

**ESTIMATING RADIAL VELOCITY AND IMPROVING LINE IDENTIFICATION
WITHIN THE HERSCHEL/SPIRE AUTOMATED SPECTRAL FEATURE
CATALOGUE**

JEREMY SCOTT
Bachelor of Science, University of Lethbridge, 2015

A Thesis
Submitted to the School of Graduate Studies
of the University of Lethbridge
in Partial Fulfillment of the
Requirements for the Degree

MASTER OF SCIENCE

Department of Physics and Astronomy
University of Lethbridge
LETHBRIDGE, ALBERTA, CANADA

© Jeremy Scott, 2017

Dedication

This thesis is dedicated to my Mother for her unwavering love and support.

Abstract

The Herschel space observatory housing the Spectral and Photometric Imaging Receiver (SPIRE) was commissioned to study the far-infrared Universe. Though *Herschel's* mission has ended, much data processing is still required to extract all the available information from its observations.

To help with data mining and preliminary analysis, the SPIRE Fourier transform spectrometer instrument team developed an automated spectral Feature Finder (FF) that extracts significant spectral features from SPIRE spectra. As part of this team I developed two modules, an automated radial velocity estimating routine which uses a pattern seeking algorithm to identify ^{12}CO rotational emission, and a separate routine which conducts a dedicated search for neutral carbon fine structure transitions.

This thesis provides an introduction to Fourier transform spectroscopy, then describes in detail the two modules I developed for the FF. Validation of these routines is presented, and their results, as they pertain to the public FF catalogue, are summarized.

Table of Contents

Table of Contents	v
List of Tables	viii
List of Figures	ix
1 Introduction	1
1.1 Far-Infrared Astronomy	2
1.2 SPIRE and the Herschel Space Observatory	5
1.3 My Contributions to SPIRE	7
1.4 Summary	10
1.5 Thesis Outline	11
2 Fourier Analysis	13
2.1 Definitions	14
2.1.1 Fourier Series	14
2.1.2 Fourier Integral	15
2.1.3 Dirac Delta Function	16
2.1.4 Useful Transforms	17
2.2 Properties of Fourier Transforms	23
2.2.1 Symmetry	23
2.2.2 Theorems	28
2.3 Discretization	31
2.3.1 Discrete Fourier Transform	33
2.3.2 Fast Fourier Transform	35
2.3.3 Sampling	39
2.4 Summary	42
3 Optical Fourier Transform Spectroscopy	43
3.1 Spectral Composition	44
3.1.1 Black-body Radiation	44
3.1.2 Radiative Transfer	49
3.2 Line Shapes	58
3.2.1 Natural Broadening	59
3.2.2 Pressure Broadening	60
3.2.3 Doppler Broadening	62
3.2.4 Signal Truncation	63
3.2.5 Combining Line shapes	64

3.3	Radiation Processes	65
3.3.1	Atomic Emission	66
3.3.2	Rotational Emission	75
3.4	FTS Instrumentation and Processing	80
3.4.1	Michelson Interferometer	80
3.4.2	Synthetic Spectral Resolution	86
3.4.3	Phase Correction	89
3.5	Summary	95
4	SPIRE FTS and the Automated Spectral Feature Finder	97
4.1	Introduction	97
4.2	SPIRE Imaging Fourier Transform Spectrometer	99
4.3	The Herschel Interactive Processing Environment	104
4.4	SPIRE Pipeline and Products	106
4.4.1	Calibration Schemes	109
4.5	SPIRE FTS Feature Finder	115
4.5.1	Feature Finder Routine	117
4.6	Summary	122
5	Redshift	125
5.1	Introduction	125
5.2	Theoretical Foundation	126
5.3	Routine Outline	131
5.3.1	Detailed Routine Description	133
5.4	Validation	138
5.4.1	Real Data	140
5.4.2	Simulated Data	146
5.5	Feature Finder Implementation	151
5.6	Summary	159
6	Neutral Carbon Detection	161
6.1	Introduction	161
6.2	Theoretical Basis	164
6.3	Neutral Carbon Detection Routine	167
6.3.1	General Overview	167
6.3.2	Detailed Outline	169
6.3.3	Secondary Feature Amplitude Estimate	172
6.4	Neutral Carbon Detection Results	175
6.5	Further Applications	182
6.6	Summary	186
7	Conclusion	187
7.1	Introduction	187
7.2	Redshift Routine	188
7.3	Neutral Carbon	189

7.4	Future Work	190
References		192
A	Visual Example of Radial Velocity Estimating Algorithm	200
B	Mapping Observation Postcards	206
B.1	^{12}CO Dominant	211
B.2	NII Dominant	213
B.3	Null Dominant	215
C	Neutral Carbon $^3\text{P}_1$-$^3\text{P}_0$ Fine Structure Line Detection	217
C.1	Unique Neutral Carbon Pairs	217
C.2	Unique $^3\text{P}_1 - ^3\text{P}_0$ Transition Detection	219

List of Tables

2.1	Symmetry properties of Fourier transform	27
2.2	Properties and theorems of Fourier transforms	32
2.3	Fast Fourier transform flowchart	36
6.1	Bulk results of NCC on SPIRE FTS sparse dataset.	177
6.2	NCC results on a manually inspected dataset of 103 observations.	178
B.1	Summary of FF mapping postcards	206

List of Figures

1.1	Spectral opacity of Earth’s atmosphere	5
1.2	The <i>Herschel</i> space observatory	6
1.3	FTS sinc line shape	8
2.1	Fourier transform of a boxcar function	19
2.2	Fourier transform of a Gaussian	20
2.3	Fourier transform of an exponential	22
2.4	Fourier transform of a sinusoid	23
2.5	Fourier transform of a real valued even function	24
2.6	Fourier transform of a real value odd function	25
2.7	Fourier transform of an asymmetric function	27
2.8	A visual representation of convolution	29
2.9	Fourier transform algorithm flowchart	37
2.10	Effects of sampling frequency on measured interferogram	40
2.11	Spectral aliasing	41
3.1	SPIRE spectrum of NGC 7027 showing spectral components	49
3.2	Collisionally induced random phase shift of dipole moment.	61
3.3	Energy level splitting of carbon atom.	76
3.4	Modified Michelson interferometer diagram	81
3.5	Effects of zero padding interferogram	87
3.6	Demonstration of how zero padding interferogram results in ILS interpolation	88
3.7	Asymmetric interferogram and phase correction	93
3.8	Symmetrized/Phase corrected interferogram and spectrum	94
3.9	Comparison of uncorrected and phase corrected spectra	95
4.1	SPIRE FTS optics	100
4.2	SPIRE FTS hexagonal close packed SLW and SSW bolometer arrays . . .	101
4.3	Frequency dependence of SPIRE FTS beams	102
4.4	SPIRE FTS spatial sampling jiggle patterns	103
4.5	SPIRE FTS data reduction pipeline	107
4.6	Source spatial extent and SPIRE FTS calibration schemes	110
4.7	Flowchart showing key steps of the FF	116
4.8	Sample postcard generated by the FF	120
4.9	FF feature detection completeness curves for 3 classes of synthetic data . .	123
5.1	Demonstration of source radial velocity on observed spectrum	128
5.2	Flowchart showing key steps of the automated radial velocity estimating routine	132

5.3	SPIRE spectrum and PSW image of SH 104 expressing strong NII emission	137
5.4	Flux integration map and spectral image of SH 104	138
5.5	Close up of SH 104 spectrum showing both NII and ^{12}CO rotational emission	139
5.6	Velocity map showing discontinuity of velocity gradient as a result of velocity estimates derived from NII or ^{12}CO emission	139
5.7	Correlation plot of FF velocity estimate and HIFI velocities	141
5.8	Residual velocity completeness curve for HIFI velocity comparison	142
5.9	Residual velocity of FF redshift routine and HIFI velocities based on flagging parameter	143
5.10	Correlation plot of FF velocity estimates and velocities derived using manual inspection	145
5.11	Residual velocity completeness curve and dependence on flagging parameter	146
5.12	Comparison of FF residual velocity and SPIRE FTS frequency resolution in velocity space	147
5.13	Correlation plot of FF velocity estimates on simulated data with known redshift	150
5.14	Residual velocity completeness curve and dependence on flagging parameter for simulated data	151
5.15	Distribution of FF velocity flagging parameters for estimates with low residual velocities on simulated data	152
5.16	Distribution of FF velocity estimates	153
5.17	On-sky location of SPIRE FTS observations and corresponding FF velocity estimate quality.	154
5.18	Example of SPIRE spectrum with many features from a source with poor FF velocity estimate	155
5.19	Example of SPIRE spectrum with many features from a source with high quality FF velocity estimate	155
5.20	Comprehensive FF mapping postcard	158
5.21	Sample FF mapping postcard included in public features catalogue	159
6.1	Model of the different components of the ISM	162
6.2	Example of ^{12}CO (7-6)/CI spectral region	164
6.3	Convolution model of ^{12}CO (7-6)/CI spectral region	166
6.4	Flowchart showing the key steps of the NCC	168
6.5	SPIRE spectra and PSW images of NCC test case spectra	174
6.6	Tolerance curves for initial amplitude estimate of CI feature	176
6.7	Resulting NCC fits on test case spectra	177
6.8	Residuals comparison of ^{12}CO (7-6)/CI region with NCC turned off and on	180
6.9	^{12}CO (7-6)/CI fitted frequency difference with NCC turned off and on	181
6.10	Example of NCC applied to neutral carbon $^3\text{P}_1$ - $^3\text{P}_0$ fine structure transition	184
6.11	SNR distribution of $^3\text{P}_1$ - $^3\text{P}_0$ detection with and without corresponding $^3\text{P}_2$ - $^3\text{P}_1$ detection	185
A.1	SPIRE spectrum of NGC 7027	200
A.2	Rotational ^{12}CO emission ladder indicating CDA	201

A.3	CDA and corresponding tolerances	202
A.4	Continuum subtracted spectrum of NGC 7027 with features identified by the FF	202
A.5	Visual representation of the inputs to the redshift routine	203
A.6	Frequency differences between one feature under investigation and all features	203
A.7	Frequency differences matching the CDA to within the permitted tolerance .	204
A.8	Total number of matches for each feature identified by the FF	205
B.1	FF Mapping Postcards, 1342192174	211
B.2	1342214827	211
B.3	1342227516	212
B.4	1342228703	212
B.5	1342228732	213
B.6	1342221711	213
B.7	1342254040	214
B.8	1342265807	214
B.9	1342212345	215
B.10	1342210860	215
B.11	1342208388	216
B.12	1342259066	216
C.1	$[CI]^3 P_1 -^3 P_0$ snr = 37.846340	217
C.2	$[CI]^3 P_1 -^3 P_0$ snr = 12.180308	218
C.3	$[CI]^3 P_1 -^3 P_0$ snr = 18.500926	218
C.4	$[CI]^3 P_1 -^3 P_0$ snr = 7.949178	218
C.5	$[CI]^3 P_1 -^3 P_0$ snr = 6.268916	219
C.6	$[CI]^3 P_1 -^3 P_0$ snr = 7.668759	219
C.7	$[CI]^3 P_1 -^3 P_0$ snr = 5.366224	220
C.8	$[CI]^3 P_1 -^3 P_0$ snr = 5.002306	220
C.9	$[CI]^3 P_1 -^3 P_0$ snr = 6.959381	220
C.10	$[CI]^3 P_1 -^3 P_0$ snr = 5.755870	221
C.11	$[CI]^3 P_1 -^3 P_0$ snr = 6.509541	221
C.12	$[CI]^3 P_1 -^3 P_0$ snr = 7.247980	221

Acronyms

- BSM** Beam Steering Mirror. 101
- CDA** Characteristic Difference Array. 130
- CNM** Cold Neutral Medium. 162
- COM** Center of Mass. 76
- DAPSAS** Canadian Data Processing and Science Analysis Software. 105
- DFT** Discrete Fourier Transform. 34
- ESA** The European Space Agency. 3
- FDA** Frequency Difference Array. 134
- FF** SPIRE FTS automated Feature Finder. 7, 115
- FFT** Fast Fourier Transform. 35
- FIR** Far-Infrared. 2, 97
- FTS** Fourier Transform Spectrometer. 7, 13
- FWHM** Full Width at Half Maximum. 60
- HIFI** Heterodyne Instrument for the Far Infrared. 98
- HIM** Hot Ionized Medium. 161
- HIPE** Herschel Interactive Processing Environment. 104
- HPDP** Highly Processed Data Products. 114
- HR** High Resolution. 103
- HSA** The *Herschel* Science Archive. 7, 104
- IFTS** Imaging Fourier Transform Spectrometer. 6, 99
- ILS** Instrument Line Shape. 164
- IR** Infrared. 2

- IRAC** Infrared Array Camera. 4
- IRAS** Infrared Astronomical Satellite. 4, 6
- IRS** Infrared Spectrograph. 4
- ISM** Interstellar Medium. 2, 161
- ISO** Infrared Space Observatory. 4, 6
- ISRF** Interstellar Radiation Fields. 161
- LR** Low Resolution. 104
- MIPS** Multiband Infrared Photometer for Spitzer. 4
- NASA** The National Aeronautics and Space Administration. 5
- NCC** Neutral Carbon Check. 167
- obsid** observation identification number. 106
- OPD** Optical Path Difference. 82
- PACS** Photodetector Array Camera and Spectrometer. 98
- PCF** Phase Correction Function. 92
- PLW** Photometer Long Wavelength. 6
- PMW** Photometer Medium Wavelength. 6
- POSTCAT** POSTcard CATalogue. 119
- PSW** Photometer Short Wavelength. 6
- RSRF** Relative Spectral Response Function. 111
- SAFECAT** Spire Automated Feature Extraction CATalogue. 8, 119
- SCAL** Spectrometer Calibration Source. 111
- SLW** SPIRE Long Wavelength Spectrometer array. 7, 100
- SMEC** Spectrometer Mechanism. 112
- SNR** Signal-to-Noise Ratio. 7
- SPG** Standard Product Generation. 109
- SPIRE** Spectral and Photometric Imaging REceiver. 6, 98

SSW SPIRE Short Wavelength Spectrometer array. 6, 100

UV Ultraviolet. 161

WIM Warm Ionized Medium. 161

WNM Warm Neutral Medium. 161

ZPD Zero Path Difference. 89

Chapter 1

Introduction

How the Universe came to be in its current form is a question which has captivated humanity throughout history. Significant progress toward this and related questions began with Galileo and his improved telescope design in 1609 [1]. This technology opened a window into the cosmos which revolutionized our picture of the Universe. In short order, it was discovered that we reside not at the center of the Universe, but around one out of hundreds of billions of stars, in a galaxy which is but one out of hundreds of billions of galaxies [2]. This highlights the importance of astronomy and technology in answering such questions.

Unfortunately the vast expanse of the Universe makes a direct study of its structure and dynamic processes impractical, if not impossible. For perspective, the man-made object most distant from Earth is Voyager 1 which was launched on September 5, 1977 [3]. As of August 11, 2017 the spacecraft has travelled $\sim 2.08 \times 10^{10}$ km, resulting in an average speed of ~ 17 km/s [3]. This speed is remarkably fast, much faster than any terrestrial vehicle. However, the nearest star to Earth is Proxima Centauri at a distance of 4.2 light years, and travelling at the speed of Voyager 1 would take $\sim 1,930$ years to reach from Earth. As such, astronomy is to a large extent limited to the study of light, where it comes from, how bright it is, and what electromagnetic frequencies it is composed of, including polarization and temporal variability.

Modern advances in astronomy come primarily from improved observational technologies and advances in theoretical models. There is also a constant effort to extract as much information as possible from existing data. This second point expresses the dominant theme

of this thesis. For the remainder of this chapter, I will attempt to outline the context and motivation for the work I have done, and how it ties into the broader goals of the astronomical community.

1.1 Far-Infrared Astronomy

The radiant energy contained within the Infrared (IR) and Far-Infrared (FIR) bands of the electromagnetic spectrum comprises nearly half of all spectral energy observed within the Universe [4]. This figure is not surprising when considering the diversity and abundance of FIR sources. Significant contributors are described in this section.

A significant portion of the baryonic matter in the Universe exists in the form of molecular gas and dust, which fills the space between stars and planets. This so called Interstellar Medium (ISM) radiates brightly in the FIR, both as a result of the low temperature thermal radiation from dust grains, and the low energy rotational emissions from molecules [5]. The ISM itself contains density variations, with regions of higher density becoming the focus of gravitational collapse. These collapsing regions become the seeds of new star formation, producing stellar nurseries with extensive filamentary structure [2]. The life cycle of stars begins with pre-stellar cores which develop into protostars after sufficient gravitational collapse [2]. These early stages of star formation exist in low temperature states ranging from tens to a few hundred Kelvin, which express black body curves peaking in the FIR [5]. In later stages of stellar evolution, a hot stellar core is embedded in a thick molecular cloud. The high energy photons emitted by the core cannot be observed directly as they are largely absorbed or scattered by the surrounding dust and gas [6]. The enveloping molecular cloud will, however, re-radiate this energy in the form of FIR radiation, contributing further to the FIR spectral content of the Universe [6]. Star formation also expresses characteristic structures like accretion disks, and bipolar outflows which by a similar process also radiate most brightly in the IR.

Hubble's law [7], the discovery of which is a foundational piece of evidence for an ex-

panding Universe, states that the further away an object is, the faster it will be travelling away from an observer. As such, a source's spectrum is Doppler shifted to longer wavelengths [8], an effect called redshift, which has significant implications on the study of the early universe. Consequentially, bright hydrogen emission from low metallicity primordial stars, the first stars to ignite in the Universe, may be detectable in the FIR [6]. As such, studying the evolution of the early Universe requires FIR astronomy. Much of the stellar and galactic formation is thought to have occurred in the first half of the Universe's lifetime [9], and therefore much of the total radiant energy has been shifted to longer wavelengths, again further increasing the FIR spectral content of the Universe.

The utility of the FIR should thus be clear in addressing some of the most fundamental questions in astronomy. "How did the Universe originate and what is it made of?", "How were the first luminous objects in the Universe ignited?", "What are the conditions for planet formation?". These are all questions posed by the The European Space Agency (ESA) in their Cosmic Vision programme [10], which can be directly investigated using the FIR. In fact, organic molecules and water have abundant emission in the FIR, and questions about the conditions for the origin of life itself may also be probed by FIR astronomy.

There are three key challenges which exist in FIR astronomy [6]. First, Earth's atmosphere is opaque to the FIR frequencies of interest mostly as a result of water in the atmosphere (Figure 1.1), so in order to get a clear, unobstructed view, a space born instrument is necessary. Second, thermal radiation peaks in the FIR band at very low temperatures ($\sim 10\text{--}100\text{ K}$). As such, FIR astronomy typically requires instruments to be cooled below these temperatures so that it is the source under observation, not the instrument itself, which dominates the radiation incident on the detectors. Third, the spatial resolution of a telescope increases with primary mirror diameter, and decreases with the wavelength of light being observed. Since FIR consists of wavelengths several orders of magnitude greater than optical light, a telescope with a proportionately bigger primary mirror is required to obtain the same spatial resolution.

The highly successful Infrared Astronomical Satellite (IRAS) mission, launched in 1983, significantly advanced the state of FIR astronomy with respect to the challenges mentioned above. IRAS was a space telescope with a primary mirror cooled to ~ 10 K, and focal plain optics cooled to ~ 3 K via a liquid helium cryostat, providing unprecedented sensitivity at the wavelengths under observation (~ 10 – $100 \mu\text{m}$) [11]. The mission goal and landmark achievement of IRAS was to produce an all-sky survey of the FIR Universe. The mission duration was January to November 1983, during this time IRAS surveyed approximately 95% of the sky. Analysis of this dataset led to the discovery of more than 250,000 new IR sources [12].

During the late stages of development of IRAS, a follow-up mission was proposed to observe the variety of new sources from IRAS. This mission was called the Infrared Space Observatory (ISO) [13], and was launched in November 1995. ISO was an observatory class mission allocating the majority of observation time to targets requested by the astronomical community. The space-based telescope and instruments were cooled to 2–8 K by a liquid helium cryostat with enough cryogen to sustain operations until the mission ended in May 1998. ISO improved on the IRAS design by using photometers with higher sensitive while covering a larger bandwidth (2.5– $240 \mu\text{m}$). Additionally, the spectrometers on-board ISO covered a larger bandwidth with 10 to 1000 times the resolving power of IRAS. A number of exciting discoveries have been made with ISO [14], but perhaps the most stunning is the abundance of water in the Universe [15].

Another ground-breaking FIR telescope was the Spitzer Space Telescope named after Lyman Spitzer [16]. *Spitzer* was actively cooled to ~ 5.5 K and made observations at wavelengths from 3.6– $160 \mu\text{m}$. *Spitzer*'s instrument suite consisted of the Infrared Array Camera (IRAC), a four-channel IR imager covering the bandwidth 3.19– $9.34 \mu\text{m}$; the Infrared Spectrograph (IRS), a diffraction grating spectrometer sensitive to 5.2– $19.6 \mu\text{m}$ wavelengths with a resolving power ranging from 64 to 600; the Multiband Infrared Photometer for Spitzer (MIPS), a three-band photometer with central wavelengths 24, 70, and

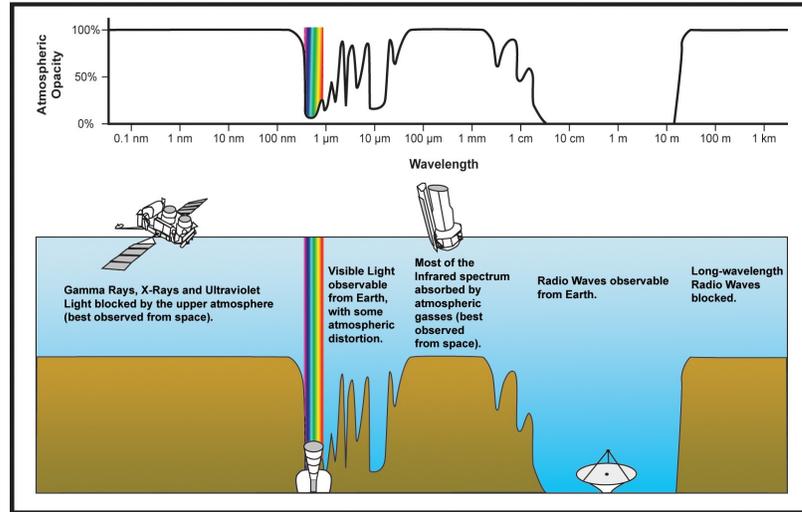


Figure 1.1: Spectral opacity of Earth’s atmosphere at various wavelengths. Image courtesy of The National Aeronautics and Space Administration (NASA).

160 μm . One of *Spitzer*’s mission objectives was to observe planetary structures orbiting distant stars [17]. In 2005 *Spitzer* recorded the first direct detection of an extrasolar planet [18], marking a historical achievement for the mission. The lifetime of most FIR telescopes is limited by the on-board cryogen supply. In the case of *Spitzer*, however, the “cold phase” of the mission lasted from August 2003 to May 2009, and this was followed by a “warm phase”. As of October 2016 *Spitzer* has been engaged in the “beyond phase”, which is on-going at the time of writing this thesis. The goal of this phase is, in part, to identify potential targets suitable for detailed observation by the James Webb Space Telescope [19].

The aforementioned missions have contributed significantly to FIR astronomy, but the quest for more sensitive measurements is ongoing, and the above telescopes do not significantly address the angular resolution limitations that come with observing at longer FIR wavelengths. The next advancement in FIR astronomy instrumentation came from the Herschel Space Observatory, which is the subject of the next section.

1.2 SPIRE and the Herschel Space Observatory

All three challenges described in Section 1.1 are addressed with the latest FIR telescope. The Herschel Space Observatory (*Herschel*; Figure 1.2, [20]) was launched in May 2009 to

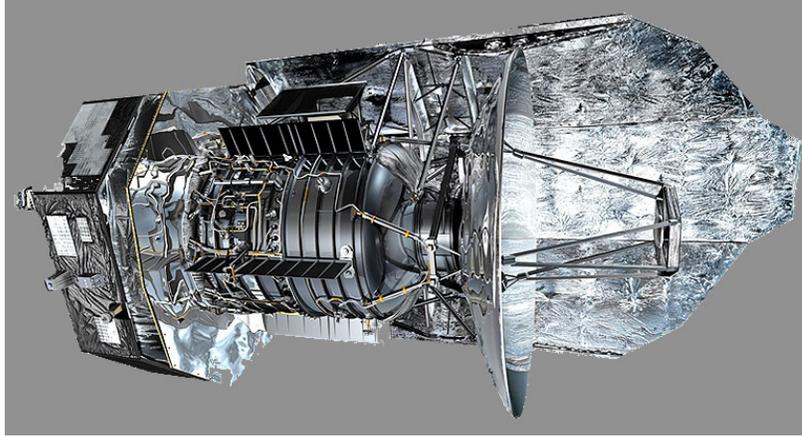


Figure 1.2: An artist's rendition of *Herschel*, courtesy of ESA/AOES Medialab.

obtain an unobstructed view of the FIR Universe in orbit around the Sun-Earth system L2 Lagrangian point. *Herschel* was equipped with a 3.5 m diameter primary mirror, the largest single piece mirror ever launched into space. Though the spatial resolution is still lower than typical optical telescopes, *Herschel* provided a significant improvement over previous IR and FIR space observatories (IRAS 57 cm, [11] 1983; ISO 60 cm, [13] 1995; Spitzer 85 cm, [16] 2003). *Herschel* also employs a heat shield to block additional thermal radiation. This configuration passively cools the telescope to ~ 85 K [20]. Inside *Herschel*, the focal plane optics and bolometer detectors were actively cooled by a closed cycle $^3\text{He}/^4\text{He}$ cryogenic system, to ~ 300 mK for optimal detector sensitivity. As such, *Herschel* was well equipped to confront the challenges of FIR astronomy.

Contained within the *Herschel* payload module was a suite of three instruments. Most important for this thesis is the Spectral and Photometric Imaging REceiver (SPIRE); [21]. SPIRE hosts three photometric arrays with broad-band coverage centered on the wavelengths 250, 350, and $500\ \mu\text{m}$ (Photometer Short Wavelength (PSW), Photometer Medium Wavelength (PMW), and Photometer Long Wavelength (PLW), respectively). The SPIRE instrument also includes an Imaging Fourier Transform Spectrometer (IFTS) consisting of two feedhorn-coupled hexagonal close packed detector arrays with spatial configuration optimized for their respective bands. The SPIRE Short Wavelength Spectrometer array

(SSW), and SPIRE Long Wavelength Spectrometer array (SLW) arrays cover the spectral ranges 194-313 and 303-671 μm , respectively. SPIRE had a maximum spectral resolution of 1.2 GHz which is sufficient to discriminate between significant spectral features in most cases. With the broad spectral coverage and high throughput characteristic of a Fourier Transform Spectrometer (FTS), SPIRE is equipped to provide high quality science observations and data products for a variety of sources.

Herschel's mission concluded in April 2013 after depletion of the on-board cryogen supply. During its 4 year mission, *Herschel* acquired a wealth of science observations with unprecedented spatial resolution and sensitivity now publicly available in the The *Herschel* Science Archive (HSA); [22]. The *Herschel* data-set is a popular one in astronomical community resulting in over 2,000 publications, at the time of writing this thesis, according to the HSA. Though the products available at the HSA are processed to the level of calibrated spectra ready to be used by astronomers, analysis of the data is inherently difficult. In addition to the large volume of data, FTS instruments produce spectra exhibiting a unique line shape which is well approximated as a sinc function [23] (Figure 1.3). The sinc function distributes its energy over a wide spectral range in the form of damped oscillating side-lobes. These side-lobes can interact with and suppress or alter the apparent position of other features. As such, a rigorous line fitting procedure is required to extract accurate spectral feature information.

1.3 My Contributions to SPIRE

SPIRE Automated Spectral Feature Finder

In order to facilitate data mining and assist with preliminary analysis, the SPIRE FTS team developed the SPIRE FTS automated Feature Finder (FF); [24]. I am a significant contributor to the FF both in terms of development and testing. The FF is a sophisticated automated fitting routine which extracts significant spectral features from SPIRE FTS data. Signal-to-Noise Ratio (SNR), as well as a quality parameter which indicate the reliability of

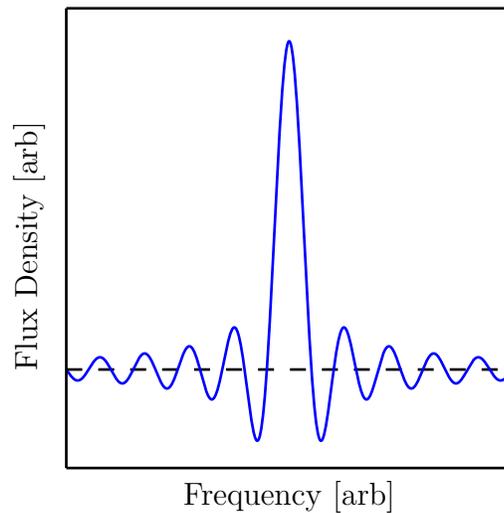


Figure 1.3: FTS sinc line shape with dashed line indicating zero flux density.

the detection, are included for each feature. Postcards showing the result of the FF are also provided alongside the FF products. All the FF products are hosted publicly on the HSA in what is formally called the Spire Automated Feature Extraction CATalogue (SAFECAAT) ¹.

Initially my contribution began with a simple validation of FF results by comparing extracted ^{12}CO rotational emission features to line fits obtained using more dedicated spectral models on Orion observations. Flaws in the FF routine were observed typically around the neighbourhood of 809 GHz, the spectral region surrounding the fine structure neutral carbon emission (see Chapter 6). I then began the task of correcting these flaws which culminated into the work presented in this thesis. The remainder of this section gives a brief overview of this work.

Automated Radial Velocity Estimating Routine

The FF products were extended to include a source radial velocity estimate for each source. The primary way in which this is accomplished is by attempting to identify ^{12}CO rotational emission features within the spectrum being analyzed. ^{12}CO is the second most abundant molecule in the ISM [13], and unlike molecular hydrogen, the most abundant,

¹<https://www.cosmos.esa.int/web/herschel/spire-spectral-feature-catalogue>

^{12}CO has a strong dipole moment resulting in bright rotational emission lines. As such, ^{12}CO is both likely to be present within SPIRE spectra and easily detected by the FF.

Like all linear/diatomic molecules with non-zero dipole moments, adjacent rotational transitions are separated by a near constant frequency. Autonomous ^{12}CO detection is based on searching for characteristic differences between ^{12}CO rotational transitions with an arbitrary line list of features identified within a given observation. Features having the most matches with expected ^{12}CO emission are the most likely ^{12}CO feature candidates. The frequency difference between emission lines is, however, dependent on the source radial velocity. As such, a velocity dependent tolerance is imposed to allow identification within sources with non-zero velocity.

The analysis becomes even more complicated with the introduction of other bright features into the spectrum or features in close proximity to ^{12}CO lines. Developing a reliable automated velocity estimating routine which is applicable to a wide range of sources is a task I put considerable effort into as part of the FF team. The resulting routine exists as a modular implementation suitable for any arbitrary line list with associated SNR within the SPIRE FTS bands, and can be easily generalised for wider application. In addition to the radial velocity itself, the routine outputs associated flagging parameters which have proven to be a robust indicator of the estimate accuracies.

Neutral Carbon Detection

With accurate parameterization of the source radial velocity for SPIRE FTS observations, it is possible to perform more dedicated searches for specific lines. One line that merits such investigation is the neutral carbon fine structure $^3\text{P}_2\text{-}^3\text{P}_1$ [CI] transition at 809 GHz. In a coupled rest frame, [CI] is located only a few resolution elements away from the typically more bright $^{12}\text{CO}(7\text{-}6)$ rotational transition at 806 GHz. The relative difference in amplitude, close spectral proximity, and sinc instrument line shape results in a situation where the [CI] feature is difficult to distinguish from the $^{12}\text{CO}(7\text{-}6)$ emission's positive

side-lobe.

With a velocity estimate based on ^{12}CO rotational emissions, and the assumption that both ^{12}CO and atomic carbon are approximately coupled in the same inertial system, it is possible to predict the expected frequency of the potential [CI] feature. Different spectral models can then be tested to assess the presence of the [CI] feature in spectra under study. This regression based approach relies on modifying fitter tolerances and careful selection of initial estimates for the potential [CI] feature. It was empirically determined that a relative amplitude estimate tied to the more prevalent feature can successfully detect faint [CI] features, while limiting false positives.

A natural extension of successful detection of the $^3\text{P}_2$ - $^3\text{P}_1$ carbon feature is a search for the $^3\text{P}_1$ - $^3\text{P}_0$ fine structure transition of the same atomic species. This feature has a rest frequency emission of 492 GHz, which approaches the noisy low frequency edge of the SLW band, and is typically of lower amplitude than the $^3\text{P}_2$ - $^3\text{P}_1$ transition. Though this extension has not been fully implemented in the current FF release, tests have shown detection of the $^3\text{P}_1$ - $^3\text{P}_0$ feature, in some cases, can be improved using a dedicated search algorithm based on previous [CI] detection.

1.4 Summary

FIR astronomy is crucial for the study of the evolution of the Universe, the formation of stars and planets, and the conditions under which life may have emerged. *Herschel* is the latest FIR telescope which has provided unfettered access to the FIR spectrum, and substantive advances in spatial resolution and sensitivity in the FIR band. On-board *Herschel* was SPIRE, an instrument containing an IFTS which collected a wealth of high quality FIR spectra for a variety of observation targets. To assist in data mining and preliminary analysis of SPIRE spectra, the SPIRE FTS team developed the FF. For the FF, I developed an automated radial velocity estimating routine, and a separate neutral carbon detection routine. My unique contributions have been included in the resulting SAFECAT hosted as part

of the HSA, and are documented in the form of conference and in-preparation publications, in addition to the content of this thesis.

1.5 Thesis Outline

This thesis begins by giving an introduction to Fourier analysis in Chapter 2. An understanding of Fourier analysis is required to interpret SPIRE spectra, and since much of the data processing involved in FTS involves the manipulation of discretized datasets, the discussion is extended to computational algorithms used for discrete Fourier transforms. This chapter also derives many of the most commonly used analytical transforms, and covers useful properties of Fourier transforms with an emphasis on symmetry. Chapter 3 is an introduction to spectroscopy. This chapter provides information on how radiation interacts with matter, and how a polychromatic radiation field can be decomposed into its constituent frequencies using FTS. The focus is on black body radiation, along with rotational and electron transition emissions, the two emission processes which produce the spectral features I have studied in depth. The information conveyed is at a level sufficient to understand the amplitude and emission frequencies of these processes, and the complex energy structure of photoelectronic atoms. Derivations for commonly encountered line shapes that result from various physical processes is presented in compliment to emission mechanisms. This is followed by a section which derives the theory of FTS, and is supplemented with a discussion on noise considerations when employing this spectroscopic technique.

My person contributions begin with Chapter 4 which introduces *Herschel* and the SPIRE instrument in more detail, including its observation modes, calibrations schemes, and publicly available products. The FF routine is then outlined including its data products and relevant validation testing. Chapter 5 presents the FF source velocity estimating routine, giving a full description of its operations. The routine is then validated using both real and synthetic datasets. Results, as they pertain to the FF public release, are then discussed. Chapter 6 presents a detailed outline of the neutral carbon check, a dedicated search for

the 3P_2 - 3P_1 neutral carbon fine structure transition employed by the FF. This is followed by a validation of the routine, and a summary of the results as they pertain to the public FF catalogue. Concluding remarks are provided in Chapter 7, with mapping observation postcards and preliminary neutral carbon 3P_1 - 3P_0 emission detection postcards included in appendix B and C, respectively. A demonstration of the FF source radial velocity estimating routine is given in appendix A.

Chapter 2

Fourier Analysis

Fourier analysis traces its origin back to 1807 when Jean-Baptiste Joseph Fourier (1768 - 1830), investigated general solutions to the partial differential heat equation. As outlined in *Mémoire sur la propagation de la chaleur dans les corps solides* [25], Fourier's general solution consisted of the linear superposition of simple sinusoidal waves. The key result of this discovery was that “any” function can be expressed as a combination of an orthogonal basis set of simple sine and cosine vectors. This has immediate implications for spectroscopy where the goal is to decompose an input signal into its constituent frequencies. Though it took well over a century to develop the technology required for practical application of Fourier analysis on large datasets, it is now an integral part of modern spectroscopy, and the central spectroscopic method discussed in this thesis. As such, it is critical to construct a solid theoretical foundation to understand the spectroscopic techniques presented in Chapter 3.

Section 2.1 defines the Fourier series, which is then extended to the continuous Fourier transform. The Dirac delta function is then defined in Fourier transform notation, followed by the derivation of a sample of commonly employed Fourier transforms in Fourier Transform Spectroscopy (FTS). Section 2.2 derives the symmetry properties, and significant theorems of Fourier transforms. Discretization is then discussed in Section 2.3, which focuses on the discrete, and fast Fourier transforms. Conditions for proper sampling of continuous signals is also explored.

2.1 Definitions

2.1.1 Fourier Series

A periodic function, $h(t)$, which varies in time, t , can be expressed as a Fourier series expansion given by [26]

$$h(t) = \frac{1}{2}a_0 + \sum_{n=1}^{\infty} [a_n \cos(n2\pi f_0 t) + b_n \sin(n2\pi f_0 t)] , \quad (2.1)$$

where

$$f_0 = \frac{1}{T} = \frac{\omega_0}{2\pi} . \quad (2.2)$$

T and f_0 are the period and fundamental frequency of $h(t)$, respectively. The functions $\cos(n2\pi f_0 t)$ and $\sin(n2\pi f_0 t)$ constitute an orthogonal set of basis vectors which are weighted by the real valued coefficients a_n and b_n , respectively. These constants are given by

$$\begin{aligned} a_n &= \frac{2}{T} \int_{-T/2}^{T/2} h(t) \cos(n2\pi f_0 t) dt , \\ b_n &= \frac{2}{T} \int_{-T/2}^{T/2} h(t) \sin(n2\pi f_0 t) dt . \end{aligned} \quad (2.3)$$

The equivalent complex Fourier series can be constructed using Euler's formula $e^{iz} = \cos(z) + i \sin(z)$ [27]

$$h(t) = \sum_{n=-\infty}^{\infty} c_n e^{in2\pi f_0 t} , \quad (2.4)$$

where

$$c_{\pm n} = \frac{1}{T} \int_{-T/2}^{T/2} h(t) e^{\mp in2\pi f_0 t} dt = \frac{1}{2}(a_n \mp ib_n) . \quad (2.5)$$

In order to express a function as a Fourier series expansion (Equations 2.1, 2.4), the function $h(t)$ must satisfy the following Dirichlet conditions [26]: $h(t)$ must be defined over the range $(-\infty, \infty)$ and is periodic, $h(t)$ must have no more than a finite number of discontinuities and extrema over one period and $h(t)$ must converge absolutely over one period (i.e., $\int_{-T/2}^{T/2} |h(t)| dt < \infty$).

2.1.2 Fourier Integral

Fourier series are well suited for expressing functions which are periodic over a finite interval. However, more general functions required a slightly different treatment. Combining Equations 2.4 and 2.5 gives

$$h(t) = \sum_{n=-\infty}^{\infty} \left\{ \frac{1}{T} \int_{-T/2}^{T/2} h(t') e^{-in2\pi f_0 t'} dt' \right\} e^{in2\pi f_0 t} . \quad (2.6)$$

A general function requires no periodicity, or alternatively a period of infinite extent. Additionally the sum of discrete frequencies must be replaced by an integral over a continuum of frequencies. Using the substitutions $nf_0 = f$, and $\lim_{T \rightarrow \infty} 1/T = df \rightarrow 0$, equation 2.6 becomes

$$\begin{aligned} h(t) &= \int_{-\infty}^{\infty} \left\{ \int_{-\infty}^{\infty} h(t') e^{-i2\pi f t'} dt' \right\} e^{i2\pi f t} df \\ &= \int_{-\infty}^{\infty} H(f) e^{i2\pi f t} df . \end{aligned} \quad (2.7)$$

The above equation expresses the general function $h(t)$ as the integration of an infinite number of sinusoids of frequency f , with amplitude $H(f)$.

Equation 2.7 is none other than the **inverse Fourier transform**

$$h(t) = \int_{-\infty}^{\infty} H(f) e^{i2\pi f t} df = \mathcal{F}^{-1}\{H(f)\} , \quad (2.8)$$

with the **forward Fourier transform** given by

$$H(f) = \int_{-\infty}^{\infty} h(t) e^{-i2\pi f t} dt = \mathcal{F}\{h(t)\} , \quad (2.9)$$

The functions $h(t)$ and $H(f)$ constitute a Fourier transform pair. Though the notation used implies the use of time for t , and frequency for f , all that is required for the conjugate variables t and f is that their units be reciprocals of one another. For example the Fourier transform of a signal recorded in time will produce a temporal frequency spectrum (i.e., [s]

\leftrightarrow [Hz]), whereas the Fourier transform of a signal record in space will produce a spatial frequency spectrum (i.e., [cm] \leftrightarrow [cm⁻¹]).

Slight variations on Equations 2.8 and 2.9 are often encountered in the literature. The variation results from a change of variables (e.g., $f = \omega/2\pi$), which requires the inclusion of a normalization constant of $1/2\pi$ in the inverse Fourier transform. Alternatively a factor of $1/\sqrt{2\pi}$ can be included in both the forward and inverse Fourier transforms. The format used above, however, is more simple and results in less confusion, and will be employed throughout this thesis unless otherwise indicated. In addition, the forward and reverse Fourier transform functions differ primarily by the sign in the exponent. As such what is considered the forward and inverse Fourier transform is a matter of convention which is not strictly adhered to in all texts. This thesis will exclusively use the definitions as provided in Equations 2.8, and 2.9.

2.1.3 Dirac Delta Function

The mathematical object known as the Dirac delta function ($\delta(x)$; [28]) consists of a distribution of real values which is everywhere zero except at the origin where it has an infinite magnitude. It is also defined to have infinitesimal width such that the area under the distribution is equal to unity. That is,

$$\delta(x) = \begin{cases} 0, & x \neq 0 \\ \infty, & x = 0 \end{cases} \quad (2.10)$$

with

$$\int_{-\infty}^{\infty} \delta(x) dx = 1. \quad (2.11)$$

Consider the product of an ordinary function $h(x)$ and $\delta(x - x_0)$. In such a situation the result will be zero everywhere on the number line except when $x = x_0$, at which point the

product is infinite. Since the integrated area of $\delta(x)$ is unity, integrating this product gives

$$\int_{-\infty}^{\infty} h(x)\delta(x-x_0)dx = h(x_0) . \quad (2.12)$$

This is referred to as the sifting property of the delta function [29]. It is useful for mathematically representing point sources and pulses in physics, and in the case of spectroscopy, is used to represent the spectrum of a monochromatic source.

The inverse Fourier transform of the delta function is

$$\mathcal{F}^{-1}\{\delta(f-f_0)\} = \int_{-\infty}^{\infty} \delta(f-f_0)e^{-i2\pi ft}df = e^{-i2\pi f_0 t} . \quad (2.13)$$

This result can be used to provide a useful Fourier transform based definition of the Dirac delta function

$$\mathcal{F}\{e^{-i2\pi f_0 t}\} = \delta(f-f_0) . \quad (2.14)$$

The implications of this result will be discussed further in § 2.1.4.4.

2.1.4 Useful Transforms

Some Fourier transforms are regularly encountered in spectroscopy and the most common ones will be derived here. A more complete discussion on the spectroscopic significance of these transforms is presented in Chapter 3.

2.1.4.1 Boxcar

The boxcar function is a piecewise defined function such that

$$\Pi(x) = \begin{cases} A, & |x| \leq L \\ 0, & |x| > L . \end{cases} \quad (2.15)$$

The Fourier transform is the

$$\begin{aligned}\mathcal{F}\{\Pi(t)\} &= \int_{-\infty}^{\infty} A e^{-i2\pi ft} dt \\ &= A \int_{-L}^L e^{-i2\pi ft} dt .\end{aligned}$$

The exponential can be expanded using Euler's equation such that

$$\begin{aligned}\mathcal{F}\{\Pi(t)\} &= A \int_{-\infty}^{\infty} [\cos(2\pi ft) - i \sin(2\pi ft)] dt \\ &= \frac{A}{2\pi f} [\sin(2\pi ft) + i \cos(2\pi ft)]_{-L}^L \\ &= \frac{A}{\pi f} \sin(2\pi fL) .\end{aligned}$$

The cardinal sine function, or sinc function is defined as $\sin(arg)/arg$. With this notation, the last line can be recast to

$$\mathcal{F}\{\Pi(t)\} = 2LA \operatorname{sinc}(2\pi fL) . \quad (2.16)$$

A graphical representation of this result is shown in Figure 2.1. This shape is very common in FTS, since all interferograms are theoretically infinite in extent and are naturally truncated by the finite scan range of physical instruments. As such, recorded signals are in effect multiplied by a boxcar function (Section 3.4).

2.1.4.2 Gaussian

The Gaussian function has the form

$$G(t) = \sqrt{\frac{A}{\pi}} e^{-At^2} . \quad (2.17)$$

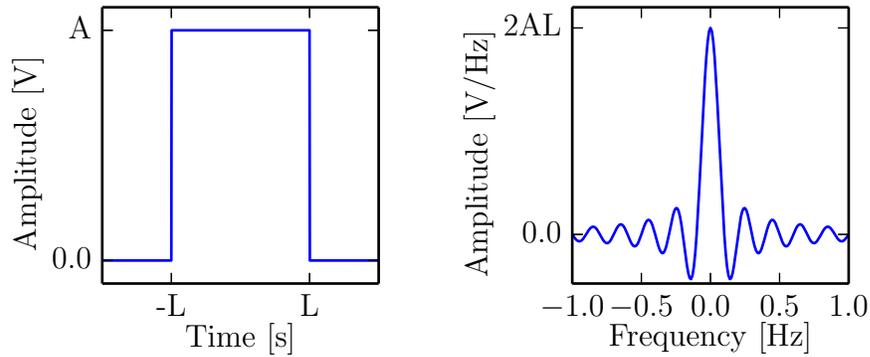


Figure 2.1: The Fourier transform of a boxcar function with amplitude A over the range $[-L, L]$ (left), is a sinc function of amplitude $2AL$ and $\text{FWHM} = 1.2067/(2AL)$.

Amplitude and exponential coefficients can vary from the specific relation given above; these values are only chosen for convenience. The Fourier transform is then

$$\begin{aligned}\mathcal{F}\{G(t)\} &= \int_{-\infty}^{\infty} \sqrt{\frac{A}{\pi}} e^{-At^2} e^{-i2\pi ft} dt \\ &= \sqrt{\frac{A}{\pi}} \int_{-\infty}^{\infty} e^{-At^2 - i2\pi ft} dt.\end{aligned}$$

The quadratic equation in the exponent is more easily dealt with upon completing the square giving

$$= \sqrt{\frac{A}{\pi}} \int_{-\infty}^{\infty} e^{-A(\gamma)^2 - n} d\gamma,$$

where $\gamma = t + i\pi f/A$, and $n = (\pi f)^2/A$. The constant exponential, e^{-n} , can be taken outside the integral leaving only the integral of a Gaussian function. The analytical solution of this integral is $\sqrt{\pi/A}$ which is easily shown using polar coordinates [30]. After cancellation, and substitution of the original variables, the result is

$$\mathcal{F}\{G(t)\} = e^{-\pi^2 f^2/A} \quad (2.18)$$

Thus the Fourier transform of a Gaussian is another Gaussian with a modified amplitude and width. The result is shown in Figure 2.2. The Gaussian distribution is a common

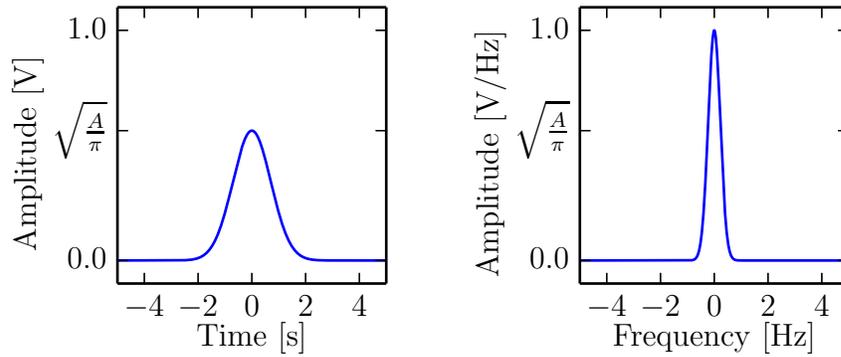


Figure 2.2: The Fourier transform of a Gaussian is another Gaussian with varying width and amplitude.

spectral profile, which results from a mechanism known as Doppler broadening [31]. It is also a popular apodizing function [32], used to smooth line shapes.

2.1.4.3 Exponential

The exponential decay function has the form

$$E(t) = E_0 e^{-2\pi A t}, \quad (2.19)$$

which begins at some real value E_0 at $t = 0$, and decays asymptotically to zero at a rate determined by the time constant $1/(2\pi A)$. The Fourier transform of the exponential decay function produces the Lorentzian function, a common line shape encountered in spectroscopy, which is discussed in more detail in Section 3.2. To derive the Fourier transform of Equation 2.19, the function is mirrored across the origin producing a symmetric function. This operation simplifies the arithmetic while preserving the physical significance of the function for time greater than or equal to zero. This operation can be represented mathematically using the symmetric exponential decay function

$$E(t) = e^{-2\pi A |t|} \quad (2.20)$$

with the absolute value of the independent variable, t , resulting in an exponential decay on either side of the coordinate origin, and a discontinuity of the function's derivative at this point. Like the Gaussian function, the factors in the exponent, and amplitude (in this case 1), have been chosen for convenience. The Fourier transform is then

$$\begin{aligned}\mathcal{F}\{E(t)\} &= \int_{-\infty}^{\infty} e^{-2\pi A|t|} e^{-i2\pi ft} dt \\ &= \int_{-\infty}^0 e^{2\pi t(A-if)} dt + \int_0^{\infty} e^{-2\pi t(A+if)} dt\end{aligned}$$

The integral has been broken up about the origin so that t can be replaced with $-t$ on the negative interval, preserving the absolute value over the whole domain. This result is readily integrated giving

$$= \left. \frac{e^{2\pi t(A-if)}}{2\pi(A-if)} \right|_{-\infty}^0 + \left. \frac{-e^{-2\pi t(A+if)}}{2\pi(A+if)} \right|_0^{\infty}$$

After evaluating this expression at the limits, one obtains the final result

$$\mathcal{F}\{E(t)\} = \frac{A}{\pi} \frac{1}{A^2 + f^2}. \quad (2.21)$$

That is the Fourier transform of the exponential function is the Lorentzian distribution [33]. This result is shown in Figure 2.3 where $2A$ corresponds to the full width at half maximum. The Lorentzian function is yet another commonly encountered spectral profile that results from the intrinsic uncertainty of energy levels in quantum systems [31]. Collisions between radiative particles can also produce this line shape, which is why this spectral profile is also associated with collision, or pressure broadening [31].

2.1.4.4 Sinusoid

A sinusoid with amplitude, A , and constant frequency, f_0 , is given by

$$\Psi(t)_{f_0} = A \cos(2\pi f_0 t) \quad (2.22)$$

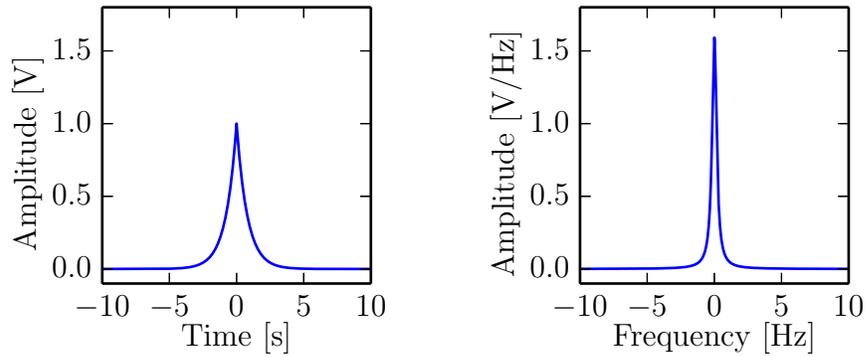


Figure 2.3: The Fourier transform of an exponential function is a Lorentzian distribution.

The Fourier transform is then

$$\begin{aligned}
 \mathcal{F}\{\Psi(t)_{f_0}\} &= \int_{-\infty}^{\infty} A \cos(2\pi f_0 t) e^{-i2\pi f t} dt \\
 &= \frac{A}{2} \int_{-\infty}^{\infty} [e^{i2\pi f_0 t} + e^{-i2\pi f_0 t}] e^{-i2\pi f t} dt
 \end{aligned} \tag{2.23}$$

In the second line, $\cos(\theta)$ has been replaced by the exponential equivalent. Comparing Equation 2.23 to Equation 2.14, it is clear that the second line consists of two delta functions centered on f_0 and $-f_0$. That is the Fourier transform of a sinusoid with amplitude, A , and constant frequency, f_0 , is

$$\mathcal{F}\{\Psi(t)_{f_0}\} = \frac{A}{2} [\delta(f - f_0) + \delta(f + f_0)] \tag{2.24}$$

This example demonstrates the utility of Fourier transforms. The input signal is decomposed into its constituent frequencies producing a spectrum at the output of the transform. In this case there are two sharp peaks at the frequency of the input signal with the energy distributed equally between the positive and negative frequency domains. This result is shown in Figure 2.4. Clearly, this is the spectral profile of a monochromatic source.

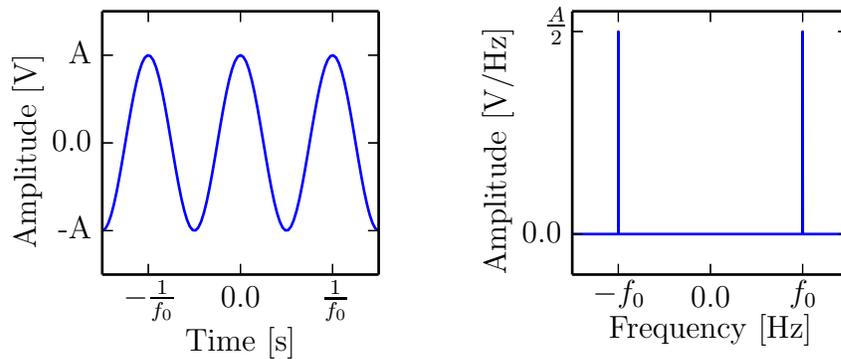


Figure 2.4: The Fourier transform of a sinusoid is two sharply peaked delta-like functions at the positive and negative frequency of the sinusoid with. With proper normalization, the amplitude of each peak is half the amplitude of the sinusoid.

2.2 Properties of Fourier Transforms

2.2.1 Symmetry

Fourier transforms map complex data onto real and imaginary components of the reciprocal domain. How the energy is distributed into real and imaginary components is dependent on the symmetry of the input function.

The cosine function is an even function (i.e., $h(x) = h(-x)$), which is symmetric across the origin. The Fourier transform of the cosine function was derived in Section 2.1.4.4, with the result consisting of two real value delta functions at the positive and negative values of the cosine function's frequency. There is no imaginary component to this transform, and so we can conclude that the Fourier transform of a real valued even function is another real valued even function (Figure 2.5).

The sine function is an example of an odd function (i.e., $h(-x) = -h(x)$), which is antisymmetric across the origin. The Fourier transform of the sine function can be derived in the same way as the Fourier transform of the cosine function. The final result is given by

$$\mathcal{F}\{A \sin(2\pi f_0 t)\} = \pm i \frac{A}{2} \delta(f \mp f_0) \quad (2.25)$$

Similar to the transform of the cosine function, the transform of the sine function consists of

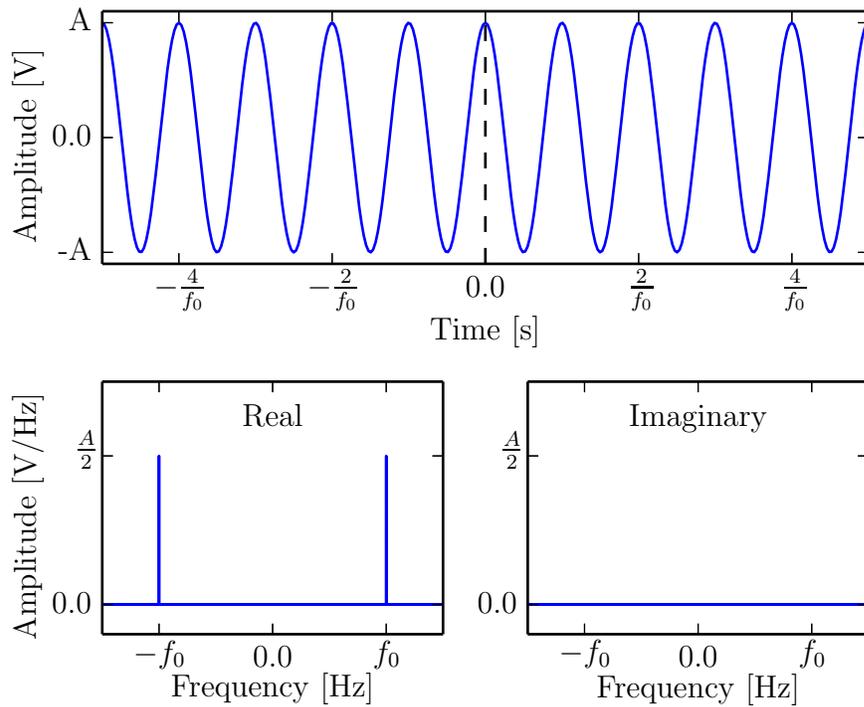


Figure 2.5: The Fourier transform of a real valued even function produces yet another real valued even function. This result is shown for a cosine function where the Fourier transform consists of a symmetric distribution in the real domain, with no imaginary component.

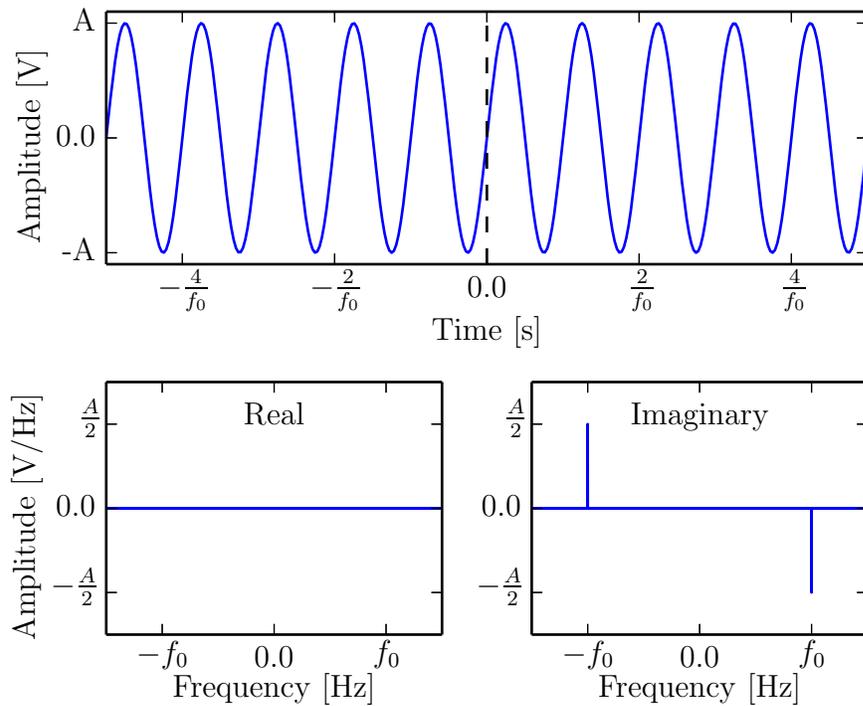


Figure 2.6: The Fourier transform of a real value odd function is an imaginary odd function. This result is shown for a sine function where the Fourier transform consists of an antisymmetric distribution in the imaginary domain, with no real components.

two delta functions at the positive and negative frequencies of the sine function. However, for the sine function these delta functions are antisymmetric and reside in the imaginary domain. That is the Fourier transform of a real valued odd function is an imaginary odd function (Figure 2.6).

Any arbitrary asymmetric function, which is neither symmetric nor antisymmetric, defined on some finite interval can be expressed as the sum of even and odd functions. These even and odd functions can be constituted by sine and cosine functions (a Fourier series). Consider an asymmetric sinusoid $h_{\text{Asym}}(x)$, composed of a linear combination of a sine and cosine function

$$h_{\text{Asym}}(x) = \alpha \cos(x) + \beta \sin(x) \quad (2.26)$$

With the use of the trigonometric identity $\cos(x \pm \phi) = \cos(x)\cos(\phi) \mp \sin(x)\sin(\phi)$, this

function is equivalent to

$$h_{\text{Asym}}(x) = \gamma \cos(x + \phi) \quad (2.27)$$

with $\alpha = \gamma \cos(\phi)$, and $\beta = \gamma \sin(\phi)$. The angular offset ϕ , is referred to as the phase, which has important implications for Fourier transforms. Note that as the phase goes to zero, so too does the sine term in Equation 2.26, along with the imaginary component of the Fourier transform. That is, the phase is a measure of asymmetry of the data being Fourier transformed with the extrema values 0 indicating complete symmetry, and $\pi/2$ indicating complete antisymmetry.

As an example, consider the asymmetric function

$$h(t) = A \cos\left(2\pi f_0 t + \frac{\pi}{3}\right) \quad (2.28)$$

This function and its Fourier transform are shown in Figure 2.7. The transform gives a complex result with peaks in both the real and imaginary domain at the frequency of the sinusoid. A complex number can be represented as a vector in the complex plane, with the real and imaginary components making up an orthogonal basis. With this representation, the phase of the complex number is calculated using

$$\phi = \arctan\left(\frac{Im}{Re}\right) \quad (2.29)$$

with Im and Re indicating the imaginary and real components, respectively. In the above example, if $A = 1.0$, then at f_0 the imaginary component is $\sqrt{3}/4$, and the real component is $1/4$. Computing the phase then gives $\pi/3$, the value which would be expected from Equation 2.28.

A full set of generalized symmetry properties of Fourier transforms are given in table 2.1.

