Cosmic Background Explorer (*COBE*)

Far Infrared Absolute Spectrophotometer (*FIRAS*)

Explanatory Supplement

Edited by S. Brodd, D. J. Fixsen, K. A. Jensen, J. C. Mather, and R. A. Shafer on behalf of the *COBE* Science Working Group

Version 4
FIRAS Project Dataset Final (Pass 4) Release
December 30, 1997

ABSTRACT

The Far Infrared Absolute Spectrophotometer (*FIRAS*) on the *COBE* satellite mapped the sky from 1 cm^{-1} to 97 cm^{-1} frequency with $\sim$ .5 cm^{-1} spectral resolution and excellent absolute calibration. The data products have been completed and delivered for the use of the scientific community. This supplement describes these products and gives instructions for their proper use, including computation of error estimates.

PREFACE

This version of the *FIRAS* Explanatory Supplement supersedes three previous versions: Version 2.1, dated July 14, 1995, Version 2.0, dated October 7, 1994, and Version 1.0, dated July 8, 1993.

FIRAS team members who contributed to the preparation of this document include:
W. J. Barnes, S. Brodd, D. A. Cottingham, R. E. Eplee, D. J. Fixsen, J. M. Gales, R. B. Isaacman, K. A. Jensen, J. C. Mather, J. S. Newmark, P. D. Noerdlinger, S. M. Read, L. P. Rosen, R. A. Shafer, F. A. Shuman, A. R. Trenholme, and K. A. Turpie.

A current version of the document will be kept on line at the National Space Science Data Center and may be obtained from the *COBE* Home Page at `http://www.gsfc.nasa.gov/astro/cobe/cobe_home.html` on the World Wide Web.

This document is to be referenced as:

> *COBE Far Infrared Absolute Spectrophotometer (FIRAS) Explanatory Supplement*, ed. S. Brodd, D. J. Fixsen, K. A. Jensen, J. C. Mather, and R. A. Shafer, *COBE* Ref. Pub. No. 97–C (Greenbelt, MD: NASA/GSFC), available in electronic form from the NSSDC.

Contents

List of Figures

List of Tables

1. Introduction

1.1. *COBE FIRAS* Data Overview

The Far Infrared Absolute Spectrophotometer (*FIRAS*) is a cryogenically cooled Michelson interferometer. It operated for 10 months of the *COBE* mission and mapped 99% of the sky, with 60% of the sky observed twice 6 months apart. It uses bolometer detectors and is calibrated absolutely by an external blackbody. The frequency coverage extends from 1 to 97 cm^{-1} (30 to 2910 GHz).

The spectral range of the *FIRAS* data includes most of the energy of the cosmic microwave background radiation, emission from interstellar dust, and the major far infrared cooling lines of the interstellar gas. The cosmological measurements to date have been reported by Mather *et al.* (1990), Mather *et al.* (1994), Fixsen *et al.* (1994a), Wright *et al.* (1994), Fixsen *et al.* (1996) (included as Appendix A), and Fixsen *et al.* (1997a), based on a calibration described by Fixsen *et al.* (1994b) (included as Appendix C). They include limits on the cosmological μ and Y parameters of 9.0×10^{-5} and 1.5×10^{-5} respectively. The lines detected include the rotational sequence of CO from the $J = 1 - 0$ transition up to the $J = 7 - 6$ transition, as well as two lines of [C I], two of [N II], and one of [C II]. These lines are very widespread, and the [C II] line is highly correlated with the dust emission. The [N II] lines are also correlated with the dust emission but the correlation is not as tight, presumably due to the higher ionization potential of the nitrogen, chemical abundance gradients, and differences in excitation conditions. The water vapor line at 37.136 cm^{-1} (1113.3 GHz) is detected in absorption against the Galactic center. Analyses of these data on dust and lines have been published by Wright *et al.* (1991), Petuchowski and Bennett (1993), Bennett *et al.* (1994) (included as Appendix D), Reach *et al.* (1995), Dwek *et al.* (1997), and Shafer *et al.* (1997).

In this release of *FIRAS* data, we include the following: skymaps of spectra that have been calibrated and that have had post-calibration offsets removed; uncertainties and errors associated with these skymaps; maps of the emission at frequencies of known emission lines, including the CO rotational ladder from the $J = 1 - 0$ transition to the $J = 7 - 6$ transition, the lines of [C I], [C II], and [N II], and several undetected lines; a map of the cosmic microwave background temperature; and maps of the Galactic dust spectra.

We refer to the process of removing post-calibration offsets as "destriping"; we also provide skymaps of spectra that have been calibrated but not destriped, together with the calibration models that were used and their associated errors.

The calibrated spectra were derived from averaged, or "coadded" interferograms. We include these interferograms in the data release as they are a suitable basis for certain kinds of recalibration strategies that would avoid the Fourier transformation into spectra. They may also be used to investigate deviations of blackbody spectra from Planck spectra, or to search for inconsistent data and other indications of improper action by the instrument or the analysis software.

We provide the original time-ordered interferograms and engineering data. These may be examined to search for time-dependent infrared phenomena and may also be used in studies of the aurora. The auroral electron current heated the *FIRAS* antenna (sky horn) quite substantially at certain times, far exceeding the temperature regulator capability for stabilization. The distribution of cosmic ray glitches is reported directly from the flight microprocessor.

1.2. Pass 4 Improvements In Data Products

This version (4) of the *FIRAS* Explanatory Supplement is specific to, and part of, the "Pass 4" data products, produced and released at the end of 1997. These products superceed all earlier released *FIRAS* products. They are the final deliveries to the community from the COBE project team.

The "Pass 4" data have significant changes in comparison to the earlier "Pass 3" data, derived from bug fixes and fundamental changes in the production algorithms. The process of using neighbors in the coaddition process (Section 4.3) allowed the inclusion of an additional 15% of the data. An error in the template formation was found and corrected. Further improvements in the calibration (Section 5), including padding which allowed a 20% improvement in resolution. Refinements in destriping and error analysis reduced the systematic errors. The net effect of all of these improvements is a factor of 2 reduction in the uncertainty in comparison to Pass 3. In addition, the Pass 4 release includes a number of ancillary data sets, which give values of some astrophyscially interesting quantities derived from the *FIRAS* data (Section 8).

1.3. Explanatory Supplement Organization

The *FIRAS* Explanatory Supplement is organized in the order of *FIRAS* data processing, beginning with the ingested data from the spacecraft telemetry, and ending with the results of the scientific analyses of the Galactic spectral lines, dust, and cosmic background temperature map. Sections 1-3 provide a general overview of the mission. Section 4 describes the processing the *FIRAS* data receive before they are calibrated, while section 5 describes the calibration and its application. Removal of post-calibration offset errors, or "destriping", is described in section 6; section 7 is a description of our error analysis. Algorithms for modeling Galactic spectral lines, Galactic dust, and cosmological parameters are described in section 8. Section 9 describes the various reference data sets that were used in the course of *FIRAS* processing, and section 10 offers advice for reading and using the *FIRAS* delivered project data sets in both FITS and native VAX binary formats.

2. Instrument Description

The Far Infrared Absolute Spectrophotometer (*FIRAS*) has been described in detail by Mather, Fixsen, and Shafer (1993), which is included in this document as Appendix B. Briefly, the *FIRAS* is a scanning, four-port (two input, two output) Michelson interferometer that uses polarizing grids as beamsplitters and creates an interferogram (*i.e.*, the Fourier transform of the source spectrum) by scanning a movable mirror platform (the "Mirror Transport Mechanism", or MTM). A dichroic filter at each output port (arbitrarily labeled "Left" and "Right") further splits each beam into low ($1 - 21$ cm^{-1}) and high ($20 - 97$ cm^{-1}) frequency bands. The four detectors are thus referred to as LH, LL, RH, and RL, for "Left High", etc.

The sky input antenna, or sky horn, of the *FIRAS* provides a 7° field of view with low side lobes. The beam pattern of the sky horn pattern is designed to be circularly symmetrical and unpolarized. The spacecraft rotates at 0.8 rpm throughout the mission, making more than 1/2 revolution while taking a single telemetered group of interferograms. This spin eliminates most of the beam asymmetry and polarization sensitivity that may appear in the data. The beam asymmetry and polarization sensitivity are further reduced by the averaging of data taken in random orientations of the instrument.

The MTM can be scanned at either of two speeds: "slow", at about 0.8 cm/s, or "fast", at about 1.2 cm/s of path difference. The interferograms ("IFG"s) are sampled at the detectors at the same clock rate in both cases (681.43 Hz), but the raw onboard samples are binned in different amounts depending on MTM scan speed so that the infrared Nyquist frequency of the telemetered data is the same. The sampling is controlled by the motion of the mirror. Data collected in the different scan speeds have been averaged ("coadded") and calibrated separately.

The MTM sweep can also be set to one of two scan lengths, "short" or "long", which limit the spectral resolution. Long scans have four times the maximum optical path difference compared to short scans. In the low frequency channels, the long scan mode thus yields four times the spectral resolution of the short. However, in the high frequency channel there is insufficient telemetry bandwidth to take advantage of the longer sweep, and only the first quarter of it is used. All high frequency data therefore have the same spectral resolution. Thus, the data can be divided into low resolution (LRES): low frequency short scan length, high frequency short scan length, and high frequency long scan length; and high resolution (HRES): low frequency long scan length. As with scan speeds, data taken in the two scan lengths have been processed separately. Table 2.1 shows the maximum optical path difference, spectral resolution, and other parameters of the high and low resolution scan modes and channels.

Table 2.1: Channel- and Scan Mode- Dependent Properties

		LRES	HRES
Maximum path difference	(cm)	1.22	5.85
Total scan length	(cm)	1.76	7.07
MTM sampling rate	Hz	681.43	681.43
Nyquist frequency	(cm^{-1})	145.21	36.30
Nyquist frequency	(GHz)	4353.3	1088.3
Unapodized resolution	(cm^{-1})	0.57	0.14
Unapodized resolution	(GHz)	17	4.2
Bin resolution	(cm^{-1})	0.45	0.11
Bin resolution	(GHz)	13.6	3.4
Apodized resolution	(cm^{-1})	0.82	0.23
Apodized resolution	(GHz)	24.6	6.9

The four possible MTM scan speed and length combinations are referred to as SS, SF, LS, and LF, for "Short Slow", etc. There are thus in principle sixteen possible detector/scan mode combinations for data processing: LLSS, for example, is "Left Low Short Slow".

Telemetered groups of interferograms take an average of about 40 seconds to collect; the range is from $\sim$32 to $\sim$46 seconds, depending on scan mode. The collection time varies with scan mode because multiple passes ("sweeps") of the MTM are averaged on board before telemetering the IFG. Either 16 or four sweeps (for short or long scans, respectively) are averaged, and the onboard sampling is set such that each telemetered IFG is 512 samples long. By design, the collection time for an IFG is about equal to the time that the instrument boresight scans across a single sky pixel.

The unapodized resolution of the calibrated spectra quoted in Table 2.1 is the reciprocal of the total scan length and is applicable if no assumption of symmetry of the interferograms is made. The bin resolution is smaller because in processing the data, we pad the 512 point IFG to 640 points with zeros. The apodized resolution of the spectra is the full width at half maximum of the absolute value of the Fourier transform of the apodization function.

One of the most important features of the *FIRAS* is that it is a differential device, wherein the measured signal is determined by the difference between the spectra at the two input ports. One of the two inputs (the sky horn) is usually open to the sky, while the other (the reference horn) is filled by an internal near-blackbody reference source called the ICAL. For the low frequency channels, the dominant celestial signal is the cosmic microwave background radiation (CMBR), which is well described by a blackbody spectrum at a temperature of 2.728 K. By adjusting the ICAL temperature to a point near that of the CMBR, a near null interferogram is produced and the derived spectrum is insensitive to errors in the instrument gain.

In order to determine the instrument gains and offsets and the deviation of the ICAL from a blackbody, the *FIRAS* is calibrated by placing a full-beam external calibrator (the XCAL) into the sky aperture. The XCAL is a highly accurate blackbody and is controlled over a range of temperatures that includes the CMBR temperature, thereby acting as a direct substitution for the sky. This provides a substantially accurate calibration for small deviations of the CMBR from a Planck spectrum. Gain uncertainties of 1% translate into uncertainties in the calibrated measurement of the CMBR to less than 0.01% of its peak intensity.

At higher frequencies the sky spectrum includes a noticeable contribution from interstellar dust emission. This is in excess of the CMBR and, having a much higher brightness temperature, is not balanced by the ICAL. For these signals, errors in the absolute instrument gain translate directly into photometric errors in the calibrated sky spectra. The calibration procedures included heating the XCAL and ICAL to a wide range of temperatures up to 22 K in order to determine the high frequency gain. (For comparison, the infrared color temperature of interstellar dust is $\sim$18 K.)

3. Mission Profile

The *COBE* was launched on 18 November 1989 into a 99° inclination, 900 km altitude sun-synchronous orbit with a 6:00 PM ascending node. The orbital period is 103 minutes, so there are almost exactly 14 orbits per day. The satellite orbit plane precesses at a rate of 1° per day so as to remain roughly perpendicular to the Earth-Sun line, and the *FIRAS* always looks away from the Earth at an angle of $91 - 94°$ from the Sun. A brief *FIRAS* mission timeline is given in Table 3.1 and Table 3.2.

As the *COBE* orbits the Earth, the *FIRAS* line of sight traces a path on the sky along a line of approximately constant ecliptic longitude. Each interferogram is a measurement of a certain swath along this line, whose width is equal to the 7° beamwidth of the instrument. The length of the swath is slightly greater than the beamwidth because of the observing time and the movement of the beam at a rate of 3.5° per minute. To organize the data, we divide the sky into 6144 equal area pixels (Figure 10.3). These are formed by constructing a cube with each face divided into $32 \times 32 = 1024$ squares, projected onto a circumscribed celestial sphere in ecliptic coordinates. The projection is adjusted to form equal area pixels having a solid angle of $4\pi/6144$ sr or 6.7 square degrees (Chan and O'Neill 1975). Because the 7° beamwidth of the sky horn is greater than the separation between pixels (2.6° average), this binning oversamples the sky.

As the *FIRAS* scans a circle almost perpendicular to the Sun direction, the coverage of the sky near the ecliptic poles is highly redundant, with some pixels being observed over a period of months. Conversely, pixels near the ecliptic plane are only observed during a two to five day period as the precessing orbit moves past them. After a six month interval *FIRAS* again scans over the same ecliptic longitudes, thereby providing redundant coverage at two widely separated epochs. Table 3.3 gives sky coverage information.

The incomplete sky coverage has two primary causes. The first cause is the 3-day calibration periods (when the XCAL filled the sky horn) that were not "made up" six months later. For the first nine months of the mission calibration periods were timed to coincide with lunar apparitions every 29 days, when contamination by lunar thermal emission where the Moon was within 22° of the horn would have meant that the sky data could not have been used. For the last six weeks of the mission the calibration periods were made shorter and more frequent, amounting to 50% of the data taken during this period.

The second cause of coverage gaps was the deliberate suspension of operation of the MTM during passage through the South Atlantic Anomaly (SAA). For the first month of the mission particle hits on the electronics caused repeated MTM malfunctions during SAA passages, when the drive motor would continue to run after the MTM had run into a stop

Table 3.1: *FIRAS* Mission Timeline

Start Time*	Stop Time	Event
893221435		*COBE* Launch
893251118		Aperture cover ejection
893261130		*FIRAS* first light
893280000		ICAL first nulled against sky
893400706		MTM toggles SS/LF scan modes
893430152		MTM held in stowed position in SAA
893531121	893552154	Calibration 1
900181432	900200255	Calibration 2
900190205		Horns commanded from 2.7 K to 2.75 K
900471432	900490255	Calibration 3
900770937	900791721	Calibration 4
900800115		MTM turned off in SAA
900810409	900840223	Galactic center crossing
901071048	901101923	Calibration 5
901290000		Start of eclipse season
901351311	901382338	Calibration 6
901391535		Horns commanded from 2.75 K to 6.0 K
901391535	901931849	6.0 K horn period
901710148		MTM SF scan mode replaced LF scan mode
901801348	901811408	Calibration 7
901931850		Horns commanded from 6.0 K to 4.0 K
901931850	902071103	4.0 K horn period
902070000		End of eclipse season
902071104		Horns commanded from 4.0 K to 2.75 K
902071104	902081119	Hot sky horn test
902141500		MTM SF scan mode discontinued
902141505	902200609	Calibration 8

*The format of the timetags is YY-DDD-HH-MM.

Table 3.2: *FIRAS* Mission Timeline, Continued

Start Time	Stop Time	Event
902200508		XCAL commanded to 2.75 K in stow
902250110		ICAL temperature toggles discontinued
902271352	902291408	Minical 1
902311344	902341455	Minical 2
902361428	902391638	Minical 3
902411552	902441356	Minical 4
902461300	902481604	Minical 5
902501540	902532300	Calibration 9
902540519	902590834	Galactic center crossing
902561414	902571040	Minical 6
902591016	902621026	Minical 7
902640949		Cryogen depletion

Table 3.3: *FIRAS* Sky Coverage

(all channels and scan modes)	
Number of pixels covered	6067 of 6144
Fractional sky coverage	99%
Total number of interferograms	1,173,514
Maximum number of IFGs in a pixel	3,433*
Minimum number of IFGs in a pixel	1
Average IFGs per pixel	191

*Maximum occurs in Pixel 5545, at ecliptic coordinates (longitude,latitude) = (135.0,-84.5) at the end of its stroke. The power dissipated as a result of these "end-of-travel" events would cause a spike in the helium bath temperature, resulting in poor quality data until the bath temperature had stabilized (typically after several minutes). One month after launch, operations were changed to force the MTM into a stowed position while in the SAA, thus minimizing the occurrence of the "end-of-travel" events. Three months later a further improvement was made when the MTM was powered off altogether during SAA passages.

Finally, the coverage gaps are due in small part to the North polar ecliptic "hole" caused by unstable horn temperatures during the eclipse period of the mission, when heating caused by the earth limb occurred.

Throughout the mission, the state of the *FIRAS* was cycled to minimize systematic errors. The scan speed of the MTM was toggled between slow and fast every other orbit to allow instrumental errors to be distinguished from sky signals. At the same time the scan length of the MTM was toggled between short and long to provide higher spectral resolution (at the expense of signal to noise). In addition, the temperature of the ICAL was toggled between a "sky null" setting to a setting 12 mK hotter every $3-4$ days to allow instrumental gain errors to be distinguished.

Table 3.4 gives the percentage of sky data taken in each scan mode over various periods of the mission. Because the LS scan mode was the most susceptible to MTM problems, it was not used after day 89343. The calibration data is 67.7% SS scan mode, 9.3% SF, and 23.0% LF. No calibration data were taken in the LS scan mode.

Table 3.4: *FIRAS* MTM Scan Modes

Time Period	SS	SF	LS	LF	All
8932611-8934301	11.1	45.4	27.4	16.1	5.2
8934301-9013915	40.7			59.3	61.9
9013915-9019318	40.6	27.6		31.7	16.6
9019318-9020711	38.3	61.7			5.9
9020711-9020811	100.0				0.1
9020811-9022004	32.5	67.5			2.6
9022005-9026409	97.8			2.2	7.8
Entire Mission	43.3	12.3	1.4	43.0	100.0

As discussed above, the *FIRAS* was calibrated photometrically on a monthly schedule in the early mission and then weekly toward the end. These calibrations are denoted as Calibration $1-9$ and Minical $1-7$ in Table 3.1. Calibration of the *FIRAS* was done by moving the XCAL into the sky horn and then adjusting the temperatures of four bodies: the XCAL, the ICAL, the sky horn, and the reference horn. The temperatures were adjusted over a range of $2 - 22$ K. The detector bias was also changed to verify the detector model parameters.

Because the *FIRAS* bolometers can operate effectively only below 2 K, the *FIRAS* was switched off when the 600 liter liquid helium supply was depleted after 10 months. During this time, the *FIRAS* scanned the sky 1.7 times.

4. Pre-Calibration Data Processing

A set of six software programs, or "facilities", convert the *FIRAS* data from raw telemetered data from the *COBE* satellite to averaged groups of interferograms, or "coadds", that can be Fourier transformed into spectra and calibrated. These facilities are, in order of their use: NFS (iNgest_FIRAS_Stripper); FPP (FIRAS_Pre_Processor); FDQ (FIRAS_Data_Qualify); FEC (FIRAS_Extract_Calibration); FSS (FIRAS_Sort_Sky); and FIL (FIRAS_Interferogram_coaddition_Long_spectra). This section describes these facilities and the delivered *FIRAS* data sets which they produce. A detailed description of this *FIRAS* data processing "pipeline" has been given by Read *et al.* (1992), and is included as Appendix E.

Note that in the following discussion, each "science record" contains one interferogram, or "IFG", together with associated information.

4.1. Ingest

Data from the *COBE* satellite were sent in a telemetry stream to ground stations and received at Goddard Space Flight Center. The *FIRAS* data were separated from other data in the telemetry stream by the *FIRAS* stripper, NFS, which produced ten data sets: one called NFS_HKP, the housekeeping data; one called NFS_ANC, the ancillary data; four called NFS_EMF_xx, the engineering mode timing data; and four called NFS_SDF_xx, the science data, including the IFGs. For the latter two data sets, the "xx" indicates one of the four instrument channels, LH, LL, RH, or RL.

Of these data sets, the data delivery includes the housekeeping, ancillary, and engineering mode timing data. The NFS_SDF data sets are further processed by FPP and FDQ, and so are not part of the delivery.

The housekeeping data, NFS_HKP, consist of time-ordered records each storing the engineering status of the *FIRAS* instrument, including the microprocessor and other status words, temperatures at various locations in the instrument and electronics, the Internal Power Distribution Unit voltages and currents, and other miscellaneous engineering quantities. (The cryogenic temperatures were taken by a set of Germanium Resistance Thermometers, or GRTs). These data are in unconverted telemetry counts. They are stripped from the *COBE* telemetry stream in pairs of major frames, and are not synchronized with the science data. The time of the first major frame is in the header section of the record and the time of the second major frame appears at the end of the record.

The raw ancillary data, NFS_ANC, consist of time-ordered records each containing the pair of telemetry major frames corresponding to the raw housekeeping data. For each major frame, the record contains 128 minor frames of status bits packed into one byte per minor frame. The status bits are the MTM scan direction bit, the external calibrator status bit, the MTM scan length, the MTM scan speed, and four microprocessor data ready bits.

The engineering mode timing data, NFS_EMF, consist of time-ordered records of Mirror Transport Mechanism timing information. These data provide a direct measurement of the time between MTM samples. For each channel, the flight microprocessors output the elapsed time between samples for one MTM sweep into a 512-point, 8-bit array. The time elapsed between sample pulses is computed in microseconds as follows:

$$time \;=\; 45 \,+\, (10.0 \cdot counts) \tag{1}$$

where *counts* is the timing data. Engineering mode timing data were taken early in the mission during the post-launch spacecraft checkout and late in the mission during post-cryogen depletion engineering tests. Analyses of these data yielded the MTM sampling rates discussed in Section 2. Fourier analysis of the data yielded the frequencies of the coherent vibrations discussed in Sections 5 and 7.

The NFS_HKP, NFS_ANC, and NFS_EMF data sets are released in their native VAX binary file formats, the record structures for which are given in Appendix H. The filename extensions of these files (*e.g.*, ED_893272157) incorporate the timetag of the earliest data record contained in the files, using the format YY-DDD-HH-MM.

4.2. Quality Checking

FPP and FDQ are the next two facilities in the *FIRAS* data processing pipeline.

The FPP facility has two main functions: to calculate the midpoint of each IFG's collection time, which is used to determine the appropriate spacecraft attitude solution for the IFG; and to perform checks on the quality of the data. To calculate the midpoint of collect time, FPP uses the data transmit time in the science record to extract the mirror transport mechanism (MTM) scan speed and stroke length from the corresponding engineering and status records. The on-board microprocessors for each detector channel telemeter 26 header words of information with every IFG collected. Several of these header words contain counters for the start of the IFG collection and the start of IFG transmission to the telemetry stream. FPP extracts the MTM synchronization flag and the number of sweeps that are averaged together to produce the IFG from the header. It uses these MTM data, the counters, and the actual transmit time for the IFG record to compute the midpoint of the collection time. If there are any anomalies in the values used in the time computation, the IFG science record is flagged as being of bad quality. The IFG is also flagged if there are telemetry data gaps or bad telemetry quality during its collection. Finally, FPP checks the values of the MTM scan speed, MTM stroke length, and the detector gain over the IFG collection time. If these values are not constant, the IFG is flagged. Any record flagged by FPP is not used in further pipeline processing.

Table 4.1 gives the number of science records collected by the *FIRAS* instrument in each of the four channels over the entire 10 month *COBE* mission, and gives the number failed by FPP. The data are divided into three groups according to the position of the external calibrator (XCAL). Calibration data records have the XCAL in the sky horn, while sky data records occur when the XCAL is in its stowed position and the *FIRAS* instrument is looking at the sky. Some data were collected while the XCAL was in transit between the sky horn and its stowed position; these records are not used. Also, a special data mode called "fake-it" was used for about 10% of the records to conduct a variety of engineering tests; these records are also not used in further processing.

Table 4.1: *FIRAS* Facility FPP

Channel	LH	LL	RH	RL	All
Total Science Records	590937	590926	591005	587637	2360505
Records Failed by FPP	22118	22865	23315	22684	90982
Fake-it Records	59637	58797	59136	58348	235918
XCAL Transit	430	430	312	291	1463
Records Eliminated Before FDQ	82185	82092	82763	81323	328363
Cal Records Passed by FPP	61769	61755	61915	61860	247299
Sky Records Passed by FPP	446983	447079	446327	444454	1784843

The FDQ facility has three main tasks: to continue the data quality checking begun by FPP; to associate the science records with appropriate engineering data (such as temperatures, voltages, and currents); and to determine the spacecraft attitude solution (and thus, a pixel number) for each IFG.

FDQ checks the microprocessor, engineering, status, and attitude quantities associated with each IFG, and sets quality flags for the records appropriately. Most of the data quality checks are for engineering quantities, but very few records failed any of these during the mission. The principal failures are for the microprocessor status word, saturated sample counts, elevated glitch rates, and attitude.

The microprocessor status word contains 16 bit flags which indicate errors or anomalies detected by on-board programs during the collection of IFGs. These errors include deglitcher math overflow, command overrun, data buffer clear error, sync and ADC pulse time-outs, illegal instruction fetch, sync errors, ADC buffer over-run, sample add and divide overflows, group coadd overflow, and group divide overflow. If any of these bits are set, FDQ gives the IFG a bad quality flag. The saturated sample count is defined as the number of ADC samples greater than or equal to 95% of the maximum ADC range. If this count is too high, the IFG is also failed. Data collection by the *FIRAS* instrument is affected by cosmic ray hits on the detectors, or “glitches”. A glitch rate is computed for each IFG from a telemetry data word containing the glitch count for the IFG. A glitch rate that is too high causes FDQ to fail the IFG.

FDQ also assigns bad data quality to science records based on attitude quantities. If no attitude solution is available for a particular IFG, that record is failed. It is also failed if the attitude is such that the IFG might have been effected by emission from the Sun, the Moon, or the Earth limb. Because the attitude had no demonstrable effect on data taken when the XCAL was in the sky horn, calibration IFGs are not checked for attitude failures.

Table 4.2 gives the number of calibration and sky data records failed by FDQ for each of these reasons.

Table 4.2: *FIRAS* Facility FDQ

Channel	LH	LL	RH	RL	All
Saturated Sample Count	4581	5499	4462	4938	19480
Microprocessor Status Word	685	317	585	307	1894
Glitch Rate	627	181	444	271	1523
Cal Records Failed by FDQ	5053	5903	4905	5339	21200
Cal Records Passed by FDQ	56716	55852	57010	56521	226099
Saturated Sample Count	26694	44098	26206	44242	141240
Microprocessor Status Word	6400	6672	6893	6421	26386
Glitch Rate	3619	3731	3541	3779	14670
Sun, Moon, or Earth Limb	7814	7807	7796	7788	31205
No Attitude Solution	569	569	566	329	2033
Sky Records Failed by FDQ	41281	59640	41206	58353	200480
Sky Records Passed by FDQ	405702	387439	405121	386101	1584363

FDQ produces five data sets: FDQ_SDF_xx, which contain the IFGs and the associated attitude information, and where "xx" is one of the instrument channels LH, LL, RH, or RL; and FDQ_ENG, which contains the engineering data.

The time-ordered FDQ_SDF records are created from the IFGs collected by the on-board microprocessors and from spacecraft attitude information corresponding to the midpoint of collect time of the IFGs. The engineering data, FDQ_ENG, consist of time-ordered records each storing the engineering status of the *FIRAS* instrument at a particular time. The status includes the microprocessor status words, temperatures, voltages, and currents. Both FDQ_SDF and FDQ_ENG records are stored in files each containing one day of *FIRAS* data.

Because the science and engineering data are collected asynchronously in order to avoid coherent crosstalk, FDQ must properly link them together. To do so, the science IFG records are grouped according to the proximity of the midpoint collection times. Each such group may have from one to four science records, up to one from each of the four channels. The two housekeeping major frames whose times bracket the average time of the science record group are found. FDQ converts the data in these housekeeping records from microprocessor counts to physical units, and then interpolates them to determine values for each engineering quantity at the average science record group time. This time is assigned to the new FDQ_ENG engineering record, and is also put into each of the FDQ_SDF science records in the group. The science record times are put into the engineering record, so that the science and engineering data are doubly cross-indexed.

The FDQ_SDF_xx and FDQ_ENG data sets are released in their native VAX binary file formats, the record structures for which are given in Appendix H. As with the NFS data, the filename extensions of these files (*e.g.*, ED_893280000) use the format YY-DDD-HH-MM; however, because the FDQ files each contain one day's worth of data, the hour and minute are always zero.

4.3. Sorting

At this stage of the *FIRAS* pipeline, the calibration data and sky data are separated and processed independently. The FEC (for calibration data) and FSS (for sky data) facilities perform two functions: to sort the individual IFGs into ensembles for coaddition; and to provide further checks for data that are unsuitable for further use.

During the *FIRAS* mission, there were 16 intervals of typically two to three days duration when calibration data were taken (Table 3.1). During calibrations, series of time-ordered IFGs were obtained; these series were separated either by: a change in the commanded temperature for at least one of the four temperature-controllable bodies (the external calibrator (XCAL), the internal calibrator (ICAL), the sky horn, or the reference horn); or by a change in the commanded detector bias voltage. Each such series constitutes a coaddable group of calibration IFG records. The averages of the temperatures of the four controllable bodies are computed for each series. An IFG within the series is removed from the ensemble if the temperatures of any of the four controllable bodies associated with that IFG differ from the average for the series by more than a specified tolerance. IFGs with deviant temperatures are classified as unstable and are removed from further pipeline processing. Many different sets of specified tolerances were used for the calibration data; as an example, the set used for the time range 900780133 to 900782047 was: XCAL .1%; ICAL .05%; sky horn .2%; and reference horn .05%.

Most of the calibration records which fail stability tests are at the beginning of a series, because a controllable body which has just experienced a commanded temperature change requires a finite time to relax to its new state. The other major cause of stability failure occurred following unplanned transient heatings of the sky horn by the bright earth limb near the end of the eclipse season, when horn temperatures were 2.7 K. The calibration data were screened to avoid these events. Some records were unstable because of transients which occurred during the initial five minutes or so of a series defined by a change in commanded bolometer bias voltage.

Each coaddable group of IFGs which pass the stability tests is subdivided if there is a change in MTM scan mode and further subdivided if there is a change in science gain. Groups containing more than 100 IFGs are divided into smaller groups of approximately equal size. The number of segments is chosen to maximize the size of the subgroups. For example, a coadd group containing 700 IFGs would be divided into seven time-ordered groups of 100 IFGs each, a coadd group containing 380 IFGs would be divided into four groups containing 95 IFGs each, etc.

Table 4.3 gives the number of calibration IFGs failed and passed by FEC, as well as the number of coadd groups the facility created.

Table 4.3: *FIRAS* Facility FEC

Channel	LH	LL	RH	RL	All
Cal Records Failed by FEC	19275	18918	19459	19187	76839
Cal Records Passed by FEC	37441	36934	37551	37334	149260
Cal Coadd Groups Passed by FEC	771	762	767	768	3068

To allow the FSS facility to form coaddable groups of sky IFGs, the FDQ data are separated into 11 mission time periods according to various events and changes in the instrument state. The time ranges and descriptions of the mission periods are given in Table 4.4.

Table 4.4: *FIRAS* Mission Periods

Start Time	Stop Time	Description of Interval
89-326-1130	89-327-2359	First light
89-328-0000	89-343-0151	First ICAL nulling
89-343-0152	90-019-0204	MTM uses position mode through SAA
90-019-0205	90-080-0114	Horns commanded from 2.70 K to 2.75 K
90-080-0115	90-128-2359	MTM uses power off through SAA
90-129-0000	90-139-1534	Eclipse season starts
90-139-1535	90-193-1849	Horns commanded to 6 K
90-193-1850	90-207-1103	Horns commanded to 4 K
90-207-1104	90-208-1119	Sky horn calibration, XCAL out
90-208-1120	90-220-0459	Horns commanded to final temperature
90-220-0500	90-264-0936	XCAL placed under temperature control

FSS is run on each mission period in turn, for all four channels. It first checks the sky IFG records for the following conditions, rejecting those which do not qualify: a solar aspect angle greater than 91.2°; an earth limb aspect angle greater than 87°; a lunar aspect angle greater than 22°; an ICAL temperature within 2 mK of one of several specified temperatures; a dihedral temperature less than or equal to 5.5 K; and a science mode of four (the standard operating science mode).

Table 4.5 gives the number of sky IFGs passed and failed by FSS, as well as the number of coadd groups it creates for each of the four channels.

FSS then sorts the acceptable sky IFG records into coaddable groups. It first sorts by pixel number (location on the sky), then by MTM scan mode (one of SS, SF, LS, or LF), then by ICAL temperature (within 2 mK of the specified ICAL temperatures, thus forming "bins"), and, finally, by several ranges of dihedral temperatures. Coadd groups of more than 100 IFGs are subdivided into smaller sized subgroups in the same way as calibration coadds are subdivided. Table 4.6 gives the ICAL temperatures (forming the center of the "bins") and the dihedral temperature ranges for each of the 11 mission periods.

Once the coadd groups have been formed, FSS forms additional groups of "neighbor" IFGs.

