

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

The Undiscovered CO: Charting the Molecular Gas of the Early Universe

Permalink

<https://escholarship.org/uc/item/9n8963t6>

Author

Keating, Garrett Kent

Publication Date

2016

Peer reviewed|Thesis/dissertation

The Undiscovered CO: Charting the Molecular Gas of the Early Universe

By

Garrett Kent Keating

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Astrophysics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Carl E. Heiles, Chair

Professor Mariska Kriek

Professor William Holzappel

Spring 2016

The Undiscovered CO: Charting the Molecular Gas of the Early Universe

Copyright 2016
by
Garrett Kent Keating

Abstract

The Undiscovered CO: Charting the Molecular Gas of the Early Universe

by

Garrett Kent Keating

Doctor of Philosophy in Astrophysics

University of California, Berkeley

Professor Carl E. Heiles, Chair

Molecular gas, observed through tracers such as CO rotational transitions, is a vital component of galactic evolution and star formation. Recent detections of the CO molecule in massive galaxies at high redshift have demonstrated its existence in the early Universe, and have motivated its use as a means of exploring large-scale structure and as a probe of galaxy evolution in the early Universe. But many questions about molecular gas and the evolution of galaxies in the early Universe still remain: its distribution at high redshift understood is so poorly that theoretical models of the mean abundance of CO for the first several billion years of cosmic history span orders of magnitude. Direct detection of molecular gas in galaxies at these redshifts have only found the largest and most luminous of galaxies in the early Universe, whereas the bulk of the molecular gas is expected to be in the unseen masses of smaller galaxies. While difficult to detect individually, these smaller galaxies are likely detectable as an integrated ensemble with the technique of “intensity mapping”. This technique, similar to those employed by HI epoch of reionization experiments, utilizes measurements of different 3D Fourier modes to construct a power spectrum.

In this thesis, I present results from the CO Power Spectrum Survey (COPSS), an intensity mapping experiment performed with the Sunyaev Zel’dovich Array (SZA). I present power spectrum constraints from the first and second phase of this project, utilizing an archival dataset (covering 44 fields in 1400 hours observing time) and a 5000-hour observing campaign (covering 19 fields) with the SZA. With these data, we are capable of observing CO(1-0) emission arising from $z = 2.3-3.3$, surveying a volume of more than ten million cubic megaparsecs. With this measurement, we place the first-ever constraints on the CO autocorrelation power spectrum, and place constraints on the CO(1-0) galaxy luminosity function and the cosmic molecular gas density at $z \sim 3$.

I also present a series of simulations designed to probe the impact of cosmic variance, continuum foregrounds, and systematic errors on our measurement with the SZA. I will also present simulations designed to probe the challenges that will face intensity mapping

experiments with the Yuan-Tseh Lee Array (YTLA). Finally, I will present some preliminary work on CO(3-2) and [CII] intensity mapping experiments.

For Grandma, Granddad, and Papa,
who are hopefully looking down while I'm staring up.

Contents

List of Figures	v
List of Tables	vii
Acknowledgments	viii
1 Introduction	1
1.1 The Cold Gas Contents of Galaxies	1
1.2 The Evolution of Cold Gas in the Early Universe	1
1.3 The Limits of Direct Detection	3
1.4 The Alternative Approach of “Intensity Mapping”	4
2 The Intensity Mapping Technique	8
2.1 Mathematical Formalism	8
2.1.1 Interferometric Measurements	10
2.1.2 Single Dish Measurements	12
2.2 Relationship to the Luminosity Function of Galaxies	13
2.2.1 The Cluster-Power Component	14
2.2.2 The Shot-Power Component	15
2.3 Spectral Line Confusion	16
2.4 Estimates of the Impact of Continuum Foregrounds	17
3 Data Analysis for Interferometric Power Spectrum Measurements	21
3.1 Data Processing Overview	21
3.2 Pipeline Overview	23
3.3 Power Spectrum Analysis	28
3.3.1 Ground Subtraction	32
3.3.2 Point Source Contamination	35
3.3.3 Jackknife Tests	36
4 The CO Power Spectrum Survey (COPSS)	41
4.1 The Sunyaev-Zel’dovich Array (SZA)	42
4.2 Observations	44

4.2.1	COPSS I	44
4.2.2	COPSS II	45
4.3	Results	48
4.3.1	CO Power Spectra	48
4.3.2	Trim Tests	49
4.4	Discussion	51
4.4.1	Constraints on the CO Power Spectrum	51
4.4.2	Constraints on A_{CO} and σ_{CO}	53
4.4.3	Constraints on the CO Luminosity Function	58
4.4.4	Constraints on Cosmic Molecular Gas Abundance	59
5	Simulations	64
5.1	Pipeline Verification	64
5.2	Impact of Cosmic Variance	66
5.3	Continuum Point Sources	68
5.4	Instrumentation Simulations for Future Experiments	71
5.4.1	Measuring the Impact of Platform Deformation	72
5.4.2	Measuring the Impact of Pointing Errors	74
5.4.3	Measured Effects versus Source Declination	77
5.4.4	Sensitivity Estimates	77
6	Preliminary Results from Future Intensity Mapping Experiments	80
6.1	The ASIAA Intensity Mapping of CO Experiment	80
6.2	The [CII] Power Spectrum Survey	81
7	Conclusion	86
A	Sensitivity Estimates for Intensity Mapping Experiments	88
A.1	Interferometric Experiments	88
A.2	Single-Dish Measurements	90
B	The Intensity Mapping Pipeline Software (IMPS) Package	91
B.1	Flagging	91
B.2	Calibration	93
B.3	Imaging and Power Spectrum Production	95
B.4	Data structure description	99
C	A Study of Calibration Challenges for a 15 GHz Intensity Mapping Interferometer	106
C.1	DACOTA System Overview	106
C.2	Astronomical Calibrators	107
C.2.1	Phase and Amplitude Calibration	107
C.2.2	Bandpass Calibration	107

C.2.3 The Sun and Moon as Calibrators	108
C.3 Unconventional Calibrator Choices	109
C.4 Conclusion	110
C.5 List of GEO Satellites	110
Bibliography	113

List of Figures

1.1	A molecular cloud of gas	2
1.2	Combined Array or Research in Millimeter-wave Astronomy (CARMA)	4
1.3	Comparison of cosmic volumes probed between blind search and intensity mapping	6
2.1	PSF of the SZA Primary Beam in the (u, v, η) domain	12
3.1	An example flagging table from the COPSS analysis	22
3.2	An example gains solutions from the COPSS analysis	24
3.3	An example SEFD correction solutions from the COPSS analysis	25
3.4	A diagram of the processing pipeline for COPSS	26
3.5	Flux monitoring history of Jupiter within the COPSS II data set	27
3.6	Measured power versus integration time	28
3.7	The covariance matrix for the signal of interest for the Q2343 field	31
3.8	An image showing ground contamination within the COPSS data	33
3.9	Power spectrum estimates of ground contributions	34
3.10	The RMS of the bandpass solutions for the COPSS II data set	35
3.11	Jackknife results from the analysis of archival SZA data	39
3.12	Jackknife results from the full COPSS data set	40
4.1	The Sunyaev-Zel'dovich Array (SZA)	42
4.2	The SZA array layout for the COPSS survey	43
4.3	The uv coverage for a sample field in the COPSS data set	44
4.4	Thumbnail images of the COPSS I fields	47
4.5	Thumbnail images of the COPSS II fields	48
4.6	The COPSS I power spectrum result, in the form $\Delta^2(k)$	50
4.7	The COPSS II power spectrum result, in the form $\Delta^2(k)$	51
4.8	The COPSS II power spectrum result, in the form of $P(k)$	52
4.9	Power spectrum measurement as a function of redshift	53
4.10	Power measured within individual groups of fields	54
4.11	Constraints on A_{CO} as a function of M_{min}	55
4.12	Constraints on A_{CO} versus σ_{CO}	56
4.13	Fitting results for CO luminosity parameters	61
4.14	Fitting results for CO luminosity function versus theoretical models	62

4.15	Constraints on the cosmic molecular gas density	63
5.1	Simulations Measuring the Fidelity of our Analysis Software	65
5.2	Cosmic Variance Simulation Results	67
5.3	Results of simulations to probe point source contributions	69
5.4	The effect of platform deformation-induced delay errors	73
5.5	The effect of platform deformation-induced pointing errors	74
5.6	The effect of antenna-based pointing errors	75
5.7	The effect of platform pointing errors	75
5.8	Measurements of Instrumental Effects versus Source Declination	76
5.9	Sensitivity estimates for a survey with the YTLA	78
6.1	The Yuan-Tseh Lee Array (YTLA)	80
6.2	First YTLA power spectrum measurement	82
6.3	Line confusion at millimeter wavelengths	83
6.4	First constraints on the CII power spectrum	85

List of Tables

2.1	Spectral Line Table	17
3.1	Jackknife Results	38
4.1	Lead fields for each of the COPSS I groups	45
4.2	COPSS II Observing Fields	46
4.3	CO Power Spectrum Results	49
B.1	Flagging statistics for the COPSS I dataset.	93
C.1	List of viable DACOTA Prototype gain calibrators	107
C.2	List of viable DACOTA Prototype bandpass calibrators	108
C.3	List of potential K-band GEO satellite calibrators	112

Acknowledgments

After six years as a graduate student, I am amazed that one of the more difficult parts of preparing for and writing my thesis has been the writing of this acknowledgments section – my first attempt reached ten pages of text before realizing that I may be forced to cut down the length. Nevertheless, the work presented in this dissertation could not have been completed without the support and counsel of a great many people.

First, I must thank my parents and family, both for their support and their encouragement with my sometimes meandering path towards astronomy, especially in those times of solitude during my many, many months of observing. I also thank the many exceptional teachers I’ve had over the years, particularly for the critical role they played in fostering my curiosity about how the world works. Elliot, Jon, and Kevin – thanks for helping me get through high school and college; there’d likely be no “Dr. Beanie” without you. I also extend my appreciation to the many other friends (including the gentlemen of PKP) who helped keep me motivated over these many, many years of school.

The primary project of my thesis – the CO Power Spectrum Survey (COPSS) – was named in honor of the many exceptional past and present members of UC Berkeley Police Department (with whom I had the honor and privilege of working with while serving as a Community Service Officer), and in remembrance of Cpl. Timothy Zuniga, who in addition to being a friend and mentor, taught me that one need not be an academic in order to be a scholar.

Throughout most of college, I had planned on a career in law enforcement, never quite finding a niche in astronomy that excited me enough to pursue it long-term. That all changed upon enrolling in the undergraduate radio lab course (AY121), taught by Carl Heiles, who I’ve been grateful to call my advisor these many years. I will never forget my excitement at having detected the 21 cm line of neutral hydrogen, with instruments that I could touch and manipulate. The experience left me with a deep appreciation for radio astronomy, instrumentation, and clever (but statistically sound!) methods for analyzing data. He has my deepest gratitude for that inspiration, and I hope that we will continue our conversations about deciphering the secrets of the ISM (and its relevant lessons for early galaxies).

After college, I miraculously found myself working with a tremendously talented group of individuals at the Hat Creek Radio Observatory, where I got to immerse myself in all things radio astronomy. I learned a tremendous amount during my stay in Hat Creek, and I owe thanks to the many people who taught me a great deal about astronomy and life: Rick and

Dave (who donned the nickname “Karto” upon me), Leo (who hired me for my first research position), Billy, Jeanette, Colby, Matt, Mel, Calvin, Don, Jack, Jill, Gerry, Jon, Sam, Steve, Casey, Peter, Chat, Katey, Amber, Susie, Mark, Mike, Jeff, Carol, and Gary (whose hard work and extra hours helped keep the ATA running in some very lean times, and whose friendship made my time at CARMA much more enjoyable).

Hat Creek was also the place where I was fortunate enough to meet Geoff Bower, who has been an advisor to me for 4 different stages of my career (undergraduate, post-baccalaureate, graduate, and soon post-doctoral), and who first motivated me to delve into the world of intensity mapping and cosmology. His creativity and drive have helped push me to meet research challenges that I might have otherwise thought insurmountable (“detection, detection, detection!”), and I am grateful for the guidance and opportunities he provided that helped transform me from a somewhat “less-than-stellar” undergraduate to a skilled radio astronomer.

There were many people who helped make my second stay at Berkeley an amazing one – so many that I could easily fill a page with names alone. To my fellow classmates – it was a pleasure getting to know you these past many years, battling over problem sets, chatting over Sunday dinners, and occasionally eating ice cream and various grilled meats. To the amazing support staff (especially Dexter, Bill, Rayna and Nina) – thank you for the immense amount of work you did to help make the research in this thesis possible. To the Berkeley PAPER/HERA group – thank you for welcoming me into the fold, and I hope there will be many pizza lunches in the future, where we will continue to argue over intensity mapping, Fourier transforms, and photon polarization states. To the CARMA family and the SZA group (especially Dick, Nikolaus, John, Dave, James, Dave, John, Erik, Stephen, Gary, Cecil, Mark, and Mikey) – I will always be thankful for your help in making COPSS II happen, and for making Cedar Flat a home away from home. To the Astro Lab kids – the opportunity to teach so many of you has been one of the most fulfilling (and often hilarious) experiences of my life; I hope I continue to see many of you (in person and in print) in the years to come.

Of the many people I met during graduate school, I am particularly glad to have crossed paths with Dan Marrone, who over the course of the last 5 years has become my “adoptive advisor”. Some of the most intellectually rewarding moments over these years have come from conversations in the living room of his home, where he was kind enough to host me during my many trips to Tucson. His patience, thoughtfulness, and discipline are among the many of his qualities I hope to emulate as a scientist and as a person. He is a scholar and a gentleman of the highest caliber, and I take great pride in being able to call myself his student (although I take less pride in having accidentally allowed his son to color on his couch with a permanent marker).

Finally, I cannot thank enough my wife, Francesca, who has been an amazing partner, friend, and scientific compatriot these last six years of graduate school. In addition to being a talented X-ray astronomer and a fantastic editor, she is my most favorite person to talk about science (and many other topics), in part because she is both incredibly knowledgeable and extraordinarily inquisitive. It has been through her support that I have been able to face the most difficult and challenging moments of my graduate student experience. Of all

the things I am grateful for, I'm most glad that I get to continue this adventure with her.

With the exclusion of Chapter 6 (which presents preliminary work of future efforts), significant material from all chapters was originally published in [Keating et al. \(2015\)](#) and [Keating et al. \(2016\)](#). This dissertation was typeset using the [ucastrothesis](#) L^AT_EX template.

Chapter 1

Introduction

*“Ad astra per aspera”
(A rough road leads to the stars)*

–Inscription of the
Apollo 1 Memorial Plaque

1.1 The Cold Gas Contents of Galaxies

Molecular gas serves a vital role in star formation as the natal material from which stars form. Though the main constituent of this gas is molecular hydrogen, the H_2 molecule lacks a permanent dipole moment, making it a poor radiator of energy and hence difficult to observe. Traditionally, the CO molecule – the next most abundant molecule after H_2 – has been used as a tracer of molecular hydrogen (e.g., [Wilson et al. 1970](#); [Young & Scoville 1982](#); [Young et al. 1995](#); [Regan et al. 2001](#); [Bolatto et al. 2013](#)). Unlike H_2 , the CO molecule possesses a permanent dipole moment, with an excitation temperature of $T_{\text{ex}} \approx 5.5\text{K}$ for the $J=1 \rightarrow 0$ transition, making it ideal for probing the cold, dense gas of molecular clouds.

Within the local Universe, the CO luminosity (L_{CO}) of galaxies – and by extension, their molecular gas mass – shows strong correlation with far-infrared luminosity (L_{FIR}), $\text{H}\alpha$ emission and $\text{Ly}\alpha$ emission; all are strong indicators of star formation within these galaxies ([Downes et al. 1993](#); [Solomon et al. 1997](#); [Kennicutt 1998](#)). This relationship between molecular gas and star formation rates (SFR) is typically referred to as the Kennicutt-Schmidt (KS) relationship ([Schmidt 1959](#); [Kennicutt 1998](#)), and demonstrates a deep connection between the abundance of molecular gas and the formation of stars.

1.2 The Evolution of Cold Gas in the Early Universe

The gas contents of galaxies play a powerful role in shaping the evolution of galaxies. In the early Universe, large gas reservoirs dominated the baryonic mass of galaxies ([Tacconi](#)



Figure 1.1: The “Pillars of Creation” – one of the most well known molecular clouds within the Milky Way Galaxy, inside the Eagle Nebula. The cloud itself is slowly being evaporated away by the intense ultraviolet light originating from young stars, and may have already been destroyed by a supernova (Flagey et al. 2011). These processes are a part of a longer cyclical process, in which cool gas condenses down into dense molecular clouds, eventually forming stars that will disrupt the cloud itself. *Image Credit: NASA*

et al. 2010), and fueled a rapid increase in cosmic star formation, peaking at a rate 10 times higher than what is observed locally (Madau & Dickinson 2014; Hopkins & Beacom 2006). Most of our knowledge of these early galaxies has come from studying stellar light and emission lines from the hot, ionized gas of the interstellar medium (ISM). However, it is the cold molecular gas that provides the natal material from which stars form. Understanding the nature and evolution of this cold gas is thus crucially important for understanding star formation in the early Universe.

Of the several tracers available for studying cold gas in the local Universe, the bright transitions of CO are particularly suitable for examining distant, high-redshift objects. The CO molecule – the second most abundant molecule in the ISM – is typically found in clouds of molecular hydrogen (e.g., Bolatto et al. 2013). Sensitive instruments such as the VLA, ALMA and PdBI have made it possible to probe cool ISM of massive galaxies out to redshifts

as high as $z = 6.4$ (e.g., [Walter et al. 2003](#); [Wang et al. 2010](#); [Riechers et al. 2013](#); [Lentati et al. 2015](#)). However, these extraordinary objects are not likely to be characteristic of the overall population of star-forming galaxies in the early Universe, made up primarily of smaller and less luminous systems ([Smit et al. 2012](#)).

It is unknown how the connection between CO, bulk molecular gas, and star formation evolves over cosmic time. Observations of distant star-forming galaxies suggest that L_{CO} -SFR correlation persists up to $z \lesssim 2$ ([Tacconi et al. 2013](#)), implying that the correlation between CO abundance and the amount of molecular gas available for star formation has remained relatively unchanged in the several billion years following peak of cosmic star formation ([Hopkins & Beacom 2006](#)). Prior to this epoch, however, early galaxies (with their short star formation histories) may not contain enough metals to form an appreciable amount of CO, or may possess too little dust to shield the CO from dissociation by UV starlight ([Genzel et al. 2012](#); [Bolatto et al. 2013](#)). Some models of the conversion factor between CO luminosity and molecular gas mass (α_{CO}), predict a steep power-law relationship between α_{CO} and the gas metallicity of galaxies (e.g., [Israel 1997](#)). Should their predictions hold true, many high-redshift galaxies may lack significant CO emission, despite the presence of molecular gas ([Wolfire et al. 2010](#); [Muñoz & Furlanetto 2014](#)). Other theoretical work (e.g., [Glover & Mac Low 2011](#); [Obreschkow et al. 2009b](#)) suggests that CO is not so strongly affected by the lower metallicity and dust masses of early galaxies, offering a more optimistic outlook for CO as a tool for exploring molecular gas at high redshift.

1.3 The Limits of Direct Detection

The makeup of the molecular gas content of star-forming galaxies in the early Universe is currently an active area of observational research. Recently, [Decarli et al. \(2014\)](#) and [Walter et al. \(2014\)](#) used the Plateau de Bure Interferometer (PdBI) to make a very deep (100 hour integration time) observation of a portion of the GOODS-N field ([Dickinson et al. 2003](#)), and probed the CO luminosity function at $z \sim 3$ to a limit of $L'_{\text{CO}} \gtrsim 10^{10} \text{ K km s}^{-1} \text{ pc}^2$ for the $J=3 \rightarrow 2$ rotational transition of CO. For $z \sim 3$, this line luminosity limit corresponds to galaxy SFRs greater than $\sim 1.2 \times 10^2 M_{\odot} \text{ yr}^{-1}$ ([Tacconi et al. 2013](#)). Clearly, such studies are limited to massive galaxies that are rapidly forming stars and miss the lower luminosity and/or less massive systems that are expected to make up a large fraction of star-forming galaxies ([Obreschkow et al. 2009b](#); [Lagos et al. 2011](#); [Smit et al. 2012](#); [Bouwens et al. 2012](#); [Sargent et al. 2014](#)).

Unfortunately, the molecular line emission from “normal” (i.e., low-mass) star-forming galaxies at $z \gtrsim 3$ is incredibly faint. Direct detection of individual low-mass galaxies is an observationally expensive proposition – a useful illustration is the first “science goal” of ALMA: “The ability to detect spectral line emission from CO or [CII] in a normal galaxy like the Milky Way at a redshift of $z \sim 3$, in less than 24 hours.”¹ While measurements of such objects are vital for characterizing normal star-forming galaxies, individual detection of

¹<https://almascience.eso.org/about-alma/alma-basics>



Figure 1.2: The Combined Array for Research in Millimeter-wave Astronomy (CARMA), formerly located in Cedar Flat, CA. *Image Credit:* Jim Keating

large numbers of low-mass galaxies is likely to be unfeasible with the current generation of radio instruments. In part, this has left us with a very weak understanding of the cool ISM in early galaxies, and is a key limitation in using galaxy formation simulations to understand even the most luminous of galaxies (e.g., [Hayward et al. 2011](#)).

1.4 The Alternative Approach of “Intensity Mapping”

One alternative method to exploring the ISM of more typical galaxies is through a technique commonly referred to as “intensity mapping”, where the emission from a multitude of galaxies (i.e., thousands or millions) over a wide range of luminosities are detected in aggregate as larger-scale fluctuations in the mean line intensity. The intensity mapping method is a valuable tool for charting the growth of large scale structure and of the gas contents of galaxies. This method has been the subject of numerous recent theoretical investigations (e.g., [Righi et al. 2008](#); [Visbal & Loeb 2010](#); [Carilli 2011](#); [Pullen et al. 2013](#); [Breyse et al. 2014](#); [Mashian et al. 2015](#); [Li et al. 2016](#)).

Some of the primary strengths of intensity mapping include:

- **Intensity mapping experiments require only modest point-source sensitivity:** Unlike direct detection experiments, intensity mapping measurements do not require significant point source sensitivity (i.e., lots of collecting area) in order to make a detection.

- **Intensity mapping measurements are complete:** As intensity mapping measures fluctuations in the mean line intensity over large volumes, all objects with spectral line emission will contribute to the measured power (although these contributions will be weighted in some way, as discussed further in Chapter 2).
- **Intensity mapping measurements are resilient against redshift errors:** As intensity mapping experiments do not require a priori knowledge of sources within a given volume, they are not affected by errors in redshift measurement that would impact targeted surveys of individual objects.

Shown in Figure 1.3 are some comparisons between a few select intensity mapping and direct detection experiments. Intensity mapping experiments make up for an inability to detect individual sources (i.e., a weak point source sensitivity) by measuring the emission arising from very large volumes, effectively providing a measurement of the incoherently averaged flux arising from within the volume probed. In this way, intensity mapping experiments are not too dissimilar from “stacking analysis”, where voxels corresponding to the positions and redshifts of known galaxies (such as those observed at other wavelengths) are averaged together to measure the mean emission of several otherwise undetectable sources. While stacking analysis is a useful tool for digging below the nominal noise threshold of a given measurement, it is prone to the various selection effects that one nominally desires to protect against during a blind survey.

As intensity mapping experiments have significant strengths in surveying large numbers of galaxies, it can serve as a complementary tool to the type of case-study analysis typically performed with large radio arrays, providing an independent ensure that the galaxies studied in direct detection efforts are reflective of the true population of galaxies within the high redshift Universe. Several theoretical models suggest that detection of such a signal is within reach of existing instruments. Intensity mapping experiments are well-suited for data sets with large survey volumes, requiring only modest point-source sensitivity to detect an aggregate signal. As such, the Sunyaev-Zel’dovich Array (SZA) – a 3.5m-diameter \times 8-element subset of the Combined Array for Research in Millimeter-wave Astronomy (CARMA, shown in Figure 1.2) – is an instrument capable of undertaking such an experiment.

The Sunyaev-Zel’dovich Array (SZA), an 8-element interferometer of 3.5 meter dishes with 1 cm receivers, is well matched to performing a CO intensity mapping experiment. The SZA is capable of observing the $J=1\rightarrow 0$ rotational transition of CO, which we will herein refer to as CO(1–0), at a redshift range of $z = 2.3\text{--}3.3$, with greatest sensitivity to comoving size scales of $0.5\text{--}2 h^{-1}$ Mpc. The wide field of view and spectral bandwidth of the SZA make capable of surveying large volumes of the high redshift Universe, exceeding the largest volumes probed for high-redshift CO emission with the VLA and ALMA in as little as a single pointing. In this work, we describe and present results from a multi-year effort with the SZA to probe the molecular gas of the early Universe, with the goal of placing constraints on the population of CO-luminous galaxies, that might otherwise go undiscovered (or otherwise be too observationally expensive to explore) with traditional direct detection techniques.

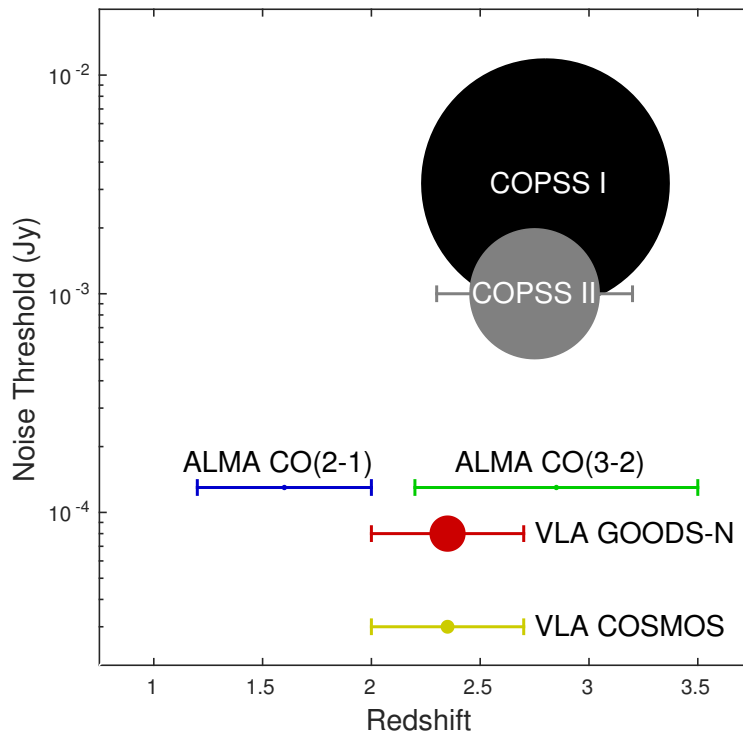


Figure 1.3: A comparison of the volumes probed and sensitivity of several deep blind-search efforts (for CO emitters) with the VLA and ALMA, versus the intensity mapping efforts with CARMA and the SZA described herein. The COPSS I measurement refers to a survey performed between 2005-2008 as part of an experiment to measure arcminute-scale anisotropies in the cosmic microwave background (CMB). The COPSS II measurement refers to a large-scale CARMA intensity mapping survey to search for $z \sim 3$. The analysis of both data sets is discussed further in Chapter 4.

This thesis is structured as follows: Chapter 2 discusses the method of intensity mapping in detail, and provides some estimates for potential contaminants within our measurements. Chapter 3 discusses some of the analysis tools developed for intensity mapping with interferometers. Chapter 4 discusses the CO Power Spectrum Survey (COPSS), a focused effort with the SZA to measure the CO power spectrum at $z \sim 3$. Chapter 5 discusses simulation work performed for in support of the COPSS analysis and for future CO intensity mapping experiments. Chapter 6 discusses preliminary results from upcoming experiments in intensity mapping with the Yeah-Tseh Lee Array (YTLA) and ALMA. Finally, we provide some conclusionary remarks in Chapter 7.

In addition to primary material, also included are three supplemental chapters on work related to the research presented here. Appendix A provides some simple tools for estimating the sensitivity of intensity mapping experiments for interferometers and single dish exper-

iments. Appendix [B](#) discusses the Intensity Mapping Pipeline Software (IMPS) package, developed for general purpose intensity mapping experiments with interferometers. Appendix [C](#) presents the result of a study of calibration challenges for a proposed 15 GHz interferometer for designed for future intensity mapping experiments.

Chapter 2

The Intensity Mapping Technique

Before proceeding with our measurement, it is important to consider the mathematical framework required for translating our native measurements – which for an interferometer, are typically measurements of flux density in units of Janskys (Jy) – into power spectrum measurements – typically in temperature squared units. Equally important is understanding how a power spectrum measurement is connected to the underlying population of luminous galaxies within a given volume. In this chapter, we seek to provide this framework, in order to better understand what it is exactly we measure in our power spectrum measurement, and to show how it compares to direct detection efforts.

Because intensity mapping measurements are entirely statistical in nature, an additional source of concern is how foregrounds and spectral lines other than the one of interest may affect our measurement. Before proceeding forward with a large observational survey or data analysis project, it is therefore useful to consider what the leading sources of contamination may be, and how significantly they may affect our result.

This chapter is structured as follows: Section 2.1 provides the mathematical framework (for both interferometric and single-dish measurements) for intensity mapping experiments. Section 2.2 describes how the power spectrum is connected to the galaxy luminosity function for a given spectral line. Section 2.3 discusses what other spectral lines (if any) may be a source of confusion or concern when attempting a power spectrum measurement of CO(1-0) emission. Finally, Section 2.4 provides some simple theoretical estimates on the impact from sources of continuum emission on our measurement.

2.1 Mathematical Formalism

The goal of intensity mapping is to use the two-point autocorrelation function, $\xi(\mathbf{x})$, to probe the underlying population of galaxies bearing molecular gas (through the tracer CO molecule) within some representative volume, V_z , of the Universe at redshift z . The final analysis product is a power spectrum, $P(k)$, which along with its dimensionless counterpart,

$\Delta^2(k)$, is related to the two-point autocorrelation function by

$$\begin{aligned}\Delta^2(k) &\equiv \frac{k^3}{2\pi^2} P(k), \\ &= \frac{k^3}{2\pi^2} \langle \mathcal{F}_{x \rightarrow k}(\xi(\mathbf{x})) \rangle_{\mathbf{k}, \mathbf{k}=k^2},\end{aligned}\quad (2.1)$$

where $\mathcal{F}_{x \rightarrow k}$ is the Fourier transform from configuration to Fourier space, and \mathbf{k} is the vector wavenumber (of magnitude k) in Fourier space. The two point autocorrelation function is further defined as:

$$\xi(\mathbf{x}) = \frac{1}{W(\mathbf{x}) \cdot W(\mathbf{x})} \left(\frac{c^2}{2k_B \nu^2} \right)^2 ((W(\mathbf{x})I(\mathbf{x}')) \star (W(\mathbf{x})I(\mathbf{x}))) \quad (2.2)$$

where $I(\mathbf{x})$ is the specific intensity observed at a given frequency, ν . $W(\mathbf{x})$ is the spatial windowing function applied over the volume in question. $W(\mathbf{x}) \cdot W(\mathbf{x})$ is the normalization term of the autocorrelation function, which is equal to the effective volume probed, V_z . The Fourier dual of the autocorrelation function is $\tilde{\xi}(\mathbf{k})$, such that $P = \langle \tilde{\xi}(\mathbf{k}) \rangle$. The function $\tilde{\xi}(\mathbf{k})$ can then be expressed as

$$\tilde{\xi}(\mathbf{k}) = \frac{1}{V_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left((\tilde{W}^*(\mathbf{k}) * \tilde{I}^*(\mathbf{k})) (\tilde{W}(\mathbf{k}) * \tilde{I}(\mathbf{k})) \right), \quad (2.3)$$

where $\tilde{I}(\mathbf{k})$ is the Fourier dual of $\tilde{I}(\mathbf{x})$. We map between comoving physical size, $\mathbf{r} = (r_x, r_y, r_z)$, and the native observing units of $(l, m, \Delta\nu)$, where l and m describe the angular position (in units of radians) and $\Delta\nu$ is the change in line frequency due to expansion of the Universe, with the following expressions:

$$l = \frac{r_x}{D_M(z)}, \quad m = \frac{r_y}{D_M(z)}, \quad \Delta\nu = \frac{H_0 E(z) \nu_{rest}}{c(1+z)^2} r_z. \quad (2.4)$$

In Equation 2.4, $D_M(z)$ is the comoving radial distance for redshift z , ν_o is the rest frequency of the line transition, H_0 is the current Hubble parameter and $E(z)$ is the dimensionless Hubble parameter. To evaluate $\tilde{\xi}(\mathbf{k})$, one converts (u, v, η) , the Fourier dual of $(l, m, \Delta\nu)$, to (k_x, k_y, k_z) :

$$u = \frac{k_x X}{2\pi}, \quad v = \frac{k_y X}{2\pi}, \quad \eta = \frac{k_z Y}{2\pi},$$

using

$$X = D_M(z), \quad Y = \frac{c(1+z)^2}{H_0 E(z) \nu_{rest}}. \quad (2.5)$$

In Equation 2.5, we have included the variables X and Y (defined in Parsons et al. 2012) for brevity only, with units of Mpc rad^{-1} and Mpc Hz^{-1} respectively. Rewriting Equation 2.3 and making the proper variable substitutions, we derive the following expression:

$$\tilde{\xi}(\mathbf{k}) = \frac{X^4 Y^2}{V_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left((\tilde{W}^*(u, v, \eta) * \tilde{I}^*(u, v, \eta)) (\tilde{W}(u, v, \eta) * \tilde{I}(u, v, \eta)) \right). \quad (2.6)$$

2.1.1 Interferometric Measurements

With our simplified expression, we now turn our attention to the interferometer response to determine its relation to $\tilde{I}(u, v, \eta)$, and thus the measurement of $\tilde{\xi}(\mathbf{k})$. Under the flat-sky approximation, the interferometer response is given by

$$\begin{aligned}\mathcal{V}(u, v, \nu) &= \mathcal{F}_{(l,m) \rightarrow (u,v)}(I(l, m, \nu) \cdot A_\nu(l, m)) \\ &= \tilde{I}(u, v, \nu) * \tilde{A}_\nu(u, v),\end{aligned}\tag{2.7}$$

where measurements for the interferometer response, $\mathcal{V}(u, v, \nu)$, are commonly referred to as “visibilities” (Thompson et al. 1986). \mathcal{F} in Equation 2.7 is the Fourier transform operator, which in this case transforms between (l, m) and (u, v) domains, and $A_\nu(l, m)$ is the primary beam response of the telescope at a given frequency ν and position in the sky (l, m) . We wish however to move from (u, v, ν) to (u, v, η) and create what we will refer to as “delay-visibility”. This requires performing one last Fourier transform over some window in frequency $\mathcal{B}(\nu)$;

$$\begin{aligned}\tilde{\mathcal{V}}(u, v, \eta) &= \mathcal{F}_{\nu \rightarrow \eta}(V(u, v, \nu) \cdot \mathcal{B}(\nu)), \\ &= \left(\tilde{I}(u, v, \eta) * A_\nu(u, v)\right) * \tilde{\mathcal{B}}(\eta) = \tilde{I}(u, v, \eta) * \tilde{W}(u, v, \eta).\end{aligned}\tag{2.8}$$

The delay-visibility in Equation 2.8 is the convolution of \tilde{I} and our windowing function, which we have defined as $\tilde{W}(u, v, \eta) = \tilde{A}_\nu(u, v) * \tilde{\mathcal{B}}(\eta)$, based on the windowing functions created by the primary beam pattern and the bandwidth window. We define $\Omega_{B,\nu} = \int A_\nu d\Omega$ to be the solid angle of the primary beam of the telescope, where A_ν is the primary beam pattern (and the Fourier dual of the aperture function, \tilde{A}), for a given frequency ν . Nominally both A and Ω are frequency dependent, but the frequency range of each window is small enough that both $\Omega_{B,\nu}$ and A_ν can be treated as constant over the volume of interest. We also define $B_z = \int \mathcal{B}(\nu) d\nu$ as the effective bandwidth for our observation of the volume of interest, where $\mathcal{B}(\nu)$ is the frequency windowing function. In practice, this windowing function is generally equal to unity over the correlator window, hence for our measurement B_z is equal to the correlator window bandwidth. We will assume that our range in redshift over which our volume V_z subsides is small enough that X and Y remain constant. With this assumption we can further simplify our expression by substituting V_z with defined instrument parameters. The effective volume can be defined as

$$V = \int W \cdot W = X^2 Y \int A\mathcal{B} \cdot A\mathcal{B} d\Omega d\nu = X^2 Y B_z \Omega_{B,\nu} / 2.\tag{2.9}$$

In Equation 2.9, we have made use of the fact that for a Gaussian beam, $\int A^2 = \Omega_B / 2$, and have assumed a uniform weighting in frequency (i.e., $\mathcal{B} = 1$ over the frequency range of the window, otherwise $\mathcal{B} = 0$), as it is the weighting scheme used in our analysis. We can now simplify Equation 2.6, and come up with an expression for Equation 2.1 that contains only

observables:

$$P(k) = \frac{2X^2Y}{\Omega_B B_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left\langle \tilde{\mathcal{V}}^*(u, v, \eta) \tilde{\mathcal{V}}(u, v, \eta) \right\rangle_{\mathbf{k} \cdot \mathbf{k} = k^2}, \quad (2.10)$$

$$\Delta^2(k) = \frac{k^3 X^2Y}{\pi^2 \Omega_B B_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left\langle \tilde{\mathcal{V}}^*(u, v, \eta) \tilde{\mathcal{V}}(u, v, \eta) \right\rangle_{\mathbf{k} \cdot \mathbf{k} = k^2}. \quad (2.11)$$

A subtle feature of Equations 2.10 and 2.11 is that each delay-visibility is a weighted mean of several values of $\tilde{I}^2(u, v, \eta)$, smeared together by the windowing function \tilde{W}^2 (shown in Figure 2.1). This smearing effect in (u, v, η) has three significant consequences worth noting.