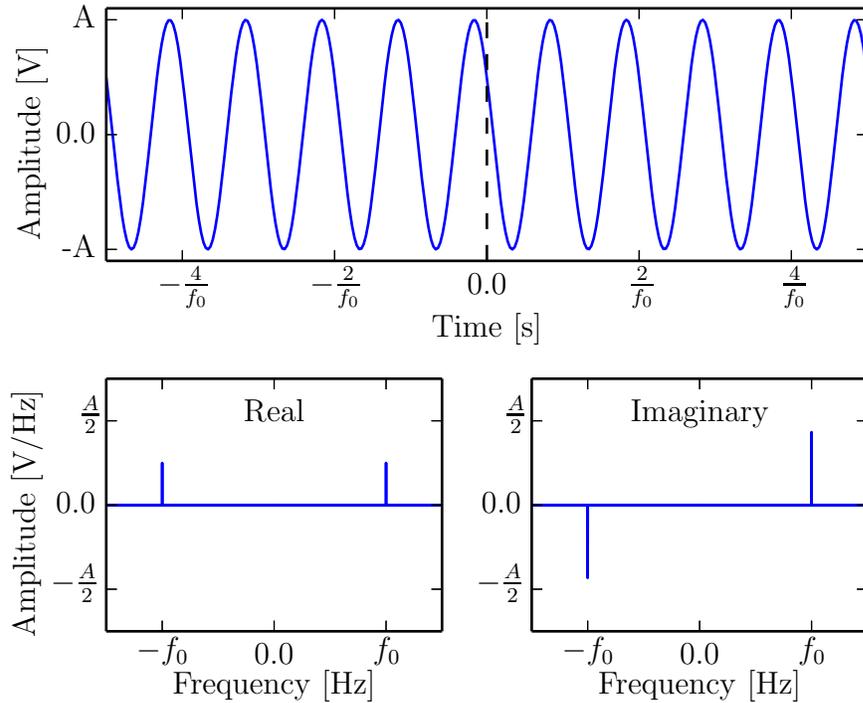


Figure 2.7: The Fourier transform of an asymmetric function is partitioned between real and imaginary components. The real component is even, and the imaginary component is odd.

Table 2.1: Generalized symmetry properties of Fourier transform pairs [34].

$h(t)$		$H(f)$	
Real	Imaginary	Real	Imaginary
even	0	even	0
odd	0	0	odd
0	even	0	even
0	odd	odd	0
asymmetric	0	even	odd
0	asymmetric	odd	even
even	odd	asymmetric	0
odd	even	0	asymmetric
even	even	even	even
odd	odd	odd	odd
asymmetric	asymmetric	asymmetric	asymmetric

2.2.2 Theorems

2.2.2.1 Superposition

The superposition principle states that when two distinct signals, $f(t)$ and $g(t)$, interact, the resulting signal will be equivalent to the algebraic sum of the two individual signals. This property is preserved in the reciprocal domain of the Fourier transform. That is, the Fourier transform of the sum of two signals is equal to the sum of the individual Fourier transforms of the two signals. This theorem can be proven as follows

$$\begin{aligned}\mathcal{F}\{af(t) + bg(t)\} &= \int_{-\infty}^{\infty} (af(t) + bg(t))e^{-i2\pi ft} dt \\ &= a \int_{-\infty}^{\infty} f(t)e^{-i2\pi ft} dt + b \int_{-\infty}^{\infty} g(t)e^{-i2\pi ft} dt \\ &= a\mathcal{F}\{f(t)\} + b\mathcal{F}\{g(t)\}\end{aligned}\quad (2.30)$$

In Equation 2.30, a and b are any complex constants. Thus, it is clear that

$$\mathcal{F}\{af(t) + bg(t)\} = a\mathcal{F}\{f(t)\} + b\mathcal{F}\{g(t)\} . \quad (2.31)$$

It is clear from Equation 2.31 that Fourier transforms express the property of linearity. This property is particularly useful in the context of spectroscopy as it allows the simultaneous processing of multiple frequencies without the need for any frequency specific corrections as a result of the transform itself.

2.2.2.2 Convolution

Convolution is a mathematical process in which two input functions are modified to produce a third output function which is a blended version of the two inputs. Convolution is defined as

$$f(z) = h(z) * g(z) = \int_{-\infty}^{\infty} h(z')g(z - z')dz' \quad (2.32)$$

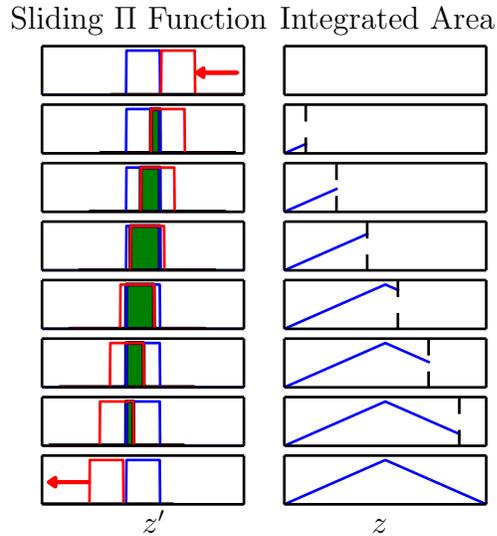


Figure 2.8: A visual representation of convolution. In the left panels is two boxcar functions with unit amplitude. The convolution kernel (red), is slid across the input function (blue), while at each point the integrated area between the two functions (green), is calculated and recoded as an element in the convolution output (right panels). In this case, the integrated area of the overlapping regions is equivalent to the integrated area of the product of the two functions.

where $h(z)$ and $g(z)$ are in general functions of a complex variable z . $h(z)$ corresponds to the initial input into the system, while $g(z)$ is the function that you are convolving the input with, and is often referred to as the convolution kernel. In effect, convolution is the process of sliding the kernel across the input function and at each point calculating the integrated area of their product. This is represented visually in Figure 2.8 for two boxcar functions.

Convolution is a useful tool for describing the physical output of optical systems (Chapter 3), and is regularly employed in graphics manipulation. It also has important properties in the context of Fourier analysis.

The Fourier transform of Equation 2.32 is

$$\begin{aligned} \mathcal{F}\{h(z') * g(z')\} &= \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} h(z')g(z - z')dz' \right] e^{-i2\pi z\sigma} dz \\ &= \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} g(z - z')e^{-i2\pi z\sigma} dz \right] h(z')dz' \end{aligned}$$

In this equation, the arbitrary conjugate pairs z and $\sigma = 1/z$ have been used to preserve

consistency with Equation 2.32. Since both integrals in the above equations must converge for physical systems, Fubini's theorem [35] has been invoked to switch the integration order in the last line. Making the substitution $\gamma = z - z'$, with $d\gamma = dz$, the equation becomes

$$\begin{aligned} &= \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} g(\gamma) e^{-i2\pi(z'+\gamma)\sigma} d\gamma \right] h(z') dz' \\ &= \int_{-\infty}^{\infty} h(z') e^{-i2\pi z'\sigma} dz' \int_{-\infty}^{\infty} g(\gamma) e^{i2\pi\gamma\sigma} d\gamma \end{aligned}$$

The last line is simply the Fourier transform of the original functions. This then proves the convolution theorem. In summary, the Fourier transform of the convolution of two functions is equal to the product of the Fourier transform of the two functions individually

$$\mathcal{F}\{h(z') * g(z')\} = \mathcal{F}\{h(z')\} \times \mathcal{F}\{g(z')\}. \quad (2.33)$$

In the context of Fourier spectroscopy, this theorem is useful for describing the line shapes of sources which are subject to different broadening mechanisms, and the coupling of source intensity distributions with finite limitations of physical instruments (Chapter 3). It is also a convenient tool for determining the Fourier transform of more complicated functions. For example, we see the convolution of two boxcar functions is a triangle function (Figure 2.8), and we know the Fourier transform of a boxcar function is a sinc function (Equation 2.16). Using the convolution theorem we can then conclude the Fourier transform of a triangle function is equal to the sinc^2 function.

2.2.2.3 Parseval's Theorem

Parseval's theorem was originally proven by Marc-Antonie Parseval in 1799 [36]. It states that any two square integrable functions $g(t)$ and $h(t)$, which are in general complex valued with Fourier transforms $G(f)$ and $H(f)$, respectively, obey the relationship

$$\int_{-\infty}^{\infty} g(t)^* h(t) dt = \int_{-\infty}^{\infty} G(f)^* H(f) df \quad (2.34)$$

This theorem was later applied in the context of energy conservation by John William Strutt [37]. Energy conservation between reciprocal domains is the most common application of Parseval's theorem in Fourier analysis.

To prove this relationship, the inverse Fourier transform definitions of $g(t)$ and $h(t)$ are substituted into the left-hand-side of Equation 2.34.

$$\begin{aligned} \int_{-\infty}^{\infty} g(t)^* h(t) dt &= \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} G(f')^* e^{-i2\pi f' t} df' \int_{-\infty}^{\infty} H(f) e^{i2\pi f t} df \right] dt \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G(f')^* H(f) \left(\int_{-\infty}^{\infty} e^{i2\pi t(f-f')} dt \right) df df' \end{aligned}$$

Taking the complex conjugate results in a negative sign in the exponent associated with the $g(t)^*$ inverse Fourier transform. The quantity in parenthesis can be identified as the definition of the delta function (Equation 2.14), and the equation becomes

$$\begin{aligned} &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} G(f')^* H(f) \delta(f' - f) df df' \\ &= \int_{-\infty}^{\infty} G(f)^* H(f) df \end{aligned}$$

The final line is equal to the right-hand-side of equation 2.34, concluding the proof.

In addition to conservation of energy, Parseval's theorem serves as a convenient check on the accuracy of a Fourier transform. If the theorem does not hold, this indicates an error in the transformed data, the units used, or potentially neglected normalization factors.

Linearity, the convolution theorem, and Parseval's theorem are the most relevant property and theorems for the purpose of this thesis. Table 2.2 summarises the full set of Fourier transform properties and theorems.

2.3 Discretization

The Fourier transform is a mathematical operation which acts on a function defined at every value of the independent variable. It is defined in terms of an integral which sums the function over infinitesimal steps of the independent variable, and is thus an operation

Table 2.2: Properties and theorems of Fourier transforms [38, 39]

Theorem	$h(t)$	$H(f)$
Superposition	$h_1(t) + h_2(t)$	$H_1(f) + H_2(f)$
Scalar multiplication	$kh(t)$	$kH(f)$
Autocorrelation	$h^*(t) * h(t)$	$ H(f) ^2$
Parseval's	$\int_{-\infty}^{\infty} h(t) ^2 dt$	$\int_{-\infty}^{\infty} H(f) ^2 df$
Reversal	$h(-t)$	$H(-f)$
Symmetry	$H(t)$	$2\pi h(-f)$
Scaling (a real)	$h(at)$	$\frac{1}{ a } H\left(\frac{f}{a}\right)$
Time shift	$h(t - t_0)$	$H(f)e^{-i2\pi ft_0}$
Frequency shift (f_0 real)	$h(t)e^{i2\pi f_0 t}$	$H(f - f_0)$
Time convolution	$h_1(t) * h_2(t)$	$H_1(f)H_2(f)$
Frequency convolution	$h_1(t)h_2(t)$	$H_1(f) * H_2(f)$
Modulation	$h(t) \cos(2\pi f_0 t)$	$\frac{1}{2}[H(f - f_0) + H(f + f_0)]$
Differentiation	$\frac{d^n h(t)}{dt^n}$	$(i2\pi f)^n H(f)$
Integration	$\int_{-\infty}^t h(t') dt'$	$\frac{H(f)}{i2\pi f} + \pi H(0)\delta(f)$

with an intrinsic continuous nature. As such Fourier transforms are not well suited for the discrete sampling and digitized intensity which is used in many practical applications. Since computers encode information in digital form, a discretized approximation of the Fourier transform must be developed for effective implementation in modern computing.

2.3.1 Discrete Fourier Transform

Assume the function $h(t)$ is discretely sampled at N equally spaced, discrete times t_k , such that

$$f_k = h(t_k), \quad t_k = k\Delta t, \quad k = 0, 1, \dots, N-1 \quad (2.35)$$

For simplicity, assume N is even so that $N/2$ is a whole number, and that Δt is a constant value representing the difference in time between each sample point.

For reasons that will be discussed shortly, the sampling frequency $f_s = 1/\Delta t$, must be twice the maximum frequency component f_M (i.e., $2f_M = 1/\Delta t$), of the discretely sampled function in order to prevent information loss [40]. With this in mind, the discrete frequencies are defined as

$$f_n = \frac{n}{N\Delta t}, \quad n = -\frac{N}{2}, \dots, -1, 0, 1, \dots, \frac{N}{2} - 1 \quad (2.36)$$

These frequencies will make up the independent variable axis in the reciprocal Fourier domain. The zero order frequency will correspond to $n = 0$, with positive and negative frequencies having $n > 0$ and $n < 0$, respectively. Note that when $n = -N/2$, $2f_{-N/2} = 1/\Delta t = -2f_M$, in fulfilment of the sampling criteria mentioned above. However, the positive frequency components have a maximum index of $N/2 - 1$, which means the resulting Fourier transform will only be defined on the interval $[-f_M, f_M - df]$.

With these values now having discrete representations, we can return to the definition

of the Fourier transform (Equation 2.9) and construct a discrete approximation.

$$\begin{aligned} H(f) &= \int_{-\infty}^{\infty} h(t) e^{-i2\pi ft} dt \\ &\approx \sum_{k=0}^{N-1} h_k e^{-i2\pi f_n t_k} \Delta t = \Delta t \sum_{k=0}^{N-1} h_k e^{-i2\pi kn/N} \end{aligned} \quad (2.37)$$

This result gives the definition of the Discrete Fourier Transform (DFT)

$$H_n = \sum_{k=0}^{N-1} h_k e^{-i2\pi kn/N} \quad (2.38)$$

The DFT is periodic in n with a period of N , thus $H_n = H_{n \pm N}$. As previously indicated, $n = -N/2$ corresponds to the negated maximum frequency component of H_n . As a result of this periodicity, $H_{-N/2} = H_{N/2}$, and the value of the DFT at $-f_M$ is equivalent to the missing value at $+f_M$. This property must be kept in mind when working with DFTs, and when interpreting the literature which often assumes the existence of this missing data point.

The discretized analogue of the inverse Fourier transform is given by

$$h_k = \frac{1}{N} \sum_{n=0}^{N-1} H_n e^{i2\pi kn/N} \quad (2.39)$$

with the key differences being the sign of the exponential, and the additional factor of $1/N$. The factor of $1/N$ is included so that $\mathcal{F}^{-1}\{\mathcal{F}\{h(t)\}\} = h(t)$. All of the symmetry properties, and theorems which apply to the Fourier transform pairs hold for their discrete counterparts. The same caveats with respect to the factors of $1/2\pi$ and units hold as well.

2.3.2 Fast Fourier Transform

A significant drawback of directly computing the DFT, as given in Equation 2.38, is computation time. This equation can be represented as

$$H_n = \sum_{k=0}^{N-1} W^{nk} h_k \quad (2.40)$$

The sum is taken over what is structurally the matrix multiplication of the $N \times N$ matrix W , and the vector h_k . The most significant computational contribution of this process consists of the N^2 complex multiplications required to produce the vector H_n . Thus the number of operations required for direct computation of the DFT is proportionate to N^2 . Roughly speaking, the computation of a DFT containing 10 million points with a 2.0 GHz processor would take ~ 14 hrs. Luckily computational efficiency can be improved using clever algorithms.

The aptly named Fast Fourier Transform (FFT), became popularized by Cooley and Tukey in the 1960's [41], however, efficient methods of calculating the DFT date back to as early as Gauss in 1805 [42]. Effectively, the FFT exploits factorization to eliminate repeated calculations in the DFT. The method is based on the work of Danielson and Lanczos in the context of X-ray scattering [43]. They showed that an N point DFT can be written as the sum of two $N/2$ point DFTs. In principle, one can divide the DFT in whichever way is desired, but it turns out that there are benefits to splitting the DFT into even and odd indexed components such that

$$\begin{aligned} H_n &= \sum_{k=0}^{N-1} h_k e^{-i2\pi kn/N} \\ &= \sum_{k=0}^{N/2-1} h_{2k} e^{-i2\pi(2k)n/N} + \sum_{k=0}^{N/2-1} h_{2k+1} e^{-i2\pi(2k+1)n/N} \\ &= \sum_{k=0}^{N/2-1} h_{2k} e^{-i2\pi kn/(N/2)} + e^{-i2\pi n/N} \sum_{k=0}^{N/2-1} h_{2k+1} e^{-i2\pi kn/(N/2)} \\ &= A_n + W_N^n B_n \end{aligned} \quad (2.41)$$

where A_n and B_n are the DFTs of the even and odd components of h_k , respectively. As such, this process has reduced the N point DFT to two $N/2$ point DFTs, or more importantly, N^2 operations to $N^2/2$ operations. The W_N^n coefficient is commonly referred to as the twiddle factor. It should also be noted that $W_N^{n+N/2} = -W_N^n$, and the twiddle factor need only be calculated explicitly $N/2$ times for a N point transform.

Fortunately this reduction can be applied recursively, dividing A_n and B_n into their “even” and “odd” indexed components, which can continue until the input array can be divided no further. This segregation of terms is made more clear in table 2.3. The utility of this process is maximized when the number of points in the DFT is $N = 2^m$, with m being a positive integer, since this results in an array that can be divided into even and odd components m times, without any leftover terms.

Table 2.3: The array input to the FFT (h_k), is divided into smaller arrays to be DFTed. This is accomplished by breaking the input array into even and odd components recursively as shown in this table for an array of $2^3 = 8$ points.

Values	Phase 1	Phase 2	Phase 3
h_0	h_0	h_0	h_0
h_1	h_2	h_4	h_4
h_2	h_4	h_2	h_2
h_3	h_6	h_6	h_6
h_4	h_1	h_1	h_1
h_5	h_3	h_5	h_5
h_6	h_5	h_3	h_3
h_7	h_7	h_7	h_7

The FFT algorithm is structured effectively as a series of DFTs computed using Equation 2.41. The composite DFTs are built up in a way which is precisely the reverse of the sorting routine shown in table 2.3, with each box indicating the array input to the DFT. For an input array of N points, the process begins with N one point DFTs (phase 3). Since $k = 0$ is the only index in a one point transform, the DFT of a single number is the number itself (Equation 2.35). Adjacent single point transforms are then added with the inclusion of the twiddle factor. This is followed by what is effectively $N/2$ two point DFTs, which

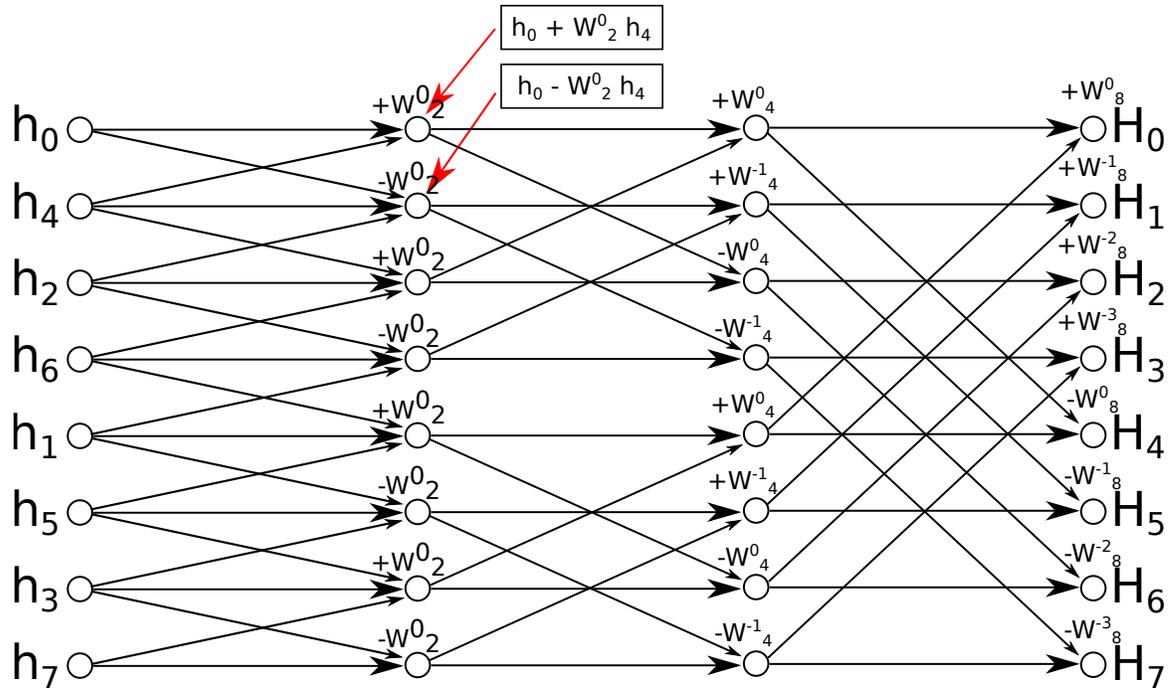


Figure 2.9: A flowchart showing the operations of the fast Fourier transform algorithm. Figure adapted from [26].

is accomplished by adding adjacent results from the previous step using the appropriate twiddle factor. This process continues until one N point DFT is calculated. This process is represented graphically in Figure 2.9 for a $2^3 = 8$ point FFT. Perhaps more clearly, the H_0 term can be calculated as

$$H_0 = \{ [(h_0 + W_2^0 h_4) + W_4^0 (h_2 + W_2^0 h_6)] + W_8^0 [(h_1 + W_2^0 h_5) + W_4^0 (h_3 + W_2^0 h_7)] \}, \quad (2.42)$$

where curly braces indicated the one effective N point DFT, square braces indicated the two $N/2$ DFTs, and circular braces indicate the four 2 point DFTs.

The efficiency of the FFT comes from factorizing the DFT into smaller transforms which can be more easily computed, and recycled when computing the various H_n terms. Each node in Figure 2.9 with a corresponding twiddle factor indicates a multiplication in the FFT routine. For this $N = 2^3 = 8$ point transforms, the “node matrix” has $N = 8$ rows and $\log_2(N) = 3$ columns. Thus computing the number of operations required for the FFT

in the same way as was done for the DFT, results in $N \log_2(N)$ operations. In theory, this makes the FFT $N / \log_2(N)$ times faster than the DFT, turning the aforementioned ~ 14 hr DFT into a ~ 0.11 sec FFT. In practice, however, FFTs take longer than this rough estimate as a result of neglected arithmetic calculations, and other back-end processes involved with code execution.

The FFT, as outlined, also exploits a property called binary reversal which expedites the sorting shown in table 2.3. Consider the fourth row (h_3), which for an 8 value array with indexing starting at zero, results in a binary memory address of 011. Binary reversal, which simply reverses the binary number, converts this to 110 which corresponds to 6 in base ten. This is precisely the element of the input array that is moved to this memory address in the proposed sorting routine (phase 3, table 2.3). In principle, the FFT algorithm can operate by breaking up the DFT in any arbitrary way. However, recursively grouping even and odd terms allows convenient use of the $W_N^{n+N/2} = -W_N^n$ property, and binary reversal, making coding the algorithm and the bookkeeping of twiddle factors much more simple.

2.3.2.1 Normalization

Though the DFT, and functionally equivalent FFT, is given by Equation 2.38, this operation does not produce the precise result obtained using the continuous Fourier transform (Equation 2.9). Inspection of Equations 2.37, and 2.38 reveals a missing factor of Δt in the latter equation. If the goal is to reproduce the continuous Fourier transform proper, then the output of the DFT algorithm must be multiplied by the time sampling interval

$$H(f_n) = H_n \Delta t . \quad (2.43)$$

When this normalisation is applied, Parseval's theorem for discretized data becomes

$$\sum_{k=0}^{N-1} |h_k|^2 \Delta t = \sum_{n=0}^{N-1} |H_n|^2 \Delta f . \quad (2.44)$$

Often the goal of performing a FFT on a function is not to produce the analytical approximation, but to extract the relative frequency content of the input data. In this case the FFT should be divided by the total number of points N . Energy will be equally distributed between the positive and negative frequency domains, so an additional factor of two may be desired if only the positive frequency content is analyzed. When this normalization is employed, Parseval's theorem becomes

$$\sum_{k=0}^{N-1} |h_k|^2 = \sum_{n=0}^{N-1} |H_n|^2 N \quad (2.45)$$

Care must be taken to ensure the correct normalization is employed for the desired task.

2.3.3 Sampling

In 1928 Harry Nyquist published a paper on telegraph transmission theory [40] which, in part, detailed the criteria for distortionless transmission of telegraph signals. In effect, in order to uniquely determine a sine wave using uniform sampling, two points must be sampled for each period. This motivates what is known as the Nyquist sampling theorem; to prevent any loss of information, the sampling frequency f_s must be at least twice as great as the maximum frequency of the signal. The Nyquist frequency, f_N , is defined as half the sampling frequency, and corresponds to the maximum frequency which can be uniquely determined based on sampling rate:

$$f_s \geq 2f_N \quad \text{or} \quad \Delta t \leq \frac{1}{2f_N} . \quad (2.46)$$

A similar criteria exists when sampling in frequency space; $\Delta f \leq 1/2T$, with T being the maximum time in the time domain. This latter definition also sets the theoretical limit on spectral resolution for FTS.

To ensure this criteria is satisfied, the signal being measured is often passed through a band-limiting filter, generally a low-pass filter. However, if frequencies greater than the

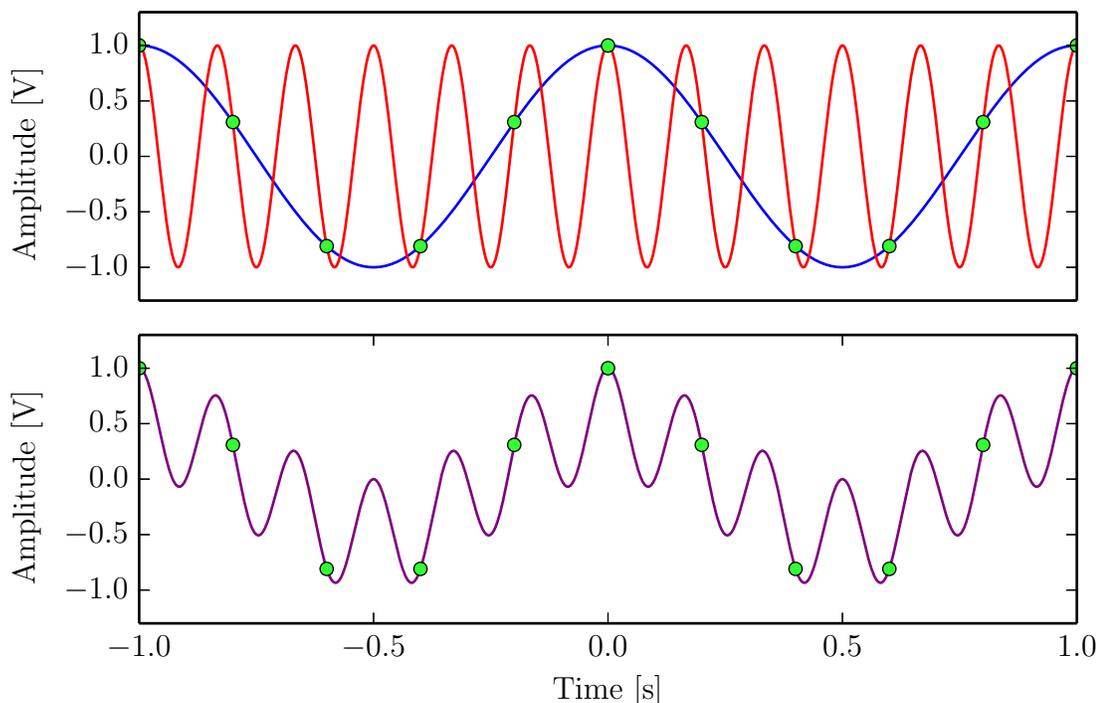


Figure 2.10: The top panel shows two sinusoids with frequencies 1 [Hz] (blue), and 6 [Hz] (red). These two sinusoids are sampled such that both waveforms coincide at each 0.2 [s] interval. The bottom panel shows the waveform produced by averaging the two sinusoids and the corresponding sampled points using the same sampling interval.

Nyquist frequency are present, the recorded signal will suffer from aliasing. An example of aliasing is shown in Figure 2.10. Sampled data is represented by the green markers, which perfectly coincide with two sinusoids of different frequency. The average of the two sinusoids is shown in the lower panel. If presented with such data with no bandwidth information, it would be impossible to determine the frequency content of the signal.

In the spectral domain, aliasing has the effect of folding the energy from frequencies above the sampling criteria across the positive and negative Nyquist frequency. This folding is repeated successively until the high frequency content resides within the Nyquist frequency limits. This phenomenon is demonstrated in Figure 2.11 where the lower panel of Figure 2.10 has been Fourier transformed. The Nyquist frequency resulting from the sampling rate of the green points is shown with a dashed line. The FFT of the over-sampled

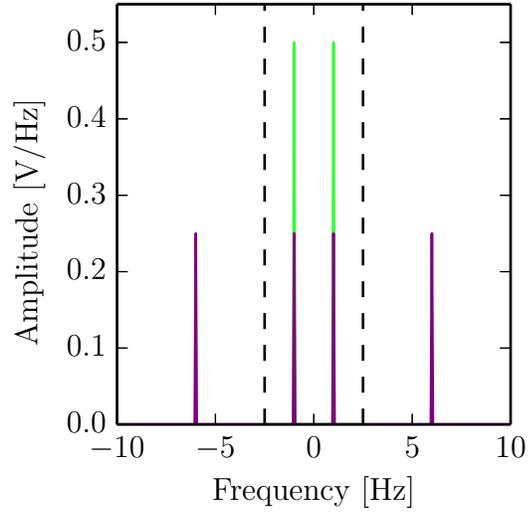


Figure 2.11: The spectra obtained from the FFT of the two datasets in the bottom panel of Figure 2.10. The over-sampled purple dataset spectrum shows the unambiguous frequency content of the original signal, whereas the under-sampled green dataset has the higher frequency components folded across its Nyquist frequency (dashed line), into the detectable frequency range.

purple signal shows four peaks at frequencies corresponding to the positive and negative frequencies of the two sinusoids used to synthesize the data. The FFT of the under-sampled green points produce only two peak which can be generated by folding the purple spectrum across the Nyquist frequency.

The relationship between two frequencies, f and f_0 , which produce this aliasing effect is given by

$$f_0 = f \pm \frac{n}{\Delta t} = f \pm 2nf_N \quad (2.47)$$

where $f > f_N > f_0$, and n is some integer. This equation is easily proven by considering the exponent in Equation 2.37.

$$\exp[i2\pi f_0 t_k] = \exp\left[i2\pi t_k \left(f \pm \frac{n}{\Delta t}\right)\right] = \exp\left[i2\pi t_k f \pm i2\pi t_k \frac{n}{\Delta t}\right] = \exp[i2\pi f t_k]$$

Where in the last equality we have used $t_k = k\Delta t$, and the fact that $e^{i2\pi\theta}$ is periodic over integer multiples of 2π . The example in Figure 2.10 uses $f = 6$ [Hz], and $f_N = 2.5$ [Hz].

Using Equation 2.47, $f_0 = 6 - (2)2.5 = 1$ [Hz], which is indeed the frequency of the blue sinusoid.

2.4 Summary

This chapter presents an introduction to Fourier analysis. An understanding of Fourier transforms is critical for understanding Fourier transform spectroscopy, which is the central spectroscopic technique employed in this thesis.

Physically realizable functions can be synthesized as a superposition of simple sinusoidal waves. The frequency of these composite waves, and their associated amplitude coefficients can be recovered from the synthesized function using the Fourier transform. In this chapter, I derived a number of useful transforms of analytic functions commonly encountered in Fourier transform spectroscopy, along with a summary of symmetry properties, and useful theorems.

In practice, Fourier transforms are applied to discretized datasets, so the discrete Fourier transform has been derived, along with its more efficient analogue, the fast Fourier transform. This was followed by a discussion on how sampling in the time domain can affect the resulting spectrum. Care must be taken to ensure that the chosen sampling rate is sufficient to resolve the frequencies of interest, and proper normalization must be employed in the analysis of resultant spectra.

Chapter 3

Optical Fourier Transform Spectroscopy

Spectroscopy is the study of the interaction between light and matter and the various observational techniques used to study the composition of light. Spectroscopy is an indispensable observational technique in astronomy and provides insight into a broad range of physical parameters of the sources under observation. Spectra from astronomical sources are generally composed of three components. First, broad continuum emission resulting from the black-body thermal emission of dust grains constitutes the underlying structure of most FIR sources. In practice, the source of thermal emission does not express unit emissivity producing a modified black-body spectrum often called grey-body emission. This grey-body radiation curve can be used to determine dust temperature, dust mass, and opacity of the sources under observation [44]. Other effects which can add to continuum emission include synchrotron radiation from rotationally accelerated electrons, along with free-free and free-bound interactions between electrons and other electrons, atoms, and ions [45]. Secondly, line emission is produced by discrete transitions of electrons from excited energy states to lower energy states. In FIR astronomy, spin state transitions of valence electrons can be of low enough energy to be observed within the wavelengths under study. Analysis of the line emission features within spectra gives information directly related to the chemical composition of the source since discrete energy transitions, and thus emission frequencies, are unique to specific atoms, ions, or molecules. Certain chemical species are in low abundance in the ISM and their spectral emission interacts weakly with the ISM as it propagates through molecular clouds. As such, analyzing different spectral features from different

molecular or atomic species can be used to probe the ISM at different depths. Spectral emissions are dependent on the available energy of their local environment, thus a variety of emission features can be used to study different energy regimes of the ISM. Analyzing the width of spectral features can provide insight into the temperature and pressure of gases under observation, while a shift in the emission frequency from the rest frame value directly probes source radial velocity, and thus line-of-sight dynamics. Thirdly, absorption features result from molecular, atomic, or ionic gas absorbing photons of a specific frequency resulting in a discrete energy transitioning of the constituent particles into an excited state. Absorption spectra give information about foreground material between a source of radiation and observer. Indeed, spectral analysis of detected radiation provides a plethora of information.

This chapter provides an introduction to the interaction of light with matter and Fourier transform spectroscopy in order to better understand the information that can be extracted from FTS spectra. In Section 3.1, the theory of black-body radiation and the propagation of light through media is presented. Spectral line shapes and their associated broadening mechanisms are introduced in Section 3.2. Electronic emission of atoms, and rotational emission for linear molecules, the emission mechanisms most relevant to this thesis are presented in Section 3.3. Finally, the spectroscopic technique known as Fourier transform spectroscopy is then outlined in Section 3.4, along with a discussion of the data reduction techniques known as zero-padding and phase correction. The goal of these sections is to cover a broad set of topics to provide sufficient detail to understand the content of the SPIRE spectra within this thesis.

3.1 Spectral Composition

3.1.1 Black-body Radiation

An object which perfectly absorbs all incident radiation would appear black to an observer, and this is precisely the defining quality of what is called a “black-body”. In practice,

a perfect black-body is difficult to obtain, but they can be approximated as a small aperture entrance to a large internal cavity. Light entering through the apertures is scattered several times by the walls of the cavity without much chance of escaping, in accordance with the black-body definition. As the radiant energy within the cavity is absorbed and re-emitted by the inner walls, the system will tend towards thermal equilibrium. Once in thermal equilibrium, radiation escaping from the small aperture will express a characteristic spectral profile known as black-body radiation determined only by the equilibrium temperature.

The study of black-body radiation played a seminal role in motivating the theory of quantum mechanics. One can imagine the internal energy of the cavity existing as a superposition of standing waves within the cavity. Each resonant mode is free to vibrate in one dimension with two polarizations and as a result of the equipartition theorem, and the well known thermodynamic expression for energy, each mode contains an average energy $E = 2 \cdot kT/2$, where k is Boltzmann's constant, and T is the absolute temperature of the cavity in Kelvin. In the classical view, there exists an infinite number of resonant modes within the cavity, and as such, one is lead to the conclusion that an infinite amount of energy exists within the cavity. This problem was first noted by Rayleigh, and later termed the ultraviolet catastrophe by Paul Ehrenfest [46]. The solution came from Max Planck's idea to restrict the energy in each mode to quantized values [47]. Planck postulated that the energy available to each mode be restricted to integer multiples of hf , with h being Planck's constant, and f the frequency of vibration of a given mode.

This revolutionary thinking culminated into what is known as Planck's law, which describes the spectral density inside a black-body cavity. The spectral density is defined as the energy per unit frequency per unit volume. In what follows, Planck's law will be obtained by summing the average energy per mode, \overline{E}_m , over all modes and then dividing the result by the cavity volume, V ,

$$E = \sum_{m=0}^{\infty} B_m = \sum_{m=0}^{\infty} \frac{\overline{E}_m}{V} . \quad (3.1)$$

B_m expresses Planck's law, and in this case is given in units of $[\text{J} \cdot \text{Hz}^{-1} \cdot \text{m}^{-3}]$. Planck's law

can be expressed in the more familiar $[\text{W} \cdot \text{m}^{-2} \cdot \text{Hz}^{-1} \cdot \text{sr}^{-1}]$ units by means of a proportionality factor which will be considered shortly. Obtaining the average energy per mode requires the introduction of a few concepts from statistical physics.

The Boltzmann factor, \mathcal{B}_s ,

$$\mathcal{B}_s = e^{-E_s/kT}, \quad (3.2)$$

when weighted by the partition function, Z ,

$$\begin{aligned} Z &= \sum_s \mathcal{B}_s = 1 + e^{-hf/kT} + e^{-2hf/kT} + \dots \\ &= \frac{1}{1 - e^{-hf/kT}}, \end{aligned} \quad (3.3)$$

gives the probability $P(s)$, of a state with energy E_s being occupied when the system is in thermal equilibrium at temperature T [48]. That is,

$$P_s = \frac{\mathcal{B}_s}{Z}. \quad (3.4)$$

The last line in Equation 3.3 has been obtained by using the binomial expansion $\sum_n (-1 \cdot x)^n = (1+x)^{-1}$. By making the substitution $\beta = 1/kT$, it is easy to show that the average energy for a given mode of vibration \overline{E}_m is given by

$$\overline{E}_m = \sum_s P_s E_s = -\frac{1}{Z} \frac{\partial Z}{\partial \beta}. \quad (3.5)$$

Substituting equation 3.3 into the above equation, the average energy for a given mode becomes

$$\overline{E}_m = \frac{hf}{e^{hf/kT} - 1}. \quad (3.6)$$

This is the desired quantity for Equation 3.1 with f representing a single vibrational mode in the closed cavity system. When Equation 3.6 is divided by hf , the quanta of energy for a photon, it becomes the Planck distribution. The Planck distribution represents the

number of photons of frequency f present on average in a system in thermal equilibrium at temperature T .

To proceed, recall that a node is required at either boundary of a closed system in order to obtain resonance. For a one dimensional system of length L , the boundary conditions require that wavelengths λ , such that $\lambda = 2L/m$, with m representing the vibrational modes and taking positive integer values. The ground state $m = 0$ is forbidden as written, but is acceptable in frequency space $f = c/\lambda = mc/2L$. Substituting this expression into Equation 3.6, and then returning to Equation 3.1 gives

$$E = \sum_{m=0}^{\infty} \frac{hmc}{2LV} \frac{1}{e^{hmc/2LkT} - 1}. \quad (3.7)$$

This equation however, is still a one dimensional representation which must be expanded to a three dimensional mode space $\{m_x, m_y, m_z\}$. In addition, a factor of 2 must be included to account for the two potential polarization modes. The calculation is made easier by exchanging the summation over Cartesian mode space for an integration over a spherical coordinate mode space

$$E = \int_0^{\infty} dm \int_0^{\pi/2} d\theta \int_0^{\pi/2} m^2 \sin(\theta) d\phi \frac{hmc}{LV} \frac{1}{e^{hmc/2LkT} - 1}. \quad (3.8)$$

Note that the integration is carried out over one eighth of the unit sphere corresponding to the positive m_i values only. This give an angular contribution of $4\pi/8 = \pi/2$. Since $hf = hcm/2L$, it follows that $dm = 2L/c \cdot df$. Making these substitutions and noting that $V = L^3$, Equation 3.8 becomes

$$E = \int_0^{\infty} \frac{8\pi h f^3}{c^3} \frac{1}{e^{hf/kT} - 1} df. \quad (3.9)$$

The integrand corresponds to Planck's law and describes the spectral energy density of a

black-body at a given absolute temperature, T ,

$$P_f(T) = \frac{8\pi hf^3}{c^3} \frac{1}{e^{hf/kT} - 1} \left[\frac{\text{J}}{\text{Hz} \cdot \text{m}^3} \right]. \quad (3.10)$$

It is worth noting that for a black-body spectrum, the amplitude at a given frequency is entirely determined by temperature alone.

Equation 3.10 gives the energy density in terms of Joules per Hertz per meter cubed, but in astronomy, an expression with appropriate units is more useful. Optical instruments typically have a limited field-of-view defined by a solid angle. Equation 3.10 is normalized for field-of-view by dividing it by the solid angle of a unit sphere 4π sr. Additionally, optical detectors integrate the incident radiation over a specific time period, and these detectors have a finite size defined by their surface area. As such, Equation 3.10 is divided by time, t , converting energy to power, then multiplied by length, ct , cancelling a unit of length from the denominator and normalizing for area. The integrand of Equation 3.10 then takes the more familiar form

$$B_f(T) = \frac{2hf^3}{c^2} \frac{1}{e^{hf/kT} - 1} \left[\frac{\text{W}}{\text{Hz} \cdot \text{sr} \cdot \text{m}^2} \right]. \quad (3.11)$$

Planck's law can be expressed in a few other convenient forms depending on the application,

$$B_\lambda(T) = \frac{2hc^2}{\lambda^5} \frac{1}{e^{hc/\lambda kT} - 1} \left[\frac{\text{W}}{\text{m} \cdot \text{sr} \cdot \text{m}^2} \right], \quad (3.12)$$

and

$$B_\sigma(T) = \frac{2 \times 10^8 hc^2 \sigma^3}{e^{100hc\sigma/kT} - 1} \left[\frac{\text{W}}{\text{cm}^{-1} \text{sr} \cdot \text{m}^2} \right], \quad (3.13)$$

where λ represents wavelength in meters [m], and σ represents wavenumber in inverse centimeters [cm^{-1}].

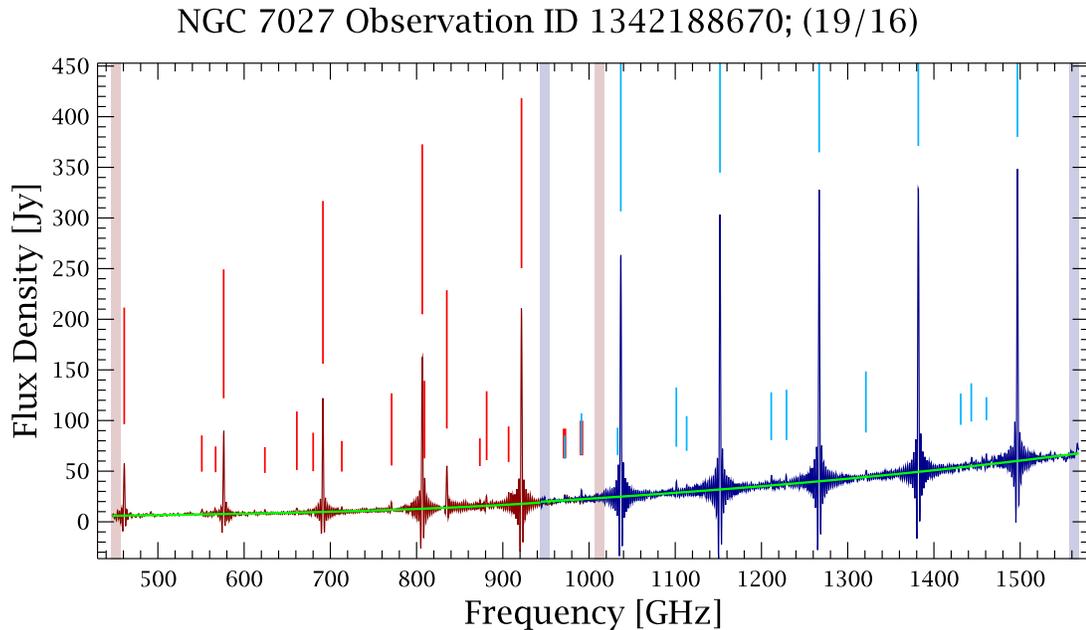


Figure 3.1: SPIRE spectrum of the young planetary nebula NGC 7027 showing the grey-body continuum with a green curve, and emission features with vertical bars. The length of a bar is proportional to the corresponding feature’s signal to noise ratio. Red and blue spectral components are derived from the central detector of the SPIRE SLW and SSW IFTS detector arrays, respectively (Section 4.2).

3.1.2 Radiative Transfer

In astronomy, black-body radiation results from the thermal radiation of objects larger than molecules (eg. dust, planets, stars), providing the broad underlying structure of most spectra. This component is often referred to as the continuum which is then modified by other radiative processes as it propagates from source to observer. These modifications result from the emission or absorption of relatively narrow bands of radiation from atoms and molecules which produce peaks and valleys in the continuum, respectively (Figure 3.1). Understanding how radiation interacts with matter, broadly referred to as radiative transfer, is thus critical to understanding emission and spectra.

3.1.2.1 Einstein Coefficients

The Einstein coefficients are material specific constants which describe the transition rates between two non-degenerate discrete energy levels. Consider a system with only two possible energy levels such that E_0 corresponds to the lower energy state, and E_1 corresponds to the upper energy state. It is assumed the system is in thermal equilibrium at temperature T , and is exposed to a radiation density $P_f(T)$ (Equation 3.10). If the system can be described by a particle density, the ratio between particle density in the upper state, N_1 , and lower state density, N_0 , is given by the Boltzmann distribution

$$\frac{N_1}{N_0} = e^{-E_{10}/kT}, \quad (3.14)$$

where $E_{10} = E_1 - E_0$ is the energy difference between the upper and lower states. In thermal equilibrium, this ratio will remain constant on average, however, particles will transit between the two states via radiative processes.

There are three radiative processes: absorption, stimulated emission, and spontaneous emission. Absorption is the result of a particle in the lower state absorbing a photon of energy $E_{10} = hf_{10}$, which transitions it to the upper state. The rate of absorption is characterized by the Einstein absorption coefficient $B_{1\leftarrow 0}$, defined implicitly as

$$\frac{dN_1}{dt} = B_{1\leftarrow 0} P_f(f_{01}, T) N_0. \quad (3.15)$$

As is expected, the absorption rate is proportional to the energy density at frequency f_{10} , and the number density of particles in the lower state N_0 . Stimulated emission occurs when a photon of energy E_{10} interacts with a particle in the upper state resulting in its transition to the lower state. The Einstein stimulated emission coefficient $B_{1\rightarrow 0}$ is implicitly defined as

$$\frac{dN_1}{dt} = -B_{1\rightarrow 0} P_f(f_{10}, T) N_1. \quad (3.16)$$

Finally, spontaneous emission is the unprovoked transition of a particle from the upper state to the lower state. The spontaneous emission coefficient is implicitly defined as

$$\frac{dN_1}{dt} = -A_{1 \rightarrow 0} N_1 . \quad (3.17)$$

In a state of equilibrium, the sum of the emission and absorption rates must be equal such that

$$N_0 B_{1 \leftarrow 0} P_f(f_{10}, T) = N_1 (B_{1 \rightarrow 0} P_f(f_{10}, T) + A_{1 \rightarrow 0}) . \quad (3.18)$$

This expression can be solved for the ratio of upper and lower states, which, when compared to Equation 3.14 gives

$$\frac{B_{1 \leftarrow 0} P_f(f_{10}, T)}{A_{1 \rightarrow 0} + B_{1 \rightarrow 0} P_f(f_{10}, T)} = e^{-E_{10}/kT} . \quad (3.19)$$

Solving for the spectral energy density $P_f(f_{10}, T)$ gives

$$P_f(f_{10}, T) = \frac{A_{1 \rightarrow 0}}{B_{1 \leftarrow 0} e^{E_{10}/kT} - B_{1 \rightarrow 0}} . \quad (3.20)$$

When comparing the above equation to Equation 3.10, it is evident that the solution requires

$$B_{1 \rightarrow 0} = B_{1 \leftarrow 0} , \quad (3.21)$$

and

$$A_{1 \rightarrow 0} = B_{1 \leftarrow 0} \frac{8\pi h f_{10}^3}{c^3} . \quad (3.22)$$

These important results show that once any Einstein coefficient has been determined, effectively all have been determined.

So far, the Einstein coefficients give statistical insight into radiative processes without any reference to the interacting particles themselves. The next subsection develops the link between matter and radiation directly.

3.1.2.2 Interaction of Light With Matter

The electromagnetic properties of matter determine their interaction with light. The electric properties are most easily parameterized using the dipole moment of the particles in question. The dipole moment in Cartesian coordinates is given by

$$\vec{\mu} = \sum \vec{r}_i q_i , \quad (3.23)$$

where the sum takes place over all charged particles with charge q_i and displacement \vec{r}_i from the origin which usually coincides with the center of mass of the system of particles.

The problem is then approached by using the semi-classical model consisting of a system with two discrete energy levels combined with classical electromagnetic waves. A key step is solving the time dependent Schrödinger equation

$$i\hbar \frac{\partial \Psi(t)}{\partial t} = (H + H'(t))\Psi = E\Psi . \quad (3.24)$$

The wave function $\Psi(t)$ is a solution to this equation and consists of a superposition of time independent states ψ_0 and ψ_1 corresponding to the two static energy levels of the system, and takes the form

$$\Psi(t) = a_0 \psi_0 e^{-i2\pi E_0 t/h} + a_1 \psi_1 e^{-i2\pi E_1 t/h} . \quad (3.25)$$

The square of the coefficients a_0 and a_1 have the physical significance of being the probability of their corresponding states being occupied. The time dependent terms result from $H'(t)$, a time dependent perturbation of the Hamiltonian that results from the interaction of the time dependent electric field of light and the dipole moment of the matter in question

$$\begin{aligned} H'(t) &= -\vec{\mu} \cdot \vec{E}(t) \\ &= -\vec{\mu} \cdot E_0 \cos(2\pi(\vec{\kappa} \cdot \vec{r} - ft)) \\ &\Rightarrow -\mu E \cos(2\pi(\kappa z - ft)) , \end{aligned} \quad (3.26)$$

where \vec{k} is the wave vector indicating the direction of propagation of the wave with $|\vec{k}| = 1/\lambda$. In the last line, the situation has been simplified to a one dimensional electric disturbance propagating in the positive z direction.

In the absence of all relaxation processes, including collisions and emission, the solution of Equation 3.24 furnishes the following results [49]:

$$\begin{aligned} |a_0(t)|^2 &= 1 - \left(\frac{\omega_R}{\Omega} \sin\left(\frac{\Omega t}{2}\right) \right)^2 \\ |a_1(t)|^2 &= \left(\frac{\omega_R}{\Omega} \sin\left(\frac{\Omega t}{2}\right) \right)^2. \end{aligned} \quad (3.27)$$

Two terms have been introduced in the above equations. The Rabi frequency [49],

$$\omega_R = \frac{M_{01}E}{\hbar}, \quad (3.28)$$

is defined in terms of E , the projection of the electric field amplitude E_0 , onto the dipole axis $\vec{\mu}$, the constant $\hbar = h/(2\pi)$, and the transition dipole,

$$M_{01} = M_{10} = \int \psi_1^* \mu \psi_2 d\tau. \quad (3.29)$$

Meanwhile the convenience term Ω is defined as

$$\Omega = \sqrt{\omega_R^2 + \Delta^2}, \quad (3.30)$$

where

$$\Delta = 2\pi(f - f_{10}). \quad (3.31)$$

The quantity Δ is often referred to as the detuning frequency as it is a measure of the difference between the incident radiation of frequency f , and the resonant frequency f_{10} .

Equation 3.27 shows that when the system is exposed to radiation in the absence of decay processes, the probability of the upper state being occupied is an oscillating sine-

squared function which is modulated by the $(\omega_R/\Omega)^2$ coefficient. If the incident radiation is in resonance with the system (ie, $\Delta = 0$), the coefficient will be unity, and the probability function will have a period of $2\pi/\omega_R$ which gives the Rabi frequency its physical significance.

Decay, or relaxation processes, include spontaneous and stimulated emission of a photon which resets the two state system back to the ground state. Collision processes can induce photon emission, or alternatively, result in a discontinuous change in phase without resetting the system to its ground state. These two types of processes are referred to as T_1 and T_2 processes, respectively [49]. If the resonance assumption is maintained, and decay processes are introduced to a large collection of these two state systems, the effect is a dampening of Equation 3.27 such that at equilibrium $|a_1|^2 = |a_0|^2 = 0.5$, and on average half of each state is occupied. The aggregate of T_1 and T_2 processes can be characterized as occurring at an average rate or frequency ω_{decay} . If $\omega_R \gg \omega_{decay}$, then dampening effects become negligible and Equation 3.27 is sufficient to describe the system. Since ω_R is proportional to E , dampening effects can be overcome simply by increasing the intensity of the incident radiation.

For the case of low radiation intensity, there is little build up of excited states. Assuming that $a_0 \sim 1$ and $a_1 \sim 0$, Equation 3.27 can be recalculated resulting in [49]

$$|a_1|^2 = \frac{\mu^2 \pi t}{\epsilon_0 \hbar^2} P_\omega(\omega_{10}), \quad (3.32)$$

where ϵ_0 is the permittivity of free space, μ is again the dipole moment of the particle in question, and $P_\omega(\omega_{10})$ is the spectral energy density with $\omega_{10} = 2\pi f_{10}$ (Equation. 3.10).

The probability of the upper state being occupied is equivalent to the average ratio of the total number of particles in the upper state to the total number of particles, N_1/N . In the low intensity approximation, the occupancy of the lower state is approximately equal to the

total number of particles, $N_0 \sim N$. By inspecting Equation 3.15, it is clear that

$$\frac{d|a_1|^2}{dt} = \frac{d(N_1/N)}{dt} = B_{1\leftarrow 0} P_f(f_{01}) . \quad (3.33)$$

Equation 3.32 was derived using an electromagnetic wave propagating in the z direction, while Equation 3.15 was derived using isotropic radiation propagating in three dimensions. As such, a valid comparison between the two equations requires the insertion of a factor of $1/3$ into Equation 3.32. Accounting for this factor, and noting that

$$\begin{aligned} \int P_f df &= \int P_\omega d\omega \\ \int P_f df &= \int P_\omega 2\pi df \\ \therefore P_f &= 2\pi P_\omega , \end{aligned} \quad (3.34)$$

Equation 3.33 becomes

$$B_{1\leftarrow 0} = \frac{2\pi^2 \mu^2}{3\epsilon_0 h^2} , \quad (3.35)$$

and, according to Equation 3.22

$$A_{1\rightarrow 0} = \frac{16\pi^3 f^3 \mu^2}{3\epsilon_0 h c^3} . \quad (3.36)$$

These modified Einstein coefficients for weak electric field interaction have been derived in the absence of any decay processes, and crucially, pertain only to non-degenerate energy states. These results are significant in that they relate photon emission rates to the dipole moment μ , a physical quantity which parameterizes the distribution of charge of an atom or molecule. They are thus fundamental in determining the intensity of spectral features.

The effect of the various decay processes results in the multiplication of Equations 3.35 and 3.36 with a spectral line shape centered on the transition frequency of the two level

system. For example, the real Einstein absorption coefficient is given by

$$B_{1\leftarrow 0} = \int_{-\infty}^{\infty} \frac{2\pi^2\mu^2}{3\epsilon_0 h^2} g(f - f_{10}) df . \quad (3.37)$$

Spectral line shapes will be discussed in Section 3.2, but for the moment the quantity $g(f - f_{10})$ can be regarded as a distribution with a particular shape, and an integrated area equal to unity.

3.1.2.3 Transmission of Radiation

The previous sections in this chapter outline a theoretical foundation for the intensity of radiating systems. In practice, the situation is further complicated by the interaction of light with matter as it propagates from source to observer.

If light is incident on a system in thermal equilibrium containing N_0 particles per cubic meter in the lower state, and N_1 particles per cubic meter in the excited states, then the change in N_1 is given by

$$\frac{dN_1}{dt} = -B_{1\rightarrow 0}P_f N_1 + B_{1\leftarrow 0}P_f N_0 , \quad (3.38)$$

where P_f is the spectral density (Equation 3.10). Equation 3.38 is effectively the radiative transfer equation and simply states that the change in excited states over time is equal to the difference between the absorption rate and the stimulated emission rate. Spontaneous emission in the system is neglected as thermal equilibrium assumes that absorption and emission processes within the isolated system are equal.

Using Equations 3.35 and 3.21, and making the substitution $P_f = I/c = hfF/c$, with I and F being the intensity and flux of the incident radiation, respectively, Equation 3.38 becomes

$$\frac{dN_1}{dt} = F(N_0 - N_1) \frac{2\pi^2\mu^2 f g(f - f_{01})}{3\epsilon_0 hc} . \quad (3.39)$$

Furthermore, the effective area that interacting particles present to the incident radiation is

given by

$$\sigma = \frac{2\pi^2\mu^2fg(f-f_{01})}{3\epsilon_0hc} \quad [\text{m}^2]. \quad (3.40)$$

This quantity is called the absorption cross-section of the interacting particles. Notice that it is frequency dependent, and defined in terms of the particle's dipole moment μ . Inspecting Equation 3.36, it is clear that the absorption cross-section can be written as

$$\sigma = \frac{A_{1\rightarrow 0}c^2g(f-f_{01})}{8\pi f^2}. \quad (3.41)$$

Real systems in general contain more than one kind of interacting particle, and each interacting particle generally contains more than two energy levels. To account for this fact, $A_{1\rightarrow 0}$ can be replaced by a sum over all spontaneous emission coefficients present in the system. Furthermore, the radiative lifetime τ for the system can be defined as [49]

$$\tau = \sum \frac{1}{A_{i\rightarrow j}}, \quad (3.42)$$

where $A_{i\rightarrow j}$ represents the spontaneous emission Einstein coefficient between arbitrary energy levels i and j , with the sum taking place over all potential transitions.

Equation 3.39 is simply the change in flux of the system. If this quantity is positive, photons have been removed from the system. Considering the propagation of the incident radiation over a distance dx , the change in flux over this interval is then

$$dF = -\sigma F(N_0 - N_1)dx. \quad (3.43)$$

Using separation of variables, and integrating over a length L , one obtains

$$\ln\left(\frac{F}{F_0}\right) = \ln\left(\frac{I}{I_0}\right) = -\sigma(N_0 - N_1)L, \quad (3.44)$$

where F_0 and I_0 are the incident flux and intensity, respectively. Rearranging this equation

gives

$$I = I_0 e^{-\sigma(N_0 - N_1)L} = I_0 e^{-\alpha L} . \quad (3.45)$$

This important result, known the as Beer-Lambert law [50], shows that the transmitted intensity is attenuated exponentially as radiation propagates through a displacement L . The decay constant α , often referred to as the absorption coefficient, is proportional to the absorption cross-section and the population difference between the lower and upper energy states.

The quantity αL is called the optical depth, τ_f . This value is frequency dependent as a result of its dependence on the absorption cross section. “Optically thick” media have an optical depth much greater than 1, resulting in significant attenuation of propagating light, and are thus essentially opaque. On the other hand, “optically thin” media have an optical depth much less than 1, and are for the most part transparent.

3.2 Line Shapes

Heisenberg’s uncertainty principle [51] states that both the position / time, and momentum / energy of a particle cannot be determined simultaneously with infinite precision. The uncertainty in both measurements are limited explicitly by

$$\Delta x \cdot \Delta p \geq \frac{\hbar}{2} , \quad (3.46)$$

or

$$\Delta t \cdot \Delta E \geq \hbar . \quad (3.47)$$

The frequency of a photon is related to its momentum as $p = hf/c$, thus the uncertainty principle imposes intrinsic limits on the minimum bandwidth of a spectral feature in the limit of well constrained position or time. The profile of spectral features is called a line shape. Though the minimum width of a spectral feature results from the uncertainty principle, other physical processes can impose additional broadening, with different frequency

dependence. Some of the most common broadening mechanisms are discussed below.

3.2.1 Natural Broadening

Natural broadening results directly from the Heisenberg time/energy uncertainty principle. The radiative lifetime $\tau = 1/A_{1 \rightarrow 0}$ (Equation 3.42), sets the time scale for spontaneous emission of a two level system, and this average decay period induces an uncertainty in the energy of the emitted photon. Using Equation 3.46, with $\Delta t = \tau$, and recognizing that $\Delta E = h\Delta f$, it can already be concluded that the bandwidth of natural broadening is

$$\Delta f \geq \frac{1}{2\pi\tau}. \quad (3.48)$$

A more detailed explanation is obtained by calculating the expectation value of the dipole moment (Equation 3.23) of a system containing a large number of identical particles in the absence of radiation

$$\begin{aligned} M(t) &= \langle \Psi(t) | \mu | \Psi(t) \rangle \\ &\rightarrow 2a_0 a_1 \mu \cos(\omega_{10} t), \end{aligned} \quad (3.49)$$

where $\Psi(t)$ is the solution to the time dependent Schrödinger equation (Equation 3.25). The last line constitutes the real valued component of the expanded calculation.

In the absence of any driving radiation, excited states are not replenished and the change in the probability of excited state occupancy is governed by the differential equation

$$\frac{d|a_1|^2}{dt} = -|a_1|^2 \gamma. \quad (3.50)$$

The solution is easily obtained as $a_1(t) = a_1(0)e^{-\gamma t/2}$. The decay constant γ is simply $1/\tau$, which when substituted into Equation 3.49 gives

$$M(t) = e^{-\gamma t/2} 2a_0 a_1 \mu \cos(\omega_{01} t). \quad (3.51)$$

This result is a time dependent function consisting of a sinusoid modulated by an exponential. The frequency content is obtained by taking the Fourier transform of this function. Invoking the convolution theorem, and recalling the Fourier transform of a sinusoid and an exponential (Section 2.1.4), the line shape is then given by a delta function at $f = \omega_{01}/2\pi$ convolved with a Lorentzian.

Intensity is obtained by squaring this result, and after converting to frequency, the line shape due to natural broadening is

$$g_{Natural}(f - f_{10}) = \frac{\gamma}{(\gamma/2)^2 + (2\pi)^2(f - f_{10})^2} \cdot \quad (3.52)$$

Bandwidth is often parameterized as the Full Width at Half Maximum (FWHM). This is the frequency width Δf subtended by the line shape at half its maximum amplitude. Solving Equation 3.52 for Δf when $g_{Natural}(\Delta f) = 2/\gamma$ (half its maximum value), gives $\Delta f = \pm\gamma/(4\pi)$. As such $\Delta f = \gamma/(2\pi) = 1/(2\pi\tau)$, in agreement with Equation 3.48.

