


A Limit on the Polarization of the Cosmic Microwave Background Radiation

Thesis by
John K. Cartwright

In Partial Fulfillment of the Requirements
for the Degree of
Doctor of Philosophy



California Institute of Technology
Pasadena, California

2002

(Defended November 12, 2002)

© 2002

John K. Cartwright

All rights Reserved

Acknowledgements

My goal for my graduate career was to contribute to all of the components of a CMBR effort—from the design of the instrument to the points on the power spectrum—and I am grateful to Tony Readhead, my advisor, for making this possible with the Cosmic Background Imager. Early in the project Tony entrusted me with several major instrumentation projects; in Chile he coordinated our observing program to yield months of polarization data for this thesis; and back in Pasadena he encouraged my effort to distill this mountain of numbers to a set of limits on the polarization of the CMBR. Perhaps most importantly, Tony endeavored to ensure that my contributions to the CBI—and my accomplishments with the CBI—would have high visibility. Without his guidance, support, and encouragement this work would not have been possible.

Tony assembled an outstanding team for the CBI. I am indebted to Steve Padin for his guidance in the design of the downconverter and noise cal (Chapter 2); Steve endured my numerous questions with great patience and equanimity. Martin Shepherd wrote CBIPOLCAL to calibrate the CBI's polarization data (Chapter 3). Martin accommodated my many requests to modify CBIPOLCAL; at each iteration he delivered the intuitive, airtight code that has garnered him so much esteem in the radio astronomy community. Martin also modified his flagship program, DIFMAP, at my request to accommodate polarization observations, and radio astronomers everywhere will sing hosannas at its public release. Tim Pearson wrote CBICAL, the primary total intensity calibration package (Chapter 3), as well POLSKY and POLFAKE, two simulation programs that were critical for testing my maximum likelihood routine (Chapter 6). Tim also authored a great body of internal literature that provided the framework for the acquisition, calibration, and analysis of the CBI data. Tony, Tim and Steve read great swaths of this sprawling leviathan, and their comments have improved the accuracy and clarity of the text.

Other members of the CBI team—Steve Myers, Brian Mason, Pat Udomprasert and Jon Sievers—contributed to this work, both in Pasadena and at the CBI site in

San Pedro de Atacama, Chile. Steve gave me a crash course in maximum likelihood as well as the nuanced arts of driving a 4×4 through the lava fields of Chajnantor. Brian served as a ready fount of radio astronomy wisdom, and his odd yet beguiling sense of humor has left its mark—for better or worse—on the entire CBI team. Pat provided many useful insights about the acquisition and analysis of the CBI data, and her good cheer ameliorated the many little nuisances of life in the Chilean hinterlands. Jon tutored me in maximum likelihood; I retained about 10% of what he told me, but it was the 10% that I needed to obtain my results. San Pedro does not offer many diversions, so we made our own in the kitchen: we dined like royalty on Tony’s beef bourguignon, Tim’s coq au vin, Steve’s Christmas pudding, and Pat’s beef panang—I only regret that we never served all of these dishes in a single, rapture-inducing feast! All members of the CBI team contributed to the observations for the 100+ nights of polarization data presented in Chapter 5.

This thesis benefited from the participation of some of the many outstanding folks at the National Radio Astronomy Observatory. I spent several months at NRAO’s Central Development Laboratory in Charlottesville, VA, to learn how to build and tune the CBI’s low-noise HEMT amplifiers (Chapter 2), and the CDL staff with whom I worked—Marian Pospieszalski, Ed Wollack, and Bill Lakatosh—were terrific. The VLA observations presented in Chapter 4 would not have been possible without fantastic assistance from Greg Taylor at NRAO in Socorro, NM. Thanks go to Barry Clark for providing six 6^h blocks of VLA time for the gain curve measurements; the first five were corrupted by high winds, but the sixth did the trick. Rick Perley, Steve Myers, Chris Carilli, Bryan Butler, Debra Shepherd and Ken Sowiński were always happy to dispense wisdom about using the VLA. Rick also provided flux models for 3C286, as well as the high dynamic range measurements of 3C273 (Chapter 5).

The social life of Robinson Hall is an outstanding feature of Caltech astronomy. I enjoyed the hospitality of Brad Behr, Chris Fassnacht, Roy Gal, Christina Hood, Brian Kern, Marc Kuchner, Ben Oppenheimer, Bob Rutledge, Dave Vakil, and John Yamasaki. Thanks to John Yamasaki in particular for letting me share his

lab for a year while I built the downconverter and the noise cal. I am particularly fortunate to have made close friendships with Pat Udomprasert, Jon Sievers, Rob Simcoe, and Alice Shapley. I shared Robinson 14 with Pat and Jon after the three of us returned from Chile, and I am grateful to them for making room 14 a loud, rambunctious, and thoroughly great place to work. I enjoyed fantastic meals from Pat and Rob's kitchen, and happy afternoons with Jon, Rob, and Pat on the links and at the batting cages. And Alice: for five wonderful years you filled my life with the sunshine of your wit, your intelligence, your charm, and your grace. I am grateful for your companionship during this short, sweet chapter; you enriched my time at Caltech beyond measure.

My family—mom and dad, Brooks and Rebecca, Kristy and Andy—has been extraordinarily supportive during this effort; your phone calls, care packages, and visits were welcome distractions and thus important contributions to this thesis. Grandma Cartwright, Auntie Jeanne, Aunt Kathy, Uncle Larry, Aunt Judy, and Uncle George all made southern California a great second home; Grandma, Judy, and George in particular labored to keep me fed and rested when I did not have the good sense to do so on my own. Grandma Piccinotti was not able to see the conclusion of this thesis; she and Grandma Cartwright have always had a keen interest in my work, so I dedicate this thesis to them.

A Limit on the Polarization of the Cosmic Microwave Background Radiation

by

John K. Cartwright

In Partial Fulfillment of the
Requirements for the Degree of
Doctor of Philosophy

Abstract

We describe polarization observations of the CMBR with the Cosmic Background Imager, a 13-element interferometer which operates in the 26-36 GHz band and is located on the Llano de Chajnantor in northern Chile. The array consists of 90 cm Cassegrain antennas mounted on a steerable platform which can be rotated about the optical axis to facilitate polarization observations. The CBI employs single-mode circularly polarized receivers and it samples multipoles from $\ell \sim 400$ to $\ell \sim 4250$. The instrumental polarization of the CBI was calibrated with 3C279, a bright polarized point source which was monitored with the VLA. Observations of two deep fields during the 2000 season yielded three limits (95% c.l.) for C_ℓ^{EE} under the assumption that $C_\ell^{BB} = 0$: $7.0 \mu\text{K}$ ($\ell = 603$); $12.8 \mu\text{K}$ ($\ell = 1144$); and $25.1 \mu\text{K}$ ($\ell = 2048$). The low- ℓ limit approaches the levels of fluctuations predicted by standard models.

This thesis also entailed the design and implementation of several major components of the CBI signal chain including the downconverter, the noise calibration system, and the low noise HEMT amplifiers. We discuss the design and performance of these critical systems.

Contents

1	Introduction	1
1.1	Overview	1
1.2	Polarization of the CMBR	4
1.3	Polarization Experiments	10
1.4	Thesis Outline	11
2	The Cosmic Background Imager	14
2.1	Introduction	14
2.2	Response of an Interferometer	17
2.3	The Cosmic Background Imager	24
2.3.1	Signal Path	25
2.3.2	Polarization Considerations for the Signal Path	27
2.3.3	Performance of the CBI	35
2.4	Instrumentation Projects Specific to this Thesis	40
2.4.1	Downconverter	41
2.4.2	Noise Calibration System	53
2.4.3	HEMT Amplifiers	63
2.5	Conclusions	67
3	Polarization Calibration	70
3.1	Introduction	70
3.2	Calibration Procedures	71
3.2.1	CBICAL: Quadrature Calibration	72

3.2.2	CBICAL: Internal Noise Source Calibration	74
3.2.3	UVSUB: Spillover Rejection	76
3.2.4	CBIPOLCAL: Leakage and Flux Calibration	77
3.3	Polarization Calibration Observations	86
3.3.1	3C279 Observations	86
3.3.2	Tau A Observations	100
3.3.3	Tau A Model 1	105
3.3.4	Tau A Model 2	113
3.3.5	Previous Tau A Polarization Observations	125
3.4	Calibration Error Budget	130
4	VLA Observations	133
4.1	Introduction	133
4.2	Observations	133
4.2.1	Observing Strategy	136
4.2.2	Antenna Selection	138
4.2.3	Referenced Pointing	139
4.2.4	Absolute Flux Calibration	140
4.2.5	Instrumental Polarization Calibration	141
4.2.6	Position Angle Calibration	142
4.2.7	Atmospheric Opacity	143
4.2.8	Gain Curves	143
4.3	Data Reduction	146
4.3.1	Interpretation	147
4.3.2	Wind	148
4.3.3	3C286	151
4.3.4	3C279 and 3C273	155
4.3.5	Summary	163
5	Polarization Observations with the CBI	166
5.1	Introduction	166

5.2	Deep Field Selection	167
5.3	08 ^h Deep Field	176
5.3.1	08 ^h Field Calibration	180
5.3.2	08 ^h Field Results	184
5.3.3	08 ^h Field Consistency Tests	186
5.4	20 ^h Deep Field	198
5.4.1	20 ^h Field Calibration	199
5.4.2	20 ^h Field Results	203
5.4.3	20 ^h Field Consistency Tests	203
5.5	Supporting Observations	210
5.5.1	3C273	211
5.5.2	3C274	212
5.5.3	Centaurus A	215
5.5.4	W44	218
5.5.5	G326.3-1.8	221
5.6	Conclusions	223
6	Likelihood Analysis of the <i>LR</i> Visibilities	227
6.1	Introduction	227
6.2	Method of Maximum Likelihood	227
6.3	Implementation Considerations	229
6.3.1	Visibility Averaging	230
6.3.2	Reals and Imaginaries	231
6.3.3	Visibility Window Function	232
6.3.4	Noise Covariance Matrix	235
6.3.5	Input Power Spectrum	236
6.3.6	Joint Fit	242
6.3.7	RX10	243
6.3.8	Explicit Calculation	244
6.4	Simulations	247

6.4.1	Baseline Likelihood Tests	248
6.4.2	Ten Channel Average	251
6.4.3	Primary Beam Pattern	252
6.4.4	Perturbations to the Calibration	256
6.5	Likelihood Analysis of the CBI Data	262
7	Summary and Conclusions	267
7.1	Results of this Thesis	267
7.2	Future Work	272
A	Phase Shifter Model	275
A.1	Transfer Matrices	275
A.2	Visibility	280

List of Figures

1.1	Post 2002 C_ℓ^{TT} , plus COBE	3
1.2	Sample spectra for C_ℓ^{TT} and C_ℓ^{EE}	6
1.3	Previous data for C_ℓ^{EE}	12
2.1	The Cosmic Background Imager	15
2.2	CBI signal chain overview	16
2.3	Quarter-wave plate	29
2.4	Downconverter overview	42
2.5	Downconverter and correlator on the array	43
2.6	L-band schematic	45
2.7	Downconverter module, <i>rf</i> side	46
2.8	Downconverter module, DC side	46
2.9	Frequency synthesizer block	49
2.10	Local oscillator assembly	49
2.11	Noise cal schematic	55
2.12	Image of the noise calibration system	56
2.13	16-way waveguide power divider	57
2.14	Noise source demonstration	58
2.15	Noise source pathology on short timescales	60
2.16	Scale factor temp dependence	62
2.17	Insertion phase error	65
2.18	CBI HEMT T_{sys} performance	69
3.1	Polarization calibration flow chart	73

3.2	Model circle fit to the gain and leakage	80
3.3	Circle fit example	84
3.4	Sample leakage comparison, RX2-RX12	92
3.5	Leakage χ^2 analysis	93
3.6	Insertion phase error	96
3.7	Leakage model fit to two baselines	98
3.8	Off-axis instrumental polarization	101
3.9	Histogram of χ^2 for the off-axis instrumental polarization	102
3.10	Sample gain calibration factors, 08 ^h field	107
3.11	Scaling demonstration for total intensity visibilities	109
3.12	Scale factor temperature dependence	110
3.13	Scaling test for <i>LR</i> baselines	112
3.14	Sample Tau A visibility comparison, 11jan00 (scan 13) vs. 06feb00	115
3.15	Tau A visibility comparison, band average	116
3.16	Tau A map for 11jan00, Model 1	120
3.17	Test of Tau A Model 1 on 3C279	122
3.18	Leakage derived from Tau A Model 1	124
3.19	Tau A, 4.885 GHz, 0.5' resolution	129
3.20	Tau A, 4.885 GHz, smoothed to 5.6' CBI beam	130
3.21	VLA data interpolation errors, 08 ^h field	132
4.1	Sample VLA gain curve	145
4.2	Effect of wind loading on the VLA antennas	150
4.3	3C286 observations with the VLA, <i>K</i> band	152
4.4	3C286 observations with the VLA, <i>Q</i> band	153
4.5	3C279 observations with the VLA, <i>K</i> band	156
4.6	3C279 observations with the VLA, <i>Q</i> band	157
4.7	3C273 observations with the VLA, <i>K</i> band.	159
4.8	3C273 observations with the VLA, <i>Q</i> band.	160
4.9	3C279/3C273 comparison with the CBI and VLA, <i>K</i> band	162

4.10	3C273 comparison with the CBI and VLA, K band	163
4.11	3C279/3C273 comparison with the CBI and VLA, Q band	164
5.1	CBI polarization observation calendar	168
5.2	CBI configuration 1	177
5.3	CBI configuration 2	178
5.4	Interpolated VLA data for 3C279, 08^h field	185
5.5	08^h field gain rms	186
5.6	08^h deep field map, Q	187
5.7	08^h deep field map, U	188
5.8	08^h deep field map, I	189
5.9	CBI T_{sys} comparison, RX12 <i>vs.</i> array mean	193
5.10	08^h field visibility uncertainties, by night	197
5.11	CBI configuration 3	200
5.12	20^h deep field map, Q	204
5.13	20^h deep field map, U	205
5.14	20^h deep field map, I	206
5.15	20^h field visibility uncertainties, by night	210
5.16	3C273 polarization comparison, CBI <i>vs.</i> VLA	213
5.17	m and χ for 3C274, CBI	215
5.18	Centaurus A, CBI map for I and P	217
5.19	Centaurus A, Parkes 64m map for I , 6.3 cm	219
5.20	Centaurus A, Parkes 64m map for P , 6.3 cm	219
5.21	W44, CBI map for I and P	220
5.22	W44, NRAO 140' map for I and P , 2.8 cm	222
5.23	G326.3-1.8, CBI	223
5.24	G326.3-1.8, Parkes 64m map for I and P , 3.5 cm	224
5.25	Tau A model cross check, revisited	226
6.1	CBI primary beam comparison	234
6.2	Visibility window functions, diagonal elements	240

6.3	(u, v) coverage, 08^h field	243
6.4	(u, v) coverage, 20^h field	244
6.5	Histogram of best fit likelihoods for baseline tests	250
6.6	CBI window function comparison, short <i>vs.</i> long	253
6.7	Likelihood curves; 08^h , 20^h , and joint fits	265
6.8	CBI limits on \mathcal{C}_ℓ^{EE} , with previous data	266
7.1	CBI configuration 7	273

List of Tables

2.1	CBI noise budget.	36
3.1	Deep 3C279 observations	90
3.2	Off-axis polarization observations	99
3.3	15jun00 amplitude scale factors	111
3.4	11jan00 scale factors	113
3.5	Test of scaling for I component of Tau A	114
3.6	Tau A Models 1 and 2	119
4.1	Dates of VLA observations	135
4.2	Sample VLA schedule	137
4.3	VLA K band data for 3C279	158
4.4	VLA Q band data for 3C279	158
5.1	Positions of the CBI fields	169
5.2	Predicted anomalous foreground in the CBI fields	171
5.3	Predicted synchrotron emission in the CBI fields	172
5.4	08 ^h field observations, configuration 1	181
5.5	08 ^h field observations, configuration 2	182
5.6	08 ^h visibility uncertainties	191
5.7	χ^2 test, 08 ^h field	194
5.8	20 ^h field observations, part 1	201
5.9	20 ^h field observations, part 2	202
5.10	20 ^h field visibility uncertainties	208

5.11	χ^2 test, 20^h field	209
6.1	Bands for the likelihood analysis, 08^h field	241
6.2	Bands for likelihood analysis, 20^h field	242
6.3	Baseline likelihood tests	249
6.4	Simulation results, 1 channel <i>vs.</i> all channels	252
6.5	Simulation results, CBI beam <i>vs.</i> Gaussian approximation	255
6.6	Simulation results, perturbations to the gain phase	258
6.7	Simulation results, perturbations to the leakage amplitude	260
6.8	Simulation results, perturbations to the leakage phase	261
6.9	CBI limits for the two deep fields	263
6.10	Likelihood biases for real data	264

Chapter 1

Introduction

1.1 Overview

The Cosmic Microwave Background has played a central role in our understanding of the universe from nearly the moment of its discovery by Penzias and Wilson in 1965 [68]. The CMB is a key component of the highly successful Standard Big Bang Model, which states that the universe began as a hot, dense plasma which expanded and cooled to its present state in the span of ~ 14 Gy. The details of the expansion—and in particular, the mechanisms by which the conditions at the surface of last scattering gave rise to the structure in the universe today—remain an area of debate which must await further observations, but all variations on the Standard Model agree that the CMB is the redshifted radiation from this initial plasma, and that as such it contains clues about the fundamental characteristics of the universe: Ω_0 , Ω_b , Ω_Λ , h_0 , n , τ , T/S [10, 40, 34]. This information resides in the minute spatial fluctuations of the intensity and polarization of the CMB. Observations have established the existence of intensity anisotropies with $\frac{\delta T}{T} \sim 10^{-5}$ on scales of $\theta \sim 0.1\text{-}0.5^\circ$, [27, 62, 58, 24, 83, 66, 50] although the polarization anisotropies, which are the focus of this work, are sufficiently small to have eluded detection until very recently [44]. The past decade has seen the emergence of low-noise detector technologies which should propel us into a new era of precision measurements of the temperature and polarization of the CMB, and the Cosmic

Background Imager is at the vanguard of these efforts.

The CBI is a 13-element interferometer array which operates in the 26-36 GHz band [63]. The array consists of 90 centimeter Cassegrain antennas mounted on a single, fully steerable platform. The antenna platform can rotate about the optical axis of the telescope, and as we will see in Chapter 2, this feature makes the CBI a very attractive instrument for polarization observations. The platform allows a range of positions for the telescopes, permitting observations of anisotropies on $\ell \sim 300 \rightarrow 4250$ scales; this range encompasses the scales over which standard models predict the bulk of the CMB power is to be found for both the total intensity and the polarization fluctuations. A downconverter that follows the receivers splits the 26-36 GHz band into 10 channels, and the spectral information in these channels allows us to discriminate against foregrounds [11]. The low-noise HEMT amplifiers in the receivers typically have $T_{sys} \sim 25$ K [76], which permit detections of $\delta T \sim 50$ μ K anisotropies with the entire array in a single night. Moreover, this sensitivity puts a detection of the polarization predicted by standard models of $\delta P \sim 5$ μ K within reach.

Figure 1.1 shows the state of our understanding of the temperature anisotropies of the CMB. While initial observations of the CMBR focused on the low ℓ region at the left of the figure, recent experiments have focused on the central region of the spectrum ($\ell \sim 10^2$) and the CBI is leading a handful of experiments which are breaking new ground in the high ℓ region at the right. This trend is driven primarily by sensitivity: cosmologists have adopted the convention of reporting $\delta T^2 \sim \ell(\ell + 1)C_\ell$, whereas observers measure $\delta T^2 \sim (2\ell + 1)C_\ell$, so the sample spectrum in Figure 1 overstates the size of the temperature anisotropies we measure by a factor $\propto \sqrt{\ell}$. The CBI is the first experiment to attack the entire $\ell \sim 300 \rightarrow 4000$ region of the power spectrum. This region is critical because it breaks the degeneracy of experiments which cover a narrow window and which must assume a shape for the initial spectrum of density fluctuations.

The CBI has made significant contributions to our understanding of the total intensity spectrum across the entire $300 < \ell < 4000$ region. We produced




Figure 1.1: This sample of points data points for C_l^{TT} taken from [86]. Sample C_l^{TT} shown for $\Omega_0 = 1.0$, $\Omega_{\text{CDM}} = 0.257$, $\Omega_\Lambda = 0.7$, $\Omega_b = 0.043$, and $h = 0.68$.

the first clear measurement of the rolloff in the spectrum between $\ell \sim 600$ and $\ell \sim 1200$ which arises from photon diffusion and the thickness of the last scattering region [62]. We have augmented this initial result with observations that sample the $\ell \sim 300 \rightarrow 3000$ region of the spectrum with resolution in ℓ of $\Delta\ell \sim 200$ [66], and these pioneering small angular scale measurements complement those of BOOMERANG [58], MAXIMA [27], DASI [24] and the VSA [83] on larger scales. The CBI has detected excess power in the $\ell \sim 2000 \rightarrow 3500$ region [50], and this detection has stimulated a considerable amount of interest in the possibility of low level secondary anisotropies on small angular scales from the Sunyaev-Zeldovich Effect (SZE) in clusters [8]. Our total intensity observations allow us to constrain Ω_0 to $\sim 15\%$ and n_s to $\sim 10\%$. When combined with data from lower ℓ and priors from SNR and surveys of Large Scale Structure, we constrain the cosmological parameters with high precision [86]. In addition, CBI maps of the SZE in resolved clusters, when combined with X-ray data, have produced a measurement of the Hubble constant, which is independent of the cosmological distance ladder [94]. Finally, the CBI's sensitivity and flexibility make it ideal for polarization observations of the CMBR, and these observations are the focus of this thesis.

1.2 Polarization of the CMBR

Standard models predict that a small fraction of the CMB photons were polarized by Thomson scattering during recombination. Rees proposed that CMBR is polarized in 1968 [79], and this initial work provided the foundation for increasingly detailed predictions of the shape of the polarization spectrum [39]. The highly evolved state of these predictions stands in stark contrast to the observational results, which until very recently consisted of limits that do not achieve the sensitivity required by standard models despite painstaking efforts [89]. Thomson scattering imposes strict constraints on the conditions that can produce net polarization which do not exist until the very last moments of recombination, when

the number of free electrons that can serve as scattering centers is small. Thus polarization anisotropies predicted by standard models are far weaker than total intensity anisotropies on the same angular scales, typically $\delta P \sim 0.1\delta T$.

In the Standard Model, the interaction of plasma physics and cosmology during recombination generates small-scale ($\theta < 1^\circ$) anisotropies in the CMB. Immediately prior to recombination, the universe is plasma in which causally connected regions of the photon-baryon fluid undergo acoustic oscillations; spatial perturbations in the gravitational potential attract the fluid, while photon pressure within the potentials resists collapse. This interplay between pressure and gravity drives the fluid through cycles of expansion and contraction, and in the absence of cosmological evolution, these oscillations would persist forever. The expansion of the universe, however, redshifts the radiation; this redshift cools the plasma to such an extent that the free electrons in the plasma combine, eliminating the coupling between radiation and matter and thereby quenching the oscillations. The absence of coupling after recombination frees the matter to collapse into the gravitational potentials, and this infall seeds the nonlinear gravitational collapse which formed the structures we see in the universe today. Recombination leaves the universe optically thin, and in the absence of reionization, the $\theta < 1^\circ$ scale features in the CMB directly reflect the distribution of radiation at recombination.

This quick sketch of the physics of recombination suffices to explain the general features of the total intensity power spectrum. The most striking characteristic of the sample spectrum for C_ℓ^{TT} shown in Figure 1 is the sequence of oscillations most clearly seen in the range $200 < \ell < 1000$. These features are due to the acoustic oscillations in the photon-baryon fluid which were frozen into the CMB at recombination. The peak on the largest angular scale represents the largest volume which had time to undergo a single collapse before recombination, while the second peak, found at half the angular size of the first, represents the largest volume which had time for a complete collapse and rarefaction, and so on. Because the matter outside the sound horizon at a particular time cannot undergo acoustic oscillations, the angular size of the first acoustic peak sets the upper limit on the

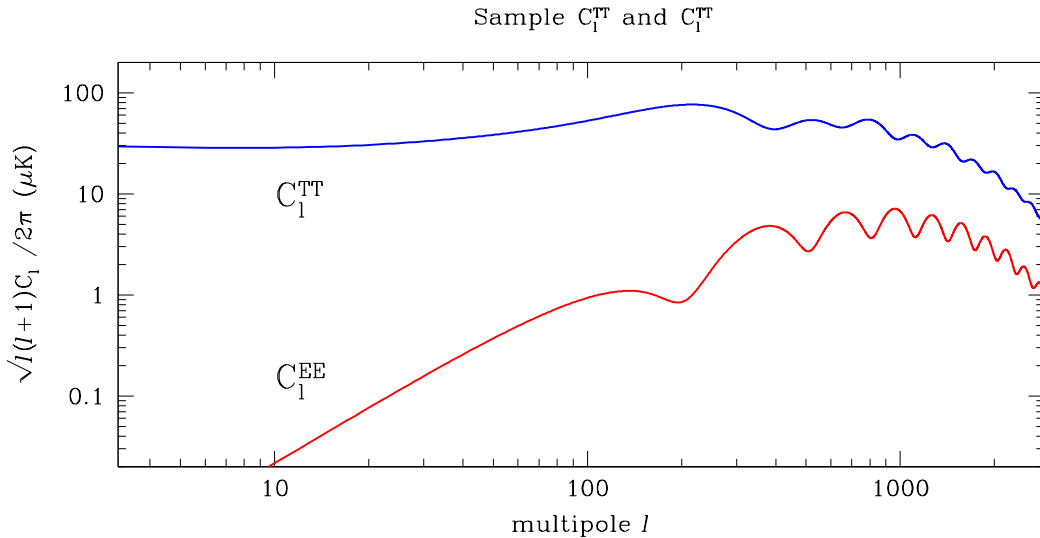


Figure 1.2: Sample spectra for C_ℓ^{TT} and C_ℓ^{EE} for a generic flat cosmology ($\Omega_b = 0.04$, $\Omega_\Lambda = 0.7$).

volume contained in the sound horizon at recombination. This volume provides a standard ruler at the surface of last scattering, so the position of the first peak in ℓ reveals the geometry of the universe in a fairly model-independent manner: $\ell_p \sim 220/\sqrt{\Omega_0}$. Figure 1.1 shows a sample of the large number of experiments that have focused on this region, including Saskatoon [59], TOCO [93], BOOMERANG, CBI, MAXIMA, DASI, and the VSA; these observations strongly favor $\Omega_0 = 1$ models.

Several effects damp the anisotropies on small angular scales ($\ell \sim 1000$). First, because the surface of last scattering is a shell of finite thickness ($\Delta z \sim 80$), it completely encompasses the anisotropies on the smallest scales, so a particular line of sight pierces multiple hot and cold regions, and the resulting averaging suppresses features from the small-scale fluctuations. Second, the diffusion of photons between hot and cold regions erases anisotropies on scales comparable to the photon mean free path [87]. The magnitude of this effect depends on Ω_b ; larger Ω_b increases the density of free electrons and therefore decreases the photon mean free path. This translates into more power on small scales. The CBI targets

this region, and our measurements of the damping tail provide a constraint on Ω_0 , which is independent of the considerations that enter on larger ($\ell \sim 200$) scales.

The polarization signal in the CMBR is generated by Thomson scattering of CMB photons from the free electrons which remain during the final moments of recombination [35]. The dependence of the Thomson scattering cross section on the polarization of the incident (ϵ) and scattered (ϵ') waves is

$$\frac{d\sigma_T}{d\Omega} \sim |\hat{\epsilon} \cdot \epsilon'|^2 \quad (1.1)$$

The radiation environment surrounding a free electron can be expressed in terms of spherical harmonics.¹ The monopole component represents the mean ambient radiation field seen by the electron, and it is clear that an isotropic distribution of CMB photons incident on the scatterer fails to produce net polarization because the scatterings from all directions cancel. Similar arguments apply to the local dipole and higher harmonics except for the quadrupole component of the ambient temperature anisotropy which gives rise to nonvanishing polarization [39].

The Y_{lm} decomposition shows that the three components of the $\ell = 2$ spherical harmonic distinguish polarization generated by three distinct mechanisms. *Scalar* modes ($m=0$) are generated by density fluctuations during recombination, while *vector* modes ($m=1$) reflect the vortical motions of the fluid about the scatterer, and *tensor* modes ($m=2$) are created by gravitational waves. Scalar modes are of particular interest to this work because standard models predict that polarization from scalar modes dwarfs the polarization from other modes on CBI angular scales.

Scalar polarization is generated by density perturbations at the surface of last scattering. In *adiabatic* models, overdensities are associated with maxima in the local gravitational potential, and in standard models the temperature fluctuations trace the underlying density perturbations on small angular scales. We can express the fluctuations in a volume as the superposition of plane wave temperature per-

¹These spherical harmonics are not to be confused with the spherical harmonics with which we characterize the CMB fluctuations on the sky.

turbations which reflect the underlying density perturbations. Consider a plane wave temperature perturbation with wavenumber \mathbf{k} , and suppose the scatterer is in the trough of the wave. Photons that flow to the scatterer from the neighboring peaks ($\mathbf{v}_\gamma \parallel \mathbf{k}$) are blueshifted relative to photons that arrive at the scatterer from within the trough ($\mathbf{v}_\gamma \perp \mathbf{k}$). The resulting intensity distribution seen by the scatterer is a quadrupole, and the scattered radiation acquires net polarization which is oriented perpendicular to \mathbf{k} . Similar arguments apply to a scatterer which is located at the peak, in which case the polarization is parallel to \mathbf{k} . The peaks and troughs, of course, are actually planes perpendicular to \mathbf{k} in this three-dimensional example, so the magnitude of this polarization has azimuthal symmetry about \mathbf{k} ; this symmetry allows us to associate this scattering mechanism with the $m=0$ component of the temperature quadrupole viewed by the scatterer. The magnitude of the polarization also depends on the line of sight $\hat{\mathbf{n}}$ to the observer; the observed polarization has a maximum at $\hat{\mathbf{n}} \perp \mathbf{k}$ and vanishes for $\hat{\mathbf{n}} \parallel \mathbf{k}$.

The surface of last scattering is a very complicated medium at the moment when the polarization is generated; it is a superposition of density perturbations which span a range of wavelengths and which are oriented in all possible directions. On small angular scales, the amplitudes of features that we measure on the sky reflect the amplitudes and coherence of the constituent modes at recombination. Despite this complexity, we can extend the preceding discussion to predict the scalar polarization patterns associated with the local temperature fluctuations at the moment of scattering. Consider the case of a hot spot; we can treat it as the coherent superposition of modes such that all are perpendicular to the observer's line of sight. We saw that the scalar polarization of radiation scattered from the crest of a mode is directed perpendicular to the crest, which implies that the polarization associated with a hot spot is directed *tangentially* to the spot. Similarly, the polarization associated with a cold spot is directed *radially* to the spot. As a practical matter, the correlation between scalar polarization patterns and temperature anisotropies is degraded by the fact that the temperature anisotropies we see on the sky are the result of many processes during and after recombination,

whereas the scalar polarization samples the density anisotropy only at the moment the polarization was generated. As a result, the correlation between temperature and scalar polarization is less than unity. The CBI measures both total intensity and polarization and can therefore address this question, although we do not do so in this work.

The sky polarization signal is a superposition of all three of the $\ell = 2$ modes, and as such we cannot distinguish between them. We can make some simplifications, however. Like any two-dimensional vector field, the polarization $\mathbf{P}(\theta, \phi)$ can be expressed as the sum of gradient and curl components, which we refer to as E and B modes:²

$$\mathbf{P}(\theta, \phi) = \vec{\nabla}E(\theta, \phi) + \vec{\nabla} \times \vec{\mathbf{B}}(\theta, \phi) \quad (1.2)$$

Intuitively, E modes can be identified as the component of the polarization for which the polarization vector is either parallel or perpendicular to the direction of greatest polarization change, whereas B modes are the modes for which the polarization is oriented $\pi/4$ from the polarization gradient. The scalar polarization around a hot spot, for example, is clearly E type; in this case, the polarization is tangential to the hot spot and therefore normal to the direction towards the center of the spot. The utility of the E/B distinction arises from the fact that after the quadrupole at the origin of scattering has been modulated by the underlying perturbation—density, vorticity, or gravity wave—the scalar polarization is exclusively E type, while vector polarization is almost entirely B type, and tensor polarization is roughly an equal combination of both. In the absence of Faraday rotation or gravitational lensing, the E and B modes are preserved, so a map of the polarization field allows us to disentangle the origins of the polarization

We have seen that standard models predict that polarization is smaller than temperature anisotropies on the same scales. The CMBFAST package computes the spectra for several combinations total intensity and polarization— C_ℓ^{TT} , C_ℓ^{EE} , C_ℓ^{TE} , and C_ℓ^{BB} —predicted by standard models for a range of cosmological para-

²These modes should not to be confused with the E and B of the underlying radiation.

meters [84]. The sample spectra in Figure 1.2 show that the E modes are $\sim 10\%$ of the intensity anisotropies found on comparable angular scales, while the B modes are smaller still. This characteristic of polarization is fairly inviolable; under the Standard Model, no reasonable combination of fundamental parameters can substantially increase this signal, and in fact, at this time the shape of the C_ℓ^{EE} spectrum is fairly tightly constrained by our understanding of the Standard Model. Because polarization is *only* generated during the final moments of recombination, causality confines polarization anisotropies to regions smaller than the horizon at that time. An inspection of Figure 1.2 reveals that this is indeed the case; the spectra all show a sharp decrease on scales larger than horizon ($\ell \sim 200$ for flat models). In this regard, the CBI, which covers the $600 < \ell < 1800$ range, is optimally suited to observe the polarization predicted by flat models.

1.3 Polarization Experiments

As with the total intensity of the CMBR, experimentalists' emphasis on the polarization power spectrum began on large angular scales and is now moving to higher ℓ as detector technology improves. Figure 3 shows the state of our understanding of the polarization power spectrum. The most striking feature of this figure is the recent detection of C_ℓ^{EE} in the $\ell \sim 200 \rightarrow 800$ range by the DASI group [44]. The fact that this detection required two years of data reflects the challenges which accompany these deep polarization observations [46]. In addition, a handful of recent experiments are notable because they achieve sensitivities that approach cosmologically interesting levels. The POLAR experiment obtained limits on large ($\ell \sim 10$) angular scales [41], and the Saskatoon experiment produced several limits on $\ell \sim 20$ scales [97, 59]. The PIQUE experiment obtained deep limits on $\ell \sim 200$ scales [30, 29], the best of which is $8.4 \mu\text{K}$; these authors also obtained the first limit on C_ℓ^{TE} with the aid of the Saskatoon maps of the total intensity of the CMBR which include their polarization fields. The best limit on very small angular scales ($\ell \sim 4000$) was obtained with the ATCA array [90]. See

Figure 1.3 for details.

Efforts to detect C_ℓ^{EE} continue apace. Many of the current polarization projects evolved from experiments dedicated to measuring total intensity. Following the conclusion of its highly successful total intensity observations, the DASI experiment retooled for polarization observations, and this work recently produced the detection discussed above; these observations will continue for at least one more season. The CBI is now dedicated to polarization observations, and in September 2002 this system saw first light (Chapter 7). BOOMERANG will fly in Antarctica with a polarized focal plane in December 2002 [49], and the MAXIPOL experiment, an adaptation of MAXIMA, will fly from Texas in 2003 with a polarized focal plane as well. In addition, the POLAR experiment has been modified to focus on smaller angular scales [74]. QUEST will target the $\ell \sim 300 \rightarrow 1800$ region [73]. Finally, two all-sky satellite missions will make polarization observations of the CMBR. The MAP mission, which launched in the summer of 2001, has several polarized channels whose sensitivity extends to $\ell \sim 500$ [1]. The Planck mission, which is scheduled to launch in 2007, will probe scales down to $\ell \sim 2000$ in polarization [20, 95]. While these satellite missions promise to make the definitive measurements of the polarization of the CMBR on these scales, they must compete with smaller, more nimble ground based experiments which are better suited to adapt their strategies as our understanding of the CMBR evolves.

1.4 Thesis Outline

The goal of this thesis is to obtain a measurement of the polarized signal, and to this end we implemented a polarization detection effort in parallel with the total intensity observations which constitute the CBI's primary objective. Chapter 2 describes the general features of the CBI and highlights several components which were developed and implemented as part of this thesis. Calibration is a critical component of the polarization observations presented in this work; the CBI derived its calibration from observations of the extragalactic source 3C279, which

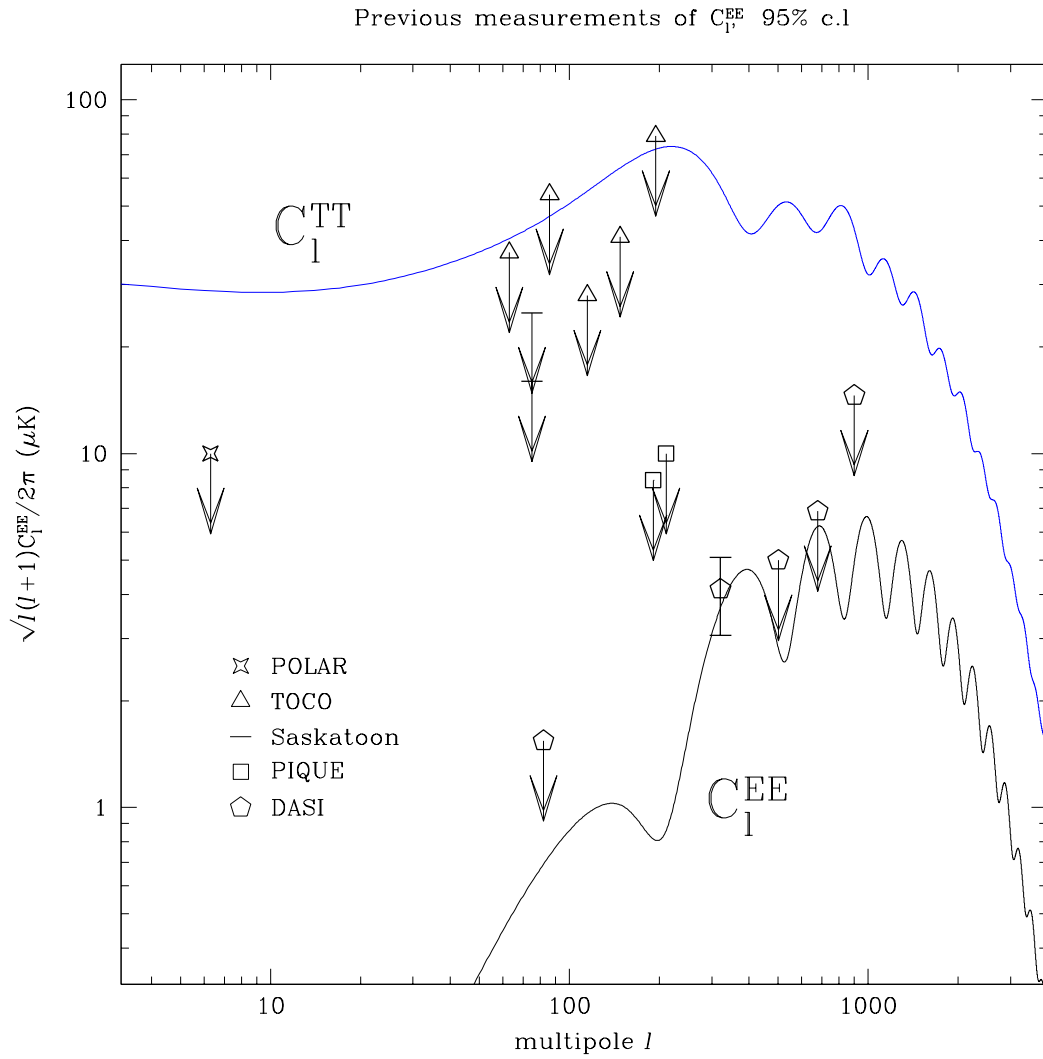


Figure 1.3: Previous data for C_ℓ^{EE} . All data are 95% upper limits except for the DASI detection at $\ell = 321$, whose error bars are 1σ .

was monitored contemporaneously with the Very Large Array (VLA). Chapter 3 discusses the polarization calibration of the CBI, and Chapter 4 presents the VLA observations. The CBI collected nearly 156 hours of polarization data on 99 nights; most of these data were obtained on deep fields which represent our best hope for detecting the polarization, while the balance was dedicated to observations of calibrators and supporting sources. Chapter 5 discusses these observations, and Chapter 6 describes the analysis of the deep field data.

Chapter 2

The Cosmic Background Imager

2.1 Introduction

This chapter describes the Cosmic Background Imager, which is shown in Figure 2.1. The platform accommodates baselines whose lengths range from 100 cm to ~ 550 cm; the minimum separation is set by the antenna diameter, and the maximum separation is limited by the size of the deck. The platform accommodates a wide range of positions for the dishes, and this flexibility allows us to target a variety of cosmological problems. The CBI was designed and built on the Caltech campus between August of 1995 and August of 1999, at which point it was deployed to the Chajnantor site in the Chilean Andes. Routine observations began in January 2000.

In Section 2 we discuss the response of an interferometer to fluctuations in the CMBR. Section 3 describes the CBI. Section 4 focuses on several specific pieces of instrumentation; a substantial component of this thesis entailed the design and implementation of major components of the CBI, and this section discusses these systems in detail. The CBI has a sister instrument, the Degree Angular Scale Interferometer [47], which is sited at the South Pole; the systems described in Section 4 were deployed with DASI as well.



Figure 2.1: The Cosmic Background Imager at an altitude of 5080 m on the Chajnantor site in northern Chile. The 13 Cassegrain antennas have cylindrical shields which reduce crosstalk between the antennas. In this picture, the array is in the initial sparse configuration (configuration 1); in April 2000, the array was reconfigured to emphasize shorter baselines. The clamshell dome, seen open at the base of the telescope, is surrounded by shipping containers which contain living spaces, a control room, and laboratory and machine shop facilities. A pair of diesel generators provides power for the facility, and a cell phone with an amplified link ties the site to the project base camp in the nearby town of San Pedro de Atacama.

CBI Signal Chain Overview

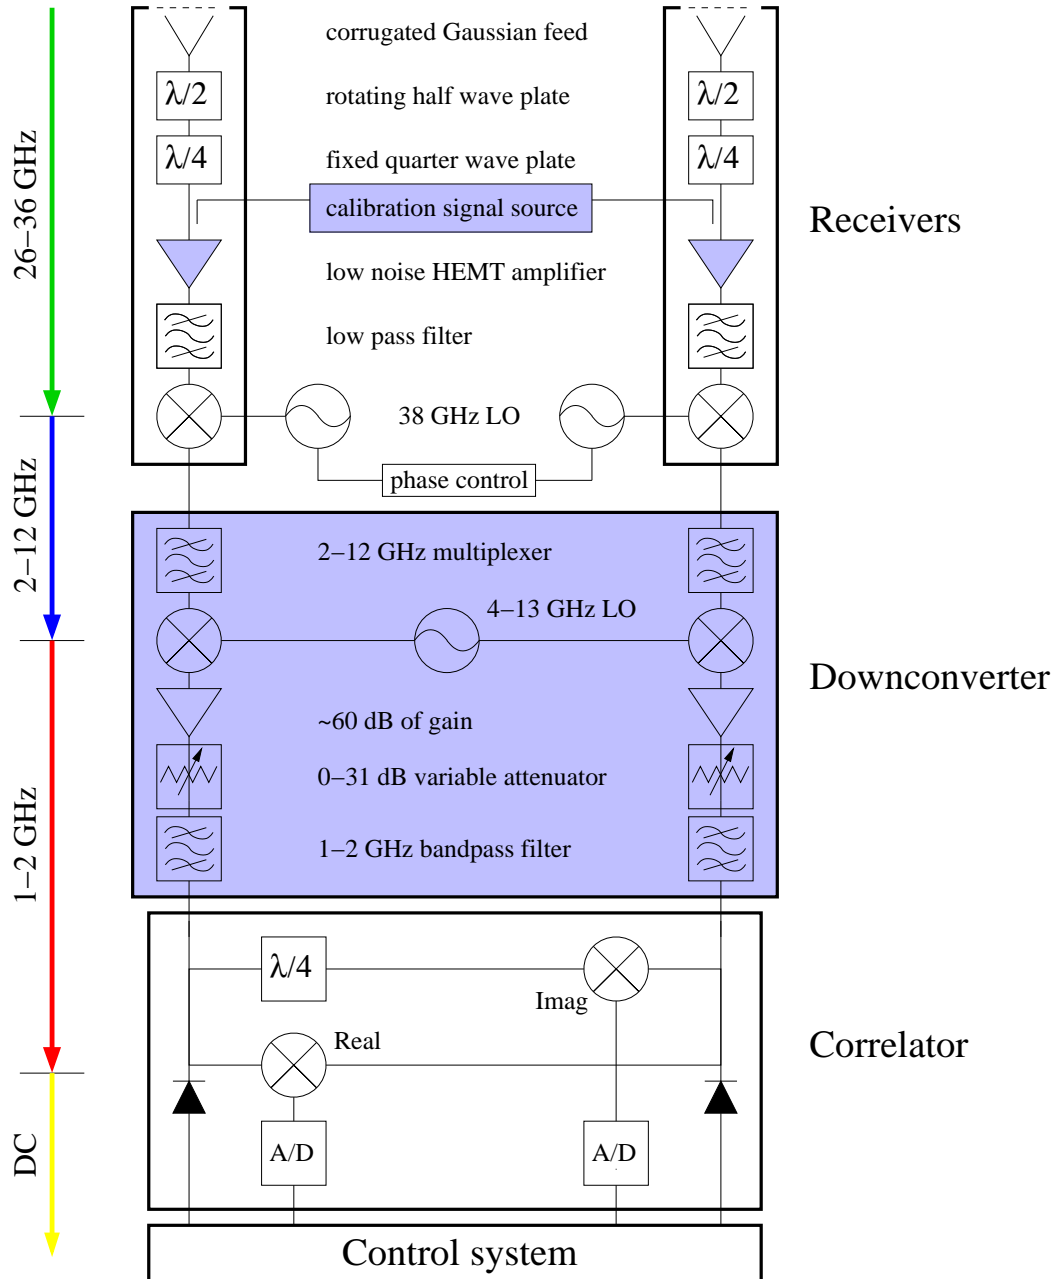


Figure 2.2: CBI signal chain overview for a single baseline. The arrows at left note the frequency changes through the signal chain. The shaded components are part of this thesis; this work discussed in Section 2.4.

2.2 Response of an Interferometer

We first review the response of an interferometer to fluctuations in the CMBR. This problem has been widely discussed, *e.g.*, [92, 32, 57, 77, 96], and this section presents only the key results.

Our goal is to characterize the spatial fluctuations in the temperature and polarization of the CMBR. The treatment which follows applies to total intensity observations. It is conventional to express the temperature fluctuations in terms of the spherical harmonics $Y_{\ell m}(\theta, \phi)$. The CBI operates on small angular scales for which the flat sky approximation applies, however, so we may express the fluctuations in terms of a simple two-dimensional Fourier transform:

$$\frac{\delta T(\mathbf{x})}{T} \sim \int a(\mathbf{u}) e^{2\pi i \mathbf{u} \cdot \mathbf{x}} d\mathbf{u} \quad (2.1)$$

The Planck function $B(\nu, T)$ relates the temperature fluctuations on the sky to the intensity fluctuations measured by an interferometer:

$$\delta I(\theta, \phi) = \frac{\partial B(\nu, T)}{\partial T} \delta T(\theta, \phi) \quad (2.2)$$

where the partial derivative is evaluated at $T = T_0 = 2.726$ K.

A single baseline interferometer measures components of the Fourier transform of $I_\nu(\mathbf{x})$, the brightness distribution on the sky. This quantity, called the *visibility* and denoted by $\mathcal{V}_\nu(\mathbf{u})$, is the time average of the complex product of the voltages (V_j, V_k) produced by the two antennas that form the interferometer. On average, the product of the receiver outputs is in turn proportional to the complex product of the \mathcal{E} -fields that impinge on the receivers:

$$\mathcal{V}_\nu(\mathbf{u}) = \langle \mathcal{E}_j \mathcal{E}_k^* \rangle = G_{jk}^{-1} \langle V_j V_k^* \rangle = \kappa \int_{4\pi} A_k(\mathbf{x} - \mathbf{x}_k) \frac{\delta T(\mathbf{x})}{T_0} e^{-2\pi i \mathbf{u} \cdot \mathbf{x}} d\mathbf{x} \quad (2.3)$$

The factor of κ is obtained from Equation 2.2; it relates the visibilities, which are measured in Janskys, to the specific intensity I_ν , or equivalently, to the brightness temperature. The factor of G_{jk}^{-1} is a complex gain term which reflects the

cumulative change in amplitude and phase of the signal as it passes through the system; we will neglect this term until the calibration discussion in Chapter 3. The visibility is a sample in the aperture domain whose size is set by the Fourier transform of the primary beam and whose position $\mathbf{u} = (u, v)$ is determined by the orientation and length of the baseline.¹

The expression for the visibility immediately lends itself to interpretation: the visibility is the superposition of two quantities—the primary beam pattern $A(\mathbf{x})$ centered at a point \mathbf{x}_k , and a corrugation whose pitch and orientation are given by $2\pi i\mathbf{u} \cdot \mathbf{x}$ —on the sky brightness distribution $I_\nu(\mathbf{x})$. The coherence between $I_\nu(\mathbf{x})$ and $2\pi i\mathbf{u} \cdot \mathbf{x}$ across $A(\mathbf{x})$ determines the amplitude and phase of $\mathcal{V}_\nu(\mathbf{u})$. The visibility encompasses two angular scales: the primary beam sets the field of view, while the corrugation determines the resolution. The primary beam of the CBI is approximated by a Gaussian of FWHM $a_p \sim \lambda/D \sim 47'$ at 1 cm, so while the integrand in Equation 2.3 is meant to be evaluated over the entire sky, in practice the primary beam confines the integral to $\Omega_p \ll 4\pi$. The corrugation is the effective pixel for the baseline; the superposition of many visibilities from baselines at different orientations results in a *synthesized beam*.

The similarity between the visibility in Equation 2.3 and Fourier transform in Equation 2.1 illustrates the power of interferometry as a tool for understanding the CMBR. A visibility provides a matched filter that selects the CMBR fluctuations which are coherent with the corrugation defined by $2\pi i\mathbf{u} \cdot \mathbf{x}$. By measuring many visibilities we can estimate their variance, which in turn provides an estimate of the underlying power spectrum of fluctuations.

Up to this point, we have regarded $I_\nu(\mathbf{x})$ as a generalized brightness which encompasses the full polarization characteristics of the celestial signal; we now discuss polarization in detail. Polarization introduces considerations beyond those which are necessary for total intensity [17, 54, 26, 82, 25]. In addition, unlike many interferometers, the CBI is a single-mode system: each receiver responds to either

¹Note that the autocorrelation functions of the primary beams for a pair of visibilities can overlap in the aperture domain, and this overlap gives rise to correlations between the visibilities.

right or left circular polarization (*RCP* or *LCP*), so under ideal circumstances each receiver measures only one component of the radiation incident on the array: either \mathcal{E}^R or \mathcal{E}^L . This characteristic of the CBI has consequences for how we acquire, calibrate, and interpret the polarized visibility data.

Combinations of *LCP* and *RCP* antennas measure four types of visibilities, which are in turn related to the four Stokes parameters. It can be shown that

$$\begin{pmatrix} \mathcal{V}^{RR} & \mathcal{V}^{RL} \\ \mathcal{V}^{LR} & \mathcal{V}^{LL} \end{pmatrix} = \begin{pmatrix} \langle \mathcal{E}_j^R \mathcal{E}_k^{R*} \rangle & \langle \mathcal{E}_j^R \mathcal{E}_k^{L*} \rangle \\ \langle \mathcal{E}_j^L \mathcal{E}_k^{R*} \rangle & \langle \mathcal{E}_j^L \mathcal{E}_k^{L*} \rangle \end{pmatrix} \Rightarrow \begin{pmatrix} \tilde{I} + \tilde{V} & [\tilde{Q} + i\tilde{U}] \\ [\tilde{Q} - i\tilde{U}] & \tilde{I} - \tilde{V} \end{pmatrix} \quad (2.4)$$

The second correspondence is not an equality because the visibilities are integrated over the primary beam, but the Stokes parameters are not. Since the CBI uses single-mode receivers, each baseline measures either \mathcal{V}^{LL} or \mathcal{V}^{LR} . The data presented in this thesis were obtained with twelve of the CBI antennas configured for *LCP* and one configured for *RCP*; this choice results in 12 cross polarized baselines (*LR*) and 66 total intensity (*LL*) baselines. Equation 2.4 suggests that these visibilities do not suffice to measure both Q and U . Because of the CBI's deck rotation, however, this is not the case.

Under most circumstances, an interferometer which lacks one of *LR* and *RL* at each (u, v) point would be unable to simultaneously isolate both \tilde{Q} and \tilde{U} ;² the Fourier transforms for Q and U are complex for extended sources, so both \mathcal{V}^{LR} and \mathcal{V}^{RL} are required to solve for \tilde{Q} and \tilde{U} . The CBI's deck rotation allows us to circumvent this problem. The visibility measured by an *LR* baseline at a point \mathbf{u} is simply

$$\mathcal{V}^{LR}(\mathbf{u}) = \int A(\mathbf{x})P^*(\mathbf{x})e^{-2\pi i\mathbf{u}\cdot\mathbf{x}}d\mathbf{x} \quad (2.5)$$

with $P^* = Q - iU$. If we repeat the measurement at $-\mathbf{u}$, we obtain

$$\mathcal{V}^{LR}(-\mathbf{u}) = \int A(\mathbf{x})P^*(\mathbf{x})e^{2\pi i\mathbf{u}\cdot\mathbf{x}}d\mathbf{x} \quad (2.6)$$

²This statement does not apply to point sources at the phase center. In this case, \tilde{Q} and \tilde{U} are both real, and they are simply the real and imaginary parts of the complex visibility, so \mathcal{V}^{LR} suffices.

Conway and Kronberg note that $[\mathcal{V}^{LR}(-\mathbf{u})]^* = \mathcal{V}^{RL}(\mathbf{u})$, which is precisely the quantity we seek [17]. With the CBI's deck rotation, the mapping from \mathbf{u} to $-\mathbf{u}$ is trivial: we simply rotate the deck through 180° . Thus, with the aid of deck rotation, we can obtain near-simultaneous measurements of \tilde{Q} and \tilde{U} for observations of extended sources with our single-mode system.

The visibilities measured by the CBI must be modified to reflect several additional considerations. We must include the effect of the rotation of the feeds about the optical axis of the array [17]. The orientation of receiver advances or retards the phase of a circularly polarized wave, so we augment the fields viewed by the receivers with a phase term to reflect this dependence: $(\mathcal{E}^R, \mathcal{E}^L) \rightarrow (\mathcal{E}^R e^{-i\theta_k}, \mathcal{E}^L e^{i\theta_k})$, where $\theta_k = \tan^{-1}(u_k/v_k)$ for visibility k . In addition, as noted in Chapter 1, the polarization of the CMBR can be expressed in terms of scalar and pseudo-scalar modes \tilde{E} and \tilde{B} . These assignments are convenient because the physical processes which polarize the CMBR give rise to orthogonal signatures in this basis; scalar fluctuations produce E -type polarization, while tensor fluctuations produce B -type. \tilde{E} and \tilde{B} are simply linear combinations of \tilde{Q} and \tilde{U} :

$$\tilde{Q}(\mathbf{v}) = \tilde{E}(\mathbf{v}) \cos 2\theta_v - \tilde{B}(\mathbf{v}) \sin 2\theta_v \quad (2.7)$$

$$\tilde{U}(\mathbf{v}) = \tilde{E}(\mathbf{v}) \sin 2\theta_v + \tilde{B}(\mathbf{v}) \cos 2\theta_v \quad (2.8)$$

where \mathbf{v} is the position in the aperture domain for the point of interest, which can be expressed in polar coordinates: $\mathbf{v} = (|v|, \theta_v)$. With these considerations in mind, the Stokes representation takes the form

$$\begin{pmatrix} \langle \mathcal{E}_j^R \mathcal{E}_k^{R*} \rangle & \langle \mathcal{E}_j^R \mathcal{E}_k^{L*} \rangle \\ \langle \mathcal{E}_j^L \mathcal{E}_k^{R*} \rangle & \langle \mathcal{E}_j^L \mathcal{E}_k^{L*} \rangle \end{pmatrix} \Rightarrow \begin{pmatrix} \tilde{I} + \tilde{V} & [\tilde{E} + i\tilde{B}]e^{2i(\theta_v - \theta_k)} \\ [\tilde{E} - i\tilde{B}]e^{-2i(\theta_v - \theta_k)} & \tilde{I} - \tilde{V} \end{pmatrix} \quad (2.9)$$

This expression relates the visibilities measured by the interferometer to \tilde{E} and \tilde{B} , the quantities of interest for cosmology. Although the CBI measures \tilde{Q} and \tilde{U} , it cannot directly distinguish between \tilde{E} and \tilde{B} because the finite extent of the primary beam mixes the two components [57]. Standard models predict vanishingly

small circular polarization [18], so for the discussion that follows we assume $V = 0$.

Now consider the visibility. The visibility is a powerful tool for measuring spatial fluctuations in the CMBR because it selects the fluctuations over the extent of the primary beam that maximize the coherence with the corrugation defined by $2\pi i \mathbf{u} \cdot \mathbf{x}$. To compare observations to theory, however, we must use our sample of visibilities to estimate the properties of the underlying distribution of fluctuations. The natural starting point is the sample mean $\langle \mathcal{V}_p \rangle$. This expectation vanishes, however, because the boundary condition enforced by the spherical (closed) sky requires that only the monopole can have a nonzero mean. The next moment is the sample variance, of which there are two varieties: the variance of a visibility $\langle \mathcal{V}_p \mathcal{V}_p^* \rangle$, and the covariance of two visibilities $\langle \mathcal{V}_p \mathcal{V}_q^* \rangle$. These quantities form the diagonal and off-diagonal elements of a *covariance matrix* \mathbf{M}_{pq} which characterizes the correlation between visibilities measured at \mathbf{u}_p and \mathbf{u}_q . We will assume that the fluctuations are Gaussian, in which case the mean and the variance suffice to fully characterize their distribution. Non-Gaussian fluctuations would require the consideration of higher moments, but an analysis of the BOOMERANG and CBI maps provides no evidence to support such a treatment [7].

The covariance matrix is a component of the maximum likelihood analysis which we employ to test the *LR* visibility data for the presence of a celestial signal. While the details of the likelihood calculation are deferred until Chapter 6, we present the calculation of the covariance matrix elements \mathbf{M}_{pq} here because this discussion demonstrates how the covariance matrix connects the visibilities to the underlying power spectrum—and the ease with which the spectrum can be extracted from the visibilities.

The convolution theorem provides a path to obtaining the elements of the covariance matrix from the visibilities. In the Fourier domain, the visibility measured at a point \mathbf{u}_p is convolution of the beam with the source brightness:

$$\mathcal{V}_p^{LR}(\mathbf{u}_p) = \kappa_p \tilde{A}_p(\mathbf{u}_p) * [\tilde{E}(\mathbf{v}) - i\tilde{B}(\mathbf{v})] e^{-2i(\theta_v - \theta_p)} \quad (2.10)$$

$$= \kappa_p \int_{4\pi} \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) [\tilde{E}(\mathbf{v}) - i\tilde{B}(\mathbf{v})] e^{-2i(\theta_v - \theta_p)} d^2\mathbf{v} \quad (2.11)$$

The autocorrelation function of the aperture distribution weights the visibility's response in the aperture plane. To a good approximation, the antenna primary beam can be modeled by a Gaussian of FWHM $\Delta_x(\nu) = 45.2'$ ($31 \text{ GHz}/\nu$), which yields in the aperture domain:

$$\tilde{A}_p(\mathbf{u}_p - \mathbf{v}) = \frac{1}{2\pi\sigma_p^2} e^{-\frac{(\mathbf{u}_p - \mathbf{v})^2}{2\sigma_p^2}} \quad (2.12)$$

where $\sigma_u(\nu) = 30.3$ ($\nu/31 \text{ GHz}$). The subscript of p reminds us that the autocorrelation function depends on the frequency of the sample at \mathbf{u}_p . In practice, each visibility is accompanied by an uncertainty ϵ_p which reflects the noise in the visibilities; the contributions to these uncertainties are considered in Chapter 6, and we ignore them here.

Let \mathbf{M}_{pq} denote the covariance between two LR visibilities $\mathcal{V}(\mathbf{u}_p)$ and $\mathcal{V}(\mathbf{u}_q)$:

$$\begin{aligned} \mathbf{M}_{pq} &= \left\langle \left[\kappa_p \tilde{A}_p(\mathbf{u}_p) * [\tilde{E}(\mathbf{v}) - i\tilde{B}(\mathbf{v})] e^{-2i(\theta_v - \theta_p)} \right] \right. \\ &\quad \left. \times \left[\kappa_q \tilde{A}_q(\mathbf{u}_q) * [\tilde{E}(\mathbf{w}) - i\tilde{B}(\mathbf{w})] e^{-2i(\theta_w - \theta_q)} \right]^* \right\rangle \end{aligned} \quad (2.13)$$

We can decouple the effect of the spatial sampling function from the fluctuations on the sky. Expand the expectation to find

$$\begin{aligned} \mathbf{M}_{pq} &= \left\langle \kappa_p \int d^2\mathbf{v} \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) [\tilde{E}(\mathbf{v}) - i\tilde{B}(\mathbf{v})] e^{-2i(\theta_v - \theta_p)} \right. \\ &\quad \left. \times \kappa_q \int d^2\mathbf{w} \tilde{A}_q^*(\mathbf{u}_q - \mathbf{w}) [\tilde{E}^*(\mathbf{w}) + i\tilde{B}^*(\mathbf{w})] e^{2i(\theta_w - \theta_q)} \right\rangle \end{aligned} \quad (2.14)$$

The antenna beams remain constant for an ensemble of visibilities, so the expectation comes into the integral; \mathbf{M}_{pq} then takes the form

$$\begin{aligned} \mathbf{M}_{pq} &= \kappa_p \kappa_q e^{2i(\theta_p - \theta_q)} \int d^2\mathbf{v} \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) e^{-2i\theta_v} \\ &\quad \times \int d^2\mathbf{w} \tilde{A}_q^*(\mathbf{u}_q - \mathbf{w}) e^{-2i\theta_w} \langle \tilde{E}(\mathbf{v}) \tilde{E}^*(\mathbf{w}) + \tilde{B}(\mathbf{v}) \tilde{B}^*(\mathbf{w}) \rangle \end{aligned} \quad (2.15)$$

which is accompanied by a pair of cross terms that cancel under the expectation. This expression isolates the characteristics of the brightness on the sky in the expectation.

Since the CMBR fluctuations are Gaussian and isotropic, the variance of fluctuations on each scale depends only on the scale and the frequency of the observations. With this consideration in mind, we can express the expectations in terms of coordinate-independent power spectra $C(v)$:

$$\langle \tilde{E}(\mathbf{v})\tilde{E}^*(\mathbf{w}) \rangle = C^{EE}(v)\delta^2(\mathbf{v} - \mathbf{w}) \quad (2.16)$$

$$\langle \tilde{B}(\mathbf{v})\tilde{B}^*(\mathbf{w}) \rangle = C^{BB}(v)\delta^2(\mathbf{v} - \mathbf{w}) \quad (2.17)$$

with $v = |\mathbf{v}|$. The delta functions eliminate one integration in Equation 2.15, so we have

$$\mathbf{M}_{pq} = \kappa_p \kappa_q e^{2i(\theta_p - \theta_q)} \int \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) \tilde{A}_q^*(\mathbf{u}_q - \mathbf{v}) [C^{EE}(v) + C^{BB}(v)] d^2\mathbf{v} \quad (2.18)$$

for a pair of LR visibilities. Standard practice calls for explicitly separating the two components of \mathbf{M}_{pq} :

$$\mathbf{M}_{pq} = \kappa_p \kappa_q e^{2i(\theta_p - \theta_q)} \int_0^\infty W_{pq}(v) [C^{EE}(v) + C^{BB}(v)] v dv \quad (2.19)$$

where $W_{pq}(v)$ is the *visibility window function*:

$$W_{pq}(v) = \int_0^{2\pi} \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) \tilde{A}_q^*(\mathbf{u}_q - \mathbf{v}) d\theta_v \quad (2.20)$$

In Chapter 6 we will see that the covariance matrix must incorporate the correlation between $\mathcal{V}(\mathbf{u}_p)$ and $\mathcal{V}(-\mathbf{u}_q)$ as well as that between $\mathcal{V}(\mathbf{u}_p)$ and $\mathcal{V}(\mathbf{u}_q)$; this additional correlation provides a means of separating $C^{EE}(v)$ and $C^{BB}(v)$ in Equation 2.18. The window function provides an explicit measure of the interferometer's sensitivity to fluctuations as a function of ℓ ; it maps the response of the visibility pair to ℓ -space. For large ℓ , the correspondence between the flat

space radius v and the spherical harmonic ℓ is straightforward: $\ell = 2\pi v$ [77, 96]. The autocorrelation for a baseline of length 100 cm, for example, samples scales of $\ell = (100 \text{ cm})/(1 \text{ cm}) \times 2\pi \simeq 630$ near the band center. The width of the primary beam in the aperture plane—the factor of A^2 in Equation 2.18—yields $\Delta\ell \simeq \pm 420$. Figure 6.2 shows a variety of window functions for the diagonal covariance matrix elements.