Table 4.5: *FIRAS* Facility FSS

Channel	LH	LL	RH	RL	All
Earth Limb Angle <87.0	10914	10763	10939	10706	43322
Wrong ICAL Temperature	7415	7099	7370	5250	27134
Sun Angle <91.2	6745	6764	6723	6768	27000
Wrong Science Mode	718	717	744	827	3006
Dihedral Temperature >5.5	381	380	371	332	1464
Sky Records Failed by FSS	23745	22586	23697	20752	90780
Sky Records Passed by FSS	381957	364853	381424	365349	1493583
Sky Coadd Groups Passed by FSS	78577	77562	78547	77592	312278

Table 4.6: *FIRAS* Facility FSS Temperatures

Start Time	Stop Time	ICAL Temperatures	Dihedral Temperatures
89-326-1130	89-327-2359	2.789	2.14,2.5,3.1,3.7,4.3,4.9,5.5
89-328-0000	89-343-0151	2.758,2.763,2.789	2.02,2.5,3.1,3.7,4.3,4.9,5.5
89-343-0152	90-019-0204	2.759,2.771	2.14,2.5,3.1,3.7,4.3,4.9,5.5
90-019-0205	90-080-0114	2.758,2.771	2.14,2.5,3.1,3.7,4.3,4.9,5.5
90-080-0115	90-128-2359	2.758,2.771	1.98,2.5,3.1,3.7,4.3,4.9,5.5
90-129-0000	90-139-1534	2.758,2.770	2.0,2.5,3.1,3.7,4.3,4.9,5.5
90-139-1535	90-193-1849	2.7455,2.755,2.768	2.03,2.5,3.1,3.7,4.3,4.9,5.5
90-193-1850	90-207-1103	2.746,2.757,2.769	2.01,2.5,3.1,3.7,4.3,4.9,5.5
90-207-1104	90-208-1119	2.757,2.769	2.01,2.5,3.1,3.7,4.3,4.9,5.5
90-208-1120	90-220-0459	2.758,2.770	2.0,2.5,3.1,3.7,4.3,4.9,5.5
90-220-0500	90-264-0936	2.758,2.771	2.0,2.5,3.1,3.7,4.3,4.9,5.5

For each coadd group, FSS finds sky IFG records with the same MTM scan mode, ICAL temperature within the same bin, and dihedral temperature within the same range from the eight surrounding, or "neighboring" pixels. It then computes the mean Galactic latitude of the original coadd group, and orders the neighbor IFGs by increasing absolute values of the difference between the Galactic latitude of each neighbor record and this mean.

Unlike calibration data, IFGs in sky coadd groups are not tested for XCAL, sky horn, or reference horn temperature stability. Analysis has shown that variations in the XCAL temperature do not have a detectable effect on the calibrated spectra. Because the Mission Periods are defined so that the horn temperatures do not vary greatly within them, no sky coadd contains records with greatly different horn temperatures. Horn temperature differences between coadds with the same sky pixel number obtained in different Mission Periods are accounted for by the calibration model (Section 5). Residual post-calibration effects are corrected by destriping (Section 6).

The FEC_SSCAL_xx and FSS_SSSKY_xx records, which are the FEC and FSS products and where "xx" denotes one of the channels LH, LL, RH, or RL, are released in their native VAX binary file formats, the record structures for which are given in Appendix H. The filename extensions of these files (*e.g.*, ED_8935312_8935323) incorporate the timetags of the earliest and latest data records contained in the files, using the format YY-DDD-HH. These files do not contain any science or engineering data. Instead, they are index files that refer back to the FDQ_SDF and FDQ_ENG files.

4.4. Coaddition

The FIL facility coadds the IFGs by performing the following functions: reading the data; checking the data for quality and instrument state consistency within each coadd group; creating and subtracting primary and, possibly, secondary templates; removing the effects of cosmic ray hits on the detectors (called "deglitching"); checking the shape of the IFGs for consistency within each group; coadding the science and engineering data using a weighting scheme based on glitch rates; creating, as required, "FS" or "FL" scan mode IFGs from SF or LF scan mode IFGs; and subtracting a baseline from each coadded IFG.

FIL first uses the FEC or FSS products to retrieve the FDQ_SDF science records (including the IFGs) and the FDQ_ENG engineering records. Members of the original coadd group (excluding neighbors) are always read. If FIL is operating on sky (FSS) data and if there are fewer than 12 IFGs in the original group, then neighbor IFGs are read (if they exist) until a group of 12 is formed.

The neighbor IFGs are not coadded together with the "original" IFGs. Instead, they are used for the template formation and deglitching process. Analysis shows that a group of at least three IFGs is needed to form a reliable template. Since many sky data coadd groups contain only one or two IFGs, the use of neighbors permits the IFGs in these groups to be included in template formation and deglitching, and thus prevents their loss. (The coadd groups of calibration data created by FEC always contain five or more IFGs).

FIL next checks that none of the IFGs in a coadd group were marked as bad by a previous pipeline facility, and that all of them are in the same instrument mode. Verification of the instrument mode includes checking that attitude quantities and operating conditions such as channel, scan mode, science mode, and temperature controller gain status bits are the same for all IFGs. For engineering quantities such as bolometer voltage, temperatures of the four controllable bodies, and bolometer temperature, FIL calculates midaverages of the coadd group. IFGs with measures not within tolerance ranges of these midaverages are discarded; these tolerances are defined in the reference data set FEX_CTH.TXT (Section 9.1). Specifically, bolometer temperature tolerances are set at the $\frac{1}{2}\%$ level so that the dispersion in the bias voltage (which is highly correlated to the bolometer temperature) will be less than 1%, while GRT tolerances are chosen to give calibrated flux changes of less than 2% in the worst case.

If in the course of performing these checks the number of remaining good IFGs (including neighbors) falls to less than three, then the one or two remaining IFGs are considered to be "too few to coadd", and are marked as bad.

Following the quality and instrument consistency checks, if the number of remaining good IFGs including neighbors is less than or equal to eight, then all of them are used in subsequent processing. If the number of good IFGs excluding neighbors is greater than or equal to eight, then no neighbor IFGs are needed and all of them are marked as bad. If, on the other hand, the number of good IFGs including neighbors is greater than eight, then enough neighbor IFGs are marked as bad to bring the total down to eight. The question of which neighbors to mark bad is decided as follows. FIL calculates the mean Galactic latitude of the good IFGs excluding neighbors. If this mean is less than 10.0, then the neighbor IFGs with the greatest absolute values of the difference of their Galactic latitudes and the mean are marked bad. If the mean is greater than or equal to 10.0, then the neighbor IFGs with the highest glitch rates are marked bad.

In order to normalize data accumulated under differing operating conditions, each IFG is next divided by the real-valued commanded preamplifier gain and the number of onboard sweeps of the MTM, since the digitized voltages are proportional to these quantities. FIL also removes the "dither", a randomly generated value, different for each IFG, which was added to each point of the IFG during its collection. It does so by finding the median value of each IFG and subtracting it.

A robust mean is then constructed as a template by calculating the pointwise midaverage of the coadd group. This mean is the average of the middle two quartiles of the coadd group values at each point. This so-called "primary template" is subtracted from each IFG.

To remove variable signals from gas and dust in the Galactic plane, a "secondary template" is then calculated from the coadd group of interferograms. The IFGs are temporarily aligned so that the signal at the zero path difference is positive in each case; after this is done the secondary template is calculated in the same way as was the primary template.

To ensure that secondary templates are removed only when necessary, secondary template subtraction occurs only if the amplitudes and signal-to-noise ratios of both the primary and secondary templates exceed reference thresholds as specified in the reference data set FEX_CTH.TXT. In that case, suitable multiples of the secondary template are subtracted from each IFG so that the total of the absolute values of the differences between the individually scaled secondary template and the IFGs is at a minimum for 21 points centered at the zero path difference. These secondary template subtraction cutoffs were deemed necessary because this process was seen to influence performance of the deglitching algorithm (see below); the deglitcher removed too few or too many supposed glitches near the peak of the IFG in some cases if the secondary template was subtracted.

A transient response from the onboard digital filters is next subtracted from the first 128 points of each IFG using a least squares fit. The transient is very small after the first few points.

Next, a pattern recognition algorithm is applied to remove contamination due to glitches on the detectors, which are the dominant noise source. For each IFG, the deglitching algorithm uses a robust noise estimator which is the maximum of 1.25 times the median absolute value of the IFG and the threshold bit noise. In an iterative process, the largest value of the 512 points of the template-subtracted IFG is found. If the absolute value of the ratio of this value to the IFG noise is greater than 3.7, a glitch is deemed to exist at that point; a three-point parabolic fit is used to determine the "true" peak position of the glitch, and the glitch profile whose peak is closest to the glitch peak is subtracted. (See Section 9.2 for a description of the glitch profiles). The height of the profile subtracted is scaled by .2 if the ratio above is greater than or equal to 5.5, and by .7 if the ratio is less than 5.5 but greater than or equal to 3.7. The process repeats until all glitches are removed (that is, no point has a ratio greater than 3.7). This technique is based on the *CLEAN* algorithm and has been described in detail by Isaacman, Read, and Barnes (1992), which is included in this document as Appendix F.

Following deglitching, "shape consistency checking" is used to reject anomalous IFGs whose overall shape or noise properties differ significantly from those of the coadd group as a whole. The noise (sigma) of each IFG is defined to be 1.25 times the median of the absolute values of the deviations from its median, and the median sigma of the coadd group is defined to be the larger of the median of these numbers and a value denoting a deviation of one bit. The ratio of each IFG's sigma to that of the ensemble is calculated, and IFGs with ratios above or below values of 1.5 and 0.5 respectively are considered to have "high noise" or "low noise", and are marked as bad. For each remaining IFG, the ratio of its value at each point to the median sigma is then calculated; it is marked as bad if more than six points have a ratio greater than six.

Table 4.7 gives the numbers of calibration and sky IFG records and coadd groups failed for various reasons, as well as the number of IFG records and coadd groups (which became coadded IFGs) that passed.

FIL now coadds the science and engineering data using a weighting based on the glitch rates of the IFGs being coadded. Even after being deglitched, the high glitch rate data have higher noise than the low glitch rate data. This is not surprising since the glitches are not perfectly determined, the glitch profiles are not perfect, and there are a large number of glitches that are too small to be detected by the deglitching algorithm but which still contribute to the overall noise.

The glitches are essentially delta functions of energy incident on the bolometers, and so add power at all frequencies. Because of the filtering of the data and the response of the bolometers, more of the glitch noise shows up at low frequencies than at high frequencies. The arrival times of the cosmic rays are not correlated with the phase of the *FIRAS* scans, and so do not bias the data in any way. The glitches add only noise. Furthermore, because the *FIRAS* pipeline handles calibration data the same way that it handles sky data, any bias introduced to the data by the deglitching algorithm occurs in both the sky and calibration data, and consequently is removed during calibration.

Since the high glitch rate data have a higher noise, these data are deweighted with respect to the low glitch rate data. Once the coadded IFGs were converted to spectra and calibrated, an average variance per frequency sample for each spectrum could be computed:

$$\langle var(i)\rangle_\nu \;=\; \frac{nifgs(i)}{nfreq}\sum_{\nu=0}^{nfreq}\left(\frac{variance(\nu,i)}{D(\nu)^2}\right) \tag{2}$$

where $nifgs(i)$ is the number of IFGs in spectrum i and $nfreq$ is the number of frequency points in the spectrum. The D is from a previous calibration (Section 7.1.1). These

Table 4.7: Records failed and passed in the *FIRAS* Facility FIL.

Channel	LH	LL	RH	RL	All
Bolometer Temperature	2383	2864	2004	2294	9545
Bolometer Voltage	914	477	528	295	2214
Noise Too High	151	468	468	612	1699
Mismatched Science Mode	135	130	147	149	561
All Other Failure Reasons	26	139	160	177	502
Cal Records Failed by FIL	3609	4078	3307	3527	14521
Cal Coadds Failed by FIL	12	11	9	8	40
Cal Records Coadded by FIL	33832	32856	34244	33807	134739
Cal Coadds Produced by FIL	759	751	758	760	3028
Bolometer Temperature	33943	40474	36230	39722	150369
Sky Horn Temperature	12205	12856	12488	12851	50400
Bolometer Voltage	17463	8231	10769	4643	41106
Reference Horn Temperature	1887	2001	1990	1944	7822
Too Few to Coadd	6120	6274	6122	5893	24409
Temp Cont Gain Status Bits	3120	3002	3244	3149	12515
Noise Too High	1608	3127	3873	5479	14087
All Other Failure Reasons	347	412	867	629	2255
Sky Records Failed by FIL	76693	76377	75583	74310	302963
Sky Coadds Failed by FIL	8833	9223	9034	8658	35748
Sky Records Coadded by FIL	305264	288476	305841	291039	1190620
Sky Coadds Produced by FIL	69744	68339	69513	68934	276530

average variances were then modeled as a linear function of the glitch rate:

$$\langle var(i) \rangle_\nu \;=\; slope \cdot glitch\_rate(i) \,+\, intercept \tag{3}$$

The results of this modeling are given in Table 4.8.

Table 4.8: Dependence of Variance on Glitch Rate

Scan Mode	Slope	Error	Intercept	Error
LHSS	1.5191	0.0954	0.8917	0.0075
LHSF	0.7267	0.1775	0.9526	0.0137
LHLF	0.2083	0.0686	0.9840	0.0054
LLSS	0.9034	0.0166	0.6037	0.0080
LLSF	1.1911	0.0595	0.6090	0.0207
LLLF	0.7500	0.0194	0.6825	0.0088
RHSS	0.3181	0.0151	0.8078	0.0096
RHSF	0.2141	0.0304	0.8748	0.0190
RHLF	0.0967	0.0123	0.9389	0.0081
RLSS	1.4353	0.0259	0.4982	0.0096
RLSF	0.8577	0.0493	0.7115	0.0181
RLLF	0.5659	0.0246	0.8027	0.0088

The slope and intercept from this fit was used to renormalize the number of IFGs in the spectra:

$$weight(i) \;=\; \frac{nifgs(i)}{slope \cdot glitch\_rate(i) \,+\, intercept} \tag{4}$$

The net effect of this renormalization is to redistribute the weight within and to decrease the total variance in the data set. This occurs because the decrease in variance of the low glitch rate data more than offsets the increase in variance of the high glitch rate data. The total decrease in variance is $\sim$2%.

The good IFGs in the coadd group are coadded by: determining the weight for each IFG; calculating the sum over the coadd group of each IFG multiplied by its weight; and dividing by the sum of the weights. The primary template and, if appropriate, a multiple of the secondary template formed by averaging the individual factors used in secondary template subtraction are added back to this result to form the "coadd".

The engineering and attitude data associated with the IFGs are coadded in the same manner as the IFG itself. Each coadd is thus accompanied by mean instrument voltages, temperatures, etc., and a mean sky position. Note that this quantity need not be at the nominal pixel center.

Because there was relatively little low frequency data in the SF scan mode taken during the mission, it proved impossible to create a calibration model for it alone (Section 5). A scheme was devised to combine "decimated" calibration coadds for low frequency SF data with "truncated" ones for low frequency LF data to permit the computation of a calibration model that could be applied to both. At this point in its processing, FIL creates both calibration and sky coadds for these two new modes, which are given the designations "FS" for decimated SF and "FL" for truncated LF. The truncated FL coadds are simply the first 128 points of the 512 point LF coadds. The decimated FS coadds are generated using the following algorithm: discard the first three points of the SF coadd; for points 1 to 127 of the FS coadd, take the next four points of the SF, add them together, and divide by four; point 128 of the FS is set to point 512 of the SF.

To remove the effects of internal defocussing of the instrument at the extrema of the MTM sweep, a fourth-order polynomial baseline is fitted to the coadd using a least squares fit and subtracted (see Section 9.2 for details of the baseline function). Last, the transient response from the onboard digital filters is subtracted from the first 128 points of the coadd using a least squares fit (as was previously done for each of the individual IFGs.)

Because so little LS scan mode data was taken during the mission, no calibration model exists for it. Thus, of the original 16 possible channel/scan mode combinations and of the four new low channel/FS or FL scan modes, FIL creates data sets for LHSS, LHSF, LHLF, LLSS, LLFS, LLFL, LLLF, RHSS, RHSF, RHLF, RLSS, RLFS, RLFL, and RLLF. The project data release includes the time-ordered coadded calibration interferograms and the pixel-ordered coadded sky interferograms for each of these channels and scan modes.

The FITS binary and extension headers for these data products are given in Appendix G; they include the names and descriptions of each of the data fields of these data sets.

5. Calibration of *FIRAS* Data

In the initial part of the calibration processing, calibration data and sky data are treated essentially identically. The coadded IFGs are apodized and padded, and then Fourier transformed to produce complex-valued spectra. The spectra are converted from counts to volts at the detector upon division by the analog-to-digital converter scale factor and by the electronics transfer function. The processing of the calibration data and the sky data then diverges. The calibration coadds (and their associated covariance matrices) are used as input to the calibration model program (FISH) to derive a solution, and both the calibration and the sky coadds are calibrated with the resulting solution. The initial apodization and padding of the calibration data are done by the FFL facility; the final apodization, padding and calibration for both the calibration and sky data is done by the FSL facility.

5.1. Apodization and Padding

The coadded interferograms (IFGs) are apodized before Fourier transformation into spectra. The apodization function $A(x)$ is an asymmetric smooth window function which enhances the signal-to-noise ratio for wide spectral features and reduces ringing from spectral lines. $A(x)$ is designed to have a value of one at the IFG peak, to correctly weight the long and short sides of the IFG, and to fall off smoothly to zero at the ends.

The piecewise expression for the apodization function $A(x_i)$ for the 512 sampled data points in the IFGs is:

$$A_i = f_i \left[1 - \left(\frac{i-c}{j-c}\right)^4\right]^2 \tag{5}$$

where $i \in [1, 512]$ is the IFG sample number, c is the IFG peak position given in Table 5.1. For the low spectral resolution data $j = 1$ and

$$f_i = \begin{cases} 0 & i \in [1, 2] \\ 2 & i \in [3, 2c-513] \\ \left\{3 - \cos\left(\frac{\pi(2c-482-i)}{30}\right)\right\}/2 & i \in [2c-512, 2c-483] \\ 1 & i \in [2c-482, 482] \\ \left\{1 - \cos\left(\frac{\pi(513-i)}{30}\right)\right\}]/2 & i \in [483, 512] \end{cases} \tag{6}$$

For the high spectral resolution data $j = 513$ and

$$f_i = \left\{ \begin{array}{ll} 0 & i \in [1,2] \\ \left\{1 - \cos\left(\frac{\pi(i-2)}{30}\right)\right\}/2 & i \in [3,32] \\ 1 & i \in [33, 2c-32] \\ \left\{3 - \cos\left(\frac{\pi(i+32-2c)}{30}\right)\right\}/2 & i \in [2c-31, 2c-2] \\ 2 & i \in [2c-1, 512] \end{array} \right\} \tag{7}$$

The Fourier transform of $A(x)$ gives the instrumental spectral window function before the calibration gain correction is applied and is the dominant factor in determining the instrument line profile and the frequency-to-frequency noise correlation. The full width half maximum of the Fourier transform of $A(x)$ is 0.8191 cm^{-1} (24.56 GHz) for the low resolution spectra and 0.2331 cm^{-1} (6.988 GHz) for the high resolution spectra. See Figure 5.1 for the low and high resolution apodization functions.

Table 5.1: *FIRAS* IFG Peak Positions

Channel and Scan Mode	Peak Position Sample	Resolution (cm^{-1})
LHSS RHSS	357	0.8191
LHSF RHSF	359	0.8191
LHLF RHLF	355	0.8191
LLSS RLSS	360	0.8191
LLFS RLFS	90	0.8191
LLFL RLFL	90	0.8191
LLLF RLLF	90	0.2331

In the *FIRAS* instrument the zero path difference is at bin ~ 355, so there are 355 independent frequencies available. There are only 512 points, so 256 complex points in the transform provide a complete description. Ideally, the response of the instrument is symmetric with respect to the zero path difference position of the MTM. Therefore, the antisymmetric part of the IFG is an indication of model errors, instrument imperfections, and unresolved spectral features. The antisymmetric part of the IFG is transformed to the imaginary part of the spectrum. The short side of the interferogram is padded by appending zeros. This padding is asymmetric, so the loss of degrees of freedom is concentrated in the imaginary part as well. The information about the sky is concentrated in the real part of the spectrum, but the imaginary part still has information relevant to the instrument performance. Therefore, the full spectra are calibrated but only the real part is subsequently analyzed.

If the imaginary part of the spectrum is ignored, only the real part of the covariance matrix need be considered. Padding the interferogram beyond 710 (2×355) points forces the real part of the matrix to become singular. The process is not abrupt: as the padding is extended the condition number of the covariance matrix rises until numerical problems in finding the inverse actually obscure the point of formal singularity. We chose to pad the 512 point interferograms with 128 zeros as a compromise between getting all of the sky information in the real part of the spectrum and retaining a well-behaved covariance matrix. This results in 320 spectral points, allowing a direct comparison to previous results in 1/5 of the new points (1/4 of the previous points). The full covariance matrix is singular, but the covariance matrix of the real part of the spectrum is easily inverted.

The complex spectra are retained until after the calibration so that all of the instrumental effects can be modeled and removed. The uncertainties used in the calibration modeling use only the diagonal parts of the covariance matrix.

5.2. Channel Specific Items

Nominally each detector-scan mode combination is calibrated separately. There are two path lengths and two speeds of the MTM, resulting in four possible modes of operation (SF, SS, LF, and LS). There are four detectors (LH, LL, RH, RL) for 16 separate calibrations. However, the LS mode was used for only a few days in orbit, so there was little point in trying to use these data. About half of the data were in SS mode and this mode has the best calibration data. The LF mode for the low frequencies has high resolution. However, the Nyquist frequency is 36 cm^{-1}, so for the high frequencies (LH and RH) the LF data are truncated on board the *COBE* spacecraft rather than averaged; thus, their format is like that of the SF data.

The high frequency LF data look like the high frequency SF data with two important differences. First, the SF IFGs are averages of 16 sweeps of the MTM while the LF data are averages of only four sweeps, so each SF IFG receives approximately four times the weight of an LF IFG in both the calibration and the sky data. Second, differing numbers of A$\rightarrow$D conversions are discarded for the LF and SF data, so that there is a 3.5 bin offset between the two data sets. An offset of three bins is removed by simple rotation, and the extra half bin of offset is removed by a phase adjustment after the Fourier transform is applied. Thus, the LHSF and LHLF data sets are combined into a single LHFA data set and the RHSF and RHLF data sets are combined into RHFA.

Only the low frequency LF mode has the higher resolution; to preserve it, there are two separate calibrations performed on the low LF data. The first calibrates the LF data with high resolution, while the second combines the SF and LF data. To calibrate the SF data, each coadd is added in groups of four to mimic the processing done on the LF data on board the *COBE* spacecraft. This reduces the Nyquist frequency to 36 cm^{-1} and the length of the coadd to 128. These "decimated" SF coadds are given the scan mode FS. The LF coadds are then truncated to 128 points which reduces the resolution to that of the other data; these are given the scan mode FL. The apodization and padding for the combined low FS and FL coadds must also be reduced by a factor of four. This is accomplished by averaging the apodization function in groups of four (Section 4.4). The LLFS and LLFL data are combined into LLFA, and the RLFS and RLFL into RLFA.

Thus there are 10 separate calibrations: LHSS, LHFA, RHSS, RHFA, LLSS, LLLF, LLFA, RLSS, RLLF and RLFA. The LL detector performed much better than the RL detector and the RH detector performed much better than the LH detector. Since most of the calibration was done in SS mode, the LLSS and RHSS provide much of the information. However, the other modes provide additional information and instrument checks. The high resolution xLLF ($x \in [L, R]$) and low resolution xLFA are not independent, as both include a portion of the xLLF data, ('a portion' because the xLFA data are truncated).

5.3. Instrument Calibration

Each calibration model consists of several thousand parameters (the exact number depends on the channel). The parameters are grouped into a set of about 20 parameters describing the bolometers and general instrument; a set of $12 \times n$ emissivities, where n is the number of frequencies; and a set of $5 \times m$ temperatures, where m is the number of coadds in the calibration. The parameters are adjusted by a simultaneous least squares fit to the calibration coadds in the program FISH. The calibration model is summarized here and discussed in Appendix C. Each set of parameters is uncertain, and the uncertainties can be propagated to the final sky spectra using the Jacobian ($\partial S/\partial p$). The uncertainties for the different sets are given separate names for convenience: JCJ, for the bolometer and general parameters; P_0EP_0, for the emissivities; and PUP, for the temperature parameters (Section 7). The *FIRAS* calibration model is composed of a nine-parameter model of the bolometer response function and a linear photometric model of the instrument throughput.

5.3.1. *Electronics Transfer Function*

The first step in the calibration process is to undo the electronics to get the potential at the bolometer rather than at the A→D converter. The electronics transfer function (ETF) is an analytic model of the instrument's post-detector analog and digital filters, including a 6-pole DC-blocking filter, a 5-pole Bessel filter, a treble boost, a digital lowpass filter, and smoothing due to data compression. This is included mainly as a convenience. Any errors in the ETF are compensated for by the optical transfer function (Section 5.3.3). Figure 5.1 gives the electronics transfer functions for LLSS, RHSS, and LLLF. The ETFs for the other channels are similar.

5.3.2. *Bolometer Parameters*

The bolometer resistance is modeled as:

$$R(T, E) = R_\circ \exp(T_\circ/T)^{1/2}(x/\sinh(x)) \qquad x = E\rho/T, \tag{8}$$

where $R_\circ$ gives the detector resistance at infinite temperature, $T_\circ$ is a characteristic temperature, T is the temperature of the bolometer, x is a dimensionless variable, E is the electrical potential across the detector, and ρ is a constant involving the electron charge, characteristic hopping length, and length of the detector. The bolometer temperature is a balance between the radiation heating it and the conduction to the housing cooling it. The cooling power is modeled by $P = G(T^{\beta+1} - T_c^{\beta+1})$ where G is a constant describing the geometry of the detector mounting and β is the index of thermal conduction for the material. The heat capacity of the bolometer is modeled as $H = C_1T + C_3T^3$ (Mather 1982, 1984a, 1984b). Although the parameters $(R_\circ, T_\circ, G, \beta, \rho, C_1, C_3)$ were measured for each bolometer before launch the calibrations refit $R_\circ, G, \rho, C_1$ with the orbital data. There is not enough temperature variation in the orbital data to fit $T_\circ$ and β or to simultaneously fit C_1 and C_3.

The bolometers were closely coupled to a JFET acting as an amplifier inside the dewar. The offset J_o and gain J_g of the JFET are also parameters of the model. The offset was fit along with everything else, but the gain was determined an analysis of the temporal response in the data immediately after a change in the bias voltage, using in orbit data

The bolometer model uses the linear small-signal approximation of the non-linear bolometer model to compute the bolometer's operating temperature, time constant, and DC responsivity from its bath temperature, commanded bias, and readout voltage. The bolometer response function, B, is computed from the detector's time constant (τ) and DC

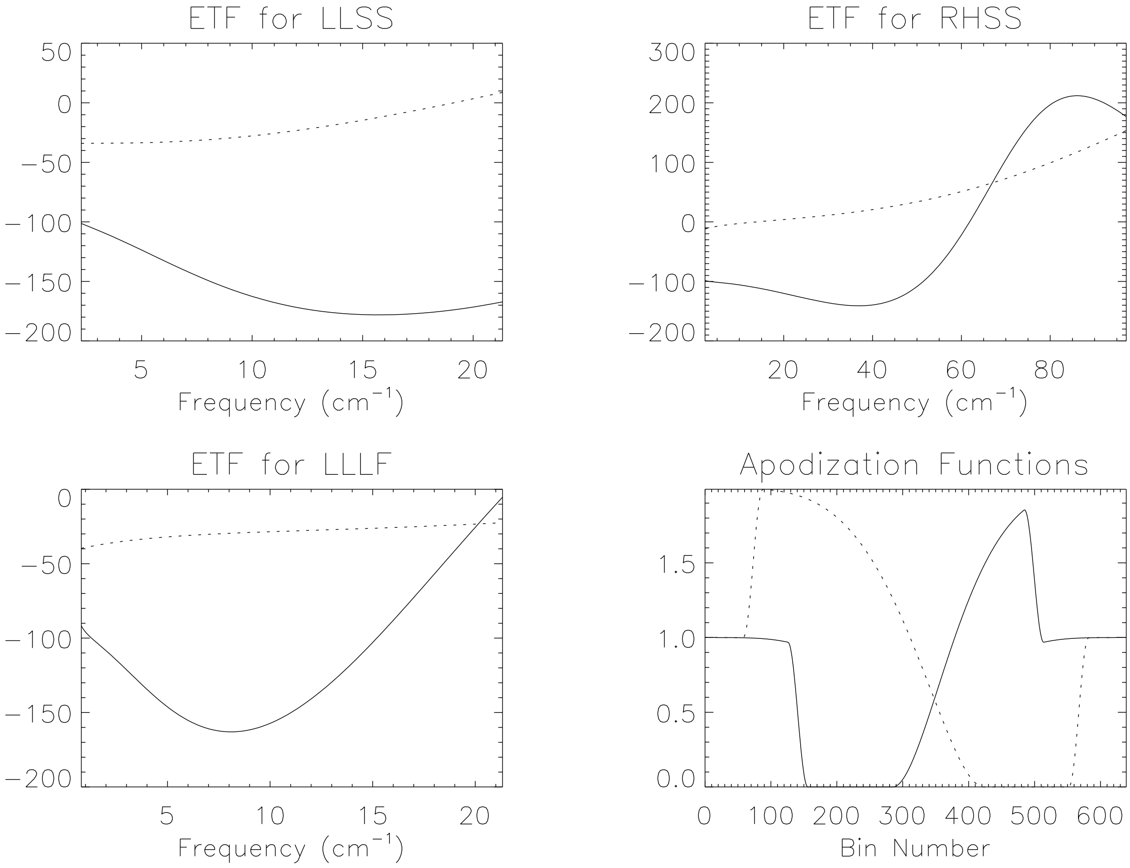

Fig. 5.1.— Electronics transfer functions for LLSS, RHSS and LLLF and apodization functions for low and high resolution data — For each of the ETF plots, the solid line is the real part and the dotted line is the imaginary part. For the apodization functions, the solid line is the low resolution (RHSS; the other low resolution apodization functions are nearly the same), and the dotted is the high resolution (LLLF and RLLF).

responsivity ($S_\circ$):

$$B(\omega) = \frac{S_\circ}{1 + i\omega\tau} \tag{9}$$

at audio frequency ω.

There are another ten parameters to model certain faults in the instrument. Three of these model a drift in the ICAL thermometry $T' = T + A\exp(t/\tau_{\mathrm{ICAL}}) + T_{\mathrm{offset}}$. T is the "raw" ICAL temperature, which is converted to T' by adding $T_{\mathrm{offset}} = -3.054$ mK and an exponential decay with amplitude $A = 4.26$ mK and time constant $\tau_{\mathrm{ICAL}} = 104.3$ days. The parameters were chosen after many fits on all of the detectors. It is not entirely clear that the problem was in the ICAL or its thermometers, but the fact that it was seen in all detectors and scan modes and in both the sky and calibration data indicates that it is truly an optical drift. All of the channels used the same parameters to preserve the comparison used in destriping.

Analysis indicates that a reflectance of $\sim 4\%$ as seen in the ICAL can lead to significant second and third harmonic response in the instrument. Hence two parameters were included to model this effect. These might be different for different detectors so they were fit along with the other parameters.

Finally, vibrations of the MTM at 57.57 and 8.32 Hz were detected, which led to ghost spectra around the vibration lines. These were modeled with a time varying amplitude determined by either a constant or a fourth order polynomial, depending on the channel and vibration frequency.

The derivative of the model was found with respect to each of these parameters both for use in the FISH program and to model the uncertainties due to the errors in estimating these parameters. These are combined and orthogonalized to form the JCJ uncertainties, which are discussed in Section 7.3.

5.3.3. Emissivities

The full optical model includes seven frequency-dependent emissions from the external calibrator, the internal calibrator, the sky horn, the reference horn, the dihedral mirrors, the structure, and the bolometers. Each of these is complex, but the total at each frequency is constrained to be zero by the Kirchoff condition. So, there are 12 free parameters at each frequency. Since the external calibrator is assumed to be good, its emission is identified as the transfer function, and the emissivities of the other elements are the ratio of their emission to the external emission. Uncertainties in these emissivities propagate to the sky in the P_0EP_0 uncertainty (Section 7.2).

The amplitudes of the emissions are shown in Figure 5.2; for the phases see Section 7.9.1. The ripples in the transfer function at low frequencies are due to channel reflection from the front and back of the filter. The deep trough at $\sim$20 cm^{-1} occurs because the low filter turns off before the reflected (high) signal turns on. Ripples at $\sim$55 and $\sim$67 cm^{-1} are due to an interference effect between the polarizers. These also appear in the ICAL emissivity.

The horns are well matched at low frequencies but diverge beyond 30 cm^{-1}. The sky horn is much larger than the reference horn: only the length to diameter is matched. The bolometer is kept at $\sim$1.5 K, so its radiation is essentially zero for $\nu > 10$. Similarly, the dihedral radiation is essentially zero for $\nu > 30$, as its temperature is less than 5.5 K. The high frequency bolometer and dihedral emissivities allow mismatch in the other emissivities (*i.e.* stray light) to be accommodated.

5.3.4. Temperature and Phase Corrections

Although the temperatures of all of the important components of the *FIRAS* were measured with a precision of approximately 1 mK, the photometric precision is better than this. The temperatures of the four controllables (XCAL, ICAL, sky horn, and reference horn) were accordingly adjusted to best fit the photometry along with all of the other parameters. These adjustments were constrained with an estimate of the precision of the temperature measurement. The errors are typically a fraction of 1 mK. Although temperature errors propagate to the other parameters in the calibration model, they do not directly propagate to the sky data, as the temperature measurements of the *FIRAS* during the calibration are not used in the sky calibration. The temperature measurements taken during the sky observations are used, however, and these do propagate to the sky data. These are called the *PUP* uncertainty (Section 7.4).

The phase of the interferograms varied by a significant amount. Some of the variation was identified as due to the amplifier having different delay for different gains, but this did not explain the entire effect. Instead a phase was fit, with constraint, for each coadd. This method was also used on the sky data. These phases are included with the temperature adjustments, as they have the same size and same correlations as the other uncertainties (Figure 5.3).