3.2.2 Pressure Broadening

The first attempt to describe emission features due to particle collisions was done by Lorentz [52]. His theoretical derivation predicts line shapes and widths in good agreement with empirical measurements of gases at low to moderate pressures. Different approaches have been used to derive analytical expressions for the line shape that results from pressure broadening [53], some producing more accurate predictions that include an empirically observed shift of the peak emission frequency. However, to first order, the Lorentzian derivation is sufficient for most physical situations.

The Lorentzian approach uses a semi-classical model where the radiation emitted from a particle is proportional to the dipole moment. In the absence of driving radiation, the dipole moment has an expectation value which oscillates at the Bohr angular frequency ω_{10} (Equation 3.49). Lorentz postulated that collisions between particles result in a random phase shift in this periodic oscillation. This assumption produces an average dipole mo-

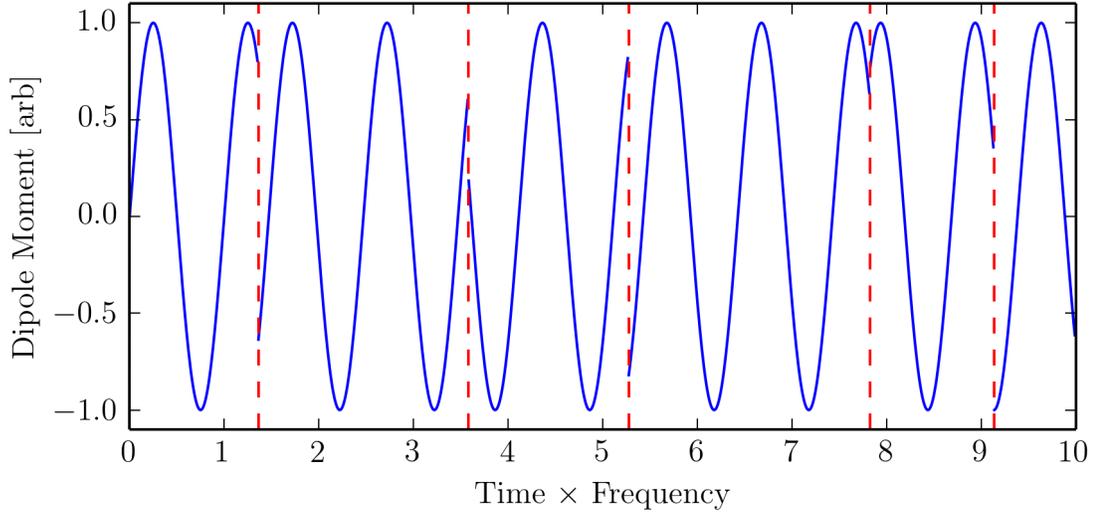


Figure 3.2: Oscillating dipole moment with random phase shifts induced by collisions indicated by the dashed red lines.

ment with the appearance of multiple discontinuous sine waves of the same frequency ω_{10} (Figure 3.2). The Fourier transform of this function gives the spectral line shape

$$g_{Pressure}(f - f_{01}) = \frac{n\pi d^2 v_a / \pi}{(n\pi d^2 v_a)^2 + 4\pi^2 (f - f_{01})^2} . \quad (3.53)$$

This Lorentzian line shape has an FWHM of $2n\pi d^2 v_a$, with n , d , and v_a being the number density [m^{-3}], collision radius [m], and mean velocity [$\text{m} \cdot \text{s}^{-1}$] of the particles in the system, respectively. Through the ideal gas law, pressure is proportional to the number density of particles and the square of their average velocity. As a result the FWHM can also be represented as

$$\Delta f = Qp , \quad (3.54)$$

where p is the pressure of the gas, and Q is the pressure broadening coefficient. Values for the pressure broadening coefficient are typically a few MHz per Torr [54].

Both pressure and natural broadening mechanisms affect all particles in the system equally, and are thus referred to as homogeneous. Properties of individual particles, such as velocity, are not in general uniform and produce inhomogeneous line shapes. This inho-

ogeneous broadening mechanism is discussed in Section 3.2.3.

3.2.3 Doppler Broadening

Doppler broadening results from the interplay of the distribution of particle velocities along the line of sight and the resultant frequency shift of their emitted radiation as seen by an observer.

If a stationary particle emits radiation at a frequency, f_0 , an observer moving radially towards the emitter at velocity, v , will detect this radiation as having a frequency, f . These quantities are related by the non-relativistic Doppler equation

$$f = \left(\frac{v}{c} + 1 \right) f_0 . \quad (3.55)$$

It is worth noting that this equation is defined such that the positive velocity corresponds to the source and observer moving towards each other and at positive velocities the emission line is shifted to higher frequencies. A different convention is used in astronomy which will be discussed in Chapter 5.

The one-dimensional distribution of particle velocities in the system is described by the Maxwell-Boltzmann velocity distribution function [55]

$$P(v) = \sqrt{\frac{m}{2\pi kT}} e^{-mv^2/2kT} dv , \quad (3.56)$$

spanning all velocities between 0 and ∞ . This distribution assumes thermodynamic equilibrium at absolute temperature, T , for a system of identical particles of mass m . Substituting Equation 3.55 into the above equation and noting that $dv = (c/f_0)df$, one obtains the spectral distribution of the system

$$g_{Doppler}(f - f_0) = \frac{1}{f_0} \sqrt{\frac{mc^2}{2\pi kT}} e^{-mc^2(f-f_0)^2/2kTf_0^2} . \quad (3.57)$$

This spectral distribution is simply the line shape associated with Doppler broadening. Finally, solving for the FWHM gives the line width as

$$\Delta f = 2f_0 \sqrt{\frac{2kT \ln(2)}{mc^2}} . \quad (3.58)$$

At a given temperature, the line width of an emission feature will vary depending on the species of particle and the specific transition, resulting from the mass and transition frequency emission dependence, respectively.

3.2.4 Signal Truncation

The general case of signal truncation results in a broadened spectral profile for a monochromatic signal. One can imagine a low intensity polychromatic beam with a diameter d interacting with a stream of particles of resonant frequency f_0 , travelling orthogonal to the beam at a velocity v . In this case, the particles have a limited time $t = d/v$ to interact with the beam. Ignoring other broadening effects, the observed transmission spectrum will not consist of an infinitely narrow absorption feature at f_0 , but a sinc-squared absorption profile.

To explain this result, consider the signal observed by the particles passing through the beam. From their perspective, they interact with an oscillating electric field with frequency f_0 over the time period $[-t/2, t/2]$. The situation is equivalent to a sinusoid of infinite extent multiplied by a boxcar function of width t . From the results developed in Chapter 2, it is clear the Fourier transform of this signal is a delta function at f_0 convolved with a sinc function.

Defining the quantity $L = t/2$, the sinc profile that results from the situation above results in an FWHM of

$$\Delta f = \frac{1.207}{2L} . \quad (3.59)$$

The intensity absorbed is, however, proportional to the electric field absorbed squared, thus

the line shape is obtained as

$$g_{Truncation}(f - f_0) = 2L \operatorname{sinc}^2(2\pi L(f - f_0)), \quad (3.60)$$

with the resulting FWHM

$$\Delta f = \frac{0.886}{2L}. \quad (3.61)$$

Though the context of this discussion was particle absorption over a finite time. A similar situation is found within Fourier transform spectroscopy, where the detected signal is truncated by the finite spatial extent of the instrument. This will be discussed further in Section 3.4.

3.2.5 Combining Line shapes

In general, multiple broadening mechanism tend to be present for a given emission feature. The net effect results from the convolution (Equation 2.32), of all constituent line shapes. The homogeneous line shapes, corresponding to pressure and lifetime broadening, take the form of Lorentzian distributions, while the inhomogeneous line shapes, corresponding to Doppler broadening, take the form a Gaussian distribution. These mechanisms are often present simultaneously in sources under observation, and their net line shape is the convolution of the two. The ubiquity of the resulting line shape merits its own name, the Voigt profile.

The Voigt profile is defined as the convolution of a Gaussian and a Lorentzian distribution

$$V(f) = \int_{-\infty}^{\infty} G(f')L(f - f')df', \quad (3.62)$$

where G , and L , are the Gaussian and Lorentzian distributions respectively. Being symmetric functions, the resulting convolution also produces a symmetric function, and it does not matter whether G or L is used as the convolution kernel. If the Gaussian or Lorentzian is sufficiently narrow compared to the other, then the narrow distribution can be approximated

as a delta function. In these limiting cases, the Voigt profile is simply equal to the broader distribution of the two.

3.3 Radiation Processes

Radiation fundamentally results from the transition of a system from a state of high energy to a state of lower energy. The systems in which this occur, and how this energy is manifest, comes in a variety of forms.

High energy transitions can occur in the nuclei of atoms including the transmutation of one nucleus into another element or isotope of the same element via the decay of nucleons (protons, neutrons), or simply from a rearrangement of nucleons into a lower energy configuration. These processes tend to produce radiation of the highest energy corresponding to photons in the X-ray and gamma-ray band ($\lambda \lesssim 1 \text{ nm}$). The next highest energy photons result from electron transitions within atoms or molecules. The inner electron transitions are generally of higher energy than the outer transitions, and these emissions span the far-infrared to the ultraviolet regions of the electromagnetic spectrum ($\lambda \sim 0.1 \text{ mm} - 1 \text{ nm}$). This is followed by the transitions between vibrational states of molecules occupying the microwave to infrared wavelengths ($\lambda \sim 10 \text{ mm} - 1 \mu\text{m}$). Molecules may also contain energy in the form of angular momentum. These rotational states span energies extending from radio to infrared wavelengths ($\lambda \sim 10 \text{ m} - 100 \mu\text{m}$). The remaining form of radiation results from the accelerations of free electrons. These tend to be of the lowest energies occupying wavelengths in the microwave band or higher ($\lambda \gtrsim 10 \text{ mm}$).

The energy scheme presented above is a very rough approximation in which exceptions can be found in abundance. None the less, the transitions typically found within the SPIRE band consist of rotational, vibrational, and electron transitions. Spectral features resulting from rotational and electron transitions are central to this thesis, and as such, will be discussed in the following sections. Models will only be developed to the extent required to determine the energies of each transition.

3.3.1 Atomic Emission

The study of the discrete nature of atomic spectra laid the foundation for quantum mechanics. Atomic spectral lines result from the transition of electrons from high to lower energy states. To understand the nature of these states it is best to start with the simplest case, the hydrogen atom.

3.3.1.1 The Hydrogen Atom

Like most questions in quantum mechanics, the problem involves finding solutions to the Schrödinger equation which describes the system in question. For the hydrogen atom, the Schrödinger equation is

$$\frac{-\hbar^2}{2\mu} \nabla^2 \psi - \frac{Ze^2 \psi}{4\pi\epsilon_0 r} = E\psi, \quad (3.63)$$

where $\mu = m_e m_p / (m_e + m_p)$ is the reduced mass of the atom, ∇^2 is the Laplace operator, Z and e are the charge of the nucleus and electron, respectively, ϵ_0 is the permittivity of free space, and r is the radial distance from the nucleus. If the wave function ψ is factored out of the left hand side, what remains is the Hamiltonian of the system. The first term in Equation 3.63 corresponds to the kinetic energy, while the second gives the $1/r$ Coulomb potential of the electron-proton interaction.

Deriving a solution to Equation 3.63 is a laborious task and presented in its entirety in numerous textbooks, including Bransden et al., [31]. The problem is best approached in spherical coordinates where the angular momentum operators take the form

$$\bar{L}^2 = -\hbar^2 \left[\frac{1}{\sin\theta} \frac{\partial}{\partial\theta} \left(\sin\theta \frac{\partial}{\partial\theta} \right) + \frac{1}{\sin^2\theta} \frac{\partial^2}{\partial\phi^2} \right], \quad (3.64)$$

and

$$L_z = -i\hbar \frac{\partial}{\partial\phi}, \quad (3.65)$$

where θ and ϕ are the polar and azimuthal angles respectively. The wave function must be a

simultaneous eigenstate of the Hamiltonian and operators \bar{L}^2 and L_z . Since the angular momentum operators do not operate on the variable r , a separable solution of the Schrödinger equation of the form

$$\psi(\vec{r}) = R(r)Y_{lm}(\theta, \phi) \quad (3.66)$$

is expected. This solution consists of a radial part $R(r)$, and an angular part $Y_{lm}(\theta, \phi)$ with subscripts which reference quantum numbers.

3.3.1.1.1 The Angular Part

The derivation of the angular component is achieved by finding solutions to the simultaneous eigenvalue equations

$$\bar{L}^2 Y_{lm}(\theta, \phi) = l(l+1)\hbar^2 Y_{lm}(\theta, \phi), \quad (3.67)$$

and

$$L_z Y_{lm}(\theta, \phi) = m\hbar Y_{lm}(\theta, \phi). \quad (3.68)$$

The solution is once again a set of separable equations

$$Y_{lm}(\theta, \phi) = \Theta_{lm}(\theta)\Phi_m(\phi). \quad (3.69)$$

Using the definition of L_z , Equation 3.67 is satisfied when

$$\Phi_m(\phi) = \frac{1}{\sqrt{2\pi}} e^{im\phi}, \quad (3.70)$$

where a normalization term has been included. The wave function is required to be invariant under the transformation $\phi \rightarrow \phi + 2\pi$. This forces m to take integer values, though it is not restricted in sign. Since m is directly associated with the z projection of the electron angular momentum vector, and thus in some sense is a measure of its magnetic dipole orientation, this variable is referred to as the magnetic quantum number.

A power series is encountered when solving for $\Theta_{lm}(\theta)$, which diverges for arbitrary values of l . However, if l takes integer values greater than or equal to zero, the series terminates after a finite number of terms, and the series is then convergent. With this information, possible values of l and m can be defined. Noting that $|L^2| \geq L_z^2$, and that Equation 3.67 requires that $l(l+1) \geq m^2$, it can be concluded that

$$l = 0, 1, 2, \dots \quad (3.71)$$

and

$$m = -l, -l+1, \dots, l-1, l. \quad (3.72)$$

These limitations are achieved by mandating that the wave function have unique values at each point in space, while being a simultaneous eigenfunction of the \bar{L}^2 and L_z operators. In this context, l is referred to as the orbital quantum number. However, in the more general case where the spatial dependence is neglected, l can take half integer values with the relationship between m and l unchanged. In the general case l is replaced by j , and the allowed values for numbers which describe angular momentum are, $j = 0, 1/2, 1, 3/2, \dots$ Electron spin is an example of a one-half integer angular momentum.

Solutions for $\Theta_{lm}(\theta)$ are identified as the associated Legendre functions [31], with appropriate normalization factors. The final solution to the angular component of the wave function (Equation 3.69) is then obtained as a set of equations called spherical harmonics. Spherical harmonics make up a complete set of orthogonal basis vectors which can be used to describe an arbitrary distribution over a sphere. They are much like the harmonics of a Fourier series, but applied to a different geometry. Plotting the square of the spherical harmonic functions traces the relative shape of the so called electron orbitals for given values of l and m . It is worth noting that $|\Psi(\vec{r})|^2 = |R(r)|^2 |Y_{lm}(\theta, \phi)|^2 = |R(r)|^2 |\Theta_{lm}(\theta)|^2 / 2\pi$, and the resulting expectation value of the wave function is independent of ϕ , and possesses cylindrical symmetry.

3.3.1.1.2 The Radial Part

Solutions for the radial component of Equation 3.66 require as much labour as the angular component, and once again only key points will be highlighted. The Schrödinger equation for the hydrogen atom can be rewritten as

$$\frac{d^2u(r)}{dr^2} + \frac{2\mu}{\hbar^2}[E - V_{eff}]u(r) = 0, \quad (3.73)$$

where $u(r) = rR(r)$, and

$$V_{eff} = -\frac{Ze^2}{4\pi\epsilon_0 r} + \frac{l(l+1)\hbar^2}{2\mu r^2}. \quad (3.74)$$

The effective potential V_{eff} consists of the central potential resulting from the Coulomb force, and a term dependent on l known as the centrifugal barrier.

Making two further substitutions,

$$\rho = r\sqrt{-8\mu E/\hbar^2} \quad (3.75)$$

and

$$n = \frac{Ze^2}{4\pi\epsilon_0\hbar}\sqrt{\frac{-\mu}{2E}}, \quad (3.76)$$

the Schrödinger equation takes the form

$$\left[\rho \frac{d^2}{d\rho^2} + (2l+2-\rho) \frac{d}{d\rho} + (n-l-1) \right] g(\rho) = 0. \quad (3.77)$$

The function $g(\rho) = \sum_{k=0}^{\infty} c_k \rho^k$ is a power series constructed to have the same behaviour as $u(\rho)$ in the limiting cases of $\rho \rightarrow 0$, or ∞ . The above differential equation has the same form as the Kummer-Laplace differential equation [56]. The general solution to this equation is the confluent hypergeometric function [31]. Physically acceptable convergent solutions only occur when $(n-l-1)$ is a negative integer or zero. As previously discussed, l must be a positive integer or zero. This new restriction limits n to a positive integer, and limits the

upper value of l to $n - 1$. Under these restrictions, the confluent hypergeometric function reduces to the associated Laguerre function. After including appropriate normalization coefficients, the associated Laguerre function constitutes the final solution to the radial part of Equation 3.66.

Returning to Equation 3.76, the energy of each bound state can be calculated as

$$E = -\frac{\mu}{2} \left(\frac{Ze^2}{4\pi\hbar\epsilon_0 n} \right)^2. \quad (3.78)$$

It is worth noting that the energy levels in the hydrogen atom, in the absence of any electric or magnetic fields, are determined only by the principle quantum number n .

3.3.1.2 Polyelectronic Atoms

The introduction of additional electrons into an atom results in the inclusion of an additional term into the Hamiltonian operator of the electron states. This new term represents the mutual repulsion between electrons and is given by

$$\bar{H}_{ee} = \frac{e^2}{4\pi\epsilon_0 d}, \quad (3.79)$$

where d is the distance between electrons. Unsurprisingly, this new term is similar to the term arising from the electron-nucleus interaction. The introduction of this term prevents the convenient separation of variables used to derive the exact solutions of the wave equation for the hydrogen atom. It is, however, still conceptually useful to adopt the orbital configuration of the hydrogen atom as a rough guide for the spatial distribution of electrons. Doing so is formally referred to as the orbital approximation.

The mutual electron interaction breaks the symmetry of the orbital states, and, as a result, electron energy states are no longer determined by the principle quantum number alone. Numerical calculations of electron energies are complicated, and it is generally more beneficial to develop a framework for understanding the orientation and labelling of

the energy levels. This approach is based on electron angular momentum coupling.

Mutual electron interaction breaks the commutation relationship between the z projection of the electron orbital angular momentum operator and the Hamiltonian. As such, a new set of good quantum numbers is required to describe electron states. A new set of operators is generated using the Russell-Saunders coupling scheme [49]:

$$\begin{aligned}
 \bar{L} &= \sum_{i=1}^N \bar{l}_i \quad (\text{Total Orbital Angular Momentum}) \\
 \bar{S} &= \sum_{i=1}^N \bar{s}_i \quad (\text{Total Spin Angular Momentum}) \\
 \bar{J} &= \bar{L} + \bar{S} \quad (\text{Total Angular Momentum}) .
 \end{aligned} \tag{3.80}$$

Bars now indicate operators, and a convention of using lower case letters when referring to individual electrons, and upper case letters when referring to coupled states has been employed. This coupling scheme takes the sum over the angular momentum components of all electrons in the atom. Similarly, the total z projection of the angular momentum is generated as would be expected from the above equations (eg, $\bar{L}_z = \sum_{i=1}^N \bar{l}_{zi}$). The new set of mutually commuting operators is then \bar{H} , \bar{L}^2 , \bar{S}^2 , \bar{L}_z , \bar{S}_z , with corresponding good quantum numbers, n , L , S , M_L , M_S .

3.3.1.3 Atomic Configuration

It is now worth introducing the notation used to describe the distribution of electrons in an atom. The principle quantum number n defines a shell, and internal to this shell are subshells defined by the orbital quantum number l . The $(2l + 1)$ potential orientations of the orbital angular momentum vector corresponding to allowed values of the magnetic quantum number m_l , and the $(2s + 1)$ orientations of the electron spin angular momentum vector corresponding to the allowed spin quantum number m_s , define all electron states within a subshell. Recall that the allowed values of these quantum numbers for an electron

are:

$$\begin{aligned}
 n &= 1, 2, 3, \dots \\
 l &= 0, 1, \dots, n-1 \\
 m_l &= l, l-1, \dots, -l \\
 m_s &= \frac{1}{2}, -\frac{1}{2}.
 \end{aligned}
 \tag{3.81}$$

Each compliment of quantum numbers defines an electron state. The Pauli exclusion principle states that no two particles can share the same quantum state, so each electron in an atom will have its own unique set of quantum numbers. The aufbau (“building-up”) principle states that the lowest energy orbital available will be filled before higher energy orbitals become occupied. Though this principle is not rigid, and different atoms have a different distributions of energy levels for each orbital, in general electron occupancy tends to build up in shells, with the subshells being filled in order of increasing l . By convention, subshells are denoted with letters such that $l = 0, 1, 2, 3, 4, 5, \dots \rightarrow s, p, d, f, g, h, \dots$

Electron configurations in atoms are built up with the following notation

$$\text{Electron Configuration} = n l^a, \tag{3.82}$$

with a being the number of electrons in the l subshell. For example, carbon has the following configuration $1s^2 2s^2 2p^2$, where the first shell is filled, and the second shell is filled up to the p subshell which contains two electrons. Notice how this notation uses lower case letters, and thus implies not accounting for angular momentum coupling.

In the case of angular momentum coupling, the following notation is used

$$\text{Coupled State} = {}^{(2S+1)} L \text{ (Term)}, \tag{3.83}$$

where S and L are the sum of the spin and orbital angular momentum quantum numbers

of the electrons in the atom respectively. The $(2S + 1)$ denotes the spin degeneracy of the state and indicates whether the state is a singlet, doublet, triplet, etc., for values of 1, 2, 3, etc. The expression in equation 3.83 is called a “term”. Understanding the notation is best accomplished with an example.

Referring back to the electron configuration of carbon, the s subshells correspond to an orbital angular momentum of 0. They also have a pair of electrons with opposite spin orientations, so only the $2p^2$ subshell need be considered. Each electron is in the p subshell and thus contains an orbital angular momentum of $l_1 = l_2 = 1$. These vectors can have multiple orientations as indicated by their magnetic quantum number, as a result, the net orbital angular momentum can take the following values $L = l_1 + l_2, l_1 + l_2 - 1, l_1 - l_2 = 2, 1, 0 \rightarrow D, P, S$. Accounting for the net spin in each case requires more work, however, since there are only two electrons, their orientations are either parallel or antiparallel giving $S = 1, 0$. As a result, the spin multiplicity is either 3 or 1, respectively.

It turns out that there are a total of 15 coupled states: five 1D , nine 3P , and one 1S . The total number of states can be calculated from the total momentum multiplicity $(2S + 1)(2L + 1)$, where S and L take their maximum values. The total multiplicity of each term can also be calculated using the same formula where S and L take their appropriate value (Equation 3.83) for the term in question.

3.3.1.4 Fine Structure

An additional perturbation term is present in the Hamiltonian of polyelectronic atoms. Electrons “orbiting” the nucleus produce a current loop, which subsequently produce magnetic dipoles. In addition, each “spinning” electron induces its own magnetic dipole. The interaction of these two dipoles is called spin-orbit coupling, and adds the following term to the Hamiltonian

$$\bar{H}_{so} = \xi \vec{l} \cdot \vec{s} . \quad (3.84)$$

To expand this term to include all electrons in the system, the Russell-Sanders coupling scheme along with the Wigner-Eckart theorem is employed to obtain

$$\bar{H}_{so} = \zeta \bar{L} \cdot \bar{S}. \quad (3.85)$$

The spin-orbit coupling factor ζ parameterizes the strength of the coupling, and tends to increase with the number of electrons. This interaction only has a significant effect on isolated terms (Equation 3.83).

Similar to angular momentum coupling, spin-orbit coupling breaks the commutation relationship between \bar{L}_z and the Hamiltonian. The new set of commuting operators are, $\bar{H}, \bar{L}^2, \bar{S}^2, \bar{J}^2, \bar{J}_z$, with corresponding good quantum numbers n, L, S, J, M_J .

Spin-orbit coupling breaks the symmetry within each term, separating each state into its own unique energy level. The notation used to express each state is an extension of the angular momentum coupling notation given by

$$\text{Coupled State} = {}^{(2S+1)}L_J, \quad (3.86)$$

where $J = L + S$ is the total angular momentum of the state. For example, consider the 3P term of the carbon atom. This state corresponds to a net spin and orbital angular momentum equal to one for both cases. Accounting for the various orientations of S and L , total angular momentum can take the values $J = S + L, S + L - 1, \dots, |S - L| \rightarrow 2, 1, 0$. That is, the 3P term of carbon can be split into the three states ${}^3P_0, {}^3P_1$, and 3P_2 .

The splitting of terms into additional energy states as a result of spin-orbit coupling is formally referred to as fine-splitting, and results in the fine structure lines observed in spectra. The energy between adjacent states can be determined by calculating the expectation values of the spin-orbit coupling perturbation term. After making the substitutions $J = J + 1$

and $J = J$, and then calculating the difference, the energy splitting is given by

$$E_{J+1} - E_J = \zeta(J+1) . \quad (3.87)$$

This approximation is called the Landé interval rule, and expresses the interesting property of being proportional to $(J+1)$. The spin-orbit coupling constant is not always positive, and its sign determines if the ground state corresponds to a high J value or a low value. Determining ζ theoretically is complex, but a few rules exist to help determine the ground state.

Hund's rules provide a useful criteria for determining the ground state.

1. The term with the highest spin multiplicity $(2S+1)$, corresponds to the ground state.
2. For terms with the same spin multiplicity, the term with the highest orbital angular momentum L , corresponds to the ground state.
3. If a subshell is half full or less, ζ will be positive. If the subshell is more than half full, ζ will be negative.

Since carbon has a maximum spin multiplicity of 3, and has a subshell that is less than half full, the ground state of carbon is 3P_0 . A diagram illustrating the energy level splitting that results from mutual electron interactions and spin-orbit coupling for carbon is shown in Figure 3.3.

Hyperfine splitting can also occur as a result of nuclear spin coupling to the total angular momentum of the electron. This effect is small, resulting in spectral line splitting on the order of 1 cm^{-1} , and will not be discussed further.

3.3.2 Rotational Emission

Atoms that bind together to form more complex molecular structures acquire additional modes of energy storage other than electron states. One mode of energy storage of great

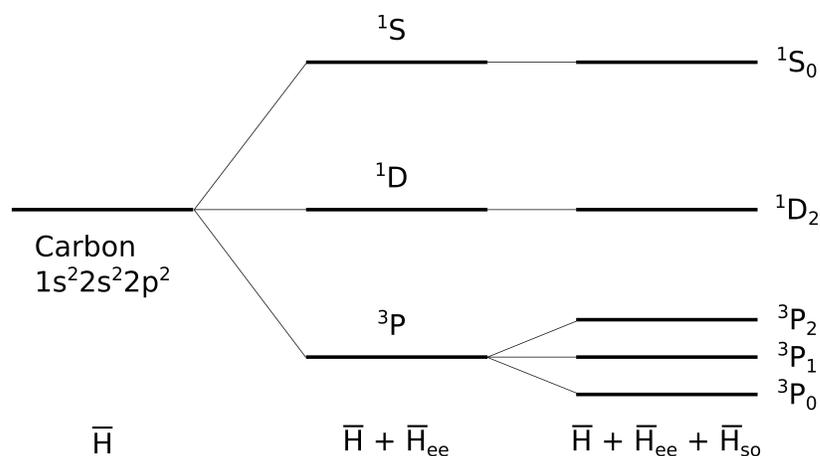


Figure 3.3: Energy level splitting resulting from electron interaction (center) and spin-orbit coupling (right) in a carbon atom.

importance to infrared spectroscopy is molecular rotation. Many rotational emissions are typical in SPIRE spectra, and an introduction to the subject is presented here.

Fundamental to the discussion of any rotating system is the moment of inertia. The moment of inertia of a system of N particles is defined with respect to a rotational axis as

$$I = \sum_{i=1}^N m_i r_i^2, \quad (3.88)$$

where m_i indicates the particle mass, and r_i indicates the orthogonal distance from the rotational axis to the particle in question. A moment of inertia can be defined for three orthogonal axes passing through the origin, much like the orthogonal unit vectors in a Cartesian coordinate system. The origin of the system is conveniently chosen to be the Center of Mass (COM), of the system of particles, which is defined as

$$COM = \frac{\sum_{i=1}^N m_i r_i}{\sum_{i=1}^N m_i}. \quad (3.89)$$

The moment of inertia describes the distribution of mass about a rotational axis, and in rotational physics plays much the same role as mass in translational physics.

Molecules can be classified based on the relative values of their moment of inertia about the three principle axes. The orientation of these axis are chosen such that the z -axis consists of the highest degree of rotational symmetry [49], with the remaining axis chosen somewhat arbitrarily while attempting to maximize rotational symmetry in one of the remaining planes. The axes can then be relabelled such that $I_A \leq I_B \leq I_C$. With this convention, the different classes of molecules are defined as follows:

- Linear: $I_C = I_B, I_A = 0$.
- Spherical: $I_A = I_B = I_C$.
- Prolate: $I_A < I_B = I_C$.
- Oblate: $I_A = I_B < I_C$.
- Asymmetric: $I_A < I_B < I_C$.

The specific case of linear molecules will be discussed in more detail as they are the most pertinent to this thesis. However, the principles presented apply equally to more complex molecules.

3.3.2.1 Rigid Rotators

The intensity of the radiation emitted is to a large extent determined by the transitional dipole moment (Equation 3.29), which is itself dependent on the dipole moment of the system or molecule. For this reason molecules with a high degree of charge symmetry, such as H_2 , do not emit strong rotational or vibrational spectral lines.

The first step in calculating the energy levels of rotational states of molecules is to approximate the molecules as rigid rotators. In classical mechanics, the rotational kinetic energy of a rotating body is given by

$$E_k = \frac{1}{2} \vec{I} \cdot \vec{\omega}^2 = \frac{J^2}{2I} . \quad (3.90)$$

In the case of a linear molecule, $I_z = 0$ as a result of the definition of the z axis given above, and the geometry of a linear molecule where all atoms reside on the z axis. The final expression is obtained by noting that angular momentum L , is the product of the moment of inertia I , and the angular velocity ω , and that $L_x^2 + L_y^2 = L^2$. The general angular momentum variable J has also been used in place of L .

The Hamiltonian can then be written as

$$\bar{H}_{rot} = \frac{\bar{J}^2}{2I}. \quad (3.91)$$

The eigenvector describing the rotational states must be the spherical harmonics derived for the general case of angular momentum. As such the eigenvalues for the \bar{J}^2 operator are $J(J+1)\hbar^2$, with the complete energy eigenvalues given by

$$E_J = \frac{J(J+1)\hbar^2}{2I}. \quad (3.92)$$

The rotational constant B can then be defined as

$$B = \frac{\hbar^2}{2I} = \frac{h^2}{8\pi^2 I}. \quad (3.93)$$

It can easily be shown that the energy difference between rotational states is then,

$$\begin{aligned} E_{(J+1) \rightarrow J} &= B[(J+1)(J+2) - J(J+1)] \\ &= 2B(J+1). \end{aligned} \quad (3.94)$$

This approximation is based on the assumption that atoms are rigidly attached to one another, and is sufficient for most cases. However, a better approximation is one where atoms are attached with springs which can stretch inducing a proportional restoring force. This effect is called centrifugal distortion.

3.3.2.2 Centrifugal Distortion

For this discussion it is easiest to envision the molecule as diatomic, though the ideas apply to all linear molecules. This approximation assumes that atoms in a molecule are attacked by a force which acts much like a spring does. Stretching of an atom away from its equilibrium position induces a restoring force described by Hooke's law

$$F_r = -k\Delta r, \quad (3.95)$$

where k is the spring constant, and Δr is the distance from equilibrium. At the same time the masses making up the rotating system tend to want to continue their linear trajectory in accordance with Newton's first law. This results in a "force" called the centrifugal force given by

$$F_c = \mu\omega^2 r = \frac{J^2}{\mu r^3}, \quad (3.96)$$

where μ is the reduced mass of the system, and r is the separation between the two atoms.

In equilibrium the restoring and centrifugal forces are equal. After equating the two forces one finds the rotational energy to be

$$E_J = BJ(J+1) - D(J(J+1))^2 + \dots \quad (3.97)$$

Higher order terms of low significance have been neglected. The centrifugal distortion constant D is given by the Kratzer relationship [49],

$$D = \frac{4B^3}{\omega_v^2}, \quad (3.98)$$

where both B and ω_v , the molecules vibrational frequency, take their equilibrium values. Vibrational emission has been neglected in this thesis, but it should be mentioned that even in the ground state, a harmonic oscillator will have a positive non-zero energy. Finally the

energy difference between rotational states is obtained as

$$E_{(J+1) \rightarrow J} = 2B(J+1) - 4D(J+1)^3. \quad (3.99)$$

It should be noted that the centrifugal distortion constant is frequency dependent and that the correction to rotational energy transitions is relatively small. For ^{12}CO in the vibrational ground state, $B \sim 57.64$ GHz, and $D \sim 0.18$ MHz.

3.4 FTS Instrumentation and Processing

3.4.1 Michelson Interferometer

The Michelson interferometer is perhaps the simplest example of a Fourier transform spectrometer. The design was invented by Michelson, and later popularized in 1887 by the famous Michelson and Morley experiment which attempted to measure the relative motion of Earth through the hypothesized Luminiferous Aether [57]. The simple design makes it ideal for introducing the theory of interferometers. A slightly more complex design is shown in Figure 3.4. The key difference between this design and the classic set-up is the use of two beam splitters (green), and double rooftop mirrors (black), instead of one beam splitter and orthogonally oriented flat mirrors.

An electromagnetic wave enters the instrument through the entrance pupil. Only one component of the wave need be considered for the theoretical treatment, and it is convention to focus on the electric component as this is the component that detectors are sensitive to. Generally this is followed by collimating optics such that the electrical disturbance propagating through the system can be approximated as a plane wave

$$E_{in} = E_0 e^{i2\pi(x\sigma_0 - f_0 t)}, \quad (3.100)$$

with E_0 being the amplitude of the electric wave with constant frequency and wavenumber f_0 and σ_0 , respectively. The propagation of this wave in space and time is given by x

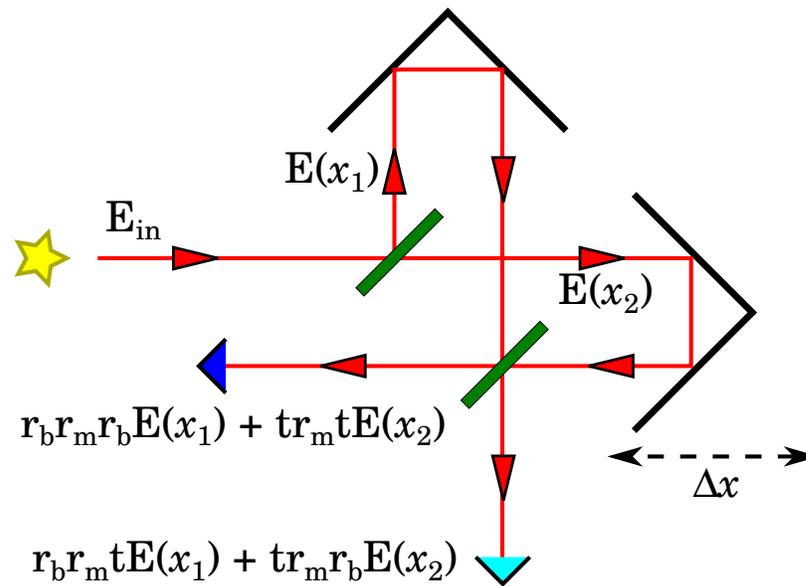


Figure 3.4: A diagram of a dual input/output Michelson interferometer. The input electric field E_{in} is split into two beams by a beam splitter (green), with then propagate along different optical paths. These beams are then re-combined by a second beam splitter which leads to the balance output (cyan), and unbalanced output (blue). Mathematically, beam interaction with optical components can be represented with amplitude coefficients which indicate the change in phase and intensity of the input beam. The electric field at either output is a combination of both constructive and destructive interference which varies as a result of the optical path difference between the two beams. The resulting signal is called an “interferogram”.

and t , respectively. The wave is then incident on a beam splitter, an optical component containing both reflective and transmissive properties such that the energy of the incoming wave is, in the ideal case, split equally between the two output paths. Both beams then propagate towards rooftop mirrors which backwards reflect the beams along a parallel path in the opposite direction. One beam, in this case $E(x_1)$, is incident on a stationary mirror, while the other beam $E(x_2)$, is incident on a movable mirror. Both reflected beams are then incident on a second beam splitter, which in this case acts as a beam combiner. The beam combiner has two outputs, either one of which can be used to record the resulting signal.

The signal recorded at either output port is the superposition of the two beams which have travelled along different optical paths. The Optical Path Difference (OPD), can induce varying degrees of constructive and destructive interference which modulates the intensity received at either output. The OPD is varied by adjusting the position of the movable mirror (Δx), and moving this mirror through its entire range of motion constitutes one complete scan. It is important to note that optical and infrared detectors do not measure the electric field directly. Instead, they measure an integrated power over the detector specific response time, which is typically observed as a voltage. The energy of the electric field is proportionate to the square of the electric field, making this the relevant quantity to be derived at either output. The interference signal measured at either output during a scan is referred to as an interferogram.

As the input electric field propagates through the optical system, it interacts with and is modified by each optical component producing the resultant beams of the interferometer. Mathematically, these interactions take the form of amplitude coefficients. For example, at the junction of the first beam splitter, $E(x_1)$ is produced by multiplying the input field, E_{in} , by r_b , the beam splitter reflection amplitude coefficient. Likewise $E(x_2)$ is the product of the input electric field and τ , the beam splitter transmission coefficient. At this point the

beams diverge which modifies x in the exponent of Equation 3.100:

$$\begin{aligned} E(x_1) &= E_0 e^{i2\pi(x_1\sigma_0 - f_0t)} \\ E(x_2) &= E_0 e^{i2\pi(x_2\sigma_0 - f_0t)} . \end{aligned} \quad (3.101)$$

After including the appropriate coefficients, the beams are given by:

$$\begin{aligned} \text{Beam 1} &= r_b E(x_1) \\ \text{Beam 2} &= \tau E(x_2) . \end{aligned} \quad (3.102)$$

At each corner reflector, both beams pick up an amplitude coefficient from interaction with the double mirror system, r_m . Finally at the second beam splitter, each beam picks up an additional τ or r_b amplitude coefficient as required. It should also be kept in mind that, implicit with every reflection coefficient, is a π radian phase shift of the electric field. This process is made more clear by considering Figure 3.4.

The complement of amplitude coefficients is what distinguishes the two outputs of the interferometer. The cyan output in Figure 3.4 is referred to as the balanced output since the two composite beams contain the same amplitude coefficients resulting in equivalent phase shifts due to the optical system. The electric field at the balanced output is

$$E_b = E_1 + E_2 = E_0 e^{-i2\pi f_0t} r_b r_m \tau (e^{i2\pi\sigma_0 x_1} + e^{i2\pi\sigma_0 x_2}) , \quad (3.103)$$

where $E_1 = E_0 r_b r_m \tau e^{i2\pi(\sigma_0 x_1 - f_0t)}$, and $E_2 = E_0 r_b r_m \tau e^{i2\pi(\sigma_0 x_2 - f_0t)}$. The detector measures time averaged intensity which is proportional to E^*E , with the asterisk indicating the complex conjugate. Thus the time averaged intensity of equation 3.103 must be calculated. The derivation is made more compact by exploiting some useful notation

$$\frac{1}{T} \int_0^T |E_b|^2 = \langle |E_b|^2 \rangle . \quad (3.104)$$

With this in mind, the time averaged intensity is given by,

$$\begin{aligned}
 I_b &\propto \langle |E_b|^2 \rangle = \langle E_b^* \cdot E_b \rangle \\
 &= \langle (E_1 + E_2) \cdot (E_1^* + E_2^*) \rangle \\
 &= \langle |E_1|^2 \rangle + \langle |E_2|^2 \rangle + \langle E_1^* \cdot E_2 \rangle + \langle E_2^* \cdot E_1 \rangle .
 \end{aligned} \tag{3.105}$$

It is easily seen that $\langle |E_1|^2 \rangle = \langle |E_2|^2 \rangle = (E_0 r_b r_m \tau)^2$, since $e^{i2\pi(\sigma_0 x_i - f_0 t)} \times e^{-i2\pi(\sigma_0 x_i - f_0 t)} = 1$.

For the remaining terms,

$$\begin{aligned}
 \langle E_1 \cdot E_2^* \rangle + \langle E_2 \cdot E_1^* \rangle &= (E_0 r_b r_m \tau)^2 \left[e^{i2\pi\sigma_0(x_1 - x_2)} + e^{-i2\pi\sigma_0(x_1 - x_2)} \right] \\
 &= 2(E_0 r_b r_m \tau)^2 \cos(2\pi\sigma_0(x_1 - x_2)) .
 \end{aligned} \tag{3.106}$$

The time dependent terms have cancelled as a result of the complex conjugate multiplication, while in the last line, the exponential values have been replaced by their cosine equivalent. Putting Equations 3.105, and 3.106 together, the final result for the intensity of the balanced output is

$$I_b = 2(E_0 r_b r_m \tau)^2 (1 + \cos(2\pi\sigma_0(x_1 - x_2))) . \tag{3.107}$$

It is common to recast this equation into a more convenient form, with optical efficiency $\eta_0 = r_m^2$, beam splitter efficiency $\eta_b = 4(r_b \tau)^2$, input intensity $E_0^2 = I_0$, and optical path difference $z = x_1 - x_2$. With these substitutions, Equation 3.107 becomes

$$I_b(z) = I_0 \eta_0 \eta_b \left[\frac{1 + \cos(2\pi\sigma_0 z)}{2} \right] . \tag{3.108}$$

With this notation, a beam splitter efficiency of 1 indicates that precisely half the incident beam intensity is reflected, and half is transmitted. Similarly, an optical efficiency of 1 indicates that the mirrors reflect 100% of the incident beam with no absorption loss.

Since energy must be conserved, the simplest way to determine the intensity at the

unbalanced output is to subtract the balanced output from the input intensity, $I_u(z) = I_0 - I_b(z)$. After a bit of algebra, one obtains

$$I_u(z) = I_0 \eta_0 \eta_b \left[\frac{1 - \cos(2\pi\sigma_0 z)}{2} \right] + I_0(1 - \eta_0 \eta_b) \quad (3.109)$$

The most significant difference between the two outputs is a π radian phase shift of the cosine term, that is the information in one output is the compliment of the other, and when one sees constructive interference, the other sees destructive interference. For this reason, it is often the case that only one output is recorded with the constant component removed by a high-pass filter.

There are however, advantages to measuring both outputs. The sum of the two output intensities is equivalent, or at least proportional to, the input beam intensity. As such, tracking time dependent intensity fluctuations of the source, which produce artefacts in the spectral domain, is possible, and can be corrected for in the interferogram. Fringe contrast or signal strength is improved by taking the difference between the two outputs

$$I(z) = I_b(z) - I_u(z) \propto I_0 \eta_0 \eta_b \cos(2\pi\sigma_0 z) , \quad (3.110)$$

where the constant offset has been ignored. This equation represents the interferogram proper of an ideal Michelson type interferometer for a monochromatic source.

Equation 3.110 is readily generalized to polychromatic sources by integrating over the constituent frequencies

$$I(z) = \int_0^{\infty} I_0(\sigma) \eta_0 \eta_b \cos(2\pi\sigma z) d\sigma . \quad (3.111)$$

In general, the optical and beam splitter efficiencies are also frequency dependent, but with only slight variations over limited bandwidths.

The quantity $I_0(\sigma)$ is none other than the spectral energy distribution, or spectrum of the

source. Defining the spectrum to be an even function $I_0(\sigma) = I_0(-\sigma)$, is mathematically convenient, allowing for the integration over the frequencies $(-\infty, \infty)$. The physical implication of this assertion is a doubling of energy in the spectral domain. To account for this, the spectral energy distribution is divided by two, $S(\sigma) = \eta_0 \eta_b I_0(\sigma)/2$, where this new variable has absorbed the efficiency terms as well. With these considerations, Equation 3.111 becomes

$$I(z) = \int_{-\infty}^{\infty} S(\sigma) e^{i2\pi\sigma z} d\sigma . \quad (3.112)$$

That is, the interferogram is equivalent to the inverse Fourier transform of the spectrum (Equation 2.8). Note that the imaginary sine term resulting from Euler's formula is negligible as the product of an even function (the spectrum), and an odd function (the sine term), is an odd function, which integrates to zero over equal intervals across the origin.

The spectrum is recovered by taking the forward Fourier transform of the recorded interferogram (Equation 2.9)

$$S(\sigma) = \int_{-\infty}^{\infty} I(z) e^{-i2\pi\sigma z} dz . \quad (3.113)$$

Note that for interferometers, the reciprocal units are position (z ; [cm]), and spatial frequency or wavenumber ($\sigma = 1/\lambda$; [cm^{-1}]).

The interferogram is a continuous signal which in practice can only be sampled discretely over a finite range. At this point all the principles developed in Chapter 2 apply, though there are a few more topics to be explored in the context of spectroscopy.

3.4.2 Synthetic Spectral Resolution

As previously discussed in Section 3.2, the finite extent of interferometers results in only a finite portion of the interferogram being measured. This limitation results in an effective truncation of the interferogram which, in theory, extends to an infinite OPD. The result of this truncation is a sinc line shape. This effect also occurs in discretely sampled

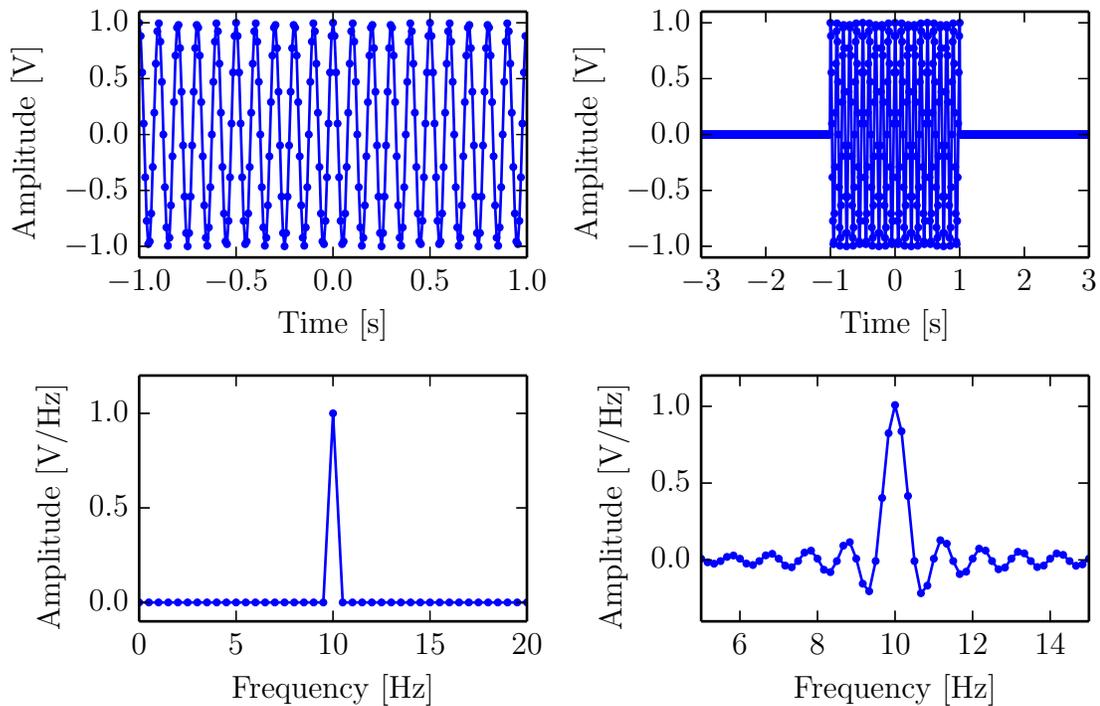


Figure 3.5: A comparison between unpadded and padded interferograms and resulting line shapes. In the unpadded case (left), the resulting line shape is approximately a delta function with a single peak at 10 Hz. The padded interferogram (right), produces the same peak at 10 Hz, with additional points being interpolated onto the expected sinc line shape. Note the difference in scale of the frequency axis between the two cases.

interferograms, however, the abrupt termination of the discrete interferogram may result in a sinc line shape which is sampled such that only the peak and zero crossings are reported. This gives the line shape the appearance of a triangular function as can be seen in the left hand panels of Figure 3.5.

Appending zeros onto both sides of the interferogram adds no additional information, but does effectively extend the maximum OPD of the signal to be processed. Since the frequency sampling Δf is inversely related to the maximum OPD, the process of appending additional zeros onto the interferogram decreases the frequency sampling. This technique, called zero padding, is shown in Figure 3.5. By superimposing the spectra of the abruptly terminated and zero padded interferogram, as shown in Figure 3.6, it is easy to see the effects of this technique. As the figure shows, zero padding merely “fills in” additional

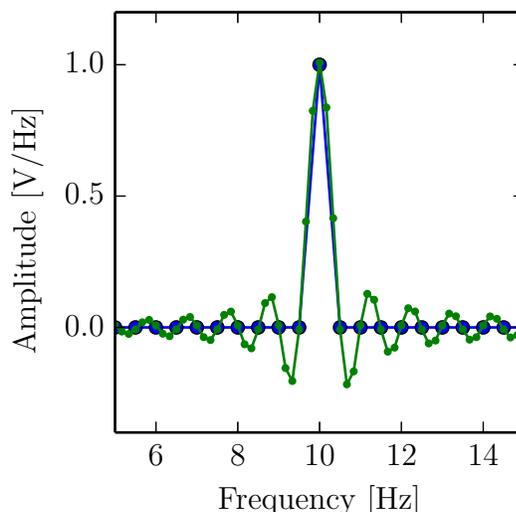


Figure 3.6: The padded (green) and unpadded (blue) line shapes resulting from the interferograms in Figure 3.5 have been plotted on top of one another. The central peak and zero crossings of the padded line intersect with the points from the unpadded line. This helps to demonstrate how no new information is created by zero-padding, additional points are merely interpolated onto the sinc line shape.

points of the sinc line shape with the zero crossings coinciding with the data points of the unpadded spectrum. This figure also makes clear the fact that spectral resolution is more properly parameterized by the line shape rather than the frequency spacing in the spectral domain.

Zero padding is thus an effective method of spectral interpolation. This technique finds convenient application when attempting to give equal spectral gridding to repeated observation of the same source when each measurement has unequal maximum optical path differences. This is the case for the off-axis detectors in the SPIRE IFTS arrays, and thus zero padding is a regular step in the SPIRE data processing pipeline [58, 59]. Zero padding can also be used to increase the number of data points to a power of 2 resulting in more efficient computation of the FFT (Section 2.3.2).

3.4.3 Phase Correction

When an FTS is used to observe a real source, one would expect a real valued spectrum after performing the Fourier transform of the interferogram. This implies a real valued symmetric interferogram (Subsection 2.2.1), which would be expected from an ideal interferometer when considering Equation 3.111. In practice, however, asymmetries may be incorporated into the interferogram from numerous sources, and produce spectra with associated deleterious effects including inaccurate amplitudes, line-shapes, and SNR. For this reason, it is worth exploring the cause of the asymmetry, and how to correct the resulting spectrum without losing spectral integrity.

As discussed in Subsection 2.2.1, spectral phase is a measure of asymmetry in the reciprocal Fourier domain. A general interferogram can be described mathematically as

$$I_{asym}(z) = \int_{-\infty}^{\infty} S(\sigma) e^{i\phi(\sigma)} e^{i2\pi\sigma z} d\sigma, \quad (3.114)$$

which consists of the inverse Fourier transform of the physical spectrum, $S(\sigma)$, multiplied by a phase term, $e^{i\phi(\sigma)}$. The total phase $\phi(\sigma)$, represents the sum of all the contributions to the phase, and is in general frequency dependent. The main sources of phase shift are: not sampling the Zero Path Difference (ZPD), dispersive properties of optical components and electronics, and random variations in the interferogram signal [29].

If the ZPD is not sampled, then interferogram sampling can be modelled as the properly sampled OPD with a constant offset, α (i.e., $z' = z + \alpha$). Using this substitution in the Fourier transform exponential, one obtains

$$e^{i2\pi\sigma z'} \rightarrow e^{i2\pi\sigma(z+\alpha)} = e^{i2\pi\sigma z} e^{i2\pi\sigma\alpha} \quad (3.115)$$

$$\therefore \phi(\sigma) = 2\pi\sigma\alpha.$$

This result indicates that improperly sampling the ZPD produces a linearly dependent phase shift.

The dispersive properties of optical components can be modelled in a similar way. The optical path of one frequency of light is dependent on the frequency dependent index of refraction n_σ of the dispersive medium. As a result, the optical path difference will be proportional to $n_\sigma d$, where d is the thickness of the dispersive medium. Performing the substitution $z' = z + n_\sigma d$, the Fourier transform exponent becomes

$$e^{i2\pi\sigma z'} \rightarrow e^{i2\pi\sigma(z+n_\sigma d)} = e^{i2\pi\sigma z} e^{i2\pi\sigma n_\sigma d}$$

$$\therefore \phi(\sigma) = 2\pi\sigma n_\sigma d .$$
(3.116)

Well chosen optics will contain dispersive media with a near constant index of refraction over the bandwidth of interest. In this case $n_\sigma = \text{constant}$, resulting in a linear phase shift just as in Equation 3.115, producing an effective shift in ZPD. In general however, the index of refraction has a more complex frequency dependence and the result is a non-linear phase shift term [33].

Random phase components result from random fluctuations in the interferogram. Random noise on average contains equal parts symmetric and antisymmetric fluctuations. As such, on average, equal noise intensity is mapped into the real and imaginary components of the spectral domain.

Performing the inverse Fourier transform on Equation 3.114, one obtains

$$\mathcal{F}^{-1}\{I_{asym}(z)\} = S_{obs}(\sigma) = S(\sigma)e^{i\phi(\sigma)}$$

$$= S(\sigma)[\cos(\phi(\sigma)) + i\sin(\phi(\sigma))] .$$
(3.117)

By taking the absolute value of the complex function above (e.g., $|f| = \sqrt{f^*f}$), and using the trigonometric identity $\cos^2(\theta) + \sin^2(\theta) = 1$, the above expression becomes

$$|S_{obs}(\sigma)| = |S(\sigma)| .$$
(3.118)

That is, by taking the absolute value of the observed spectrum, one ignores the phase and

recovers the absolute value of the physical spectrum. The downside of this simple approach is that all the noise present in the imaginary domain is folded over into the real domain. This transformation unnecessarily increases the spectral noise by a factor of $\sqrt{2}$, which can be avoided with a more sophisticated approach.

More sophisticated approaches involve calculating the phase term directly, then multiplying its reciprocal, $e^{-i\phi(\sigma)}$, with the observed spectrum. Alternatively, by employing the convolution theorem (Equation 2.33), phase correction can be implemented within the interferogram. This approach is called the Forman method of phase correction. What follows is a demonstration of the Forman method of phase correction.

Since phase is a measure of asymmetry, it should be clear that an interferogram with information on both sides of the ZPD is required to calculate the phase. In the following example, a double-sided interferogram which extends from $-L_{max}$ to L_{max} is assumed. In practice, it is more efficient to use a so called single-sided interferogram which extends only slightly into the negative OPD domain, with the remainder of the scan range allocated to positive OPD. Since spectral resolution is limited by the maximum OPD, a single-sided interferogram allows for the highest spectral resolution offered by the limited scan range of an FTS while preserving sufficient phase information. Since phase is typically a slowly varying quantity in regions of high spectral intensity, it is sufficient to interpolate low resolution phase information obtained from the central burst region of the interferogram to the required resolution for point-by-point correction of the interferogram or spectrum [29]. The final variant is the one-sided interferogram which extends from ZPD to L_{max} maximizing the spectral resolution at the cost of no phase information.

To outline the Forman method of phase correction, consider Equation 3.117 multiplied by the inverse phase term

$$S_{obs}(\sigma) \times e^{-i\phi(\sigma)} = S(\sigma)e^{i\phi(\sigma)} \times e^{-i\phi(\sigma)} = S(\sigma) . \quad (3.119)$$

However, taking the forward Fourier transform of this expression gives

$$\begin{aligned} \mathcal{F}\{S_{obs}(\sigma) \times e^{-i\phi(\sigma)}\} &= \mathcal{F}\{S_{obs}(\sigma)\} * \left[\int_{-\infty}^{\infty} e^{-i\phi(\sigma)} e^{i2\pi\sigma z} d\sigma \right] \\ &= \mathcal{F}\{S(\sigma)\} = I_{sym}(z) . \end{aligned} \quad (3.120)$$

The first line in Equation 3.120 employs the convolution theorem (Equation 2.33). The term in square parenthesis is simply the Fourier transform of the phase factor, and is formally known as the Phase Correction Function (PCF),

$$PCF(z) = \int_{-\infty}^{\infty} e^{-i\phi(\sigma)} e^{i2\pi\sigma z} d\sigma , \quad (3.121)$$

which acts as a convolution kernel in the interferogram domain. The convolution in Equation 3.120 restores the symmetry of the interferogram, and as such, the inverse Fourier transform of the convolution of the asymmetric interferogram and the PCF returns the real valued physical spectrum,

$$\mathcal{F}^{-1}\{I_{asym}(z) * PCF(z)\} = S(\sigma) . \quad (3.122)$$

The top left panel of Figure 3.7 shows the asymmetric interferogram which results from a spectrum containing an induced second degree polynomial phase distortion. The bottom panel shows the inverse Fourier transform of the asymmetric interferogram, with the blue and red curves showing the real and imaginary components of the spectrum, respectively. The green phase spectrum is also shown, which was calculated using Equation 2.29. The top right panel contains the PCF obtained by computing Equation 3.121 with ϕ , given by the green curve in the center panel.

After convolving the PCF with the asymmetric interferogram, the symmetric interferogram is recovered as shown in the left panel of Figure 3.8. To show the phase has been corrected properly, the inverse Fourier transform has been recalculated and the resulting real

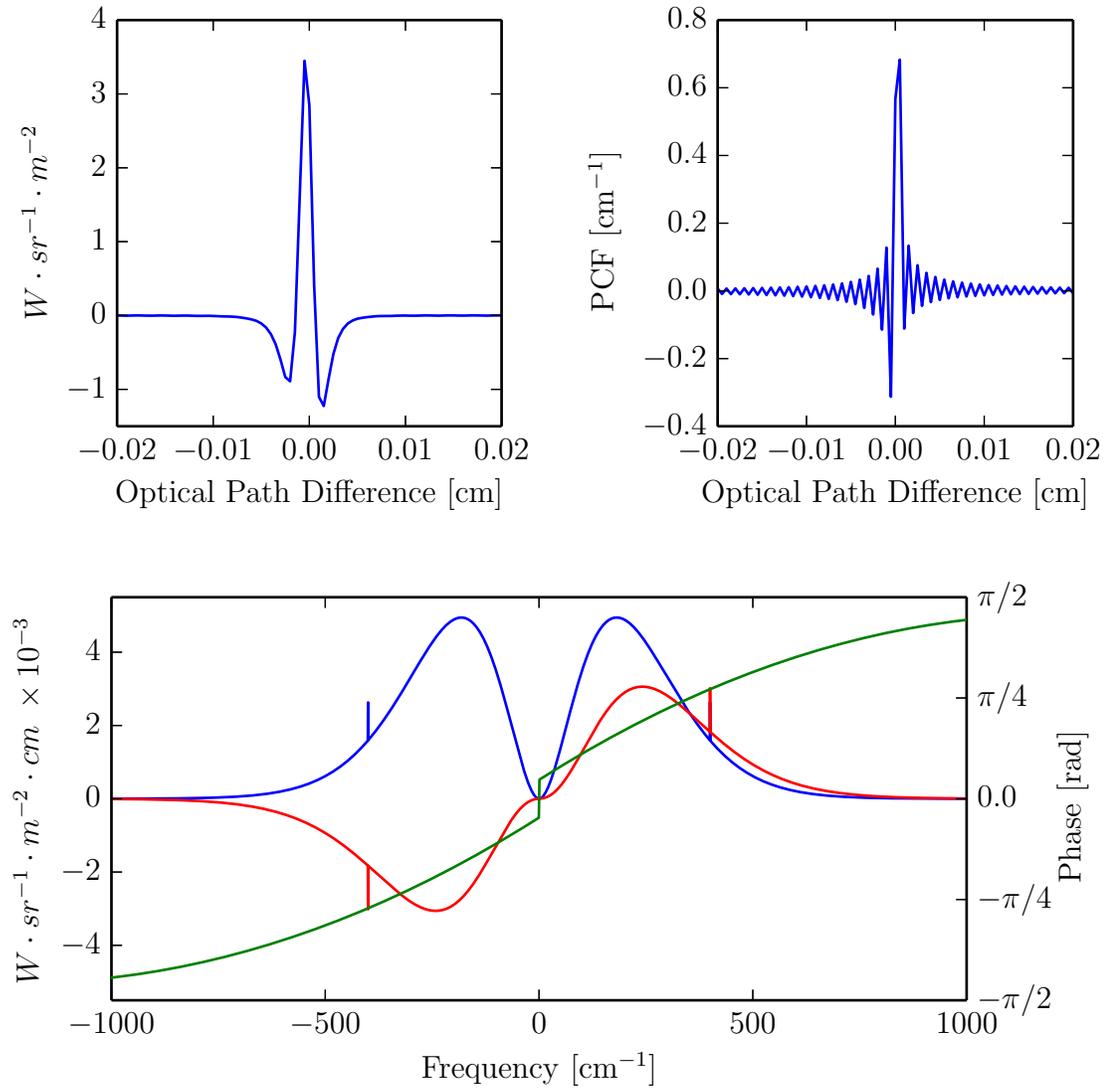


Figure 3.7: An asymmetric interferogram has been plotted in the top left panel, with the resulting Fourier transform shown in the bottom panel. The uncorrected spectrum in the bottom panel shows the real (symmetric) and imaginary (antisymmetric) components of the Fourier transform in blue and red, respectively. The resulting phase has also been calculated using Equation 2.29, is plotted in green, and exhibits a discontinuity at the origin. The phase correction function has been calculated using Equation 3.121 and is shown in the top right panel.