We can now describe the temperature and polarization fluctuations in terms of the spectra which scale the covariance matrix elements. Expressing the spatial fluctuations in terms of $\delta T/T$, we have

$$(2\pi v)^2 C^{\alpha\alpha}(v) = \ell(\ell + 1) C_\ell^{\alpha\alpha} = 2\pi \left(\frac{\delta T}{T} \right)^2 \quad (2.21)$$

for $\alpha\alpha = TT$, TE , or EE . Other correlations yield C_ℓ^{TB} and C_ℓ^{EB} , and while standard models predict that these combinations should vanish, these spectra, like $C^{EE}(v)$ and $C^{TE}(v)$, remain to be measured with the sensitivity required to explore cosmologically interesting levels. In addition to C_ℓ^{TT} and C_ℓ^{EE} , the CBI measures C_ℓ^{TE} , but this spectrum is acutely susceptible to errors in the instrumental polarization calibration, and thus we do not consider it in this work [57].

2.3 The Cosmic Background Imager

We are particularly concerned with the aspects of the interferometer which affect its polarization performance, so we emphasize the signal paths through the system and their effects on the polarization response. We will see that imperfections in the signal chain hardware contaminate the cross polarized visibilities with spurious instrumental polarization; in this section we derive a model for the instrumental polarization and in Chapter 3 we compare the model to measurements of the instrumental polarization derived from observations of celestial sources of known polarization. This section also discusses several aspects of the global performance of the instrument—such as its sensitivity and its pointing—because these considerations affect the polarization performance as well.

2.3.1 Signal Path

The CBI signal chain converts the electric fields viewed by the antennas, \mathcal{E}_j and \mathcal{E}_k , into DC levels $\langle V_j V_k^* \rangle$ at the correlator output. The signal chain consists of three stages: the receivers, the downconverter, and the correlator; very simply, the receivers and the downconverter condition the signals to be multiplied by the correlator. Figure 2.2 shows the signal paths associated with a single baseline. The shaded components are parts which represent work presented in this thesis, and these components are discussed in detail in Section 2.4. The entire radiofrequency signal path shown here is located in the CBI deck.

Receivers

In Figure 2.2 signals travel from the top to the bottom. Radiation incident on the array enters the signal chain through receivers at the Cassegrain foci of the antennas. A phase shifter assembly follows the horn; the phase shifter consists of a rotating $\lambda/2$ plate, which modulates crosstalk between the receivers; and a fixed $\lambda/4$ plate, which transforms circular polarization at the receiver inputs to linear polarization for the backend system. The mode of circular polarization to which the receiver responds is determined by the orientation of the quarter-wave plate relative to the rectangular waveguide that follows it. The phase shifter assembly is the dominant source of instrumental polarization (Section 2.3.4).

Downstream of the quarter-wave plates all thirteen signal paths are identical. The components in the signal chain are susceptible to gain variations, so the CBI has an internal calibration source which injects a signal into the signal chain to monitor gain stability. This calibration system is a component of this thesis (Section 2.4.2). The first stage of amplification is the most significant source of noise, so the CBI employs low-noise HEMT³ amplifiers at the front end which typically have $T_{sys} \sim 15$ K and ~ 30 dB of gain across the 26-36 GHz band. The production of the HEMT amplifiers is also part of this work (Section 2.4.3). The HEMT

³High Electron Mobility Transistor

is followed by a downconverter which combines a 38 GHz Local Oscillator (LO) signal with the 26-36 GHz HEMT output to mix the receiver band down to 2-12 GHz. The phase of each LO is modulated by a Walsh function that is unique to the antenna, and each visibility is demodulated at the correlator output by the product of the corresponding Walsh functions; this modulation scheme rejects spurious correlated signals from sources downstream of the LOs. The HEMTs require cryogenic conditions to achieve this low noise performance, so a mechanical helium gas cryocooler cools the cryostat to ~ 10 K.

Warm Signal Path

The receiver outputs must undergo several additional transformations before they can be multiplied by the correlator. The correlator has strict input requirements; it requires +16 dBm signals because powers below this level will starve the correlator, while levels above it will saturate it. The correlator operates in the 1-2 GHz band, and the rejection outside of this bandpass must be good to avoid contaminating the visibility with out-of-band power. The receivers, however, produce signals in the 2-12 GHz band with a wide range of powers—the cold downconverter, for example, can impart a slope as large as 10 dB to the signals which pass through it. A significant degree of manipulation is required to bring the receiver outputs into alignment with the correlator inputs; the warm downconverter accomplishes this task. The downconverter was designed for this thesis (Section 2.4.1).

A filter bank at the downconverter input splits the 2-12 GHz band into ten bands with $\Delta\nu=1$ GHz and passes each band to the input of one of 130 identical channels.⁴ A mixer at the input of each channel combines one of ten local oscillator (LO) tones with the filter bank output to transform the signal to the 1-2 GHz band. This signal passes to a series of amplifiers that provide the necessary ~ 70 dB of gain, and to several filters that improve the definition of the band. Each channel has a variable attenuator that sets the output power to the level required by the correlator with resolution of ~ 1 dB.

⁴13 antennas \times 10 bands/antenna

The downconverter outputs terminate in the correlator. The CBI employs an analog complex correlator [61]. Each baseline has a pair of Gilbert cell multipliers for each band that multiply the signals out of the receivers to obtain the real (cosine) and imaginary (sine) correlation products. As Figure 2.2 shows, the real and imaginary branches of each correlator channel are identical save a 90° offset in one branch. The CBI requires 1560 multipliers: real and imaginary multipliers for each of 78 baselines in each of the ten CBI bands. The multiplier outputs are DC levels which are sampled, digitized, and written to the archive by the online control system. In addition, each receiver output terminates at a detector which measures the total power of the channel; the control system compares the 130 total powers to the +16 dBm required by the correlator and commands the downconverter to set the power levels for each channel accordingly.

2.3.2 Polarization Considerations for the Signal Path

We saw that the signals paths for the *LCP* and *RCP* antennas are largely identical; only the phase shifters distinguish between them. The phase shifter plays a central role in the CBI's polarization performance because it defines the polarization of the receivers and is a significant source of instrumental polarization contamination.

The phase shifter consists of a pair of teflon slabs inserted into the circular waveguide following the feed. The first slab resides in a section of rotating waveguide; it is a rotating half-wave plate that was incorporated into the signal path to modulate crosstalk between adjacent antennas. The second slab is a fixed quarter-wave plate which selects the circular polarization of the receiver. Both plates are solid teflon sections whose lengths and shapes are designed to produce the desired phase delay at the center of the 26-36 GHz band. We have developed a model for the plates which shows that the half-wave plates change the combination of the Stokes parameters to which a baseline responds; that the half-wave plates modulate the instrumental polarization; and that the total intensity of a $\sim 10\%$ polarized source is contaminated at the 1% level by instrumental polarization whose contribution varies with deck angle position. The model shows that

all of the CBI's large ($\sim 10\%$) instrumental polarization can be attributed to the phase shifter assembly, as discussed in Chapter 3.

The model considers three types of errors in the quarter and half-wave plates. First, both plates contribute a complex insertion loss to the incident radiation. The imaginary part of this loss corresponds to the insertion phase of the plate, which we use to obtain the desired quarter and half-wave phase shifts. The plates have substantial bandpass errors at the edges of the band; benchtop measurements of a sample of quarter-wave plates show that its insertion phase departs from a quarter wave by $\sim -2^\circ$ at 26 GHz to $\sim +5^\circ$ at 36 GHz. These bandpass errors are a significant source of instrumental polarization contamination. Second, the real part of the insertion loss attenuates the signal, so the model must also consider the instrumental polarization which arises from degradation of the sky signal in the phase shifter. Third, both plates can have orientation errors—for the rotating half-wave plates these errors arise from the resolution of the feedback loop which sets the plate position, while for the fixed quarter-wave plates, orientation errors were introduced when the plates were installed. In the former case, the orientation error contributes noise to the instrumental polarization, while the latter case introduces a fixed systematic offset. In both cases, these orientation errors are a few degrees peak to peak, or $\sim 1^\circ$ rms.

The quarter-wave plate defines the polarization of the receiver, so we first focus on this component. The quarter-wave plate is oriented at $\pm 45^\circ$ with respect to the rectangular waveguide that follows it (Figure 2.3). To illustrate how the plate determines the polarization of the receiver, consider an incident wave parallel to the \hat{e}_x -axis which propagates in the $+\hat{e}_z$ -direction. This analysis assumes that the viewer is looking in the direction of propagation—into the feed, as well as a standard right-handed coordinate system, so positive angles sweep from \hat{e}_x to \hat{e}_y . Orient the quarter-wave plate at an angle of -45° , and inject a test wave parallel to the \hat{e}_x -axis into the quarter wave plate. The test signal can be expressed in terms of components parallel to and perpendicular to the quarter-wave plate (the \hat{e}'_x and \hat{e}'_y directions). An ideal quarter-wave plate imparts a $\pi/2$ lag to

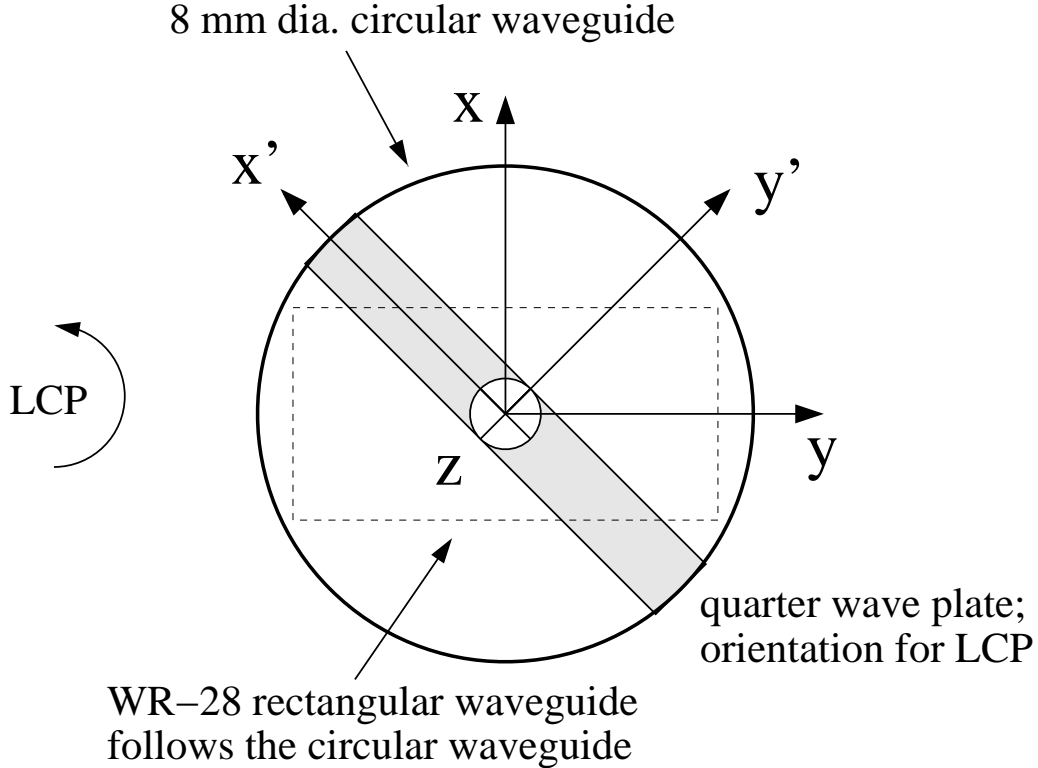


Figure 2.3: Quarter-wave plate, view looking into the plate. The wave propagates into the page (away from the reader).

the parallel component, while the perpendicular component passes unperturbed.⁵ Suppose that the \mathcal{E} -field in the \hat{e}'_y direction peaks at the output of the quarter-wave plate, so that the parallel component, which lags, is at zero on its way to its maximum in the $+\hat{e}'_x$ direction. As the test wave propagates, the resultant vector sweeps counterclockwise from the $+\hat{e}'_y$ direction to the $+\hat{e}'_x$ direction: this temporal evolution of the wave corresponds to left circular polarization according to the IEEE definition adopted by the IAU [25]. By the same arguments, if the quarter-wave plate is oriented at $+45^\circ$ with respect to the waveguide frame, the resulting polarization is RCP.

We wish to characterize the effect of the phase shifter components on the

⁵In practice the plate contributes an insertion phase to the parallel component as well, but we only require that the insertion phase *difference* between the two components is a quarter wave.

polarization of the incident wave. Express the incident field in the Cartesian basis: $\mathcal{E} = \mathcal{E}^X \hat{\mathbf{e}}_x + \mathcal{E}^Y \hat{\mathbf{e}}_y$, with the sign convention $\mathcal{E} \sim e^{i(\mathbf{k}\cdot\mathbf{x}-\omega t)}$. Subsume the temporal modulation ($e^{-i\omega t}$) into the field coefficients and neglect the $\hat{\mathbf{e}}_z$ term ($e^{ik_z z}$). We will consider a single cross polarized baseline; choose an LR baseline, which requires that receiver $j \rightarrow$ LCP and receiver $k \rightarrow$ RCP. The voltages at the correlator inputs, V_j^L and V_k^R , are related to the \mathcal{E} -fields at the receiver inputs by complex gain factors:

$$V_j^L = g_j \mathcal{E}_j; \quad V_k^R = g_k \mathcal{E}_k \quad (2.22)$$

where the \mathcal{E} -terms contain both LCP and RCP components. The gain terms do not affect the calculations that follow, so we will neglect them until the calibration discussion of Chapter 3. The effects of the receiver components on the input field \mathcal{E} can be expressed as a series of matrix operations: $\mathcal{E}' = \mathbf{m}_n \cdots \mathbf{m}_0 \mathcal{E} = \mathbf{M} \mathcal{E}$, and an understanding of the receivers permits the elements of the 2×2 transfer matrices \mathbf{m}_i to be derived. We expect to find that in the absence of errors, this visibility yields

$$\mathcal{V}^{LR} \sim \langle V_j^L V_k^{R*} \rangle \sim \langle \mathcal{E}'_j \mathcal{E}'_k{}^* \rangle \sim \langle \mathbf{M} \mathcal{E}_j (\mathbf{N} \mathcal{E}_k)^* \rangle \sim Q - iU \quad (2.23)$$

based on the standard representation of the Stokes parameters in terms of the \mathcal{E} -field coherence terms [92]. We will show that the half-wave plates change this result.

The calculation consists of several steps. We first derive \mathbf{M} , the transfer matrix for receiver j , and \mathbf{N} , the transfer matrix for receiver k ; we then compute the complex product of the modified fields to obtain the visibility \mathcal{V}^{LR} . Rotations form many of the building blocks for \mathbf{M} and \mathbf{N} ; to derive the transfer matrix for the half-wave plate, for example, we rotate into the frame of the half-wave plate, add a complex insertion loss to the component parallel to the plate, and then rotate back to the original frame. We repeat this procedure for the quarter-wave plate. The algebra which leads to the visibility is quite complex—in the presence of first order approximations to the errors discussed above, the elements of \mathbf{M} and \mathbf{N}

each have $\sim 10^2$ terms, so the visibility has $\sim 10^4$ terms; fortunately, *Mathematica* can plow through the calculation to a concise solution. Appendix A discusses the components of the calculation. The intermediate results, such as \mathbf{M} and \mathbf{N} , are far too cumbersome to reproduce there, so this section considers several simple test cases to verify that the model behaves as desired.

Appendix A introduces several terms to quantify the characteristics of the quarter and half-wave plates. The half-wave plate for receiver j has an insertion phase of $\alpha_j = \pi + \delta\alpha_j$, where $\delta\alpha$ represents the phase error. Similarly, the half-wave plate for receiver j is attenuated; we quantify this loss by scaling the input wave by $A_j = 1 - \delta A_j$. Both terms are real. The half-wave plate orientation is given by ϕ_j , which can take any value. The insertion errors for the quarter-wave plate are similar; they are given by β_j and B_j for the phase and loss, respectively. For proper performance, the quarter-wave plate must be oriented at $\pm 45^\circ$ with respect to the rectangular waveguide which follows it; in the presence of errors, $\theta_j = \pi/2 + \delta\theta_j$. We require a similar set of parameters for antenna k . Several simple test cases follow below.

1. Neglect the half-wave plates, and let the error terms vanish. This case tests our expectation that $LR \sim Q - iU$. We find

$$\mathcal{V}^{LR} = (Q - iU)e^{2i\psi} \quad (2.24)$$

This case yields exactly the result which we expect, with an additional factor of ψ which modulates the response. The argument of the phase term scales as twice the deck position ψ because a rotation which advances the phase seen by an *LCP* antenna by an amount η will retard the phase seen by an *RCP* antenna by the same amount, so the phase of the visibility doubles. Equivalently, if a wave produces pure $+Q$ radiation at a particular deck position, a 180° rotation of the deck will also produce pure $+Q$.

2. Neglect the half-wave plates, include the quarter-wave plates, but align the plates to the *same* angle (e.g., $\theta = -\pi/4$), so that both receivers measure the

same mode of circular polarization. In this case, we measure an LL visibility:

$$\mathcal{V}^{LL} = (I - V) \quad (2.25)$$

This is exactly the result that we expect.

- 3.** Repeat case 2., but include the half-wave plates:

$$\mathcal{V}^{LL} = (I + V)e^{-2i(\psi_1 - \psi_2)} \quad (2.26)$$

This result is a bit of a surprise: the visibility is proportional to \mathcal{V}^{RR} , not \mathcal{V}^{LL} , and it is modulated by a phase factor. The half-wave plates in the receivers cause this change; they convert an RCP component to LCP , and an LCP component to RCP . A simple example illustrates this result; suppose that the incident wave has a small RCP component, so that \mathcal{E}^X leads \mathcal{E}^Y . Insert a half-wave plate, and orient it to be parallel to $\hat{\mathbf{e}}_x$; the half-wave plate flips the $\hat{\mathbf{e}}_x$ component: $(\mathcal{E}^X, \mathcal{E}^Y) \Rightarrow (-\mathcal{E}^X, \mathcal{E}^Y)$, causing \mathcal{E}^X to lag \mathcal{E}^Y (LCP). This is the first indication that the half-wave plates change the outputs of the system, and this feature recurs in the results that follow. The half-wave plates do not affect the CBI's total intensity calibration because we expect $V = 0$ for all sources.

- 4.** Insert all the plates, assign the quarter-wave plates to their correct orientations, and ignore the errors to find

$$\mathcal{V}^{LR} = (Q + iU)e^{-2i(\psi + \phi_j + \phi_k)} \quad (2.27)$$

As in case 2, the half-wave plates map \mathcal{V}^{LR} to what we conventionally regard as \mathcal{V}^{RL} . Note that there is *always* an ambiguity between \mathcal{V}^{LR} and \mathcal{V}^{RL} ; the polarization assignments of the antennas are fixed, but for a given antenna pair with $(j, k) = (R, L)$, one can choose between assigning that baseline to LR by identifying the correlator output as \mathcal{V}^{LR} or to RL by identify-

ing the correlator output as \mathcal{V}^{RL} . The critical ingredient is that the same assignments—the same complex gain factors—are used for calibrations and for program observations.

5. For the purpose of illuminating the behavior of the plates in the presence of errors, we will briefly assume that two half-wave plate orientations are the same: $\phi_k \rightarrow \phi_j$. Insert both the quarter and half-wave plates, and include errors to find

$$\begin{aligned}
\mathcal{V}^{LR} = & \quad I/2 \quad \left[-(\Sigma A + i\Delta\alpha)e^{-2i\phi_j} + i\Delta B - \Sigma\beta + 2i\Sigma\theta \right] \\
& \quad Q/2 \quad \left[2 - \Sigma A - \Sigma B - i\Delta\alpha - i\Delta\beta + 2i\Sigma\theta \right] e^{-2i\psi - 4i\phi_j} \\
& \quad iU/2 \quad \left[2 - \Sigma A - \Sigma B - i\Delta\alpha - i\Delta\beta + 2i\Sigma\theta \right] e^{-2i\psi - 4i\phi_j} \\
& \quad V/2 \quad \left[-(\Delta A + i\Sigma\alpha)e^{-2i\phi_j} + i\Sigma B - \Delta\beta - 2i\Delta\theta \right]
\end{aligned} \tag{2.28}$$

where $\Sigma A = \delta A_j + \delta A_k$, $\Delta\alpha = \delta\alpha_k - \delta\alpha_j$, etc.

We can simplify this expression with some assumptions about the manufacturing and assembly tolerances for the phase shifter. The wave plates were made from a single batch of teflon, so the insertion phase and the insertion loss per unit volume should be uniform from part to part, and only variations in the dimensions of the plates will produce variations in the electrical performance. The machining tolerances of the quarter and half-wave plates are on the order of thousandths of an inch, which corresponds to a few tenths of a degree of insertion phase. Since the bandpass errors of a few degrees dominate these machining errors, we may assume that the insertion phase errors repeat with a high degree of accuracy from unit to unit: $\delta\alpha_j \sim \delta\alpha_k$ for the half-wave plates, and $\delta\beta_j \sim \delta\beta_k$ for the quarter-wave plates. Similarly, since the entire length of the half-wave plate attenuates the signal by only $\sim 5\%$, variations in loss from machining errors may be neglected: $\delta A_j \sim \delta A_k$, and $\delta B_j \sim \delta B_k$. We can therefore delete all of the Δ terms in the visibility which arise from the plate transfer matrices. Benchtop tests of the plates show that $\delta A \cong 2\delta B$, and $\delta\alpha \cong 2\delta\beta$. Finally, the orientation errors for the half-wave

plate ϕ and the quarter-wave plate θ are random from unit to unit, so these terms do not cancel. Under these assumptions, Equation 2.28 becomes

$$\begin{aligned} \mathcal{V}^{LR} \cong & I \left[-\delta A e^{-2i\phi_j} - \delta\beta + i\Sigma\theta \right] \\ & + (Q + iU) \left[1 - \frac{3}{2}\delta B - i\Sigma\theta \right] e^{-2i\psi - 4i\phi_j} \end{aligned} \quad (2.29)$$

to first order in the errors. Since most astronomical sources have vanishingly small circular polarization, we have neglected the V term in Equation 2.28.

Equation 2.29 has several interesting features. First, the half-wave plate orientation ϕ modulates the contribution to the instrumental polarization from the insertion loss of the half-wave plates. Since $\delta A \sim \delta\beta \sim 5\%$, this effect can be significant: if we had rotated the half-wave plate throughout a series of observations of the polarization calibrator, each observation would be contaminated by a different amount of instrumental polarization. During the CBI observations, however, the half-wave plate orientation remained constant, so this dependence is not a concern for the analysis of the data. The origin of this contribution comes as no surprise; the half-wave plate attenuates radiation parallel to the plate, so unpolarized radiation acquires a linear polarization in the orthogonal direction, the magnitude of which is proportional to the loss in the plate. Of course, the orientation of the plate modulates the direction of the spurious linear polarization, so this component varies with ϕ . Second, Equation 2.29 shows that the orientation errors of the quarter-wave plates can cancel for the correct combination of orientation errors: $\Sigma\theta = \delta\theta_k - \delta\theta_j = 0$. This result implies that under these contrived circumstances, the orientation of the plates relative to the rectangular guide does not matter—even if the error is large, the correct combination of offsets can cancel under the first order approximation. Of course, it is unlikely that the orientation errors for the quarter-wave plates for all 13 antennas will combine to produce cancellation for all baselines. Finally, all of the errors in this expression are on the order of 5%, so a net instrumental polarization of

$\sim 10\%$ is not beyond reach.

6. Let $\phi_j \neq \phi_k$. After making the assumptions described in the preceding section, the general expression for the leakage has the form

$$\begin{aligned} \mathcal{V}^{LR} = & I \left[(-\delta A \cos[\Delta\phi] + \delta\alpha \sin[\Delta\phi])e^{-i\Sigma\phi} \right. \\ & -\delta B \sin[2\Delta\phi] - \delta\beta \cos[2\Delta\phi] \\ & \left. + i\Sigma\theta \cos[2\Delta\phi] + \Delta\theta \sin[2\Delta\phi] \right] \\ & + (Q + iU) \left[1 - \delta A - \delta B + i\Sigma\theta \right] e^{-2i(\psi + \phi_j + \phi_k)} \end{aligned} \quad (2.30)$$

The calibration procedure assumes that all spurious effects from the phase shifter appear as leakage, so to compare Equation 2.30 with the measured leakage, we first divide the expression by the coefficient of $(Q + iU)$. We will revisit this expression in Chapter 3 when we discuss measurements of the CBI's large instrumental polarization. Note that the coefficients for Q and U are the same; were the coefficients to diverge, the leakage model for the instrumental polarization—the model which forms the basis for the calibration procedure—would fail.

2.3.3 Performance of the CBI

The CBI is a complex instrument, but its critical performance characteristics depend on a few key factors. The sensitivity determines the time required to measure the temperature and polarization signatures of the CMBR. The stability of the system affects the sensitivity because noiselike gain variations augment the thermal noise in the visibilities, and systematic gain variations can bias the results. Pointing errors can bias the phases. External sources of emission, such as spillover, the sun, and the moon, can corrupt the visibilities, and any restrictions which these sources place on our observational duty cycle degrades our effective sensitivity. This section focuses on these factors.

source	T (K)
CMBR	3
atmosphere	1
ground spillover	2
optics	2
HEMT	18
cold downconverter	1
total	27

Table 2.1: CBI noise budget.

Sensitivity

The time τ required to achieve a particular noise level with a single baseline depends on T_0 , the system temperature; and $\Delta\nu$, the bandwidth:

$$\sigma_{\nu}(\text{Jy}) = \frac{\sqrt{2}k_B T_{\text{sys}}}{A_{\text{eff}}\eta\sqrt{\Delta\nu\tau}} \quad (2.31)$$

where A_{eff} ($= 4860 \text{ cm}^2$) is the effective antenna collecting area and η ($= 0.86$) is the correlator efficiency. The system temperature arises from all of the sources which deposit power in the system, and since it forms the backdrop against which we measure fluctuations in the CMBR, a key design goal is to make T_{sys} as small as possible. While the contributions to T_{sys} can change, the bandwidth of the system is fixed by the bandwidth of the HEMT amplifiers at the receiver inputs, and the filters which follow are selected to match the input bandpass. It is important to establish that the known sources of T_{sys} agree with the noise in the data. Table 2.3.3 provides an *a priori* noise budget.

Several external sources contribute to T_{sys} . The CMB monopole contributes $\sim 3 \text{ K}$ regardless of the telescope orientation. Atmospheric emission contributes an additional $\sim 1 \text{ K}$ which depends on elevation; NRAO's site testing data suggest low opacities ($\tau \sim 0.004$) [33] and these provided one of the motivations for siting

the CBI on the Llano de Chajnantor. In contrast, ground spillover is source of loading at all terrestrial sites; the far sidelobes of the primary beam illuminate the ground, so ground spillover contributes another elevation dependent source of emission to the array.⁶

The remaining sources of noise are found within the system. Benchtop tests show that the optics which precede the receiver—the primary, the secondary, and the antenna shields—contribute ~ 1 K, while the phase shifter assembly contributes another ~ 1 K to the total power. The low-noise HEMT amplifiers dominate the noise budget. As we see in Section 2.4, the HEMT noise across the 26-36 GHz band is meniscus-shaped; the noise has a broad minimum of ~ 15 K and rises sharply to more than 20 K at the band edges. The noise given in the table represents the 26-36 GHz average for all of the receivers. The cold downconverter is the first active component which follows the HEMT; it has $T_0 \sim 10^3$ K, but the ~ 30 dB of gain in the HEMT suppresses its contribution to ~ 1 K. A cold low-noise amplifier which follows provides sufficient gain to reduce the noise from all of the downstream components to negligible levels. The external and internal noise sources combine to produce $T_0 \sim 27$ K at the antenna inputs.

Several factors diminish the sensitivity given by Equation 2.23. The notch filters at the output of the downconverter (Section 2.3.1) reduce the effective bandwidth of a single channel to 900 MHz, and the differencing strategy adopted to eliminate ground spillover imposes a penalty of $\sim \sqrt{2}$.⁷ Taken together, these considerations imply a sensitivity of 4.7 Jy in 1 s. In Chapter 5, we explore the noise in detail; at that point we report that the mean noise for the *LR* visibilities is ~ 6.1 Jy, while that for the *LL* baselines is ~ 5.8 Jy. Both values exceed the expected noise by $\sim 20\%$. The excess for *LL* is not understood, and the additional 5-10% excess for *LR* is a consequence of higher noise in RX12 coupled with an amplitude calibration error (Chapter 5).

⁶Ground emission can scatter from the insides of the antenna shields into the feeds as well.

⁷In practice we neglect the $\sqrt{2}$ and simply recompute the effective integration time to reflect the penalty from differencing.

Pointing

The CBI's pointing affects the accuracy with which we can reconstruct the phases of the visibilities. For total intensity visibilities, phase errors affect the inferred position of the flux on the sky, while for cross polarized visibilities, phase errors can corrupt the inferred polarization position angle as well. Consider an unpolarized point source: let $I(\mathbf{x}) = S_0\delta(\mathbf{x}_0)$ in Equation 2.1 and assume that the pointing error $\delta_{\mathbf{x}}$ is small ($A(\mathbf{x} - \mathbf{x}_0) \sim 1$) to obtain $\mathcal{V}^{LL}(\mathbf{u}) = S_0e^{2\pi i\mathbf{u}\cdot\mathbf{x}_0}$. In the small error approximation, pointing errors do not affect the amplitude, but they do affect the phase of the visibility ($\delta_{\mathbf{x}}\mathcal{V}_{\phi} = 2\pi u\delta x$). A pointing error of $20''$ produces a phase error of $\sim 20^\circ$ on a 500 cm baseline, for example. Now suppose the source is polarized; in this case $P(\mathbf{x}) = mI(\mathbf{x})e^{2i\chi}$, where m is the fractional polarization, and $\mathcal{V}^{LR}(\mathbf{u}) = mS_0e^{2\pi i\mathbf{u}\cdot\mathbf{x}_0+2i\chi}$. Since the phase of the cross polarized visibility conveys information about the source's position angle, a pointing error can transfer power between Q and U .

Pointing errors are a concern for the CBI observations, but we have implemented strategies to mitigate their effect on the data. To minimize its shipping weight, the CBI was designed to operate without a counterweight about the elevation axis. The unbalanced mount introduces an elevation dependence to the pointing errors; the raw pointing offsets can be as high as $40''$. Through a combination of enhanced telescope telemetry in parallel with improvements in the pointing model, however, the pointing rms was reduced to $\sim 20''$, which is well within the limit required to obtain a good phase calibration for the polarization observations. The orientation of the deck about the pointing axis is an additional dimension of the pointing which can have a deleterious effect on the polarization observations. The deck position measurements are stable and accurate to better than 1° , however, so these errors are well below the uncertainties in the position angle data from the VLA observations (Chapter 4).

Ground Spillover

Ground spillover is far more insidious than its 1 K contribution to the noise budget suggests. The ground near the CBI contributes a uniform ~ 300 K, but while these patches of ground are a significant source of total power, the corrugations which modulate the antenna sidelobes cancel their contributions to the visibilities. At the horizon, however, the discontinuity between the ~ 300 K ground and the ~ 3 K sky is an ideal source of structure, and while the CBI's elevation limit of 45° keeps the bulk of the power in the primary beam away from the horizon, the far sidelobes are sufficiently high to permit ground spillover signals as large as a few Janskys. This contamination can dominate the celestial contributions to the visibilities, so it must be eliminated.

While the spillover signature changes from visibility to visibility, the contamination conforms to a number of global characteristics. The amplitude of the spillover falls sharply with increasing baseline length because longer baselines fit more fringes into a fixed area on the horizon, thereby improving the cancellation—in fact, we do not detect spillover on baselines longer than 104 cm. Similarly, since the CBI has $\sim 30\%$ bandwidth, the spillover is a strong function of frequency: the channels near the high end of the CBI band squeeze $\sim 30\%$ more fringes into a given patch than do the low frequency channels, resulting in a discernible enhancement in the cancellation.⁸ Spillover is also a strong function of the orientation of the fringes relative to the horizon; the contamination peaks when the fringes are parallel to the horizon and it decreases rapidly as the the fringes rotate through 90° . This feature can be demonstrated by simply rotating the deck while viewing a fixed point above the horizon. The observation that the spillover corrupts the cross polarized baselines yielded an important insight about the paths through which spillover enters the system; since the ground is unpolarized, the cross polarized visibilities should be immune to spillover. In fact, we find that spillover contaminates the total intensity visibilities and cross polarized visibilities with nearly

⁸The change in beamsize with frequency does not affect this conclusion because we are concerned with the far sidelobes.

equal force, suggesting that the spillover enters the system along a path which involves some reflections, and the obvious candidate is scattering from the insides of the antenna shields. The fact that spillover persists as the boresight approaches the zenith—well beyond the elevation at which the dishes have a direct view of the horizon—supports this hypothesis.

Contamination from spillover with \sim Jy amplitudes must be addressed before any information about the CMBR can be extracted from the visibilities. The snow and high winds at the CBI site precluded the installation of a ground shield of the necessary size, so we modified our observing technique to reject the spillover. We adopted a conservative subtraction strategy to cancel the spillover from the visibilities: we observed the deep fields in pairs that were separated by 8^m in right ascension, and then differenced the pairs of matching visibilities offline. To a very good approximation the temperature of the ground remains constant over the few minutes which separate the matching samples in the lead and the trail, so the constant spillover contribution cancels from the differenced visibilities. All tests of the efficacy of this technique suggest that it performs as desired, but it comes at a price: differencing requires that we effectively spend half of our time measuring the ground, so it imposes a penalty of $\sqrt{2}$ on our sensitivity.

2.4 Instrumentation Projects Specific to this Thesis

We now turn to several hardware projects which constitute a significant component of this thesis: the downconverter, the noise calibration system, and the low-noise HEMT amplifiers. The benefits of this work extended beyond the CBI: DASI implemented the CBI designs for the downconverter and the noise calibration system, and two of the HEMT amplifiers were installed in the Owens Valley 40-meter telescope for the point source monitoring campaign.

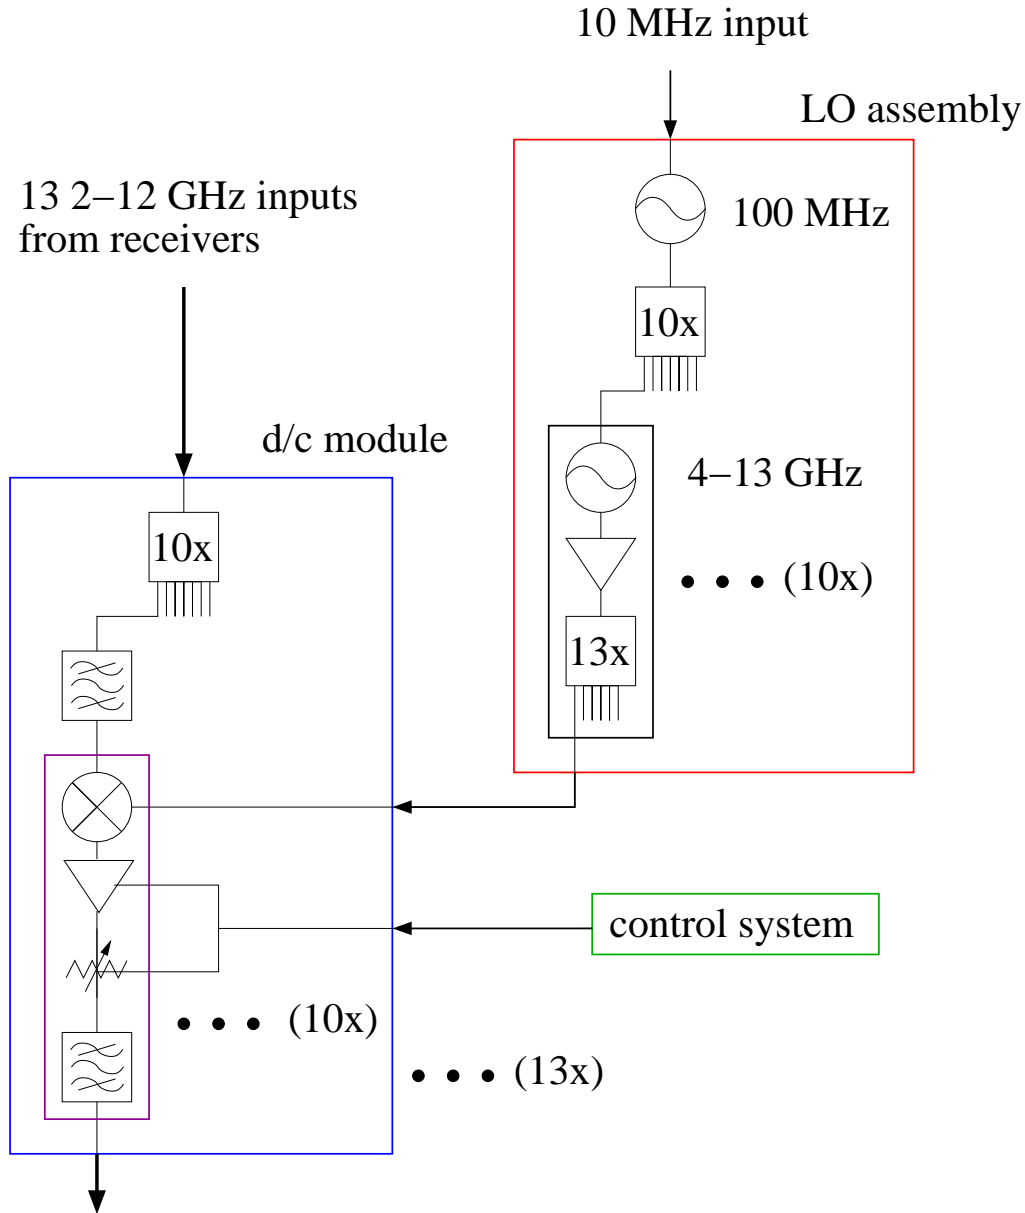
2.4.1 Downconverter

The CBI downconverter performs two functions: it converts the 2-12 GHz output from each of the 13 CBI receivers into ten 1-2 GHz (L-band) signals, and it sets the power level of each of the 130 L-band channels to the +16 dBm required by the correlator. The downconverter must accommodate a variety of inputs; the receivers outputs can have tilts across their 2-12 GHz band which are as large as ~ 10 dB, and have uniform offsets of several dB. The downconverter must fit in the small volume of a standard $17 \times 21 \times 21$ in³ electronics crate, and the packaging and cooling challenges resulting from this requirement contribute a substantial degree of complexity to the design. The downconverter is a major subsystem of the CBI; it required ~ 60 mechanical drawings, ~ 700 machined parts, ~ 600 circuit boards, and costs \sim \$400k. This section describes the downconverter and its performance.

Figure 2.4 provides an overview of the downconverter, and Figure 2.5 shows a picture of the downconverter installed on the underside of the CBI deck adjacent to the correlator. The downconverter is divided vertically into three regions, by function: the power supplies are on top; the local oscillators are in the middle; and the downconverter modules are on the bottom. An aluminum plate between the local oscillators (LO) and the modules provides an electrical, thermal, and mechanical interface for the LOs and the modules. Each downconverter module attaches to the underside of the plate with a single pair of bolts—this simple mounting scheme allows the modules to be installed and removed quickly. The plate contains quick release *rf* connectors for the 130 LO inputs, as well as a custom backplane which supplies power and control to the 13 modules.

The downconverter operates in conjunction with the correlator and the control system. Power levels measured by detectors in the correlator are used by the control system to set the attenuators in the downconverter to provide +16 dBm at the correlator input. The downconverter cannot eliminate amplitude slopes across a single $\Delta\nu \sim 1$ GHz band, although it can set the power levels of all of the bands to +16 dBm with 1 dB precision. The process of setting the gains to the correct level takes seconds; the procedure is performed at the beginning of each observing

CBI Downconverter Overview



10x13 1-2 GHz outputs to correlator

Figure 2.4: CBI downconverter overview. Many of the downconverter components come in multiple quantities; there are 10 frequency synthesizer blocks in the LO assembly, 10 L-band channels in each downconverter module, and 13 downconverter modules. The LO assembly requires a 10 MHz source to set the frequencies of the ten LO outputs.




Figure 2.5: The CBI downconverter and correlator installed on the array.

session, and the settings are left unchanged for the duration of the session. A comparison of attenuator settings across periods spanning several weeks provides a crude measure of the stability of the system; we should expect that the gains remain constant from night to night to within the 1 dB ($\sim 25\%$) resolution of the system. Inspection of the attenuator settings across these timescales shows that many of the gains remain constant at this level for these periods.

Downconverter Module

Each receiver is assigned to one of the thirteen identical downconverter modules. Each module consists of a filter bank, ten downconverter circuits, and a control board. Figure 2.6 provides a circuit diagram, and Figures 2.7 and 2.8 provide two views of a single module. The receiver signal enters at the filter bank input, and the ten module outputs pass directly to the correlator inputs. Each module dissipates ~ 50 W, and since the dense packing of components precludes the use of forced air for cooling, the bottom of the module contains a cooling loop which connects to the chilled water supply for the downconverter through a pair of self-sealing quick release connectors.

The filter bank at the module input splits the 2-12 GHz receiver output into ten bands with $\Delta\nu \sim 1$ GHz. The filter bank is a custom part, designed to satisfy the insertion characteristics required for good correlator performance. In addition, the rejection must increase rapidly outside of the desired band; we chose a rejection limit of 20 dB at points 100 MHz beyond the band edges. To ensure uniform bandpass characteristics from unit to unit, the vendor first made a master filter bank and then adjusted the insertion characteristics of the production filters to match the master unit.

The downconversion schematic is shown in Figure 2.6. The ten filter bank outputs are centered on different frequencies, so a mixer at the input of each channel combines one of ten LO signals with each filter output to mix the power to L-band. The circuits which follow the mixers consist of a series of amplifiers, filters, and a single variable attenuator. The amplifiers provide ~ 65 dB of gain, and the

CBI Downconverter L-band Circuit

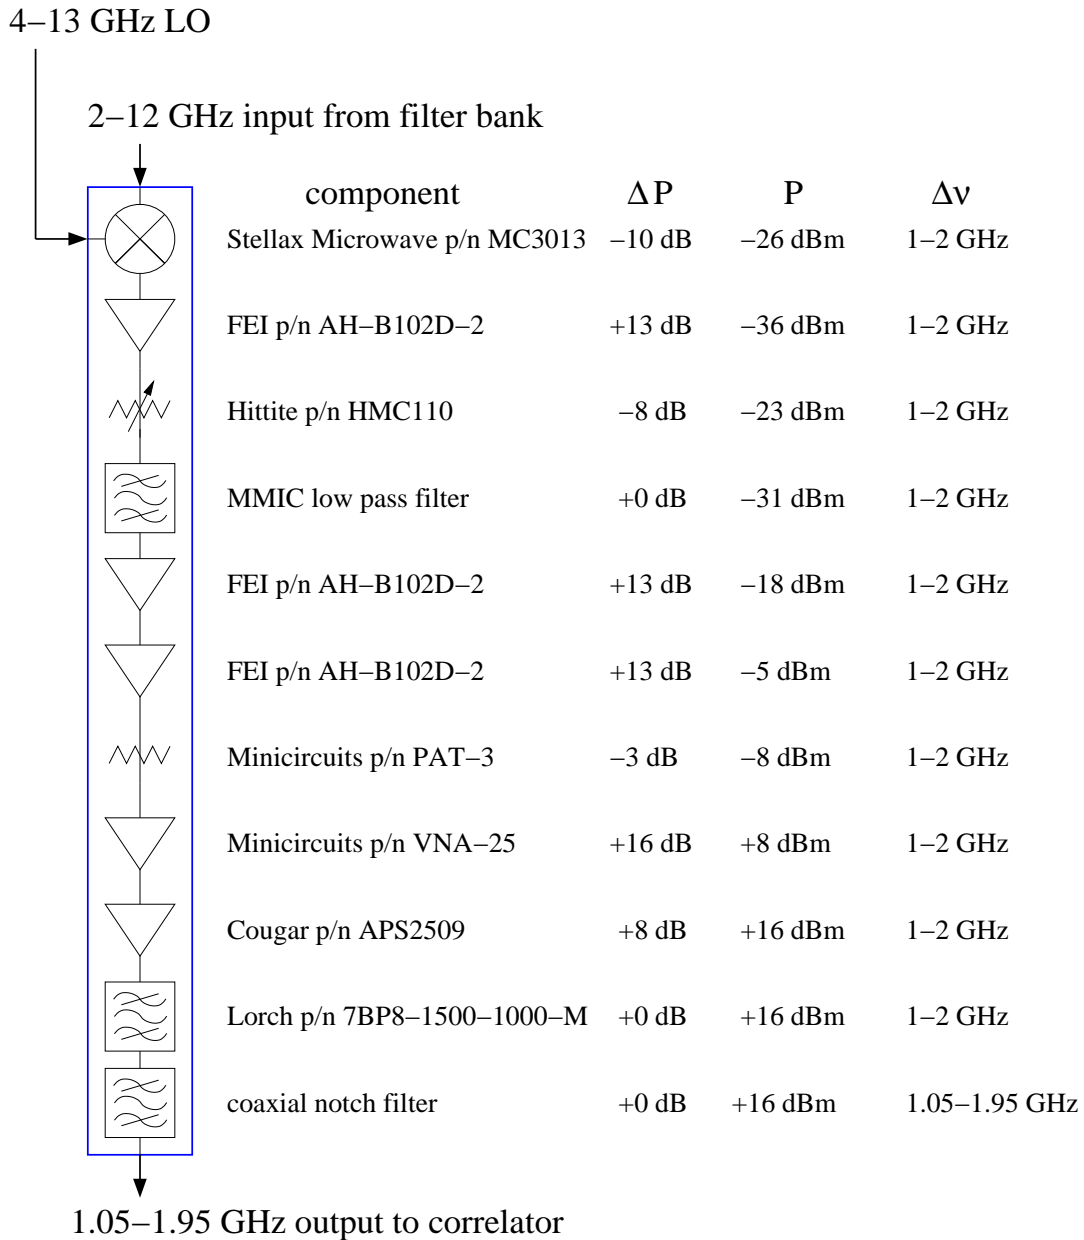


Figure 2.6: L-band schematic. Each module contains ten L-band circuits, so there are 130 of these circuits in the downconverter.




Figure 2.7: *rf* side of a downconverter module.




Figure 2.8: *DC* side of a downconverter module.

variable attenuator provides up to 31 dB of attenuation with 1 dB resolution. The gain in the downconverter channels is such that the correlator typically requires ~ 8 dB of attenuation, although receivers roll the 35-36 GHz band off to such a degree that the 35-36 GHz channel requires an additional stage of amplification on all modules to avoid starving the correlator. The L-band circuits employ a combination of surface mount and hybrid components, all mounted on thin alumina (TMM10) boards. Alumina is fragile, so each circuit board is glued to a 1/16" thick Al plate to provide the mechanical strength required for installation and removal. Each L-band circuit occupies a machined slot in the module which isolates the circuit from its neighbors.

The L-band circuits receive instructions and power from a control card mounted on the opposite side of the module. The control board mates with the backplane with a set of D connectors, and it connects to the L-band circuits with a set of quick-release feedthrough connections which are capacitively coupled to ground. The heart of the control circuit is a field-programmable gate array which decodes and implements the commands sent to the 13 modules over the downconverter backplane. The FPGA is a device which can be programmed to perform a wide variety of logic functions; the user designs and compiles the logic for the FPGA with Altera's Max+II software, and then writes this command set to a memory chip which is installed with the FPGA on the control board. The FPGA is designed to set the attenuators to a benign low power state on startup to avoid saturating the correlator before the attenuators can be set. The cards for all 13 receivers are identical, and a set of jumpers on each card facilitates a unique mapping between each receiver and each module.

The downconverter LO system radiates power which can enter the receivers at their feeds and corrupt the visibilities. These signals arise from the harmonics of the LO signals which fall in the 26-36 GHz band, and they appear at the downconverter output as spikes at the edges of the output bandpass—the harmonics produce spikes at exactly 1 and 2 GHz, with amplitudes as high as ~ 20 dB. To reject these signals, a notch filter was added to the downconverter output; it con-

sists of a pair of stubs in series whose lengths are selected to short signals at 1 GHz and its harmonics (both stubs are $\lambda/2$ long). The filter is made from coiled segments of semi-rigid coaxial cable, and it attaches to the outputs of each of the 130 downconverter channels. The filter required some tuning by hand to obtain at least 40 dB of rejection at exactly 1 and 2 GHz. The filters decrease the effective bandwidth of the receivers by 100 MHz, or 10%, which in turn degrades the sensitivity by 5%.

Local Oscillator Assembly

The LO assembly provides the ten LO signals required to map the ten filter bank outputs to L-band. The LO system consists of ten frequency synthesizer modules and a distribution network which routes each LO to the 13 downconverter modules. Figure 2.9 shows a single frequency synthesizer module, and Figure 2.10 shows the entire LO assembly. The synthesizer module consists of a frequency source, power splitters, and amplifiers which raise the LO power to the +13 dBm required by the mixers in the downconverter. Each frequency synthesizer requires a 100 MHz reference, which is supplied by a single 100 MHz synthesizer located in the LO assembly. The 100 MHz source in turn requires a 10 MHz reference, which is supplied by the CBI.

Power Supply, Cooling, and LO Plate

The power supply assembly provides the -5V, 5V, and 15V levels required by the amplifiers and synthesizers in the downconverter. The supplies consume ~ 1 kW of power, over half of which is dissipated as heat in the downconverter's components. Chilled water from the CBI supply circulates through the unit; a pair of manifolds at the back of the downconverter, seen at the far right in Figure 2.10, route the water to the LO/module interface plate and to each of the 13 downconverter modules. The plumbing connections are made with self-sealing quick release connectors which allow the modules to be removed without disabling the entire cooling network.




Figure 2.9: 4 GHz frequency synthesizer block.

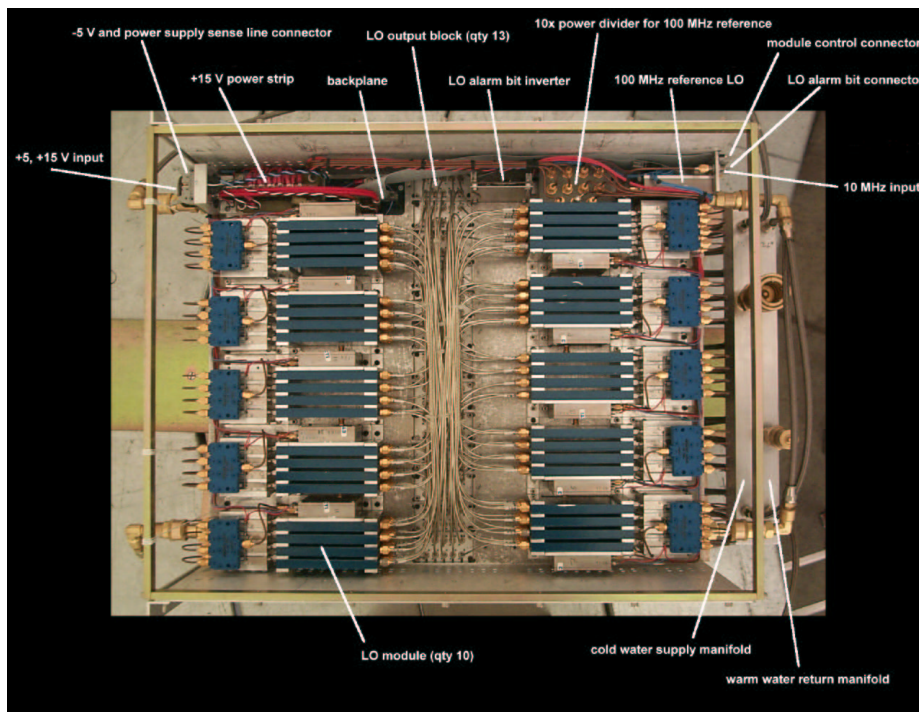


Figure 2.10: Top view of the local oscillator assembly; all ten frequency synthesizer blocks shown.

Downconverter Performance

The performance requirements for the downconverter are simple: it must deliver 130 receiver signals with the appropriate power and bandpass to the correlator. The loss of a channel, band, or module of the downconverter decreases the amount of data through the system, which in turn degrades the CBI's sensitivity, so the downconverter must also have a high duty cycle. These requirements are difficult to achieve, however, because the downconverter is a complex part; it has nearly a thousand active *rf* components which are susceptible to electrostatic damage and thousands of wirebonds which can be compromised by thermal and mechanical stresses. In light of these concerns, the downconverter was designed to permit easy access to the most fragile components for repair, but as we will see, the components which caused persistent problems were among the least accessible. We lost few nights to downconverter problems, but there were long periods during which we operated with one or several damaged channels.

Variable Attenuator

The insertion loss of the variable attenuator was usually less than the amount commanded by the system; the measured insertion loss could be as little as 27 dB, for example, for a requested value of 31 dB.⁹ This is not a problem, however, because the channels typically require ~ 10 dB of attenuation. Since the control system simply searches for the combination of attenuator settings which set the power at the correlator input to +16 dBm, the failure of the attenuator to insert the required loss does not impair the performance of the system.

Downconverter Bandpass

The shape of the downconverter bandpass is the consequence of the individual sources of gain and loss in the filter bank, mixer, and L-band circuit. High correlator efficiency requires that the amplitude fluctuations in the circuits be as small as

⁹This may be a consequence of parasitic coupling on the attenuator.

possible, and 2 dB ripple p-p was selected to be the design specification. Concerns about flatness drove many of the component choices in the L-band circuits. The Minicircuits amplifiers, for example, cost only a few dollars per unit, but since they have a 2 dB slope across the 1-2 GHz band, we cannot use more than one in the circuit. The more expensive FEI amplifiers have far better flatness characteristics, and these devices provided most of the gain. Generally, although all of the components are broadband in the sense that they have no sharp features across the 1-2 GHz band, the net result of the sum of the components and all of their small matching errors is to create variations on the order of 1-2 dB across the band. All of the L-band circuits were tested for flatness at 1.0, 1.1, 1.5, 1.9, and 2.0 GHz before they were accepted for use on the array, and none were installed if the flatness exceeded 2 dB p-p in the central 900 MHz of the band. Bandpass errors which exceeded this level were often the result of poor assembly, and were thus easily rectified. Larger variations at the band edges were tolerated because the notch filters suppress power at those points.

Gain Stability

The amplifiers in the L-band circuits provide ~ 65 dB of gain, and the small volume in which this gain is developed fomented oscillations in some circuits. One type of oscillation arose between the FEI and the Minicircuits amplifiers, and it was eliminated by a single 3 dB surface mount attenuator between these components (Figure 2.6). In other cases the power radiated by the circuit gave rise to instabilities, and a strip of Eccosorb in each slot eliminated the resulting oscillations.

Gain Compression

The gain of the L-band circuits must remain linear over the range of likely input power. The gain departs from linearity when an amplifier ceases to provide the output power necessary for a given level of input power; when this occurs, the amplifier has *compressed*. The L-band circuits were designed to compress well above

the +16 dBm required at the output—the high power Cougar amplifier provides the last stage of amplification. Direct measurements of the L-band gain with a signal source show that the power compresses by 1 dB when $P_{out} \sim 27$ dBm, which is consistent with the expected performance from the last stage amplifier. This result agrees with system temperature measurements on the array that show $\sim 3\%$ compression for load temperatures of ~ 300 K. Since the antenna temperature of ~ 30 K is 10 dB below the 1 dB compression point, the standard operating regime is well below the range for which compression is a consideration.

Duty Cycle

Failures in the downconverter can degrade the CBI's sensitivity through lost bandwidth and observing time. Since the downconverter modules contain hundreds of active components and thousands of wirebonds, they were designed to permit quick removal for repair. Failures in the LO distribution network were not anticipated, and these required disassembly of the downconverter. Nonetheless, problems with the downconverter resulted in only a marginal loss of observing time.

During the first year of operations the downconverter suffered from several failures, most of which were confined to the LO assembly. First, several frequency synthesizers and amplifiers in the LO assembly failed, so these parts were returned to the vendor for warranty repairs. Second, the spring-loaded coaxial connectors with which the LO assembly mates with the modules were a source of recurring difficulties: six of the connectors failed, thereby starving the mixers of power and disabling the associated channels.¹⁰ These connectors are buried deep within the downconverter, so for these repairs the downconverter was removed from the array and disassembled. These connectors were a poor choice for this particular application. These failures did not result in a significant loss in observing time.

¹⁰This type of connection is inferior to one in which the inner conductor is *soldered* in place. The manufacturer's suggested technique for assembling these connectors leaves the connector prone to failure; the user solders the jacket of the coaxial cable to shell of the connector, and the inner conductor of the cable mates with the inner conductor of the connector with a simple mechanical (friction) fit.

Summary

The downconverter, a key component of the CBI signal chain, was designed and implemented for this thesis. It meets the size and performance requirements for the CBI, although on a few occasions connector failures in the LO distribution network precipitated major repairs. Very few nights were lost to downconverter failures.

2.4.2 Noise Calibration System

The active components in the CBI signal chain can succumb to gain drifts which introduce errors in the calibration over the course of an observing session. Gain variations come in two flavors; if the variations are incoherent from baseline to baseline, they give rise to scatter which augments the noise introduced by T_{sys} . If the variations are coherent, they produce systematic errors in the measured visibilities. In either case, good performance requires that these gain variations are measured and eliminated. The CBI employs a system which injects a bright calibration signal into each receiver; the resulting visibilities are recorded in the data stream, and the offline calibration software uses these calibration source visibilities to correct the program data. During routine observations, the calibration source is fired several times an hour, and the calibration software uses the source visibilities to interpolate a correction for the intervening visibilities. The correlator gain calibration (Section 3.2.2) also requires the internal noise source. This section describes this calibration system.

The noise calibration system is very simple: it consists of a noise source, a power meter, and a distribution network (Figures 2.11 and 2.12). A series of amplifiers combine their intrinsic noise with that of a load to produce a bright signal at the receiver inputs; the amplifier noise dominates the load, so the noise source stability tracks that of the amplifiers, which are thermally coupled to the site chilled water supply. Since the amplitude of the noise source can vary with time, a power meter measures the noise and reports these values to the data stream, and

the noise source visibilities are corrected for these variations before they are used to interpolate the gain variations in the program data. The noise signal passes to a 16-way power splitter and a series of switches. The power splitter was quickly and inexpensively realized from a block of aluminum; it consists of a sequence of four consecutive two-way splittings (Figure 2.13), and this design provides good temperature stability. The shape of each division is based on a tangent curve, which ensures that the derivatives at the boundaries between successive divisions—*i.e.*, the shape of the waveguide—agree. The outputs of the splittings are not matched. The splitter behaves as desired; the power out of all the ports is $\sim 1/16$ of that at the input, and the input return loss—more than 20 dB—is well behaved across the band.

The amplifier gain ramps up sharply when the amplifiers turn on, so the amplifiers run continuously during observations. We therefore require a considerable degree of isolation between the amplifiers and the receivers, and two solid state waveguide switches isolate each receiver from the source when it is not in use; one switch is at the input of the 16-way power splitter, and the others are at 13 of the 16 outputs. Together the switches provide 80 dB of isolation which ensures that the correlator does not see the noise source during celestial observations, nor that the receivers couple to each other through the noise calibration system. The placement of the switches also provides some control over the choice of receivers that see the source.

The 13 noise cal outputs are distributed to the receivers via coaxial cable. The insertion characteristics of coaxial cable vary with temperature, so to maintain a stable temperature the cables were enclosed in an insulated shroud with cooling water which was fanned out to the receivers. In addition, a cooling manifold attached to the amplifiers maintains the temperature of the active components to within several degrees Celsius. Despite these efforts to control the temperature of the noise cal, the system succumbed to temperature-dependent drifts which complicated the application of the noise source visibilities. This problem is discussed in the following section.

CBI Noise Calibration System

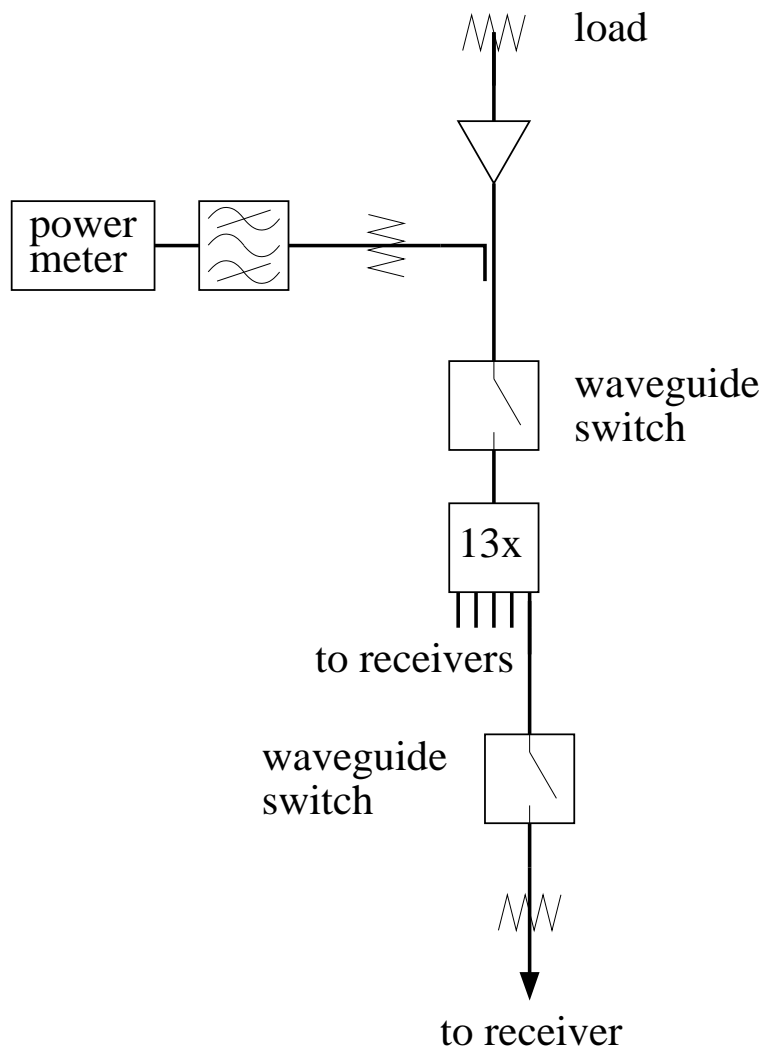


Figure 2.11: Schematic of the CBI noise calibration system.

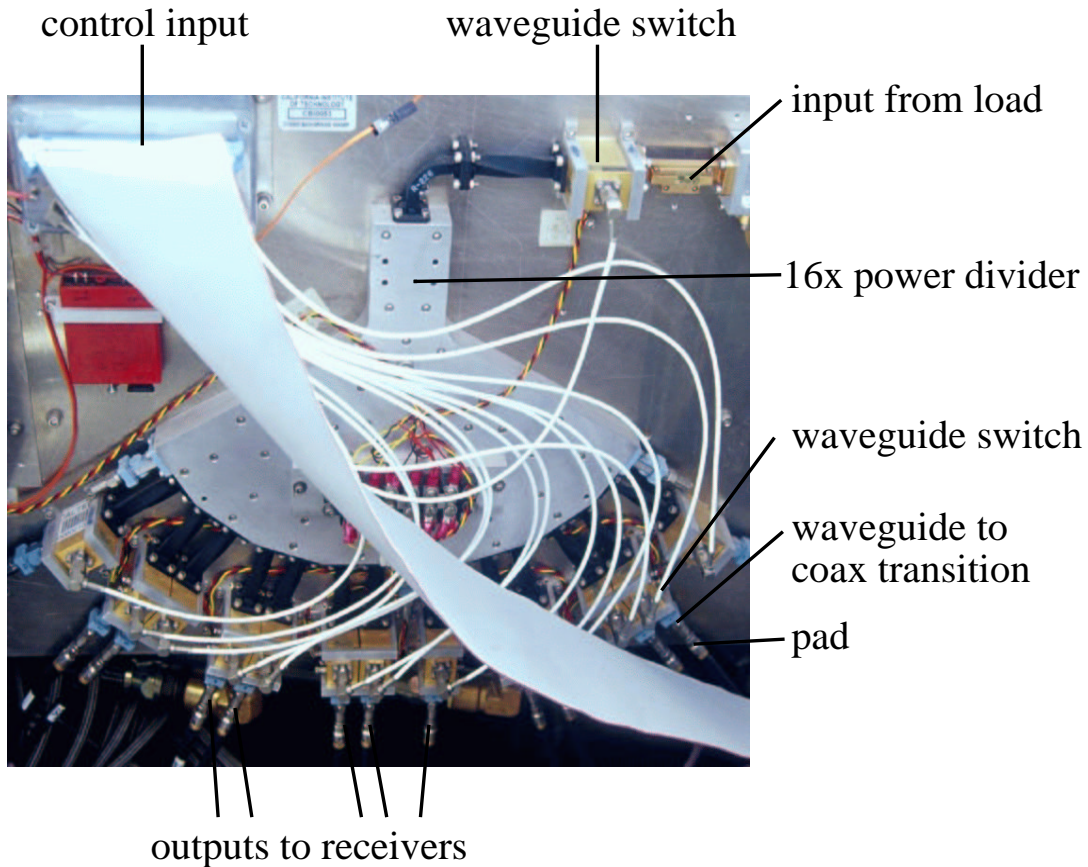


Figure 2.12: Image of the noise cal system. The load, amplifiers, and power meter are located on the other side of the plate on which the noise calibration system resides.