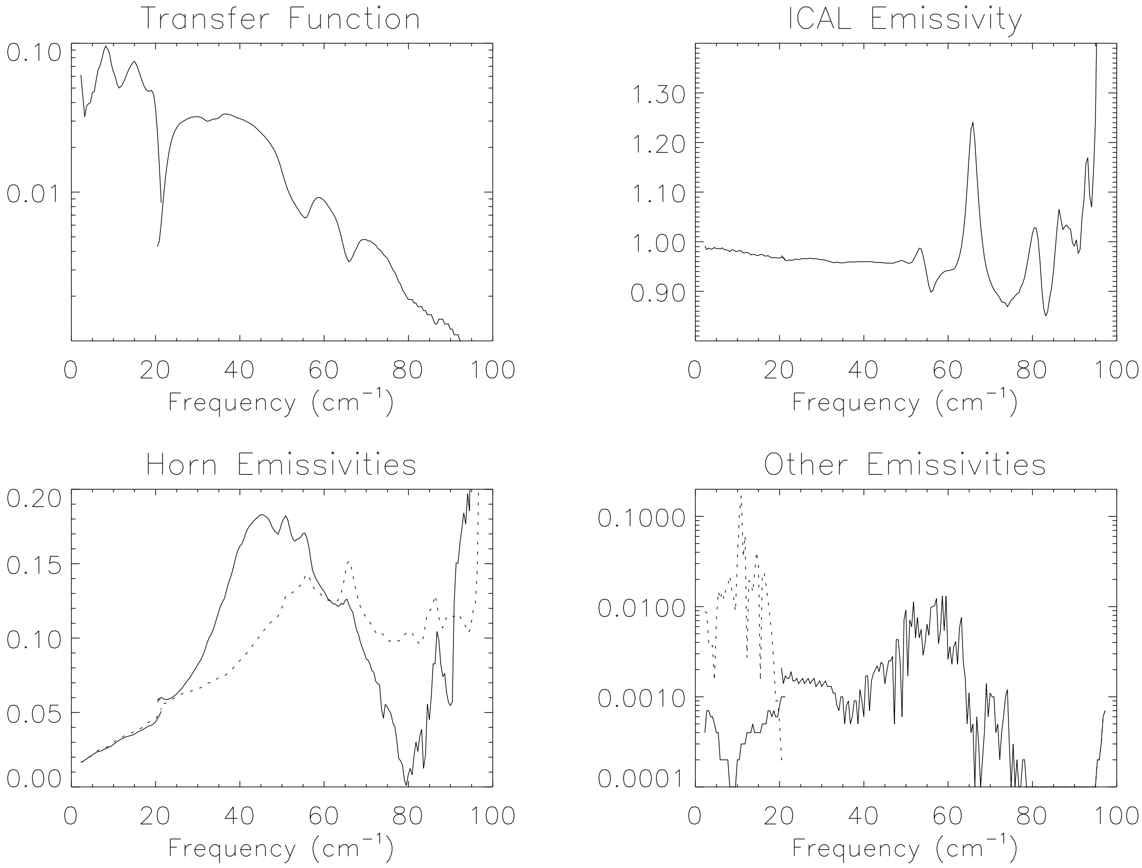

Fig. 5.2.— Optical transfer function and emissivities — The values for LLSS (2–21 cm^{-1}) and RHSS (20–97 cm^{-1}) are shown here. (a) Transfer function (logarithmic scale). (b) ICAL emissivity. (c) The sky (solid curve) and reference horn (dotted curve) emissivities. (d) Dihedral (solid) and bolometer (dotted) emissivities; the bolometer emissivity is only relevant for low frequencies. See text for discussion.

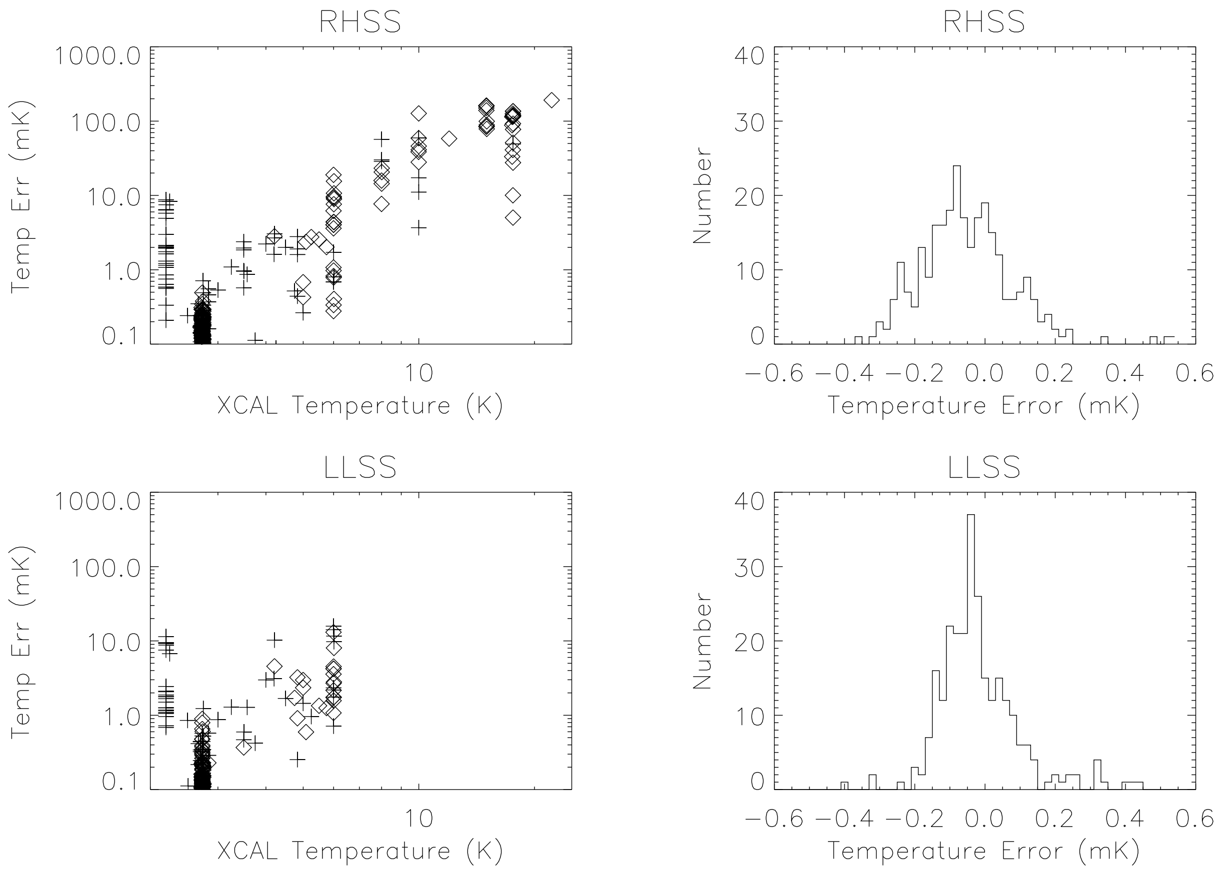

Fig. 5.3.— Photometric XCAL temperature adjustments — The values for RHSS (top plots) and LLSS (bottom plots) calibrations are shown here. The left hand plots are the adjustments (in mK), plotted as a function of XCAL temperature (in K). The symbol + indicates a positive correction, and the symbol ⋄ indicates a negative correction. Temperatures greater than 7K are not used in LLSS. On the right are histograms of temperature adjustments for cold nulls (i.e. all controllables ∼2.7K).

5.4. Temperature Scale Calibration

To first order the errors in the temperatures other than the XCAL are removed by the calibration process, as the same errors are repeated in the sky measurements. This section covers the XCAL thermometry calibration.

5.4.1. Frequency Scale Calibration

The frequency scale calibration of the *FIRAS* was determined from observations of the interstellar emission lines. Analysis of the [C II] and [N II] lines in the high frequency spectra, using the line profiles described in Section 8.4, shows that the frequency calibration is self-consistent to 0.005% at 60 cm^{-1} (1800 GHz). Analysis of the CO lines in the low frequency spectra has shown that the frequency calibration is self-consistent to 0.03% at 15 cm^{-1} (450 GHz). It is possible that there may be non-linear variations in the frequency scale due to beam divergence within the instrument. Thus, even though the [C II] and [N II] lines have better resolution the low frequency scale was used to determine the CMBR cosmic parameters. The frequency scale is set by the CO and [C I] lines observed in that data.

5.4.2. Temperature Calibration

Twenty-four germanium resistance thermometers (GRTs) are mounted throughout the *FIRAS* to measure the temperatures of various components of the instrument. These GRTs were calibrated before flight relative to a NIST-calibrated standard GRT, nominally to 1 mK accuracy. Eight calibration resistors are used to convert the GRT readings from counts to ohms. Readings from the 24 GRTs are combined to yield temperatures of ten components of the instrument: the external and internal calibrators, the sky and reference horns, the four detectors, the mirrors on the mirror transport mechanism, and the instrument structure. The weights for the combination of GRT readings are discussed in Section 9.1.

The temperature scale of the *FIRAS* was calibrated in two ways. In the first method, temperature readings from three GRTs attached to different locations on the external calibrator were averaged to produce a measured temperature of the calibrator. In the second method the calibration model included an additive correction to the measured temperature of the external calibrator. This parameter resulted in a determination of the color temperature scale of the XCAL, which ties the calibrated frequency scale to temperatures through the form of the Planck function. This analysis gives an estimated adjustment to the *FIRAS* thermometer scale of -4.5 ± 1.0 mK. We split the difference to make a 2.25 mK temperature correction (Section 7.5) and claim a systematic precision of 2 mK (one σ).

5.5. Calibration Performance

The validity of the χ^2 calculations depend on the variance that the calibration program assigns to the input data. Determining this for the low temperature null data (*i.e.* those data taken with all four controllables at ~ 2.7K) is straightforward. There are ~ 18000 IFG's taken in this state, so the observed measurement variation determines the variance to $\sim 1\%$. For other temperature combinations there are only a few tens of IFG's in each combination, so the variance using just a particular temperature condition's data can only be determined to $\sim$ 10–20%. For higher temperatures the thermometers don't read as precisely, and the temperature variations (which are fit in the calibration program) can artificially inflate the uncertainty estimates. Also, the gain is turned down, so the true noise can be hidden within a single bit.

To estimate the variance of the higher temperature data, all of the data are used to fit a phenomenological model for the data variance:

$$\sigma_\nu(T) = D_\nu(T_x^2 + T_i^2 + 0.1T_s^2 + 0.1T_r^2)/(2.2 \cdot 2.7^2) \tag{10}$$

which gives approximately the right scale to the variance in the input data. However, the χ^2 for the high temperature data is still high (Table 5.2 and Figure 5.4), particularly for the high frequency data.

Table 5.2 shows the χ^2 per degree of freedom. The “Null” column is for the cold null data, where all controllables are $\sim$2.7K. The “Hot” column is for the data where at least one element is greater than 2.8K. In addition, there are χ^2 and degree of freedom errors introduced when determining the temperature adjustments phase, bolometer parameters, *etc.*, which account for the discrepancy between the “Detector” χ^2 and the “Final” χ^2.

For the low frequency calibrations we exclude all of the data with any of the controllable temperatures above 7K. This reduces the requirements on the calibration model, as the detector radiation loading scales approximately as T^4, and the total power is a key issue in the detector model. For the high frequency detectors we must include the data with temperatures up to 22K in order to get sufficient power at 95 cm^{-1} to calibrate the instrument. This leads to distortions in the high frequency detector calibration models, so they are not as precise as the low frequency models.

Table 5.2: χ^2 per Degree of Freedom

Data set	Real	Imaginary	Final	Detector	Null	Hot
LHSS	1.77	2.58	2.20	1.75	1.39	2.40
LHSF	2.24	2.34	2.37	2.20	1.10	2.93
LLSS	1.53	2.20	1.96	1.49	1.34	1.79
LLSF	1.30	1.66	1.52	1.25	1.29	1.21
LLLF	1.13	1.11	1.13	1.10	1.14	1.05
RHSS	2.10	3.44	2.79	2.08	1.12	3.90
RHSF	2.28	2.42	2.40	2.24	1.06	3.05
RLSS	1.46	1.34	1.53	1.41	1.40	1.44
RLSF	1.18	1.15	1.20	1.14	1.11	1.16
RLLF	1.13	1.12	1.15	1.10	1.10	1.10

See text for details.

In all of the data sets a few coadds still have high χ^2 and these are removed from the data. The reason for the high χ^2 was understood for some of these but for others no apparent reason was found. There were more of them at high temperatures but some were found in the low temperature null data as well. Less than 10% of the data were eliminated in this way.

5.6. Application of Calibration Models

The application of the calibration models to the sky data is summarized here and is described in detail in Appendix C. The input interferogram is defined as a function of optical path difference by the equation:

$$I(x) = \int_{-\infty}^{+\infty} d\nu\, Z(\omega) B(\omega) \sum_{i=1}^{N} E_i(\nu) S_i(\nu) \cos(2\pi\nu x) \qquad (11)$$

where x is the optical path difference, ν is the optical frequency in cm^{-1}, ω is the audio frequency in rad/s, $Z(\omega)$ is the electronics transfer function, $B(\omega)$ is the bolometer response function, (see equation 9) and the sum over i represents the contribution from N sources with emissivity $E_i(\nu)$ and spectrum $S_i(\nu)$. The N inputs are the sky or external calibrator and various instrument components (primarily the internal calibrator).

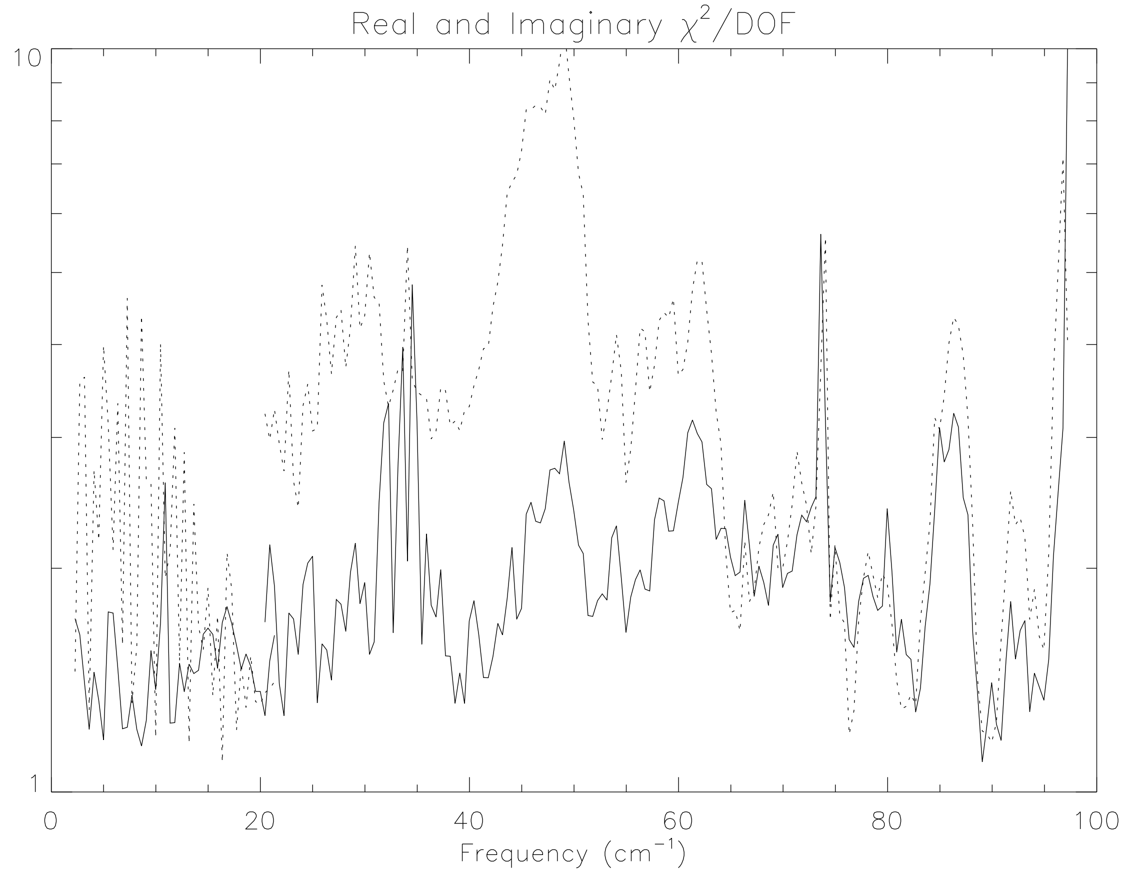

Fig. 5.4.— χ^2/DOF as a function of frequency — The values for the real part of the spectrum (solid line) and the imaginary part (dotted line) for LLSS ($\nu < 22$) and RHSS ($\nu > 21$) are plotted. Note vibration lines at 10.6 and 73.5, and bands at $73.5 - 2 \cdot 10.6 = 52.3$, $73.5 - 10.6 = 62.9$, $73.5 + 10.6 = 84.1$ and $73.5 + 2 \cdot 10.6 = 94.7$. The three-spike feature from 30 to 35 cm^{-1} is a robust artifact from an unknown source.

A coadded IFG is apodized and Fourier transformed into a spectrum in counts:

$$Y(\nu) \;=\; \int_{-\infty}^{+\infty} A(x)\, I(x)\, e^{2\pi i\nu x}\, dx \;=\; a(\nu) * \left(Z(\omega)B(\omega) \sum_{i=1}^{N} E_i(\nu)S_i(\nu) \right) \tag{12}$$

where $*$ is the convolution operator, $A(x)$ is the apodization function discussed in Section 5.1, and $a(\nu)$ is the Fourier transform of the apodization function including padding.

The spectrum in counts is normalized by the electronics transfer function to yield a spectrum in volts. The bolometer response function is computed for the spectrum based on its commanded bias, readout voltage, and bolometer temperature. The bolometer response function and the optical transfer function $H(\nu)$ (from the XCAL emission) are used to convert the spectrum from volts to MJy/sr, yielding a differential calibrated spectrum of the sky or XCAL (a spectrum of the difference in signal between the sky or XCAL and the ICAL):

$$D(\nu) \;=\; e^{i\nu\psi}\, H(\nu)^{-1} \left(\frac{1}{Z(\omega)B(\omega)} Y(\nu) \;-\; C(\nu) \right) \tag{13}$$

The vibration and harmonic corrections $C(\nu)$ are then applied to the differential spectrum. There is a small phase jitter in the *FIRAS* data which is corrected in each differential spectrum by the application of a linear autophase correction ψ which is computed to minimize the imaginary part of the spectrum.

The instrument emissivities are used to compute a spectrum of the internal calibrator using Planck functions $P(T,\nu)$ of the measured temperatures of the instrument components for the source spectra:

$$R(\nu) \;=\; H(\nu)^{-1} \sum_{i=1}^{N-1} E_i(\nu) P(T_i, \nu) \tag{14}$$

where $N-1$ signifies that the source is not included in the spectrum.

The differential spectrum and the internal calibrator spectrum are combined to produce a calibrated spectrum of the sky:

$$S_{sky}(\nu) \;=\; e^{i\nu\psi}\, H(\nu)^{-1} \left(\frac{Y(\nu)}{Z(\omega)B(\omega)} \;-\; C(\nu) \right) \;-\; H(\nu)^{-1} \sum_{i=1}^{N-1} E_i(\nu) P(T_i, \nu) \tag{15}$$

The calibration models are applied to the calibration data using this algorithm to provide the zero-point reference for the post-calibration destriper processing described in Section 6. The sky data are similarly calibrated, and are then Doppler-shifted from the spacecraft frame-of-reference to the solar system barycentric frame-of-reference. The number of spectra in each calibration data set is given in Table 5.3 while the number of spectra in each sky data set is given in Table 5.4.

The project data release includes: the ten calibration models; the time-ordered calibrated and differential calibration spectra for each of the 14 individual channels; and the pixel-ordered calibrated and differential sky spectra for each of the 14 individual channels. The FITS headers for all of these data sets, which are given in Appendix G, include the names and descriptions of each of the data fields.

Table 5.3: Number of Spectra in *FIRAS* Calibration Data

	LHSS	LHSF		LHLF	RHSS	RHSF		RHLF
Spectra	512	70		177	510	71		177
Spectra Used	495	58		176	482	55		176
IFGs	24485	2186		7161	24647	2199		7398
	LLSS	LLFS	LLFL	LLLF	RLSS	RLFS	RLFL	RLLF
Spectra	512	69	170	170	515	71	174	174
Spectra Used	418	28	138	140	423	34	140	142
IFGs	24300	2146	6410	6410	24456	2267	7084	7084

Table 5.4: Number of Spectra in *FIRAS* Skymaps

	LHSS	LHSF		LHLF	RHSS	RHSF		RHLF
Pixels	5967	2092		5962	5972	2089		5954
Spectra	35284	5099		28183	34957	5086		28185
IFGs	130857	36846		133149	130102	36982		134183
	LLSS	LLFS	LLFL	LLLF	RLSS	RLFS	RLFL	RLLF
Pixels	5970	2090	5956	5956	5976	2089	5952	5952
Spectra	34819	5048	27313	27313	34881	5140	27697	27697
IFGs	127098	35989	121406	121406	127047	36931	122924	122924

6. Destriping of *FIRAS* Data

The *FIRAS* destriper, FMD, designed to correct for residual instrumental effects in the calibrated coadd spectra, is the facility which produces the pixel spectra and associated errors. If the calibration model does not sufficiently account for all the "sources" in the *FIRAS* instrument, the propagation of their effects into the calibrated spectra will produce discrepancies between different observations within the same sky pixel. The destriper attempts to model these as a set of offset (subtractive) correction spectra. Unlike the primary calibration program, the destriper uses both sky and calibration data, producing the correction spectra by minimizing the discrepancy between multiple observations within the same sky pixel for the former and deviations from a Planck (blackbody) spectrum for the latter. Also, unlike the primary calibration, the destriper is linear.

6.1. Destriper Model

The destriper minimizes:

$$\chi^2 \;=\; \sum_i \left[S_i^{(sky)} - A_{p(i)} - \sum_k J_k f_{ki} \right]^2 w_i \;+\; \sum_i \left[S_i^{(cal)} - P_i - \sum_k J_k f_{ki} \right]^2 w_i \qquad (16)$$

with respect to A_p, the destriped spectrum for the pth sky pixel and J_k, the offset spectra for the kth correction. These quantities are computed independently for each frequency so we have suppressed the frequency index. $S_i^{(sky)}$ is the ith calibrated sky spectrum, $S_i^{(cal)}$ is the ith calibrated calibration spectrum, P_i is a Planck spectrum at the external calibrator (XCAL) temperature of the ith calibration observation, f_{ki} is the value of the kth model function for the ith sky or calibration observation and w_i is the corresponding effective weight for this observation. The first sum is over all sky observations, excluding those within pixels whose center is located within a mask of the Galactic center. The Galactic exclusion is made to minimize errors due to spatial gradients. The use of spatial gradient correction functions (Section 6.3.3) allows us to use more of the Galactic region in the destriper. Analysis of χ^2 as a function of Galactic coordinates showed that some masking is still needed. The Galactic masks selected are :

LOWF (freq $<$ 630 GHz) : $[b \;<\; 5, l \;<\; 30]$
HIGH (600 GHz $<$ freq) : $[b \;<\; 8, l \;<\; 100]$

The second sum is over the "cold-null" calibration spectra taken with all four major controllables, external calibrator (XCAL), internal calibrator (ICAL), sky horn (SKYH), and reference horn (REFH), set between 2.6 K and 2.8 K. The calibrated coadd spectra S_i

are acquired from the calibrated spectra records. The Planck spectra P_i are computed from the XCAL temperatures of the calibration observations. The A_p are the derived pixel spectra, with the offset corrections applied. To operate the destriper, it is necessary to determine the observation weights w_i and to select the model functions f_{ki} .

6.2. Observation Weights

Ideally, the weight for each observation would be the inverse of its variance, which is calculated from the dispersions of the individual interferograms which make up the observation (Section 7.1.1) . In general, the measured variance of an observation deviates from the variance of the underlying process. The effect of using measured variances overweights those observations whose measured variance is smaller than its long term average variance. The destriper minimization fit then is incorrectly dominated by observations whose measured variances are much smaller than their long term average variances.

The solution is to derive a parametric model for the weights. The simplest model is to use the number of interferograms in the coadd as the weight, adopting the premise that all interferograms are equal. Since we know from theory and from our analysis of the data that the interferogram noise is dependent on the glitch rate, we have already incorporated a glitch rate correction into our model (Section 4.4 and Table 4.8).

Prior to deriving a final model for the weights, it is necessary to identify and eliminate any bad observations which have managed to elude the quality checks applied upstream in the pipeline, since they will corrupt the fit. The destriper is initially run for the 10 individual channels: the low resolution low frequency data sets LLSS, LLFA, RLSS, and RLFA; the high resolution low frequency data sets LLLF and RLLF; and the low resolution high frequency data sets LHSS, LHFA, RHSS, and RHFA. The initial weights are the weights determined at the coaddition stage. For each observation, the χ^2 combined over the entire frequency band is examined. Observations with anomalously high χ^2 per degree of freedom are identified and eliminated by having their weights set to zero.

The weights for each of the 10 individual channels are derived by fitting a parametric model to the measured variances of the surviving observations. Each individual channel is fitted with its own model. Model parameters include glitch rate, bolometer responsivity, ecliptic latitude, and time. Time dependences are modelled by a set of Legendre polynomials. The maximum order of the time polynomials is seven, capable of resolving time scales greater than 10 days. Variations on shorter time scales are modelled by a running filter function of width three days or less. The selection of the optimum parametric model is made as follows. A trial model is applied to derive tentative observation weights. The destriper is run with the trial weights, and the χ^2 combined over the frequency band is computed for each observation from its dispersion from the pixel mean, for sky data, or its dispersion from a Planck spectrum, for calibration data. The distribution of combined χ^2 is fit to a theoretical distribution to determine the RMS scatter of the actual variance from the model variance. We iteratively converge on a model which minimizes the RMS scatter. The optimal model is then applied in production. Using this method, we are able to derive weights with an RMS error of 15% for low channel spectra and 8% for high channel spectra.

Note that we use a dimensionless weight, so that the expression minimized by our fit had dimensions of $(\textit{spectral intensity})^2$. Implicit in the model is a common variance, σ_ν, that applies to all observations and provides a scale for the χ^2. In fact, a variance is computed from the residual variation among multiple observations within the same pixel. The corresponding σ is called the *C-Vector* and is discussed in Section 7.1.

A few more steps are taken to determine the final set of observation weights. First, a more sensitive test for bad observations can be applied by screening out those with combined χ^2 per degree of freedom greater than a specified threshold. The destriper is run again with these observations de-weighted. The relative weights for the individual channel combinations are then derived from their *C-Vectors*. The destriper is run again for the combined data set. A final test for anomalous χ^2 per degree of freedom is made to screen out any remaining bad observations. Finally, the observation weights are scaled so that the total weight is equal to the number of sky pixels, so that the average total weight per pixel is one.

Table 6.1 gives the sky coverage for each data set and weight of an average IFG in each channel. Comparison with Table 5.4 shows the number of IFGs eliminated during the destriping process.

6.3. Destriper Model Functions

We have broken the high frequency data into four frequency sections. The lowest section, (HI1), covers the same frequency range as the LOWF. After careful inspection we decided that the difficulty in correcting for the systematic effects in the low frequency end of the high frequency data exceeded the possible improvements to the LOWF data. The other 3 ranges (HI2, HI3, and HI4) were destriped separately with separate model functions. These were later rejoined to form the HIGH data set.

Table 6.1: Number of IFGs and Relative IFG Weightings in Destriped *FIRAS* Skymaps

	LHSS	LHFA		RHSS	RHFA	
Number of pixels	5962	6002		5969	5999	
Number of IFGs	130784	169875		129595	170917	
IFG Weight*	0.998	0.265		35.201	9.491	
	LLSS	LLFA	LLLF	RLSS	RLFA	RLLF
Number of pixels	5960	5995	5899	5970	5993	5936
IFG Weight*	126591	157118	117944	126791	159461	122168
Obs/1000IFGs	32.250	4.161	38.165	9.952	1.279	11.743

*Weight per 1000 IFGs. The weight has been scaled so that the total weight of the entire dataset is equal to the number of pixels, 6144.

The art of destriping lies in the selection of the model functions f_{ki} . The set of model functions for previous releases of *FIRAS* data was limited to two exponentially decaying time functions and tophats corresponding to the entire mission period and those periods when the temperatures of the sky horn and reference horn were raised to 6K and 4K. The decaying exponentials were used to model an unidentified source of excess emission which was present at the start of the mission. With our improved sensitivity, additional correlates for the residuals were found, allowing us to expand our set of model functions.

6.3.1. Special Time Period (Tophat) Functions

We use the same three tophat functions as in the previous releases of *FIRAS* data. One change has been to restrict their use to the frequency bands where they reduce χ^2 by a significant amount as measured by an F-statistic (Bevington 1969).

LOWF (freq $<$ 630 GHz)	:	[6K,4K,Mission]
HIF2 (600 GHz $<$ freq $<$ 1350 GHz)	:	[6K,Mission]
HIF3 (1350 GHz $<$ freq $<$ 2100 GHz)	:	[Mission]
HIF4 (2100 GHz $<$ freq)	:	[Mission]

Separate tophat corrections are computed for each of the individual channel combinations.

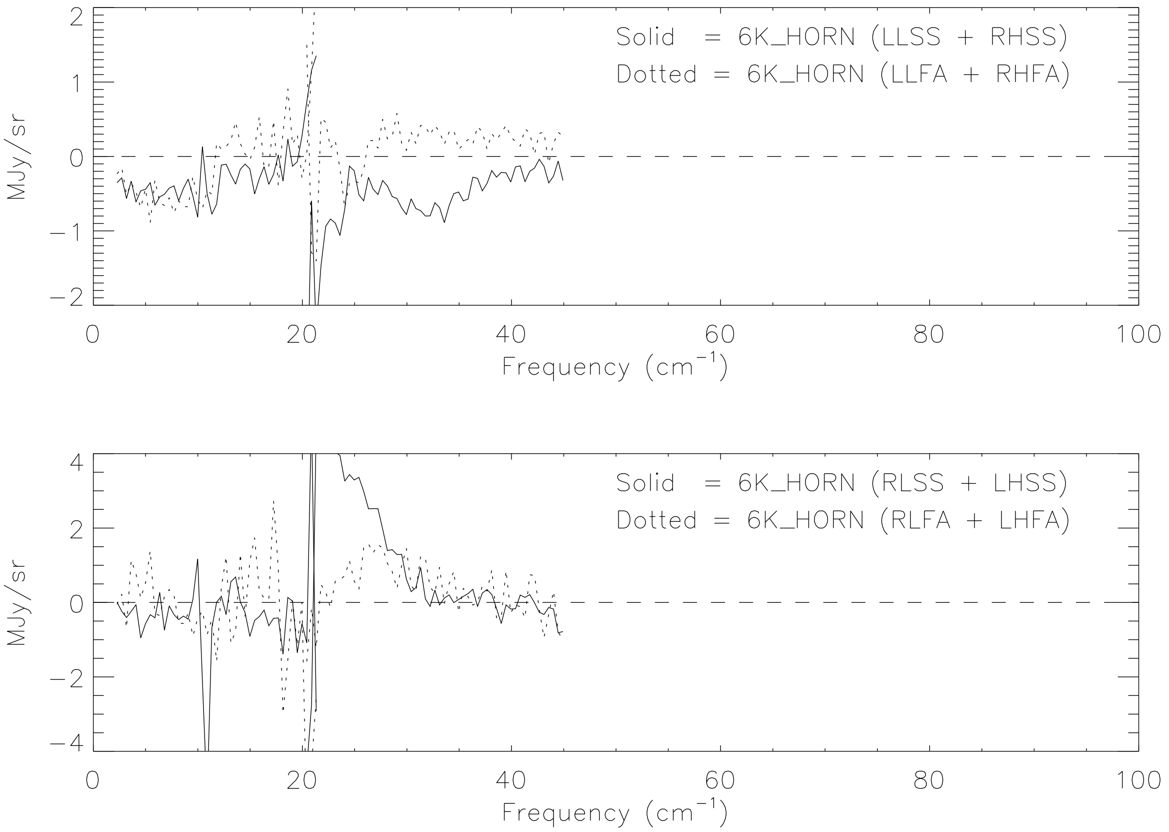

Fig. 6.1.— ‘6K Horn’ offset correction spectra — Applied to the observations for the 6K horn period.

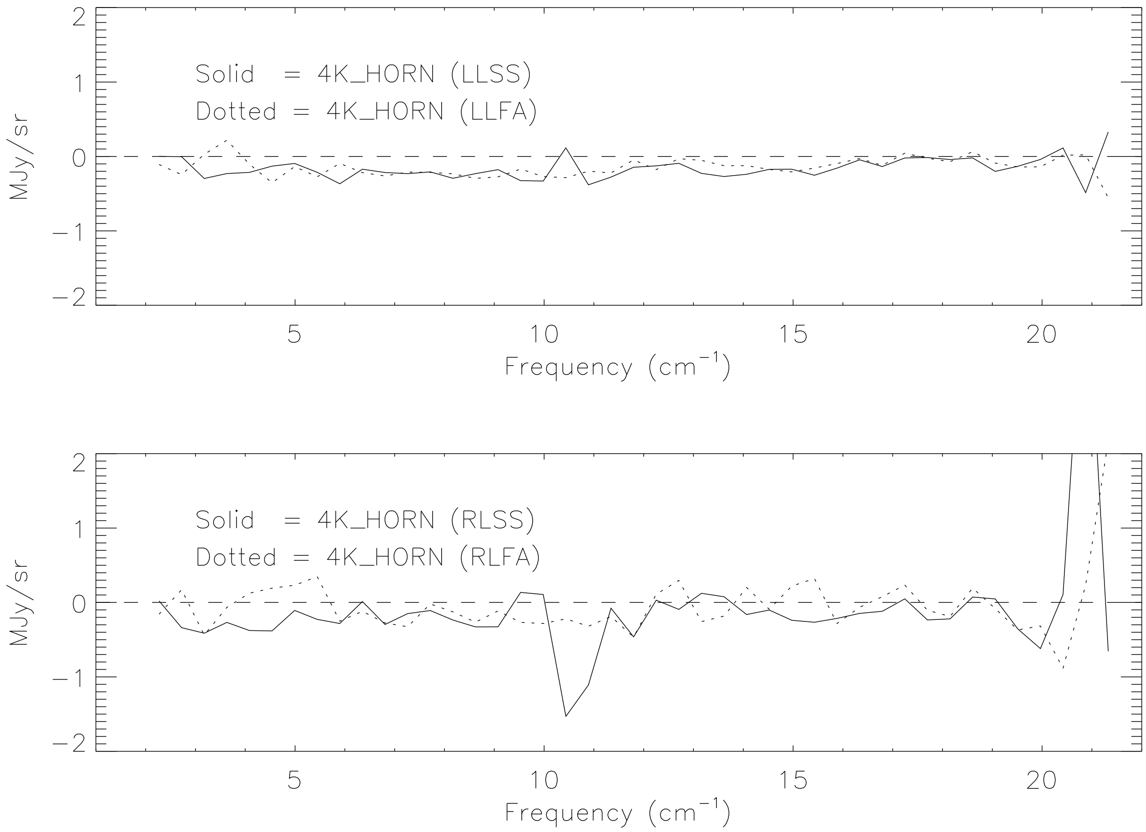

Fig. 6.2.— ‘4K Horn’ offset correction spectra — Applied to the observations for the 4 K horn period.