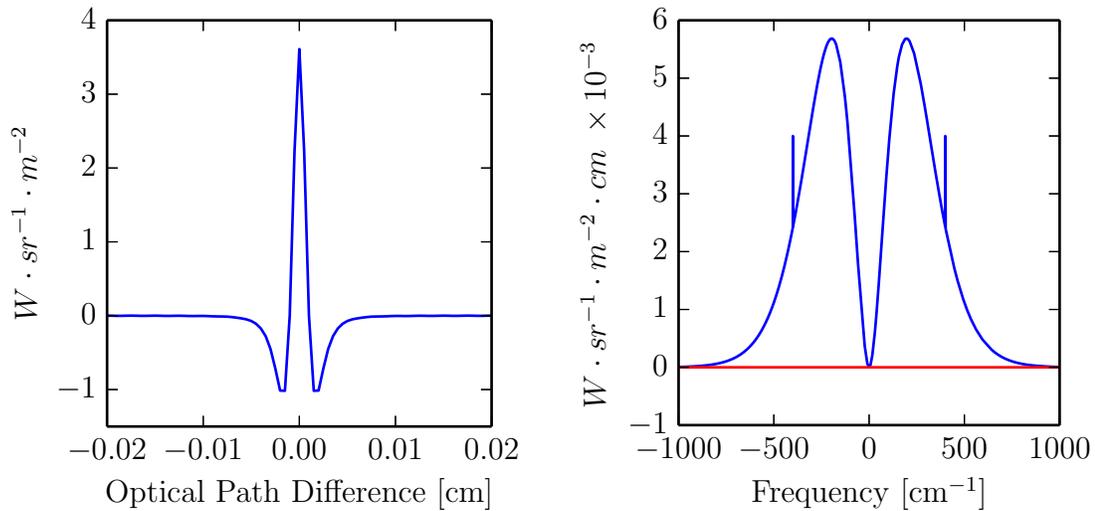


Figure 3.8: Phase corrected interferogram (left) and spectrum (right). After performing the Forman method of phase correction, the resulting interferogram has been symmetrized (left). The Fourier transform is plotted in the right panel, and it can be seen that the spectrum is completely real (blue), with a negligible imaginary component (red).

and imaginary components of the spectrum, blue and red respectively, have been plotted in the right panel. It can be seen that the imaginary component is zero to within numerical error. It should be cautioned that evaluating the phase of this corrected spectrum does not yield a uniform value of zero across the spectral band. In regions with low spectral amplitude, the phase is equal to the arctan of a small number divided by a small number, which can yield misleading values.

As a final comparison, the positive frequencies of the original, phase shifted, and corrected spectra are plotted in Figure 3.9. The original spectrum, plotted as green “X”s, consists of a black-body continuum at 100 K with an emission feature at 400 cm^{-1} . The phase shifted spectrum is shown in red, with the corrected spectrum shown in blue, which as expected, coincides with the original spectrum. It can be seen that the shifted spectrum produces a black-body curve with both a lower amplitude and shifted peak frequency. In addition, the emission feature is of lower amplitude, and depending on position of the emission feature and form of the phase shift, can also be altered in shape and position [29]. These

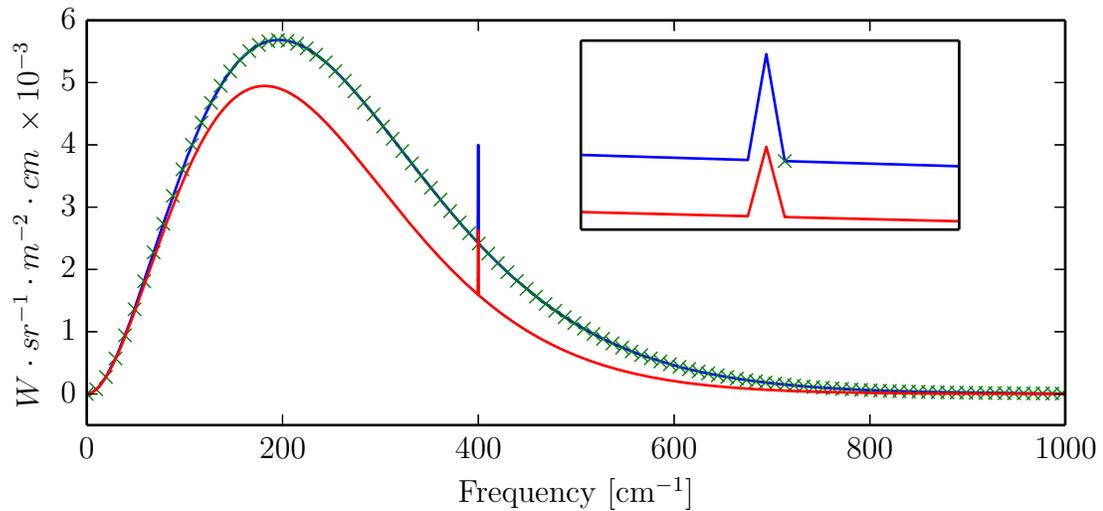


Figure 3.9: Comparison of uncorrected and phase corrected spectra. To compare the results of the Forman method of phase correction, the corrected (blue), and uncorrected (red) spectra have been plotted. The Original spectrum used before a phase shift was induced is plotted with green markers. The inset shows a close up of the spectral feature at 400 cm^{-1} . It can be seen that an uncorrected spectrum produces both a shift in peak frequency, and a change in amplitude of the black-body curve. The spectral feature has also been modified highlighting the importance of phase correction.

alterations demonstrate the importance of proper phase correction of data before drawing scientific conclusions.

3.5 Summary

Spectroscopy is an important field of study for astronomy, and much can be learned about a source's structure, energy, dynamics, and chemical composition by analyzing the source's spectrum. In this chapter I introduced the theory of black-body radiation, and derived the models by which rotational and electronic energy transitions are predicted. Furthermore, it was shown how line intensities can be calculated using Einstein coefficients, by considering the populations and transition probabilities between non-degenerate energy states. As radiation propagates from source to observer, it often interacts with matter during the journey. The attenuation effects of these interactions were considered, along with other physical processes which can modify the shape of spectral features. The information

presented in this chapter is sufficient for modelling and interpreting real spectra, and was oriented towards the kinds of spectral emissions often encountered in SPIRE spectra, which are the focus of this thesis.

Following the discussion on the radiative processes mentioned above was an introduction to Fourier transform spectroscopy. The theory was derived for an ideal dual input/output Michelson interferometer, along with how the interferogram is converted into a science ready spectrum. This process includes phase correction, and in many cases zero-padding of the interferogram to obtain the desired real valued spectrum and spectral gridding. These processing steps are employed in the reduction of SPIRE spectral data used in this thesis.

Chapter 4

SPIRE FTS Spectral Feature Finder

4.1 Introduction

As introduced in Section 1.2, the Herschel Space Observatory (*Herschel*; [20]), was the fourth cornerstone mission in the ESA's, Horizon 2000 programme [60]. Launched on May 14, 2009, with its broad frequency coverage and position above Earth's atmosphere, *Herschel* provided an unobstructed view of the Far-Infrared (FIR) universe in orbit around the Sun-Earth system second Lagrangian point (L2). Approximately 37,000 successful science observations were obtained during the mission which terminated on April 29, 2013 [22]. *Herschel* was designed to build on previous infrared observatories, and extend spectral coverage into the relatively unexplored the FIR and submillimeter bands ($\sim 10\ \mu\text{m}$ – $1.0\ \text{mm}$). As such, *Herschel* contained three science instruments with overlapping operational bands between ~ 55 – $672\ \mu\text{m}$ [61]. Observing these longer wavelengths poses certain technological challenges in terms of spatial resolution and the thermal profile of the optical systems.

An optical systems ability to distinguish between two points in space is called the limit of spatial resolution. Typically, this limit is defined in terms of Rayleigh's criterion [33] whereby two points on an object plane are considered resolved if the center of one point's diffraction pattern coincides with the first minima of the other's. Using this criteria, the limit of resolution for an optical system with an entrance aperture of diameter, D , is given by [33]

$$\Delta l_{min} = 1.22 \frac{\lambda}{D}, \quad (4.1)$$

where λ is the wavelength of the light under study. Importantly, this relationship shows

that resolution is degraded with increasing wavelength, and the only instrumental solution is to increase the diameter of the aperture. Compared to the optical band, FIR wavelengths are 100 – 10,000 times longer, and require an equally larger aperture to obtain comparable spatial resolution. For this reason, *Herschel* was outfitted with a 3.5 m diameter primary mirror, roughly corresponding to the aperture of the telescope. This constitutes the largest single dish telescope ever put into orbit, and although lacking the resolution of typical optical telescopes, is still a significant improvement over previous IR/FIR telescopes (e.g., Spitzer 0.85 m [16]; IRAS 0.57 m [11]; ISO 0.6 m [13]). As such, *Herschel* offers the highest spatial resolution of any FIR instrument to date.

The wavelength bands covered by the *Herschel* instruments correspond to the peak radiation of a black body ranging in temperature between 5 and 50 K. Since all solids give off thermal emission with spectral profiles approximated by black body radiation, this imposes significant thermal limitations on the optics and detectors of a telescope. For this reason, *Herschel* was equipped with a heat shield, offering protection from sources of thermal radiation like the Sun, passively cooling the telescope to ~ 85 K. Furthermore, a closed cycle $^3\text{He}/^4\text{He}$ cryogenic system actively cooled the focal plane optics and bolometer detectors to ~ 300 mK in order to obtain optimal sensitivity. The cryogen supply was in fact the limiting factor of the mission, which terminated after the controlled thermal environment of the payload module could no longer be maintained.

The three science instruments aboard *Herschel* were: the Heterodyne Instrument for the Far Infrared (HIFI); [62], the Photodetector Array Camera and Spectrometer (PACS); [63], and the Spectral and Photometric Imaging Receiver SPIRE; [21]. HIFI utilizes the heterodyne technique where an input signal is multiplied by a local oscillator which produces a new modified signal equal to the difference between the local oscillator and the input frequency. This significantly lowers the input frequency and equally improves spectral resolution [64]. Thus HIFI can obtain a high spectral resolving power typically on the order of $R = \lambda/\Delta\lambda \approx 10^6$, at the operational bands 157-212 and 240-625 μm . PACS consists of two

bolometer arrays and a grating spectrometer. Simultaneous photometry can be obtained using the dual mode array sensitive to wavelengths 60-85 or 85-130 μm , and a second array with a 130-210 μm bandwidth. The grating or integrated field spectrometer operates with a moderate resolving power between $R = 1000$ -4000. Finally the SPIRE instrument consists of an Imaging Fourier Transform Spectrometer (IFTS) and 3 band photometer. The photometer comprise three bolometer arrays with central detection wavelengths of 250, 350, and 500 μm , while the IFTS consists of two bolometer arrays with bandwidths 194-324 and 316-671 μm .

The data obtained from the SPIRE IFTS is central to this thesis. This chapter will describe the instrument in more detail (Section 4.2), and introduce the data analysis tools and environment for SPIRE (Section 4.3). The SPIRE data processing pipeline and products will be outlined (Section 4.4), with the SPIRE FTS Spectral Feature Finder, a high level data mining tool developed for post-processing data products, discussed in Section 4.5.

4.2 SPIRE Imaging Fourier Transform Spectrometer

A model of the physical SPIRE IFTS, along with a simplified optical path layout, is shown in Figure 4.1. The Mach-Zehnder configuration is based on the classic interferometer of the same name which shares the critical feature of dual input and output ports [65]. This optical design has two key advantages. First, the dual input ports allow for absolute flux calibration via a differential measurement of a black body at two different temperatures located at one input port (Section 4.4). Secondly, translation of the double rooftop mirrors induces the optical path difference between the two beams. Every one unit of linear displacement of the rooftop mirrors results in four units of relative optical path difference, significantly improving the scan efficiency of the Mach-Zehnder design over the classical Michelson.

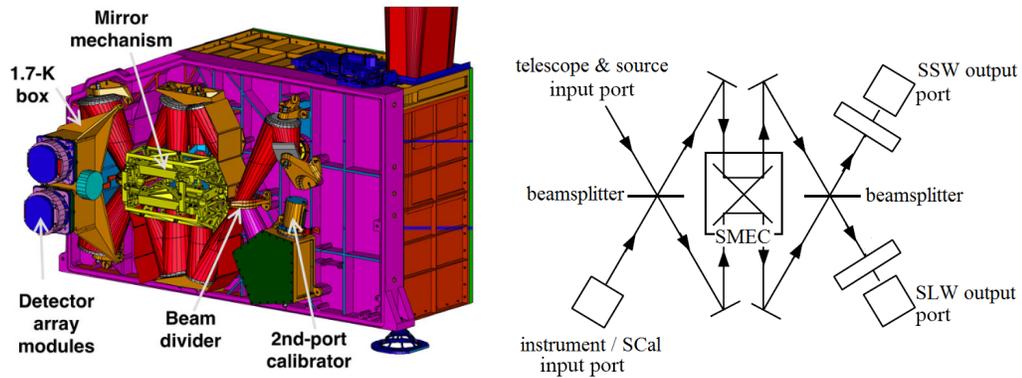


Figure 4.1: A three dimensional model of the SPIRE IFTS instrument (left), with a simplified schematic (right). Figures courtesy of SPIRE observers manual (left; [59]) and the HIPE documentation (right; [66]).

4.2.0.1 Bolometer Arrays

Both outputs of the IFTS terminate on an array of bolometers. The bandwidth of the instrument is controlled by a series of metal mesh filters, and a wave guide cut-off specific to each bolometer array which produces the final bandwidths of $194\text{-}324\ \mu\text{m}$ for the Spectrometer Short Wavelength array (SSW), and $316\text{-}671\ \mu\text{m}$ for the Spectrometer Long Wavelength array (SLW). The two bolometer arrays consist of pixels oriented in a hexagonal close packing scheme.

Figure 4.2 shows the detector layout for both bolometer arrays. Each detector is labelled with a letter next to a number, resulting in a row-column like notation. The central detector for SSW and SLW are D4 and C3, respectively, and are located along the optical axis of the IFTS. Grey detectors in the SSW array were not functional for the duration of *Herschel's* mission. Detectors shaded in blue have co-aligned spatial projections on-sky. An overlapping projection of the SLW and SSW array footprint is shown on the right of Figure 4.2, with red and blue circles corresponding to SSW and SLW components, respectively. The diameter of each circle corresponds to the nominal beam FWHM for the corresponding detector. The nominal beam FWHM for SSW and SLW detectors is $19''$ and $35''$, respectively. The red dashed and bold circles indicates the $\sim 2'$ unvignetted field of view of the SPIRE FTS [59].

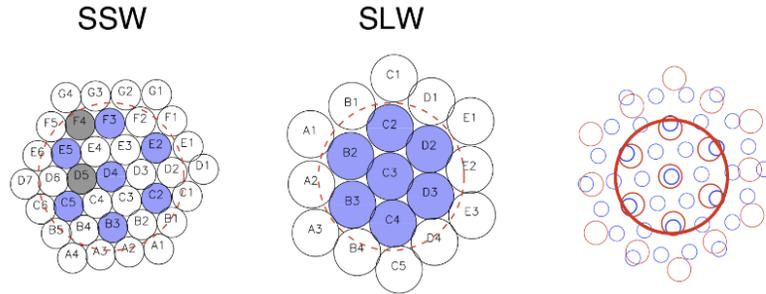


Figure 4.2: The orientation of the SSW and SLW SPIRE bolometer arrays. In circle annotations indicate the conventional detector labelling scheme. Blue shaded detectors between arrays are co-aligned on-sky, while the grey shaded detectors of the SSW array are not operational. On the left is a footprint of the overlapping array footprint on-sky with small red and blue circles indicating SSW and SLW detector positions, respectively. The red dashed circles in the first two images, and the bold red circle in the final image indicate the instrument's unvignetted field of view. Image courtesy of the SPIRE observers manual [59].

4.2.0.2 Beam Profile

Before reaching the bolometer arrays, radiation passes through feed-horns consisting of a conical entrance which transitions to a cylindrical wave-guide [67]. The feed-horns are manufactured to dimensions which allow multimodal coupling for the bandwidth of interest and in effect focuses incoming radiation onto the bolometer detector. The feed-horns have different coupling efficiencies with the various permitted modes, and as a result, varies with frequency. This property of the feed-horns, and the other effects of the complicated SPIRE optical design results in a frequency dependent beam profile. This dependence was determined experimentally by scanning a point source in a raster pattern, and fitting the resultant intensity maps with a superposition of Hermite-Gaussian functions [68]. The frequency dependence of the beam size is shown in Figure 4.3.

4.2.0.3 Spatial Coverage

SPIRE operates in different spatial sampling modes providing various levels of spatial coverage on-sky. Each observing mode is dictated by the operation of the SPIRE Beam Steering Mirror (BSM), which is positioned in the SPIRE optical system before the incoming radiation is directed to either the IFTS or photometer modules. Various levels of spatial

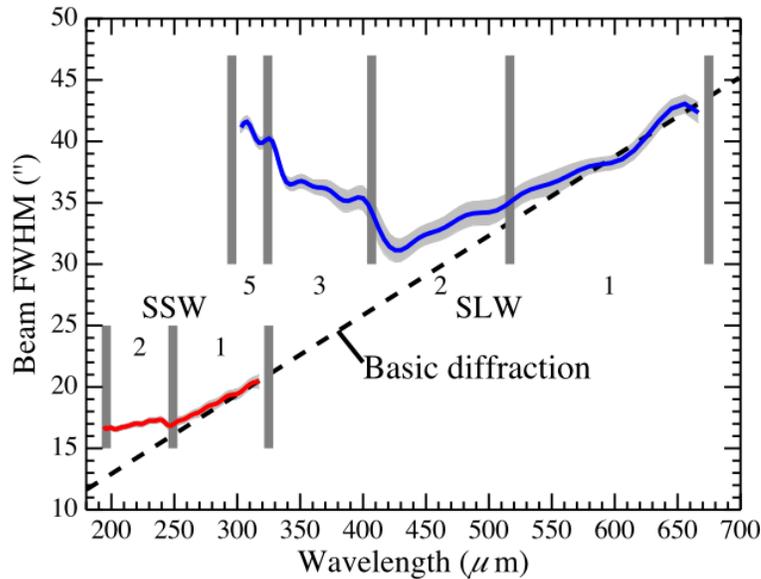


Figure 4.3: The frequency dependence of the SPIRE beam. The diffraction limited case is represented by the dashed line, with the grey shading around the red SSW, and blue SLW beam profile indicating the 3σ error on the fitted result. Vertical grey bars partition the different modes coupled in the feed-horns. Image courtesy of Makiwa et al., (2013) [68].

coverage are obtained by adjusting the orientation of the BSM which affects the portion of the optical axis without adjusting the orientation of *Herschel*. The motion of the BSM is formally called “jiggle”.

The sparse observing mode employs a single position of the BSM, and results in a single array footprint on-sky as shown on the left of Figure 4.4. The spacing between the center of the detector beams for a given array is two times the beam FWHM. Intermediate sampling is achieved using 4 jiggle positions. The resulting footprint on-sky is shown in the center of Figure 4.4. In this mode the spacing between beams is one beam FWHM. Finally, fully sampled images consist of 16 jiggle positions providing full Nyquist spatial coverage with each beam spaced by $1/2$ beam FWHM. The beam footprint for full image sampling is shown in the right of Figure 4.4.

In addition to spatial sampling modes which correspond to BSM motion, SPIRE also operated in different telescope pointing modes. In single pointing mode, *Herschel* maintained a stationary orientation, and could employ any of the spatial coverage modes. To

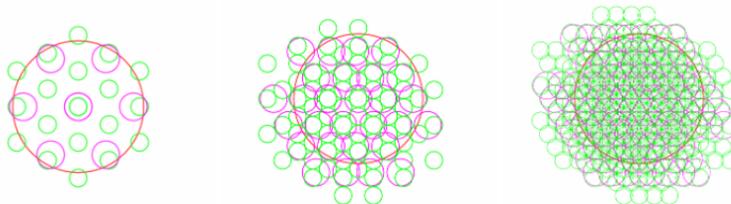


Figure 4.4: The on-sky beam footprint for sparse (left), intermediate (center), and full (right) spatial sampling. The red circle shows the unvignetted field of view of the initial jiggle position, and provides a spatial reference for the three sampling modes. Image courtesy of the SPIRE observers manual [59].

create images extending beyond the SPIRE instrument field of view, *Herschel* had to be reoriented. This was accomplished using a raster scanning pattern with pointing positions separated by 116'' horizontally, and 110'' vertically, with respect to the telescope orientation. This raster pointing mode was used to generate large maps and could employ any of the spatial sampling methods.

4.2.0.4 Spectral Resolution

Though the actual spectral resolution for unresolved spectral features is limited by the sinc instrument line shape of the IFTS, which results from the finite travel of the translation stage, the resolution modes of SPIRE are reported in terms of the frequency sampling. As outlined in Chapter 2, for an FTS, the frequency resolution in the spectral domain is dictated by the maximum OPD given by, $\Delta\sigma = 1/(2L_{max})$. SPIRE was configured to operate in 4 different spectral resolution mode.

In High Resolution (HR) mode, a single-sided interferogram was obtained extending an OPD from -0.555 to 12.645 cm about the ZPD. The symmetric sampling about the ZPD ([-0.555, 0.555] cm), can be used for phase correction as discussed in Chapter 3. The maximum OPD results in a frequency sampling of 1.184 GHz, with a corresponding FWHM of the instrument line shape of 1.429 GHz. This frequency spacing results in a spectral resolving power of about 1,000 making SPIRE a moderate spectral resolution imaging spectrometer. This mode is useful for spectral surveys since SPIRE provides broad spectral coverage

with a relatively short scan time of 66.6 s per iteration.

In Low Resolution (LR) mode, a double-sided interferogram is obtained, spanning the OPD range -0.555 to 0.560 cm about the ZPD. The maximum OPD gives this mode a frequency sampling of 24.98 GHz, with an instrument line shape FWHM of 30.14 GHz. This sampling results in a very low resolving power of about 48. The utility of this mode comes from the short scan time of 6.4 s, and its application towards obtaining the black body continuum of faint or featureless sources. Since HR mode overlaps with this scan range, the same information can be obtained from truncated HR observations. However, LR scans can be combined with HR scans in order to improve the continuum signal. This advantage is employed by the HR+LR operation modes [69] which record a variable number of HR and LR scans to obtain a desired continuum SNR.

The final two modes are medium resolution and calibration resolution. Neither resolution mode was used for science observations and thus do not warrant further explanation here.

4.3 The Herschel Interactive Processing Environment

The native environment for *Herschel* data analysis is the Herschel Interactive Processing Environment (HIPE); [70]. This science tool began its development in 2002 under the original title of the Herschel Data Processing System. It was initially used to test *Herschel* instruments in the early stages of development and was utilized in pre-flight instrument validation and characterization testing. From the outset, the goal of this software was to provide astronomers with a robust, user friendly, and fully integrated system complete with extensive documentation and built in analysis tools, which is freely available and distributable.

HIPE possesses the novel characteristic of being a complete astronomy tool facilitating the user's work-flow from beginning to end (i.e., from raw data to publication ready figures). HIPE is directly linked to the Herschel Science Archive (HSA); [22], so data acquisition can take place directly within the user interface. The user has access to data at various

levels of processing, and is provided with general data reduction scripts for all observation modes which can be modified to suit specific science goals. HIPE also offers a diverse set of graphical user interfaces, which allow real time interaction and manipulation of *Herschel* products, in addition to a wide range of plotting features for more specific data analysis and presentation.

HIPE is coded in the Jython computer language [70]. This is essentially a Python based language configured to run within Java platforms. Python is open source, and Java is implemented in a wide range of operating systems. As a result, HIPE is a free, easily accessible, cross-platform scientific analysis tool available to a broad user base. In addition to its accessibility, Jython is also a high level, object oriented programming language. This means that the data structures being used already have a number of useful predefined internal operations, and external functions which are immediately accessible to the user, without the need to write them from scratch. Furthermore, Jython is a dynamic language, which means that variables, functions, and data structures can be manipulated live from within a terminal session. This is compared to scripted languages with data structures which can only be accessed within the script itself. Dynamic languages are very useful for debugging, and testing small changes without requiring the time consuming process of recompiling and executing the whole script to test modifications.

HIPE is currently on its fifteenth full release version, though it is possible to access earlier versions and builds if required. The development of HIPE is the result of a large international effort, and has been subject of rigorous testing from hundreds of community members including myself and members of the Canadian Data Processing and Science Analysis Software (DAPSAS) centre. HIPE 15.0.0 is the programming environment used for the vast majority of what follows in this thesis.

4.4 SPIRE Pipeline and Products

The HSA is a repository for data obtained from the three *Herschel* instruments, SPIRE, PACS, and HIFI. Within the HSA, the highest level data structure of SPIRE data products are observations, which are discriminated based on a 10 digit observation identification number (obsid). Each observation contains an extensive hierarchy of datasets including: auxiliary data consisting of supplementary housekeeping information about the telescope including pointing, and temperature readouts, calibration data containing numerous parameters within the calibration tree used when processing the specific observations, and science data. Science FTS data is offered in various levels of processing [66],

- level 0 contains the raw unprocessed timeline data for each observation.
- level 0.5 expresses the level 0 data in a form that is appropriate for inspection.
- level 1 data is obtained after significant processing resulting in an interferogram with physical units. Some observations may contain extended source corrected spectra for each detector and jiggle position.
- level 2 contains the fully processed, science ready data. Results obtained using both extended and point source calibration schemes can be found here, along with apodized variations.

The major processing steps from level 0.5 to level 2 are illustrated in Figure 4.5.

A detailed outline of the SPIRE standard processing pipeline is given in [58], but one point is worth mentioning here. Comparing spectra of the same source with different spectral resolutions is made easier when the spectra share the same frequency gridding. For this reason each interferogram is zero-padded such that the frequency spacing of the corresponding spectrum has the desired spacing. For high spectral resolution, the frequency spacing is 0.3 GHz, which should be compared with the actual spectral resolution of 1.2 GHz. As such, each SPIRE spectra obtained from any detector in either bolometer array is presented in the data products with the same spectral gridding through this artificial interpolation.



Figure 4.5: Key processing steps in the SPIRE FTS data reduction pipeline [66]. Data reduction begins with level 0.5 data consisting of raw timeline ordered instrument readouts. The steps between level 0.5 and 1 data products are focused on correcting for spurious data points, instrumental effects, and flux calibration appropriate for extended sources. Level 1 products are then further processed into level 2 products by applying a flux conversion appropriate for point sources. In the case of mapping observations, the level 1 products are projected onto a uniform, rectilinear spatial grid.

4.4.0.5 Spatial Projection

The SPIRE bolometers are oriented in a hexagonal closed packed configuration (Figure 4.2), yet the spectral maps in HIPE are presented on a uniform rectilinear grid. Accomplishing this translation requires some sort of projection algorithm. The two main projection algorithms employed for SPIRE observations are naive, and convolved, spatial projection. The resulting spectral maps are formally known as hyper-spectral cubes.

Naive spatial projection initiates by defining an on-sky Cartesian like grid with equally spaced rows and columns. This grid extends over the sky covered by the total mapping observation, which itself may consist of numerous telescope pointings and/or jiggle positions. The routine then identifies all spectra residing within a given element of the grid which is then populated by the average of all of the identified spectra. SSW and SLW grids are kept separate due to the difference in nominal beam size. This straight forward averaging projection method is reasonable, though tends to leave coverage gaps in the SSW grids due to the dead detectors in the bolometer array.

The convolution method begins by defining a grid in the same way the naive projection task does. In this case, however, each element is populated by convolving a kernel across a spectral map. Typically the kernel is a Gaussian with an FWHM equal to the FWHM of the SLW or SSW beam for the SLW and SSW maps, respectively. The spectral map can be either the raw raster scan map, or a naive projection map. Conceptually, the process can be regarded as a naive projection where a weighted average is used instead of a uniform average. One can imagine a Gaussian with unit amplitude centered on the grid element being populated. All spectra in the input map are weighted based on the amplitude of the Gaussian they would coincide with based on their displacement from the Gaussian peak. A key advantage of this convolution projection is that it fills any gaps present in the associated naive projections.

4.4.1 Calibration Schemes

For Standard Product Generation (SPG), two different calibration schemes are employed. Extended calibration assumes the source being observed fills the entire SPIRE beam and can be approximated as a uniform intensity distribution over the instrument field of view. The second calibration scheme assumes the source being observed is point-like compared to the SPIRE beam and thus requires a different approach. Given the frequency dependent beam sizes (Figure 4.3), there exist cases which are point sources from the perspective of the SLW beam, but are extended with respect to the SSW beam. These sources are semi-extended and require additional processing outside of SPG.

The different calibration schemes affect the overall flux density of the spectrum, and most notably the continuity between the SLW and SSW bands. Figure 4.6 shows the effect of the SPG processing schemes on three different sources of varying spatial extent. The top row consists of spectra processed as extended sources, with the bottom representing point source calibration. The source on the left is point like with respect to both SLW and SSW beams, while the source on the right is extended with respect to both beams. The appropriate calibration scheme can be determined by looking at the overlapping region between both bands with continuity being the desired characteristic. The source in the center however, is semi-extended and neither calibration is sufficient in preserving the continuity between the two bands [71].

Accurate parameterization of the continuum is critical when attempting to model the spectrum as a superposition of a polynomial continuum and sinc line shapes. An inaccurate continuum can easily lead to spurious line fits in local regions where the polynomial continuum residual is high. Due to their effect on the continuum, it is worth briefly introducing the SPG calibration schemes. A detailed description of the processing steps required to convert raw timeline bolometer voltages into a spectrum with units of VGHz^{-1} is given in [58]. What follows focuses on the absolute flux conversion to the standard spectroscopic units of $[\text{Wm}^{-2}\text{Hz}^{-1}\text{sr}^{-1}]$ and Jy for extended and point sources, respectively.

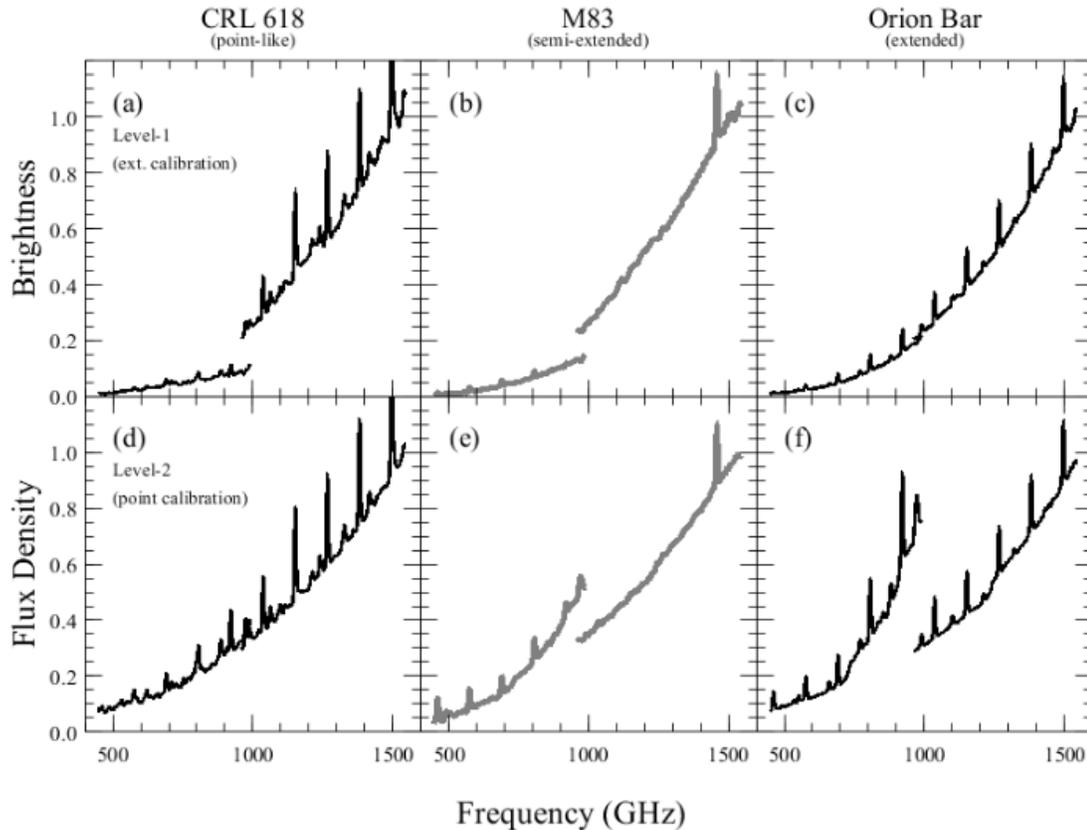


Figure 4.6: The effects of calibration scheme on sources with different spatial extent. Sources with different spatial extent are presented in each column, with extended and point-source calibration schemes employed when generating the figures in the top and bottom rows, respectively. Point-source calibration successfully preserves the continuity of the spectral continuum for sources like CRL 618, which appear point-like in both SPIRE FTS beams. Sources which fill both beams, like the Orion Bar, require extended calibration to preserve continuity between the SLW and SSW bands. Neither calibration scheme is sufficient for sources with intermediate extent which are point-like in the SLW beam, but extended in the SSW beam, like M83. Figure courtesy of Wu et al., (2013) [71]

4.4.1.1 Extended Calibration in SPIRE

The observed low level spectrum can be represented as

$$V_{obs} = R_{ext}(f)I_{ext}(f) + R_{tel}(f)M_{tel}(T_{tel}, f) + R_{inst}(f)M_{inst}(T_{inst}, f) \quad [VGH_z^{-1}], \quad (4.2)$$

where $R_i(f)$ is the Relative Spectral Response Function (RSRF): $[VGH_z^{-1}Wm^{-2}sr^{-1}Hz^{-1}]$, of the instrument with respect to the source component $R_{ext}(f)$, telescope emission $R_{tel}(f)$, and instrument emission $R_{inst}(f)$. The source intensity $I_{ext}(f)$ $[Wm^{-2}sr^{-1}Hz^{-1}]$ corresponds to the spectrum of the source and is the quantity of interest. The telescope and instrument both radiate as grey bodies, and these emissions contribute to the overall detected signal. The emission profile of the telescope and instrument are given by $M_{tel}(T_{tel}, f)$, and $M_{inst}(T_{inst}, f)$, respectively, with units $[Wm^{-2}sr^{-1}Hz^{-1}]$.

The telescope emission intensity is assumed to be uniform and fills the entire SPIRE beam. If the source is assumed to be fully extended, then it also fills the SPIRE beam, and is assumed to be uniform. Under these conditions $R_{ext}(f) = R_{tel}(f)$. Instrumental emission enter the system primarily through the second input port of the Mach-Zhender design, and corresponds to the thermal radiation emitted from the Spectrometer Calibration Source (SCAL); Figure 4.1. With the source/telescope radiation entering through one port and the instrument radiation entering through another, it is assumed that the instrument requires a different RSRF. With these assumptions, the source intensity is given by

$$I_{ext} = \frac{(V_{obs} - M_{inst}R_{inst})}{R_{tel}} - M_{tel} \quad [Wm^{-2}Hz^{-1}sr^{-1}], \quad (4.3)$$

where the dependencies have been removed for convenience.

The telescope emission results from the thermal radiation of both the primary and secondary telescope mirrors. The spectral profile of the telescope is thus modelled as the sum

of two modified black bodies corresponding to each mirror

$$M_{tel}(T, f) = (1 - \epsilon_{M2})\epsilon_{M1}E_{corr}B_f(T_{M1}) + \epsilon_{M2}B_f(T_{M2}) . \quad (4.4)$$

The Planck function (Equation 3.11), is indicated by $B_f(T_i)$, with emissivity given by ϵ_i . The subscript $M1$ and $M2$ represent contributions from the primary and secondary mirror of the telescope, respectively. A small empirical correction term E_{corr} , is included to account for the slowly varying temperature of the mirrors over the course of the mission [72].

The primary and secondary mirror are manufactured from the same material and are thus expected to have equal emissivities. The emissivity of the mirrors was determine by ground based experiments on a dusty mirror sample. It was determined that the emissivity is approximated by [73]

$$\epsilon_{M1} = \epsilon_{M2} = 6.1366 \times 10^{-5} \sqrt{f} + 9.1063 \times 10^{-7} f . \quad (4.5)$$

The instrument is modelled more simply as a perfect black body. The spectral profile is determined by the Planck function with a temperature is given by the average temperature of the SCal over the duration of the scan.

The final task is to determine the telescope and instrument RSRF, which can conveniently be accomplished simultaneously. To begin, the telescope is pointed at a region of sky with no significant sources of radiation. This provides a dark field with no features above the detection limit of 40 mJy, thus the bolometer readout consists of telescope and instrument components only. Heat is generated by the Spectrometer Mechanism (SMEC) as the stage moves, which is partially dissipated by the SCal. That is, the instrumental contribution to the recorded signal varies between simultaneous scans of the SMEC while the telescope contribution remains relatively constant during a given observation. By taking differential measurements of sequential scans, and possessing models for the telescope and instrument spectra, it is possible to decouple the two RSRFs. The final RSRFs are pro-

vided in the observation's calibration data, and are the product of averaging all dark field calculations over *Herschel's* mission. Since forward and reverse scans of the SMEC have marginally different temperature profiles due to the position dependent power required to drive the SMEC, RSRFs are separately derived for each scanning direction.

With the relevant models and RSRFs determined, the extended source calibrated intensity can be obtained from Equation 4.3.

4.4.1.2 Point Source Calibration

Point source calibration builds off the extended source intensity through the introduction of a frequency dependent conversion factor

$$I_{pnt}(f) = I_{ext}(f)C(f). \quad (4.6)$$

The conversion factor $C(f)$ was determined empirically using SPIRE observations of the planet Uranus.

Uranus has the advantage of being point like to both SLW and SSW beams, and whose spectrum has been extensively studied. A sophisticated ESA model spectrum [74], was used as the expected spectrum of Uranus, which was used as a starting point for what SPIRE would detect. The model had to be modified to account for the spatial extent of Uranus from *Herschel's* location, which provides a solid angle normalization factor. A beam correction factor was also used to describe the coupling of Uranus' finite size with SPIRE's Gaussian like beam.

With an appropriate model of Uranus, $M_{Uranus}(f)$, the point source correction factor is obtained as the ratio of the model spectrum and the observed intensity, $I_{Uranus}(f)$,

$$C(f) = \frac{M_{Uranus}(f)}{I_{Uranus}(f)}. \quad (4.7)$$

Here the observed intensity is the extended source calibration intensity of Uranus observa-

tions. In some cases, dark field observations were subtracted from the observed intensity in order to reduce statistically independent sources of noise.

4.4.1.3 Highly Processed Data Products

The SPG outlined above produces high quality, science-ready data in most cases. However, exceptions exist which require special attention. Semi-extended sources which appear point-like to the SLW beam but are extended over the SSW beam produce discontinuities between the SPIRE bands in the overlapping region using either calibration scheme (Figure 4.6). Correcting for the semi-extended nature of sources is similar to the point source correction in that a frequency dependent correction factor is multiplied into the extended calibration intensity [71]. Determining this correction factor involves assuming a radially symmetric source intensity distribution, then varying the diameter until the absolute difference between the overlapping SPIRE bands is minimized.

Since the SPIRE beam is approximated as a Gaussian, a source centered off the optical axis will result in a detected intensity lower than would be expected. As such an appropriate normalization factor is required to recover the true intensity. Additionally, sources embedded in a high emission background will have a higher than expected intensity and potentially spurious spectral features. In this case background subtraction is required to obtain science-ready spectra.

These corrections require personal intervention and are not available as part of the default HIPE pipeline. Highly Processed Data Products (HPDP) of this kind, in general, must be created by the user, though a publicly available repository exists for calibration sources [75], which form a subset of SPIRE observations used to derive the processing steps outlined above.

4.5 SPIRE FTS Feature Finder

Analysis of FTS, and specifically SPIRE FTS, spectra poses unique challenges. The SPIRE instrument line shape is well approximated as a sinc function [23] as a result of the finite extent of the interferometer translation stage travel (Sections 3.2, 3.4.2). The sinc line shape disperses spectral energy over an extended range of frequencies as a result of its characteristic side-lobes. These side-lobes can interact with or even nullify other spectral features and, as a result, identifying spectral features can be challenging using manual inspection. With a relatively wide ~ 1.2 GHz spectral resolution, line blending may also be of concern. Furthermore, SPIRE covers a broad bandwidth ($\sim 450 - 1550$ GHz), and observed a variety of sources leading to numerous molecular and fine structure atomic emission lines being present in SPIRE spectra. The diverse variety of spectra and complicated line shape makes FTS data analysis both difficult and time consuming.

Over the course of *Herschel's* mission phase, SPIRE recorded 1850 FTS observations. Most of these consist of single pointing sparse observations, producing one spectrum for each of the 54 live detectors including both the SLW and SSW bolometer arrays. Roughly 16% of SPIRE observations were conducted in mapping mode which produce hyper-spectral cubes with two different projections. The number of pixels in each map can vary, both as a function of source and detector array, but a 15x15 grid for each detector array, averaged together, is typical. A rough calculation based on this information suggest the HSA stores $> 350,000$ science ready SPIRE spectra. Searching through this mountain of data for specific observations for a given research objective is impractical. As such, the SPIRE FTS team has developed the automated *Herschel* SPIRE FTS Feature Finder (FF) to assist in data mining and preliminary data analysis of SPIRE spectra.

A flowchart outlining the main steps of the FF routine is given in Figure 4.7 with a comprehensive explanation of the routine given in [76]. The remainder of this chapter will be dedicated to summarizing this routine, its products, and performance. Further discussion on specific subroutines are the subject of Chapters 5, and 6.

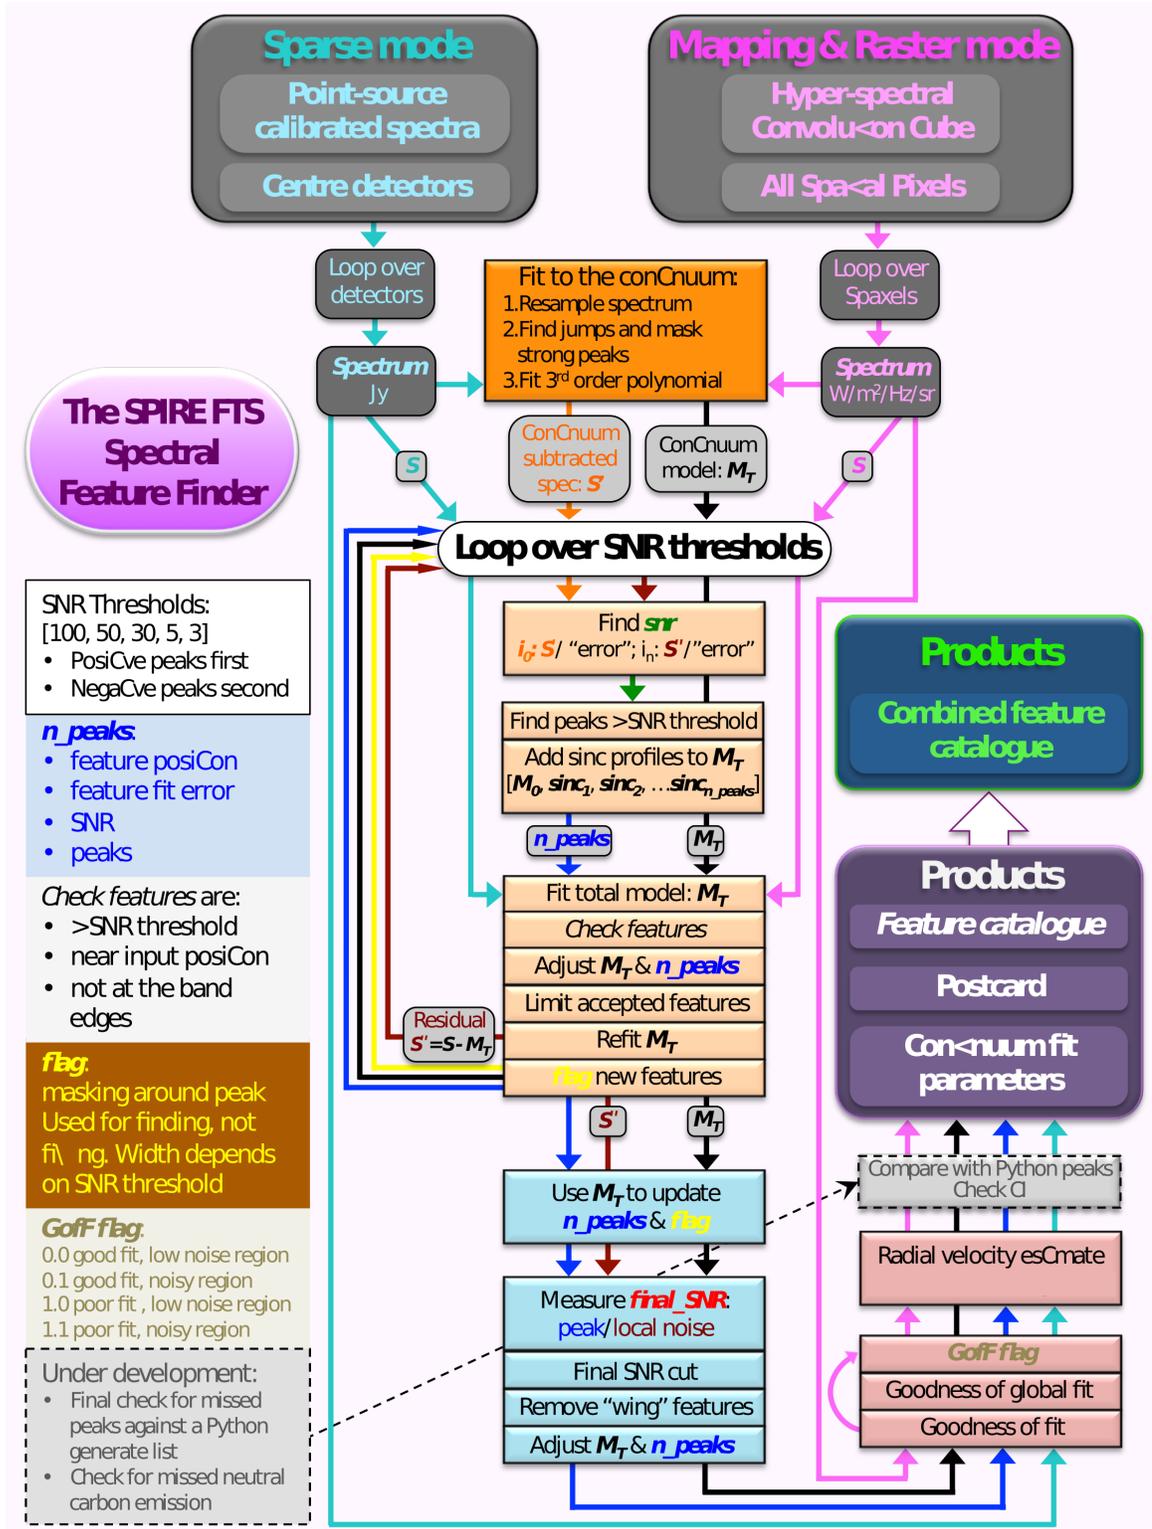


Figure 4.7: Caption on next page.

Figure 4.7: Flowchart outlining the main steps in the SPIRE FTS Feature Finder [76]. Outlined in this figure are the key steps within the automated SPIRE Feature Finder routine. The core of the routine is an iterative loop which seeks to detect and remove features meeting or exceeding the SNR threshold of the current iteration. Before a feature is accepted, additional checks are conducted to insure the quality of the candidate feature. A full description of the routine is presented by Hopwood et al., (2018, in-preparation). This thesis focuses on two subroutines, the radial velocity estimate, and the neutral carbon check, which take place just before product generation (purple block).

4.5.1 Feature Finder Routine

The FF attempts to parameterize the continuum, and detect significant spectral features within SPIRE FTS data. All observation modes are accommodated by the script, though the outputs vary depending on which mode is being analyzed (4.5.1.1). The routine operates in a similar way, regardless of observation mode.

For each of the SLW and SSW bands, the process initiates by calculating an approximate model of the continuum. Strong peaks are masked by searching for discontinuities within a differential measurement between the HR, and assumed to be feature free, LR spectra within the same observation. A third order polynomial is then fit to the HR spectrum which ignores any masked regions which have been identified as strong spectral features. The LR spectrum is subsequently fit using a second order polynomial, again separately for SLW and SSW spectra.

From this point the FF proceeds in an iterative fashion. Iterations are determined by a SNR threshold which decreases from 100 to 3 over six iterations ($\pm 100, \pm 50, \pm 30, \pm 10, 5, 3$). The first pass searches for emission features with a positive SNR, while absorption features with a negative SNR, are the focus of a second pass once all positive SNR iterations have been complete. For each iteration, a continuum subtracted SNR spectrum is generated by dividing the spectral flux array by the corresponding error array. Peaks greater than or equal to the current threshold are identified, and grouped over a range of 10 GHz. For each peak, a sinc model is generated with a central frequency and amplitude, corresponding to the SNR peak's frequency and continuum subtracted flux, respectively. The vast majority of SPIRE

spectral features are assumed to be spectrally unresolved. As such, the width of the sinc function is set to a constant value (i.e., the width of the instrument line shape). In HIPE, this width is defined as the distance between the first zero-crossings divided by 2π .

A global fit is then performed on the original unmodified HR spectrum consisting of the most recent continuum parameters and accepted spectral features, in addition to the newly identified spectral feature candidates from the current iteration. The results of each new feature are inspected in an automated way to limit spurious detections. These checks include: significant drift from the initial line position, incorrect amplitude sign, SNR below the iteration threshold, and features found within the noisy edges of both spectral bands. Additional checks are performed to reduce multiple fittings on partially resolved spectral features. Features that do not meet the above criteria are removed, with surviving features confined within 2 GHz of their current position in future iterations. A subsequent global fit is performed and the remaining lines are added to the current line list if they do not reside within a SNR threshold dependent masking region imposed on previously detected lines. This helps prevent spurious fitting to strong feature side-lobes. The current spectral model is then subtracted from the original HR spectrum, which is then used to generate the SNR spectrum for the next iteration. Once the process is complete for emission features, the process repeats for absorption features as previously indicated, with negative SNR thresholds.

For each feature, the amplitude of the line model parameterizes the “signal”. The residual is defined as the difference between the spectrum and total model, and is used to estimate the noise of each detected feature. The standard deviation of the residual within a 25 GHz window on either side of the line, excluding the central 5 GHz, parameterizes the features noise. The final SNR is given as the ratio of the amplitude and noise parameters. Additional checks are performed for each line, based on spectral proximity to other lines and their respective SNR ratio, which attempt to remove potential spurious side-lobe fitting which evaded previous checks.

After the SNR and other quality assessment parameters have been calculated for each

feature, the FF attempts to estimate a source radial velocity. Two methods exist, one based on searching for ^{12}CO rotational transitions by looking for a characteristic pattern between the frequencies of detected features, and one based on maximising the correlation between the spectral data and a model spectral template [77]. The first method, currently employed within the FF, will be discussed in detail in Chapter 5.

Following the source velocity estimate, the FF performs an additional search for the often elusive neutral carbon fine structure atomic emission feature. This is a linear regression based technique which builds off of the radial velocity estimate and employs careful empirical selection of initial parameter estimates. This subroutine is the subject of Chapter 6.

4.5.1.1 Feature Finder Products

The final product of the FF is the SPIRE Automated Feature Extraction Catalogue (SAFECAT²), a publicly available repository, offered as part of the HSA [22]. SAFECAT contains entries for each SPIRE obsid. For sparse observations, each entry consists of a file containing a list of frequencies for detected features, the corresponding SNR, and a two digit binary quality parameter indicating the reliability of the fitted feature. The file also includes meta data which gives information about the source, observation mode, calibration scheme, units, parameter keys, processing statistics and radial velocity with associated quality parameters. Also present is a file containing the fitted continuum parameters.

A companion dataset provided in complement to SAFECAT is the POSTcard Catalogue (POSTCAT) which graphically presents the result of the FF for each obsid. Figure 4.8 shows a sample postcard for NGC 7027, a planetary nebula used as one of the SPIRE spectral calibration sources. The postcard combines the spectral data from both SLW (red) and SSW (blue) bands, and indicates the noisy band edges ignored by the FF with vertical bars extending the height of the figure. The fitted continuum is shown with a green curve, and detected features are marked with vertical lines at the frequencies of the detected features. The length of the lines indicates the relative SNR estimate of the detected feature.

²<https://www.cosmos.esa.int/web/herschel/spire-spectral-feature-catalogue>

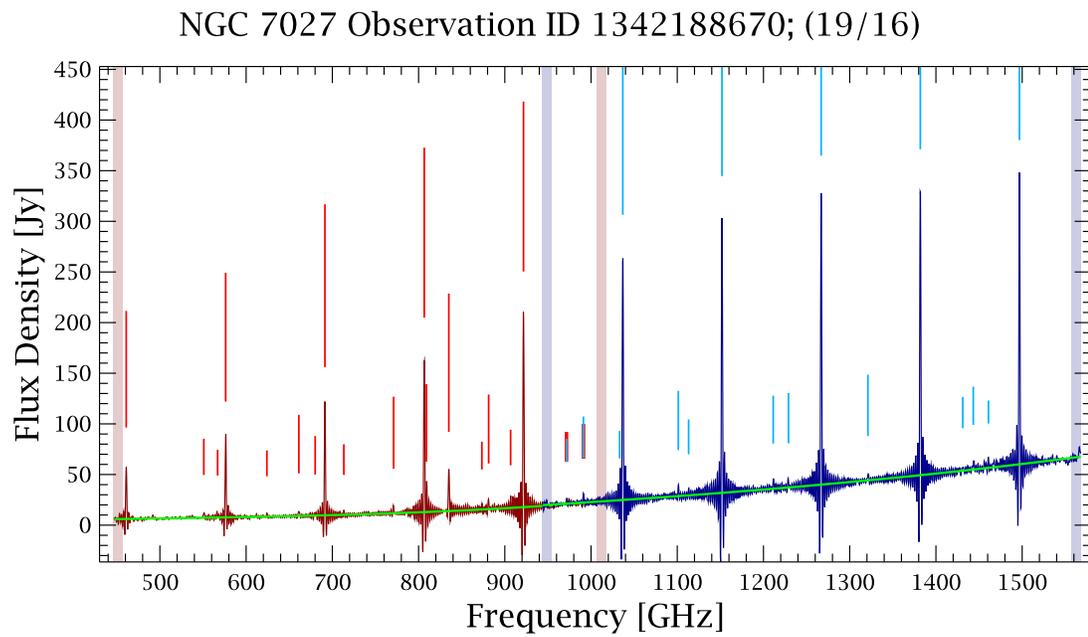


Figure 4.8: A sample postcard of NGC 7027. SLW (red), and SSW (blue) spectra are plotted with a green curve indicating the fitted continuum. Wide vertical bars extending the height of the figure indicate the band edges ignored by the FF. Vertical lines indicate the positions of detected features with their lengths representing their relative SNR.

Mapping observations produce similar products. Tables must, however, include entries from each pixel in the map being analysed. The format for mapping postcards includes a figure much like figure 4.8 for the brightest and faintest pixel, and for the pixel with the most detected features. Additional visuals such as velocity maps, and integrated intensity maps are also included. Examples of these visuals are provided in Chapter 5, with an additional selection presented in appendix B.

Datasets for low spectral resolution observations consist of merely the fitted continuum parameters, and generally, although not always, have a high resolution counterpart.

4.5.1.2 Feature Finder Performance

The performance of the FF was determined in two ways. First, the results of the FF operating on standard calibration sources was compared to the feature list derived from more rigorous analysis [78]. Second, and more importantly for this thesis, the FF was pointed towards a large body of synthetic spectra with known feature frequencies and amplitudes. The second method will be discussed further below.

The testing pool consists of three different sets of synthetic spectra, with each set containing 20,000 generated spectra. Each spectrum is built on top of a randomly selected dark sky observation to provide a realistic noise profile. In the first dataset, 4 random spectral features were added to the base spectrum. The SPIRE instrument line shape is well approximated by a sinc, though there is a slight asymmetry which manifests as a slightly lower negative side-lobe to the lower frequency side of the line shape [79]. For this reason, features were added using an empirical line shape which accounts for this irregularity. In the second dataset, 10 random spectral features were added to the base spectrum. In the final dataset, 4 random lines were added in addition to a rotational ^{12}CO emission template. This template was derived from co-added NGC 7027 observations. The template was scaled down in flux density to provide a more typical situation, since ^{12}CO features in NGC 7027 co-added data are brighter than in most SPIRE spectra. A random redshift was induced on

the template using a radial velocity derived from a random Gaussian distribution. Finally the template was randomly scaled from a uniform distribution between 0.1 and 1.0 before being added to the synthetic spectrum. The SNR for all randomly added lines was randomly selected from a uniform distribution with a maximum value of 100.

The completeness curves, which show what fraction of lines are detected with a given SNR, for all three datasets are shown in Figure 4.9. These curves all show a detection rate of $\geq 98\%$ for features with a SNR > 10 . Most importantly for this thesis are the results for the ^{12}CO dataset. This simulation indicates that significant ^{12}CO features can be reliably detected if present in SPIRE spectra.

4.6 Summary

During *Herschel's* mission phase, SPIRE recorded many observations in both single pointing and mapping modes, with either sparse, intermediate, or full spatial coverage. Raw data from these observations is processed in the Herschel Interactive Processing Environment which converts time-line recorded voltages from the SPIRE FTS into science-ready spectra. The processing pipeline calibrates the spectra first as an extended observation appropriate for sources with a spatial extent which fill both SPIRE SLW and SSW beams. This process is focused on removing telescope and instrumental contributions from the SPIRE FTS beam, and produces the level 1 dataset. Data is further modified using the point source calibration scheme, which multiplies level 1 spectra with a point-source correction factor derived from observations of Uranus. This calibration is appropriate for sources which are point-like in both SLW and SSW beams. Point source calibrated spectra comprise the level 2 dataset for SPIRE FTS observations.

With the large number of SPIRE spectra and the difficulties associated with FTS data analysis, including diversity of sources and the complicated nature of the line shape, a preliminary analysis and data mining tool was developed by the SPIRE FTS team. The SPIRE FTS Automated Feature Finder attempts to fit the source continuum and extract significant

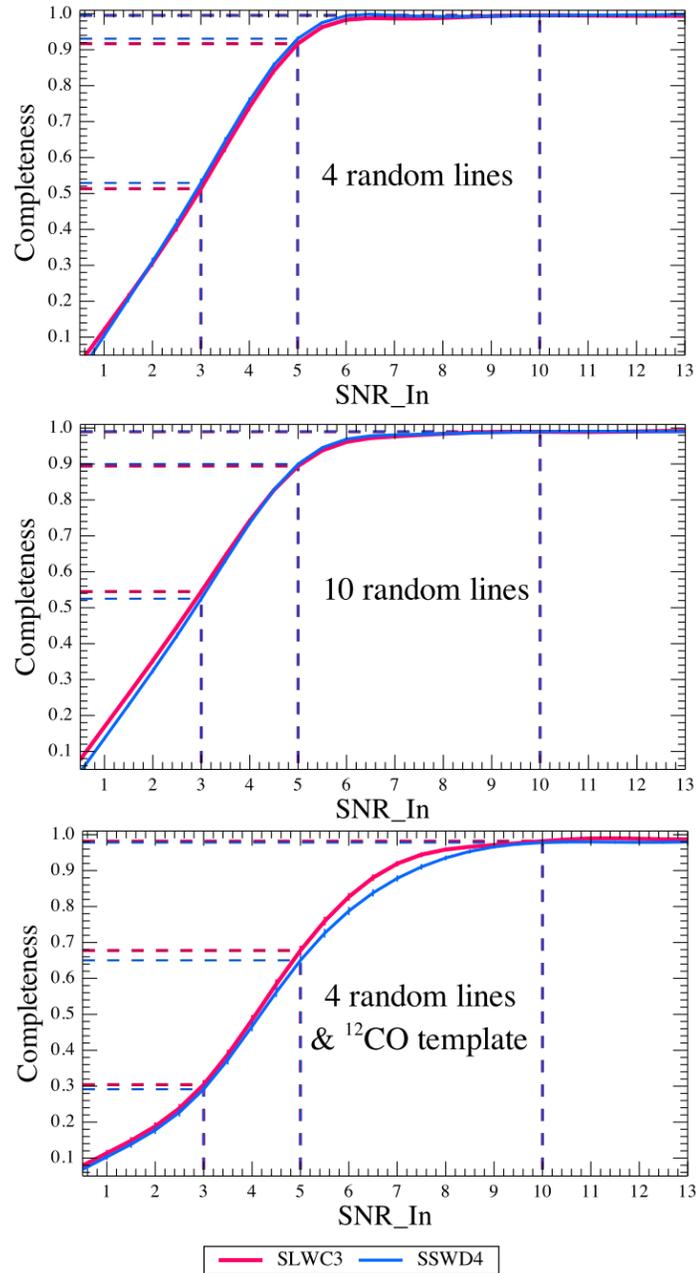


Figure 4.9: Completeness curves showing the detection reliability of the FF for synthetic spectra in terms of feature SNR. The three datasets presented are constructed with 4 random features (left), 10 random features (center), and 4 random features with an added ¹²CO template. Dashed lines indicate the intersection of the curves with simulated input SNRs of 3, 5, and 10 where both SLW (red) and SSW (blue) have been kept separate. Figure courtesy of Hopwood et al., (2018, in-preparation) [76].

features from SPIRE spectra. The features catalogue SAFECATE lists significant spectral features along with SNR and other quality parameters, with these products publicly available as part of the HSA. Each observation is accompanied by a complementary postcard offering quick visual inspection of the FF results, and a source radial velocity as another source parameter of interest.

The FF products provide a useful tool for astronomers, and I have provided significant contributions towards the development and testing of the FF routine. I have developed a module which is the primary method of estimating source velocity for the FF, and is the subject of Chapter 5. These results have been extended into a check for the often elusive neutral carbon fine structure emission line, which is the subject of Chapter 6.