Noise Source Performance

Figure 2.14 provides an example of the correction that the noise source applies to real data. The night of 10mar00 was bracketed by two total intensity calibrators which together provide a check on the efficacy of the noise calibration procedure; Tau A was observed for 5^m at the beginning of the night, and 3C274 for 5^m at the end. Gain drifts which occur over the course of the night will produce flux errors for one source as calibrated on the other. For this demonstration, the flux calibration was obtained from Tau A, and this scale was transferred to the 3C274 observation at the end of the night. Figure 2.14 shows the results of this test; the




Figure 2.13: Autocad rendering of the 16-way waveguide power divider.

heavy solid line is the expected flux for 3C274, and the dashed line shows the flux obtained with the noise source correction. In the absence of a correction to the gain variations the derived flux differs from the expected values by ~ 0.8 Jy, or 5%. The noise source correction improves the error to less than 0.3 Jy, or $\sim 2\%$.

Although this example provides a demonstration of the efficacy of the noise cal, the performance of the system often fell short of desired specifications. First, on some occasions the variations in the noise cal visibilities associated with individual receivers exceeded the variations in the underlying receiver gains. An inspection of the raw visibilities shows that on rare occasions the noise source amplitude can vary by as much as 10% on timescales of an hour, and under these circumstances

Comparison of 3C274 LL flux densities, two calibration procedures, 10mar00

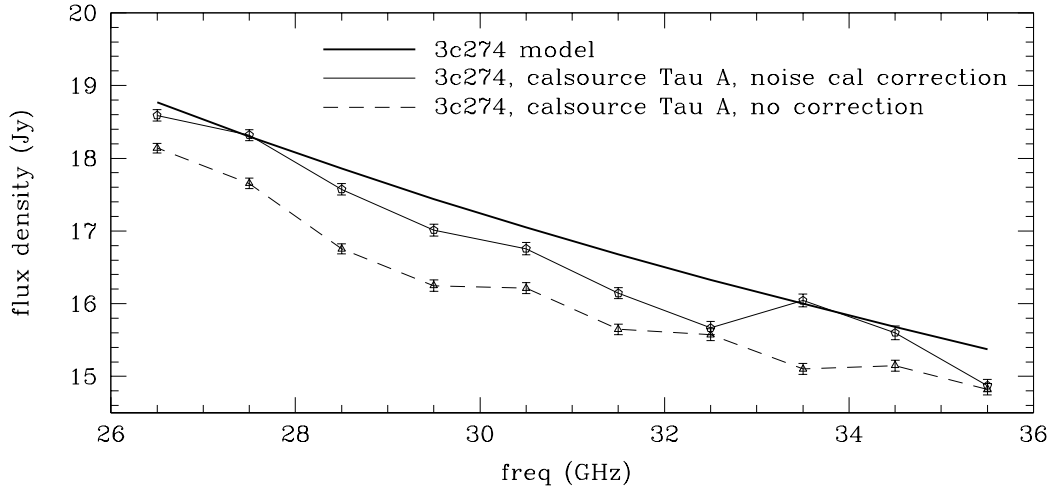


Figure 2.14: 3C274 flux densities with and without the noise source correction.

the application of the noise cal correction introduces more errors than it removes. Second, the power measured by the power meter can diverge from the power injected into the receivers by as much as 15% peak to peak. While this range is quite large, changes of this magnitude are only observed over timescales spanning several months; the discrepancies in a single night are much smaller. Observations of supporting sources spanning many months suggest that these variations introduce spurious variations in the measured flux densities which obviated the benefits of the noise cal. As we will see, these discrepancies have a strong dependence on the physical temperature of the system. Finally, the amplifier output has a downward slope of 10 dB across the 26-36 GHz band, so the power read by the power meter is dominated by the power at the low end of the band. These problems, and their remedies, are discussed below.

Large Noise Cal Variations

On rare occasions the variations in the noise source visibilities exceeded the variations seen in the contemporaneous program data. The poor S/N in the CMBR data on short timescales precludes its use as a proxy for the receiver gains, but on several occasions we performed deep observations of the polarization calibrators at

a variety of deck positions, and the individual scans have sufficient S/N in LL to permit an assessment of the noise source stability. On some dates the noise source flux variations exceeded the underlying gain changes; Figure 2.15 provides a particularly egregious example of this problem. In this case the noise source power injected into RX10 is far more unstable than the underlying gain of the system; the figure compares noise source and 3C279 visibilities for bands 2 and 7 on baseline RX0-RX10. An inspection of *all* of the visibilities reveals that this problem is confined to RX10, and since the 3C279 fluxes provide an upper limit on the gain changes for RX10 which is well below the variations in the noise source, the problem probably resides in the noise source distribution network between the power source and RX10—one possible candidate is a loose coaxial connection which suffers from a microphonic response during the slews that separate the observations. Errors of this sort are rare, and easily identified by eye in the data.

The observation of occasional pathologies such as these, coupled with tests of the accuracy of the noise source correction for observations spanning timescales of months or longer, prompted a modification to the calibration procedure for which the noise source was designed. The original design called for the noise source visibilities to be used to interpolate gain corrections for the program data. These problems motivated a second procedure in which the amplitude correction is derived from an average of all of the noise source visibilities during the session. This procedure, discussed in detail in Chapter 3, applies a uniform correction to all of the visibilities based on the mean of the noise source amplitude for the session; this procedure effectively *eliminates* the noise cal from the calibration procedure. While this second procedure was used for the total intensity observations, however, the polarization program prevailed on the noise cal system on nights which lacked a suitable observation of a polarization calibrator.

Power Meter Discrepancy

The noise source power measured by the power meter can depart from the power which the system injects into the array, and these flux scale errors can complicate

Comparison of 3C279 LL fluxes and interleaved ncals; RX0–RX10 on 23apr00

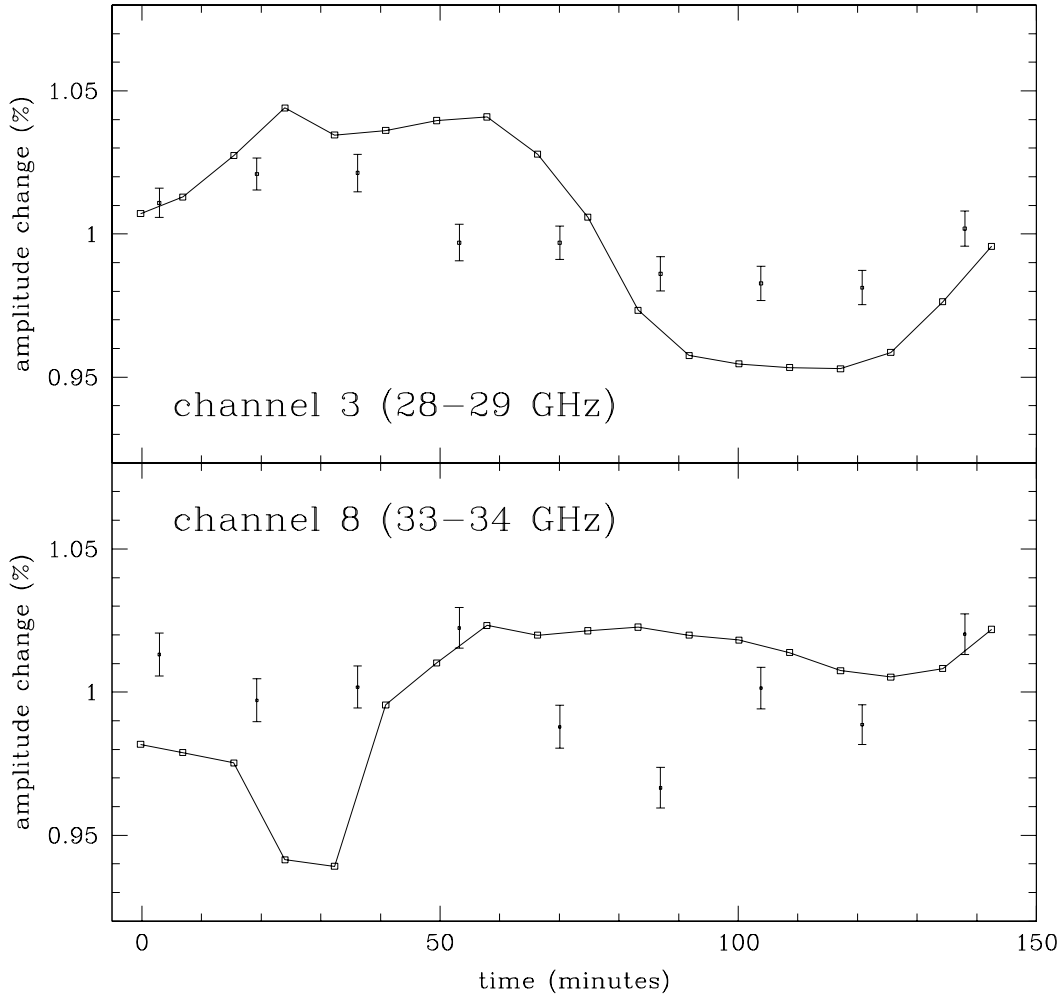


Figure 2.15: Internal noise source amplitude errors for two baselines. The points denote 3C279 fluxes, and lines denote noise source fluxes where the boxes are the actual noise source pulses. All data have been normalized to the mean for the period. The standard noise source calibration would interpolate the 3C279 fluxes to the noise source amplitudes, thereby exacerbating the noise in the data. An inspection of other baselines suggests that this problem is confined to RX10; one possible candidate is a microphonic response in the noise source connection to RX10. Pathologies this extreme are rare, however.

the application of the noise cal correction. Unlike the large noise source variations discussed above, these errors escape a visual inspection of the data; they can only be detected by using astrophysical sources of known flux density to track changes in the noise source flux scale. The details of this technique require an understanding of the CBI polarization calibration routine, which must await the discussion of Chapter 3, but this section describes the results of this work. We show that the power meter discrepancy has a quadratic dependence on the ambient temperature.

The two deep field observations both span at least three months, and these periods provided the long baselines necessary to explore the temperature dependence of the power meter discrepancy. Figure 2.16 shows the noise cal amplitude error as a function of the temperature of the noise cal amplifiers for the 31-32 GHz band during the 01aug00-03oct00 period;¹¹ the error is a flat quadratic function of temperature spanning $\pm 10\%$ for most amplifier temperatures in the vicinity of -5° . The other nine bands show the same quadratic behavior, although the fit coefficients differ. Intervals during the 08^h field observations have the same quadratic dependence, albeit with a different offset; the constant terms in the quadratic fit depend on the date that provides the reference observation against which the corrections are determined. We exploit the high correlation between the power meter error and the amplifier temperature in Chapters 3 and 5.

Noise Cal Band Slope

The noise cal power suffers from ~ 10 dB slope across the band. This steep slope has a deleterious effect on the behavior of the noise calibration system; the power meter measures the power across the entire 26-36 GHz band, and since the high end of the band is much weaker than the lower end, the power measured by the system does not accurately reflect the variations at the high-frequency end. To remedy this problem, a pair of bandpass filters and a second power meter were added to the system; the filters break the 26-36 GHz band into two 5 GHz wide bands, and the two power meters report the noise in the two halves of the band to

¹¹After 03oct00 the source underwent maintenance which changed the zero point of the fit.

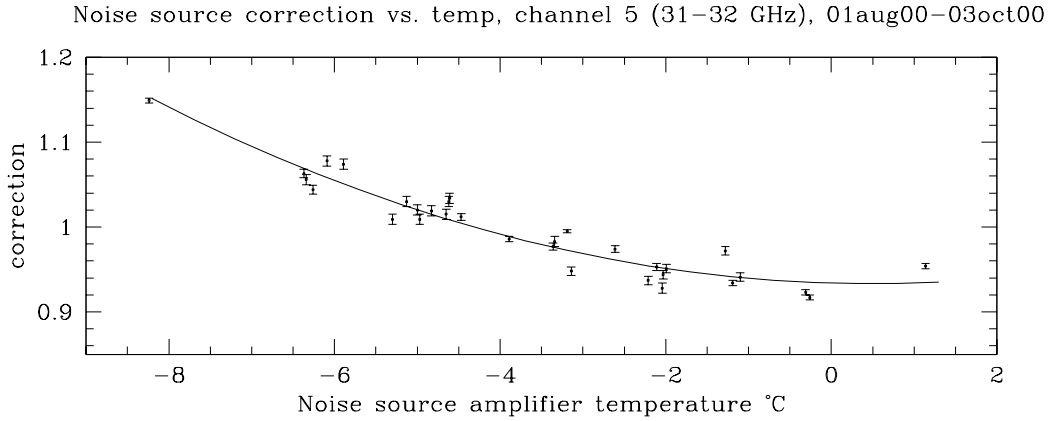


Figure 2.16: Temperature dependence of the noise source flux scale error for the 31jul00-03oct00 period.

the data stream. This modification provides some improvement, particularly for the upper end of the band, but this addition is not part of the scope of this thesis, nor were any thesis data collected after this modification was performed.

Summary

The noise calibration system is a subsystem of the signal chain that was designed and implemented for this thesis. The noise cal suffered from a variety of problems which compromised its efficacy for the total intensity observations; these problems resulted primarily from poor control of the temperature of the system, and greater attention to the thermal stability of the amplifiers, switches, and coaxial distribution network would improve its performance. The bandpass slope was addressed by dividing the 26–36 GHz band into two adjacent bands, each of whose power was measured separately; a filter at the amplifier output to flatten the signal seen by the receivers would achieve the same result at a smaller cost. In Chapters 3 and 5 we show that despite these problems, the noise cal aids the calibration of polarization data for dates which lack direct observations of a polarization calibrator.

2.4.3 HEMT Amplifiers

The CBI is among the first telescopes to employ a new generation of low-noise, high-bandwidth Indium Phosphide HEMT amplifiers for the 26-36 GHz band. These amplifiers provide the high bandwidth and the low noise which are necessary to achieve the high sensitivity required for our science goals. The CBI HEMT design was developed by at NRAO's Central Development Laboratory [76], and the author went to NRAO to learn how to build and optimize these critical components. Our program requires 17 HEMT amplifiers: 13 + 2 spares for the CBI, and 1 + 1 spare for the Owens Valley 40-meter telescope.¹² The amplifiers were assembled at NRAO's Central Development Laboratory in Charlottesville, VA, and in the CBI labs on the Caltech campus. HEMT production requires several specialized tools: a wirebonder and a microscope for assembly, and a network analyzer and a cryostat for testing; both labs had these tools on hand. The Caltech lab was shipped with the CBI to Chile to facilitate HEMT repairs at the Chajnantor site, and since several HEMTs were damaged during the shipping of the instrument, the CBI benefitted from having these facilities on site. On average, each HEMT requires up to two days for assembly and two to four days for testing.

A HEMT amplifier consists of four Indium Phosphide Field-Effect Transistors (FETs) in series (Figure 2.17). Each FET provides ~ 8 dB of gain at wavelengths of ~ 1 cm, and when cooled to a physical temperature of ~ 10 K, the devices have a total noise temperature of ~ 15 K near the band center and ~ 20 K at the band edges. The *rf* signal travels along the vertical stripline from the input at top of the unit to the output at the bottom. The stripline couples to waveguide (not shown) with probes that are soldered to the substrate; the silver tip of the input probe can be seen at the top of Figure 2.17. The *rf* stripline is punctuated by the FETs at the intersections of the *rf* channel and the four horizontal DC bias channels. The bias channels accommodate the DC networks which set and

¹²The 40-meter telescope was an essential component of the point source strategy that was necessary for the total intensity observations at high ℓ . A Ka-band receiver was built for the 40-meter telescope for the CBI program.

monitor the gains of the FETs. The FETs, stripline, and chip components in the bias network are glued to the body of the amplifier with conductive epoxy and connected with 0.7 mil (0.0007") gold wire. The figure notes the physical scale; the FETs are small, ~ 10 mil on a side, and consistent electrical performance requires that they are positioned to an accuracy of a few mil. The primary challenge to producing these devices is satisfying these tight mechanical assembly requirements with a production process in which components are installed by hand.

HEMT Tuning

All of the amplifiers required some tuning to achieve the noise and gain specifications of the CBI. The CBI requires 30-35 dB of gain and as little noise as possible across the 26-36 GHz band. The gain and noise can vary dramatically across the 26-36 GHz band—the noise by 50% and the gain by 5 dB—so the tuning required measurements across the band with at least 1 GHz resolution. There are four primary degrees of freedom for tuning at our disposal: the lengths of the gate bonds for each of the four stages, which in turn determine the inductance and parasitic capacitance looking into each device. The first stage, for example, requires a gate bond which is 23 mil long; increasing the length of this bond increases the inductance at the device input and thereby attenuates the power at the high end of the band. Unfortunately, the tuning process is somewhat destructive; the FETs can sustain only a few repeated bonds, so to avoid installing a new bond with each tuning iteration, the first bonds are made ~ 5 mil too long, and then shortened by crimping the end of the bond to the *rf* stripline. The optimum set of tuning parameters required to obtain good T_{sys} performance across the band is not the same as that which is required for the gain; in practice we compromise between the two constraints, with a bias in favor of flat T_{sys} across the 26-36 GHz band.

The time required to tune a HEMT is set by the cooling cycles that are required to evaluate the HEMT performance under cryogenic conditions. Each cryogenic cycle requires $\sim 7^h$, and while measurements at cryogenic temperatures provide the best indicator of how the HEMT will perform on the array, the cost in time

CBI Ka-band HEMT amplifier




Figure 2.17: Closeup view of an NRAO HEMT. Signals enter from the top, pass through the four stages at the junctions of the horizontal and vertical channels, and exit at the bottom. The two holes accommodate two of several 2-56 screws which hold the cover on the amplifier. The key at the bottom left shows the scale of the device.

for many cycles can be considerable. Fortunately, experience shows that there is a mapping between the warm performance and that at cryogenic temperatures; most importantly, the gain and noise bandpasses tend to shift ~ 1 GHz lower when the HEMT is cooled. This behavior was exploited for the tests during the latter part of the CBI production run. For these HEMTs, warm measurements of the insertion characteristics, combined with warm T_{sys} measurements, circumvented the initial cryogenic testing cycles, and thereby dramatically reduced the production time. The warm T_{sys} , which is typically ~ 100 K, provides little guidance about the absolute minimum noise under cryogenic conditions, but since the process of cooling the HEMTs preserves the shape of the noise curve (to within a ~ 1 GHz shift), the warm T_{sys} curves retain some value for the purpose of tuning the HEMTs.

CBI HEMT Performance

The 17 CBI HEMT amplifiers showed minor variations in performance. Figure 2.18 shows the range of noise performance across the 26-36 GHz band for the CBI HEMTs. The figure also provides a model prediction for comparison, as well as T_{sys} values for a sample unit built by NRAO [75].¹³ No HEMTs have demonstrated the noise performance predicted by the model, although the NRAO prototype comes within 2 K of achieving this spectacular performance near the center of the band. The CBI HEMTs approximate the NRAO performance at the low end of the band, and most of the CBI HEMTs occupy the lower half of the shaded region in Figure 2.18; for the 15 best CBI amplifiers, for example, $\langle T_{sys} \rangle \sim 16$ K for the central eight channels, while that for the NRAO sample is ~ 12 K across the same frequencies.

For a properly tuned HEMT, the noise performance is a strong function of the quality of the FETs which are installed in the HEMT. The first stage device in particular dominates the noise performance of the entire unit. A comparison of gain and T_{sys} curves for the CBI HEMTs shows that the tuning is fairly consistent, so in light of the iterations required to obtain the best T_{sys} performance, the range

¹³NRAO A23

in T_{sys} seen in Figure 2.18 reflects the intrinsic noise characteristics of the first stage devices.

Summary

Part of this thesis focused on the production and implementation of the HEMT amplifiers based on an NRAO design. These components that are critical to the performance of the CBI; their ~ 15 K of noise dominates the noise budget. The HEMT amplifiers achieved nearly the performance to be expected based on the NRAO test unit; since most of the units have the same performance across the band, the tuning is similar. Most of the difference arises from device variations for the first stage FETs.

2.5 Conclusions

This chapter describes the Cosmic Background Imager. Section 2.2 reviews the response of an interferometer; the main results of this discussion are that for small scales a visibility measured by a baseline of length $|b|$ samples a scale ℓ given by $\ell = 2\pi|b|/\lambda$, and that the covariance of a pair of visibilities provides an estimate of C_ℓ . Section 2.3 discusses the Cosmic Background Imager in detail. We are particularly concerned with its polarization characteristics; the CBI employs single-mode circularly polarized receivers, and we saw that the quarter-wave plates, which are used to define the mode to which the receiver responds, are also a source of instrumental polarization. Section 2.3.3 presents a model for the instrumental polarization from the CBI signal chain, and in Chapter 3 we compare the model to measurements of the instrumental polarization.

This thesis entailed design and implementation of a number of major systems for the CBI, and the latter part of Chapter 2 focuses on these projects. The downconverter was the largest of these efforts; it met the CBI's requirements for size and performance, and its occasional failures did not have a deleterious effect on our observing time. The noise calibration system suffered from a variety of

instabilities which precluded its use for the total intensity observations. One of the major problems is a failure for the power meter in the system to accurately report the power distributed to the receivers; we show that these errors are strongly temperature dependent, and in Chapters 3 and 5 we exploit this temperature dependence to improve the noise cal performance to a level that is acceptable for the polarization observations. Finally, the low-noise HEMT amplifiers were produced for this thesis; the HEMTs are critical components of the CBI because they dominate the considerations which affect the CBI's sensitivity. The CBI benefitted from having the capacity to build and repair these parts on site.

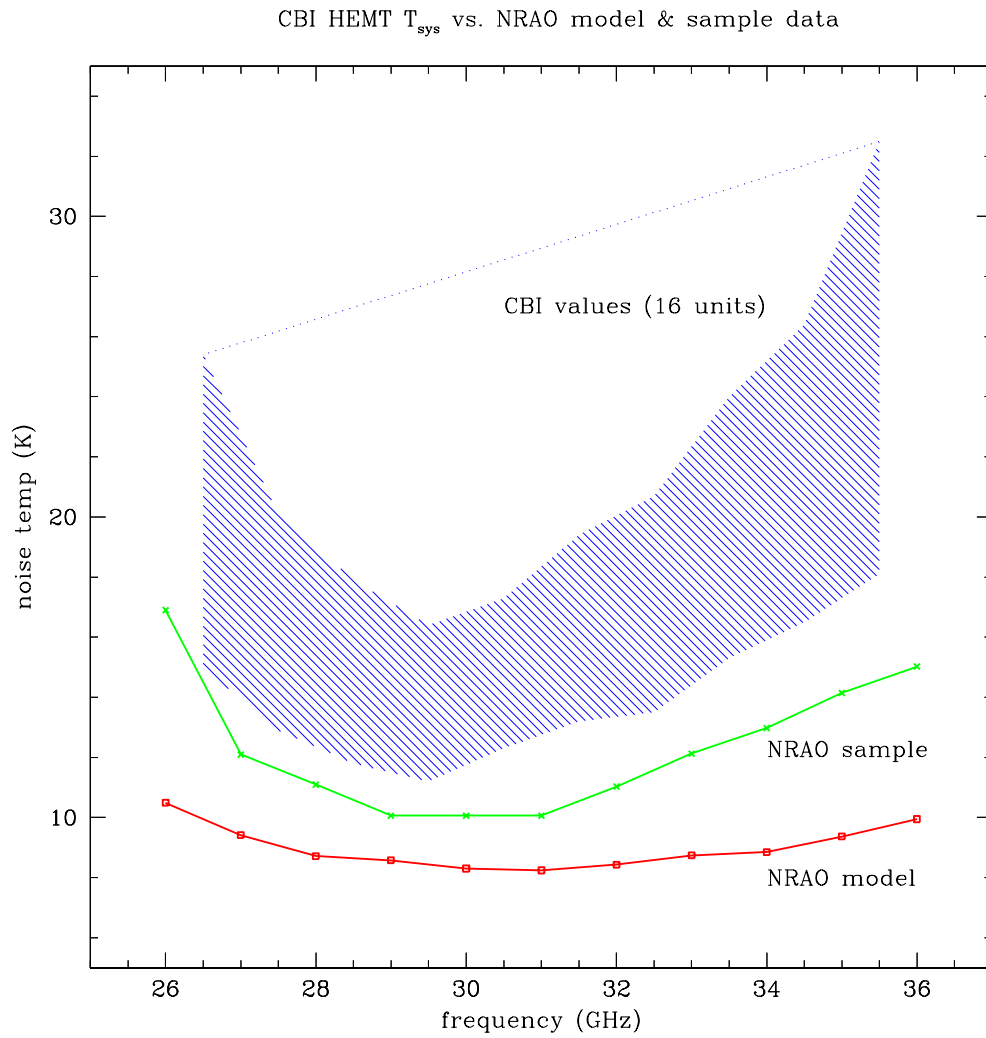


Figure 2.18: CBI HEMT results with NRAO model and data from single NRAO sample amplifier.

Chapter 3

Polarization Calibration

3.1 Introduction

The outcome of our polarization detection effort hinges on the quality of the polarization calibration, and the substantial amount of time reserved for calibration observations—roughly 15% of each observing session—reflects the importance of this concern. We incorporated a number of calibration procedures into each night of observations. Each session began at sunset with an optical pointing procedure to create a pointing model for the session; good pointing is a requirement for a good phase calibration, and the models typically yielded pointing accuracy of $\sim 20''$ p-p. The pointing calibration was followed by a quadrature calibration, which is necessary to measure the gains of the correlator’s real and imaginary channels. The program observations—CMB fields, the SZ effect in clusters, and supporting observations—followed the quadrature calibration. These observations were interleaved with observations of the CBI’s internal calibration source, to assess receiver gain variations; and astrophysical calibrators, to refer the internal amplitudes and phases to an astronomical scale. The astronomical calibration is also necessary to correct for the CBI’s instrumental polarization, or *leakage*.

This chapter discusses the considerations which affect the polarization calibration of the CBI. Section 3.2 reviews the calibration procedures; this section develops the leakage model for the instrumental polarization, which states that

polarization observations of a source are contaminated by a fraction of its total intensity. The CBI's deck rotation about the pointing center plays a critical role in the polarization calibration procedures; we will see that we can isolate the leakage by observing a source of known polarization at a variety of deck positions. This procedure also permits us to measure the gains of the receivers.

We observed 3C279, our primary polarization calibrator, under a range of conditions to evaluate the polarization performance of the CBI, and Section 3.3 discusses these observations. We are particularly interested in the stability of the instrumental polarization, as well as the off-axis behavior of the CBI's polarization response. Section 2.3 presents the results of tests of these characteristics; we show that the instrumental polarization is somewhat less stable than the uncertainties on the measurements would suggest, and that the off-axis polarization is consistent with the instrumental polarization at the telescope boresight. We also compare the measurements of the leakage with the leakage to be expected from the phase shifter model developed in Chapter 2, and show that the model can account for all of the leakage.

3C279 and Tau A were the primary polarization calibrators for the CBI polarization observations. While the 3C279 polarization was established from the contemporaneous monitoring campaign with the VLA, we did not have independent measurements of Tau A's polarization of sufficient quality to calibrate the CBI. Tau A is resolved by the CBI, so observations in band and with the CBI's resolution offer the best hope of obtaining an accurate model. Tau A and 3C279 overlapped briefly during the 2000 observing season, and Section 3.3 discusses efforts to transfer the calibration from 3C279 to Tau A; the techniques discussed in this section yield a model which is good to $\sim 10\%$.

3.2 Calibration Procedures

The CBI developed several custom software packages to calibrate the LL and LR visibilities. Figure 3.2 provides an overview of the CBI polarization calibration

pipeline. CBICAL reads the raw visibilities from the CBI archive, along with various pieces of housekeeping and telemetry data which aid the analysis. CBICAL provides a suite of tools for inspecting and flagging visibilities, as well as algorithms which correct some irregularities, such as pointing errors which escape the pointing calibration. CBICAL performs the `quad` and `ncal` calibrations (described below), and writes the partially calibrated data to `.wvf` (*wv*-fits) files.¹ Most observations were accompanied by observations of trailing fields for the rejection of spillover; UVSUB differences the lead and trail visibilities written by CBICAL, and passes the subtracted data in the form of a *wv*-fits file to CBIPOLCAL. CBIPOLCAL employs observations of sources of known polarization to isolate the instrumental polarization, and it removes this spurious signal from the program data before placing the visibilities on an astrophysical flux scale.

We deal with several different kinds of files. CBICAL reads the data from the archive and writes the visibilities to `.wvf` files; this is the standard format for visibility data. The polarization calibrator characteristics are stored in `.par` files, which are read by CBIPOLCAL. The polarization calibration results—the gain and leakage for each baseline and band—are saved by CBIPOLCAL in `.cal` files, where they are stored for application to subsequent observations of program sources.

3.2.1 CBICAL: Quadrature Calibration

The quadrature calibration is necessary to balance the gains of the real and imaginary branches of each correlator channel. The CBI employs a complex correlator, and under ideal circumstances the real and imaginary branches for each baseline have identical amplitudes and a 90° phase offset. In practice, however, bandpass errors and component variations give rise to mixing between the channels which degrades their orthogonality, and these errors must be eliminated. The *quadrature calibration* consists of a 5^m procedure in which the CBI’s internal noise source

¹CBICAL performs the astronomical calibration for the *LL* visibilities, but the *LR* visibilities bypass this procedure because CBICAL does not isolate or eliminate the instrumental polarization.

Standard CBI polarization calibration pipeline

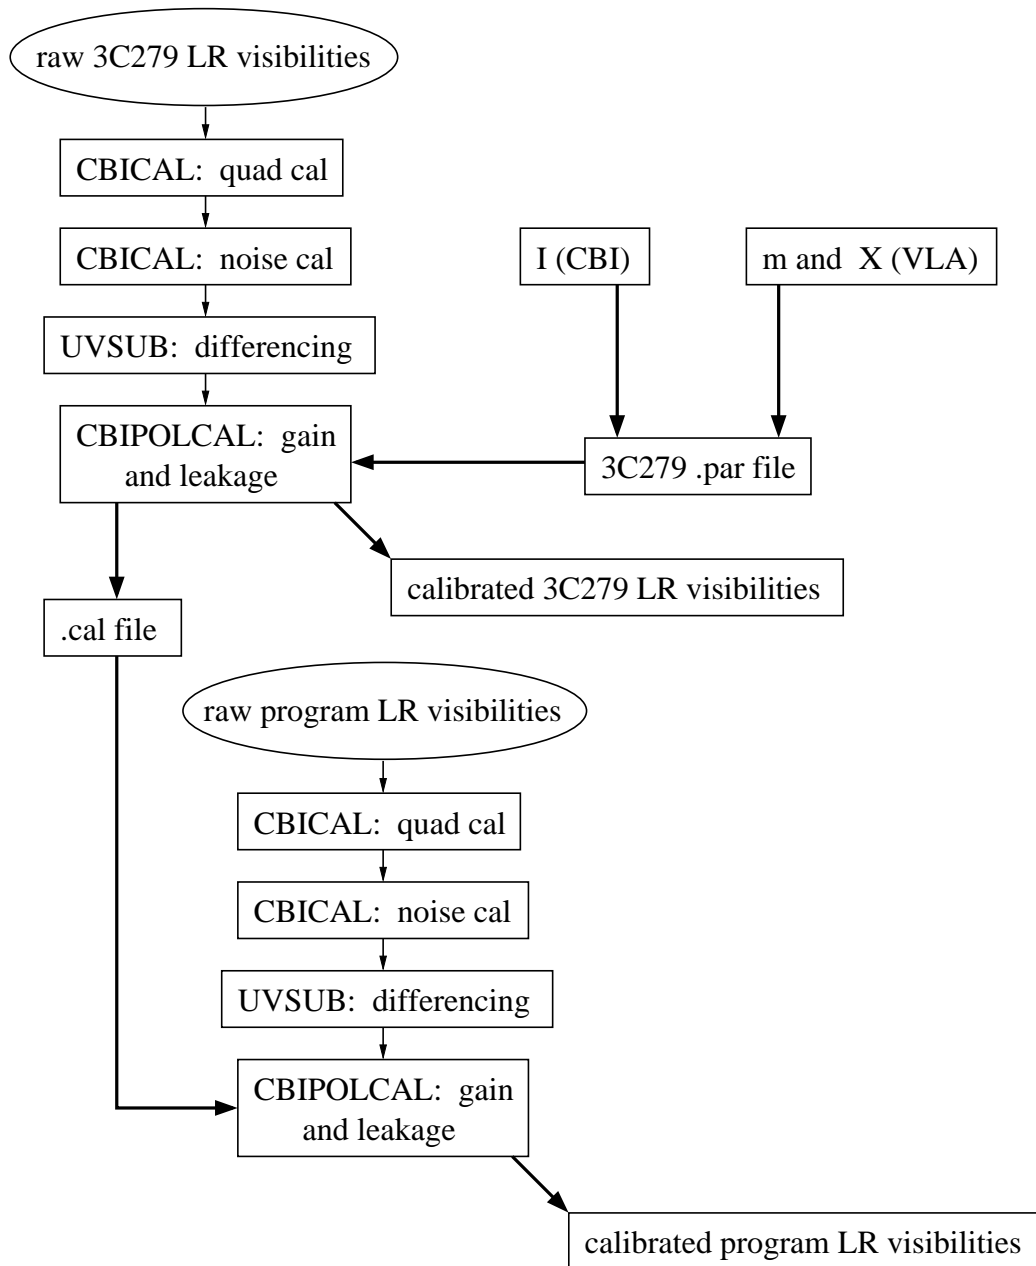


Figure 3.1: Polarization calibration flow chart. The upper half shows the procedure which is required to derive the calibration factors—gain and leakage for each band and channel—using the fractional polarization m and position angle χ data from the VLA. The calibration factors are stored in the *.cal* file. The lower half shows the application of these factors to the program data. The calibration procedure for *LL* baselines in CBIPOLCAL differs in only one respect: no correction is made for the leakage.

injects a pair of orthogonal signals into each receiver in succession.² In the offline analysis, CBICAL’s quad task compares the real and imaginary visibilities at each correlator output, computes the phase offset and amplitude mismatch between the two branches, and adjusts all of the visibilities for the session accordingly. The quadrature corrections are typically $\sim 5\%$ in amplitude and $\sim 3^\circ$ in phase, and quadrature calibrations which bracket the observing sessions demonstrate that these corrections are effectively constant ($\delta A < 1\%$, $\delta\phi < 1^\circ$) over the duration of the observing sessions. Although the quad corrections can occasionally be quite large—10% in amplitude, and 6° in phase—they are stable and easy to measure.

3.2.2 CBICAL: Internal Noise Source Calibration

The CBI employs an internal broadband noise source that provides a reference against which changes in the receiver amplitude gains during an observing session can be measured. The noise cal system distributes noise power from a thermal load to the 13 receivers. The noise power can drift with time, so a power meter at the amplifier output reports the amplitude of the signal to the data stream, and CBICAL uses these readings to correct the noise source visibilities for the amplitude variations which are intrinsic to the noise cal system. Temperature probes on components in the noise source provide diagnostic data which are crucial for characterizing the stability of the noise source.

The noise source observations were brief but frequent; all slews and deck rotations for program and calibration sources were bracketed by 10^8 integrations on the source. The noise cal flux is equivalent to a ~ 2000 Jy source, so these short observations suffice to obtain high S/N measurements of the noise cal visibilities. The noise power is injected into the receivers immediately after the quarter-wave plates which define the polarizations of the receivers, so both the noise calibration

²Figure 2.2 shows that a 38 GHz local oscillator is required to multiply the 26-36 GHz band to the 2-12 GHz band; the phase shifts for the quad calibration are implemented by changing the phase of this LO by 90° . Since the LO is a single tone, passband errors do not degrade the orthogonality of the 0° and 90° LO signals, and the signals seen by the correlator during the quad routine are orthogonal to a high degree of accuracy.

and the quadrature calibration apply to *all* baselines.

Noise Calibration Stability during Single Sessions

The utility of the noise cal system rests with its stability; the noise power injected into the receivers must be at least as stable as the underlying receiver gains if the system is to improve the raw data. As noted in Chapter 2, there is evidence that the noise cal instabilities can exceed the gain instabilities in the receivers, and there are two classes of problems which demonstrate the shortcomings in the noise cal system. First, comparisons of total intensity calibrators which bracket an observing session suggest that the noise cal procedure introduces spurious errors in the flux calibration. Second, a comparison of noise cal fluxes with interleaved observations of bright sources show that the noise cal does not always track the intrinsic receiver gain variations.

In light of these problems, CBICAL provides two procedures for incorporating the noise cal data: `ncal` and `ncal1`. The `ncal` task assumes that noise cal flux variations reflect real gain variations in the system, so it uses the noise cal visibilities to interpolate a gain correction for the data between the noise cal pulses; this is the mode in which the analysis was originally intended to proceed. The `ncal1` task assumes that the variations in the noise source dominate the gain variations, so it averages all of the noise visibilities in a session and computes a *mean* gain correction for the entire session based on the averaged noise source flux; this procedure effectively removes the noise cal from the calibration procedure. Chapter 2 discusses the performance of the noise cal, and notes that although there are nights for which `ncal` provides a superior correction to the data, on whole the data suggest that `ncal1` results in a slightly better solution. Based on this work, `ncal1` was used for both the total intensity and the polarization data.

Noise Calibration Stability across Multiple Nights

The CBI polarization program differs from the intensity program in the extent to which the polarization observations rely on the noise cal. The paucity of primary

polarization calibrators often compelled us to transfer calibrations across many observing sessions, and in these instances, the noise calibration system provided the only guide to the gain stability of RX12. We must compensate for gain variations during these intervals, and the precision with which we can do so is limited by the stability of the noise cal.

We can estimate the stability of the noise source by comparing the noise cal amplitudes to known total intensities for astronomical sources. An analysis of the noise cal visibilities during the 01aug00 to 31oct00 period demonstrated that the amplitude of the noise cal has a quadratic temperature dependence; for ambient temperatures between -20 and 10 °C, the noise cal amplitude has a typical temperature coefficient of $-4\%/^{\circ}\text{C}$ to first order, although the variations are best fit by a quadratic. Figure 2.8 shows a sample of this temperature dependence for the 31-32 GHz channel; these data are representative of all of the channels. As discussed in Section 5.4.1, these coefficients are used to calibrate observations which occurred during a 2.5 week span in August-September 2000 during which there was no primary polarization calibrator. The total intensity observations never require this approach because nearly all observations included at least one total intensity calibrator.

3.2.3 UVSUB: Spillover Rejection

The presence of ground spillover in the visibilities forced us to incorporate a spillover rejection strategy into our observing scheme. We implemented a procedure in which we tracked fields in pairs—a lead and a trail—and differenced the two fields to reject the common spillover contamination. UVSUB performs the subtraction; it matches (u, v) points in the lead with the corresponding points in the trail, and differences the real and imaginary parts of the associated visibilities. UVSUB also computes the errors on the visibilities from the scatter in the subtracted data; a 5^m observation, for example, permits an estimate of the noise from the sample variance of ~ 40 8^s visibilities that constitute the scan. Since the cross polarized visibilities arrive at UVSUB in correlator units, the errors computed by

UVSUB cannot be compared to estimates based on the antenna T_{sys} ; this comparison must await the flux calibration performed by CBIPOLCAL. In cases where observations are not accompanied by trails, UVSUB simply averages the many 8.4° visibilities in the scan, and estimates the visibility errors from the scatter in the 8.4° integrations.

3.2.4 CBIPOLCAL: Leakage and Flux Calibration

The CBI imparts artificial polarization to the visibilities, and since this instrumental polarization can compete with the intrinsic polarization of the sources of interest, we require an additional calibration procedure. To complete the calibration, we must correct the cross polarized visibilities for instrumental polarization, and refer all of the visibilities— \mathcal{V}^{LL} and \mathcal{V}^{LR} —to an astrophysical flux scale. These tasks cannot be disentangled, so a new package, CBIPOLCAL, was written to perform both procedures [85]. These tasks are considered in detail below, beginning with a model for the instrumental polarization.

CBIPOLCAL employs the *leakage model* for the instrumental polarization, which assumes that the ideal pure polarization response of each receiver is corrupted by a small contribution from the orthogonal mode.³ The derivation of the leakage model is straightforward. The plane wave incident on the interferometer can be expressed in terms of right circularly polarized (RCP) and left circularly polarized (LCP) components:

$$\mathcal{E}(\mathbf{x}, \nu; t) = \mathcal{E}^R(\mathbf{x}, \nu; t)e^{i\phi} + \mathcal{E}^L(\mathbf{x}, \nu; t)e^{-i\phi} \quad (3.1)$$

The factor of $e^{\pm i\phi}$ enters the expression because the baseline orientation advances or retards the phase of the circularly polarized components of the wavefront, depending on the mode. The position of the baseline in the aperture plane determines the phase: $\phi = \tan^{-1}[v/u]$. Because the baselines are fixed to the deck, the base-

³The leakage model ignores the loss in signal for the desired polarization which arises, for example, from the attenuation in the wave plates in the CBI's phase shifter assembly.

line orientation and the deck orientation are interchangeable to within an offset determined by the array geometry.

An ideal circularly polarized receiver responds to only a single mode of circular polarization. In practice, however, a variety of effects conspire to contaminate the ideal response with a small contribution from the total intensity of the source. Consider two imperfect receivers (j, k) which combine to form a cross polarized baseline; configure receiver j for LCP and k for RCP. Instrumental polarization allows one mode of polarization to leak into the receiver configured for the other; we characterize this contamination with the complex leakage term ϵ . The signals at the receiver outputs are simply voltages,

$$V_j^L(\mathbf{x}, \nu; t) = g_j \left[\mathcal{E}^L(\mathbf{x}, \nu; t) e^{-i\phi} + \epsilon_j \mathcal{E}^R(\mathbf{x}, \nu; t) e^{i\phi} \right] \quad (3.2)$$

$$V_k^R(\mathbf{x}, \nu; t) = g_k \left[\mathcal{E}^R(\mathbf{x}, \nu; t) e^{i\phi} + \epsilon_k \mathcal{E}^L(\mathbf{x}, \nu; t) e^{-i\phi} \right] \quad (3.3)$$

which the correlator multiplies to obtain the visibility $\mathcal{V}^{LR} = \langle V_j^L V_k^{R*} \rangle$:

$$\mathcal{V}^{LR} = g_j g_k^* \left[\langle \mathcal{E}^L \mathcal{E}^{R*} \rangle e^{-2i\phi} + \epsilon_k^* \langle \mathcal{E}^L \mathcal{E}^{L*} \rangle + \epsilon_j \langle \mathcal{E}^R \mathcal{E}^{R*} \rangle + \epsilon_j \epsilon_k^* \langle \mathcal{E}^R \mathcal{E}^{L*} \rangle e^{2i\phi} \right] \quad (3.4)$$

Apply Equation 2.4 with $V = 0$ to find

$$\mathcal{V}^{LR}(\mathbf{u}, \nu) = g_j g_k^* \left[\tilde{P}^*(\mathbf{u}, \nu) e^{-2i\phi} + \tilde{I}(\mathbf{u}, \nu) (\epsilon_j + \epsilon_k^*) + \epsilon_j \epsilon_k^* \tilde{P}(\mathbf{u}, \nu) e^{2i\phi} \right] \quad (3.5)$$

with $P^* = (Q - iU)$. We can make some assumptions to simplify this expression. For typical sources, $P \sim 0.1I$, and for the CBI, $\epsilon \sim 10\%$, so that $P:\epsilon I:\epsilon^2 P^*$ scale as $0.1:0.1:10^{-3}$. We therefore ignore the $\epsilon^2 P$ term. In addition, we have a leakage term for each of the 13 antennas, but we have only 12 cross polarized baselines, so we can never solve for each leakage—only the sum of the two terms associated with each baseline.⁴ We therefore regard the leakages (ϵ_j, ϵ_k) associated with a pair of

⁴In principle, we can use the instrumental polarization in the LL visibilities (see Equation 3.20) to solve for all 13 leakage terms to obtain an antenna-based solution, but the high S/N necessary to do this is prohibitive.

antennas as a baseline-based parameter ϵ_{jk} . For consistency we will do the same with the gain, letting $G_{jk} = g_j g_k^*$, and although we have the S/N to determine the antenna-based gains, we do not do so in this work. With these assumptions we obtain the following:

$$\mathcal{V}^{LR}(\mathbf{u}, \nu) = G_{jk;\nu} \left[\tilde{P}^*(\mathbf{u}, \nu) e^{-2i\phi} + \epsilon_{jk;\nu} \tilde{I}(\mathbf{u}, \nu) \right] \quad (3.6)$$

The goal of polarization calibration is to determine the gains G_{jk} and the leakages ϵ_{jk} for each of the ten bands and the twelve cross polarized baselines.

The polarization calibration procedure capitalizes on the fact that the CBI's deck rotation modulates the source polarization term $\tilde{P}^*(\mathbf{u}, \nu) e^{-2i\phi}$ relative to the instrumental polarization $\epsilon \tilde{I}(\mathbf{u}, \nu)$. This relationship has a simple graphical representation: as the deck angle ϕ changes, the cross polarized visibilities for a polarized point source trace a circle in the (Q, U) plane, such that the circle is centered on the instrumental polarization, and the radius is proportional to the source polarization (Figure 3.2.4). As the instrumental polarization increases, the circle develops ellipticity, although this ellipticity is negligible for the CBI. CBIPOLCAL combines multi-deck angle observations of polarization calibrators with values for $\tilde{I}(\mathbf{u}, \nu)$ and $\tilde{P}^*(\mathbf{u}, \nu)$ supplied by the user in the *.par* file to solve for G_{jk} and ϵ_{jk} for each band.

Calibration: Two Visibility Example

It is illustrative to apply Equation 3.6 to a sample calibration observation. A single observation of the calibrator suffices to determine the gain or the leakage, but not both. Observations at three deck positions define a circle in the (Q, U) plane, although two observations suffice if we *know* that the visibilities span the diameter of the circle—i.e., the visibilities are measured at deck positions that are 90° apart. This case has an analytic solution, which is presented below.

CBIPOLCAL inspects the (u, v) points of the calibrator data and identifies pairs of visibilities which correspond to deck angles that are separated by 90° . CBIPOL-

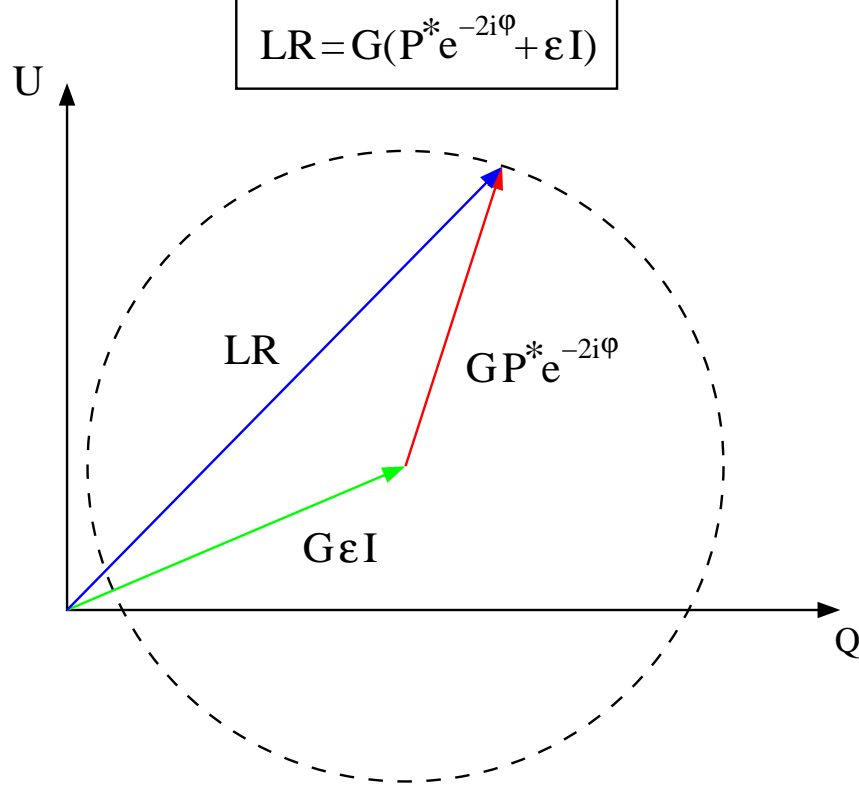


Figure 3.2: In the presence of instrumental polarization, the cross polarized visibilities measured at a variety of deck positions ϕ trace a circle in the (Q, U) plane which is centered on the leakage ϵ and whose diameter provides a measurement of the gain G .

CAL fits these points to a circle, the diameter of which is proportional to the gain, and the center of which is proportional to the instrumental polarization. For simplicity, assume that the calibrator is a point source, which means that its characteristics are uniform across the (u, v) plane:

$$\tilde{I}(\mathbf{u}, \nu) \Rightarrow \tilde{I}_\nu; \quad \tilde{P}^*(\mathbf{u}, \nu) \Rightarrow \tilde{P}_\nu^* = m_\nu \tilde{I}_\nu e^{-2i\chi_\nu} \quad (3.7)$$

where m_ν is the fractional polarization of the source, and χ_ν is the position angle of the source polarization on the sky. For notational clarity we will dispense with the ν that denotes the channel, and bear in mind that the following discussion

applies to each of the channels. For the same reason, we will also discard the indices j and k that refer G_{jk} and ϵ_{jk} to a particular baseline.

CBIPOLCAL first uses the diameter of the circle to obtain the gain for the baseline. In the absence of noise, the diameter is simply the vector difference of two visibilities separated by a 90° deck rotation. Let the first visibility be at \mathbf{u}_1 :

$$\mathcal{V}^{LR}(\mathbf{u}_1, \nu, \phi_1) = G[\tilde{P}^*(\mathbf{u}_1, \nu)e^{-2i\phi_1} + \epsilon\tilde{I}(\mathbf{u}_1, \nu)] \quad (3.8)$$

and let $\phi_2 = \phi_1 + 90^\circ$:

$$\mathcal{V}^{LR}(\mathbf{u}_2, \nu, \phi_2) = G[\tilde{P}^*(\mathbf{u}_2, \nu)e^{-2i(\phi_1+90^\circ)} + \epsilon\tilde{I}(\mathbf{u}_2, \nu)] \quad (3.9)$$

Compute the vector difference of the two visibilities to cancel the instrumental polarization terms, or equivalently, to isolate the diameter of the circle:

$$\mathcal{V}^{LR}(\mathbf{u}_2, \nu) - \mathcal{V}^{LR}(\mathbf{u}_1, \nu) = \Delta\mathcal{V}^{LR} = 2G\tilde{P}^*e^{-2i\phi_1} \Rightarrow G = \frac{1}{2\tilde{P}^2}\tilde{P}\Delta\mathcal{V}^{LR}e^{2i\phi_1} \quad (3.10)$$

Explicitly,

$$G_r = \frac{1}{2m\tilde{I}} \left[\Delta\mathcal{V}_r^{LR} \cos(2\chi + 2\phi_1) - \Delta\mathcal{V}_i^{LR} \sin(2\chi + 2\phi_1) \right] \quad (3.11)$$

$$G_i = \frac{1}{2m\tilde{I}} \left[\Delta\mathcal{V}_r^{LR} \sin(2\chi + 2\phi_1) + \Delta\mathcal{V}_i^{LR} \cos(2\chi + 2\phi_1) \right] \quad (3.12)$$

With G in hand, we can solve for the instrumental polarization. The instrumental polarization can be isolated by computing the vector average of $\mathcal{V}^{LR}(\mathbf{u}_2)$ and $\mathcal{V}^{LR}(\mathbf{u}_1)$, which forces the source polarization terms to cancel:

$$\frac{1}{2}[\mathcal{V}^{LR}(\mathbf{u}_2) + \mathcal{V}^{LR}(\mathbf{u}_1)] = \frac{1}{2}\Sigma\mathcal{V}^{LR} = G\tilde{I}\epsilon \Rightarrow \epsilon = \frac{1}{\tilde{I}}\frac{G^*}{G^2}\Sigma\mathcal{V}^{LR} \quad (3.13)$$

Express ϵ explicitly in terms of the gains G :

$$\epsilon_r = \frac{1}{2\tilde{I}G^2} \left[G_r\Sigma\mathcal{V}_r^{LR} + G_i\Sigma\mathcal{V}_i^{LR} \right] \quad (3.14)$$

$$\epsilon_i = \frac{1}{2\tilde{I}G^2} [G_r \Sigma \mathcal{V}_i^{LR} - G_i \Sigma \mathcal{V}_r^{LR}] \quad (3.15)$$

with $\Sigma \mathcal{V}_r^{LR} = \mathcal{V}_{r,1}^{LR} + \mathcal{V}_{r,2}^{LR}$, etc. Note that since $G \propto 1/I$, the leakage terms do not depend on the source intensity I . On the other hand, since $G \propto 1/m$, $\epsilon \propto m$; if we overestimate the fractional polarization of the calibrator, then we overestimate the instrumental polarization.

The leakage terms vary with ϕ , the baseline orientation, through its influence on G . This dependence becomes an important consideration when we reconfigure the array; reconfiguration changes the orientations of the cross polarized baselines with respect to their previous orientations, so a set of leakage terms measured in one configuration acquires phase offsets relative to the same terms measured in a different configuration. To compare leakage terms measured in different configurations, we use the geometries of the arrays to refer one set of leakage terms to the other. Reconfiguration does not change the leakage amplitudes.

The errors for G and ϵ follow directly from Equations 3.11, 3.12, 3.14, and 3.15. Based on these expressions, we have

$$\sigma_{G_r}^2 = 2 \left(\frac{1}{2mI} \right)^2 \sigma_{\mathcal{V}}^2 + \left(\frac{G_r}{m} \right)^2 \sigma_m^2 + \left(\frac{G_r}{I} \right)^2 \sigma_I^2 \quad (3.16)$$

$$\sigma_{G_i}^2 = 2 \left(\frac{1}{2mI} \right)^2 \sigma_{\mathcal{V}}^2 + \left(\frac{G_i}{m} \right)^2 \sigma_m^2 + \left(\frac{G_i}{I} \right)^2 \sigma_I^2 \quad (3.17)$$

for the gains, and

$$\sigma_{\epsilon_r}^2 = 2 \frac{m^2 (\Delta \mathcal{V}^2 + \Sigma \mathcal{V}^2)}{\Delta \mathcal{V}^4} \sigma_{\mathcal{V}}^2 + \left(\frac{\epsilon_r}{m} \right)^2 \sigma_m^2 + 4 \epsilon_r^2 \sigma_{\chi}^2 \quad (3.18)$$

$$\sigma_{\epsilon_i}^2 = 2 \frac{m^2 (\Delta \mathcal{V}^2 + \Sigma \mathcal{V}^2)}{\Delta \mathcal{V}^4} \sigma_{\mathcal{V}}^2 + \left(\frac{\epsilon_i}{m} \right)^2 \sigma_m^2 + 4 \epsilon_i^2 \sigma_{\chi}^2 \quad (3.19)$$

for the leakages. $\sigma_{\mathcal{V}}$ denotes the error on the LR visibilities, which is assumed to be the same for the real and imaginary visibilities, as well as for \mathcal{V}_1 and \mathcal{V}_2 if the corresponding integration times are the same.

Calibration with More than Two Visibilities

We occasionally observed the polarization calibrator at more than two deck positions. When the number of visibilities overdetermines a fit to a circle, CBIPOLCAL employs a Levenberg-Marquardt algorithm to solve for the diameter and center, both of which are nonlinear in the parameters of interest. The errors on the gain and leakage are derived from the errors on the visibilities which go into the fit, and they do not reflect the quality of the Levenberg-Marquardt solution. Figure 3.2.4 shows a fit for one band and one baseline which employs eight visibilities separated by 45° deck steps. The factor of 2 in Equation 3.6 maps the 45° deck steps to 90° separations on the circle. The fit for this band and baseline has $\chi_\nu^2 = 0.6$; the mean for all 110 fits on 06feb00 is $\langle \chi_\nu^2 \rangle = 1.06$,⁵ and the other deep measurements have similar results.

Total Intensity Calibration with CBIPOLCAL

CBIPOLCAL also calibrates the total intensity baselines. In this case, the instrumental polarization makes a negligible contribution to the visibilities, so we must simply isolate the gain term.⁶ The arguments that lead to Equation 3.6 also yield

$$\mathcal{V}_{jk}^{LL}(\mathbf{u}, \nu) = G \left[I(\mathbf{u}, \nu) + \epsilon_k^* P(\mathbf{u}, \nu) e^{-2i\phi} + \epsilon_j P^*(\mathbf{u}, \nu) e^{2i\phi} + \mathcal{O}(\epsilon^2) I \right] \quad (3.20)$$

where ϵ_j and ϵ_k are the instrumental polarization factors associated with antennas j and k . Since $\epsilon \sim P \sim 0.1$, the instrumental polarization is a 1% effect for the total intensity observations, so both CBICAL and CBIPOLCAL neglect the last three terms, and a single visibility suffices to obtain the LL calibration for both programs. The LL calibration precedes the LR calibration; CBIPOLCAL sorts the LL baselines to find simultaneous or nearly simultaneous matches for the LR baselines on the array, and performs the LL calibration to obtain the values for I which are necessary to remove the instrumental polarization from the LR visibilities. If the

⁵ $\nu = 10$ bands \times 11 LR baselines; RX5 was out for repair.

⁶ $\delta\mathcal{V} \sim 2\%$ for $\epsilon \sim 0.2$ and $P/I \sim 0.1$

RX11-RX12 (IF 9)




Figure 3.3: Circle fit example for RX11-RX12, 27-28 GHz on 23apr00

source total intensity flux density is known, the user can supply a model for that flux in lieu of forcing CBIPOLCAL to perform the *LL* calibration with the *imodel* flag. This feature suppresses the uncertainty in the leakage correction from the *LL* contribution; it is particularly useful for observations of calibrators, because we can simply use the CBI's *LL* observations to obtain the model for *imodel*. Modeling the *LL* distribution of visibilities for the deep fields is a significant challenge, however, so we do not apply this feature to the calibration of the deep field data.

Calibration with Extended Sources

At the outset of the CBI polarization campaign, 3C279 served as the primary calibrator. 3C279 is a point source, so its salient characteristics—its total intensity I , fractional polarization m , and position angle χ —are all uniform across the (u, v) plane, and the source can be completely characterized by these three quantities for each of the CBI’s ten channels. Early in the 20^h deep field observations 3C279 was no longer visible at night, but Tau A became visible 2.5 weeks later, at which point it became our primary polarization calibrator. Tau A is bright, highly polarized supernova remnant which is slightly resolved by the CBI, so CBIPOLCAL was modified to accommodate a model for Tau A’s morphology. Generally, any extended source can serve as a calibrator provided that the user supplies a model for $I(\mathbf{x}, \nu)$, $Q(\mathbf{x}, \nu)$, and $U(\mathbf{x}, \nu)$ which CBIPOLCAL can invert for application in the (u, v) plane. Sources which are comparable in size to the CBI’s $\sim 45'$ primary beam require an additional correction for the taper in the primary beam across the source. Since Tau A is only 4-5' across, however, failure to correct for the beam introduces an amplitude error of less than 1%, so CBIPOLCAL neglects this correction. The latter part of this chapter discusses efforts to derive a model for Tau A and to apply the model to the uncalibrated data.

Application of Calibration Factors to the Program Data

CBIPOLCAL stores the gains for all baselines and leakage terms for the cross polarized baselines in a *.cal* file. To calibrate program visibilities, the user supplies CBIPOLCAL with the program data and the appropriate *.cal* inputs. The program data must be corrected for the presence of instrumental polarization, so CBIPOLCAL first uses the *LL* gains to calibrate the *LL* baselines which are parallel to *LR* baselines. With $\tilde{I}(\mathbf{u}, \nu)$ in hand, CBIPOLCAL employs the gains and leakages for the cross polarized baselines from the *.cal* file to calibrate the cross polarized baselines. CBIPOLCAL inserts the G_{jk} and ϵ_{jk} into Equation 3.6, and inverts the expression to solve for the quantity of interest, $\tilde{P}(\mathbf{u}, \nu)$.

While we generally measure $\tilde{I}(\mathbf{u}, \nu)$ with high S/N for the calibrator sources,

the S/N per (u, v) point for the deep CMBR measurements is generally comparable to unity, so even with perfect measurements of the leakage factors, the instrumental polarization correction can introduce significant errors into the LR visibilities. We can reduce the uncertainty in $\tilde{I}(\mathbf{u}, \nu)$ by modelling the distribution of total intensity visibilities for the deep fields; the size of the autocorrelation of the primary beam is such that we generally oversample the aperture domain and this tends to mitigate the poor S/N per (u, v) point. The (u, v) plane model for I would also provide LL counterparts for the baselines which lack LL counterparts on the array, so we could incorporate all of the LR baselines in the analysis. Although this approach offers great promise, it was not applied to the data presented in this work.

3.3 Polarization Calibration Observations

We employed two primary polarization calibrators—3C279 and Tau A—during the CBI polarization campaign. Figure 5.1 provides a timeline of the CBI polarization calibration source observations. The 3C279 observations encompass the entire 08^h deep field, as well as many supporting observations which demonstrated the capabilities of the system. As noted above, during the 20^h field observations 3C279 ceased to be in view at night, and Tau A rose following a 2.5 week gap during which no primary polarization calibrator was visible.⁷ This section discusses the details of the polarization calibration observations of 3C279 and Tau A.

3.3.1 3C279 Observations

3C279 provided the primary polarization calibration from 06feb00 until it set below the CBI’s elevation limit on 21aug00. While this interval encompassed less than half of the deep field data which go into the final result, this eight month period spanned a number of important demonstrations of the CBI’s polarization

⁷3C279 overlapped with Jupiter, and although we hoped to transfer the calibration to Jupiter, our observations showed that Jupiter has less than 1 Jy of polarized emission at 1 cm; this flux density falls well short of the level of polarization necessary to permit fast calibrations, so we did not pursue this tack.

capabilities. The length of this period sufficed to provide an understanding of the long-term behavior of the instrumental polarization. This interval also included several tests: cross polarized beammaps, and supporting observations of galactic and extragalactic sources of known polarization.

3C279 has several qualities which make it a desirable calibrator; with $I \sim 25$ Jy and $P \sim 2$ Jy, 3C279 permits relatively fast calibrations. A sequence of 5^m observations suffices to obtain a several σ detection of $\sim 10\%$ instrumental polarization. In addition, the CBI cannot resolve 3C279, so uncertainties in the source's morphology do not undermine the calibration. The 3C279 observations consisted of at least one 5^m integration on the source followed by a trail to reject spillover. On rare occasions the calibration consisted of single scans which allow a measurement of either the gain or the leakage, but not both; on these occasions leakage measurements from other dates were supplied to obtain a measurement of the gain. Routine calibration observations were performed on a daily basis during deep field observations; these observations usually consisted of two scans on the source separated by 90° , which together permit a fit to a circle in the (Q, U) plane and thus a determination of the gain *and* the leakage. On several occasions we performed extended observations of 3C279 which consisted of as many as twelve integrations to obtain high S/N measurements of the instrumental polarization.

The reference polarization characteristics required by CBIPOLCAL were derived from a combination of VLA and CBI data for 3C279. The VLA observation campaign is described in Chapter 4 of this work; as discussed in that chapter, these high-frequency observations were often severely hampered by factors such as wind loading which reduced the efficiencies of the VLA antennas and ultimately left the reliability of the absolute flux density measurements from VLA data open to doubt. To first order, efficiency losses should affect I and P equally, so to circumvent these problems, only $m = P/I$ and $\chi = \frac{1}{2} \tan^{-1}(U/Q)$ were transferred to the CBI. The CBI's excellent absolute flux density calibration for total intensity aids the polarization effort at this juncture: the polarization flux density scale was obtained from the total intensity observations of 3C279 with the CBI, and this

flux scale, which is based on *cm*-wave observations of Jupiter, has an absolute uncertainty of better than 5%.

The uncertainties for the VLA data are discussed in Section 5.3, at which point the interpolations to the CBI observations are presented in detail. To summarize, two interpolations are necessary to transfer the VLA values for m and χ to the CBI. First, the CBI observed 3C279 nearly every session during the 06feb00-10aug00 interval, while the VLA observations yielded only eight good measurements interspersed throughout this period from which we can transfer the calibration.⁸ We must therefore transfer the VLA calibration to the CBI calibration across spans of ~ 1 month. Second, the two frequencies at which the VLA observations were performed—22.46 GHz (*K* band) and 43.34 GHz (*Q* band)—straddle the CBI’s 26-36 GHz band, so the values for m and χ from the VLA must be interpolated to the CBI bands.