The first consequence is that the number of independent k -modes of the power spectrum, N_k , will be limited by the size of the windowing function. The size of the ‘footprint’ of a measurement is $\Delta\tilde{V} = \int W^2$. While this value does not directly affect the estimators \hat{P} and $\hat{\Delta}^2$, it will affect the statistical significance of our measurement by limiting the number of independent modes one can measure.

The second consequence is that $|\tilde{I}^2|$ may change appreciably over the range in (u, v, η) that our delay-visibility is spread over, particularly in cases where the delay-visibility probes a large range in spatial scales (a scenario not applicable to the COPSS data set, where $\Delta k < k$). If unaccounted for, this may bias estimates of the power spectrum and lead to misinterpretation of results.

The final consequence is that delay-visibilitys centered at different positions in (u, v, η) may measure partially overlapping regions of $\tilde{I}(u, v, \eta)$, and can produce a measurement of the power spectrum. Invoking the assumption that the Universe is homogeneous and isotropic on sufficiently large scales, orthogonal spatial modes for a randomized brightness temperature fluctuation field will be incoherent with one another, such that $\langle \tilde{I}^*(k) \tilde{I}(k') \rangle_{k \neq k'} = 0$. The product of two delay-visibilitys offset by some distance in (u, v, η) produces a coherent measurement of the specific intensity variance, I^2 , weighted by the inner product of the two offset windowing functions:

$$C(\Delta u, \Delta v, \Delta \eta) = \frac{\int W(u, v, \eta) \cdot W(u + \Delta u, v + \Delta v, \eta + \Delta \eta) du dv d\eta}{\int W(u, v, \eta) \cdot W(u, v, \eta) du dv d\eta}. \quad (2.12)$$

$C(\Delta u, \Delta v, \Delta \eta)$ is the normalized covariance (with respect to the astronomical signal being measured) of two delay-visibilitys separated by some distance $(\Delta u, \Delta v, \Delta \eta)$, normalized such that for two delay-visibilitys with zero separation have a normalized covariance of $C(0, 0, 0) = 1$. With this we can now define our estimator, $|\hat{I}^2|$, for the true specific intensity variance, $|\tilde{I}^2|$, as

$$|\hat{I}^2(u, v, \eta)| = \frac{\sum_{n \neq m} \tilde{\mathcal{V}}_m^*(u, v, \eta) \tilde{\mathcal{V}}_n(u', v', \eta') C(u - u', v - v', \eta - \eta') (\sigma_m^2 \sigma_n^2)^{-2}}{\sum_n (C(u - u', v - v', \eta - \eta') / \sigma_m^2 \sigma_n^2)^2} \quad (2.13)$$




Figure 2.1: *Left*: A single cut through the function \tilde{W}^2 , where $\eta = 0$. Different contours correspond to 0.2 dex differences in sensitivity, from unity (dark red) to 0.01 (dark blue). *Right*: Another slice through \tilde{W}^2 , where $u = 0$. Dashed lines mark the discrete values of η sampled by our analysis.

where σ_m^2 is the noise estimate in a given visibility, \mathcal{V}_m . The sum is carried out for each visibility $\mathcal{V}_m(u, v, \eta)$ over all other independent visibilities $\mathcal{V}_n(u', v', \eta')$. We discard any terms that arise from the product of a delay-visibility with its own complex conjugate in order to prevent our estimator from being positively biased by instrumental noise. Our sum in Equation 2.13 has been naturally weighted so as to maximize the sensitivity of our measurement (though other weighting schemes can be used). With this, we define our estimators for the power spectrum, \hat{P} and $\hat{\Delta}^2$, as

$$\hat{P}(k) = \frac{2X^2Y}{\Omega_B B_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left\langle \left| \hat{I}^2 \right| \right\rangle_{\mathbf{k} \cdot \mathbf{k} = k^2}, \quad (2.14)$$

$$\hat{\Delta}^2(k) = \frac{k^3 X^2Y}{\pi^2 \Omega_B B_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left\langle \left| \hat{I}^2 \right| \right\rangle_{\mathbf{k} \cdot \mathbf{k} = k^2}. \quad (2.15)$$

2.1.2 Single Dish Measurements

Single dish observations require a slightly different derivation, as the fundamental measurement is one that measures real power (versus the complex power that the interferometer measures). The most significant difference between interferometer and single dish intensity mapping experiments is that Fourier dual of the windowing function, $\tilde{W}(u, v, \eta)$, presented in Equation 2.6 is no longer strictly set by the primary beam pattern, but rather more generally is set by the mosaicking pattern used in the observing. One can rewrite the expression for

the windowing function as $W = S(l, m, \nu) * A_\nu(l, m)$, where S is the sampling function used for the mosaicked observations. This change in the windowing function will change the the normalization term in Equations 2.2 Equation 2.9, $W(x) \cdot W(x) = V_z$, such that

$$V = \int W \cdot W = X^2 Y \int (S * A) \mathcal{B} \cdot (S * A) \mathcal{B} d\Omega d\nu \approx X^2 Y B_z \Omega_{B,\nu} N_{\text{points}} / \mathcal{R}_\Omega, \quad (2.16)$$

where N_{points} is the number of independent sky positions observed, and \mathcal{R}_Ω is a measure of the overlap/redundancy of those different pointings (which arises from the need to enforce that $W(\mathbf{x}) \leq 1$). An additional difference between single dish and interferometric measurements is that that the primary beam will act as a spatial filter, which will suppress modes beyond the edge of the aperture function, \tilde{A} , in the uv plane. As such, our estimator for the Fourier dual of specific intensity, \hat{I} , can be written as

$$|\hat{I}^2(u, v, \eta)| = |\tilde{T}_m^2(u, v, \eta) / A_\nu^2(u, v)|, \quad (2.17)$$

where \tilde{T}_m is the measured values from the Fourier transform of the (l, m, ν) cube of brightness temperature, which we will refer to as “image modes”. We can now modify Equation 2.6 to be suitable for single dish measurements:

$$\tilde{\xi}(\mathbf{k}) = \frac{X^2 Y}{\Omega_{\text{survey}} B_z} \left(\frac{\tilde{T}_m^2(u, v, \eta)}{A_\nu^2(u, v)} \right), \quad (2.18)$$

where Ω_{survey} is the effective survey area, equal to $\Omega_{B,\nu} * N_{\text{points}} / \mathcal{R}_\Omega$.

2.2 Relationship to the Luminosity Function of Galaxies

The power spectrum for CO as a function of wavenumber and redshift is given by

$$P(k, z) = \langle T_{\text{CO}} \rangle^2 b^2(z) P_{\text{lin}}(k, z) + P_{\text{shot}}(z), \quad (2.19)$$

where P_{lin} is the linear matter power spectrum, $b(z)$ is the halo bias, where $b(z)$ is the halo bias, and P_{lin} is the linear matter power spectrum. The $\langle T_{\text{CO}} \rangle^2(z) b^2(z) P_{\text{lin}}(k, z)$ term is the contribution to the power spectrum from the clustering of CO emitters (sometimes referred to as “clustering power”). P_{shot} is the shot contribution to the power spectrum (sometimes referred to as “shot power”). One can further define the halo bias as

$$b(z) = \frac{\int_{M_{\text{CO},\text{min}}}^{\infty} dM L(M) \frac{dn}{dM} b(M, z)}{\int_{M_{\text{CO},\text{min}}}^{\infty} dM L(M) \frac{dn}{dM}}, \quad (2.20)$$

where $b(M, z)$ is the mass-dependent halo bias, $L(M)$ is the CO luminosity of a halo of mass M . Halos with masses below the low-mass limit, $M_{\text{CO},\text{min}}$, are assumed to be deficient in CO (either due to the suppression of star formation or because of a breakdown in the simple

linear $M-L_{\text{CO}}$ relationship assumed), such that they do not appreciably contribute to our power spectrum measurement (e.g. Pullen et al. 2013; Li et al. 2016).

Our goal is to connect these two components of the power spectrum (the cluster and shot power) to properties of the population of CO luminous, molecular gas-bearing galaxies within our measurement. Let us assume that the number density of these objects is well described by the Schechter function (Schechter 1976), which has the general form

$$\Phi(L)dL = \phi_* \left(\frac{L}{L_*} \right)^\alpha e^{-L/L_*} dL/L_* \quad L > L_{\text{min}}. \quad (2.21)$$

Equation 2.21 is normally parameterized by a high-luminosity cutoff, L_* , a low-luminosity power law index, α , a normalization factor for the overall density of luminous sources ϕ_* . We have added an additional parameter, L_{min} , which describes the low-end luminosity limit (below which the number of sources drops significantly).

2.2.1 The Cluster-Power Component

In order to calculate the cluster-power component of the power spectrum, one needs to know the mean brightness temperature for the spectral line of interest (i.e., CO(1-0)), $\langle T_{\text{CO}} \rangle$. We can relate $\langle T_{\text{CO}} \rangle$ to the volume emissivity of the CO(1-0) line, $\varepsilon_{\text{CO}(1-0)}$, using the following expression:

$$\langle T_{\text{CO}} \rangle = \frac{1}{8\pi} \frac{(c/\nu)^3}{2k_B H(z)} \varepsilon_{\text{CO}(1-0)}. \quad (2.22)$$

Using the Schechter parameterization, the volume emissivity is equal to

$$\varepsilon_{\text{CO}(1-0)} = \phi_* L_* \Gamma(\alpha + 2, L_{\text{min}}/L_*), \quad (2.23)$$

where Γ is the incomplete gamma function. If there exists a correlation between CO luminosity and some other well-studied galaxy property (e.g., star formation rate), then we can also relate the emissivity to another galaxy function. Let us then generally describe the correlation between CO luminosity and this other property as

$$L_{\text{CO}} = C(m)^\beta, \quad (2.24)$$

where L_{CO} is CO luminosity, C and β are the linear and power-law components of the correlation, and m is the value of the other property our correlation is measured against. The volume emissivity of the CO(1-0) transition can then be calculated as,

$$\varepsilon_{\text{CO}(1-0)} = \phi_* C^{1+\beta} \left(\frac{L_*}{C} \right)^\beta \Gamma(\alpha + \beta + 1, L_{\text{min}}/L_*), \quad (2.25)$$

where ϕ_* , L_* , and α in Equation 2.25 are the Schechter parameters from the “well-measured” galaxy function under consideration. We note that Pullen et al. (2013) use this method to convert the star formation rate (SFR) functions of galaxies out to $z \sim 6$ to provide estimates

for the volume emissivity of the CO(1-0) line, although in principle, other functions (e.g., bolometric IR luminosity) could be used to provide estimates as well.

From Equation 2.23, we note two items of potential interest for intensity mapping experiments. First, if attempting to evaluate the Schechter parameters, one cannot separately constrain the values of ϕ_* , L_* , and α with the measurements of the cluster power alone – there exist three unknowns but only one measured value. This degeneracy emphasizes that for constraining galaxy populations, intensity mapping measurements should be regarded as complementary (rather than redundant) with direct detection efforts, as individual galaxy detections (notably, blind detections) help to break this degeneracy. Second, a less obvious consequence of Equation 2.23 is that L_{\min} has an impact on the cluster-power measurement only when $\alpha \lesssim -2$. While within the local Universe, Keres et al. (2003) found that for the CO luminosity function $\alpha \sim -1$, observations of high redshift galaxies (e.g., Smit et al. 2012) suggest that this value may decrease with increasing redshift (i.e., further back in cosmic time).

We note that the cluster-power portion of the power spectrum is also dependent on the halo bias, which requires knowledge of the relationship between CO luminosity and host halo mass. While some work has gone into exploring the connection between molecular gas and host-halo mass (e.g., Popping et al. 2015), present observational data provide only weak constraints for high redshift objects. As the work presented in Chapter 4 is focused primarily on the shot-power component of the power spectrum, we allow more detailed calculations of $b(z)$ to be left to future work.

2.2.2 The Shot-Power Component

We now consider the shot-power component of the power spectrum (which we expect to be sensitive to with an instrument like the SZA). The shot power, P_{shot} is proportional to the galaxy luminosity function, $\Phi(L)$, in the following way:

$$P_{\text{shot}}(z) = \left(\frac{c^3(1+z)^2}{8\pi\nu_o^3 k_B H(z)} \right)^2 \int L^2 \Phi(L) dL. \quad (2.26)$$

In Equation 4.4, $\int L^2 \Phi(L) dL$ is the second moment of the luminosity function, which we define as being equal to the volume emissivity variance, ε^2 . With the Schechter parameterization, we can further define ε^2 as

$$\varepsilon^2 = \phi_* L_*^2 \gamma(\alpha + 3, L_{\min}/L_*). \quad (2.27)$$

Once again, we consider two important consequences of Equation 4.4. First, shot power measurement is strongly dependent on L_* , moderately dependent on ϕ_* and only very weakly dependent on α . While this implies that a shot power measurement is only weakly sensitive to sub- L_* galaxies (and thus can only weakly constrain α), the strong dependence on L_* and ϕ_* make such a measurement sensitive to the “knee” of the luminosity function (beyond which the number density of objects drops precipitously). As such, one needs only moderate

number of blindly detected galaxies (versus what might be required for disentangling the cluster-power measurement) in order to separately constrain ϕ_* and L_* . Second, as the shot and cluster power components are differentially sensitive to different parts of the luminosity function, a power spectrum measurement that probes both regimes (in combination with a limited number of direct detections) would be sufficient to break the degeneracies for fits of the different Schechter parameters, allowing for (at least to first-order) a characterization of the broad population of CO-luminous galaxies – both bright and dim – contained within the survey volume.

2.3 Spectral Line Confusion

The analysis presented in this thesis rests on the assumption the dominant source of spectral line emission detected at 30 GHz arises from the CO(1–0) transition. However, the method of intensity mapping does not discriminate between different sources of emission; verification is required to confirm that any detected signal is originating from the CO(1–0) transition. In particular, some care is required to ensure that other spectral lines – which we refer to here as “spectral interlopers” – do not dominate the CO(1-0) signal that we are trying to detect.

Table 2.1 shows a listing of the strongest spectral lines between 30 and 300 GHz, corresponding spectral lines emanating from redshifts $z < 10$. While not exhaustive, it contains lines with fluxes that might be comparable to the CO(1-0) signal. For comparative purposes, we have listed luminosity of transition in question to that of the CO(1-0) transition for the galaxy Arp 157, an interacting galaxy undergoing starburst activity that may be representative of some fraction of galaxies in the early Universe (Visbal & Loeb 2010), or otherwise may be representative of objects from which we might expect relatively enhanced spectral interloper emission (Breysse et al. 2015). Values for the comparative line luminosities are derived from observations and analysis performed by Davis et al. (2013). As shown in Table 2.1, the CO(1–0) dominates the line luminosity of most other transitions. Transitions beyond $z \gtrsim 10$ are not expected to be significantly impact our measurement due the limited amount of cosmic star formation and subsequent low abundance of metals (Vonlanthen et al. 2009). Finally, different isotope species of CO – such as ^{13}CO and C^{18}O – occur in limited enough abundance such that their presence will not significantly impact the measurement of $^{12}\text{C}^{16}\text{O}$ (Danielson et al. 2013), and higher order transitions of CO at higher redshift are expected to be much weaker than the $z \sim 3$ signal (Pullen et al. 2013; Muñoz & Furlanetto 2014). Similar studies on other analogs to high redshift galaxies, such as M82, show similar dominance of the CO(1-0) transition over other lines nearby in frequency (Aladro et al. 2011).

While technically contaminants, the presence of additional transitions of the CO molecule provides a useful tool for verification of the signal detected through intensity mapping. Several theoretical studies have suggested that cross-correlation of different rotational states (e.g. CO(1-0) with CO(2-1)) could be used to verify a detection (Righi et al. 2008). Addi-

Transition	Rest Freq (GHz)	Redshift Range	$L'_{\text{line}}/L'_{\text{CO}}$ in Arp 157
C ₂ H	87.316	1.49 – 2.23	0.035
HCN	88.631	1.53 – 2.28	0.046
HCO ⁺	89.189	1.55 – 2.30	0.049
HNC	90.663	1.59 – 2.36	0.019
CS	97.981	1.80 – 2.63	0.019
C ¹⁸ O(1 – 0)	109.782	2.14 – 3.07	0.007
¹³ CO(1 – 0)	110.201	2.14 – 3.08	0.052
CN	113.157	2.23 – 3.19	0.012
CN	113.499	2.24 – 3.20	0.026
CO(1 – 0)	115.271	2.29 – 3.27	1.000
H ₂ CO	140.840	3.02 – 4.22	0.003
0 CS	146.969	3.20 – 4.44	0.009
C ¹⁸ O(2 – 1)	219.564	5.27 – 7.13	0.006
¹³ CO(2 – 1)	220.403	5.30 – 7.16	0.028
CO(2 – 1)	230.538	5.59 – 7.54	0.468

Table 2.1: A listing of significant spectral lines potentially detectable in the COPSS dataset.

tionally, different isotopes of CO (e.g. ¹³CO(1–0)) may also provide verification, with the advantage that these isotopes are observable over a nearly identical redshift range as the CO(1-0) transition. Values from Table 2.1 suggest that abundances of these isotopes in high redshift are low enough to limit contaminating the CO(1 – 0) result, while strong enough to allow for a detection (or confirmation of a detection) in cross-correlation with the CO(1-0) transition.

We note here that theoretical estimates from [Breysse et al. \(2015\)](#) suggest that the CO signal at moderate redshifts ($z \sim 3$) may be dominated HCN emission arising from rare ($\sim 10^{-7} \text{ Mpc}^{-3}$) halos hosting ultra-luminous and hyper-luminous infrared galaxies (ULIRGS and HLIRGS, respectively). While there are large uncertainties in this model of HCN, we note that such luminous halos are likely to be above the modest direct detection threshold of an intensity mapping experiment. We therefore conclude that in the absence of a detection of such objects, the threat of line contamination from HCN is likely much less than what [Breysse et al.](#) predict.

2.4 Estimates of the Impact of Continuum Foregrounds

Our analysis is contingent on matching the the frequency of an emitter to its radial comoving distance by determining its redshift. However, continuum emission can be mistaken for an extended structure along the line-of-sight direction, contaminating data across all

frequencies/redshifts. Fortunately, under the Fourier transform in Equation 2.7, emission from smooth-spectrum sources will primarily be contained to modes with $k_z = 0$, where k_z is the wavenumber along the line-of-sight direction. Excluding these modes from analysis presents a simple and straightforward way eliminate contamination from continuum sources. This method is not perfect; there are two ways in which this simple approximation fails. First, only sources that are perfectly flat spectrum (i.e. with a spectral index of $\alpha = 0$) – non-flat sources, such as synchrotron, free-free and thermal sources, will leak some small amount of power into modes other than $k_z = 0$. Second, the response of an interferometer is a function of (among other things) the distance between telescopes measured in numbers of wavelengths; this value changes as a function of frequency, creating a problem akin to bandwidth smearing (Bridle & Schwab 1999) that is commonly referred to as “mode-mixing”. In order to quantify the effect continuum emitters will have on our analysis, we consider three different sources of emission: bright point-sources, faint point-sources and wide-field.

We define bright sources with fluxes bright enough to be well above a 3σ detection limit (~ 1 mJy). We will consider the general case of a point source, whose flux is described by $S(\nu) \propto \nu^\alpha$. The primary beam of the telescope, which can be approximated by a 2D Gaussian, attenuates the source as by a factor,

$$A(\nu) = e^{-\ln(2)\theta_s^2/(\theta_P/2)^2} = e^{\ln(16)\theta_s^2/(\theta_{B,o}\nu^2/\nu_o)^2}. \quad (2.28)$$

In Equation 2.28, θ_s is the angular offset of the source from field center, $\theta_{B,o}$ is the FWHM at a particular default frequency ν_o . At the field edge, this induces a spectral index of ~ -2 on a flat spectrum source. Though the induced spectral index can be larger at positions further from the field center, the attenuation of the primary beam will greatly decrease the received flux from sources well outside of the field of view of the telescope; the point at which the induced spectral index from primary beam attenuation is -2 , the primary beam attenuation itself is 10dB. The amount of power leaked into modes other than the $k_z = 0$ mode is

$$P_{\text{leak}} \approx \frac{2X^2Y}{N_{\text{ch}}\Omega_B B_z} (10^{-23} \frac{c^2}{2k\nu^2} \Delta S_{\text{Jy}} B_z)^2, \quad (2.29)$$

$$\Delta S_{\text{mJy}} \approx S_{\nu, \text{mJy}} \left(1 - \left(\frac{\Delta\nu}{\nu} \right)^\alpha \right). \quad (2.30)$$

Equation 2.29 represents the worst case scenario, where all of the leakage power has been consolidated into a single delay bin. Under this scenario, the leakage will still be diminished by roughly a factor of $1/N_{\text{ch}}$, since the remaining unpolluted channels will lower the mean level of contamination. In practice this leakage will not quite be so strong, though a smooth-spectrum source will have most of this leakage power in the lowest k_z bins. Plugging in nominal values for Equation 2.29, and assuming a worst case scenario of a steep synchrotron source ($\alpha \sim -2$) at field edge – giving a total effective spectral index of $\alpha \sim -4$ – we get $P_{\text{leak}} \approx 10^3 S_{\text{mJy}}^2 \mu\text{K Mpc}^3$, $\Delta_{\text{leak}}^2 (k = 1h \text{ Mpc}^{-1}) \approx 50 S_{\text{mJy}}^2 \mu\text{K}^2$. Under this scenario, the leakage from a non-flat spectrum generally exceeds that from bandwidth smearing, hence we limit ourselves in looking at solely the leakage as described in Equation 2.29. For sources of

flux density $S_{30\text{GHz}} \lesssim 5$ mJy, this worst-case residual error will reside below our sensitivity threshold, and can be safely ignored. In practice, the typical residuals are expected to be much lower (discussed further in Sections 3.3.2 and 5.3).

We now evaluate the impact of the unseen background of continuum point sources, which may be undetectable individually but detected as an integrated sum. As before, the leakage flux will be heavily dependent upon the spectral index of the ‘average’ point source. The average point source is much more likely to have a spectral index closer to zero than our discussion of bright point sources allowed for, hence we cannot ignore the contribution of mode mixing as we did in the previous discussion.

To establish the limits of the contamination in our measurement, we assume that the two effects are additive. Analysis of SZA and VLA data suggests that the density of point sources at 30 GHz is given by $dN/dS = 32.1(S_{\text{mJy}})^{-2.18} \text{deg}^{-2} \text{mJy}^{-1}$ for sources with fluxes below 15mJy (above which, the number of point sources drops precipitously), with spectral indexes ranging from $\alpha \approx -1 \rightarrow 2$ (Muchovej et al. 2010). These point sources add together in an incoherent fashion, such that the total flux density in a single measurement from the interferometer, \mathcal{V} (commonly referred as a visibility), will be given by $\mathcal{V}_{\text{bg}} = \sqrt{\int \Omega_{\text{B}} N(S_{\text{mJy}}) \cdot S_{\text{mJy}}^2}$, which is dominated by the brightest of the point sources. The leakage power of this temperature fluctuation will be dependent upon how much the UV footprint for a single baseline overlaps with itself as the frequency changes:

$$P_{\text{leak}} \approx \frac{2X^2Y}{\Omega_{\text{B}}B_z} \left(10^{-23} \frac{c^2}{2k\nu^2} \tilde{\mathcal{V}}_{\text{bg}} \right)^2 \left(\left(\frac{\tilde{A}_\nu(0) \cdot \tilde{A}_{\nu+\Delta\nu}(\Delta B_\lambda)}{\sqrt{\tilde{A}_\nu^2 \tilde{A}_{\nu+\Delta\nu}^2}} \right)^2 + \frac{1}{N_{\text{ch}}} \left(\alpha \frac{\Delta\nu}{\nu} \right)^2 \right). \quad (2.31)$$

In Equation 2.31, \tilde{A}_ν refers to the aperture function of a baseline at frequency ν , ΔB_λ is the change in baseline length (in wavelengths) over the frequency width of a single channel, and $\tilde{A}(0) \cdot \tilde{A}(D)$ is the inner product of two aperture functions offset by some distance D (where $\tilde{A}^2 = \tilde{A}(0) \cdot \tilde{A}(0)$). Plugging in appropriate values, we find that $P_{\text{leak}} \approx 50 \mu\text{K Mpc}^3$, $\Delta_{\text{leak}}^2(k = 1h \text{ Mpc}^{-1}) \approx 2 \mu\text{K}^2$ for the shorter baselines between the antennas within the inner cluster, and $P_{\text{leak}} \approx 2 \times 10^3 \mu\text{K Mpc}^3$, $\Delta_{\text{leak}}^2(k = 1h \text{ Mpc}^{-1}) \approx 10^2 \mu\text{K}^2$ for the longer baselines associated with the outrigger antennas. Once again, we assume that all leakage power has been collapsed into a single $k_z \neq 0$ mode – in practice, the leakage power is likely to be distributed across multiple bins, diluting the contamination. The estimated effect is in relatively good agreement with the estimate presented here with simulations (presented in Chapter 5).

We finally turn our attention to wide-field contaminants, such as galactic dust, free-free, synchrotron, and CMB anisotropies. As with the background radio source population, there will be two contributions to the leakage – the spectral index of the source emission, and the changing position of the baseline within the UV plane. There is one additional factor to consider with galactic emission sources, in that the expected flux for a given baseline depends on baseline length, and is not solely driven by the shot power arising from the random distribution of sources (as is the case for the radio source population). Galactic

foregrounds are significantly diminished on the small angular scales that the SZA is sensitive to, hence they are extremely unlikely to contaminate our measurement (White et al. 1999; Planck Collaboration et al. 2013). Analysis from Sharp et al. (2010) (hereafter S10) shows that for the SZA, the most significant source of contamination will be CMB anisotropies (both primary and those generated by the SZ effect), with an expected contribution of $P \approx 100 \mu\text{K}^2$. However, discarding the $k_z = 0$ modes from our analysis considerably reduces the power from the data to $P_{\text{leak}} \lesssim 1 \mu\text{K}^2$, $\Delta_{\text{leak}}^2(k = 1h \text{ Mpc}^{-1}) \lesssim 0.1 \mu\text{K}^2$.

Chapter 3

Data Analysis for Interferometric Power Spectrum Measurements

The analysis of interferometer data for power spectrum analysis required the development of an analysis package capable of performing the various calibration and data verification steps required for an intensity mapping experiment. This section provides a broad overview of the data processing and calibration software used in our analysis.

This chapter is structured as follows. Section 3.1 provides a brief description of the basic algorithms used in data processing, including those used for flagging, calibration, and imaging. Section 3.2 discusses the higher-level pipeline used to process and accumulate data to produce aggregate results. Section 3.3 discusses power spectrum analysis methods, and null tests are described in Section and 3.3.3. Methods of eliminating contributions from ground and point source contributions, and estimates on the residual power remaining in our measurement, are presented in Sections 3.3.1 and 3.3.2, respectively. All of the routines used here were developed within MATLAB¹. A full description of the algorithms used can be found in Appendix B.

3.1 Data Processing Overview

The calibration of the raw SZA data follows similar procedures to those described in Muchovej et al. (2007) and S10. The raw data are recorded as complex correlation amplitudes with associated time-tagged status information. They are converted to a physical power scale using system temperature measurements that are made during every source-calibrator cycle. Absolute telescope and system efficiencies, derived from Mars via the Rudy et al. (1987) brightness temperature model, are applied. As these factors are identical to those used in S10, we expect they are accurate to 10%. The data are flagged to remove bad data, bandpass calibration is determined from a strong point source in each track, and relative gain calibration is determined from the gain calibrator observed on a 20-minute cycle. Some

¹Mathworks, Version 2013b, <http://www.mathworks.com/products/matlab/>

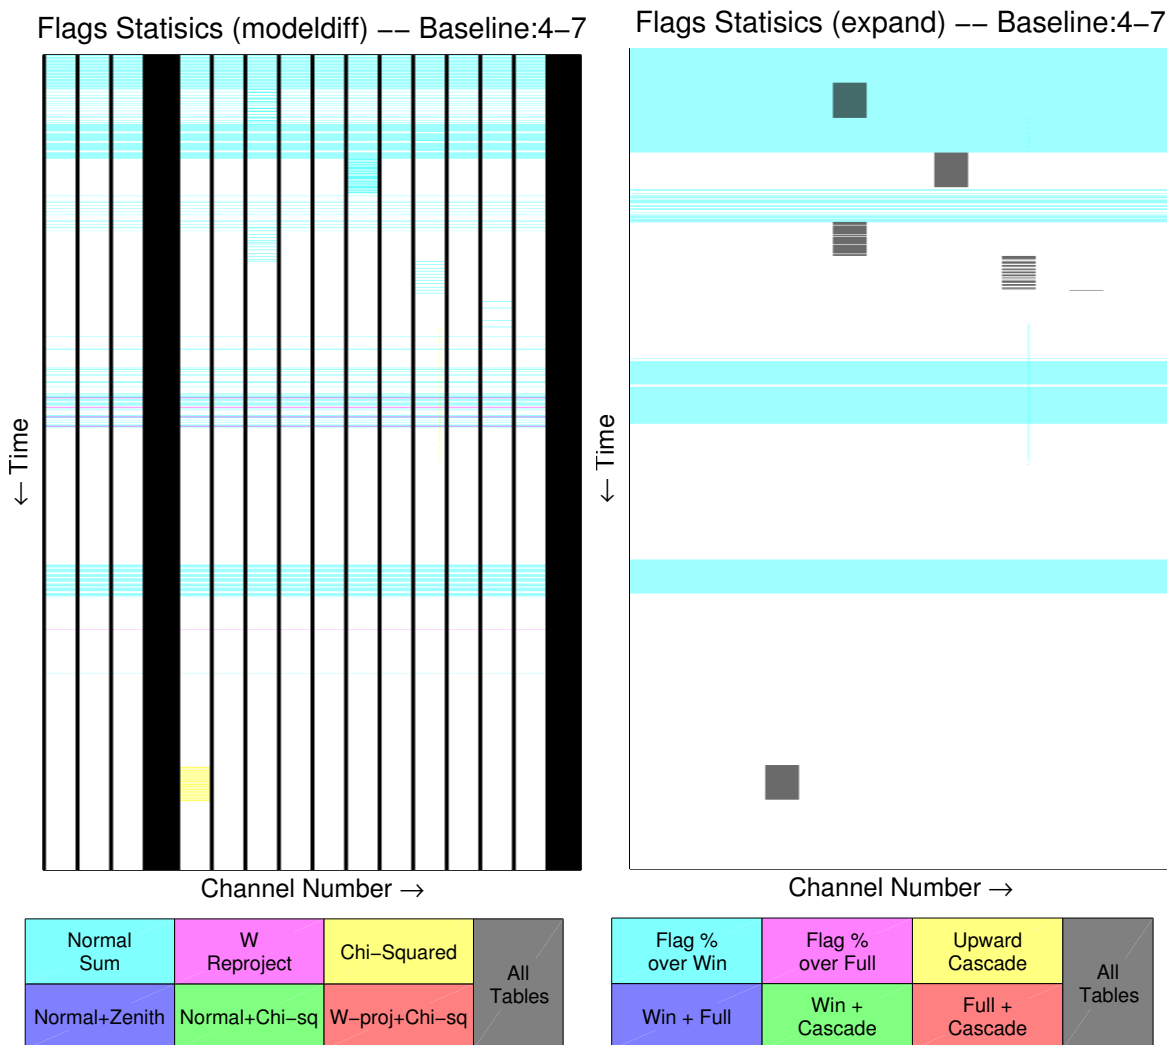


Figure 3.1: An example of the flags table generated for FLANK1.294 – from a single obsblock within the COPSS dataset. *Left*: Results from statistics-based flagging of the data, with color denoting the reason for the flag being set. Color opacity corresponds to the table in which the flag is set (where lighter colors correspond to higher-level tables). Data marked black have been flagged by a priori flagging routines (which capture problems like hardware problems and antenna showing). *Right*: Results from the flagging based on bad adjacent data.

features of these steps are outlined below. Flagging of data is done using three principal methods. First, data affected by known hardware issues (e.g., an antenna fails to point correctly) are marked as bad. Second, data are passed through various statistical tests and checked against theoretical estimates to see if they behave in a Gaussian fashion (data that exceed estimated noise thresholds are flagged as bad). Finally, data adjacent (i.e., belonging to preceding or subsequent frequency channels or integrations) to bad data are removed as

well. Flagging of data is typically done with an iterative approach, identifying outliers in small groups of data (and removing them when appropriate) before reevaluating the data over larger groupings. This approach helps limit the impact that a few bad data points may have on an otherwise good data set, at the expense of additional processing time. We remove 28.7% of the data for known hardware problems (including shadowing of antennas), 4.8% of data for exceeding noise thresholds, and 6.5% of data for bad neighboring data. We also discard all data from our shortest baseline (amounting to 3.6% of our data additionally removed) due to known systematic noise issues. In total, 43.6% of all data are flagged.

The SZA system produces very stable calibration across frequency and time. For the first phase of the COPSS experiment, gain amplitude and phase typically vary by 2% and $< 10^\circ$ across a track, respectively. With the second phase of the COPSS survey, gain amplitude and phase typically vary by 3% and $< 20^\circ$ over the course of a 24-hour period, respectively. When large gain shifts are observed ($> 10\%$ in amplitude or $> 30^\circ$ in phase) between gain calibrator observations, data are marked as bad and excluded from later analysis. Phase solutions are linearly interpolated, while the very stable gain amplitudes are averaged over the track. Bandpass solutions are also typically stable to better than 1% between days, with solutions showing an RMS variability of 1.3% over the course of the entire survey. An example for the gains solution for a single observing block is shown in Figure 3.2. One additional calibration step is performed to account for discrepancies between the expected and actual noise within our measurements. System temperatures are measured for each window, but RF/IF features and quantization effects can introduce variability in the system temperature on a channel-by-channel basis. To account for this effect, the variance within each channel of each baseline is calculated for all data within a single observing block (excluding calibrator data, and after subtraction of known sources in each field), and the “system equivalent flux density (SEFD) correction” is determined as the difference between the theoretical and measured noise. These differences are believed to be due to antenna-based effects (primarily due to standing waves in the receiver bandpass that introduce a spectrally-varying signal level at the digitizers, resulting in varying quantization noise) hence an antenna-based correction factor is determined (using a χ^2 fit) for each frequency channel. The correction factors are determined once per track and are seen to be consistent to 1% between observing blocks. An example SEFD correction spectrum is shown in Figure 3.3.

3.2 Pipeline Overview

Flagging, calibration and imaging/power spectrum production tools are arranged together into a pair of pipelines for the processing of data. The data pipeline can be broken into “daily processing” and “accumulation” stages. The daily processing pipeline analyzes a single contiguous observing block at a time. This stage of the pipeline can be broken down into the following operations:

1. Flag bandpass calibrator data, and generate bandpass solutions




Figure 3.2: An example of the gains solutions generated for BB20070906 – a single obsblock within the COPSS dataset. The gains solutions are typical of those derived for our analysis; we typically observe very small deviations in amplitude (consistent with noise-dominated errors), and slow drifts in phase of $\sim 30^\circ$ over the course of 4-6 hours. Gain stability is important, as gain errors may produce artifacts capable of spoiling our power spectrum measurement.

2. Flag gain calibrator data, and generate gains solutions
3. Perform a first round flagging of source data
4. Calculate SEFD-correction solutions
5. Perform a second round of flagging of source data
6. Repeat the above steps 1 – 5 one additional time (with tighter constraints for flagging)




Figure 3.3: An example of the SEFD correction solution generated for BB20070906 – a single observing block within the COPSS I data set. SEFD corrections show good day-to-day agreement, similar to the bandpass calibration solutions.