Chapter 5

Redshift

5.1 Introduction

The FF returns a list of significant features in SPIRE spectra along with their corresponding SNR, intentionally omitting other parameters such as integrated flux to avoid complications with partially resolved spectral features, and to encourage follow up analysis by users. With this information in hand, it is convenient to derive a source radial velocity in a similar automated fashion for each of the analyzed spectra, since radial velocity has a number of useful applications in astronomy.

In 1929, Edwin Hubble observed a set of extra-galactic nebula, and established an empirical relationship between source radial velocity and distance from the observer [7]. This relationship, in confirmation of Georges Lemaître’s prediction two years earlier [80], states that the further a source is from the observer, the faster it will be travelling away from the observer [7]. This dependence is known as Hubble’s law, and the measurement of an extra-galactic source’s radial velocity is simultaneously a measure of its distance from Earth. Since light travels at a constant finite speed, sources under observation which are far away are actually being observed as they were in the past. Thus radial velocity is a useful search criteria when looking for sources from certain epochs in the history of the Universe. SPIRE recorded 440 extra-galactic FTS observations, and the HSA currently has no source velocity or redshift search criteria apart from that provided through the FF and the subroutine introduced in this chapter. Though the redshift routine outlined in this chapter can only produce accurate source velocity estimates for low redshift ($z \sim 0.05$) extra-galactic sources, it may

still be useful for identifying nearby galaxies.

A larger portion of the SPIRE FTS observations (661), are primarily targeted on galactic sources, and some of the most interesting FIR galactic sources are the sites of star formation. The role of molecular cloud dynamics in star formation is still uncertain, but it is well established that turbulence and other dynamic processes play a critical role in the early stages of star formation [81]. Models of star formation are generally constructed on idealized scenarios of gravitational collapse, yet observations exist which contradict predictions derived from these simple models, which in part motivates the significance of turbulence in star formation [81]. In addition, common star formation structures such as accretion disks, and bipolar outflows [82] express readily identifiable velocity profiles. As such, radial velocity has more immediate science implications than simply as a search criteria.

The utility of radial velocity is clear, though in terms of the historical development of the FF algorithm, a radial velocity estimating routine was developed in order to provide a dedicated search for the neutral carbon fine structure line, which is the subject of Chapter 6. In this chapter the theoretical foundation of the radial velocity estimating routine is described in Section 5.2, with a detailed outline of the routine itself in Section 5.3. This is followed by a validation of the routine using both real and simulated data in Section 5.4. Finally the results of the routine in the context of the FF and its data products are presented in Section 5.4.2.2.

5.2 Theoretical Foundation

5.2.0.3 Redshift and Its Effects On Spectra

When a wave emitted at a rest frame frequency f_0 is detected by an observer with an associated radial velocity v relative to the emitter, the observer will detect a wave with a different frequency f . This wave phenomenon is broadly referred to as the Doppler effect, named after Christian Doppler who gave the effect a theoretical foundation [8]. This phenomenon is experienced in every day life in the form of sirens passing people on the

street. However, this property applies to all waves including light. For sources which travel at velocities comparable to the speed of light, the original Doppler shift equation must be modified to account for time dilation. The relativistic Doppler shift equation is

$$\frac{f_0}{f} = \sqrt{\frac{1+v/c}{1-v/c}}, \quad (5.1)$$

where c is the speed of light, and v , f , and f_0 as defined above. This particular form of the equation is defined such that a positive velocity corresponds to the source moving away from the stationary observer. If $v/c \ll 1$, then Equation 5.1 can be simplified using the binomial theorem. The first two terms in the binomial expansion are given by

$$(1+x)^n \approx 1+nx, \quad x \ll 1. \quad (5.2)$$

Using this result to expand Equation 5.1, and ignoring emergent quadratic v/c terms, the non-relativistic Doppler equation is then,

$$\begin{aligned} f &= \frac{f_0}{1+v/c}, \\ v &= \left(\frac{f_0}{f} - 1 \right) c. \end{aligned} \quad (5.3)$$

In the case of a positive radial velocity, $f < f_0$ and spectral lines are shifted towards the lower frequency, “red”, side of the electromagnetic spectrum. As such, in the case of light this effect is called redshift, and is synonymous with radial velocity. For negative velocities where the detected frequency is shifted to higher frequency, the effect is called blueshift.

Figure 5.1 shows the co-added continuum subtracted spectrum of NGC 7027, a young planetary nebula used as one on the standard SPIRE calibration sources. In the figure a series of redshifts has been induced to show the effects of source velocity on the position of spectral features. It is worth noting that features at a higher frequency are shifted more than their lower frequency counterparts. The prominent features in the spectrum are the

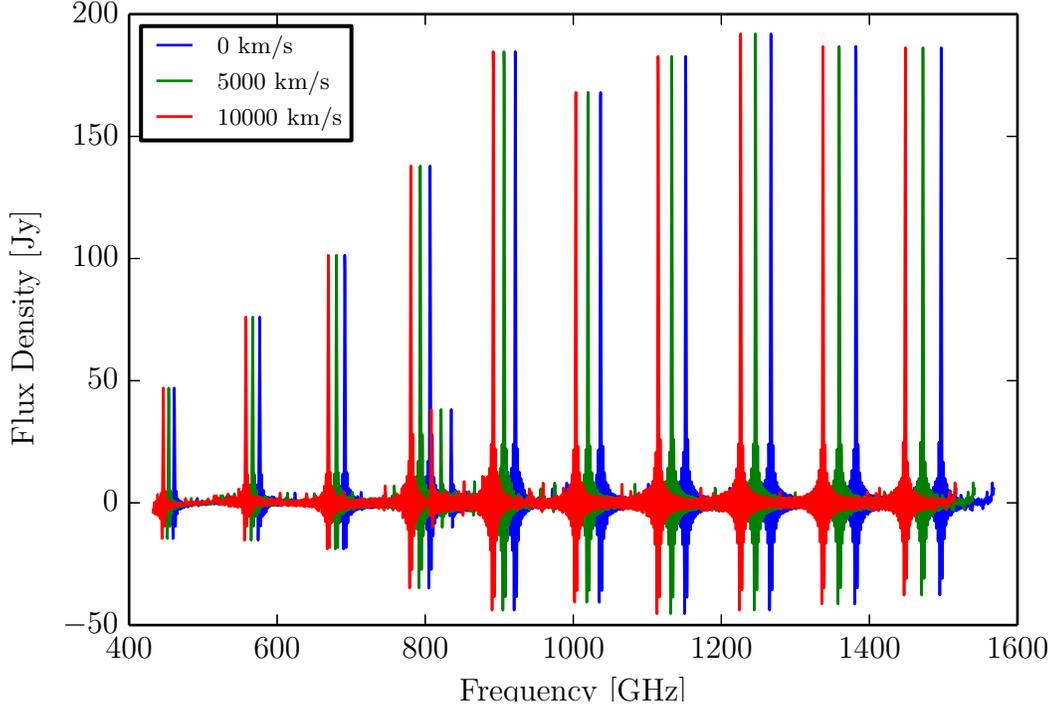


Figure 5.1: The co-added continuum subtracted spectrum of NGC 7027 has been plotted with various induced redshifts to show how source velocity affects the resulting spectrum.

rotational transitions of ^{12}CO , and this sequence of emissions is referred to as the ^{12}CO ladder.

A few relationships between shifted spectral features are derived here which will be of use later in this chapter. Solving for the shifted frequency in Equation 5.3, one obtains

$$f = \left(\frac{v}{c} + 1\right)^{-1} f_0. \quad (5.4)$$

Now considering the shifted frequencies of two different features f_1 and f_2 , the frequency difference between these features is

$$f_2 - f_1 = \left(\frac{v}{c} + 1\right)^{-1} (f_{02} - f_{01}), \quad (5.5)$$

where f_{01} , and f_{02} are the rest frequencies of f_1 and f_2 respectively. The important aspect of

this result is that the difference between the shifted frequencies is related to the difference between the rest frequencies by a velocity dependent scaling term. It is then instructive to examine to what extent the shifted and rest frequency differences change as a function of source velocity. This can be done simply by subtracting the rest frequency difference from both sides of Equation 5.5. Performing this operation one obtains

$$|(f_2 - f_1) - (f_{02} - f_{01})| = \left(1 + \frac{c}{|v|}\right)^{-1} |f_{02} - f_{01}|. \quad (5.6)$$

In the case of a positive source velocity (redshift), spectral features are compressed, as shown in Figure 5.1. In the converse situation where the source has a negative velocity (blueshift), the spectral features become more spread out. Equation 5.6 describes the absolute difference between shifted and rest frequency differences so as to not discriminate between these two cases. The utility of these equations will become clear in the next subsection. The most direct way of calculating a sources radial velocity is to identify a shifted spectral feature and match it with its corresponding rest frequency. Equation 5.3 can then be used to calculate the source radial velocity. The challenge then becomes the accurate determination of a shifted spectral feature's frequency position, and correctly matching this feature to it's rest frequency. This task is not trivial, but there are some factors which can be considered which make this process easier.

5.2.0.4 Why ^{12}CO ?

The most abundant molecule in the Universe is molecular hydrogen (H_2 , [83]). This is no surprize with hydrogen representing roughly three quarters of all baryonic matter in the Universe [84], while simultaneously possessing an unpaired valance electron. When hydrogen forms a covalent bond with another hydrogen atom, the result is a more stable energy configuration with electric charge distributed uniformly about the molecule's center of mass. From Equation 3.23, it can be seen that H_2 has no significant dipole moment. For both rotational and vibrational emission, the intensity of transitions is dependent on

the system's transition dipole moment (Equation 3.29), which is itself proportional to the dipole moment. As such, H₂ has no significant rotational or vibrational transitions, thus there are no H₂ lines visible in the low energy FIR band.

The second most abundant molecule in the Universe is carbon monoxide (¹²CO) [83], which in contrast to H₂, has a strong electric dipole moment. As a result ¹²CO rotational and vibrational emission is of high intensity and is particularly bright in the FIR band. An example of a FIR ¹²CO containing spectrum is shown in Figure 5.1.

As a result of its abundance and strong dipole moment, ¹²CO is not only likely to be present in most FIR sources, but also easily detected. As described in Chapter 4, the FF can successfully detect 98% of spectral features with a SNR ≥ 10 within ¹²CO containing spectra. Thus, in the vast majority of cases where a SPIRE spectrum contains ¹²CO, it is a safe assumption that the FF will return a line list containing the desired ¹²CO features.

Another advantage of ¹²CO is that it is a diatomic (linear) molecule whose rotational emission profile is easily described. As outlined in Section 3.3.2, the rotational frequency transitions of a linear molecule are to first order described by the rigid rotor model [49]. Using this approximation, emission occurs at integer multiples of $\Delta f = h/(4\pi^2 I)$, where h is Planck's constant and I is the moment of inertia of the linear molecule. For ¹²CO, this quantity is ~ 115.1 GHz, and simultaneously corresponds to the difference between adjacent ¹²CO rotational emissions. Centrifugal distortion does produce deviations from this constant value, though over the whole SPIRE band, this value changes by < 0.2 GHz which is trivial compared to the 1.2 GHz spectral resolution of high resolution SPIRE observations.

It is then possible to construct an array consisting of the difference between one ¹²CO line and all ¹²CO lines nominally within the SPIRE band. I call this the Characteristic Difference Array (CDA), which is given by

$$CDA = \{115.1 \times n\} \quad n = 0, 1, \dots, 9. \quad (5.7)$$

It is worth emphasizing that $n = 0$, is included in the CDA, which corresponds to the fre-

quency difference between the feature under investigation and itself. As such there are a total of 10 elements in the CDA corresponding to the 10 ^{12}CO rotational emission features (J=4–3 to J=13–12) nominally within the SPIRE band (Figure 5.1).

The CDA can be identified as $f_{0i} - f_{0j}$ in Equation 5.5, where i and j represent any ^{12}CO rotational transitions within the SPIRE band as long as $i \geq j$. As such, this equation can be recast to

$$CDA' = \left(\frac{v}{c} + 1\right)^{-1} CDA, \quad (5.8)$$

where CDA' is the difference between shifted ^{12}CO transitions. This relation applies equally well to cases where $j > i$, though this results in a negative sign. Similarly, Equation 5.6 can be recast in terms of the CDA as

$$\textit{tolerance} = |CDA' - CDA| = \left(1 + \frac{c}{v}\right)^{-1} \times CDA. \quad (5.9)$$

I refer to this result as a tolerance, the significance of which will become clear in the next section.

To summarize the key points so far, the objective is to find spectral features suitable for an automated search, which can be reliably identified for the purpose of estimating source radial velocity. ^{12}CO is both abundant in the ISM, and reliably detected within SPIRE spectra by the FF, while also expressing a periodic and distinct spectral profile. As such, ^{12}CO was chosen as the primary candidate for the automated radial velocity estimating routine implemented in the FF.

5.3 Routine Outline

5.3.0.5 General Overview

In order to limit confusion from the detailed description that follows, I will begin by presenting a very general overview of the main body of the redshift routine. A flowchart showing the key steps in the main algorithm is shown in Figure 5.2, with a detailed example

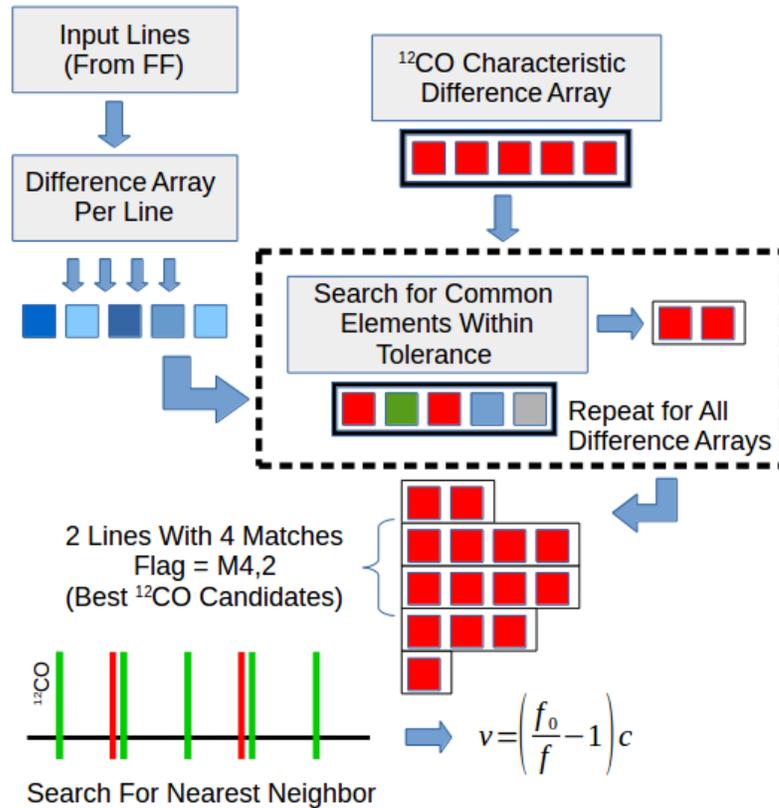


Figure 5.2: A simplified flowchart showing the key steps of the radial velocity estimating routine. A line list of emission features is input to the routine. For each feature, a frequency difference array is calculated and searched for elements matching the CDA to within a velocity dependent tolerance. Features with the most number of matches are the most probably ¹²CO candidates. These features are paired with the nearest ¹²CO rest frequency, and a radial velocity is calculated. See main text for a more detailed description.

of the routine given in appendix A.

A line list derived from earlier steps in the FF routine are used as inputs. For each feature, a difference array is calculated between the frequency of the feature in question and that of all other lines. For each difference array, a pattern seeking search then initiates that attempts to find elements matching the CDA to within a predetermined tolerance. Features which all have the maximum number of matches are the most compelling ¹²CO candidates. The candidate ¹²CO features are then paired with their nearest ¹²CO rest frequency neighbour. Having obtained a shifted frequency and rest frequency for specific transitions, a radial velocity is calculated using Equation 5.3.

In addition to the radial velocity estimate, the routine returns a velocity error and two flagging values which parameterize the reliability of the estimate. The first parameter, denoted as “M”, indicates the maximum number of matches between the difference arrays of the ^{12}CO candidate features and the CDA. The second parameter, denoted as “N”, indicates the total number of lines identified as ^{12}CO candidates (lines which have the maximum number of matches between their respective difference array and the CDA). In Figure 5.2 the maximum number of matches is 4, and there are two features with this maximum number of matches. As such this hypothetical estimate would receive a classification of $M, N = 4, 2$.

Determining a binary value in regards to the reliability of each estimate based on the flagging parameters is not straight forward. This subject will be addressed in Section 5.4.2 though at this time a few rules of thumb can be developed. Since there are nominally 10 ^{12}CO features in the SPIRE band, and the frequency difference between these emissions should all match the CDA, the best possible flag corresponds to $M, N = 10, 10$. In general, the higher the values of M or N the more reliable the estimate, and it is especially encouraging when both values are equal. As such, an assessment of these flags is a balance between magnitude and consistency.

After analyzing the results of the redshift routine acting on a synthetic dataset, it was determined that the N flagging parameter was well suited to play the role of a binary threshold on the quality of the estimate. In what follows, all references to “the flagging parameter” should be interpreted as a reference to N, unless otherwise stated, and “high quality” estimates are those with $N > 6$.

5.3.1 Detailed Routine Description

5.3.1.1 A Dedicated Search For ^{12}CO

The routine initializes by first assuming some maximum source radial velocity which is used to calculate a velocity dependent tolerance on the agreement between the CDA and

elements from the difference array of individual features. The initial maximum velocity is 6,000 km/s which is high enough to include expected radial velocities of most galactic sources, and becomes the variable v in Equation 5.9. The higher the maximum velocity, the greater the tolerance, which should ideally be as small as possible to limit features in close spectral proximity to true ^{12}CO features from being interpreted as ^{12}CO . Sources with a radial velocity greater than the maximum velocity will inevitably return poor estimates. As such, the routine proceeds in an iterative fashion increasing the maximum velocity by 2,000 km/s until an upper limit of 14,000 km/s or until a reliable estimate is obtained. This upper velocity limit, and the validation for what constitutes a reliable estimate is discussed in Section 5.4.2. For now it will suffice to know that a reliable estimate is based on the flagging parameters.

The redshift routine receives as inputs a line list corresponding to features detected from within the main body of the FF routine, along with associated SNR, and feature frequency errors, which are products of the linear regression algorithm employed by the FF. Since ^{12}CO emission lines are the features of interest, the line list is first filtered by SNR, removing any absorption features.

To initiate the ^{12}CO feature search, the first feature in the remaining list is selected, and a Frequency Difference Array (FDA) is calculate between the feature in question and all the frequencies in the line list (including itself). The routine then iterates over the elements of the CDA checking for elements within the FDA which match within the frequency dependent tolerance of the specific CDA entry. If at least one match is found, a “match” counter, which is initialized at zero, increments by one. Iterating over the CDA in this way results in a maximum match counter of 10, and since zero is included in the CDA, a minimum value of 1 is anticipated. This process then repeats for every feature in the line list. Generally, features will have FDAs with negative entries, while the CDA is defined to be positive. As such, the iterations over the CDA occurs twice, with negative FDA values counted separately from positive values, though the final match counter contains both. This

process results in a double counting of the zero term, which is subsequently corrected.

At this point in the routine, each feature in the line list has its own match counter variable. The maximum value from the set of features is determined, and all features with this value are selected as potential ^{12}CO candidates. These features are then paired to their closest ^{12}CO rest frequency neighbour. Due to the tolerance parameter, it is possible that multiple features will be paired to the same ^{12}CO rest frequency. This tends to occur when multiple features are clustered around true ^{12}CO lines. When multiple features are paired to the same ^{12}CO rest frequency, the set is filtered by SNR with all but the highest SNR line being removed.

With multiple ^{12}CO candidates paired to their respective rest frequencies in hand, a radial velocity is calculated for each line using Equation 5.3. All features should produce a similar velocity estimate as they are derived from the same molecular species. Under this assumption, if the standard deviation of the velocity estimates is ≥ 100 km/s the line which deviates the most from the median velocity estimate is removed. This process repeats until the standard deviation is below the threshold. At this point the flagging parameters are determined; the “M” parameter corresponds to the maximum match counter mentioned above, and “N” corresponds to the total number of surviving ^{12}CO candidates. The radial velocity estimate is calculated as the median value of the velocities corresponding to the ^{12}CO candidates. If $N > 1$, the velocity error is calculated as the standard deviation of the individual velocity estimates, otherwise the velocity error is calculated using error propagation of the uncertainty in the line position of the single spectral feature.

If a reliable radial velocity estimate has been obtained based on rotational ^{12}CO emission, the routine terminates. If not, a separate routine begins which performs a much simpler search for the fine structure $^3\text{P}_1$ - $^3\text{P}_0$ ionized atomic nitrogen line [NII], with a corresponding 1461.1 GHz rest frequency [85]. The execution of this search requires 10 or fewer features to be detected by the FF in the spectrum under analysis. The search begins by identifying all detected features within 60 GHz of the [NII] rest frequency. If any features

are found in this range, the one with the highest SNR is assumed to be [NII], and a radial velocity is calculated if this feature has an $\text{SNR} \geq 10$. If this test is successful, the estimate is assigned the flag $M, N = \text{NII}, 1$. Justification for this final check is provided in the next subsection. The threshold of 10 or fewer detected features was chosen empirically to accommodate a subset of spectra which express more than just the [NII] feature, while still being restrictive enough to provide some confidence when using the broad 60 GHz search window when blindly looking for the [NII] line.

At this point in the routine, it is possible that no acceptable radial velocity estimate has been obtained. As a final attempt to estimate the redshift, a 50 GHz search window on either side of every ^{12}CO rotational emission rest frequency, with the highest SNR feature within each window ascribed to their corresponding rest frequency. The resulting median radial velocity estimate is assigned the flag $M, N = 1, 1$, the least reliable flag possible. If no features are found in any of these windows, the redshift routine returns “NaN”.

5.3.1.2 Justification and Considerations for NII Exception

Observations exist which express a strong [NII] feature, with no detectable ^{12}CO . One such example is the massive star forming region SH-104 shown in Figure 5.3. Such observations can be explained by the fact that ^{12}CO has a dissociation energy of ~ 11.11 eV [86], while nitrogen ionizes at ~ 14.53 eV [87]. As such, the ^{12}CO rotational transitions and ionized nitrogen lines probe different energy regimes of the ISM. Though the transitions are not associated with the same components of the ISM, both rotational ^{12}CO and [NII] lines can appear in the same field of view depending on the distribution of warm and cold molecular gas. The barred spiral galaxy NGC 1068 is an example which shows such characteristics.

Figure 5.4 shows a SPIRE spectral image for NGC 1068. On the left is an integrated flux map for the SSW band, while on the right is a subset of corresponding SSW spectra about the central region. The brightest pixel in the integration map is shown with a magenta

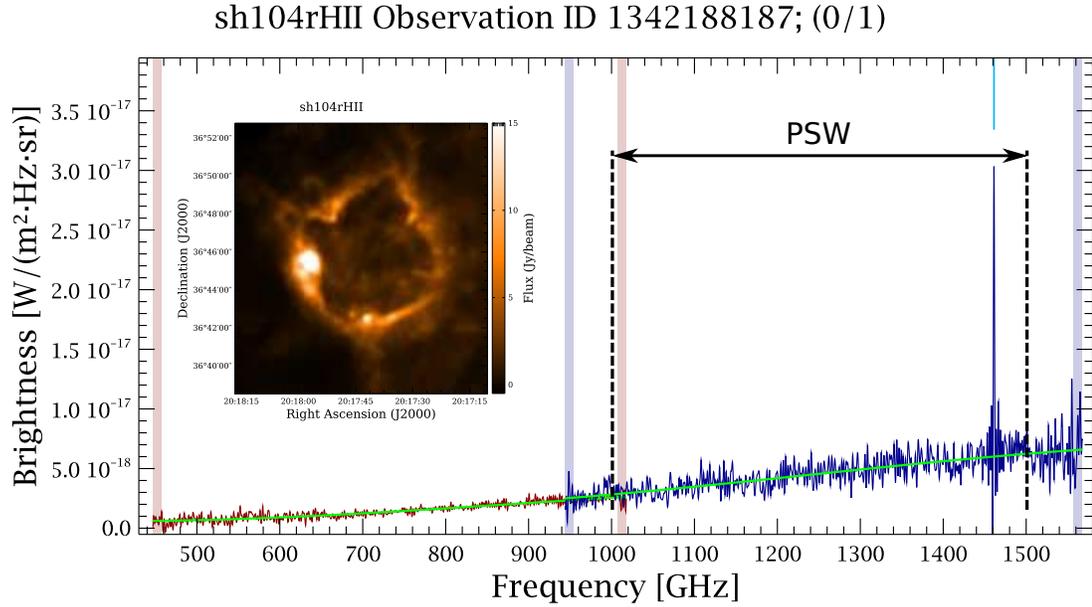


Figure 5.3: The SPIRE spectrum and PSW image of Sh 104. Sh 104 is an ultracompact ionized HII region with embedded massive stars driving the ionization front [88]. As such the presence of an ionized nitrogen line is not unexpected. Dashed lines indicate the approximate PSW band edges [59].

outline, while the faintest pixel is outlined in green. For reference, the brightest pixel occupies the 13th row and 12th column in the integration map, while the spectral image on the right contains annotations with the notation (row,column).

Figure 5.5 shows the spectrum of the brightest pixel in Figure 5.4 with more detail. The faint ^{12}CO ladder is evident, indicating the presence of cold molecular gas, which is contained in the same field of view as the warm molecular gas expressing the [NII] line. Physically, this is explained by a warm outer molecular cloud ejected by the active galactic nucleus at the center of NGC 1068 [89]. The lower energy ^{12}CO rotational emission is attributed to star forming activity in the galactic disk residing in the background of the image [90].

For the spectrum in Figure 5.5, a velocity estimate derived from the ^{12}CO features gives an average value of 1114.2 km/s with a standard deviation of 16.2 km/s, while an estimate based on the [NII] feature returns a velocity of 1175.1 ± 5.0 km/s. This 61 km/s differ-

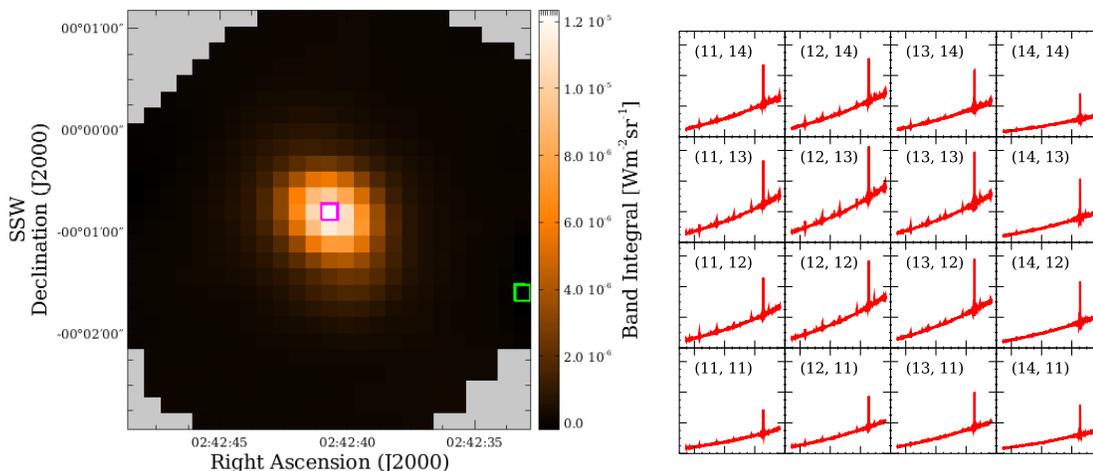


Figure 5.4: Flux integration map (left) of the SPIRE SSW spectral cube with the faintest pixel indicated with a green square, while the brightest pixel is indicated with a magenta square. The brightest pixel (enlarged in Figure 5.5) corresponds to the (13, 12) element in the corresponding SSW spectral grid (right). The spectral grid shows the distribution and intensity of ¹²CO (faint) and [NII] (bright) around the central region of the flux integration map.

ence indicates that the dynamics of the cold and warm ISM components are different, and an awareness of this fact is required when using velocity estimates derived from different spectral transitions. The significance of this effect is clearly shown in Figure 5.6. This figure show a velocity map of NGC 1068 generated from the radial velocity estimating routine. Estimates derived from [NII] are marked with a “N”, with all others derived from ¹²CO with the number indicating how many ¹²CO candidates were used to derive the estimate. The special case where all 10 ¹²CO line candidates were used is indicated with an “A”. Focusing on the bottom right where the estimate notation changes from 5 to N, one can see a clear discontinuity in the velocity profile along with a gradient reversal.

5.4 Validation

This section presents evidence which can be used to assess the performance of the automated source velocity estimating routine. This is accomplished by analyzing the results of the routine operating on both real and synthetic data.

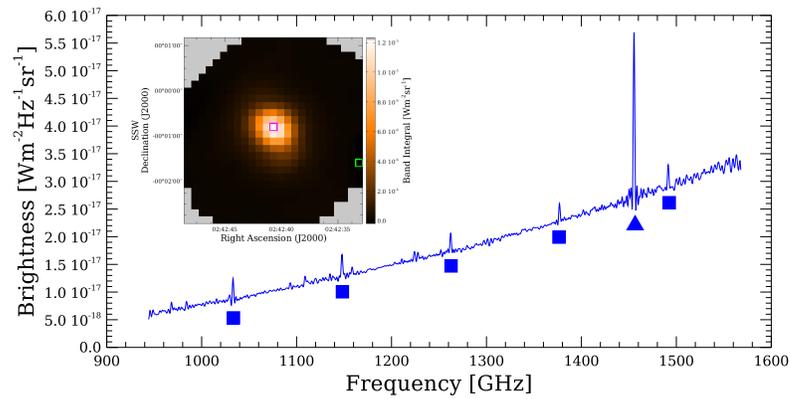


Figure 5.5: SSW spectrum of the (13, 12) element of the spectral grid in Figure 5.4 showing, in more detail, the presence of ^{12}CO rotational emission features (square) and the [NII] emission (triangle).

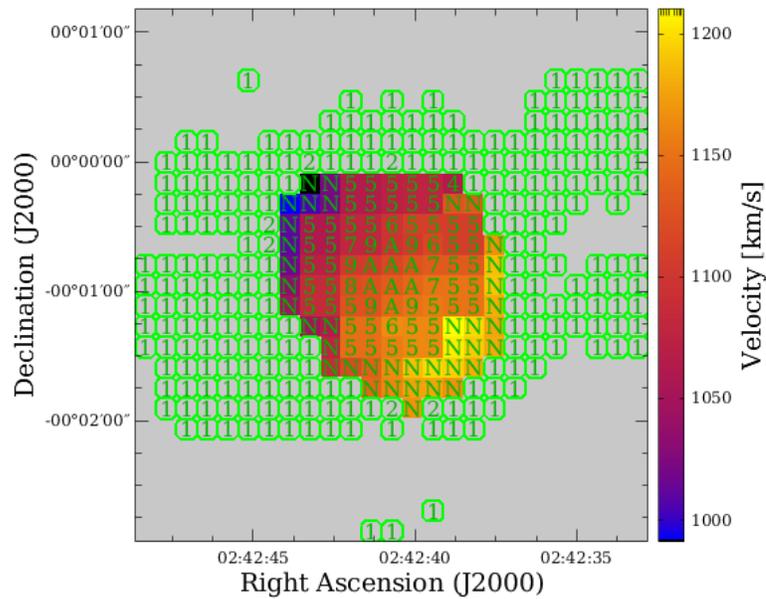


Figure 5.6: FF velocity map of NGC 1068. Annotations indicate which features were used for the velocity estimate, with “N” indicating estimates based on [NII] emission. The remaining estimates are based on ^{12}CO rotational emission with the number indicating how many features were used to generate the estimate (“A” for All 10 features). Grey squares indicate pixels without a valid estimate. An abrupt discontinuity is seen in the bottom right where estimates transition from those based on ^{12}CO to those based on [NII]. This shows that the dynamics of the cold components of the ISM do not always match the dynamics of the warm ionized components, and that velocity estimates based on the two species can produce different results.

5.4.1 Real Data

The real data used for validation consists of two independent datasets. The first dataset is composed of velocity estimates derived from high spectral resolution HIFI observations. The velocity estimates themselves were calculated by Benamatti et al. (in-preparation) and were shared with the SPIRE FTS instrument team. The second dataset was compiled by myself through manual inspection of SPIRE spectra and comparing the physical data with the results from the FF. Details on how this dataset was constructed are presented below. All real datasets are compared with suitable HR sparse SPIRE spectra.

5.4.1.1 HIFI Velocity Estimate Comparison

The list of HIFI velocity estimates used for comparison consists of a subset of the total list produced by Benamatti et al. For each source, SPIRE and HIFI results were paired by searching for the HIFI entry with the nearest on-sky pointing to the corresponding SPIRE observation. This step is used to ensure that both instruments are observing the same region of the same source, maximizing the potential correlation between velocity estimates. The dataset used for comparison contained 258 velocity estimates.

A correlation plot showing the agreement between HIFI and FF estimates is shown in Figure 5.7. The panel on the left shows the full velocity range, while the panel on the right is a close up of the low velocity cluster. A coloring and marker convention has been adopted in these plots which is used throughout the remainder of this chapter. Colors are chosen according to the estimate's corresponding N flagging parameter as follows: 10, 9, are blue; 8, 7, are green; 6, 5, are cyan; 4, 3, are orange; 2, 1 are red. Estimates based on the [NII] feature are presented as black triangles, while estimates with matching flag parameters are marked with an "x", or an "o" otherwise.

It can be seen that there is good agreement between both methods for all estimates with a flag greater than 2. Estimates with poor flags are red and show significant deviations from the dashed one-to-one line, which is expected. Estimates based on the [NII] feature

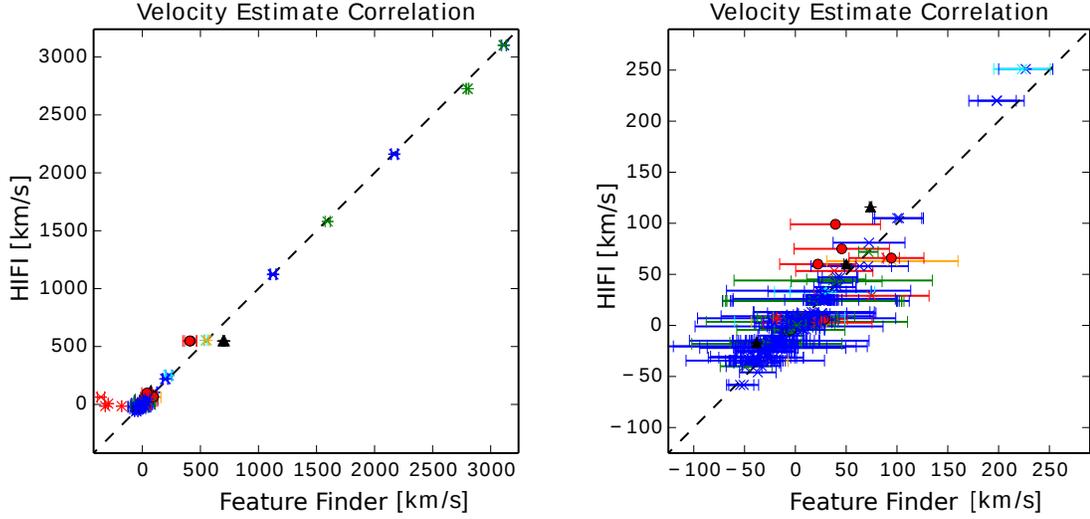


Figure 5.7: Correlation plot of velocity estimates derived from HIFI and those derived from the FF automated velocity estimating routine. The right panel shows a zoomed in version of the panel to the right. Color and maker conventions are described on page 140.

also tend to deviate slightly from the one-to-one line, but this can be explained based on the different energy regimes probed by ^{12}CO and $[\text{NII}]$ features, as described in Subsection 5.3.1.2. Every estimate based on $[\text{NII}]$ was manually inspected, and without exception the redshift routine correctly identified the $[\text{NII}]$ feature.

Figure 5.9 shows the performance of the FF velocity estimates in terms of the N flagging parameter. The quantity being presented is the mean residual velocity, calculated using

$$\bar{\Delta} = \sum_{i=1}^n \frac{|HIFI_i - FF_i|}{n}, \quad (5.10)$$

where n is the total number of estimates with a given flag value, with $HIFI_i$ and FF_i representing complementary HIFI and FF velocity estimates, respectively. The error is calculated simply as the standard deviation of all the residuals. The figure expresses an expected trend; both the mean absolute residual and the associated error decrease with increased flag value.

Finally, Figure 5.8 shows a completeness curve indicating what percentage of estimates have an absolute residual less than values shown on the x-axis. The figure contains an inset which is a histogram of the absolute residuals. For both of these plots, only estimates

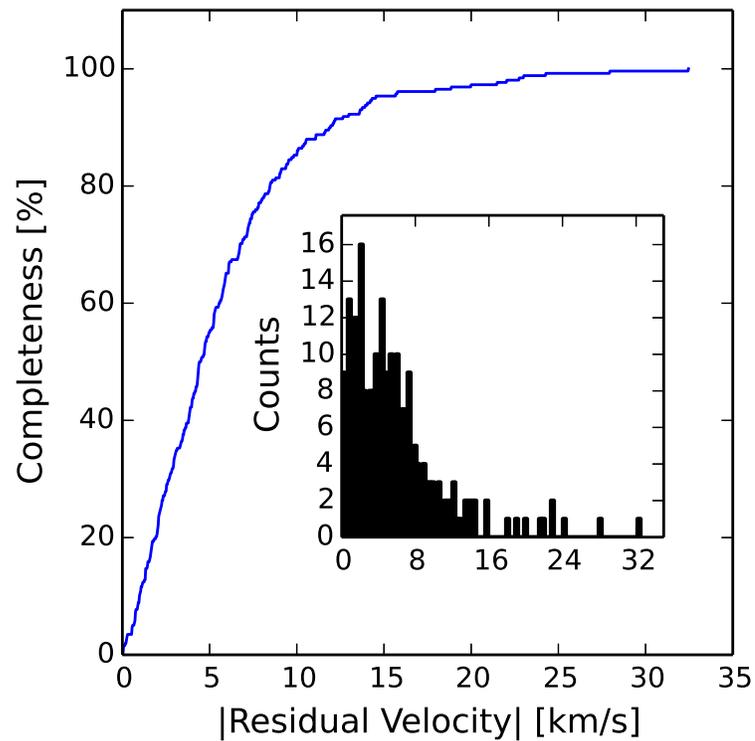


Figure 5.8: A completeness curve and associated histogram of velocity residuals between velocity estimates based on HIFI data and those derived from the automated FF velocity estimating routine. Only estimates with $N > 3$ have been included.

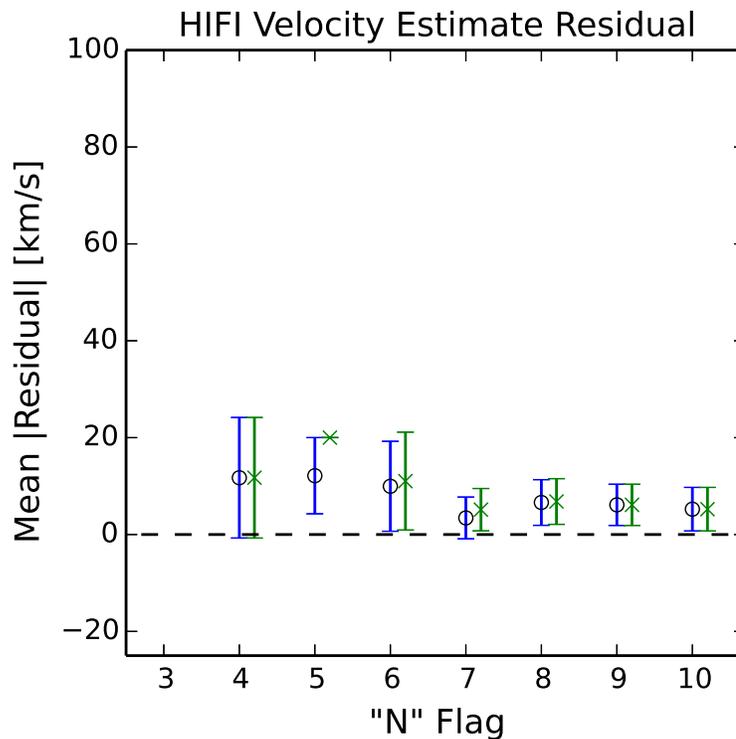


Figure 5.9: Mean absolute residual velocities for the HIFI/FF comparison in terms of the N flagging parameter. Estimates where $N=M$ are indicated with “x”, otherwise “o” is used. It can be seen the mean absolute residual decreased with increased N value.

with a flag value greater than 3 have been used. It can be seen that the completeness curve begins to saturate around 20 km/s, and that all estimates have an absolute residual less than 35 km/s.

5.4.1.2 Manual Inspection Velocity Comparison

The second real dataset was constructed from a manual comparison between SPIRE spectral data and the FF output products for the spectrum in question. Velocities were determined in a way that is very similar to how the redshift routine nominally generates estimates. Each spectrum was plotted with all features detected by the FF overlaid onto the spectral data. The results were inspected manually, and FF detections which unambiguously corresponded to ^{12}CO rotational emission were identified. In an automated fashion, these identified features were then paired to their nearest ^{12}CO rest frequency. Radial velocity

estimates were then calculated following the standard redshift routine protocols. In some cases the [NII] feature was used when ^{12}CO identification was difficult or impossible. The complete dataset used for this validation consisted of 220 spectra.

Repeating this process for hundreds of spectra was a monotonous prospect. I expedited the full workload by developing a user interface which automated most of the repetitive tasks. The software included zoom features, SNR filtering, inspection and editing of active ^{12}CO line lists, and progress saving features allowing this process to take place over multiple sessions.

This manual inspection process to a large extent mimics the operation of the redshift routine, though a few subtle distinctions are worth mentioning. This dataset was generated using the results from a previous iteration of the FF script. In principle this could result in slight variations between the frequencies of detected features used in the manual comparison, and the features used by the redshift routine. For the manual inspection, only strong and/or isolated ^{12}CO features were selected, and since the same linear regression algorithm was used for all FF versions, no significant change in frequency is expected for the line used. Due to the time required to generate this dataset, no subsequent dataset was generated for the updated FF products. Unlike the redshift routine, no effort was made to identify all of the ^{12}CO features. In general, only a few easily identifiable transitions were selected, both for the sake of time, and to ensure only valid ^{12}CO features were used. Both of these factors contribute to the overall uncertainty of the manual inspection dataset results. The resulting error presented, is a combination of uncertainty in the fitted line center, and the distribution of velocities from the ^{12}CO features, which are the same sources of error in both datasets, so the comparison remains valid.

Correlation plots of the results are shown in Figure 5.10. A result similar to the HIFI comparison is obtained with some low flagging parameter estimates deviating significantly from the one-to-one line. Again, estimates based on [NII] in some cases deviate appreciably from the one-to-one line. Each estimate based on [NII] with a deviation higher than

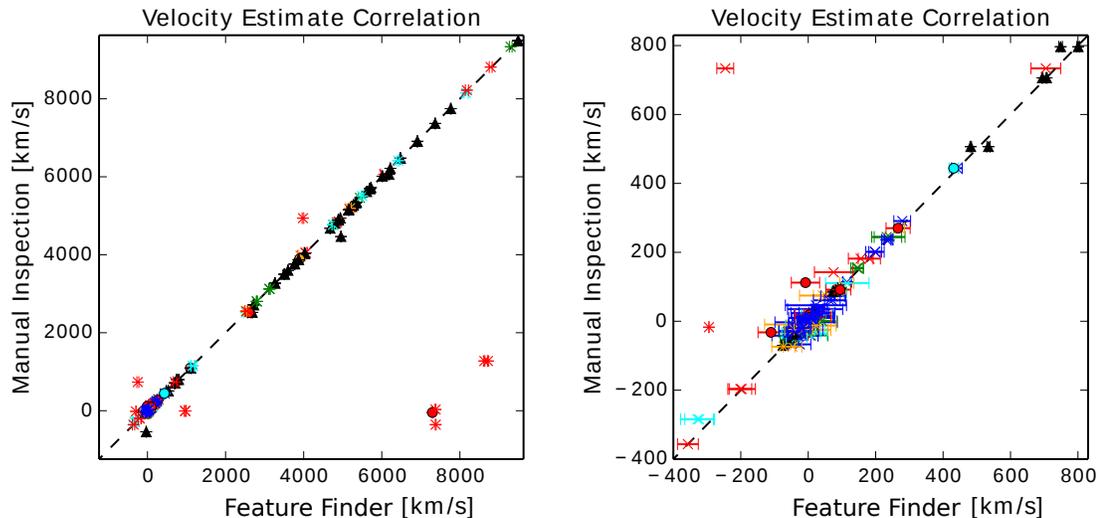


Figure 5.10: Correlation plot of velocity estimates derived from manual inspection and those derived from the FF automated velocity estimating routine. The right panel shows a zoomed in version of the panel to the left. Color and maker conventions are described on page 140.

20 km/s was inspected in more detail, and it was determined that in all cases the correct [NII] feature was identified by the redshift routine. The deviations are explained by the different components of the ISM probed by ^{12}CO and ionized nitrogen (Section 5.3.1.2).

Plots showing the mean absolute residual in terms of flagging parameter, and residual completeness for flags greater than 3 are presented together in Figure 5.11. These plots show a similar result to the HIFI comparison with improved estimate accuracy for higher flagging values, and roughly 90% of residual velocities having a value ≤ 20 km/s.

5.4.1.3 Combined Results

Both the HIFI and manual inspection datasets were combined, choosing the manual inspection based estimate when competing estimates occurred. Out of the combined set, the redshift routine returned 251 high quality velocity estimates based on ^{12}CO ($N > 6$). The mean residual of all of these observations was 5.80 km/s, with a standard deviation of 7.12 km/s.

Estimates based on [NII] were analyzed separately, with the combined dataset contain-

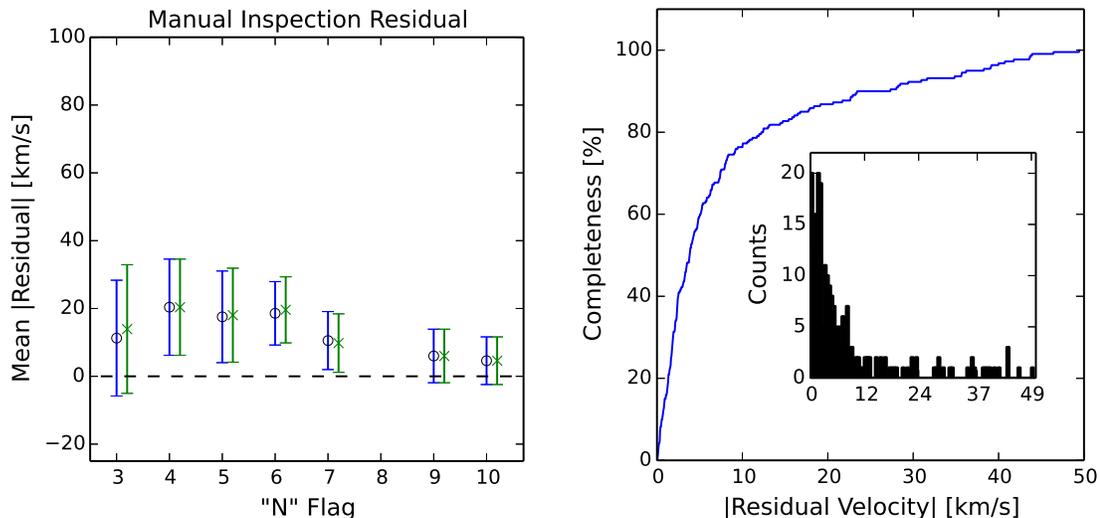


Figure 5.11: Mean absolute residuals for different N values of the manual inspection/FF comparison is shown in the left panel. The resulting completeness curve and associated residual histogram for velocity estimates with $N > 3$ is shown on the right.

ing 49 such estimates. Of these, 4 had absolute residuals > 100 km/s. All 4 were manually inspected, and it was determined that each one was based on the correct identification of a [NII] feature within the routine. After masking these outliers, the mean absolute residual was 4.14 km/s, with a standard deviation of 5.66 km/s.

Figure 5.12 puts these values in context with the theoretical velocity resolution of the SPIRE FTS instrument for HR observations. Note that the SPIRE velocity resolution curve has been scaled down by a factor of 30. It can be seen that the redshift estimating routine produces velocity estimates between ~ 50 to 100 times below the resolution limit. These results are in line with the SPIRE FTS line center repeatability which is quoted as having a systematic offset of ~ 5 km/s with a spread of ~ 7 km/s [91].

5.4.2 Simulated Data

To further assess the quality and limitations of the redshift routine, a large set of 20,000 simulated FF results were generated. The goal of the synthetic data was to represent realistic cases of SPIRE spectra containing ^{12}CO rotational emissions within a continuum of

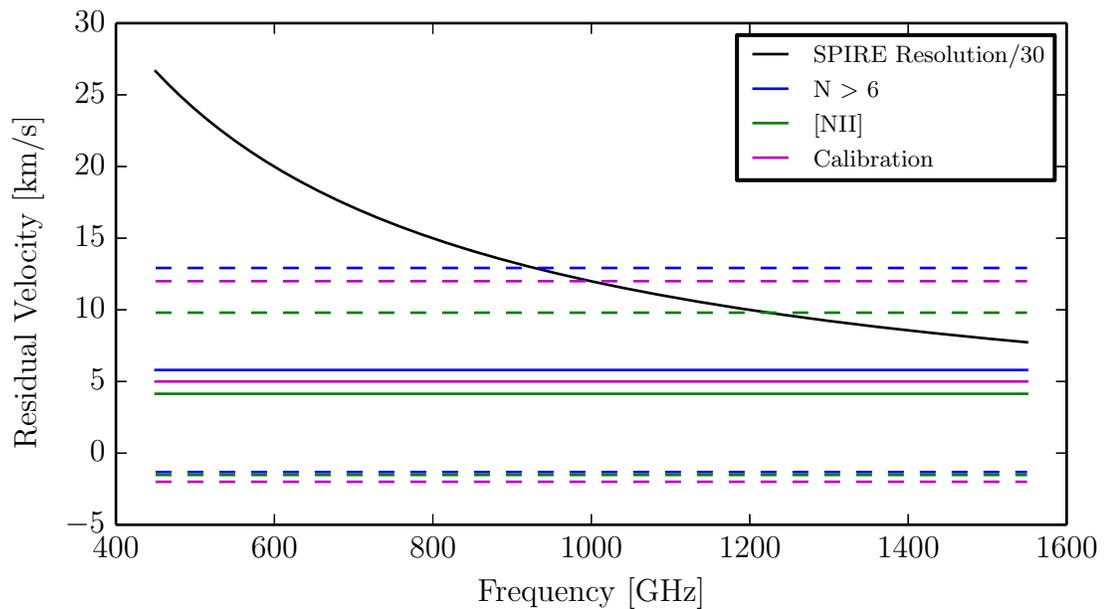


Figure 5.12: The velocity-equivalent spectral resolution of the SPIRE instrument divided by 30 is shown as a black curve. The mean absolute residual velocity for the combined manual inspection/HIFI comparison dataset is also presented with estimates based on ^{12}CO with $N > 6$, and based on [NII] emission shown with solid blue and green lines, respectively. The velocity equivalent of the SPIRE FTS line center repeatability is shown in purple. Dashed lines indicated the associated standard deviation of each dataset.

radial velocities between $\pm 15,000$ km/s. Recall that the redshift routine acts on a line list with corresponding SNR and frequency error, so these are the products that must first be generated by the simulation.

5.4.2.1 Generating FF Data

The routine begins by generating the required quantities for the ^{12}CO ladder. A template for the ladder's SNRs is obtained using the results of the FF operating on the co-added spectrum of NGC 7027. The co-added dataset is essentially the average spectrum from all 31 HR sparse NGC 7027 observations [78]. The SNR returned by the FF for each ^{12}CO feature is used to construct the SNR template. NCG 7027 expresses particularly bright ^{12}CO features with little line blending [78]. In addition, the averaging of multiple spectra cancels some of the uncorrelated noise. As such, the SNR template expresses values higher than would be expected in most spectra. In generating the ^{12}CO SNR for each simulated spectrum, the SNR template was scaled by a random factor between 0.1 and 1.1 chosen from a uniform distribution. The asymmetric scaling factor about unity was chosen to account for the uncharacteristically high SNR of the NCG 7027 co-added data. The ^{12}CO frequencies were then shifted using a randomly selected radial velocity between $\pm 15,000$ km/s from a uniform distribution. Finally, there was a 25% chance that ^{12}CO features would be removed. In cases where lines were removed, there was an equal probability that between 1 and 4 lines would be removed that were selected at random.

The simulation then constructs parameters for additional lines. The number of additional lines was randomly selected between 10 and 25, and each line was given a random SNR between 5 and 100. The minimum value of 5 was imposed to follow the FF minimum detection protocol, whereas the maximum value was chosen to mimic the simulations used for testing the FF detection accuracy. The additional lines were then positioned randomly between 447 and 1547 GHz. The number of lines, SNRs, and positions of the additional lines were all selected from a uniform distribution.

The ^{12}CO ladder and additional lines were then combined into a single list. Additional realistic white noise was imposed by varying the frequency of all added features. The frequency perturbation for each feature was chosen randomly from a normal distribution with a standard deviation of 0.11 GHz. This is the average uncertainty of the line centers returned from fitting the NGC 7027 co-added data. The frequency error list for each synthetic dataset was uniformly populated by this value as well.

5.4.2.2 Simulation Results

A correlation plot showing the relationship between input and estimated velocities is shown in Figure 5.13. The left panel shows the full dataset, and a few characteristics are immediately apparent. The redshift routine performs very well within the velocity range $\pm 14,000$ km/s with only a handful of estimates deviating from the one-to-one line. Outside of this range, however, is an immediate breakdown of the routine's ability to determine accurate estimates. Though the breakdown at high velocities is not desirable, it is encouraging that not a single estimate which has been marked as a high quality estimate (blue or green), has returned an inaccurate velocity estimate; the only points which deviate significantly from the one-to-one line are those which receive a flag of $N < 5$. The right panel shows a close up view of the positive velocity edge where the breakdown begins to occur. This region is marked by a rather smooth transition from the highest quality estimates (blue), to the lowest quality estimates (red). This figure demonstrates a great deal of confidence in the flagging parameter's accuracy in determining the reliability of velocity estimates.

The breakdown at $\sim 14,000$ km/s is expected. In one of the final steps of the redshift routine, ^{12}CO emission candidates are paired to their nearest ^{12}CO rest frequency. Referring back to Figure 5.1, it can be seen that at high radial velocities, high J ^{12}CO rotational transitions approach the J-1 rest frequency. If the source velocity is high enough, then a redshifted high J feature will be erroneously associated with a blueshift corresponding to the J-1 ^{12}CO rest frequency. A similar effect occurs with spectra from sources with

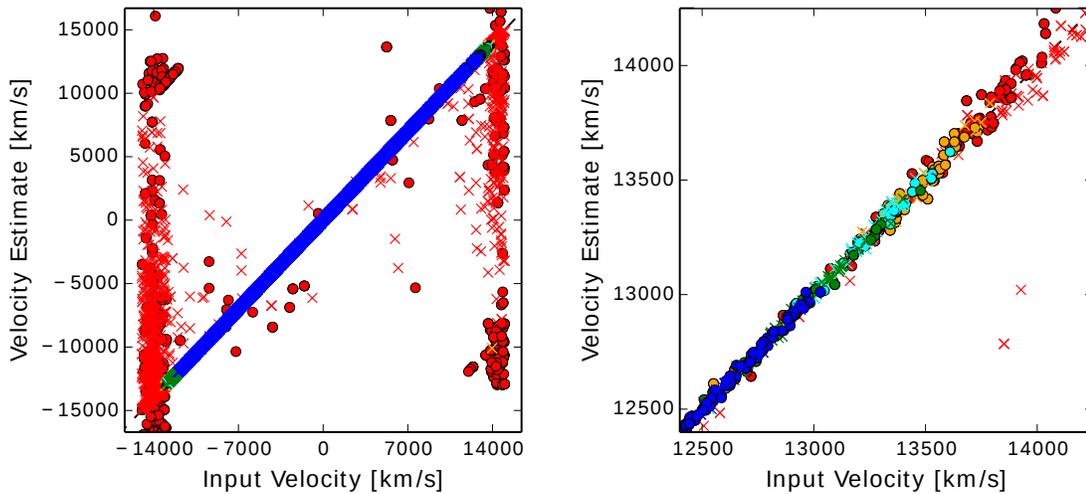


Figure 5.13: Correlation plot of velocities input to a simulated dataset and the resulting estimates from the FF automated velocity estimating routine operating on the synthetic line lists. The right panel shows a zoomed in version of the panel to the right. Color and maker conventions are described on page 140.

radial velocities $\sim -14,000$ km/s. Since the threshold for a high quality estimate is $N > 6$, the redshift routine can tolerate up to 3 such miss-associations while still producing a valid estimate. Using the rigid rotor approximation, this places the velocity limit at $\sim 14,255$ km/s in good agreement with Figure 5.13. For this reason the maximum assumed velocity, which dictates the tolerances within the redshift routine, is capped at 14,000 km/s.

The left panel of Figure 5.14 shows the mean residual of the simulation broken up by the flagging parameter. The right panel shows the completeness curve and residual histogram for all velocity estimates with a flag > 3 . By themselves, these plots only give a visual representation of the noise level of the simulation itself, but comparing these plots to their real data analogues (Figures 5.9, 5.11), its can be seen that the simulation results are consistent with real data. Similar to the results for real data, $\sim 90\%$ of estimates have a residual less than 20 km/s, and accuracy of velocity estimates is dependent on the flagging parameter, with higher values returning better estimates.

The results of this simulation can be used to motivate a binary criteria for “high quality estimates”. Visual inspection of Figure 5.13 reveals that estimate quality does not begin

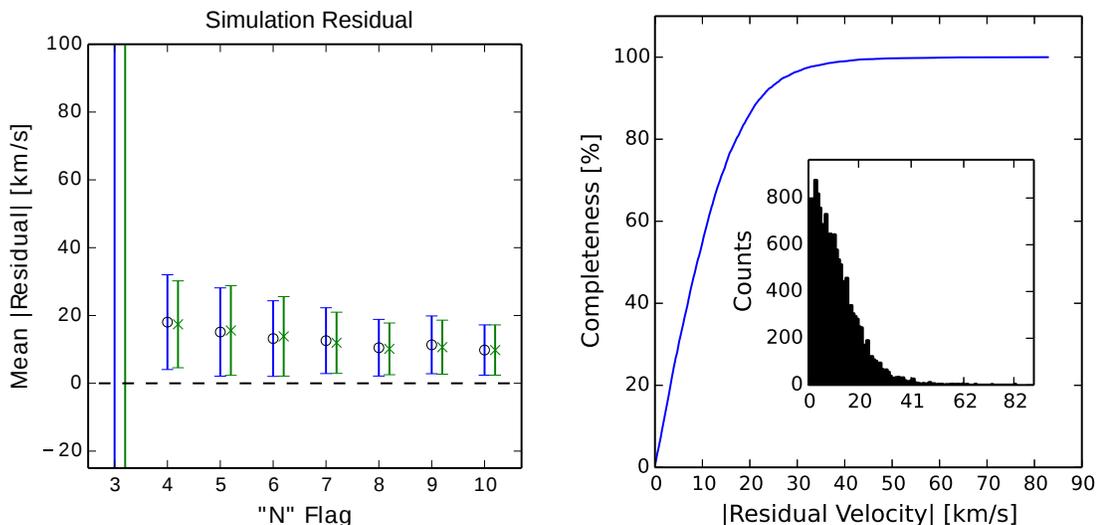


Figure 5.14: Mean absolute residuals for different N values of the simulated dataset is shown in the left panel. The resulting completeness curve and associated residual histogram for velocity estimates with $N > 3$ is shown on the right.

to degrade significantly until N is lower than 5. The FF team selected $N > 6$ with $M > 2$ as sufficiently reliable, and this is the standard for estimates included in FF products and public HSA catalogue.

To analyze the implications of this selection criteria, estimates with an input velocity between $\pm 14,000$ which returned an absolute residual ≤ 50 km/s were isolated. A histogram of the resulting N flag parameters was then generated which is shown in Figure 5.15. The y-axis is presented in terms of percent of total estimates meeting the selection criteria. 82.1% of the resulting estimates are labelled with a flag > 6 , and little is gained from including lower order flag values. As such, the selection criteria implemented in the FF is well optimized according to this simulation.

5.5 Feature Finder Implementation

At the time of writing this thesis, the FF team has made publicly available only the feature catalogue for SPIRE sparse observations.³ The first iteration of the public release

³www.cosmos.esa.int/web/herschel/spire-spectral-feature-catalogue

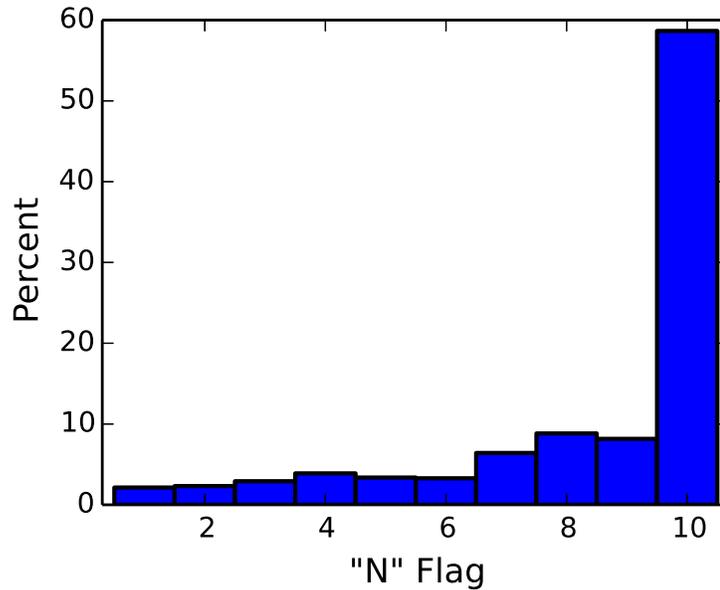


Figure 5.15: A histogram of the N flagging parameter with input velocities between $\pm 14,000$, with a residual velocity ≤ 50 km/s. Effectively, this is the distribution of flagging parameters with high accuracy estimates.

of the companion catalogue for mapping observations is also available at the same address. This section will summarize the results of the reshift routine as they pertain to the sparse observation catalogue, and will comment on the current developments for the mapping observations.