We merged the m_ν and χ_ν data for 3C279 from the VLA with values for I_ν from the CBI. There are several reasons for this approach. First, the CBI measures $I(\nu)$ for 3C279 at the bands of interest, in contrast to the VLA, which provides only I_K and I_Q . Second, the CBI is optimized to operate at centimeter wavelengths, so the CBI observations are more reliable than those with the VLA. The absolute uncertainties of the CBI and VLA flux density scales are $\sim 5\%$, and coupled with the systematic errors from factors such as wind loading, the uncertainties which accumulate while transferring the VLA total intensities to the CBI inflate the VLA uncertainties to values much greater than those for the CBI.

Instrumental Polarization Measurements with 3C279

A solid understanding of the instrumental polarization is the keystone to a good polarization calibration, so we performed a variety of measurements to better understand its behavior. These observations consisted of multiple scans—between four and twelve—over a range of deck orientations separated by 45° ; Table 3.1

⁸A third of the VLA observations were lost to poor observing conditions and other site-specific factors at the VLA (Chapter 4).

lists the dates of these measurements and the interpolated characteristics that were applied to the calibration; the changes in m and χ reflect intrinsic changes in the source. The deep CBI observations do not necessarily coincide with the VLA observations, so the VLA data were interpolated to the CBI dates with a simple linear fit to the pair of VLA observations which straddle the CBI observations of interest.⁹ The errors on the polarization characteristics in the table reflect the uncertainties in the K and Q band data from the VLA, as well as the additional uncertainties due to the interpolation of the VLA data to the CBI channels. The deep observations yielded high S/N measurements of the CBI’s instrumental polarization, a sample of which is shown in Figure 3.4. The leakage terms can be quite large at the band edges—up to 20% in amplitude. We will see in the following section that this large instrumental polarization can result from bandpass errors in the phase shifter assembly.

The qualitative consistency between the sets of leakage terms shown in Figure 3.4 suggests a high degree of repeatability in the instrumental polarization. In particular, the array reconfiguration in the middle of April 2000 did not significantly change the leakage. The two week reconfiguration occurred midway through the observations of the 08^h deep field, and the reconfiguration was bracketed by careful measurements of the instrumental polarization: the observation on 10apr00 preceded the reconfiguration, and the observations on 23apr00 and 24apr00 followed it. A comparison of these leakage terms, provided by the two middle points in each cluster of points in Figure 3.4, demonstrates that the reconfiguration did not have a significant effect on the instrumental polarization.

A χ^2 analysis of the scatter in the leakages provides a test of the stability of the leakage terms. To obtain χ^2 , the real and imaginary components of the leakage for the eight deep observations in the 06feb00-17jun00 period were averaged with weights to derive a set of mean leakages. The real and imaginary components are

⁹The deep observations on 06feb00 and 09feb00 were not preceded by a VLA observation, so these dates use only the data from the VLA observation of 18feb00 with an assumed 5% uncertainty for m and χ .

date	conf	duration	VLA observations		m	σ_m	χ ($^\circ$)	σ_χ ($^\circ$)
06feb00	1	8×5^m	18feb00	—	0.094	0.005	62.0	3.1
09feb00	1	8×5^m	18feb00	—	0.094	0.005	62.0	3.1
10apr00	1	8×5^m	12mar00	05apr00	0.089	0.002	54.6	0.9
23apr00	2	8×5^m	23apr00	—	0.092	0.001	52.6	1.1
24apr00	2	13×5^m	23apr00	—	0.092	0.001	52.4	1.1
15jun00	2	6×5^m	02jun00	30jun00	0.114	0.004	52.8	3.6
16jun00	2	6×5^m	02jun00	30jun00	0.113	0.003	52.7	3.4
17jun00	2	6×5^m	02jun00	30jun00	0.113	0.003	52.6	3.2
13jul00	3	5×5^m	30jun00	08aug00	0.113	0.003	50.2	1.4
10aug00	3	4×5^m	08aug00	—	0.115	0.002	47.7	3.5
11aug00	3	3×5^m	08aug00	—	0.115	0.002	47.6	3.6
12aug00	3	4×5^m	08aug00	—	0.114	0.002	47.8	3.5

Table 3.1: Deep 3C279 observations. Columns 1, 2, and 3 list the dates, configurations, and durations of the deep CBI observations of 3C279. Columns 4 and 5 list the bracketing VLA observations, and columns 6-9 list the interpolated values for m and χ , with uncertainties. These errors reflect the uncertainties in the VLA observations after the interpolation to 31 GHz on the date in question. The deep observations of 15jun00-17jun00 were part of a series of beammap observations.

independent; we compute χ^2 for the real and imaginary components separately as

$$\chi^2 = \sum_{i=1}^8 \left(\frac{\langle \epsilon_i \rangle - \epsilon_i}{\sigma_{\epsilon_i}} \right)^2 \quad (3.21)$$

for each baseline and each channel. Several antennas (RX5, RX11) were not operational for the entire period, and another (RX4) underwent many cryogenic cycles which changed its instrumental polarization; these three antennas do not permit a meaningful comparison, so they were culled from the analysis. We augmented the uncertainties reported by CBIPOLCAL for the leakages with the errors on the VLA data listed in Table 3.1 prior to computing the weighted average: $\sigma_\epsilon^2 \Rightarrow \sigma_\epsilon^2 + \sigma_{VLA}^2$. The leakage measurement does not require a flux scale, so the CBI’s absolute flux density calibration uncertainty of 5% does *not* contribute to the uncertainties.

Figure 3.5 shows a histogram of the results of the χ^2 calculation. The χ^2 analysis provides strong, but perhaps not compelling, evidence for the repeatability of

the instrumental polarization during the 06feb00-17jun00 period. The centroid of the values of χ^2 is shifted well past the peak of $P_\chi(\chi^2; \nu = 7)$; the data redistribute power from the peak to the high- χ^2 shoulder of the distribution. The mean value of χ^2 for the nine antennas is 77.8, or $\chi_\nu^2 \sim 1.2$ per antenna. The χ^2 test shows that the uncertainties cannot account for the scatter in the data. An inspection of the leakage terms for the nine cross polarized baselines in this analysis does not reveal a particularly offensive baseline or band. The excess in χ^2 may reflect an error in the interpolation of the VLA data; an error in the uncertainties on the interpolations listed in Table 3.1; or a real systematic change in the instrumental polarization over time. The first two uncertainties can be reduced by matching the dates of the deep 3C279 observations with those of the VLA observations, although this strategy requires the cooperation of the weather at both sites. Systematic changes in the leakage can be better isolated with more deep observations.

The July and August deep observations listed in Table 3.1 are absent from the comparisons of Figures 3.4 and 3.5. These two sets of leakage terms are marginally discrepant with the means from the 06feb00-17jun00 period, and one possible cause for this systematic change in the leakage is a thermal cycle of RX12 on 30jun00 due to a compressor failure. While we are very concerned about a putative change in the leakage, this particular event does not directly affect the polarization observations presented in this work; however, because it occurred during a lull in the polarization observations following the completion of the first deep field.¹⁰ When the deep field observations resumed in August, we measured the leakage terms with 3C279 on 10aug00-12aug00 and applied this set to the 20^h observations. One lesson from this experience is that thermal cycles should be followed by careful measurements of the instrumental polarization; since the instrumental polarization is treated as a baseline-based quantity by CBIPOLCAL—that is, we let $(\epsilon_j + \epsilon_k^* = \epsilon_{jk})$ —this admonition applies to *all* of the receivers on the array, not just RX12.

¹⁰The maintenance history of the CBI during the 2000 season does not note any other thermal cycles for RX12.

CBI Instrumental Polarization Comparison, RX2-RX12




Figure 3.4: Instrumental polarization comparison for deep observations of 3C279. The ten clusters of points along the horizontal axis represent the CBI’s ten channels. The four points in each cluster denote the four sets of deep observations; neighboring dates were averaged to simplify the presentation. The phases of the latter two sets of leakages were adjusted to reference the factors to configuration 1. The error bars include the additional uncertainties introduced by the VLA measurements of m and χ . A visual inspection of the four points in each cluster shows that they are in good agreement, and a χ^2 analysis of the eight deep observations provides some quantitative support for this conclusion.

χ^2 for 8 deep measurements of the CBI instrumental polarization




Figure 3.5: Histogram of χ^2 for the eight deep instrumental polarization measurements in the 06feb00-17jun00 interval. The real and imaginary components of ϵ are independent and were treated separately but histogrammed *together* in this figure. The histogram is compared to a plot of $P_\chi(\chi^2; \nu = 7)$.

Comparison with Phase Shifter Model

The CBI's high instrumental polarization is a concern for the polarization observations. The instrumental polarization is comparable to the polarization of our calibrators and program sources, so small changes in the leakage can have a profound effect on our results. The model for the instrumental polarization developed in Chapter 2 guides our understanding. The model shows, for example, that the leakage varies with the orientations of the half-wave plates in the two phase shifters in each baseline; during the 20^h field observations, the orientation for the phase shifter in RX9 changed, and this resulted in a significant change in the leakage for baseline RX9-RX12. The model shows why we should expect this result.

Chapter 2 presented a model for the instrumental polarization which the phase shifter assembly can generate. In Section 2.3.3, we showed that

$$\begin{aligned}
 \mathcal{V}^{LR} = & \quad I \left[(-\delta A \cos[\Delta\phi] + \delta\alpha \sin[\Delta\phi])e^{-i\Sigma\phi} - \delta B \sin[2\Delta\phi] \right. \\
 & \quad \left. - \delta\beta \cos[2\Delta\phi] + i\Sigma\theta \cos[2\Delta\phi] - \Delta\theta \sin[\Delta\phi] \right] \\
 & + (Q + iU) \left[1 - \delta A - \delta B + i\Sigma\theta \right] e^{-2i(\psi + \phi_j + \phi_k)}
 \end{aligned} \tag{3.22}$$

The coefficient of the total intensity in Equation 3.22 shows the sources of the leakage: it arises from errors in the quarter and half-wave plates. Instrumental polarization is produced by insertion losses ($\delta A, \delta B$) and insertion phase errors ($\delta\alpha, \delta\beta$) in the quarter and half-wave plates, respectively, as well as orientation errors for the quarter-wave plates ($\Delta\theta, \Sigma\theta$). These contributions are weighted by the relative half-wave plate orientations ($\Delta\phi, \Sigma\phi$). The insertion characteristics are intrinsic to the plates and can be measured on the benchtop, while the quarter-wave plate orientation errors arise from random, unknown assembly variations and the half-wave plate positions are set by the control system. To assess the accuracy of the model in Equation 3.22, we assume values for the insertion characteristics of the plates derived from lab measurements, and fit the data for the four unknown plate orientations. Since the leakage is nonlinear in the unknowns, a gradient search

method was applied to find the best fit values for (ϕ_i, ϕ_j) and (θ_i, θ_j) . The leakage model used by CBIPOLCAL assumes that all spurious effects appear as leakage—it does not isolate the degradation to the polarized signal from the insertion loss—to compare the model to real data we divide the coefficient of I in Equation 3.22 by the coefficient of $(Q + iU)$. The errors are on the order of a few percent, so this factor is slightly less than unity.

The model requires accurate inputs for the wave plate insertion loss $(\delta A, \delta B)$ and insertion phase $(\delta\alpha, \delta\beta)$. System temperature measurements provide an estimate of the loss: the phase shifter contributes ~ 1 K to the ~ 15 K of noise seen in the receivers, so since the assembly has a physical temperature of ~ 10 K, its loss must be $\sim 4\%$. This loss is uniform across the band, and to first order it is apportioned in a 2:1 ratio between the half wave plate and the quarter-wave plate. The wave plate insertion phase contributes a frequency dependent error which gives rise to the strong achromaticity seen in many of the measurements of the instrumental polarization. Benchtop measurements of the quarter-wave plate show that at room temperature its insertion phase varies roughly linearly from -5° to $+4^\circ$ across the 26-36 GHz band (Figure 3.6). As the plate cools to cryogenic temperatures, two competing effects *increase* its insertion phase by $\sim 3\%$. Cooling causes the plate to contract by $\sim 2\%$, which decreases the electrical length encountered by the wave a similar amount. The contraction increases the density of the plate, however, which in turn increases the effective dielectric constant by $\sim 5\%$. The net result is to increase the insertion phase by $\sim 2.5^\circ$, so the values in Figure 3.6 were adjusted accordingly for the calculation; this change can be computed from the Kramers-Kronig relation [42, 60, 36].

Figure 3.7 provides a comparison between the phase shifter model and the leakage data for two baselines. The two baselines shown in the figure represent the extremes of the bandpass shapes seen in the data; RX3-RX12 has the shape seen on many of the cross polarized baselines; the leakage rises sharply from nearly zero at the low end to more than 20% at the high end, and in this regard it parallels the shape of the insertion phase error seen in Figure 3.6. The leakage for RX7-RX12

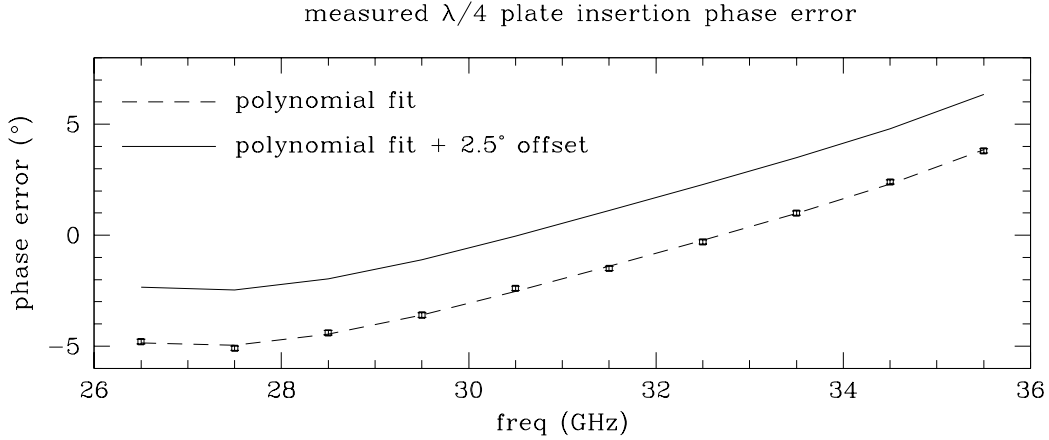


Figure 3.6: Insertion phase error for a quarter-wave plate at room and cryogenic temperatures.

is one of several exceptions to this trend, however, but as Figure 3.7 shows, the shifter model yields a good fit to the data for this baseline as well. The best-fit values for the half-wave plate positions, ϕ_1 and ϕ_2 , span a range of orientations for both baselines, while the best-fit values for the quarter-wave plate orientation $\delta\theta_1$ and $\delta\theta_2$ errors are all small, typically $\sim 1^\circ$, with a rare maximum of $\sim 4^\circ$. The former values are all well within the assembly tolerances for the quarter-wave plate sections, while the latter is perhaps slightly large. Nonetheless, the best fit quarter-wave plate orientation errors are modest, and the fact that the model does not require extreme values for the quarter-wave plate orientation errors provides confidence in this treatment.

This exercise demonstrates that the measured instrumental polarization is well within reach of the known phase shifter errors, so it provides insight into the factors which affect the repeatability of the instrumental polarization. If the shifter characteristics—the plate lengths, orientations, and the complex insertion loss of the teflon sections—remain constant, it is reasonable to assume that the instrumental polarization will as well.¹¹ The preceding section notes that the instrumental

¹¹The half-wave plate positions are archived with the CBI’s housekeeping data, and these registers were inspected to confirm that the shifter positions remained constant for the entire

polarization appears to have changed when RX12 warmed in late June 2000. The change in the leakage may arise from a shift in the zero point of the half-wave plate orientation during the thermal cycle.

Off-Axis Instrumental Polarization

The instrumental polarization measurements described in the preceding section sample the response of the system at the boresight of the primary beam. Observations of extended polarized emission, however, require an understanding of the response of the system across the entire primary beam; any departures from uniform behavior would affect the interpretation of the cross polarized visibilities for extended sources. These abnormalities might arise, for example, from the illumination pattern of the feeds. We measured the off-axis characteristics of the beams with a series of instrumental polarization observations of 3C279 at the beam half-power points in the four cardinal directions and compared these values to the instrumental polarization measured at the beam center. The initial incarnation of CBIPOLCAL operated under the assumption that the calibrator is a point source at the phase center, so we derived the leakages at the half-power points by forcing CBIPOLCAL to *assume* that the calibrator is at the phase center. Under this assumption, pointings at the beam half-power positions introduce phase errors in the visibilities, so a new command, `offset`, was added to CBIPOLCAL which allows the user to supply a position offset to correct the phases.¹² No analogous correction is required for the amplitude; since CBIPOLCAL employs the fractional polarization rather than to absolute polarization, the rolloff in the primary beam should affect P and I with equal force. In fact, the use of the fractional polarization eliminates an additional layer of complication: it allows us to ignore the $\sim 30\%$ change in width of the primary beam across the 26-36 GHz band. This approach relegates period between 11jan00 and 31oct00. During the 11sep00 to 31oct00 interval, an encoder error in RX9 changed the orientation of the half-wave plate, so this baseline was excised from the 20^h data set.

¹²To confirm the positions of the off-axis pointings, the total intensity data for these observations were mapped in DIFMAP.

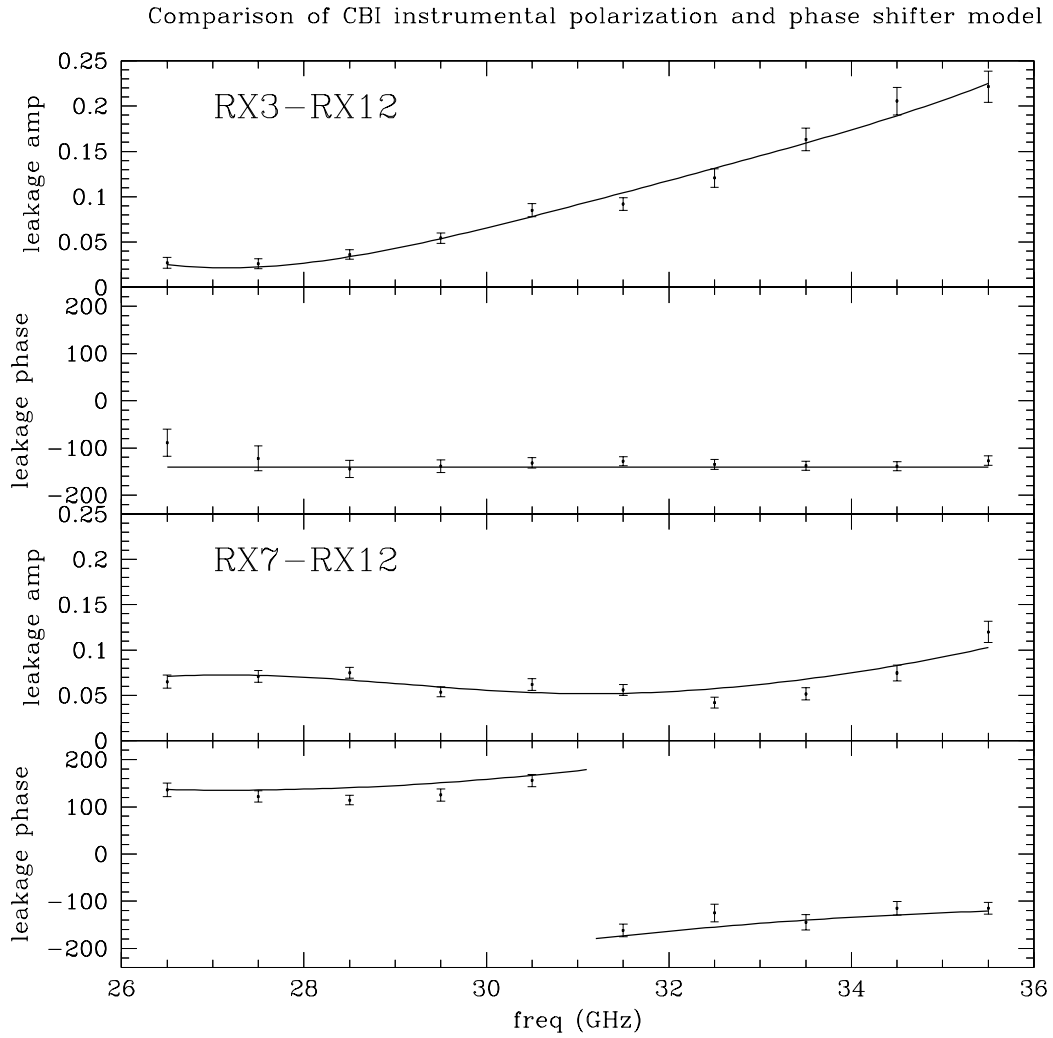


Figure 3.7: Comparison of the leakage model with the instrumental polarization measured on 06feb00 for two baselines. For each of the two baselines, the upper frame shows the leakage amplitude, while the lower shows the leakage phase. The amplitude for the leakage for RX3-RX12 resembles the insertion phase error shown in Figure 3.6, while that for RX7-RX12 does not. In both cases, however, the model yields a good fit.

position	15jun00	16jun00	17jun00	total
center	7	8	3	17
north	7	—	—	7
south	5	—	—	5
west	—	8	3	10
east	—	7	3	10

Table 3.2: Off-axis polarization observations of 3C279 in June 2000. The table summarizes the number of 4.5^m integrations at each position; the integrations were separated by 45° deck steps.

any variations in the polarization characteristics across the beam to the instrumental polarization; spurious polarization introduced by the beam will change the leakage relative to that at the boresight.

We performed a series of beammap observations in the middle of June 2000. The observations consisted of $\sim 4.5^m$ pointings at the beam center for eight deck steps followed by an identical sequence of observations at the half-power points; the observations are listed in Table 3.2. To reject spillover, the north, south, and east pointings were followed by trails, while the west pointing was preceded by a lead. Figure 3.8 shows the off-axis instrumental polarization for a sample pair of baselines, and Figure 3.9 presents a histogram of values of χ^2 for the real and imaginary parts of the leakage for all of the baselines. The data shown in Figure 3.8 are representative of all of the baselines; the leakage factors measured at the four cardinal half-power points are in good agreement with the leakage measured at the boresight. This conclusion is supported by the χ^2 test: the values of χ^2 agree with $P_\chi(\chi^2; \nu = 4)$, which is shown in the figure for comparison. The errors in Figure 3.8 and in the χ^2 analysis combine the measurement errors on the CBI observations with the uncertainties on m and χ from the VLA data; the two sets of errors are added in quadrature as shown in Equations 3.18 and 3.19. The measurement and interpolation uncertainties for the VLA data on 15jun00-17jun00 are $\sim 3\%$ for m and $\sim 6^\circ$ for χ (Table 3.1).

The preceding approach interprets a change in the polarization response across the beam as a change in the instrumental polarization. The beams are antenna-based quantities while the leakages are baseline-based quantities, however, so this is not a satisfactory approach for quantifying beam pattern anomalies. Changes in the polarization characteristics of the primary beam would change the $A_j(\mathbf{x} - \mathbf{x}_0)A_k(\mathbf{x} - \mathbf{x}_0)$ term in Equation 2.5 from a simple product of Gaussians to a shape which reflects deviations in amplitude and phase across the beams for antennas j and k , although in the absence of very high S/N measurements of the beamshape pathologies we would treat this quantity as a baseline-based correction. In the case at hand, however, there are no significant anomalies in the off-axis polarization characteristics of the beams, so we can avoid these complexities.

3.3.2 Tau A Observations

Tau A provided the primary polarization calibration for subsets of the 08^h and 20^h deep field observations. Tau A has the ingredients of an exceptional polarization calibrator; with $I \sim 350$ Jy and $P \sim 28$ Jy, Tau A permits fast calibration observations, and the physical extent of the nebula ensures that the emission is reliably constant at 1 cm. In spite of these advantages, however, the use of Tau A is complicated by the fact that the source is slightly resolved by the CBI; total intensity observations have shown that the spatial extent of Tau A’s total intensity is roughly Gaussian with a FWHM of $\sigma \sim 3.7'$ along its major axis. We might initially assume that Tau A’s polarized emission tracks the underlying structure in total intensity, but we will see that this is not the case—even at the CBI’s resolution limit. Since the accuracy of the gain and leakage calibration depends heavily on the accuracy of the model for the calibrator, the Tau A model is a paramount concern for the polarization calibration.

This section discusses the details of our efforts to obtain an accurate Tau A model from the CBI data. The primary impediment to this effort is the limited overlap between Tau A and 3C279; Figure 5.1 shows that on 11jan00 we performed a deep observation of Tau A, while our first deep observation of 3C279 followed




Figure 3.8: Off-axis instrumental polarization at four beam half-power points for a pair of baselines. The error bars include the uncertainties for the VLA interpolation.

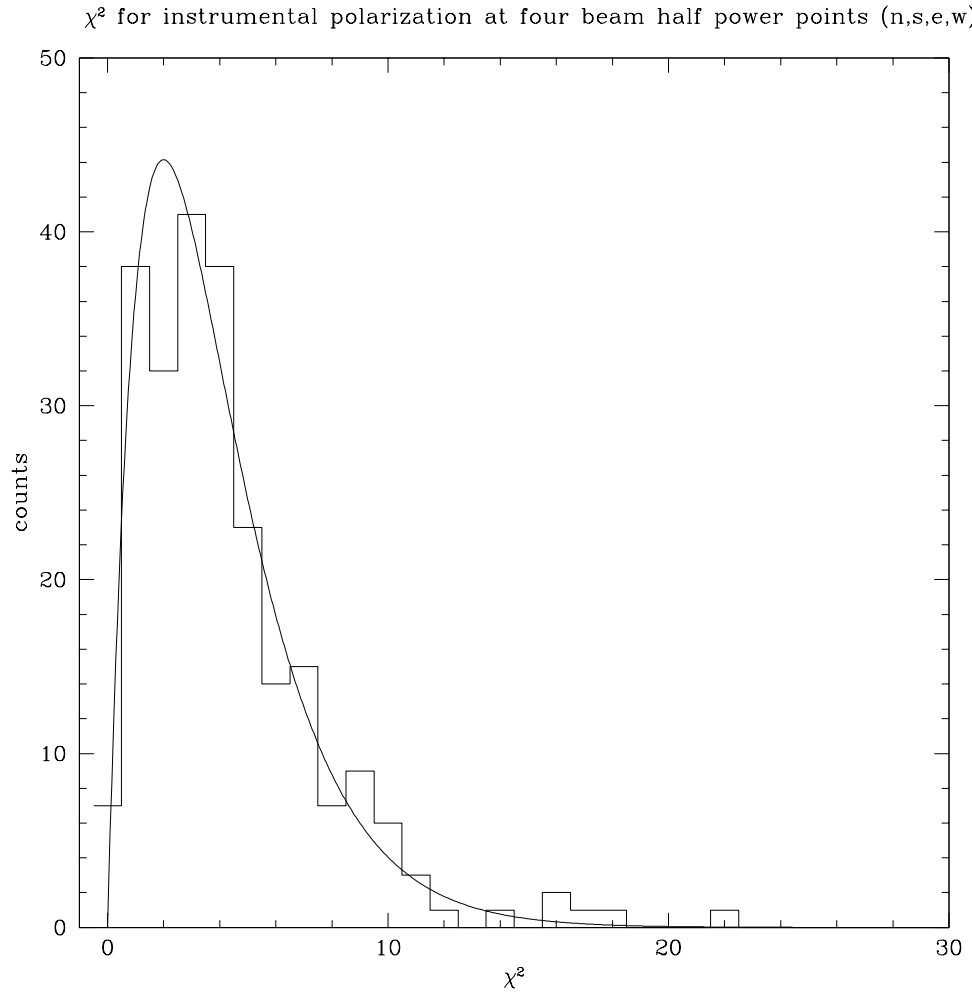


Figure 3.9: Histogram of χ^2 for the four off-axis beam half-power points. $P_\chi(\chi^2; \nu = 4)$ is shown for comparison. Although χ^2 was evaluated for the real and imaginary components separately, the resulting values are plotted together in this figure. The agreement with $P_\chi(\chi^2; \nu = 4)$ supports the hypothesis of a null change in the polarization characteristics of the beam at the four half-power points.

nearly a month later on 06feb00. There are several single integration observations of Tau A after 06feb00 which coincide with 3C279, but we require the deep observation at multiple deck angles—that of 11jan00—to derive a model.¹³ In this section we present the details of two techniques which were employed to transfer the calibration from 06feb00 to 11jan00; the two techniques yield the models shown in Table 3.6. The models agree in the polarization characteristics of the source by $\sim 4\%$, and while tests of the model presented in the latter part of this section suggest that the overall uncertainty is roughly the same, tests presented in Chapter 5 suggest that the overall uncertainty is $\sim 10\%$, and we adopt the latter uncertainty for the model. We begin this section with a discussion of changes to CBIPOLCAL which were required to use Tau A as a calibrator.

The first incarnation of CBIPOLCAL accommodated only unresolved calibrators such as 3C279, so our efforts to converge on an accurate Tau A model required several modifications to the software to accommodate the increasing degrees of freedom for the model. The initial Tau A model required minor changes to CBIPOLCAL. This model augmented the single Gaussian component for total intensity with a uniform fractional polarization m and a uniform position angle χ across the source. The Gaussian model for Tau A’s total intensity which CBICAL employs has seven degrees of freedom: integrated flux density (1), position (2: x and y), Gaussian shape (3: width, axial ratio, and position angle), and spectral index (1), to which we added m and χ . In practice we set the spectral index to $\alpha = -0.3$ for the total intensity model [52], and we do the same for the polarization model. This simple model for Tau A’s polarization failed spectacularly on two basic tests: the polarization characteristics of cross-check sources *and* the instrumental polarization inferred from the model both differed by as much as a factor of two from the expected values. These errors arise because the model incorrectly assumes

¹³At the time of these observations, the expectation was that the CBI deep polarization observations would be complete well before 3C279 set, so no effort was made to transfer 3C279’s calibration to Tau A. As it happened, however, poor weather pushed the schedule for the second deep observation to the beginning of August, at which time 3C279 set below the CBI’s elevation limits at night and Tau A was soon to rise.

that the maxima of the total intensity and polarized emission coincide; as we will see, observations at lower frequencies show that while the polarization is uniform across much of the source, it is relatively disordered at the northwest corner of the source, so the net polarization in that region is suppressed when averaged over the CBI's $\sim 4'$ synthesized beam.

The effect of the position offset between the I and P centroids can be demonstrated analytically. Consider a single component Gaussian with major and minor axes a and b ($a||\hat{\mathbf{x}}$) with offsets of (δ_x, δ_y) for the polarized emission. In the presence of instrumental polarization ϵ , the visibility is the Fourier transform of the model on the sky:

$$\mathcal{V}^{LR}(\mathbf{u}) \sim I_0 e^{[-\pi^2(u^2 a^2 + v^2 b^2)]} [m e^{i\pi(u\delta_x + v\delta_y)} e^{2i\chi} e^{2i\theta} + \epsilon] \quad (3.23)$$

We can estimate the fractional change of a cross polarized visibility in the presence of a typical offset. Subsume the deck orientation and field orientation into a single term χ , and consider just the part of the visibility which is sensitive to a position offset:

$$\mathcal{V}^{LR*} \propto [m e^{2i\chi} e^{i\phi} + \epsilon] \quad \text{with} \quad \phi = \pi(u\delta_x + v\delta_y) \quad (3.24)$$

$$= [m \cos(2\chi + \phi) + \epsilon_r] + i[m \sin(2\chi + \phi) + \epsilon_i] \quad (3.25)$$

The length L and orientation ψ of this vector are

$$L = m^2 + \epsilon_r^2 + \epsilon_i^2 + 2m[\epsilon_r \cos(2\chi + \phi) + \epsilon_i \sin(2\chi + \phi)] \quad (3.26)$$

$$\psi = \tan^{-1} \left[\frac{m \sin(2\chi + \phi) + \epsilon_i}{m \cos(2\chi + \phi) + \epsilon_r} \right] \quad (3.27)$$

Insert some typical numbers to get a sense of the magnitude of this effect. Let $m=0.06$, $\chi=30^\circ$, $\epsilon_r=0.08$, and $\epsilon_i=-0.02$. Consider a baseline of intermediate length, so let $(u, v) = (250, -300)$, and assume that the position offset of the centroid of

polarized emission is $\delta_x = 0.4'$, $\delta_y = -0.8'$, so that $\phi = 0.311$ rad. With these inputs, we find $L=0.010$ and $\psi=22.6^\circ$. In the absence of a position offset ($\phi \equiv 0$), however, $L_o=0.013$ and $\psi_o=16.2^\circ$. A typical position offset for the polarized source can clearly introduce a substantial ($\sim 30\%$) error—which translates into a comparable error in the gain and the leakage—so the Tau A model requires additional degrees of freedom.

In an early version of the Tau A model we tried shifting the polarized component relative to the centroid of total intensity. Shifts suggested by the literature failed to bring the leakage terms into agreement to better than $\sim 2\sigma$,¹⁴ so this approach was discarded in favor of a strategy which specified separate models for each of I , Q , and U .¹⁵ This approach proved to be satisfactory (see below).

3.3.3 Tau A Model 1

Figure 5.1 shows that on 11jan00 we performed a deep observation of Tau A, while our first deep observation of 3C279 followed nearly a month later on 06feb00. After 06feb00 there are only eight other dates in February and March of 2000 for which contemporaneous observations of Tau A and 3C279 permit a comparison of visibilities on all baselines. And as we will see, the structure of these observations is such that most are useful only for *checking* the calibration between the two sources, however; the one date for which the calibration can be transferred between the two sources is 06feb00, which also happens to be the date of the first deep 3C279 observation. The effort to reconcile Tau A with 3C279 hinges on this date, but fortunately we can ascribe a high degree of confidence to that calibration; because this is the date of a deep observation of 3C279, the uncertainty for the CBI calibration is limited only by the extrapolation of the VLA data from 18feb00 ($\sim 5\%$). As we see below, a blind calibration of the 11jan00 data using the calibration factors from 06feb00 yields a calibration which is surprisingly good in light of the

¹⁴In part because the beamshapes for the extant measurements did not match that for the CBI.

¹⁵While this strategy always loomed on the horizon, the modifications to CBIPOLCAL required to accommodate separate models for each parameter were sufficiently draconian that simple changes were explored exhaustively first.

nearly one month interval that separates them.

The CBI is generally a fairly stable instrument: the gain changes from night to night, as demonstrated by the evolution of the complex gains for the system, are typically 15% in amplitude and 8° in phase. Figure 3.10 shows the amplitude and phase corrections for a sample baseline for the 08^h field observations; in this case, the rms of the gain is $\sim 11\%$, while that for the phase is $\sim 6^\circ$. These changes are sufficiently small to permit the transfer of calibrations from one night to adjacent nights within this accuracy; this is particularly true if the gain variations are noiselike, in which case the errors add in quadrature. On long timescales, however, such as periods of several weeks, the gains have been observed to drift, and a naive application of calibration factors across these intervals will systematically corrupt the visibilities.

To calibrate the 11jan00 data with the 06feb00 calibration we must correct for the drifts in the gain during the intervening month; the noise calibration system aids us in this regard. The noise cal system refers all flux measurements to the power measured by the power meter, so under ideal circumstances, the noise calibration system should enable us to eliminate gain drifts to the accuracy of the power meter reading. In practice, however, the noise cal correction suffers from additional errors; the power measurement reported to the data stream by the power meter can diverge from the power which the system injects into the receivers because the insertion characteristics of the components downstream of the power meter change. These discrepancies force CBICAL to apply an erroneous correction to the visibility amplitudes during the ncal and ncal1 procedures. These gain errors, confined to the amplitude of the noise cal signal, must be rectified before the calibration factors from a particular night can be transferred across long spans.¹⁶

The errors in the power meter response can be measured by applying a set of

¹⁶There are also phase errors which are induced by temperature changes—driven by changes in the electrical lengths of cables with thermal expansion, for example—but these changes cannot be measured by the power meter, which only reports total power to the data stream. CBICAL's ncal procedure cannot use the power meter reading to fine tune the noise source visibility phases.

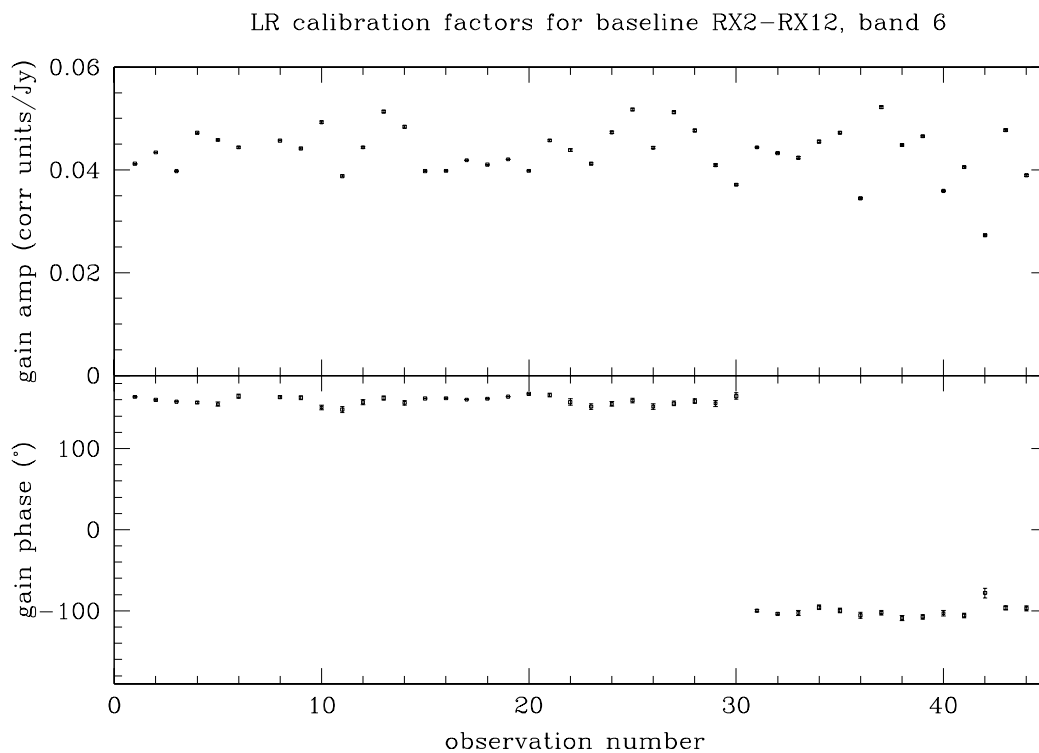


Figure 3.10: LR calibration factors for amplitude and gain for the 08^h field. These observations spanned many months, but for simplicity the dates have been condensed to chronologically increasing observation numbers which can be cross referenced with the dates given in column 2 of Tables 5.4 and 5.5. The average change in amplitude and phase for all baselines for this period is $\sim 15\%$. The break in the phases at day 30 (08apr00) is due to the reconfiguration; the reconfiguration changed the baseline length and orientation, and thus the phase.

calibration factors derived from an absolute calibration on one night to another night's observations. This approach should suffice to produce a reliable calibration *after* the ncal correction has been applied to both sets of data, but in practice this is not the case. The extent to which the calibrated visibilities depart from the correct values provides a direct measure of the error in the noise cal power measurement, and in this regard, the observations of the total intensity calibrators played a central role in isolating the noise cal errors. The scale factors deduced from the noise cal errors are then used to correct the polarization calibration factors.

The initial experiments with the techniques which evolved from this insight focused on a period in early August 2000 which suffered from a lack of polarization calibrators. The CBI obtained its total intensity calibration from Jupiter and Saturn during this interval until Tau A rose at the end of the month. Most of the dates in August which lacked direct observations of a polarization calibrator had a 5^m observation of at least one of Jupiter and Saturn. The deep observations of 3C279 on 12aug00 were used to derive a benchmark calibration for *LR* and *LL*; the factors for *LL* were applied backward to the Jupiter observations through the beginning of the month, and forward to Jupiter observations through the end of the month. The disparity between the total intensities for Jupiter as deduced from the 12aug00 benchmark calibration and Jupiter's true flux yields a correction to be applied to the calibration factors for all baselines for that night—the implicit assumption is that the error is confined to the noise cal power reported to the data stream by the power meter. In other words, the 66 *LL* baselines are used to measure the discrepancy between the noise source power injected into the system and the power reported to the data stream, and this correction is then applied to all 78 baselines.

Figure 3.11 shows the errors in the noise calibration system based on an *LL* observation of Saturn in early August. In this case, contemporaneous *LL* observations of Jupiter were used to infer the error in the power meter reading, and the resulting scale factors were applied to the Saturn observations. The two frames in the figure show that the blind calibration yields fluxes which are a few Jy ($\sim 10\%$)

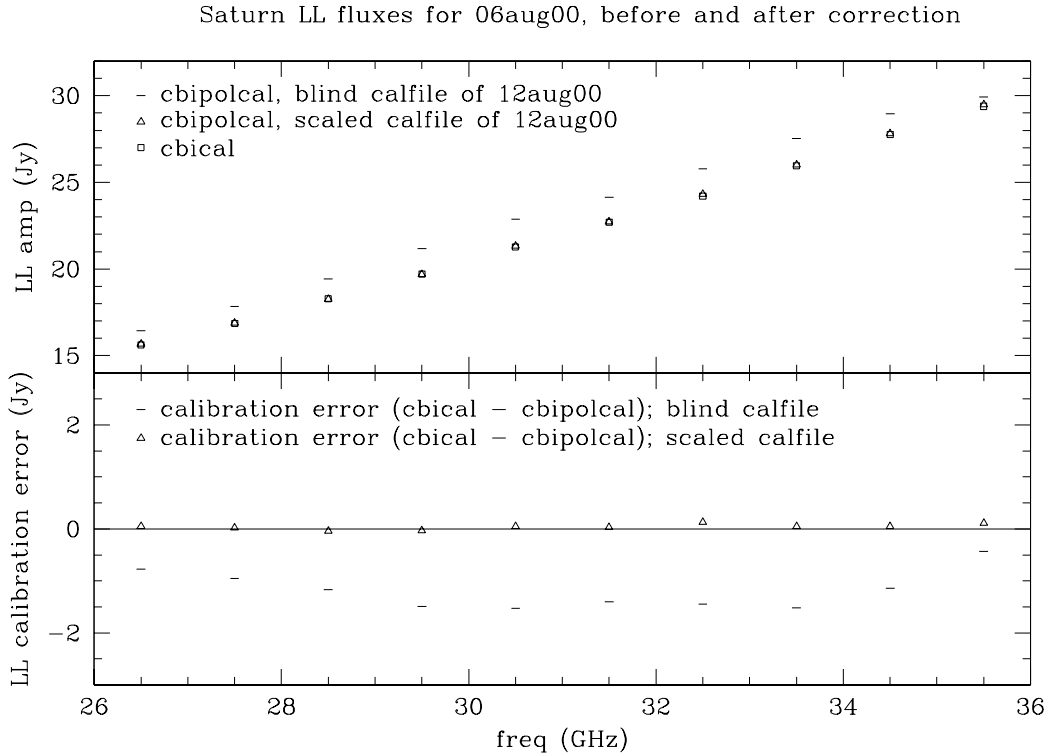


Figure 3.11: Calibration errors for 06aug00 LL visibilities. Errors are all ~ 100 mJy; they were left off of the figure for clarity. This result is typical.

more than the expected values. The scaling technique brings the LL flux densities into agreement at a level which is better than the errors on the observations; this is *precisely* the result we expect because the total intensity flux densities were used to infer the calibration error. The key feature of this analysis is that the scale factors show a tight quadratic dependence on temperature; this dependence is shown in Figure 3.12 for one channel during the 01aug00-03oct00 period. This dependence on the ambient temperature—rather than on antenna-dependent factors—provides confidence that the corrections to the LR gains will obey the same dependence.

Our goal is to recover the correct LR calibration. To test the scaling procedure we require two dates with deep 3C279 observations which are separated by enough time to allow the gains to drift. In the demonstration that follows, the deep 3C279 observation of 24apr00 is projected to the next deep observation, that of




Figure 3.12: Temperature dependence of the noise cal amplitude errors for the 31jul00–03oct00 period.

15jun00.¹⁷ A blind application of the 24apr00 calibration to the 3C274 observation of 15jun00 yields the total intensity flux densities shown in the first set of columns of Table 3.3. The second set of columns of the table shows the correct flux densities derived with CBICAL for comparison; the flux ratios show that the noise cal power error has undergone a change of as much as $\sim 20\%$ at the low end of the band. We performed deep observations of 3C279 on 16jun00 and 17jun00 as well, and the flux comparisons for these dates confirm these changes.

We are particularly concerned with the efficacy with which the scaling technique improves the polarization calibration. Upon applying the scale factors to all of the gains—including those for RX12—we find that the procedure improves the calibration for the band averaged values for I and $|P|$. Figure 3.13 shows these results. The top frame demonstrates that the scaling rectifies the large drift in the low end of the band, which in turn enables us to recover the total intensity to ~ 0.5 Jy, or $\sim 2\%$.¹⁸ The second frame shows the result of scaling on the polarization

¹⁷One other candidate pair of dates is 06feb00–10apr00, but the uncertainties in the interpolation of the 18feb00 VLA observations to the 06feb00 3C279 observations complicate the interpretation of the result: the 06feb00 observation of 3C279 with the CBI is not bracketed by VLA observations.

¹⁸The 2% discrepancy results from gain drifts which occur between the observations of 3C274 and 3C279 on each night.

$\Delta\nu$ (GHz)	CBIPOLCAL (Jy)	CBICAL (Jy)	ratio
26-27	16.70 \pm 0.07	18.45 \pm 0.05	1.105 \pm 0.005
27-28	15.78 \pm 0.06	18.02 \pm 0.06	1.142 \pm 0.005
28-29	14.85 \pm 0.07	17.57 \pm 0.05	1.183 \pm 0.007
29-30	14.31 \pm 0.07	17.06 \pm 0.05	1.192 \pm 0.007
30-31	14.12 \pm 0.06	16.70 \pm 0.06	1.182 \pm 0.006
31-32	14.14 \pm 0.06	16.27 \pm 0.06	1.151 \pm 0.006
32-33	14.59 \pm 0.08	15.93 \pm 0.07	1.092 \pm 0.008
33-34	14.68 \pm 0.06	15.59 \pm 0.09	1.062 \pm 0.008
34-35	14.62 \pm 0.07	15.22 \pm 0.08	1.041 \pm 0.007
35-36	14.63 \pm 0.08	14.91 \pm 0.06	1.019 \pm 0.007

Table 3.3: Amplitude scale factors, by band, required to map the 24apr00 calibration to 15jun00. First two sets of columns show the flux inferred for 3C274 by CBIPOLCAL and CBICAL, while the last shows the ratio of the two, which provides the scaling necessary to bring the two calibrations into agreement.

amplitude $|P|$; the scale factors bring the polarization into fair agreement with the benchmark, although they introduce a marginally significant tilt across the band. The virtue of the procedure is most pronounced in the band averaged quantities; based on the VLA observations, the mean polarization for 3C279 on 15jun00 is 3.09 Jy. The CBI measures 2.78 ± 0.04 Jy and 3.10 ± 0.04 Jy before and after scaling; the procedure reduces the band-averaged error from 11% to less than 1%. The polarization position angle χ is not affected, of course, because the flux ratios scale Q and U equally; the χ data are included in the figure as a cross-check, and to highlight an important point: the small position angle change of $\sim 6^\circ$ over a span of nearly two months is remarkable because it suggests that the cumulative phase drift over this time is of this size.

Tau A Model from Total Intensity Scaling

We now use the scaling technique to map the 06feb00 calibration on 3C279 to the deep Tau A observation of 11jan00. Table 3.4 presents scaling factors deduced from 3C274 observations on both dates; this table shows that the scale factors

Scaling demonstrating for 3C279: 24apr00 calibration applied to 15jun00




Figure 3.13: Scaling example. The scale factors were obtained from 3C274 observations on 24apr00 and 15jun00 and applied to the gains to map the 24apr00 calibration to 15jun00. The values for χ overlap because the scaling factors only affect the amplitudes. The error bars do not reflect the uncertainty in the VLA measurements, which are typically $\sim 5\%$.

required to bring the two calibrations into agreement are not large. Table 3.5 demonstrates the effect of the scaling on the total intensity component of the Tau A model; the table shows that without the scaling, the 06feb00 calibration recovers the total intensity to $\sim 3\%$, and that scaling reduces the error by half. While this improvement is modest, a substantial benefit appears in the inferred spectral index, which is shown in the last column of the table. In both cases the peak to peak span of the flux in the residual map after fitting to a single Gaussian component is $\sim \pm 2.5$ Jy/beam, which suggests that the fits—independent of systematic errors—are accurate to $\sim 1\%$. The first row of Table 3.6 shows the model components for

$\Delta\nu$ (GHz)	CBIPOLCAL (Jy)	CBICAL (Jy)	ratio
26-27	19.10 \pm 0.09	18.98 \pm 0.20	0.993 \pm 0.012
27-82	18.46 \pm 0.07	18.94 \pm 0.11	1.026 \pm 0.007
29-30	17.34 \pm 0.08	17.78 \pm 0.08	1.025 \pm 0.006
30-31	16.97 \pm 0.10	17.42 \pm 0.10	1.026 \pm 0.008
31-32	16.30 \pm 0.06	16.84 \pm 0.08	1.033 \pm 0.007
32-33	15.87 \pm 0.09	16.80 \pm 0.13	1.058 \pm 0.010
34-35	14.64 \pm 0.10	15.94 \pm 0.12	1.089 \pm 0.011
35-36	13.80 \pm 0.08	15.65 \pm 0.11	1.134 \pm 0.010

Table 3.4: Scale factors, by band, required to map the 06feb00 calibration to 11jan00. The first two columns show the flux densities inferred for 3C274 by CBIPOLCAL (blind calibration) and CBICAL (correct calibration). The ratios of these flux densities provide scale factors that have the tight quadratic temperature dependence seen in Figure 3.12.

all three Stokes parameters for Tau A, after scaling all of the gains. This model, *Model 1*, will be our working model for Tau A. We will focus on the errors in this model after presenting a second approach to deriving the Tau A model. foop

3.3.4 Tau A Model 2

We have a second avenue for obtaining a model for Tau A from the CBI data. While the only deep observation of Tau A occurs well prior to the deep observations of 3C279, we can use the handful of brief observations of Tau A early in the 08^h field campaign to infer the correct calibration for 11jan00. That we can even make these comparisons is by no means a foregone conclusion, however; since Tau A is an extended source, visibilities from different dates cannot be compared unless they sample the same (u, v) points. The deep 11jan00 observation consisted of 19 scans, starting at deck angle -95° and increasing in deck steps of 20° . In contrast, the nine subsequent Tau A observations consisted of single scans was performed at single deck positions, and only one was performed at a deck angle which matched one of the those on 11jan00. This exception is 06feb00, at which time Tau A was observed at d.a. 145° : a position which was sampled by the thirteenth scan on 11jan00.

component	F_0 (Jy)	x_0 (')	y_0 (')	σ (')	b/a	ϕ (°)	α
baseline	355.3	0.0	-102	3.58	0.66	-50.0	-0.30
before scaling	346.3	0.1	170	3.56	0.67	-49.8	-0.65
after scaling	360.8	0.1	83.9	3.55	0.66	-50.1	-0.28

Table 3.5: Gaussian model components for the total intensity of Tau A at 31 GHz based on the observations of 11jan00. The top row (*baseline*) shows the model deduced from the total intensity observations with the CBI; this model is accurate to better than 5%. The third set of columns shows the position of the centroid of the Gaussian (measured relative to the catalog position), while the fourth set of columns shows the FWHM, axial ratio, and the orientation of the Gaussian, and the fifth shows the spectral index of the component.

06feb00 is also the date of a deep 3C279 observation, so the polarization calibration on that date is among the most reliable. We can use these visibilities to infer the gain change for the cross polarized baselines between 06feb00 and 11jan00. These changes will differ from those determined from 3C274 in the previous section, but the results of the application of the inferred gain changes should be the same.

We would first like to assess the degree to which the 11jan00 visibilities for scan 13 on Tau A agree with those for the single scan on Tau A on 06feb00; this comparison will provide a sense of how much work is necessary to bring the 11jan00 calibration into agreement with that for 06feb00. The scaling procedure discussed in the previous section provides a gross estimate which suggests a slope in the power meter error across the band. Figure 3.14 shows a side-by-side comparison of a sample of the Tau A visibilities;¹⁹ this figure compares the visibilities for RX1-RX12, and the change between the two sets of visibilities reflects the evolution of the gains during the 11jan00-06feb00 period. The changes are clearly small: at

¹⁹The blind calibration first requires a correction to the phase for the relative *pointing* offset between the observations on 11jan00 and 06feb00, i.e., the two observations must be placed on the same intrinsic pointing scale. This correction is obtained by applying the 06feb00 calibration to 11jan00 to infer the relative position offset for I . The 11jan00 data are then exported from `cbical` a second time after this correction has been applied with `cbical`'s `shift` command.




Figure 3.14: Comparison of Tau A *LR* visibilities for 11jan00 (scan 13) and 06feb00, the former of which was calibrated with a simple application of the 06feb00 calibration. Channels 3 and 8 (29 and 33 GHz) were not installed during the 11jan00 observations, so these points are missing from the figure. This figure shows that on average, the blind calibration comes relatively close to the correct calibration.

most $\sim 20\%$ in amplitude and 10° in phase. They also depart from the scaling suggested by the band-averaged scale factors presented in Table 3.4 for several reasons: they do not employ 3C274 as an intermediate reference for the calibration; the scaled noise cal correction cannot rectify gain errors downstream of the power meter; the 06feb00 calibration has errors; and there is noise in the measurements on both dates. The change shown in this figure is representative of the changes seen on other baselines. Figure 3.15 compares the band-averaged visibilities for the nine baselines which are common to the two sets of data²⁰ plotted as a function of (u, v) radius to provide a sense of the change in the visibilities between the two dates. While the Gaussian shape of the visibilities is apparent, the points do not trace out a single Gaussian because they represent samples at a variety of

²⁰RX10 was added between the two dates, and thus is not included in the figure.




Figure 3.15: Tau A visibility comparison, 11jan00 (scan 13) vs. 06feb00, band ave. The upper frame shows the amplitude and the lower frame shows the phase. This figure demonstrates that the blind application of the 06feb00 calibration to the 11jan00 data comes within 15% of correct calibration for each baseline.

baseline orientations, and thus reflect different “slices” through the source; the two points at $|u|=200$, for example, which correspond to RX5-RX12 and RX7-RX12, were obtained from baselines which differ in orientation by 120° . On average, the 11jan00 visibilities at deck angle 145° differ in amplitude from those on 06feb00 by ~ 1 Jy, or $\sim 5\%$, and in phase by $\sim 2^\circ$. This good agreement—in the absence of any correction, apart from pointing—reflects the high stability of the CBI.

Tau A Model from Visibility Scaling

We can improve on the simple application of the 06feb00 calibration factors to 11jan00 by comparing the uncalibrated visibilities which sample identical (u, v) points to obtain an estimate of the change in gain for each baseline and channel.

In the notation which follows, we will let 11jan00 $\rightarrow \mathcal{V}_1$ and 06feb00 $\rightarrow \mathcal{V}_2$. Under ordinary circumstances we have the gains for both dates— G_1 and G_2 —and we can solve for the two polarizations P_1 and P_2 separately for each (u, v) point:

$$\mathcal{V}_1 = G_1(P_1 + \epsilon I_1); \quad \mathcal{V}_2 = G_2(P_2 + \epsilon I_2) \quad (3.28)$$

By selecting visibilities which sample the same (u, v) points, we force $(P_1 + \epsilon I_1) = (P_2 + \epsilon I_2)$, and we can then solve for G_1 , the gain for 11jan00.²¹ In the case of interest, we do not know G_1 , but we have the raw correlator output for both dates $(\mathcal{V}_1, \mathcal{V}_2)$, and since the deep 3C279 observations on 06feb00 provide G_2 , we can solve Equation 3.28 for G_1 :

$$G_1 = G_2 \frac{\mathcal{V}_1}{\mathcal{V}_2} \quad (3.29)$$

for each channel and each baseline. This approach transfers the noise in the raw visibilities to G_2 , and while a ten channel average would suppress the noise and thus provide a substantial improvement in the solution—to first order the gain drifts are uniform across the entire band, so a piston correction can compensate for much of the drift—in practice, the steep slope in the uncalibrated visibilities across the band negates any improvement because the correction is dominated by a few channels at the band edge. The slope across the band is such that the visibilities at the band edges ($|\mathcal{V}_{\nu=35.5 \text{ GHz}}|$ vs. $|\mathcal{V}_{\nu=26.5 \text{ GHz}}|$) differ in amplitude by a factor of ~ 10 , and often more.²² As a consequence of working in correlator units, the band averaged values for \mathcal{V} prior to calibration are dominated by the visibilities at the high-frequency end of the band. We are thus forced to perform the scaling on a channel-by-channel basis.

The complex ratio of the 06feb00 and 11jan00 visibilities was used to scale the

²¹We assume that the instrumental polarization ϵ is constant, and that the total intensity which shines through ϵ remains constant as well. The `imodel` command in `CBIPOLCAL` ensures that the correct total intensity model is applied to the instrumental polarization subtraction.

²²This steep slope is an artifact of the rolloff in the noise cal power; since the noise cal is much weaker in the highest frequency band, the `ncal` correction in `CBICAL` boosts the highest frequency amplitudes by a larger factor to set all of the noise cal visibility amplitudes to the same value.

06feb00 gains. The mean ratios of the correct and incorrect gains for the reals and imaginaries are 0.96 and 0.98, respectively. These ratios are qualitatively inconsistent with the finding of the previous section that the gains had drifted downward by 5% across the band; since these ratios focus on the baselines and channels of interest, they are more precise than the band averages which were used to obtain Model 1, but this precision is degraded by the higher noise in the single channels. This approach offers an independent check on the results of the scaling technique employed to obtain Model 1, so we pursue it below.

The calibration scaling procedure enables us to obtain a second Tau A model from the deep 11jan00 observations. Upon fitting the I , Q , and U components of the 11jan00 Tau A observation to single Gaussians in DIFMAP, we obtain the model components for *Model 2* shown in the second row of Table 3.6. The spectral index was held to $\alpha = -0.3$ for all the components, although relaxing this requirement did not improve the fit substantially. The total intensity component shown in the table is the model derived from the total intensity calibration for the CBI—the fit to the total intensity obtained from the scaling has $\sim 3\%$ less integrated flux, and this discrepancy provides one measure of the accuracy of the scaling procedure. The position offsets for the Q and U components (columns 3 and 4) are qualitatively consistent with each other as well as the inferred position offsets which originally precipitated this exercise. The centroids for Q and U differ because the position angle of the polarization changes across the source. The model components for Q and U are essentially unresolved. Several types of models were tried, including multiple Gaussians for each of Q and U , but these more complicated models did not substantially improve the fit; this is to be expected, as Tau A is just at the edge of the CBI’s resolution limit.

This section presented two models for Tau A, both of which were obtained from different calibrations of the 11jan00 Tau A observations with the CBI. The difference in the polarization of the two models is marginal; Model 1 has $P \sim \sqrt{Q^2 + U^2} \sim 27.5$ Jy and $\chi \sim -28^\circ$, while Model 2 has $P \sim 28.2$ Jy and $\chi \sim -29^\circ$. The two approaches share some of the same systematic errors; both approaches

model	component	S (Jy)	x_0 (')	y_0 (')	σ (')	b/a	ϕ (°)	α
1	I	355.3	0.0	-102	3.58	0.66	-50.0	-0.3
	Q	15.5	-48.1	15.7	3.56	0.67	-49.8	-0.3
	U	-22.7	-29.8	25.5	3.55	0.66	-50.1	-0.3
2	I	355.3	0.1	-102.0	3.58	0.66	-50	-0.3
	Q	14.9	-48.8	14.9	2.93	0	83	-0.3
	U	-23.9	-30.1	26.2	2.28	0.52	56	-0.3

Table 3.6: Gaussian model components for two Tau A models at 31 GHz based on the observations of 11jan00.

hinge on the 3C279 calibration on 06feb00, which in turn is calibrated using the 18feb00 VLA observation, an observation whose extrapolated uncertainty we take to be 5% at 31 GHz. The methods used to derive the calibrations on these dates cause the errors to diverge, however, because the gain scaling technique (Model 1) relies on the extent to which the scale factors derived on LL observations can be applied to the LR baselines, while the visibility scaling technique (Model 2) is limited by the uncertainties in the visibilities. As a practical matter, the models are very similar—smaller than the errors in the techniques—so we will use Model 1 as the working model for Tau A. Figure 3.16 shows the CBI’s map of Tau A for the calibration which yielded Model 1. In the following section we explore the errors in this model.

Tests of the Tau A Model

The magnitude of the scaling factors required to map the 06feb00 calibration to 11jan00—typically a few to 20%—suggests a rough upper limit on the uncertainty of the Tau A model of $\sim 20\%$. We can refine the accuracy of this estimate of the uncertainty by testing the Tau A model against known calibration metrics. The best test of the model is to compare 3C279 observations calibrated on Tau A with 3C279’s own calibration. The limited overlap of these two sources limits the




Figure 3.16: Polarization map of Tau A based on the CBI data of 11jan00 (Model 1). The polarization at the peak of the total intensity is 22.3 Jy/beam (9.2%) and the position angle is -30° . The fractional polarization peaks at 12% 2' south of the centroid.

application of this test, but there are nine dates early in the 08^h field observations which can be used.²³ We can also use the measurements of the instrumental polarization to estimate the error in the model; the deep Tau A observation of 11jan00 suffices to produce a very high S/N measurement of the instrumental polarization, which can be compared to the careful leakage measurement with 3C279 on 06feb00. This section discusses these tests.

Nine nights early in the 08^h field observations included both Tau A and 3C279, and we used these observations to test the model. On these nights the Tau A observations consisted of single 5^m scans at the beginning of each night, while the 3C279 observations consisted of two scans separated by a deck rotation of 90° at the end of each night. The 3C279 data were extracted and calibrated in CBIPOLCAL with the Tau A model;²⁴ these data were compared to the polarization characteristics interpolated from the VLA observations. The total intensity calibration for the CBI is accurate to 5%, so we focus on the the scale-free quantities m and χ . Figure 3.17 shows the results of the test on 3C279; the upper frame compares the fractional polarization amplitudes, while the bottom frame compares the position angles from the two sets of data. The error bars on the data reflect only the uncertainties in the polarization characteristics of the source; for the VLA data, the uncertainties include raw measurement errors and additional uncertainties from the two interpolations necessary to reach the CBI channels on the dates in question. For the CBI data calibrated with Tau A, the uncertainties are simply the measurement uncertainties. For clarity, the global systematic uncertainty of the CBI’s Tau A model—as deduced from the 11jan00 observation—has been neglected.

The figure shows that the Tau A model recovers the source polarization obtained with the VLA to a very high degree of precision. The model recovers the position angle to $\langle \delta\chi \rangle \sim 3^\circ$, or $\sim 5\%$ of a radian. The observations of Tau A were all made at a variety of deck positions: at $\phi = 145^\circ$ for 06feb00; at $\phi = -90^\circ$ for

²³There several more dates during the 08^h field observations during which both sources were observed, but these data were struck due to small lunar elongations or the lack of trails.

²⁴We assumed a set of leakage factors from the deep 06feb00 observation.


m and χ for 3C279; comparison of VLA results with Tau A model

Figure 3.17: Comparison of 3C279 values for m and χ from the VLA with those obtained from band-averaged CBI observations after calibration with the Tau A model. The error bars reflect only the uncertainties in the polarization characteristics of the measurements; the VLA uncertainties are derived from the intrinsic uncertainties in the VLA measurements, along with the uncertainties incurred while interpolating to the CBI band. There are additional uncertainties which stem from the LL calibration of $\sim 5\%$ which affect the overall polarization calibration, but these uncertainties do not enter this comparison. Since the measurement of m is in the high S/N limit, the correction for the bias in P from the noise (i.e., $P' = \sqrt{P^2 - \sigma_P^2}$) has been neglected. For clarity, the two sets of data are offset with respect to each other; the data calibrated with the Tau A model follow the VLA data. The larger error bars for the later set of VLA data reflects the greater uncertainty in the 12mar00 observations.

09feb00; at $\phi = 130^\circ$ for 10feb00; at $\phi = -95^\circ$ for 11feb00; and at $\phi = -90^\circ$ for 01mar00-05mar00. These measurements therefore sample different components of the model, so a model error would not necessarily produce a uniform systematic offset for the dates shown in the figure. Since the offsets are generally smaller than the overall uncertainties in the calibration deduced from the VLA data, which tend to be $\sim 10\%$, no effort was made to use this comparison to fine tune the model. In practice, we cannot achieve the level of precision suggested by this figure because the calibration loop does not always recover the flux in total intensity, which in turn affects the inferred polarized flux. This uncertainty is in addition to the 5% absolute error in the in the LL calibration. Considering that the receiver gains can drift by as much as $\sim 5\%$ between the Tau A observation at the beginning of the night and the 3C279 observation at the end, the agreement is good.