7. Cross-field subtraction performed
8. Data are weighted and gridded for imaging or power spectrum production

Bandpass calibrator data are processed first, followed by gain calibrator data. Once calibrator data has been processed, a first round of flagging is performed on all sources, excluding flagging methods that perform χ^2 tests of the data. Proper analysis of χ^2 values requires good estimates of the noise, hence they are best evaluated after SEFD corrections are performed. The first round of flagging (in step three) helps to eliminate any data which would significantly affect the noise correction solutions. SEFD correction solutions are then calculated, and the flagging process on source data is repeated once again, this time with the

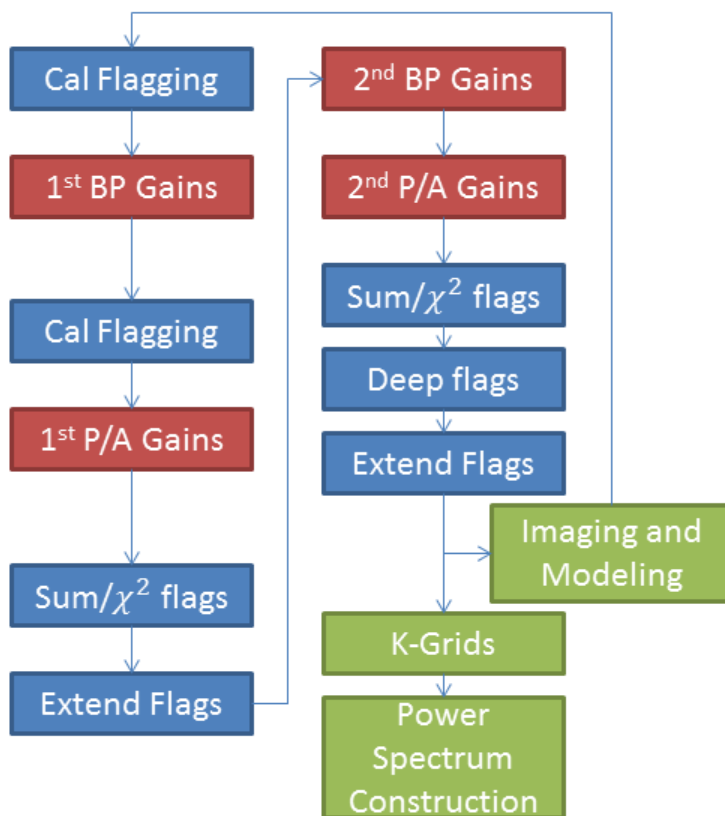


Figure 3.4: A diagram of the pipeline used in the processing of the COPSS data sets. An iterative approach is used with the pipeline, as early gains and SEFD correction solutions are expected to be at least partially contaminated by bad data during the earlier iterations.

χ^2 analysis enabled and flagging constraints set tighter. Next, the above steps are repeated once more, with refined gains solutions and noise corrections allowing for even tighter flagging constraints for the data. A copy of the data is gridded for imaging, while another copy undergoes ground subtraction prior to producing gridded data for power spectrum analysis.

After all days are processed, the accumulation stage of the pipeline begins. At this stage, images and gridded power spectrum data from individual days are stacked together to produce a single image and set of power spectra (one for each redshift window) for each field, with interday jackknife tests (also discussed further in Section 3.3.3) performed at this stage as well. Power spectra are summed together across different fields and redshifts to produce a cumulative measurement.

As a final check for systematics (in addition to the jackknife tests described in Section 3.3.3), a series of tests are performed on the accumulated datasets. First, images for each field are generated (shown in Chapter 4), and the noise statistics for all spectral channels

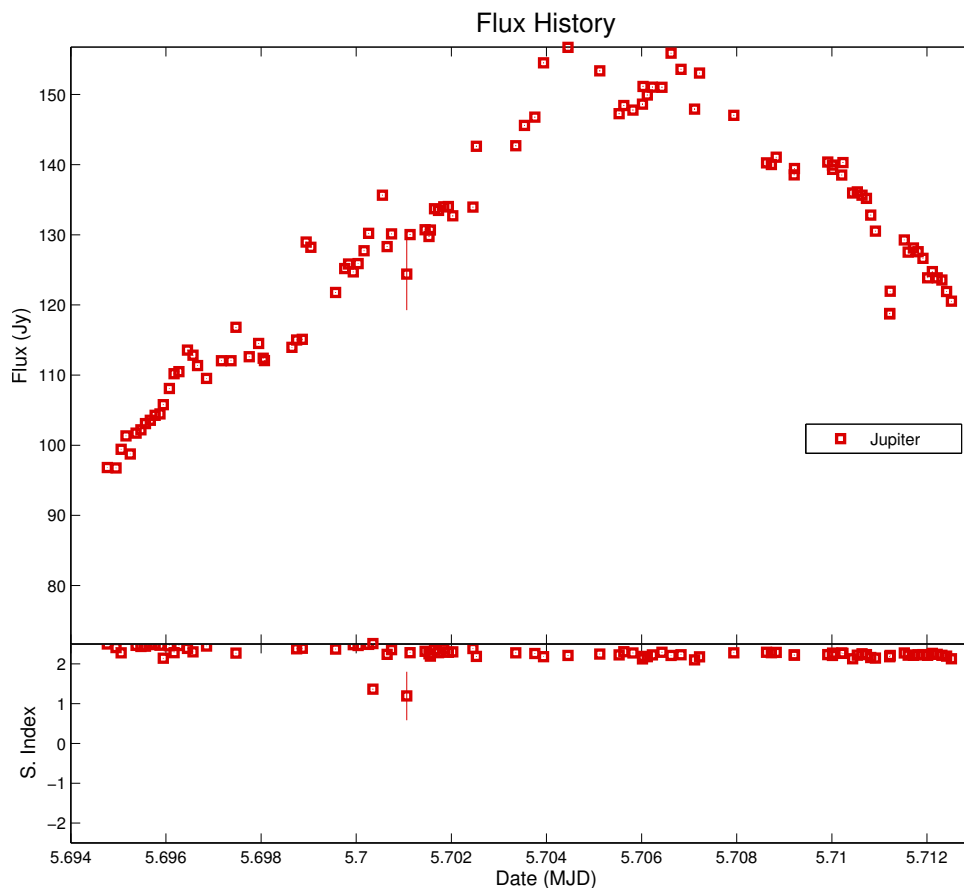


Figure 3.5: Flux monitoring history of Jupiter, which serves as a secondary flux calibrator for the COPSS data set. The observed flux matches well with theoretical estimates for 30 GHz, indicating that our absolute flux scaling for our data appear to be well calibrated. Flux history plots like these are used for quality control with processed data.

for all fields are calculated after deconvolution, in order to verify that the measured residual noise agrees with theoretical estimates. Second, the fluxes of gain and bandpass calibrators are measured as a function of time (across all days and obsblocks; an example of which is shown in Figure 3.5). Obsblocks that show large deviations in calibrator flux are visually inspected to look for evidence of bad data or other contaminants not caught by automated processing, and are subsequently reprocessed. Obsblocks that continue to show unexplained discrepancies in measured flux are excluded from further analyses. Finally, a power spectrum (described further in Section 3.3) is generated for each obsblock, and is accumulated as a running sum across the length of our observation. The end result of this analysis is a measure of the power in a single wavenumber-bin within a single redshift window for each field, as a function of integration time. As the power within a single redshift window is expected to be well below the noise threshold of our experiment, one expects the power to integrate down

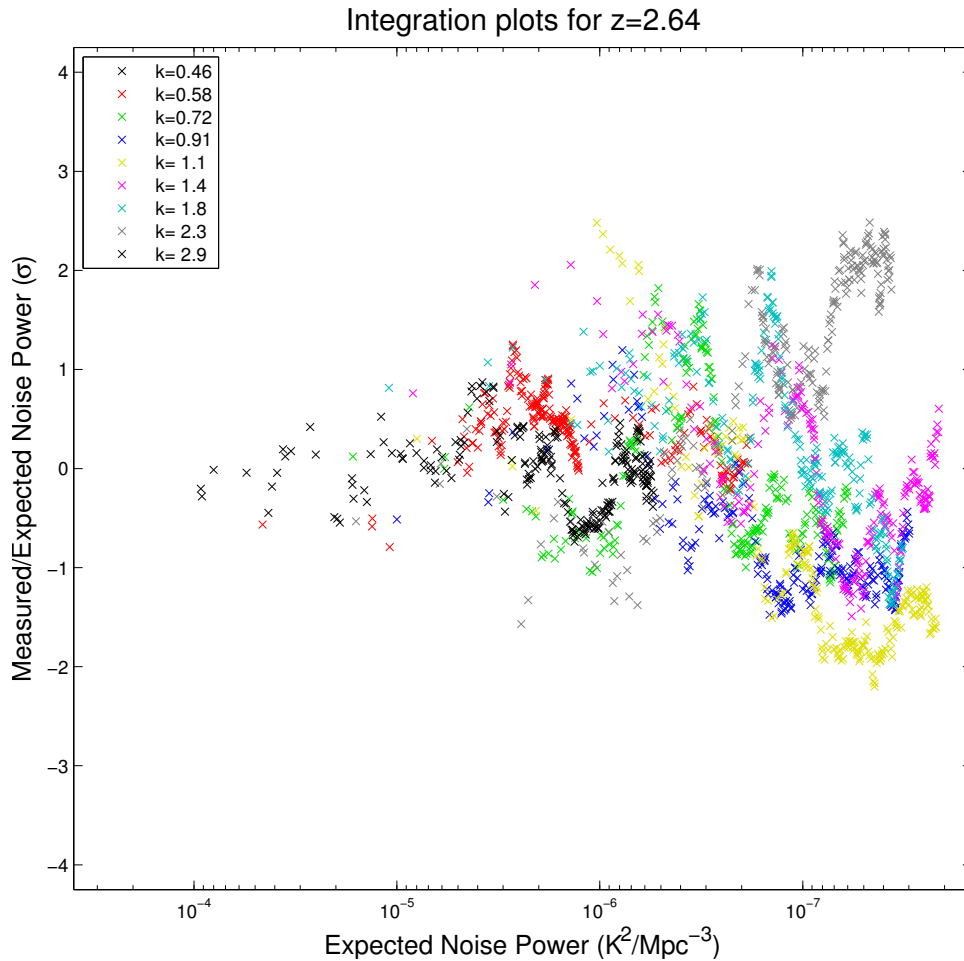


Figure 3.6: Measured power within individual modes for a single redshift window in GOODS-N, a field observed as part of the COPSS II survey. As the astrophysical signal is expected to be very weak within a single mode-redshift window pairing, one expects to see power levels consistent with noise. Such checks help to identify systematics that may only appear over short time intervals (and may be buried within the data over the long period of our measurement).

like noise if our measurement is free of systematics. An example measurement is shown in Figure 3.6 – we find that the measured power within individual modes integrate down in a noise-like fashion.

3.3 Power Spectrum Analysis

The goal of our power spectrum analysis is to transform the frequency-resolved interferometric visibilities into a measurement of the CO power spectrum. We have provided the

full details of this method in Chapter 2, but we provide a brief overview of the method here. The power spectrum measurement is performed in a three-dimensional spatial frequency domain that is the Fourier transform of the real-space universe. The interferometer naturally provides measurements of spatial frequencies in the two dimensions of the sky, while the third is constructed by Fourier transforming the frequency dimension of the data, which maps to the line of sight distance via the redshift of the CO line. The data are Fourier transformed one 500 MHz window at a time, providing redshift-segmented measurements of the Fourier-transformed intensity field, $\tilde{I}(u, v, \eta, z)$, with z representing the median redshift of the window. We refer to these as “delay-visibility”. The delay-visibility are then converted to temperature units, weighted by their variance estimates and gridded in the (u, v, η, z) space. Here u and v are the standard spatial frequency variables for interferometry, and η represents the Fourier transform of the frequency axis. The (u, v, η) domain is closely related to the \mathbf{k} space, differing only by the conversion factors X and Y , which convert between comoving physical size and angular distance or frequency with units Mpc rad^{-1} and Mpc Hz^{-1} , respectively (e.g., [Parsons et al. 2012](#)). With the application of the X and Y conversion factors, we find the Fourier dual of the specific intensity in the three-dimensional k -vector space, $\tilde{I}(\mathbf{k})$. Within an individual window our coverage of \mathbf{k} space is similar in the η direction and the u and v directions, considering only the compact portion of the array.

For our analysis, we define

$$\begin{aligned} \Delta^2(k, z) &\equiv \frac{k^3}{2\pi^2} P(k, z), \\ &= \frac{k^3}{2\pi^2} \frac{X^4 Y^2}{V_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \left\langle \left| \tilde{I}^2 \right| \right\rangle_{\mathbf{k} \cdot \mathbf{k} = k^2}. \end{aligned} \quad (3.1)$$

In Equation 3.1, the 1D power spectrum, $P(k, z)$, is a measure of the variance in brightness temperature over a given comoving volume, expressed in this paper in units of $\mu\text{K}^2 h^{-3} \text{Mpc}^3$ (where $h = H_0/100 \text{ km/s/Mpc}$; H_0 represents the current Hubble parameter). The power spectrum is given as a function of comoving spatial frequency k (with units of $h \text{Mpc}^{-1}$) and redshift z , and is proportional to the Fourier transform of the autocorrelation function of the intensity field under consideration (with \tilde{I} defined as the Fourier transform of this intensity field). The 1D power spectrum is produced by spherically averaging over logarithmically sized shells of $\mathcal{P}(\mathbf{k}, z)$, which is the three dimensional power spectrum, as a function of vector wavenumber, \mathbf{k} , and redshift. Given the solid angle of the telescope primary beam, Ω_B , and bandwidth, B_z , the volume surveyed by our measurement at a given redshift is $V_z = X^2 Y B_z \Omega_B / 2$. X and Y are conversion factors between comoving distance and angular/frequency separation, respectively (e.g., [Parsons et al. 2012](#)). For a single pointing, the volume surveyed by a single correlator window (in this case, the central one) is a cylinder of diameter $12 \text{ Mpc}/h$ and length $45 \text{ Mpc}/h$. The Boltzmann constant is represented by k_B , and the speed of light by c .

$\Delta^2(k)$ is the variance in brightness temperature per $\ln(k)$, expressed in this paper in units of μK^2 (often called the dimensionless power spectrum; e.g., [Furlanetto et al. 2006](#)). As a dimensionless quantity, $\Delta^2(k)$ is sometimes favored over $P(k)$ ([Dodelson 2003](#)). However,

$P(k)$ has the advantage that it maintains the same value for different values of k when in the cosmological shot noise limit; for both the CO measurement and some foreground contaminants, our experiment resides firmly in the shot noise regime. As such, values in both conventions are presented throughout this paper.

For the analysis presented here, we use two different methods to calculate power spectrum values: measuring the product of semi-correlated pairs (PSCP) of gridded delay-visibilitys, and performing a maximum-likelihood evaluation (MLE) of the band-averaged power. Values for the PSCP method using the following equation:

$$\begin{aligned} \mathcal{P}(\mathbf{k}, z) &= \frac{\sum_{\mathbf{k}'} \sigma_k^{-2} \sigma_{k'}^{-2} C(\mathbf{k} - \mathbf{k}') \left(|\tilde{I}^*(\mathbf{k}, z) \tilde{I}(\mathbf{k}', z)| \right)}{\sum_{\mathbf{k}'} \sigma_k^{-2} \sigma_{k'}^{-2} C^2(\mathbf{k} - \mathbf{k}')} - \mathcal{A}_{\mathbf{k}} \\ P(k, z) &= \langle \mathcal{P}(\mathbf{k}, z) \rangle_{\mathbf{k} \cdot \mathbf{k} = k^2}. \end{aligned} \quad (3.2)$$

In Equation 3.2, gridded data within a single redshift window, $\tilde{I}(\mathbf{k}, z)$, are cross-multiplied against one another, weighted by their estimated thermal noise variance, σ_k^2 , and their normalized covariance (for the signal of interest), $C(\delta\mathbf{k})$ for data separated by $\delta\mathbf{k}$ in the (u, v, η) domain. The covariance is the analytically calculated correlation between adjacent points in the (u, v, η) space. For the uv plane, this amounts to a consideration of the overlap of the visibilitys being multiplied, each of which samples an area in the uv plane described by the autocorrelation of the telescope illumination pattern. The correlation is similarly calculated between η channels. To eliminate noise bias from our measurement, the autocorrelations of individual delay-visibilitys are averaged over each grid cell ($\mathcal{A}_{\mathbf{k}}$) and removed from the power spectrum measurement. As a result, the autocorrelation measured in our power spectrum analysis may produce negative values when the result is noise-dominated. The power spectrum measurement is collapsed down to a single dimension, averaging over spheres in the k space. A power spectrum is created for each redshift window, within each source field.

The PSCP method is computationally fast, requiring a few seconds of CPU time to calculate a power spectrum for an individual field. The primary limitation of this method is that it presumes the power spectrum errors are normally distributed (by way of the central limit theorem), when the distribution is actually a χ^2 distribution with N_k degrees of freedom (where N_k is the number of independent measurements contained within each bin of the power spectrum).

Following the prescription from Bond et al. (1998) and Hobson & Maisinger (2002), under the assumption that our signal of interest is comprised Gaussian fluctuations in brightness temperature, the likelihood of a given model for these fluctuations can be expressed as

$$\mathcal{L}(C) = \frac{1}{2\pi^{N/2} |C|^{1/2}} \exp\left(-\frac{1}{2} \tilde{\mathbf{I}}^T C^{-1} \tilde{\mathbf{I}}\right) \quad (3.3)$$

In Equation 3.3, $\tilde{\mathbf{I}}$ is the vector containing the renormalized delay-visibilitys, C is the covariance matrix of the data set, and N is the number of data within the data set. The




Figure 3.7: The covariance matrix for the signal of interest for the Q2343 field, one of the fields observed during the second phase of the COPSS experiment. The partial correlation between large groups of gridded delay-visibilitys make it computationally expensive to calculate the inverse of this matrix, which is required for the MLE analysis. Larger grid cells would reduce the computational burden, at the cost of some loss in sensitivity (due to the averaging together of only partially coherent data).

covariance matrix can further be expressed as

$$C = P(z)C_{\text{signal}} + C_{\text{noise}} + C_{\text{const}}, \quad (3.4)$$

Where C_{signal} (an example of which is shown in Figure 3.7) is the covariance induced by the signal of interest (calculated in the same fashion as the PSCP method), and $P(z)$ is the band-averaged power of the power spectrum at a given redshift (i.e., our simple model assumes that all modes measured within a single redshift window have power $P(z)$). C_{noise} is a diagonal matrix containing the contribution of instrumental noise to the measurement, and C_{const} is the “constraint matrix”, used to downweight mode(s) with known unwanted contributions to the power spectrum (e.g., ground contaminants).

The primary advantage of the MLE method is that it allows for the direct calculation of uncertainties in the power spectrum constraints, rather than relying on the central limit theorem approximations of the PSCP results. However, due to the the rotation of the Earth, our measurement samples many partially correlated positions within the uv -plane. As demonstrated in Figure 3.7, this makes our covariance matrix non-sparse, making the MLE method significantly more computationally expensive to calculate versus the PSCP method.

In our analysis, we find that PSCP and MLE methods produce values that are generally within 0.1σ of each other, with errors that agree to within a few percent. Except where otherwise noted, we use the PSCP method to produce power spectrum figures, and otherwise use the MLE values in our analysis.

3.3.1 Ground Subtraction

It was found in S10 that without ground subtraction there was weak evidence for a ground-correlated contamination of the CMB power spectrum measurement. Similar tests, described in the next section, suggest that removal of such contamination is important to cleaning our data. The image shown in Figure 3.8 is a composite of the GOODS-N and FLANK1 fields prior to ground subtraction (after removal of known point sources). The horizontal striping in the image is indicative of ground-contaminated baselines in the north-south direction. The power spectrum in Figure 3.8 shows the impact of this ground contribution compared to the sensitivity achieved in COPSS I – without ground subtraction, the power spectrum amplitude is $\sim 100\times$ the noise level in our measurement.

The lead-trail design of the S10 data set and the primary observations of the COPSS II campaign (discussed in further detail in Chapter 4) allows the removal of any contributions to the power spectrum that are correlated with antenna position, such as ground contamination or antenna cross-talk. We create a model for the contamination using a variance-weighted average of all fields within a group (4 fields per group for the S10 data set, 3 for the COPSS II data set). The average is generated visibility by visibility, so that the first integrations on each field in the group are averaged together, preserving the individual baselines and frequency channels, as are the second integrations, and so on. This model is subtracted from the individual visibilities in each of the fields within a group, reducing the number of independent measurements in our experiment by 25% for the S10 survey and 33% for the COPSS II survey, which degrades the sensitivity by approximately 13% and 18% respectively.

To evaluate the efficacy of our ground subtraction method, we evaluate the maximum

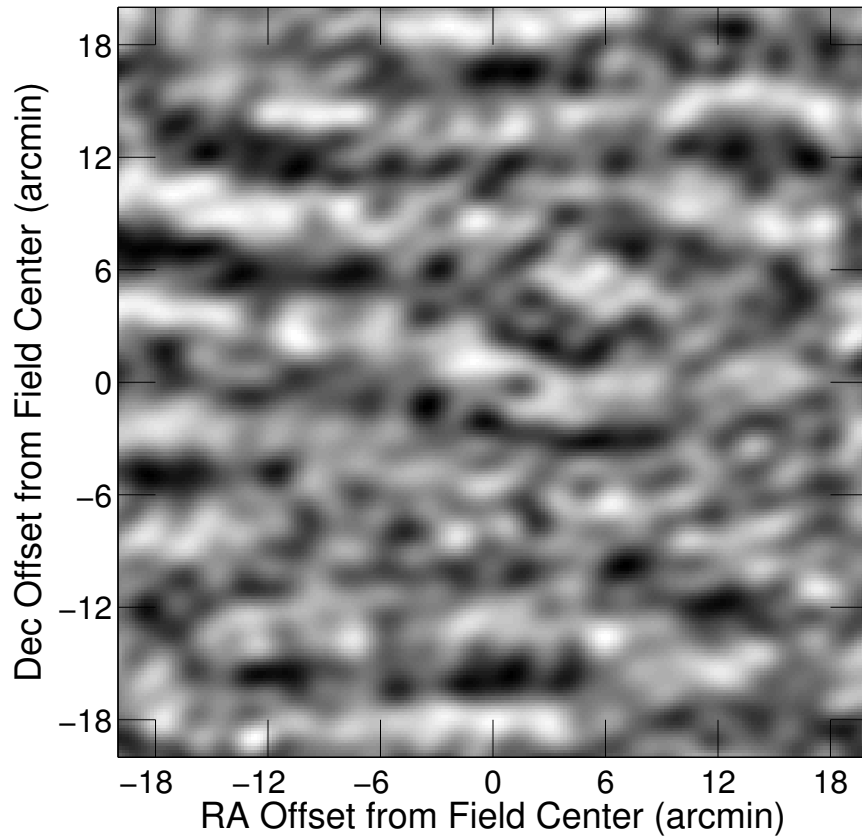


Figure 3.8: An image of the combined data from the FLANK1 and GOODS-N groups, after subtraction of known point sources. The horizontal striping in the image is indicative of ground contamination, with peak amplitude of $\sim 100 \mu\text{Jy}$ (5 times the theoretical noise limit).

power measured within our 3D power spectrum (i.e., $\mathcal{P}(\mathbf{k}, z)$), before and after the ground subtraction operation. Prior to ground subtraction, we find 2280 individual modes (0.2% of all modes sampled) exceed a 5σ threshold, with the greatest outliers containing power $\sim 10^3$ times greater than expected. After ground subtraction, we find no values in excess of 5σ , and of the smaller subset of 2280 aforementioned modes with known contamination, we find none in excess of 3.5σ . The reduction in maximum power measured within a single mode suggests that our subtraction step is suppressing ground-contributed power by at least a factor of $\sim 10^2$.

For the pilot data of the COPSS II survey, observations of different fields were sometimes separated by several days, preventing direct subtraction of the ground contribution. However, a jackknife analysis of the data (discussed further in Section 3.3.3) indicates that the longer baselines to the outrigger antennas are free of ground contamination and can therefore be included without this subtraction step.

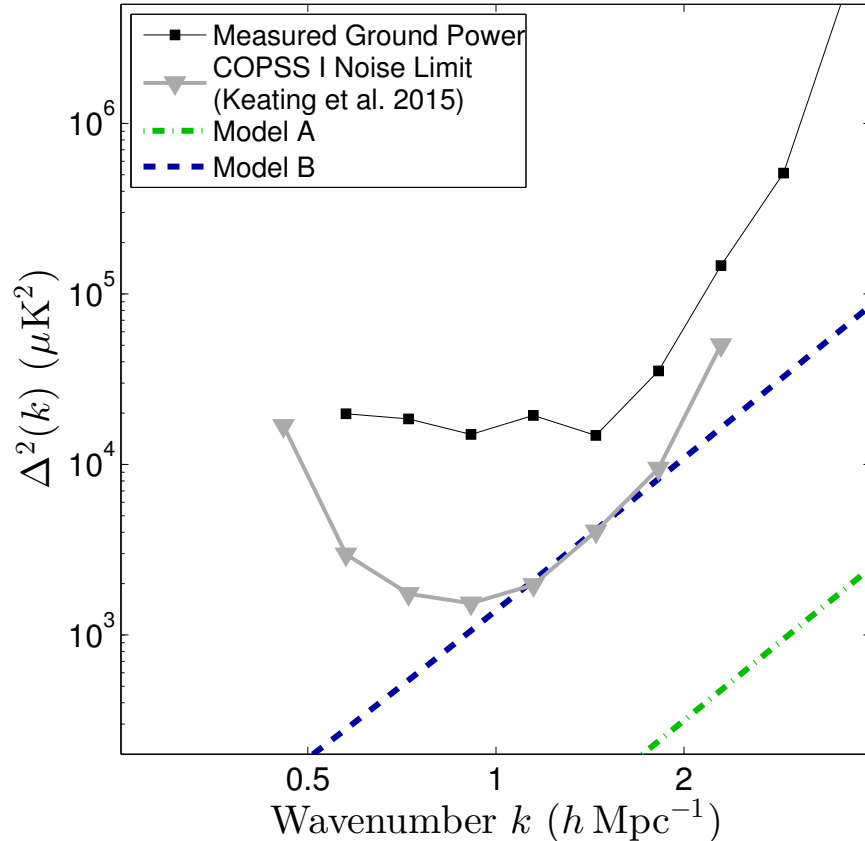


Figure 3.9: A power spectrum of the data prior to ground subtraction. The level of power in the ground-contaminated data resides two orders of magnitude higher than the noise limit presented in Section 4.3. Jackknife tests (further discussed in Section 3.3.3) of the unsubtracted data show similar or higher levels of power.

To recover the short-baseline data taken during the pilot COPSS II observations, the data are summed (in the (u, v, η, z) domain) across the fields of the pilot survey. A χ^2 test is then performed on these data for each position in u and v , evaluated across all redshift windows and delay channels (excluding the $\eta = 0$ channel). We expect individual data to be thermal-noise dominated; we therefore assume data that exceed the 4σ confidence threshold for our χ^2 test are irreparably contaminated by systematics, and exclude them from further analysis. The remaining data is presumed to be only weakly contaminated by the ground, such that cross-correlation between the pilot data and the ground-subtracted primary data is not expected to be significantly contaminated. With the PSCP method, this cross-correlation only requires using a slightly modified version of Equation 3.2, using the product of the pilot and primary datasets (and dropping the \mathcal{A}_k term). For the MLE method, we employ the constraint matrix to downweight the autocorrelation of the pilot




Figure 3.10: The RMS of the bandpass solutions of the data over the course of the observations. Most show variation of order a few percent, with edge channels typically showing an increased variability. Antenna 4 shows higher-than-average variability in a few windows due to a known intermittent hardware problem – affected data are automatically flagged in survey field datasets.

data, such that only the autocorrelations of primary survey data and the cross-correlation of pilot and primary survey data contribute to our measurement.

3.3.2 Point Source Contamination

The primary source of contamination in our measurement is expected to arise from continuum point source emission. Our primary means of rejecting such contributions is to remove the $\eta = 0$ channel, although bandpass calibration errors and non-zero spectral indices of sources will lead to contributions to channels other than $\eta = 0$.

From simulations (discussed further in Chapter 5), we expect point sources to contribute $20 \mu\text{K}^2 h^{-3} \text{Mpc}^3$ to data belonging to short baselines, and $1000 \mu\text{K}^2 h^{-3} \text{Mpc}^3$ to long

baselines. However, the inclusion of point source subtraction reduces these contributions. As the median RMS noise in our images is 0.05 mJy, point sources brighter than 0.5 mJy are expected to be removed completely by point source subtraction. Based on simulations, the removal of these sources suppresses point source contributions by a factor of 20, such that we expect point sources to contribute $1 \mu\text{K}^2 h^{-3} \text{Mpc}^3$ and $50 \mu\text{K}^2 h^{-3} \text{Mpc}^3$ to short and long baselines, respectively. These contributions are well below the sensitivity achieved in the COPSS measurement (discussed in Chapter 4).

Contributions from calibration errors do not benefit from point source subtraction, and thus must be considered as an additional source of contamination for our measurement. As shown in Figure 3.10, the day-to-day stability of the SZA is excellent – typically showing an RMS variability of 1.3% for individual channels. If the errors in bandpass calibration were correlated across days, we expect this level of error to add $100 \mu\text{K}^2 h^{-3} \text{Mpc}^3$ of power to our measurement. However, our bandpass errors are thermal-noise dominated, and we expect these errors to average down when adding together multiple days worth of data. As bandpass error-related contributions are expected to scale with as the square of the fractional bandpass error, with approximately 400 different bandpass solutions, we expect bandpass errors to only add $1 \mu\text{K}^2 h^{-3} \text{Mpc}^3$ of power to our present measurement. While this contribution is comparable to the primary contribution from continuum point source for the shortest of baselines, it is well below the sensitivity threshold of our experiment.

3.3.3 Jackknife Tests

To determine the impact of systematic errors on our measurements, we conduct a set of null tests in which we remove the astronomical signal from our data via linear combination or randomization and search for residual power. These tests, commonly referred to as jackknife tests, are split into four different categories: intraday, interday, “cross-window” and noise tests. Results from our jackknife analysis can be found in Section 4.3.1.

Intraday tests take temporally adjacent visibilities (i.e., subsequent integrations of the same channel within the same baseline) and modulates them in such a way that the sky signals are canceled out. Two tests are performed in this manner: “couplet” tests (visibilities adjacent in time in a given baseline are simply subtracted from one another) and “triplet” tests (three adjacent visibilities are multiplied by a phase offset of 0° , 120° , and 240° respectively before being summed together). Visibilities are gain-corrected and flagged prior to this operation. In the event that one visibility of the couplet/triplet group is flagged as bad, the entire group is thrown out and not considered in subsequent jackknife analysis. Both couplet and triplet tests are sensitive to high cadence systematics within our data set, with the couplet and triplet tests evaluating the data set over respectively shorter and longer periods.

Interday tests utilize gridded data rather than visibilities. Data from different observing blocks (i.e., days) are summed into two separate stacks and subtracted from one another. Two different interday tests are performed. The “even-odd” test sorts observations from alternating days into two stacks, and then subtracts them. The “first-last” test sums together

data from all days belonging to the first half of the set of days and subtracts them from the sum of the second half.

The cross-window test operates by correlating different redshift windows with one another. The CO signal is not expected to correlate between different redshift windows, hence no subtraction-like step is necessary. However, continuum contamination that is not rejected by the removal of the $\eta = 0$ channel will correlate between windows, so this jackknife is an empirical check on our continuum rejection.

To further test the analysis software and pipeline, we perform a pair of tests to measure and confirm our noise estimates. In both tests, values are multiplied by a random complex number with absolute value of 1 and random phase. These “randomized phase” tests gives us our best verification of noise estimates, as any potentially coherent emission detected by the interferometer should be scattered and thus not affect the results of this test. The “random raw” test randomizes the phases of ungridded data, while the “random grid” randomizes the phases of gridded data. While the former is much more thorough, the latter requires much less processing time to complete (while still being a useful tool for the verification of noise estimates).

We show the results of our jackknife analysis in Figures 3.11 and 3.12, and in Table 3.1. Listed in the table are the measured values (and associated errors) for each of the tests. Additionally listed for each test is the probability to exceed (PTE) – the likelihood for a noise-like event to produce a measurement of equal or greater statistical significance. Our jackknife results appear to be consistent with noise, suggesting that our analysis is not dominated by systematics. We note that our cross-window test correlates windows that are two steps apart (e.g., window 1 with window 3, window 2 with window 4). When correlating adjacent windows with one another (e.g., window 1 with window 2, window 3 with window 4), the cross-window test finds power at a level of 2.3σ significance within the COPSS I results. We attribute this power to signal leakage between windows (brought on by minor imperfections in analog bandpass filtering), and do not expect it to impact our analyses.

Additionally, we validated our error estimates for the PCSP method using the random grid test. The RMS power is estimated by utilizing the measured system temperatures, aperture efficiency, and SEFD corrections to estimate the amount of noise power in a given delay-visibility (and propagating that noise estimate forward using Equation 3.1 to estimate the noise power in the power spectrum). The calculated uncertainty and the uncertainty derived from the random grid test agree to 10%, as we would expect for the 100 trials used in the test. We do find some minor differences between the two estimates in bins with lowest sensitivity; these bins have a small number of independent modes used to measure the power in the bin (hence their errors will not be normally distributed), therefore some discrepancy is to be expected.

Table 3.1: Jackknife Results

Jackknife Test	Data Set	Result		PTE
Couplet		PCPS	MLE	
	COPSS I	2.0 ± 1.7	—	0.25
	Pilot	2.1 ± 3.0	$2.4^{+3.0}_{-3.1}$	0.48
	COPSS II	-1.1 ± 1.4	$-1.7^{+1.4}_{-1.4}$	0.44
	Total	-0.6 ± 1.3	$-0.9^{+1.3}_{-1.3}$	0.62
Even-Odd	COPSS I	1.2 ± 1.3	—	0.39
	Pilot	4.4 ± 3.0	$4.1^{+3.0}_{-3.0}$	0.15
	COPSS II	-2.4 ± 1.4	$-2.6^{+1.4}_{-1.4}$	0.09
	Total	-1.1 ± 1.3	$-1.3^{+1.3}_{-1.3}$	0.38
	First-Last	COPSS I	0.2 ± 1.5	—
Pilot		-1.4 ± 3.4	$-0.9^{+3.4}_{-3.4}$	0.66
COPSS II		0.6 ± 1.6	$0.5^{+1.6}_{-1.7}$	0.65
Total		0.3 ± 1.4	$0.2^{+1.4}_{-1.4}$	0.82
Cross-Win		COPSS I	-0.2 ± 1.6	—
	Pilot	1.9 ± 2.4	$2.0^{+2.4}_{-2.4}$	0.43
	COPSS II	2.5 ± 1.5	$2.3^{+1.5}_{-1.5}$	0.10
	Total	2.1 ± 1.3	$1.9^{+1.3}_{-1.3}$	0.11

Note. — All power spectrum values are in units of $10^3 \mu\text{K}^2 h^{-3} \text{Mpc}^3$. PTE values are calculated using the PCPS values. “Pilot” refers to pilot observations performed with CARMA in preparation for the larger observing campaign. “COPSS II” refers to the primary observing campaign with CARMA.




Figure 3.11: Jackknife results for our analysis of the S10 data set, alongside the 1σ noise threshold for our measurement (solid gray). For each bin, we find the results to be noise-like in their distribution, consistent with our assumption that the data are predominately free of systematics that may affect our result. Our largest outlier has a significance of 2.0σ , consistent with what one would expect for a set of ~ 60 values that have been normally distributed.




Figure 3.12: Jackknife results for our analysis of the full COPSS data set, along with the estimated noise threshold of our measurement (solid gray). We find the jackknife results to be noise-like in their distribution, consistent with the assumption that our data are predominately free of systematics. The largest outlier has 2.4σ significance, consistent with what one would expect from a normally distributed set of data given the ~ 60 values produced by our jackknife analysis.

Chapter 4

The CO Power Spectrum Survey (COPSS)

In this chapter, we discuss our search for measurable anisotropy in the three-dimensional distribution of molecular gas by characterizing the CO variance spectrum, using the SZA. This analysis resembles the two-dimensional power spectrum measurement of that work, extended to a third dimension using the frequency channels recorded in the data. Sensitive measurements may eventually image the intensity variations due to the large-scale structure of CO-emitting galaxies, but at the depth of these data we expect, at best, only a statistical detection of CO fluctuations via variance that exceeds that expected from the thermal noise of the data set.

The COPSS experiment is divided into separate phases. The first phase (which we refer to as “COPSS I”) makes use of archival SZA data to place the first-ever constraints on the CO auto-power spectrum at $z \sim 3$. The data set for our analysis was previously used by [S10](#) as a measurement of the cosmic microwave background (CMB) power spectrum on arcminute angular scales. The second phase (which we refer to as “COPSS II”) makes use of data from a focused observing campaign with the SZA. These newer data offer an improved ability to constrain the CO power spectrum at $z \sim 3$, with a greater number of hours focused on fewer fields (versus the archival data analyzed in COPSS I), as well as a more compact array configuration with improved instantaneous sensitivity.