5.5.0.3 Sparse Observations Catalogue

The redshift routine produced 624 high quality radial velocity estimates which were incorporated into the sparse catalogue. The distribution of these velocity estimates is shown in left panel of Figure 5.16. An on-sky plot showing the location of all sparse observations and their corresponding velocity estimate quality is shown in Figure 5.17. The background of this image is a composite of the Planck 857 GHz integration map, and the ^{12}CO (J=2-1) emission map [92]. There is a higher density of observations on the galactic plan compared to the rest of the map. This is because, with the exception of redshifted galaxies, the majority of FIR sources observed by SPIRE (cold molecular clouds, stellar nurseries, etc), are

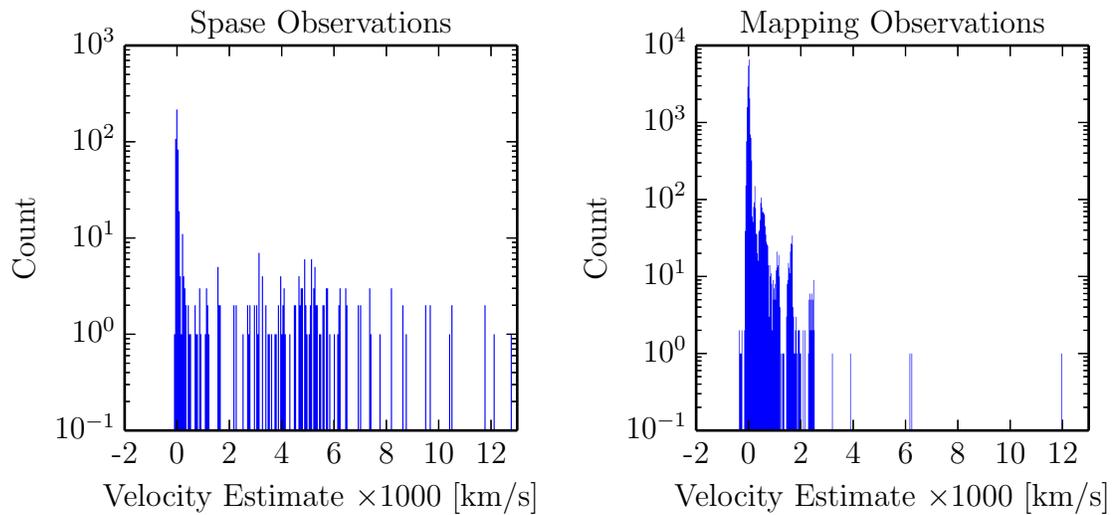


Figure 5.16: The distribution of radial velocity estimates for SPIRE sparse (left), and mapping (right) observations. An expected bias of source moving away from observer (i.e., positive velocity) is present in these histograms. It is also worth noting that sparse observations, which are more likely to be of extra-galactic sources, tend to have higher radial velocities than mapping observations, which are generally within the Milky Way itself. This is also an expected result.

located within the Milky-Way itself.

It was predicted that sources containing a large number of features would result in poor estimates, however, the redshift routine has proven to be robust in terms of the number of identified features. Of the 90 spectra with > 40 identified features, only 5 returned poor flag parameters. Of these 5 sources, 1 returned a correct velocity estimate, but did not meet the baseline flagging criteria. The remaining 4 have complicated spectra where the ^{12}CO features, if present, are difficult to identify even under careful manual inspection. Figure 5.18 shows an example of one of the spectra which did not produce a high quality estimate. On the other hand, Figure 5.19 shows one of the spectra with a large number of lines which did produce an accurate estimate. As such, it seems the presence and SNR of ^{12}CO features is the dominant factor in the success of the redshift routine.

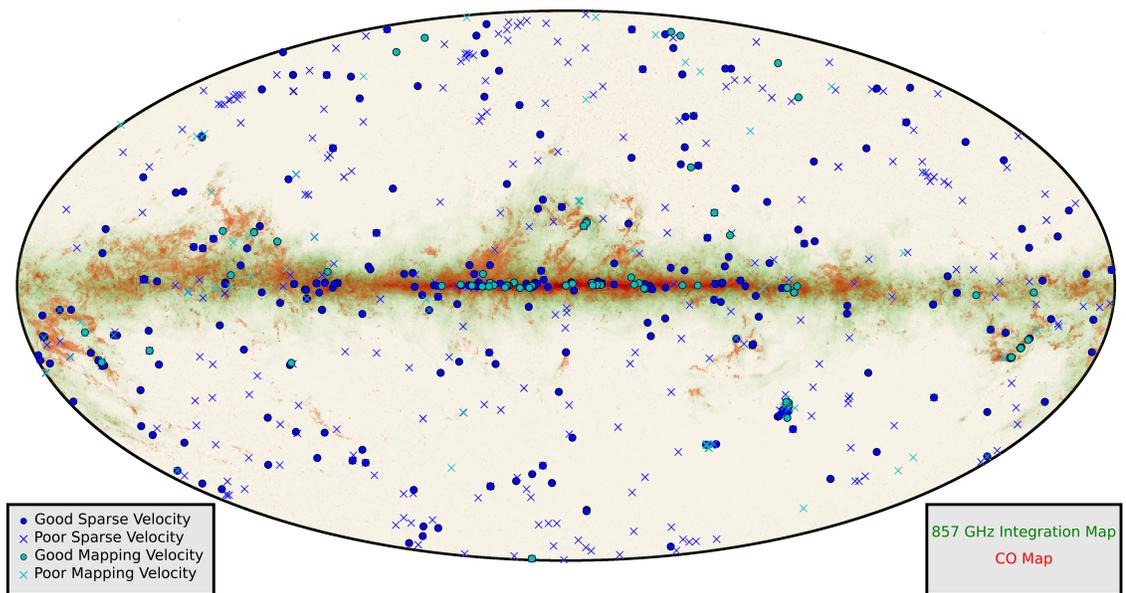


Figure 5.17: Full-sky composite image derived from the Planck 857 GHz integration map (green) and the Planck ^{12}CO (J=2-1) emission map (red) [92], showing SPIRE FTS observations coordinates. Markers indicating the position of SPIRE FTS sparse (blue) and mapping (cyan) observations on-sky. Sparse observations receiving a velocity estimate meeting the FF quality criteria are indicated with circles, and crosses otherwise. Mapping observations with $\geq 50\%$ of pixels receiving a velocity estimate meeting the FF quality criteria are indicated with circles, and crosses otherwise.

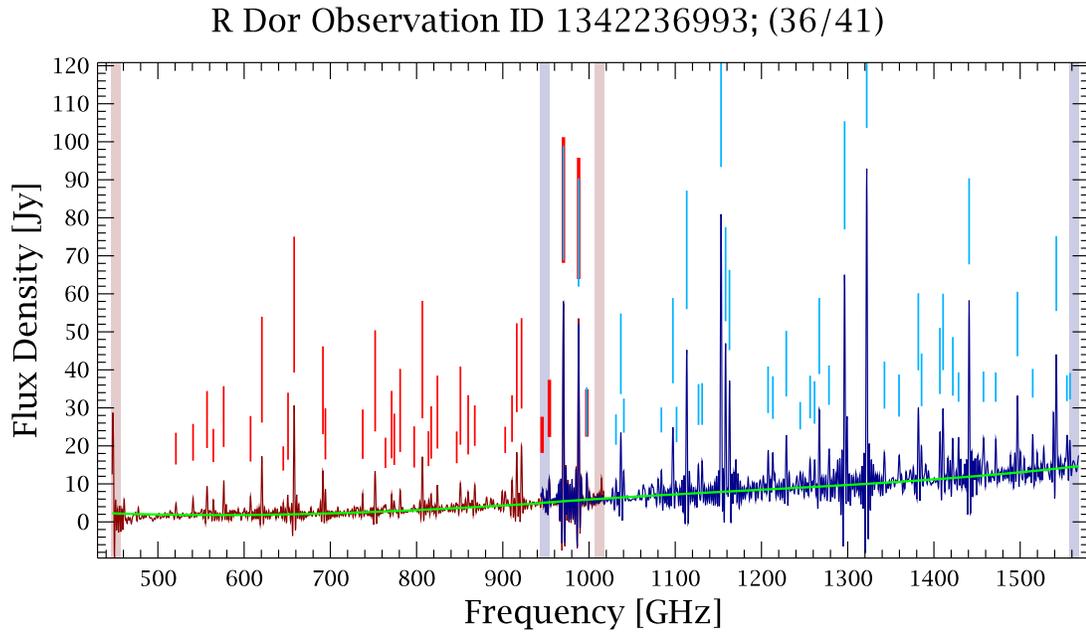


Figure 5.18: An example spectrum with more than 40 detected spectral features which did not receive an accurate velocity estimate using the redshift routine. Even with manual inspection, the ^{12}CO ladder, if present, is difficult to identify.

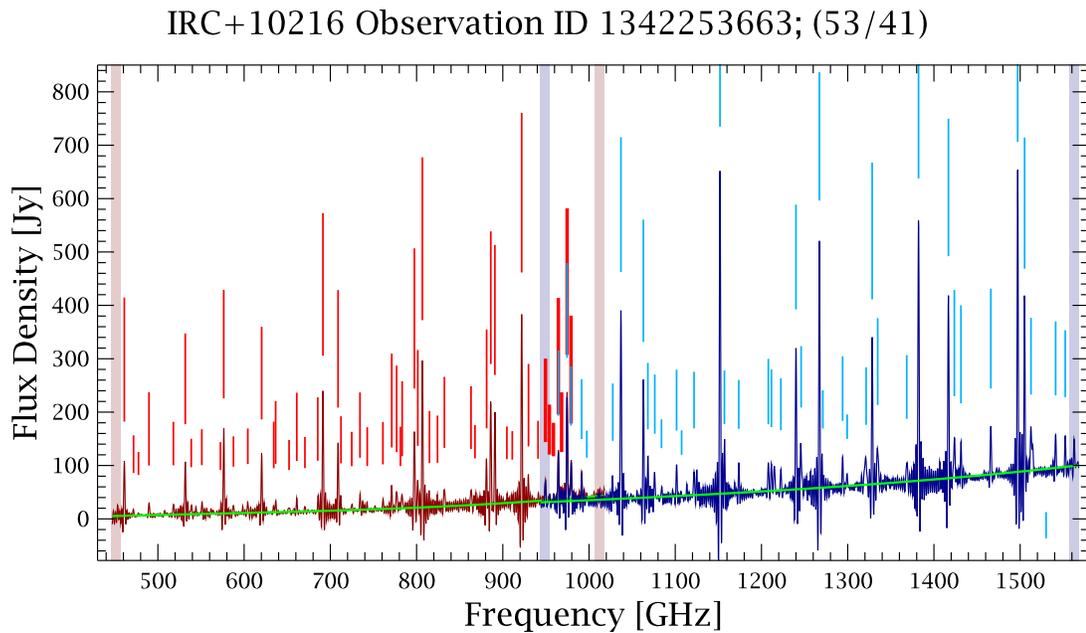


Figure 5.19: An example spectrum with more than 40 detected spectral features which did receive an accurate velocity estimate.

5.5.0.4 Mapping Observations

Efforts to extend the FF catalogue and thus redshift routine to SPIRE mapping observations are currently in progress. One of the principle advantages of the redshift routine is its modular implementation. All that is required for valid execution of the redshift routine is a line list and associated SNR (i.e., it does not require raw spectral data). The routine takes frequency error as another input, but this information is only used when calculating velocity error for estimates based on a single feature such as [NII]. As such the application of the redshift routine to mapping observations, or indeed any other spectral line catalogue is simple. The challenge comes from combining lines from SSW and SLW maps which have different footprints on-sky with pixels of different spatial extent. This effect, which is due to the different nominal beam sizes between detector arrays, can be seen in the integrated intensity maps (first two rows of the first column) of Figure 5.6.

The redshift routine operates best on line lists with more ^{12}CO features, and thus keeping the SLW and SSW arrays separate results in a maximum N value of 5 which can in principle occur by chance more easily than a N value of 10 for a combined spectrum. Thus, the current approach is to combine the lines from SLW and SSW pixels. To do so, the SLW pixels are projected onto the SSW grid. This projection is accomplished by taking the SLW pixel with the most overlap on a given SSW pixel and combining the two complementary line lists. This is a simple approach, though it is still somewhat unsatisfying in that line lists centered on different spatial locations are being combined. In effect, this can smooth the velocity gradient of the resulting maps. It is worth noting however, that the spectral maps themselves are generated from a convolution process described in Chapter 4, and this spatial smoothing is already somewhat inherent in the maps themselves. The only rigorous solution would be to regenerate the SLW maps to match the SSW projection, and refit all spectra within the SLW map. However, such an approach is impractical, and physically unjustifiable given the different SLW and SSW beam sizes. As such, this alternative approach is not being considered.

The distribution of preliminary high quality radial velocity estimates for mapping observations is shown in Figure 5.16. The on-sky location of these observations are shown in Figure 5.17. This distribution consists of 22,897 velocity estimates, with each estimate representing the radial velocity attributed to a single pixel within the spectral maps. $\sim 82\%$ of velocity estimates returned by the redshift routine for mapping observations that did not return a “NaN” value were considered high quality. The standard for high quality estimates is modified for mapping observations. Projecting the SLW results onto the SSW grid, in some cases, results in pixels around the perimeter of the combined maps containing lines from only one array. For these pixels, the modified standard for a high quality estimate is $N > 3$. This new standard, although not as reliable as that for sparse observations, is still validated by the test conducted on real and simulated data as outlined in Section 5.4

Figure 5.20 shows an example of the most comprehensive mapping postcard. The first two rows consist of SLW (top), and SSW (middle) maps. The first column shows integrated intensity maps from which the complications of spatial projection become clear. These maps have purple and green boxes around the brightest and darkest pixels respectively, from a region with the same overlapping spatial coverage between maps. The second column shows the same gridding, but each pixel value represents the number of lines found by the FF for that pixel. In these maps a red box indicates the pixel with the most lines. The third column shows the resulting velocity maps. Each pixel contains an annotation indicating the flagging parameter associated with the estimate. The bottom row shows the combined maps, with the spectral plot showing the spectra from the highlighted pixels in the previous separate maps. It is worth noting the good consistency between the three velocity maps after accounting for the difference in scale between colorbars.

An example of the public version of the mapping postcards is shown in Figure 5.21. The public FF velocity estimates for sparse observations are based on SLW/SSW combined data, so individual SLW and SSW velocity maps, as shown in Figure 5.20, are inconsistent with this approach. Sample postcards using the public format for a selection of mapping

Object: HD37041, Obsid: 1342192173

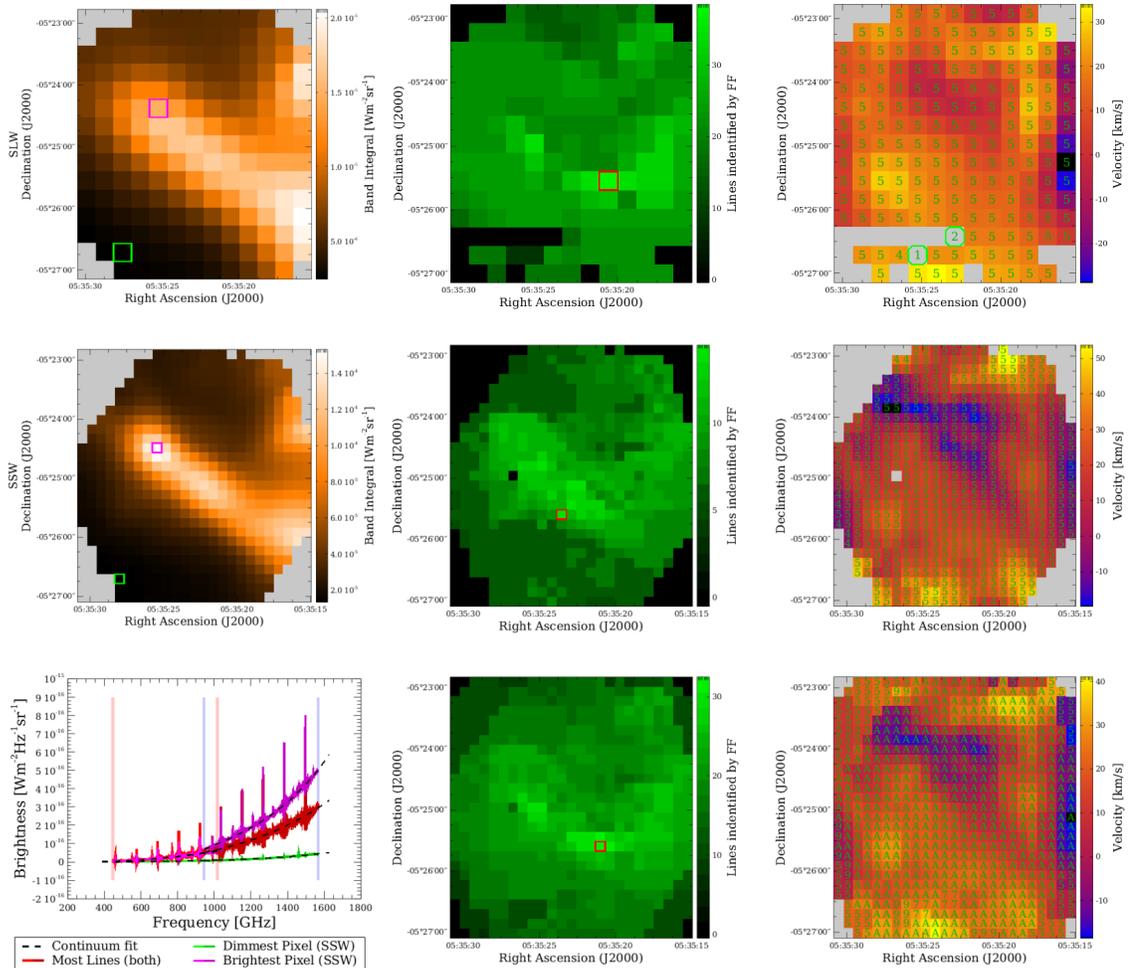


Figure 5.20: An overview of the FF velocity routine for mapping observations presenting a comprehensive dataset for the observation under study. The top and center rows show results specific to SLW and SSW, respectively. The bottom row shows combined SLW/SSW results. With the exception of the bottom left panel which shows combined spectral plots, the first column shows integrated intensity maps. The center column indicates the number of features found within each pixels, and the column on the right presents the calculated velocity maps. Magenta, and green squares within the integrated intensity maps indicate the brightest and faintest pixels from overlapping regions respectively. Red squares within the center column indicate the pixels with the most number of detected lines. Pixels indicated with squares within the SLW/SSW separated maps have their corresponding spectra plotted in the bottom left panel. Spectra are color coded to match their corresponding square's color.

Object: HD37041, Obsid: 1342192173

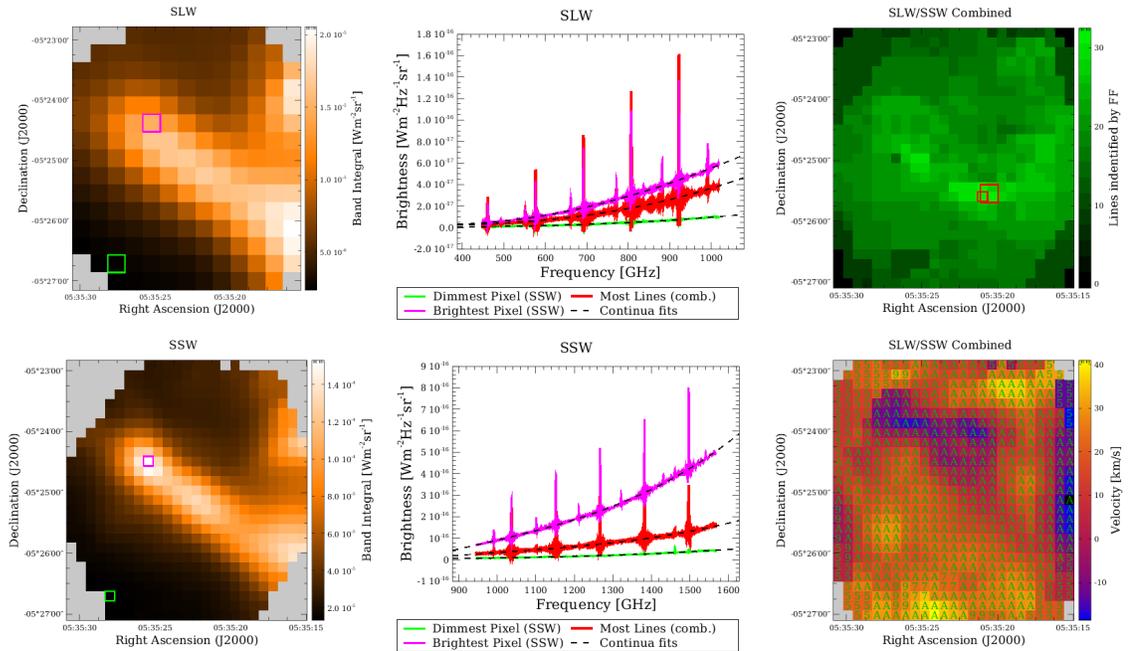


Figure 5.21: The public FF mapping catalogue postcards. The first column shows integrated flux maps for SLW and SSW spectral cubes, with green and magenta squared indicating the faintest and brightest pixels, respectively, from overlapping SLW/SSW footprints. Center column shows spectra from spectral cube pixels highlighted by squares from other plots. The top right figure indicates the total number of detected feature in each pixel with SLW and SSW features combined into one SSW equivalent grid. Red squares indicate the pixels with the most identified spectral features (SLW and SSW kept separate). Finally the bottom right panel shows the resulting velocity map after combining detected spectral features from both SLW and SSW. Annotations indicate the number of features used to generate the velocity estimate as previously described.

observations are given in appendix B.

5.6 Summary

Radial velocity is a useful parameter for probing the dynamics of the ISM, and also for estimating the distance to extragalactic sources. Since the FF already determines the positions of spectral features, and the process of determining source radial velocity involves matching shifted frequencies to rest frequencies, I decided to develop an automated radial velocity estimating routine. Due to the abundance of ^{12}CO in the ISM and the characteristic brightness or its rotational emission in the FIR, ^{12}CO was selected as the primary candidate

for a dedicated search. The routine attempts to find ^{12}CO rotational transitions by searching for the characteristic difference pattern expected for ^{12}CO within the features detected by earlier steps in the FF. The routine returns a set of flagging parameters which indicate the accuracy of the estimate. When no ^{12}CO emission is detected, the routine attempts a simple search for [NII] emission. As a last resort, spectral windows about the rest frequencies of ^{12}CO emission are defined, and the highest SNR feature within each window is assumed to be an ^{12}CO emission. A radial velocity estimate is then calculated as the average velocity resulting from these features, though these estimates were never used in any FF products.

Tests on both real and simulated data show the redshift routine is successful at performing its task. The flagging parameter is a robust indicator of velocity estimate quality with $\sim 90\%$ of high quality estimates having an absolute residual < 20 km/s while estimates flagged as high quality contain $\sim 82\%$ of all estimates with a residual velocity < 50 km/s. The main limitations of the algorithm is the presence of high SNR ^{12}CO features, and the routine without further modification, has no applicability to spectra with source velocities $> 14,000$ km/s. At the same time, the routine benefits from a simple modular implementation that is applicable to virtually any line list containing ^{12}CO , and can be easily modified to search for the rotational emission of other linear molecules. The routine is also resilient to line lists containing large numbers of detected features in excess of the desired ^{12}CO features.

This redshift routine has been implemented as the primary radial velocity estimating routine of the FF, and currently has contributed 624 estimates to the publicly available sparse feature catalogue. Work has been done to implement this routine on mapping observations, and has generated 22,897 additional estimates. The mapping velocity estimates are presented in the form of easy to inspect postcards, a sample of which are included in appendix B.

Chapter 6

Neutral Carbon Detection

6.1 Introduction

The Interstellar Medium (ISM) is segregated into different components which are defined in terms of hydrogen nucleus density n , temperature T , and ionization state [93]. To a large extent these components are mediated by embedded or external stellar objects, and most notably hot O/B type stars which produce sufficient quantities of dissociating and ionizing photons [94]. Other shock processes from supernova can heat the interstellar medium to temperatures in excess of millions of Kelvin, and is often referred to as the Hot Ionized Medium (HIM) [93].

Figure 6.1 illustrates the typical structure of the ISM which is illuminated from the left by hot stars or other Interstellar Radiation Fields (ISRF). The region most near the source consists of ionized hydrogen resulting from an abundance of Ultraviolet (UV) photons with energy in excess of 13.6 eV per photon, the ionization potential of hydrogen. This region is referred to as the Warm Ionized Medium (WIM) which expresses a temperature and hydrogen density of $T \sim 8,000$ K, and $n \sim 0.3 \text{ cm}^{-3}$, respectively [93]. Tracers for this region include the H-alpha line and some ionized lines like [NII]. This is followed by a transition zone to molecular hydrogen once the high energy UV photons have been sufficiently attenuated. The transition leads into the Warm Neutral Medium (WNM) which expresses a similar temperature and hydrogen density as the WIM, with the key distinction being the abundance of neutral hydrogen, oxygen, and nitrogen, with some ionized carbon [93]. The WNM is primarily probed by the 21 cm neutral hydrogen electron spin transition emission

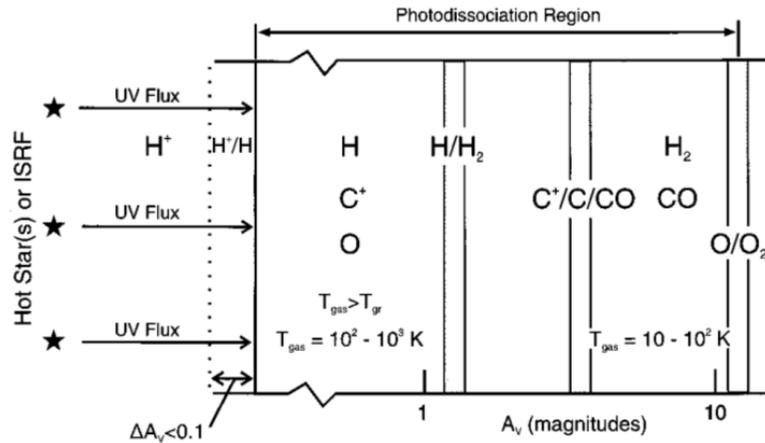


Figure 6.1: A model of the different ISM components described in Section 6.1. The visual extinction of light, A_v , is a measure of radiation attenuation as a result of interactions with dust. The ratio of gas and dust is assumed to be constant, so A_v is both a measure of depth and gas density in molecular clouds. A 2.5 fold decrease in visual flux is obtained for $A_v = 1$, and is typically measured at 555 nm. Figure courtesy of Hollenbach et al., (1999) [93].

line. Deeper into the molecular cloud is the Cold Neutral Medium (CNM) with a temperature and hydrogen density of $T \sim 100$ K, and $n \sim 30 \text{ cm}^{-3}$, respectively [93]. This region is marked by an increased abundance of H_2 . Tracers for this region are typically optical and UV absorption lines resulting from the illumination of bright stars. Finally, this is followed by the cold molecular cloud region of the ISM, which is often folded into the CNM. Cold molecular clouds are marked by a temperature and hydrogen density of $T \sim 20$ K, and $n > 10^3 \text{ cm}^{-3}$, respectively [93]. This component expresses an abundance of H_2 and CO , and is marked by a sharp transition region where ionized carbon transitions to neutral carbon which transitions to CO or CH [93]. The main tracer for this region is rotational CO emission.

The region covered by the WIM to the CNM is collectively called the photodissociation region due to the presence of far-UV photons with energies ranging between < 13.6 to ~ 6 eV. These low energy UV photons are of sufficient energy to break some molecular bonds, but cannot ionize hydrogen. This region is of particular interest in the study of star

formation as it covers a range of associated ISM components. The thin neutral carbon shell is of particular interest in some studies, as it marks the boundary between cold and warm components of molecular clouds.

Neutral carbon has an orbital configuration $1s^2 2s^2 2p^2$ [95]. As discussed in Section 3.3.1.4, once mutual electron interactions and spin-orbit coupling is accounted for, the ground state configuration of neutral carbon is split into 5 fine structure lines. Most relevant to this thesis are the ground state configuration 3P_0 , and the excited states 3P_1 and 3P_2 . The states 3P_1 and 3P_2 correspond to the energy levels 2.04 meV and 5.38 meV [96], respectively, with the transitions $^3P_1-^3P_0$ and $^3P_2-^3P_1$ resulting in emissions at 492.16 GHz and 809.34 GHz, respectively. Both these emissions reside within the SPIRE SLW band.

The primary mode of neutral carbon generation is the photodissociation of carbon containing molecules, most significantly CO and CH. Most relevant to this thesis is the production channel through CO so I will focus on this molecule. The narrow band of neutral carbon in Figure 6.1 can be explained in that CO has a binding energy of ~ 11.100 eV [97], while the first ionization of carbon requires ~ 11.260 eV [96]. As such, a narrow shell of neutral carbon is indicative of the restrictive 0.16 eV energy range at which it is stable. This also explains the rather large disparity in emission amplitudes between rotational ^{12}CO and neutral carbon emission commonly seen in SPIRE spectra. Figure 6.2 shows an example of the disparity between the prominent ^{12}CO (7-6) and faint carbon $^3P_2-^3P_1$ transitions for NGC 7027, which are indicated by vertical dashed lines.

Upon examining Figure 6.2 one can appreciate the difficulties of detecting the neutral carbon $^3P_2-^3P_1$ feature (hereafter referred to as [CI]), in an automated fashion. To improve detection of this feature, I developed a FF subroutine dedicated to finding this feature if present. This chapter will focus on outlining the theoretical foundation of the subroutine in Section 6.2, followed by a detailed description of the routine in Section 6.3. Results of the neutral carbon check are shown in Section 6.4, and future applications with respect to the FF are discussed in Section 6.5.

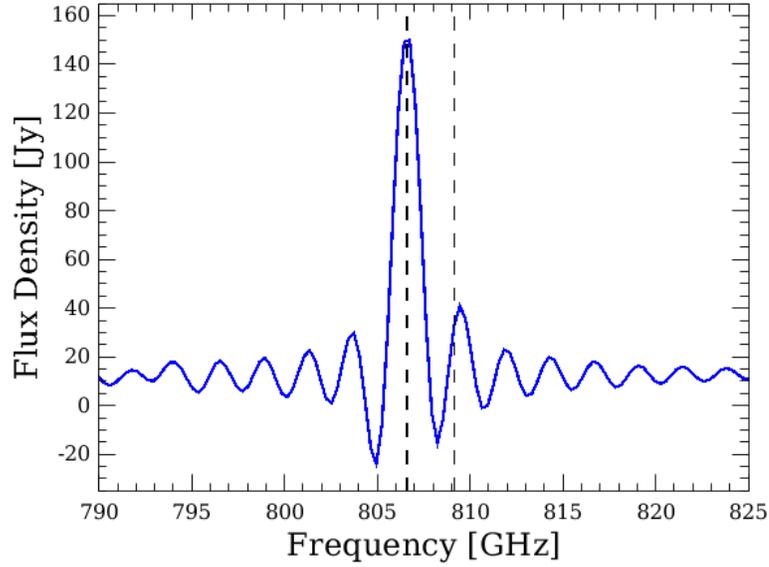


Figure 6.2: Neutral carbon region isolated from a SPIRE FTS observation of NGC 7027, a young planetary nebula. Vertical dashed lines indicated the fitted position of the ^{12}CO (7-6) rotational feature (left) and $[\text{CI}]$ $^3\text{P}_2$ - $^3\text{P}_1$ emission (right). The close spectral proximity, and low amplitude of $[\text{CI}]$ makes detection of this $[\text{CI}]$ feature difficult.

6.2 Theoretical Basis

SPIRE FTS spectra can be modelled with two component classes. First, a continuum largely derived from thermal emission from dust grains, and second, spectral lines, in both absorption and emission, derived from transitions between discrete energy states within atoms and molecules. The continuum can be approximated with a Planck curve (Equation 3.11), often requiring a correction for emissivity, whereas the emission and absorption lines typically take the form of the Instrument Line Shape (ILS). For SPIRE and most FTS instruments, the ILS is well approximated as a sinc function given by

$$\begin{aligned} \text{arg} &= \frac{f - p_1}{p_2} \\ \text{sinc}(\text{arg}) &= p_0 \frac{\sin(\text{arg})}{\text{arg}}. \end{aligned} \quad (6.1)$$

This representation is equivalent to the sinc function presented earlier in this thesis and is the exact format used by the fitter models in HIPE [66]. The model is a function of frequency f ,

and is dependent on the three variables p_0 , the amplitude at the emission frequency, p_1 , the frequency of the emission, and p_2 , the spectral distance between the peak and the first zero crossing divided by π . For unresolved features, the spectral distance between the peak and first zero crossing is equivalent to the intrinsic frequency sampling of independent spectral data points $\Delta f = 1/2L_{max}$, where L_{max} is the maximum optical path difference of the FTS. As such, in high resolution mode

$$p_2 = \frac{\Delta f}{\pi} \sim 0.377 \quad [\text{GHz} \cdot \text{rad}^{-1}]. \quad (6.2)$$

It is important to stress that Equation 6.2 applies only to unresolved spectral features. This means that from the instrument perspective, the spectral feature is infinitely narrow, much like how a spatially unresolved point source is effectively of infinitesimal spatial extent. The unresolved spectral feature is still broadened by the ILS, and modelling such a situation can be accomplished with the convolution of a delta function with the ILS. This process is demonstrated for the ^{12}CO (7-6)/[CI] pair in Figure 6.3. With SPIRE being an instrument of only moderate spectral resolution, the vast majority of spectral features are assumed to be unresolved. In the few cases of partially resolved features, it is in fact more accurate to model the spectral data as a convolution of a Gaussian with the ILS. This sinc-convolved Gaussian is incorporated into HIPE, but is useful in a minority of cases, and is thus not explored here. Furthermore, the FF is not dedicated to determining accurate line width or the integrated flux of features, only their emission frequency and SNR as a starting point for individual follow-up analysis. As such, Equation 6.1, and 6.2 are adequate for all cases in this work. To find a spectral feature meeting the criteria of the FF, the problem then becomes accurate parameterization of the detected emission frequency and the amplitude alone.

It has already been established that neutral carbon probes a different component of the ISM than ^{12}CO rotational emission, however, the two components are usually in relatively close spatial proximity and do have similar energies. The spectral resolution in

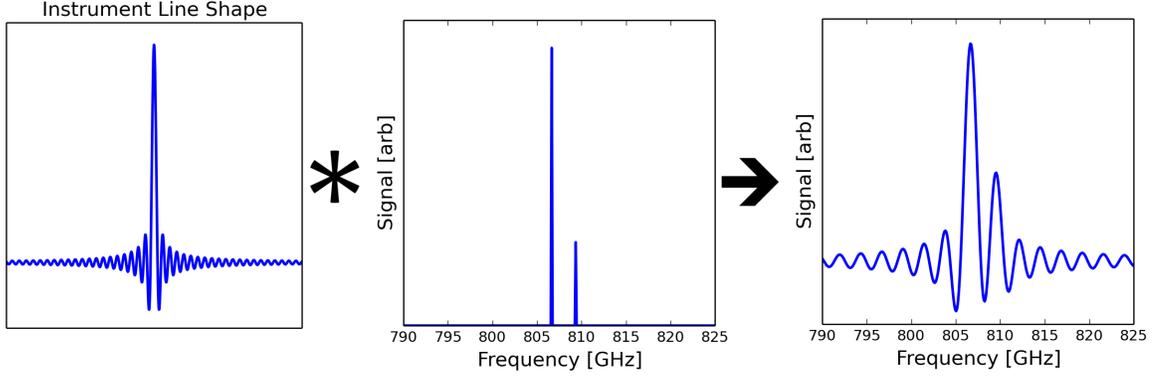


Figure 6.3: SPIRE FTS spectra with unresolved spectral features can be modelled (right) as the convolution of the SPIRE FTS ILS (left), and delta functions at the emission frequency of spectral features (center).

velocity units of the SPIRE HR observations at 809 GHz (the rest frequency of [CI]), is ~ 444.7 km/s. As such, the assumption that both ^{12}CO and neutral carbon are coupled in the same inertial frame, within resolution, is not unreasonable. This suggests that a velocity estimate based on ^{12}CO can be used to calculate a redshift that is approximately valid for the [CI] transition.

In the rest frame, the ^{12}CO (7-6) rotational transition and the [CI] fine structure transition are separated by ~ 2.69 GHz. Much like the CDA discussed in Chapter 5, the separation between adjacent transitions is altered with any non-zero radial velocity v . Under such circumstances, this separation is modified to

$$\Delta = \left(\frac{v}{c} + 1\right)^{-1} \times 2.69 \quad [\text{GHz}]. \quad (6.3)$$

With this information, a search for [CI] adjacent to ^{12}CO (7-6) rotational transitions can be constructed.

6.3 Neutral Carbon Detection Routine

6.3.1 General Overview

Neutral carbon detection in the FF is accomplished using a linear regression based approach called the Neutral Carbon Check (NCC). The ^{12}CO (7-6)/[CI] region is tested using 3 different models, and the model which produces the smallest difference between the spectral data and the model is taken as the most accurate representation for the region under study. Key steps of the routine are presented in Figure 6.4. The radial velocity estimate calculated from the redshift routine described in Chapter 5 is used to isolate the ^{12}CO (7-6)/[CI] region. This spectral sample is then compared against 3 test models. The first model attempts to fit only the most prominent feature within the sample using a continuum provided by the FF, and one sinc function, and is indicative of there being only ^{12}CO (7-6) or only [CI] in the sample. The next two models are similar to the first with the only difference being the introduction of a second sinc function. The position of the second feature is at either a higher or lower frequency with respect to the primary feature, with the frequency difference between the primary and secondary features determined by Equation 6.3. The model which produces the smallest residual is propagated as the best representation of the spectral sample. The “residual” in this case is a single number defined as

$$res = \sum_{i=1}^N |M_i - S_i|, \quad (6.4)$$

where the sum is taken over all points in the sample, while M_i and S_i represent the values of the model and spectrum at point i , respectively.

The main challenge of this routine is determining accurate initial parameter estimates for the model. The parameter estimates of the primary feature are initially provided by or derived from the previous steps in the FF. The frequency estimate of the secondary feature is derived using the radial velocity estimate and the frequency of the primary feature. The initial amplitude estimate for the secondary feature in practice is set to be 30% of the am-

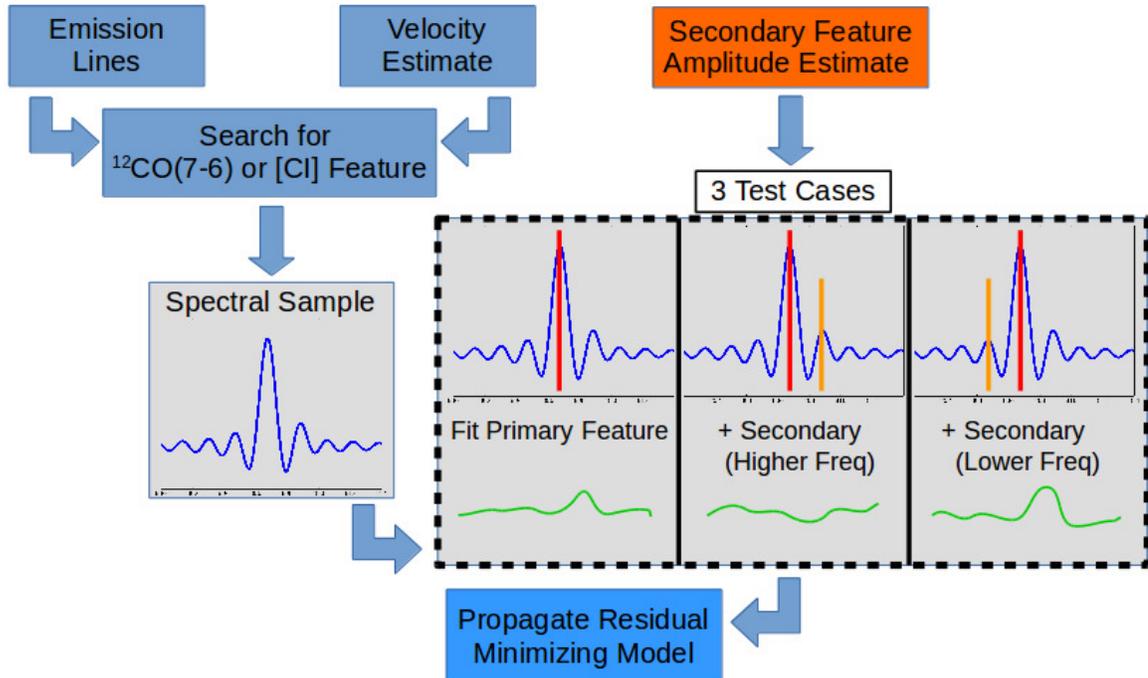


Figure 6.4: Flowchart showing the key steps of the NCC. A list of emission features, and the source radial velocity estimate are input to the routine. The expected frequency of the potential ^{12}CO (7-6) and [CI] emission is calculated based on the radial velocity estimate. If only one of these features is found within the input line list (nominally the ^{12}CO (7-6) feature), a spectral sample is extracted about the frequency of the detected feature. The spectral sample is tested with three test case models: only the primary feature (representing the detected ^{12}CO (7-6) feature), the primary feature and a secondary feature fitted to a higher frequency (nominally representing the [CI] feature), the primary feature and a secondary feature fitted to a lower frequency. The model which produces the lowest residual is considered to be the best representation of the spectral data.

plitude of the primary feature. The amplitude estimate ends up being a critical piece of information, which, if properly selected in conjunction with regression routine tolerances, can limit false detections. The value chosen was derived from careful empirical analysis of how the initial estimate affects the residual of representative test case spectra.

The remainder of this section is dedicated to describing the neutral carbon detection routine in detail, and how the nominal amplitude estimate for the secondary feature was determined.

6.3.2 Detailed Outline

Neutral carbon detection takes place immediately after a radial velocity estimate has been calculated. The routine takes as inputs the spectrum container (a HIPE object containing the spectrum and all relevant information about the spectrum), the current line list and associated SNR values, and the radial velocity estimate, all derived from previous steps of the FF. Optional inputs include the SLW continuum parameters which were determined earlier in the FF routine, and a “window” variable which determines the frequency range of the spectral sample. In practice, the previously calculated continuum is used as an input, while the window parameter is set to 30 GHz by default.

Since the routine is concerned with emission features, all absorption features, identified by the FF as having a negative SNR, are removed. At this point, the routine will only proceed if a high quality velocity estimate is present (“high quality estimate” as outlined in Section 5.3). Using the radial velocity estimate, the shifted frequencies of both the ^{12}CO (7-6) and [CI] features are calculated. A search begins which extract any features within 0.8 GHz of the calculated frequencies of ^{12}CO (7-6) and [CI]. This threshold was selected as being sufficiently restricted to discriminate between the features of interest, while still allowing for a generous velocity uncertainty ($\sim \pm 300$ km/s). If one and only one of the two features is identified a search begins for the other. In most cases, it is the ^{12}CO (7-6) feature which is most prominent and is found in the earlier stages of the FF with the [CI] feature

being missed. However, depending on which part of the ISM is being observed, this can be reversed with the [CI] feature being more prominent and the ^{12}CO (7-6) transition being missed, or even removed as a spurious detection by the FF. The neutral carbon check in fact attempts to find either missing feature, and the title of this routine is perhaps a misnomer. However, in the majority of cases it is the [CI] feature which is neglected thus the title remains. For simplicity, the remainder of this chapter will be presented as if ^{12}CO (7-6) was the detected primary feature, with [CI] being the missed secondary feature, though it should be kept in mind that the routine works if the situation is reversed, and in fact does check both scenarios. There is no attempt to refine detection when both features are already identified in previous FF steps. The neutral carbon routine operates on an isolated spectral sample, and as previously mentioned, the sinc function distributes energy over a wide range of frequencies. As such, strong spectral features outside of the isolated region can have an effect on the spectral sample. Such features are not accounted for within the test case models in this routine, thus a model for the entire spectrum is generally more accurate. The parameters for detected ^{12}CO (7-6) and [CI] in previous FF steps are determined using such full spectrum models, thus no attempt is made to improve them with tests proposed by the neutral carbon check.

After ^{12}CO (7-6) is identified within the emission line list, the routine may attempt to fit the SLW spectrum continuum if no continuum parameters are provided as optional inputs. This fitting process follows the prescription of the default continuum extraction routine offered within HIPE where a 3rd order polynomial is fit to the spectral data without any strong feature masking. This option is not recommended as the available FF derived continuum results from a simultaneous fit of both significant spectral features and the continuum.

A spectral sample is then isolated, centered on the shifted ^{12}CO (7-6) frequency, and extended on either side by a frequency given by the window parameter, nominally set to 30 GHz, resulting in a 60 GHz spectral sample. For the first pass, a fitter object is constructed with a tolerance that is intentionally increased significantly in order to obtain a

certain mode of operation. This will be discussed further in Section 6.3.3.

The spectral sample is then modelled in 3 different ways. The first model consists of a continuum, the primary feature, and any extra lines that were identified by the FF which reside within the spectral sample. Parameter estimates for the primary feature are determined by examining the continuum subtracted flux with the peak flux becoming the amplitude estimate, and corresponding frequency becoming the line center estimate. These parameters are allowed to vary during the fitting process. For the extra features, line center is determined by FF detection frequency, with amplitude given by the product of SNR and corresponding spectral error at the detection frequency. This amplitude estimate is not very accurate since SNR is calculated in a more complicated way, as such, this parameter is allowed to vary during the fitting process. However, the frequency center of each extra line is held fixed. All features within this model have their sinc width held fixed for the reasons outlined in Section 6.2. A global fit is then initiated with the resulting residual and the fitted model is recorded for later comparison.

The second and third models are similar to the first. The distinction comes from the introduction of an additional sinc function which represents the potential missing [CI] feature. This secondary feature has an estimated frequency center at $\pm\Delta$ (Equation 6.3), relative to the primary feature, with an initial amplitude estimate equal to 30% the amplitude of the primary feature. In these models, the line center and amplitude of the primary feature are derived from the fitted results of the first model. These parameters are allowed to vary during the fitting process. Again, the continuum and extra features are added, and a global fit is initiated with the resultant residual and fitted model recorded.

The model with the lowest residual is then selected as the best representation of the spectral data. If this happens to be the second or third model, the fitted ^{12}CO (7-6) and [CI] features are checked to insure they still reside within 0.8 GHz of their expected frequency based on the radial velocity estimate. This check is implemented to account for the frequency drift that can occur in fitting routines with a large parameter space. If both lines

meet this threshold, a second pass fit is initiated. The second pass takes the parameters from the minimum residual model and transfers them into another fitting object which has its tolerances restored to the default values. A global fit is then performed, and the resulting ^{12}CO (7-6) and [CI] line positions are checked for drift. If both features pass this final test, their corresponding SNRs are calculated using the default FF SNR function. The resulting model parameters and SNR for the ^{12}CO (7-6) and [CI] features are then passed back to the main script to be incorporated into the FF output products. FF products contain a meta-data entry which indicates if the NCC found a new feature.

6.3.3 Secondary Feature Amplitude Estimate

The two critical pieces of information required for successful operation of the neutral carbon check are the source radial velocity and initial amplitude estimate of the secondary feature. Chapter 5 gives a detailed outline of how source velocity is estimated from rotational ^{12}CO emission. What follows here is the justification for the standardized secondary feature amplitude estimate.

The initial parameter estimates input to a fitting algorithm have a significant impact on the resulting fit when attempting to fit data to a predefined function with multiple variables. The neutral carbon check exploits this property in order to help prevent false detections. In order to optimize the input parameter estimates, three test cases were studied to see how initial estimates of the secondary feature amplitude effect the resulting residual.

The three test cases under study are,

1. Isolated ^{12}CO rotational emission feature with no [CI] feature.
2. Prominent ^{12}CO (7-6) rotational emission with faint [CI] feature.
3. Both ^{12}CO (7-6) and [CI] emissions expressed as prominent features.

The spectra used for these test cases are shown in Figure 6.5. The isolated region shown in the third row, first column, expresses an asymmetry which is most easily seen by looking

at the first negative side-lobes. This might lead one to believe that [CI] is present to some extent, however, this asymmetry is a well known feature of the SPIRE empirical line shape [79]. The appearance is further perturbed by a relatively high noise floor. To avoid any potential for [CI] contamination in the region under study, an equivalent bandwidth was isolated about the ^{12}CO (6-5) feature where no [CI] exists. The sample region in the third row, second column, comes from NGC 7027 which is known to possess a faint [CI] feature [78]. The third case is represented by the isolated region in the third row, third column, in which both ^{12}CO (7-6) and [CI] emissions are clearly evident.

Recall that the criteria for [CI] detection is when the model containing a secondary sinc function positioned at the expected frequency of the [CI] emission reduces the residual compared to fitting the primary feature alone. To begin the investigation of how the secondary feature amplitude estimate affects the residual, the primary feature alone in all three cases was fit along with the continuum. The result of this initial step is shown in the first row of Figure 6.6. Blue curves represent the spectral data with the orange curves showing the resulting fit. The green curves indicate the difference between data and model, the absolute sum of which become the residuals. The residual for each case was calculated and is plotted in the second row of the figure as a red dashed line. Though this line extends over the whole x-axis, it is not truly a function but a constant value distributed across the x-axis for reference. The fit process was then repeated with the addition of a sinc function for the potential [CI] feature. The initial amplitude estimate of this secondary feature was varied, but the initial line center estimate was kept constant between trials ($\pm\Delta$ GHz with respect to the primary feature; Equation 6.3). The resulting residual from varying the secondary feature amplitude estimate for each case is shown in the second row of Figure 6.6. The green curve shows the residual when placing the secondary feature to the right of the main feature, while the blue curve shows the same thing when placed to the left. The x-axis shows the initial amplitude estimate of the secondary feature presented in terms of fractional amplitude of the primary feature.

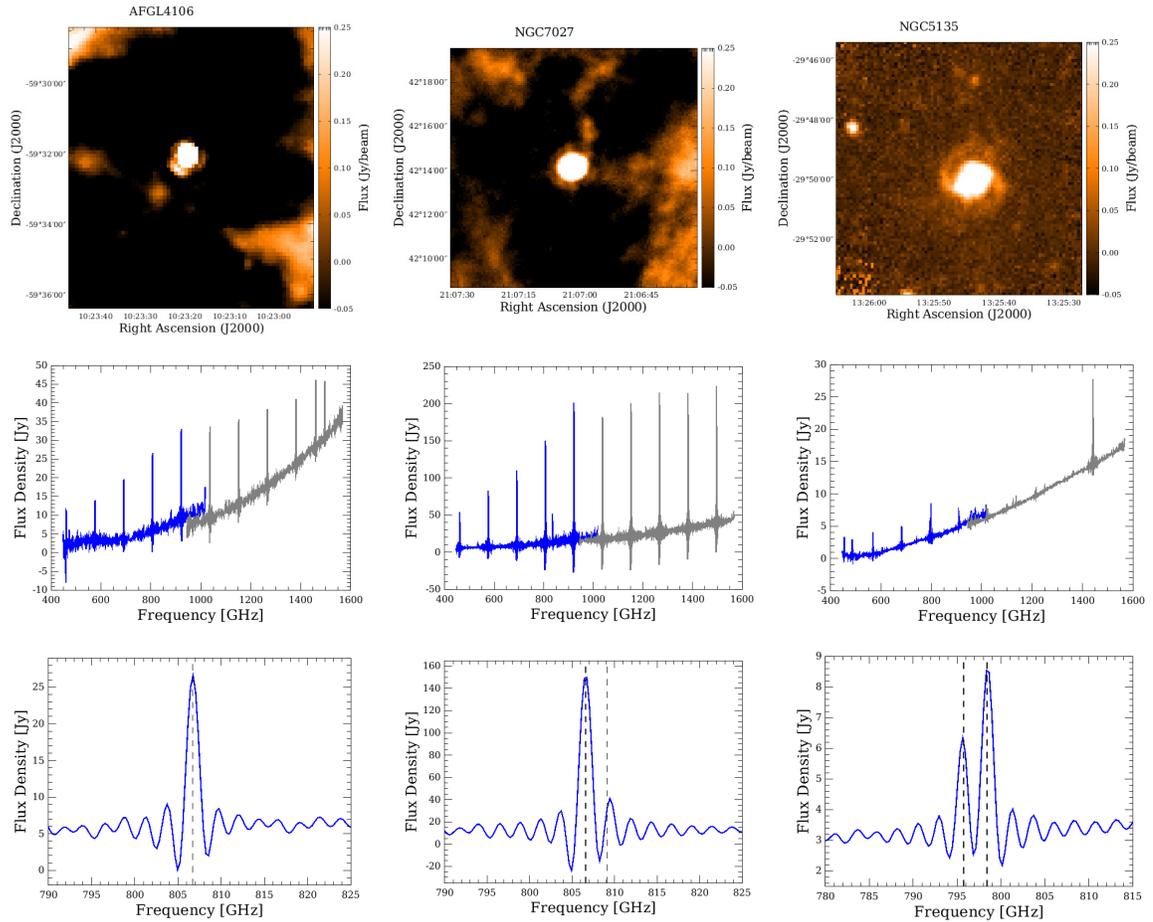


Figure 6.5: The top row shows short wavelength (250 μm) SPIRE photometer maps of carbon rich sources. The second row shows their corresponding full band SPIRE spectra. The third row is a close up view of the $^{12}\text{CO}(7-6)$ region with present features indicated by dashed lines. The first column contains the source AFGL 4106, a post-red supergiant binary, that exhibits no detectable neutral carbon. The young planetary nebula NGC 7027 in the second column and, the starburst galaxy NGC 5135 in the third column contain faint and strong neutral carbon lines respectively.

Focusing on the no [CI] case, it can be seen that fitting the secondary feature to the left or the right does not result in drastically different residual curves. This is to be expected from the high degree of symmetry of the spectral sample about the primary feature. For initial amplitude estimates less than $\sim 10\%$ the residual is reduced compared to fitting the primary feature alone. Based on the selection criteria for [CI], a secondary feature amplitude estimate greater than 10% the amplitude of the primary feature should be used to prevent such false detections. For the faint [CI] case, there is a distinct difference between placing the secondary feature to the left or the right of the primary. The critical value at which the amplitude estimate switches from producing a smaller residual to a larger residual is $\sim 20\%$ and $\sim 35\%$ when placing the secondary feature to the left and right, respectively. Both of these thresholds are greater than the no [CI] case, but in order to ensure the secondary feature is fit to the correct side of the primary, an initial amplitude estimate between these two critical values should be adopted. As such the initial amplitude estimate of 30% the amplitude of the primary feature was adopted. This value is indicated in the bottom row of Figure 6.6 with a yellow diamond. For completeness, we can see that this standard also produces a lower residual in the case of strong [CI] emission (third column).

It is important that the fitter tolerance be increased when executing this check. With default settings, the residual will be reduced for all three cases regardless of amplitude estimate as the routine will simply fit noise if that is all that is present.

The results of applying this standard and running the neutral carbon check on the test case spectral samples are shown in Figure 6.7. Fits to only the primary feature are shown as insets for reference. As can be seen, no additional feature was included in the no [CI] case while still detecting the faint [CI] feature.

6.4 Neutral Carbon Detection Results

The NCC was applied to all 818 HR SPIRE FTS sparse observations processed by the FF with corresponding continuum parameters. Though no attempt was made to improve

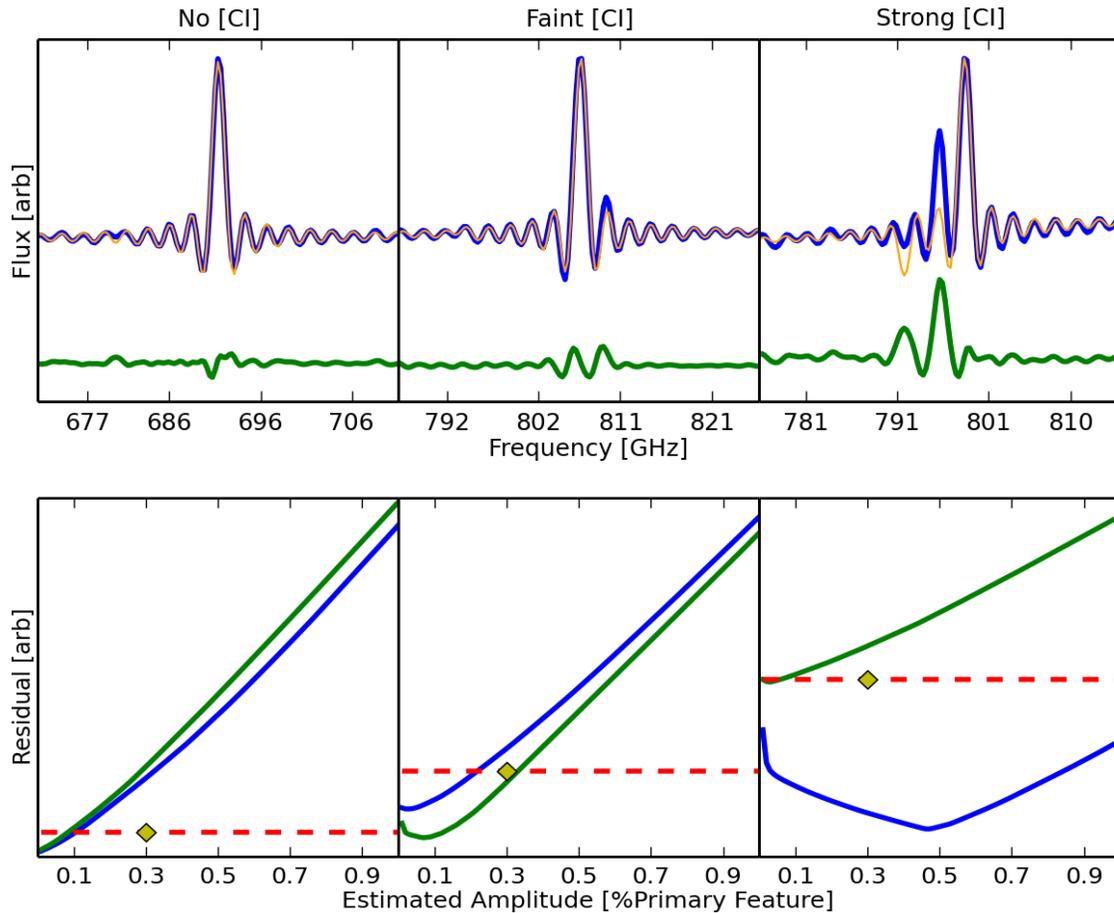


Figure 6.6: The top row illustrates three test case spectra (dark blue), and the models obtained from fitting the primary features (light orange) alone, with green residuals. The bottom row shows the resulting residuals using different initial amplitude estimates for the secondary features. Red dashed lines show the residual for the corresponding primary feature fit in the top row. Blue/green curves show the resulting residuals after including secondary sinc features with line center frequencies less/greater than the primary features. The yellow diamonds indicate the amplitude estimates used in the NCC (30% primary feature amplitude).

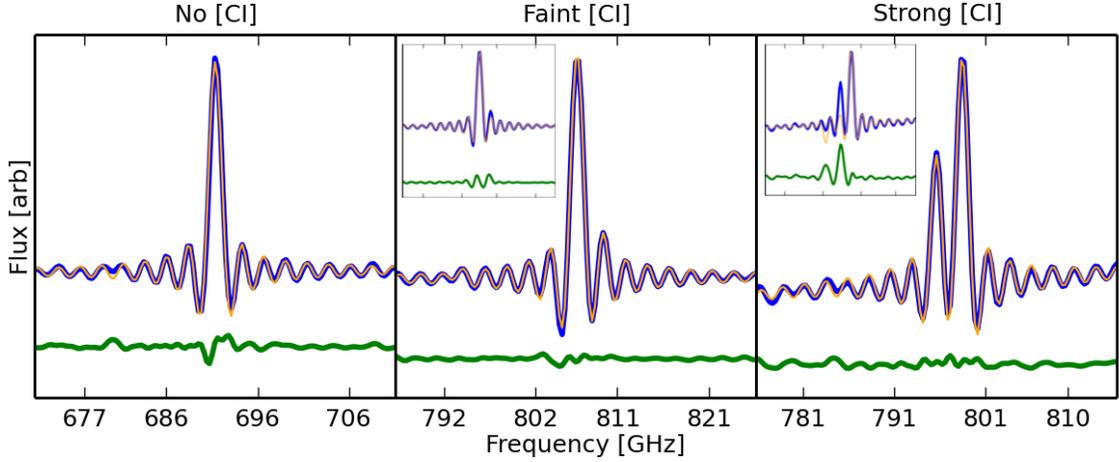


Figure 6.7: Results obtained using the NCC for the no [CI] case (left), weak [CI] case, (centre), and strong [CI] case (right). For comparison, insets show the fit to the primary feature only with the no [CI] case unchanged as expected.

Table 6.1: Bulk results of NCC on SPIRE FTS sparse dataset.

[CI] Detected	Redundant	132
	Unique	54
No [CI]	Trivial	505
	Substantive	127
Total		818

[CI] which was detected by the FF in advance of the NCC, these features were masked in the input data to see if they could be recovered by the NCC routine. Results are shown in table 6.1. For the spectra in which [CI] was detected, the results are further subdivided into redundant and unique detections. Redundant detections are those which were already found by the FF before applying the neutral carbon check. Unique detection are those for which a new [CI] feature was found. The spectra for which no [CI] was detected is also subdivided into trivial and substantive. Trivial results are those which found no [CI] simply because of poor velocity estimates, or lack of ^{12}CO (7-6) features. The substantive results are those which performed the full neutral carbon check but found no [CI] feature.

In order to assess the reliability of the neutral carbon check, 103 spectra were manually inspected to determine the presence of [CI] emission. These findings were then compared

Table 6.2: NCC results on a manually inspected dataset of 103 observations.

Criteria	Correct	Incorrect
[CI] Present	72	1
No [CI]	27	3
Total	99	4

to the results of the neutral carbon check. The outcome of this comparison is shown in table 6.2. Of the set of 103 spectra, the NCC correctly determined the presence of [CI] in 99 cases giving a success rate of $\sim 96\%$.