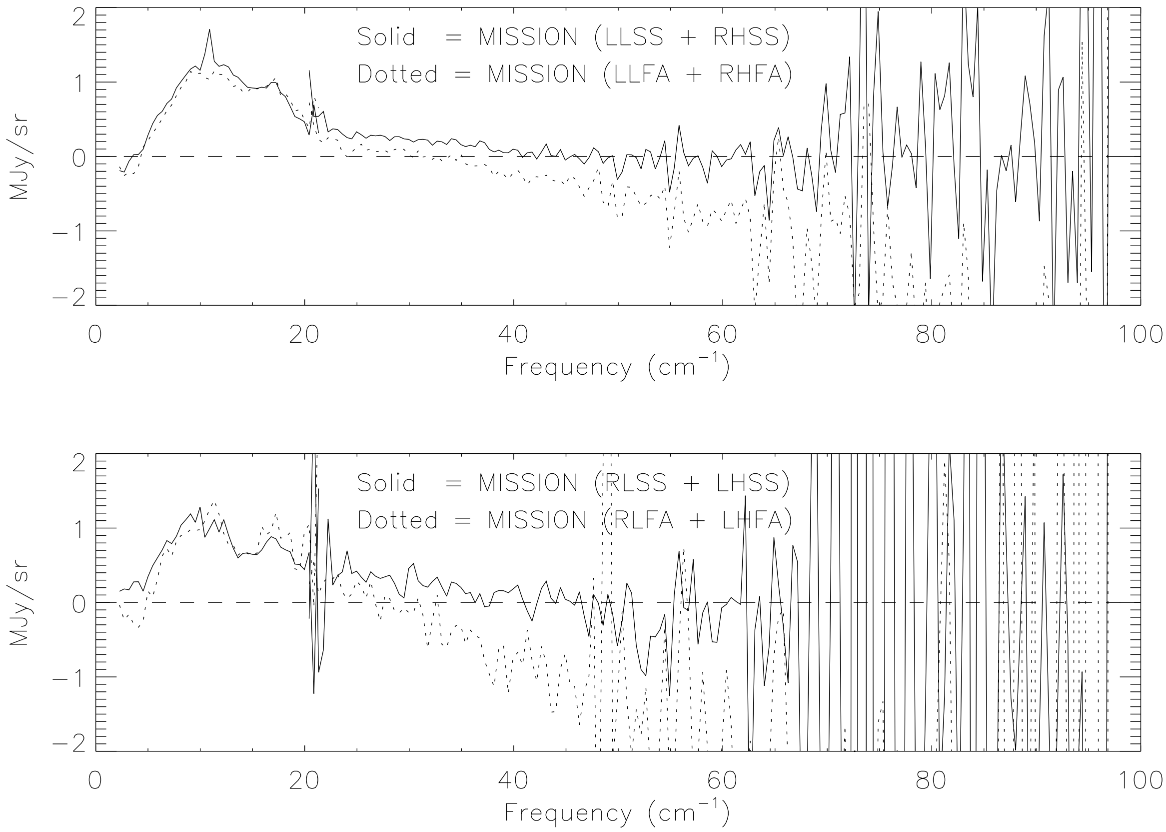

Fig. 6.3.— ‘Mission’ offset correction spectra — Applied to the observations for the entire mission.

6.3.2. *Time Functions*

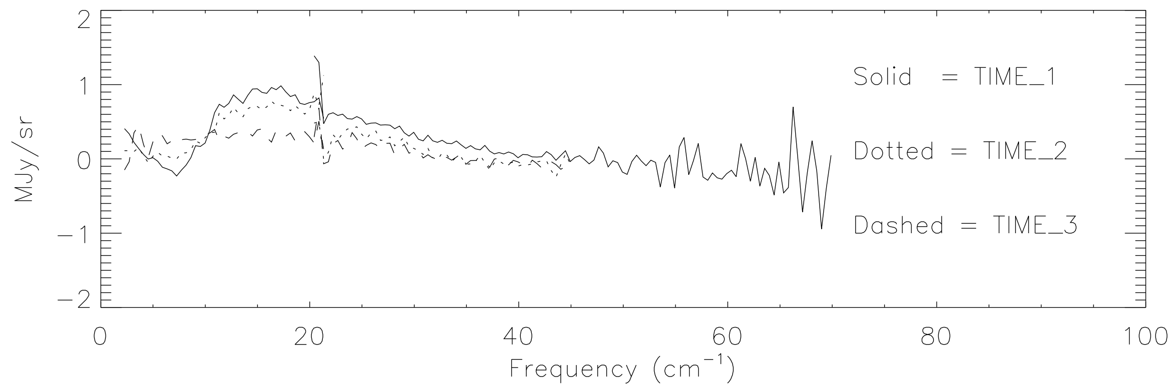

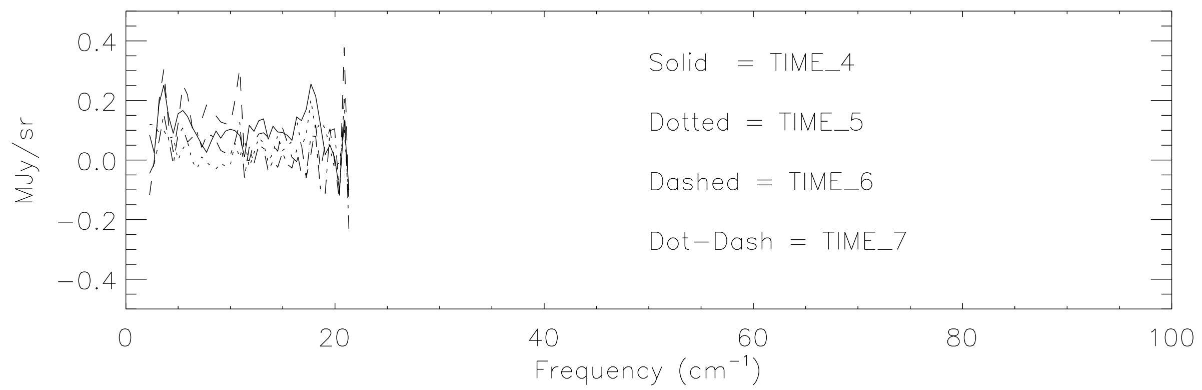

Fig. 6.4.— Time varying correction spectra — Based on Legengre polynomial stripe kernal

---

We have replaced the decaying exponentials of the previous destriper with Legendre polynomials for this destriper. This more general form is preferred for several reasons. The time constant for the decay is frequency-dependent. The previous selection of two time constants represented a best compromise among the frequencies. With our improved sensitivity, this selection was found to be inadequate. As we expanded the number of time constants, the selection of optimum values became cumbersome. Also, the non-orthogonality of the functions resulted in a poorly conditioned matrix inversion. The Legendre polynomials, defined so that they are orthogonal over the mission interval, eliminated inversion problems, eliminated the need to select time constants, and allowed for flexibility in destriping at different frequency bands by selecting an independent set of polynomial orders for each frequency band. The polynomial orders were determined by trial. For each frequency band, increasing orders of polynomials were tried until the improvement in χ^2 became insignificant as measured by an F-statistic. The following sets of time functions were determined :

LOWF : [1,2,3,4,5,6,7] HIF2 : [1,2,3] HIF3 : [1] HIF4 : None

The argument of the Legendre functions is defined to vary linearly with time from +1 at the beginning of the mission to -1 at the end to satisfy orthogonality. The DC term is not included in the set of time functions, because it is degenerate with the mission tophat function, as described above. Unlike the tophat functions, the time functions are derived in common for all of the channel combinations, under the assumption that a common source is being modelled.

6.3.3. Source Gradient Functions

We compare observations grouped into the same sky pixel. Because of the finite sizes of the pixel and the *FIRAS* beam (Section 7.9.4), spatial gradients in the source spectra in the vicinity of the pixel will produce variations which are not intrinsic to the *FIRAS* instrument. Crosstalk between these variations and the model functions will introduce errors to the offset spectra, hereafter referred to as "stripes" . Previously, this source of error was handled by excluding data from pixels within 10° of the Galactic plane, under the assumption that spatial gradients at higher Galactic latitudes were negligibly small. For the final release, we have added model functions to account for the spatial gradients, enabling us to include more data in the destriper, and assuring that gradient errors at higher Galactic latitudes truly are suppressed.

The gradient functions were derived from the *DIRBE* data. First the high resolution *DIRBE* data were convolved with the *FIRAS* beam (Section 7.9.4) centered on each of the 6144 skycube pixels. Then a six parameter quadratic model was used to fit each pixel and its nearest eight neighbors. That model was used to estimate the *DIRBE* intensity for each coadd. Finally the weighted average of the coadd estimates for each pixel were subtracted from the estimates for that pixel. This was done so that there is no net signal in the model. Functions were formed from *DIRBE* bands 8, 9, and 10 (100, 140 and 240 μm) although only the functions for bands 9 and 10 were used in the released destriper models. F-statistic analysis indicated that two *DIRBE* gradient stripes were optimal. We selected the Band 10 and Band 9 functions.

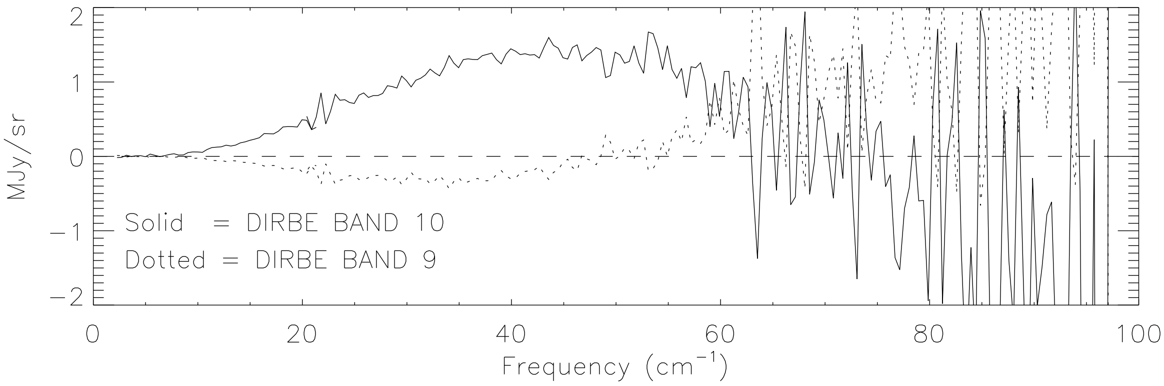

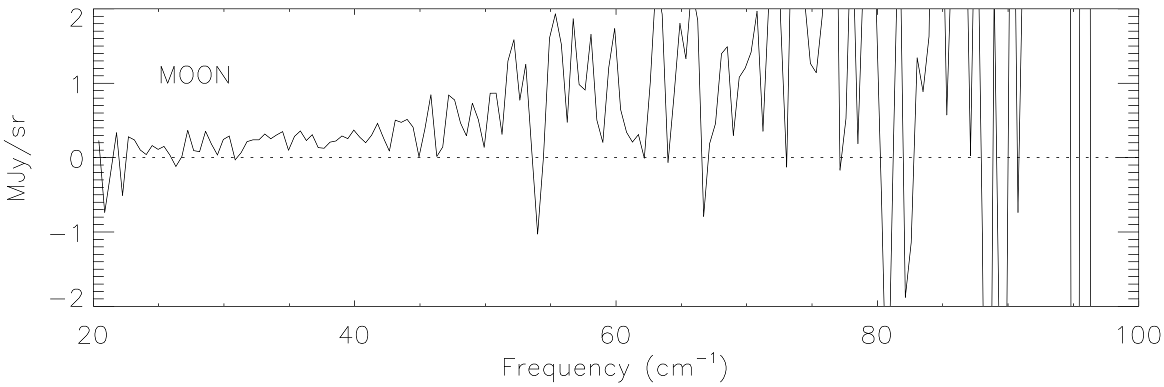

Fig. 6.5.— Sky gradient and moon contamination offset correction spectra — By construction the net contribution of the sky gradient to the pixel spectra is zero.

6.3.4. Moon Contamination

Since the beam profile is not ideal, there is some residual contamination from the Moon, even at angles greater than the 22° data quality cutoff. For the previous data release, we reported 3σ upper limits for the moon contamination. The improved sensitivity for this release has resulted in a detection of moon contamination. We have added a model function to allow the destriper to remove it by an offset correction.

For each sky observation, the moon angles and weights for its component interferograms were recovered, and the fraction of the coadd weight with moon angles within given ranges was computed. The residuals of the coadd spectra from the pixel mean were recovered from the results of the previous destriper runs, and examined as a function of observation weight at the various moon angles. Positive residuals were found for lower moon angles for all high channel frequency bands, indicating the presence of detectable moon contamination at frequencies greater than 600 GHz. The residuals were fit by a second order polynomial of moon angle to produce a simple but adequate model of moon contamination as a function of moon angle. This function was sufficiently similar for all high channel frequency bands that a single model was used for all bands. The moon contamination model function for each observation was then constructed as the weighted mean function of its component interferograms.

6.3.5. Dihedral Functions

Previous releases disqualified interferograms with associated dihedral temperatures greater than 3.5 ° K, eliminating a large amount of data taken during the first four months of the mission, when MTM power dissipation events caused dihedral heating spikes (Section 3). For this release, we have retrieved these data by grouping sky coadds by dihedral temperature. This advance was made possible by the use of interferograms from neighboring pixels in template formation, thereby allowing the formation of coaddable groups for each of six dihedral temperature ranges (Section 4.3) . Trial destriper runs showed that there were significant residuals correlated with dihedral temperature for this expanded set of sky data. The source of these residuals could be calibration errors in the dihedral emissivity and/or other instrumental effects associated with the MTM power dissipation. The residuals appeared to roughly fit a dependence on the square of the dihedral temperature. To correct for effects associated with dihedral heating events, we have constructed a function equal to the square of the dihedral temperature minus the square of the median dihedral temperature for all calibration observations used in the destriper. The use of the dihedral function was also determined by an F-statistic analysis.

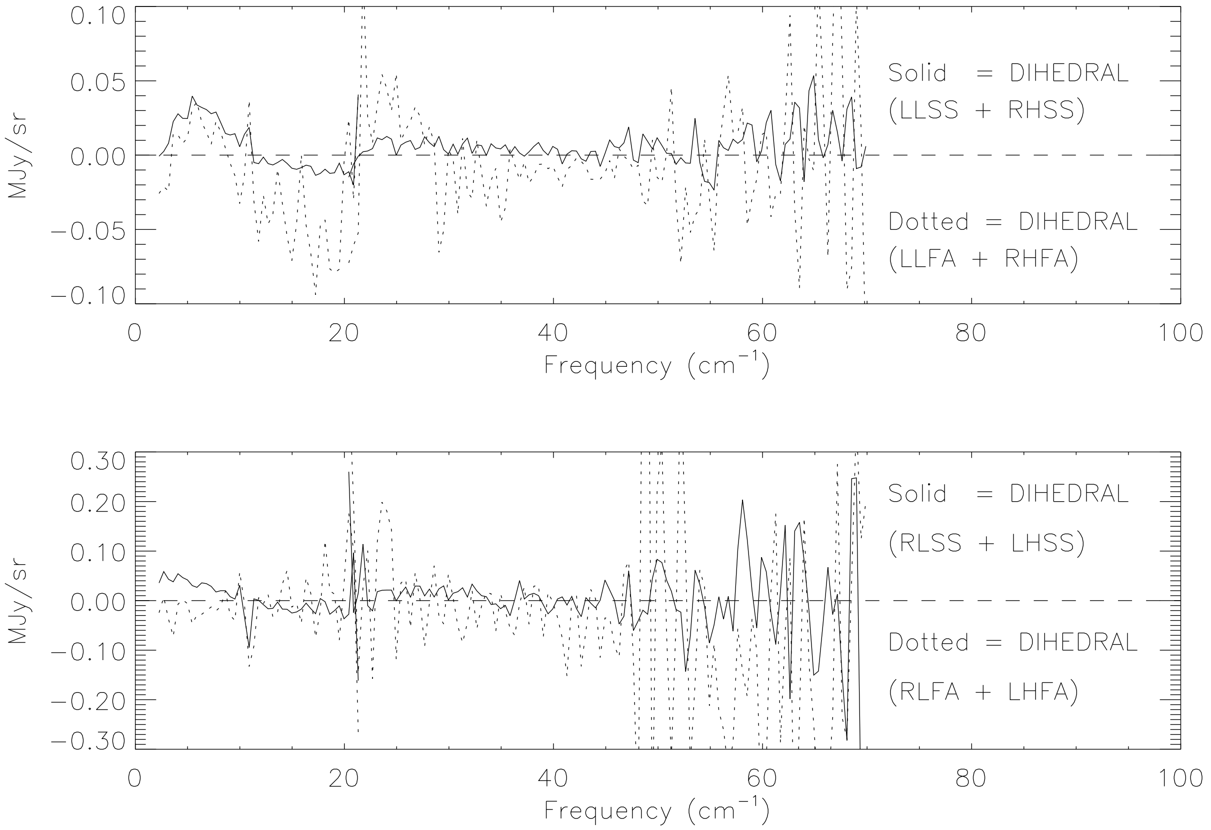

Fig. 6.6.— 'Dihedral' offset correction spectra — Based on a model where the correction is proportional to changes in the square of the dihedral temperature, relative to a nominal temperature.

LOWF : Yes HIF2 : Yes HIF3 : Yes HIF4 : No

The definition of the function means that no dihedral correction is applied for observations with dihedral temperature equal to the reference temperature of 2.063 K. Dihedral effects at this temperature are absorbed into the mission stripes.

6.3.6. Bolometer Functions

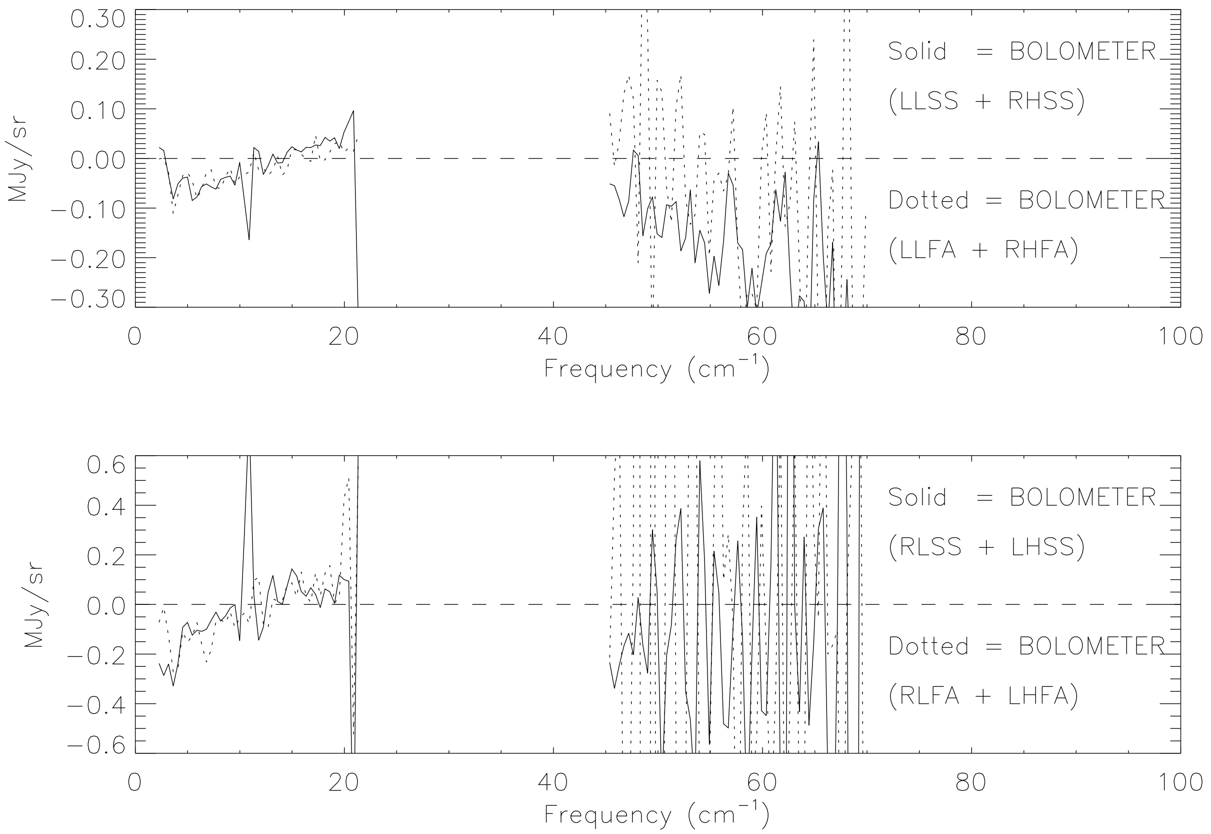

Fig. 6.7.— 'Bolometer' offset correction spectra — Based on a model where the correction is proportional to changes in the bolometer responsivity, a function of the bolometer temperature and bias voltage, relative to a nominal responsivity.

Trial destriper runs also showed residuals correlated with bolometer temperature, as expressed by the derived bolometer responsivity (Section 5.3.2) . The correlation was roughly linear with responsivity. A bolometer function was defined to be the deviation of responsivity from the mean, normalized by the RMS variation of responsivity. The bolometer function was used for those frequency bands for which the F-statistic indicated significant reduction in variance :

LOWF : Yes HIF2 : No HIF3 : Yes HIF4 : No

The definition of the function means that no bolometer correction is applied for observations with responsivity equal to the reference values of -1.974, -1.6261 for the LL detector and RL detector respectively. The referance values for the LH and RH detectors are -.5235 and -.5607 in the slow scan mode and -.4919 and -.5312 in the fast scan mode. These were chosen as the median values of the responsivity. Bolometer effects at these responsivities are absorbed into the mission stripes.

6.3.7. Vibrations

Examination of the destriper residuals showed that the largest residuals were found at frequencies associated with major MTM resonances. These frequencies are at 318 GHz and 2203 GHz for slow scans, and at 212 GHz and 1469 GHz for fast scans (Sections 5.3.2 and 7.9.1). We correct for these by constructing a model function for the residual at each vibration frequency. Our model is an empirical time function derived by smoothing the observed residuals in time. The intent is to reduce the contribution from detector noise by averaging over a sufficiently large number of time-ordered observations so that the dominant remaining residual is due to vibration effects. We use weighted mean residuals for the observations nearest in time for each of the individual channel observations. The residuals have opposite signs for the left and right sides, so we treat the left channel data separately from the right channel data. This indicates a voltage effect rather than a photometric effect. It was also necessary to separate the fast scan data from the slow scan, as the vibration effects are at different frequencies for the different scan speeds.

For the vibration effect at 212 GHz, there was no excess variance in the residuals. No model function was constructed for this effect. For the vibration effect at 318 GHz, we averaged the residuals for the nearest 200 observations for LLSS and RLSS, setting the functions to zero for LLFA and RLFA. The function is called VIBRATION_10 . For the vibration effect at 1469 GHz, we averaged the residuals for the nearest 200 observations for LHFA and RHFA, setting the functions to zero for LHSS and RHSS. The function is called VIBRATION_49 . For the vibration effect at 2203 GHz, we averaged the residuals for the nearest 100 observations for LHSS and RHSS, setting the functions to zero for LHFA and

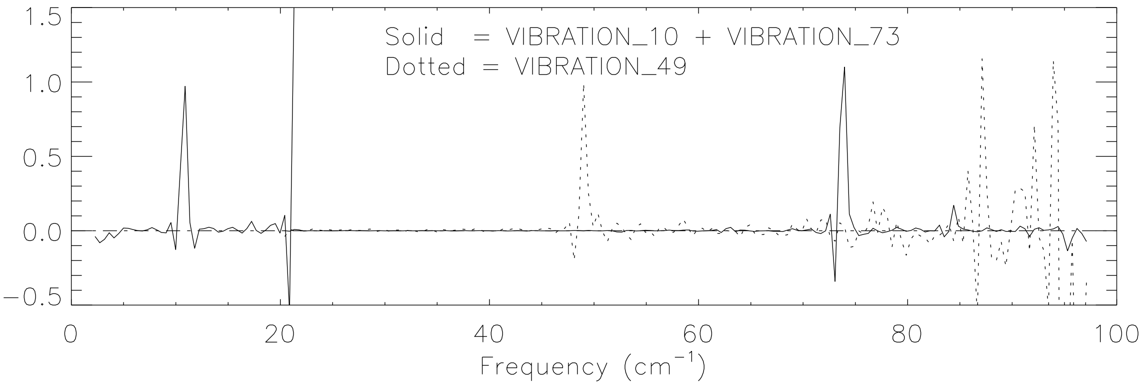

Fig. 6.8.— 'Vibration' offset correction spectra — Derived from a model proportional to the residual signal at the MTM resonaance frequencie.

RHFA. We were able to use a shorter time scale filter for this vibration because its relative effect was larger compared to the detector noise, thus requiring less averaging of the detector noise contribution to the residuals. The function is called VIBRATION_73 .

Examinations of the C-Vectors before and after destriping with the new vibration functions showed that the excess variances at 318 GHz and 1469 GHz were eliminated, as was about 85% of the excess variance at 2203 GHz. Evidence that we really were eliminating vibration effects was seen by the fact that the only features in the correction spectra were found at the vibration frequencies, at their harmonic frequencies, and at the beat frequencies between the 318 GHz and 2203 GHz vibrations. About 15% of the excess variance at 2203 GHz remained. We attempted to eliminate this by constructing a second model function for this vibration from the new residuals. The second function was also a smooth function of time constructed by fitting a 10th order Legendre function to the correlation function between the new residual spectra and the mean residual spectrum for a specified time range, chosen as the 6K horn interval. We found that the fitted model functions were anti-correlated for LHSS and RHSS, and uncorrelated for LHFA and RHFA, giving us confidence that we were modelling the 2203 GHz vibration effect. The function is called VIBRATION_73_X. Comparison of C-Vectors before and after destriping with the secondary vibration function showed that most of the remaining 15% extra variance was eliminated.

6.4. Algorithmic Considerations

The least-squared fit model is linear, therefore in principle it can be solved through matrix inversion. However, because the pixel spectra are fit parameters, this matrix is quite large ($> 6000 \times 6000$) and a direct inversion is not practical. Fortunately, a large portion of the matrix is diagonal and can be inverted simply if partitioned. (We will henceforth call this the *curvature matrix* after Bevington (1969).)

If we write:

$$\mathbf{M} \equiv \begin{pmatrix} \mathbf{S} & \mathbf{R} \\ \mathbf{R^T} & \mathbf{D} \end{pmatrix} \tag{17}$$

where $\mathbf{S}$ is a square matrix with $\mathbf{S}_{km} = \sum_i f_{ki} w_i f_{mi}$, $\mathbf{R}$ is a rectangular matrix with $\mathbf{R}_{pk} = \sum_{i\in p} f_{ki} w_i$, and $\mathbf{D}$ is a diagonal matrix with $\mathbf{D}_{pp'} = \delta_{pp'} \sum_{i\in p} w_i$, then we can express its inverse as:

$$\mathbf{M^{-1}} = \begin{pmatrix} \mathbf{Q} & -\mathbf{QRD^{-1}} \\ -\mathbf{D^{-1}R^TQ} & \mathbf{\Upsilon} \end{pmatrix} \tag{18}$$

where

$$\mathbf{Q} \equiv (\mathbf{S} - \mathbf{R}\mathbf{D}^{-1}\mathbf{R}^{\mathbf{T}})^{-1} \tag{19}$$

and

$$\mathbf{\Upsilon} \equiv \mathbf{D}^{-1} + \mathbf{D}^{-1}\mathbf{R}^{\mathbf{T}}\mathbf{Q}\mathbf{R}\mathbf{D}^{-1}. \tag{20}$$

Thus we see that the largest non-diagonal matrix that must be inverted is the **Q** matrix, whose maximum dimension is 30 × 30.

6.5. Destriper Project Data Sets

The derived quantities of the destriper fit are the correction offset spectra J_ν for each model function, hereafter referred to as "stripes", the spectra A_ν for each sky pixel, hereafter referred to as the "skymap" spectra, and the destriper errors. These quantities are included in the project data sets, as indicated in Table 6.1 . A description of the contents of these data sets can be found in Appendix G.

The A-Vector, C-Matrix, Pixel χ^2, and C-Vector are discussed in Chapter 7. The coadd ZODI model is discussed in Chapter 8. The coadd *DIRBE* gradient functions were discussed in Section 6.3. The orthogonalized and physical stripe errors are discussed in Chapter 7. The pixel ZODI model is the weighted mean of the coadd ZODI model. The ZODI models are included in the destriper data sets because the coadd ZODI spectra were removed prior to destriping. This was done to prevent time variability of the zodiacal light from affecting the destriper. After destriping, the ZODI was added back in to the spectra. The spectra in the FMD_SKY data set therefore include the zodiacal light.

The physical stripes, derived from the FMD_PST data sets, are presented in Figures 6.1 - 6.8. Figures derived from the FMD_AVC, FMD_CSQ, and FMD_CVC data sets will be presented in Chapter 7. The FITS headers for the FMD data sets are given in Appendix G.

7. Errors and Uncertainty Estimation

The estimation of the *FIRAS* errors and uncertainties was a major part of the *FIRAS* data reduction, calibration and destriping. This chapter is an attempt to distill the output of that effort into a concise yet complete description of what is known about the *FIRAS* errors and uncertainties.

Uncertainties in the *FIRAS* data come from many different sources and manifest themselves in many different ways. A full covariance matrix is $210 * 6067 \times 210 * 6067$ which is too large to be convenient. Fortunately many of the uncertainties are quite well described by a few dominant terms that show the source of the uncertainty.

We concentrate on the HIGH, LOWF and HRES data sets as these are the best. The individual channels and scan modes have many of the same uncertainty estimates but post-destriper data sets have lower systematic errors.

7.0.1. Sensitivity

The *FIRAS* sensitivity is characterized by the C and D vectors, which are estimates for the random errors. Over most of the sky and most of the frequency range there is a positive signal. However, there are systematic errors of several types. The differences are important, as some of the systematic errors can be reduced by averaging in some ways but not in others. The uncertainties are both frequency and position dependent. The systematic uncertainties are frequency and position dependent in different ways.

7.0.2. Formal Calibration Uncertainty

In the calibration program (Section 5 and Appendix C) the derivatives of the results with respect to the calibration parameters are calculated. Although the derivatives are expressly calculated for the calibration data the same derivatives can be used for the sky data to propagate the uncertainties of the sky data and the calibration model to the calibrated sky. For convenience the derivatives and the propagation are divided into groups: the detector measurements, the calibration emissivities, the bolometer model parameters, the temperature measurements of all but the XCAL, and the temperature measurement of the XCAL. In earlier discussions (Fixsen *et al.* 1994b), this covariance matrix was expressed:

$$V = D + P_0EP_0 + JCJ + PUP + PTP \tag{21}$$

D is due to random uncertainties in the readings on the detectors. P_0EP_0 is due to uncertainties in the emissivities derived in the calibration. JCJ is due to uncertainties in the other parameters derived from the calibration. PUP is due to random errors in reading the thermometers while observing the sky. PTP is due to uncertainty in the absolute temperature scale of the external calibrator.

The description is complicated by the destriping process, but it is still useful to follow these terms. At the end of this chapter we add sections on tests for other errors and a description of how these uncertainty estimates can be used in practical problems.

7.1. The Detector Noise

Conceptually the detector noise is simple. At each readout there is some noise on the detector which then gets processed and ends up in the final output. The detector outputs are contaminated by noise from the resistance of the detectors, the thermal conductance, cosmic rays, and the readout noise of the electronics. The Johnson noise, and intrinsic noise on the thermal link can be well approximated by a white noise source followed by a filter. The electronic readout noise is small and can also be well approximated by a white noise source followed by a (different) filter. The cosmic rays can be approximated by delta functions applied at random times with some distribution. Although all of these sources lead to noise that is correlated in the time domain, because the noise processes are *stationary* they are uncorrelated in the frequency domain. Thus the only frequency correlations are introduced by the apodization.

7.1.1. The D Matrix

There are two ways to estimate the matrices. As the coadds are formed from the averages of a number of IFGs, the mean dispersion is calculated and from it, the covariance in the final data. This matrix is termed the "D Matrix". For historical reasons we write the covariance as D_ν^2/N_p where D^2 is a diagonal matrix and N_p is some number of "observations" at each pixel, even though the matrix is not really diagonal. By keeping only the square root of the diagonal part we store it in a vector space ("D Vector"). As such we have estimates for the uncertainty of an IFG in any scan mode. For weighted averages we use a standard weight and estimate a "D Vector" "per observation".

This is done for the calibration and the sky data. The calibration data taken in the sky-like near null condition with all controllables near 2.7 K show the same noise characteristics as the sky data. The calibration data taken with one or more of the controllables at a high temperature had higher noise, so these data were deweighted with respect to the sky-like null calibration data. But, the higher signals meant in general they had higher signal to noise ratios (Section 5.5).

The data near the Galactic plane also had higher noise because the data had significant variation over a small change in position. These variations were fit before coadding by a second template to reduce their impact on the data processing. Data that had many glitches removed had more noise that those that had fewer glitches removed so the weights were adjusted to account for the number of glitches (Section 4.4). Although the glitch noise has a redder spectrum than the other bolometer noise we used a single spectrum to allow a simple weight for each IFG, coadd and pixel.