This chapter is structured as follows: Section [4.2](#) provides a brief description of SZA instrument. Section [4.2](#) describes the observational data used for both phases of the experiment. Section [4.3](#) presents the results of our analysis, and Section [4.4](#) discusses these results in the context of theoretical expectations for the CO power spectrum, CO galaxy luminosity function and the cosmic density of molecular gas.



Figure 4.1: An intrepid young radio astronomer, standing in front of two of the 3.5 meter antennas, both part of the Sunyaev-Zel'dovich Array. The SZA was originally located at the Owens Valley Radio Observatory, but was later moved up to Cedar Flat to join CARMA. *Image Credit: Jim Keating*

4.1 The Sunyaev-Zel'dovich Array (SZA)

The Sunyaev-Zel'dovich Array (SZA) is an 8-element interferometer designed for measurements of the Sunyaev-Zel'dovich (SZ) effect signature of galaxy clusters on arcminute angular scales (Muchovej et al. 2007). At the time of data collection, the SZA was located at the Owens Valley Radio Observatory (OVRO), it was later incorporated into the nearby Combined Array for Research in Millimeter-wave Astronomy.

Each SZA antenna had a single polarization (left-hand circular) 1-cm receiver, capable of observing between 27–35 GHz, which corresponds to $z \approx 2.3$ –3.3 for CO(1–0). At 30 GHz, an SZA antenna has a typical system temperature $T_{\text{sys}} \approx 40$ K, and a typical aperture

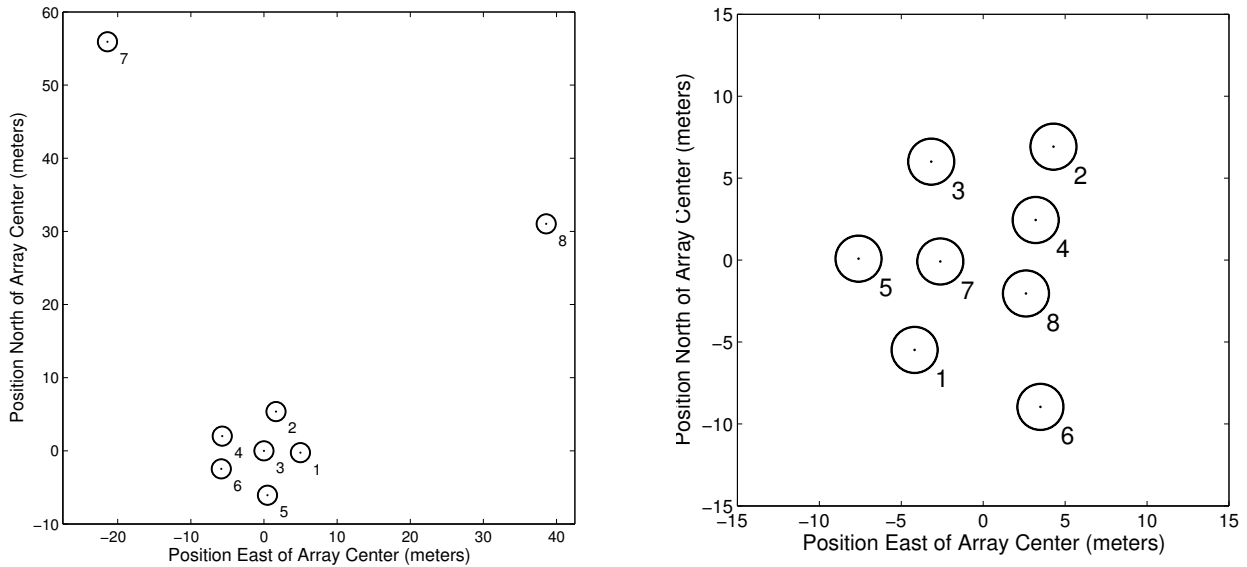


Figure 4.2: *Left:* SZA antenna positions for the S10 observations. The antennas are drawn as 3.5 meter diameter circles. *Right:* SZA antenna positions for the COPSS II observations. The more compact configuration allows for better instantaneous sensitivity of the SZA for power spectrum measurements.

efficiency of $\eta_{\text{eff}} \approx 0.6$. Each antenna is 3.5 m in diameter, which corresponds to a FWHM $\theta_B \approx 11'$ at 30 GHz and covers an effective solid angle of $\Omega_B = 0.03 \text{ deg}^2$. The SZA antennas were arranged in a compact group of 6 antennas (4.5-11.5 m spacings), with two outrigger antennas ~ 50 meters away to separate the extended SZ signal from compact radio continuum sources. Figure 4.3 shows both the layout of the SZA as well as the uv coverage of a single field during a typical observation. While pilot observations were performed with the SZA in a similar configuration to K15, the bulk of the observations were performed with the SZA antennas in a special, compact configuration (4.5-16m spacings) without antennas in outrigger positions.

The SZA used a 2-bit digital XF correlator (with efficiency $\eta_{\text{corr}} = 0.87$), which provided 8 autocorrelations and 28 cross-correlations. The full correlator had a total bandwidth of 8 GHz split across 16 spectral windows of 500 MHz bandwidth. Within each window are 17 spectral channels, with all but the first and last channel having a width of 31.25 MHz (the first and last channel effectively have half this width). The first and last channels are automatically removed by our pipeline, providing 468.75 MHz of bandwidth per spectral window for our analysis.




Figure 4.3: *Left:* the uv coverage of the full array for one spectral window (30.2 GHz) for one of the fields in the S10 survey, with darker shades corresponding to areas with greater sensitivity. The baselines to the outrigger antennas, 7 and 8, are well-separated from the baselines within the compact portion of the array. *Right:* The inner region of the uv plane.

4.2 Observations

4.2.1 COPSS I

The S10 data were obtained between 2005 and 2008. The data consist of 44 telescope pointings, arranged as 11 groups of 4 fields. Each group is composed of 4 pointings at constant declination separated by 4 minutes in RA, such that each field is observed over the same hour angle over a series of sequential 4 minute observations. The duration of the loop through the 4 fields, consisting of twenty 20-second integrations per field, plus several minutes on a gain calibrator is approximately 20 minutes. A bandpass calibrator was typically observed for 5 minutes at the beginning or end of the track. We refer to a contiguous track of data, typically 6 hours in length, as an “observing block”. Each group of 4 fields was observed for approximately 45 days, providing an hour a day of integration time, with an average total of 20 hours of integration time per field after taking into account data flagging. A listing of the position of the lead field for each group is provided in Table 4.2.

Figure 4.5 presents images of all 44 fields, deconvolved via the CLEAN algorithm (Högbom 1974). No primary beam correction is applied to these images. The resolution of the images is $\sim 0.5'$ and the median theoretical noise is 0.15 mJy. We find that the measured RMS residual noise in our images following deconvolution is consistent with theoretical estimates based on thermal noise, with the exception of cmbI17 (due to the presence of a bright point source at the field edge). We detect 24 sources above 10σ across the 44 fields,

Table 4.1: A listing of the lead field for each of the 11 groups observed in the S10 data set.

Field Name	RA	Dec	Gain Cal
cmbA1	02 ^h 12 ^m 00 ^s .0	+33°00′00″	J0237+288
cmbAA1	21 ^h 24 ^m 38 ^s .7	+25°29′37″	J2139+143
cmbBB1	21 ^h 24 ^m 38 ^s .1	+25°59′24″	J2025+337
cmbCC1	02 ^h 11 ^m 31 ^s .3	+33°27′43″	J0237+288
cmbDD1	13 ^h 18 ^m 40 ^s .1	+35°01′42″	J1131+305
cmbEE1	14 ^h 18 ^m 39 ^s .2	+35°31′52″	J1331+305
cmbI1	02 ^h 12 ^m 00 ^s .0	+32°37′08″	J0237+288
cmbR1	02 ^h 12 ^m 15 ^s .6	+32°11′24″	J0237+288
cmbY1	02 ^h 12 ^m 00 ^s .0	+31°51′24″	J0237+288
cmbXX1	21 ^h 24 ^m 38 ^s .7	+24°59′37″	J2139+143
cmb07	02 ^h 07 ^m 37 ^s .0	+34°00′00″	J0237+288

with a median flux of 2.5 mJy and a maximum (primary beam corrected) flux of 100 mJy. These sources are subsequently removed from our data prior to our power spectrum analysis. Sources below the 10σ threshold are ignored because of their limited impact on our power spectrum measurement (see Section 3.3.2). These results are consistent with those reported by Sharp et al. (2010).

4.2.2 COPSS II

Data were collected during two different phases: pilot observations and primary survey observations. Primary survey observations were conducted between October 2014 and April 2015, and consist of 19 telescope pointings arranged into 6 groups. Pilot observations were conducted between April 2013 and April 2014, and consist of 5 telescope pointings. Primary survey fields – FLANK1, GOODS-N (Dickinson et al. 2003), AEGIS (Davis et al. 2007), Q2343 (Steidel et al. 2004), and SXDS (Furusawa et al. 2008) – for both phases were selected to allow for continuous 24-hour observations. GOODS-N, AEGIS, Q2343 and SXDS were also selected based on present and future availability of optical spectroscopic data (e.g., Reddy et al. 2006; Brammer et al. 2012; Steidel et al. 2014; Kriek et al. 2015), to enable potential future cross-correlation experiments.

During pilot observations of GOODS-N, a series of flanking fields with the same declination but differing right ascension – FLANK1, FLANK2, FLANK3 and FLANK4 – were observed to help remove contamination dependent on the telescope pointing position. During the primary survey, primary fields (i.e, Q2343, SXDS, FLANK1, GOODS-N, and AEGIS) were observed in conjunction with a “leading” and “trailing” field, separated by approximately 5 minutes in RA, such that each field is observed over the same hour angle over a series of sequential 5 minute observations. The one exception was GOODS-N, where two trailing fields were observed due to the presence of a strong ~ 40 mJy point source in the

Table 4.2: COPSS II Observing Fields

Field Name	RA	Dec	Gain Cal	Obs Time (hours)	Notes Notes
FLANK1-L	06 ^h 33 ^m 28 ^s .0	+62°13'53"	3C147	221.6	a
FLANK1	06 ^h 38 ^m 50 ^s .0	+62°14'00"	3C147	270.3	a,b
FLANK1-T	06 ^h 44 ^m 11 ^s .0	+62°14'07"	3C147	226.8	a
GOODS-NL	12 ^h 31 ^m 28 ^s .3	+62°14'01"	J1153+495	13.5	a
GOODS-N	12 ^h 36 ^m 50 ^s .0	+62°14'00"	J1153+495	592.9	a,b
GOODS-NT	12 ^h 42 ^m 11 ^s .8	+62°13'59"	J1153+495	280.7	a
GOODS-T2	12 ^h 47 ^m 33 ^s .5	+62°13'57"	J1153+495	268.8	a
AEGIS-L	14 ^h 14 ^m 09 ^s .6	+52°51'04"	J1419+543	95.8	a
AEGIS	14 ^h 19 ^m 31 ^s .0	+52°51'00"	J1419+543	101.8	a
AEGIS-T	14 ^h 24 ^m 52 ^s .4	+52°50'56"	J1419+543	92.3	a
Q2343-L	23 ^h 40 ^m 44 ^s .2	+12°49'13"	3C454.3	146.4	a
Q2343	23 ^h 46 ^m 05 ^s .0	+12°49'12"	3C454.3	167.4	a
Q2343-T	23 ^h 51 ^m 25 ^s .8	+12°49'12"	3C454.3	158.6	a
SXDS-L	02 ^h 17 ^m 12 ^s .0	-04°59'59"	J0224+069	27.5	a
SXDS	02 ^h 18 ^m 00 ^s .0	-05°00'00"	J0224+069	11.5	a
SXDS-T	02 ^h 18 ^m 48 ^s .0	-05°00'01"	J0224+069	21.6	a
FLANK2	09 ^h 33 ^m 00 ^s .0	+62°14'00"	J0841+708	145.2	b
FLANK3	15 ^h 39 ^m 02 ^s .0	+62°14'00"	J1642+689	38.7	b
FLANK4	18 ^h 11 ^m 02 ^s .0	+62°14'00"	3C371	103.4	b

^aObserved during primary survey

^bObserved during pilot survey




Figure 4.4: Thumbnail images of the deconvolved maps for all 44 fields included in the S10 survey. On average, we find there to be a rough average of one bright point source ($S \gtrsim 1$ mJy) per pointing. Most fields after deconvolution show noise-like residuals, though cmbI17 shows some artifacts due to the presence of 100 mJy point source at the edge of the primary beam – arguably the worst-case scenario for contamination of our measurements (see Section 3.3.3).

leading field that was discovered during the first week of observations. For both phases of observations, fields were observed for a total of 180 5-second integrations (spent on one field for pilot observations, split between three fields during the primary survey), after which a gain calibrator was observed for several minutes, for a total of 20 minutes per observing loop. A bandpass calibrator and flux calibrator was typically observed between groups of fields, for 10 and 5 minutes respectively.

A listing of the position, observing and integration time for each field is provided in Table 4.2. Figure 4.5 presents images of all fields (with the exception of GOODS-NL), deconvolved using the CLEAN algorithm (Högbom 1974). No primary beam correction is applied to these images. The typical synthesized beam for these images is $2'$, with a

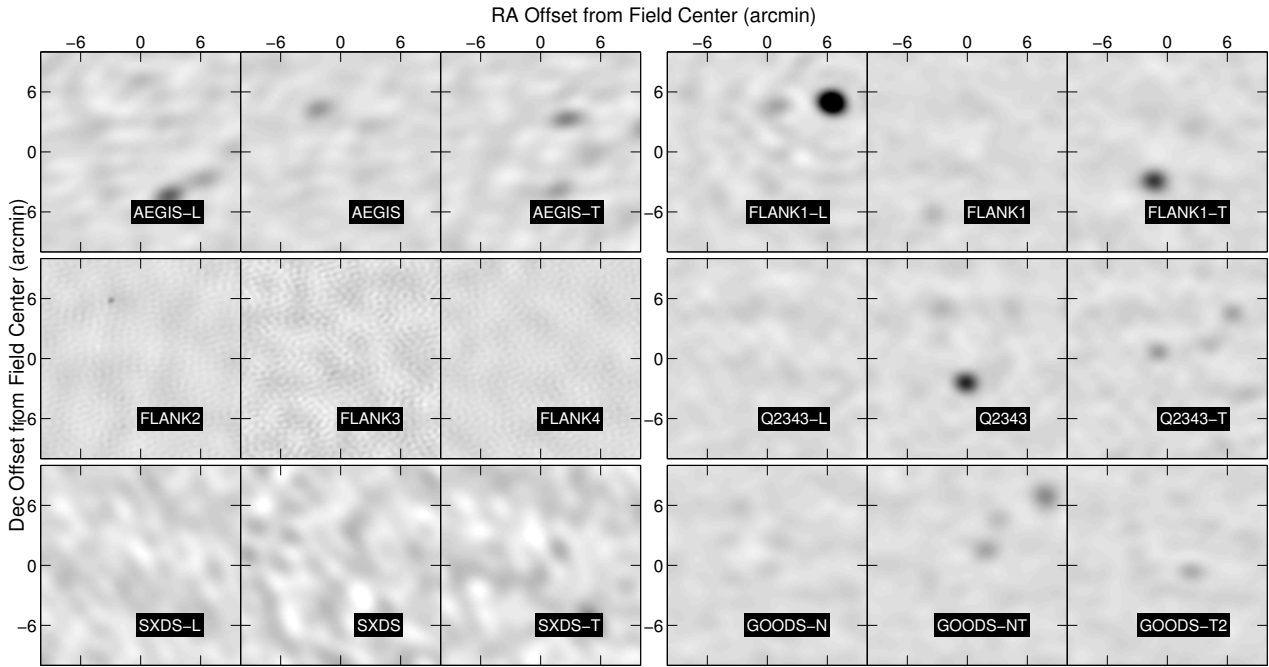


Figure 4.5: Thumbnail images of the deconvolved maps for all 18 fields included in the survey. We find there to be a rough average of two sources per pointing above our detection threshold (~ 0.25 mJy). After deconvolution, all fields show noise-like residuals. As pilot observations included longer baseline data (generated by antennas occupying the outrigger positions), the FLANK2, FLANK3, and FLANK4 images show higher resolution than other fields.

median theoretical noise is 0.05 mJy/beam. The RMS residual noise measured in our images following deconvolution is consistent with theoretical estimates based on thermal noise. As part of our analysis, we fit for the position and flux of all sources above a 5σ noise threshold in our images. We also fit for a spectral index for all sources detected above 10σ . Once fitted, point sources are removed in subsequent flagging and power spectrum analysis.

4.3 Results

4.3.1 CO Power Spectra

Presented in Figure 4.6 are the results of our analysis of the S10 data set. Our measurement has peak sensitivity at $k = 1.3 h \text{Mpc}^{-1}$, with best sensitivity between $k = 0.5\text{--}2 h \text{Mpc}^{-1}$. Integrating over all redshift windows and wavenumbers, we see no evidence for excess power from CO at $z \sim 3$ above a 2σ noise threshold of $P_N = 2.6 \times 10^4 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$. Placing this measurement into Δ_N^2 units, where Poisson power grows like k^3 , requires choosing a k value. At $k = 1 h \text{Mpc}^{-1}$, $\Delta_N^2 = 1.3 \times 10^3 \mu\text{K}^2$.

Table 4.3: CO Power Spectrum Results

Data Set	Result		PTE
	PCPS	MLE	
COPSS I	6.4 ± 13.3	—	0.32
Pilot	-0.4 ± 3.7	$0.1^{+3.7}_{-3.8}$	0.54
COPSS II	3.6 ± 1.4	$3.3^{+1.5}_{-1.5}$	0.01
Total	3.1 ± 1.3	$3.0^{+1.3}_{-1.3}$	0.01

Note. — All power spectrum values are in units of $10^3 \mu\text{K}^2 h^{-3} \text{Mpc}^3$. PTE values are calculated using the PCPS values.

Theoretical models (e.g., Pullen et al. 2013, hereafter P13) suggest that there may be significant evolution between the redshift range sampled by these data ($z = 2.3\text{--}3.3$). With a mean 2σ noise limit of $P_N = 10^5 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$, $\Delta^2(k = 1 h \text{Mpc}^{-1}) = 5 \times 10^3 \mu\text{K}^2$ within each redshift bin (of characteristic width $\Delta z = 0.06$), we see no evidence for excess power from CO within any of our individual redshift intervals within the COPSS I data set. Presented in Table 4.3 and Figures 4.7 and 4.8 are the final results of our analysis of the complete COPSS data set. Our measurement has peak sensitivity at $k = 1.3 h \text{Mpc}^{-1}$, with best sensitivity between $k = 0.5\text{--}2 h \text{Mpc}^{-1}$ (and marginal sensitivity between between $k = 2\text{--}10 h \text{Mpc}^{-1}$). Integrating over all redshift windows and wavenumbers, we detect power of $P_{\text{CO}} = 3.0^{+1.3}_{-1.3} \times 10^3 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$, and rejecting the null hypothesis ($P_{\text{CO}} > 0$) to 98.9% confidence. Placing this measurement into Δ_N^2 units, where Poisson power grows like k^3 , requires choosing a k value. At $k = 1 h \text{Mpc}^{-1}$, $\Delta_{\text{CO}}^2 = 1.5^{+0.7}_{-0.7} \times 10^3 \mu\text{K}^2$. We show in Figure 4.9 the results of our analysis for each redshift bin. We find weak evidence for increasing power with decreasing redshift: we measure $P_{\text{CO}} = 4.1^{+1.6}_{-1.6} \times 10^3 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$ for the low-redshift half of the data ($z = 2.3\text{--}2.8$), and $P_{\text{CO}} = 1.0^{+2.4}_{-2.4} \times 10^3 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$ for the high-redshift half of the data ($z = 2.8\text{--}3.3$). While this trend is not of high enough significance to demonstrate any evolution with redshift, it does match expectations of most models that the measured power should decrease with increasing redshift over the redshift range of our measurement.

4.3.2 Trim Tests

As an additional probe to search for systematics or other problems within the data, we perform a series of tests where we divide the dataset into smaller subsets and look at the total power measured – we refer to these as “trim tests”. Trim tests differ from jackknife tests in that the astrophysical signal of interest is not subtracted out (or otherwise removed), affording them an additional degree of flexibility (i.e., such tests do not require two groups

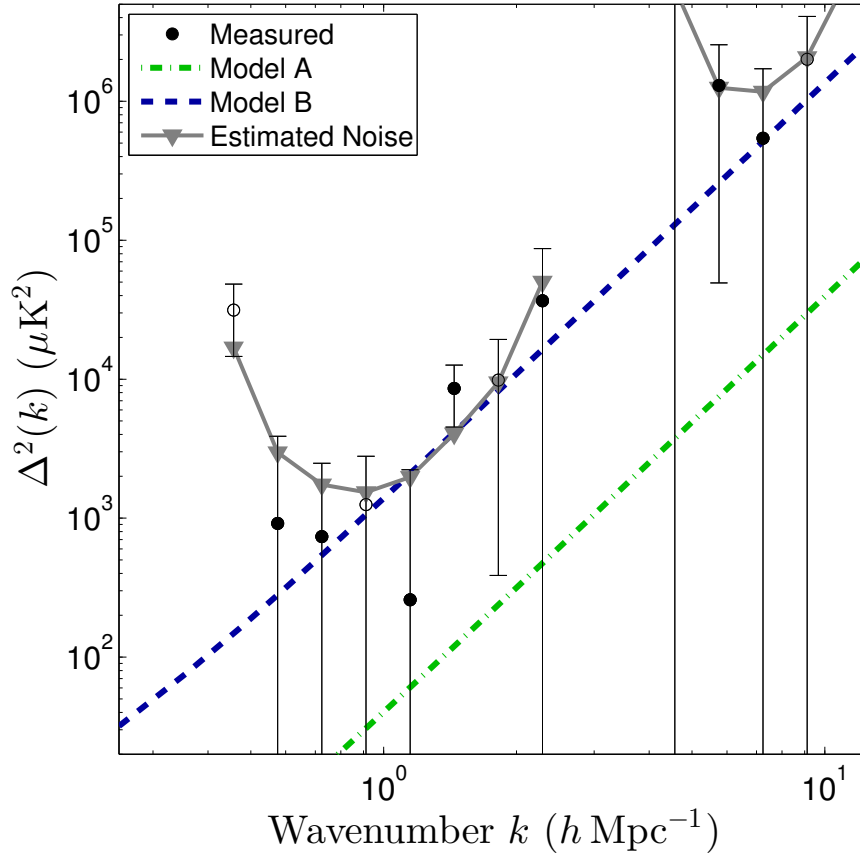


Figure 4.6: The power spectrum result from our analysis of the S10 data, in the form $\Delta^2(k)$. Filled circles correspond to positive values for $\Delta^2(k)$, while open circles correspond to negative values, and the error bars corresponding to the 1σ errors on our measured values. There exists a gap in sensitivity around $k \sim 4 \text{ h Mpc}^{-1}$, owing to the separation between baselines to the outrigger antennas and baselines within the compact portion of the array (see Figure 4.3). For reference, model A (dot-dashed green) and model B (dashed blue) from Pullen et al. (2013) are shown, along with the estimated RMS noise power (gray triangle), absent any astrophysical signal.

of data with identical sky signal to be differenced). Our primary trim test results are shown in Figure 4.10, which measures the power within a single group of fields (i.e., the primary field plus its leading and trailing fields), amounting to the removal of order 70% of the data. We find the results of this test to be in agreement with the results shown in Section 4.3.

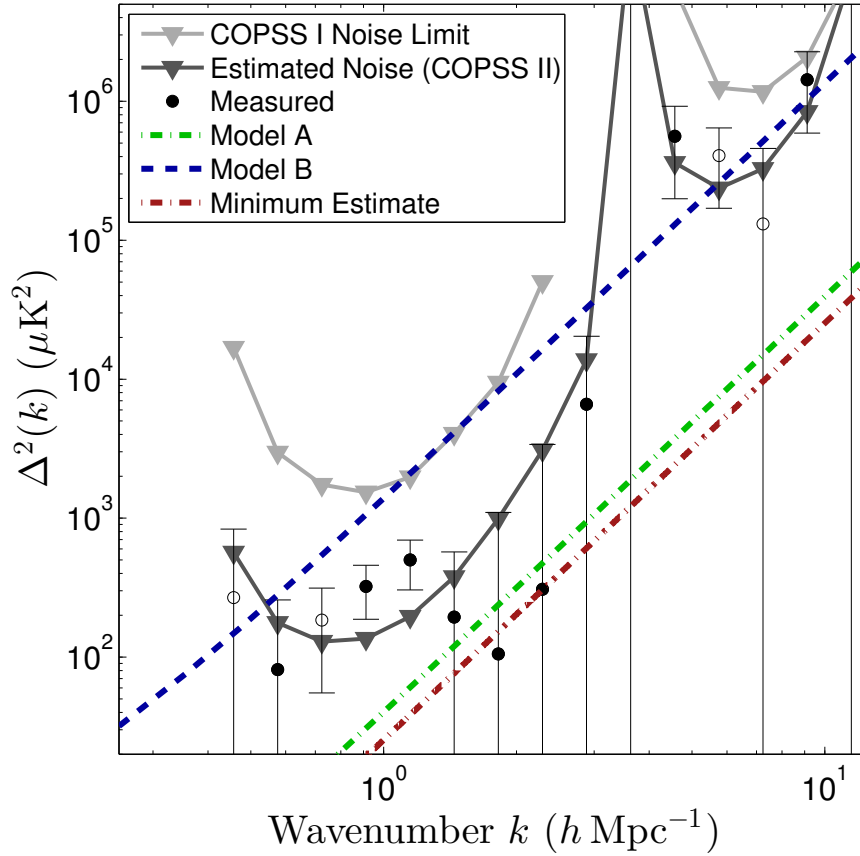


Figure 4.7: The power spectrum result from our analysis of the full COPSS data set, in the form $\Delta^2(k)$. Open circles represent negative values for $\Delta^2(k)$, while filled circles represent positive values, with error bars corresponding to the 1σ errors on our measured values. Model A (dot-dashed green) and model B (dashed blue) from Pullen et al. (2013) are shown for reference (discussed further in Section 4.4.2), along with the estimated RMS noise power (gray triangle), absent any astrophysical signal. Also shown the estimated power, based on detections of CO emission (with optical counterparts) found in Decarli et al. (2014). *Right:* The power spectrum result, in the form of $P(k)$.

4.4 Discussion

4.4.1 Constraints on the CO Power Spectrum

As previously discussed in Chapter 2, the power spectrum for CO as a function of wavenumber and redshift is given by

$$P(k, z) = \langle T_{\text{CO}} \rangle^2 b^2(z) P_{\text{lin}}(k, z) + P_{\text{shot}}(z), \quad (4.1)$$




Figure 4.8: The final power spectrum result, in the form of $P(k)$. Model A (dot-dashed green) and model B (dashed blue) from Pullen et al. (2013) are shown for reference, along with the estimated RMS noise power (gray triangle), absent any astrophysical signal.

where T_{CO} is the mean brightness temperature, $b(z)$ is the halo bias, P_{lin} is the linear matter power spectrum, and P_{shot} is the shot contribution to the power spectrum. Assuming a linear relationship between CO luminosity (L_{CO}) and halo mass, P_{shot} can further be defined as

$$P_{\text{shot}}(z) = \left(A_{\text{CO}} \frac{c^3(1+z)^2}{8\pi\nu_o^3 k_B H(z)} \frac{L_{\odot}}{M_{\odot}} \right)^2 f_{\text{duty}} \int_{M_{\text{min}}}^{\infty} M^2 \frac{dn(z)}{dM} dM, \quad (4.2)$$

where ν_o is the rest frequency of the line, $H(z)$ is the Hubble parameter $dn(z)/dM$ is the number of halos per unit mass as a function of redshift, f_{duty} is the duty cycle of CO emitters (i.e., the fraction of time a halo hosts CO-emitting galaxies), and A_{CO} is the ratio of CO(1-0) luminosity to host halo mass for CO-luminous halos, with units of $L_{\odot} M_{\odot}^{-1}$ (Lidz et al. 2011; Breyse et al. 2014). Halos with masses below the low-mass limit, M_{min} , are assumed to lack sufficient CO, and thus do not appreciably contribute to the larger-scale emission detectable




Figure 4.9: Measured variation of $P(k)$, averaged over all k , as a function of redshift, with the 1σ error bars shown for each value. We find that the results for individual bins are consistent with the mean power measured across all bins.

in our power spectrum measurement (e.g. Visbal & Loeb 2010; Lidz et al. 2011). While some uncertainty in M_{\min} exists, the shot power component is only weakly dependent on the choice of M_{\min} (as demonstrated in Figure 4.11).

Using Equation 4.2 and provided with an appropriate halo mass function (Tinker et al. 2008), one can use the upper limit presented in Section 4.3.1 to constrain the product $A_{\text{CO}} f_{\text{duty}}^{1/2}$. This constraint is shown in Figure 4.11 as a function of M_{\min} , compared to several estimates from theoretical models.

4.4.2 Constraints on A_{CO} and σ_{CO}

We note that values of f_{duty} and M_{\min} are not well-constrained; in particular, the value of f_{duty} differs by more than an order of magnitude between different models. P13 argues

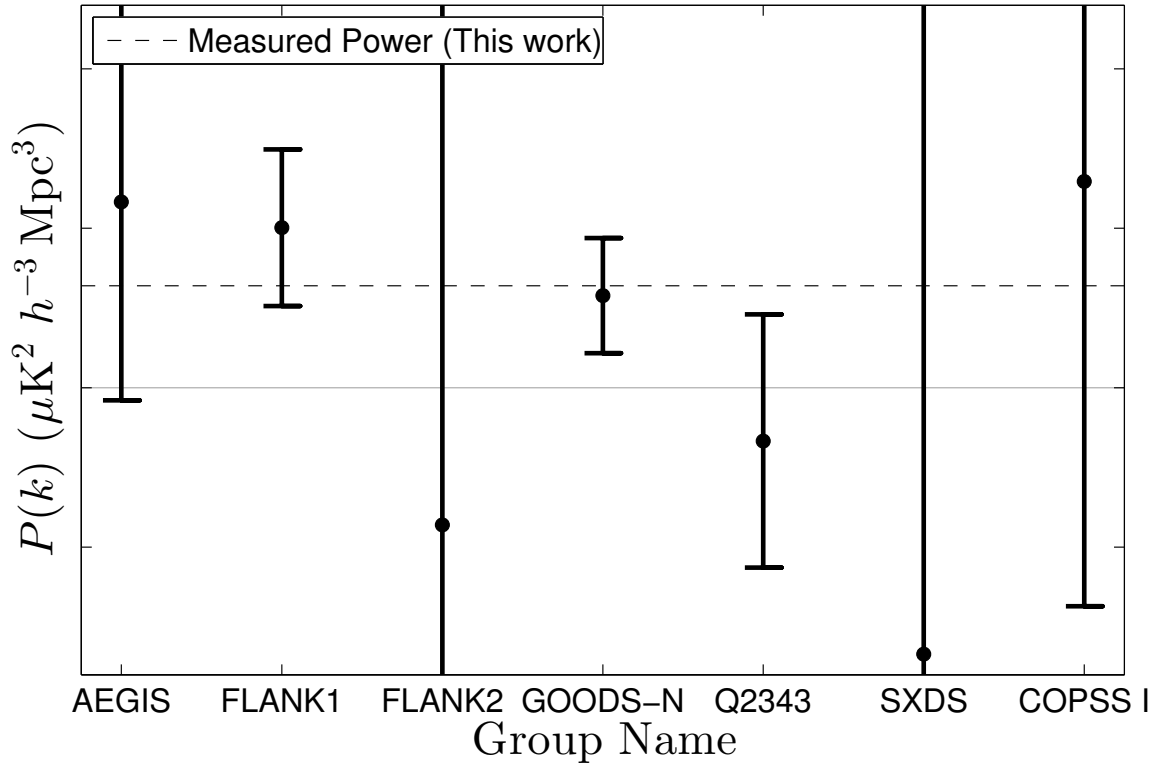


Figure 4.10: Power measured within individual groups of fields. The largest outlier (though still within the statistically nominal range) in the trim tests was the Q2343 group, which has power 1.2σ below what was measured across the entire survey.

$f_{\text{duty}} = t_s/t_{\text{age}}(z)$, which is ~ 0.05 for the redshift range of our measurement for their assumed $t_s = 10^8$ yr. This choice of t_s is based on arguments that the dominant source of CO emission at high redshift ($z \gtrsim 6$) will arise from galaxies undergoing extreme starburst events (Lidz et al. 2011; Righi et al. 2008). However, there are multiple observations indicating higher star formation duty cycles at $z \sim 1-4$, with f_{duty} approaching 100% (Noeske et al. 2007; Lee et al. 2009; Tacconi et al. 2013). The Li et al. (2016) model implicitly assumes a 100% duty cycle, though they introduce intrinsic scatter model parameters for the star formation rate to L_{CO} and halo mass to star formation rate relations to allow for observed variations in halo activity. Following Equation 4.2, our constraint on A_{CO} depends on $f_{\text{duty}}^{-0.5}$, so increasing it to unity would drop our limit on A_{CO} by a factor of 4, though the unaccounted scatter in halo properties noted by Li et al. (2016) would mimic a lower f_{duty} . We use the low value of f_{duty} from P13 for our conservative upper limit on A_{CO} , and show the combination of A_{CO} and f_{duty} in Figure 4.11. Variations in values of M_{min} have smaller effects on the model constraints, as can be seen in Figure 4.11. Uncertainties in the halo mass function (i.e., dn/dM) contribute insignificantly to the uncertainty. Tinker et al. (2008) found 1%




Figure 4.11: Constraints on A_{CO} as a function of M_{min} . The 1σ constraints from our analysis (gray) are shown versus several theoretical expectations for A_{CO} (Breysse et al. 2014), multiplied by the square root of their adopted values for f_{duty} ($f_{\text{duty}} = 0.1$ for Visbal et al. (2011), $f_{\text{duty}} = t_s/t_H$ for all others). As Righi et al. (2008) do not explicitly supply a value for f_{duty} or A_{CO} , for this model we adopt value for f_{duty} from Pullen et al. (2013) and employ the value of A_{CO} calculated for this model by Breysse et al. (2014).

accuracy for their fitting function over the range of masses of interest for this problem, which translates into a 1% uncertainty in the quantity of interest for determining A_{CO} , the second moment of dn/dM .

Because of the aforementioned problems with the f_{duty} parameter over the redshift range considered in our analysis, we instead introduce a term analogous to that found in Li et al. (2016) (hereafter L16): σ_{CO} , the log-scatter (in dex) of the correlation between halo mass and L_{CO} . We note that L16 disaggregate this term into the scatter in the underlying relationships




Figure 4.12: Constraints on A_{CO} versus σ_{CO} , with the 25% (red), 50% (orange), 68.3% (yellow), 90% (green) and 95.4% (blue) confidence limits shown. Also shown are the theoretical range of expectations from Li et al. (2016) (with 1σ errors) – as this model predicts a mass-dependent value for A_{CO} , we have used a mean value for A_{CO} , weighted by $M^2 dN/dM$ (i.e., the shot power contribution from halos of a given mass).

in the halo mass to CO luminosity correlation. As we are unable to separate those different sources of scatter in our measurement, we have chosen to use a single, aggregate term in our analysis. The scatter enters Equation 4.2 in place of f_{duty} as p_σ , defined as the fractional change in shot power induced by σ_{CO} . We further define p_σ as

$$p_\sigma = \int_{-\infty}^{\infty} \frac{10^{2x}}{\sqrt{2\pi\sigma_{\text{CO}}^2}} e^{(-x^2/2\sigma_{\text{CO}}^2)} dx. \quad (4.3)$$

We show our constraints on A_{CO} in Figure 4.11, as a function of M_{min} (for those models dependent upon f_{duty}). With the results from the analysis of the S10 dataset, and adopting

the values for f_{duty} and M_{min} found in P13 ($M_{\text{min}} = 10^9 M_{\odot}$; $f_{\text{duty}} = t_s/t_{\text{age}}(z)$, where t_s is the star formation timescale, of order 10^8 yr, and $t_{\text{age}}(z)$ is the Hubble time at a given redshift), we constrain $A_{\text{CO}} < 1.2 \times 10^{-5} L_{\odot} M_{\odot}^{-1}$. With analysis of the full dataset (under the same assumptions stated previously), we constrain $A_{\text{CO}} = 3.8_{-0.9}^{+0.8} \times 10^{-6} L_{\odot} M_{\odot}^{-1}$.