The single false positive comes from an observation of AFGL 4106, a source which is considered to be [CI] free. However, a noise feature induces a $\sim 15\%$ amplitude difference between the first positive side-lobes on either side of the central peak of the ^{12}CO (7-6) feature. This perturbation makes the detection understandable even if not physically correct. The 3 false negatives come from sources which are particularly noisy and have poorly defined continua in the region under investigation. NGC 6302 is one of these sources, whose SPIRE observations are known to suffer from pointing errors and contamination from background radiation. These factors have been corrected for in the set of HPDPs discussed in Chapter 4. After running the neutral carbon check on the corrected spectrum, the neutral carbon check returns the correct result for this source. Another false positive comes from the source L1448-B2. This source is semi-extended in the SPIRE beams resulting in a poorly defined continuum. The extended calibration of this observation improves the continuum somewhat, and after subjecting this data to the NCC, no [CI] is found, as expected. The final false negative observation suffers no apparent errors resulting from calibration. As such, there is no straight-forward explanation for the discrepancy. Discarding the FF continuum for this source and determining the parameters from within the NCC itself did produce the correct result. This approach is not generally recommended as strong features are not masked, and the resulting fit is for most cases sub-optimal. However, this does emphasize the importance of accurate continuum parametrization when using the NCC.

Without comprehensive manual inspection, further validation if the NCC results are difficult. There are however, two characteristics one would expect to see from good [CI] detection: an overall reduction in residual within the spectral sample, and a general preservation of the expected separation between the ^{12}CO (7-6) and [CI] features.

To test the spectral sample modelling accuracy of the NCC, the observations corresponding to the 132 redundant [CI] detections from table 6.1 were examined in more detail. For each source, a fitter object was constructed for the spectral sample used in the NCC. This fitter object was composed of sinc models and a continuum, similar to the NCC. The parameter estimates for each sinc function were derived from the results of the FF with the NCC turned off. Line centers and widths were held fixed with only the amplitude allowed to vary. The resulting residuals were then calculated and recorded for later comparison. A plot showing the ratio of the NCC derived residual and the residuals obtained with the NCC turned off (FF), is shown in Figure 6.8 (blue), with lower values indicating an improved result using the NCC. The spread in the results is quite large, but overall there is an average improvement in the residual of $\sim 24\%$. Also included in this plot is the ratio of the NCC residual and the residual obtained without fitting the [CI] feature (^{12}CO (7-6)). For the sources under study, including the NCC improves the residual by $\sim 57\%$ over models with no [CI] over the 132 observation dataset.

For the same 132 observations, the spectral difference between the ^{12}CO (7-6) and [CI] feature was calculated for both the NCC case and the NCC turned off case (FF). This difference was corrected for radial velocity, making the ideal separation 2.69 GHz for velocity coupled ^{12}CO and carbon. The resulting frequency differences are shown in Figure 6.9. The spread of frequency differences is quite a bit tighter in the NCC dataset (0.04 GHz) than the FF dataset (0.24 GHz), with the average velocity corrected separation residing ~ 0.11 GHz closer to the nominal 2.69 GHz.

Though Figures 6.8, and 6.9 are not a direct verification of the NCC operation, they do provide statistical evidence that the NCC produces models with the characteristics expected

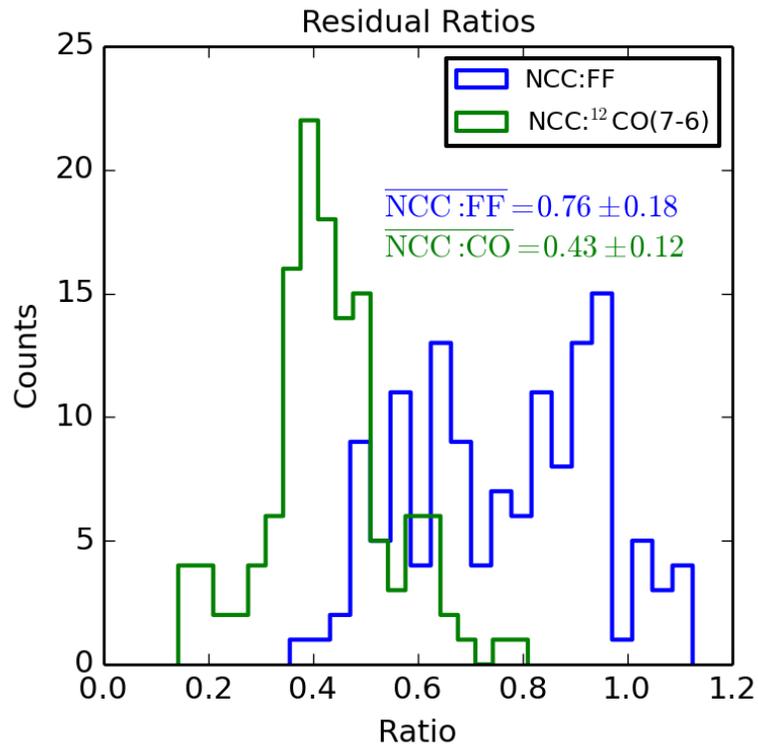


Figure 6.8: Histograms showing the residual ratios of the sampled spectrum using various models. The blue data is the ratio of the NCC model and the modelled FF results with the NCC turned off. Green data is the ratio of the NCC model and the model without attempting to fit the [CI] emission.

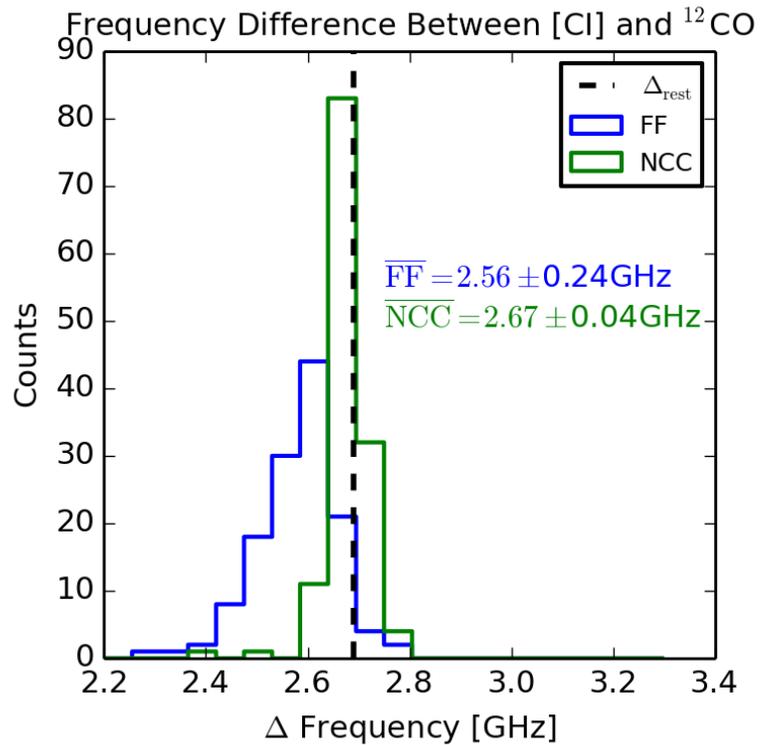


Figure 6.9: Histogram of the radial velocity corrected frequency difference between detected ^{12}CO (7-6) rotational emission and [CI] emission features. Data from the FF with the NCC turned off is blue, with green showing the results using the NCC.

for accurate [CI] detection.

The NCC was applied to SPIRE FTS mapping observations, and detected an additional 2,402 [CI] features. These features are in compliment to the 3,978 [CI] features found with the NCC turned off. Though the same level of analysis has not yet been conducted on these results, the fraction of new [CI] features to [CI] features detected with the NCC turned off is similar between the sparse (~ 0.4) and mapping ~ 0.6 datasets.

6.5 Further Applications

As discussed in the introduction of this chapter, neutral carbon has two transitions which reside within the SPIRE FTS band. The NCC is dedicated to detecting the 3P_2 - 3P_1 transition at ~ 809 GHz near the ^{12}CO (7-6) transition, but successful detection of this transition provides greater confidence in the existence of the 3P_1 - 3P_0 transition at ~ 492 GHz. As such, the results of the NCC can be used to perform a further check for the 3P_1 - 3P_0 transition. Though the 3P_1 - 3P_0 transition is not generally in close spectral proximity to any prominent features, the feature is often faint and located in one of the noisy edges of the SLW band. The 3P_1 - 3P_0 transition is roughly 3 times fainter than the 3P_2 - 3P_1 transition at 300 K, further emphasizing how faint the feature can be. These factors make detection difficult, but detection can potentially be improved with the combination of an identified 3P_2 - 3P_1 [CI] line and a dedicated search, much like the 3P_2 - 3P_1 transition.

For this task a fitter object with default tolerances is constructed in much the same way as the NCC consisting of a model including the feature of interest, the continuum, and all other features detected within the spectral sample. The initial estimate for the 3P_1 - 3P_0 line center is given by

$$f_{1-0} = f_{2-1} - 316.97 \left(\frac{v}{c} + 1 \right)^{-1} \quad [\text{GHz}], \quad (6.5)$$

where f_{2-1} is the measured frequency of the 3P_2 - 3P_1 transition. This is simultaneously

the frequency about which the spectral sample is taken. The amplitude estimate is set to one-third the amplitude of the ${}^3\text{P}_2\text{-}{}^3\text{P}_1$ transition, though the amplitude estimate is not as important in this case as it was in the NCC routine. The fitting algorithm then operates on this model, but instead of focusing on the residual, the SNR is used as the parameter of interest.

The result of this process for one observation is shown in Figure 6.10. The top panel shows the nominal FF postcard for this spectrum with dashed boxes around the neutral carbon ${}^3\text{P}_1\text{-}{}^3\text{P}_0$ and ${}^3\text{P}_2\text{-}{}^3\text{P}_1$ transitions. This spectrum was chosen for demonstration because of the unambiguous presence of the ${}^3\text{P}_2\text{-}{}^3\text{P}_1$ transition shown in the bottom right panel. The solid lines show the position of detected ${}^{12}\text{CO}$ (7-6) and [CI] features. For this spectrum the ${}^3\text{P}_1\text{-}{}^3\text{P}_0$ transition was not detected, the predicted position of which is shown by a dashed line in the bottom left panel. The resulting fit is shown in this panel with an orange curve. This example shows the clear utility of an additional check based on the NCC results.

The proposed ${}^3\text{P}_1\text{-}{}^3\text{P}_0$ check was applied to both spectra with positive and negative [CI] detection. For spectra with no [CI] detection, the line center estimate for the ${}^3\text{P}_1\text{-}{}^3\text{P}_0$ line was determined simply from the radial velocity estimate. The resulting SNRs were recorded and plotted in Figure 6.11. Features with an $\text{SNR} < 5$ were discarded in accordance with the standard FF protocol. The ${}^3\text{P}_1\text{-}{}^3\text{P}_0$ transition SNRs with corresponding ${}^3\text{P}_2\text{-}{}^3\text{P}_1$ transitions (red), are clearly of much higher SNR than those with no corresponding ${}^3\text{P}_2\text{-}{}^3\text{P}_1$ detection (blue). A dedicated ${}^3\text{P}_1\text{-}{}^3\text{P}_0$ check was not implemented in the FF, but Figures 6.10 and 6.11 make a compelling case for its inclusion in future release after further testing.

Work is currently being done to publicly release a FF catalogue for SPIRE FTS mapping observations. Naturally, the NCC will contribute to the resulting products. No significant changes are required to extend the functionality of this routine to mapping observations, other than a proportionate change in the first-pass fitter tolerance to account for the change in flux units between point-source and extended calibration of SPIRE data.

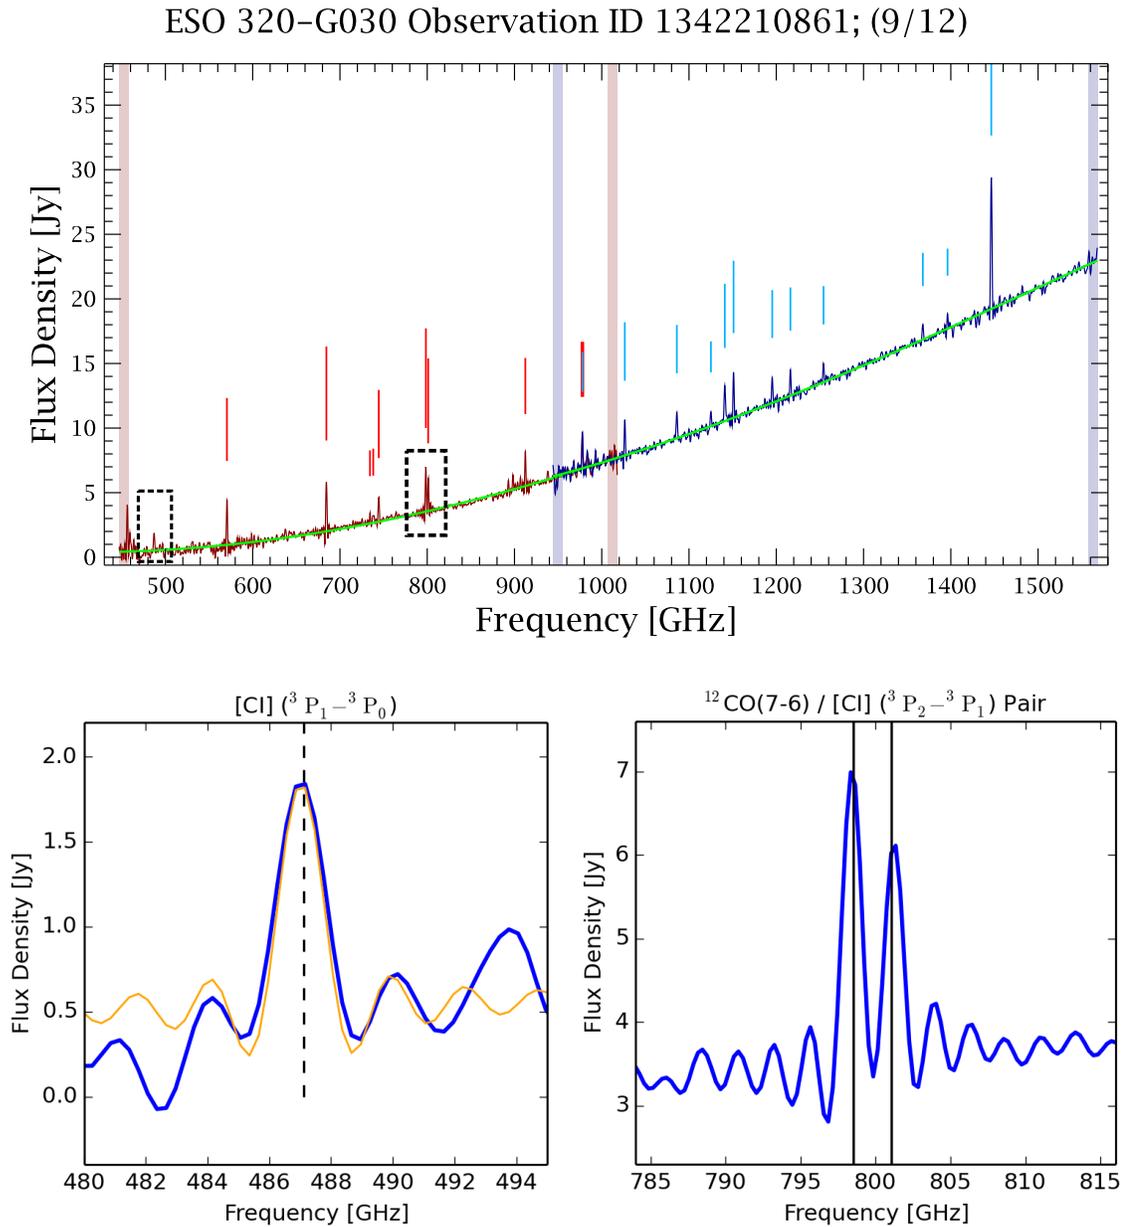


Figure 6.10: The top panel shows the full SPIRE FTS spectrum of a source containing moderately bright neutral carbon fine structure emission. The $^3P_1 - ^3P_0$ and $^3P_2 - ^3P_1$ transition regions are indicated in the full spectrum by dashed rectangles. Close-ups of these regions are shown in the bottom two panels. Bottom right panel shows the unambiguous detection of $^{12}\text{CO}(7-6)$ and $^3P_2 - ^3P_1$ [CI] emission features, indicated with solid lines. The expected position of the $^3P_1 - ^3P_0$ transition based on the [CI] detection is shown in the bottom left panel with a dashed line. This feature was not detected using the nominal FF, though the fit obtained with a $\text{SNR} = 6.5$, shown as an orange curve, along with the $^3P_2 - ^3P_1$ detection make a compelling case for the existence of the $^3P_1 - ^3P_0$ feature.

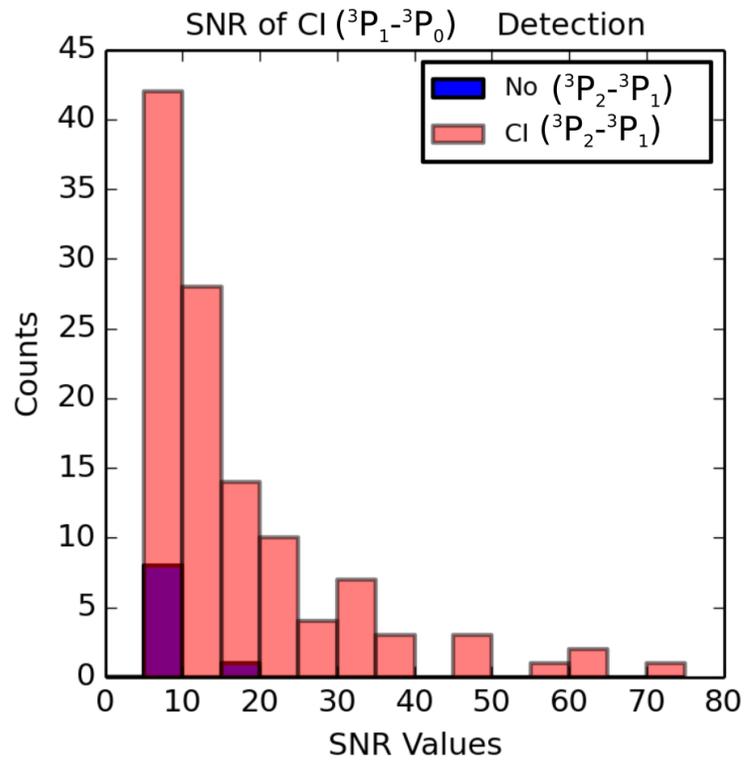


Figure 6.11: Histogram showing the SNR of fits to potential ${}^3P_1-{}^3P_0$ lines based on the previous detection (red) of the ${}^3P_2-{}^3P_1$ transition, or lack thereof (blue).

6.6 Summary

Regions with high neutral carbon mark the transition between the WNM and the CNM of the ISM. These regions are of interest in the study of star formation, and can be studied by analysing the distribution of neutral carbon fine structure emission. Detecting these transitions in SPIRE FTS data using an automated fitting routine, such as the FF, is difficult due to their typically low amplitude, and in the case of the 3P_2 - 3P_1 transition, the close spectral proximity to the often more prominent ^{12}CO (7-6) rotational emission. As such I have developed a module which performs a dedicated search for this feature.

The NCC is a regression based approach which tests three different models on an isolated spectral sample with potential [CI] emission. Two of these models contain the [CI] and ^{12}CO (7-6) emission, and one model contains only one spectral feature. The model which produces the lowest residual is considered to be the best representation of the spectral data. The sample is isolated based on the predicted [CI] emission frequency resulting from the radial velocity estimate described in Chapter 5. In the test case models containing a sinc functions for [CI] emission, the emission frequency estimate is derived from the radial velocity estimate. The initial amplitude estimate is based on a careful empirical selection which helps to reduce false detections.

The NCC has added 54 unique [CI] features to the publicly available FF sparse features catalogue and confirmed 132 cases found by the nominal FF. Tests on real SPIRE spectra show a $\sim 97\%$ accuracy on the NCC results, with the main source of error being poorly defined continua. This routine is also being extended to SPIRE mapping observations, and has detected an additional 2,402 [CI] features. Another potential extension is a dedicated search for the neutral carbon 3P_1 - 3P_0 transition based on successful detection of 3P_2 - 3P_1 emission. This extension requires more detailed study, but the evidence presented thus far makes a compelling case for its utility.

Chapter 7

Conclusion

7.1 Introduction

The FIR is one of the least explored regions of the electromagnetic spectrum in astronomy and astrophysics, yet accounts for much of the radiant energy within the Universe. This band probes astronomical sources both near and far, including molecular clouds in the ISM, the site of star birth, and redshifted galaxies. Questions about the early stages of the Universe, galaxy and star formation, and the chemical and structural evolution of the ISM can be investigated with FIR astronomy. The work presented in this thesis helps to advance this field of study by extracting information in an automated fashion from existing FIR datasets in the form of the freely available SPIRE spectra within the HSA.

This thesis began by introducing the relevant theory for understanding FTS and the structure of SPIRE spectra. This was accomplished in Chapters 2 and 3 which derive many of the important features of Fourier transforms, both continuous and discrete, and outlines the radiative processes that produce the continuum and emission components of SPIRE spectra. The SPIRE instrument was discussed in detail in Chapter 4, including observation modes, and key steps in data product generation. Motivations for the FF are presented along with an outline of the routine's operation, and validation. My unique contributions to the FF routine are introduced in Chapters 5 and 6 in the form of an automated source velocity estimating routine, and a separate subroutine which performs a dedicated search for the often elusive neutral carbon fine structure 3P_2 - 3P_1 transition.

7.2 Redshift Routine

The automated radial velocity estimating routine implemented in the FF attempts to parameterize source radial velocity by matching spectral features detected by the FF with their corresponding rest frequency, and using Equation 5.3. The features of interest are ^{12}CO rotational emission lines. ^{12}CO rotational emission is particularly bright in the FIR and the molecule is abundant in the ISM making it an ideal search candidate. These features are identified in the FF features list by searching for a frequency difference pattern characteristic of the ^{12}CO rotational emission ladder as outlined in detail in Chapter 5. The routine output is accompanied by a robust set of flagging parameters which have proven to be reliable indicators of the estimate accuracy.

Tests on both real and large synthetic datasets show that all estimates with a flag value $N > 6$ produce accurate estimates. Roughly 90% of estimates with $N > 3$ result in velocity estimates with a residual velocity less than 20 km/s. Based on the standard flagging criteria used in the public FF release, it is expected that $\sim 82\%$ of the velocity estimates with accurate radial velocity parameterization are included in the FF products (i.e., $\sim 18\%$ of accurate estimates have poor flags and have not been included). A total of 624 high quality velocity estimates have been added to the public FF catalogue for sparse-pointed observations, with another 22,897 high quality velocity estimates added to the FF mapping catalogue.

The routine has a few well understood limitations, the most limiting of which being the upper velocity limit of $\sim \pm 14,000$ km/s. At high red/blueshifts, spectral features are shifted closer to the rest frequency of neighbouring transitions than their own rest frequency, and as a result, are erroneously paired with the wrong rest frequency transition. Clearly this routine requires ^{12}CO rotational transitions for reliable velocity estimates, but performance can also be compromised when spectra contain low SNR ^{12}CO features in conjunction with many other features with equal or greater SNR. The adopted flagging convention reliably removes poor estimates which result from these limitations. Conveniently, the routine is

resilient to the total number of lines used within the routine as long as ^{12}CO rotational transitions are present with higher than average SNR. The routine benefits from a modular implementation which is applicable to any arbitrary line list, and can be easily adapted to search for rotational transitions from other linear molecules.

7.3 Neutral Carbon

Neutral carbon is abundant in the transition region between the WNM and CNM of the ISM. Outside of this zone it is highly unstable being either easily ionized or bonded to other atoms to form molecules. Thus, emission from neutral carbon such as the fine structure $^3\text{P}_2\text{-}^3\text{P}_1$ [CI] transition are often of low amplitude in SPIRE spectra. The close spectral proximity of the [CI] emission to the ^{12}CO (7-6) rotational emission, in addition to the amplitude disparity between the two features, makes [CI] detection difficult using a generic automated process. As such, I developed a dedicated neutral carbon check to search for this feature.

The NCC is a regression based approach that tests three different models on a spectral sample which potentially contains [CI] emission, and propagates the model that produces the lowest residual. The sample is isolated by employing the previously calculated source radial velocity, using the rest frequency of [CI] and the velocity estimate to predict the shifted line position. One of the important aspects of this routine is the initial parameter estimates that are used in the [CI] sinc model. The emission frequency is derived from the radial velocity estimate, and the amplitude estimate is carefully selected from empirical test results on a variety of representative test case spectra. It has been demonstrated that [CI] feature amplitude estimate, along with variations in the fitter routine tolerances, can be used to prevent false detections.

Tests show that the NCC has a $\sim 97\%$ success rate in determining the presence of [CI] in SPIRE spectra. The routine is most sensitive to the correct parameterization of the continuum, and for the few cases where the NCC failed during automated use, the

results showed improvement with a more accurate continuum input. In total, 54 unique [CI] features have been detected by the NCC and included in the FF sparse features catalogue, with another 2,402 [CI] features added to the mapping features catalogue.

Potential exists for extending the NCC to detect the neutral carbon 3P_1 - 3P_0 fine structure transition. Successful detection of [CI] provides confidence in the existence of the companion fine structure transition from the same element. More analysis is required, however, attempts to detect 3P_1 - 3P_0 features based on corresponding [CI] detection results in a systematically higher SNR than attempts to detect 3P_1 - 3P_0 features with no corresponding [CI] detection. This makes a compelling case for the efficacy of this approach.

7.4 Future Work

The work done in this thesis is in contribution to the FF project within the SPIRE instrument team. This project is nearing completion, yet there are still some elements of the redshift and NCC routines which I would like to improve. Foremost is extending the redshift method beyond its current limitations, and writing the routine in such a way so as to make it more immediately applicable to other linear molecule rotational transition ladders. Extending beyond the $\sim 14,000$ km/s limit is difficult, but can perhaps be accomplished by testing different pairing schemes between ^{12}CO rotational emission candidates and ^{12}CO rest frequencies. A successful pairing scheme can be identified when the resulting velocity estimates for the ^{12}CO candidate produces a low standard deviation.

Extension of the NCC to search for the neutral carbon 3P_1 - 3P_0 fine structure line is the main focus of future work for this routine. The underlying structure is already in place, yet more testing is required to validate the results. Further success criteria beyond a SNR limit will likely be required to ensure valid detection. Feature emission frequency is likely to drift as a result of the fitting process for spectral regions with a high noise floor, so a criteria based on maximum drift from expected emission frequency is likely required. Beyond this, manual inspection of a subset of test case spectra will be required to validate the routine's

performance similar to the [CI] validation.

In broader academic terms, I am excited to begin work on a cutting edge technology for FIR astronomy as part of my Ph.D. project. As outlined in Chapter 1, one of the main limitations to FIR astronomy is spatial resolution. At present, the FIR community is very interested in the development of spectral-spatial interferometry. Such a technique combines the spectroscopy advantages of FTS with the spatial resolution advantages of aperture synthesis commonly employed in radio astronomy. I look forward to participating in the development of this new technology as the main focus of my Ph.D. programme.

References

- [1] Stillman Drake and J Bruce Brackenridge. Galileo: Pioneer scientist. *American Journal of Physics*, 59(4):378–378, 1991.
- [2] Jeffrey O Bennett, Megan Donahue, Nicholas Schneider, and Mark Voit. *The cosmic perspective*, volume 3. Addison-Wesley, 2016.
- [3] NASA. Voyager mission status. <https://www.jpl.nasa.gov/voyager/mission/status/>, 2017.
- [4] Andrew W Blain, Ian Smail, RJ Ivison, J-P Kneib, and David T Frayer. Submillimeter galaxies. *Physics Reports*, 369(2):111–176, 2002.
- [5] Gibion Makiwa. *A Study of the Physics of the Interstellar Medium using the Herschel-SPIRE Instrument*. PhD thesis, University of Lethbridge (Canada), 2014.
- [6] F. P. Helmich and R. J. Ivison. FIRI - A far-infrared interferometer. *Experimental Astronomy*, 23:245–276, March 2009.
- [7] Edwin Hubble. A relation between distance and radial velocity among extra-galactic nebulae. *Proceedings of the National Academy of Sciences*, 15(3):168–173, 1929.
- [8] Christian Doppler. Beitrage zur fixsternenkunde. *Prag, Druck von G. Haase sohne, 1846.*, 1846.
- [9] SC Chapman, AW Blain, Ian Smail, and RJ Ivison. A redshift survey of the submillimeter galaxy population. *The Astrophysical Journal*, 622(2):772, 2005.
- [10] ESA. ESA Cosmic Vision. <http://sci.esa.int/cosmic-vision/46510-cosmic-vision/>, 2016.
- [11] G Neugebauer, HJ Habing, R Van Duinen, HH Aumann, B Baud, CA Beichman, DA Beintema, N Boggess, PE Clegg, T De Jong, et al. The infrared astronomical satellite (iras) mission. *The Astrophysical Journal*, 278:L1–L6, 1984.
- [12] NASA/IPAC Infrared Science Archive. IRAS Explanatory Supplement. <http://irsa.ipac.caltech.edu/IRASdocs/exp.sup/>, 2017.
- [13] MF Kessler, JA Steinz, ME Anderegg, J Clavel, G Drechsel, P Estaria, J Faelker, JR Riedinger, A Robson, BG Taylor, et al. The Infrared Space Observatory (ISO) mission. *Astronomy and Astrophysics*, 315:L27–L31, 1996.

- [14] Catherine Cesarsky and Alberto Salama. *ISO Science Legacy: A Compact Review of ISO Major Achievements*, volume 119. Springer Science & Business Media, 2006.
- [15] José Cernicharo and Jacques Crovisier. Water in space: the water world of ISO. In *ISO Science Legacy*, pages 29–69. Springer, 2005.
- [16] MW Werner, TL Roellig, FJ Low, GH Rieke, M Rieke, WF Hoffmann, E Young, JR Houck, B Brandl, GG Fazio, et al. The spitzer space telescope mission. *The Astrophysical Journal Supplement Series*, 154(1):1, 2004.
- [17] RD Gehrz, TL Roellig, MW Werner, GG Fazio, JR Houck, FJ Low, GH Rieke, BT Soifer, DA Levine, and EA Romana. The nasa spitzer space telescope. *Review of scientific instruments*, 78(1):011302, 2007.
- [18] David Charbonneau, Lori E Allen, S Thomas Megeath, Guillermo Torres, Roi Alonso, Timothy M Brown, Ronald L Gilliland, David W Latham, Georgi Mandushev, Francis T ODonovan, et al. Detection of thermal emission from an extrasolar planet. *The Astrophysical Journal*, 626(1):523, 2005.
- [19] Jonathan P Gardner, John C Mather, Mark Clampin, Rene Doyon, Matthew A Greenhouse, Heidi B Hammel, John B Hutchings, Peter Jakobsen, Simon J Lilly, Knox S Long, et al. The James Webb Space Telescope. *Space Science Reviews*, 123(4):485–606, 2006.
- [20] GL Pilbratt, JR Riedinger, T Passvogel, G Crone, D Doyle, U Gageur, AM Heras, C Jewell, L Metcalfe, S Ott, et al. Herschel space observatory—an esa facility for far-infrared and submillimetre astronomy. *Astronomy & Astrophysics*, 518:L1, 2010.
- [21] Matthew Joseph Griffin, A Abergel, As Abreu, Peter AR Ade, P André, J-L Augueres, T Babbedge, Y Bae, T Baillie, J-P Baluteau, et al. The Herschel-SPIRE instrument and its in-flight performance. *Astronomy & Astrophysics*, 518:L3, 2010.
- [22] ESA. Herschel science archive. <http://archives.esac.esa.int/hsa/whsa/>, 2017.
- [23] David A Naylor. Astronomical far-infrared imaging fourier transform spectroscopy: the challenges and the potential. In *Fourier Transform Spectroscopy*, pages FT4A–2. Optical Society of America, 2015.
- [24] Rosalind Hopwood, Ivan Valtchanov, Locke D Spencer, Jeremy Scott, Gibion Makiwa, Trevor Fulton, Gavin Noble, Chris S Benson, Nichola Marchili, Edward Polehampton, et al. The Herschel SPIRE spectral Feature Finder. In *Fourier Transform Spectroscopy*, pages FM2D–3. Optical Society of America, 2016.
- [25] Ivor Grattan-Guinness, Jean Baptiste Joseph Fourier, et al. *Joseph Fourier, 1768-1830; a survey of his life and work, based on a critical edition of his monograph on the propagation of heat, presented to the Institut de France in 1807*. MIT Press, 1972.

- [26] Jyrki Kauppinen and Jari Partanen. *Fourier transforms in spectroscopy*. John Wiley & Sons, 2011.
- [27] Leonhard Euler. *Introductio in analysin infinitorum, auctore Leonhardo Eulero...* apud Marcum-Michaelem Bousquet, 1748.
- [28] P. A. M. Dirac. *The principles of quantum mechanics*. 1947.
- [29] Sumner P Davis, Mark C Abrams, and James W Brault. *Fourier transform spectrometry*. Academic Press, 2001.
- [30] Eric W Weisstein. Gaussian integral. 2004.
- [31] BH Bransden and CJ Joachain. Quantum mechanics 2 nd edition, 2000.
- [32] AS Filler. Apodization and interpolation in fourier-transform spectroscopy. *JOSA*, 54(6):762–767, 1964.
- [33] E. Hecht. *Optics 4th edition*. 2001.
- [34] Locke Dean Spencer et al. *Imaging Fourier transform spectroscopy from a space based platform: the Herschel/SPIRE Fourier transform spectrometer*. PhD thesis, Lethbridge, Alta.: University of Lethbridge, Dept. of Physics and Astronomy, 2009.
- [35] Ren L Schilling. *Measures, integrals and martingales*. Cambridge University Press, 2017.
- [36] Marc-Antoine Parseval. Mémoire sur les séries et sur l'intégration complète d'une équation aux différences partielles linéaires du second ordre, à coefficients constants. *Mém. prés. par divers savants, Acad. des Sciences, Paris, (1)*, 1:638–648, 1806.
- [37] Lord Rayleigh. Liii. on the character of the complete radiation at a given temperature. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 27(169):460–469, 1889.
- [38] Graham Woan. *The Cambridge handbook of physics formulas*. Cambridge University Press, 2000.
- [39] David Millar, John Millar, Margaret Millar, and Ian Millar. *The Cambridge Dictionary of Scientists*. Cambridge University Press, 2002.
- [40] Harry Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, 1928.
- [41] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301, 1965.
- [42] Michael Heideman, Don Johnson, and C Burrus. Gauss and the history of the fast fourier transform. *IEEE ASSP Magazine*, 1(4):14–21, 1984.

- [43] Gordon Charles Danielson and Cornelius Lanczos. Some improvements in practical fourier analysis and their application to x-ray scattering from liquids. *Journal of the Franklin Institute*, 233(5):435–452, 1942.
- [44] Gibion Makiwa. *A Study of the Physics of the Interstellar Medium using the Herschel-SPIRE Instrument*. PhD thesis, University of Lethbridge (Canada), 2014.
- [45] John Edward Dyson and David Arnold Williams. *The physics of the interstellar medium*. CRC Press, 1997.
- [46] Luis J Boya. The thermal radiation formula of planck (1900). *arXiv preprint physics/0402064*, 2004.
- [47] Max Planck. Ueber das gesetz der energieverteilung im normalspectrum. *Annalen der physik*, 309(3):553–563, 1901.
- [48] Daniel V Schroeder. *An introduction to thermal physics*, 1999.
- [49] Peter F Bernath. *Spectra of atoms and molecules*. Oxford university press, 2015.
- [50] JD Ingle Jr and SR Crouch. *Spectrochemical analysis* prentice hall. *New Jersey*, 1988.
- [51] Werner Heisenberg. Über den anschaulichen inhalt der quantentheoretischen kinematik und mechanik. In *Original Scientific Papers Wissenschaftliche Originalarbeiten*, pages 478–504. Springer, 1985.
- [52] HA Lorentz. The absorption and emission lines of gaseous bodies. In *Knaw, Proceedings*, volume 8, pages 1905–1906, 1906.
- [53] HM Foley. The pressure broadening of spectral lines. *Physical Review*, 69(11-12):616, 1946.
- [54] Laurence S Rothman, Iouli E Gordon, Alain Barbe, D Chris Benner, Peter F Bernath, Manfred Birk, Vincent Boudon, Linda R Brown, Alain Campargue, J-P Champion, et al. The hitran 2008 molecular spectroscopic database. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 110(9):533–572, 2009.
- [55] James Clerk Maxwell. V. illustrations of the dynamical theory of gases. part i. on the motions and collisions of perfectly elastic spheres. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 19(124):19–32, 1860.
- [56] LJ Slater. *Confluent hypergeometric functions*. 1960.
- [57] Albert A Michelson and Edward W Morley. On the relative motion of the earth and of the luminiferous ether. *Sidereal Messenger*, vol. 6, pp. 306-310, 6:306–310, 1887.
- [58] T Fulton, David A Naylor, ET Polehampton, I Valtchanov, R Hopwood, N Lu, J-P Baluteau, G Mainetti, C Pearson, Andreas Papageorgiou, et al. The data processing pipeline for the Herschel SPIRE Fourier Transform Spectrometer. *Monthly Notices of the Royal Astronomical Society*, 458(2):1977–1989, 2016.

- [59] I Valtchanov et al. SPIRE Observers Manual. Technical report, HERSCHEL-DOC-0798, 2011.
- [60] Norman Longdon. European space science: Horizon 2000. In *ESA Special Publication*, volume 1070, 1984.
- [61] Herschel ObserversManual. Herschel-hsc-doc-0876, 2014.
- [62] Th De Graauw, FP Helmich, TG Phillips, J Stutzki, E Caux, ND Whyborn, P Dieleman, PR Roelfsema, H Aarts, R Assendorp, et al. The Herschel-heterodyne instrument for the far-infrared (HIFI). *Astronomy & Astrophysics*, 518:L6, 2010.
- [63] Albrecht Poglitsch, Christoffel Waelkens, Norbert Geis, Helmut Feuchtgruber, Bart Vandenbussche, Louis Rodriguez, Oliver Krause, Etienne Renotte, Christiaan Van Hoof, P Saraceno, et al. The photodetector array camera and spectrometer (pacs) on the herchel space observatory. *Astronomy & astrophysics*, 518:L2, 2010.
- [64] AP Marston and D Teyssier. HIFI Observers Manual. Technical report, HERSCHEL-HSC-DOC-0784, 2011.
- [65] Peter AR Ade, Peter A Hamilton, and David A Naylor. An absolute dual beam emission spectrometer. In *Fourier Transform Spectroscopy*, page FWE3. Optical Society of America, 1999.
- [66] ESA. Herschel interactive processing environment help system. <http://herschel.esac.esa.int/hcss-doc-12.0/>, 2016.
- [67] Goutam Chattopadhyay, Jason Glenn, James J Bock, Brooks K Rownd, Martin Caldwell, and Matthew J Griffin. Feed horn coupled bolometer arrays for SPIRE-design, simulations, and measurements. *IEEE Transactions on Microwave Theory and Techniques*, 51(10):2139–2146, 2003.
- [68] Gibion Makiwa, David A Naylor, Marc Ferlet, Carl Salji, Bruce Swinyard, Edward Polehampton, and Matthijs HD van der Wiel. Beam profile for the Herschel–SPIRE Fourier transform spectrometer. *Applied optics*, 52(16):3864–3875, 2013.
- [69] N Marchili, R Hopwood, T Fulton, ET Polehampton, I Valtchanov, J Zaretski, DA Naylor, MJ Griffin, P Imhof, T Lim, et al. Calibration of Herschel SPIRE FTS observations at different spectral resolutions. *Monthly Notices of the Royal Astronomical Society*, 464(3):3331–3342, 2016.
- [70] S Ott, J Bakker, J Brumfitt, N de Candussio, L Dwedari, AM Heras, S Leeks, AP Marston, J-J Mathieu, J Pizarro, et al. The herchel data processing system. In *Astronomical Data Analysis Software and Systems XV*, volume 351, page 516, 2006.
- [71] Ronin Wu, Edward T Polehampton, Mireya Etxaluze, Gibion Makiwa, David A Naylor, Carl Salji, Bruce M Swinyard, Marc Ferlet, Matthijs HD van der Wiel, Anthony J Smith, et al. Observing extended sources with the Herschel SPIRE Fourier transform spectrometer. *Astronomy & Astrophysics*, 556:A116, 2013.

- [72] Rosalind Hopwood, Trevor Fulton, Edward T Polehampton, Ivan Valtchanov, Dominique Benielli, Peter Imhof, Tanya Lim, Nanyao Lu, Nicola Marchili, Chris P Pearson, et al. Herschel SPIRE FTS telescope model correction. *Experimental Astronomy*, 37(2):195–205, 2014.
- [73] Jacqueline Fischer, Tjeerd Klaassen, Niels Hovenier, Gerd Jakob, Albrecht Poglitsch, and Oren Sternberg. Cryogenic far-infrared laser absorptivity measurements of the herchel space observatory telescope mirror coatings. *Applied Optics*, 43(19):3765–3771, 2004.
- [74] Glenn S Orton, Leigh N Fletcher, Julianne I Moses, Amy K Mainzer, Dean Hines, Heidi B Hammel, F Javier Martin-Torres, Martin Burgdorf, Cecile Merlet, and Michael R Line. Mid-infrared spectroscopy of uranus from the spitzer infrared spectrometer: 1. determination of the mean temperature structure of the upper troposphere and stratosphere. *Icarus*, 243:494–513, 2014.
- [75] Rosalind Hopwood. SPIRE Spectrometer HPDP: Observations of the Spectrometer calibrators. http://archives.esac.esa.int/hsa/legacy/HPDP/SPIRE/SPIRE-S/cal_targets/.
- [76] N. Marchili L. D. Spencer J. Scott C. Benson N. Hładczuk E. T. Polehampton N. Lu G. Makiwa D. A. Naylor G. Noble M. J. Griffin R. Hopwood, I. Valtchanov. The Herschel SPIRE Fourier Transform Spectrometer Spectral Feature Finder and The Spectral Feature Catalogue. 2018. In-preparation.
- [77] Shay Zucker. Cross-correlation and maximum-likelihood analysis: a new approach to combining cross-correlation functions. *Monthly Notices of the Royal Astronomical Society*, 342(4):1291–1298, 2003.
- [78] R. Hopwood, E. T. Polehampton, I. Valtchanov, B. M. Swinyard, T. Fulton, N. Lu, N. Marchili, M. H. D. van der Wiel, D. Benielli, P. Imhof, J.-P. Baluteau, C. Pearson, D. L. Clements, M. J. Griffin, T. L. Lim, G. Makiwa, D. A. Naylor, G. Noble, E. Puga, and L. D. Spencer. Systematic characterization of the Herschel SPIRE Fourier Transform Spectrometer. *MNRAS*, 449:2274–2303, May 2015.
- [79] David A Naylor, Gibion Makiwa, Ian Veenendaal, Jeremy Scott, and Locke D Spencer. Line Asymmetry in Herschel SPIRE FTS Spectra. In *Fourier Transform Spectroscopy*, pages FM2D–2. Optical Society of America, 2016.
- [80] Georges Lemaître. Un univers homogène de masse constante et de rayon croissant rendant compte de la vitesse radiale des nébuleuses extra-galactiques. In *Annales de la Société scientifique de Bruxelles*, volume 47, pages 49–59, 1927.
- [81] Richard B Larson. Turbulence and star formation in molecular clouds. *Monthly Notices of the Royal Astronomical Society*, 194(4):809–826, 1981.
- [82] Frank H Shu, Fred C Adams, and Susana Lizano. Star formation in molecular clouds: observation and theory. *Annual review of astronomy and astrophysics*, 25(1):23–81, 1987.

- [83] Eric Herbst. Chemistry in the interstellar medium. *Annual Review of Physical Chemistry*, 46(1):27–54, 1995.
- [84] David John Adams. *An introduction to galaxies and cosmology*. Cambridge University Press, 2004.
- [85] Anthony J Remijan and Andrew J Markwick-Kemper. Splatalogue: Database for astronomical spectroscopy. 2008.
- [86] Linus Pauling and William F Sheehan. The dissociation energy of carbon monoxide and the heat of sublimation of graphite. *Proceedings of the National Academy of Sciences*, 35(7):359–363, 1949.
- [87] Yinghe Zhao, Nanyao Lu, C Kevin Xu, Yu Gao, Steven D Lord, Vassilis Charmandaris, Tanio Diaz-Santos, Aaron Evans, Justin Howell, Andreea O Petric, et al. The [n ii] 205 μm emission in local luminous infrared galaxies. *The Astrophysical Journal*, 819(1):69, 2016.
- [88] L Deharveng, B Lefloch, A Zavagno, J Caplan, Anthony Peter Whitworth, D Nadeau, and S Martin. Triggered massive-star formation at the border of the hii region sh 104. *Astronomy & Astrophysics*, 408(3):L25–L28, 2003.
- [89] W Jaffe, K Meisenheimer, HJA Rottgerine, Ch Leinert, et al. The central dusty torus in the active nucleus of ngc 1068. *Nature*, 429(6987):47, 2004.
- [90] N Lu, Y Zhao, CK Xu, Y Gao, L Armus, JM Mazzarella, KG Isaak, AO Petric, V Charmandaris, T Díaz-Santos, et al. Warm molecular gas in luminous infrared galaxies. *The Astrophysical Journal Letters*, 787(2):L23, 2014.
- [91] B. M. Swinyard, E. T. Polehampton, R. Hopwood, I. Valtchanov, N. Lu, T. Fulton, D. Benielli, P. Imhof, N. Marchili, J.-P. Baluteau, G. J. Bendo, M. Ferlet, M. J. Griffin, T. L. Lim, G. Makiwa, D. A. Naylor, G. S. Orton, A. Papageorgiou, C. P. Pearson, B. Schulz, S. D. Sidher, L. D. Spencer, M. H. D. van der Wiel, and R. Wu. Calibration of the Herschel SPIRE Fourier Transform Spectrometer. *MNRAS*, 440:3658–3674, June 2014.
- [92] ESA. Planck maps. http://irsa.ipac.caltech.edu/data/Planck/release_2/all-sky-maps/, 2017.
- [93] David J Hollenbach and AGGM Tielens. Photodissociation regions in the interstellar medium of galaxies. *Reviews of Modern Physics*, 71(1):173, 1999.
- [94] Glenn J White and Rachael Padman. Images of atomic carbon in the interstellar medium. *Nature*, 354(6354):511–513, 1991.
- [95] Alan D McNaught and Alan D McNaught. *Compendium of chemical terminology*, volume 1669. Blackwell Science Oxford, 1997.
- [96] PJ Linstrom and WG Mallard. National institute of standards and technology: Gaithersburg. MD, March, 20899, 2003.

[97] NSR DS-NES. Bond dissociation energies in simple molecules. 1970.

Appendix A

Visual Example of Radial Velocity Estimating Algorithm

This appendix is intended to give a visual representation of how the redshift routine identifies ^{12}CO candidate features and calculates a corresponding radial velocity estimate. This is one of the more simple examples which only requires an input of detected feature emission frequencies and does not require the subsequent checks described in Section 5.3.1. The source under consideration is NGC 7027 (Figure A.1, obsid 1342188670).

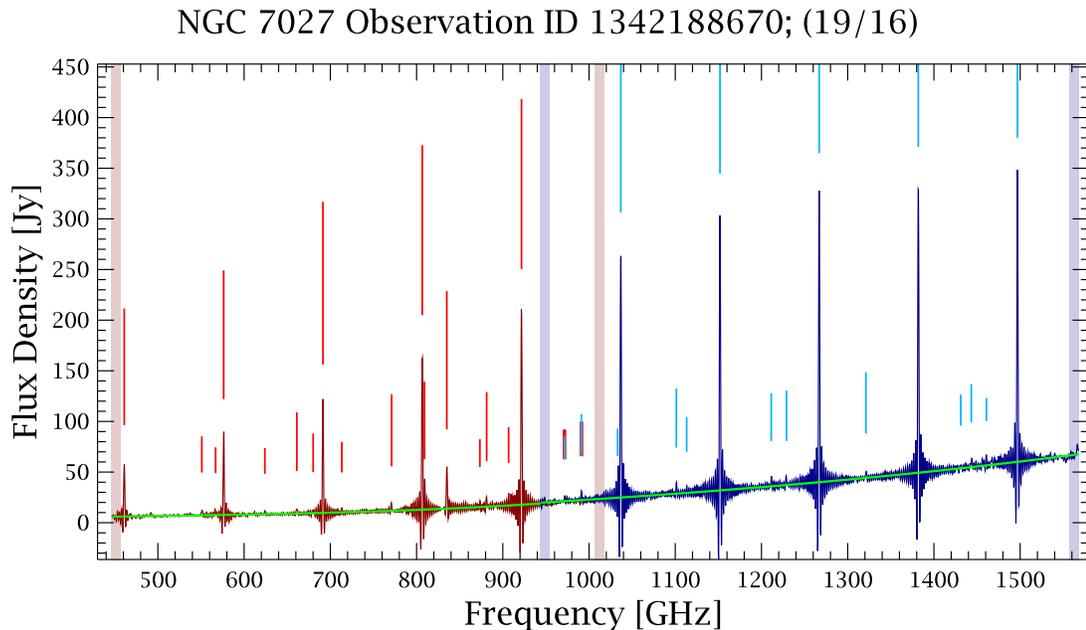


Figure A.1: The source under consideration is NGC 7027. Spectral features are indicated with vertical lines with the length of the lines indicating the feature's corresponding relative SNR.

To begin, we recall that identification of ^{12}CO emission is based on searching the frequency differences between detected features for elements matching the characteristic difference array (CDA) of ^{12}CO . An example of the ^{12}CO rotational emission ladder, along with lines indicating the values of the CDA, is shown in Figure A.2.

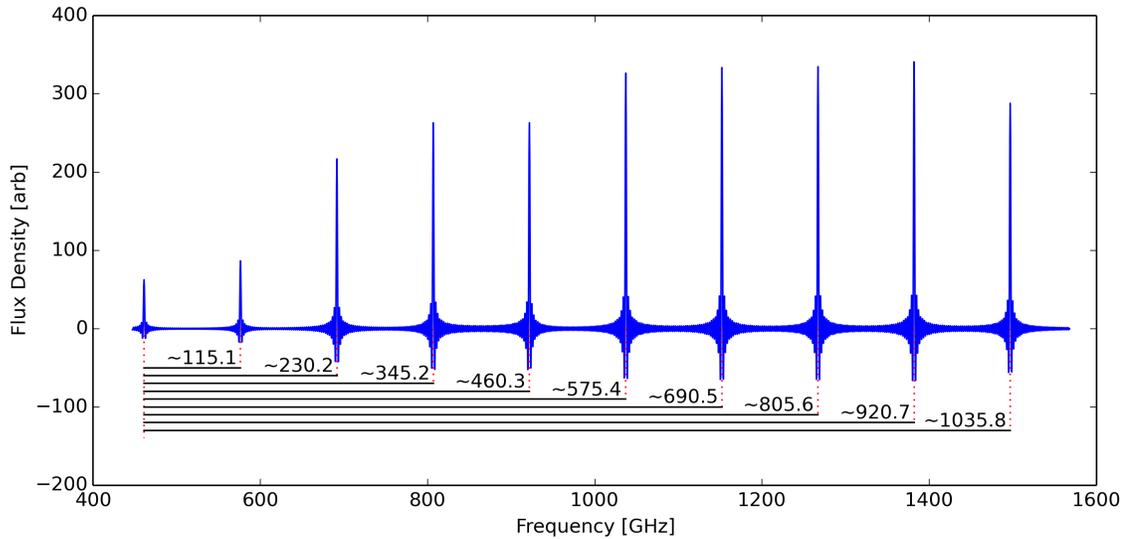


Figure A.2: The rotational ^{12}CO emission ladder with the CDA indicated by horizontal black lines. Annotations indicate the approximate value of the corresponding CDA element.

Since the CDA is somewhat altered when there is a non-zero relative velocity between source and observer, a velocity dependent tolerance must be included when searching for frequency differences matching elements of the CDA. If one assumes some maximum radial velocity, a tolerance for each CDA element can be calculated using Equation 5.6. Assuming a maximum radial velocity of 6,000 km/s, the resulting CDA and its associated tolerances can be represented as in Figure A.3.

We can now return to the source being analyzed. The continuum subtracted spectrum of NGC 7027 is shown in Figure A.4, with identified features indicated by vertical green lines, with the length of each line indicating the relative SNR of the associated feature. Since the redshift routine is only concerned with the frequencies of detected features and their associated SNR, the spectrum is disregarded and the remaining information is plotted in Figure A.5.

For each feature identified by the FF, a frequency difference array is calculated between the feature under inspection and all features identified by the FF. This process is represented graphically in Figure A.6 for the feature detected at ~ 460 GHz.

Each frequency difference is then compared with the elements of the CDA. Frequency differences which match the CDA to within the permitted tolerance are identified as in Figure A.7. The total number of CDA values with at least one corresponding match are counted. This produces the total number of "matches" for the feature under inspection as indicated in Figure A.7. Note that this is not the same as counting the total number of red lines in Figure A.7. The two closely spaced features at ~ 800 GHz together only increment the "matches" counter by one.

This process then repeats for each identified features. The total number of matches for each features in the spectrum of NGC 7027 is shown in Figure A.8. The maximum number of matches in this case is 10, and all features receiving this maximum value are identified as potential ^{12}CO emission feature candidates. The candidate features are then paired with

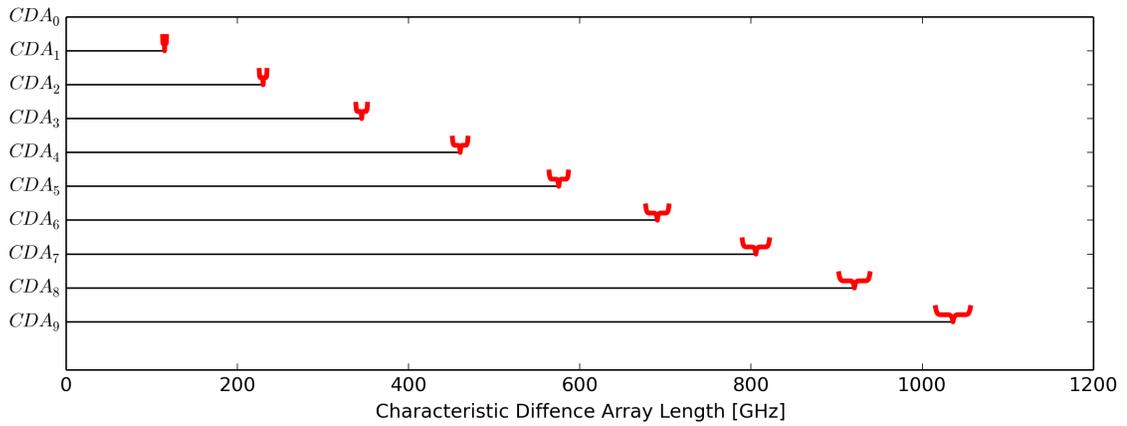


Figure A.3: The frequency length of each element of the CDA is presented as horizontal black lines with their corresponding tolerance indicated by red braces. The tolerances in this case have been calculated for a maximum assumed radial velocity of 6,000 km/s.

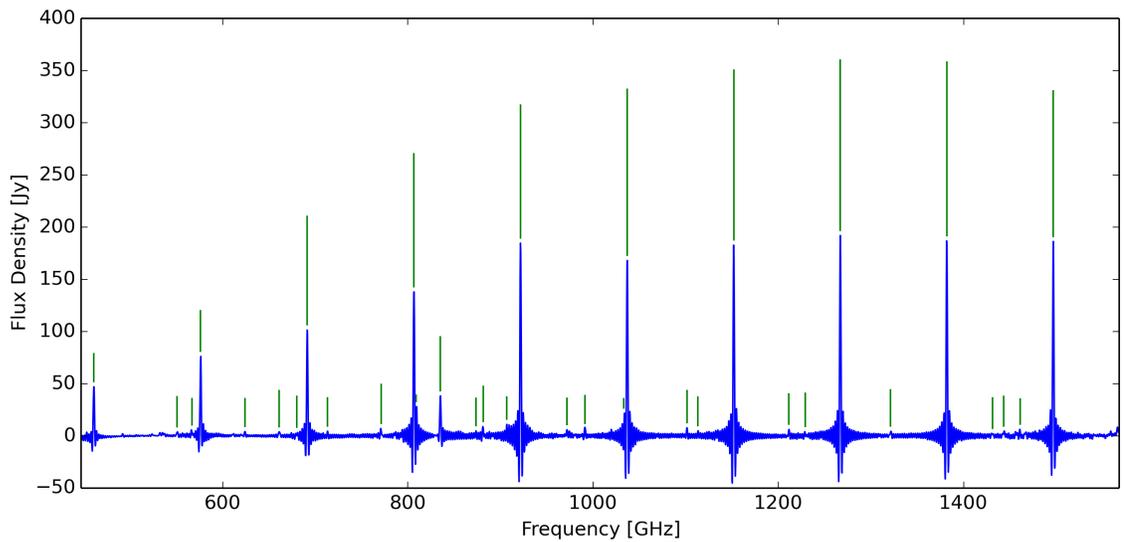


Figure A.4: The continuum subtracted spectrum of NGC 7027 with features identified by the FF indicated with vertical green lines. The length of each vertical line indicated the relative SNR of the identified feature.

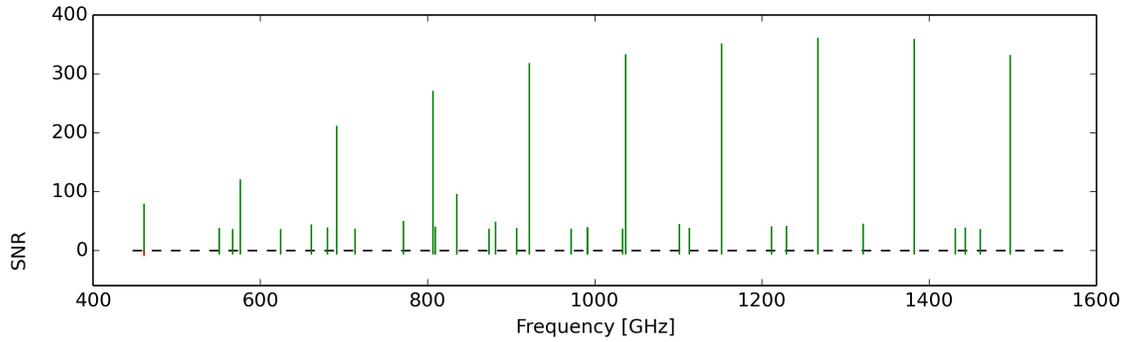


Figure A.5: A visual representation of the inputs to the redshift routine. Features identified by the FF are indicated with vertical green lines with the length of each vertical line indicated the relative SNR of the identified feature. The spectrum itself is not required for this analysis.

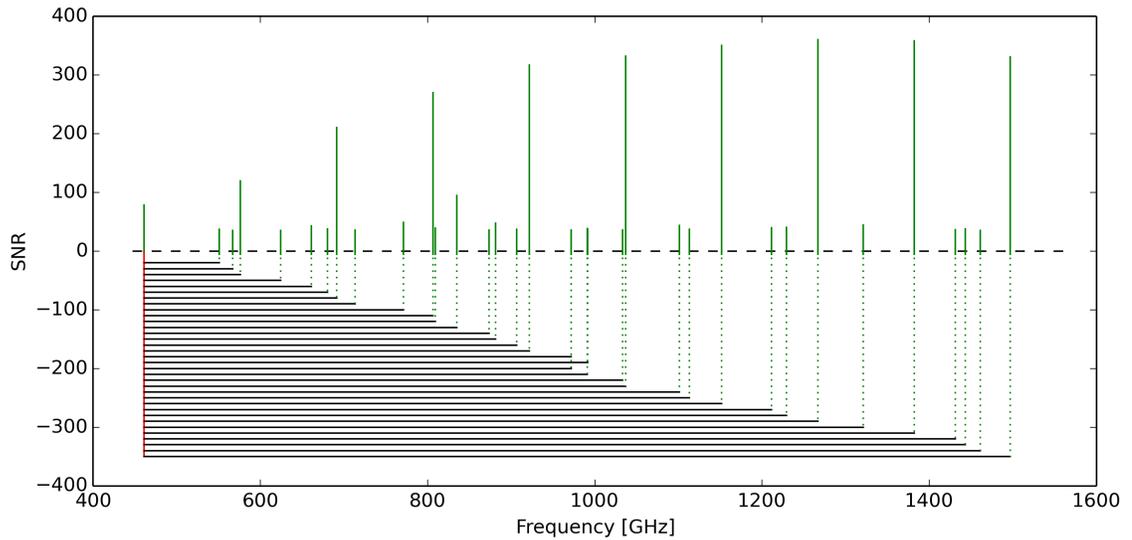


Figure A.6: The frequency differences between one feature under investigation and all features identified by the FF.

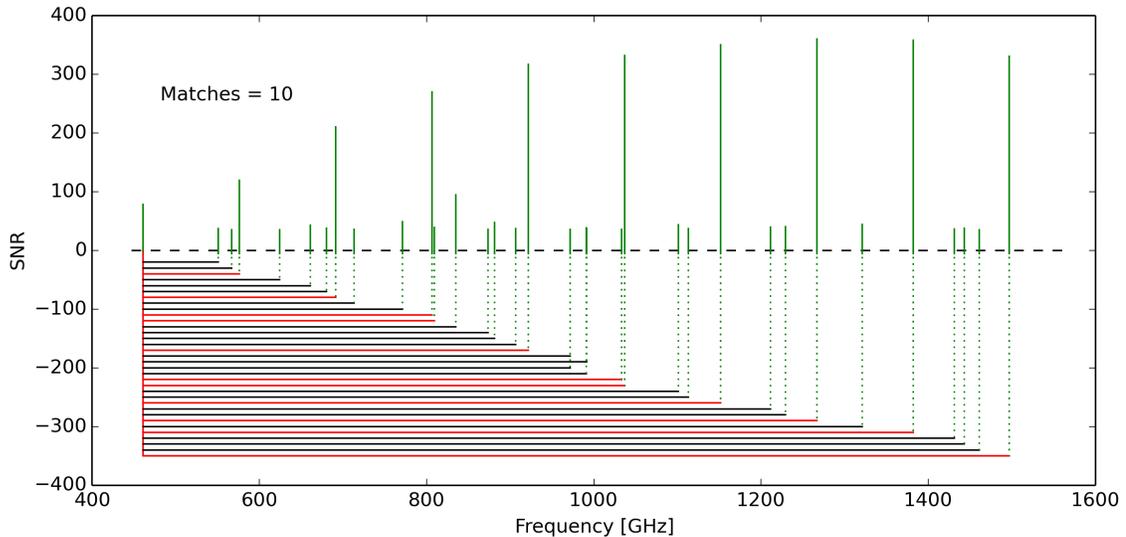


Figure A.7: Frequency differences matching the CDA to within the permitted tolerance are highlighted in red. The total number of CDA elements with corresponding matches is indicated by the “matches” parameter. Note that zero, the frequency difference between a feature and itself, is an element of the CDA which is highlighted in red and contributes to the “matches” count.

their nearest ^{12}CO emission rest frequency. For each candidate feature, a velocity estimate is calculated using Equation 5.3 producing a velocity estimate array. The median value and standard deviation of this array are used to parameterize the source’s radial velocity and velocity error, respectively.