Another test at our disposal is a comparison of the instrumental polarization inferred from the Tau A model on 11jan00 with that which is derived directly from the 3C279 observations on 06feb00. Figure 3.18 shows this comparison for two baselines. The figure shows that for the two baselines under consideration, the amplitude and phase of the instrumental polarization agree; the results for other baselines are similar. The two baselines shown in the figure, RX2-RX12 and RX6-RX12, were selected because they represent a variety of lengths; RX2-RX12 is a 173 cm baseline, so it samples $\sim 16'$ scales, while RX6-RX12 is a 458 cm baseline, so it resolves structure at the $\sim 4'$ level. These resolution considerations are important because the shorter baselines are more likely to mask errors in the model.

The discussion in Section 3.3.2 shows that large errors in the Tau A model would cause the leakage inferred from the model to depart from the benchmark values determined by the deep 3C279 observation on 06feb00; this figure, and the corresponding comparisons for the other baselines, show that this is not the case. Since the model errors would appear as errors in the leakage, we may infer from Figure 3.18 that the Tau A model is a good approximation to the source.

Taken together, these two tests provide confidence in the Tau A model obtained




Figure 3.18: Comparison of the instrumental polarization obtained with 3C279 on 06feb00 with the leakage measured on 11jan00 with the Tau A model for all ten bands. Two sets of leakage factors are shown; the upper pair of frames shows the amplitude and phase for RX6-RX12, while the lower pair shows data for RX2-RX12. Channels 3 and 8 had not been installed on the array at the time of the 11jan00 observations, so these data are missing from the comparison. The two sets of data are offset for clarity. The figure shows the high consistency between the leakage factors determined with the two approaches; comparisons for other baselines show a similar level of agreement.

from the scaling procedure; the tests suggest that the Tau A model is within striking distance of the correct model, and certainly, the errors in the model are comparable to the errors which stem from the application of the VLA data to the CBI observations. The cross-check of the 3C279 observations in early February and early March show that the Tau A model suffices to recover the expected polarization for 3C279 to within the overall uncertainty of the 3C279 values from the VLA, although this test glosses over errors for individual bands and baselines. The instrumental polarization comparison provides a more detailed test of the model, and it shows that the channel-by-channel discrepancies are small—certainly well within the errors on the measurements. The Tau A leakage test measures an instrumental polarization which produces the equivalent of a sky polarization of $5\% \times 350 \sim 20$ Jy, which is an order of magnitude larger than the polarization of 3C279—we therefore obtain a higher S/N measurement of the leakage than we do of the intrinsic polarization of 3C279; this favors the second test. These tests provide confidence in the CBI’s Tau A model, but we will see in Chapter 5 that some significant uncertainties remain in the calibration which may be a consequence of the Tau A model.

3.3.5 Previous Tau A Polarization Observations

The literature contains many discussions of Tau A’s polarization at *cm* wavelengths. Most of these data are not readily accessible to comparisons with the CBI data, however, because differences in the observations affect the outcome of the observations. We will see that resolution effects, for example, can rapidly change the inferred polarization. Polarization observations are quite difficult to interpret in isolation, even in the best circumstances. Bietenholz and Kronberg, for example, have made exquisite VLA high resolution images of Tau A in which they measure a peak fractional polarization of 20% at 4.9 GHz; they note that this polarization falls well below the theoretical maximum of 65% for a source of synchrotron emission with spectral index $\alpha = -0.26$ [6]. This finding is typical of the difficulties in interpreting observations of Tau A’s polarization. Further com-

plicating the matter, the Tau A observations with the CBI must contend with a poor calibration. We consider these issues in detail in this section.

Hobbs and Hollinger [31] have observed Tau A with a $2'$ beam at 2 cm at the NRAO 140' telescope. They find that the peak fractional polarization is $11.6 \pm 0.5\%$, and it coincides with the peak of the total intensity; at that point the position angle is $\chi = 152 \pm 2^\circ$, although within a beamwidth of the peak the position angle swings between 117° and 159° . These orientations are in good agreement with those measured with the CBI, which found a slightly lower fractional polarization (11%) and a similar position angle (150°). In contrast, the peak of polarized emission in the CBI observations is offset from that for the total intensity, and perhaps the superior resolution of the Hobbs and Hollinger measurements failed to reveal this feature because it does not persist at their lower frequency.

Mayer and Hollinger [51] mapped Tau A's polarization at 1.5 cm with $1.7'$ resolution with the NRAO's 140' telescope. The authors find that the total intensity spans a region which is roughly elliptical in extent with half-widths of $4'.2 \times 3'.0$, while the polarized emission is confined to a much smaller area; the polarization extends across a region only $2'.0 \times 2'.5$ across which peaks $0'.25$ east and $0'.5$ south of the centroid of total intensity. The authors measure a peak fractional polarization of $\sim 16\%$, at which point the position angle is $\sim 154^\circ$. Mayer and Hollinger report the polarization for a matrix of $0'.8$ pixels around the centroid, and from these values the offset for the polarization is apparent; according to their measurements, the polarization rises from 2.8% through 13.2% to 16.0% along a diagonal of three pixels through the center of the source from the northwest to the southeast. Here once again we see that the position angle agrees with the angle measured by the CBI, while the fractional polarization does not. The offset for the peak of polarization agrees with the offset measured by the CBI. The peak fractional polarization measured by the two telescopes, however, differs by over 50%.

Flett and Henderson [22] observed Tau A with a 1.5° beam at 9 mm with the Chilbolton 25 meter dish. Their results are not presented with the same

detail as those discussed above, but they do confirm the position offset. Since this measurement is closest in frequency to the CBI observations, this agreement may well be significant. The authors find that the centroid of polarized emission is at a position $\sim 0.5'$ southeast of that for the total intensity, at which point the fractional polarization is $\sim 17\%$ and the position angle is $\sim 155^\circ$. This fractional polarization is nearly twice as large as the peak observed in the CBI maps; much of this difference arises from the contrast in beamsize between the CBI ($\Omega_s \sim 5.6$) and the Chilbolton telescope ($\Omega_s \sim 1.5$). The Chilbolton observations provide qualitative support for some of the CBI results, but these observations are not amenable to a quantitative comparison.

The factors which complicate the comparison of Tau A polarization observations reported in the literature to the CBI data prompted us to seek calibrated visibility data which we could smooth to the CBI beam. Bietenholz and collaborators [5] have made an extensive study of Tau A at GHz frequencies with the VLA, and they have generously made their 4.885 GHz visibility data available for this work. These data were obtained in all four of the VLA configurations, but we excised the long baseline data because the corresponding visibilities make only a marginal contribution to the total flux. The Bietenholz data were used to make maps for I , Q , and U ; Figure 3.19 shows the map of the source restored with a $0.5'$ beam. The Bietenholz maps confirm the offset between the total intensity and the polarized emission; their maps show that the fractional polarization peaks at $\sim 20\%$ roughly $0.25'$ east and $0.5'$ south of the centroid. This position agrees with the centroid found by Mayer and Hollinger, but the peak fractional polarization is significantly higher—no doubt due in large measure to the VLA's smaller beam. At the peak of polarized emission the position angle is 138° , which is $\sim 15^\circ$ smaller than the position angle measured at the peak by Mayer and Hollinger. Perhaps most importantly, the high resolution map provides a clear picture of the change in polarization position angle across the source; the polarization position angle wraps from -40° near the peak to 40° in a region $\sim 1'$ northwest of the center. The CBI beam averages the polarization over the entire source.

The Bietenholz data allow us to explore the effect of smoothing with the CBI beam. The cleaned I , Q , and U data were restored in DIFMAP with a Gaussian $5.6'$ beam to approximate the beam produced by the LR baselines in configuration 1; the resulting map is shown in Figure 3.20. Smoothing reduces the peak polarization from 20% to $\sim 7\%$, while the position angle at the peak rises to $\sim 142^\circ$. The larger beam spreads the peak polarization around a plateau $\sim 1'$ east and $\sim 2'$ south of the centroid of total intensity. On the opposite side of the total intensity centroid, the polarization rises to $\sim 4\%$. This exercise demonstrates the effect of beamsize on the inferred polarization: smoothing over the beam reduces the peak fractional polarization in the high resolution Bietenholz maps by nearly a factor of three, while position angle remains nearly unchanged.

The smoothed Bietenholz maps permit a comparison with the CBI observations of 11jan00. The fractional polarization in the CBI maps is generally higher than that in the Bietenholz maps. At the peak of the emission in total intensity, the CBI measures a fractional polarization of 8.9% and a position angle of 151° , while the corresponding values for the smoothed Bietenholz data are 6% and 142° . In the vicinity of the peak of polarized emission, the fractional polarization in the CBI data rises to $\sim 11\%$, while the position angle remains nearly unchanged. The fractional polarization in the smoothed Bietenholz maps also increases by 1% at the polarization peak, while the position angle remains roughly constant at 142° . Thus, even when the beamsizes agree, the 4.885 GHz VLA data yield a fractional polarization which is $\sim 50\%$ lower than that which is measured by the CBI. There must therefore be additional frequency-dependent effects which change the polarization between 4.885 GHz and 31 GHz. At the outset of the experiments with the Bietenholz 4.885 GHz observations we had hoped that these high quality maps would yield a model which could be applied to 31 GHz CBI data after the appropriate scaling. The fractional polarization seen in the smoothed data show that this is not the case.

We are confident that the Tau A model is substantially correct. We presented two techniques for extrapolating the calibration of 06feb00 to 11jan00, and the

Clean LL map. Array: 1987 May 26
 CRAB at 4.885 GHz 1987 May 26




Figure 3.19: Tau A, 4.885 GHz, resolved. Note that the scale of the image differs from that of 3.20 by a factor of two.

models that result from these techniques differ in polarized flux by $\sim 2\%$. The comparison of the instrumental polarization derived from Model 1 with known data suggests that errors in the model are not significant, and the tests on observations of 3C279 support this conclusion. The literature do not contain data which can confirm our measurements of Tau A's polarization. We will revisit the Tau A model in Chapter 5, at which point we will see that the model is accurate to $\sim 10\%$.




Figure 3.20: Same as 3.19 after smoothing with a 5.6° beam to approximate the CBI maps.

3.4 Calibration Error Budget

The deep CMBR observations are calibrated on 3C279 and Tau A. Figure 3.21 shows the uncertainties in the VLA data; these uncertainties reflect the raw VLA uncertainties as well as the additional effects of the interpolation to 31 GHz on the dates of the CBI observations. The shaded regions denote times during which the 08^h deep field was observed; the figure shows that the uncertainties in the VLA data are typically a few percent for most of the 08^h deep field observations. The Tau A model has an uncertainty at least 5%, the uncertainty we ascribe to the

06feb00 3C279 observation, and the techniques used to determine the calibration on 11jan00 increase this uncertainty further. In Chapter 5 we use supporting observations to better constrain the errors on this model.




Figure 3.21: VLA interpolation errors during the 08^h field observations. The dates start with 01jan00=1. The shaded regions denote times during which the 08^h field was observed. The sharp rise at the end of May is troubling, but it affects only $\sim 3\%$ of the 08^h deep field data.

Chapter 4

VLA Observations

4.1 Introduction

Observations of 3C279 with the VLA provided the foundation for the polarization calibration of the CBI. 3C279 combines all of the major requirements for a polarization calibrator: it is bright, highly polarized, and unresolved by the CBI and most configurations of the VLA at centimeter wavelengths; it is our good fortune that 3C279's declination of $\delta = -5^\circ$ puts the source within view of both telescopes for much of the year. 3C279 is variable, however, so a calibration program with 3C279 requires regular VLA observations to characterize the changes in the source. We were awarded VLA time to monitor 3C279 for the duration of the CBI polarization campaign. The monitoring program combined two challenging frontiers of VLA performance—high-frequency photometry and polarization—and much of the discussion in this chapter focuses on efforts to unravel the complications which accompanied these observations. Section 4.1 describes the observations, and Section 4.2 discusses the analysis and interpretation of the VLA data.

4.2 Observations

The VLA is a 27 element interferometer sited near Socorro, NM. The VLA is an extraordinarily flexible instrument; the array can be configured to permit baselines from tens of meters to tens of kilometers, and users can choose from any of 8 bands

between ~ 70 MHz and nearly 50 GHz. The VLA lacks channels which overlap with the CBI's 26-36 GHz band; the nearest VLA bands are a pair of bracketing channels, 22.46 and 43.34 GHz (K and Q band). These are the highest bands at which the VLA operates, and the monitoring observations were subject to the gain and pointing problems which arise while observing at the limit of the VLA's capabilities. The VLA cycles through four major configurations every 16 months, from the sparse A array to the compact D array; our monitoring observations encompassed B , C , and D configurations. At K and Q band, the VLA employs dual circular polarization, dual sideband receivers, and the backend computes all possible correlations— RR , RL , LR , & LL —for each of the two IFs simultaneously. These receivers reside at the Cassegrain focus, and when the user commands an observation at a particular frequency, the subreflector at the prime focus nutates to illuminate the receiver which is tuned to the desired band.

The VLA observations spanned two periods: September 1998 to December 1998, and December 1999 to August 2000. The 1998 epoch consisted of five observations over 90 days; these observations provided a test run for the monitoring program. The changes in 3C279's flux density during this period suggested that the program would require biweekly observations, and complications which arose during the analysis of these data provided insights into how to structure the observations. The second period consisted of thirteen observations over ten months. These observations encompassed most of the CBI polarization campaign, and in so doing provided the reference for all of the polarization observations with the CBI. Table 4.1 lists the observations for the second period. Four runs were lost to high winds at the site; the rejection of these data is discussed in Section 4.2.2.

Given the paucity of bright, compact, and polarized sources at centimeter wavelengths, it is fortunate that 3C279 has a close neighbor, 3C273, which also fills these requirements. 3C273 has $I_{31} \sim 20$ Jy and $m \sim 5\%$, and although its lower polarized emission makes it less desirable than 3C279 as the primary polarization calibrator, it is bright enough to permit fast observations to verify the CBI's polarization calibration. We therefore included a 3C273 observation in

date	configuration	comment
18feb00	BC	
12mar00	BC	
05apr00	C	
25apr00	C	
04may00	C	
19may00	C	rejected: high wind
02jun00	C	
15jun00	C	rejected: high wind
30jun00	C	
13jul00	CD	rejected: control system crash
31jul00	D	rejected: high wind
10aug00	D	
24aug00	D	rejected: high wind

Table 4.1: BC is a hybrid array: north arm in B array, east-west arms in C. CD is a also hybrid array: north arm in C, east-west arms in D.

each VLA run. The VLA and the CBI observations encompassed numerous other calibration sources, but these two sources represent the only overlap between the two programs.

4.2.1 Observing Strategy

The time awarded for our program consisted of 2^h blocks in the 08:30-12:30 LST range. Most of this time was consumed by calibration observations (40%) and overhead for antenna slews and subreflector nutations (40%); the balance of the time was filled by flux measurements of diagnostic sources (14%) and program sources (6%). The preponderance of time dedicated to calibrations and diagnostic observations reflects the demands that high-frequency polarization observations place on the system.

Table 4.2 provides a schematic overview of a typical VLA schedule. 3C279 and 3C273 are the program sources; all other sources facilitate the analysis and interpretation of the observations of 3C279 and 3C273. We included a handful of diagnostic sources, J0713+438, J1310+323, and J1400+621, to assess the repeatability of the VLA flux measurements. J1310+323 and J1400+621 are sources drawn from the VLBA polarization calibrator monitoring program which have a history of long-term stability [55]. J0713+438 is a Compact Symmetric Object [78], and as such, it is unpolarized [71]; this source provides a test of the quality of the polarization calibration of the VLA. Each epoch required an observation of a flux calibrator (3C286) and at least four observations of an instrumental polarization calibration (0927+390 or 1159+292). Nearly every observation was preceded by a pointing calibration procedure. Time comes at a premium during these observations, and we exploited some redundancies among the calibration sources. 3C286, for example, doubles as the absolute polarization position angle reference. In addition, all of the sources are unresolved under nearly all possible combinations of VLA configurations and frequencies; this eliminates the need for additional phase references.

The introduction alludes to the complications inherent in high-frequency po-

source	band	τ (m)	el ($^\circ$)	p.a. ($^\circ$)	purpose
0927+390	<i>X</i>	2.5	—	—	primary pointing
	<i>Q, K</i>	5.0	80	-115	polarization calibration
0713+438	<i>X</i>	2.5	—	—	primary pointing
	<i>Q, K</i>	2.5	68	—	diagnostic flux measurement
0927+390	<i>X</i>	2.5	—	—	primary pointing
	<i>Q, K</i>	5.0	83	-135	polarization calibration
1400+621	<i>X</i>	2.5	—	—	primary pointing
	<i>Q, K</i>	2.5	38	—	diagnostic flux measurement
0927+390	<i>X</i>	2.5	—	—	primary pointing
	<i>Q, K</i>	5.0	85	-176	polarization calibration
3C286	<i>X</i>	2.5	—	—	primary pointing
	<i>Q</i>	2.5	—	—	secondary pointing
	<i>Q, K</i>	5.0	40	—	flux calibration
1310+323	<i>X</i>	2.5	48	—	diagnostic flux measurement
0927+390	<i>X</i>	2.5	—	—	primary pointing
	<i>Q, K</i>	5.0	82	128	polarization calibration
1224+035	<i>X</i>	2.5	—	—	primary pointing
3C273	<i>Q</i>	2.5	—	—	secondary pointing
	<i>Q, K</i>	5.0	45	—	program flux measurement
3C279	<i>Q</i>	2.5	—	—	secondary pointing
	<i>Q, K</i>	5.0	35	—	program flux measurement
0927+390	<i>X</i>	2.5	—	—	primary pointing
	<i>Q, K</i>	5.0	77	108	polarization calibration

Table 4.2: Sample VLA schedule.

larization observations with the VLA. Some of the more pernicious effects can be mitigated through the organization of the observing runs. As we will see, high-frequency observations suffer from severe elevation-dependent gain variations. To first order, these errors can be suppressed by scheduling the program sources to occur at elevations which are close to that of the flux calibrator. In competition with this requirement, however, polarization calibration observations with the VLA require multiple scans of the polarization calibrator over uniform steps in parallactic angle (Section 4.1.5). The final schedule is usually a compromise between these considerations. In the schedule shown in Table 4.2, for example, the third and fourth polarization calibrator observations are separated by a large step in parallactic angle in order to permit observations of 3C286, 3C279, and 3C273 at similar elevations. These considerations, as well as the components of the schedule, are discussed in detail in the following sections.

4.2.2 Antenna Selection

Several years ago the VLA initiated an ongoing effort to equip all 27 antennas with Q band receivers. At the outset of our VLA observing program, 19 antennas had been fitted with Q band receivers, but a subset of these antennas suffered from problems which compromised their performance at this high frequency. Only 15 had surfaces which had been figured to accommodate 0.7 cm observations, for example; antennas 10, 21, 24, and 28 had surface errors which substantially degraded their efficiencies, and thus had $\sim 3\times$ less gain than the other antennas. In addition, antennas 19 and 21 had random $\sim 0.2'$ pointing errors arising from servo control problems. We culled the weaker antennas to obtain a subset of antennas with uniform characteristics, and these antennas were used for the duration of the program. The five rejected antennas were used as a subarray dedicated to measuring the atmospheric opacity during each run (Section 4.2.7).

4.2.3 Referenced Pointing

The VLA antennas have raw antenna pointing errors of $\sim 10''$, and errors of this magnitude are comparable to the VLA's primary beam of $\Omega_p \sim 1'$ at Q band. In the absence of a corrective measure, this pointing error produces an amplitude error of $\sim 10\%$. This error can bias the measured program flux up or down, depending on whether the pointing error occurs during a flux calibration or a program observation. To remedy this problem, the VLA provides a procedure, *referenced pointing*, which corrects the pointing during the observing session. Referenced pointing is a standard practice at K and Q band; VLA experience demonstrates that at these frequencies the user should schedule a pointing calibration at least once an hour or after slews larger than $\sim 10^\circ$ [81]. The pointing calibrator must be bright and unresolved; in the case of the schedule in Table 4.2, 3C273's extended emission at X band confuses the system, so the nearby source J1224+020 provided the initial pointing solution for the program sources.

The referenced pointing system generates real-time corrections to the antenna pointing. When the user commands a pointing procedure, the VLA measures the pointing calibrator's flux at the commanded position and at offsets in the four cardinal directions. The telescope control system immediately fits the five fluxes for each antenna to a model of the antenna beam, and solves for a set of antenna-based pointing corrections which are to be applied to subsequent observations. The control system discards these corrections when the user commands a new calibration; an inspection of the corrections shows that successive corrections bear no resemblance to each other. Each iteration requires $\sim 2.5^m$, and repeated pointing calibrations cost a significant amount of time; any arrangement of the observing schedule which eliminates pointing calibrations pays large dividends.

The VLA provides two types of pointing correction: *primary* and *secondary* referenced pointing. These procedures differ in the VLA band at which the correction is made. All primary referenced pointing observations are carried out at 8 GHz (X band), the VLA's most sensitive channel, and these corrections can be applied to all other observations at all frequencies. Secondary referenced pointing

builds on the primary referenced pointing solutions by repeating the procedure at the frequency of interest, usually K or Q band. The primary beam is smaller at these higher frequencies, so the secondary solution refines the primary solution. Secondary referenced pointing also eliminates any errors arising from collimation differences between X band and K or Q band.

We employed both primary and secondary referenced pointing throughout the VLA observations. As Table 4.2 shows, all sources received a primary pointing correction, while secondary corrections were derived for the sources of greatest interest: the flux calibrator and the program sources. VLA pundits debate the efficacy of secondary referenced pointing; we chose to err on the side of caution and use both procedures when time permitted. The secondary referenced pointing solutions were generally on the order of a few arcseconds, which is $\sim 5\%$ of the primary beam at Q band; these corrections are smaller than the primary referenced pointing solutions, which suggests that the procedure yields a refinement to the primary solutions. To reduce the overhead required for pointing calibrations, the solutions were recycled where possible; 3C273 and 3C279 are separated by 10.4° on the sky, for example, so in the schedule in Table 4.2, the primary solution for 3C273 was applied to 3C279, and thus 3C279 required only a secondary solution. As Table 4.2 shows, secondary referenced pointing solutions derived at Q band were applied to the subsequent Q and K band observations.

4.2.4 Absolute Flux Calibration

3C286, a compact steep spectrum QSO [64], is one of several primary flux calibrators used by the VLA for observations at K and Q band. 3C286 has a flux density $I_\nu \sim 2$ Jy at both K and Q band, so a 2^m observation yields a $\sim 10^3\sigma$ detection at both frequencies. 3C286 is a rare example of a stable, compact extragalactic radio source; most of the source's radio emission originates from a region whose light crossing time is tens of kyr [37]. While the VLA's C and D configurations cannot resolve 3C286 at centimeter wavelengths, it is slightly resolved by the larger B & BC configurations, and under these circumstances a good model is required for

the analysis. The K and Q band flux scales for 3C286 are derived from comparisons of D array observations of Mars to a model for the planet’s emission [72]. The VLA recommends canonical flux densities of $I_K=2.52$ Jy and $I_Q=1.45$ Jy for 3C286, although for this work we used a different set of values: $I_K=2.589$ Jy and $I_Q=1.562$ Jy; these flux densities were obtained from recent comparisons of 3C286 and Mars at high frequencies and and extrapolations of the Baars scale at low frequencies [69, 2]. Comparisons of Mars and the optically thin planetary nebula NGC7027 with the VLA provide an estimate of the uncertainties on these values: 2% at K band and 8% at Q band. The flux calibration observations are among the most critical in each run, so both primary and secondary referenced pointing procedures preceded these observations.

4.2.5 Instrumental Polarization Calibration

Imperfections in the VLA receivers produce instrumental polarization which can be comparable to the polarization of the program sources. The technique that the VLA recommends to calibrate the VLA’s instrumental polarization resembles that for the CBI: the user observes a calibrator over a range of parallactic angles and lets changes in the parallactic angle modulate the source polarization relative to the instrumental contribution. The VLA differs from the CBI, however, in several fundamental ways which affect this comparison. First, since the VLA receivers respond to both hands of circular polarization, the absolute flux calibration suffices to determine the receiver gains. An inspection of the calibration equation reveals that this characteristic obviates the need for prior knowledge of the calibrator polarization:

$$\mathcal{V}^{RL} = g_j g_k^* [P e^{2i\phi} + I(\epsilon_j + \epsilon_k^*)] \quad (4.1)$$

For the VLA, the absolute flux calibration (RR and LL) on 3C286 isolates the gains g_j and g_k^* , so the remaining unknowns— ϵ_j and ϵ_k^* —can be deduced from the fit to the center of the distribution of cross polarized visibilities. This feature

of the VLA has clear virtue that the uncertainty in the intrinsic polarization of the calibrator plays no role in the reliability of the calibration. Second, while the CBI can force a change in parallactic angle with deck rotation, the VLA must rely on the diurnal motion of the sky to change the parallactic angle of the source. To obtain a good solution with a polarized source, the source parallactic angle must swing through at least 90° in at least four uniform steps during the run; this requirement favors sources which transit during the 2^h run. The source must also be bright and unresolved. These requirements confine our choice to one of two polarization calibrators: 1159+292 and 0927+390, both of which have several Jy of total intensity at *cm* wavelengths. We used single referenced pointing to calibrate the pointing observations, and while double referenced pointing would have been desirable, the additional 12.5^m required is prohibitive. Section 3 discusses efforts to assess the quality of the polarization calibration.

The VLA leakage terms (“D terms” in VLA parlance) are typically on the order of 5% at *K* and *Q* band. The leakage terms can be measured to better than 0.5% relative to the total intensity of the calibrators; this uncertainty limits the quality of the polarization measurements. Some evidence suggests that the instrumental polarization varies on week to month timescales [14], so the D terms must be determined anew for each polarization observation.

4.2.6 Position Angle Calibration

The technique described above isolates the instrumental polarization. The phases of the gains g_i and g_j^* , however, also affect the measured polarization; the phase difference between the *R* and *L* channels is equivalent to the inferred position angle of the source polarization. To obtain an absolute measurement of this phase, we observe a point source of known polarization, and insert a phase offset derived from the position angle error into the phases of the gain solutions. 3C286 is the primary position angle reference for the VLA; the position angle of 3C286 is 33 degrees at all VLA frequencies [88].

4.2.7 Atmospheric Opacity

Atmospheric extinction at K and Q band can bias the measured fluxes. Atmospheric opacity at K band arises from the 22 GHz water vapor rotational transition, and at Q band from the broad wings of the multiple ~ 60 GHz O_2 rotational transitions. The opacities at both frequencies can be on the order of ~ 0.1 . While this opacity is large, it affects the flux measurements differentially; the opacity introduces a systematic amplitude error which varies with the elevation difference between the flux calibrator and the program source. There are several ways to mitigate the effect of atmospheric extinction; we can measure the opacity during the observation and compensate for the extinction in the analysis, and we can structure the observations to minimize the elevation difference between the flux calibrator and program sources.

The atmospheric opacity at the VLA site changes with weather, so each observing session required a fresh opacity measurement. The VLA provides a standard skydip procedure which requires at least $\sim 12^m$ for nine elevation steps from the zenith to the horizon. Opacity measurements for both K and Q band would clearly require a prohibitive amount of time. Skydips do not require the entire array, however, so the five antennas which had serious efficiency problems at Q band were reassigned to a subarray which performed skydips during each run. In the compact VLA configurations, the azimuth of the skydip was chosen to prevent shadowing at low elevations. The skydip data were analyzed with *Tiptool* [13], and the errors on the opacities are roughly 1%. For an opacity of 0.2, this 1% error results in an extinction error of less than 1% for an elevation separation of 10° about an elevation of 45° . The opacities computed by *Tiptool* were added to the analysis pipeline (Section 4.2.1).

4.2.8 Gain Curves

The VLA antennas suffer from gravitational deformations that degrade their performance at high frequencies. Most of the distortion is confined to flexure in the

strut assembly which supports the subreflector; as the elevation of the telescope changes, the struts sag and thereby change the forward gain. At 43 GHz, for example, an elevation slew of $\sim 40^\circ$ can alter the forward gain by as much as $\sim 20\%$. If left uncorrected, these deformations will bias the flux measurements. Fortunately, however, the gain variations repeat over long timescales; a careful measurement of this effect yields antenna *gain curves* which can be used to correct data spanning many months.

Elevation-induced gain variations of 10% were seen in the Aug-Dec 1998 Q band data. At that time the extant gain curves covered a subset of the Q band antennas, so we measured the gain curves again. A gain curve measurement requires observations of a source through a large range in elevation, at least between the horizon and the zenith. The gain curves generally fit a quadratic function of the elevation, so a gain curve measurement requires at least three flux measurements, and preferably many more. We obtained $\sim 6^h$ of VLA engineering time for gain curve measurements on each of five occasions in spring and summer 2000; the first four observations were rejected because of high wind at the site, but the weather for the fifth attempt was favorable. These data were analyzed by Steve Myers for his VLBA polarization calibrator monitoring program, and he has made these curves available for the use of the entire VLA community [56]. Figure 4.1 shows Q and K band gain curves for antenna 6. The importance of this gain correction is apparent from this plot: if the flux calibrator is observed at 40° elevation, and the program source at 70° elevation, this antenna will underestimate the source intensity by 15% at Q band.

Gain curves are not the only line of defense against elevation-dependent errors. Elevation effects can be suppressed by scheduling the program source observations at elevations which are near that of the flux calibrator. Given the scheduling requirements for the polarization calibration, however, this can be difficult. In the case of the 25apr00 schedule shown in Table 4.2, we observed 3C286 at 40° elevation and 3C279 and 3C273 at 35° and 45° , respectively. Most of the antennas are designed to have maximum forward gain at intermediate elevations, and these




Figure 4.1: Sample gain curves at *K* and *Q* band for a VLA antenna.

observations benefit from their positions near the flat region of the gain curves. In the absence of a gain correction, elevation-dependent errors produce a $\sim 4\%$ bias in the program fluxes at Q band for these observations.

4.3 Data Reduction

The VLA data were edited and calibrated with standard AIPS tools. Data edits substantially reduced the body of monitoring data; these edits can be grouped based on the reach of the underlying problem—some problems affected the entire the array, while others affected individual antennas or baselines. External factors, particularly bad weather, compelled the rejection of all data for five of the thirteen runs: high winds at the site rendered four runs useless, and a control system crash during the run eliminated another. The wind-based edits resulted in a substantial loss of data, and these edits are discussed in detail in Section 4.3.2. Most antenna-based edits arose from pointing calibration failures, which are noted in the referenced pointing solutions; and from gain instabilities, which were apparent in the uncalibrated visibilities. Failure to reject the antennas with no pointing solutions for 3C286 on 05apr00, for example, biases the K band 3C279 flux upward by 10%. Finally, bad correlations in one of RR , RL , LR & LL for either IF prompted the rejection of all visibilities for that integration. Together these edits reduced the data from the subset of good antennas in the eight remaining runs by 30% at K band and 15% at Q band.

The edited data were passed to the AIPS calibration pipeline. The AIPS task CALIB computed the antenna-based complex gains, and these solutions were transferred to all the sources. The BC configurations of February and March 2000 resolved 3C286, so the calibration package employed a clean component model for each band; the models were supplied by the VLA [15], and the total fluxes in the models are those given in Section 4.3.4. For consistency, these models were used for all the observations. The AIPS task CLCOR combined the gain curves and the opacities with the source elevations to correct the gain amplitudes for elevation-

dependent biases. Once the parallel hand gain terms had been established, the AIPS task PCAL averaged the five observations of the polarization calibrator to isolate the instrumental polarization contribution to the cross polarized visibilities. The R-L phase difference was adjusted with CLCOR to recover the known position angle for 3C286, and the data were written to (u,v) -fits files for analysis in DIFMAP.

DIFMAP provides a variety of tools for extracting flux from visibilities [67]. For unresolved sources, the radially averaged visibilities provide an accurate measure of the source flux. For sources with simple morphologies, this average provides an estimate of the flux; a better measurement requires DIFMAP's visibility fitting tools. An inspection of the visibilities for the sources suggested that the sources were at most only marginally resolved by all configurations at *cm* wavelengths, so all visibilities were fitted to a single Gaussian. In cases where the source was unresolved, the modelfitting routine forced the Gaussian to be smaller than a resolution element. This technique was applied to I , Q , and U for all sources.

4.3.1 Interpretation

As anticipated, 3C279 demonstrated significant variations in both intensity and polarization throughout the Feb-Aug 2000 period. Some of the variations are, of course, real; the central issue for the interpretation of these data is the extent to which the variations reflect artifacts of the observing technique and the reduction pipeline. We can attack this problem on several fronts. First, the telescope diagnostic data provide rough guidance about the bias introduced by wind loading (Section 4.3.2). Second, we can prevail upon the stability of 3C286 to identify intensity and polarization variations of instrumental origin (Section 4.3.3). Third, contemporaneous measurements with the CBI of the total intensities of the program sources can verify that the observed variability is intrinsic to the sources (Section 4.3.4). Finally, the observations of the diagnostic sources provide some handle on the quality of the data, although the conclusions which these data support are mixed (Sections 4.3.5, 4.3.6).

Section 4.2 notes that the VLA has two IFs at each of K and Q band. These IFs are separated by 50 MHz for both bands. This difference between the IFs is too small to subsume significant changes in flux density due to the source spectra, so the two IFs may be regarded as independent channels which measure the same flux density. The differences in total intensity and polarized flux for the two IFs tend to be much larger than the thermal errors in the associated visibilities, so the flux densities from two IFs, which do not show any systematic differences, offer a more realistic estimate of the noise in the measurements. The error bars in the figures which follow are the differences between the measurements at the two IFs.

4.3.2 Wind

Wind loading on the VLA antennas is a paramount concern for observations at high frequencies. Winds in excess of 6-8 m/s can cause the referenced pointing system to fail to find solutions. In addition, we found that winds above ~ 5 m/s can degrade the antenna efficiencies to such an extent that the measured fluxes are biased by as much as 40% at Q band. Wind is particularly severe during the summer, and we lost four epochs of VLA data due to high winds at the site.

The VLA provides several types of wind data for users, and these data aided the efforts to evaluate the quality of the VLA fluxes. The observer's log notes wind conditions on hour timescales, and the resulting 2-3 points per 2^h run provide a gross estimate of the wind conditions during the run. On two dates with very high winds, these data supported the rejection of the entire run. In addition, the referenced pointing solutions contain time-stamped wind data for each pointing calibration; as we will see below, these data provided greater insight into the anomalous fluxes measured during some epochs.

Figure 4.2 demonstrates the effect of wind on the flux measurements for 3C279 at Q band. The top two frames show the source flux I and fractional polarization $m = |P|/I$, while the third shows several measures of the wind. The boxes represent the limits on the wind reported in the operator's log. The points denote the windspeeds during observations of 3C286 and 3C279. High winds on 19may00

and 15jun00 clearly undermine the reliability of the Q band fluxes on both dates; in both cases, wind loading suppresses 3C279’s flux by nearly 50%. High winds can bias fluxes up *or* down. On 31jul00, 3C286 was observed during a ~ 5 m/s wind, while 3C279 was observed during a ~ 8 m/s wind, and the relative loading depresses the 3C279 flux. On 24aug00 the relative wind loading between these two sources reverses, and 3C279 acquires a boost. Despite these simple considerations, wind loading generally defies characterization; the magnitude of loading depends on the interplay between telescope orientation and the wind direction, and as such, the strongest statement we can make is that winds in excess of ~ 5 m/s can explain some variations in total intensity. Based on the considerations discussed above, we rejected the data for these four epochs—all dates for which windspeeds exceeded 5 m/s. Contemporaneous observations of I_ν of 3C279 with the CBI on these dates confirmed that these variations are not intrinsic to the source.

An inspection of the fractional polarization in the middle frame of Figure 4.2 suggests that the wind loading threshold is too severe; the fractional polarization m , the quantity which is transferred to the CBI, appears relatively immune to the wind loading which can so strongly bias the total intensity. Indeed, even on 15jun00, the windiest day, m does not depart from the neighboring values. Given that the intensity and polarization both scale with the overall antenna efficiency η_A , the fact that $m = \eta_A p / \eta_A I = p / I$ remains unchanged on this date is not a great surprise. Were this true for all the dates, however, perhaps the wind loading criterion could be relaxed—the 31jul00 data contradict this conclusion; high winds during the 3C279 observation on this date boost the fractional polarization to twice the neighboring values. An inspection of the time-stamped wind data revealed that this was a gusty day, and the five observations of the polarization calibrator 1159+292 in particular experienced a wide range of windspeeds, from ~ 3 to ~ 9 m/s. A comparison of the Q band leakage terms for 31jul00 and those of neighboring dates confirmed that the 31jul00 polarization calibration was fundamentally different; the two sets of leakage terms differed in amplitude by as much as a factor of two—far larger than the normal. Based on these considerations, the




Figure 4.2: Demonstration of the effect of wind loading on the VLA antennas. The bottom frame shows windspeeds noted during the run; the boxes denote windspeeds reported in the observer's log for regular intervals throughout the run, while the points show windspeeds at the times of the observations of the sources shown in the key. Note that we have only incomplete data for some of the dates (e.g., 18feb00)

31jul00 data were rejected along with those of 19may00, 15jun00, and 24aug00. The polarization position angle χ was not included in Figure 4.2 because wind loading appeared to have little bearing on the repeatability of the values for χ .

4.3.3 3C286

Figures 4.3 and 4.4 show the intensities and polarized fluxes for 3C286 at K and Q band for the remaining eight dates. The total intensity I and polarization position angle χ are supplied to the calibration pipeline; the dashed lines in frames 1 and 3 denote the 1999.2 values recommended by the VLA, and properly calibrated data should agree with these values. The fractional polarization m of 3C286 is also constant [91]. Since m is not an input for the calibration, however, departures from the expected value reflect either intrinsic polarization variations or spurious effects introduced by the observations or the calibration. There is no precedent for intrinsic variations, so we will assume that variations in m reflect errors in the acquisition or the analysis of the data.

The measured total intensity for 3C286 often falls short of the given values for both K and Q band. The maximum differences are 0.5% for K band (12mar00) and 1% for Q band (30jun00). The total intensities were obtained from a single component Gaussian fit to the visibilities, and these discrepancies arise because the Gaussian model fails to represent all of the flux. An inspection of the angle-averaged visibilities provides some insight into the cause of the missing flux; the mean of the amplitudes of short spacing visibilities agrees with the expected value, which suggests that phase errors have depleted the flux at the center of the map. While a phase self-calibration recovers the remaining flux, we chose *not* to self-calibrate the data because a consistent treatment of all of the data would require self-calibration of sources with poor S/N, which in turn produces erroneous flux determinations for those sources. This is an important consideration for the weak components of polarized emission in some sources.

The polarization for 3C286 provides a test of the polarization calibration. The Gaussian model recovers nearly all the expected flux for both Q and K band,




Figure 4.3: 3C286 observations with the VLA, K band.




Figure 4.4: 3C286 observations with the VLA, Q band.

although there are a number of dates for which the model flux falls short of the expected value. In the case of the total intensity, these discrepancies arose from phase errors, and self-calibration recovered the missing flux. At Q band, we expect $(Q, U) = (79, 176)$ mJy for 3C286, and the map noise can be as high as 20 mJy. Upon self-calibration, the measured fluxes exceed the expected values by 10%-15%. This excess flux may arise from the application of self-calibration to low S/N data, or it could be a symptom of a poor polarization calibration; since the latter affects the interpretation of the polarized flux from the program sources, we must establish whether this excess flux is real.

The radially averaged visibilities for Stokes Q on the dates in question can have a significant noise component, so the mean amplitude of the visibilities provides little guidance about the amplitude of the underlying flux. 3C286's position angle of $\chi = 33^\circ$ splits the polarized flux between Q and U with a ratio of 0.41:0.91, and a change in polarization position angle can isolate all of the polarized flux in either Q or U . Of course, to determine the angle which forces the polarized flux into one component, we must establish the current position angle χ , which in turn requires accurate measurements of Q and U . If the excess flux is an artifact of self-calibration, the magnitude of the error should be inversely proportional to the S/N; parameterize this S/N dependence with η , so that $F' = \eta F$, where F is Q or U . The input position angle χ , which balances the measured flux between Q and U , should force the prefactors to be approximately equal, and at that position angle, the position angle χ is a reliable number: $\chi \sim \tan^{-1}(\eta U / \eta Q) \sim \tan^{-1}(U / Q)$. To isolate the polarized flux, add the phase offset necessary to force the flux into U . At that position angle, the S/N is sufficient to extract the polarized flux from the radially averaged visibilities. In most cases where the fitted fluxes are discrepant with the expected values, the mean of the visibilities is equivalent to the expected value. This circuitous technique demonstrates that the 3C286 visibilities contain the expected polarized flux, so we may assume that the polarization calibration for these dates is valid despite the fact that the simple fits to the visibilities fall short of the expected values.

4.3.4 3C279 and 3C273

Figures 4.5 and 4.6 show the 3C279 data that were employed for the CBI polarization calibration. 3C279 appears quiescent in total intensity through early April, at which point it undergoes an outburst that tapers off through August. During this period χ decreases smoothly while m undergoes changes which, with the exception 12mar00 Q band value, track the changes in I . The S/N for the polarized components of 3C279 is high enough to permit self-calibration, which in turn shows that the amount of flux lost to phase errors is small. As we will see below, the contemporaneous CBI observations confirm the variations seen in the VLA total intensities. Tables 4.3.4 and 4.3.4 list the VLA values for I , m , and χ for K and Q band; these data, particularly the values for m and χ , provide the foundation for the CBI polarization calibration. Chapter 5 discusses the application of these data to the CBI observations.

Figures 4.7 and 4.8 show the total intensities and polarized flux densities for 3C273. At K band, the total intensity decreases linearly while the fractional polarization increases monotonically for nearly the entire period. The Q band data show similar variations for all but 10aug00, at which time the polarized flux increases. The sharp upturn in the 10aug00 3C273 Q band polarization is not seen in the 3C279 data, which suggests that the change is not an artifact of the observation.

Figure 4.9 provides a comparisons of the VLA fluxes for K band with extrapolated values from the CBI. The top two frames of this figure show comparisons for 3C279 and 3C273. The extrapolated CBI fluxes were derived from a least squares fit to the fluxes in the ten CBI channels. The differences for 3C279 and 3C273 imply a systematic $\sim 10\%$ calibration offset between the CBI and the VLA; at K band, the extrapolated CBI flux exceeds the VLA flux by $\sim 8\text{-}10\%$ for 3C279 and $\sim 10\%$ for 3C273. Much of the discrepancy for 3C273 can be attributed to the contribution from the source's extended jet which the VLA does not resolve, however; high dynamic range measurements of the jet suggest that it should contribute ~ 1.5 Jy at K band [70]. Upon adding this to the flux of the source, we obtain the comparison shown in Figure 4.10. After augmenting 3C273's flux density to




Figure 4.5: 3C297 observations with the VLA, K band.




Figure 4.6: 3C279 observations with the VLA, Q band.

date	I (Jy)	σ_I	m	σ_m	χ ($^\circ$)	σ_χ ($^\circ$)
18feb00	24.61	0.04	0.089	0.001	60.1	0.8
12mar00	24.58	0.03	0.095	0.006	57.6	4.3
05apr00	24.58	0.05	0.085	0.005	58.0	1.8
25apr00	26.22	0.02	0.089	0.001	52.7	1.2
04may00	26.00	0.03	0.090	0.001	52.3	0.2
02jun00	25.75	0.15	0.107	0.004	54.0	4.2
30jun00	25.39	0.02	0.107	0.007	51.7	1.2
10aug00	24.70	0.14	0.114	0.002	49.1	1.4
mean	25.23	0.06	0.097	0.003	54.4	1.9

Table 4.3: K band polarization measurements for 3C279 with the VLA. To provide a sense of the precision with which we measure I , m , and χ , the bottom row shows the arithmetic means of the entries in the table. The mean values show that apart from systematic errors, we typically measure m and χ to $\sim 3\%$ at K band. The uncertainties for m are consistent with the offsets seen for the K band observations of 3C286 shown in Figure 4.3.

date	I (Jy)	σ_I	m	σ_m	χ ($^\circ$)	σ_χ ($^\circ$)
18feb00	22.77	0.08	0.101	0.004	64.7	2.3
12mar00	22.62	0.12	0.117	0.006	60.6	4.0
05apr00	24.03	0.03	0.090	0.001	51.9	1.3
25apr00	26.17	0.04	0.097	0.001	52.2	2.2
04may00	25.19	0.13	0.100	0.001	52.9	0.7
02jun00	24.70	0.17	0.127	0.012	54.8	15.0
30jun00	25.16	0.10	0.119	0.002	50.8	2.4
10aug00	22.11	0.14	0.116	0.003	45.4	8.6
mean	24.09	0.10	0.108	0.004	54.2	4.6

Table 4.4: Q band polarization measurements for 3C279 with the VLA.




Figure 4.7: 3C273 observations with the VLA, *K* band.




Figure 4.8: 3C273 observations with the VLA, Q band.

account for the jet, the agreement between the VLA observations and the CBI observations is remarkable; the two sets of data cannot be distinguished.

Figure 4.11 repeats the source flux density comparison for Q band. At Q band, the discrepancies between the VLA and the CBI are larger, particularly for 3C273. This increase for 3C273 may reflect some steepening in the source's spectrum between 36 GHz and 43 GHz. The jet adds an additional 0.7 Jy at Q band, but this flux density does not suffice to bring the two sets of data into agreement. These observations show that the flux density scales used by the CBI and the VLA are in good agreement, particularly in the neighborhood of K band. The gain irregularities which plague VLA observations at Q band affect the two sources with equal force, so the offset seen in the third frame of Figure 4.11 is probably a feature of the spectral shapes of the two sources.

The VLA observations provide confidence in the CBI's absolute flux density calibration. The absolute calibration error for the CBI is $\sim 5\%$, while those for the VLA at K and Q band are $\sim 2\%$ and $\sim 8\%$. These uncertainties exceed the striking agreement shown in Figure 4.10, although they can marginally account for the discrepancies in the Q band data.

Although the two telescopes differ by an absolute calibration offset, both should recover similar ratios of flux for the two sources at each band (Figures 4.9 and 4.11, frame 3). After accounting for the additional flux density of 3C273's jet, the ratios of the K band VLA observations agree to $\sim 5\%$ with those from the CBI after extrapolation to this lower frequency, although some flattening in the 10aug00 total intensity for 3C273 causes the VLA ratios to diverge from the CBI ratios on that date. At Q band, the agreement between the ratios is weaker; the offset between the two sets of ratios is $\sim 10\%$. While this larger offset may reflect spectral changes between 36 GHz and 43 GHz for one or both of the sources, the VLA and the CBI ratios also differ in the shape of the evolution of the ratios over time. This suggests that either the sources are undergoing temporal changes in the shape of their spectra between 36 GHz and 43 GHz, or that the VLA Q band measurements are more uncertain than their errors suggest. One peculiarity of this ratio analysis




Figure 4.9: 3C279/3C273 comparison with the CBI and VLA, K band. The scatter in the CBI data suggests that the errors on the flux densities, which are derived from the scatter in the visibilities, underestimate the real errors in the data. This discrepancy may arise from errors introduced when the total intensity calibration is transferred to the program sources during the `ncal1` procedure. The 3C279 observations with the CBI were typically performed at several deck positions to suit the requirements of the polarization calibration, whereas the 3C273 observations were performed at a single deck position; Section 2.4.2 discusses a noise source pathology which is observed during observations at multiple deck positions, and the superior consistency of the 3C273 flux densities over that for 3C279 may result from the same underlying effect.




Figure 4.10: 3C273 flux density comparison for the CBI and VLA, K band, after including the interpolated flux for the jet in the VLA data.

is that the ratios for the telescopes appear to diverge in early August for different reasons at different bands; at K band the divergence is caused by flattening in the 3C273 VLA flux, while for Q band the divergence is caused by a low 3C279 VLA flux on that date.

4.3.5 Summary

This chapter describes efforts to obtain accurate measurements of the total intensity and polarized flux densities of 3C279 and 3C273 with the VLA. These measurements are susceptible to a variety of effects which can bias the flux densities by tens of percent at centimeter wavelengths; these effects are sufficiently large to mask or mimic intrinsic changes in the source characteristics. While the VLA data were winnowed to a subset which is believed to be free of significant systematic errors, the final data still contain several puzzles. Why is the sharp $\sim 20\%$ increase in m for the 12mar00 3C279 Q band data not seen in the observations at K band? Why do the ratios of the program sources as measured with the CBI and the VLA diverge near 10aug00 for both bands? While these variations may be real, the magnitude of the many systematic errors which plague these observations




Figure 4.11: 3C279/3C273 comparison with the CBI and VLA, Q band.

suggests that the data still contain spurious variations at levels which are greater than the formal errors.

Chapter 5

Polarization Observations with the CBI

5.1 Introduction

The CBI polarization observations presented in this work targeted a variety of sources. Most of these data we consider were obtained from observations of two deep fields, and in Chapter 6 these data are applied to the primary scientific result of this work: a limit on the polarization of the CMBR. The observing strategy required to obtain an interesting limit on the polarization of the CMBR is simple—we require months of nights on a single pointing. Fortunately, this requirement dovetails with that for a component of the CBI’s total intensity program; one of the CBI’s objectives is to measure intensity fluctuations on small ($\ell > 1500$) scales, and long integrations on single pointings are needed to achieve the $\delta T < 10 \mu\text{K}$ levels predicted by standard models on these scales. These observing requirements are commensurate with those to detect the polarization predicted by standard models for $\ell \sim 600$ scales. These deep field observations were augmented by observations of galactic and extragalactic polarized sources to evaluate the CBI’s polarization capabilities.

The observations discussed in this chapter spanned a ten-month period in 2000. Figure 5.1 provides a timeline of the CBI’s polarization observations. The first of the two deep fields, located at $\alpha \sim 08^h$ (Table 5.1), was the focus of our attention soon after the CBI was commissioned in December 1999, and observations on

this field concluded when the field set in May 2000. We then embarked on an observing strategy consisting of mosaiced observations to improve the resolution in ℓ with which we sample the power spectrum in the $\ell \sim 300 \rightarrow 2000$ range. The mosaics require a lattice of pointings roughly a beamwidth apart on the sky, and since the single pointings which compose the lattice do not have sufficient S/N to obtain a useful limit on the polarization CMBR, the mosaic observations are not considered in this work. In August 2000, the total intensity effort returned to the high- ℓ program with a second deep field at $\alpha \sim 20^h$; these observations concluded in 29oct00, at which point RX12 was reconfigured for LCP, thus bringing the polarization observations to a close.

This chapter discusses the polarization observations with the CBI. Section 5.2 presents the considerations which led to our choice of deep fields. The 08^h and 20^h deep fields are described in Sections 5.3 and 5.4, respectively. These discussions focus on the calibration of these data as well as a number of consistency checks, although the likelihood analysis of the deep field data must await Chapter 6. The deep field observations were regularly interrupted by dates for which small lunar elongations precluded observations of the CMBR, and these nights provided the opportunity to explore the polarization performance of the CBI through observations of polarized galactic and extragalactic sources. These supporting observations are discussed in Section 5.5.

5.2 Deep Field Selection

The ease with which we interpret the total intensity and polarization bandpowers from the CMBR observations hinges on a judicious choice of fields, and a small part of this work entailed the search for a set of fields which would minimize confusion from astrophysical foregrounds. Several factors constrained the candidate fields to a subset of the sky. The CBI's elevation limit of 43° confines the time on source to at most 6^h , so the natural spacing for the fields in right ascension is $\sim 6^h$ provided that the fields transit near the zenith at the CBI's latitude of -23° . In addition,




Figure 5.1: CBI polarization observation calendar. The top row lists the three CBI configurations; the next row shows the duration of the 3C279 observations with the VLA; the next two show the polarization observations with the CBI.

the fields must be accessible to the VLA (latitude $+34^\circ$) and the OVRO 40 meter telescope (latitude $+37^\circ$), both of which have effective elevation limits of $\sim 20^\circ$ above the horizon at 30 GHz.

The goal of the search was to find four fields at equal intervals in right ascension which minimize contamination from diffuse galactic emission and extragalactic point sources. The galactic foreground templates were used to find regions of low diffuse emission, and the point source data were used to fine tune the positions of the fields to exclude the brightest point sources. This approach eliminates all but two free parameters: the right ascension zero point and the declination of the set of fields. In practice, the declination was confined to a band between $\delta \sim -23^\circ$ and the equator to keep the fields in view of the Northern hemisphere telescopes while

field	α_l (h, m, s)	α_t (h, m, s)	δ ($^\circ, ', ''$)	l ($^\circ$)	b ($^\circ$)
02^h	02:44	—	-03:30	176.54	-54.21
08^h	08:44:40	08:52:40	-03:10:00	230.30	24.19
14^h	14:44	—	-03:30	348.87	49.07
20^h	20:48:40	20:56:40	-03:30:00	44.44	-28.43

Table 5.1: The CBI fields. During the 01jan00 to 01jan02 period, the 02^h , 14^h , and 20^h fields were used for mosaiced observations, while the 08^h , 14^h , and 20^h fields were used for deep observations. The deep polarization observations presented in this work concentrated on the 08^h and the 20^h fields, so the positions of both the leads and the trails for these fields are given. The last two columns list the galactic coordinates of the fields. The 08^h field is relatively close to the galactic plane, but since it is far from the galactic center, the contamination in this field is relatively low. This is not the case for the 20^h field, however, which is relatively close to the galactic center.

maximizing the time on source. Concerns about foreground confusion for the total intensity observations drove the choice of fields, but the considerations which favor low foreground regions for total intensity are generally concerns for polarization observations as well. At 1 cm, contamination to total intensity observations can arise from galactic free-free emission, galactic synchrotron emission, extragalactic radio sources, and any anomalous sources of emission. High galactic latitudes are preferred, but our observing strategy limits the extent to which we can avoid the galactic plane for all four fields. These foregrounds—and the templates which guided the search—are considered in succession below. This effort resulted in the four fields listed in Table 5.1.

Anomalous Foreground

Observations at centimeter wavelengths with the 5-meter and 40-meter telescopes at OVRO demonstrated the presence of a diffuse source of emission of unknown origin in the vicinity of the North Celestial Pole [48]. The emission has a spectral

signature which resembles that of free-free, and it is strongly correlated with dust emission found in the *IRAS* 100 μm data [4]. Unlike normal free-free emission, however, this anomalous emission is considerably brighter than the level predicted by the $\text{H}\alpha$ emission in the field. Using observations which sample $7'$ - $22'$ scales at 14.5 and 32 GHz, the authors derive a simple scaling between the *IRAS* 100 μm emission and the GHz brightness temperature of the foreground:

$$T_f = 7.5 \times 10^{-2} \nu_{\text{GHz}}^{-2.2} I_{100\mu} (\text{MJy/sr})^{-1} \text{K} \quad (5.1)$$

The $\nu^{-2.2}$ spectral dependence indicates that foreground shares the spectrum of free-free or very shallow synchrotron. The 100 μm *IRAS* maps show that regions of the galaxy with brightnesses comparable to a MJy/sr are relatively common, so concerns about contamination from this foreground set a tight constraint on the set of candidate fields. Further observations by the OVRO group at $\delta = -5^\circ$ and spanning $\alpha = 0^h - 24^h$ suggest that the correlation with 100 μm *IRAS* emission at other points on the sky is not nearly as strong as that in the NCP region, and indeed, the NCP region has a considerable amount of structure which may contribute to local pathologies in the emission. The observation at $\delta = -5^\circ$ suggested contamination of $\sim 20 \mu\text{K}$, which, when added in quadrature to $\delta T \sim 60 \mu\text{K}$ which is expected for $\ell \sim 600$, results in a $\sim 5\%$ effect. These observations indicate that the anomalous foreground is perhaps less of a threat than the scaling in Equation 5.1 suggests.

In light of the uncertainties about the anomalous foreground, we confined the field search to regions of low, smooth emission in the *IRAS* 100 μm maps. Table 5.2 lists the mean *IRAS* emission for the the central $1^\circ \times 1^\circ$ of the four fields, along with the anomalous emission obtained from the scaling in Equation 5.1. The predicted foreground levels are quite high—they can exceed the CMBR signal on CBI scales—but the follow up observations that showed that the NCP scaling does not persist near the equator give reason to believe that these estimates of the contamination are very conservative. Despite the fact that the *IRAS* emission in

field	$\bar{T}_{100\mu}$ (MJy/sr)	$\bar{T}_{31 \text{ GHz}}$ (μK)
02 ^h	2.74	110
08 ^h	0.75	30
14 ^h	6.09	244
20 ^h	4.39	172

Table 5.2: Anomalous foreground expected for the four CBI fields based on *IRAS* 100 μ emission and the scaling of Leitch [48].

the fields differs by as much as a factor of eight, the levels of foreground emission listed in Table 5.2 represent the best compromise between the emission in the four fields. In the event that this foreground competes with the CMBR in the CBI fields, the CBI’s ten channels provide some leverage in disentangling the foregrounds from the CMBR [11], although in so doing we incur a penalty in sensitivity because this technique effectively relegates some channels exclusively to foreground measurements. The polarization characteristics of the anomalous foreground are not known.

Galactic Synchrotron Emission

Galactic synchrotron emission is a concern for both the total intensity and the polarization observations with the CBI. As in the case of other foregrounds, however, surveys provide templates which aid the search. Haslam *et al.* have mapped the entire sky at 408 MHz with a resolution of 0.85° , and the distribution and fractional polarization of the emission seen in these maps suggests that the continuum emission arises from galactic synchrotron processes [28]. We can extrapolate the Haslam maps to 31 GHz to estimate the mean synchrotron emission in CBI bands. Table 5.3 lists the mean brightness temperature for the four CBI fields at 408 MHz, as well as the extrapolated brightness temperature at 31 GHz based on a conservative spectral index of $\alpha = -2.7$.

field	$\bar{T}_{408 \text{ MHz}}$ (K)	$\bar{T}_{31 \text{ GHz}}$ (μK)	$\delta T_{31 \text{ GHz}}$ (μK)	$\delta P_{31 \text{ GHz}}$ (μK)
02 ^h	18.8	157	2.4	1.7
08 ^h	14.4	120	1.8	1.3
14 ^h	30.9	260	3.9	2.8
20 ^h	26.1	220	3.3	2.3

Table 5.3: Galactic synchrotron emission based on 408 MHz maps. The scaling from 408 MHz to 31 GHz was performed based on a spectral index of $\alpha = -2.7$.

While the Haslam maps provide a measure of the mean emission in the field, the CBI observations are contaminated only by the emission which fluctuates across the primary beam, and the Haslam maps do not sample these small scales. The 325 MHz WENSS survey provides useful data about the small scale fluctuations in the synchrotron emission [80], and when combined with the extrapolated Haslam maps these data allow us to estimate the expected level of fluctuations at 31 GHz. The WENSS maps show that the fluctuations on 5-20' scales are $\sim 1.5\%$ of the mean, so we adopt this level for these scales, and report the results for the four fields in column 4 of Table 5.3. These values are well below the tens of μK expected for the total intensity of the CMBR on the same scales.

Since the synchrotron emission can be highly polarized, this foreground is an important concern for the CBI's polarization observations. If we assume the maximum possible allowed fractional polarization for synchrotron of $\sim 70\%$, the small scale fluctuations inferred from the Haslam maps suggest that the polarization fluctuations will be on the order of a few μK (column 5 of Table 5.3). While these levels are small, they are not insignificant compared to the expected polarization fluctuations in the CMBR of $\delta P \sim 5 \mu\text{K}$ on the same scales.

Our understanding of foregrounds on the angular scales and at the frequencies of interest to cosmology remains remarkably sparse, although the amount of literature on foregrounds has increased since the time that the CBI fields were selected.

In particular, the MAP and Planck missions have accelerated the interest in foregrounds; since these missions will perform polarization observations, much of this work has focused on polarized foregrounds as well. We pause here to consider what a sample of these efforts has to say about polarized synchrotron emission in the CBI fields.

Baccigalupi *et al.* have used a variety of ~ 3 GHz maps to estimate the power spectrum of polarized synchrotron emission one might encounter at high frequencies and a range of angular scales [3]. Their analysis considers two cases: small scales ($\ell < 1000$) near the galactic plane ($b = \pm 5^\circ$) and large scales ($\ell < 200$) far from the plane. While they do not address the intermediate galactic latitudes which the CBI samples in the 08^h and 20^h deep fields, their analysis for fields near the galactic plane provides a conservative upper limit on the synchrotron contamination we might expect to encounter in the CBI fields.

Baccigalupi *et al.* find that the small scale synchrotron fluctuations near the galactic plane can be significant. The authors derive a scaling based on the Parkes survey of the equatorial strip at 2.4 and 2.7 GHz for which the spectrum of the polarized component of synchrotron emission falls with ℓ :

$$C_\ell^P = 1.2 \times 10^{-9} \left(\frac{\ell}{450}\right)^{-1.8} \left(\frac{\nu}{2.4 \text{ GHz}}\right)^{-5.8} \text{ K}^2 \quad (5.2)$$

This scaling suggests a contamination of $\ell(\ell + 1)C_\ell^P \sim 10 \mu\text{K}$ at $\ell \sim 600$, and this level exceeds the expected CMBR polarization fluctuations on the same scales by nearly a factor of two. The two CBI polarization deep fields are far from the $b = -5^\circ$ threshold from which this scaling was inferred, however, and while the authors do not consider latitudes near $b \sim 25^\circ$, they find that at galactic latitudes near 50° the polarization spectrum has steepened to $C_\ell^P \sim \ell^{-2.9}$. This steeper spectrum, which appears below this high threshold in b , should provide relief from the synchrotron contamination at the intermediate latitudes of the CBI fields. For comparison, the authors also report the results of a calculation based on the Haslam 408 MHz maps which suggest that the spectrum of polarized synchrotron

is similar: $\ell(\ell + 1)C_\ell^P \sim 2.5 \times 10^{-13} \text{ K}^2$ at $\ell = 600$ and 100 GHz; scaling this result to 31 GHz yields $\sim 15 \mu\text{K}$ on the same scales. Unfortunately, the authors do not discuss this calculation, nor do they comment on its range of applicability, so it cannot be compared to the inference from the Haslam maps discussed above. Taken together, these calculations suggest that the synchrotron contamination to be expected in the CBI fields is significant—particularly for polarization—but *all* of these predictions remain to be verified on small angular scales and at high frequencies.

We have one additional tool at our disposal: the CBI deep total intensity observations on the same fields. These data, taken at the same time as the polarization data, suggest that that any synchrotron in the deep fields is negligible. The initial total intensity result on the 08^h field rules out a 15% synchrotron component with a spectral index of $\beta = -2.7$ at 2σ [62], while the joint analysis of the 08^h and 20^h fields permits a 2σ limit on the synchrotron contamination of 21% on the short $\ell \sim 600$ baselines where the contamination is expected to be most severe. While these levels are not insignificant, they are well below the levels predicted by Baccigalupi *et al.*

The intersection of constraints provided by the anomalous foreground as traced by *IRAS* 100 μm emission and 408 MHz continuum emission resulted in a set of fields on $\alpha = 6^h$ intervals which start in the neighborhood of $\alpha = 02^h30^m, \delta = -05^\circ$. Since the extragalactic point sources are distributed uniformly on the sky, they were considered *after* the galactic constraints produced a set of candidate regions which minimized the diffuse galactic foregrounds. The point source considerations are discussed below.

Extragalactic Point Sources

Point sources are a source of severe contamination for CMBR observations in the high- ℓ region targeted by the CBI. We cannot, of course, choose fields which are free of point sources down to an arbitrary flux density threshold, but we can choose fields which exclude the brightest sources. We adopted a 31 GHz threshold of 100

mJy in a $\sim 2^\circ \times 2^\circ$ region. The point source filter consisted of several sets of data. The 1.4 GHz NVSS survey was used to identify sources whose flux densities, when extrapolated to 31 GHz with a reasonable spectral index ($\alpha = -0.5$), exceed the ~ 100 mJy limit [16]. The vast majority of these sources were also detected by the PMN survey [23] at 4.85 GHz and Parkes Survey survey [98] at 8.4 GHz, so their 31 GHz flux densities could be estimated with a reasonable degree of confidence. The source fluxes were extrapolated to 31 GHz and compiled to form a master list of bright sources for the four candidate regions. The bright point source data for the four fields were stacked to create a single map which contained all of the potentially confusing sources. The stacked data suggested several possible pointing centers which limited the number of > 100 mJy sources; we settled on one set of fields for the initial observations; these are listed in Table 5.1.

While some of the point sources in the CBI fields are quite bright, two factors mitigate their effect on the polarization observations. First, according to the standard models, the polarization signal peaks in the $\ell \sim 600$ region of the spectrum, and the likelihood analysis of the total intensity observations of the 08^h field show that on these relatively large scales, point sources with flux densities of tens of mJy carry very little weight relative to the CMBR fluctuations of $\delta T \sim 50 \mu\text{K}$. Second, point sources are generally very weakly polarized—they are rarely polarized by more than 10%, so given our experience with point sources in the total intensity observations, as well as the relative amplitudes of the temperature and polarization fluctuations ($\delta T : \delta P :: 10 : 1$), the error introduced by point sources should not exceed the source contribution to the total intensity which is negligible at $\ell \sim 600$. And of course, the polarization effort benefits from the CBI observations of the total intensity: CBI observations of the 08^h and 20^h fields show that the point sources are a minor consideration at $\ell \sim 600$; the brightest source within the FWHM of the CBI beam for either the lead or the trail for both deep fields has a total intensity flux density of ~ 40 mJy. In summary, point sources are only of marginal concern for the polarization observations of the 08^h and 20^h fields.