We also show in Figure 4.12 these constraints as a function of σ_{CO} . Under their model, L16 adopt a value of $\sigma_{\text{CO}} \approx 0.37 \pm 0.12$ dex – over the 95% confidence interval of σ_{CO} , the constraint on A_{CO} changes by as much as a factor of 40. This is the manifestation of the degeneracy between A_{CO} and p_{σ} in the combination of Equations 4.2 and 4.3 (although we weakly break this degeneracy by measuring the power of multiple groups of widely separated fields). Additional data at smaller k , where the clustering component of the power spectrum dominates, would help break this degeneracy because the clustering component is less affected by σ_{CO} . In the absence of constraints on σ_{CO} , we are restricted to placing only an upper limit on the host-halo mass to CO luminosity ratio, with $A_{\text{CO}} < 1.5 \times 10^{-6} L_{\odot} M_{\odot}^{-1}$ (corresponding to the 68% confidence limit where $\sigma_{\text{CO}} = 0$). However, if we adopt the value of σ_{CO} from L16 (marginalizing over the uncertainty in this parameter), we determine $A_{\text{CO}} = 6.3_{-2.5}^{+4.0} \times 10^{-7} L_{\odot} M_{\odot}^{-1}$.

One can also use a variant of Equation 4.2 to convert “blind” detections of CO emitters into an estimate for the minimum power. For this estimate, we will utilize those detections presented in Decarli et al. (2014). We consider only those galaxies with optical counterparts, as emitters without counterparts are more prone to being either spurious detections or incorrectly ascribed to the wrong redshift (by incorrectly identifying which rotational transition is being observed). Assuming $L'_{\text{CO}(3-2)}/L'_{\text{CO}(1-0)} = 0.5$ (Walter et al. 2014), we estimate the minimum shot power to be $P_{\text{CO},\text{min}} = 3.8_{-1.1}^{+3.7} \times 10^2 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$. We note that this minimum estimate resides below the power detected in our analysis, and resides near the bottom edge of the 95% confidence range for P_{CO} (as presented in Section 4.3).

We consider the models of P13 in a bit more detail, as their model B is closest to our upper limit. P13 present two models (A and B) that predict the CO power spectrum based on locally observed correlations between star formation rates, far-infrared luminosity and CO luminosity (Kennicutt 1998; Wang et al. 2011). Model A utilizes a predicted relationship between halo mass and SFR for power spectrum estimates, while model B uses SFR functions based on UV and IR observations (Smit et al. 2012).

As shown in Figure 4.11, our constraint on A_{CO} falls below estimates from Lidz et al. (2011), though those estimates are tailored for $z \gtrsim 6$. Our present constraints also lie below predictions of Model B from P13. These constraints also lie well above predictions made by Visbal & Loeb (2010), as does our estimate for the minimum shot power. We now briefly consider what we can learn from the exclusion of these models. At $z \sim 3$, the Visbal & Loeb model likely suffers from the fact that it is tailored for $z \geq 6$, and predicts global star formation rates (SFRs) that are a factor of a few different than what is observed at $z \sim 3$, although this difference alone does not completely account for the discrepancy between this model and our lower limit. Visbal & Loeb also use M82 as a template to calibrate their SFR- L_{CO} relationship (Weiß et al. 2005), which predicts CO luminosities that are a factor of a few below what is observed in massive main-sequence galaxies at $z \sim 2$ (Tacconi et al. 2013). Our

present constraints on P_{CO} suggest that molecular gas properties of local starburst galaxies (like M82) are not well-matched to that of normal galaxies at $z \sim 3$ (though they may still hold for higher redshift galaxies).

Model B of P13 uses the SFR function parameters from Smit et al. (2012) and a prescription for the SFR- L_{CO} relationship observed in local galaxies (Kennicutt 1998; Wang et al. 2011) to calculate an estimate for the mean brightness temperature of CO, and then use this to scale A_{CO} (versus what is predicted in their Model A). However, this adjustment overpredicts the number of extremely CO-luminous objects (i.e., $L_{\text{CO}} \geq 10^8 L_{\odot}$) – likely due to the inclusion of f_{duty} within their model (which here has the effect of concentrating CO emission within fewer, more luminous halos).

The CO abundance may change significantly over the redshift range of our measurement, owing in part to increasing metallicity and dust masses within galaxies (Valiante et al. 2009), rising feedback and quenching processes (Kereš et al. 2009; Wheeler et al. 2014) and the rapid depletion of neutral gas for star formation (Bauermeister et al. 2010) over cosmic time. Model B of P13 effectively makes such a prediction, with A_{CO} rising with decreasing redshift as one approaches the peak of cosmic star formation. As previously discussed in Section 4.3.1, we find weak evidence of decreasing power with increasing redshift.

4.4.3 Constraints on the CO Luminosity Function

Theoretical models indicate that our measurement should sample the shot noise portion of the CO power spectrum. Accordingly, we measure the second moment of the CO luminosity function at $z \sim 3$. The second moment of the luminosity function, $\int L^2 \Phi(L) dL$, is related to the shot power by

$$P_{\text{shot}}(z) = \left(\frac{c^3(1+z)^2}{8\pi\nu_o^3 k_B H(z)} \right)^2 \int L^2 \Phi(L) dL. \quad (4.4)$$

One can use the value for the second moment, in combination with data from direct detection efforts, to place constraints on the shape of the luminosity function. To do so, we will assume that the luminosity function is (to first order) well-described by the Schechter function (Schechter 1976), which is nominally parameterized by a high-luminosity cutoff, L_* , a low-luminosity power law index, α , and a normalization factor for the overall density of luminous sources ϕ_* . For our analysis, we evaluate the likelihood of the combined choice of these three parameters by evaluating the second moment of the luminosity function produced. We will further weight this likelihood by

1. the galaxies detected in CO(3-2) with optical counterparts by Decarli et al. (2014),
2. a lack of detections of individual emitters within the COPSS dataset of $\geq 5\sigma$ significance with twice the FWHM of the primary beam, and
3. a prior on the slope of low-luminosity end of the luminosity function.

The search for individual emitters within our data was performed assuming a Gaussian emission profile (of width $\Delta v = 300 \text{ km s}^{-1}$, consistent with observations of Decarli et al. 2014). Due to the relatively coarse channelization, the search for individual emitters was conducted by searching only single and 2-channel averaged maps for any points above the threshold of 5σ . Under the Schechter parameterization, our measurement is generally more sensitive to changes in ϕ_* and L_* , and less sensitive to changes in α . To provide some constraint on the α parameter, we consider that given the SFR to CO luminosity relationship appears roughly linear for high redshift main-sequence galaxies at $z \sim 2$ (Tacconi et al. 2013). In this case, the SFR and CO luminosity functions should share similar values of α . Based on measured SFR function parameters presented in Smit et al. (2012), we adopt a conservative prior of $\alpha = -1.5 \pm 0.75$.

In including data from Decarli et al. (2014), we consider only those galaxies with optical counterparts, as emitters without counterparts are more prone to being either spurious detections or incorrectly ascribed to the wrong redshift (by incorrectly identifying which rotational transition is being observed). In evaluating the likelihood of any set of parameters for the luminosity function, we weight each particular parameter by \mathcal{L}_{gal} , the likelihood of observing *at least* the number of objects detected in any particular survey. We further define \mathcal{L}_{gal} as

$$\mathcal{L}_{\text{gal}} = 1 - \sum_{n=n_{\text{gal}}+1}^{\infty} \text{Pois}(n; V_z \rho_{\text{gal}}) \quad (4.5)$$

In Equation 4.5, n_{gal} is the number of galaxies detected within a particular bin, ρ_{gal} is the expected number density of galaxies (based on the set Schechter parameters being evaluated), and $\text{Pois}(k; \lambda)$ is the probability of detecting k objects given a Poisson distribution with mean λ .

The results of our likelihood analysis are shown in Figure 4.14. With our data (along with the constraints and priors mentioned earlier), we constrain $\phi_* = 1.3_{-0.7}^{+0.6} \times 10^{-3} L_{\odot}^{-1} \text{ Mpc}^{-3}$ and $L_* = 4.5_{-1.9}^{+1.4} \times 10^{10} \text{ K km s}^{-1} \text{ pc}^{-2}$ to 68% confidence. We use these constraints, along with importance sampling, to generate a fit (and 1σ errors) on the CO luminosity function at $z \sim 3$. Our fitted luminosity function agrees with earlier constraints made by Walter et al. (2014), as well as model predictions made in L16 and Sargent et al. (2014). Our fitted function also appears to disfavor those model predictions made by Obreschkow et al. (2009b), Lagos et al. (2011), and Popping et al. (2016), in that all three appear to underpredict the number of higher luminosity objects at $z \sim 3$.

4.4.4 Constraints on Cosmic Molecular Gas Abundance

Assuming a linear relationship between L_{CO} and M_{H_2} , one can use the CO luminosity to molecular gas mass conversion factor, α_{CO} , to determine the mass fraction (with respect to the host halo mass) of the molecular gas within galaxies, f_{H_2} . For a Milky Way-like $\alpha_{\text{CO}} = 4.3 M_{\odot} (\text{K km s}^{-1} \text{ pc}^{-2})^{-1}$ (Frerking et al. 1982; Dame et al. 2001) – equivalent to $8.7 \times 10^4 M_{\odot} L_{\odot}^{-1}$ (Solomon et al. 1992) – our constraint on A_{CO} translates to a limit on

the molecular gas mass fraction of $f_{\text{H}_2} = 5.5_{-2.2}^{+3.4} \times 10^{-2}$. As galaxies within “massive-but-common” host halos are expected to dominate our shot power measurement, we note that our constraint on f_{H_2} is primarily applicable to halos with masses of order $\sim 10^{12} M_\odot$ over our redshift range.

Our measurement can also be used to infer the cosmic density of H_2 ($\rho(\text{H}_2)$). These constraints are presented in Figure 4.15. Theoretical models from Popping et al. (2015) suggest that f_{H_2} has some dependence on halo mass; we therefore assume that for halos below some characteristic mass, $M_0 = 5 \times 10^{11}$, f_{H_2} scales linearly as M/M_0 . Adopting the fiducial value from L16 of $\sigma_{\text{CO}} = 0.37 \pm 0.12$, we constrain the cosmic density of H_2 to $\rho_{z \sim 3}(\text{H}_2) = 1.1_{-0.4}^{+0.7} \times 10^8 M_\odot \text{Mpc}^{-3}$. Several theoretical predictions for the cosmic molecular gas density at $z \sim 3$ lie within the 68% confidence interval of our constraint. Compared to previous work, our constraint is in agreement with previous work presented in Walter et al. (2014), who constrain $\rho_{z \sim 3}(\text{H}_2)$ between $5 \times 10^6 - 4 \times 10^9$ (68% confidence interval). As Walter et al. only account for galaxies above their detection threshold, we note that their measurement should be taken as a low-limit alone. In this way, our measurement offers an improved constraint on the maximum cosmic CO(1-0) luminosity at $z \sim 3$, as it is an integrated measurement across the entire population range of galaxies.

We note that our choice α_{CO} , M_0 , σ_{CO} , and the precise scaling relationship between halo mass and molecular gas mass fraction have a non-negligible impact on this estimate. Our analysis suggests that of these factors, the uncertainty in value of σ_{CO} has the strongest impact on our estimate (as demonstrated in Figure 4.15). We also note that models from Popping et al. (2015) suggest that for halos above some limit ($\gtrsim 10^{13} M_\odot$), the predicted value of f_{H_2} drops precipitously. Halos of such mass are expected to be rare, such that they do not meaningfully contribute to either the shot power or overall cosmic molecular gas density. Given the uncertainty of our measurement, this accounting has a negligible impact on our constraint.




Figure 4.13: The constraints on the individual Schechter parameters of the luminosity function: ϕ_* , α and L_* .




Figure 4.14: *Left:* The constraints on the individual Schechter parameters of the luminosity function: ϕ_* , α and L_* . *Right:* The fitted luminosity function and 1σ errors (red), along with the direct detection constraints from the COPSS data set (dashed), and the constraint on the CO luminosity function constraint at $z = 2.75$ (light blue) from Walter et al. (2014). Shown for comparison are models from Obreschkow et al. (2009a) (orange), Lagos et al. (2011) (dark green), Sargent et al. (2014) (dark blue), and Popping et al. (2016) (yellow).

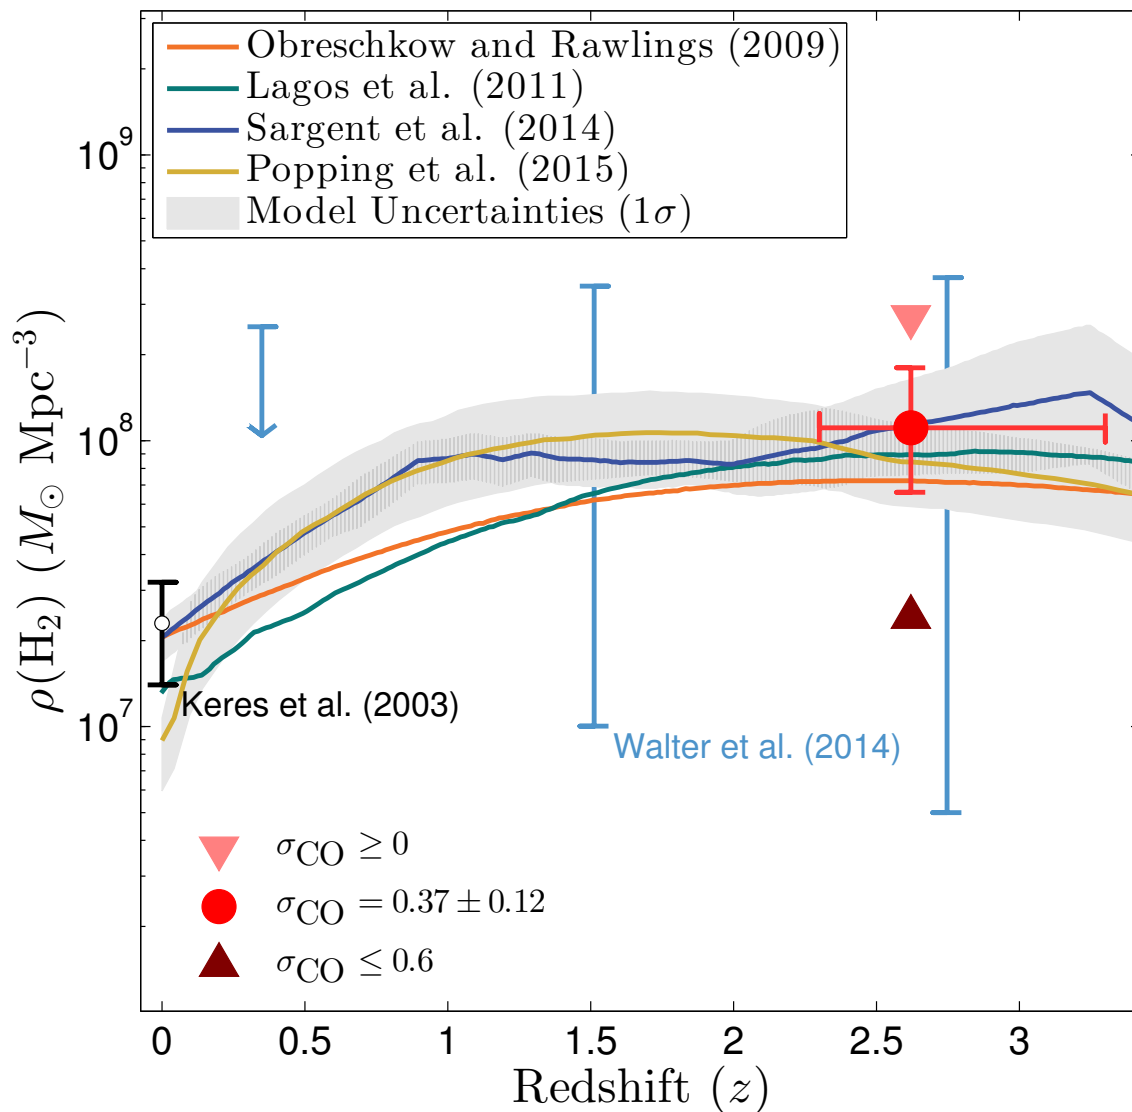


Figure 4.15: The constraint on the cosmic molecular gas density. The point labeled as $\sigma_{\text{CO}} \geq 0$ marks the upper limit of the 68% confidence region under the assumption that $\sigma_{\text{CO}} = 0$. The point labeled as $\sigma_{\text{CO}} \leq 0.6$ similarly marks the low limit of the 68% confidence region under the assumption that $\sigma_{\text{CO}} = 0.6$. Shown for comparison are the constraints from [Walter et al. \(2014\)](#) (light blue). Also shown are the theoretical expectations from [Obreschkow & Rawlings \(2009\)](#) (orange) and [Lagos et al. \(2011\)](#) (green), [Sargent et al. \(2014\)](#) (dark blue), and [Popping et al. \(2015\)](#) (yellow). Shown in gray are the 68% confidence limits on the [Sargent et al.](#) and [Popping et al.](#) models.

Chapter 5

Simulations

As intensity mapping measurements are entirely statistical in nature, simulations play an important role in verifying our analysis techniques, probing for sources of statistical bias, and measuring the potential impact of systematics and other containments on our measurement. To test the fidelity of our pipeline for a power spectrum measurement, as well to measure the impact of continuum foregrounds and systematics, we perform a series of simulated observations with synthetic data.

For this work, we use computer-generated mock fields, generating visibilities based on a randomly distributed set of sources. Effects such as primary beam attenuation, pointing errors, gain errors and mode-mixing (i.e., bandwidth smearing) are taken into account when creating visibilities of the mock fields, in order to better understand the limits of our methods of analysis. When appropriate, instrumental noise is added to the visibilities assuming a Gaussian distribution, using estimates based on system temperature and instrument parameters. We create simulated data for several fields (typically of order 10^4) for each set of simulations to help reduce sampling variance. Data from these mock fields are then processed through the same pipeline described in Chapter 3, and the resultant power spectra for the fields are then summed together to provide the final simulation results.

This chapter is structured as follows: Section 5.1 describes simulations used to verify the the analysis pipeline. Section 5.2 describes estimates of the impact of cosmic variance for a survey like COPSS. Section 5.3 describes simulations used to estimate the impact of sources of continuum emission (point sources and otherwise) on our measurement. Section 5.4 discusses instrument simulation work, performed to examine the impact of various systematic errors on a power spectrum measurement.

5.1 Pipeline Verification

We first look to verify that our analysis pipeline and processing methods do not significantly affect our measurements. For these simulations, we adopt the instrument configuration for the SZA as it was for the S10, simulating observations of groups of four lead-trail fields in order to test the impact of our method of ground subtraction. Mock fields are set to

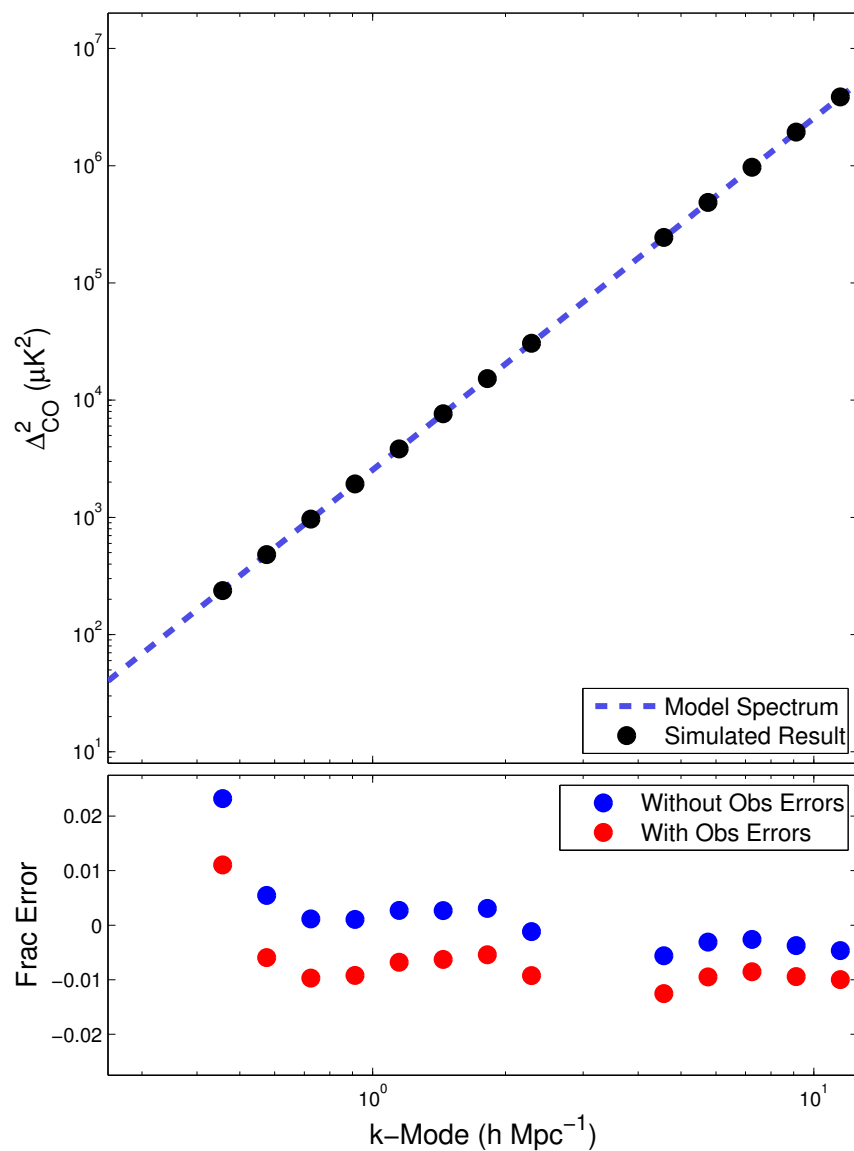


Figure 5.1: *Top:* The expected power spectrum (dashed) versus the mean measured power from our simulated datasets (circles). *Bottom:* The fractional difference between the mean measured power and the expected power, with basic instrumental errors (red circles) and without (blue circles). For both cases, we find that our analysis software agrees with the expected power to within a percent.

a declination of $\delta = 0\text{d}00'$, with simulated observations running over an hour angle $\pm 3\text{h}$, roughly matching those observations described in Chapter 4. In these tests, we run two different sets of simulations: one with pointing and gain errors, and one without. For those simulations that include errors, we assume RMS pointing errors of $0.3'$ and gain errors of

0.5%, consistent with the estimated errors for the second phase of the COPSS survey. From these simulations we additionally find three results that are noteworthy. For our source population, we include only spectral line sources of equal luminosity ($10^6 L_\odot$) with mean number density $0.1 (\text{Mpc}/h)^{-3}$, randomly distributed over a volume covered over 0.25 square degrees between $z = 2.2 - 3.4$. We assume for simplicity that line width is significantly narrower than the channel width of the simulated data. Under this configuration, our simulated measurement should be equal to $P_{\text{CO}} = 5.3 \times 10^4 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$

The results of these simulations are shown in Figure 5.1; we find excellent agreement between the expected and measured power from our analysis. First, our use of the Fourier transform (along the frequency axis) prior to converting the data into temperature units does result in a small overestimation of the power spectrum. Since the Fourier transform is implemented before the temperature unit conversion, a single frequency for all channels in the window must be used when converting between flux density and temperature units. As the flux density scales as $S \propto \nu^2$ at fixed brightness temperature, this introduces a bias in our measurement, although the fairly narrow frequency range of a single spectral window of the SZA ($\Delta\nu/\nu \sim 10^{-2}$) makes this effect negligible. Second, our choice to grid the data for the sake of computational expediency results in some minor power spectrum estimation errors, as the value of C in Equation 3.2 will be affected by the rounding of \mathbf{k} for each visibility. Though we have chosen our to minimize the grid cell size to reduce this effect, some errors are still seen in our measurement. This effect is strongest for delay-visibilitys that sample the (u, v, η) domain with only marginal overlap, affecting baselines that are located in more sparsely sampled regions of the uv -plane (such as those longest baselines of the SZA). Finally, we find that small pointing and gain errors have only a modest effect on the measured power, generally leading to $\sim 1\%$ underestimates of the measured by our analysis. After additional simulations, we note that calibration errors of all varieties will reduce the sensitivity of our measurement (as these errors will “wash out” the signal of interest), by roughly a factor equal to the fractional error in the calibration. Given that one expects calibration errors are of order one part in a hundred, we expect this loss of sensitivity is likely to be minimal. Ultimately, we estimate the difference between the measured and true power in our analysis to be of order one percent – inconsequential in the case of a null result or weak detection, though it may be limiting in the case of a very strong ($\gg 10\sigma$) detection. We otherwise find that our simulated results are in agreement with what theory predicts.

5.2 Impact of Cosmic Variance

To measure the impact of cosmic variance, we simulate a series of fields to match those observed as part of the COPSS survey (weighting the power spectrum from each field by the estimated sensitivity achieved in Chapter 4). For all simulations, we assume a linear relationship between CO luminosity and host-halo mass (the ratio of which is given by A_{CO}). We perform three sets of simulations: one with a fixed value of A_{CO} , one where emitters have a duty cycle f_{duty} (i.e., only this fraction of halos is CO luminous), and one

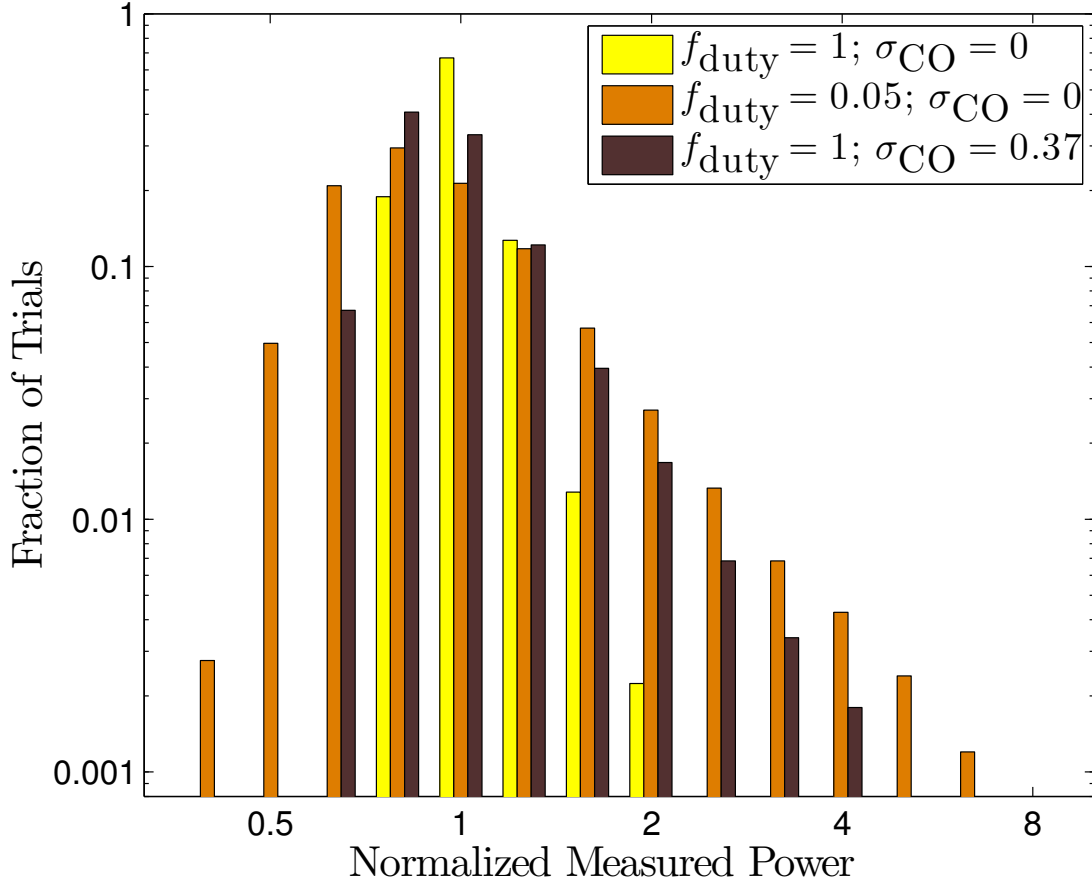


Figure 5.2: The relative likelihood of a randomly distributed group of emitters dispersed over a volume probed by the COPSS survey to produce a given amount of shot power, normalized by the “true power” (i.e., what would be measured over an infinitely large volume), under different assumptions of f_{duty} and σ_{CO} .

where there exists some scatter in the halo mass to CO luminosity relationship (σ_{CO}). Halos over a volume covered over 0.25 square degrees between $z = 2.2 - 3.4$, with number densities following theoretical models from [Tinker et al. \(2008\)](#). Results of the simulations are shown in Figure 5.2.

In the shot-power regime, there are two sources of cosmic variance: large-scale structure inducing a local over/underdensity over the area of our measurement ([Tegmark et al. 1998](#)), and the Poisson noise associated with the limited number of massive halos (which are the primary contributors to the shot power). The power measured is roughly proportional to the number density of emitters within the volume measured – specifically, the number density of

“luminous-but-common” emitters that are contributing most to the shot component of the power spectrum. Assuming that CO luminosity is linearly proportional with host halo mass (e.g., Lidz et al. 2011; Pullen et al. 2013), at $z \sim 3$ halos with masses of order $10^{12} M_{\odot}$ will dominate the shot power measurement (Tinker et al. 2008). For the COPSS measurement – with a total survey volume of $4.9 \times 10^6 h^{-3} \text{Mpc}^3$ – a total of $\sim 10^4$ halos with masses $\gtrsim 10^{12} M_{\odot}$ are expected to contribute to the measurement, translating to cosmic variance-induced errors of a few percent in our measurement (which is a factor of several greater than that induced by large-scale structure over our survey area).

However, if one assumes some duty cycle of CO emission ($f_{\text{duty}} < 1$), then cosmic variance will have a much greater impact on our measurement. As the expected number of contributing halos scales with f_{duty} , theoretical estimates of $f_{\text{duty}} \approx 0.05$ would reduce the expected number of luminous objects within the COPSS survey volume to only a few hundred. Under this scenario, from our simulations we estimate cosmic variance-induced errors of $\Delta P/P \approx 0.36$. However, as previously discussed, observational data suggests that f_{duty} is likely close to unity, which would imply that this estimate overpredicts the impact of cosmic variance, and thus regard it strictly as an upper limit for our measurement. For a more accurate estimate, we instead consider the case that some scatter in the host-halo mass to CO luminosity relationship ($\sigma_{\text{CO}} \neq 0$) exists. This scatter will have the effect of reducing the number of halos that appreciably contribute to the power measurement (i.e., a subset of this population will become slightly more luminous, and hence will contribute more to the shot power). Adopting a fiducial value of $\sigma_{\text{CO}} = 0.37$ and assuming a linear scaling between halo mass and L_{CO} , we estimate cosmic variance-induced errors of $\Delta P/P \approx 0.17$ in the COPSS measurement. In the case of a moderate or weak detection ($\lesssim 4\sigma$), this contribution is insignificant in comparison to thermal noise, and the impact cosmic variance will be minor. We have therefore neglected cosmic variance in our power spectrum estimates, though have included its impact in our calculations for A_{CO} , $\rho(\text{H}_2)$, and the CO luminosity function parameters.

We note that some models (e.g., Pullen et al. 2013; Li et al. 2016) effectively predict that above some characteristic mass (of order $10^{12} M_{\odot}$), CO luminosity may remain flat or decrease with increasing host halo mass. Such a relationship would decrease the estimated contribution of very massive halos ($\sim 10^{13} - 10^{14} M_{\odot}$), which do not add significant power overall but may otherwise be the dominant driver of cosmic variance in smaller volumes $\lesssim 10^6 \text{Mpc}^3$. We therefore consider our estimate of $\Delta P/P \approx 0.17$ to be a reasonable but conservative estimate for cosmic variance.

5.3 Continuum Point Sources

Continuum point sources have the potential to significantly contaminate our measurements. Point sources provide of order 10^3 more power (in units of μK^2) than CMB anisotropies or Galactic foregrounds over the range of multipoles ($\ell \sim 3000$) that the SZA is sensitive to ((S10); Bennett et al. 2013; Planck Collaboration et al. 2015). As such, we expect point

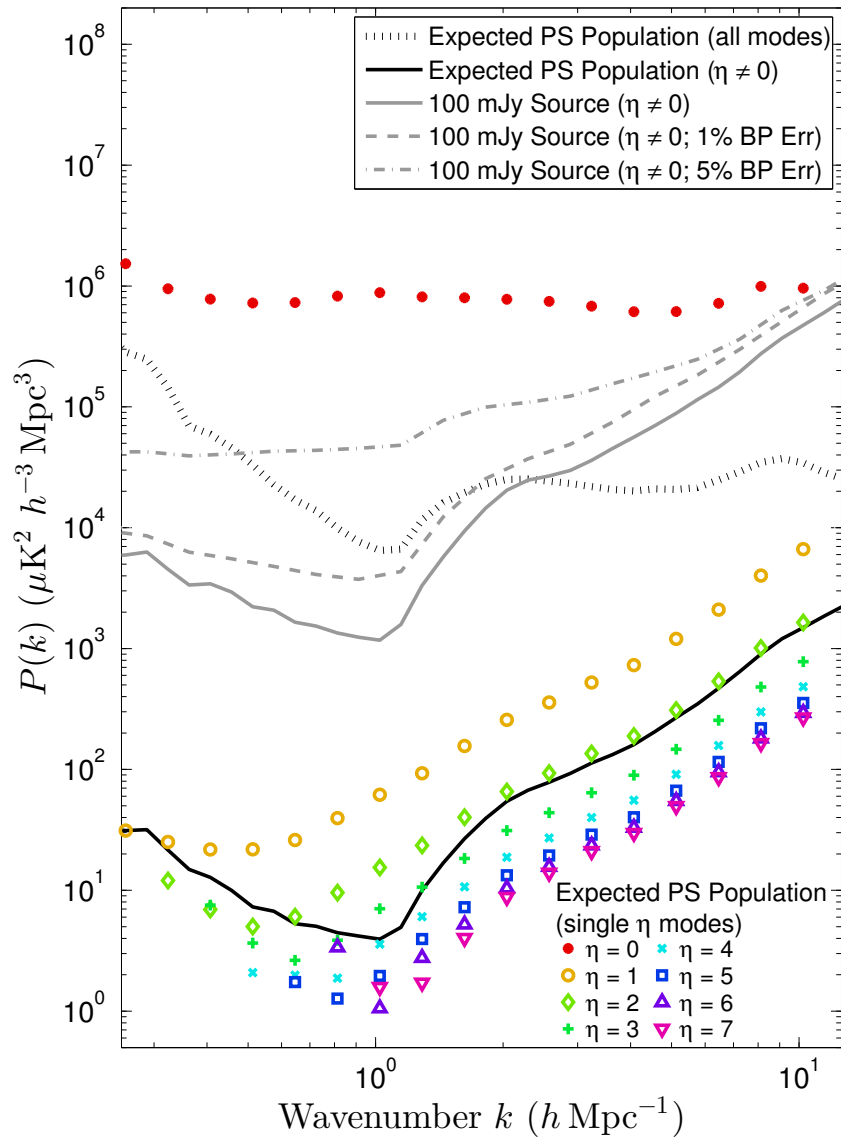


Figure 5.3: Results of simulations to probe the impact of point sources on our analysis. Separate simulations were used to probe the impact of a bright 100 mJy source (gray solid) and of the expected population of point sources. The measured power, averaged over three-dimensional shells in k , is shown before (black dotted) and after (black solid) removal of the $\eta = 0$ modes. For the 100 mJy source simulations, we additionally show the impact that bandpass errors have on the measured residual power, for RMS errors of 1% (gray dashed) and 5% (gray dot-dashed). For the expected point source population simulations, we additionally show the power measured when using data from a single η channel.

sources to be the dominant astrophysical contaminant in our measurements. Fortunately, most sources at 30 GHz will be close to spectrally flat over the frequency range of one window (Muchovej et al. 2010), and therefore will contribute power primarily to Fourier modes with $\eta = 0$. To reject this contamination, we drop the $\eta = 0$ channel. Dropping this channel does not significantly affect the k range the experiment, as the length of the smallest three dimensional k vector is determined primarily by the shortest baselines, rather than the smallest η mode.

We test the effectiveness of our method of rejecting continuum source power with three sets of simulations. First, we simulated observations of mock fields with a 100 mJy point source of varying spectral index, inserted at a random position within twice the FWHM primary beam. Such a source is much brighter than average, and is somewhat brighter than the brightest single source found in our data. Second, we simulated observations of mock fields following the 30 GHz point source counts and spectral index distributions measured by the SZA in Muchovej et al. (2010) (which we refer to as the “expected point source population”). We simulated sources of flux density up to 15 mJy, above which observations show a steep drop in point source counts. Third, we introduced calibration errors (i.e., baseline length, gains and bandpass) into both our expected point source population and 100 mJy source simulations.

The results of these simulations are shown in Figure 5.3. For simulations probing the impact of the expected point source population, we additionally present a breakdown of the power measured in Fourier modes for individual η channels. Power from continuum sources is predominantly found in the $\eta = 0$ channel; as the one-dimensional power spectrum is an average over all remaining η , the effect of contamination is an average of the η channels shown.