7.1.2. The C Matrix

As the data are destriped, coadds from different times and different channels are compared. The *rms* dispersions of the coadds are accumulated to find the mean covariance matrix (the "C" matrix). In the destriper each coadd is assigned a weight depending on the channel, scan mode, and glitch rate. A single noise spectrum is then determined from the dispersions among the coadds. The noise spectra of the different channels are different in shape as well as amplitude so the use of a single noise spectrum for the combined data is not as efficient as it could have been. However the use of a frequency dependent weight would only reduce the uncertainties $\sim 1\%$ at the price of requiring a frequency dependent weight, and position.

The sky covariance is given by:

$$C_{\nu\nu'} \;=\; \frac{\sum_i \left[S_{i\nu}^{(sky)} - \sum_k J_{k\nu} f_{ki} - A_{p(i)\nu}\right] w_i \left[S_{i\nu'}^{(sky)} - \sum_k J_{k\nu'} f_{ki} - A_{p(i)\nu'}\right] / (1 - f_i)}{N - n_J} \tag{22}$$

where $N \equiv \sum_{i \in sky} 1$, w_i is the weight (Section 6.2), f_i is the fraction of pixel weight supplied by the ith coadd, and n_J is the number of model functions (between 10 and 30). It gives an estimate of the variance of a single interferogram that includes both the random error due to the detector plus any remaining systematic variance. It also provides the scale of the other destriper errors.

The square root of the diagonal is termed the "C vector". Because the dispersion can include errors associated with long term drift and channel to channel errors it might be expected to be larger than the "D vector". But because the weights have been adjusted for time and position the "C vector" is actually lower than the "D vector" for some frequencies. The "C vector" is more inclusive and it is the preferred uncertainty estimator. The C and D "vectors" are only a few percent different (Figure 7.1). In the following text C_ν refers to the "vector" and $C_{\nu\nu'}$ refers to the matrix.

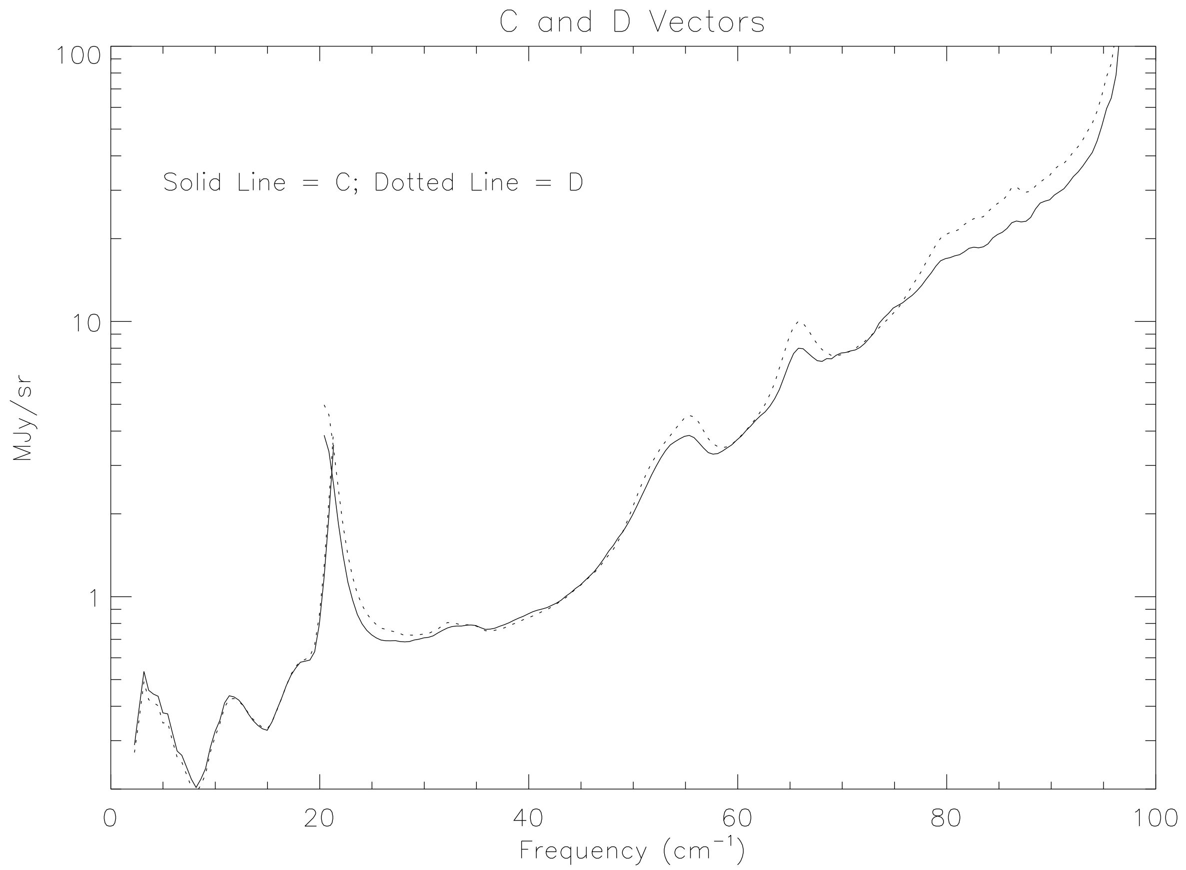

Fig. 7.1.— Relative noise verses frequency — C and D vectors for LOWF (2 to 21cm^{-1}) and HIGH (20 to 100 cm^{-1}).

The off diagonal elements are estimated from the data exactly as the diagonal parts are. This yields a 210×210 covariance matrix. It is convenient to separate it into a correlation part A and the diagonal part C which is the "C vector" identically, $C_{\nu\nu'} = C_\nu C_{\nu'} A_{\nu\nu'}$. A brief look at the $A_{\nu\nu'}$, shows it is very regular. Within 1% it is summarized as $A_{\delta\nu}$, where $\delta\nu = |\nu - \nu'|$. This is called the "A vector".

The expected apodization effect can be calculated from $A = FFT(apod^2)$. This nearly reproduces the measured "A vector" (Figure 7.2). The differences can be attributed to glitch correction effects, baseline subtraction, and perturbations like the vibrations and

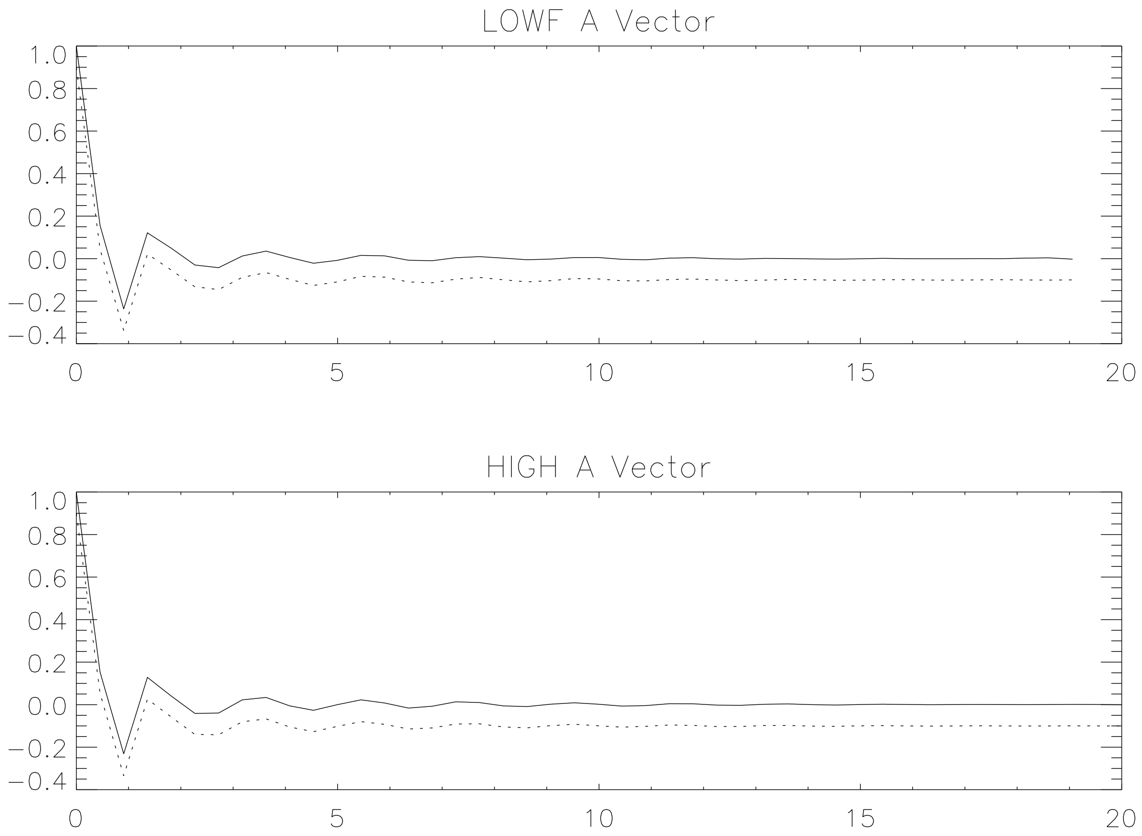

Fig. 7.2.— Frequency-frequency correlations — A vectors for LOWF and HIGH; The solid lines are the "A Vector" estimated from the correlations in the C matrix. The dotted line is the "A Vector" estimated from the apodization, and offset by 0.1 for clarity.

harmonics. The fact that final sky dispersions have the correlations implied by the apodization function alone to within $\sim 1\%$ gives strong support to this uncertainty description.

7.2. The Emissivity Uncertainty

The P_0EP_0 term is the uncertainty which results from the "emissivity" terms of the calibration model in the calibration program, FISH, (*e.g.* errors in the emissivity of the internal calibrator). These errors are propagated to the final sky data by multiplying by the Planck function, P_0, associated with the temperature of the emitting object: in the covariance matrix this term takes the form $P_0 \cdot E \cdot P_0^T$. Since the critical objects are kept at 2.73 K during the observations the P_0EP_0 is summarized with P_0 =Planck(2.73) with the emissivity uncertainty matrix E. But the P_x (P for the external calibrator) is the actual sky spectrum.

For the low frequencies ($\nu < 20$ cm^{-1} or 600 GHz) the sky spectrum is well approximated by a 2.73 K Planck spectrum and the full calculation can be done. The resulting error estimate is contained in the PEP_Offset field of the error term file. For high frequencies ($\nu > 25$ cm^{-1} or 750 GHz) the sky is quite variable and the calculation must be done at each pixel. Most of the terms of the calculation are essentially zero because Planck$\rightarrow$0 at high frequencies. Only the single term P_x, is significant and it assumes the form of a gain error. This gain error, P_gEP_g, is contained in the PEP_GAIN field of the error term file. For frequencies in the range $20 - 25$ cm^{-1} this situation is more complicated but is reasonably well modeled by using both P_0EP_0 and P_gEP_g.

Since the emissivities are calculated independently at each frequency, P_0EP_0 and P_gEP_g have no correlations across frequencies. However, because the whole sky is calibrated with the same instrument emissivity model all pixels are correlated. The terms become important whenever spectra from a large area of the sky are averaged to form a single spectrum.

7.2.1. *The Destriper and P_0EP_0 Offset Errors*

The uncertainty in the parameters of a linear least-squares fit is given by the diagonal elements of the inverse of the *curvature matrix* (Section 6.4). Thus, the uncertainties in the correction spectra of the destriper are given by:

$$\Delta J_{k\nu} = C_\nu \sqrt{[\mathbf{Q}]_{kk}} \tag{23}$$

These quantities are useful for determining whether a particular function of the model is significant. The off-diagonal elements of $\mathbf{Q}$ give the covariance among the different model functions.

The destriper introduces correlations between pixels through the correction spectra, $J_k(\nu)$. These correlations and the uncertainties on the pixel spectra, $A_p(\nu)$, are contained in the lower right sub-matrix $\boldsymbol{\Upsilon}$. In principle this matrix is easy to compute, however, in practice it is not easy to store. However, if we perform a Cholesky decomposition on the real symmetric matrix $\mathbf{Q}$, then we can express this matrix as:

$$\boldsymbol{\Upsilon} \;=\; \mathbf{D}^{-1} + \beta\beta^{\mathbf{T}} \tag{24}$$

where $\beta \;\equiv\; \mathbf{D}^{-1}\mathbf{R}^{\mathbf{T}}\mathbf{L}$ and $\mathbf{Q} = \mathbf{L}\ \mathbf{L}^{\mathbf{T}}$. Note that β has the same dimensions as $\mathbf{R}$, where $\mathbf{R}$ is the matrix of stripe kernels (Section 6.2).

The excluded Galactic pixels are not included in $\boldsymbol{\Upsilon}$, however, it can be shown that correlations among the pixel spectra:

$$A_{p\nu} \;=\; \frac{\sum_{i\in p}\left[S_{i\nu}^{(sky)} \;-\; \sum_k J_{k\nu} f_{ki}\right] n_i}{\sum_{i\in p} n_i} \tag{25}$$

introduced by the destriper corrections, are given by the same expression: $\mathbf{D}^{-1} + \mathbf{D}^{-1}\mathbf{R}^{\mathbf{T}}\mathbf{Q}\mathbf{R}\mathbf{D}^{-1}$, where $\mathbf{D}^{-1}$ and $\mathbf{R}$ now include all observed pixels. ($\mathbf{Q}$) still excludes the pixels from the Galactic center. The diagonal elements of $\boldsymbol{\Upsilon}$ give the pixel spectra variance in units of $C_{\nu\nu'}$.

When modeling the pixel spectra, whether the entire sky or a more localized area, it is necessary to compute the corresponding weights matrix given by $\left[\boldsymbol{\Upsilon}(p, p' \in R)\right]^{-\mathbf{1}}$ where p and p' are pixels within the region R. If the number of pixels within the region is small then this can be done directly, however it is also possible to perform the inverse so that only a small 10×10 matrix need be inverted explicitly. If we express the covariance matrix (or the relevant sub-matrix) as:

$$\boldsymbol{\Upsilon}(p, p' \in R) = \widetilde{\mathbf{D}}^{-\mathbf{1/2}}\ \left[\mathbf{I_m} + \widetilde{\mathbf{D}}^{\mathbf{1/2}}\tilde{\beta}\tilde{\beta}^{\mathbf{T}}\widetilde{\mathbf{D}}^{\mathbf{1/2}}\right]\widetilde{\mathbf{D}}^{-\mathbf{1/2}} \tag{26}$$

$$= \;\widetilde{\mathbf{D}}^{-\mathbf{1/2}}\ \left[\mathbf{I_m} + \mathbf{U}\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathbf{T}}\mathbf{U}^{\mathbf{T}}\right]\ \widetilde{\mathbf{D}}^{-\mathbf{1/2}} \tag{27}$$

where $\widetilde{\mathbf{D}} \;\equiv\; \mathbf{D}(p \in R)$, $\mathbf{I_m}$ is an $m \times m$ identity matrix (m = number of pixels in region), $\tilde{\beta} \;\equiv\; \beta(p \in R)$ and $\mathbf{U}$ is an $m \times n$ row orthogonal matrix (n = number of model functions) then:

$$\mathbf{W} \equiv \left[\boldsymbol{\Upsilon}(\mathbf{p}, \mathbf{p}' \in R)\right]^{-\mathbf{1}} \;=\; \widetilde{\mathbf{D}}^{\mathbf{1/2}}\left[\mathbf{I_m} - \mathbf{U}\boldsymbol{\Lambda}\boldsymbol{\Lambda}^{\mathbf{T}}\mathbf{U}^{\mathbf{T}}\right]\widetilde{\mathbf{D}}^{\mathbf{1/2}} \;=\; \widetilde{\mathbf{D}} - \boldsymbol{\Omega}\boldsymbol{\Omega}^{\mathbf{T}} \tag{28}$$

where

$$\mathbf{\Lambda\Lambda^T} \equiv \mathbf{I_n} - (\mathbf{I_n}+\mathbf{\Phi\Phi^T})^{-1}. \tag{29}$$

and

$$\mathbf{\Omega} \equiv \widetilde{\mathbf{D}}^{\mathbf{1/2}}\,\mathbf{U\Lambda} \tag{30}$$

We can compute $\mathbf{U}$ and $\mathbf{\Phi}$ through a QR decomposition of $\widetilde{D}^{1/2}\widetilde{\beta}$ and obtain $\mathbf{\Lambda}$ via a Cholesky decomposition of $\mathbf{I_n} - (\mathbf{I_n}+\mathbf{\Phi\Phi^T})^{-1}$. Note that $\mathbf{\Omega}$ is an $m \times n$ matrix and thus easy to store and manipulate.

Because the destriper model includes the difference between the calibration and Planck spectra, the destriper uncertainties, β, include "improved" P_0EP_0 and JCJ offset uncertainties. These P_0EP_0 uncertainties are most important when modeling a large portion of the sky, for in this case the error is dominated by the uncertainty of the calibration rather than that of the sky. This can be illustrated as follows. Consider a destriper model with a single entire mission offset with the sky described by a single pixel of N observations. If the total number of calibration observations is given by n, then:

$$\mathbf{M} = \begin{pmatrix} N+n & N \\ N & N \end{pmatrix} \Longrightarrow \mathbf{M^{-1}} = \frac{1}{Nn}\begin{pmatrix} N & -N \\ -N & N+n \end{pmatrix}. \tag{31}$$

The inverse of the lower right element of $\mathbf{M^{-1}}$ is the effective number of observations or weight and is given by $n_{eff} = (Nn/(N+n))$, thus as $N \longrightarrow \infty$, $n_{eff} \longrightarrow n$. Therefore, no matter how precisely the sky spectra are known, the absolute sky is determined no better than the calibration.

7.2.2. β Uncertainties

The destriping effectively recalibrates the instrument eliminating the P_0EP_0 offset errors (and replacing them with the destriper errors). By fitting the sky with stripes, the offset errors in the emissivities of the *FIRAS* internal objects are eliminated. These are the P_0EP_0 offset errors. However the gain uncertainties remain.

The stripe determinations are not perfect and the uncertainty in their determination still leaves uncertainty in the sky maps. The HIGH data are destriped in three frequency segments (Section 6.3). These have varying numbers and types of stripes. We approximate all of the stripes with 10 stripes which are then good for all HIGH frequencies. There are two approximations in this proceedure. First, we concatinate the stripes for the three frequency segments weighted by the total weight of each frequency segment. The pixel weight is divided out of each pixel to give a fair comparison of the data. The matrix is then decomposed into eigen values and eigen vectors. Then we select the eigen vectors and values for the largest 10 eigen values. This set contains $\sim 99\%$ of the total power in the matrix. Finally the pixel weight is multiplied back to give the same form as the original matrix.

This is obviously a simplification, but since the original stripe kernels are empirically motivated, even the full set of stripes is a simplification. Further, by including the other segments' uncertainties we allow for variation in kernels that are used in the other segments, but do not have enough signal to be significant. Finally, the reduction to 10 stripes merely mathematically encodes what we already know: many combinations are well known, and the errors in them are unimportant. For example, the LHSS-RHSS is well determined, since virtually every LHSS coadd has an RHSS coadd to compare. However, since the HIGH uses the average, the LHSS+RHSS has the important uncertainty, although it is much less well determined.

For the combined LOWF there are 30 stripes. When these stripes are orthogonalized the largest few dominate the uncertainties. Therefore only the largest 10 need be used. This effectively summarizes our knowledge and uncertainty of the stripes that we applied. However, the stripes were largely motivated empirically, and the true form of the drifts remains hidden. For this reason we recommend including an uncertainty of $\sim .04\ C_\nu$ when making measurements on the scale of 10° and larger.

7.2.3. Gain Uncertainties

Another way of estimating the gain errors is to compare the results of different channels and scan modes (Section 7.7). The gain can be determined by weighting by the signal and the pixel weight of each channel and scan mode. This shows the gain variations have a random (in frequency) and correlated part. If the random part is identified as the P_gEP_g then it is larger by 1.7 than the P_gEP_g estimated from the calibration results. We recommend using this result because it is larger and it is derived from the final product.

7.3. The Bolometer Model Uncertainties

The bolometer model parameters that are fit in the calibration program have uncertainties which give rise to uncertainties in the sky spectra. These can be conveniently expressed collectively as the product of the Jacobian, J, and the covariance matrix, C, (not the C of Section 7.1) of the parameter uncertainties: $J \cdot C \cdot J^T$. A linear transformation to a different basis set is applied to the parameters to find the principal components of the covariance.

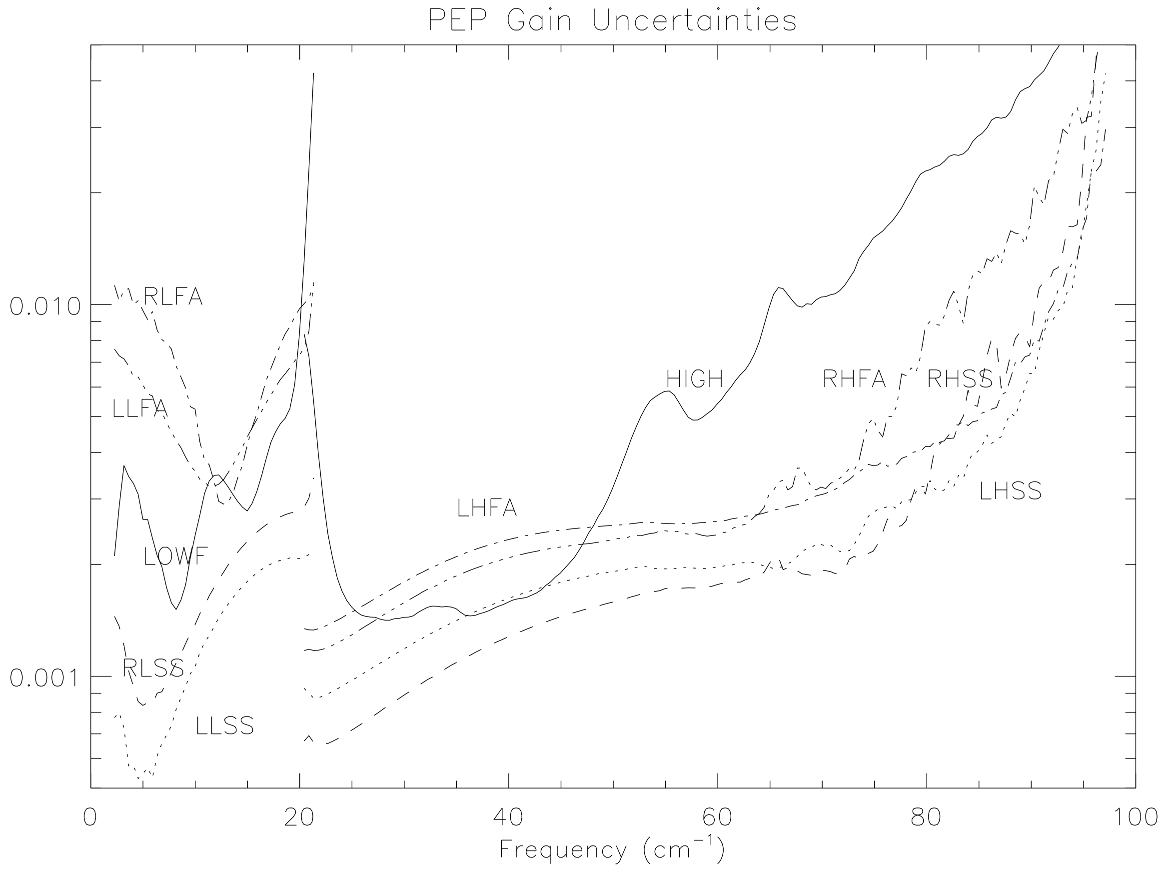

Fig. 7.3.— $P_g E P_g$ uncertainties — Estimated from the calibration models (LLSS, RLSS, LLFA, and RLFA) and the variation in the calibrated data. Gain uncertainty is dimensionless. The shape of the LOWF estimate is from the "C vector" and so should not be taken too seriously, but the scale is determined from the calibrated data variation.

Some errors can be well approximated as gain errors while others are better approximated as offset errors. To find these approximations the calibration parameters are perturbed and the difference in the calibrated sky is noted for 19 selected pixels. The pixels at high latitude, where the signal is small, are used to determine the "typical" offset while those at low latitude, where the signal is large, are used to determine the "typical" gain variation.

These uncertainty estimates are contained in the vectors *JCJ_OFFSET* and *JCJ_GAIN* fields of the uncertainty terms file. While all of the uncertainty terms are included for completeness, most of the uncertainty is included in a single uncertainty term.

7.3.1. JCJ Offset

Offset errors are corrected to first order by the destriper (and replaced by the destriper errors) so these are largely covered in the β matrix of the destriper. The β matrix is assumed to be independent of frequency, however, the *JCJ* offset errors are correlated. This issue will be addressed again below in the "FEF" errors. Detailed study has not shown the *JCJ* offset errors to be a problem either in theory or in fact for either the low frequency (where any *JCJ* is effectively overshadowed by the β) or the high frequency where the β errors are larger and the sky variation can mask the effects.

7.3.2. JCJ Gain

For many of the *JCJ* gain terms the largest effect is a broadband error. These are particularly significant for the high frequency data. The raw terms tend to be flat or slowly rising functions of frequency, sometimes with a steep increase at the high frequency end. This can be easily understood as the highest frequencies remain in the Wien portion of the spectrum even for the highest temperatures (~ 20 K), so a small change in these temperatures in the calibration can lead to a large change in the spectrum and hence in the gain. Over most of the spectrum however, there are at least some calibration data with the frequency in the Rayleigh-Jeans part of the spectrum. Over this range changes in the performance of the bolometer lead to changes in response over the whole spectrum, hence the slow functions of frequency. Even changes in the heat capacity or electrical capacitance effects lead to at most a linear dependence on frequency.

But *JCJ* gain uncertainties are a difficult issue. The formal solution covers the uncertainty in the model parameters, however it does not address the errors in the model $form$. The variance of the data is strongly dependent on the temperature, and higher temperatures

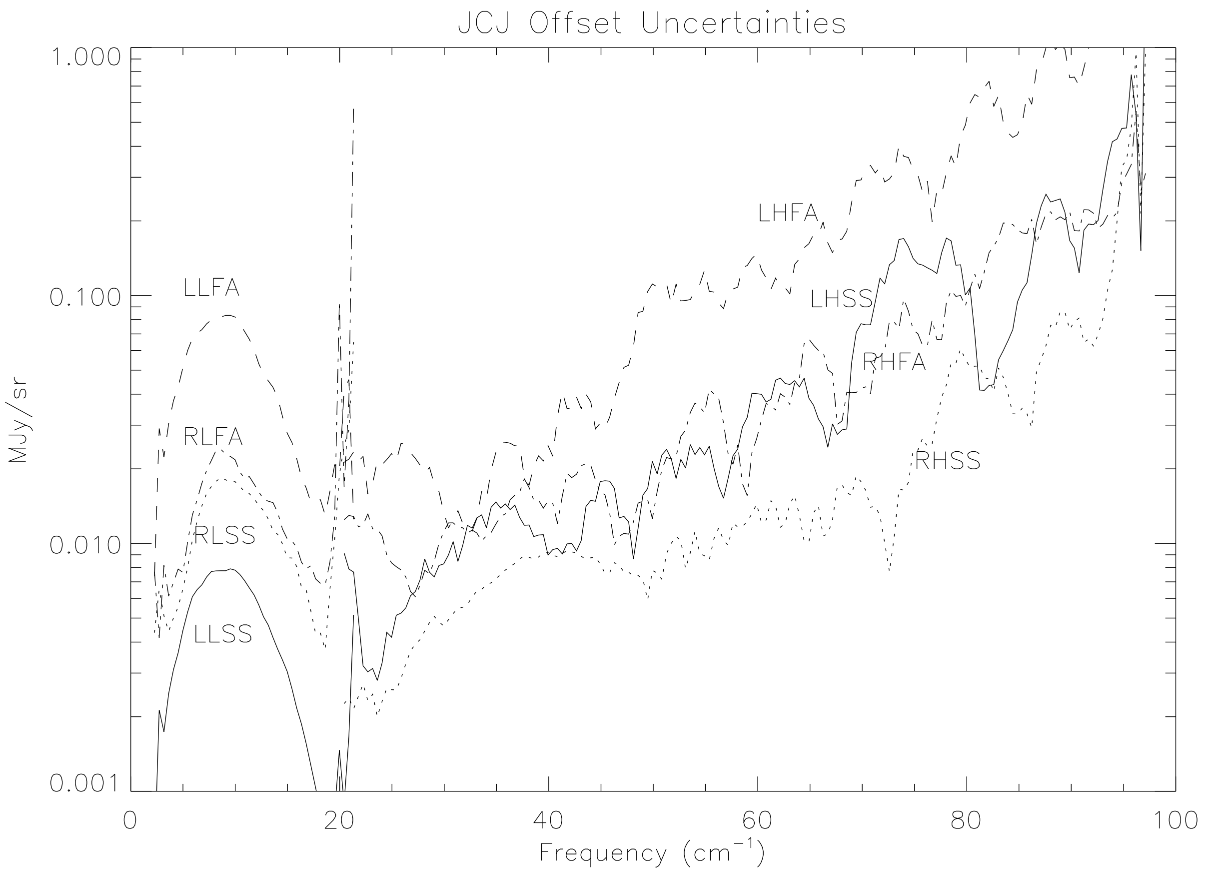

Fig. 7.4.— The *JCJ* offset uncertainties — Derived from the calibration models. The uncertainties reflect the amount and quality of the calibration data.

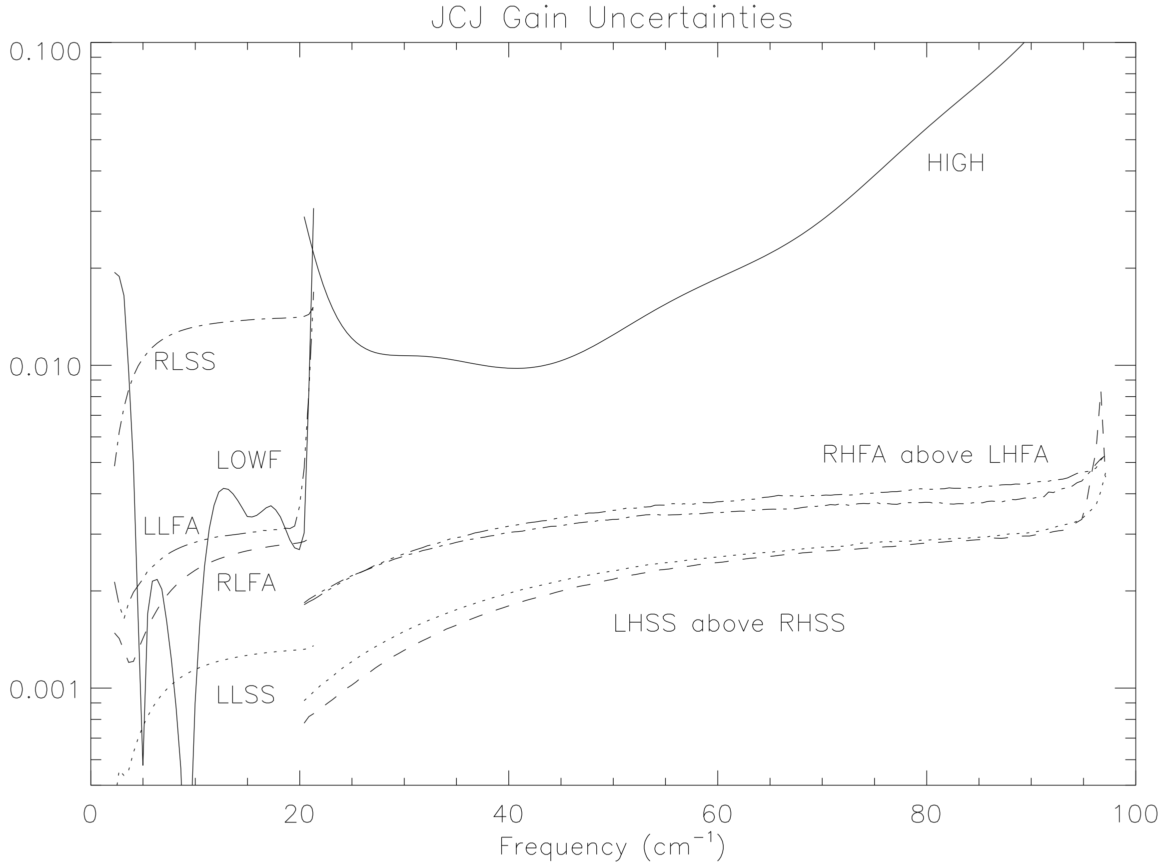

Fig. 7.5.— *JCJ* gain uncertainties — Derived from the calibration models except for LOWF and HIGH which are derived from the calibrated data differences. The dips in the LOWF are probably the result of a chance alignment of the gains.

lead to larger variance within the coadds. Even correcting for this, the χ^2 is larger for the high temperature data (Section 5.5). This is probably an indication of an inadequate bolometer model. Still, even a much more flexible bolometer model is likely to have similarly shaped Jacobians, because the major effects of the bolometer are in the responsivity and the time constant. Further the structure of the χ^2 shows only the features of the vibration lines and the transfer function. Since the radiative loading from the high temperature calibrations is much higher than any place on the sky, the calibration data stresses the model more than the sky data.

Barring the introduction of different parameters in the model we have only a limited means to estimate the true uncertainties. One method is to compare different channels. The random variations were identified as P_gEP_g errors. Identifying the correlated variations in gain as JCJ errors leads to the conclusion that the calibration JCJ determination underestimates the JCJ gain error by a factor of 5. This suggests errors of 1%, however, the basic model is the same in all cases. A missing parameter "X" could well have a similar effect for all channels. If such a factor exists, it probably depends either on the radiative loading on the dectector or on the time constant of the detector. The implied power is different by a factor of ~ 4 between the left and right and the output frequency is in the ratio 3:2 for fast and slow. So, although "X" may push all of the channels in the same direction, it should push left or right ~ 4 times further or push fast or slow 1.5 times further. In either case, the correlated gain errors are the right scale. The results from FISH typically estimate the high frequency gain uncertainties a factor of 2-4 larger than the low frequency uncertainties, and, in fact, they appear a factor of ~ 4 larger. We recommend the combined JCJ gain uncertainties plotted in Figure 7.5 or 0.4% for LOWF.