5.3 08^h Deep Field

The 08^h field was overhead at midnight when the CBI began routine observations in January 2000, and we dedicated a considerable number of nights to this field to evaluate the performance of the system. These observations produced our first detection of the total intensity fluctuations in the CMBR; this measurement demonstrated a significant decrease in the power spectrum between bins centered on $\ell \sim 600$ and $\ell \sim 1200$ [62]. The 08^h field observations spanned two configurations; Table 5.4 lists dates of the 08^h field observations in configuration 1, and Table 5.5 lists the observations for configuration 2, as well as a summary of the total time on source. The period between 11jan00 and 30may00 totals 141 nights, but the number of nights on the field is far smaller: 59 nights were lost to weather; 9 were required for the reconfiguration of the array; and 29 were lost to a variety of other factors, predominantly low ($\theta < 60^\circ$) lunar elongation relative to the deep field. And as the summary of the scans per night in the table demonstrates, the time on source per night peaked early in the 08^h field campaign.

The 08^h field observations encompassed two array configurations. Configuration 1 (Figure 5.2) was an initial sparse configuration which provided uniform (u, v) coverage, and is thus neutral to the shape of the underlying spectrum, but its power to measure polarization was limited because few of its baselines had the LL counterparts necessary to correct the LR visibilities for instrumental polarization. Only two of these baselines (RX0-RX12 and RX2-RX12) sampled the $\ell \sim 800$ scales on which the polarization signal predicted by standard models peaks, although a third (RX10-RX12) can be included if we relax the requirements for the instrumental polarization calibration; we explore this possibility in Chapter 6. To improve our sensitivity to large scale polarized emission, RX12 was moved to the center of the array midway through the 08^h field observations; this configuration (Figure 5.3) emphasized short cross polarized baselines while permitting access to all of the receivers for repair.

The observing strategy remained roughly the same for all of the deep field ob-

CBI configuration 1: 01jan00 to 14apr00




Figure 5.2: CBI configuration 1: 01jan00 to 14apr00. RX12 (hatched) is the orthogonally polarized receiver. Light solid lines denote LR baselines which have simultaneous LL counterparts, the latter of which are shown with dark solid lines. The light dotted lines show LR baselines which can be calibrated after several 20° deck rotations until they are parallel with the LL counterpart. The dashed line to RX10 shows an additional baseline which cannot be calibrated under ordinary circumstances, but which can be calibrated if we tolerate a small 7° error in the phase calibration for that baseline.

CBI configuration 2: 22apr00 to 01jul00




Figure 5.3: CBI configuration 2: 23apr00 to 01jul00. The solid shaded lines denote the *LR* baselines which have simultaneous *LL* counterparts. The dotted lines denote baselines which can be calibrated after some multiple of a 30° deck rotation.

servations. Chapter 2 notes that the spillover in the visibilities forced us to observe the deep fields in pairs which were differenced offline to excise the spillover. Each cycle of lead and trail required $\sim 16^m$. After each cycle the deck position was stepped 20° increments (configuration 1) or 30° increments (configuration 2); this increase in the stepsize followed the realization that adequate coverage of the (u, v) plane did not require the denser sampling provided by the 20° steps. The choice of 30° increments, as compared to 40° or 45° increments, is particularly conducive to polarization observations in configuration 2 (and configuration 3 for the 20^h field) because it permits *nearly* simultaneous LL - LR matches for LR baselines which lack simultaneous LL counterparts. Baselines RX12-RX6 and RX12-RX7 in configuration 2, for example, cannot be calibrated until baseline RX0-RX7 has rotated through 60° . To exploit this strategy, CBIPOLCAL was modified to search all of the LL visibilities to find the nearly simultaneous matches for the LR visibilities which lack simultaneous LL counterparts. The change in gain over the intervals spanned by a few scans is typically far smaller than the overall gain uncertainties, and is not a concern.

In Tables 5.4 and 5.5, the number of 16^m scans is shown arranged in columns of increasing baseline length for the 08^h field observations. This sorting allows us to quantify the data which go into the limits on each scale. In Chapter 6 we assign the visibilities to three bins in ℓ -space for the likelihood analysis; band 1 ($446 < \ell < 779$) corresponds to columns 5-6, band 2 ($930 < \ell < 1395$) corresponds to columns 7-8, and band 3 ($1539 < \ell < 2702$) corresponds to columns 9-12. After deducting overhead for slews, each lead or trail lasts $\sim 6.5^m$, and this conversion is used to obtain the sums at the bottom of Table 5.5. These times denote the total time on either the lead *or* the trail, and since differencing spreads twice as much time between two the fields, the net result is a penalty to the total sensitivity of $\sqrt{2}$ relative to the sensitivity implied by the integration times in the table. The sums at the bottom of Table 5.5 show that the integration times between the three bins scale as $\sim 1:0.7:1$. The expected uncertainties in the maps (second to last row) scale accordingly, while the real uncertainties obtained directly from the

visibilities show greater parity. The origin of this change, which is caused by a relative increase in the real noise for the first bin, is not known. All of the measured uncertainties exceed the expected values; we examine this point in detail in Section 5.3.2. The goal of the likelihood analysis is to determine the best fit bandpowers for each bin, and since the sensitivity per unit area in the aperture plane falls like $1/\ell$, the sensitivity to the first bin exceeds that for the other bins.

One might infer that the sensitivity is diminished further still because some of the scans will lack the 180° counterparts which are necessary to form both RL and LR , and thus both Q and U . While only a subset of the scans have the 180° counterparts, this deficiency does not affect the likelihood analysis because the calculation is built around LR visibilities; in this regard the limiting factor is thus the number of matching LL baselines, since these baselines determine which LR baselines can be calibrated.

5.3.1 08^h Field Calibration

The 3C279 observations with the CBI form the backbone of the 08^h field polarization calibration. The majority of the 08^h field observations were calibrated on 3C279, and while some were calibrated on Tau A, the Tau A model was calibrated with 3C279 using the techniques discussed in Chapter 3. Chapter 4 discusses the absolute uncertainties in the VLA values for 3C279's polarization. Two additional factors limit the accuracy with which the VLA observations can be applied to the observations with the CBI: the variability of 3C279 and the differences in the frequencies at which the CBI and the VLA observe; these factors result in a pair of interpolations which are necessary to transfer the VLA observations to the CBI. In addition, some of the 08^h field observations lacked observations of the calibrator, and on these dates we bootstrap the calibration from neighboring dates. This section discusses the variety of techniques which were used to calibrate the 08^h field.

The first uncertainty in the VLA calibration arises from the variability in 3C279. Figures 4.5 and 4.6 show that the VLA observations undersample the

#	date	calsource	caldate	scans per baseline length (cm)							
				100	104	173	200	300	346	400	458
1	11jan00	Tau A	—	20	20	0	—	0	—	0	0
2	12jan00	Tau A	—	16	16	16	—	0	—	0	0
3	13jan00	Tau A	—	21	21	21	—	0	—	0	0
4	09feb00	3C279	—	17	17	17	—	0	—	17	17
5	10feb00	3C279	—	21	22	21	—	0	—	22	22
6	11feb00	3C279	—	22	22	22	—	0	—	22	22
7	12feb00	3C279	14feb00	19	19	19	—	0	—	19	19
8	13feb00	3C279	14feb00	22	22	22	—	0	—	22	22
9	14feb00	3C279	—	19	19	19	—	0	—	19	20
10	01mar00	3C279	—	20	20	20	—	20	—	20	20
11	02mar00	3C279	—	16	16	16	—	16	—	16	16
12	03mar00	3C279	—	19	18	19	—	18	—	19	19
13	04mar00	3C279	—	17	16	17	—	16	—	17	17
14	05mar00	3C279	—	19	18	19	—	18	—	19	19
15	20mar00	Tau A	22mar00	17	17	17	—	0	—	0	0
16	21mar00	Tau A	22mar00	17	17	17	—	0	—	0	0
17	22mar00	Tau A	—	18	18	18	—	0	—	0	0
18	23mar00	Tau A	—	17	17	17	—	0	—	0	0
19	24mar00	Tau A	—	17	17	17	—	0	—	0	0
20	25mar00	Tau A	—	17	17	17	—	0	—	0	0
21	26mar00	3C279	28mar00	18	17	18	—	17	—	18	18
22	27mar00	3C279	28mar00	18	17	18	—	17	—	18	18
23	28mar00	3C279	—	17	16	17	—	16	—	17	17
24	30mar00	3C279	—	16	15	16	—	15	—	16	16
25	31mar00	3C279	—	15	14	15	—	14	—	15	15
26	04apr00	3C279	—	15	14	15	—	14	—	15	15
27	05apr00	3C279	—	13	12	13	—	12	—	13	13
28	06apr00	3C279	—	17	16	17	—	16	—	17	17
29	07apr00	3C279	—	14	13	14	—	13	—	14	14
30	08apr00	3C279	—	17	16	17	—	16	—	17	17
total, configuration 1				531	519	511	—	238	—	372	373

Table 5.4: 08^h field observations, configuration 1.

#	date	calsource	caldate	scans per baseline length (cm)							
				100	104	173	200	300	346	400	458
31	23apr00	3C279	—	22	—	33	11	11	11	—	—
32	24apr00	3C279	—	30	—	45	15	15	14	—	—
33	25apr00	3C279	—	24	—	36	12	12	12	—	—
34	26apr00	3C279	—	20	—	30	10	10	10	—	—
35	27apr00	3C279	—	18	—	27	9	9	9	—	—
36	28apr00	3C279	—	14	—	21	7	7	7	—	—
37	29apr00	3C279	—	24	—	36	12	12	12	—	—
38	30apr00	3C279	—	22	—	33	11	11	11	—	—
39	01may00	3C279	—	18	—	27	9	9	9	—	—
40	02may00	3C279	—	4	—	2	2	2	0	—	—
41	03may00	3C279	—	10	—	15	5	5	4	—	—
42	23may00	3C279	—	8	—	15	0	3	4	—	—
43	29may00	3C279	28may00	8	—	12	0	4	2	—	—
44	30may00	3C279	28may00	8	—	12	0	4	2	—	—
total, configuration 2				230	—	344	103	114	107	—	—
total, both configurations				761	519	855	103	352	107	372	373
expected σ_{LR} (mJy)				3.07		3.55		3.17			
measured σ_{LR} (mJy)				3.84		3.96		3.57			

Table 5.5: 08^h field observations, configuration 2. The sensitivity is computed based on 6.5^m per scan.

dramatic changes in 3C279’s emission over the duration of the 08^h field observations. In the most extreme case—the interval between 25apr00 and 19may00—the fractional polarization changed by nearly 20%, although in this case the damage to the overall calibration of the 08^h field is limited because this period encompassed only $\sim 10\%$ of the data. In the absence of additional information data we chose to linearly interpolate the source characteristics from the VLA observations to the intervening dates. We can estimate a reasonable upper limit on the error in the interpolation from the changes in 3C279’s characteristics; the variations in 3C279’s polarization from epoch to epoch are $\delta m \sim 0.01$ and $\delta\chi \sim 5^\circ$: both are $\sim 10\%$.

In practice, the errors in the linear interpolation are probably smaller than these values: the total intensity observations of 3C279 with the CBI show no evidence for wild excursions during the intervening dates, so we may plausibly conclude that the polarization is well behaved as well. The analytic errors in the interpolation tend to be $\sim 3\text{-}5\%$, but these estimates are valid only in the linear approximation of the variations.

The second uncertainty arises from the interpolation from the VLA channels to the CBI channels. The VLA K (22.46 GHz) and Q (43.34 GHz) band channels bracket the CBI's 26-36 GHz band, so the VLA values for m and χ were interpolated to the intervening frequencies. Again, we chose a linear interpolation. There is no guarantee that the linear interpolation is valid, and indeed, spot inspections of the flux densities in the ten CBI channels suggests that the bracketing data do not fully characterize the behavior of the source across the band. The interpolation errors tend to be $\sim 3\text{-}5\%$ in amplitude and phase; the error for a particular date depends heavily its proximity to the bracketing VLA observations. Figure 5.4 shows the interpolated VLA data at 31 GHz, as well as the uncertainties for these data; the figure shows that the net uncertainty in the interpolations is $\sim 3\text{-}5\%$.

Several nights were calibrated with Tau A. On these nights we applied Tau A Model 1 (Chapter 3) which is uncertain to 10%. While this uncertainty is larger than that for the 3C279 calibration, Tau A was required for only 9 nights, so the weight which this uncertainty carries in the overall error budget is small.

While $\sim 80\%$ of the nights were calibrated with direct observations of 3C279 or Tau A, there were occasions for which the weather, the moon, or schedule errors left us with no polarization calibrator for the night.¹ We therefore implemented a

¹The Tau A cross polarized visibilities are generally more susceptible to corruption from small lunar elongations than those for 3C279 scans because Tau A was never observed with trails. Comparison of 3C279 observations calibrated on Tau A to direct calibrations of the source on 11feb00 and 12feb00 demonstrate that lunar elongations of $\sim 30^\circ$ relative to Tau A can corrupt the inferred fractional polarization of 3C279 by 2%, a 20% error for this 10% polarized source. The moon moves $\sim 5'$ during the course of a 10^m lead/trail pair, so while the cancellation obtained with trails is not ideal, it would improve upon the undifferenced visibilities. Evidence of spillover

second approach to the calibration of the 08^h field: *indirect* calibration, in which the calibration is bootstrapped from a neighboring date. As we will see below, this approach introduces different errors than the direct calibration above. Column 3 of Table 5.4 notes the dates calibrated with this method.

We would like to estimate the error in the bootstrapped calibration, and the distribution of LR calibration factors for dates with good calibrations provides perhaps the best guidance for this problem. Figure 3.10 shows the amplitude and phase components of the gain for the 08^h observations for one baseline and one channel; this figure suggests that the night to night variations in the calibration are $\sim 10\%$. A quantitative analysis of the amplitudes and phases of the gains for all channels and all LR baselines for the 11jan00 to 10apr00 period shows that the *mean rms* of the gains is $\sim 15\%$ for both amplitude and phase; Figure 5.5 shows the distribution of the rms for all channels and LR baselines. The means were obtained by computing the variance of fractional change in amplitude between sessions for each channel and each baseline for the configuration 1 period, 11jan00 to 10apr00. The variance of the gain phase was measured relative to a radian. This 15% mean uncertainty is our best estimate of the error in the bootstrapped calibration for a single night. Figure 3.10 shows that there is no systematic drift in the calibration, and in this regard it is representative of all of the baselines, so the real error in the bootstrapped calibration for n nights should fall like $1/\sqrt{n}$.

5.3.2 08^h Field Results

Figures 5.6 and 5.7 show maps of Q and U for the short baselines. These maps show no sign of a signal; for both maps the rms of the central region of the maps is ~ 12 mJy/beam, which is indistinguishable from the rms away from the primary beam. For comparison, Figure 5.8 shows a map for I ; the total intensity signal is obvious in the center of the map, and the rms in the central region of 28 mJy/beam exceeds that of ~ 16 mJy/beam outside of the primary beam. These maps have on Tau A's ~ 25 Jy cross polarized visibilities would also argue in favor of the use of trails for polarization calibration observations on this bright source.




Figure 5.4: Interpolated VLA data for 3C279 at 31 GHz, 08^h field. The dates in the horizontal axis start with 01jan00 = 1. The top two frames show m and χ , while the bottom shows the fractional uncertainties for both. The errors include the VLA errors as well as the errors from the two interpolations. The shaded regions denote times during which the 08^h deep field was observed. The 11jan00-13jan00 08^h field observations are not shown; these were calibrated on Tau A, the model for which is good to $\sim 10\%$.




Figure 5.5: Histogram of the rms of the LR gains for the 08^h field for each baseline and channel. The data reflect the configuration 1 period, 11jan00- 10apr00. The phase errors are measured in terms of a radian. The means of rms for both the amplitudes and the phases is $\sim 15\%$; these means are a good measure of the change in the calibration from night to night.

not been corrected for the shape of the primary beam. The Q and U maps for bands 2 and 3—the intermediate and long baselines—are similar in that they show no evidence for a signal.

5.3.3 08^h Field Consistency Tests

Inspection of the deep field visibilities suggests that the data are dominated by noise. In the absence of a strong signal, the consistency tests must rely almost exclusively on the noise. In this section, we explore several diagnostics of the 08^h field calibration. We first compare the noise in the data to the performance predicted by the system characteristics, and then perform a χ^2 test to assess the quality of the signal. Finally, the noise is used to estimate changes in the amplitude calibration. The χ^2 test provides limited guidance as to whether the LR data contain a signal; the more rigorous maximum likelihood analysis is presented in

Residual Q map. Array: CBI
 C0844-0310 at 31.000 GHz 2000 Jan 12




Figure 5.6: Map of Q for the 08^h deep field, short baselines ($|u| < 150\lambda$). The circle shows the extent of the primary beam at the band center: $45.2'$ (FWHM) at 31 GHz. The noise in the central region is identical to that outside of the area defined by the primary beam.

Residual U map. Array: CBI
C0844-0310 at 31.000 GHz 2000 Jan 12




Figure 5.7: Map of U for the 08^h deep field, short baselines ($|u| < 150\lambda$). See Figure 5.6 for details.

Residual LL map. Array: CBI
C0844-0310 at 31.000 GHz 2000 Jan 12




Figure 5.8: Map of LL for the 08^h deep field, short baselines ($|u| < 150\lambda$). See Figure 5.6 for details. The rms in the region encompassed by the primary beam is nearly twice that outside the beam (see text).

Chapter 6.

08^h Field Visibility Uncertainties

Our first task is to establish that the noise in the visibilities agrees with our expectation based on the system characteristics. Section 2.3 reviews the contributions to T_{sys} ; the factors discussed there translate directly to the scatter in the visibilities, so we can compare the noise computed *a priori* with the real data.

The visibility uncertainties must be corrected for a bias which enters when the data are averaged. The uncertainties are estimated from the scatter in the ~ 50 or so 8.4^s integrations which form each scan: during routine observations, the control system records 8.4^s integrations to the archive, and in the offline analysis, UVSUB combines these data to form $\sim 6^m$ averages of the lead-trail difference visibility \mathcal{V}_j . UVSUB computes the sample variance of the differenced visibilities and records the corresponding uncertainty s_j with \mathcal{V}_j in the output *.uvf* file. We average the visibilities at each (u, v) point to obtain $\hat{\mathcal{V}}(\mathbf{u}) \pm \hat{\sigma}$. These mean quantities are obtained from the weighted averages of the N observations at each (u, v) point:

$$\hat{\mathcal{V}} = \frac{\sum_{j=1}^N \mathcal{V}_j / s_j^2}{\sum_{j=1}^N 1 / s_j^2} \quad (5.3)$$

for the visibilities, and

$$\hat{\sigma}^2 = \frac{1}{\sum_{j=1}^N 1 / s_j^2} \quad (5.4)$$

for the uncertainties.

Equation 5.3 is an unbiased estimator of the mean of the visibilities, but Equation 5.4 introduces a bias in the combined uncertainty. This bias enters because the uncertainties from each scan are each drawn from an underlying distribution; the scatter in the integrations follows from the intrinsic distribution of errors, so the scatter will differ from scan to scan in a way which does not necessarily reflect the true character of the noise. We correct for the bias by computing it and scaling

\mathcal{V}	calibration	$\langle\sigma\rangle$ (Jy)
\mathcal{V}^{LR}	CBIPOLCAL	6.10
\mathcal{V}^{LL}	CBICAL	5.89
\mathcal{V}^{LL}	CBIPOLCAL	5.84
\mathcal{V}^{LR}	CBIPOLCAL	5.86

Table 5.6: Means of the 08^h visibility uncertainties for LR and LL . Column 2 lists the calibration package used to calibrate the data. The LL calibration with CBIPOLCAL is included as a cross-check; the LL flux scale determined by CBIPOLCAL sets the flux scale for the subsequent LR calibration. The bottom entry shows the mean LR uncertainty after excising the edge channels from the CBI band (see text); these uncertainties are in good agreement with those for \mathcal{V}^{LL} .

the variance accordingly. The bias on the variance is given by

$$\beta \simeq 1 + \frac{4}{n} \quad (5.5)$$

where n is the number of visibilities which go into the $\sim 6^m$ average [50]. For routine CMBR observations, $n \sim 50$, so $\beta = 1.08$. Simulations that reflect lower order effects as well as numerical artifacts of the data pipeline show that β tends to be slightly lower: $\beta = 1.06 \pm 0.01$. We adopt this value of β and scale all of the visibility variances by this value.

Table 5.6 lists the means of the visibility uncertainties for the 08^h field observations. All uncertainties have been normalized to 1 s. The means were obtained from averaging over all of the visibility uncertainties for all channels and all nights. The uncertainty means have been corrected for the bias discussed above.

The 08^h visibility uncertainties all exceed the value computed from T_{sys} of 4.7 Jy by at least $\sim 20\%$. The source of this additional noise remains unknown. CBIPOLCAL performs the total intensity calibration for the LL baselines which match the LR baselines, so the table lists the values for $\langle\sigma^{LL}\rangle$ from CBIPOLCAL as a cross-check. The LL uncertainties obtained from CBICAL (row 2) and

CBIPOLCAL (row 3) agree. This result is important because CBIPOLCAL derives the absolute LR flux scale from the LL baselines it calibrates; an amplitude error in these LL visibilities, as indicated by a significant excess in $\langle\sigma^{LL}\rangle$ from CBIPOLCAL, would propagate to the LR calibration. The LR uncertainties are $\sim 5\%$ higher than those for LL . Since the two values of $\langle\sigma^{LL}\rangle$ agree, we infer that the excess in $\langle\sigma^{LR}\rangle$ is *not* a consequence of an error in the flux scale applied by CBIPOLCAL. It may reflect a calibration error in the values of m for 3C279 derived from the VLA.

The LR uncertainties in the top row are a bit higher than those for LL . RX12 is common to all of the LR visibilities, so one possible explanation is a higher system temperature for RX12. Since $\sigma^{jk} \sim \sqrt{T_j T_k}$, this explanation requires $T_{\text{RX12}} \sim (1.04)^2 \langle T_{\text{sys}} \rangle$, and indeed, as Figure 5.9 demonstrates, direct T_{sys} measurements for each of the receivers show that RX12's noise is ~ 1 K higher than the array average of 27.4 K. The excess T_{sys} for RX12 can account for much of the divergence between $\langle\sigma^{LL}\rangle$ and $\langle\sigma^{LR}\rangle$. Figure 5.9 suggests a quick test of this hypothesis: the higher T_{sys} for RX12 appears to arise largely from the two channels at the band edges. If we repeat the σ^{LR} analysis without these two channels, we find that $\langle\sigma^{LR}\rangle = 5.86 \pm 0.06$ Jy, which is consistent with $\langle\sigma^{LL}\rangle$. In a later section we look at the errors in greater detail, at which point we will see that the high T_{sys} for RX12 explains only part of the excess in $\langle\sigma^{LR}\rangle$.

χ^2 Test

χ^2 provides an important test of the data. The total integration time on the 08^h field is such that if the standard models are to be believed, the S/N per LR visibility should be unity for baselines of all lengths; we may assume that values of χ_ν^2 that exceed one indicate some combination of errors in our estimates of the noise and spurious contamination to the signal. The short baselines may contain a small contribution from polarization of the CMBR, but χ^2 does not have the sensitivity to detect these signals in the CBI data. We compute χ_ν^2 for the reals




Figure 5.9: T_{sys} comparison between RX12 and the array average, by band. The 1 K errorbars on the RX12 values are estimates which include both measurement uncertainties and systematic errors; the systematic errors dominate the measurement errors, and since the systematic errors enter both the RX12 measurement and the mean of the other 12 receivers, the errors in this figure understate the significance of the T_{sys} excess for RX12.

and imaginaries as

$$\chi_\nu^2 = \frac{1}{2N\beta} \sum_i^N \left[\left(\frac{\mathcal{V}_i^R}{\sigma_i^R} \right)^2 + \left(\frac{\mathcal{V}_i^L}{\sigma_i^L} \right)^2 \right] \quad (5.6)$$

where N subsumes all channels and unique visibilities per night for each of the 41 nights of data in the 08^h field. The factor of two accounts for the real and the imaginary components of \mathcal{V} ; in this case, $\nu = 2N$. The σ_i are the uncertainties on the visibilities; the quadrature calibration ensures that $\sigma_i^R = \sigma_i^L$. The factor of β removes the bias in the noise.

The χ^2 test demonstrates that the LR visibilities are generally consistent with noise. Table 5.7 reports the mean values of χ_ν^2 for all of the baselines in the 08^h dataset. The LR baselines were calibrated on 3C279 and Tau A in CBIPOLCAL, and the LL baselines were calibrated on 3C274 in CBICAL. The values of χ_ν^2 were sorted by baseline length to demonstrate the effect of a signal on χ_ν^2 for the short LL baselines: χ_ν^2 exceeds unity because the CBI's short LL baselines detect the CMBR with high S/N. The LR visibilities are inconsistent with noise for the short baselines, while the noise estimates for the long LR baselines appear to be 1-2%

		$ u < 150 \lambda$		$ u > 150 \lambda$		all $ u $	
\mathcal{V}	calibration	χ_ν^2	ν	χ_ν^2	ν	χ_ν^2	ν
\mathcal{V}^{LR}	CBIPOLCAL	1.025	14276	0.983	45060	0.993	59336
\mathcal{V}^{LL}	CBICAL	1.153	29698	0.987	291080	1.003	320778
\mathcal{V}^{LL}	CBIPOLCAL	1.101	28132	0.981	313032	0.991	341164

Table 5.7: χ^2 tests on the 08^h deep field data for *LL* and *LR*. The *LL* data were obtained from CBICAL, while the *LR* data were obtained from CBIPOLCAL. $\nu = 2N_d$ to account for the real and imaginary components of the visibilities.

too high.² The origin of this excess in the short *LR* baselines is not known, and given the number of degrees of freedom, it is significant; the likelihood analysis on these data in Chapter 6 can establish whether the signal is of celestial origin. We detect a signal on the long *LL* baselines, but the simple χ^2 as computed above does not have the sensitivity to reveal this signal. While the values of χ_ν^2 are generally well behaved given the few percent uncertainty in the bias correction, they do not tell us about our primary concern—the calibration—because both \mathcal{V} and σ scale with $|g_j g_k^*|$.

Jackknife Test

χ_ν^2 for the short *LR* baselines shows a highly significant excess over unity, so we performed a jackknife test on these data to test whether the excess represents a celestial signal. The 08^h data were split into two groups by alternating date (odd *vs.* even observation number in column 1 of Tables 5.4 and 5.5), and the two sets of visibilities were averaged by (u, v) point for each channel. Upon differencing the (u, v) point averages, we find that $\chi_\nu^2 = 1.03$ for the reals and $\chi_\nu^2 = 1.11$ for the imaginaries with $\nu = 590$ for both; the probabilities to exceed these values of χ_ν^2

²One night was an exception to this, however; the night of 13jan00 had a $\chi^2 \sim 1.5$ for all baselines, so it was deleted from the *LR* data set. No other nights were excised based on the χ^2 test.

under the null hypothesis are 30% and 3%, respectively. We cannot rule out the possibility that the signal is celestial, but the test shows that the troubling excess in Table 5.7 diminishes in significance as the visibilities are averaged.

Noise as a Proxy for the Amplitude Calibration

The 08^h deep field observations were often accompanied by observations of 3C273, which provide a measure of the consistency of the calibration from night to night. There are several nights during the 08^h field observations that lacked a suitable polarization calibration, however, because the calibrator, usually 3C279, was corrupted by small lunar elongations, and on these nights, neighboring 3C273 suffered as well. We would like to have tracers of the calibration that are consistently in view for all of the nights, and only one source satisfies this requirement: the 08^h deep field. In particular, we will use the deep field visibility *uncertainties* to infer changes in the amplitude calibration. Since $\sigma \propto |g_j g_k^*|$, the visibility uncertainties provide no insight into the phase calibration.

The visibility uncertainties offer several advantages over traditional supporting observations. The uncertainties and the visibilities undergo the same gain calibration, but since the uncertainties do not contain instrumental polarization, they are not affected by systematic errors in the leakage correction. In addition, unlike the supporting observations of polarized sources which typically last $\sim 5^m$ each night, the visibility uncertainties are measured throughout the entire session, and thus they provide us with an extensive, high S/N body of data on a quantity which is directly proportional to the gain amplitude. This use of the uncertainties is an unorthodox technique, but we will see that the errors provide a useful diagnostic of the calibration, and in particular, this analysis provides a check on the accuracy of the Tau A model developed in Chapter 3. We must guard against sources of confusion because factors which affect T_{sys} will change the uncertainties; these include clouds, spillover,³ and ice on the antenna windows. The *LL* uncertainties provide some guidance in this regard, because the considerations which affect the

³Spillover in the total power; this spillover affects *all* baselines regardless of length.

noise characteristics of the LL visibilities affect those for LR with equal force.⁴

The LR visibility uncertainties were normalized by the associated integration times to place the uncertainties on a uniform scale. After averaging by (u, v) point the number of visibilities per night can be as high as $\sim 10^3$ or more;⁵ these data were then averaged with weights to obtain a mean uncertainty for the night. The errors on the mean uncertainties were obtained from the variance of the $\sim 10^3$ visibilities for each date; Figure 5.10 shows the standard error on the mean. This procedure was repeated for the LL data calibrated with CBICAL to provide a baseline for comparison. Figure 5.10 presents a time series of the LR uncertainties (top frame) and the LL uncertainties (middle frame) as a function of the chronologically sorted observation number (corresponding to column 1 of Tables 5.4 and 5.5). The bottom frame shows the ratio $\langle\sigma^{LR}\rangle/\langle\sigma^{LL}\rangle$ for each night.

The LR visibility uncertainties provide a measure of the consistency of the amplitude calibration from night to night. Figure 5.10 shows that the variations in the mean LR uncertainties exceed those for the LL data; the rms of the uncertainty means for $\langle\sigma^{LR}\rangle$ is $\sim 7\%$, while that for $\langle\sigma^{LL}\rangle$ is $\sim 4\%$; since some of the variance seen in $\langle\sigma^{LR}\rangle$ is due to factors which give rise to that in $\langle\sigma^{LL}\rangle$ (rms $\sim 4\%$), we infer a $\sim 6\%$ rms in the gain amplitude calibration for LR . The variations in $\langle\sigma^{LR}\rangle$ in Figure 5.10 are clearly not Gaussian, however: the mean uncertainties undergo systematic drifts (dates 20-26), and data calibrated on Tau A are consistently high. In the former case, the data were calibrated on 3C279, and these drifts may reflect secular changes in the source polarization. In the latter case, they suggest an error in the polarized flux of the Tau A model; the excess in $\langle\sigma^{LR}\rangle/\langle\sigma^{LL}\rangle$ for dates 14-20 (22mar00-25mar00) is $\sim 10\%$, or 9% after correcting for the 4% increase in T_{sys} for RX12. This excess may be the consequence of $\sim 10\%$ too much polarized flux in the Tau A model. We will revisit this point in the analysis of the 20^h field uncertainties, for which the calibrations are divided evenly between 3C279 and Tau A (Section 5.4). The stability of the uncertainties for these data suggests that

⁴RX12's above average T_{sys} notwithstanding.

⁵ ~ 10 bands $\times \sim 20$ scans $\times \sim 5$ LR baselines.


Comparison of mean LR and LL visibility uncertainties; 08^h deep field

Figure 5.10: Means of visibility errors for the 08^h field observations. The dates, listed in Table 5.4, have been compressed to indices from 1 to 43. The top frame shows the visibility errors for the *LR* baselines, which were calibrated with CBIPOLCAL, while the middle frame shows the errors for the *LL* baselines calibrated with CBICAL. Since the *LL* calibration has a $\sim 5\%$ error, the *LL* errors provide our best measure of the uncertainty in the CBI’s measurements, and thus provide the benchmark which the *LR* uncertainties are to be compared. The light shaded regions denote dates which required calibration on Tau A instead of 3C279, and the dark shaded regions denote dates for which the calibration was transferred from a nearby 3C279 calibration. The third frame shows the ratio of the errors. The early January dates (days 1-2) are contaminated by ice on the dishes, and as a consequence, the noise is higher for both *LR* and *LL* on these dates, although this does not fully explain the large disparity between $\langle \sigma_{LR} \rangle$ and $\langle \sigma_{LL} \rangle$ on these dates.

Tau A can be an excellent polarization calibrator with the correct model.

Summary

The consistency tests for the 08^h field data show that the *LR* data are free of serious errors. The *LR* uncertainties are $\sim 4\%$ higher than the *LL* uncertainties, but much of this excess results from the high T_{sys} in RX12. The χ^2 test shows that the *LR* data are largely consistent with noise, and while the χ^2 for the short baseline *LR* data is $\sim 2.5\%$ too high to support the null hypothesis, the jackknife test shows that the source of the excess averages down to a negligible level of significance. The scatter in the *LR* uncertainty means suggests an overall rms of $\sim 6\%$ in the gain amplitude calibration for the 08^h field; these errors are generally noiselike, so their contribution to the total calibration uncertainty for the 43 nights of data on the 08^h field should fall by $1/\sqrt{43}$. The uncertainties for visibilities calibrated on Tau A appear to be consistently $\sim 10\%$ higher than those for data calibrated on 3C279, and this excess suggests an error of the same magnitude in the polarized component of the Tau A model. This latter conclusion conflicts with the results of Chapter 3, which suggested that the Tau A model has the appropriate amount of polarized flux. Nearly 60% of the 20^h deep field data were calibrated on Tau A, however, and in the following section we will see that this 10% discrepancy persists in these data.

5.4 20^h Deep Field

After committing the austral winter of 2000 to observations of mosaiced fields, the CBI returned to a second deep field observation in August. The 20^h field represents a rich vein of polarization data because the array configuration at that time, configuration 3 (Figure 5.11), emphasized short cross polarized baselines, and because the beginning of observations in August permitted several months of time on source. Nonetheless, the analysis of the 20^h data is complicated by the irregular coverage of the polarization calibrators during this period and by the

handoff between the 3C279 calibration in early August and the Tau A calibration from late August onward. These considerations are discussed in detail in this section.

The 20^h deep field observations spanned the three month interval between 31jul00 and 29oct00. Table 5.8 lists the observations on this field. The early data were calibrated on 3C279, while the later data were calibrated on Tau A. As the table shows, the 20^h observations include many dates for which a direct polarization calibration was not possible; the period included several intervals of a few weeks' duration during which low lunar elongations precluded observations of the polarization calibrator.⁶ Since Tau A was not due to rise above the CBI's elevations limits at night until the end of the August, the calibration for the latter part of August was transferred from the mid-August 3C279 observations. Similarly, in the middle of September and the middle of October Tau A was again obscured by the moon, and each of these intervals required an indirect calibration.

5.4.1 20^h Field Calibration

The 20^h deep field observations suffered from a paucity of polarization calibrators. At the end of July, 3C273 and 3C274 ceased to rise above the CBI's elevation limits at night, and 3C279, the prime polarization calibrator for the previous eight months, would soon follow. On 10aug00-12aug00, we performed deep observations of 3C279 with the intention of transferring 3C279's polarization calibration to Jupiter to fill the gap until Tau A rose at the end of the month. On these dates 3C279 was above the CBI's elevation limit of 45° for little more than a half an hour before it set, so these deep observations yielded at most four 5^m integrations, plus trails, per night. These observations demonstrated that Jupiter was not sufficiently polarized to provide an expedient polarization calibration for the dates to follow.

The August observations in particular lacked calibrator observations. We used

⁶August, for example, included two such intervals; in early August, 3C279 was obscured by the moon, and shortly after the lunar elongation exceeded the 60° threshold, 3C279 ceased to rise above the CBI's elevation limit at night.

CBI configuration 3: 01jul00 to 01nov00

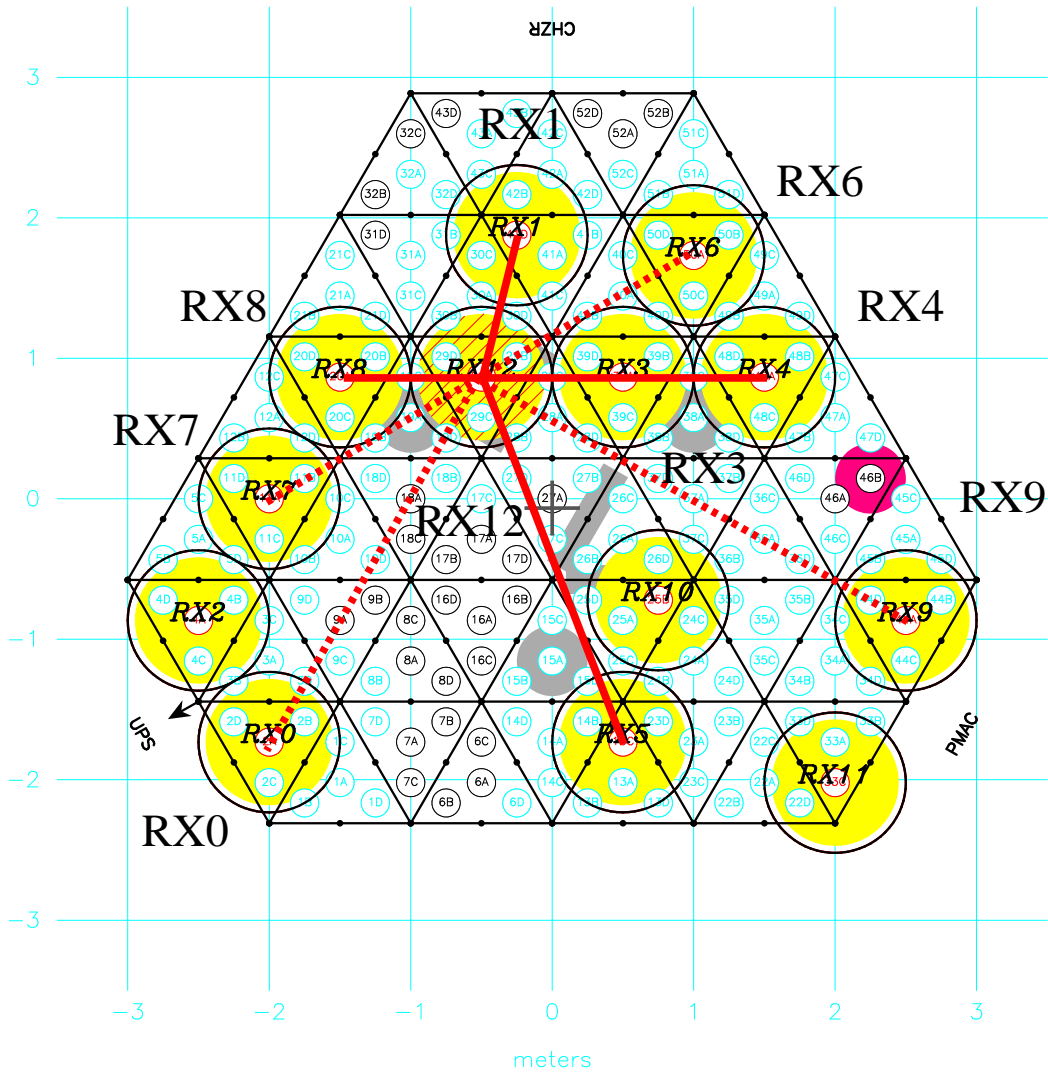


Figure 5.11: CBI configuration 3. The solid shaded lines denote the *LR* baselines which have simultaneous *LL* counterparts. The dotted dashed lines denote baselines which can be calibrated after some multiple of a 30° rotation.

#	date	calsource	caldate	scans per baseline length (cm)					
				100	104	173	200	278.4	300
1	31jul00	3C279	10aug00	44	22	44	22	22	22
2	01aug00	3C279	10aug00	44	22	44	22	22	22
3	02aug00	3C279	10aug00	44	22	44	22	22	22
4	03aug00	3C279	10aug00	44	22	44	22	22	22
5	04aug00	3C279	10aug00	32	16	24	16	16	12
6	05aug00	3C279	10aug00	44	22	44	22	22	22
7	06aug00	3C279	10aug00	40	20	32	20	20	16
8	07aug00	3C279	10aug00	40	20	32	20	20	16
9	08aug00	3C279	10aug00	38	19	30	19	19	15
10	09aug00	3C279	10aug00	44	22	40	22	22	20
11	18aug00	3C279	10aug00	40	20	32	20	20	16
12	19aug00	3C279	10aug00	38	19	30	19	19	15
13	21aug00	3C279	10aug00	24	12	8	12	12	4
14	22aug00	3C279	10aug00	32	16	24	16	16	12
15	23aug00	3C279	10aug00	36	18	28	18	18	14
16	29aug00	Tau A	—	38	19	30	19	19	16
17	30aug00	Tau A	—	34	17	26	17	17	14
18	31aug00	Tau A	01sep00	38	19	30	19	19	16
19	01sep00	Tau A	—	28	14	24	14	14	12
20	02sep00	Tau A	01sep00	34	17	26	17	17	14
21	16sep00	Tau A	24sep00	36	18	28	18	18	14
22	18sep00	Tau A	24sep00	36	18	28	18	18	14
23	19sep00	Tau A	24sep00	36	18	28	18	18	14
24	20sep00	Tau A	24sep00	32	16	24	16	16	12
25	21sep00	Tau A	24sep00	30	15	24	15	15	12
26	22sep00	Tau A	24sep00	28	14	24	14	14	12
28	23sep00	Tau A	24sep00	18	9	14	9	9	7
29	24sep00	Tau A	—	28	14	24	14	14	12
30	25sep00	Tau A	—	28	14	24	14	14	12
31	26sep00	Tau A	—	30	15	24	15	15	13
32	27sep00	Tau A	—	32	16	28	16	16	14
33	28sep00	Tau A	—	22	11	22	11	11	11
34	29sep00	Tau A	—	26	13	24	13	13	12
35	30sep00	Tau A	—	22	11	22	11	11	11
36	01oct00	Tau A	—	26	13	24	13	13	12
37	02oct00	Tau A	—	26	13	24	13	13	12
38	12oct00	Tau A	20oct00	16	8	8	8	8	4
39	13oct00	Tau A	20oct00	18	9	14	9	9	7
40	14oct00	Tau A	20oct00	20	10	16	10	10	8
41	15oct00	Tau A	20oct00	18	9	14	9	9	7

Table 5.8: 20^h field observations, configuration 3. Table continued on next page.

#	date	calsource	caldate	scans per baseline length (cm)					
				100	104	173	200	278.4	300
42	16oct00	Tau A	20oct00	10	5	2	5	5	2
43	17oct00	Tau A	20oct00	18	9	14	9	9	7
44	18oct00	Tau A	20oct00	22	11	22	11	11	11
45	19oct00	Tau A	20oct00	18	9	14	9	9	7
46	20oct00	Tau A	—	18	9	14	9	9	7
47	21oct00	Tau A	—	18	9	14	9	9	7
48	22oct00	Tau A	—	12	6	4	6	6	2
49	23oct00	Tau A	—	8	4	0	4	4	0
50	24oct00	Tau A	—	14	7	6	7	7	4
51	25oct00	Tau A	—	12	6	4	6	6	2
52	26oct00	Tau A	—	14	7	6	7	7	4
53	27oct00	Tau A	—	6	3	0	3	3	0
54	28oct00	Tau A	—	10	5	2	5	5	2
55	29oct00	Tau A	—	14	7	6	7	7	4
total				1478	739	1182	739	739	601
expected σ_{LR} (mJy/beam)				2.33		2.51		3.00	
measured σ_{LR} (mJy/beam)				2.70		2.87		3.60	

Table 5.9: 20^h field observations, configuration 3, continued. Sensitivity is computed based on 6.5^m per scan.

the deep observations of 3C279 on 10aug00-12aug00 to derive a benchmark calibration which was transferred to the preceding and following nights with the noise source scaling technique described in Chapter 3; the total intensity observations on these dates were used to infer the noise cal amplitude errors, and these corrections were applied to the gain amplitudes for the cross polarized baselines. The uncertainties on the VLA data for 3C279 on 10aug00 are $\delta m/m \sim 2\%$ and $\delta\chi/\chi \sim 8\%$. The discussion of Chapter 3 demonstrated that the scaling techniques tend to improve the calibration to $\sim 10\%$ over the $\sim 15\%$ intrinsic gain variations from night to night. Tau A rose above our elevation limits at night for the first time at the end of August, and it was the only polarization calibrator for the remainder of

the 20^h field observations. We used the Tau A model presented in Chapter 3 for these calibrations. Tau A is at $\alpha \sim 5.5^h$, so there are regular gaps of about a week during which the moon obscured Tau A but not the 20^h field; during these dates we transferred the calibration with the noise cal scaling technique using Jupiter as the noise cal error reference source. We assess the errors on these calibrations in the next section.

5.4.2 20^h Field Results

Figures 5.12 and 5.13 present maps of Q and U for the 20^h field. The map noises for Q and U are typically 4-6 mJy/beam both inside and outside of the central region of the maps, so the maps contain no obvious evidence for a signal. In contrast, the I map, which shows a clear signal, has an rms of ~ 31 mJy/beam in the central part of the beam, and ~ 10 mJy/beam elsewhere (Figure 5.14). The maps for Q and U include fewer visibilities than we can calibrate for LR because not all visibilities had the RL counterparts required to form Q and U . Table 5.10 lists the values of σ^{LR} derived from the 20^h visibilities. These values are $\sim 15\%$ higher than those predicted by the integration time on the field; we will see that this excess represents a combination of calibration errors and a real excess for the LR visibilities due to the higher T_{sys} for RX12. The Q and U maps for bands 2 and 3 (the intermediate and long baselines) are similar in that they show no evidence for a signal.

5.4.3 20^h Field Consistency Tests

The consistency tests for the 20^h field are the same as those for the 08^h field. We first compare the measured visibility uncertainties to expected values, and then compute χ^2 to establish that the visibilities are consistent with noise. Finally, we use the uncertainties on the visibilities to assess amplitude calibration. The latter test confirms the 10% amplitude error in the Tau A model that was suggested by 08^h visibility uncertainties.

Residual Q map. Array: CBI
 C2048-0330 at 31.000 GHz 2000 Aug 01




Figure 5.12: Map of Q for the 20^h deep field, short baselines ($|u| < 150\lambda$). The circle shows the extent of the primary beam at the band center: $45.2'$ (FWHM) at 31 GHz. The noise in the central region is identical to that outside of the area defined by the primary beam.

Residual U map. Array: CBI
C2048-0330 at 31.000 GHz 2000 Aug 01




Figure 5.13: Map of U for the 20^h deep field, short baselines ($|u| < 150\lambda$). See Figure 5.12 for details.

Residual LL map. Array: CBI
 C2048-0330 at 31.000 GHz 2000 Aug 01




Figure 5.14: Map of LL for the 20^h deep field, short baselines ($|u| < 150\lambda$). See Figure 5.12 for details. The noise in the central region of the map is nearly three times that outside of the primary beam.

20^h Field Visibility Uncertainties

Table 5.10 lists the mean visibility uncertainties for the 20^h field observations. Nearly 40% of the 20^h field data were calibrated on 3C279, while the balance was calibrated on Tau A, so the table sorts the uncertainties by calibrator: Set 1 corresponds to the first 15 nights (3C279), and Set 2 corresponds to the remaining 40 nights (Tau A). The values in the table are the means of the uncertainties for all of the visibilities for the sets in question.

The *LR* data in the first column of Table 5.10 were calibrated on 3C279 and Jupiter; 3C279 sets the polarization scale through the fractional polarization m , and Jupiter sets the absolute flux scale through the noise cal scaling. The *LL* uncertainties (calibrated in CBICAL on Jupiter) are $\sim 4\%$ larger than those for *LR*. Given that RX12 has $\sim 4\%$ more noise than the array average, this difference suggests a systematic amplitude calibration error of at least $\sim 6 - 8\%$ for the data that were calibrated on 3C279. An error of this size comes as no surprise, however, as most of the dates during this period required some scaling, and we do not expect the scaling procedure to improve the calibration to better than this level.

The most striking feature of Table 5.10 is that the *LR* uncertainties for the data calibrated on Tau A are significantly higher than the other values. This attribute, which was seen in the 22mar00-25mar00 08^h field data, is almost certainly the result of $\sim 10\%$ excess polarized flux in the Tau A model.

Row three of Table 5.10 lists the mean uncertainties for the *LL* data calibrated with CBIPOLCAL. We report these values to confirm that they agree with those for the *LL* data calibrated with CBICAL; since these *LL* data set the flux scale for the *LR* calibration, we are encouraged that they agree. Finally, Table 5.6 shows that the *LL* uncertainties for the 20^h field agree with those for the 08^h field.

χ^2 Test

Table 5.11 reports the results of the χ^2 analysis of the 20^h visibilities. As expected, the short *LL* baselines show a slight excess due to the presence of a real signal, while the long baselines are slightly less than unity, perhaps reflecting an error in

\mathcal{V}	calibration	Set 1 $\langle\sigma\rangle$ (Jy)	Set 2 $\langle\sigma\rangle$ (Jy)	all $\langle\sigma\rangle$ (Jy)
\mathcal{V}^{LR}	CBIPOLCAL	5.47 ± 0.04	6.32 ± 0.03	6.01 ± 0.06
\mathcal{V}^{LL}	CBICAL	5.68 ± 0.05	5.81 ± 0.02	5.76 ± 0.02
\mathcal{V}^{LL}	CBIPOLCAL	5.78 ± 0.05	5.84 ± 0.04	5.82 ± 0.03

Table 5.10: Average noise in the 20^h visibility data. These mean values are weighted averages whose weights are obtained from the time on source per night. The Set 1 *LR* data were calibrated with 3C279 to obtain m and χ , although the absolute *LR* flux scale was derived from contemporaneous observations of Jupiter. Row 2 shows the *LL* benchmark set by the calibration in CBICAL; these data were calibrated on Jupiter. Row 3 shows the *LL* calibration in CBIPOLCAL; the Set 1 data were calibrated on 3C279, whose model for I was obtained from the *LL* calibration in CBICAL, so rows 2 and 3 in column 1 should agree. The Set 2 *LL* and *LR* data calibrated in CBIPOLCAL used Tau A as the calibrator for the entire calibration, so the flux scale for the *LR* calibration was obtained directly from the well tested I model for Tau A.

our estimate of the noise. The χ_ν^2 for the short *LR* baselines are perhaps a bit high given that the integration time on this field falls short of the level necessary to detect a signal, and in Chapter 6 we explore whether this excess reflects a real signal. The χ_ν^2 for the long baselines are all 1-2% low; this is our best measure of χ_ν^2 since the values for the all baseline average are boosted by the short baseline data. The error in the bias factor β (Section 5.3.3: $\beta = 1.06 \pm 0.01$) is not quite enough to bring the long baseline χ_ν^2 into agreement with unity; these data suggest that the estimates of the noise are a few percent too large. We saw this attribute in the χ_ν^2 test for the long baseline 08^h data.

Jackknife Test

The high χ_ν^2 for the short *LR* baselines for $\nu = 33682$ contradicts the null hypothesis, so we performed a jackknife test on these visibilities. As with the 08^h data, the data were sorted into two groups based on interleaved dates. The two sets of data were averaged and differenced. We find that χ_ν^2 is 1.003 for the reals and 0.974 for

		$ u < 150 \lambda$		$ u > 150 \lambda$		all $ u $	
\mathcal{V}	calibration	χ_ν^2	ν	χ_ν^2	ν	χ_ν^2	ν
\mathcal{V}^{LR}	CBIPOLCAL	1.029	33682	0.977	42272	1.000	75954
\mathcal{V}^{LL}	CBICAL	1.079	61562	0.985	520696	0.995	582258
\mathcal{V}^{LL}	CBIPOLCAL	1.072	67042	0.985	493508	1.000	560550

Table 5.11: χ^2 tests on the 20^h deep field data. $\nu = 2N_d$ to account for the real and imaginary components of the visibilities.

the imaginaries with $\nu = 720$; these values are highly consistent with unity. These results suggest that the source of the excess in Table 5.11 loses significance as the data are averaged, and it is thus not a great concern.

Noise as a Proxy for the Gain Calibration

Section 5.3.2 presents an analysis of the gain calibration in which the visibility uncertainties are used to track changes in the amplitude of the gain calibration. We apply the same technique to the 20^h data. Figure 5.15 shows a time series of the LR and LL uncertainties for the 20^h field observations; the first and third frames of the figure show a clear demarcation in the calibration at the transition between Set 1 (3C279) and Set 2 (Tau A). After sorting the ratios by set, we find that $\langle \sigma^{LR} \rangle / \langle \sigma^{LL} \rangle = 0.96$ for Set 1, while $\langle \sigma^{LR} \rangle / \langle \sigma^{LL} \rangle = 1.09$ for Set 2, in agreement with the data reported in Table 5.10. We interpret the discrepancy for the Set 2 data as evidence for a 10% excess in the polarized flux for the Tau A model. The rms of the uncertainties is quite small: 3% for the LL data and 7% for the LR data; these scatters include intrinsic variations in T_{sys} , so they provide conservative estimates of the repeatability of the amplitude of the gain calibration. If the interpretation of the mean uncertainties in Table 5.10 is to be believed, these calibration offsets cancel, for a net systematic amplitude calibration error for the entire 20^h LR data set of $\sim +4\%$.


Comparison of LR and LL visibility uncertainties; 20^h deep field

Figure 5.15: Top: *LR* baselines; middle: *LL* baselines; bottom: ratio of the two. Some dates do not have *LL* observations; these dates lacked calibrator observations, and the calibration requirements for the total intensity observations ($\leq 5\%$) cannot be met with a bootstrapped calibration. Heavy shading: bootstrapped (scaled) calibration on 3C279. Light shading: bootstrapped calibration on Tau A.

5.5 Supporting Observations

We observed a variety of polarized sources to test the polarization performance of the CBI. The 08^h deep field observations were nearly always accompanied by observations of 3C274 and 3C273, the former because it provided the total intensity calibration, and the latter because it provides a consistency check on the calibrations for *LL* and *LR*; both of these sources proved to be useful in evaluating the polarization calibration. 3C273 is a 20 Jy, $\sim 5\%$ polarized quasar which was included in the VLA monitoring program, and it provides an important check of

the internal consistency of the calibration. 3C274 is a 17 Jy, $\sim 4\%$ polarized radio galaxy whose stability allows us to assess the repeatability of the polarization calibration. 3C273 and 3C274 are unresolved by the CBI, however, so observations of these sources do not test the CBI’s mapping capabilities. To evaluate the CBI’s response to extended polarized emission, we observed several resolved polarized sources: the double inner lobes of the active radio source NGC 5128 (Centaurus A), and the emission nebulae of two galactic SNR, W44 and G326.3-1.8. We will see that the 3C273 and 3C274 observations permit a quantitative assessment of the calibration, while the extended sources merely provide qualitative evidence for the CBI’s mapping capabilities. Since our goal is to evaluate the polarization performance of the system, we focus on comparisons of m and χ to the exclusion of I .

5.5.1 3C273

The polarized radio galaxy 3C273 provides our best test of the internal consistency of the polarization calibration. 3C273 is a bright, moderately polarized extragalactic radio source whose region of polarized emission is unresolved by both the CBI and the VLA. 3C273 was observed on a routine basis during the 08^h field observations with the CBI, and throughout the concurrent VLA polarization monitoring campaign on 3C279. A comparison of the CBI measurements of 3C273 with the data from the VLA tests the precision of the reduction pipeline.

The analysis of the CBI observations of 3C273 required several departures from the standard technique for 3C279. Most of the CBI observations of 3C273 consisted of single 5^m scans, without trails, so the short baseline visibilities were corrupted by spillover. The lack of trails does not prevent the use of these data, however, because spillover has no discernable effect on baselines longer than ~ 150 cm; we simply deleted the short baselines. In addition, the 3C273 observations consisted of single scans, which for extended structure would preclude the use of the data for polarization; as noted in Section 2.2, observations separated by a deck angle of 180° are required to obtain the full polarization characteristics— LR and

RL , or equivalently, Q and U —of the source. For unresolved sources at the phase center, however, \tilde{Q} and \tilde{U} are both real, so the polarization of 3C273 is readily obtained from the calibrated LR visibilities. Because 3C273 is unresolved, all of the baselines view the same LL structure, so we can apply the leakage corrections to all LR baselines, even those which lack LL matches.

Figure 5.16 provides a comparison of the CBI and the VLA results for the polarization of 3C273. This comparison required some interpolations. The shaded regions denote the periods during which the 08^h field was observed; some deep field observations were not accompanied by 3C273 observations simply due to low lunar elongations or truncated observing schedules.

The comparison in Figure 5.16 shows that the CBI does a fair job of recovering the polarization characteristics measured by the VLA for 3C273. If we exclude the 22mar00-25mar00 observations, the mean fractional polarization discrepancy between the CBI values and the VLA values for m is $\langle \Delta m \rangle \sim 0.0025$; this systematic offset corresponds to $\sim 5\%$ of m or $< 1\%$ of the total intensity. The mean discrepancy for χ is $\langle \Delta \chi \rangle \sim 2^\circ$, or $\sim 5\%$. These offsets agree, and they are marginally consistent with the cumulative uncertainties in the interpolations of the VLA data for 3C279 (3-5%) and 3C273 (2-5%).

We excluded the 22mar00-25mar00 data from the analysis above; these observations were calibrated on Tau A, and the corresponding values for m and χ are in poor agreement with the VLA values. The mean fractional polarization discrepancy for Tau A is $\langle \Delta m \rangle = 0.0012$, or $\sim 20\%$, while that for χ is $\langle \Delta \chi \rangle = 7^\circ$, or 18%. These errors, which are significant at the few σ level, point to an error in the Tau A model which provided the calibration. The visibility errors suggested that the Tau A model had $\sim 10\%$ too much LR flux, and this comparison supports this hypothesis given the combined errors in the interpolations for 3C279 and 3C273.

5.5.2 3C274

3C274 (M87, Virgo A) is a nearby giant elliptical galaxy which served as the primary total intensity calibrator throughout the 08^h field observations. 3C274 is an


Comparison of VLA and CBI values for m and χ for 3C273 at 31 GHz

Figure 5.16: Comparison of 3C273 as measured with the CBI (points) and the VLA (lines with $\pm 1 \sigma$ bands). Shaded regions denote times during which the deep fields were in view. The first four dates in late March are dates whose data were calibrated on Tau A; the fractional polarization on these dates shows a clear systematic offset relative to the expected values, a problem which is no doubt a consequence of the Tau A model. Since m and χ are in the high S/N limit, no correction was made for bias from the noise in P .

outstanding total intensity calibrator because its large emission region limits the temporal variations in its total intensity. 3C274 is also marginally polarized at centimeter wavelengths, and this characteristic, coupled with its high stability, makes it a good source on which to test the repeatability of the polarization calibration for the duration of the 08^h field observations. 3C274 was observed each night for a single 5^m scan, without trails, immediately prior to the 3C273 and 3C279 observations discussed above. The reduction of the 3C274 data was identical to that for the 3C273 observations. 3C274 was observed on $\sim 75\%$ of the nights of the 08^h deep field observations, and none of the nights of the 20^h deep field observations.

The 3C274 observations provide an upper limit on the combined variability of the calibration and the polarization of 3C274. Figure 5.17 shows the 3C274 polarization data as measured by the CBI. The fractional polarization varies about a mean of 4.1% with a fractional rms of $\sigma_m/\langle m \rangle = 15\%$, although the figure shows that the scatter is not noiselike. In fact, the fractional polarization tends to cluster in groups in a manner that suggests that the polarization calibration or the source undergoes systematic errors from epoch to epoch. The peak to peak scatter in the groups is $\Delta m \sim 0.015$, which is larger than the 5% uncertainty in the VLA calibration for m ($0.05 \times 0.1 \sim 0.005$) and the error seen in the 3C273 data (5% in m). The values for χ are a bit more stable than those for m ; they vary about a mean of 46° with an rms of 3°, or $\sim 7\%$. It is not unusual to find that the position angles are in better agreement than measurements of the fractional polarization; the position angle is the easier of the two quantities to calibrate.

The analysis of the 3C273 observations suggests that the excess polarized flux for the dates calibrated on Tau A arises from an error in the Tau A model. It is therefore striking that the 3C274 data for the same dates show no excess (Figure 5.17); in fact, on these dates (days 82-85) 3C274 appears to have a slightly *lower* fractional polarization. This feature, which may be intrinsic to 3C274, complicates our attempts to assess the accuracy of the Tau A model.




Figure 5.17: CBI observations of 3C274. This source was not monitored with the VLA.

5.5.3 Centaurus A

Centaurus A (NGC 5128) is a nearby active galaxy whose polarized emission provides another test of the polarization imaging capabilities of the CBI. Centaurus A has a rich variety of structure at centimeter wavelengths that extends to angular scales as large as $\sim 10^\circ$; we will focus on the small, central region which contains components referred to as the double inner lobes [38]. The DILs are $\sim 10'$ long and highly polarized at centimeter wavelengths, so they provide a laboratory for exploring the polarization mapping capabilities of the CBI.

Centaurus A has been the subject of many radio studies, but few permit a direct comparison with CBI observations because these observations differ in resolution and frequency from those of the CBI. The observations which permit the best

comparison with the CBI data are those of Junkes *et al.* [38], who used the Parkes 64 meter telescope at 6.3 cm to map the DILs with 4.3' resolution. While this resolution approaches that of the CBI observations, the frequency at which the Parkes observations were performed is substantially lower, so a comparison with the CBI data could be complicated by frequency dependent changes to the polarization.


CBI Observations of Centaurus A

Centaurus A was observed for 6.8^h on 08jun00, 13jun00, and 14jun00. We are particularly interested in the $\sim 10'$ core, so to maximize the time on source these observations exclude trails, and the short baselines were excised from the data. Figure 5.18 shows the CBI map of the double inner lobes, along with the southernmost edge of the northern middle lobe. The image is centered on the northern end of the double inner lobe, at which point the total intensity peaks at 20.1 Jy, the fractional polarization reaches 12% and the position angle is -36 degrees. The CBI's resolution of 6.5' FWHM allows us to discern the polarization gradient along the two lobes. While the total intensity of the southern lobe agrees with that in the northern lobe—it peaks at 18.7 Jy/beam—the polarization characteristics of the southern lobe are strikingly different; the fractional polarization reaches 3.6% at the total intensity peak, at which point the PA is $\sim -37^\circ$. Continuing to the south, the position angle winds around to $\sim 10^\circ$.

Parkes Observations of Centaurus A

The Junkes observations at 6.3 cm are qualitatively consistent with the CBI maps. Figure 5.19 shows the Parkes map of the total intensity of the DILs with the polarization vectors superposed, and Figure 5.20 shows the contours of the polarized emission. The authors find that over the northern inner lobe the fractional polarization peaks at 13%, while at the peak of the total intensity of the southern inner lobe the polarization rises to only $\sim 5\%$ at the southernmost edge of the lobe. The position angle across the two inner lobes is $-70^\circ < \chi < -33^\circ$, and it wraps around to $\sim +5^\circ$ along the southern slope of the southern inner lobe. These polar-

Clean LL map. Array: CBI
CenA at 31.000 GHz 2000 Jun 08



Map center: RA: 13 25 27.600, Dec: -43 01 08.800 (2000.0)
Map peak: 23.6 Jy/beam
Contours %: 1 2 4 8 16 32 64
Beam FWHM: 6.59 x 6.48 (arcmin) at -11.7°

Figure 5.18: CBI map of the double inner lobes of Centaurus A at 31 GHz.

ization findings are in good agreement with the characteristics inferred from the CBI observations of Centaurus A; the fractional polarization measurements for the northern inner lobe agree, while those for the southern inner lobe differ by at most a factor of two. The authors compare their 6.3 cm observations with their 3.5 cm maps to be released in a later paper to obtain a rotation measure of -60 rad/m^2 towards the central region of Centaurus A; this small rotation measure suggests that Faraday rotation between these observations and 1 cm will be on the order of a few degrees, or negligible.⁷ Given the difficulties in comparing different polarization observations, the agreement between the CBI and the Parkes observations is good.

5.5.4 W44

Supernova remnant W44 provides another test of the polarization mapping capabilities of the CBI; it has several janskys of polarized emission at 1 cm, and its size of $\sim 30'$ matches the CBI's primary beam. The remnant has a pear-shaped shell, with a distinct asymmetry arising from the steep density gradient in the immediate neighborhood of the remnant [19]. W44 is a composite remnant because its emission is divided between a steep spectrum shell and a flat spectrum core. There is little quantitative discussion in the literature of its polarization on arcminute scales and at GHz frequencies.

CBI Observations of W44

The CBI observed W44 for a total of 2.6^h on 22jun00 and 24jun00. These observations were accompanied by trails. The data were calibrated on 3C279. Figure 5.21 shows the CBI map of W44 after being restored with a $8.6' \times 7'$ beam. This polarization map was obtained by cleaning the maps for I , Q , and U down to the a noise level of $\sim 50 \text{ mJy/beam}$ peak to peak, so the S/N in I is ~ 50 and the S/N in P is ~ 8 over much of the source. Since we are primarily concerned with the frac-

⁷Burns suggests that the SW lobe is more strongly depolarized because it is behind the source, so radiation from the Southern Lobe encounters more depolarization from the ISM [12].