In the absence of calibration errors, we find that rejection of the $\eta = 0$ mode reduces the contribution of continuum sources by factors of 10^1 to 10^4 , depending on k . We find that the residuals rise roughly as k^2 for $k \gtrsim 2 h \text{Mpc}^{-1}$. This is due to the appreciable change of the baseline length over the bandwidth of a window for the long baselines of the array – an effect sometimes referred to as ‘mode-mixing’ (Datta et al. 2010; Morales et al. 2012). Our experiment does not achieve interesting sensitivity to these modes, so we do not expect that this effect requires more thorough mitigation. We find that rejection of the $\eta = 0$ channel reduces the power spectrum contributions of a 100 mJy point source to levels comparable to the sensitivity of our experiment. In the real data, the suppression of power from bright point sources should be greater, as we remove such sources from the visibilities before calculating power spectra. We also find that rejection of the $\eta = 0$ channel reduces the power spectrum contributions of the expected source population to levels well below our sensitivity threshold. Further testing reveals that the point source residual power decreases by an additional factor of ~ 10 after removing sources of flux density greater than 1.5 mJy (roughly corresponding to the flux limit for source subtraction for the S10 data set).

In simulations that include calibration errors, we find that bandpass errors can significantly increase the residual power from continuum point sources. We find that for RMS fractional bandpass errors, σ_{BP} , between 1% and 10%, the residual power scales as σ_{BP}^2 . We

expect that our bandpass errors are only 1% (discussed in Section 3.2); our simulations indicate that this will increase point source residuals by 30-300%, depending on k . This effect is negligible for the expected point source population, though it may allow residuals from bright point sources to exceed our detection threshold. We note these simulations represent a worst-case scenario, as only a single set of bandpass errors are incorporated into each mock field. In real data, a new bandpass solution is derived for each observing block; hence, data for a typical field implements 45 different bandpass solutions. Assuming that the errors in these bandpass solutions are thermal-noise dominated, we expect the increase in residual power to be only a few percent of what we have found in these simulations. We further find that gain and baseline errors do not significantly increase the residual point source power, though they do impact the fidelity with which we are able to remove bright sources from our data set.

5.4 Instrumentation Simulations for Future Experiments

As noted previously, intensity mapping experiments require good characterization of instruments to ensure for proper control of systematics and other contaminants. While experiments with the SZA benefited from previous characterization efforts, the development of the next generation of instruments requires more sophisticated simulation work in order to understand what features and design specifications are most critical for future experiments. Of particular interest for CO intensity mapping is ASIAA Intensity Mapping CO (AIM-CO) project, with the Y. T. Lee Array (Ho et al. 2009). The YTLA is a 3mm instrument, with frequency coverage between 86 and 102 GHz, making it sensitive to both CO(3-2) ($z = 2.3 - 3.0$) and CO(2-1) ($z = 1.2 - 1.7$) transitions. Presented here are the results of a series of sophisticated simulations for a wide array of observing scenarios with the YTLA. We note that while these results are most directly applicable for future observations of the YTLA, the software developed for these simulations is flexible enough to produce sensitivity estimates and explore systematics of both existing and hypothetical interferometers.

For these simulations, we test the YTLA under two different 7-element hexagonal configurations: 60 cm dishes with 0.6 meter spacings (YTLA-7 \times 60cm), and 120 cm dishes with 1.4 meter spacings (YTLA-7 \times 120cm). For both series of tests, we simulate dual polarization measurements, with system temperature 80 K and an aperture efficiency of $\eta = 0.5$ with 16 GHz bandwidth (channelized such that the mean velocity width for each channel is 300 km/s). For these test, we simulate three different scenarios. First, we simulated observations of mock fields with a 1 Jy point source of varying spectral index, inserted at a random position within twice the FWHM primary beam (much brighter than the typical point source at 90 GHz). Second, we simulated observations of mock fields following the 90 GHz point source counts presented in Davies et al. (2013). Finally, we simulate fields with spectral line emitters of mean number density $0.1 \text{ (Mpc}/h)^{-3}$, randomly distributed over a volume covered over 1 square degree between $z = 2.3 - 3.1$ (simulating emission from the CO(3-2) transition). These simulated data are processed through the same analysis pipeline used for

the SZA simulations described in Sections 5.1 and 5.3.

For both array configurations, we find that rejection of the $\eta = 0$ mode leaves residual power of $10 \mu\text{K}^2 h^{-3} \text{Mpc}^3$ from the expected point source population, well below the sensitivity achievable in a year long survey (assuming 4000 hours integration time spread across 12 fields). However, at 90 GHz a 0.6 meter baseline has primary sensitivity to $\ell \sim 1200$ (which for a 60-cm antenna, will provide some sensitivity to multipoles as low as $\ell \sim 600$), where primary CMB anisotropies will contribute approximately ~ 100 times more power than the background point source, of order $P(k) \sim 10^7 \mu\text{K}^2 h \text{Mpc}^{-1}$. Our SZA simulations found that generally, the maximum suppression achieved by simply rejecting the $\eta = 0$ channel was 10^5 , implying that the residual power from the CMB on the shortest baselines will be $P(k) \sim 102 \mu\text{K}^2 h \text{Mpc}^{-1}$ – approaching the power predicted by Model A from Pullen et al. (2013) (assuming $L'_{\text{CO}(3-2)}/L'_{\text{CO}(1-0)} = 0.5$). Under nominal conditions, this level of contamination is still generally below the sensitivity one could achieve with the YTLA $7 \times 60\text{cm}$ configuration with a year long observing campaign. However, bandpass errors have the potential to increase the residual power in all channels; a 1% (per 300 km/s channel) error would increase the residual power by as much as a factor of 10. Assuming these errors are noise-dominated, this effect would be somewhat mitigated as bandpass errors between different observing blocks would be uncorrelated, though the small collecting area of the 60 cm dishes will make proper bandpass calibration difficult with typical astronomical sources. Therefore, some attention will need to be paid to ensure that bandpass solution errors do not lead residual CMB power to overtake the power from the signal of interest.

5.4.1 Measuring the Impact of Platform Deformation

The YTLA is a platform-based interferometer (with all antennas sitting on a planar surface), using a hexapod mount system to tip, tilt and rotate the platform to different positions on the sky. The advantage of a platform-based system is that keeps all elements coplanar to the sky, permitting baselines to stay in the same position of the uv -plane, maximizing redundancy and thus sensitivity (e.g., Parsons et al. 2012; see Appendix A for further details). One disadvantage is that mechanical stresses (such as those imposed by gravity) may lead the platform to warp. With the YTLA, this generally leads to three-fold axisymmetric vertical displacements and pointing errors for individual antennas. Based on previous work performed by Liao et al. (2013), we estimate the vertical displacement, dz , for a given position on the platform (x, y) is given by

$$dz = A(\phi_{\text{az}}, \phi_{\text{el}})[(x^2 - y^2)\cos 2\theta + 2xy\sin 2\theta], \quad (5.1)$$

where θ is the phase of the saddle pattern, and A is the amplitude of the pattern as a function of azimuth (ϕ_{az}) and elevation (ϕ_{el}) angles. For our simulations, A is further defined as

$$A(\phi_{\text{az}}, \phi_{\text{el}}) = (3 \times 10^{-5} + (7 \times 10^{-5} \cos \phi_{\text{el}} (1 + \cos 3\phi_{\text{az}}))) \text{ m}^{-1}, \quad (5.2)$$

in approximate agreement with measurements performed by Liao et al.. For these tests, we measure three properties. First, we measure the effective loss in integration time due to

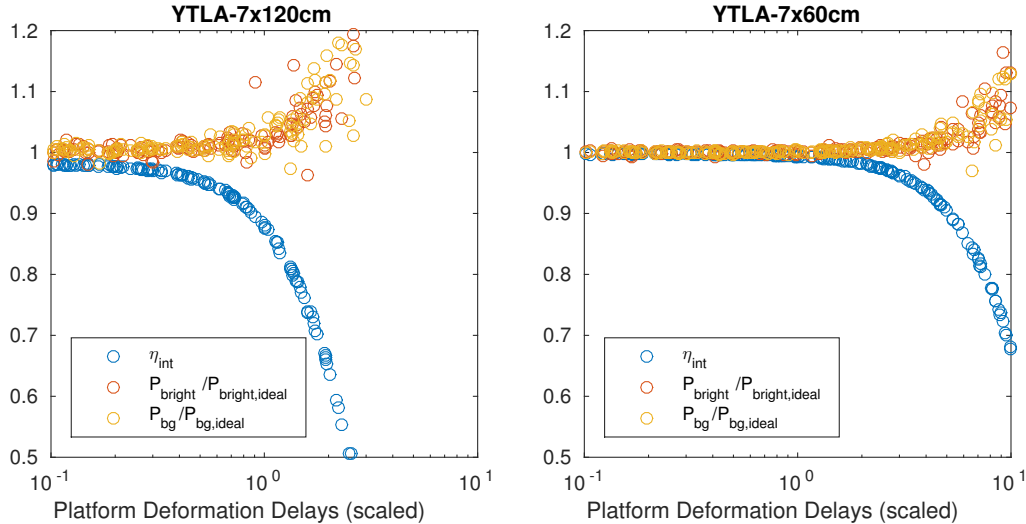


Figure 5.4: The impact of deformation-induced delay errors on a hypothetical measurement with the YTLA. For both configurations, we also show the effect of minimizing and exaggerating the delay errors on our simulated measurement. Shifts in delay will primarily produce changes in phase for the interferometer, appearing as a cyclic variation in gain phase with time. Regular gain calibrator observations would help to somewhat reduce this affect (as calibrator and target fields should be coincident in the sky, and thus should have similar delay errors), although proper modeling of the platform deformation is likely to be much more effective in reducing these errors.

systematics (reported as the observing efficiency η_{int}) by taking the ratio of the measured power with and without systematic effects added. Second, we measure the relative increase in contaminant power due to bright point sources ($P_{\text{bright}}/P_{\text{bright,ideal}}$), after correcting for η_{int} . Finally, we measure the relative increase in contaminant power due to background point sources ($P_{\text{bg}}/P_{\text{bg,ideal}}$, once again correcting for η_{int}). For platform deformation, we are primarily concerned with delay and pointing related errors. Delay errors are linearly proportional to dz , whereas pointing errors are calculated taking the inverse tangent of the partial derivatives $\partial(dz)/\partial x$ and $\partial(dz)/\partial y$ to find the elevation and cross-elevation pointing errors, respectively. As these effects may be correctable in post-processing (e.g., delay errors could be calibrated out using models of the platform deformation), we measure the effect of scaling either error separately. We simulate 100 mock fields for each choice in error scaling, choosing 200 randomly chosen (logarithmically distributed) scaling factors, applied to the estimated delay or pointing errors.

Shown in Figures 5.4 and 5.5 are measured effects of platform deformation-related errors on our simulated measurement. For YTLA-7 \times 60cm, we find that the effects of platform deformation are negligible, due in large part to the extremely compact nature of this configuration. For nominal pointing errors with YTLA-7 \times 120cm, we find that $\eta_{\text{int}} \approx 0.9$, with point source contributions increasing by a few percent. These effects are largely due to

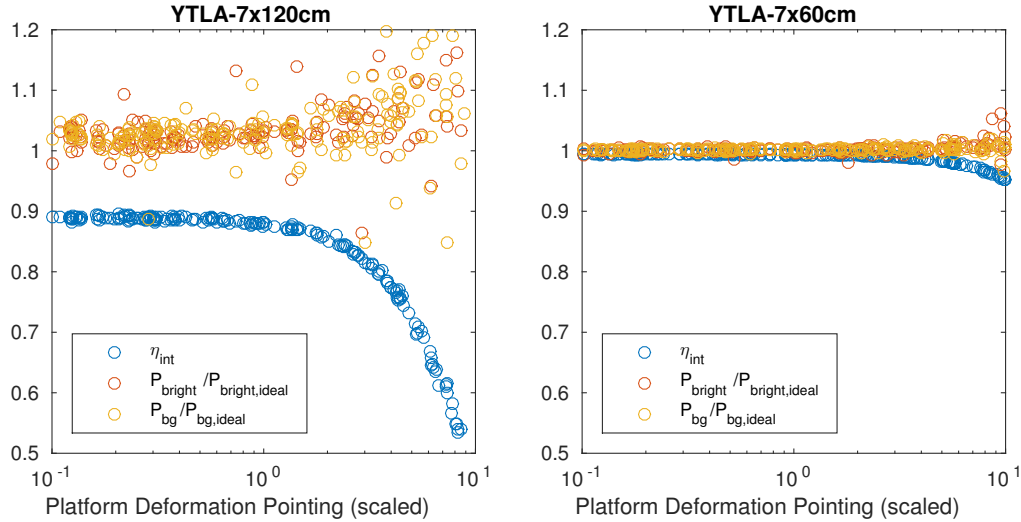


Figure 5.5: The impact of deformation-induced pointing errors on a hypothetical measurement with the YTLA. For both configurations, we also show the effect of minimizing and exaggerating pointing errors on our simulated measurement. As the deformation pattern is bisymmetrical, the effect of is most significant for antennas on opposite sides of the platform, as their pointing errors will be anticorrelated.

delay-related errors in the simulated measurement; our data show that deformation-induced pointing errors have only a marginal effect on η_{int} and continuum contributions.

5.4.2 Measuring the Impact of Pointing Errors

To further explore the potential effect of pointing systematics, we ran an addition series of tests to measure the impact of platform and antenna-based pointing errors on an intensity mapping measurement. Because antennas are fixed to the YTLA platform, we have modeled antenna-based pointing errors to be fixed with respect to the local tangent plane of the platform at the position of the antenna (to account for platform deformation). Similar to the simulations in Section 5.4.1, we assume nominal pointing errors (for both platform and individual antennas) of $0.3'$, measuring the effect of scaling either error separately. We simulate 100 tracks of a single mock field for each choice in pointing error, choosing 200 randomly chosen (logarithmically distributed) errors, applied to the platform or antenna pointing errors.

Shown in Figures 5.6 and 5.7 are the results of simulated platform and antenna pointing errors on our measurement. We find that platform pointing errors of less than an arcminute have generally little effect on either η_{int} or on residual point source power. Individual antenna pointing errors on average also have a small effect, although we note that the effect is much more significant when these errors align with deformation-induced pointing errors.

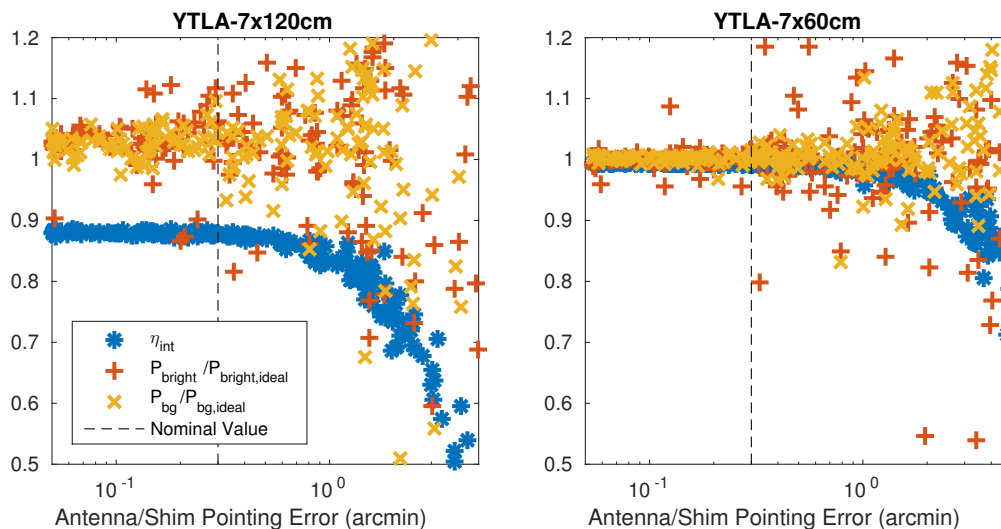


Figure 5.6: The impact of antenna-based pointing errors on a hypothetical measurement with the YTLA, versus the size scale of the pointing error. For a typical interferometer (i.e., not platform based), the effect is similar to a pointing model error, producing correlated position offsets as a function of pointing position. The effect of is most significant when the antenna-based and deformation-induced errors are aligned.

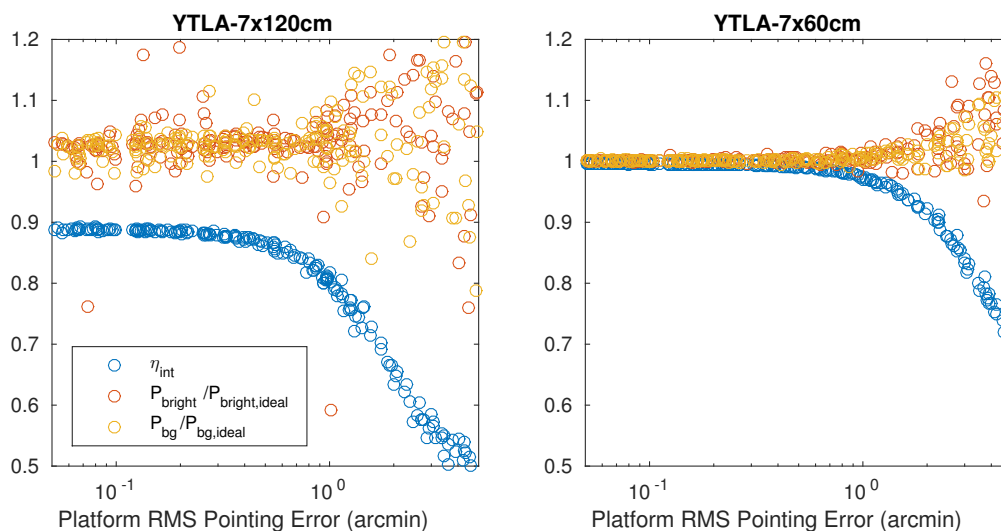


Figure 5.7: The impact of platform pointing errors on a hypothetical measurement with the YTLA, versus the size scale of the pointing error. We find that these errors have only a minimal impact on our measurement.

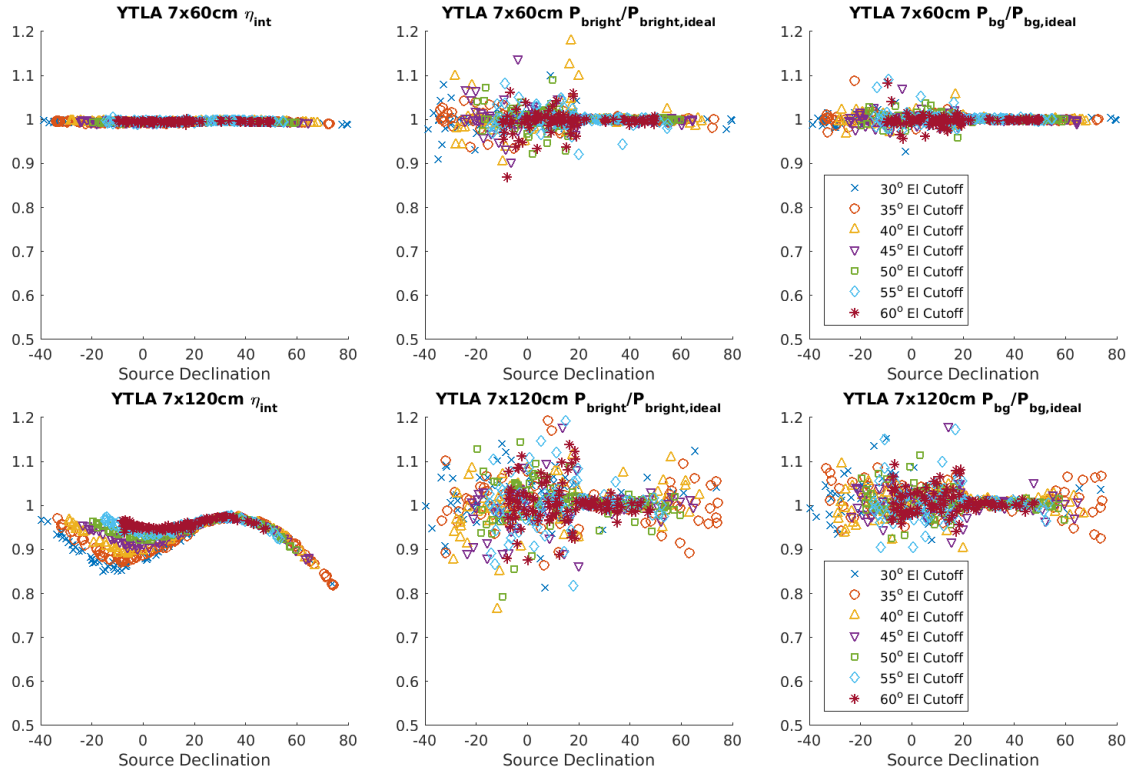


Figure 5.8: Measurements of instrumental effects versus source declination, versus several choices in cutoff elevation (below which the simulated targets are not observed). The top series of panels shows the results for the YTLA7×60cm configuration, while the bottom panel shows the YTLA7×120cm configuration. The leftmost panels show changes in integration efficiency η_{int} , the center panels show relative changes in residual power from bright continuum point sources, and the rightmost panel shows relative changes in residual power from the expected point source population. For the YTLA7×60cm configuration, we find that the source declination and elevation cutoff have little effect on either η_{int} or on the expected residual power from point sources, whereas the YTLA7×120cm appears to suffer from significant systematics below an elevation cutoff of 45 degrees (corresponding to $\delta \leq -25^\circ$ in the southern half of the sky, and $\delta \geq 65^\circ$ in the northern half).

From these additional tests, we conclude that pointing errors of 10% or less of the diameter of the primary beam are generally negligible for intensity mapping experiments. While the simulations discussed here were focused on the YTLA, additional simulation work suggests similar findings for the SZA, where pointing errors of less than an arcminute are insignificant for results with $\lesssim 10\sigma$ significance.

5.4.3 Measured Effects versus Source Declination

Having now studied each of the independent instrumental effects that we expect to dominate systematics within our measurement, we now look at the sum total of these effects versus source declination in the sky. Platform deformation is particularly sensitive to the elevation angle of the instrument at the time of observation, thus we may be required to restrict observations below (or above) some declination limit in order to minimize systematics and maximize the sensitivity of our measurement. As cross-correlation measurements are part of the goal of the AIM-CO experiment, such information is critical for understanding what observationally-rich targets offer the best opportunities for such analysis.

For these simulations, we have assumed nominal instrumental errors (mentioned in Sections 5.4.1 and 5.4.2). We simulate 100 tracks of a single mock field for each choice in source declination, choosing 100 randomly chosen (uniformly distributed) declinations between $-40^\circ \leq \delta \leq 80^\circ$ (chosen due to the fact that the YTLA cannot observe below an elevation angle of 30 degrees), over an hour angle range of ± 5 hours. As systematics are expected to be significantly worse at lower elevations, we also test the effect of excluding data below a given elevation cutoff. We generally find that the YTLA7 \times 60cm does not show a large increase in systematics as a function of declination or elevation cutoff. We do find that there is increased scatter in the residual power from point sources in these simulations for fields below $\delta \approx 20^\circ$. Further testing reveals this scatter to in part be due to the limited number of simulations run per declination, as well as smaller rotation in position angle for southern sources (versus higher declination sources) resulting in particular offset positions from field center being more vulnerable for introducing residual power into our measurement. Barring the presence of a bright point source (> 100 mJy), we expect such contributions to be minimal compared to the sensitivity one would achieve with a year-long survey with the YTLA. With the YTLA7 \times 120cm configuration, we find that observations at lower elevations have a much stronger impact on the fidelity and systematics of our measurement, with η_{int} dropping as low as 0.5 at the very lowest elevations of 30 degrees. Based on these simulations, we find that increasing the low elevation cutoff for observations to 45 degrees will ensure that $\eta_{\text{int}} \gtrsim 0.9$ throughout the track. Due to the platform design and layout, systematics vary more slowly for observations targeting fields in the northern hemisphere, and thus may be better controlled for in post-processing. Therefore, a more lenient cutoff (of 35 or 40 degrees in elevation) may be applied to fields in the northern half of the sky. We therefore conclude that fields between $-25^\circ \leq \delta \leq 70^\circ$ are suitable for observation with the YTLA.

5.4.4 Sensitivity Estimates

We now calculate sensitivity estimates for a realistic year-long observing campaign with the YTLA. As one of the goals for such a survey would be cross-correlation with the COPSS II dataset, we simulate observations of the 12 fields within the FLANK1, GOODS-N, AEGIS, and Q2343 groups. We assume an identical observing cadence and efficiency as was achieved




Figure 5.9: The estimated survey of a year-long intensity mapping experiment with the YTLA targeting CO(3-2) emission at $z = 2.3\text{--}3.0$, versus the estimated sensitivity from the first and second phases of the COPSS experiment. To see how additional antennas might improve the sensitivity of our measurement, we also simulate a YTLA- $13\times 120\text{cm}$ configuration (arranged in a six-pointed star) with half the bandwidth. We find that such a configuration would improve sensitivity by roughly a factor of two over what one could achieve with YTLA- $7\times 120\text{cm}$.

with the SZA (e.g., 20 minute observing loops with 15 minutes of on-source integration time per loop), and employ the same method of ground subtraction as discussed in Section 3.3.1. As platform-deformation effects are strongest at low elevation, we assume data collected while fields below 40 degrees elevation will be flagged or otherwise excluded in our analysis.

The results of these simulations are shown in Figure 5.9. With a year-long observing program with the YTLA- $7\times 120\text{cm}$ configuration (assuming 16 GHz of bandwidth), for the CO(3-2) transition we estimate that one could achieve a 1σ noise threshold of $1.2 \times 10^3 \mu\text{K}^2 h^{-3} \text{Mpc}^3$, similar to that achieved with the analysis of the COPSS dataset.

For the CO(2-1) transition, we estimate that one could achieve a 1σ noise threshold of $0.6 \times 10^3 \mu\text{K}^2 h^{-3} \text{Mpc}^3$. Such a measurement would improve upon our existing sensitivity threshold by a factor of two, though we note that the CO(2-1) transition would probe galaxies over a completely separate redshift range than that surveyed by COPSS.

Chapter 6

Preliminary Results from Future Intensity Mapping Experiments

In this chapter, we discuss preliminary work performed in support of future intensity mapping experiments. Section 6.1 discusses work completed in support of the AIM-CO experiment with the YTLA (also discussed in Chapter 5). Section 6.2 discusses early results from a [CII] intensity mapping experiment, targeting emission at $z \sim 4$.

6.1 The ASIAA Intensity Mapping of CO Experiment

The AIM-CO experiment with the YTLA has substantial potential for exploring CO in the early Universe. The YTLA has frequency coverage between 86 and 102 GHz, making it sensitive to both CO(3-2) ($z = 2.3 - 3.0$) and CO(2-1) ($z = 1.2 - 1.7$) transitions (Ho et al. 2009). Separating the contributions of these two transitions will require a cross-correlation analysis. Fortunately, the YTLA is very well matched in resolution and redshift coverage



Figure 6.1: The Yuan-Tseh Lee Array (YTLA), located on Muana Loa, HI. Muana Kea is visible in the background. *Image Credit:* ASIAA

of the COPSS survey, making it an ideal choice for cross-correlation with COPSS. Such an analysis would allow us to more directly confirm a potential detection, as well as allow us to probe the evolution of molecular gas from $z \sim 1.2 - 3.3$. As the YTLA and SZA are somewhat sensitive to larger-scale emission, cross-correlation of these two datasets may also additionally allow exploration galaxy clustering and large-scale structure.

As part of early science commissioning, observations were conducted in February 2016, using a prototype correlator with 2 GHz bandwidth and 4 120cm antennas (single polarization). Observations targeted a “blank field”, using Jupiter as a bandpass calibrator and QSO 0851202 as a gain calibrator. Observations ran for a total of six hours, and included approximately 2 hours of YTLA data on the blank field. The resultant data was processed through a similar pipeline used for the SZA (discussed in Chapter 3). The power spectrum result of these early data are shown in Figure 6.2. Our measurement is consistent with noise, yielding an upper limit of $P_{\text{CO}(3-2)} \lesssim 10^8 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$. As this threshold is well above that set by the COPSS measurement (and nearly a factor of 10^6 higher than the threshold we wish to achieve in future experiments), there is little that we additionally can constrain about CO emission in early galaxies from this result. However, as a commissioning experiment, there are two important findings from these data. First, our data show no signs of systematics, at levels that may otherwise stymie future intensity mapping efforts with the YTLA. Second, the IMPS package appears to be capable of readily reducing future data with the YTLA.

New science observations with the YTLA are slated to begin in 2016/17 [Bower et al. \(2015\)](#), with a focus on improving upon our current constraints on the CO power spectrum by at least a factor of a few.

6.2 The [CII] Power Spectrum Survey

In the local Universe, CO has often been used as the tracer of choice for studying the bulk of molecular gas in galaxies. However, galaxies in the early Universe may not possess enough dust to shield CO from the dissociating UV starlight within star-forming galaxies, leading to molecular clouds which are “CO-dark” ([Wolfire et al. 2010](#)). Such environments could mean a rapid decrease in CO luminosity at high redshift ($z \gtrsim 4$), and would make constraints on molecular gas abundance difficult with CO alone. However, this lack of dust may also lead to an increase in the abundance of ionized carbon ([CII]). [CII] is an important coolant for the ISM, typically tracing the cool and neutral components of the ISM, with significant emission arising in the photon-dominated regions (PDRs) surrounding molecular clouds ([Pineda et al. 2013](#)). Theoretical models suggest that in the absence of dust, [CII] emission in molecular gas would be significantly enhanced ([Muñoz & Furlanetto 2014](#)), making it a better tool for studying molecular gas in the very early Universe. The combination of CO and [CII] intensity mapping results would allow for more complete view of high redshift molecular gas, less sensitive to the particular environments of individual galaxies ([Stacey et al. 2010](#)). Furthermore, [CII] intensity mapping efforts are extremely complementary to those for CO. The use of these two tracers will allow for more robust estimates for the

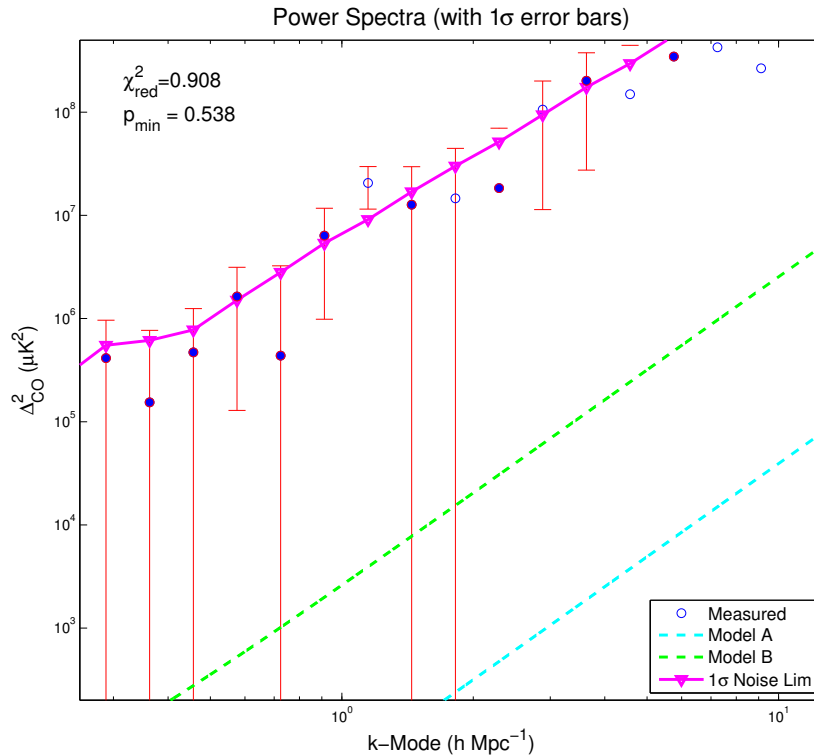


Figure 6.2: The first CO(3-2) power spectrum measurement produced using YTLA data, in the form $\Delta^2(k)$. Filled circles correspond to positive values for $\Delta^2(k)$, while open circles correspond to negative values, and the error bars corresponding to the 1σ errors on our measured values. For comparison, the CO power spectrum estimates from model A (dashed cyan) and model B (dashed green) from Pullen et al. (2013) are shown, along with the estimated RMS noise power (magenta triangle), absent any astrophysical signal. Based on χ^2 statistics of the data, along with evaluation of the PTE for the summed measurement (p_{\min}), the resulting spectra appears to be consistent with noise (as one would expect for a limited integration). This measurement is an important step forward in preparing the instrument for the AIM-CO experiment.

abundance of molecular gas (Bolatto et al. 2013), much in the same way that one might use combined UV and IR luminosities of galaxies to make robust estimates of the star formation rate. The fact that CO and [CII] are differentially affected by metal and dust abundances will additionally allow for constraints on the dust and metal contents of normal star-forming galaxies in the early Universe, that can be compared against current and future results from optical instruments (e.g. JWST, MOSFIRE and HST) to connect the cold gas and star formation properties of galaxies in the early Universe.

The challenges facing [CII] intensity mapping are somewhat different than those afflicting CO intensity mapping efforts. As shown in Figure 6.3, the bright rotational transitions of CO are a formidable foreground for [CII] intensity mapping efforts (e.g. Breyse et al. 2015).




Figure 6.3: The relative power of the several different transitions expected at millimeter wavelengths. At most frequencies within the regime shown here, emission is expected to arise from at least two different spectral lines/ redshift ranges. Such confusion may limit the ability to put precise constraints on molecular gas at high redshift.

For measurements of the [CII] power spectrum, one must either find a way of removing the interloper emission (by way of cross-correlation, masking, etc.), or move to high enough frequencies such that the interloping transitions of CO are relatively weak compared to [CII]. Some theoretical work suggests that at ~ 350 GHz, the [CII] emission may be dominant contributor (in the shot-power regime) to a power spectrum measurement (Cheng et al. 2016).

ALMA is a revolutionary telescope for submillimeter astronomy, providing sensitivity orders of magnitude higher than any similar instrument. The first cycle of ALMA observations – already publicly available – contained several data sets collected at ~ 350 GHz, matching the [CII] fine-structure line at $z \sim 4 - 5$. Matsuda et al. (2015) used these data sets in a

blind search for [CII] emitters. No detections were observed, due in part to the fact that the “wide and shallow” nature of the survey is not conducive to direct imaging efforts. However, this type of survey is optimal for an intensity mapping experiment to detect the aggregate emission of the underlying undetected galaxy population.

Shown in Figure 6.4 are the first results of power spectrum analysis with archival data. These results are consistent with a null detection, though we note that additional jackknife analysis is necessary in order to verify that our results are free of systematics. This additional development in jackknife analysis tools is required, as both of our “interday” tests are unsuitable for the typical archival data sets available from ALMA, which typically include mere minutes per field (for those observations with a sufficient number of pointings to reduce cosmic variance).




Figure 6.4: The power spectrum result from preliminary analysis of ALMA archival data, in the form $\Delta^2(k)$. Filled circles correspond to positive values for $\Delta^2(k)$, while open circles correspond to negative values, and the error bars corresponding to the 1σ errors on our measured values. For comparison, the CO power spectrum estimates from model A (dot-dashed green) and model B (dashed blue) from Pullen et al. (2013) are shown, along with the estimated RMS noise power (gray triangle), absent any astrophysical signal.

Chapter 7

Conclusion

We have used archival and new data with the SZA to place the first-ever constraints on the CO autocorrelation power spectrum at high redshift. With over 6000 hours of observations spread across 63 fields, we have constrained the power spectrum for CO at $z \sim 3$ to $P_{\text{CO}} = 3.0_{-1.3}^{+1.3} \times 10^3 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$, or $\Delta_{\text{CO}}^2(k = 1 h \text{Mpc}^{-1}) = 1.5_{-0.7}^{+0.7} \times 10^3 \mu\text{K}^2$, and find that $P_{\text{CO}} > 0$ to 98.9% confidence. We have used this to constrain the relationship between host halo mass and CO luminosity, and exclude Model B from Pullen et al. (2013), as well as the models from Visbal & Loeb (2010) and Lidz et al. (2011). With these data, we have also demonstrated the efficacy of the intensity mapping method we have been able to survey the molecular gas contents of large volumes (a total of $4.9 \times 10^6 h^{-3} \text{Mpc}^3$ for the COPSS experiment), and for setting constraints on the CO luminosity function and the cosmic molecular gas density of the early Universe. We find that many theoretical models appear to under-predict the number of moderately luminous objects ($\sim 10^{10} \text{K km s}^{-1} \text{pc}^{-2}$). We also find that our estimate for the cosmic molecular gas density agrees with many theoretical models, and offers constraints that improve upon previous work by more than an order of magnitude.

We have also demonstrated our ability to control for systematics and reject foreground contributions within our measurement. From simulations and additional data analysis, we find that we are able to suppress these contributions by at least a factor of 10^3 , below the estimated noise threshold of our experiment. We find that for compact arrays, observing methods that permit ground subtraction are critically important for intensity mapping surveys, as ground contributions can easily exceed those from the signal of interest. We also find that proper bandpass calibration is equally important for intensity mapping experiments, as bandpass errors in the presence of continuum emission can mimic the signal we wish to measure.