The differences between LLLF and RLLF are too noisy to use to estimate the gain uncertainty. The HRES P_0EP_0 uncertainty is estimated from the FISH estimates by multiplying by 1.7, the ratio observed in LOWF. The HRES JCJ estimate is just two times the LOWF estimate. The factor of two comes from the ratio of the amount of calibration data in HRES to that of LOWF. The correlated gain errors are treated next.

7.4. The Temperature Noise

Errors in the temperature of the various components of the *FIRAS* while observing the sky get wrongly attributed to the sky. Each error δT_k of course is multiplied by the emissivity of object e_k and the Planck function, $P_k = \partial \mathrm{Planck}/\partial T|_{T_k}$. The errors and emissivities can be summed into a convenient matrix $U_{kk'}$. So the variance can be expressed as $\sum_{kk'} P_{k\nu} U_{kk'pp'} P_{k'\nu'}$. We assume that the off diagonal terms (that is the correlations between thermometers) are negligible, but it is not required for the formalism. The emissivities show that U is dominated by the ICAL, $e_I \sim .98$ (the XCAL is not in the horn for sky observations), except during the "hot horn" season which is treated separately. Since all of the critical elements are kept at 2.7 K for most of the flight the subscript of the P_k is dropped. If we make the assumption (not entirely warranted) that the ICAL temperature readings are independent then the pp' part of the U becomes diagonal and can be conveniently summarized as $1/N_p$ where N_p is the number of observations of the pixel. Then the PUP problem is reduced to estimating a single constant: the uncertainty of reading the ICAL temperature in an IFG. There are three separate ways of estimating this uncertainty.

7.4.1. Resistor Noise

To calibrate the Germanium Resistance Thermometers (GRTs), *FIRAS* uses a set of 4 fixed resistors in the circuit. These were measured along with the GRTs every 32 seconds during the 10 month mission. These $\sim$ 3 million measurements can also be used to estimate the readout noise of the electronics *in situ* to high precision. The *rms* estimate, 16.6 bits is good to 1% for all of the resistors. Multiplying this by an estimate of 130 μK per bit gives an *rms* uncertainty of 2158 μK per readout.

Since the critical ICAL temperatures are derived from two thermometers on two sides as the interpolation between two times, the temperatures are derived from eight separate readouts. This implies an *rms* uncertainty of 880 μK on each IFG and a similar uncertainty on each sky measurement.

This estimate ignores possible correlations and effects that might make the GRT readings less stable than the resistor measurements, so this can be treated as a lower limit to the PUP uncertainty.

7.4.2. Calibration Noise

In the calibration program (FISH) we explicitly solve for the temperatures of the ICAL, XCAL etc. This list of corrections can be used as an estimator of the *rms* errors of the temperature readout and hence an estimator of the PUP uncertainty.

The *rms* weighted noise estimated in this way is 837 μK per IFG. However there are two corrections which must be applied to get an estimate of the PUP uncertainty. First it must be multiplied by $\sqrt{2}$ because FISH is only sensitive to the difference in temperatures in the cold data set (all controllables between 2.6 K and 2.8 K), which is divided between the ICAL and XCAL corrections. Second in the calibration data the IFGs are taken "back to back" so there are correlations introduced by the interpolation to get the temperatures. On average these inflate the final noise by 1.09. So after correcting for these effects the estimated PUP uncertainty is 1086 μK .

At first blush it might seem that this is a perfect estimator of the PUP uncertainty, however there are stripe errors (which we later find and remove in the destriper) and these are included in the estimate. Also there are only of order 300 cold coadds from which to make the estimate. The first item means that this estimator may be biased high and the second suggests uncertainties in the PUP estimate of $\sim 3\%$.

7.4.3. Resulting Noise

The final answer is always a good place to look for noise. In this case we can use several techniques to extract the CMB temperature (Fixsen 1997a) and that temperature can be compared to the known temperature. In the comparison to the *DMR* temperatures we note there is "extra noise" in the *FIRAS* data at $\sim$ 45 μK per *DMR* beam not accounted for by other uncertainties. On average there are 767 IFGs/beam. A *DMR* beam is 16.4 pixels. There are 126382 LLSS IFGs and 157207 LLFA IFGs in the sky data. The RL data are not counted, as their PUP error is fully correlated with the corresponding LL PUP error. This gives an estimate of 1246 μK per IFG.

Since this is an estimate of PUP from the final data it leaves little chance that there is something else "hiding" in the data. However, this discrepancy may be due, at least in part, to other errors. In particular the stripe errors may make a significant contribution to this uncertainty. Still this provides a firm upper limit to the PUP uncertainty.

7.4.4. Hot Horns

For 68 days during the mission the horns (both the sky horn and the reference horn) were held at an elevated temperature, first 6K and then 4K. Although the horn emissivities are much smaller than the ICAL the higher temperatures more than overcome the emissivity and the the net PUP is larger than during normal times. However the PUP in this case is $\nu dP/dT$ at 6K rather than dP/dT at 2.7K, because the horn $e \propto \nu$.

7.4.5. Recommended PUP

The three quite different methods of estimating the PUP uncertainty vary by only 30%. Further the order is what is expected given the possibility of contamination by other uncertainties. We recommend the use of 1 mK per IFG or 150 μK per observation as a good estimate of the PUP uncertainty, with an additional 8μK $\nu dP/dT$ at 6K during the hot horn time.

7.5. The Temperature Uncertainty

Thermometry errors during calibration induce systematic errors in the calibration model parameters. These are collectively called the PTP error. Since there are a large number of measurements in the calibration data ($\sim$ 20000), small random errors largely cancel out. Furthermore, systematic errors in the thermometers of the internal calibrator, the horns, and other internal components of the instrument will be repeated while taking sky data. Hence, the "errors" induced in the model will later be removed by the same "errors" when observing the sky. This is not so for the external calibrator, for which any systematic errors are reproduced in the final sky spectra. Thus, the dominant contribution to the PTP uncertainty is the systematic error in the external calibrator. The form of the PTP error is

$$\frac{\partial P}{\partial T} \cdot T \cdot \frac{\partial P}{\partial T} \tag{32}$$

where the outer factor is the partial derivative of the Planck function of the CMBR with respect to temperature, and T is the variance of the thermometer uncertainty for XCAL. The estimate of the absolute thermometry uncertainty is based primarily on the disagreement between the photometric and thermometric models, which differed by $\sim$4.5 mK at 2.7 K. We have split the difference and adopted 2 mK as the one σ uncertainty in the XCAL thermometer. This uncertainty estimate is contained in the PTP_TEMP and PTP_SPEC fields of the uncertainty term file. These were checked by a third method which uses the Dipole of the CMB. This uses the same frequency scale; however, if we use the DMR to recalibrate the amplitude, we have an independent check, which gives a correction of -7 ± 9mK to the XCAL thermometry.

These three methods give answers that are within three sigma. Apparently there are additional systematic errors in one or more of these methods. We recommend a -2.25 mK temperature adjustment to the XCAL diode temperature scale. This has been applied to the calibrated data. We recommend a 2 mK estimate of the PTP uncertainty. This allows covering both the color temperature calibration and the thermometer calibration. While this is not a true statistical uncertainty it is a useful summary of the uncertainty in the result. The *PTP* error is the dominant error for the absolute temperature of the CMBR and is important for comparisons of *FIRAS* measurements to other experiments.

7.6. Full Uncertainties

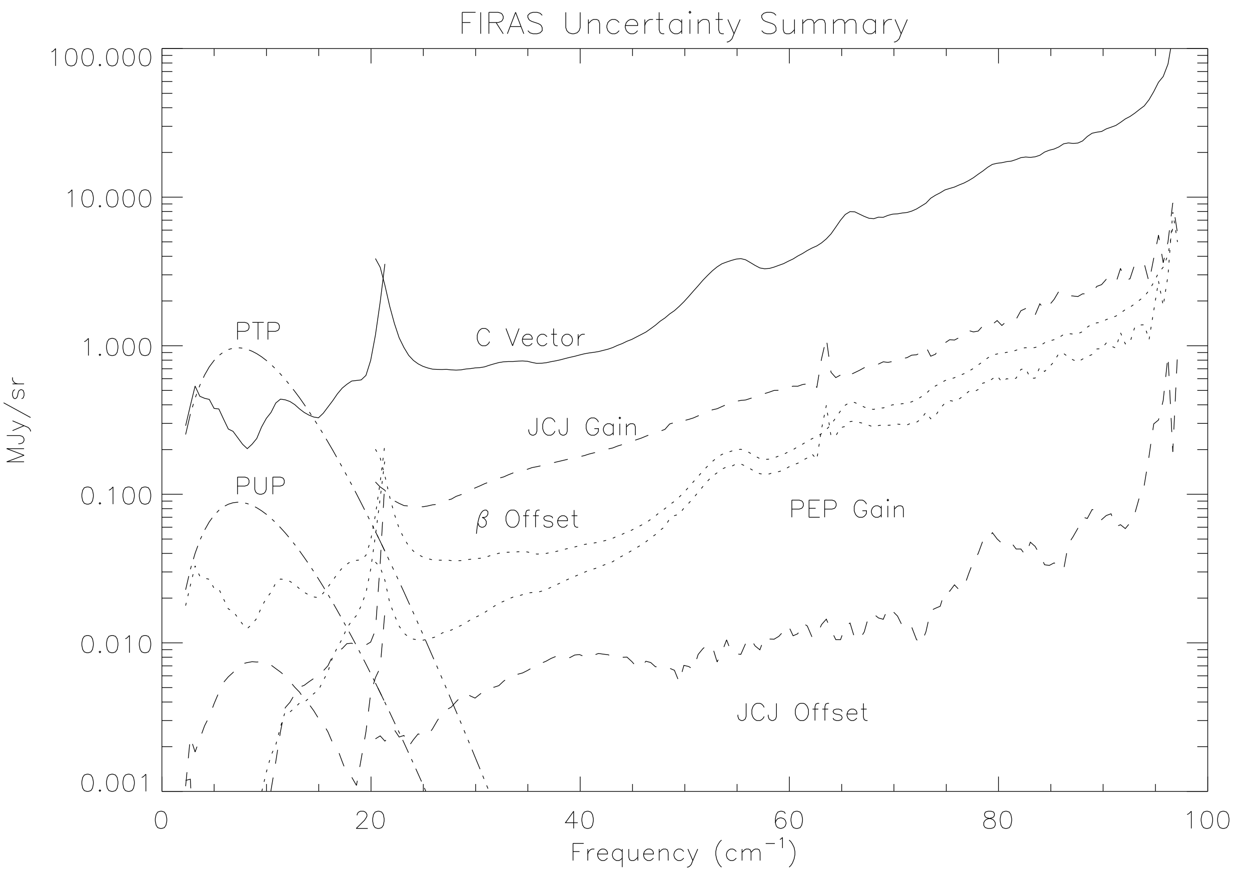

Fig. 7.6.— *FIRAS* Uncertainty Summary — Plots of calibration model solution uncertainty estimates from the C vector, the P_0EP_0, the β, the JCJ, the PUP, and the PTP uncertainties. The uncertainties for JCJ gain and P_gEP_g have been multiplied by the average sky.

The calibration model uncertainties are shown in Figure 7.6. The offset uncertainties have contributions from all five error terms, while the gain uncertainties only have contributions from the P_gEP_g and JCJ uncertainties. The gain uncertainties are multiplied by the average sky in order to make a meaningful comparison. Examination of the offset uncertainties shows that for analyses using less then ~200 observations, the *C-Vector* uncertainties dominate, but for larger data sets other uncertainties must be taken into consideration. Examination of the gain uncertainties shows that the *JCJ* gain error dominates the P_gEP_g uncertainties. The effect of these uncertainties on particular analyses is discussed in Section 7.10.

The largest uncertainty (at least over part of the range) is the *PTP* so when it applies, that is, when the question is what is the absolute temperature, it is the dominant term. Next is the *C-Vector*, but this is for a single average pixel. By averaging or fitting over either pixels or frequencies it can be suppressed. When averaging or fitting over pixels the *PUP* is also suppressed, so it is only important when finding temperature variations and then it is smaller although on the the same order as the *C-Vector*. The *JCJ* gain uncertainty is important at higher frequencies particularly if fitting or averaging over frequencies. The β offset takes the place of the P_0EP_0 which dominates when averaging more than about 400 "observations". The β shown here is an "average" over the sky; it can be larger or smaller depending on the sky sample. The P_gEP_g uncertainty is only important when looking at differential measurements over frequency, e.g. finding lines. We have not found the *JCJ* offset to be a significant uncertainty but we include it here for completeness.

7.7. χ^2 Distributions of Combined Skymaps

We have compared the actual χ^2 distributions to expected distributions for HIGH, LOWF, and HRES. In the HIGH and LOWF combinations the Galaxy stands out as a region of high χ^2. This is due to the fast and slow data sets being taken at different times and hence the registration difference is of the order of a pixel (2.5°). This has only a small effect where there is only a slowly varying signal, but near the Galactic plane it is significant. The errors show the largest χ^2 is not on the Galactic plane but approximately 3° to either side where the beam is just starting to pick up the Galactic plane.

The HIGH skymap displays an excess in χ^2 due to position variation at Galactic latitudes less than 10° (Figure 7.7). The effect is smaller for the LOWF because the high frequency band is more sensitive to the Galactic signal. As with LOWF, sets of IFGs for HIGH were obtained at different times.

The combined χ^2 *vs.* frequency for HIGH averaged for all pixels greater than 8° from the Galactic plane shows the χ^2 to be anomalously large at the seven highest frequencies. The actual and expected χ^2 distributions agree well, however, if one assumes an 8% random

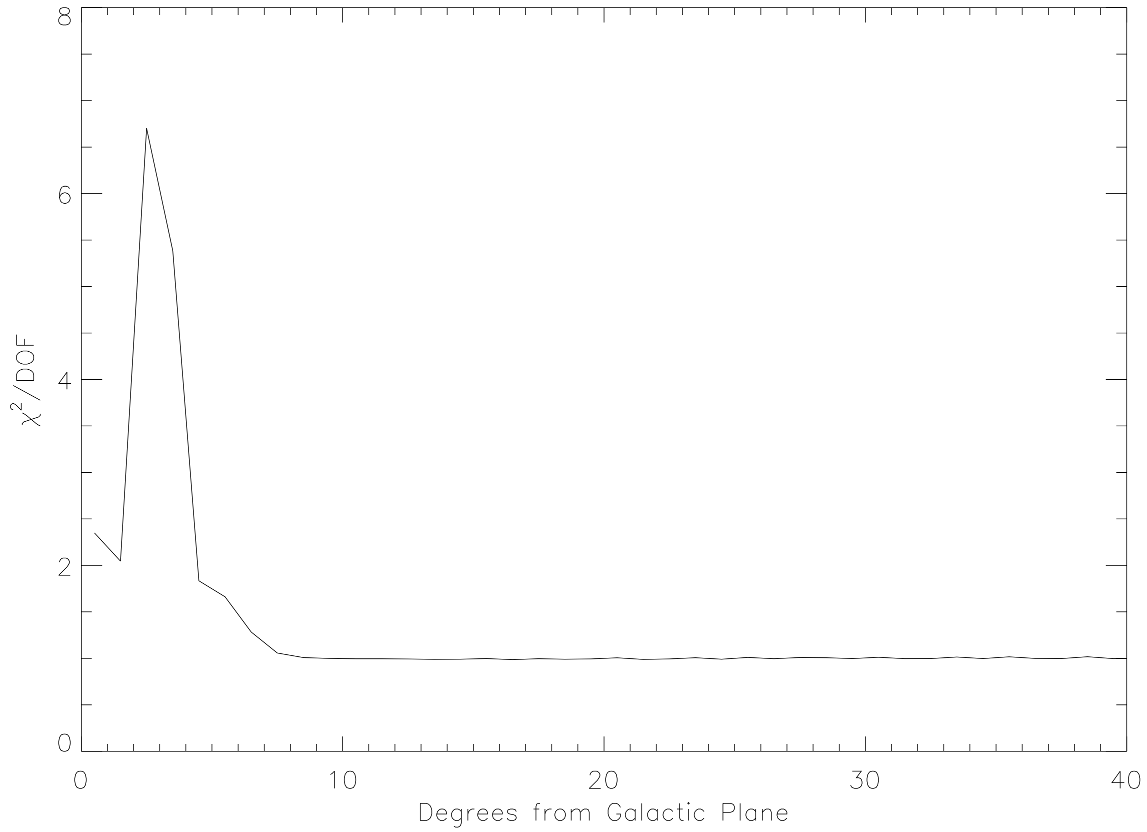

Fig. 7.7.— χ^2/DOF as a function of Galactic latitude — The gradient induces variation in the data larger than the intrinsic variance near the Galactic plane.

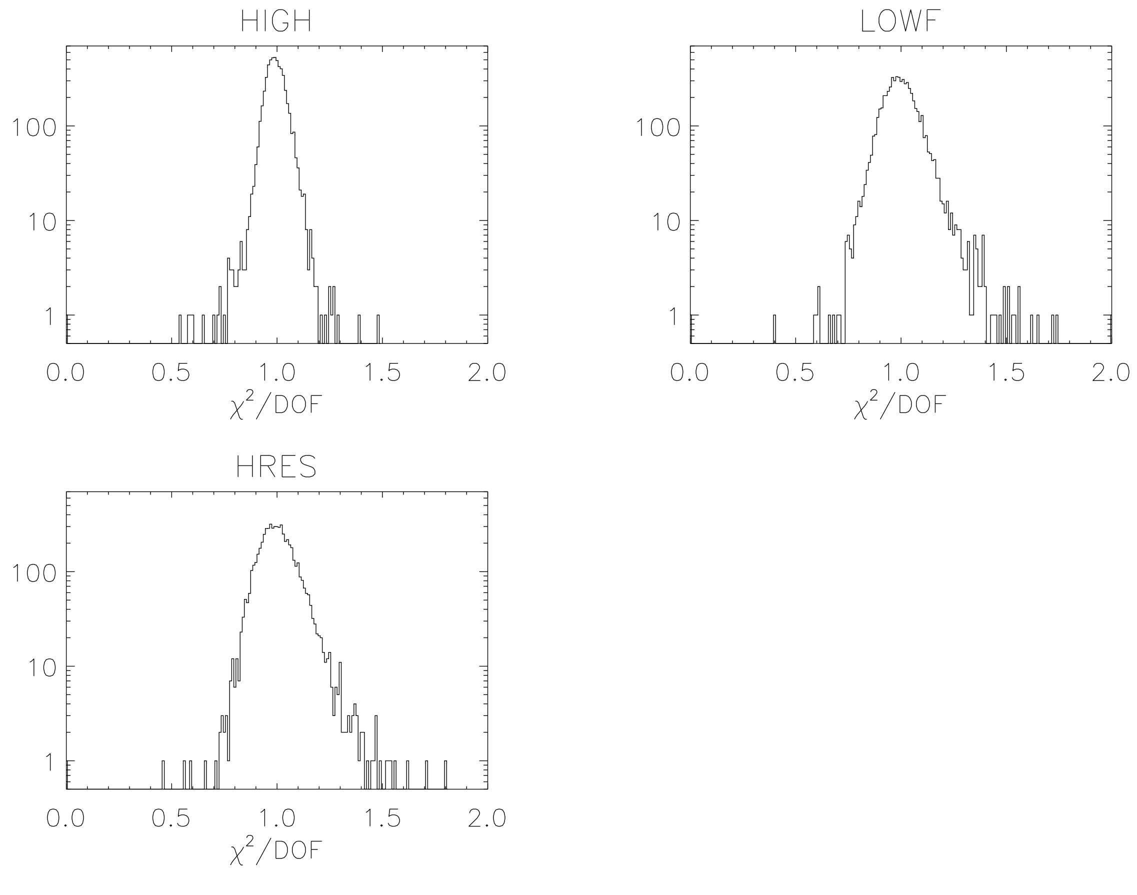

Fig. 7.8.— Histograms of the χ^2/DOF of the HIGH, LOWF, and HRES combinations — The DOF for these combinations are 170, 43, and 182, respectively. The Galactic center is omitted from these plots.

error in the pixel weight (Figure 7.8).

The combined χ^2 averaged for LOWF for all pixels greater than 5°away from the Galactic plane has a broader distribution than expected. The actual and expected χ^2 distributions agree well, if one assumes a 15% random error in the pixel weight (Figure 7.8).

The HRES map shows no evidence of the Galaxy. This is expected, since the input LLLF and RLLF maps have minimal registration difference. The HRES map does display stripes of anomalously low and high χ^2 . The relation of these stripes to mission intervals when the sky and reference horns were raised to higher temperatures suggests that errors in the calibration model for LLLF and/or RLLF may not have been entirely removed by destriping. The other maps show high χ^2 associated with the Galaxy, but are otherwise free of obvious systematic effects.

7.8. "*FIRAS* Extra Factor" Errors

A way to investigate possible errors is to change the processing in "modest and reasonable" ways and to note the variation in the output. Although not particularly quantitative, these variations indicate the type and scale of errors introduced by the restrictions of the models used in processing the *FIRAS* data. These *FIRAS* Extra Factors (FEF) are not rigorous errors in the statistical sense, but are included to point out possible pitfalls for users of the *FIRAS* data.

7.8.1. The Destriper Model

The final results from the *FIRAS* data depend on several choices concerning the destriper model. The choices made in the data processing were the "best" choices in our estimation. Other choices are possible, so we have provided the change in offset for several different parameterizations (Figures 7.9 to 7.14). These changes can be used to estimate an uncertainty in a final result. This is not straight forward, since a one σ result is not defined for a change in the model. Rather than a true estimation of errors, these changes should be used to allay fears (in many cases) or to emphasize caution. In all cases, we have estimated an offset by subtracting the nominal result from a variant averaged over $|b| > 60°$, which measures the offset over a cosmologically interesting region of the sky.

The nominal destriper model has from 10 to 30 terms (Section 6.2). The project data sets contain five sets of comparisons of destriper models: three frequency band comparisons, a low frequency low resolution comparison, and a low frequency high resolution comparison.

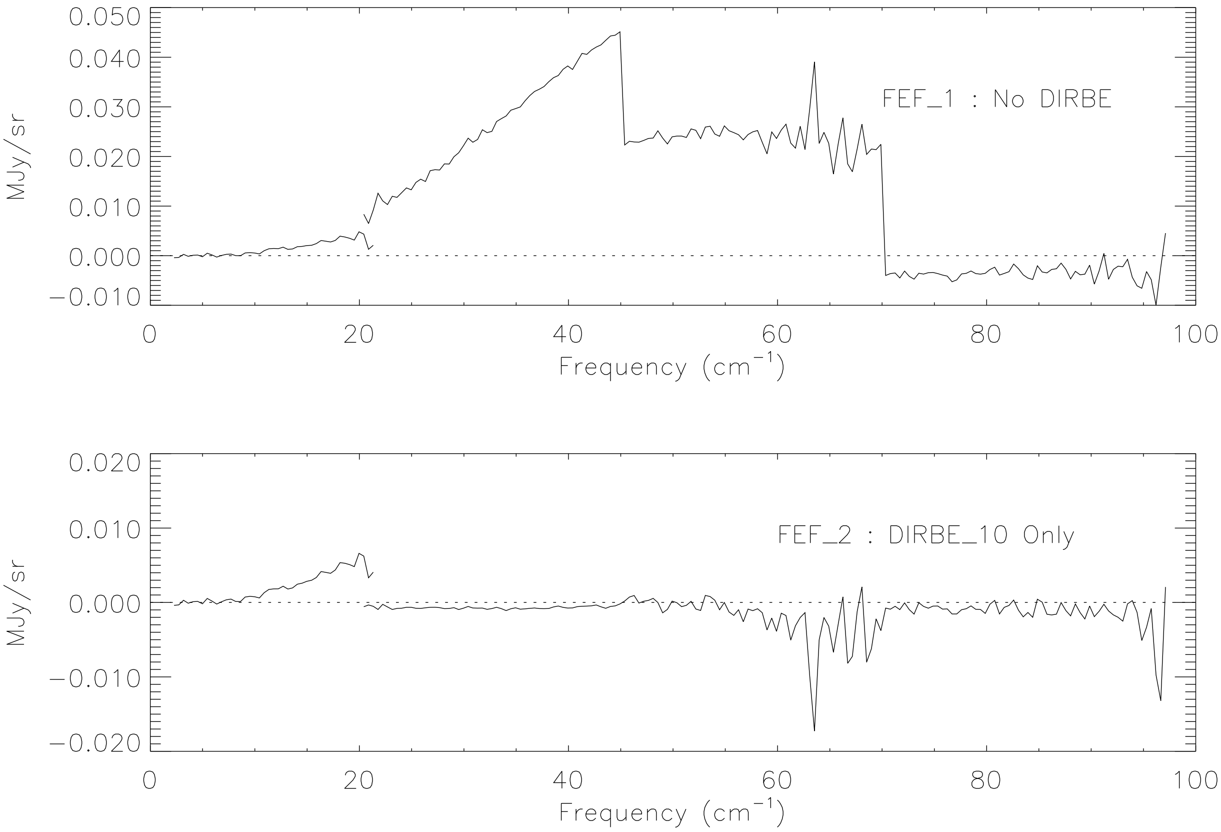

Fig. 7.9.— Galactic gradient model difference spectra — Average variations in the sky spectra for changes in the number or form of the gradient models.

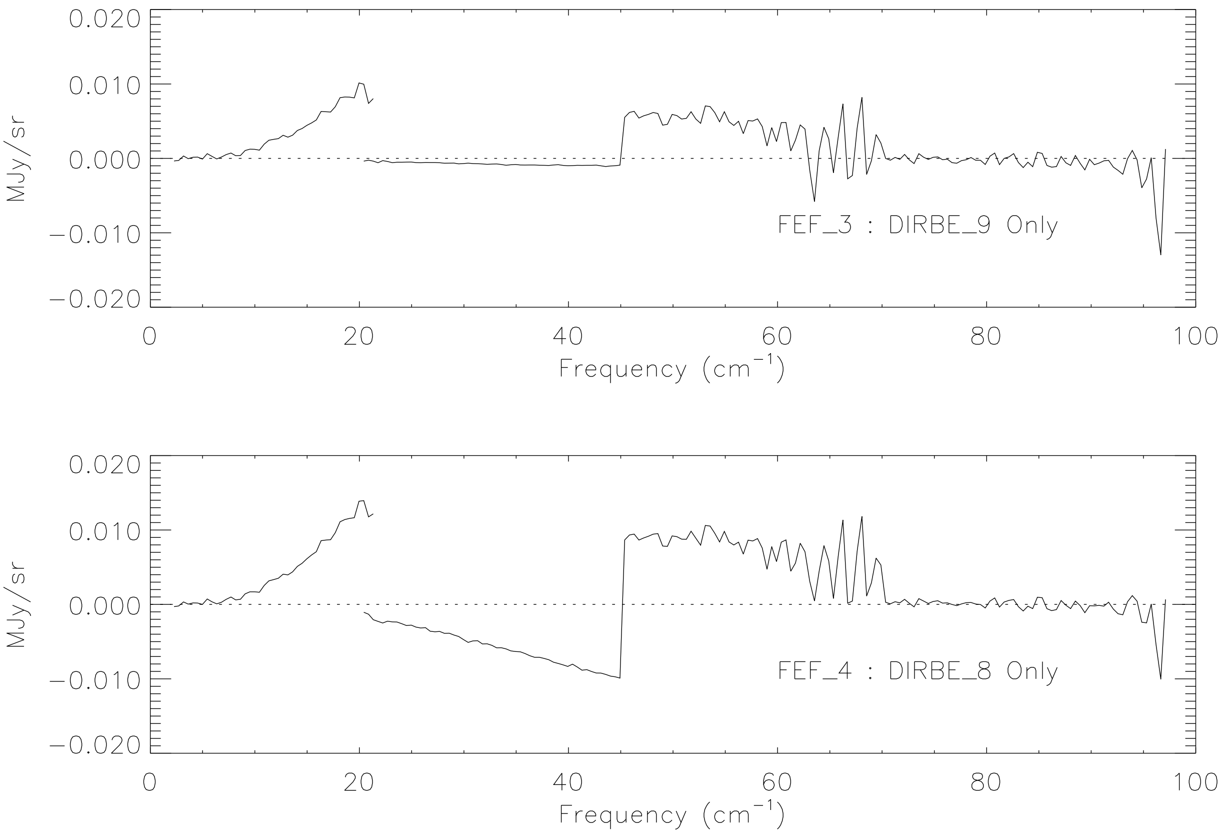

Fig. 7.10.— Galactic gradient model difference spectra (continued) — Q.v. Figure 7.9

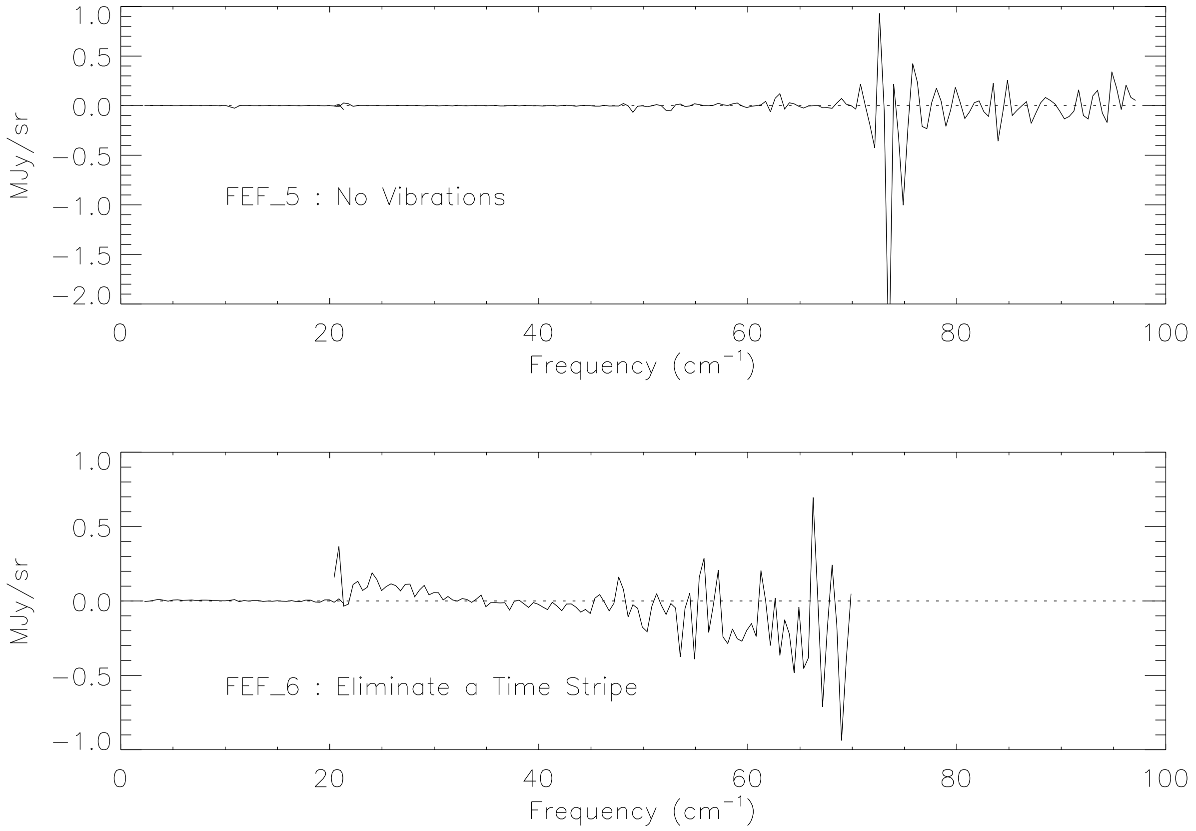

Fig. 7.11.— Vibration and time stripe difference spectra — Average variations in the sky spectra if no vibration stripe was included or if one less time stripe was included. Even though the effect is small in the LOWF, it is still significant.

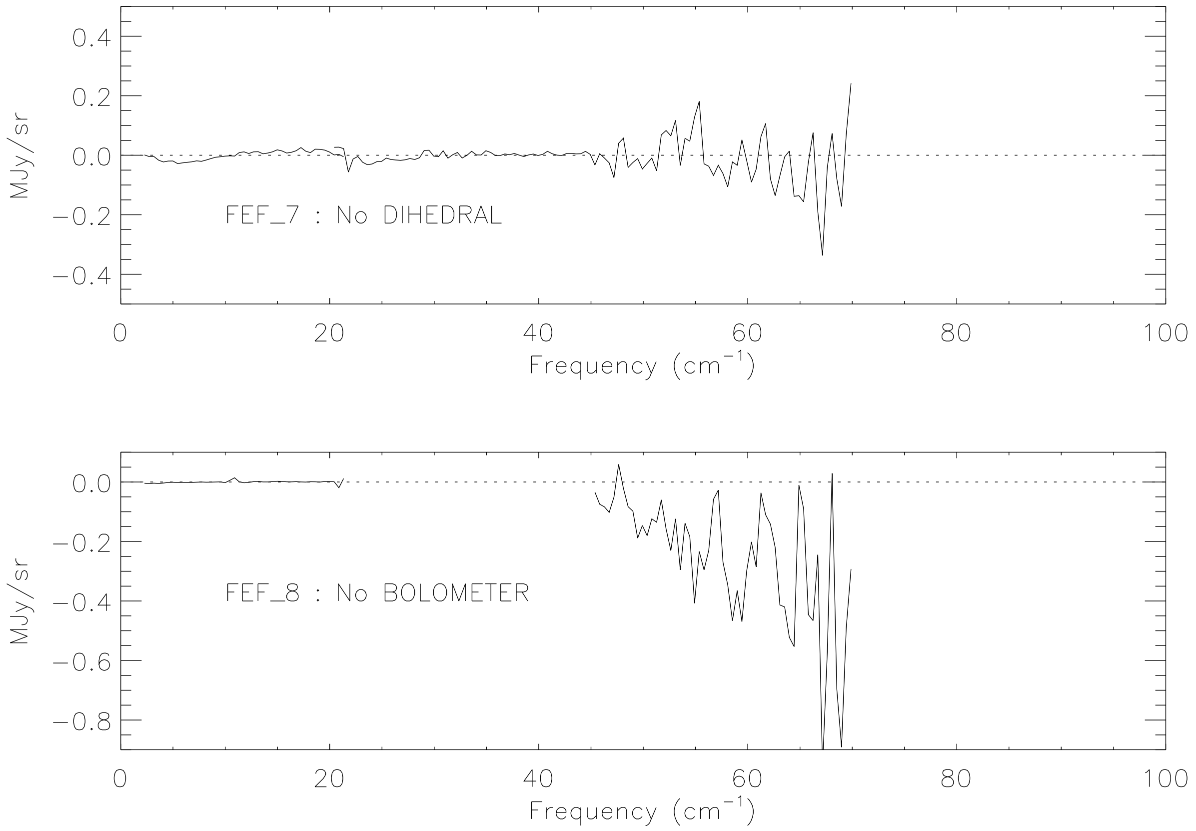

Fig. 7.12.— Dihedral and bolometer stripe difference spectra — Average variations in the sky spectra if the dihedral and bolometer stripes were not included. The effects in LOWF are significant, if small.