Figure 5.19: Parkes 6.3 cm map of the total intensity of central region of Cen A; contours are in mJy/beam.




Figure 5.20: Parkes 6.3 cm map of the polarization of central region of Cen A.




Figure 5.21: CBI map of supernova remnant W44 and the galactic HII region G34.3+0.1

tional polarization, no correction was made for the CBI primary beam, although removing the beam would have the effect of adding additional total intensity flux to the edge of the remnant. The CBI maps show that the fractional polarization peaks at $\sim 33\%$ on the Northwestern slope of the source, and across the center of the source it is relatively uniform at 10-12%. While the position angle varies across the source, it is roughly uniform at $\sim 60^\circ$ across most of the emission in total intensity.

The neighborhood of W44 contains a galactic HII region which provides an additional demonstration of the CBI's polarization capabilities. The emission from

this source, G34.3+0.1, is due to free-free emission, so the source should be unpolarized. The fractional polarization at the total intensity of the emission is $\sim 0.5\%$, so we may conclude that the CBI is not creating spurious polarization at more than the 1% level. It is probably not correct to infer that the leakage correction is good to this level, however, because the image is the sum of visibilities measured at a variety of deck positions, and any leakage errors will add incoherently and will thus be suppressed. Nor does this test allow us to argue that the CBI is not *suppressing* real polarized emission on the sky. At a minimum, however, this map suggests that qualitatively, the CBI's mapping capabilities are consistent with expectation.

NRAO 140' Observations of W44

Kundu and Velusamy used the NRAO 140' telescope to map W44 at 10.7 GHz with a 3' beam [45]. Figure 5.22 shows their map of the source. The authors find that the fractional polarization peaks at $\sim 20\%$ along the NE edge, and it remains uniform over the dominant region of emission along the east side of the source. At the peak of emission in total intensity the authors find that the fractional polarization $m \sim 20\%$. They do not note a position angle, although the figure shows that the position angle agrees qualitatively with the position angles measured by the CBI. The maps also agree qualitatively in the extent of the polarized emission across the source. No other data in the literature come nearly as close to matching the CBI's resolution and frequency coverage.

5.5.5 G326.3-1.8

SNR G326.3-1.8 is a composite SNR; it consists of a bright central source and a low surface brightness shell [21]. G326.3-1.8 has been studied by a number of authors, although, as with W44 and Centaurus A, comparisons are hampered by resolution and wavelength-dependent effects.




Figure 5.22: Map by Kundu and Velusamy of W44 at 2.8 cm with a 3' beam.

CBI Observations of G326.3-1.8

The CBI observed G326.3-1.8 for a total of 1.3^h on 07jul00 and 12aug00. The CBI images show a smooth plateau of emission punctuated by a bright region of emission from the plerion near the center of the remnant. The emission across the plateau is typically ~ 0.8 Jy/beam, and the emission from the plerion rises sharply to 7.6 Jy/beam at the peak. The fractional polarization across the plateau is quite large—it averages 20-30%, while the fractional polarization across the plerion drops to $\sim 8\%$. The EVPA differs between the two components as well; on the shell the EVPA is typically -20° to -30° , while at the plerion it winds around to $\sim 10^\circ$.

Parkes Observations of G326.3-1.8

Milne *et al.* have mapped G326.3-1.8 at 8.4 GHz with the Parkes 64 meter telescope; the authors provide maps of the source which have been smoothed to 4.6' [53]. Their 8.4 GHz map (Figure 5.24) does not permit a quantitative comparison with that from the CBI, but the qualitative agreement between the two is immediately apparent. Milne finds that the fractional polarization at the peak of the total intensity is $\sim 10\%$, while on the surface of the shell it rises to $\sim 20\%$.




Figure 5.23: G326.3-1.8 map obtained with the CBI.

These values generally agree with those from the CBI. The change in position angle across the source in the Parkes map resembles that for the CBI map as well.

5.6 Conclusions

Our paramount concern throughout this work is the calibration, and the most important result of this chapter is an estimate of the uncertainty of the deep field calibration. Our efforts to assess this uncertainty are hampered by the paucity of calibrators, the differences in calibration techniques, and the small number of supporting observations which parallel the deep observations. We attacked this problem from two directions: for the data calibrated on 3C279 we adopted a bottom-up approach in which we propagate the errors in all of the factors which affect the cal-




Figure 5.24: G326.3-1.8 map from the parkes 64 meter telescope at 8.4 GHz [53] The map has been smoothed to $4.6'$ by the authors.

ibration, while for the data calibrated on Tau A we employed a top-down approach in which we used the supporting observations to infer the calibration uncertainty. The limited amount of data is such that neither approach is ideal, so we choose a compromise between the two.

The 3C279 calibration is the most amenable to analysis. The mean uncertainty in the amplitude and phase for the VLA values of m and χ can be as high as 10%, although for most dates the errors are far smaller. The analysis of the values for m and χ from the CBI observations of 3C273 implies 3% uncertainties in both amplitude and phase, while the scatter in the uncertainty ratio σ_{LR}/σ_{LL} suggests a 6% amplitude calibration error. The former test does not reflect the uncertainty in the absolute calibration of the CBI, while the latter does. We adopt a conservative 5-10% uncertainty for the 08^h field data calibrated on 3C279. The analysis of the visibility uncertainties for the 20^h field data calibrated on 3C279 is a somewhat more problematic, as these data suggest a consistent 4% deficit in polarized flux

for 3C279; we adopt a conservative 10% uncertainty for these data.

In Chapter 3 we used dead reckoning to calibrate the deep Tau A observation of 11jan00, and the model obtained on this date was applied to many of the subsequent 08^h and 20^h deep field observations. One important result of this chapter is that a variety of tests suggest that the Tau A model has $\sim 10\%$ too much polarized flux; this excess is seen in the visibility uncertainties and the supporting observations of 3C273. And yet the CBI data remain inconclusive; the test on 3C279 on 01mar00-05mar00 discussed in Chapter 3, coupled with the tests on the leakage, suggest that the model is better than 10%, and the 3C274 *LR* data show none of the behavior seen in the 3C273 *LR* data. For the sake of argument we can assume that the Tau A model has $\sim 10\%$ too much polarized flux, and evaluate the leakage after scaling Q and U components by 0.9. Figure 5.25 shows the results of this analysis for the pair of baselines considered in Chapter 3; the figure shows that the the effect of scaling the polarized flux in the model by 0.9 is at best marginal for the two baselines in question. A quantitative comparison of all of the leakages shows that while the leakage amplitudes obtained with the standard Tau A model exceed the values obtained from the deep 06feb00 observations by $\langle \Delta \epsilon_A \rangle = 3.1\%$, the amplitudes obtained with the scaled Tau A model fall short of the 06feb00 values by $\langle \Delta \epsilon_A \rangle = 7.2\%$. We see that while the 10% change in Tau A's polarized flux translates into a $\sim 10\%$ change in the leakage amplitude, both Tau A models straddle the expected value. The phase offset for the leakage $\langle \Delta \epsilon_\phi \rangle \sim 10\%$, and the 90% amplitude scaling does not affect this result. In light of these problems, we adopt a uniform 10% uncertainty for the amplitude and the phase of the Tau A model. In Chapter 6 we explore the effects of systematic calibration errors on the best fit bandpowers for the CMBR data.


CBI Instrumental Polarization Comparison; 3C279 vs. Tau A model $\times 0.9$ 

Figure 5.25: Cross check on the Tau A model using leakages. Triangles denote leakage derived with Tau A Model 1, and circles denote leakage obtained after Model 1 has been scaled by 0.9. The comparison shows that for these two baselines the leakage is fairly insensitive to a 10% change in the fractional polarization of the model, although for all baselines the scaling of 0.9 widens the discrepancy with respect to the reference leakage (3C279; squares) from $\sim 3\%$ to $\sim 7\%$ (see text).

Chapter 6

Likelihood Analysis of the *LR* Visibilities

6.1 Introduction

The simple χ^2 test discussed in Chapter 5 suggests that the *LR* visibilities are dominated by noise, and the images of the deep fields support this conclusion. The χ^2 test makes a gross approximation, however, because it fails to consider the high degree to which the visibility data oversample the underlying fluctuations, and a complete treatment must accommodate these correlations. We adopt a maximum likelihood power spectrum estimator to test the *LR* visibilities for the presence of a polarization signal and evaluate the likelihood for a range of trial polarization spectra. The trial spectrum that maximizes the likelihood is the best estimate for the underlying spectrum. This chapter discusses the details of the maximum likelihood calculation for the *LR* visibilities, and presents limits on the polarization of the CMBR in the 08^h and the 20^h deep fields.

6.2 Method of Maximum Likelihood

The method of Maximum Likelihood rests on Bayes' theorem of conditional probability. The signal reflects an underlying theory, and we want to estimate the most likely spectrum of fluctuations, \mathbf{q} , given the data, \mathbf{x} . Unfortunately, this quantity cannot be computed directly, but Bayes' theorem tells us that this unknown is proportional to a quantity that we can compute: the theory which is most likely

to produce our data. Bayes' theorem states that $P(\mathbf{x} \cap \mathbf{q})$ —the joint probability of \mathbf{x} and \mathbf{q} —can be expressed in terms of conditional probabilities:

$$P(\mathbf{x} \cap \mathbf{q}) = P(\mathbf{q}|\mathbf{x})P(\mathbf{x}) = P(\mathbf{x}|\mathbf{q})P(\mathbf{q}) \quad (6.1)$$

where $P(\mathbf{x})$ and $P(\mathbf{q})$ are the priors for \mathbf{x} and \mathbf{q} .

This simple statement is the basis of the method of maximum likelihood. In the case at hand, our goal is to determine the peak of $P(\mathbf{q}|\mathbf{x})$: the most likely spectrum \mathbf{q} given the data \mathbf{x} . We cannot directly compute this quantity, but since it is proportional to $P(\mathbf{x}|\mathbf{q})$, we simply search for the theory which is most likely to produce the data, and that theory provides our best estimate. $P(\mathbf{x}|\mathbf{q})$ is the *likelihood* $\mathcal{L}(\mathbf{x}|\mathbf{q})$ of the data given the theory. The prior probabilities $P(\mathbf{x})$ and $P(\mathbf{q})$ are constant—we have one realization of the data, and we assume a uniform prior for the theory—so these factors do not affect the maximization procedure.¹

The first task is to compute the likelihood $\mathcal{L}(\mathbf{x}|\mathbf{q})$. The likelihood is simply the joint probability distribution function of the N_d visibilities, the real and imaginary components of which are assumed to be Gaussian. The width of the autocorrelation of the illumination pattern in the aperture domain, σ_p , coupled with our standard observing strategy, results in correlations between the visibilities given by the covariance matrix \mathbf{C} . For Gaussian random variables \mathbf{x} , the likelihood has the form

$$\mathcal{L}(\mathbf{x}|\mathbf{q}) = \frac{1}{\pi^{N_d} |\mathbf{C}|} \exp \left[-\mathbf{x}^t \mathbf{C}^{-1} \mathbf{x} \right] \quad (6.2)$$

$|\mathbf{C}|$ denotes the determinant of \mathbf{C} . The vector \mathbf{x} contains N_d complex visibilities: $\mathbf{x} = \{\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_{N_d}\}$. For computational simplicity, we follow the standard practice and work with the log-likelihood:

$$\ln[\mathcal{L}(\mathbf{x}|\mathbf{q})] = -N_d \ln[\pi] - \ln[|\mathbf{C}|] - \mathbf{x}^t \mathbf{C}^{-1} \mathbf{x} \quad (6.3)$$

¹In principle, we can use existing measurements of $\{\Omega_0, \Omega_\Lambda, \Omega_b h^2, \dots\}$ to weight regions of parameter space and thereby assign probabilities to the corresponding values for C_ℓ , but for this work we assume a template for the spectrum that is neutral to the underlying physics.

The spectrum which maximizes the log-likelihood provides the best estimate of the spectrum that is most consistent with our data.

The covariance matrix includes the correlations between all pairs of visibilities; it is the sum of the theoretical correlation \mathbf{M} and the measurement uncertainty \mathbf{N} : $\mathbf{C} = \mathbf{M}(\mathbf{q}) + \mathbf{N}$. The first term embodies our *a priori* understanding of the correlations between the visibilities; it is obtained from the window matrix \mathbf{W} , which describes the correlations produced by the observing strategy, and the spectrum \mathbf{q} . The spectrum represents our hypothesis about the nature of the correlations. \mathbf{N} is obtained from the uncertainties which accompany the visibilities. Despite the simplicity of this prescription, there are a number of ways in which the components of the likelihood can be assembled, and these choices affect the performance of the computational algorithm; we review these issues in the section that follows.

6.3 Implementation Considerations

We have considerable latitude in carrying out the likelihood calculation. There are many compromises between speed and precision; as one endeavors to improve the accuracy of the analysis, the computation can balloon to unmanageable proportions. The polarization likelihood analysis routine for this thesis was implemented in *Mathematica*, which sacrifices speed for transparency. A single iteration of the routine does not require a significant amount of time, but the time required to test the procedure with many simulations becomes prohibitive for all but the simplest implementations. This section discusses several assumptions which were necessary to keep the likelihood calculation tractable. The likelihood calculation requires the manipulation of N_d dimensional matrices, and since the time required for these procedures can scale as N_d^3 , we must keep the covariance matrix as small as possible. These choices affect all parts of the likelihood calculation: the data vector \mathbf{x} , the spectrum \mathbf{q} , the window matrix \mathbf{W} , and the noise matrix \mathbf{N} . We characterize the effects of these approximations with simulations in Section 6.4.

6.3.1 Visibility Averaging

The hundreds of hours of observations of the two deep fields have generated thousands of visibilities in each of the ten CBI channels. Since the dimension of the covariance matrix scales with the number of visibilities, the visibilities must be combined to minimize N_d . We averaged the LR visibilities in two passes. We first averaged all of the visibilities at each (u, v) point. The 08^h field observations were performed at equally spaced deck positions; first 20° steps (configuration 1), then 30° steps (configuration 2), so there are 18 and 12 LR visibilities, respectively, for each baseline and each channel for the 08^h observations.² The strategy for the 20^h field observations was the same as that for the configuration 2 observations on the 08^h field. Each of these points has been sampled hundreds of times during the observations, and the first average is taken over the n_j samples at each unique (u, v) point j :

$$\hat{V}_j = \frac{\sum_i^{n_j} v_i w_i}{\sum_i^{n_j} w_i} \quad (6.4)$$

The real and imaginary data are averaged separately. UVSUB computes the weight w_i for each scan from the ~ 50 integrations which constitute the scan, and after averaging all scans at each (u, v) point, the uncertainty for the visibility in Equation 6.4 is

$$\hat{\sigma}_j^2 = \frac{\sum_i^{n_j} \sigma_i^2 w_i^2}{\sum_i^{n_j} w_i^2} = \frac{1}{\sum_i^{n_j} \sigma_i^2} \quad (6.5)$$

As discussed in Chapter 5, we scale the variance in the real data by the factor of $\beta = 1.06$ to account for the bias in this estimate of the uncertainties. The simulated visibilities have uniform noise, so this correction is not applied to the simulated data.

The second average is taken over the ten CBI channels, and in this regard the LR analysis departs from that for LL . The ten channel average decreases the dimension of the covariance matrix by a factor of ten, but does so at the cost of lost information because the channels at the band edges sample different

²Antennas which formed identical baselines with RX12 given the choice of deck stepsize, such as RX3 and RX8 in configuration 2, sample the same (u, v) point and are thus averaged together.

scales.³ We might expect that this loss is most acute on the long baselines because the separation between the band edges relative to the size of the autocorrelation function of the illumination pattern increases with increasing ℓ ; in Section 6.4 we explore the effect of the ten channel average, at which point we will see that this hypothesis is correct.

These two averages permit a significant improvement in the speed of the calculation. The unaveraged data set consists of 33980 visibilities for the 08^h field and 54577 visibilities for the 20^h field. The two averages reduce these quantities to 185 complex visibilities for the 08^h field and 149 visibilities for the 20^h field—one band average for each unique (u, v) point—each of which is accompanied by a weight. Later we will truncate the covariance matrix further by binning the visibilities by baseline length.

6.3.2 Reals and Imaginaries

We have several options for how we construct the covariance matrix. Since the visibilities are complex,

$$\mathcal{V}_p = \mathcal{V}_p^R + i\mathcal{V}_p^I \quad (6.6)$$

for visibility p . We have two choices: we can compute the covariance directly from the complex visibilities, or we can consider the real and imaginary components separately. In the former case, the covariance matrix is complex and has dimension N_d , while in the latter it is real and has dimension $2N_d$. The former approach is the faster of the two, but the additional book-keeping introduced by the manipulations of complex quantities negates many of the enhancements derived from the smaller matrix. Additionally, while the real and imaginary components of the visibilities are Gaussian, the complex visibilities occasionally are not, so the complex covariance matrix, when computed from $\langle \mathcal{V}_j \mathcal{V}_k^* \rangle$, does not always fully describe the distribution of the visibilities.

The method by which we treat the real and imaginary components of the

³In addition, this average erases the spectral information about the signal.

visibilities has additional ramifications for the covariance matrix. The extent of the autocorrelation function of the illumination pattern in the aperture plane is sufficiently large that for short baselines the visibility at a point \mathbf{u}_p overlaps with those at \mathbf{u}_q and $-\mathbf{u}_q$; we must therefore consider both correlations: $\langle \mathcal{V}_p \mathcal{V}_q^* \rangle$ and $\langle \mathcal{V}_p \mathcal{V}_q \rangle$.⁴ Adopting the notation used by Hobson [32], we express the covariance matrix elements in terms of the real and imaginary components:

$$\mathbf{M}_{pq} = \langle \mathcal{V}_p \mathcal{V}_q^* \rangle = \langle (\mathcal{V}_p^R + i\mathcal{V}_p^I)(\mathcal{V}_q^R - i\mathcal{V}_q^I) \rangle \quad (6.7)$$

$$\overline{\mathbf{M}}_{pq} = \langle \mathcal{V}_p \mathcal{V}_q \rangle = \langle (\mathcal{V}_p^R + i\mathcal{V}_p^I)(\mathcal{V}_q^R + i\mathcal{V}_q^I) \rangle \quad (6.8)$$

so that,

$$\mathbf{C} = \begin{pmatrix} \langle \mathcal{V}^R \mathcal{V}^R \rangle & \langle \mathcal{V}^R \mathcal{V}^I \rangle \\ \langle \mathcal{V}^I \mathcal{V}^R \rangle & \langle \mathcal{V}^I \mathcal{V}^I \rangle \end{pmatrix} = \frac{1}{2} \begin{pmatrix} \text{Re}[\mathbf{M} + \overline{\mathbf{M}}] & -\text{Im}[\mathbf{M} - \overline{\mathbf{M}}] \\ \text{Im}[\mathbf{M} + \overline{\mathbf{M}}] & \text{Re}[\mathbf{M} - \overline{\mathbf{M}}] \end{pmatrix} + \mathbf{N} \quad (6.9)$$

which has dimension $2N_d$. In this approach, the data vector is a real vector of length $2N_d$: $\mathbf{x} = (\mathcal{V}^R, \mathcal{V}^I)$.

6.3.3 Visibility Window Function

Our instrument and observing strategy result in a high degree of correlation between the visibilities, and the likelihood routine must take these correlations into account. This covariance is quantified by the window matrix of Chapter 2. The observing strategy determines the shape of the window matrix: the point \mathbf{u}_p determines the sampling in the aperture domain, and the autocorrelation function of the illumination pattern spreads the response at \mathbf{u}_p across an area of size σ_p in the aperture domain. If the visibilities are more than σ_p apart, for example, the covariance matrix will be very nearly diagonal. For the CBI observing strategy σ_p is sufficiently large to encompass neighboring visibilities, and this gives rise to a high degree of correlation which is most pronounced on the short baselines. Since

⁴Beyond visibility separations which are greater than σ_p the conjugate correlations are negligible.

σ_p affects the coupling between the antennas, it is an integral component of the window matrix. The window matrix elements are given by

$$\mathbf{W}_{pq} = \kappa_p \kappa_q \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) \tilde{A}_q^*(\mathbf{u}_q - \mathbf{v}) d^2 \mathbf{v} \quad (6.10)$$

$$\overline{\mathbf{W}}_{pq} = \kappa_p \kappa_q \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) \tilde{A}_q(\mathbf{u}_q + \mathbf{v}) d^2 \mathbf{v} \quad (6.11)$$

where $\tilde{\mathbf{A}}_p(\mathbf{u}_p - \mathbf{v})$ is the autocorrelation of the illumination pattern at the point \mathbf{u}_p ; the subscript p on $\tilde{\mathbf{A}}$ reminds us that σ_p is a function of the frequency of the sample at \mathbf{u}_p .

The shape of the autocorrelation of the illumination pattern is an important component of the window function calculation. The CBI's primary beam is nearly Gaussian; the feed illuminates the primary reflector with a Gaussian pattern, but the secondary reflector obscures the central part of the illumination pattern, and the edges of the primary reflector truncate the wings of pattern. We performed beam pattern measurements on 3C274 for the total intensity analysis; Figure 6.1 compares the best fit beam pattern with a Gaussian approximation. These measurements focused on the amplitude of the beam; we assume that the beam does not impart spurious changes to the phase, and the lack of significant change in phase of the leakage measured at the half-power points (Section 3.3.1) supports this conclusion.

The analysis presented in this work assumes that the beam is Gaussian. The autocorrelation function of the beam centered at a point \mathbf{u}_p has the form

$$\begin{aligned} \tilde{A}_p(\mathbf{u}_p - \mathbf{v}) &= \frac{1}{2\pi\sigma_p^2} \exp\left[\frac{-1}{2\sigma_p^2}(\mathbf{u}_p - \mathbf{v})^2\right] \\ &= \frac{1}{2\pi\sigma_p^2} \exp\left[\frac{-1}{2\sigma_p^2}(u_p^2 + v^2 - 2|u_p||v|\cos[\theta_p - \theta_v])\right] \end{aligned} \quad (6.12)$$

and that at the conjugate point is similar,

$$\tilde{A}_p(\mathbf{u}_p + \mathbf{v}) = \frac{1}{2\pi\sigma_p^2} \exp\left[\frac{-1}{2\sigma_p^2}(\mathbf{u}_p + \mathbf{v})^2\right]$$

CBI Primary Beam Comparison




Figure 6.1: CBI primary beam comparison. The best fit curve is obtained from beam map observations of 3C274. The beammap data are best fit by a Gaussian illumination pattern between radii of 7.5 cm and 45 cm which is then inverted to obtain the beam. The Gaussian approximation has a FWHM of $45.2'$ ($31 \text{ GHz}/\nu$); this approximation is chosen to have a FWHM which matches that of the best fit to the real data. By definition, the beams have unit response at the boresight. POLFAKE can generate simulated skies with both beams seen in the figure.

$$= \frac{1}{2\pi\sigma_p^2} \exp\left[\frac{-1}{2\sigma_p^2} \left(u_p^2 + v^2 + 2|u_p||v|\cos[\theta_p - \theta_v]\right)\right] \quad (6.13)$$

The sign change follows from mapping of $\mathbf{u}_p \rightarrow -\mathbf{u}_p$. The denominator in the argument of the Gaussian sets the size of the autocorrelation of the beam:

$$\sigma_p^2 = \frac{2\ln[2]}{\pi^2 a_p^2} \quad (6.14)$$

The factor of a_p is the FWHM of the beam on the sky, which we obtain from the best fit to the beam measurements: $a_p = 45.2'(31 \text{ GHz}/\nu)$. The standard practice is to integrate the window function over θ_v to obtain the azimuthally averaged function $W_{pq}(v)$; this form provides a direct measure of the response in ℓ space under the flat sky approximation. In Section 6.4 we use simulations to explore the

validity of the Gaussian approximation.

Figure 6.2 shows the diagonal visibility window functions $W_{pp}(v)$ for both fields. The autocorrelation of the illumination pattern sets the widths of the functions in ℓ space: $\Delta\ell \sim 420$ (FWHM). The most striking feature of this figure is the decrease in the window functions with increasing ℓ . Loosely speaking, a baseline of length $|v|$ samples scales $\theta_{\ell=2\pi v}$ under the flat sky approximation, so a baseline which corresponds to a length ℓ fills the primary beam with $\sim \Omega_p/\theta_\ell^2$ independent estimates of C_ℓ . The rms of the fluctuations within the primary beam decreases with the number of fluctuations, which in turn scales as $\sqrt{\Omega_p/\theta_\ell^2} \sim \ell$, so the window functions, which measure the response to fluctuations of fixed brightness, decrease as $1/\ell$.

6.3.4 Noise Covariance Matrix

We obtain the noise from the visibility weights, which are determined from the scatter in the visibilities during each scan. As we saw, the noise in the LL visibilities can be $\sim 20\%$ higher than the values predicted, and the LR visibilities are $\sim 5-10\%$ higher still due to the lower sensitivity in our RCP receiver. We assume that the noise between different visibilities is uncorrelated, so the noise covariance matrix is diagonal.⁵

Section 6.3.2 discusses two approaches for constructing the covariance matrix given that the visibilities are complex; in parallel with these choices, there are two options for the noise covariance matrix. If we treat the elements of the data vector as complex quantities, the corresponding noise matrix has the form

$$\mathbf{N}_{pq} = \delta_{pq}(\sigma_p^R \sigma_q^R + \sigma_p^I \sigma_q^I) = 2\delta_{pq} \sigma_p^R \sigma_q^R \quad (6.15)$$

The cross terms between the reals and imaginaries are negligible because the quadrature calibration ensures that the two components are orthogonal at the

⁵Coupling between the antennas, for example, would produce off-diagonal matrix elements in the noise matrix, but there is no evidence for a false coupled signal in the data.

1% level; the quadrature calibration also ensures the real and imaginary uncertainties are equal. Alternatively, if we separate the reals and imaginaries to form a covariance matrix of size $2N_d$, we separate the reals and imaginaries in Equation 6.15 to form identical diagonal submatrices for the $\langle \mathcal{V}^R \mathcal{V}^R \rangle$ and the $\langle \mathcal{V}^I \mathcal{V}^I \rangle$ quadrants of the covariance matrix. The terms of the noise covariance matrix are given by

$$\mathbf{N}_{pq} = \delta_{pq} \sigma_p \sigma_q \quad (6.16)$$

This is the expression which we apply to the covariance calculation.

6.3.5 Input Power Spectrum

The source spectrum has more degrees of freedom than any other component of $\mathcal{L}(\mathbf{x}|\mathbf{q})$. CMBFAST generates values for C_ℓ^{EE} and C_ℓ^{BB} over our region of interest in ℓ , and we are free to insert these values into Equation 6.2 after the appropriate flat sky mapping to v . This approach adds a significant degree of complexity to the likelihood calculation, however, because it introduces a large space of trial parameters to be explored; we can incorporate physical considerations in the analysis, for example, by computing C_ℓ over a grid of cosmological parameters and evaluating the likelihood of each model. In this case we would also weight the likelihoods by the priors for the model parameters. This approach introduces a bias into the analysis, however, because it imposes a matched filter on the data, and it precludes transparent comparisons with other experiments, so we do not apply this approach for this work.

To simplify the problem we adopt the common convention and apply a trial spectrum that is constant across specific bands in ℓ . Under this assumption of *flat band power* we introduce a spectrum which is piecewise flat ($q_i \sim \ell(\ell+1)C_\ell \sim \text{constant}$) for each of the N_b bands [9]. This template also permits direct comparisons with flat bandpowers from other experiments.⁶ The trial amplitudes $q_i^{\alpha\alpha}$ for each

⁶The DASI polarization effort demonstrated the virtues of considering a shaped spectrum; they improved the significance of their detection by applying a spectrum whose shape was determined by the cosmology obtained from a concordance fit to the total intensity data [44]. Shaped trial

band i are related to the brightness fluctuations by

$$q_i^{\alpha\alpha} = \left(\frac{\delta T}{T}\right)^2 = \frac{\ell(\ell+1)}{2\pi} C_\ell^{\alpha\alpha} = 2\pi v^2 C^{\alpha\alpha}(v) \Rightarrow C^{\alpha\alpha}(v) = \frac{1}{2\pi v^2} q_i^{\alpha\alpha} \quad (6.17)$$

where $\alpha\alpha = TT, EE$, or BB , and we have used the flat sky approximation to relate ℓ to v : $\ell = 2\pi v$ for large ℓ . We insert $C^{\alpha\alpha}(v)$ into the theory covariance matrix. The spectrum is often expressed as \mathcal{C}_ℓ , which is related to C_ℓ by

$$\mathcal{C}_\ell = \frac{\ell(\ell+1)}{2\pi} C_\ell \quad (6.18)$$

This notation is convenient because \mathcal{C}_ℓ is constant for flat band power.

B Modes

The full description of the polarization of the CMBR requires both $E(\mathbf{x})$ and $B(\mathbf{x})$.

Recall that

$$\mathbf{M}_{pq} = \langle \mathcal{V}_p^{LR} \mathcal{V}_q^{LR*} \rangle \sim (\tilde{E} - i\tilde{B})(\tilde{E}^* + i\tilde{B}^*) \sim (C_\ell^{EE} + C_\ell^{BB}) \quad (6.19)$$

$$\overline{\mathbf{M}}_{pq} = \langle \mathcal{V}_p^{LR} \mathcal{V}_q^{LR} \rangle \sim (\tilde{E} - i\tilde{B})(\tilde{E} - i\tilde{B}) \sim (C_\ell^{EE} - C_\ell^{BB} - 2iC_\ell^{EB}) \quad (6.20)$$

Symmetry requires that $C_\ell^{EB} = 0$ [ref]. In the presence of E and B modes, the theoretical component of the covariance has the form

$$\mathbf{M} = \frac{1}{2} \begin{pmatrix} \text{Re}[\Sigma_C \mathbf{W} + \Delta_C \overline{\mathbf{W}}] & -\text{Im}[\Sigma_C \mathbf{W} - \Delta_C \overline{\mathbf{W}}] \\ \text{Im}[\Sigma_C \mathbf{W} + \Delta_C \overline{\mathbf{W}}] & \text{Re}[\Sigma_C \mathbf{W} - \Delta_C \overline{\mathbf{W}}] \end{pmatrix} \quad (6.21)$$

with

$$\Sigma_C = C^{EE}(v) + C^{BB}(v); \quad \Delta_C = C^{EE}(v) - C^{BB}(v) \quad (6.22)$$

This is the expression for the covariance that we insert in the likelihood.

spectra have great importance for polarization in particular because the peak to peak variations in C_ℓ^{EE} and C_ℓ^{TE} are greater than those for C_ℓ^{TT} ; C_ℓ^{TE} , in fact, is expected to have multiple zero crossings on $\ell = 200 \rightarrow 2000$ scales, and flat band power would artificially suppress this power.

In the analysis that follows we will assume that $C_\ell^{BB} = 0$. Given the absence of direct measurements of C_ℓ^{BB} , this assumption is not ideal, but it is forced upon us by the computational limitations, and it does have some observational support. Increasingly precise observations of the total intensity of the CMBR are providing compelling evidence for families of models which predict $C_\ell^{EE} \gg C_\ell^{BB}$ on CBI scales. In addition, none of the existing polarization observations have detected the presence of B modes [89].

Trial Spectra

While the assumption of flat band power permits vast simplifications in the computation of the likelihood, it still provides discretion in how the problem is cast. In particular, we must decide how to sort the data in ℓ ; we must still choose band centers and band widths, and these choices affect the outcome. Two competing factors affect the band assignments. We could define a single band of sufficient width ($\ell \sim 3000$) to include all of the baselines and thus put all of the visibilities in the service of a single limit. The loss of information in this approach, however, outweighs the benefits; since the number of independent fluctuations in the (u, v) plane increases with increasing $|v|$, or equivalently, for increasing ℓ , the S/N per fluctuation falls like $1/\ell$ —sufficiently rapidly that the likelihood soon becomes dominated by noise. This consideration favors narrow bins. And the CBI introduces an innate scale size; the CBI’s natural sampling scale is the size of the autocorrelation function of the illumination pattern which, when measured in ℓ , has FWHM $\Delta\ell \sim 560$.

We adopted a simple approach to binning the data in ℓ . With this approach the (u, v) data are sorted into bands based on $|v|$, and the likelihoods for the bands are evaluated separately without regard for the correlations between the neighboring bands. This method is equivalent to sorting the LR baselines by length and treating each set as an isolated experiment, and while this approach is computationally desirable, it has a serious shortcoming. The tails of the window functions for a particular band extend into the neighboring bands, and the resulting correlation

must be taken into account with a joint fit to all bands. The latter approach is the more accurate of the two, but the attendant increase in computational expense is considerable. We discuss both approaches in this section.

Independent Bands

The independent binning scheme adopted for the polarization limit follows from the geometry of the array configurations. The minimum antenna separation of 100 cm provides a natural scale for the band demarcations because the antenna separations for the LR baselines tend to be grouped in multiples of this length. We defined three bins centered on $|b| \sim 100, 200, \text{ and } 300$ cm, which correspond to $\ell \sim 630, 1260, \text{ and } 1890$. When the likelihood is evaluated, the limits of integration—the bounds on $|v|$ —are taken to be sufficiently large to encompass all of the power in the window function, so in this case the band boundaries serve only to sort the baselines which sample the scales of interest for the bin.

Table 6.1 lists the bins used for the analysis of the 08^h field, and Table 6.2 lists the bins used for the analysis of the 20^h field. Figure 6.2 shows the diagonal elements of the visibility window functions for this choice of bands. We use the same bands for both fields to enable a joint fit to both fields after the fields have been fit separately. These bins are highly correlated, and in the absence of a correction for the correlation, it is not valid to report all three limits from the bins together; with this approach, we can state that the data support a limit for band 1 or band 2, for example, but not both.⁷ Figures 6.3 and 6.4 show the (u, v) coverage for the two deep fields; the heavy circles in the figures denote the demarcations between the three bands.

Correlated Bands

The remedy for the high correlations described above is to expand the covariance matrix to include the correlations. We construct N_b window matrices, each of

⁷Bins 1 and 3 are sufficiently far apart that we can report limits for those two bins together without fear of making a serious error.

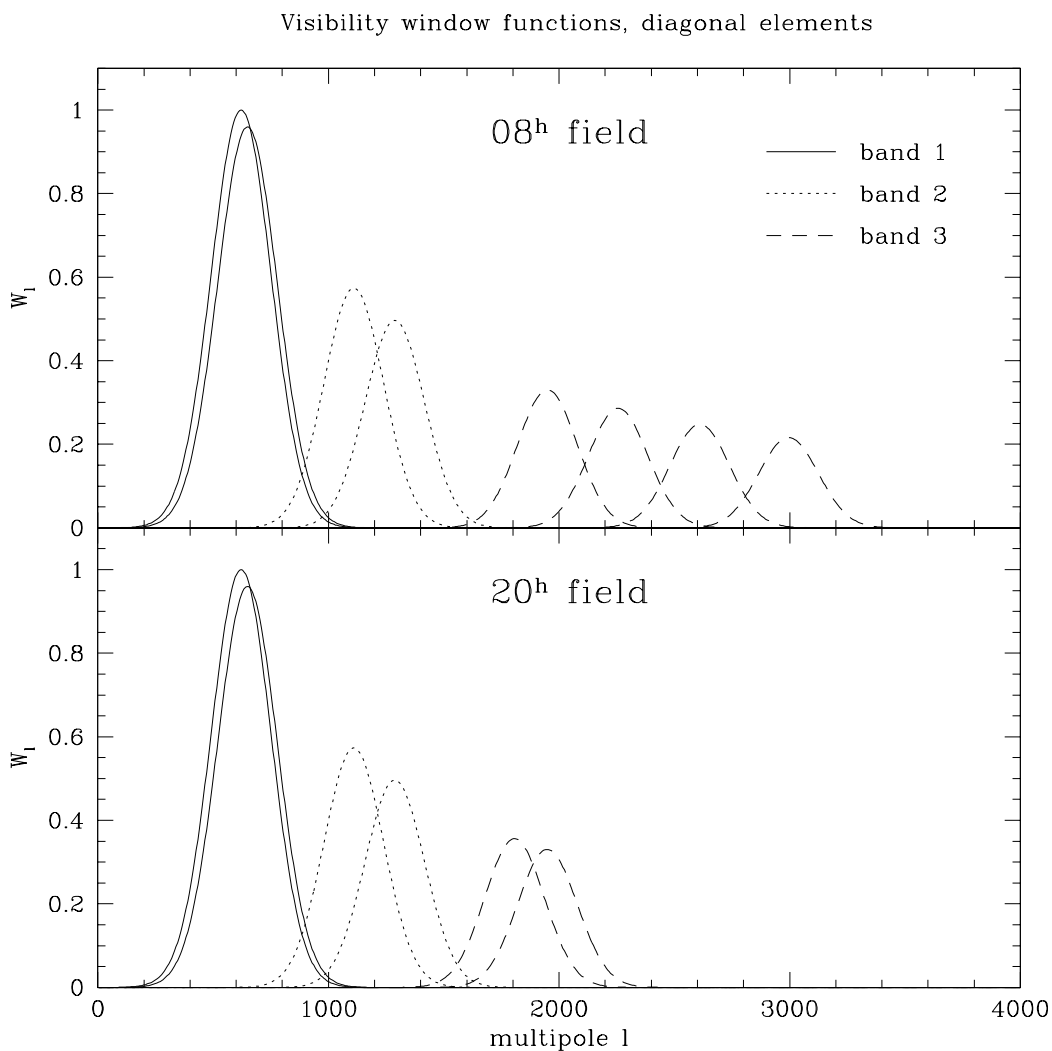


Figure 6.2: Visibility window functions, diagonal elements. This figure shows the band scheme adopted for this work; the baselines were sorted into three bands based on the natural demarcations between baseline lengths. The figure hints at the degree of correlation between the fields; bands 1 and 2 are clearly strongly correlated for both fields, while for the 08^h field, band 3 is only marginally correlated with band 2. The real correlations are higher because the 30% bandwidth spreads the functions in ℓ space.

$0 < \ell < 945$			$945 < \ell < 1575$			$1575 < \ell < 2205$		
RX	b (cm)	τ (h)	RX	b (cm)	τ (h)	RX	b (cm)	τ (h)
0	100.0	55.3	2	173.2	53.2	11	300.0	24.8
10	104.1	54.1	—	—	—	3	400.0	38.8
—	—	—	—	—	—	6	458.3	38.8
8	100.0	12.0	1	173.2	11.9	0	300.0	11.9
3	100.0	12.0	6	173.2	11.9	9	346.4	11.1
—	—	—	7	173.2	11.9	—	—	—
—	—	—	4	200.0	10.7	—	—	—
133.4			99.6			125.4		

Table 6.1: Bands for the likelihood analysis, 08^h field. The RX column reports the receiver with which RX12 forms the baseline, and the b column notes the length of the baseline. The three vertical groupings are the bands, and the two horizontal groupings correspond to configuration 1 and configuration 2. The bottom row tallies the total integration time in hours for each band.

which has dimension $2N_d$ and for which the window matrix elements are evaluated over the range applicable for the bin, i.e., the ranges in v (or equivalently, in ℓ) shown in the top row of Tables 6.1 and 6.2. We then scale each matrix by a trial spectrum q_i :

$$\mathbf{C} = \sum_i^{N_b} q_i \mathbf{W}_i + \mathbf{N} \quad (6.23)$$

and search for the likelihood peak within the N_b dimensional space of possible spectra (q_1, q_2, \dots, q_i) . With this approach, \mathbf{C} is largely block-diagonal; most of the power resides in the blocks which correspond to baselines of similar length. These blocks resemble the covariance matrices for the independent bands, and they are surrounded by off-diagonal elements which represent the correlations between the bands. While this approach is more accurate than the use of independent bins, it comes at a considerable computational expense; the additional correlations increase

$0 < \ell < 945$			$945 < \ell < 1575$			$1575 < \ell < 2205$		
RX	b (cm)	τ (h)	RX	b (cm)	τ (h)	RX	b (cm)	τ (h)
3	100.0	77.0	6	173.4	61.6	5	278.4	77.0
8	100.0	77.0	7	173.4	61.6	0	300.0	62.6
1	104.1	77.0	4	200.0	77.0	—		
231.0			200.2			139.6		

Table 6.2: Bands for likelihood analysis, 20^h field. All data were taken with the array in configuration 3.

the dimension of the covariance matrix to $2N_d$, where N_d is the size of the entire data set, and we must now search an N_b dimensional space—rather than a one dimensional space—for the best fit values for the components of \mathbf{q} . Because of the prohibitive CPU requirements, we did not explore this avenue for this work.

6.3.6 Joint Fit

We observed two deep fields, and the combined data offer our best hope of a detection, so we performed a joint fit to the two fields. For two nonoverlapping fields, the covariance takes the form

$$\mathbf{C}_T = \begin{pmatrix} \mathbf{M}_1 & 0 \\ 0 & \mathbf{M}_2 \end{pmatrix} + \begin{pmatrix} \mathbf{N}_1 & 0 \\ 0 & \mathbf{N}_2 \end{pmatrix} \quad (6.24)$$

where the data vector $\mathbf{x} = (\mathbf{x}_1, \mathbf{x}_2)$. The off-diagonal submatrices vanish because the fields are uncorrelated, so rather than contend with a block-diagonal covariance matrix of size $2N_{d_1} + 2N_{d_2}$ we simply evaluate $\ln[\mathcal{L}_i] \sim -\ln[|\mathbf{C}_i|] - \mathbf{x}_i^t \mathbf{C}_i^{-1} \mathbf{x}_i$ for each field i and sum the log-likelihoods to find the best fit bandpower for the two fields.




Figure 6.3: (u, v) coverage at the band center, 08^h field. Heavy circles show the regions of the (u, v) plane defined by the three bands shown in Table 6.1. Light circles show the extent of the autocorrelation function of the illumination pattern, which has width $\Delta_p = 2\sqrt{2\ln[2]}\sigma_p$ with $\sigma_p = 28.6$ (ν/GHz) at 1 cm. The two sets of circled points represent baselines whose orientations differ by 60° (dashed line). The figure shows that for fixed baseline separation, the short baseline visibilities are more correlated than the long baseline visibilities.

6.3.7 RX10

During the configuration 1 observations of the 08^h field there were only two baselines which sample the ($\ell \sim 600$) scales at which the polarization predicted by standard models peaks: RX0-RX12 and RX10-RX12. Of these, only RX0-RX12 could be calibrated because RX10-RX12 lacked an LL counterpart under our routine observing strategy.⁸ We can approximate the correct calibration for RX10-RX12, however, because with $\theta = -16.1^\circ$ and $|b| = 104.1$ cm,⁹ it is very close to RX1-RX6 after the latter has undergone a 20° rotation; at that point the two

⁸RX10-RX12 is parallel to RX7-RX8 after a deck rotation of 87.8° , for example, but irregular deck rotations such as this were not part of the routine observations during which we stepped the deck angle in increments of 20° .

⁹For a coordinate system for which the orientation of baseline RX0-RX12 is 0° .




Figure 6.4: (u, v) coverage at the band center, 20^h field. Heavy circles show the regions of the (u, v) plane defined by the three bands shown in Table 6.2.

baselines differ in length by 4% and angle by 4° . By displacing RX10's position by a distance of ~ 4 cm and an angle of $\sim 4^\circ$ relative to its original position, we place it at a point at which it can be calibrated. We do not, of course, perform this change physically on the array, but we can reposition the antenna in CBICAL before calibration in CBIPOLCAL, and use simulations to estimate the error introduced in the best fit bandpowers by this approximation. This position offset introduces a systematic $\delta\theta \sim 8^\circ$ error in the phase calibration for this baseline. 3C279 is unresolved by the CBI, so this change does not corrupt the amplitude calibration. Since CBIPOLCAL applies baseline-based calibrations, RX12-RX10 is the only baseline which is affected by this change. We explore the effect of this change on the best fit likelihoods with simulations in Section 6.4.

6.3.8 Explicit Calculation

We now combine all of the assumptions of this section to obtain a concise expression for the covariance matrix. The primary anisotropy contribution to the covariance

matrix is given by Equation 6.21, and the window functions \mathbf{W}_{pq} and $\overline{\mathbf{W}}_{pq}$ are given by 6.10 and 6.11, respectively.

Our template for the underlying spectrum permits several simplifications in the likelihood calculation. The assumption that $C_\ell^{BB} \equiv 0$ allows us to write $\boldsymbol{\Sigma}_C = \boldsymbol{\Delta}_C = C^{EE}(v)$. We also assume a spectrum which is piecewise flat in ℓ , so $C^{EE}(v)$ has the form

$$C^{EE}(v) = \frac{1}{2\pi} \frac{1}{v^2} \left(\frac{\delta T}{T} \right)^2 \quad (6.25)$$

and we step the trial spectrum across a range of values in units of $(\delta T/T)^2$. This term is the sole degree of freedom in the maximization routine. We simplify the calculation further by sorting the baselines into three bins as function of $|v|$, or equivalently ℓ , and evaluating the covariance separately for each band.

The assumption of a Gaussian primary beam simplifies the covariance calculation because it allows us to express the window matrix elements in terms of Bessel functions. Upon inserting the flat spectra into the covariance, we obtain

$$q_i \mathbf{W}_{pq} = \frac{1}{2\pi} \left(\frac{\delta T}{T} \right)^2 \int_0^\infty \frac{W_{pq}(v)}{v} dv \quad (6.26)$$

and

$$q_i \overline{\mathbf{W}}_{pq} = \frac{1}{2\pi} \left(\frac{\delta T}{T} \right)^2 \int_0^\infty \frac{\overline{W}_{pq}(v)}{v} dv \quad (6.27)$$

for each band i . After multiplying by a factor of two to account for differencing, integration over θ_v in 6.10 yields

$$\begin{aligned} q_i \mathbf{W}_{pq} &= \frac{1}{2\pi} \left(\frac{\delta T}{T} \right)^2 \frac{\kappa_p \kappa_q}{\pi \sigma_p^2 \sigma_q^2} e^{2i(\theta_p - \theta_q)} \\ &\times \int_0^\infty \exp \left[-\frac{u_p^2 + v^2}{2\sigma_p^2} - \frac{u_q^2 + v^2}{2\sigma_q^2} \right] I_0[vC_{pq}] \frac{dv}{v} \end{aligned} \quad (6.28)$$

where $I_0[vC_{pq}]$ is a modified Bessel function of the zeroth kind:

$$I_0[vC_{pq}] = \frac{1}{2\pi} \int_0^{2\pi} \exp[vC_{pq} \cos[\theta]] d\theta \quad (6.29)$$

and the constant C_{pq} —not to be confused with the spectrum $C^{EE}(v)$ —is given by

$$C_{pq}^2 = \frac{u_p^2}{\sigma_p^2} + \frac{u_q^2}{\sigma_q^2} + 2 \frac{|u_p|}{\sigma_p} \frac{|u_q|}{\sigma_q} \cos[\theta_p - \theta_q] \quad (6.30)$$

and the factor of κ_p converts from intensity to brightness temperature:

$$\kappa_p = \frac{2\nu_p^2 k_B T_0}{c^2} \frac{x^2 e^x}{(e^x - 1)^2} \quad (6.31)$$

with $x = h\nu_p/k_B T_0$. The covariance matrix elements for the conjugate points resemble those for $q_i \mathbf{W}$:

$$\begin{aligned} q_i \overline{\mathbf{W}}_{pq} &= \frac{1}{2\pi} \left(\frac{\delta T}{T} \right)^2 \frac{\kappa_p \kappa_q}{\pi \sigma_p^2 \sigma_q^2} e^{2i(\theta_p + \theta_q)} e^{-4i\Theta_{pq}} \\ &\times \int_0^\infty \exp \left[-\frac{u_p^2 + v^2}{2\sigma_p^2} - \frac{u_q^2 + v^2}{2\sigma_q^2} \right] I_4[v \overline{C}_{pq}] \frac{dv}{v} \end{aligned} \quad (6.32)$$

and the constant \overline{C}_{pq} is identical to C_{pq} , save a single sign change:

$$\overline{C}_{pq}^2 = \frac{u_p^2}{\sigma_p^2} + \frac{u_q^2}{\sigma_q^2} - 2 \frac{|u_p|}{\sigma_p} \frac{|u_q|}{\sigma_q} \cos[\theta_p - \theta_q] \quad (6.33)$$

The change in sign comes from mapping $\Delta\theta \rightarrow \Delta\theta + \pi$ for the conjugate point. $\overline{\mathbf{W}}_{pq}$ also introduces Θ_{pq} , a constant phase:

$$\Theta_{pq} = \tan^{-1} \left[\frac{|u_p| \sin(\theta_p) - |u_q| \sin(\theta_q)}{|u_p| \cos(\theta_p) - |u_q| \cos(\theta_q)} \right] \quad (6.34)$$

One must take care in evaluating the expression for Θ because some diagonal terms are singular.

The preceding discussion presents all of the factors which enter the covariance calculation. We now make a few final assumptions. CBIPOLCAL applies the parallactic angle correction, so we let $\theta_j = \theta_k = 0$. We also combine the ten CBI channels to form a single band centered at $\nu = 31$ GHz; under this assumption, $\sigma_p = \sigma_q = \sigma$, and $\kappa_p = \kappa_q = \kappa$, both of which are evaluated at 31 GHz. Thus we

compute the covariance as,

$$\mathbf{C} = \frac{q_i}{2} \begin{pmatrix} \text{Re}[\mathbf{W} + \overline{\mathbf{W}}] & \text{Im}[\mathbf{W} + \overline{\mathbf{W}}] \\ -\text{Im}[\mathbf{W} - \overline{\mathbf{W}}] & \text{Re}[\mathbf{W} - \overline{\mathbf{W}}] \end{pmatrix} + \frac{1}{2} \begin{pmatrix} \mathbf{N} & 0 \\ 0 & \mathbf{N} \end{pmatrix} \quad (6.35)$$

with $q_i \mathbf{W}$ and $q_i \overline{\mathbf{W}}$ given by Equations 6.32 and 6.33, and \mathbf{N} given by Equation 6.16.

6.4 Simulations

Simulated data are a critical tool for assessing the performance of the likelihood routine, and they provide insights into effects such as calibration errors which cannot be understood analytically. The polarization simulation pipeline consists of a sequence of programs: POLSKY, POLFAKE [65], and CBIPOLCAL. The user supplies spectra (C_ℓ^{TT} , C_ℓ^{EE} , C_ℓ^{BB} , C_ℓ^{TE}) to POLSKY, which generates maps of $I(\mathbf{x})$, $Q(\mathbf{x})$, and $U(\mathbf{x})$. The user then feeds the maps, an observing strategy, a primary beam, and an array configuration to POLFAKE, which simulates observations of these artificial skies to produce LL , LR , RL , and RR visibilities. POLFAKE cannot introduce instrumental polarization, but we circumvent this shortcoming downstream with CBIPOLCAL.

The simulation pipeline was designed to recreate the observations of the 08^h and 20^h CBI deep fields. To simplify the interpretation of the simulations, flat spectra were assumed for C_ℓ^{TT} ($\sqrt{C_\ell^{TT}} = 60 \mu\text{K}$) and C_ℓ^{EE} ($\sqrt{C_\ell^{EE}} = 6 \mu\text{K}$), while the other two spectra were set to zero. Two sets of maps were simulated for each iteration. The first field was passed to POLFAKE twice to generate visibilities for the configuration 1 and configuration 2 observations of the 08^h field, while the second field was passed to POLFAKE a single time to generate a set of visibilities for the configuration 3 observations of the 20^h field. The deck rotation increments were chosen to recreate the strategy used for the real observations; they were set to 20° , 30° , and 30° , respectively, for the three configurations. Although POLFAKE generates calibrated visibilities, the visibilities were passed to CBIPOLCAL for a

null calibration. This step was included for consistency, as later simulations were recalibrated in CBIPOLCAL to explore the effects of calibration errors on the best fit bandpowers. The configuration 1 and configuration 2 visibilities for the 08^h field were combined with UVSUB, and the data for the two fields were written to text files for the likelihood analysis in *Mathematica*.

The simulations depart from the real data in a number of significant ways. First, the real data span the entire 26-36 GHz band, but to expedite the production of simulated visibilities many of the simulations focused on a central 1 GHz band. Section 6.4.2 discusses this aspect of the simulations, and at that point we explore the question of frequency coverage in greater detail. Second, we compute the theory covariance under the assumption of a Gaussian primary beam, but we have seen that the Gaussian approximation disagrees with the real beam at the few percent level. We investigate the primary beam approximation in Section 6.4.3. Finally, in Section 6.4.4 we introduce calibration errors into the simulated data, but we do not approximate the cumulative effect of the $\sim 10^2$ separate calibrations for the deep field observations; we only explore the effects of systematic calibration errors on the best fit bandpowers.

6.4.1 Baseline Likelihood Tests

The first test consisted of a set of simulations to establish that the likelihood routine behaves as desired under ideal circumstances. For these tests we simulated 10^3 sets of data, each of which represents 5000 nights of integration, for a single 31-32 GHz channel. The large number of simulations was necessary to suppress the standard error on the mean of the distribution of peak likelihoods to levels well below one μK , and the long integration time ensures that the signal remains well above the noise for all baselines. These data were calibrated in CBIPOLCAL with a set of null calibration factors: unity gain ($G_r \equiv 1, G_i \equiv 0$) and zero leakage. Although the CBI beam is not a perfect Gaussian, the simulations and the likelihood routine both used the same beam—a Gaussian with FWHM $45.2'(31.0 \text{ GHz}/\nu)$ —so the two procedures are internally consistent. If the routine performs as desired, the

		08 ^h field (μK)		20 ^h field (μK)		joint fit (μK)	
\mathcal{C}_ℓ^{EE}	band	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}
6 μK	1	5.99	0.03	5.95	0.03	5.97	0.02
	2	5.99	0.03	6.02	0.03	5.97	0.02
	3	5.98	0.04	6.03	0.05	6.00	0.03
9 μK	1	9.00	0.04	9.00	0.05	8.99	0.03
	2	8.93	0.04	8.97	0.04	8.94	0.03
	3	8.92	0.04	8.97	0.04	8.94	0.03
12 μK	1	11.95	0.05	11.93	0.06	11.93	0.04
	2	11.89	0.04	11.97	0.05	11.92	0.03
	3	11.87	0.05	11.97	0.06	11.91	0.04

Table 6.3: Table of baseline simulation results for null calibration (unity gain, zero leakage) for a variety of input spectra. In all cases we set $\mathcal{C}_\ell^{BB} = \mathcal{C}_\ell^{TE} \equiv 0$. The three different choices for the spectra demonstrate the linearity of the routine.

mean of the 10^3 best fit amplitudes should equal the input amplitudes for all three of the bins listed in Tables 6.1 and 6.2.

The first group of rows in Table 6.3 reports the results of the baseline tests of the likelihood routine for a single channel centered on 31.5 GHz and an input spectrum of $\mathcal{C}_\ell^{EE} = 6 \mu\text{K}$. The uncertainties listed in the table were obtained from the variance of the distribution of best fit bandpowers; the table reports the amplitude of the mean of the best fit bandpowers ($= \sqrt{\langle q_i \rangle}$), and the uncertainties are the standard errors on the mean. The routine recovers the input spectrum to within the uncertainties; in all cases the amplitudes are within 1% of the input values for both the single fields and the joint fit. Figure 6.5 shows the distribution of the best fit amplitudes for band 1.

The preceding tests established that the likelihood routine returns the input spectrum to a fairly high degree of accuracy. We augmented this initial test with




Figure 6.5: Histogram of 10^3 best fit likelihoods for baseline tests of the joint analysis pipeline.

several tests on data reflecting progressively larger spectra— $C_{\ell}^{EE} = 9 \mu\text{K}$ and $C_{\ell}^{EE} = 12 \mu\text{K}$ —to test the linearity of the routine. The second two rows of Table 6.3 show the results of these tests; the output amplitudes fall short of the input spectra by $\sim 1\%$ in the worst cases, but these discrepancies are not significant. These tests provide great confidence in the accuracy of the routine.

For the real data, the deep field visibilities for the lead and trail fields are differenced to excise ground spillover. For very noisy data, as we have here, visibility subtraction increases the rms of the visibilities by $\sqrt{2}$, so for the likelihood analysis of the real data we scale the window matrix elements by a factor of two over those which are required for the undifferenced data, but this is a trivial change and we carried out our tests on single pointings. To halve the time required to generate the simulated data, however, most of the data presented in this chapter consisted of single (undifferenced) pointings. After performing the likelihood tests discussed above, a set of 10^3 lead and trail fields were generated to assess the effect of the field subtraction on the best fit likelihoods; these data were differenced in UVSUB and passed to the likelihood routine which was modified to reflect the additional

scatter. These simulations demonstrated that the factor of two fully accounts for the additional scatter introduced by the visibility subtraction in UVSUB.

6.4.2 Ten Channel Average

We average the ten CBI channels at each (u, v) point to expedite the likelihood calculation. This assumption reduces the dimension of the covariance matrix by a factor of ten, but the attendant increase in speed comes at the cost of lost information because the 30% bandwidth for each visibility subsumes fluctuations on different scales. In this section we use simulations to quantify the effects of the ten channel average in the real data. For the channel average test, each set of simulated maps was observed twice; once with all ten CBI channels, and then again with only a single central channel. As with the real data, the ten channels of simulated data for the former case were averaged at each (u, v) point. The integration times for the two sets of observations were adjusted to force the visibilities to have identical S/N; they differ by a factor of 10.

Table 6.4 provides the results of this test; the ratios in the bottom row of the table show that while the error introduced by the band average is marginal ($\sim 4\%$) on the short baselines, it rapidly becomes significant for increasing baseline length. This effect can be understood in terms of the window functions; Figure 6.6 compares the diagonal elements of the visibility window functions for the shortest and longest baselines in the 08^h field observations. The figure shows that while the visibilities at the band edges for the short baselines sample very similar scales and thus nearly identical fluctuations, those for the longest baselines sample very different scales, and the additional incoherence introduced in the latter case suppresses the measured amplitude by $\sim 20\%$. The magnitude of this effect is evident in a comparison of the long baseline amplitudes for the 20^h and the 08^h data (row 6, columns 1 and 3 of Table 6.4). The longest baseline during the 20^h observations is 300 cm, while that for the 08^h observations is 458.3 cm; as a consequence, the band average over the high ℓ band suppresses more power for the 08^h field than for the 20^h field.

		08 ^h field (μK)		20 ^h field (μK)		joint fit (μK)	
test	band	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}
$\Delta\nu = 1$ GHz	1	6.04	0.03	5.99	0.03	6.01	0.02
	2	5.94	0.02	6.04	0.03	5.99	0.02
	3	5.93	0.03	6.01	0.04	5.94	0.02
$\Delta\nu = 10$ GHz	1	5.81	0.03	5.79	0.03	5.78	0.02
	2	5.25	0.02	5.31	0.02	5.27	0.01
	3	4.54	0.02	4.97	0.02	4.72	0.02
ratio	1	0.962	0.007	0.967	0.007	0.962	0.005
	2	0.884	0.007	0.880	0.007	0.880	0.005
	3	0.766	0.007	0.827	0.008	0.795	0.005

Table 6.4: Table of simulation results: single central channel *vs.* all channels. The bottom row shows ratios.

If the standard models are to be believed, the integration times for the two deep fields are such that none of the bands are expected to yield a detection. The preceding test reveals the errors introduced by the ten channel average in the presence of a strong signal. Since we do not expect to detect a signal, we should consider the effect of the ten channel average on data which are pure noise. Simulations with a vanishingly small signal show that the distributions of the best fit bandpowers—nearly all of which are nondetections at the 2σ level—do not change when the data are averaged. In light of this result, we must use great caution when applying the scaling in the bottom row of Table 6.4 to the real data. We will revisit this point in the discussion of the real data in Section 6.5.

6.4.3 Primary Beam Pattern

The likelihood analysis routine assumes that the primary beam pattern for the CBI is a perfect Gaussian, but the beam pattern measurements (Figure 6.1) show




Figure 6.6: Comparison of the diagonal window functions for a short baseline and a long baseline. The two sets of window functions represent the low and high-frequency channels for each of the two baselines, and the arrows denote the positions of the band centers. The figure shows that because of the interplay between the $\sim 30\%$ fractional change in baseline length with the fixed size of the autocorrelation function of the illumination pattern, the edge channels for the long baselines measure virtually independent fluctuations while those for the short baselines are highly correlated.

that the real beam only resembles a Gaussian. The Gaussian approximation vastly expedites the likelihood calculation because under this assumption the integrals over the beam in the covariance matrix become Bessel functions, which *Mathematica* obtains from installed functions rather than from time-consuming numerical integrations. The shape of the primary beam affects the likelihood calculation because the autocorrelation of the illumination pattern determines the correlation between visibilities. Intuitively, as we increase the size of the beam, we improve the precision with which we reject all Fourier modes except that which corresponds to $\ell \sim 2\pi u \sim 2\pi|b|/\lambda$; under these circumstances, the covariance matrix becomes more diagonal because the $\tilde{A}_p(\mathbf{u}_p - \mathbf{v})$ approach delta functions. In this regard, our beam approximation is a second order effect: it affects the way in which power is distributed between the central part of the beam and the wings, which in turn

affects the likelihood calculation through the degree of correlation in covariance matrix. We might therefore expect that since the visibilities for the shortest baselines have the greatest overlap in the aperture domain, the effect of the beamshape is most pronounced for the low ℓ modes; we will show that this is the case.

We can use the simulated data to quantify the effect of the beam pattern on the best fit bandpowers. We performed two sets of simulations for comparison. The first set confined the analysis to a single channel to isolate the effect of the beam; for this analysis, 10^3 sets of maps were simulated in POLSKY with the standard $C_\ell^{EE} = 6 \mu\text{K}$ spectrum. These maps were imaged with POLFAKE twice; first with the Gaussian approximation to the real beam, and again with POLFAKE’s interpolation of the real CBI beam. Both sets of data passed through the likelihood analysis pipeline, which assumes a 45.2’ Gaussian beam at 31 GHz for both cases. The first set of data provides a benchmark for the algorithm’s ability to recover the input power, while the second mimics the conditions under which the real data are obtained. The extent to which the best fit bandpowers for the second case depart from those for the first provides a measure of the effect of the Gaussian beam assumption on the analysis.

Table 6.5 provides the results of beam pattern tests. The second set of entries shows the results of imaging the data with a real beam while extracting bandpowers under the Gaussian approximation. These amplitudes are smaller than those in the first row by as much as 3%, and while these differences are marginally significant given the number of simulations, an inspection of the last set of entries suggests that there is a trend for which longer baselines are less susceptible to beam shape effects than shorter baselines. Since the illumination autocorrelation function determines the correlation between visibilities in the aperture plane, this result depends on the (u, v) coverage of the observations. These tests show that when we compute the correlation using a Gaussian beam for data which have been observed with the CBI beam—an assumption that is forced upon us by the considerable improvement in speed—we underestimate the likelihoods by at $\sim 3\%$ for the shortest baselines, and less for the longer baselines.

		08 ^h field (μK)		20 ^h field (μK)		joint fit (μK)	
test	band	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}
Gauss beam	1	5.99	0.03	5.95	0.03	5.97	0.02
	2	5.97	0.03	6.02	0.03	5.97	0.02
	3	5.98	0.04	6.03	0.05	6.04	0.04
CBI beam	1	5.83	0.03	5.80	0.03	5.81	0.02
	2	5.87	0.03	5.97	0.03	5.89	0.02
	3	5.93	0.04	5.97	0.05	5.95	0.03
CBI/Gauss	1	0.973	0.007	0.975	0.07	0.973	0.005
	2	0.983	0.007	0.992	0.07	0.987	0.005
	3	0.992	0.010	0.990	0.010	0.985	0.008

Table 6.5: Table of simulation results for two beam patterns, $\Delta\nu = 1$ GHz, 10^3 nights. This table shows that when we compute the likelihood under the assumption of a Gaussian beam for data which have been observed with the real CBI beam, we underestimate the amplitudes by at most 3%. This effect is most pronounced on the shortest baselines.

This single channel test ignores the effect of the 30% change in the CBI’s beamwidth across the CBI band. The beammap data span all ten channels, and the beams for the ten channels were fit separately; POLFAKE uses these fits for each of the ten bands when the data are simulated with the CBI beam. Similarly, POLFAKE applies a simple $1/\nu$ scaling to the Gaussian beam using the beamsize supplied by the user for 30 GHz as a reference. Based on the results of the single channel beam test, we might reasonably expect that the effect of the beamshape is small for all ten channels. A second set of beam simulations were performed to test this hypothesis, and these results show that the additional frequency dependent effects in the beamshape are marginal; the dominant effect is the suppression of signal from the ten channel average at each (u, v) point which we saw in Section 6.4.2.

6.4.4 Perturbations to the Calibration

The preceding tests explored the effects of the approximations which enter the data reduction pipeline. Real data suffer from errors in the gain calibration and the instrumental polarization correction, however, so we must also consider the effects of these errors on the best fit likelihoods. Recall that the cross polarized visibility is given by,

$$\mathcal{V}^{LR} = G \left[P^* e^{-2i\phi} + \epsilon I \right] \quad (6.36)$$

Calibration errors enter the visibility through G and ϵ , so we consider four quantities: G_A , G_ϕ , ϵ_A , and ϵ_ϕ . For the purpose of this analysis, we will subsume systematic errors in I into the errors in ϵ . The calibration errors can bias the best fit bandpowers up or down, but we are particularly concerned with the latter, because effects which *reduce* our limit can cause us to misinterpret the underlying physics.