Upcoming 3mm observations with the Yuan-Tseh Lee Array (Bower et al. 2015; Ho et al. 2009) will offer increased sensitivity and will be capable of deeply probing the CO power spectrum at $z \sim 3$. As the Lee Array and the SZA share similar spatial frequency and redshift coverage, these new observations will also be an opportunity for cross-correlation between the CO(1-0) and CO(3-2) transitions, offering both improved sensitivity and a more

complete probe into the physical properties of the molecular gas fueling early star formation. These observations will also be sensitive to lower wavenumbers, where contributions from the clustering of galaxies are more likely to dominate the power spectrum. The added constraints on the cluster-power contributions to the power spectrum will be vital in constraining A_{CO} and σ_{CO} , and will offer added insight into the population sub- L_* CO emitters.

Appendix A

Sensitivity Estimates for Intensity Mapping Experiments

A.1 Interferometric Experiments

Utilizing equations 2.10 and 2.11, one may construct a simple estimate for the sensitivity one may achieve for a given choice of instruments and observations. The estimated noise for the power spectrum may be defined as

$$P_N = \frac{2X^2Y}{\Omega_B B_z} \left(\frac{c^2}{2k_B \nu^2} \right)^2 \frac{\tilde{\mathcal{V}}_N^2}{N_V^{1/2} N_k^{1/2} N_{rd}}, \quad (\text{A.1})$$

where $\tilde{\mathcal{V}}_N$ is the noise within a single delay-visibility, N_V is the number of independent volumes measured, N_k is the number of independent measurements made in the k-space volume, and N_{rd} is the number of redundant measurements made at a given position of (u, v, η) . Note that $N_{\text{meas}} = N_V N_k$. $\tilde{\mathcal{V}}_N$ can be further define as

$$\tilde{\mathcal{V}}_N = \frac{2k_B \nu^2}{c^2} \frac{\Omega_b T_{\text{sys}}}{\eta_{\text{eff}}} \sqrt{\frac{B}{\tau_{\text{int}}}}, \quad (\text{A.2})$$

where T_{sys} is the system temperature, η_{eff} is the aperture efficiency, B is the bandwidth over which the measurement is being made and τ is the integration time (Parsons et al. 2012; Thompson et al. 1986).

The number of independent volumes probed is simply the product of the number of fields, N_f , and the number of redshift windows, N_z . The number of redundant measurements, N_{rd} and number of independent measurements in the k-space volume, N_k , are heavily intertwined and influenced by both the array and observing configuration. Consider an array with N_{base} number of baselines. In a minimally redundant array, each of these baselines produce an independent measurement. However, in a redundant array (where elements are placed equally spaced grid) or in a extremely compact array, some baselines will overlap in the uv plane. If we define N_{ob} as the typical number of baselines that instantaneously share

the same position in the uv plane (where $N_{\text{ob}} = 1$ corresponds to a minimally redundant array), then the number of independent measurements will decrease as N_{ob}^{-1} , but the number of redundant measurements will rise as N_{ob} . While compact, the SZA is laid out in such a way that the instantaneous baseline coverage produces only a few, partially overlapping baselines; hence, $N_{\text{ob,SZA}} \approx 1$.

When the array is not coplanar with the sky, the projected baselines will change position over time – an effect commonly referred to as Earth rotation aperture synthesis (ERAS) (Thompson et al. 1986). ERAS reduces the amount of time a given baseline is found at a particular position in the uv plane. Defining f_r as the fraction of the integration time baselines are found at already sampled positions in the uv -plane (either because a baseline remains in a single position for that time, or the baseline moves into a new position in the uv plane previously sampled by another baseline), then the number of redundant measurements scales linearly with f_r , though the number of independent measurements will scale as f_r^{-1} . The calculation of f_r will be heavily affected by the position of the source(s) of interest, as well as the range in hour angle over which observations are conducted, but a simple estimate for redundant fraction is $f_r = 2d_a^2/d_{\text{B,med}}^2$, where d_a is the diameter of the antenna and $d_{\text{B,med}}$ is the median baseline length. The justification for this estimate is that the interferometer will sweep out an area roughly equal to $\pi d_{\text{B,med}}^2$ in the uv plane, whereas the area instantaneously sampled by a single baseline is $2\pi d_a^2$ (the factor of 2 arises from the Hermitian nature of the uv plane). Therefore, the interferometer measures approximately $d_{\text{B,med}}^2/2d_a^2$ independent positions. Excluding baselines with the outrigger antennas, the median SZA baseline is approximately 6.5 meters, leading to $f_{r,\text{SZA}} \approx 0.6$.

Finally, if one has N_{ch} frequency channels across a particular redshift window, then there are an equal number of delay-channels, each of which provide an independent measurement. One caveat to this statement is that if the frequency resolution is too fine, then the flux from an individual emitter may become spread across several channels, leaving some delay-visibility to resolve out the emission and not meaningfully contribute to the measurement (at a velocity resolution of $\sim 300\text{km/s}$, this should not be an issue for the SZA). Combining all of the above pieces of information, the estimates for the number independent and redundant measurements may be written as,

$$N_{\text{rd}} \approx N_{\text{ob}} f_r, \quad (\text{A.3})$$

$$N_k \approx N_{\text{base}} N_{\text{ob}}^{-1} f_r^{-1} N_{\text{ch}}. \quad (\text{A.4})$$

Combining Equations A.3 and A.4 with Equation A.1:

$$P_N \sim 2X^2 Y \Omega_{\text{B}} \frac{(T_{\text{sys}}/\eta_{\text{eff}})^2}{\tau_{\text{int}} (2N_{\text{fields}} N_z N_{\text{base}} N_{\text{ob}} N_{\text{ch}})^{1/2}} \frac{d_{\text{B,med}}}{d_a}. \quad (\text{A.5})$$

Plugging in values from the S10 data set, we obtain an estimated sensitivity of $P_N \sim 1.3 \times 10^4 \mu\text{K}^2 (h^{-1} \text{Mpc})^3$, very close to the actual sensitivity achieved in Section 4.3.

A.2 Single-Dish Measurements

For a single voxel within an image cube, we define estimated the noise within that voxel as

$$T_N = \frac{T_{\text{sys}}/\eta_{\text{eff}}}{\sqrt{\Delta\nu\tau_{\text{int}}}}, \quad (\text{A.6})$$

where $\Delta\nu$ is the frequency channel width, and τ_{int} is the integration time on a single pointing. Given a bandwidth, B , and an effective total survey area, Ω_{survey} with N_{points} independent pointings, then the estimated noise within a single image mode is given by

$$\tilde{T}_N = \frac{T_{\text{sys}}}{\eta_{\text{eff}}} \Omega_{\text{survey}} \sqrt{\frac{B}{N_{\text{points}}\tau_{\text{int}}}}. \quad (\text{A.7})$$

Plugging this expression into Equation 2.18, we can determine the noise within a single Fourier mode in the 3D power spectrum:

$$\tilde{\xi}(\mathbf{k}) = \frac{X^2 Y \Omega_{\text{survey}}}{N_{\text{points}} \tau_{\text{int}}} \left(\frac{T_{\text{sys}}/\eta_{\text{eff}}}{\tilde{A}_\nu(u, v)} \right)^2. \quad (\text{A.8})$$

We note that this is close to the approximation $\tilde{\xi}(\mathbf{k}) = T_N^2 V_{\text{vox}}$ (where V_{vox} is the volume of a single volume), though we note two important consequences of Equation A.8. First, the spatial frequency attenuation will generally filter spatial with perpendicular components greater than $\sim X\lambda_{\text{obs}}/d_a$. Second, as even sensitivity of a survey area will generally require overlapping pointings, it is important to note that the sensitivity is proportional to \mathcal{R}_ω , as highly overlapped pointings effectively allow for coherent integration on a single patch of sky.

Appendix B

The Intensity Mapping Pipeline Software (IMPS) Package

Presented here is a brief description of the algorithms used in processing of the COPSS data sets. The software described herein was written in MATLAB.

B.1 Flagging

There are four primary methods for flagging the data as bad. The first method, which we refer to as “outlier flagging” is for use with calibrator data without the use of gains solutions (which the calibrator data itself must be used to generate). The second method, which we refer to as “coherence flagging” is also for use with calibrator data, without the need for gains solutions. The third method, which we refer to as “PTE flagging” is for use with calibrated data, after gains and passband solutions have been determined and is typically used with source data. The fourth and final method, which we refer to as “expansion flagging” is for use with data that has been previously flagged to expand flags to data that surround groups of visibilities that have been predominately marked as bad.

Flags are set based on 5 different groupings of visibilities: individual visibilities (or baseline-channel pairs), frequency-channel, baseline-window, baseline, and time. When a group is flagged, all constituent visibilities within that group are marked as bad, and are normally excluded from further analysis.

Our first method of flagging, outlier flagging, is typically used in processing of calibrator data prior to the creation of a gains solution. This step is important as contaminated calibrator data will affect the gains solution, which in turn will corrupt source data. Clean calibrator data can be described as being normally scattered about some value, such that it appears as an offset 2D Gaussian in the real-imaginary plane. Because of the lack of gains solutions, one does not know this offset position or the standard deviation of the distribution a priori. The goal of outlier flagging is to determine this information with no additional information in place, favoring the use of medians over means in determining these values (in order to

increase the robustness of the derived values against small groups of outliers). For a given set of visibilities, the median real and imaginary values are calculated and subtracted off, and the absolute value of these modified values are taken. For a distribution with Gaussian-like noise, the absolute value of the difference between the median and each point within that distribution is well described by a Rayleigh distribution, whose median is given as $\sigma\sqrt{\ln(4)}$. The derived value for σ is used as an estimate for the standard deviation. Mathematically we can describe the above operation as

$$\sigma'(\mathcal{V}_m) = \frac{|\mathcal{V}_m - \mu_{0.5}(\mathcal{V})|}{(\ln 4)^{-0.5} \mu_{0.5}(|\mathcal{V} - \mu_{0.5}(\mathcal{V})|)}, \quad (\text{B.1})$$

where $\mu_{0.5}$ is the median function and $\sigma'(\mathcal{V}_m)$ is number of standard deviations away that the visibility \mathcal{V}_m is from the median of the distribution. Typically data that are 3σ away are marked as bad. This style of flagging is most robust against small numbers of bad visibilities, employed on a channel-by-channel basis for calibrator data.

The second method of flagging, coherence flagging, is also used in flagging calibrator data prior to formulation of gains solutions. For a given set of visibilities, we define their coherence ratio to be

$$R_C = \frac{|\sum \mathcal{V}|}{\sum |\mathcal{V}|}, \quad (\text{B.2})$$

where $|\sum \mathcal{V}|$ is the absolute value of the vector sum of all visibilities within a group, and $\sum |\mathcal{V}|$ is the scalar sum. R_C is a metric by which one can evaluate the signal-to-noise within a group of visibilities; the ratio will approach zero for noise-like data and one for noiseless data. While unable to detect bad individual visibilities, this method is excellent for detecting large groups of bad data (that may be missed by outlier flagging). This metrics are evaluated for each individual track of data (3 minutes for gain calibrators, ~ 5 minutes for bandpass calibrators). Bandpass calibrator data is commonly evaluated for each channel within each baseline, while gain calibrator data is evaluated channel-by-channel, while gain calibrator data is evaluated window-by-window (after the application of bandpass solutions). Bandpass data data are flagged when $R_C < 0.95$, and gains data are flagged when $R_C < 0.75$.

Our third method of flagging, PTE flagging, is used for source data after gains have been applied. PTE flagging evaluates the statistics of various sums over different groups of visibilities. In total, 5 statistical tests are used, 4 of which take a variance-weighted vector sum of the data and evaluate it against the expected noise level, the last of which evaluates the chi-squared sum of the data. The first sum is evaluated with the data phased up to the field center, the second sum is evaluated with the data phased to zenith, and the final two sums are evaluated with the data phased to the positions of the Sun and Moon. These sums are calculated after the removal of known sources in the field(s) under evaluation. To evaluate the data, visibilities are broken up into different groups, and the various sums and statistics are calculated over increasingly large time intervals. This process is repeated for larger and larger groups, first starting with individual visibilities, then for all visibilities in a particular frequency channel, then baseline-window groups, then all visibilities in a single baseline, and finally visibilities for a given integration. In some instances, we expect

Flag Group	Percentage of Data Marked as Bad								Total
	Data Statistics			Flags Statistics			Apriori		
	Normal Sum	Zenith Sum	Solar Sum	χ^2	Bad Neighbor	Bad Track	Bad Group	Hardware	
Chan	< 0.1	< 0.1	< 0.1	< 0.1	< 0.1	< 0.1	0	21.3	21.3
Freq	< 0.1	< 0.1	< 0.1	< 0.1	< 0.1	< 0.1	< 0.1	0	< 0.1
Win	< 0.1	< 0.1	< 0.1	< 0.1	0.1	< 0.1	< 0.1	0	0.1
Base	0.6	0.4	0.1	1.4	4.7	3.3	0.1	0	5.5
Ant	0	0	0	0	2.2	0.3	0.8	0	2.2
Time	0.2	0.1	0.3	1.5	4.1	0	0.2	0	5.4
All	0.7	0.4	0.3	2.6	7.9	3.4	0.8	21.3	30.7

Table B.1: Flagging statistics for the COPSS I dataset.

contamination to be broad-band in frequency, but narrow in delay-space. To increase our sensitivity to such contamination, we Fourier transform data across the frequency axis when evaluating data in the baseline-window and baseline groupings (looking for outliers in any of the delay-channels). Thresholds for are set such that in the absence of any contamination, noise-like events will cause < 0.1% of data to be marked as bad.

Our final method of flagging evaluates the flags themselves, finding where large of data have been marked as bad by other methods. There are three measurements considered when flagging data as bad in this method. First is the fraction of visibilities within a group that have been flagged (e.g. if all but a few channels are marked as bad for a given baseline in a single integration, that baseline is marked as bad for that integration). Second is the fraction of time-adjacent integrations a given group is marked as bad (e.g. if a single baseline is flagged for all but a couple of integrations are flagged during a single 4 minute track, those few integrations are marked as bad). Third is fraction of integrations over the entire observation a given group is marked as bad (e.g. if a given baseline is flagged for most of the observation, that baseline is marked as bad for the entire observation).

A tabulation of the fraction of data flagged for the COPSS I experiment is shown in Table B.1, with different rows corresponding to different levels of flagging and different columns corresponding to different reasons for the data being marked as bad.

B.2 Calibration

The output of the interferometer is given here by Equation B.3:

$$B_1^*(\nu)g_1^*B_2(\nu)g_2(\mathcal{V}_{T,12} + \mathcal{N}_{12}) = \mathcal{V}_{M,12} \quad (\text{B.3})$$

where $B(\nu)$ is the frequency-dependent bandpass solution, g the wideband gains solution, N the noise in the measurement, $\mathcal{V}_{T,12}$ the true Fourier transform of the sky as measured by

the interferometer created by antennas 1 and 2, and $\mathcal{V}_{M,12}$ the measured output of the interferometer. To then calibrate our data, we must solve for the bandpass solutions, gains solutions and noise estimates for our data. Methods for all three are presented within this section.

To solve for both the bandpass and gains solutions, we use a non-linear chi-squared solver, using the first-order Taylor expansion of a slightly modified version of Equation B.3. The problem with using this equation outright is that the derivative of the complex conjugate is not defined, which makes generating the first-order Taylor expansion impossible with this form of the equation. To get around this problem, we split the gains coefficients into real and imaginary components, as shown in Equation B.4.

$$(g_{\Re,1} - \imath g_{\Im,1})(g_{\Re,2} + \imath g_{\Im,2})\mathcal{V}_{T,12} = \mathcal{V}_{M,12} \quad (\text{B.4})$$

This differs from the usual way in which gains are solved for (amplitude and phase), though our method offers a significant advantage: the errors in the real-imaginary plane to our fit should be Gaussian. In contrast, the errors in the dimensions of phase-amplitude are not explicitly Gaussian. They do approach a Gaussian in the limit of the SNR becoming very large – however, this case cannot be assumed for all calibrator data (and it is indeed false for our gain calibrators, where the SNR for a single visibility is roughly 1). The bandpass solutions are solved on a channel-by-channel basis within each window for each antenna, while the gains solutions are solved for each window for each antenna. The bandpass solutions are normalized so that the mean-square-inverse amplitude across all antennas for a given window is equal to one, so as to disentangle them from the more wideband gains solutions.

Calibrator observations are interspersed between observations of sources, hence it is necessary to interpolate gains solutions between the calibrator observations these solutions are derived from. For the bandpass solutions, any available solutions are simply averaged together. Nominally, only one bandpass calibrator observation is done per observation, so no averaging is necessary. In the case that two bandpass solutions exist, but one solution is missing gains corrections for a particular channel, window or antenna, then the gains are effectively blanked. For the phase gains solutions, the interpolation is done through a simple linear interpolation. The amplitude gains solutions are stable enough that they can be averaged over the entire track, providing one solution per obsblock. Should one of the two gains tables being interpolated between be missing a gains solution for a particular window or antenna, the gains solution for that window or antenna will be effectively blanked, and the underlying data will not be used. The interpolation between or averaging of these solutions is done between the phases and amplitudes of the solutions, rather than the native real and imaginary components. Additionally, any shift of more than 15% in the gains in either amplitudes or phases for any given antenna will be interpreted as a bad solution, and that gains solution will be blanked during the interpolation process.

The primary source of information on our noise estimates arises from system temperature measurements performed at the time of observation. However, these measurements are performed over entire windows, whereas system temperature variations can be seen across channels within a single window. The source of this variation is not from the sky or from the receiver – which typically dominates the system noise contribution to the signal – but rather

lies in the quantization noise created by digital sampling. These effects are antenna based and are assumed to remain constant over the course of the track. To estimate the SEFD-correction factor, data from the source fields have flags, gains and bandpass solutions applied and known point sources subtracted from the data. A non-linear least squares method is used to solve for the first-order Taylor expansion of the following equation:

$$\mathcal{N}_{M,1-2}(\nu) = (\mathcal{N}_{1,E}(\nu)\mathcal{N}_{2,E}(\nu)C_{1,\nu}C_{2,\nu})^{1/2}. \quad (\text{B.5})$$

In Equation B.5, $\mathcal{N}_{M,1-2}(\nu)$ is the measured noise of the baseline between antennas 1 and 2 (measured at frequency ν), $\mathcal{N}_{1,E}$ is the estimated noise for antenna 1 and $C_{1,\nu}$ is the noise-correction factor which we are attempting to solve for.

B.3 Imaging and Power Spectrum Production

While COPSS is not an imaging experiment, we are reliant upon images for both quality control and as an integral part in measuring point-source positions and fluxes. There are two fundamental algorithms in use for imaging: gridding and deconvolution.

The gridding algorithm takes the visibilities and plugs them into a regular grid of (u, v, ν) . Prior to gridding, gains solutions and flags are applied to the data. There are a variety of weighting options for gridding, though the default option is to variance weight the data based on noise estimates (provided by the system temperatures and noise calibration). For each frequency channel, a pair of 2D grids in (u, v) is constructed – one for visibilities and one for weights. Then, for each visibility, the algorithm finds the closest grid cell to the position of the visibility u and v , adding the weighted visibility to that position in the grid (and the appropriate weights to the weights grid). This continues until all visibilities have been processed, after which a 2D FFT is performed on the grids to create the dirty map and beam/PSF maps.

The deconvolution algorithm uses a simple CLEAN algorithm (Högbom 1974) to create a clean map, with a loop gain of 0.1 and a minimum of 25 cycles before determining whether or not to stop. After those 25 cycles, the RMS and peak flux of the residual map is measured. A threshold is determined by calculating the expected peak noise based on the number of pixels in the image, the measured RMS and assuming a normal distribution of noise. When the measured peak falls below the expected peak, deconvolution is stopped and the map is considered clean.

The construction of the power spectrum begins with with the analysis of a visibilities for a single field on a single day. We defining $\mathcal{V}_m(z_i, \nu_j)$ as the group of visibilities corresponding to a single baseline-integration pairing (specified by the index m), arranged as a function of redshift window, z_i (corresponding to a single correlator window), and frequency, ν_j (corresponding to a single channel), with corresponding noise estimates of $\mathcal{N}_m(z_i, \nu_j)$. Our measurement requires us to Fourier transform our data along the frequency axis axis in order to move our measurement into a space analogous to (k_x, k_y, k_z) . We do so using the

following equations:

$$\tilde{\mathcal{V}}_m(z_i, \eta_k) = 10^{26} \frac{\lambda_i^2}{2k} \mathcal{F}_{\nu \rightarrow \eta}(\mathcal{V}_m(z_i, \nu_j) \cdot \mathcal{B}_m(z_i, \nu_j)) \langle \mathcal{B} \rangle_{z=z_i}^{-1} \quad (\text{B.6})$$

$$\tilde{\mathcal{B}}_m(z_i, \eta_k) = \mathcal{F}_{\nu \rightarrow \eta}(\mathcal{B}(z_i, \nu_j)) \langle \mathcal{B} \rangle_{z=z_i}^{-1} \quad (\text{B.7})$$

$$\tilde{\sigma}^2(z_i) = \left(10^{23} \frac{\lambda_i^2}{2k} \right)^2 \cdot \sum_{\nu} (\Delta\nu \cdot N(z_i, \nu_j))^2 \langle \mathcal{B} \rangle_{z=z_i}^{-2} \quad (\text{B.8})$$

Equation B.6 describes the conversion of the flagged, gains-corrected visibilities to delay-visibilitys $\tilde{\mathcal{V}}_m(z_w, \eta_c)$, where η_c now describes the delay of a given channel of the Fourier-transformed visibilities and where $\mathcal{B}_m(z_i, \nu_j)$ is the windowing function in frequency space applied to the data. In practice, this windowing function incorporates the effect flagging has on the data (with values for this function either ‘0’ or ‘1’). Equation B.7 describes the calculation of the PSF for each group of visibilities within a window. Equation B.8 describes the calculation of the variance estimate for each delay-channel, where $\Delta\nu = 31.125\text{MHz}$ is required for proper Fourier normalization. Since we expect all delays within a given redshift window to have equal sensitivity, Equation B.8 has no dependence on η .

Equations B.6, B.7 and B.8 all feature a factor of $\langle \mathcal{B} \rangle_{z=z_i}$, which accounts for the loss in sensitivity for a given delay-visibility due to channels being flagged in the frequency domain (such that $\tilde{\mathcal{B}}(\eta = 0)$, the center of the PSF for the delay-visibility is always equal to 1). Equations B.6 and B.8 both have a conversion factor of $10^{26} \frac{\lambda_{\text{win}}^2}{2k}$ (where λ_{win} is the wavelength that corresponds to the center of the redshift window, z_{win}) to convert the visibilities from units of flux density (Jy) to units of brightness temperature (K).

Once $\tilde{\mathcal{V}}$, $\tilde{\mathcal{W}}$ and $\tilde{\sigma}^2$ have been calculated, two addition steps take place, one can now grid the data into what we refer to as “k-grids”, which are in coordinates of (u, v, η, z) . K-grids are formulated using the following expression:

$$\mathcal{K}(u, v, \eta, z) = \sum_m \delta(|u - u_m| < R, |v - v_m| < R) \tilde{\mathcal{V}}_m(\eta, z) (\tilde{\sigma}_m^2(\eta, z))^{-1} \quad (\text{B.9})$$

$$\mathcal{W}(u, v, \eta, z) = \sum_m \delta(|u - u_m| < R, |v - v_m| < R) (\tilde{\sigma}_m^2(z))^{-1} \quad (\text{B.10})$$

$$\mathcal{A}(u, v, \eta, z) = \mathcal{K}^*(u, v, \eta, z) \mathcal{K}(u, v, \eta, z) \quad (\text{B.11})$$

$$\mathcal{W}_{\mathcal{A}}(u, v, \eta, z) = \mathcal{W}^*(u, v, \eta, z) \mathcal{W}(u, v, \eta, z) \quad (\text{B.12})$$

Equation B.9 describes the gridding and natural weighting of the data by the estimated variance in delay-visibility (with those weights recorded in \mathcal{W} , as shown in Equation B.10), while Equation B.11 describes the gridding of the autocorrelations of individual delay-visibilitys (and the corresponding weights for those autocorrelations in Equation B.12). The PSF and noise variance estimates are also gridded in this manner, such that for a given position in (u, v, z_{win}) , there are a series of delay-visibilitys for each η , a PSF, and a noise variance estimate. This operation is repeated for all sources, for each day of observing, such that in total ~ 2000 such k-grids are constructed.

K-grids are constructed in units of meters (i.e. physical distance) rather than the canonical wavelengths. This offers a significant advantage in computational simplicity, as the aperture function of the SZA antennas does not change substantially in units of meters, but will change appreciably over the 8GHz of observing bandwidth. Hence, later in our analysis we will be able to use a single convolving kernel instead of several when constructing our power spectrum estimates. However, in gridding the data in this fashion, we introduce the possibility of mode-mixing into our analysis (Di Matteo et al. 2002, Oh & Mack 2003, Bowman et al. 2009). Mode-mixing is a consequence of the changing resolution with the interferometer as a function of frequency, which can be thought of as the position of a baseline in the UV plane over the window which we perform our Fourier transform over to convert the data into delay-space measurements. This mode-mixing can introduce errors into our calculation, as it can cause spectrally-smooth foregrounds to show up in delay bins other than $k_z = 0$ by virtue of the interferometer response changing due to the movement of the baseline through the UV plane, rather than the amplitude of the source changing as a function of frequency. Mode-mixing can be ignored so long as the change in baseline length (in wavelengths), over the frequency interval in question is small compared to the size of the diameter of the antenna, as expressed in Equation B.13

$$\frac{B \Delta\nu}{D \nu} \ll 1 \quad (\text{B.13})$$

where B is the length of the baseline and D the diameter of the dish (both in meters). For the SZA, Equation B.13 maxes out at $\sim 1/20$ for baselines between the compact-group antennas, and $\sim 1/3$ for the baselines of the outrigger antennas. Hence, we expect that mode-mixing will have a negligible impact on our measurements.

The value R in Equations B.9 and B.11 sets the size of the gridding cell for our k-grids. Setting the cell size too large has the consequence of averaging together data that are not necessarily coherent with one another (i.e. they extend over slightly different regions of the UV plane), however setting the size too small will unnecessarily increase the computational burden of processing without an improvement of sensitivity. We set our cell size to 0.03 meters, which will result in a loss of coherence of $\sim 1\%$ while not significantly increasing the computational processing time within our pipeline. This choice in cell size in effect cause data from a 4-minute track into to be averaged together into a single cell (with different tracks generally averaged into different cells).

Once the individual k-grids are constructed, all of the k-grids for a single field are summed together to create a cumulative k-grid for the field, \mathcal{K}_{tot} . Once this final k-grid has been constructed, the power spectrum for the field can be constructed. In calculating an estimate for the value of the brightness temperature variance at each position (u, v, η, z) , there are two things we must make note of. First, as described in Equation 2.13, we will in effect be multiplying delay-visibilitys against the complex-conjugate of the delay-visibilitys, weighted in a way related to the PSF. When multiplying a delay-visibility by the complex conjugate of itself, we run into a problem by virtue the noise within the measurement will correlate with itself, creating a positive bias in the value. One can throw away the auto-correlated

pairs, however each cell in our k-grid contains an appreciable amount of data, so this is a less than optimal option. However, using the sum of the autocorrelations of the delay visibilities that went into each cell, we can use this sum to de-bias our value, effectively leaving behind only the cross-multiplied pairs of visibilities for different days. The second consideration we must remember is that adjacent cells in u and v will be correlated, courtesy of the aperture function of the telescope. Turning to Equation 2.13, we can calculate how correlated adjacent cells are based on how much PSFs for two cells overlap. As in 2.13, we will weight by this overlap ratio quantity squared, in keeping with our natural weighting scheme for the data, and divide out the sum total of the weights, as shown in Equation B.14.

$$\mathcal{P}(u, v, \eta, z) = \left(\sum_{u'} \sum_{v'} \mathcal{K}_{\text{tot}}^*(u, v, \eta, z) \mathcal{K}_{\text{tot}}(u', v', \eta, z) (W^2(u - u', v - v'))^2 - \mathcal{A}_{\text{tot}}(u, v, \eta, z) \right) \times \left(\sum_{u'} \sum_{v'} \mathcal{W}_{\text{tot}}(u, v, \eta, z) \mathcal{W}_{\text{tot}}(u', v', \eta, z) (W^2(u - u', v - v'))^2 - \mathcal{W}_{\mathcal{A}, \text{tot}}(u, v, \eta, z) \right)^{-1} \quad (\text{B.14})$$

$$\mathcal{P}(u, v, \eta, z) = (\mathcal{K}^* \cdot (\mathcal{K} * (W^2 \cdot W^2)) - \mathcal{A}_{\text{tot}}) / (\mathcal{W}^* \cdot (\mathcal{W} * (W^2 \cdot W^2)) - \mathcal{W}_{\mathcal{A}}) \quad (\text{B.15})$$

Equation B.14 is the most explicit implementation of our technique for calculating \mathcal{P} , the brightness temperature variance at each position in (u, v, η, z) , but is computationally expensive and inefficient. Equation B.15 is the formula implemented in our analysis – its mathematically equivalent, but is much more computationally efficient. In Equation B.15, we have assumed that $A^* = A$, and have normalized our weighting by $|A^2|$ so that 1 is the maximum weight given to product of two cells multiplied together. Once $\mathcal{P}(u, v, \eta, z)$ is calculated, we go to the final step of collapsing our 3D distribution of measurements down to a one-dimensional power spectrum based on wavenumber and redshift.

$$P_{\text{CO}}(k, z) = \frac{X^2 Y}{\Omega_{\text{B}} B_z} \sum \sum \sum \mathcal{P}(u, v, \eta, z) \cdot \delta(|X^2(u^2 + v^2) + Y^2 \eta^2 - k^2| < \Delta k) \quad (\text{B.16})$$

Equation B.16 shows this collapse down to a single dimension, with Δk representing the space in between bins. In our analysis, our data are logarithmically binned, with the total width of each bin set at 0.1dex.

We note that Equation B.15 does not yield a measurement with Gaussian-distributed errors, but rather in effect a distribution that is the difference between two χ^2 distributions. The end result will be that for each independent “voxel” in our k-grid, the \mathcal{P}_{CO} estimate will be a χ^2 -like distribution with a mean of zero. Nominally we would desire a method of analysis that would give Gaussian errors. Fortunately, the sheer number of independent measurements made from the first and second phases of COPSS implies that we can make use of the central limit theorem to assume that the errors in our measurement will be Gaussian. To test this assumption, we make use of a number of simulations, which are discussed in 5.1.

B.4 Data structure description

Here is a short description of the structure used in the IMPS package. A structure is generated for each unique source within a single observing block. All of the below items are contained within the visibility structure. In general, data are arranged so that the first dimension is the time axis, the second dimension is the antenna/baseline number axis, the third dimension is the window number axis and the fourth dimension is the channel number axis. Where a given array only consists of a 1D set of values (that are not values with respect to time), then those values are populated across the second dimension of the array (the MATLAB standard).

- **auto (4D double)**: The autocorrelations within the data. These values are not currently used in processing. (Dimensions: number of integrations x number of antennas x number of windows x number of channels)
- **cross (4D complex double)**: Cross-correlations for all antennas/baselines. Zeros and NaNs are interpreted as being bad data, and will be excluded from analysis. (Dimensions: number of integrations x number of baselines x number of windows x number of channels)
- **antData (structure)**: This structure contains information related to the antennas.
 - **apEff (2D double)**: Aperture efficiency for each antenna, used for calculating the estimated noise. (Dimensions: number of antennas x number of windows).
 - **corrEff (1D double)**: Correlator efficiency for each window, used for calculating the estimated noise. (Dimensions: 1 x number of windows)
 - **antXYZ (3D double)**: Position (in meters) of each antenna in the XYZ position system, where the z direction is parallel axis of Earth's rotation, and the xy plane lies along the meridian. Positions are relative to the array center, in units of meters. (Dimensions: number of integrations x number of antennas x 3)
 - **tsys (3D double)**: System temperature for each window of each baseline, in units of Kelvin. (Dimensions: number of integrations x number of antennas x number of windows)
 - **azErr (2D double)**: Azimuth pointing error of the telescopes, in units of radians. These values are not currently used in processing. (Dimensions: number of integrations x number of antennas)
 - **elErr (2D double)**: Elevation pointing error of the telescopes, in units of radians. These values are not currently used in processing. (Dimensions: number of integrations x number of antennas)
 - **antID (1D cell)**: Name of antenna (string). These values are not currently used in processing. (Dimensions: 1 x number of antennas)

- **antType (1D cell)**: The type of antenna used in the array. Currently, there are 4 accepted values (string): ‘OVRO’, ‘BIMA’, ‘SZA’ and ‘YTLA’. (Dimensions: 1 x number of antennas)
- **antScale (1D double)**: Scale factor to multiply the antenna dimensions by – useful for quickly simulating larger/smaller dish diameters. (Dimensions: 1 x number of antennas)
- **antPol (1D signed 8 bit int)**: Polarization of the antenna pol, expanding on the standard FITS polarization code list. -9 = Circular R, -10 = Circular L, -11 = Linear X, -12 = Linear Y. (Dimensions: 1 x number of antennas)
- **antPlatNum (1D double)**: The platform number that each antenna belongs to. This value is generally not used, and can be set to zero for conventional data. (Dimensions: 1 x number of antennas)
- **antPlatOffset (2D double)**: Position (in meters) of the antenna relative to platform center in the XYZ coordinate system, with the x direction lies perpendicular plane of the platform. This value is generally not used, and can be set to zero for conventional data. (Dimensions: 3 x number of antennas)
- **baseCoords (structure)**: This structure contains all of the baseline coordinates for the visibilities within the dataset.
 - **u (4D double)**: U coordinate for each visibility (calculated for each channel in each baseline), in units of wavelengths. (Dimensions: number of integrations x number of baselines x number of windows x number of channels)
 - **v (4D double)**: V coordinate for each visibility (calculated for each channel in each baseline), in units of wavelengths. (Dimensions: number of integrations x number of baselines x number of windows x number of channels)
 - **w (4D double)**: W coordinate for each visibility (calculated for each channel in each baseline), in units of wavelengths. (Dimensions: number of integrations x number of baselines x number of windows x number of channels)
- **gains (structure)**: This structure contains all of the information for the gains solutions.
 - **gainAntSoln (2D complex double)**: The ‘wideband’ gains solutions, used to correct all channels across the band. These values should be set to 1 if these gains were not solved for. (Dimensions: number of gain solns x number of antennas)
 - **gainWinSoln (3D complex double)**: The gains solutions for each window. These values should be set to 1 if these gains were not solved for (Dimensions: number of gain solns x number of antennas x number of windows)
 - **gainChanSoln (4D complex double)**: Effectively the bandpass solutions for the data. These values should be set to 1 if these gains were not solved for.

- (Dimensions: number of gain solns x number of antennas x number of windows x number of channels)
- **gainTime (1D complex double):** The time range over which data was collected for the generation of the gains solution, with the real component being the start MJD time of the solution interval, and the imaginary component being the end MJD time. (Dimensions: number of gain solns x 1)
 - **gainType (1D unsigned 8 bit int):** A mask to record what kind of gains solution was generated. The only important bits to set are bit 6 (set to 1 if the gainChanSoln table in this solution should be used), bit 7 (set to 1 if the gainWinSoln table should be used) and bit 8 (set to 1 if the gainAntSoln table should be used). (Dimensions: number of gain solns x 1)
 - **gainCovar (5D double):** A measurement of the covariance between the fitted gains solutions. Not used in processing, can be set to 0. (Dimensions: number of gain solns x number of antennas x number of antennas x number of windows x number of channels)
 - **gainRefAnt (1D double):** Reference antenna number that was used for generating the solution, only used with website generation. (Dimensions: number of gain solns x 1)
 - **gainSou (1D cell):** Name of calibrator (string) used for generating the gains solution, only used in (Dimensions: number of gain solns x 1)
 - **noiseSpec (4D double):** Noise estimate corrections for each channel. For new data (or data without solns available), these values should be set to 1. (Dimensions: 1 x number of ants x number of windows x number of channels)
 - **apEffCorr (3D double):** Correction factor for the nominal aperture efficiency recorded in the ‘antData’ structure. For new data (or data without solns available), these values should be set to 1. (Dimensions: 1 x number of ants x number of windows)
 - **hostName (string):** Name of the computer that the gains solutions was generated on, any dummy string can be used here. (Dimensions: 1 x length of host name)
 - **timeStamp (1D double):** Date when gains values were generated, arranged in MM-DD-YY HH:MM:SS format. (Dimensions: 1 x 6)
- **flags (structure):** This structure contains all of the flags for the data. For all of the flags tables, different bits correspond to different reasons for a flag being set within the pipeline. For translating data into the SZA pipeline format, bit 8 for all tables is reserved as a ‘special’ bit for data apriori known as bad (though one can accomplish the same thing by substituting bad data with NaN).

