for Astronomy (MUSCAT) on-sky commissioning: focal plane performance". In: *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy XI*. Ed. by Jonas Zmuidzinas and Jian-Rong Gao. Vol. 12190. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1219017, p. 1219017. DOI: 10.1117/12.2630377. URL: https://ui.adsabs.harvard.edu/abs/2022SPIE12190E..17T.

Tegmark, Max (Oct. 1997). "CMB mapping experiments: A designer's guide". In: *Physical Review D* 56.8, pp. 4514–4529. DOI: 10.1103/PhysRevD.56.4514. arXiv: astro-ph/9705188 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/1997PhRvD..56.4514T.

Temim, Tea and Eli Dwek (Sept. 2013). "The Importance of Physical Models for Deriving Dust Masses and Grain Size Distributions in Supernova Ejecta. I. Radiatively Heated Dust in the Crab Nebula". In: *Astrophysical Journal* 774.1, 8, p. 8. DOI: 10.1088/0004-637X/774/1/8. arXiv: 1302.5452 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2013ApJ...774....8T.

Thirlwall, Jordan J. et al. (June 2020). "A radiative transfer model for the spiral galaxy M33". In: *Monthly Notices of the Royal Astronomical Society* 495.1, pp. 835–863. DOI: 10.1093/mnras/staa905. arXiv: 2004.00400 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2020MNRAS.495..835T.

Thomas, Daniel et al. (June 2010). "Environment and self-regulation in galaxy formation". In: *Monthly Notices of the Royal Astronomical Society* 404.4, pp. 1775–1789. DOI: 10.1111/j.1365-2966.2010.16427.x. arXiv: 0912.0259 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2010MNRAS.404.1775T.

Thornton, R. J. et al. (Dec. 2016). "The Atacama Cosmology Telescope: The Polarization-sensitive ACTPol Instrument". In: *Astrophysical Journal Supplement* 227.2, 21, p. 21. DOI: 10.3847/1538-4365/227/2/21. arXiv: 1605.06569 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2016ApJS..227...21T.

Tielens, A. G. G. M. et al. (Apr. 1996). "The Infrared Spectrum of the Galactic Center and the Composition of Interstellar Dust". In: *Astrophysical Journal* 461, p. 210. DOI: 10.1086/177049. URL: https://ui.adsabs.harvard.edu/abs/1996ApJ...461..210T.

Todini, Paolo and Andrea Ferrara (Aug. 2001). "Dust formation in primordial Type II supernovae". In: *Monthly Notices of the Royal Astronomical Society* 325.2, pp. 726–736. DOI: 10.1046/j.1365-8711.2001.04486.x. arXiv: astro-ph/0009176 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2001MNRAS.325..726T.

Tolstoy, Eline, Vanessa Hill, and Monica Tosi (Sept. 2009). "Star-Formation Histories, Abundances, and Kinematics of Dwarf Galaxies in the Local Group". In: *Annual Review of Astronomy and Astrophysics* 47.1, pp. 371–425. DOI: 10.1146/annurev-astro-082708-101650. arXiv: 0904.4505 [astro-ph.CO]. URL: https://ui.adsabs.harvard.edu/abs/2009ARA%5C&A..47..371T.

Veneziani, M. et al. (July 2013). "A Bayesian Method for the Analysis of the Dust Emission in the Far-infrared and Submillimeter". In: *Astrophysical Journal* 772.1, 56, p. 56. DOI: 10.1088/0004-637X/772/1/56. arXiv: 1305.1339 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2013ApJ...772...56V.

Ventura, P. et al. (Aug. 2012). "Dust formation around AGB and SAGB stars: a trend with metallicity?" In: *Monthly Notices of the Royal Astronomical Society* 424.3, pp. 2345–2357. DOI: 10.1111/j.1365-2966.2012.21403.x. arXiv: 1205.6216 [astro-ph.SR]. URL: https://ui.adsabs.harvard.edu/abs/2012MNRAS.424.2345V.

Wakelam, V. et al. (Mar. 2017). "Binding energies: New values and impact on the efficiency of chemical desorption". In: *Molecular Astrophysics* 6, pp. 22–35. DOI: 10.1016/j.molap.2017.01.002. arXiv: 1701.06492 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2017MolAs...6...22W.

Warren, S. J. and S. Dye (June 2003). "Semilinear Gravitational Lens Inversion". In: *Astrophysical Journal* 590.2, pp. 673–682. DOI: 10.1086/375132. arXiv: astro-ph/0302587 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2003ApJ...590..673W.

Watson, Darach et al. (Mar. 2015). "A dusty, normal galaxy in the epoch of reionization". In: *Nature* 519.7543, pp. 327–330. DOI: 10.1038/nature14164. arXiv: 1503.00002 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015Natur.519..327W.

Wells, D. C., E. W. Greisen, and R. H. Harten (June 1981). "FITS - a Flexible Image Transport System". In: *Astronomy and Astrophysics Supplement* 44, p. 363. URL: https://ui.adsabs.harvard.edu/abs/1981A&AS...44..363W.

Werner, M. W. et al. (Sept. 2004). "The Spitzer Space Telescope Mission". In: *Astrophysical Journal Supplement* 154.1, pp. 1–9. DOI: 10.1086/422992. arXiv: astro-ph/0406223 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2004ApJS..154....1W.