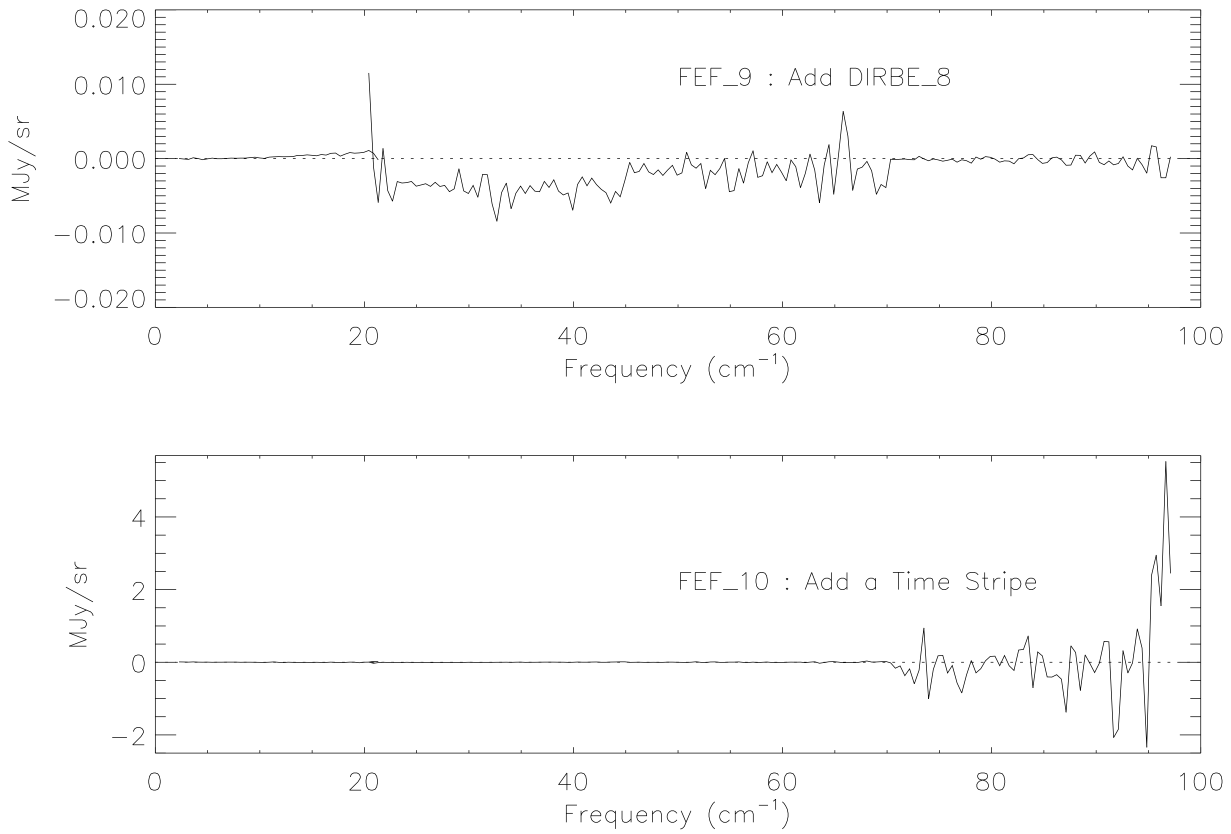

Fig. 7.13.— Additional galactic gradient and time model difference spectra — Average variations in the sky spectra when an additional DIRBE based stripe or an additional time stripe are used. The magnitude of the time stripe difference appears large at high frequencies, but the uncertainties are large there as well. The stripe is not significant by the F test.

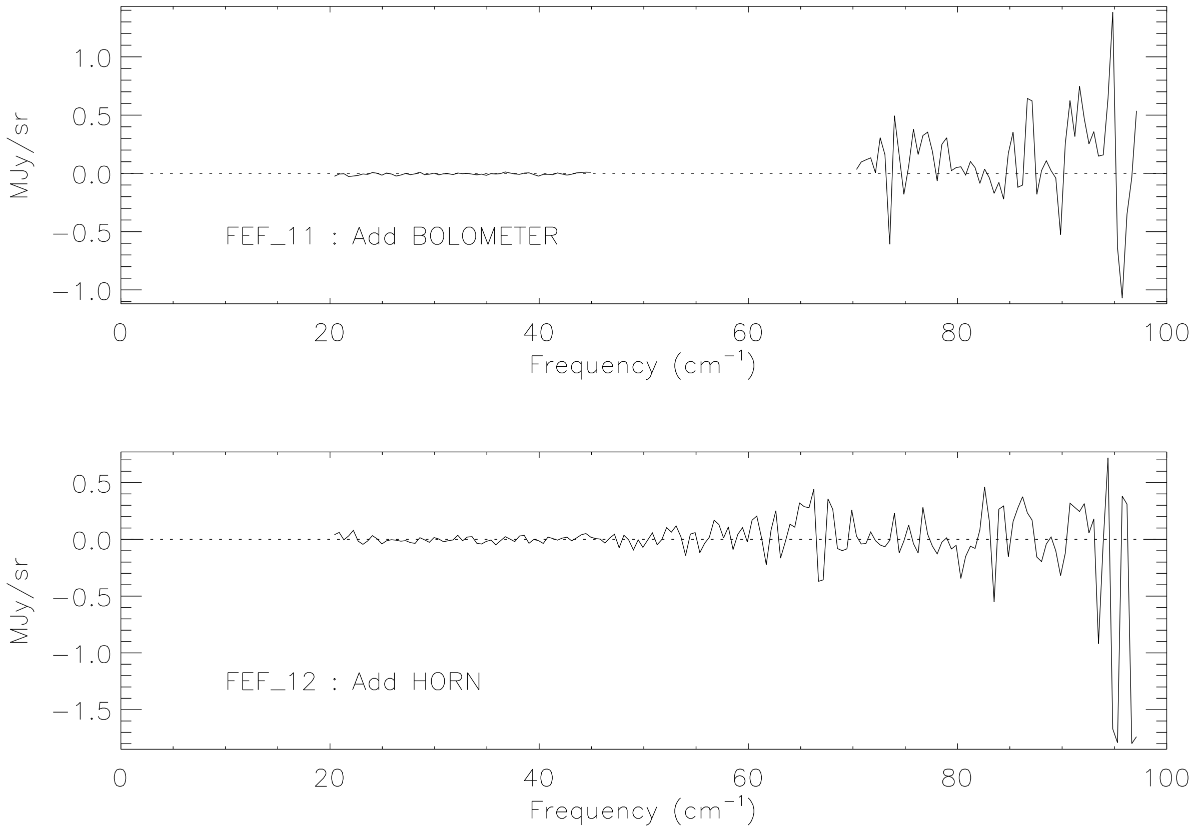

Fig. 7.14.— Additional bolometer and 'hot horn' difference spectra — Average variations in the sky spectra when a bolometer stripe, a 4K horn, or 6K horn stripe, for data where such stripes were not used. The additions fail the F test.

Table 7.1: *FEF* Comparisons

Comparison	Description
FEF_1	No *DIRBE*
FEF_2	*DIRBE* 10 Only
FEF_3	*DIRBE* 9 Only
FEF_4	*DIRBE* 8 Only
FEF_5	No Vibration
FEF_6	Eliminating a Time Stripe
FEF_7	No Dihedral
FEF_8	No Bolometer
FEF_9	Three *DIRBE* Channels (Adding Band 8)
FEF_10	Adding a Time Stripe
FEF_11	Add Bolometer
FEF_12	Add Horn

Each of these data sets contains comparisons of destriper models, which are shown in Table 7.1. The comparisons for LOWF and HIGH are shown in Figures 7.9 through 7.14. The HRES comparisons are available, but are not plotted here.

Together, these comparisons show that final results are quite robust under variations of the destriper model. The LOWF variations are typically on the scale of a few kJy/sr ($\sim 10^{-5}$ of the CMBR peak). These are $< 0.1\ \sigma$ for a large fraction of the data, so we can be confident that any reasonable destriper will give a comparable answer. Since the same data are going into both destripers the errors could in principle be zero. A point at 10.6 cm^{-1} (318 GHz) often sticks out in the LLSS and RLSS data; this is the vibration discussed in Sections 6.3.7 and 7.9.1.

More problems are visible in the high frequency data. The vibration stands out at 73.5 cm^{-1} (2203 GHz) and 49 cm^{-1} (1469 GHz) in the slow and fast data respectively. The high frequency (> 70 cm^{-1}or 2100 GHz) fall off in the performance of the interferometer is evident in the larger errors there. Still the high frequency data are robust under the various choices of destriper, with variations $< 0.1\ \sigma$.

Since the gain variation is an issue for the high frequency channels (Fixsen *et al.*, 1994b) we have included the gain under changes in the destriper model. Not surprisingly the gain changes show little variation since the destriper does not adjust the gain.

7.8.2. Channel Consistency

One of the strengths of the *FIRAS* is that it has multiple channels which cover the same parts of the spectrum. These allow cross checks on the results at many points. If we look at the differences between the channels we see they generally agree to within a few σ. These differences are statistically significant, but we have no clear explanation of their causes. This test is not as powerful as one might hope since the LH detector has much more noise than the RH detector and the RL detector has more noise than the LL detector.

Comparing the gain of the various channels also indicates the possibility of problems. The gain ratio test is only valid for $\nu >\sim 25$ cm^{-1} (750 GHz) since absolute spectra were compared rather than differential spectra (Section 7.3).

To estimate the variations in gain, each channel combination was used to construct a separate map. Each of these maps was then averaged with a weight that included a factor for the noise of each of the maps and a factor for the amplitude of the signal. The result is a spectrum for each of the channels, and a comparison spectrum for the full map. The ratios of these should ideally be one, however, noise and gain errors will induce variations. The variations are fit with a polynomial (10th order for HIGH and 8th order for LOWF). The polynomials are identified as the systematic variations (*JCJ*), while the residuals are identified as the random (in frequency) variations (P_gEP_g errors). These are larger than expected by the gain uncertainties propagated from the FISH results, particularly in the case of *JCJ*. So, the results were combined to give an estimate for the P_gEP_g and *JCJ* uncertainties (Sections 7.2.3, 7.3.2) in the final HIGH and LOWF data sets. Comparison noise dominates in the HRES case.

We have compared the *FIRAS* HIGH data to the *DIRBE* bands 8, 9, and 10 data. This is a robust check since the *DIRBE* instrument is quite different, has a different calibration standard, and looks at the same region of the sky at different times from the *FIRAS* observations. The results for band 10 ($\nu = 42 \pm 8$ cm^{-1}) and band 9 ($\nu = 71 \pm 10$ cm^{-1}) show that to within $\sim$ 5% *DIRBE* and *FIRAS* HIGH data agree. While the absolute *DIRBE* gain calibration is only known to 10%, the agreement with *FIRAS* suggests that both are correct to ~5%. At band 8 ($\nu = 100 \pm 20$ cm^{-1}) the agreement is good to 20%. Noise on the *FIRAS* data and the necessity of a color correction limit the accuracy of the band 8 comparison.

The comparison with *DMR* (Fixsen *et al.* 1997a) gives less than 0.5% difference in the amplitude of the dipole. This suggests that the gain error is <0.5% for the low frequencies, in agreement with our estimate of 0.4%..

7.9. Unmodeled Systematic Errors

In addition to the modeled uncertainties other discrepancies have appeared in the data analysis. These problems are difficult to quantify and their consideration can be awkward but in some applications they dominate the error picture and so they must be considered.

7.9.1. Vibration

There are major MTM resonances at 57.57 Hz and 8.32 Hz, and there are effects on the spectra at harmonics and sum and difference frequencies associated with these resonances. A vibration correction term in the calibration model corrects most effects. However, at the vibration frequencies themselves, the data cannot be entirely corrected. Destriping has greatly suppressed this problem, but one should be wary of any result that depends on the vibration frequencies (Section 6.3.7). The data frequencies translate to optical frequencies 73.5 and 10.6 cm^{-1} (2203 and 318 GHz) for slow scans or 49.0 and 7.1 cm^{-1} (1469 and 212 GHz) for fast scans.

For the high fast data the second harmonic of the vibration (115.14 Hz) occurs at an optical frequency of 98.0 cm^{-1}. In addition, there are sidebands of this harmonic that can contaminate the data above 90 cm^{-1}. These sidebands have not been modeled in the calibration, but have been suppressed by destriping. The second harmonic should be small relative to the main vibration but the signals at high frequencies are also small. In the high frequency slow data the second harmonic is pushed to 147 cm^{-1} and even the sidebands will be strongly attenuated by 120 cm^{-1} so the slow data are not contaminated in this way.

7.9.2. Phase Corrections

There are some significant phase shifts in some of the data. We provide undestriped complex spectra as they are derived from the Fourier transformation. Ideally, the imaginary parts of these spectra would be zero except in regions of the spectra where there are high resolution spectral features, *e.g.*, lines or bandpass filter edges. The line-like features in the imaginary parts are important in the determination of flux and frequency of real astrophysical line emission, but the wideband features in the imaginary parts are presumed to indicate instrument or data analysis faults.

The principal sources of phase variation in the complex spectra are resonant vibrations in the MTM and a small phase shift of unknown origin that is largely linear in frequency (*i.e.*

equivalent to a small displacement of the interferogram).

In addition to the MTM, an unknown source induced both short term and long term interferogram phase variations. It is important for high signal (*i.e.* far from null) data, which occurs primarily in calibration and in the high frequency channels in the Galactic plane. To account for this, we have introduced a phase corrector of the form $e^{i\nu\psi}$, where ψ is typically $< 10^{-4}$ cm but is statistically significant. Because the phase is time variable we fit a ψ to each calibration spectrum. It is not possible to follow an identical phase correction procedure for "cold" sky data, most of which are near null, so an analogous correction is applied wherein the phase is determined by minimizing the imaginary component of the sky spectrum. When applied to the calibration data as well, the results of this algorithm match the results of the more sophisticated fit. Phase errors are most important for the imaginary parts of the spectrum because $e^{i\nu\psi} \sim (1 - (\nu\psi)^2/2 + i\nu\psi)$ for small errors. The effect on the imaginary part is linear in the error, while the effect on the real part is second order as most of the information is in the real part of the spectrum.

7.9.3. Pointing Errors

The pointing solution is from the *DIRBE* data which have uncertainties smaller than 0.1° (Wright 1990). Since the *FIRAS* beam is by definition aligned with the spin axis of the *COBE* spacecraft, the *FIRAS* absolute and relative pointing is good to this same level. This is much better than is required for the 7° beam.

Each IFG is assigned to a pixel according to the line of sight at the center of the observing time interval, but the line of sight moves during that time. In addition, each individual IFG in a coadd group will have sampled its own distinct swath of the sky about that pixel. The individual swaths will largely overlap, since the pixel is oversampled, but the overlap will not be exact. In regions of the sky with large gradients in brightness across the field of view, especially the Galactic plane, we can expect the mean measured spectrum for a pixel to depend upon the details of the pixel sampling (Section 7.6).

The data preparation averages all the observations within a given pixel that are sufficiently consistent with one another and computes a mean line of sight for these consistent interferograms. The process of consistency determination is nonlinear, since it is intended to eliminate seriously defective IFGs. It is therefore difficult to make a quantitative estimate of these errors from the dispersion of the position coordinates alone.

A reasonably reliable estimate can be made from the dispersion of the spectra of individual coadds from the pixel mean. For most pixels, this dispersion is dominated by the random noise described by the "C Vector". Near the Galactic plane, the pointing dispersion dominates (Figure 7.7).

7.9.4. The Beam Profile

The sky horn antenna of the *FIRAS* is imperfect, especially at the higher frequencies. Observations of the Moon show both radial and azimuthal deviations from the nominal 7° circular top hat beam profile (Figure 7.15).

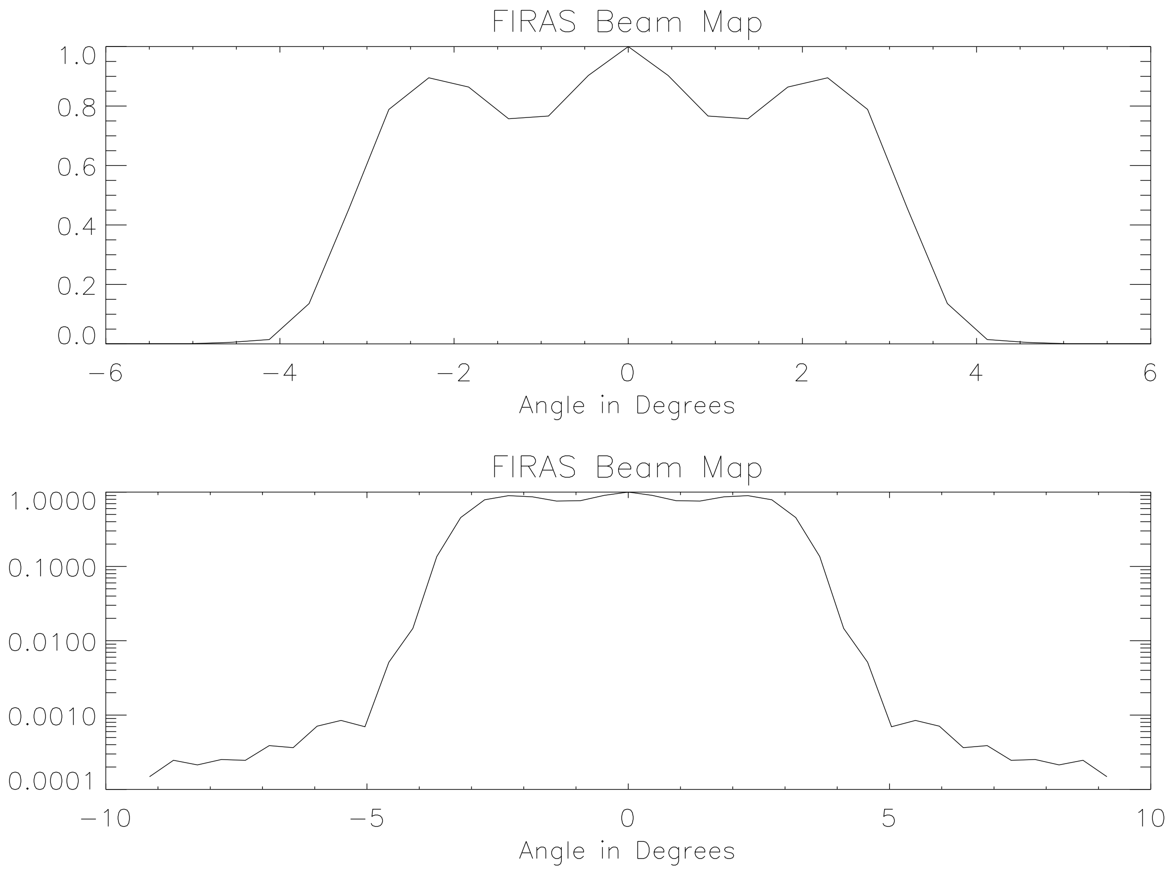

Fig. 7.15.— *FIRAS* beam map — Determined by moon observations. Both linear and logarithmic plots are displayed to show the structure in the beam and the sidelobes.

The measured profile is uniform with respect to radial angle for lunar aspect angles less than 2.8° . About 73% of a 25 cm^{-1} point source flux is contained within a circle of 5.6° diameter. An additional 24% of the flux is contained within an annulus of 5.6° to 7° , with about 3% leaking out of the nominal beam. Within the central 5.6° circle, we observe azimuthal variations of 35% at 25 cm^{-1} and 65% at 70 cm^{-1}.

Although these variations appear to be large, it is important to note that each sky pixel contains data taken from on average 100 IFGs, and each of those is the average of 16 strokes of the interferometer taken while the spacecraft spins about half a revolution and moves its line of sight. Therefore, the effective beam profile is the average of many profiles taken in a large number of possible orientations and positions, and the residual errors are expected to be very small. It is sufficient to note that the beam profile is reflected in the pointing variance.

7.10. Propagation of Errors

The data can be thought of as a matrix with each row as a map at a particular frequency, and each column as a spectrum of a particular pixel. Thus the data can be labeled $S_{p\nu}$. A full covariance matrix would be of the form $V_{pp'\nu\nu'}$ which has $(210 \cdot 6144)^2 \sim 10^{12}$ components. This is too large to store or use easily, so we store five parts which must be added to get the covariance matrix:

$$V \;=\; D\,(or\; C) \;+\; P_0EP_0 \;+\; JCJ \;+\; PUP \;+\; PTP \tag{33}$$

Since the covariances of each uncertainty term add linearly, the uncertainty vectors add in quadrature. Normally, only one or two of the dominant uncertainties need be considered when analyzing the data, but which terms dominate depends on how the data are used. In general, averaging the spectra over different parameters will reduce one or more of the uncertainties leaving the others dominant.

The effect of averaging or fitting spectra across pixels reduces the D and PUP uncertainties. The effect of averaging or fitting spectra across frequencies reduces the D and P_0EP_0 uncertainties. The PTP and JCJ uncertainties remain unchanged no matter how the data are averaged.

To explicitly state the full propagation of uncertainties, suppose a solution of interest, A^a, can be obtained from the *FIRAS* data $S^{p\nu}$ by some linear operator H:

$$A^a = H^a_{p\nu} S^{p\nu} \tag{34}$$

where we have adopted the Einstein summation convention. The index p runs over the list of pixels and ν runs over the frequencies. Now the uncertainty terms that have been discussed (and provided with the data) are: the C matrix= $C^{\nu\nu'}$, the beta matrix=β^p_k, the pixel weight=N_p, the JCJ gain term=J^ν, the P_gEP_g term=G^ν, the PUP temperature= uncertainty U, and the PTP absolute temperature uncertainty=T. Then the variance in the result A^a is a matrix V^{ab} and

$$V^{ab} = H^a_{p\nu} H^b_{p'\nu'} [C^{\nu\nu'} (\delta^{pp'}/N_p + \beta^p_k \beta_{p'k} + .04^2) + S^{p\nu} S^{p'\nu'} (J^\nu J^{\nu'} + G^\nu G^\nu \delta^{\nu\nu'}) + P^\nu P^{\nu'} (U^2 \delta^{pp'}/N_p + T^2)] \tag{35}$$

with some abuse of the summation convention. The $\delta^{pp'}$ is the familiar delta function which is one if the indecies match and zero otherwise and $P^\nu = \partial \mathrm{Planck}(2.728, \nu)/\partial T$.

After reading all of the error analysis and data devoted to error estimation one might come to the erroneous conclusion that the error estimation is difficult and complex. All of the processing has maintained the separability of the pixels and frequencies. This means that in operations involving fits (an average is a particular type of fit) in the pixel domain alone, only the $1/N_p$ and the β terms need to be treated. The $C^{\nu\nu'}$, G^ν, J^ν, and P^ν merely tag along and multiply the uncertainty of the result. Similarly when making a fit on the frequencies only (as in our line and temperature fits) only the $C^{\nu\nu'}$, G^ν, and J^ν need to be treated and the $1/N_p$ and the β terms are carried along and multiply the uncertainties of the results. If a fit involves both pixels and frequencies in a way that can be expressed as a product of operators then the pixels and frequencies can be treated separately.

Further simplification can be realized by noting that in any particular example only one or two terms may completely dominate the uncertainty. A few rules of thumb can be applied in determining the uncertainty terms that need to be used in analyzing the *FIRAS* data. Five examples of the use of the uncertainty terms in particular situations are provided. The rules of thumb are:

1. When only a few pixels are used (with less than 200 observations) and the signal is small, the detector noise alone is adequate.
2. When the signal levels are large, gain uncertainties (P_gEP_g, JCJ) dominate.
3. When data are averaged over large areas (more than $\sim$400 observations) the destriper uncertainties dominate.
4. When large scale structure in the spectra is analyzed JCJ is important.
5. When the absolute temperature is desired (in almost any data set) ONLY the PTP need be considered.

In any case when the final residual from any fit or model is made one should look over the list of FEF errors for one final sanity check.

These examples are given as questions to be asked of the data:

1. Question: What is the spectrum of the north Galactic pole?

   Dominant Uncertainty: C

   Explanation: In order to answer this question we select all of the pixels with Galactic latitudes greater than 87° and average them weighted by the number of pixels. We notice that the total number of observations which went into these spectra is 6. Since this number is less than 200, we expect that only the C uncertainty is important. The only part of the analysis where this is not true is for the absolute temperature of the CMBR.

2. Question: Is there H_2O absorption at the Galactic center?

   Dominant Uncertainty: P_gEP_g

   Explanation: Here we average together spectra within 5° of the center of the Galaxy. If the C uncertainty is the only uncertainty included in the analysis, it would appear that there is a very significant dip in the average spectra at the H_2O line. However, the dip is only a small fraction of the total power at the line. Because the signal level is high the gain fluctuations as a function of frequency (the P_gEP_g gain) are important. Taking these uncertainties into account shows that the dip is a $\sim 3\ \sigma$ effect.

3. Question: What is the temperature of the Galactic dust?

   Dominant Uncertainty: JCJ

   Explanation: We make a weighted average of the Galaxy spectra. In this weighted average we note there are approximately 760 observations. This tells us that C is unlikely to be a good estimate of the uncertainty. The P_gEP_g uncertainty is more important than the C, but the JCJ is still more important. Although the P_gEP_g uncertainty is big at any given frequency, fitting the temperature over frequency reduces the effect of the P_gEP_g. Since the JCJ shows a systematic trend correlated from frequency to frequency, such an error reduction does not occur for the temperature fit. Consequently, the JCJ uncertainty dominates.

4. Question: What is the absolute temperature of the CMBR?

   Dominant Uncertainty: PTP

Explanation: The PTP is the largest error, but it is only important when estimating the absolute temperature. This uncertainty is derived from the uncertainty in the absolute temperature of the external calibration thermometers. Since the sky on average looks VERY much like the external calibrator, we only need to know what the external calibrator temperature is to get the CMBR temperature. But there is a 2 mK uncertainty in the calibrator temperature, hence there is a 2 mK error in the absolute temperature, even though there is enough signal-to-noise to see temperature variations of 10 μK.

5. Question: Can we see the temperature anisotropy that is seen by the *COBE DMR*?

   Dominant Uncertainties: PUP, D

   Explanation: The PUP uncertainty becomes important when looking for anisotropy in the CMBR. Since we are looking for variation in the absolute temperature across the sky, JCJ and PTP are unimportant. Fitting the CMBR temperature over frequency reduces the P_0EP_0 and C uncertainties. Thus the dominant uncertainty is C, but PUP is of the same order so it must be included too. This (and related questions) is the only type of modeling where the PUP uncertainty is important.

Note that in four of the five examples only one of the error terms is needed. And in Question 5 only two terms need to be considered. Further the full matrix was not needed as the error can be calculated after the final result was obtained.

8. Modeling *FIRAS* Spectra

There are four known astrophysical contributors to the *FIRAS* spectra: the cosmic microwave background radiation (CMBR), thermal Galactic dust emission, narrow atomic and molecular lines from interstellar gas, and solar system thermal dust emission (zodiacal emission). The determination of these components is the subject of a number of papers from the *FIRAS* analysis team (Section 8.2 and Appendices A and D). The project data sets include a determination of the contributors' intensities and properties, though there may be slight differences in the techniques used from those of the refereed published results, which should be considered definitive. In this section we discuss how the delivered products were produced. They can be duplicated using the delivered material. Except where explicitly stated otherwise the delivered products' production did not include the appropriate "off-diagonal" weights for the pixel-pixel and frequency-frequency correlations. The derived uncertainties can be much smaller than what the more detailed analysis would derive.

The delivered products were based on a four part analysis of the combined HIGH and LOWF data sets: (1) A model, provided by the *COBE DIRBE* analysis team, for zodiacal emission was subtracted from the HIGH; (2) A model CMBR (monopole and dipole) was fit to the LOWF data, and then subtracted from both the LOWF and HIGH data; (3) A model galaxy was fit to the LOWF and HIGH data and subtracted; (4) The residual was fit to a polynomial baseline and line profiles to determine the line intensity maps.

8.1. Zodiacal Emission Model

The zodiacal emssion is thermal emission from dust, with temperatures of order 250 K, that is in a cloud generally centered on the sun. There are strong variations in the dust quantity and temperature as a function of distance to the sun, and the projected distance out of the ecliptic plane. The peak of the emission is at frequencies higher than covered by *FIRAS* though its low frequency tail can impact the *FIRAS* data, particularly for $\nu > 20$ cm^{-1}, attaining intensities comparable to the galactic dust emission at 100 cm^{-1} for the lowest intensity pixels.

The existence of zodiacal emission can be confirmed with the *FIRAS* data, but they lack the signal to noise to be able model the emission in detail. For that we have relied on a model derived by the *COBE DIRBE* team based on a parametric fit of a geometric model to the time variation in their data. The particular model and parameters used (*DIRBE* Science Pipeline Run 6) are to be described in detail by Kelsall *et al.* (1998).

The *DIRBE* zodiacal emission model is explicitly calculated only for the particular *DIRBE* band wavelengths. The most relevant ones for the *FIRAS* are the two longest wavelength *DIRBE* bands at 1260 and 2128 GHz (240 and 140 μm). The *DIRBE* model was interpolated/extrapolated to the entire *FIRAS* band by fitting a power law emissivity that connected the two *DIRBE* points.

This *DIRBE* based model can be compared in detail to the *FIRAS* sky, at least at the lower flux levels seen outside the galactic plane. Though the *FIRAS* data set is itself insufficient to derive the the details of the zodiacal model, it can constrain the general normalization and the emissivity power law index. Earlier versions of the *DIRBE* model were clearly improved if the emissivity was steepened, and the overall intensity reduced for the *FIRAS* frequencies. These changes were clearly important in that they could induce systematic shifts in the determined galactic dust properties by as much as 5% in temperature and 25% in amplitude.

In addition to its use in the astrophysical fits of the *FIRAS* spectra, the zodiacal model was used by the destriper (Section 6.5). The destriper depends on the constancy of all the spectra within a pixel. Temporal variations are induced in the spectra as the earth moves through the zodiacal dust cloud. For this purpose the unadjusted *DIRBE* model was used to calculate the expected *FIRAS* emission as a function of time for each coadd. The destriper runs for the high frequency channels then subtracted these spectra before the destriping. After the stripe amplitudes were determined, and the destriped pixel spectra calculated, the zodiacal model intensities were then added back to the pixel spectra so that the project data sets *still contain the entire contribution from zodiacal emission.*

For the analysis data sets, whose production are described in the following sections (8.2–8.4), the *DIRBE* based zodiacal emission model, without alteration, was first subtracted from the HIGH spectra.

8.2. CMBR Determination

8.2.1. Temperature Maps

The low-frequency spectra (LOWF) are dominated by the cosmic background, but they contain a significant Galactic spectral component which must be well-modeled in order to derive accurate CMBR temperatures. We do so by computing an empirical Galactic spectrum at each frequency by modeling the sky with three spatial templates: a uniform component, a dipolar (Doppler-shifted) component, and a Galactic component that we assume correlates spatially with the average high-frequency intensity at each pixel. That is

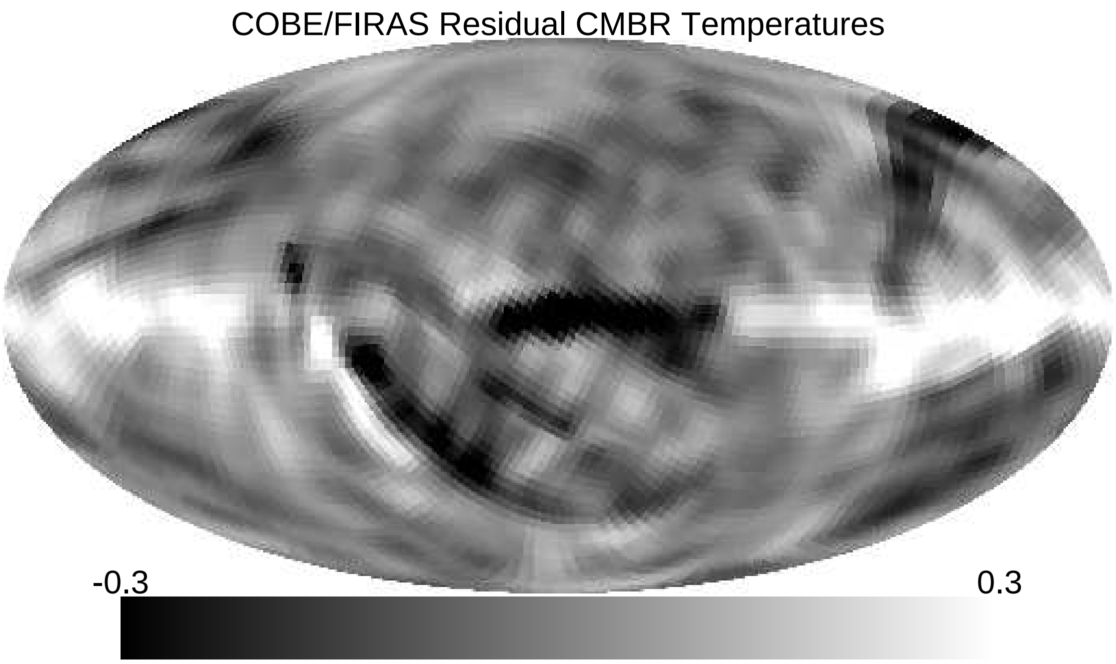

Fig. 8.1.— CMBR residual temperature — (mK). Both the monopole and dipole terms have been removed. This and all other maps in the explanatory supplement have been projected in Galactic coordinates. In addition, a color postscript version of this and all other maps are available, see Appendix K.

we minimize:

$$\chi^2_\nu = \sum_p \left[S_{p\nu} - u_\nu - D^x_p d^x_\nu - D^y_p d^y_\nu - D^z_p d^z_\nu - G_p g_\nu\right]^2 n_p \tag{36}$$

where

$$D^x_p \equiv \cos(b_p)\cos(l_p),\ \ D^y_p \equiv \cos(b_p)\sin(l_p),\ \ D^z_p \equiv \sin(b_p) \tag{37}$$

provide the dipole spatial pattern based on the pixel position (l_p, b_p), and G_p is the weighted average of the combined high-frequency skymap for frequencies above 25 cm^{-1} and n_p is the pixel weight. Pixels within 15° of the Galactic plane are excluded from the fit. The derived Galactic spectrum, g_ν, is next used to compute the CMBR temperature using the model:

$$\chi^2_p = \sum_\nu \left[S_{p\nu} - P(\nu;T_0) - \Delta T_p \frac{\partial P(\nu;T_0)}{\partial T} - G_p g_\nu\right]^2 \frac{n_p}{C^2_\nu} \tag{38}$$

where $P(\nu;T_0)$ is a Planck spectrum at $T_0 \equiv 2.728°$, ΔT_p are the temperature corrections and we fit a new value for G_p, the "amount" of Galaxy in each pixel. The fractional expression at the end gives the spectral weight at each frequency. The final CMBR temperatures are given by:

$$T_p = 2.728° + \Delta T_p. \tag{39}$$

These temperatures are computed at each pixel for the low frequency destriped skymap. We can compute an overall CMBR temperature and dipole temperature (amplitude) and direction by spatially decomposing the temperature map much as before. That is:

$$T_p \ = \ T_{cmbr} \ + \ D^x_p\, T_x + D^y_p\, T_y + D^z_p\, T_z \tag{40}$$

where T_{cmbr} is the CMBR temperature and T_x, T_y, T_z are the dipole temperature components from which we compute the dipole amplitude and direction.