- **spur (4D unsigned 8 bit int)**: Channel flagging table, default should be 0. (Dimensions: number of integrations x number of baselines x number of windows x number of channels)
- **rfi (4D unsigned 8 bit int)**: Frequency flagging table (for excluding all channels at a single frequency), default should be 0. (Dimensions: number of integrations x 1 x number of windows x number of channels)
- **win (3D unsigned 8 bit int)**: Window flagging table (for excluding all data in a single window for a given baseline), default should be 0. (Dimensions: number of integrations x number of baselines x number of windows)
- **base (2D unsigned 8 bit int)**: Baseline flagging table, default should be 0. (Dimensions: number of integrations x number of baselines)
- **ant (2D unsigned 8 bit int)**: Antenna flagging table, default should be 0. (Dimensions: number of integrations x number of antennas)
- **time (1D unsigned 8 bit int)**: Time flagging table (for flagging all data within a single integration), default should be 0. (Dimensions: number of integrations x 1)
- **meta (4D unsigned 8 bit int)**: A special flagging table for recording when a bandpass or noise correction solution is missing. Should be set to 0. (Dimensions: 1 x number of antennas x number of windows x number of channels)
- **hostName (string)**: Name of the computer that the gains solutions was generated on, any dummy string can be used here. (Dimensions: 1 x length of host name)
- **timeStamp (1D double)**: Date when gains values were generated, arranged in MM-DD-YY HH:MM:SS format. (Dimensions: 1 x 6)
- **metaData (structure)**: This structure contains various pieces of meta information on the dataset.
 - **nTime (double)**: Number of integrations in the dataset (Dimensions: 1 x 1)
 - **nBase (double)**: Number of baselines in the dataset. (Dimensions: 1 x 1)
 - **nWin (double)**: Number of different frequency windows. (Dimensions: 1 x 1)
 - **nChan (double)**: Number of channels within a window. (Dimensions: 1 x 1)
 - **nAnts (double)**: Number of antennas that make up the baselines in the dataset. (Dimensions: 1 x 1)
 - **nPlat (double)**: Number of platforms that make up the baselines in the dataset. This value is generally not used, and can be set to zero for conventional data. (Dimensions: 1 x 1)
 - **freq (2D double)**: RF frequency for the center of each channel within each window, in units of Hz. (Dimensions: number of windows x number of channels)

- **delFreq (double)**: Width of a single channel in Hz. (Dimensions: 1 x 1)
- **edgeChan (double)**: Number of channels at the edge of each window to discard during processing. (Dimensions: 1 x 1)
- **ant1Arr (1D double)**: A listing of the first antenna for each baseline; e.g. – if the 9th baseline is between antennas 2 and 3, then the 9th element in this array would have the value ‘2’. (Dimensions: 1 x number of baselines)
- **ant2Arr (1D double)**: A listing of the first antenna for each baseline; e.g. – if the 9th baseline is between antennas 2 and 3, then the 9th element in this array would have the value ‘3’. (Dimensions: 1 x number of baselines)
- **baseToAnts (2D double)**: A listing of which antennas belong with each baseline. 0 is used for antennas that aren’t part of the baseline, -1 for the antenna that is complex conjugated prior to correlation, and 1 to the other member antenna of the baseline; e.g. – if the 9th baseline is between antennas 2 and 3, then “baseToAnts(9,:) = [0 -1 1 0 0 0 0]”. (Dimensions: number of baselines x number of antennas)
- **basePol (1D signed 8 bit int)**: Polarization of the baseline, using the standard FITS polarization code list (-1 = Circular RR, -2 = Circular LL, -3 = Circular RL, -4 = Circular LR, -5 = Linear XX, -6 = Linear YY, -7 = Linear XY, -8 = Linear YX, 1 = Stokes I, 2 = Stokes Q, 3 = Stokes U, 4 = Stokes V). (Dimensions: 1 x number of baselines)
- **omegaBeam1cm (1D double)**: Solid angle for the baseline (*not antenna!*), as measured at 1 cm. Equivalent to the inverse of the geometric mean of for the effective collecting area of the two antennas that make up the baseline, measured in units of cm². (Dimensions: 1 x number of baselines).
- **omegaBeamSq1cm (1D double)**: Total sum of the squared aperture function for the baseline (*not antenna!*), as measured at 1 cm. (Dimensions: 1 x number of baselines).
- **platData (structure)**: This structure contains information related to antenna platforms. (n.b., this structure is generally not used for processing – it was added to help support simulations with YTLA).
 - **platXYZ (2D double)**: Position (in meters) of each platform in the XYZ position system, where the z direction is parallel axis of Earth’s rotation, and the xy plane lies along the meridian. Positions are relative to the array center, in units of meters. (Dimensions: number of integrations x number of platforms x 3)
 - **azErr (2D double)**: Azimuth pointing error of the telescopes, in units of radians. These values are not currently used in processing. (Dimensions: number of integrations x number of platforms)

- **elErr (2D double)**: Elevation pointing error of the telescopes, in units of radians. These values are not currently used in processing. (Dimensions: number of integrations x number of platforms)
- **azRot (2D double)**: Azimuthal rotational position of the platform structure, in units of radians. In effect, this is the first Euler angle of the platform rotation. (Dimensions: number of integrations x number of platforms)
- **tipDir (2D double)**: Altitudinal rotational direction of the platform, in units of radians. In effect, this is an alternate specification (e.g., for hexapod mounts) for the first Euler angle of the platform rotation. (Dimensions: number of integrations x number of platforms)
- **tipRot (2D double)**: Altitudinal rotational direction of the platform, in units of radians. In effect, this is the second Euler angles of the platform rotation. (Dimensions: number of integrations x number of platforms)
- **zRot (2D double)**: Azimuthal rotation of the platform with respect to the structure, in units of radians. These values are not currently used in processing. In effect, this is the third Euler angle of the platform rotation. (Dimensions: number of integrations x number of platforms)
- **time (structure)**: This structure contains various information related to time.
 - **mjd (1D double)**: Modified Julian date for the integration. (Dimensions: number of integrations x 1)
 - **inttime (1D single)**: Integration time for each integration, in units of seconds. (Dimensions: number of integrations x 1)
 - **lst (1D double)**: LST at the time of the integration, in decimal hours. (Dimensions: number of integrations x 1)
- **siteData (structure)**: This structure contains information about the observatory site.
 - **tau (3D double)**: Atmospheric opacity at the time of the observation. Not used in processing. (Dimensions: number of measurements x number of antennas x number of windows)
 - **lat (double)**: Latitude of the observatory, in decimal degrees. Not used in processing. (Dimensions: 1 x 1)
 - **long (double)**: Longitude of the observatory, in decimal degrees. Not used in processing. (Dimensions: 1 x 1)
 - **elev (double)**: Elevation of the site in meters. Not used in processing. (Dimensions: 1 x 1)

-
- **obsName (string):** Name of the observing block that this data belongs to. Only used in website generation, only needs a unique identifier for each block of data, e.g. the date at the time of the observation. (Dimensions: 1 x length of string)
 - **source (structure):** This structure contains information on the target of the observations.
 - **name (string):** Name of the source in the observation – since a single file is created for each source, this is just one string. (Dimensions: 1 x length of string)
 - **raApp (1D double):** The RA of the source at the epoch of the integration, in units of decimal hours. (Dimensions: number of integrations x 1)
 - **decApp (1D double):** The Dec of the source at the epoch of the integration, in units of decimal degrees. (Dimensions: number of integrations x 1)
 - **ra (1D double):** The RA of the source at the some standardized epoch (e.g. J2000) listed the source structure. (Dimensions: number of integrations x 1)
 - **dec (1D double):** The Dec of the source at the some standardized epoch (e.g. J2000) listed the source structure. (Dimensions: number of integrations x 1)
 - **epoch (1D double):** The standard epoch for the RA and Dec of the sources, the nominal value being 2000. (Dimensions: number of integrations x 1)
 - **purpose (1D unsigned 32 bit int):** A mask to designate the purpose of the integration. Bit 1 designates that the integration contains source data, bit 2 designates that the integration contains gain calibrator data, bit 4 designates that the integration contains bandpass calibrator data, bit 5 designates that the integration contains flux calibrator data and bit 6 designates that the integration contains data with the noise source on. (Dimensions: number of integrations x 1)

Appendix C

A Study of Calibration Challenges for a 15 GHz Intensity Mapping Interferometer

Here we explore some of the challenges and requirements for calibration of the 19-element prototype of the Discovery Array for Cosmology and Transient Astrophysics (DACoTA). Overall, we find that there are 10 suitable astronomical calibrators for wideband phase and amplitude solutions, and 4 suitable astronomical calibrators for bandpass solutions. The distribution of these calibrators across the sky ensures that there will be at least one of each calibrator type available for any LST. Additionally, we explore the potential exploitation of Ku-band transmitters for the purposes of phase and primary beam calibration.

C.1 DACOTA System Overview

In current design specifications, a single DACOTA element for the prototype will consist of a 1.8 m primary, with a single polarization Ku-band feed (15-18 GHz). A single element is expected to have an overall efficiency of $\eta = 0.6$, and a system temperature of 100 K. The system equivalent flux density (SEFD), S_{sys} , depends on both of these values and is given by Eqn C.1.

$$S_{\text{sys}} = \frac{1}{2} G_{\text{ant}} \eta T_{\text{sys}} \quad (\text{C.1})$$

The factor of one-half in Equation C.1 is required because only one polarization is detected in the feed. G_{ant} is the gain of the antenna, which for a 1.8 meter primary with $\eta = 0.6$ is 1809 Jy/K. Thus, the SEFD is $S_{\text{sys}} = 181$ kJy. For a single baseline, assuming six minutes of integration time with 1 GHz of bandwidth, we expect to see a signal to noise ratio of $R_{NS} = \frac{S_{\text{sou}}}{0.213 \text{ Jy}}$, where S_{sou} is the flux density of the target source. Assuming an array of 19 elements packed into a hex configuration, integrating over six minutes with a bandwidth of 1 GHz will produce an image with a theoretical RMS noise of 23.05 mJy/beam and a

maximum resolution of 7.64 arcminutes. An individual channel of 16 MHz under the same constraints will produce an image with a theoretical RMS noise of 182 mJy/beam.

C.2 Astronomical Calibrators

C.2.1 Phase and Amplitude Calibration

To meet a requirement of $\leq 1\%$ errors for phase and amplitude gain solutions, we are restricted to using calibrators whose fluxes are above 6.3 Jy if we are limited to six minutes of integration and 1 GHz of bandwidth with the prototype. Presented in Table C.1 is a list of viable calibrator choices for the prototype. Overall, we find that there are a total of 10 calibrators suitable for wideband phase and amplitude calibration.

Name	IAU Name	RA (J2000)	Dec (J2000)	2-cm Flux (Jy)	Rise (LST)	Set (LST)
3C84	0319+415	03h19m48s	+41d30m42s	20.7	18.34	12.32
N/A	0609-157	06h09m41s	-15d42m41s	9.0	1.04	11.44
N/A	0927+390	09h27m03s	+39d02m21s	6.60	0.81	18.10
3C273	1229+020	12h29m07s	+02d03m09s	34.0	6.37	18.60
3C279	1256-057	12h56m11s	-05d47m22s	21.8	7.26	18.62
3C345	1642+398	16h42m59s	+39d48m37s	13.0	7.95	1.48
N/A	1733-130	17h33m03s	-13d04m50s	11.0	12.29	22.81
N/A	1924-292	19h24m51s	-29d14m30s	17.0	15.17	23.66
3C446	2225-049	22h25m47s	-04d57m01s	6.60	16.69	4.17
3C454.3	2253+161	22h53m58s	+16d08m54s	15.0	15.98	5.82

Table C.1: List of viable DACOTA Prototype calibrators. Values taken from the 2003 VLA calibrator database. Rise and set times are for Berkeley, California.

From the list provided by the VLA Calibrator Database (Perley & Taylor 2003), we find that there are at least three viable calibrators above the horizon at any given time. There are likely to be additional constraints on pointing (solar avoidance, terrain and pointing limits, etc.), but this analysis suggests that there should be at least one viable calibrator choice at any given time for the DACOTA prototype.

C.2.2 Bandpass Calibration

To ensure continuum residuals do not spoil our measurement, we will need to select calibrators that are bright enough to allow for bandpass errors of $\leq 1\%$. Fortunately, we do not expect the bandpass solutions to vary drastically over the course of the day, allowing for additional integration time for bandpass calibration observations. An integration time

of 20 minutes restricts using calibrators with fluxes of 28 Jy and above. Additionally, we constrain ourselves to calibrators which are unresolved by the DACOTA prototype (under 6 arcminutes in size), as resolved objects would need to be brighter in order to match our SNR requirements. Overall, we find that there are a total of four calibrators suitable for bandpass calibration. Presented in Table C.2 is a list of viable calibrator choices for the prototype.

Name	RA (J2000)	Dec (J2000)	2-cm Flux (Jy)	Rise (LST)	Set (LST)	Size (arcmin)
Cassiopeia A	23h23m24s	+58d48m54s	307	N/A	N/A	4.3×4.3
Cygnus A	19h59m28s	+40d44m02s	102	11.18	4.80	2.0×2.0
Taurus A	05h34m32s	+22d00m52s	499	22.30	12.85	4.2×2.6
Virgo A	12h30m49s	+12d23m28s	30	5.83	19.19	1.0×1.5

Table C.2: List of viable DACOTA Prototype bandpass calibrators. Values taken from [Baars et al. \(1965\)](#). Rise and set times are for Berkeley, California.

With the exception of a one hour block of time, there are at least two viable bandpass calibrators above the horizon at any given time. Cassiopeia A is a circumpolar target, and is small enough that it should not be significantly resolved by the DACOTA prototype. As such, it would be an ideal primary choice for bandpass calibration.

C.2.3 The Sun and Moon as Calibrators

In addition to the normal suite of calibration options, there are two astronomical targets that deserve further consideration as calibration choices - the Sun and the Moon. Both are large, bright sources that fit within the primary beam of the DACOTA prototype. At DACOTA's highest frequency of 40 GHz, both the Sun and Moon extend to twice the size of the primary beam - imaging this far outside of the primary beam will be challenging, but should be feasible.

Imaging the Sun poses two different challenges with the DACOTA prototype. The first problem is that the Sun is an extremely bright source; its average brightness temperature is 20000 K at 3.7 cm, and 10000 at 0.94 cm ([Ho et al. 2008](#)). When pointed at the Sun, the DACOTA prototype receivers will collect -68 dBm of solar radio power - 16 dB above the nominal noise floor of the receivers. This power level is likely to exceed the nominal linear range of the analog components, making imaging and calibration of solar data difficult. Even during the solar quiet phase, the Sun shows daily variations on the order of 50% of its average flux ([Ho et al. 2008](#)), which will be difficult to model accurately enough for the purposes of calibration. Therefore, the Sun is not a suitable calibrator for DACOTA.

In contrast, the Moon may be a nearly ideal calibrator for the DACOTA system. At 15 GHz, the mean lunar brightness temperature of the center of the disk is 239 K ([Linsky 1973](#)) and a total flux of roughly 100 kJy. We expect the DACOTA dish to collect approximately -86 dBm of lunar radio energy over 3 GHz. This power is comparable to the receiver noise

power of both the DACOTA prototype and the final DACOTA design, and it should not push the analog system into a non-linear regime. While not a completely homogeneous disk, observations of the moon at 3.55 cm suggest that one should be able to construct a phase-dependent model of the lunar disk with RMS errors smaller than 1% (Keihm & Gary 1979).

It should also be noted that the moon may serve as an excellent wide-field polarization calibrator as well. Results from observations at 8.3 GHz suggest that accurate phase-dependent models for Stokes Q and U can be constructed for centimeter wavelength observations of the Moon. Furthermore, with the outer edge of the lunar disk showing as much as 10% polarization, the Moon is one of the brightest polarized sources in the sky (Poppi et al. 2002). Proper polarization calibration for DACOTA will be extremely important, since results from WMAP show that the brightness temperature of polarized foregrounds will exceed 100 μ K at Ku and K bands (Gold et al. 2011).

C.3 Unconventional Calibrator Choices

There are several satellite transmitters that operate in the Ku, K and Ka bands. Many of these targets have spectrally narrow¹ transponders, which collectively transmit over a range of a few hundred MHz. Many of the brightest of these satellites are Geosynchronous Earth Orbit (GEO) satellites, and a listing of the primary GEO satellites is given in Section C.5². Of the GEO satellites, the brightest satellite would impart -77 dBW of power on a DACOTA receiver. Though this power is 7 dB above the nominal noise power of the receiver, it should be within the linear range of the analog system (although some additional precautions may be required to prevent saturation of the digital system). Additionally, several other Low Earth Orbit (LEO) satellite constellations exist with transponders in DACOTA's frequency range. Assuming that accurate ephemerides could be produced, LEOs could also serve as excellent potential sources of calibration.

Satellites could be used for both phase and primary beam calibration. Usage as phase calibrators should be relatively straightforward for LEO satellites, although it may be a little more difficult for GEO satellites due to their relatively tight clustering and similar frequency allocations. This limitation may require some mapping of the GEO belt to determine which satellites are suitable for various frequency ranges. Relative amplitude calibration may be possible using satellites, however absolute calibration may be difficult due to the variable nature of both the frequency band usage and drift in pointing/position of satellite antennas. Given the abundance of LEO and GEO satellites at any given moment across the sky, they may serve as effective primary beam calibrators, particularly in monitoring pointing errors.

Unfortunately, there are not many downlinks in the 15-18 GHz range of the DACOTA prototype. Two potential groups of satellites with detectable emissions in this range are the

¹Many transponders broadcast over a few MHz - a much smaller range than can be resolved by the prototype channel resolution of 16 MHz).

²Data for this table was collected from www.lyngsat.com and www.satbeams.com.

TDRS group and the DirecTV group. TDRS has transmitting capabilities at 15.0 GHz, which are used for communication with other satellites. However, it is unknown if these transmitters are strong enough to be detectable by the DACOTA prototype. DirecTV satellites may also have transmitting capacity in the 17.25 to 17.35 GHz range, although further follow-up will be necessary to determine the usage of this band for space-to-ground communications.

Using satellites for calibration presents two significant challenges. First, the van Cittert-Zernike theorem requires that emission from objects being imaged is spatially incoherent (i.e. the signal from satellite A should not correlate with the signal from satellite B). Emissions from satellites broadcasting data are likely to be primarily noise-like, although some analysis will be required to verify spatial incoherence at frequencies that might be considered good for calibration. Second, the use of any satellite as a calibrator will require accurate ephemerides in order to ensure both good astrometry capabilities and good mosaicking capabilities (the latter of which is a definitive requirement of DACOTA’s science goals).

C.4 Conclusion

Successful calibration of the prototype will represent a significant milestone on the path to developing DACOTA. From our analysis, there appears to be no shortage of calibration targets - both conventional and unconventional - well suited for the DACOTA prototype. Furthermore, as DACOTA moves out of the prototype stage and grows in collecting area and sensitivity, the number of suitable calibrators will greatly increase.

Beyond the prototype, additional calibration challenges will need to be addressed during the development of the system. As suggested in section C.2.3, proper polarization calibration will likely be needed in order to integrate down to a sensitivity of 1 microKelvin. Furthermore, as additional sensitivity allows for weaker calibrators to be used, the flux from additional background sources in the field of view of the calibrator will become more significant. Proper modeling of these fields will be required to meet imaging requirements. Finally, good characterization of the primary beam and of cross-talk between antenna elements will be required for high-fidelity imaging and should be considered a goal during the early development phases of the project.

C.5 List of GEO Satellites

Sat Name	EIRP	Downlink Range (GHz)	Location	Nom Az	Nom El
Intelsat 8	30 dBW	12.29–12.73	166 E	258.5	5.7
AMC 23	46 dBW	12.72–12.73	172 E	254.5	10.4
Intelsat 701	49 dBW	10.96–12.65	180 W	248.8	16.6
Ciel 2	55 dBW	12.22–12.68	129 W	190.8	45.5
Galaxy 13	46 dBW	11.73–12.18	127 W	187.6	45.8
AMC 21	47 dBW	12.10–12.18	125 W	184.4	46

Galaxy 18	41 dBW	11.73–12.18	123 W	181.1	46.1
Galaxy 23	46 dBW	11.72–12.18	121 W	177.9	46.1
Echostar 7	54 dBW	12.22–12.52	119 W	174.6	46
Echostar 14	57 dBW	12.22–12.46	119 W	174.6	46
Anik F3	49 dBW	11.72–12.16, Ka	119 W	174.6	46
DirecTV 7S	51 dBW	Ku-Band	119 W	174.6	46
Satmex 5	51 dBW	11.74–12.19	117 W	171.4	45.7
Solidaridad 2	49 dBW	Ku-Band	115 W	168.2	45.4
Satmex 6	50 dBW	11.99–12.19	113 W	165.1	45
Anik F2	44 dBW	11.75–12.19, 19.78–20, Ka	111 W	162	44.5
Echostar 10	57 dBW	12.27–12.66	110 W	160.4	44.2
Echostar 11	52 dBW	12.27–12.66	110 W	160.4	44.2
DirecTV 5	41 dBW	Ku-Band	110 W	160.4	44.2
Anik F1R	46 dBW	11.72–12.19	107 W	156	43.2
AMC 15	46 dBW	11.71–11.95	105 W	153.1	42.5
DirecTV 12	57 dBW	Ka-Band	103 W	150.3	41.7
Spaceway 1	54 dBW	Ka-Band	103 W	150.3	41.7
DirecTV 10	50 dBW	Ka-Band	103 W	150.3	41.7
AMC 1	44 dBW	11.76–12.18	103 W	150.3	41.7
DirecTV 8	42 dBW	Ku-Band	101 W	147.6	40.8
DirecTV 9S	54 dBW	Ku-Band	101 W	147.6	40.8
AMC 4R	48 dBW	12.02–12.12	101 W	147.6	40.8
Spaceway 2	54 dBW	Ka-Band	99 W	144.9	39.8
DirecTV 11	50 dBW	Ka-Band	99 W	144.9	39.8
Galaxy 16	49 dBW	11.77–12.19	99 W	144.9	39.8
Galaxy 19	48 dBW	11.71–12.18	97 W	142.2	38.7
Galaxy 3C	45 dBW	11.78–12.11	95 W	139.9	37.7
Spaceway 3	41 dBW	Ka-Band	95 W	139.9	37.7
Galaxy 25	39 dBW	11.78–12.00	93 W	137.6	36.5
Galaxy 17	44 dBW	11.74–12.14	91 W	135.3	35.3
Nimiq 1	42 dBW	12.22–12.68	91 W	135.3	35.3
Nimiq 2	38 dBW	12.20–12.70	91 W	135.3	35.3
Galaxy 28	49 dBW	11.71–12.20	89 W	133.1	34.1
AMC 3	42 dBW	11.72–12.06	87 W	130.9	32.8
AMC 16	40 dBW	11.97–12.16	85 W	128.8	31.5
AMC 9	40 dBW	11.74–12.18	83 W	126.9	30.1
AMC 2	47 dBW	11.86	79 W	123.1	27.3
AMC 5	46 dBW	11.78–12.18	79 W	123.1	27.3
Echostar 1	45 dBW	12.22–12.66	77 W	121.3	25.9
Echostar 4	45 dBW	12.25–12.66	77 W	121.3	25.9
Echostar 6	49 dBW	12.22	77 W	121.3	25.9
Echostar 8	51 dBW	12.22–12.68	77 W	121.3	25.9

Horizons 2	49 dBW	13.24	74 W	118.7	23.7
Nimiq 5	50 dBW	12.23–12.27	73 W	117.8	23.0
AMC 6	50 dBW	11.70–12.19	72 W	117.0	22.3
AMC 4	46 dBW	11.66–12.12	67 W	113.0	18.5
Telstar 14	46 dBW	11.96	63 W	110.1	15.4
Echostar 3	47 dBW	12.24–12.68	62 W	109.0	14.3
Echostar 12	58 dBW	12.22–12.56	62 W	109.0	14.3
Echostar 15	56 dBW	12.24–12.68	62 W	109.0	14.3
Amazonas 1	47 dBW	12.05–12.06	61 W	108.6	13.9
Amazonas 2	44 dBW	11.73–11.81	61 W	108.6	13.9
Intelsat 9	47 dBW	11.50–11.65	58 W	106.5	11.6
Intelsat 16	42 dBW	11.72–11.86	58 W	106.6	11.6
Intelsat 707	46 dBW	11.59–11.63	53 W	103.1	7.6
TDRS 5	35 dBW	13.40–14.05, 15.00	174 W		
TDRS 6	35 dBW	13.40–14.05, 15.00	174 W		
TDRS 8	35 dBW	13.40–14.05, 15, 25.25–27.50	171 W		

Table C.3: List of GEO satellites visible from Berkeley, CA.

Bibliography

- Aladro, R., Martín, S., Martín-Pintado, J., et al. 2011, [A&A](#), 535, A84
- Baars, J. W. N., Mezger, P. G., & Wendker, H. 1965, [ApJ](#), 142, 122
- Bauermeister, A., Blitz, L., & Ma, C.-P. 2010, [ApJ](#), 717, 323
- Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, [ApJS](#), 208, 20
- Bolatto, A. D., Wolfire, M., & Leroy, A. K. 2013, [ARA&A](#), 51, 207
- Bond, J. R., Jaffe, A. H., & Knox, L. 1998, [Phys. Rev. D](#), 57, 2117
- Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2012, [ApJ](#), 752, L5
- Bower, G., Keating, G., Marrone, D., et al. 2015, IAU General Assembly, 22, 2250511
- Bowman, J. D., Morales, M. F., & Hewitt, J. N. 2009, [ApJ](#), 695, 183
- Brammer, G. B., van Dokkum, P. G., Franx, M., et al. 2012, [ApJS](#), 200, 13
- Breysse, P. C., Kovetz, E. D., & Kamionkowski, M. 2014, [MNRAS](#), 443, 3506
- . 2015, [MNRAS](#), 452, 3408
- Bridle, A. H., & Schwab, F. R. 1999, in *Astronomical Society of the Pacific Conference Series*, Vol. 180, *Synthesis Imaging in Radio Astronomy II*, ed. G. B. Taylor, C. L. Carilli, & R. A. Perley, 371
- Carilli, C. L. 2011, [ApJ](#), 730, L30
- Cheng, Y.-T., Chang, T.-C., Bock, J., Bradford, C. M., & Cooray, A. 2016, ArXiv e-prints, [arXiv:1604.07833](#)
- Dame, T. M., Hartmann, D., & Thaddeus, P. 2001, [ApJ](#), 547, 792
- Danielson, A. L. R., Swinbank, A. M., Smail, I., et al. 2013, [MNRAS](#), 436, 2793
- Datta, A., Bowman, J. D., & Carilli, C. L. 2010, [ApJ](#), 724, 526
- Davies, M. L., Stefan, I. I., Bolton, R. C., et al. 2013, [MNRAS](#), 430, 1961
- Davis, M., Guhathakurta, P., Konidakis, N. P., et al. 2007, [ApJ](#), 660, L1
- Davis, T. A., Heiderman, A., Evans, N. J., & Iono, D. 2013, [MNRAS](#), 436, 570
- Decarli, R., Walter, F., Carilli, C., et al. 2014, [ApJ](#), 782, 78
- Di Matteo, T., Perna, R., Abel, T., & Rees, M. J. 2002, [ApJ](#), 564, 576
- Dickinson, M., Giavalisco, M., & GOODS Team. 2003, in *The Mass of Galaxies at Low and High Redshift*, ed. R. Bender & A. Renzini, 324
- Dodson, S. 2003, *Modern cosmology* (Amsterdam (Netherlands): Academic Press)
- Downes, D., Solomon, P. M., & Radford, S. J. E. 1993, [ApJ](#), 414, L13
- Flagey, N., Boulanger, F., Noriega-Crespo, A., et al. 2011, [A&A](#), 531, A51
- Frerking, M. A., Langer, W. D., & Wilson, R. W. 1982, [ApJ](#), 262, 590

- Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Phys. Rep.*, 433, 181
- Furusawa, H., Kosugi, G., Akiyama, M., et al. 2008, *ApJS*, 176, 1
- Genzel, R., Tacconi, L. J., Combes, F., et al. 2012, *ApJ*, 746, 69
- Glover, S. C. O., & Mac Low, M.-M. 2011, *MNRAS*, 412, 337
- Gold, B., Odegard, N., Weiland, J. L., et al. 2011, *ApJS*, 192, 15
- Hayward, C. C., Kereš, D., Jonsson, P., et al. 2011, *ApJ*, 743, 159
- Ho, C., Slobin, S., Kantak, A., & Asmar, S. 2008, The Interplanetary Network Progress Report, 42
- Ho, P. T. P., Altamirano, P., Chang, C.-H., et al. 2009, *ApJ*, 694, 1610
- Hobson, M. P., & Maisinger, K. 2002, *MNRAS*, 334, 569
- Högbom, J. A. 1974, *A&AS*, 15, 417
- Hopkins, A. M., & Beacom, J. F. 2006, *ApJ*, 651, 142
- Israel, F. P. 1997, *A&A*, 328, 471
- Keating, G. K., Bower, G. C., Marrone, D. P., et al. 2016, *ApJ*, submitted
- . 2015, *ApJ*, 814, 140
- Keihm, S., & Gary, B. 1979, in Lunar and Planetary Science Conference Proceedings, Vol. 10, Lunar and Planetary Science Conference Proceedings, 2311
- Kennicutt, Jr., R. C. 1998, *ApJ*, 498, 541
- Keres, D., Yun, M. S., & Young, J. S. 2003, *ApJ*, 582, 659
- Kereš, D., Katz, N., Fardal, M., Davé, R., & Weinberg, D. H. 2009, *MNRAS*, 395, 160
- Kriek, M., Shapley, A. E., Reddy, N. A., et al. 2015, *ApJS*, 218, 15
- Lagos, C. D. P., Baugh, C. M., Lacey, C. G., et al. 2011, *MNRAS*, 418, 1649
- Lee, K.-S., Giavalisco, M., Conroy, C., et al. 2009, *ApJ*, 695, 368
- Lentati, L., Wagg, J., Carilli, C. L., et al. 2015, *ApJ*, 800, 67
- Li, T. Y., Wechsler, R. H., Devaraj, K., & Church, S. E. 2016, *ApJ*, 817, 169
- Liao, Y.-W., Lin, K.-Y., Huang, Y.-D., et al. 2013, *ApJ*, 769, 71
- Lidz, A., Furlanetto, S. R., Oh, S. P., et al. 2011, *ApJ*, 741, 70
- Linsky, J. L. 1973, *ApJS*, 25, 163
- Madau, P., & Dickinson, M. 2014, *ARA&A*, 52, 415
- Mashian, N., Sternberg, A., & Loeb, A. 2015, *J. Cosmology Astropart. Phys.*, 11, 028
- Matsuda, Y., Nagao, T., Iono, D., et al. 2015, *MNRAS*, 451, 1141
- Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, 752, 137
- Muñoz, J. A., & Furlanetto, S. R. 2014, *MNRAS*, 438, 2483
- Muchovej, S., Mroczkowski, T., Carlstrom, J. E., et al. 2007, *ApJ*, 663, 708
- Muchovej, S., Leitch, E., Carlstrom, J. E., et al. 2010, *ApJ*, 716, 521
- Noeske, K. G., Weiner, B. J., Faber, S. M., et al. 2007, *ApJ*, 660, L43
- Obreschkow, D., Croton, D., De Lucia, G., Khochfar, S., & Rawlings, S. 2009a, *ApJ*, 698, 1467
- Obreschkow, D., Heywood, I., Klöckner, H.-R., & Rawlings, S. 2009b, *ApJ*, 702, 1321
- Obreschkow, D., & Rawlings, S. 2009, *ApJ*, 696, L129
- Oh, S. P., & Mack, K. J. 2003, *MNRAS*, 346, 871
- Parsons, A., Pober, J., McQuinn, M., Jacobs, D., & Aguirre, J. 2012, *ApJ*, 753, 81

- Perley, R. A., & Taylor, G. 2003, The VLA Calibrator Manual, Tech. rep., National Radio Astronomy Observatory
- Pineda, J. L., Langer, W. D., Velusamy, T., & Goldsmith, P. F. 2013, *A&A*, **554**, A103
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, ArXiv e-prints, [arXiv:1303.5072 \[astro-ph.CO\]](https://arxiv.org/abs/1303.5072)
- Planck Collaboration, Aghanim, N., Arnaud, M., et al. 2015, ArXiv e-prints, [arXiv:1507.02704](https://arxiv.org/abs/1507.02704)
- Poppi, S., Carretti, E., Cortiglioni, S., Krotikov, V., & Vinyajkin, E. N. 2002, in American Institute of Physics Conference Series, Vol. 609, Astrophysical Polarized Backgrounds, 187
- Popping, G., Behroozi, P. S., & Peebles, M. S. 2015, *MNRAS*, **449**, 477
- Popping, G., van Kampen, E., Decarli, R., et al. 2016, ArXiv e-prints, [arXiv:1602.02761](https://arxiv.org/abs/1602.02761)
- Pullen, A. R., Chang, T.-C., Doré, O., & Lidz, A. 2013, *ApJ*, **768**, 15
- Reddy, N. A., Steidel, C. C., Erb, D. K., Shapley, A. E., & Pettini, M. 2006, *ApJ*, **653**, 1004
- Regan, M. W., Thornley, M. D., Helfer, T. T., et al. 2001, *ApJ*, **561**, 218
- Riechers, D. A., Bradford, C. M., Clements, D. L., et al. 2013, *Nature*, **496**, 329
- Righi, M., Hernández-Monteagudo, C., & Sunyaev, R. A. 2008, *A&A*, **489**, 489
- Rudy, D. J., Muhleman, D. O., Berge, G. L., Jakosky, B. M., & Christensen, P. R. 1987, *Icarus*, **71**, 159
- Sargent, M. T., Daddi, E., Béthermin, M., et al. 2014, *ApJ*, **793**, 19
- Schechter, P. 1976, *ApJ*, **203**, 297
- Schmidt, M. 1959, *ApJ*, **129**, 243
- Sharp, M. K., Marrone, D. P., Carlstrom, J. E., et al. 2010, *ApJ*, **713**, 82
- Smit, R., Bouwens, R. J., Franx, M., et al. 2012, *ApJ*, **756**, 14
- Solomon, P. M., Downes, D., & Radford, S. J. E. 1992, *ApJ*, **398**, L29
- Solomon, P. M., Downes, D., Radford, S. J. E., & Barrett, J. W. 1997, *ApJ*, **478**, 144
- Stacey, G. J., Hailey-Dunsheath, S., Ferkinhoff, C., et al. 2010, *ApJ*, **724**, 957
- Steidel, C. C., Shapley, A. E., Pettini, M., et al. 2004, *ApJ*, **604**, 534
- Steidel, C. C., Rudie, G. C., Strom, A. L., et al. 2014, *ApJ*, **795**, 165
- Tacconi, L. J., Genzel, R., Neri, R., et al. 2010, *Nature*, **463**, 781
- Tacconi, L. J., Neri, R., Genzel, R., et al. 2013, *ApJ*, **768**, 74
- Tegmark, M., Hamilton, A. J. S., Strauss, M. A., Vogeley, M. S., & Szalay, A. S. 1998, *ApJ*, **499**, 555
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 1986, Interferometry and synthesis in radio astronomy (New York, Wiley-Interscience)
- Tinker, J., Kravtsov, A. V., Klypin, A., et al. 2008, *ApJ*, **688**, 709
- Valiante, R., Schneider, R., Bianchi, S., & Andersen, A. C. 2009, *MNRAS*, **397**, 1661
- Visbal, E., & Loeb, A. 2010, *J. Cosmology Astropart. Phys.*, **11**, 16
- Visbal, E., Trac, H., & Loeb, A. 2011, *J. Cosmology Astropart. Phys.*, **8**, 10
- Vonlanthen, P., Rauscher, T., Winteler, C., et al. 2009, *A&A*, **503**, 47
- Walter, F., Bertoldi, F., Carilli, C., et al. 2003, *Nature*, **424**, 406
- Walter, F., Decarli, R., Sargent, M., et al. 2014, *ApJ*, **782**, 79
- Wang, R., Carilli, C. L., Neri, R., et al. 2010, *ApJ*, **714**, 699

- Wang, R., Wagg, J., Carilli, C. L., et al. 2011, [AJ](#), **142**, 101
- Weiß, A., Walter, F., & Scoville, N. Z. 2005, [A&A](#), **438**, 533
- Wheeler, C., Phillips, J. I., Cooper, M. C., Boylan-Kolchin, M., & Bullock, J. S. 2014, [MNRAS](#), **442**, 1396
- White, M., Carlstrom, J. E., Dragovan, M., & Holzappel, W. L. 1999, [ApJ](#), **514**, 12
- Wilson, R. W., Jefferts, K. B., & Penzias, A. A. 1970, [ApJ](#), **161**, L43
- Wolfire, M. G., Hollenbach, D., & McKee, C. F. 2010, [ApJ](#), **716**, 1191
- Young, J. S., & Scoville, N. 1982, [ApJ](#), **258**, 467
- Young, J. S., Xie, S., Tacconi, L., et al. 1995, [ApJS](#), **98**, 219