White, S. D. M. and M. J. Rees (May 1978). "Core condensation in heavy halos: a two-stage theory for galaxy formation and clustering." In: *Monthly Notices of the Royal Astronomical Society* 183, pp. 341–358. DOI: 10.1093/mnras/183.3.341. URL: https://ui.adsabs.harvard.edu/abs/1978MNRAS.183..341W.

Williams, Thomas G. et al. (Aug. 2019). "High-resolution radiative transfer modelling of M33". In: *Monthly Notices of the Royal Astronomical Society* 487.2, pp. 2753–2770. DOI: 10.1093/mnras/stz1441. arXiv: 1905.09838 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2019MNRAS.487.2753W.

Wilson, C. D. et al. (Mar. 2009). "The James Clerk Maxwell Telescope Nearby Galaxies Legacy Survey. I. Star-Forming Molecular Gas in Virgo Cluster Spiral Galaxies". In: *Astrophysical Journal* 693.2, pp. 1736–1748. DOI: 10.1088/0004-637X/693/2/1736. arXiv: 0812.1718 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2009ApJ...693.1736W.

Wilson, G. W. et al. (May 2008). “The AzTEC mm-wavelength camera”. In: *Monthly Notices of the Royal Astronomical Society* 386.2, pp. 807–818. DOI: 10.1111/j.1365-2966.2008.12980.x. arXiv: 0801.2783 [astro-ph].

Wilson, Grant W. et al. (Dec. 2020). “The TolTEC camera: an overview of the instrument and in-lab testing results”. In: *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*. Vol. 11453. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 1145302, p. 1145302. DOI: 10.1117/12.2562331.

Wolfire, M. G. et al. (Apr. 1995). “The Neutral Atomic Phases of the Interstellar Medium”. In: *Astrophysical Journal* 443, p. 152. DOI: 10.1086/175510. URL: https://ui.adsabs.harvard.edu/abs/1995ApJ...443..152W.

Woolf, N. J. and E. P. Ney (Mar. 1969). “Circumstellar Infrared Emission from Cool Stars”. In: *Astrophysical Journal Letters* 155, p. L181. DOI: 10.1086/180331. URL: https://ui.adsabs.harvard.edu/abs/1969ApJ...155L.181W.

Wright, Edward L. et al. (Dec. 2010). “The Wide-field Infrared Survey Explorer (WISE): Mission Description and Initial On-orbit Performance”. In: *The Astronomical Journal* 140.6, pp. 1868–1881. DOI: 10.1088/0004-6256/140/6/1868. arXiv: 1008.0031 [astro-ph.IM]. URL: https://ui.adsabs.harvard.edu/abs/2010AJ....140.1868W.

Young, Andrew T. and William M. Irvine (Oct. 1967). “Multicolor photoelectric photometry of the brighter planets. I. Program and Procedure”. In: *The Astronomical Journal* 72, p. 945. DOI: 10.1086/110366. URL: https://ui.adsabs.harvard.edu/abs/1967AJ.....72..945Y.

Ysard, N. et al. (May 2015). “Dust variations in the diffuse interstellar medium: constraints on Milky Way dust from Planck-HFI observations”. In: *Astronomy and Astrophysics* 577, A110, A110. DOI: 10.1051/0004-6361/201425523. arXiv: 1503.07435 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2015A&A...577A.110Y.

Yu, Yaming and Xiao-Li Meng (2011). “To Center or Not to Center: That Is Not the Question—An Ancillarity—Sufficiency Interweaving Strategy (ASIS) for Boosting MCMC Efficiency”. In: *Journal of Computational and Graphical Statistics* 20.3, pp. 531–570. ISSN: 10618600. URL: http://www.jstor.org/stable/23248837 (visited on 04/03/2024).

Zeballos, M. et al. (July 2016). “Reporting the first 3 years of 225-GHz opacity measurements at the site of the Large Millimeter Telescope Alfonso Serrano”. In: *Ground-based and Airborne Telescopes VI*. Ed. by Helen J. Hall, Roberto Gilmozzi, and Heather K. Marshall. Vol. 9906. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 99064U, 99064U. DOI: 10.1117/12.2232168. URL: https://ui.adsabs.harvard.edu/abs/2016SPIE.9906E..4UZ.

Zhukovska, Svitlana (Feb. 2014). "Dust origin in late-type dwarf galaxies: ISM growth vs. type II supernovae". In: *Astronomy and Astrophysics* 562, A76, A76. DOI: 10.1051/0004-6361/201322989. arXiv: 1401.1675 [astro-ph.GA].

Zhukovska, Svitlana et al. (Nov. 2016). "Modeling Dust Evolution in Galaxies with a Multiphase, Inhomogeneous ISM". In: *Astrophysical Journal* 831.2, 147, p. 147. DOI: 10.3847/0004-637X/831/2/147. arXiv: 1608.04781 [astro-ph.GA]. URL: https://ui.adsabs.harvard.edu/abs/2016ApJ...831..147Z.

Zubko, Viktor, Eli Dwek, and Richard G. Arendt (June 2004). "Interstellar Dust Models Consistent with Extinction, Emission, and Abundance Constraints". In: *Astrophysical Journal Supplement* 152.2, pp. 211–249. DOI: 10.1086/382351. arXiv: astro-ph/0312641 [astro-ph]. URL: https://ui.adsabs.harvard.edu/abs/2004ApJS..152..211Z.