Simulations are necessary to understand the effects of calibration errors on the best fit bandpowers. Intuitively, we might expect that since $\ln[\mathcal{L}] \sim P^2$, the best fit bandpowers scale as G_A^{-2} . The gain phase G_ϕ does not lend itself to such a simple analysis, and the leakages are more complicated still, so we must rely on simulations to quantify their effects. As noted above, we have four degrees of freedom, but these four parameters can take a continuous range of values. We do not want to map this entire space, so we explore the effects of these parameters separately for a few points which approximate our understanding of the calibration errors, and if necessary, we apply the attendant corrections to the best fit bandpowers for the real data.

This simple approach belies an important layer of complexity. Both of the deep fields contain many tens of nights, each of which has a different calibration which is accompanied by an error, and the ideal approach would simulate each night of data with a separate calibration which is drawn from the known distribution of calibrations, such as the one shown in Figure 5.4. We chose to forgo this analysis in favor of an approach which isolated the components of the gain and leakage.

Noiselike errors will sum incoherently, so we are most concerned with the effects of systematic calibration errors; we believe, for example, that the Tau A model has a 10% amplitude error, and we know that the inclusion of baseline RX10-RX12 in the 08^h data introduces a $\sim 8^\circ$ phase error for that baseline. The simulations presented in this section will allow us to quantify the effects of these calibration errors on the real data.

Perturbations to the Gain

We first explored the effect of errors in the gain calibration. Simulations with perturbations to G_A show that the best fit amplitudes vary as G_A^{-1} , and this result confirms that the simulation-calibration pipeline performs as required. In contrast with the case of the gain amplitude errors, the gain phase errors are not amenable to a simple analysis, but the simulations provide some guidance about their effects on the bandpowers. The simulated data were calibrated with a variety of gain phase errors; the gain phases were rotated in steps of 6° , or $\sim 10\%$ of a radian, about $G_\phi = 0^\circ$. Table 6.6 lists the results of these tests. The table shows that small phase errors ($\pm 6^\circ$) have a marginal effect on the best fit amplitudes. The phase errors have a more deleterious effect on the short baselines than on the long baselines. These tests also suggest that the systematic phase error of $\sim 8^\circ$ for RX10 during the configuration 1 observations of the 08^h field will *increase* the best fit amplitude by a few percent, but weight which this increase carries is mitigated by the contribution to the band 1 likelihood from RX0-RX12, which does not have an analogous systematic phase calibration error.

Perturbations to the Leakage

The polarization analysis must consider errors in the instrumental polarization calibration. In the 08^h field observations, for example, the instrumental polarization inferred on a given night can differ from the expected value by $\sim 15\%$, and while the resulting error in the leakage correction contributes to the rms of the visibilities, these variations are noiselike, so these errors fall as $1/\sqrt{n} \sim 2\%$ for the $\sim 10^2$

		08 ^h field (μK)		20 ^h field (μK)		joint fit (μK)	
test	band	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}
$G_\phi = +18.0^\circ$	1	8.50	0.10	7.03	0.06	7.45	0.06
	2	5.99	0.03	6.49	0.05	6.15	0.03
	3	5.31	0.04	5.73	0.05	5.68	0.03
$G_\phi = +12.0^\circ$	1	6.70	0.05	6.37	0.04	6.43	0.03
	2	5.96	0.03	6.16	0.04	6.00	0.02
	3	5.81	0.04	5.86	0.05	5.82	0.03
$G_\phi = +6.0^\circ$	1	6.09	0.04	6.13	0.04	6.08	0.03
	2	5.95	0.03	5.98	0.03	5.95	0.02
	3	5.91	0.04	5.93	0.05	5.91	0.03
$G_\phi = 0^\circ$	1	5.89	0.03	6.00	0.03	5.94	0.02
	2	5.94	0.03	5.93	0.03	5.92	0.02
	3	5.94	0.04	5.95	0.05	5.93	0.03
$G_\phi = -6.0^\circ$	1	6.01	0.03	6.12	0.04	6.04	0.02
	2	5.96	0.03	6.03	0.04	5.95	0.02
	3	5.92	0.04	5.92	0.05	5.91	0.03
$G_\phi = -12.0^\circ$	1	6.63	0.05	6.35	0.04	6.40	0.03
	2	5.98	0.03	6.17	0.04	6.01	0.02
	3	5.83	0.04	5.84	0.05	5.77	0.0n
$G_\phi = -18.0^\circ$	1	8.39	0.10	7.01	0.06	7.30	0.05
	2	6.03	0.03	6.50	0.05	6.16	0.03
	3	5.68	0.04	5.70	0.05	5.68	0.03

Table 6.6: Table of simulation results for perturbations to the gain phase G_ϕ . The table shows that errors of $\sim 10\%$ of a radian have little effect on the measured amplitudes.

nights of data on the two deep fields. We used simulations to explore the effects of leakage correction errors on the inferred likelihoods. POLSKY and POLFAKE cannot contaminate the simulated LR visibilities with instrumental polarization, but we can force instrumental polarization into the visibilities with CBIPOLCAL. For the preceding tests we zeroed the leakage entries in the *.cal* file; we now set these entries to values which approximate the real errors, and CBIPOLCAL oversubtracts a component of LL from the LR visibilities based on this erroneous leakage. Since the likelihood routine measures the rms of the fluctuations, oversubtracted and undersubtracted leakage have the same effect on the best fit bandpowers. In the tests that follow, we perturb the amplitudes and the phases of the leakage separately.

For the leakage error tests we assumed flat spectra for both $\sqrt{C_\ell^{TT}}$ ($= 60 \mu\text{K}$) and $\sqrt{C_\ell^{EE}}$ ($= 6 \mu\text{K}$). The leakage amplitude was stepped in 1% increments, each of which corresponds to a $\sim 10\%$ error given the CBI's instrumental polarization of $\epsilon \sim 10\%$. Table 6.7 lists the results of these tests; these results show that typical errors in the instrumental polarization correction have a negligible effect on the best fit bandpowers. A systematic absolute leakage amplitude error of 2% ($\epsilon_A = 0.02$)—which corresponds to a 20% error in the measurement of the leakage term—results in a 5% change in the amplitude of the best fit bandpower for the worst case. The phase of the leakage has no effect on the best fit bandpowers. Table 6.8 shows the results for a sample of tests for which $\epsilon_A = 0.02$, $\epsilon_\phi = 0^\circ, 45^\circ$ and 90° ; both sets of amplitudes agree with those for data for which $\epsilon_\phi = 0^\circ$.

We can use simple arguments to estimate the effect of the leakage amplitude errors. If the leakage error is 2% in amplitude, the contribution to the rms from the error is $\epsilon\sqrt{C_\ell^{TT}} \sim 0.02 \times 60 \mu\text{K} \sim 1.2 \mu\text{K}$. This signal adds in quadrature with C_ℓ^{EE} , so $\sqrt{q'} \sim \sqrt{6^2 + 1.2^2} \sim 6.12 \mu\text{K}$ for a 2% effect. This is within a factor of ~ 2 of the results seen in Table 6.7. The most important result of this analysis is that on average the leakage amplitude errors *degrade* the best fit bandpowers: they only increase the bandpowers.

The shapes of the spectra C_ℓ^{TT} and C_ℓ^{EE} affect the weight which the contamination from C_ℓ^{TT} carries in our estimate of C_ℓ^{EE} , so we repeated the preceding

		08 ^h field (μK)		20 ^h field (μK)		joint fit (μK)	
test	band	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}
$\epsilon_A = 0.00$	1	5.89	0.03	6.00	0.03	5.94	0.02
	2	5.94	0.03	5.93	0.03	5.92	0.02
	3	5.94	0.04	5.95	0.05	5.93	0.03
$\epsilon_A = 0.01$	1	5.98	0.03	6.05	0.03	6.01	0.02
	2	5.98	0.03	6.03	0.04	5.96	0.02
	3	5.96	0.04	5.96	0.05	5.95	0.03
$\epsilon_A = 0.02$	1	6.34	0.04	6.31	0.04	6.30	0.03
	2	6.07	0.03	6.20	0.04	6.09	0.02
	3	5.96	0.04	5.96	0.04	5.95	0.03

Table 6.7: Table of simulation results for perturbations to the leakage amplitude ϵ_A . These results show that systematic errors in $\epsilon_A \sim 0.02$, which corresponds to a systematic $\sim 20\%$ error in the measurement of $\sim 10\%$ leakage, increase the amplitudes by a few percent.

analysis with data generated from our concordance cosmology.¹⁰ These simulations also included correlations from C_ℓ^{TE} . This analysis showed that an error in $\epsilon_A = 0.01$ boosts the best fit amplitudes by 2-3% for band 1 while leaving bands 2 and 3 unperturbed, while an error of $\epsilon_A = 0.02$ boosts the band 1 amplitudes by $\sim 10\%$ and the band 2 and 3 amplitudes by 2-3%. As with the tests with flat spectra, on average the leakage errors increase the best fit bandpowers for realistic C_ℓ^{TT} and C_ℓ^{EE} .

Conclusions

As we will see shortly, we have only upper limits to report, so our primary concern is that systematic calibration errors do not cause us to *underestimate* these limits. We know that errors in the gain amplitude simply scale the bandpowers, so these errors

¹⁰ $\Omega_0 = 1.0$, $h = 0.68$, $\Omega_b = 0.043$, $\Omega_{CDM} = 0.257$, $\Omega_\Lambda = 0.7$, $n_s = 0.95$, $\tau = 0$

		08 ^h field (μK)		20 ^h field (μK)		joint fit (μK)	
test	band	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}	$\langle q_i \rangle^{1/2}$	σ_{q_i}
$\epsilon_\phi = 0^\circ$	1	6.34	0.04	6.31	0.04	6.30	0.03
	2	6.07	0.03	6.20	0.04	6.09	0.02
	3	5.96	0.04	5.96	0.04	5.95	0.03
$\epsilon_\phi = 45^\circ$	1	6.33	0.04	6.31	0.04	6.29	0.03
	2	6.08	0.03	6.18	0.04	6.08	0.02
	3	6.01	0.04	6.02	0.05	6.00	0.03
$\epsilon_\phi = 90^\circ$	1	6.34	0.04	6.31	0.04	6.26	0.02
	2	6.07	0.03	6.20	0.04	6.09	0.02
	3	5.96	0.04	6.03	0.05	5.95	0.03

Table 6.8: Table of simulation results for perturbations to the leakage phase ϵ_ϕ . The leakage amplitude $\epsilon_A = 0.02$ for all cases. These simulations show that the absolute leakage phase ϵ_ϕ does *not* affect the best fit bandpowers, so we only require a measurement of the leakage amplitude ϵ_A .

can increase or decrease our limits. This section shows that on average systematic errors in the gain phase tend to increase the bandpowers, but modest errors of $\sim 10\%$ of a radian have little effect. The results of the likelihood analysis are fairly robust to systematic errors in the leakage correction; on average $\sim 10\%$ errors in the leakage amplitude produce only few percent increases in the bandpowers for flat spectra for C_ℓ^{TT} and C_ℓ^{EE} , and errors in the leakage phase do not significantly change the bandpowers. The simulations presented in this section suggest that the effects of calibration errors in the CBI data will be small and additive; in the next section we use the results of these simulations, combined with the known systematic errors in the calibration, to scale the limits from the deep observations.

6.5 Likelihood Analysis of the CBI Data

The simulations discussed in the preceding section demonstrate that the likelihood routine performs as desired, and they provide insights into the effects of the approximations which are necessary to make the calculation tractable. In this section we analyze real data with the routine. The real data require several changes: we scale the covariance matrix by a factor of two to account for the $\sqrt{2}$ increase in rms from differencing, and we scale the visibility variances by 1.06 to remove the bias in the noise estimates.

The likelihood analysis of the two deep fields yields limits for C_ℓ^{EE} for both fields and the joint fit. Table 6.9 lists these results; the levels represent 95% confidence limits; they are obtained by integrating the likelihood from $q_i = 0$. Figure 6.7 shows likelihood curves for all 9 limits in Table 6.9. The limits on the $\ell \sim 600$ scales approach the levels for C_ℓ^{EE} predicted by standard models, while those for the upper two bands are significantly higher. The effective bin widths are measured from the extreme edges of the FWHM of the diagonal window functions.

With a few small modifications we can apply the *LR* likelihood analysis routine to the *LL* data. To remove the polarization dependence, we let $I_4[v\overline{C}_{pq}] \rightarrow I_0[v\overline{C}_{pq}]$ (Equation 6.32) and set $\Theta_{pq} = 0$ (Equation 6.34). Upon making these changes, we find at $\ell = 603$ that $C_\ell^{TT} = 66.8_{-11.1}^{+14.1} \mu\text{K}$ for the 08^h field; this level is consistent with that of Padin *et al.* for the 08^h field at the same ℓ : $C_\ell^{TT} = 62.9_{-7.9}^{+11.3} \mu\text{K}$ [62]. The routines will not produce *identical* results because the algorithm presented in this thesis averages over all channels, and because the (u, v) sampling is different.¹¹ Nonetheless, the results are statistically equivalent.

Section 6.4 discusses a variety of factors which bias the best fit bandpowers, and some of these factors have been applied to the limits shown in Table 6.9. The average over the ten channels introduces a downward bias to the bandpowers which is most pronounced on long baselines, but since this bias affects the high S/N likelihoods, we do not apply it to our limits. The Gaussian beam approximation

¹¹The result above reflects only the subset of the short *LL* baselines which match the *LR* baselines.

				08 ^h	20 ^h	joint fit
band	ℓ_{min}	ℓ_c	ℓ_{max}	$\langle q_i \rangle^{1/2}$ (μK)	$\langle q_i \rangle^{1/2}$ (μK)	$\langle q_i \rangle^{1/2}$ (μK)
1	446	603	779	14.1	8.1	7.0
2	930	1144	1395	21.2	15.9	12.8
3	1539	2048	2702	45.3	27.7	25.1

Table 6.9: CBI limits on \mathcal{C}_ℓ^{EE} for the two deep fields and the joint fit; 95% confidence. The joint fit provides our deepest limit for each band.

in the likelihood routine decreases the best fit amplitudes by 3% for the short baselines and less for the longer baselines, and although this correction is somewhat model dependent, we make the conservative assumption that these factors affect our limits and thus increase our limits by the same amounts. Chapter 5 argues that the visibility uncertainties, when viewed as a proxy for the amplitude of the gain calibration, suggest that the calibration for the 08^h field is accurate to a few percent, while that for the 20^h field gives rise to flux densities which are systematically a few percent high due to the $\sim 8 - 10\%$ error in the polarized flux for the Tau A model. Since the 20^h visibilities are biased upward by $\sim 4\%$ by this model error, the 20^h limits were scaled downward by the same amount, and the joint fit was scaled downward by $\sim 2\%$. We saw that the phase error introduced by including baseline RX10-RX12 in the 08^h field dataset increases the amplitudes by 3% on average, but we do not correct the CBI limits for this error because this phase error only affects $\sim 1/3$ of the 08^h data for band 1, and the simulations show that the scaling for a 6° phase error for *all* the data in the band is only marginally significant. Table 6.10 summarizes the biases which were applied to the data in Table 6.9. As the table shows, many of these factors cancel.

Figure 6.8 compares the three CBI limits to the previous polarization measurements. The three CBI limits are correlated, so they must be considered separately; when these correlations are removed the limits will certainly improve, as the bands

		08 ^h field	20 ^h field	joint fit
bias	band	b_i	b_i	b_i
beam	1	1.03	1.03	1.03
	2	1.02	1.01	1.01
	3	1.01	1.01	1.02
G_A	all	1.00	0.96	0.98
net	1	1.03	0.99	1.01
	2	1.02	0.97	0.99
	3	1.01	0.97	1.00

Table 6.10: Likelihood biases for the real data based on known systematic errors in the calibration and the analysis.

will better reflect the fluctuations between ℓ_{min} and ℓ_{max} for the bin. The CBI limits are comparable to the deeper limits from previous measurements, and they complement the DASI measurements at lower ℓ . These limits are consistent with the concordance model¹² shown in the figure, but the CBI limits do not provide the power to discriminate between perturbations on this model. C_ℓ^{EE} is a weak function of the cosmological model over the windows of currently favored parameters, so we require far greater sensitivity if we are to use the EE spectrum in the CBI ℓ range to discriminate between models.

¹² $\Omega_0 = 1.0$, $h = 0.68$, $\Omega_b = 0.043$, $\Omega_{CDM} = 0.257$, $\Omega_\Lambda = 0.7$, $n_s = 0.95$, $\tau = 0$




Figure 6.7: Likelihood curves.


Comparison of CBI results to previous measurements of C_l^{EE} 95% c.l

Figure 6.8: CBI limits on C_l^{EE} with previous data. All limits upper limits are 95% confidence. The error bars on the single DASI detection are 68% confidence limits.

Chapter 7

Summary and Conclusions

7.1 Results of this Thesis

This thesis describes the major aspects of my contributions to the development and operation of the Cosmic Background Imager. These projects include the design and implementation of the downconverter and noise calibration system; the production of the low-noise HEMT amplifiers; the observation and analysis of ~ 50 hours of monitoring data with the VLA; the calibration and analysis of over 100 nights of CBI polarization data; and the likelihood analysis of these data to obtain limits on \mathcal{C}_ℓ^{EE} . In this chapter we summarize these efforts, and discuss how they affect our plans for the future of the CBI.

Chapter 2 presents an overview of the CBI and the factors which affect its performance. The chapter opens with a discussion of how visibilities provide measurements of C_ℓ , and then focuses on the aspects of the CBI which affect its performance for polarization. The polarization of each receiver is determined by the orientation of the quarter-wave plate. The CBI was deployed with all 13 receivers configured to measure *LCP*; we adjusted one receiver to measure *RCP* for the program described in this thesis. This change results in 12 cross polarized (*LR*) baselines and 66 total intensity (*LL*) baselines. The quarter-wave plate is part of a phase shifter assembly, and Chapter 2 presents a model for the instrumental polarization (leakage) which can arise from the phase shifter.

The latter part of Chapter 2 focuses on a number of instrumentation projects for the CBI which are part of this thesis: the downconverter, the noise calibration system, and the low-noise HEMT amplifiers. We describe the design and performance of the downconverter, which is a major component in the *rf* signal chain; it meets the CBI's requirements for size and performance. This downconverter design was also used by DASI. The CBI's internal noise calibration system was designed and implemented for this thesis, and while it has temperature-dependent instabilities which limit its efficacy for the total intensity observations, it proved to be very useful for transferring the polarization calibration between nights with $\sim 10\%$ accuracy. The last section of Chapter 2 discusses the production of the NRAO-designed low-noise HEMT amplifiers for the CBI. These critical components largely determine the sensitivity of the CBI. The amplifiers require specialized assembly and testing techniques, and the author learned these techniques at NRAO's Central Development Laboratory. On numerous occasions the CBI benefitted from having these in-house capabilities.

Chapter 3 discusses the polarization calibration of the CBI. Polarization observations must contend with contamination from instrumental polarization. We characterize the instrumental polarization with the leakage factor ϵ ; the leakage makes a negligible contribution to the total intensity visibilities ($\epsilon^2 P \ll I$), but it can be comparable to the polarized flux ($P \sim \epsilon I$). Chapter 3 demonstrates how the CBI's deck rotation, when applied to observations of sources of known polarization, can be used to determine the leakage and the gain for the cross polarized baselines. The CBI developed a polarization calibration package, CBIPOLCAL, which can derive and apply the leakage and the gain. By applying this procedure to observations of 3C279, we show that the phase shifter model developed in Chapter 2 can account for all of the instrumental polarization ($\epsilon \sim 10\%$).

We are concerned about the stability of the instrumental polarization, so we performed high S/N measurements of the leakage on 3C279 at regular intervals during the 2000 observing season. These observations show that the leakage is slightly less stable than the uncertainties on the measurements would suggest; we

can attribute these discrepancies to errors in the interpolation of the VLA data for 3C279 to the CBI, or to intrinsic changes in the instrumental polarization. We also show that the CBI's response to extended polarized emission is well behaved, so we measured the instrumental polarization at the beam half-power points. These measurements demonstrated that the off-axis polarization response is consistent with that at the center of the beam.

Although 3C279 was our primary polarization calibration for much of the 2000 season, we enlisted Tau A for the calibration on the many occasions for which 3C279 was not in view at night. With $P \sim 28$ Jy, Tau A has ten times the polarized flux of 3C279, but since the CBI resolves this source, we require a model of its extended emission. Unfortunately, we do not have deep observations of Tau A which are accompanied by observations of 3C279, nor does the literature contain discussions of Tau A's polarization at the CBI's resolution and band, so we must derive a model. Chapter 3 presents two different techniques for transferring the polarization calibration from the deep 3C279 observation on 06feb00 to the deep Tau A observation of 11jan00. The two techniques yield models which agree to within $\sim 4\%$ in amplitude and phase; they show that $|P| \sim 28$ Jy and $\chi \sim 28^\circ$ at 31 GHz. Chapter 3 presents a variety of tests that suggest that this model is good to this level, although tests in Chapter 5 show that the polarized component of the model is accurate to 10%, and we adopt the latter uncertainty for the model.

3C279 was the primary polarization calibrator for the polarization observations, and since it is a variable source, we monitored its polarization with the VLA for much of the 2000 season. Chapter 4 presents the results of the VLA monitoring campaign. These observations were performed with the VLA's *K* and *Q* band (22.46 and 43.34 GHz) channels which straddle the CBI's 26-36 GHz band, but at these high frequencies we had to deal with a number of difficulties associated with observations at the limits of the VLA's capabilities. A third of the data were lost to high (> 4 m/s) winds at the VLA site, and a significant number of antennas were excised from the analysis because of pointing and efficiency problems. The remaining data consisted of eight observations between 18feb00 and 10aug00 with

uncertainties of $\sim 5\%$ in amplitude and phase. During the VLA campaign 3C279's fractional polarization changed by $\sim 30\%$ ($m \sim 0.09 \rightarrow 0.12$) and its polarization position angle changed by $\sim 20^\circ$ at both K and Q band. We combined these measurements of m_ν and χ_ν with measurements of I_ν from the CBI to calibrate the cross polarized CBI baselines.

We routinely scheduled observations of 3C273 during the VLA campaign, and these data played an important role for both the polarization *and* the total intensity calibration. Comparisons of 3C273's polarization as measured with the CBI with that which was measured with the VLA demonstrated that the internal consistency of the CBI's polarization calibration is better than $\sim 5\%$. The total intensities of 3C273 from the VLA provided an independent test of the CBI's total intensity calibration; comparisons of the two sets of data demonstrated that the CBI's total intensity calibration is considerably better than $\sim 5\%$. Since the VLA and the CBI use different total intensity flux density calibrators, the agreement between the two sets of data provides great confidence in the accuracy of both flux density scales.

Chapter 5 presents the results of the polarization observations with the CBI. These observations are divided between deep field observations of the CMBR and supporting observations of polarized sources to test the CBI's polarization capabilities. The 08^h and 20^h deep field observations consist of 235 hours¹ of data spanning 99 nights, and a visual inspection of the maps suggests that these data do not contain a detection of the polarization of the CMBR. A χ^2 test on the LR visibilities does not allow us to rule out the possibility that the data contain a celestial signal, however. An analysis of the visibility uncertainties suggests that the gain amplitude calibration for the cross polarized baselines is accurate to better than $\sim 10\%$. A comparison of the uncertainties for deep field data calibrated on 3C279 with those for data calibrated on Tau A show a 10% excess in the Tau A model's polarized flux, so we adopt a 10% uncertainty for the Tau A model. The scatter in the visibility uncertainties indicates that the gain amplitude of the

¹This sum includes leads and trails.

calibration has an uncertainty of $\sim 6\%$, which is twice that for the total intensity calibration.

Chapter 5 also presents the results of polarization observations of a number of other polarized sources: the unresolved sources 3C273 and 3C274, the resolved double inner lobes of Centaurus A, and the supernova remnants W44 and G326.3-1.8. As noted above, the 3C273 observations confirm the accuracy of the polarization calibration at the $\sim 5\%$ level for data calibrated on 3C279. The polarization observations of 3C274 suggest that the combined variation in the source polarization and the CBI calibration is $m \sim 0.02$, or 2%. The observations of the extended source show good qualitative agreement with data in the literature, although the CBI operates at a combination of resolution and frequency for which there are few previous observations that permit quantitative comparisons.

Chapter 6 presents limits on the polarization of the CMBR in the 08^h and 20^h deep fields. This chapter opens with a review of the method of maximum likelihood for the analysis of the cross polarized visibilities—which differs from that for total intensity—and it discusses several approximations which are necessary to expedite the *Mathematica* likelihood algorithm that was implemented for this thesis. This algorithm does not, for example, account for the correlations between bands, nor does it consider B modes. We make extensive use of simulations in this chapter; we use simulations to verify that the algorithm performs as desired; to quantify the effects of several assumptions in the analysis; and to explore the effects of calibration errors on the best fit bandpowers. The latter set of tests is unique to this analysis; for total intensity observations we must consider the effects of errors in the gain amplitude G_A and gain phase G_ϕ on the results, but for polarization we must also contend with errors in the leakage amplitude ϵ_A and leakage phase ϵ_ϕ . We verify that the bandpowers scale with G_A , and then show that modest systematic errors in G_ϕ do not significantly affect the bandpowers. The simulations with errors in ϵ_A and ϵ_ϕ show that on average spurious contamination from C_ℓ^{TT} introduced by errors in ϵ_A tends to add in quadrature with the power in C_ℓ^{EE} , while errors in ϵ_ϕ do not change the bandpowers beyond the errors in ϵ_A . We are

most concerned with calibration errors that can bias the bandpowers downward. We know that errors in G_A can have this effect; the simulations demonstrate that on average errors in G_ϕ and ϵ_A tend to bias the bandpowers upward, while errors in ϵ_ϕ are negligible. Most importantly, these simulations show that the known systematic errors in the CBI calibration should not bias our limits downward.

Having demonstrated that the likelihood algorithm performs as desired, we apply the procedure to the 08^h and 20^h data to obtain limits on the polarization of the CMBR. We sort the cross polarized baselines into bins centered on $\ell = \{603, 1144, 2048\}$; the associated 95% confidence limits for the joint fits to the two fields are $C_\ell^{EE} < \{7.0, 12.8, 25.1\} \mu\text{K}$ under the assumption that $B = 0$. These limits approach the levels predicted by standard models for these scales; the limit at $\ell = 603$ in particular is comparable to the polarization signal on that scale, and the use of a shaped spectrum, coupled with corrections for the correlations between bands, will improve the scientific significance of this result.

7.2 Future Work

The thesis demonstrates that the CBI can perform accurate polarization observations. The success of the polarization program presented in this work, coupled with the success of the CBI's total intensity program, led to the upgrade of the CBI for polarization observations. Many of the challenges for the polarization calibration presented in this work are a consequence of the limited number of polarization calibrators, and this condition will persist for the dedicated polarization observations. This thesis demonstrated that secondary diagnostics, such as the visibility uncertainties and the leakage terms, can provide quantitative insight into the accuracy of the calibration in the absence of secondary sources.

After acquiring 2 years of total intensity data, the CBI was reconfigured to maximize its sensitivity to polarization, and these observations began in September 2002. To implement this change, the half-wave plates were removed from the rotating sections of the phase shifter assemblies and replaced with a multi-element

CBI Configuration 7




Figure 7.1: CBI configuration 7. This configuration was adopted for the dedicated polarization observations which began in September 2002. The hatched antennas denote the *LCP* antennas.

wave plate assembly employed by DASI that has uniform 1-2% instrumental polarization across the 26-36 GHz band [43]. These plates can be rotated, so the polarization assignments of the receivers can be changed by the control system. The assignments are divided nearly evenly between the 13 receivers, resulting in 42 cross polarized (*RL* and *LR*) baselines and 36 total intensity (*LL* and *RR*) baselines. In addition, the array was reconfigured to target scales of $\ell \sim 400 \rightarrow 2000$ (Figure 7.1), and with mosaicking we expect to extend this window to $\ell \sim 300$; this is the region over which C_ℓ^{EE} is expected to peak for standard models (Figure 6.8).

Many of the results of this work will play a role in the CBI's polarization program. Perhaps the most important result of this thesis is that it demonstrates that the CBI is capable of polarization observations: polarization imposes a unique and strict set of demands on the telescope, so it is not a foregone conclusion that this

should be the case. We demonstrated that the CBI's instrumental polarization, which is unimportant for total intensity observations, is not an impediment to an accurate polarization calibration. We showed that the polarization characteristics of the primary beams behave as expected, so the dedicated polarization observations will not require changes to the optics. We found that spillover and crosstalk between the antennas do not contaminate the polarization data at the level of sensitivity achieved in this work. These results have enabled us to convert the CBI to polarization observations with only minor changes to the instrument and the control and analysis software, and thus to expedite the CBI's dedicated polarization program.

Appendix A

Phase Shifter Model

A.1 Transfer Matrices

In this appendix we present the details of the calculation of the phase shifter model presented in Section 2.3.2. The effects of the receiver components on the input field \mathcal{E} can be expressed as a series of matrix operations: $\mathcal{E}' = \mathbf{m}_n \cdots \mathbf{m}_0 \mathcal{E} = \mathbf{M} \mathcal{E}$. As noted in Chapter 2, we express the field as $\mathcal{E} = \mathcal{E}^X \hat{\mathbf{e}}_x + \mathcal{E}^Y \hat{\mathbf{e}}_y$, so the transfer matrices are computed in the $(\hat{\mathbf{e}}_x, \hat{\mathbf{e}}_y)$ basis. In the discussion that follows, we associate \mathbf{M} with receiver j (*LCP*) and \mathbf{N} with receiver k (*RCP*).

The input field undergoes multiple transformations, each of which is characterized by a 2×2 transfer matrix. The transfer matrices are presented in the order in which they affect the incident radiation.

0. CBI deck orientation (or equivalently, the baseline orientation) modulates the position angle of the input polarization vector. To incorporate this effect, we must choose a reference frame, and there are two natural frames for this problem: the frame of the sky, and the frame defined by the rectangular waveguide in the receivers. The guide frame is fixed to the deck, so let the deck frame $(\hat{\mathbf{e}}_x, \hat{\mathbf{e}}_y)$ be defined by the minor axis of the rectangular guide in the receivers; the minor axis of the guide is preferred because the major axis is too large to support 26-36 GHz radiation—only radiation whose E mode is parallel to the minor axis propagates through the waveguide to the

components that follow. Nearly all 13 receivers are oriented in the same direction, and although the more general case of different orientations for each receiver is not considered here, this change would not have a significant effect on the outcome of this calculation. The transfer matrix which ties the sky frame to the guide frame has the form of a simple rotation:

$$\mathbf{m}_0 = \begin{pmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{pmatrix}; \quad \mathbf{n}_0 = \begin{pmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{pmatrix} \quad (\text{A.1})$$

where ψ is the deck orientation.

1. The first component which the incident radiation encounters after it passes through the feed is the half-wave plate. Assume that the half-wave plate is oriented at a position of ϕ with respect to the waveguide frame $(\hat{\mathbf{e}}_x, \hat{\mathbf{e}}_y)$. The second transfer matrix rotates the frame of the radiation into that of the half-wave plate, so that the half-wave plate is parallel to the $\hat{\mathbf{e}}'_x$ direction:

$$\mathbf{m}_1 = \begin{pmatrix} \cos \phi_j & \sin \phi_j \\ -\sin \phi_j & \cos \phi_j \end{pmatrix}; \quad \mathbf{n}_1 = \begin{pmatrix} \cos \phi_k & \sin \phi_k \\ -\sin \phi_k & \cos \phi_k \end{pmatrix} \quad (\text{A.2})$$

Although the control system assigns the half-wave plates at positions which remain constant for the duration of the observations, their zeros are not necessarily the same, so $\phi_j \neq \phi_k$. The resolution of the half-wave plate encoder limits the accuracy with which the half-wave plates realize the position requested by the control system: for a perfect system, ϕ is identical to the commanded position, while in practice, $\phi \rightarrow \phi + \delta\phi$ with $\langle \delta\phi \rangle \sim 1^\circ$. The rotating sections are exercised prior to each observing session, so $\delta\phi$ contributes a random error to successive measurements of the instrumental polarization.

In the presence of errors,

$$\begin{aligned}\mathbf{m}_1 &= \begin{pmatrix} \cos(\phi_j + \delta\phi_j) & \sin(\phi_j + \delta\phi_j) \\ -\sin\phi_j & \cos\phi_j \end{pmatrix}; \\ \mathbf{n}_1 &= \begin{pmatrix} \cos(\phi_k + \delta\phi_k) & \sin(\phi_k + \delta\phi_k) \\ -\sin\phi_k & \cos\phi_k \end{pmatrix}\end{aligned}\quad (\text{A.3})$$

2. The half-wave plate imparts an insertion phase and an insertion loss to the component of the field which is parallel to the plate. Insert a phase delay of α to the component of the rotated wave which is parallel to $\hat{\mathbf{e}}'_x$. Under ideal circumstances α is a half wave and there is no signal degradation, so the amplitude of the half-wave plate transfer matrix element is unity:

$$\mathbf{m}_2 = \begin{pmatrix} A_j e^{i\alpha_j} & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}\quad (\text{A.4})$$

In practice, simple bandpass errors across the 26-36 GHz band cause the insertion phase to depart from a half wave; call those errors $\delta\alpha_j$. The half-wave plate also attenuates the component of the wave which is parallel to the plate, so $A_j < 1$. To first order, the $\hat{\mathbf{e}}'_y$ component—which is perpendicular to the plate—suffers no complex insertion loss. Let $\alpha \rightarrow \pi + \delta\alpha$ and $A \rightarrow 1 - \delta A$ to obtain:

$$\mathbf{m}_2 = \begin{pmatrix} -1 + \delta A_j - i\delta\alpha_j & 0 \\ 0 & 1 \end{pmatrix}; \quad \mathbf{n}_2 = \begin{pmatrix} -1 + \delta A_k - i\delta\alpha_k & 0 \\ 0 & 1 \end{pmatrix}\quad (\text{A.5})$$

to first order in the errors.

3. Finally, reverse the rotation of \mathbf{m}_1 to return the incident field to the original reference frame:

$$\mathbf{m}_3 = \begin{pmatrix} \cos\phi_j & -\sin\phi_j \\ \sin\phi_j & \cos\phi_j \end{pmatrix}; \quad \mathbf{n}_3 = \begin{pmatrix} \cos\phi_k & -\sin\phi_k \\ \sin\phi_k & \cos\phi_k \end{pmatrix}\quad (\text{A.6})$$

4. After passing through the half-wave plate the radiation encounters the quarter-wave plate. Rotate the input polarization vector by an angle θ from the waveguide frame $(\hat{\mathbf{e}}_x, \hat{\mathbf{e}}_y)$ to the reference frame of the quarter-wave plate $(\hat{\mathbf{e}}''_x, \hat{\mathbf{e}}''_y)$:

$$\mathbf{m}_4 = \begin{pmatrix} \cos \theta_j & \sin \theta_j \\ -\sin \theta_j & \cos \theta_j \end{pmatrix}$$

Any errors in the orientation of the quarter-wave plate enter at this step: for a perfect system, $\theta_j = \pm\pi/4$, while in the presence of an error, $\theta_j = \pm\pi/4 + \delta\theta_j$. The polarization of the receiver determines the sign of the $\pi/4$ term: since receiver j is LCP, $\theta_j = -\pi/4 + \delta\theta_j$; similarly, $\theta_k = \pi/4 + \delta\theta_k$. The rotation into the quarter-wave plate frame takes the form

$$\begin{aligned} \mathbf{m}_4 &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 + \delta\theta_j & -1 + \delta\theta_j \\ 1 - \delta\theta_j & 1 + \delta\theta_j \end{pmatrix}; \\ \mathbf{n}_4 &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 - \delta\theta_k & 1 + \delta\theta_k \\ -1 - \delta\theta_k & 1 - \delta\theta_k \end{pmatrix} \end{aligned} \quad (\text{A.7})$$

5. The quarter-wave plate imparts a complex insertion loss to the signal. The transfer matrix for an ideal quarter-wave plate is simple:

$$\mathbf{m}_5 = \begin{pmatrix} B_j e^{i\beta_j} & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} i & 0 \\ 0 & 1 \end{pmatrix}$$

Of course, the quarter-wave plate introduces errors as well: let $B \rightarrow 1 - \delta B$, and $\beta \rightarrow \pi/2 + \delta\beta$. To first order the quarter-wave plate affects just the $\hat{\mathbf{e}}''_x$ component:

$$\mathbf{m}_5 = \begin{pmatrix} i - i\delta B_j - \delta\beta_j & 0 \\ 0 & 1 \end{pmatrix}; \quad \mathbf{n}_5 = \begin{pmatrix} i - i\delta B_k - \delta\beta_k & 0 \\ 0 & 1 \end{pmatrix} \quad (\text{A.8})$$

6. Reverse the rotation of \mathbf{m}_4 :

$$\begin{aligned}\mathbf{m}_6 &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 + \delta\theta_j & 1 - \delta\theta_j \\ -1 + \delta\theta_j & 1 + \delta\theta_j \end{pmatrix}; \\ \mathbf{n}_6 &= \frac{1}{\sqrt{2}} \begin{pmatrix} 1 - \delta\theta_k & -1 - \delta\theta_k \\ 1 + \delta\theta_k & 1 - \delta\theta_k \end{pmatrix}\end{aligned}\tag{A.9}$$

Note that we obtain the same result for $\mathbf{n}_6\mathbf{n}_5\mathbf{n}_4$ by letting $\theta = \pi/4$ instead of $\theta = -\pi/4$ in \mathbf{n}_4 and \mathbf{n}_6 and inserting the phase lag of i in the $\hat{\mathbf{e}}_y''$ direction in \mathbf{n}_5 .

7. The phase shifter assembly is followed by a transition from circular to rectangular waveguide. The quarter-wave plate passes a wave of the form $\mathcal{E}' = \mathcal{E}'_x\hat{\mathbf{e}}_x + \mathcal{E}'_y\hat{\mathbf{e}}_y$ to rectangular waveguide, and under ordinary circumstances, the waveguide cutoff in the $\hat{\mathbf{e}}_y$ direction of the rectangular guide would reflect $\mathcal{E}'_y\hat{\mathbf{e}}_y$ back through to the quarter-wave plate and create a hodgepodge of unruly reflections. To prevent this reflection, the signal path includes a lossy dielectric card parallel to $\hat{\mathbf{e}}_y$ immediately prior to the transition to absorb the unwanted $\mathcal{E}'_y\hat{\mathbf{e}}_y$ mode. The transfer matrix for the mode suppressor is trivial:

$$\mathbf{m}_7 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}; \quad \mathbf{n}_7 = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}\tag{A.10}$$

The \mathcal{E} -field at the mode suppressor output is simply the $\mathcal{E}'_x\hat{\mathbf{e}}_x$ component—a complex scalar—so we can dispense with the vector characterization for \mathcal{E}'_j for the remainder of this calculation.

At this point we compute the entire transfer matrices for both receivers: $\mathbf{M} = \mathbf{m}_7\mathbf{m}_6 \cdots \mathbf{m}_0$ and $\mathbf{N} = \mathbf{n}_7\mathbf{n}_6 \cdots \mathbf{n}_0$, and in turn the modified fields $\mathcal{E}'_j = \mathbf{M}\mathcal{E}_j$ and $\mathcal{E}'_k = \mathbf{N}\mathcal{E}_k$. In the presence of first order errors, the matrix elements of \mathbf{M} and \mathbf{N} each have $\sim 10^2$ terms, so these expressions will not be reproduced here. The

modified fields allow us to compute the visibility.

A.2 Visibility

The visibility is the time average of the complex product of the voltages produced by the two receivers:

$$\mathcal{V}^{LR} \sim \langle \mathcal{E}'_j \mathcal{E}_k^{*l} \rangle \sim \langle \mathbf{M} \mathcal{E}_j (\mathbf{N} \mathcal{E}_k)^* \rangle \quad (\text{A.11})$$

The visibility contains thousands of terms in varying orders of δ , but we will neglect all but those which are linear in δ . We identify \mathcal{E} -field coherence terms in the visibility with the Stokes parameters given by the standard definitions

$$\begin{aligned} \langle \mathcal{E}_j^X \mathcal{E}_k^{X*} \rangle &= \frac{1}{2}(I + Q) \\ \langle \mathcal{E}_j^X \mathcal{E}_k^{Y*} \rangle &= \frac{1}{2}(U + iV) \\ \langle \mathcal{E}_j^Y \mathcal{E}_k^{X*} \rangle &= \frac{1}{2}(U - iV) \\ \langle \mathcal{E}_j^Y \mathcal{E}_k^{Y*} \rangle &= \frac{1}{2}(I - Q) \end{aligned} \quad (\text{A.12})$$

and solve for I , Q , U , and V . The factor of $1/2$ preserves power; it ensures that the power in the orthogonal correlation terms sums to I , for example. The algebra for \mathcal{V}^{LR} is far too complicated to check by hand—the visibility contains $\sim 10^4$ terms—but in Section 2.3.2 we consider simple test cases to assess whether the calculation produces the expected results.

Bibliography

- [1] <http://map.gsfc.nasa.gov/index.html>.
- [2] J. W. M. Baars, R. Genzel, I. I. K. Pauliny-Toth, and A. Witzel. The absolute spectrum of CAS A—an Accurate Flux Density Scale and a Set of Secondary Calibrators. *A. & A.*, 61:99–106, October 1977.
- [3] C. Baccigalupi, C. Burigana, F. Perrotta, G. De Zotti, L. La Porta, D. Maino, M. Maris, and R. Paladini. Power spectrum of the Polarized Diffuse Galactic Radio Emission. *A. & A.*, 372:8–21, June 2001.
- [4] C. A. Beichman, G. Neugebauer, H. J. Habing, P. E. Clegg, and T. J. Chester, editors. *Infrared Astronomical Satellite (IRAS) Catalogs and Atlases. Volume 1: Explanatory Supplement*, volume 1, 1988.
- [5] M. F. Bietenholz. private communication.
- [6] M. F. Bietenholz, N. Kassim, D. A. Frail, R. A. Perley, W. C. Erickson, and A. R. Hajian. The Radio Spectral Index of the Crab Nebula. *Ap.J.*, 490:291, November 1997.
- [7] J. R. Bond. private communication.
- [8] J. R. Bond, C. R. Contaldi, U. L. Pen, D. Pogosyan, S. Prunet, M. I. Ruetalo, J. W. Wadsley, P. Zhang, B. S. Mason, S. T. Myers, T. J. Pearson, A. C. S. Readhead, J. L. Sievers, and P. S. Udomprasert. The Sunyaev-Zeldovich Effect in CMB-Calibrated Theories Applied to the Cosmic Background Imager Anisotropy Power at $\ell > 2000$. *astro-ph/0205386*, 2002.

- [9] J. R. Bond and G. Efstathiou. The Statistics of Cosmic Background Radiation Fluctuations. *M.N.R.A.S.*, 226:655–687, June 1987.
- [10] J. R. Bond, A. H. Jaffe, and L. Knox. Estimating the Power Spectrum of the Cosmic Microwave Background. *Phys. Rev. D*, 57:2117–2137, February 1998.
- [11] W. N. Brandt, C. R. Lawrence, A. C. S. Readhead, J. N. Pakianathan, and T. M. Fiola. Separation of Foreground Radiation from Cosmic Microwave Background Anisotropy using Multifrequency Measurements. *Ap.J.*, 424:1–21, March 1994.
- [12] J. O. Burns, E. D. Feigelson, and E. J. Schreier. The Inner Radio Structure of Centaurus A—Clues to the Origin of the Jet X-Ray Emission. *Ap.J.*, 273:128–153, October 1983.
- [13] B. Butler. Tipping Considerations at the VLA. *VLA Scientific Memo #170*, 1996.
- [14] C. Carilli and M. Holdaway. A Simple Minded Approach to Polarization Mosaics. *VLA Test Memo #163*, 1992.
- [15] C. Chandler. <http://www.aoc.nrao.edu/~cchandle/cal/cal.html>.
- [16] J. J. Condon, W. D. Cotton, E. W. Greisen, Q. F. Yin, R. A. Perley, G. B. Taylor, and J. J. Broderick. The NRAO VLA Sky Survey. *A.J.*, 115:1693–1716, May 1998.
- [17] R. G. Conway and P. P. Kronberg. Interferometric Measurement of Polarization Distribution in Radio Sources. *M.N.R.A.S.*, 142:11, 1969.
- [18] A. Cooray, A. Melchiorri, and J. Silk. Is the Cosmic Microwave Background Circularly Polarized? *Phys. Rev. Lett.*, *submitted*, March 2002.
- [19] D. P. Cox, R. L. Shelton, W. Maciejewski, R. K. Smith, T. Plewa, A. Pawl, and M. Różyczka. Modeling W44 as a Supernova Remnant in a Density Gradient

- with a Partially Formed Dense Shell and Thermal Conduction in the Hot Interior. I. The Analytical Model. *Ap.J.*, 524:179–191, October 1999.
- [20] J. Delabrouille, J. Kaplan, and The Planck HFI Consortium. Measuring CMB Polarisation with the Planck HFI. In *AIP Conf. Proc. 609: Astrophysical Polarized Backgrounds*, page 135, 2002.
- [21] J. R. Dickel, D. K. Milne, and R. G. Strom. Radio Emission from the Composite Supernova Remnant G326.3-1.8 (MSH 15-56). *Ap.J.*, 543:840–849, November 2000.
- [22] A. M. Flett and C. Henderson. Observations of the polarized emission of Taurus A, CAS A and Cygnus A at 9-mm wavelength. *M.N.R.A.S.*, 189:867–872, December 1979.
- [23] M. R. Griffith, A. E. Wright, B. F. Burke, and R. D. Ekers. The Parkes-MIT-NRAO (PMN) Surveys. 3: Source Catalog for the Tropical Survey ($-29^\circ < \delta < -9.5^\circ$). *Ap.J.s.*, 90:179–295, January 1994.
- [24] N. W. Halverson, E. M. Leitch, C. Pryke, J. Kovac, J. E. Carlstrom, W. L. Holzapfel, M. Dragovan, J. K. Cartwright, B. S. Mason, S. Padin, T. J. Pearson, A. C. S. Readhead, and M. C. Shepherd. Degree Angular Scale Interferometer First Results: A Measurement of the Cosmic Microwave Background Angular Power Spectrum. *Ap.J.*, 568:38–45, March 2002.
- [25] J. P. Hamaker and J. D. Bregman. Understanding Radio Polarimetry III. Interpreting the IAU/IEEE Definitions of the Stokes Parameters. *A. & A.S.*, 117:161–165, May 1996.
- [26] J. P. Hamaker, J. D. Bregman, and R. J. Sault. Understanding Radio Polarimetry I. Mathematical Foundations. *A. & A.S.*, 117:137–147, May 1996.
- [27] S. Hanany, P. Ade, A. Balbi, J. Bock, J. Borrill, A. Boscaleri, P. de Bernardis, P. G. Ferreira, V. V. Hristov, A. H. Jaffe, A. E. Lange, A. T. Lee, P. D. Mauskopf, C. B. Netterfield, S. Oh, E. Pascale, B. Rabbii, P. L. Richards,

- G. F. Smoot, R. Stompor, C. D. Winant, and J. H. P. Wu. MAXIMA-1: A Measurement of the Cosmic Microwave Background Anisotropy on Angular Scales of $10' - 5^\circ$. *Ap.J.L.*, 545:L5–L9, December 2000.
- [28] C. G. T. Haslam, H. Stoffel, C. J. Salter, and W. E. Wilson. A 408 MHz All-Sky Continuum Survey. II - The Atlas of Contour Maps. *A.&A.S.*, 47:1, January 1982.
- [29] M. M. Hedman, D. Barkats, J. O. Gundersen, J. J. McMahon, S. T. Staggs, and B. Winstein. New Limits on the Polarized Anisotropy of the Cosmic Microwave Background at Subdegree Angular Scales. *Ap.J.L.*, 573:L73–LL76, July 2002.
- [30] M. M. Hedman, D. Barkats, J. O. Gundersen, S. T. Staggs, and B. Winstein. A Limit on the Polarized Anisotropy of the Cosmic Microwave Background at Subdegree Angular Scales. *Ap.J.L.*, 548:L111–LL114, February 2001.
- [31] R. W. Hobbs and J. P. Hollinger. Linear Polarization of Radio Sources at 2.07 cm Wavelength. *Ap.J.*, 154:423, November 1968.
- [32] M. P. Hobson, A. N. Lasenby, and M. Jones. A Bayesian Method for Analysing Interferometer Observations of Cosmic Microwave Background Fluctuations. *M.N.R.A.S.*, 275:863, August 1995.
- [33] M. A. Holdaway and J. R. Pardo. Modeling of the Submillimeter Opacity on Chajnantor. *MMA Memo #187*.
- [34] W. Hu and S. Dodelson. Cosmic Microwave Background Anisotropies. *A.R.A.A.*, 40:171–216, 2002.
- [35] W. Hu and M. White. A CMB Polarization Primer. *New Astronomy*, 2:323–344, September 1997.
- [36] J. D. Jackson. *Classical Electrodynamics*. John Wiley and Sons, 1975.

- [37] D. R. Jiang, D. Dallacasa, R. T. Schilizzi, E. Luedke, H. S. Sanghera, and W. D. Cotton. 5 GHz EVN polarization of 3C 286. *A.&A.*, 312:380–386, August 1996.
- [38] N. Junkes, R. F. Haynes, J. I. Harnett, and D. L. Jauncey. Radio Polarization Surveys of Centaurus A (NGC 5128). I—The Complete Radio Source at 6.3 cm. *A.&A.*, 269:29–38, March 1993.
- [39] M. Kamionkowski, A. Kosowsky, and A. Stebbins. Statistics of Cosmic Microwave Background Polarization. *Phys. Rev. D*, 55:7368–7388, June 1997.
- [40] W. Kamionkowski and M. Kosowsky. The Cosmic Microwave Background and Particle Physics. *Ann. Rev. Nucl. Part. Sci.*, 2:77–123, 1999.
- [41] B. G. Keating, C. W. O’Dell, A. de Oliveira-Costa, S. Klawikowski, N. Stebor, L. Piccirillo, M. Tegmark, and P. T. Timbie. A Limit on the Large Angular Scale Polarization of the Cosmic Microwave Background. *Ap.J.L.*, 560:L1–L4, October 2001.
- [42] T. Kerr. private communication.
- [43] J. Kovac and J. E. Carlstrom. *I.E.E.E. M.T.T.*, in prep., 2002.
- [44] J. Kovac, E. M. Leitch, C. Pryke, J. E. Carlstrom, N. W. Halverson, and W. L. Holzappel. Detection of Polarization in the Cosmic Microwave Background using DASI. *astro-ph/0209478*, 2002.
- [45] M. R. Kundu and T. Velusamy. Brightness and Polarization Structure of Four Supernova Remnants 3C 58, IC 443, W 28 and W 44 at 2.8 Centimeter Wavelength. *A.&A.*, 20:237–244, August 1972.
- [46] E. M. Leitch, J. Kovac, C. Pryke, B. Reddall, E. S. Sandberg, M. Dragovan, J. E. Carlstrom, N. W. Halverson, and W. L. Holzappel. Measuring Polarization with DASI. *astro-ph/0209476*, 2002.

- [47] E. M. Leitch, C. Pryke, N. W. Halverson, J. Kovac, G. Davidson, S. LaRoque, E. Schartman, J. Yamasaki, J. E. Carlstrom, W. L. Holzapfel, M. Dragovan, J. K. Cartwright, B. S. Mason, S. Padin, T. J. Pearson, A. C. S. Readhead, and M. C. Shepherd. Experiment Design and First Season Observations with the Degree Angular Scale Interferometer. *Ap.J.*, 568:28–37, March 2002.
- [48] E. M. Leitch, A. C. S. Readhead, T. J. Pearson, and S. T. Myers. An Anomalous Component of Galactic Emission. *Ap.J.L.*, 486:L23, September 1997.
- [49] S. Masi, P. A. R. Ade, J. J. Bock, A. Boscaleri, P. de Bernardis, G. de Troia, G. di Stefano, V. V. Hristov, A. Iacoangeli, W. C. Jones, T. Kisner, A. E. Lange, P. D. Mauskopf, C. Mac Tavish, T. Montroy, C. B. Netterfield, E. Pascale, F. Piacentini, F. Pongetti, G. Romeo, J. E. Ruhl, E. Torbet, and J. Watt. B2K: The Polarization-sensitive BOOMERanG Experiment. In *AIP Conf. Proc. 609: Astrophysical Polarized Backgrounds*, page 122, 2002.
- [50] B. S. Mason, T. J. Pearson, A. C. S. Readhead, M. C. Shepherd, J. L. Sievers, P. S. Udomprasert, K. Cartwright, J. A. J. Farmer, S. Padin, S. T. Myers, J. R. Bond, C. R. Contaldi, U.-L. Pen, S. Prunet, D. Pogosyan, J. E. Carlstrom, J. Kovac, E. M. Leitch, C. Pryke, N. W. Halverson, W. L. Holzapfel, P. Altamirano, L. Bronfman, S. Casassus, May. J., and M. Joy. The Anisotropy of the Microwave Background to $\ell = 3500$: Deep Field Observations with the Cosmic Background Imager. *astro-ph/0205384*, 2002.
- [51] C. H. Mayer, T. P. McCullough, and R. M. Sloanaker. Polarization of the Radio Emission of Taurus A, Cygnus A, and Centaurus A. *A.J.*, 67:581, March 1962.
- [52] P. G. Mezger, R. J. Tuffs, R. Chini, E. Kreysa, and H.-P. Gemuend. Maps of Cassiopeia A and the Crab Nebula at $\lambda 1.2$ mm. *A.&A.*, 167:145–150, October 1986.
- [53] D. K. Milne, J. L. Caswell, M. J. Kesteven, R. F. Haynes, and R. S. Roger. Ra-

- dio Polarization and Magnetic Fields in Six Supernova Remnants. *Proceedings of the Astronomical Society of Australia*, 8:187–194, 1989.
- [54] D. Morris, V. Radhakrishnan, and G. A. Seielstad. On the Measurement of Polarization Distributions Over Radio Sources. *Ap.J.*, 139:551, February 1964.
- [55] S. T. Myers. <http://www.aoc.nrao.edu/~smyers/calibration/>.
- [56] S. T. Myers. <http://www.aoc.nrao.edu/~smyers/calibration/Gains/>.
- [57] S. T. Myers, C. R. Contaldi, J. R. Bond, U.-L. Pen, D. Pogosyan, S. Prunet, J. L. Sievers, B. S. Mason, T. J. Pearson, A. C. S. Readhead, and M. C. Shepherd. A Fast Gridded Method for the Estimation of the Power Spectrum of the CMB from Interferometer Data with Application to the Cosmic Background Imager. *astro-ph/0205385*, 2002.
- [58] C. B. Netterfield, P. A. R. Ade, J. J. Bock, J. R. Bond, J. Borrill, A. Boscaleri, K. Coble, C. R. Contaldi, B. P. Crill, P. de Bernardis, P. Farese, K. Ganga, M. Giacometti, E. Hivon, V. V. Hristov, A. Iacoangeli, A. H. Jaffe, W. C. Jones, A. E. Lange, L. Martinis, S. Masi, P. Mason, P. D. Mauskopf, A. Melchiorri, T. Montroy, E. Pascale, F. Piacentini, D. Pogosyan, F. Pongetti, S. Prunet, G. Romeo, J. E. Ruhl, and F. Scaramuzzi. A Measurement by BOOMERANG of Multiple Peaks in the Angular Power Spectrum of the Cosmic Microwave Background. *Ap.J.*, 571:604–614, June 2002.
- [59] C. B. Netterfield, M. J. Devlin, N. Jarolik, L. Page, and E. J. Wollack. A Measurement of the Angular Power Spectrum of the Anisotropy in the Cosmic Microwave Background. *Ap.J.*, 474:47, January 1997.
- [60] S. Padin. private communication.
- [61] S. Padin, J. Cartwright, M. C. Shepherd, J. K. Yamasaki, and W. L. Holzapfel. A Wideband Analog Correlator for Microwave Background Observations. *I.E.E.E. Trans. Instrum. Meas.*, 50:1234–1240, 2001.

- [62] S. Padin, J. K. Cartwright, B. S. Mason, T. J. Pearson, A. C. S. Readhead, M. C. Shepherd, J. Sievers, P. S. Udomprasert, W. L. Holzapfel, S. T. Myers, J. E. Carlstrom, E. M. Leitch, M. Joy, L. Bronfman, and J. May. First Intrinsic Anisotropy Observations with the Cosmic Background Imager. *Ap.J.L.*, 549:L1–LL5, March 2001.
- [63] S. Padin, M. C. Shepherd, J. K. Cartwright, R. G. Keeney, B. S. Mason, T. J. Pearson, A. C. S. Readhead, W. A. Schaal, J. Sievers, P. S. Udomprasert, J. K. Yamasaki, W. L. Holzapfel, J. E. Carlstrom, M. Joy, S. T. Myers, and A. Otarola. The Cosmic Background Imager. *P.A.S.P.*, 114:83–97, January 2002.
- [64] J. A. Peacock and J. V. Wall. Bright extragalactic radio sources at 2.7 GHz. I—The Northern Hemisphere Catalogue. *M.N.R.A.S.*, 194:331–349, January 1981.
- [65] T. J. Pearson. Simulations of CMB Polarization. *CBI Memo Series*, 2002.
- [66] T. J. Pearson, B. S. Mason, A. C. S. Readhead, M. C. Shepherd, J. L. Sievers, P. S. Udomprasert, J. K. Cartwright, A. J. Farmer, S. Padin, S. T. Myers, J. R. Bond, C. R. Contaldi, U.-L. Pen, S. Prunet, D. Pogosyan, J. E. Carlstrom, J. Kovac, E. M. Leitch, C. Pryke, N. W. Halverson, W. L. Holzapfel, P. Altamirano, L. Bronfman, S. Casassus, J. May, and M. Joy. The Anisotropy of the Microwave Background to $\ell = 3500$: Mosaic Observations with the Cosmic Background Imager. *astro-ph/0205388*, 2002.
- [67] T. J. Pearson, M. C. Shepherd, G. B. Taylor, and S. T. Myers. Automatic Synthesis Imaging with Difmap. *Bulletin of the American Astronomical Society*, 26:1318, December 1994.
- [68] A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Ap.J.*, 142:419–421, July 1965.
- [69] R. A. Perley. private communication.

- [70] R. A. Perley. private communication.
- [71] R. A. Perley. The Positions, Structures, and Polarizations of 404 Compact Radio Sources. *A.J.*, 87:859–880, June 1982.
- [72] R. A. Perley. VLA Observational Status Summary. 2000.
- [73] L. Piccirillo, P. A. R. Ade, J. J. Bock, M. Bowden, S. W. Church, K. Ganga, W. K. Gear, J. Hinderks, B. G. Keating, A. E. Lange, B. Maffei, O. Mallié, S. J. Melhuish, J. A. Murphy, G. Pisano, B. Rusholme, A. Taylor, and K. Thompson. QUEST-A 2.6 m mm-Wave Telescope for CMB Polarization Studies. In *AIP Conf. Proc. 609: Astrophysical Polarized Backgrounds*, page 159, 2002.
- [74] L. Piccirillo, G. Dall’Oglio, P. Farese, J. Gundersen, B. Keating, S. Klawikowski, L. Knox, A. Levy, P. Lubin, C. O’Dell, P. Timbie, and J. Ruhl. COMPASS: a 2.6 m Telescope for CMBR Polarization Studies. In *Experimental Cosmology at Centimeter Wavelengths*, page 183, 2002.
- [75] M. W. Pospieszalski. private communication.
- [76] M. W. Pospieszalski. Ultra-low-noise receivers for the 1 to 120 GHz frequency range. In *Proceedings of the 23rd European Microwave Conference*, 1993.
- [77] A. C. S. Readhead. Determination of C_ℓ from CBI Observations. *CBI Memo Series*, 1999.
- [78] A. C. S. Readhead, G. B. Taylor, W. Xu, T. J. Pearson, P. N. Wilkinson, and A. G. Polatidis. The Statistics and Ages of Compact Symmetric Objects. *Ap.J.*, 460:612, April 1996.
- [79] M. J. Rees. Polarization and Spectrum of the Primeval Radiation in an Anisotropic Universe. *Ap.J.L.*, 153:L1, July 1968.
- [80] R. B. Rengelink, Y. Tang, A. G. de Bruyn, G. K. Miley, M. N. Bremer, H. J. A. Roettgering, and M. A. R. Bremer. The Westerbork Northern Sky

- Survey (WENSS), I. A 570 Square Degree Mini-Survey Around the North Ecliptic Pole. *A.&A.S.*, 124:259–280, August 1997.
- [81] M. Rupen. Referenced Pointing at the VLA. *VLA Test Memo #202*, 1997.
- [82] R. J. Sault, J. P. Hamaker, and J. D. Bregman. Understanding Radio Polarimetry II. Instrumental Calibration of an Interferometer Array. *A.&A.S.*, 117:149–159, May 1996.
- [83] P. F. Scott, P. Carreira, K. Cleary, R. D. Davies, R. J. Davis, Dickinson. C., K. Grainge, C. M. Gutierrez, M. P. Hobson, M. E. Jones, R. Kneissl, A. Lasenby, J. P. Leahy, K. Maisinger, G. G. Pooley, R. Rebolo, J. A. Rubino-Martin, B. Rusholme, R. D. E. Saunders, R. Savage, A. Slosar, P. J. Sosa Molina, A. C. Taylor, D. Titterington, E. Waldram, R. A. Watson, and A. Wilkinson. First results from the Very Small Array III. The CMB Power Spectrum. *astro-ph/0205380*, 2002.
- [84] U. Seljak and M. Zaldarriaga. A Line-of-Sight Integration Approach to Cosmic Microwave Background Anisotropies. *Ap.J.*, 469:437, October 1996.
- [85] M. C. Shepherd. CBI Polarization Calibration with CBIPOLCAL. *CBI Memo Series*, 2002.
- [86] J. L. Sievers, J. R. Bond, K. Cartwright, J. C. R. Contaldi, B. S. Mason, S. T. Myers, S. Padin, T. J. Pearson, U.-L. Pen, D. Pogosyan, S. Prunet, A. C. S. Readhead, M. C. Shepherd, P. S. Udomprasert, L. Bronfman, W. L. Holzappel, and J. May. Cosmological Parameters from Cosmic Background Imager Observations and Comparisons with BOOMERANG, DASI, and MAXIMA. *astro-ph/0205387*, 2002.
- [87] J. Silk. Cosmic Black-Body Radiation and Galaxy Formation. *Ap.J.*, 151:459, February 1968.
- [88] M. Simard-Normandin, P. P. Kronberg, and S. Button. Integrated Linear

- Polarization of Extragalactic Radio Sources at 10.5 GHz ($\lambda 2.86$ cm) II. *A.A.S.*, 48:137, April 1982.
- [89] S. T. Staggs and et al. CMB Polarization Experiments. In *ASP Conf. Ser. 181: Microwave Foregrounds*, page 299, 1999.
- [90] R. Subrahmanyan, M. J. Kesteven, R. D. Ekers, M. Sinclair, and J. Silk. An Australia Telescope Survey for CMB Anisotropies. *astro-ph/0002467*, 2002.
- [91] G. B. Taylor. <http://www.aoc.nrao.edu/~gtaylor/calman/polcal.html>.
- [92] A. R. Thompson, J. M. Moran, and G. W. Swenson. *Interferometry and Synthesis in Radio Astronomy*. Krieger Publishing Company, 1986.
- [93] E. Torbet, M. J. Devlin, W. B. Dorwart, T. Herbig, A. D. Miller, M. R. Nolta, L. Page, J. Puchalla, and H. T. Tran. A Measurement of the Angular Power Spectrum of the Microwave Background Made from the High Chilean Andes. *Ap.J.L.*, 521:L79–L82, August 1999.
- [94] P. S. Udomprasert, B. S. Mason, and A. C. S. Readhead. Determining H_0 with the Cosmic Background Imager. In *Proceedings of the XXXVIIth Moriond Astrophysics Meeting: The Cosmological Model*, 2002.
- [95] F. Villa, N. Mandolesi, M. Bersanelli, R. C. Butler, C. Burigana, A. Mennella, G. Morgante, M. Sandri, and the LFI Consortium. The Low Frequency Instrument of the Planck Mission. In *AIP Conf. Proc. 609: Astrophysical Polarized Backgrounds*, page 144, 2002.
- [96] M. White, J. E. Carlstrom, M. Dragovan, and W. L. Holzapfel. Interferometric Observation of Cosmic Microwave Background Anisotropies. *Ap.J.*, 514:12–24, March 1999.
- [97] E. J. Wollack, N. C. Jarosik, C. B. Netterfield, L. A. Page, and D. Wilkinson. A Measurement of the Anisotropy in the Cosmic Microwave Background Radiation at Degree Angular Scales. *Ap.J.L.*, 419:L49, December 1993.

- [98] A. E. Wright, R. M. Wark, E. Troup, R. Otrupcek, A. Hunt, and D. J. Cooke.
Flux Densities at 8400 MHz for a Large Sample of Parkes Radio Sources.
Proceedings of the Astronomical Society of Australia, 8:261, 1990.