Indeed in this example the ^{12}CO emission features were correctly identified without complication. However, one can appreciate how two closely spaced lines might receive the same “matches” count if the tolerances are high enough. For example, the ^{12}CO (7-6) and [CI] features at ~ 806 GHz in Figure A.8 could both reasonably receive a “matches” count of 10 in some situations. When this happens, the feature with the highest SNR is taken as the most probable ^{12}CO emission candidate.

In rare cases, a non- ^{12}CO feature may be identified as a probable ^{12}CO emission candidate along with other true ^{12}CO emission features. This has only been observed in spectra with a high feature density. Misidentified features such as this produce a velocity estimate which deviates significantly from the median value of the velocity estimate array. As such, the velocity estimate deviating most from the median value will be removed while the resulting velocity estimate array produces a high standard deviation (> 100 km/s).

These additional checks, which are not required for the example spectrum presented, help to ensure the quality of radial velocity estimates included in the SPIRE Spectral Feature Finder catalogue.

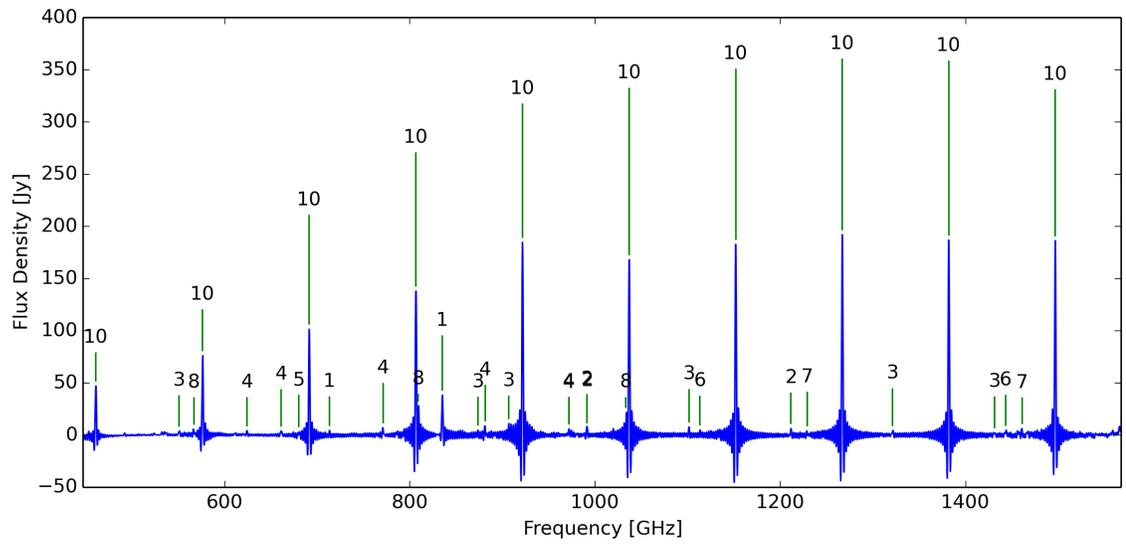


Figure A.8: The total number of matches for each feature identified by the FF. Features with the maximum number of matches are the most probable ^{12}CO emission feature candidates.

Appendix B

Mapping Observation Postcards

This appendix summarizes the FF mapping results as they pertain to velocity estimates, along with presenting a selection of mapping postcards from the public FF catalogue.⁴

Table B.1: “Obsid indicates the ten digit identification number corresponding to the SPIRE FTS mapping observation of the source given in the “Object column. The combined velocity “Map Size in terms of pixels is presented along with “Coverage indicating what percentage of pixels have estimates with $N > 3$ or are based on the [NII] feature (Section 5.3). The “Species column indicates which molecule or atom was used the most in generating the velocity map. When the velocity map contains more null pixels than estimates derived from ^{12}CO or [NII], the column entry is ordered by which molecule or atom contributed the most in generating the velocity map, with N/A indicating pixels with a null estimate.

Obsid	Object	Map Size [pix]	Coverage [%]	Species
1342231066	W43MM3	13x12	98.7	^{12}CO
1342231067	W43MM4	13x12	98.7	^{12}CO
1342268320	Haro 5a IRS	12x12	98.6	^{12}CO
1342249465	HD37041	27x24	97.2	^{12}CO
1342204882	M16-finger5-1	12x12	97.2	^{12}CO
1342231048	W43MM1	13x12	96.8	^{12}CO
1342231065	W43MM2	13x12	96.8	^{12}CO
1342227516	NGC 2023-1-1	27x25	96.7	^{12}CO
1342214846	HD37041	27x24	96.3	^{12}CO
1342192175	OrionBar_map3	27x25	96.0	^{12}CO
1342192174	OrionBar_map2	27x25	96.0	^{12}CO
1342192173	HD37041	27x25	94.5	^{12}CO
1342204891	rho_oph_fts_on	26x25	94.3	^{12}CO
1342243631	SGRA_f1	26x25	94.2	^{12}CO
1342251326	NGC6334I_on	26x25	94.0	^{12}CO
1342213373	N1333-IRAS4_on	27x26	93.6	^{12}CO
1342242591	HD37041	26x25	93.5	^{12}CO
1342204920	HD 37022 (towards)	26x25	93.5	^{12}CO
1342228734	HD37041	26x25	93.4	^{12}CO
1342204919	HD37041	26x25	93.2	^{12}CO

Continued on next page

⁴<https://www.cosmos.esa.int/web/herschel/spire-spectral-feature-catalogue>

B. MAPPING OBSERVATION POSTCARDS

Obsid	Object	Map Size [pix]	Coverage [%]	Species
1342204895	EL29_on	26x25	93.1	¹² CO
1342204883	HH216	12x12	93.1	¹² CO
1342204892	rho_oph_h2	26x25	92.5	¹² CO
1342214841	NGC6334IN_on	26x25	92.2	¹² CO
1342228703	M 17-2	26x25	92.2	¹² CO
1342214827	NGC6334I_on	26x25	92.0	¹² CO
1342242624	HOPS 108-1	26x25	91.7	¹² CO
1342204922	n2023_fts_2	26x25	90.2	¹² CO
1342204921	n2023_fts_1	26x25	90.2	¹² CO
1342262914	Carina_PosC'	12x13	89.7	¹² CO
1342204893	rho_oph_fts_off	12x12	89.6	¹² CO
1342251289	IRS46	12x12	88.9	¹² CO
1342204894	rho_oph_fts_off_2	12x12	87.5	¹² CO
1342251314	IRDC316.76-0.01_HMPO	13x13	87.0	¹² CO
1342204880	M16-finger2-1	12x12	86.1	¹² CO
1342262915	Carina_SWS	12x13	85.9	¹² CO
1342204898	IRAS16293_on	26x25	85.4	¹² CO
1342252900	G11.92-0.61	12x12	84.7	¹² CO
1342216881	HH_COpeak	26x24	84.1	¹² CO
1342268312	IRAS18264-1152-1	12x12	82.6	¹² CO
1342243604	IRAS20126+4104-4	12x12	81.9	¹² CO
1342256377	Carina_posB	12x12	79.2	¹² CO
1342216882	HH_IRpeak	26x24	79.2	¹² CO
1342204884	M16-finger1-1	12x12	78.5	¹² CO
1342216880	HH_dense_core	26x24	75.0	¹² CO
1342256376	Carina_posA	12x12	70.1	¹² CO
1342245117	UYSO 1	13x12	66.0	¹² CO
1342251325	G349.7	12x12	64.6	¹² CO
1342270044	IRAS18507+0121-1	13x12	62.2	¹² CO
1342268292	V1647 Ori	12x12	60.4	¹² CO
1342252898	G19.01-0.03	13x12	59.6	¹² CO
1342259070	mol160_3	26x26	58.9	¹² CO
1342259069	mol160_2	26x26	55.8	¹² CO
1342259068	mol1601_1	26x26	55.6	¹² CO
1342268310	G23.60+0.0M1-1	12x12	55.6	¹² CO
1342262926	MGE_4111	12x12	91.0	[NII]
1342265807	MGE_4485	12x12	91.0	[NII]
1342254041	MGE_3739	13x12	88.5	[NII]
1342262924	MGE_4204	12x12	88.2	[NII]
1342221711	RCW49	26x25	86.5	[NII]
1342254042	MGE_3736	14x12	85.1	[NII]
1342204881	M16-finger n-1	12x12	84.7	[NII]
1342254040	MGE_3280	14x12	84.5	[NII]

Continued on next page

B. MAPPING OBSERVATION POSTCARDS

Obsid	Object	Map Size [pix]	Coverage [%]	Species
1342228732	I434-2	26x25	76.8	[NII]
1342242622	Aquila-sub-fil-1	12x12	72.9	[NII]
1342262927	MGE_4121	12x12	70.8	[NII]
1342251315	IRDC321.73+0.05_P2	13x11	69.9	[NII]
1342270043	IRAS19035+0641-1	13x12	64.7	[NII]
1342243626	IRDC011.11-0.12_P1	12x12	59.7	[NII]
1342201202	ngc5194	12x12	57.6	[NII]
1342251328	G357.7B-IRS	12x12	64.6	N/A- ¹² CO -NII
1342210846	NGC253	13x12	62.8	N/A-NII- ¹² CO
1342219550	30dor-spire	13x12	58.3	N/A-NII- ¹² CO
1342228731	Rosette-pillar-1-1	12x12	55.6	N/A- ¹² CO -NII
1342243615	IRDC18223-1243_P3-1	12x12	54.2	N/A-NII- ¹² CO
1342259066	lmc n159-1-1-1	12x13	48.7	N/A- ¹² CO -NII
1342243614	IRDC019.30+0.07_P1	12x12	48.6	N/A-NII- ¹² CO
1342251321	CTB37A-N	12x12	47.9	N/A- ¹² CO -NII
1342201204	n7023E_fts_1	25x26	46.2	N/A- ¹² CO -NII
1342212343	ngc4945	13x12	44.2	N/A-NII- ¹² CO
1342253648	G19.36-0.03	13x12	43.6	N/A- ¹² CO -NII
1342212345	m83	28x24	41.5	N/A-NII- ¹² CO
1342227778	AFGL4106	11x11	41.3	N/A-NII- ¹² CO
1342243603	ngc6946	13x12	41.0	N/A-NII- ¹² CO
1342247604	ngc3627	13x12	40.4	N/A-NII- ¹² CO
1342270045	G25.04-0.2M1-1	13x12	39.1	N/A-NII- ¹² CO
1342201205	n7023E_fts_2	25x26	37.5	N/A- ¹² CO -NII
1342198923	n7023_fts	25x28	36.4	N/A- ¹² CO -NII
1342237026	ngc5055	13x12	35.3	N/A-NII- ¹² CO
1342202261	california-1	13x12	35.3	N/A-NII- ¹² CO
1342262928	V346 Nor	12x12	33.3	N/A- ¹² CO -NII
1342204036	cen a	29x25	31.9	N/A-NII- ¹² CO
1342268303	L1251B	12x13	31.4	N/A- ¹² CO -NII
1342257932	30 Doradus/N157-1	12x13	30.8	N/A- ¹² CO -NII
1342224769	ngc6946_05	12x12	29.2	N/A-NII- ¹² CO
1342254039	MGE_3269	13x12	28.2	N/A-NII- ¹² CO
1342268293	HH111-Position-A	26x25	27.8	N/A- ¹² CO -NII
1342257914	lmc n158-2	12x12	27.8	N/A-NII- ¹² CO
1342236997	ngc4254	13x12	27.6	N/A-NII- ¹² CO
1342257915	N11	12x13	27.6	N/A-NII- ¹² CO
1342247743	ngc3521	13x12	26.9	N/A-NII- ¹² CO
1342247572	ngc4321	13x12	25.6	N/A-NII- ¹² CO
1342213376	ngc 891-1	13x12	25.6	N/A-NII- ¹² CO
1342208388	m82	28x24	25.4	N/A- ¹² CO -NII
1342246253	ISOSSJ23053+5953	12x12	22.2	N/A- ¹² CO -NII
1342245871	ngc7331	13x12	21.8	N/A-NII- ¹² CO

Continued on next page

B. MAPPING OBSERVATION POSTCARDS

Obsid	Object	Map Size [pix]	Coverage [%]	Species
1342221698	NGC7027	26x24	19.4	N/A- ¹² CO -NII
1342246992	ngc4826	13x11	18.9	N/A- ¹² CO -NII
1342262822	V733 Cep	12x13	18.6	N/A- ¹² CO -NII
1342265844	Re 50 N NIRS 1	13x12	17.3	N/A- ¹² CO -NII
1342213444	ngc1068	26x25	16.2	N/A- ¹² CO -NII
1342227785	CRL 618	11x11	15.7	N/A- ¹² CO -NII
1342257933	HH54 B	26x25	14.8	N/A- ¹² CO -NII
1342247573	ngc4631	13x12	14.7	N/A-NII- ¹² CO
1342268294	HH111-Position-B	26x25	14.3	N/A- ¹² CO -NII
1342262905	N44	13x12	12.8	N/A- ¹² CO -NII
1342249472	TMC1B1	13x12	10.9	N/A- ¹² CO -NII
1342262919	MGE_4384	12x13	10.9	N/A-NII- ¹² CO
1342210859	ngc4038overlap	28x25	10.1	N/A-NII- ¹² CO
1342243609	B335	12x12	9.0	N/A- ¹² CO -NII
1342210860	ngc4038	28x25	8.4	N/A-NII- ¹² CO
1342248251	ngc4569	13x12	8.3	N/A-NII- ¹² CO
1342198922	IC63	27x24	7.7	N/A- ¹² CO -NII
1342245091	N22-1	12x13	7.7	N/A- ¹² CO -NII
1342239337	ngc1097	12x12	6.9	N/A-NII- ¹² CO
1342214849	IRAM04191_on	26x25	6.3	N/A- ¹² CO -NII
1342256104	N83-2	12x13	5.8	N/A- ¹² CO -NII
1342256101	SWbarS-6	13x13	5.3	N/A- ¹² CO -NII
1342247117	ngc3351	13x12	5.1	N/A-NII- ¹² CO
1342248250	ngc5713	13x12	5.1	N/A-NII- ¹² CO
1342257913	N180-1	12x12	4.2	N/A- ¹² CO -NII
1342265863	PP 13 S	14x12	4.2	N/A- ¹² CO -NII
1342256102	SWbarN-2	13x13	4.1	N/A- ¹² CO -NII
1342214850	IRAS04191_on	26x25	3.8	N/A- ¹² CO -NII
1342256080	N66/NGC 346-1	13x13	3.6	N/A- ¹² CO -NII
1342228745	ngc3077	12x13	3.2	N/A-NII- ¹² CO
1342248233	ngc1482	13x12	3.2	N/A-NII- ¹² CO
1342214854	L1489-IRS_on	26x26	3.1	N/A- ¹² CO -NII
1342262907	N157	13x13	3.0	N/A- ¹² CO -NII
1342245083	He2-10	12x12	2.8	N/A- ¹² CO -NII
1342237025	ngc4536	13x12	1.9	N/A-NII- ¹² CO
1342210848	ngc4418	13x12	1.3	N/A- ¹² CO -NII
1342256081	DiffuseReg1-1	12x13	1.3	N/A-NII- ¹² CO
1342239353	ngc1266	13x12	1.3	N/A- ¹² CO -NII
1342268295	HH111-Position-C	26x25	0.8	N/A- ¹² CO -NII
1342252892	ngc2798	13x11	0.7	N/A- ¹² CO -NII
1342230417	m101_02	12x13	0.6	N/A-NII- ¹² CO
1342228706	ngc2976_00	12x13	0.6	N/A-NII- ¹² CO
1342197488	Ced 201-2	14x12	0.6	N/A- ¹² CO -NII

Continued on next page

B. MAPPING OBSERVATION POSTCARDS

Obsid	Object	Map Size [pix]	Coverage [%]	Species
1342204026	IC59	25x27	0.6	N/A-NII- ¹² CO
1342221687	NGC 6826 Mid	25x27	0.0	N/A-NII- ¹² CO
1342270032	N76Wrim-5	12x12	0.0	N/A-NII- ¹² CO
1342204890	11721-6	12x12	0.0	N/A-NII- ¹² CO
1342249055	NGC1275-nucleus	13x12	0.0	N/A-NII- ¹² CO
1342243601	CB26_Class-I-1	13x12	0.0	N/A-NII- ¹² CO
1342262855	EX Lup	12x12	0.0	N/A-NII- ¹² CO
1342219574	Polaris-cvi	12x13	0.0	N/A-NII- ¹² CO
1342204887	11721-3	12x12	0.0	N/A-NII- ¹² CO
1342227520	OFF-postion2	13x12	0.0	N/A-NII- ¹² CO
1342219573	HH 381 IRS	12x12	0.0	N/A-NII- ¹² CO
1342204885	11721-1	12x12	0.0	N/A-NII- ¹² CO
1342262906	N132D-1	13x12	0.0	N/A-NII- ¹² CO
1342246280	Neptune	11x10	0.0	N/A-NII- ¹² CO
1342227787	NGC 6543 W Knot	28x24	0.0	N/A-NII- ¹² CO
1342255263	Parsamian 21	13x12	0.0	N/A-NII- ¹² CO
1342204886	11721-2	12x12	0.0	N/A-NII- ¹² CO
1342258700	HH 354 IRS	12x13	0.0	N/A-NII- ¹² CO
1342197489	Ced 201-3	14x12	0.0	N/A-NII- ¹² CO
1342204888	11721-4	12x12	0.0	N/A-NII- ¹² CO
1342262904	N49-1	13x12	0.0	N/A-NII- ¹² CO
1342270031	NGC249-1	12x13	0.0	N/A-NII- ¹² CO
1342202262	california-3	13x12	0.0	N/A-NII- ¹² CO
1342227789	NGC 6543	28x25	0.0	N/A-NII- ¹² CO
1342249464	V883 Ori	13x12	0.0	N/A-NII- ¹² CO
1342239352	ngc1377	13x12	0.0	N/A-NII- ¹² CO
1342202260	california-2	13x12	0.0	N/A-NII- ¹² CO
1342197487	Ced 201-1	13x12	0.0	N/A-NII- ¹² CO
1342204889	11721-5	12x12	0.0	N/A-NII- ¹² CO
1342243602	CB230	13x11	0.0	N/A-NII- ¹² CO
1342221681	NGC 7009	27x25	0.0	N/A-NII- ¹² CO

B.1 ^{12}CO Dominant

Object: OrionBar_map2, Obsid: 1342192174

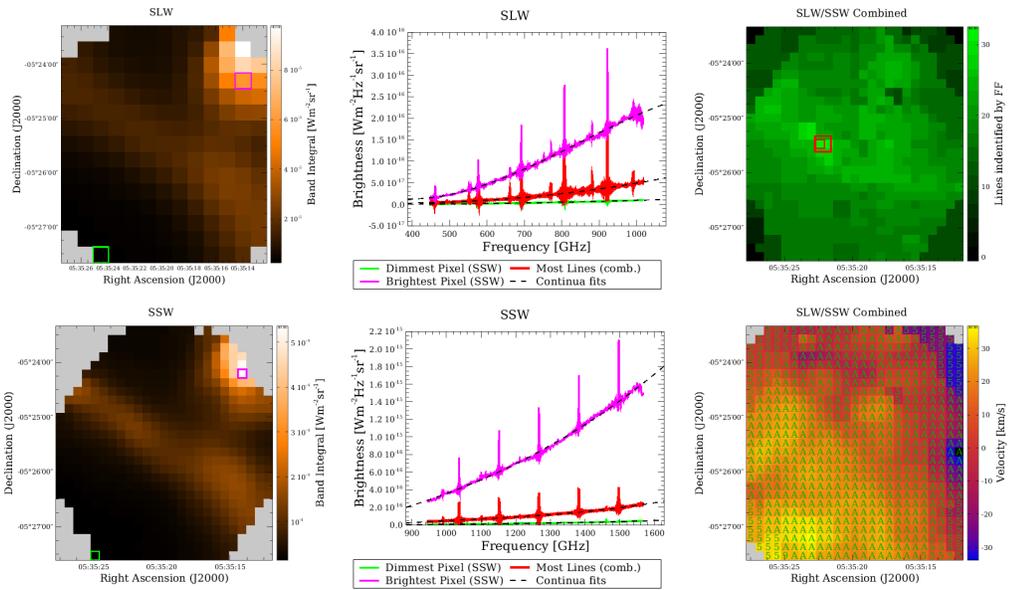


Figure B.1: FF Mapping Postcards, 1342192174

Object: NGC6334I_on, Obsid: 1342214827

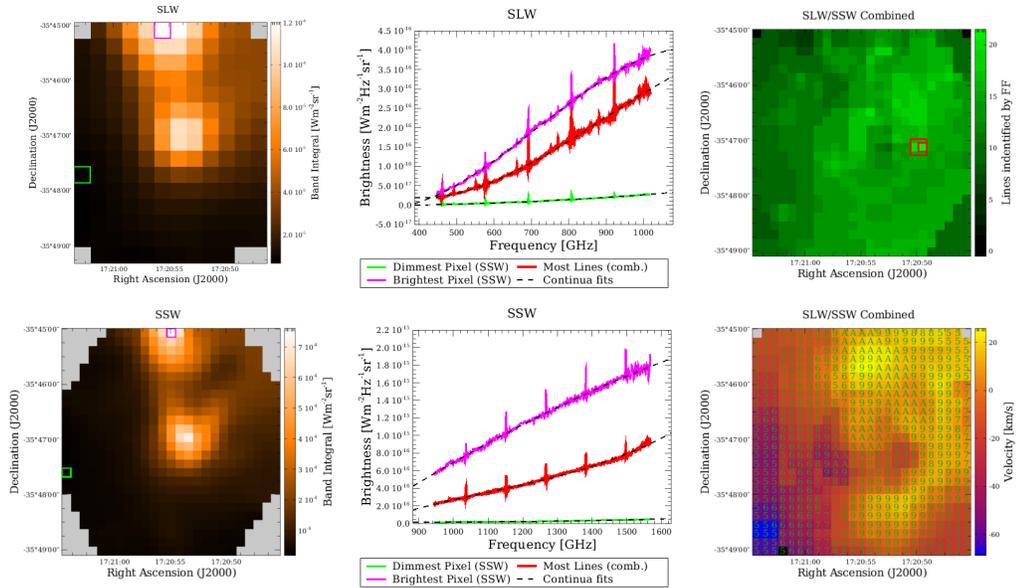


Figure B.2: 1342214827

Object: NGC 2023-1-1, Obsid: 1342227516

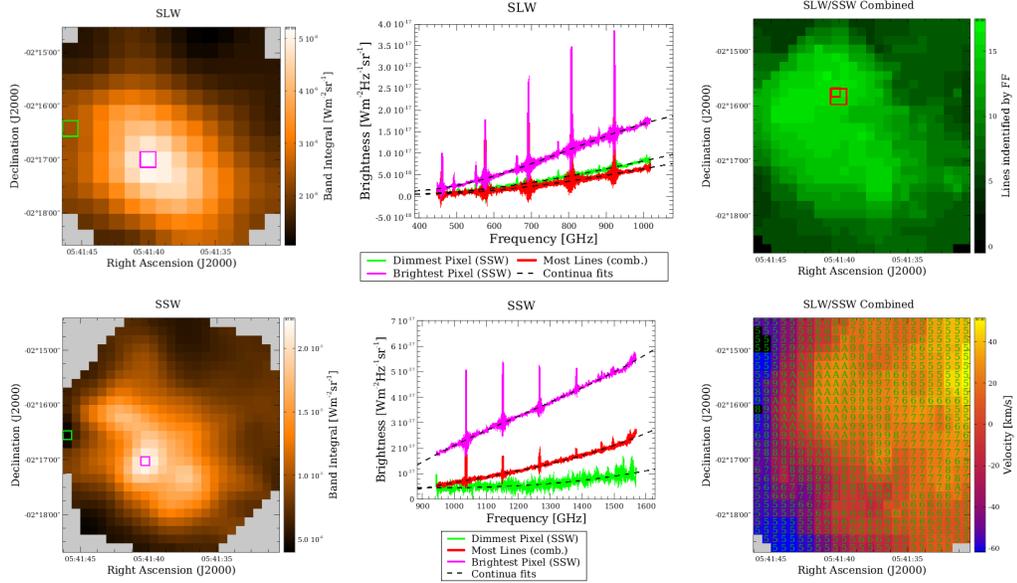


Figure B.3: 1342227516

Object: M 17-2, Obsid: 1342228703

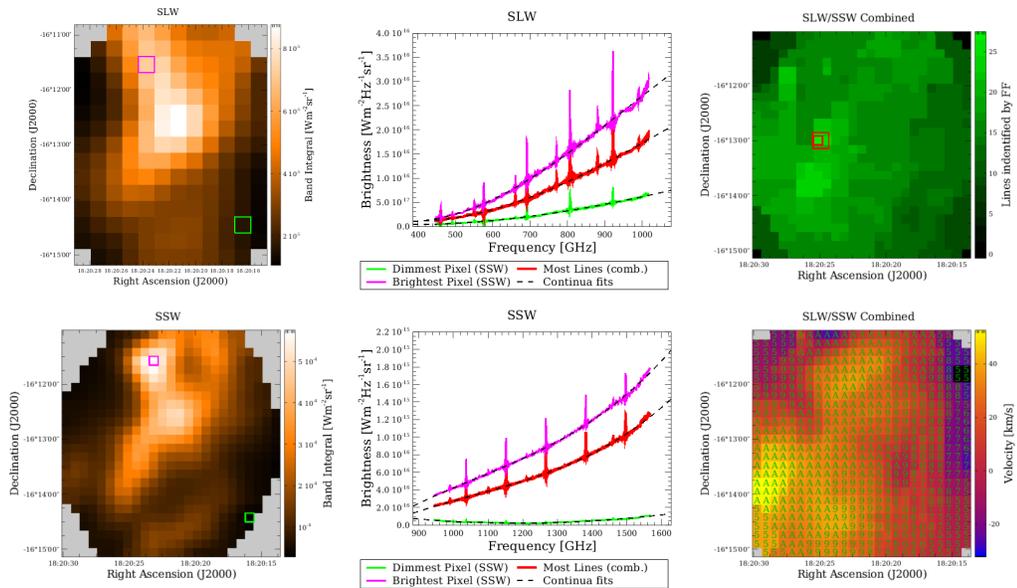


Figure B.4: 1342228703

B.2 NII Dominant

Object: I434-2, Obsid: 1342228732

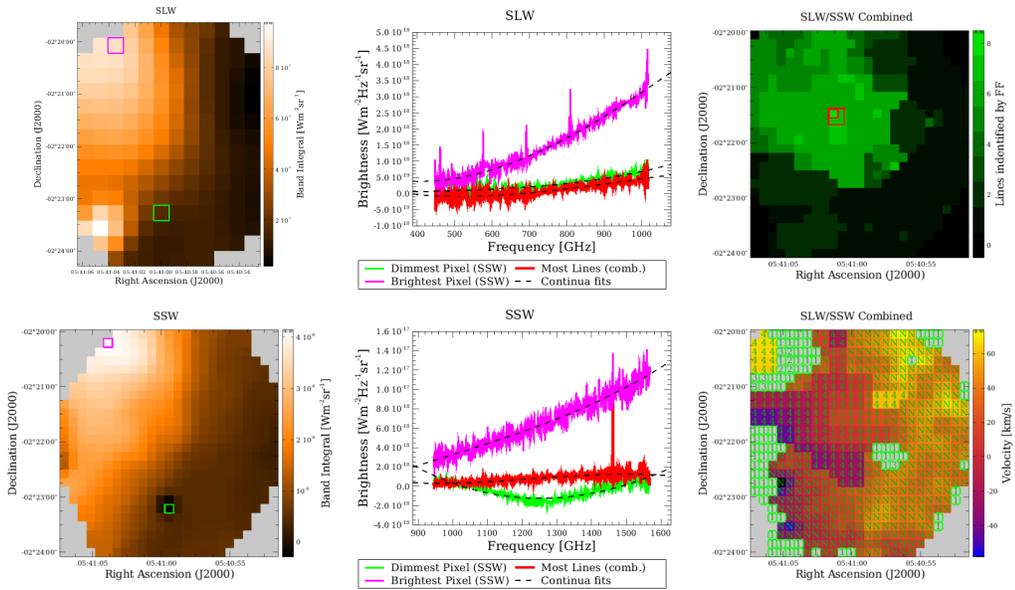


Figure B.5: 1342228732

Object: RCW49, Obsid: 1342221711

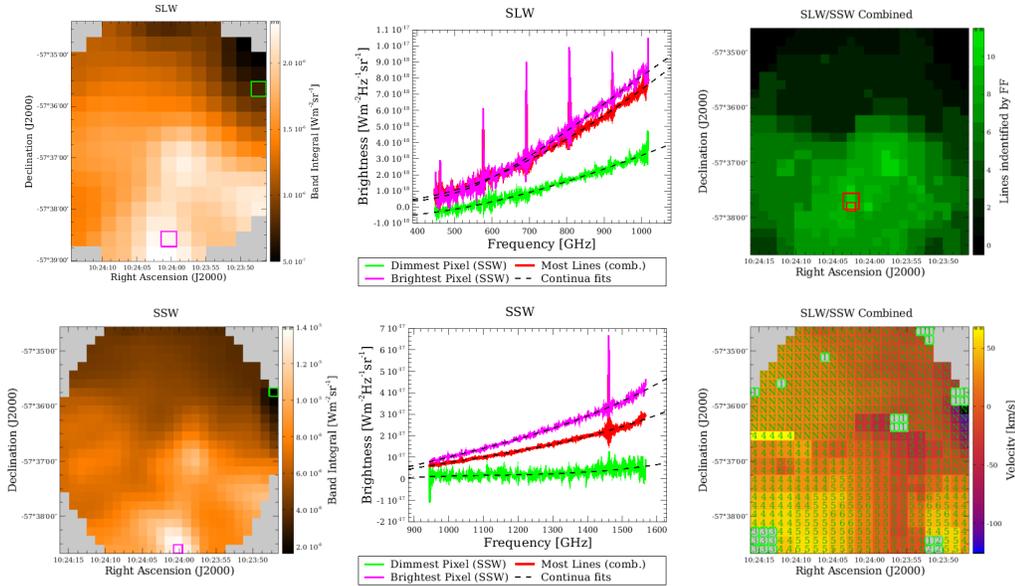


Figure B.6: 1342221711

Object: MGE_3280, Obsid: 1342254040

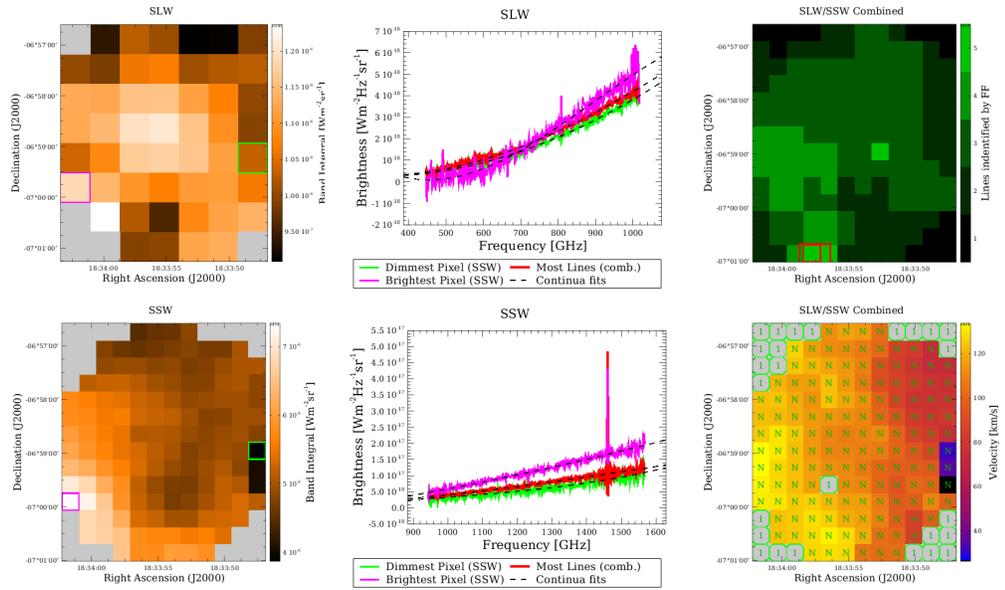


Figure B.7: 1342254040

Object: MGE_4485, Obsid: 1342265807

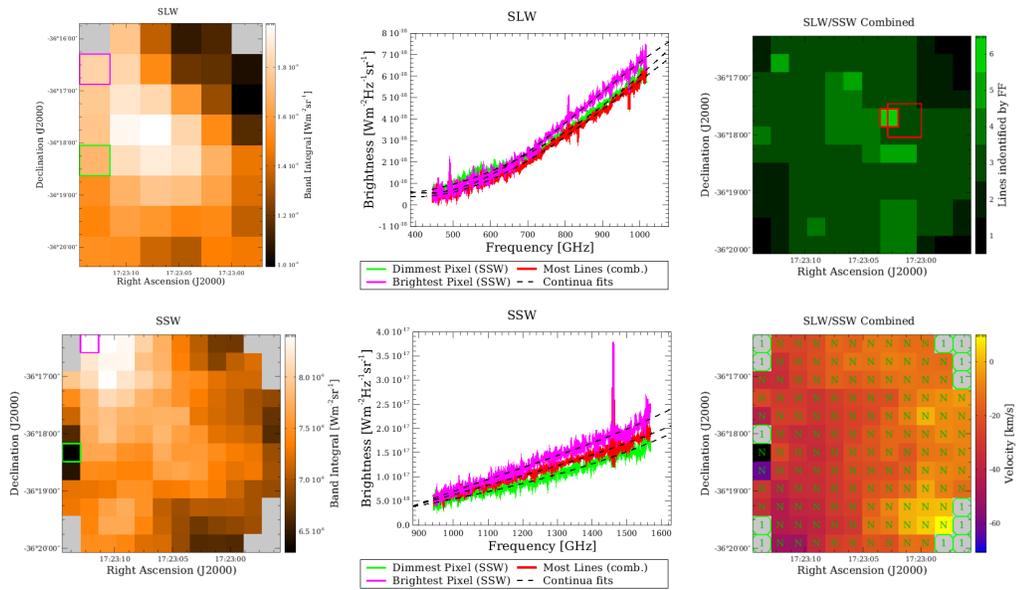


Figure B.8: 1342265807

B.3 Null Dominant

Object: m83, Obsid: 1342212345

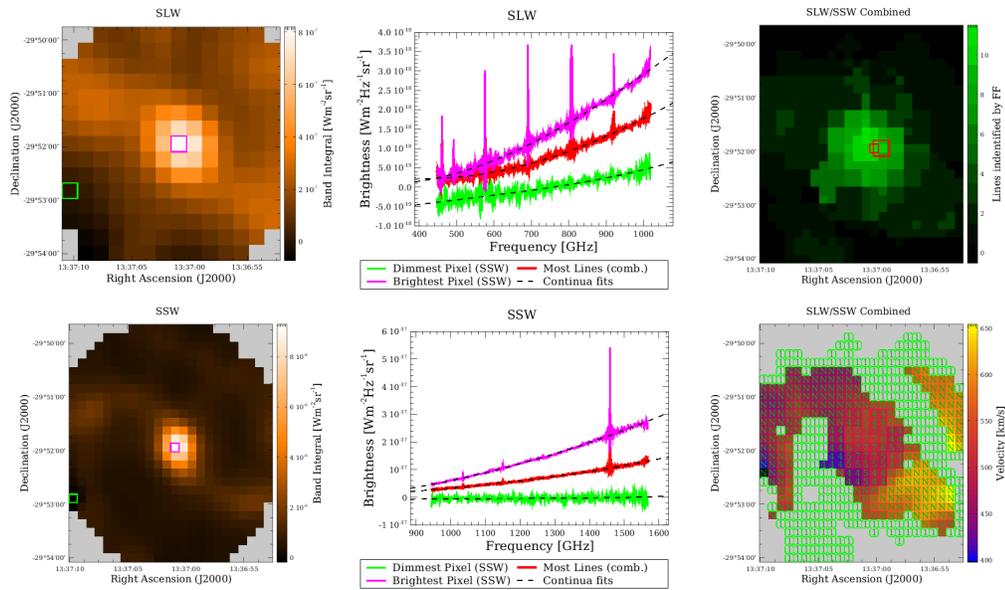


Figure B.9: 1342212345

Object: ngc4038, Obsid: 1342210860

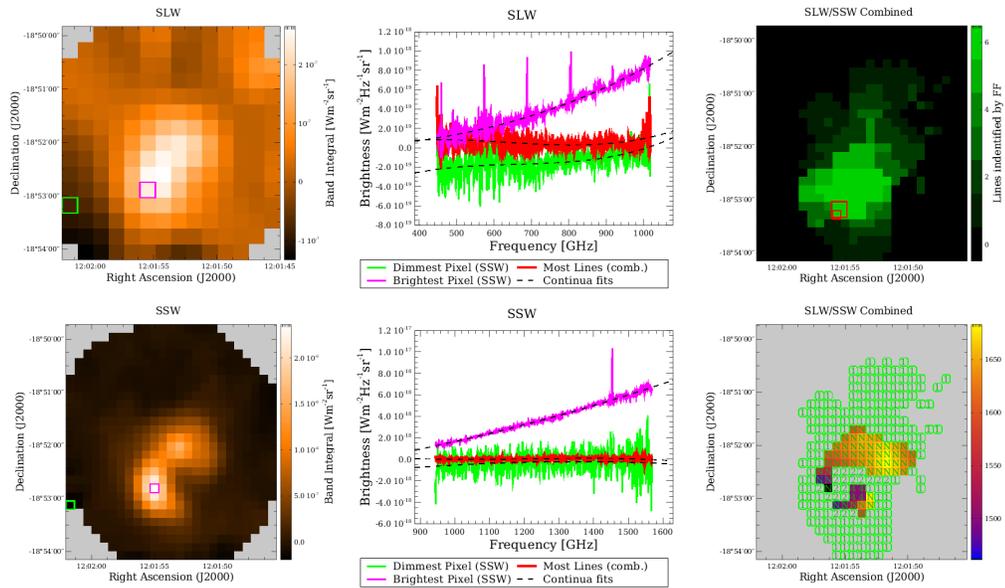


Figure B.10: 1342210860

Object: m82, Obsid: 1342208388

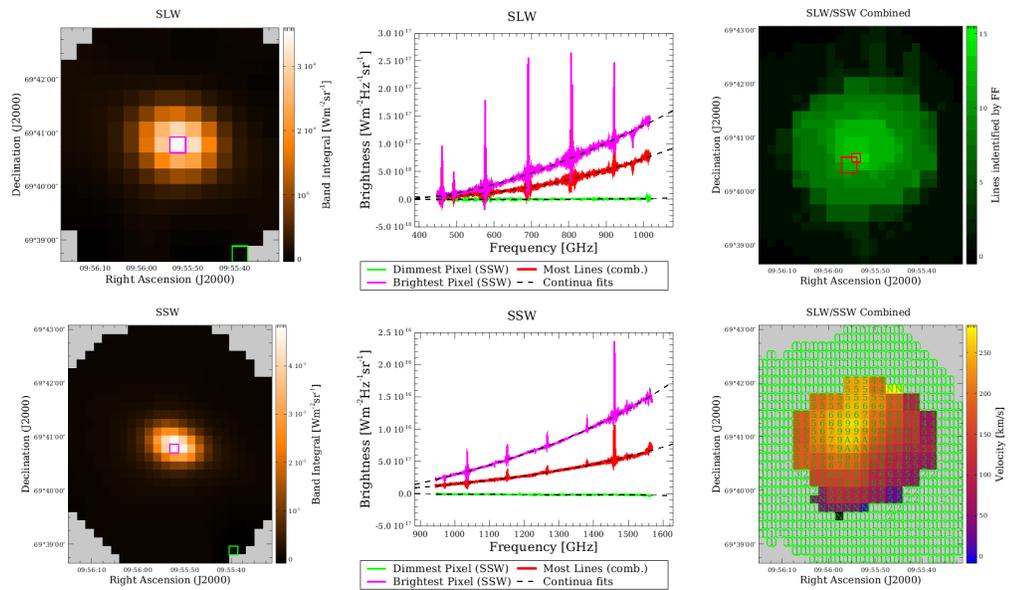


Figure B.11: 1342208388

Object: lmc n159-1-1-1, Obsid: 1342259066

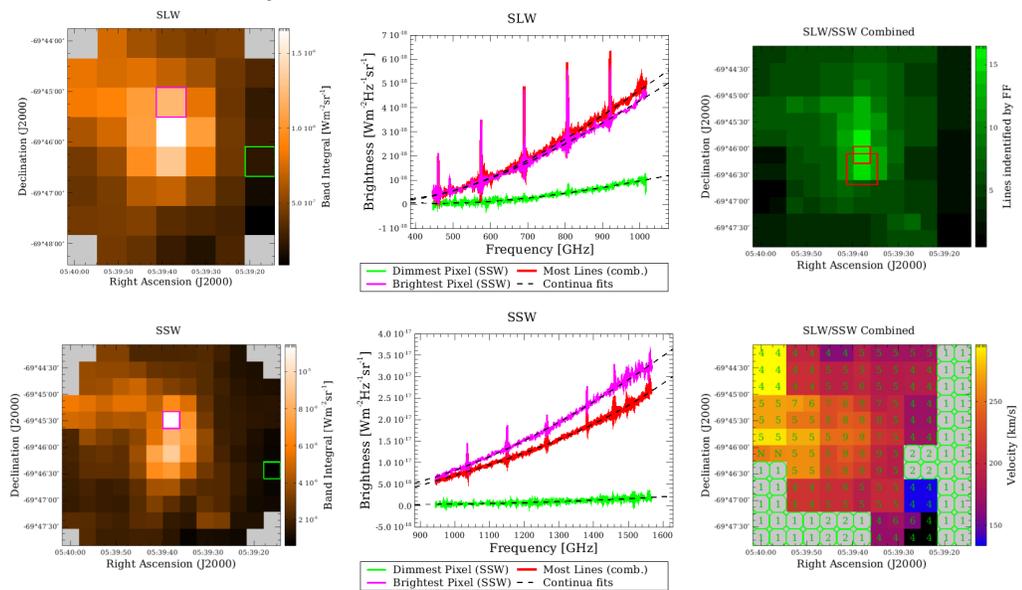


Figure B.12: 1342259066

Appendix C

Neutral Carbon 3P_1 - 3P_0 Fine Structure Line Detection

This appendix illustrates the utility of extending the NCC (Chapter 6) functionality towards detection of the neutral carbon 3P_1 - 3P_0 fine structure line within the FF. What follows is a series of plots which show postcards from the sparse FF catalogue (left), and a close up of the suspected 3P_1 - 3P_0 transition (right). SPIRE spectral data is presented in blue, the fitted model is shown with a dashed sinc profile, while the expected position of the 3P_1 - 3P_0 transition based on [CI] detection is indicated with a cyan line.

C.1 Unique Neutral Carbon Pairs

This section shows unique detection of both neutral carbon fine structure transitions using the NCC which were not detected within the main body of the FF routine.

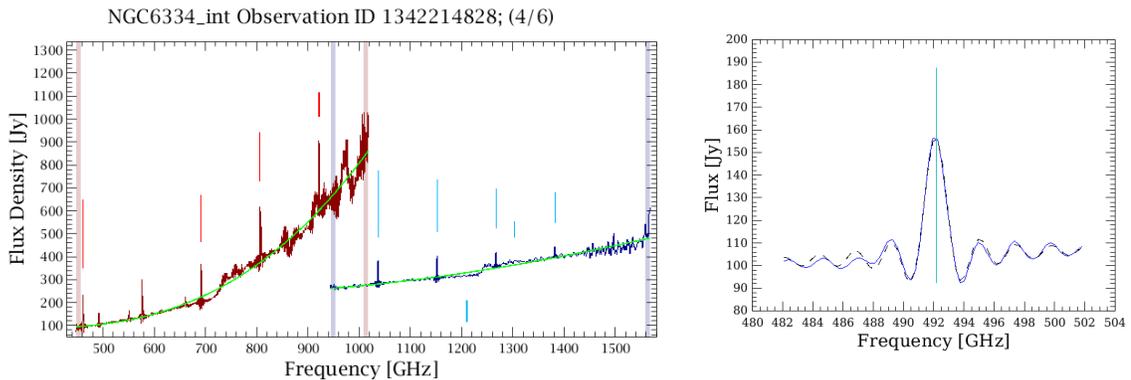
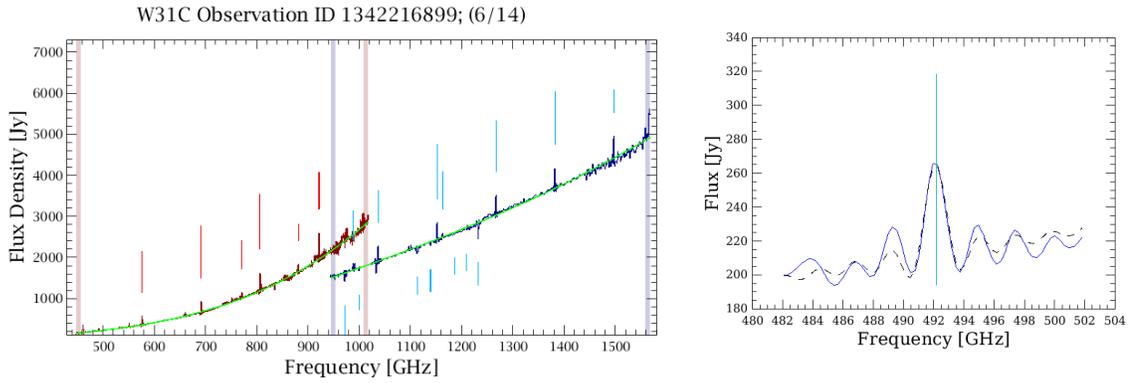
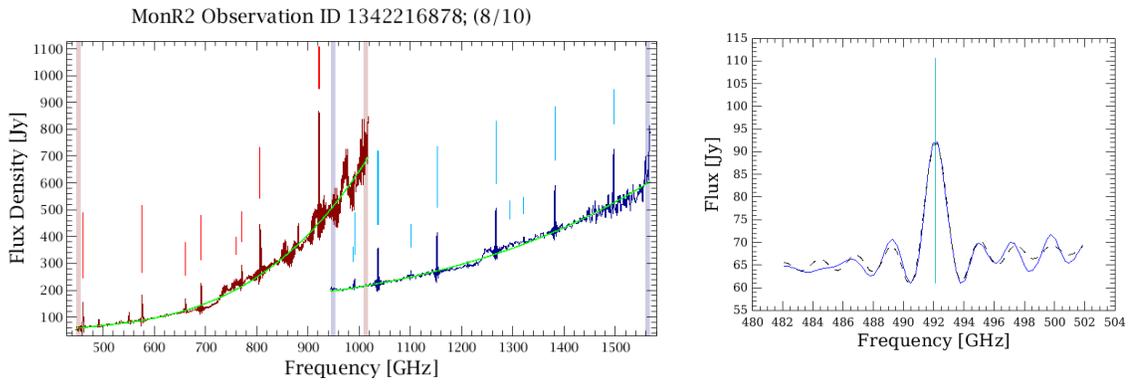
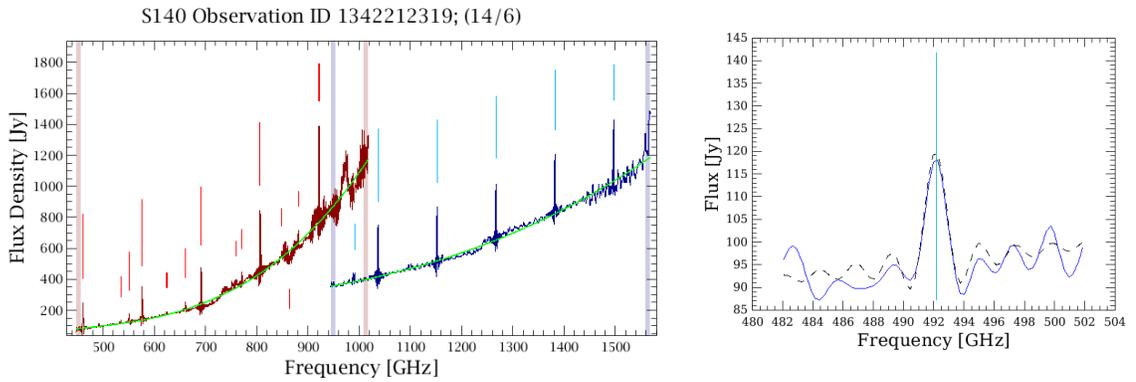


Figure C.1: $[CI]^3P_1 - ^3P_0$ snr = 37.846340

Figure C.2: $[\text{CI}]^3 P_1 - ^3 P_0$ snr = 12.180308Figure C.3: $[\text{CI}]^3 P_1 - ^3 P_0$ snr = 18.500926Figure C.4: $[\text{CI}]^3 P_1 - ^3 P_0$ snr = 7.949178

C.2 Unique ${}^3P_1 - {}^3P_0$ Transition Detection

This section shows spectra with unique ${}^3P_1 - {}^3P_0$ fine structure feature detection based on the NCC routine results. However, in these cases the ${}^3P_2 - {}^3P_1$ transition was detected by the FF without invoking the NCC.

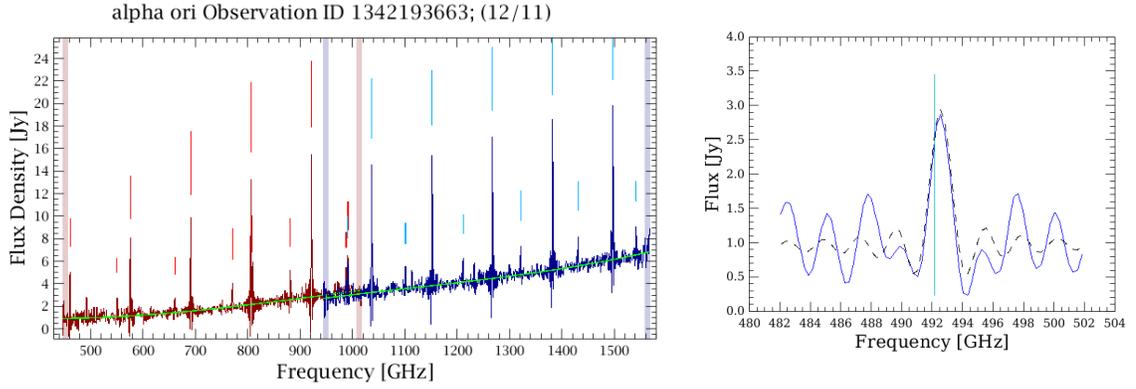


Figure C.5: $[CI]{}^3P_1 - {}^3P_0$ snr = 6.268916

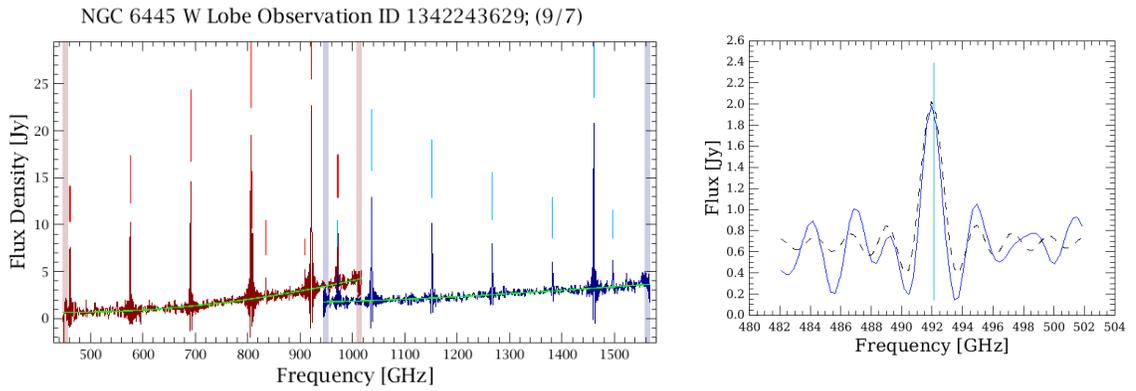


Figure C.6: $[CI]{}^3P_1 - {}^3P_0$ snr = 7.668759

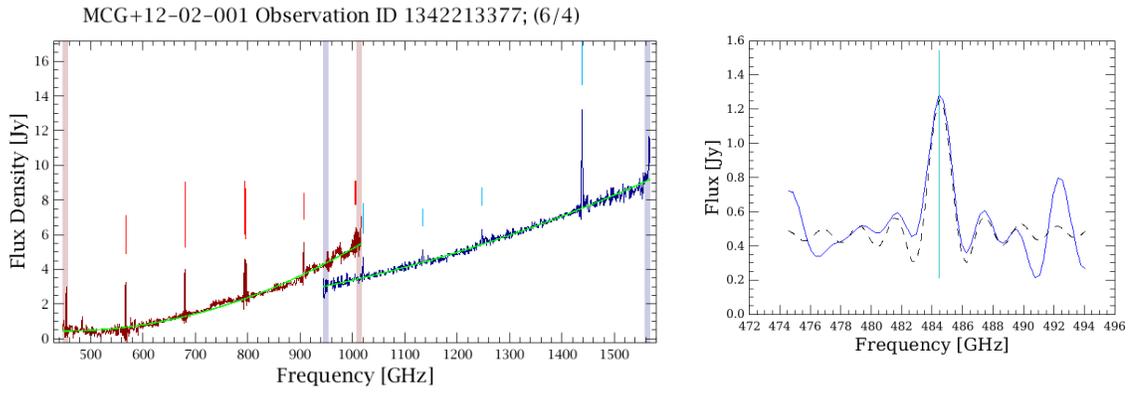


Figure C.7: $[\text{CI}]^3P_1 - {}^3P_0$ snr = 5.366224

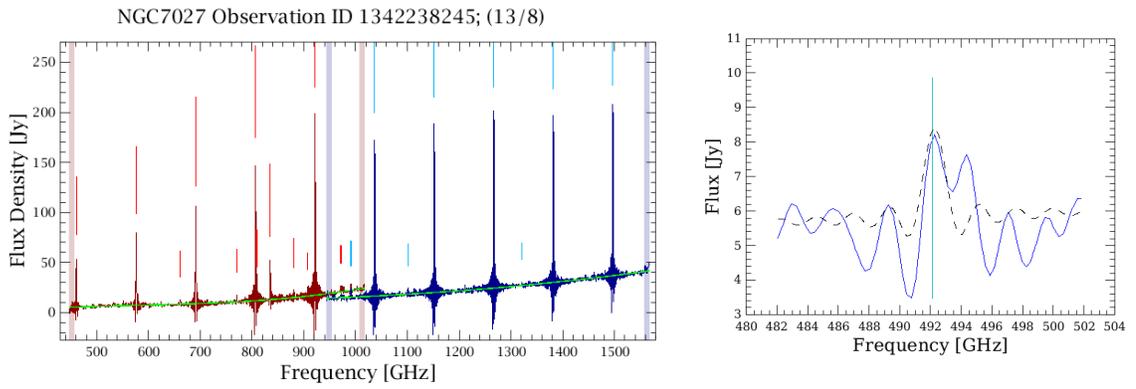


Figure C.8: $[\text{CI}]^3P_1 - {}^3P_0$ snr = 5.002306

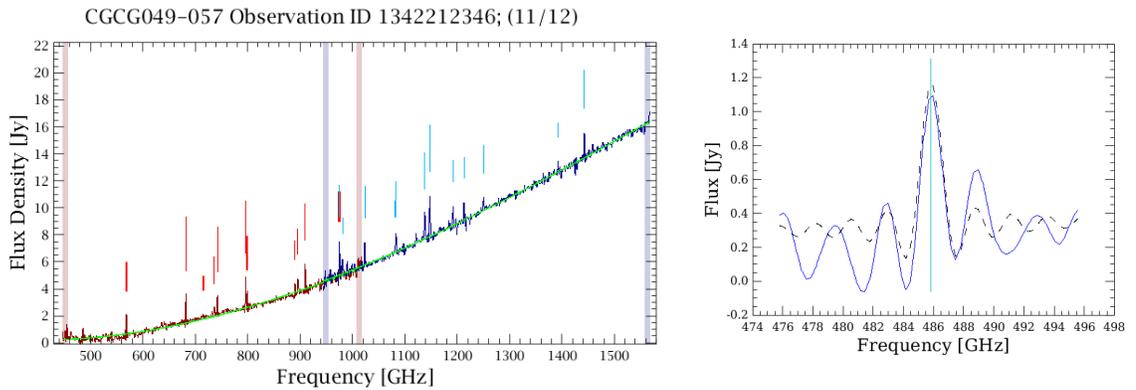


Figure C.9: $[\text{CI}]^3P_1 - {}^3P_0$ snr = 6.959381

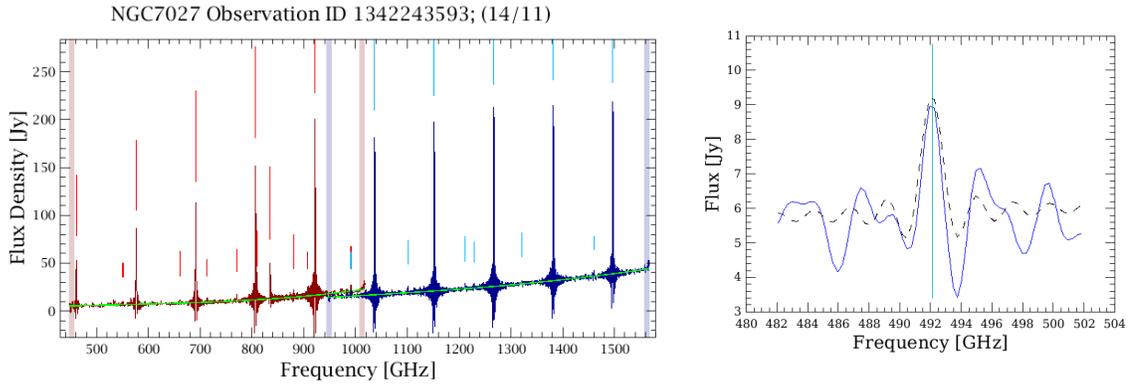


Figure C.10: $[\text{CI}]^3P_1 - {}^3P_0$ snr = 5.755870

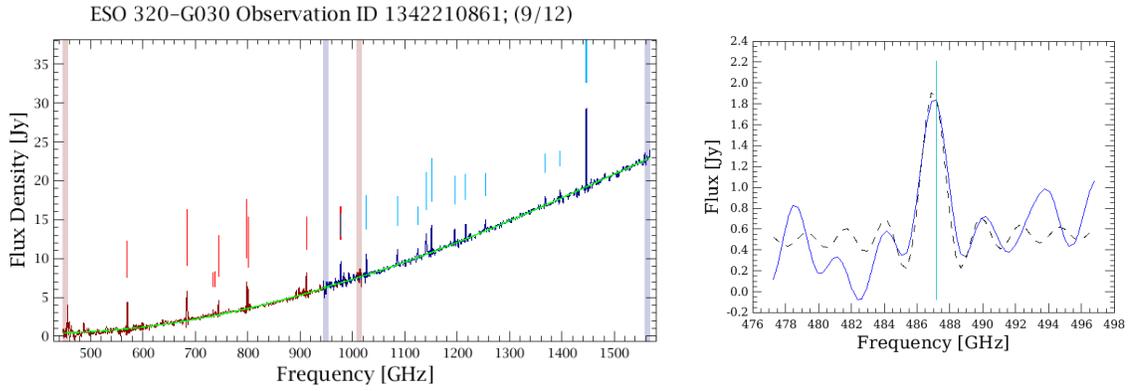


Figure C.11: $[\text{CI}]^3P_1 - {}^3P_0$ snr = 6.509541

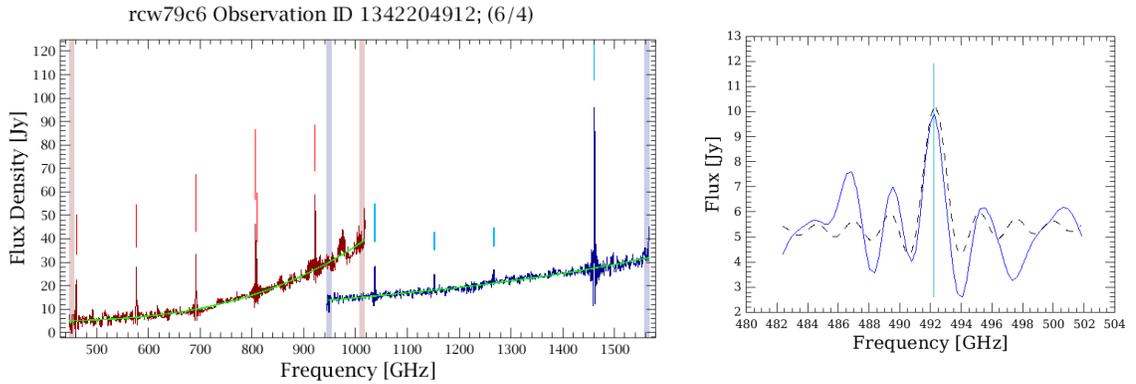


Figure C.12: $[\text{CI}]^3P_1 - {}^3P_0$ snr = 7.247980