The CMBR temperature and the CMBR residual temperature for each pixel in the input skymaps make up the CMBR analyzed science data set. The FITS header for the *FIRAS* CMBR temperature map is given in Appendix G.

Figure 8.1 is a map of the CMBR residual temperature projected in Galactic coordinates. Residual effects of 0.5 - 1.0 mK can be seen in the Galactic plane. Otherwise, the map is free of obvious systematic effects, and is roughly consistent with a normal distribution with an RMS scatter of 214 μK for Galactic latitudes greater than 15°. The uncertainty in the CMBR temperature for any given pixel is dominated by the detector noise and *PUP* uncertainty.

We have performed an autocorrelation analysis of the CMBR temperature residuals as a function of angular separation. The results are consistent with the *DMR* anisotropy results (Fixsen *et al.*, 1997a) .

8.2.2. Considerations When Modeling Cosmological Parameters

The cosmic background as measured by *FIRAS* is very well described by a Planck spectrum, therefore in order to detect any significant deviation from this, it is necessary to compute a full-sky average to obtain sufficient signal to noise. We calculate this spectrum, along with the CMBR dipole and "average" Galactic spectrum by finding the "projections" of the sky along three spatial templates. The first two are a uniform template with value one everywhere and a dipole template, $\cos(\theta)$, where θ is the angle between the given pixel and the dipole direction. The Galactic template is more problematical but is approximated with an average Galactic intensity derived from the high frequency channels (cf. section 8.2.1). To derive the spectra we minimize:

$$\chi^2_\nu \;=\; \sum_{pp'} [(S_{p\nu} - u_\nu - D_p d_\nu - G_p g_\nu) W_{pp'} (S_{p'\nu} - u_\nu - D_{p'} d_\nu - G_{p'} g_\nu)]^2 \qquad (41)$$

with respect to u_ν, d_ν, and g_ν where D_p and G_p are the dipole and Galactic templates and $w_{pp'}$ are the pixel weights described in the previous section, with the sums performed over the relevant pixels. Here we explicitly include the pixel-pixel correlations. The spectra are computed frequency by frequency.

The amount of non-thermal distortion in the CMBR is generally shown by the residuals of a fit of u_ν to a Planck spectrum, $P(\nu;T)$. The parameters of specific cosmological models are fit directly in equation 41, expanded to include a $\Sigma_{\nu\nu'} w_{\nu\nu'}$ weighted sum over frequency. Alternatively we fit a model to the u_ν derived from the general fit of equation 41. Examples of the cosmological models used are: a Bose-Einstein spectrum to compute μ, the chemical potential; or a Planck plus Zel'dovich-Sunyaev spectrum to compute y, the Comptonization (Kompaneets) parameter. The cosmic background temperature, T_{cmbr}, is computed. The two models are given by:

$$M(\nu, T_{cmbr}, \mu) \;=\; \frac{2hc^2\nu^3}{\exp(x+\mu) - 1} \qquad (42)$$

and

$$M(\nu, T_{cmbr}, y) \;=\; P(\nu, T_{cmbr}) \,+\, T_{cmbr} \left[\frac{x(1+\exp(-x))}{1-\exp(-x)} - 4 \right] \frac{\partial B}{\partial T} y \qquad (43)$$

respectively, where $x \equiv hc\nu/kT$. In fits of u_ν the effective weight must include its statistical uncertainty, which including the frequency dependent weights is:

$$w_{u;\nu\nu'} \;=\; w_{\nu\nu'} \sum_{pp'} w_{pp'} C_\nu^2. \tag{44}$$

The small allowed level of distortion generally assures that the model can be linearized in terms of the parameters, maintaining the linear nature of the problem. Although this fit is not technically difficult, it demands care to ensure that all systematic errors are properly accounted for. For example, some "Galactic" signature "leaks" in to the uniform spectrum, so it is usually necessary to include a supplementary Galactic spectrum in the fit along with the cosmological models (cf. equation 38). It is necessary to consider the sensitivity of the parameter values to Galactic latitude cutoff.

8.3. Galactic Dust Parameter Maps

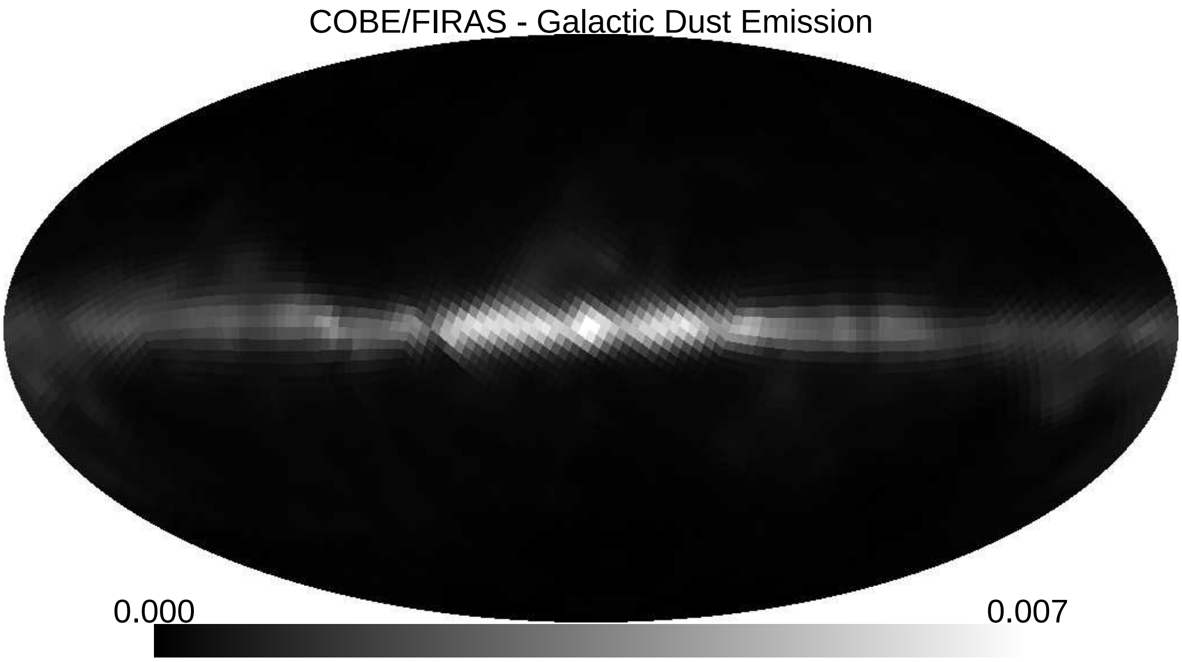

Fig. 8.2.— Galactic dust emission — τ. See equation 45, $\nu_0 = 1800$ GHz. (qv. Figure 8.4)

We use the map of CMBR temperatures derived above to subtract the CMBR, $P(\nu, T_{\mathrm{cmbr};p})$ from both the LOWF and HIGH spectra. As we have already removed the

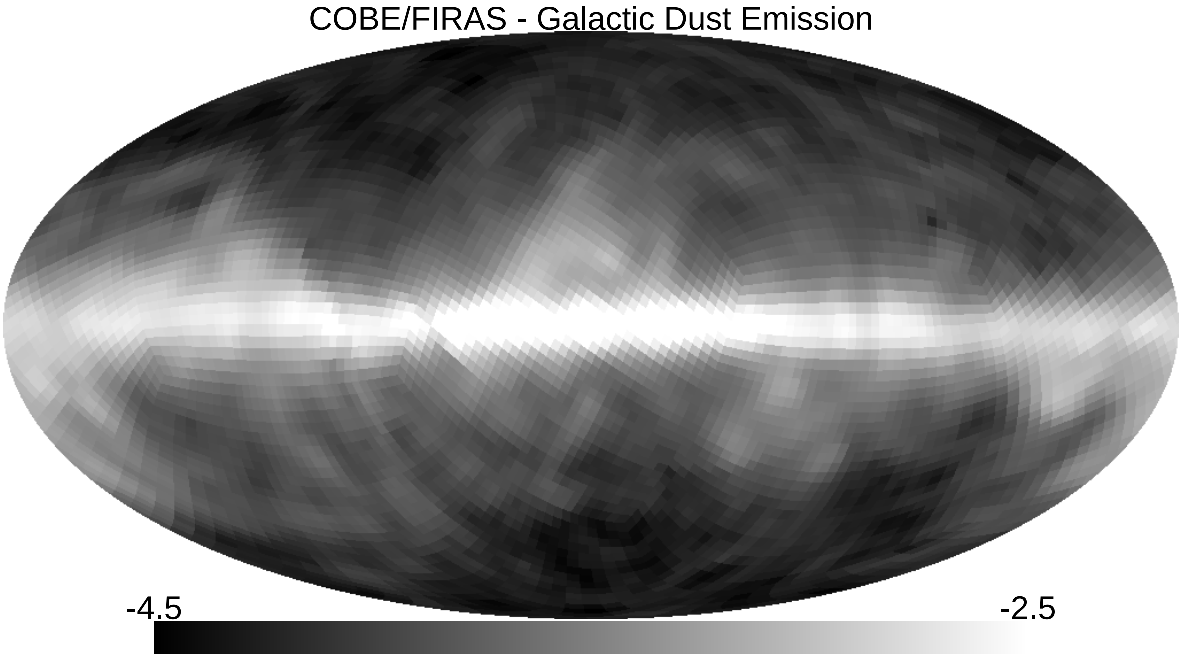

Fig. 8.3.— Galactic dust emission, logarithmic stretch — (qv. Figure 8.2)

zodiacal emission, the remaining signal is expected to be a continuum from galactic dust; narrow lines from the interstellar gas; and some possible contribution from an extragalactic "far infrared background" (FIRB). We made no explicit allowance for a FIRB in the delivered analysis data sets. (For a discussion of the FIRB and *FIRAS* data see Shafer *et al.* 1998ab; Fixsen *et al.* 1998; Hauser *et al.* 1998). The galactic dust is well characterized by a modified Planck spectrum:

$$M_{\nu p} \;=\; \tau_p \left(\frac{\nu}{\nu_0}\right)^{\alpha_p} P(\nu, T_{\mathrm{dust};p}). \qquad (45)$$

The value of τ_p is a measure of the relative dust column for each pixel, though its numerical value depends on the value used for the scaling frequency, ν_0, which we have set to 1800 GHz. The galactic spectral form is not uniform over the entire sky, which can be modeled with pixel to pixel variations in temperature, $T_{\mathrm{dust};p}$, and/or emissivity index, α_p (Reach *et al.* 1995). At higher galactic latitudes, the signal to noise is inadequate to fit both parameters, though if α is held at a fixed value for most pixels a moderately precise number for $T_{\mathrm{dust};p}$ is derived, within a few Kelvin.

The delivered analysis data sets include a pixel by pixel fit of the dust properties. For the LOWF we only determined the dust column, τ_p, fixing the temperature at 18° K. For the

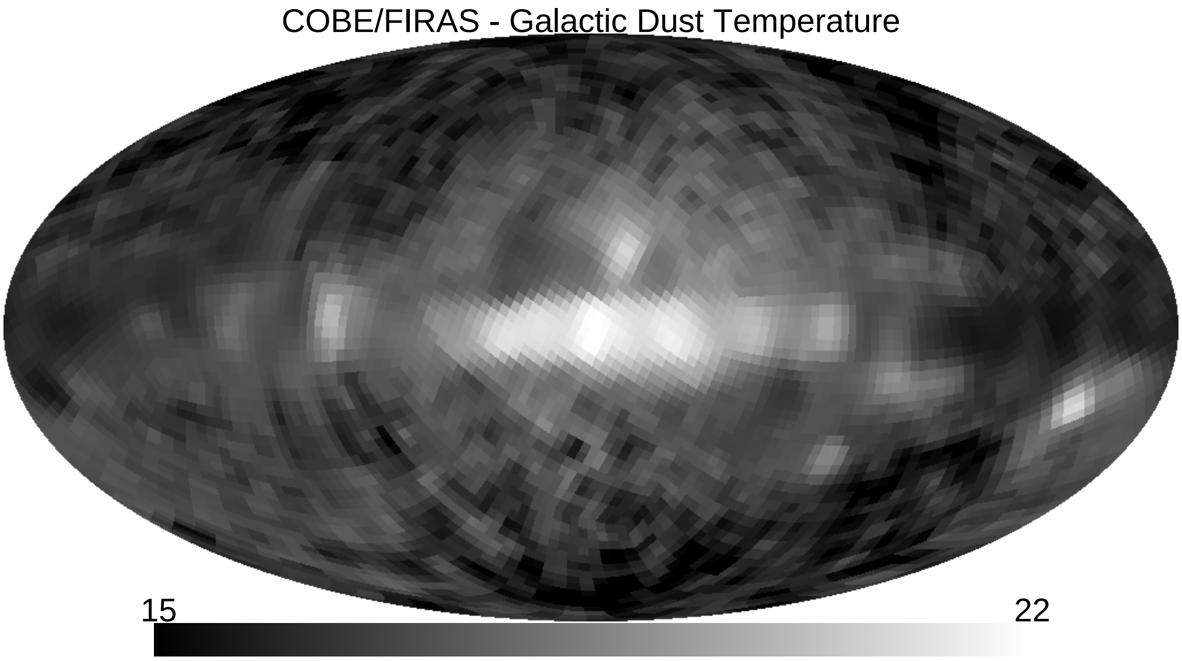

Fig. 8.4.— Galactic dust temperature — (K). From a single temperature model fit, with no removal of any isotropic far infrared cosmic background. It should be regarded as illustrative. (qv. Figure 8.1)

HIGH frequencies we fit both $T_{\mathrm{dust};p}$ and τ_p. The emissivity index was fixed at $\alpha = 2$, which was determined by a fit to the average high galactic latitude ($|b| > 20°$) spectrum. The fits did not include the frequency to frequency covariance. The delivered uncertainties include only the contribution from the detector noise. Near the Galactic plane the *FIRAS* gain uncertainties are significant.

8.4. Spectral Lines

In the *FIRAS* frequency range there are a number of narrow molecular and forbidden atomic lines produced by the interstellar gas (Table 8.1, Wright *et al.* 1991; Bennett *et al.* 1994). The *FIRAS* instrumental resolution is large compared to the expected width of the features so that they can be treated as delta function sources. Using a model for the instrument, we have produced synthetic profiles of the lines' expected contribution to the *FIRAS* spectra. These profiles, besides illustrating the intrinsic resolution of the instrument, are used in determining the intensity and significance of the line. A set of these profiles is included with the project data sets (Table 8.1).

8.4.1. Generating Line Profiles

The line profile model assumes that the interferogram of each line is a simple cosine subject to self-apodization arising from beam divergence effects. For data near the galactic center, the intensity of the 157.7 μm [C II] and 205 μm [N II] lines was so high that their signal was visible in the interferogram (as the sum of two cosines) over most of the short stroke interferogram. The self-apodization was observed in the modulation of the cosine amplitudes versus path difference. This showed that for the [C II] line, the self apodization is best described by a Gaussian with a width of 156 sample points. The [N II] line was similarly apodized with a width that scaled inversely with the line frequency. The model line interferogram is:

$$I(x) \;=\; \cos(2\pi\nu_L x)\,\exp[-\frac{1}{2}(\frac{\nu_L}{\nu_\circ}\frac{x}{\sigma})^2] \tag{46}$$

where x is the optical path difference, ν_L is the line frequency, $\nu_\circ$ is the Nyquist frequency of the instrument, and $\sigma = 5.37$ mm is the width of the Gaussian self-apodization function.

The synthetic interferogram is processed in a manner similar to that used to generate sky spectra from coadded IFGs. The interferogram is further apodized with the apodization function described in Section 5.1, then Fourier transformed to produce a synthetic single-line spectrum for an ideal instrument. At this point, the processing of the synthetic

Table 8.1: Spectral Lines Fit in *FIRAS* Data

Species	ν (cm^{-1})	ν (GHz)	λ (μm)	Excluded Frequencies‡ (cm^{-1})
CO (J=1-0)*	3.845	115.27	2600.7	–
CO (J=2-1)	7.690	230.54	1300.4	–
CO (J=3-2)	11.535	345.80	867.0	–
O_2*	14.168	424.75	705.8	–
CO (J=4-3)	15.378	461.04	650.3	–
[C I]	16.419	492.23	609.1	–
H_2O*	18.576	556.89	538.3	–
CO (J=5-4)	19.222	576.27	520.2	–
CO (J=6-5)	23.065	691.47	433.6	22.9–23.4
CO (J=7-6)**	26.908	806.68	371.6	26.5–27.5
[C I]**	27.00	809.44	370.4	
H_2O†	37.136	1113.3	269.3	36.5–37.4
[N II]	48.738	1461.1	205.2	47.5–49.2
H_2O (J=2-1)*	57.26	1716.6	174.6	–
[C II]	63.395	1900.5	157.7	62.8–64.2
[O I]	68.716	2060.1	145.5	68.3–69.2
[Si I]*	77.11	2311.7	129.7	76.9–77.4
[N II]	82.036	2459.4	121.9	81.9–82.4
CH (J=2-1)	86.38	2589.6	115.77	–

*Only upper limits determined.

**Only a single line profile with $\nu = 27$cm^{-1}is in the delivered data sets.

†Observed in absorption.

‡Frequencies (cm^{-1}) removed during continuum thermal dust spectral fits.

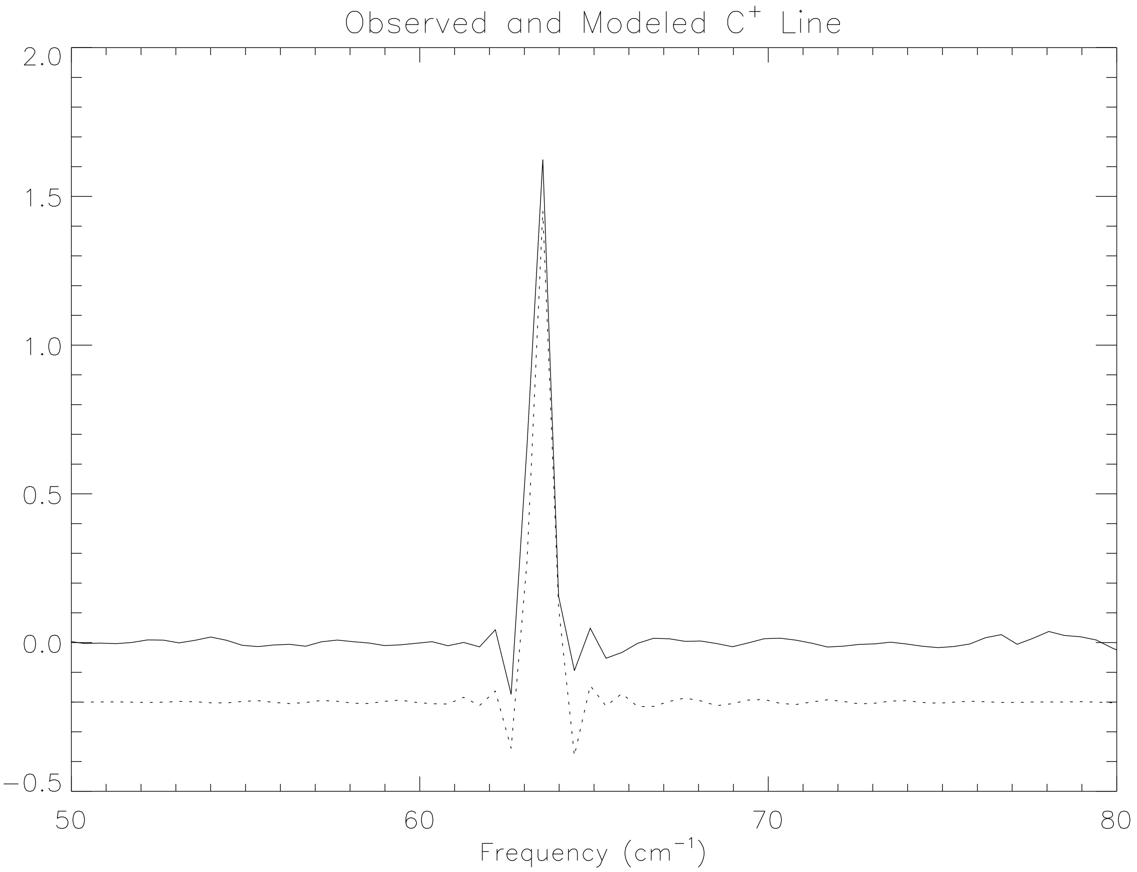

Fig. 8.5.— Observed and model [CII] line profiles — The shape of the modeled [CII] line (dotted) matches the observed (solid) shape of the [CII] line.

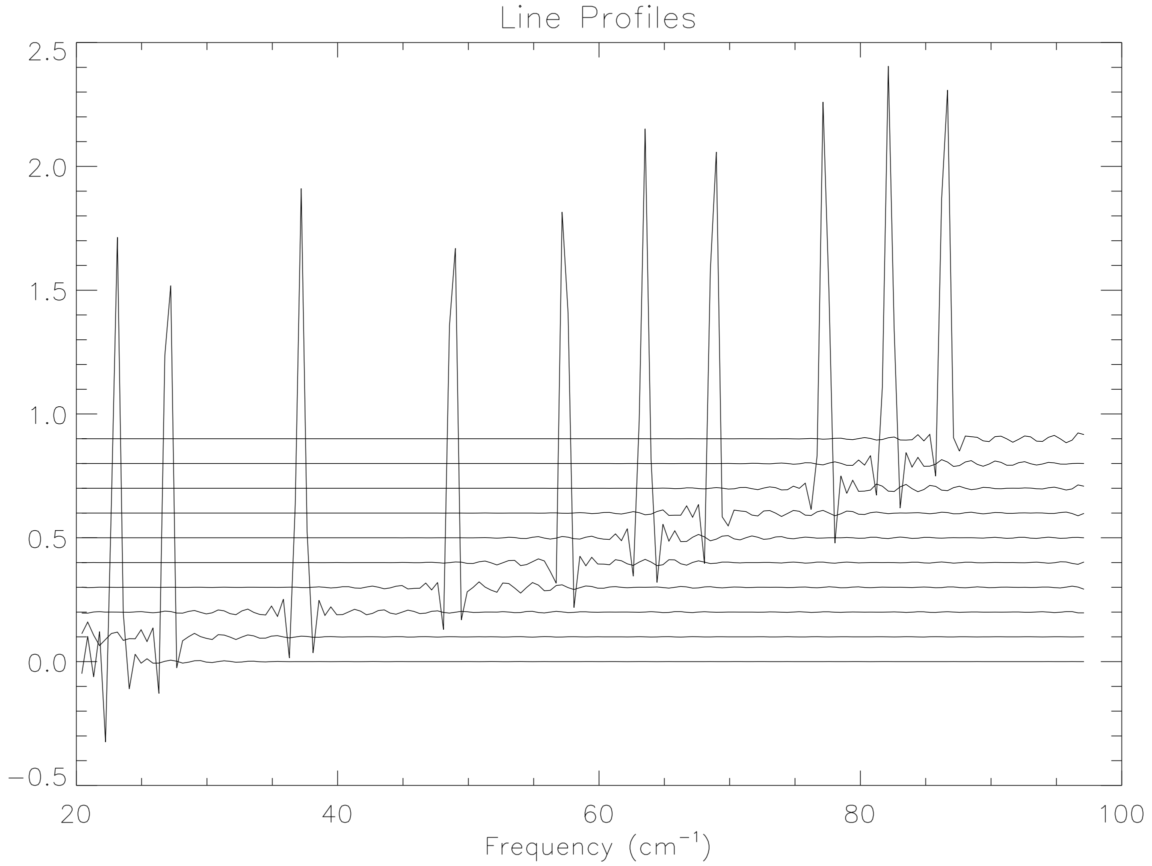

Fig. 8.6.— Model line profiles — For all lines fit in the HIGH data. Each line is offset 0.1 from the previous line for clarity.

spectrum diverges from that of the sky spectra. The synthetic spectrum is "calibrated" by multiplying the spectrum on a pointwise basis by the gain function, $G(\nu)$, described in Section 5.6 for the appropriate scan mode. The resulting calibrated synthetic line spectrum or line profile is

$$\zeta(\nu) \;=\; \Omega^{-1} G(\nu) \int_{-\infty}^{+\infty} dx \, \exp[2\pi i \nu x] \, A(x) \, I(x) \tag{47}$$

or

$$\zeta(\nu) \;=\; \Omega^{-1} G(\nu) \int_{-\infty}^{+\infty} dx \, \exp[2\pi i \nu x] \, A(x) \, \cos(2\pi \nu_L x) \, \exp[-\frac{1}{2}(\frac{\nu_L}{\nu_\circ}\frac{x}{\sigma})^2] \tag{48}$$

where $\zeta(\nu)$ is the line profile at the frequency ν, ν_L is the line frequency, Ω is a normalization factor, G is the gain function, and $A(x)$ is the apodization function.

The normalization is set so that the resulting profile in MJy sr^{-1} is for a line with incident surface brightness of one nW m^{-2} sr^{-1}. A detailed discussion of the *FIRAS* line profiles is given in Section 3 of the second *FIRAS* lines paper (Bennett *et al.* 1994), which is included as Appendix D. The delivered data sets include line profiles for each of the lines in Table 8.1. The line profiles for each low frequency channel and scan mode are contained in a single FITS file. Similarly, the line profiles for each high frequency channel and scan mode are contained in a single FITS file; the headers for these FITS files are given in Appendix G. The HIGH line profiles are plotted in Figure 8.3.

8.4.2. Modeling Spectral Lines

To determine maps of line intensity, we used data that has the zodiacal emission, the CMBR, and a model for the galactic dust removed. We fit the residual spectrum to a set of line intensities, F_m. We allow for residual continuum variations using a polynomial representation. So, we minimize for each pixel, p:

$$\chi_p^2 = \sum_\nu \left[S_{\nu p} - \sum_m F_{mp}\zeta_{m\nu} - \sum_k a_{kp} L_{k\nu} \right]^2 \frac{n_p}{C_\nu^2} \tag{49}$$

with respect to F_{mp} (the line fluxes) and the continuum residual coefficients a_{kp}, where $\zeta_{m\nu}$ are the line profiles, and $L_{k\nu}$ are the Legendre polynomials. The low frequency spectral fit uses a 4th order polynomial baseline, while the high frequency fit uses a 10th order baseline.

The line-subtracted continuum spectra are computed at each pixel using:

$$S_{\mathrm{line-sub};\nu p} = S_{\nu p} - \sum_m F_{mp}\zeta_{m\nu} \tag{50}$$

where F_{mp} are the line maps.

The *FIRAS* line emission and continuum spectrum maps are part of the data release. Their FITS headers are given in Appendix G.

The LOWF line emission skymap contains fluxes for eight low frequency lines, and the HIGH line emission skymap contains fluxes for ten high frequency lines The HIGH line emission skymap contains the fitted Galactic dust parameters. (Table 8.1). The uncertainties in the line fluxes and dust parameters are included in the line emission skymap records. These are formal uncertainties derived from the detector noise. In regions near the Galactic plane, the uncertainties introduced by the *FIRAS* gain uncertainties dominate (Section 7.2).

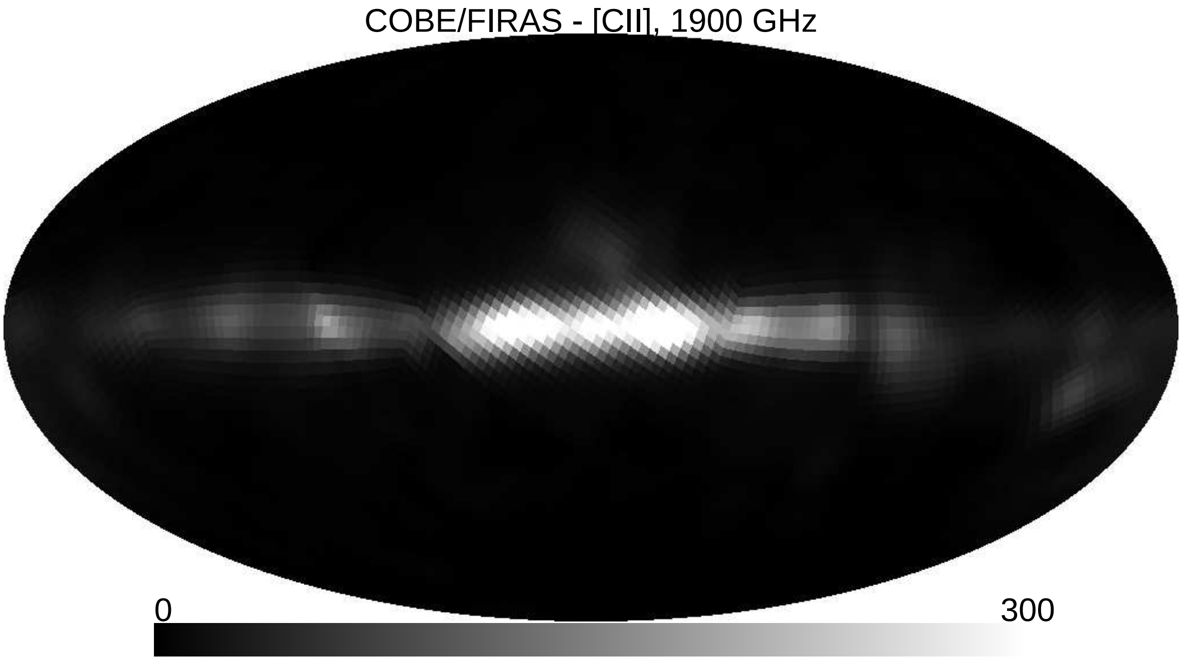

Fig. 8.7.— [CII] 1900 GHz flux map, linear stretch — ($nWm^{-2}sr^{-1}$). The strongest line detected by *FIRAS*. (qv. Figure 8.1)

Figures 8.7 and 8.8 map the [CII] 1900 GHz line flux. This is the strongest line detected by *FIRAS* (Bennett *et al.*, 1994a). Figure 8.6 maps the [NII] 1461 GHz line flux.

Figure 8.4 is a map of the fitted dust temperature. Figure 8.2 is a map of the dust emission normalization, τ. Figure 8.3 is a similar map on a log scale chosen to show the variation in the dust emission at higher Galactic latitudes.

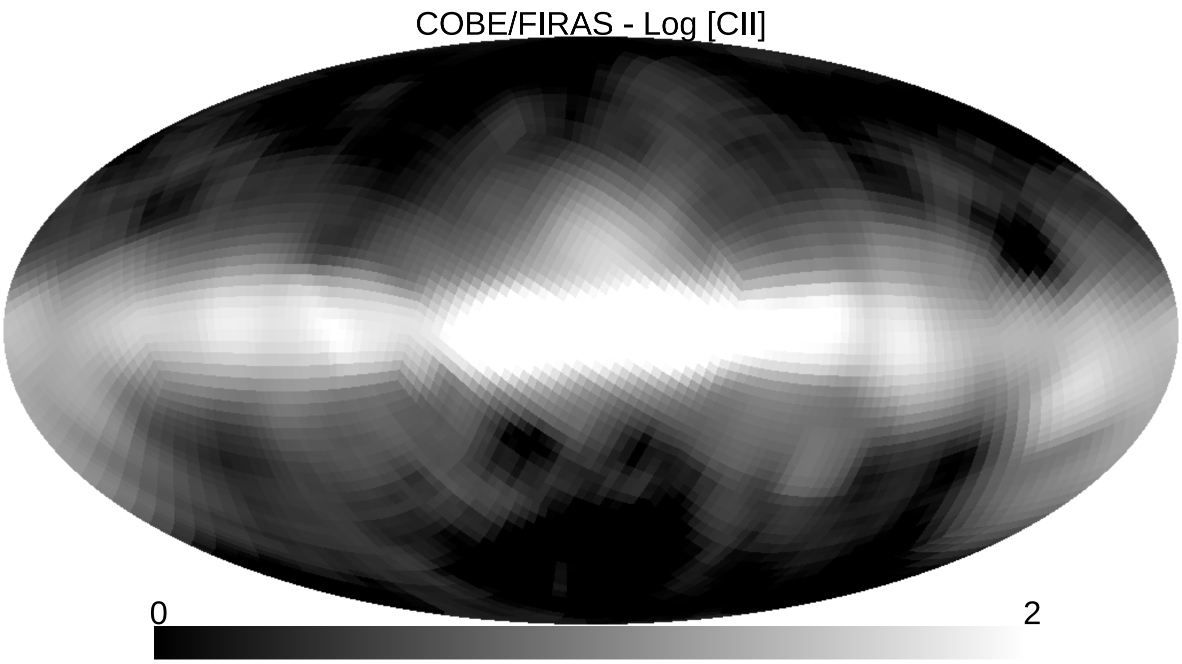

Fig. 8.8.— [CII] flux map, logarithmic stretch — (cf. Figure 8.7)

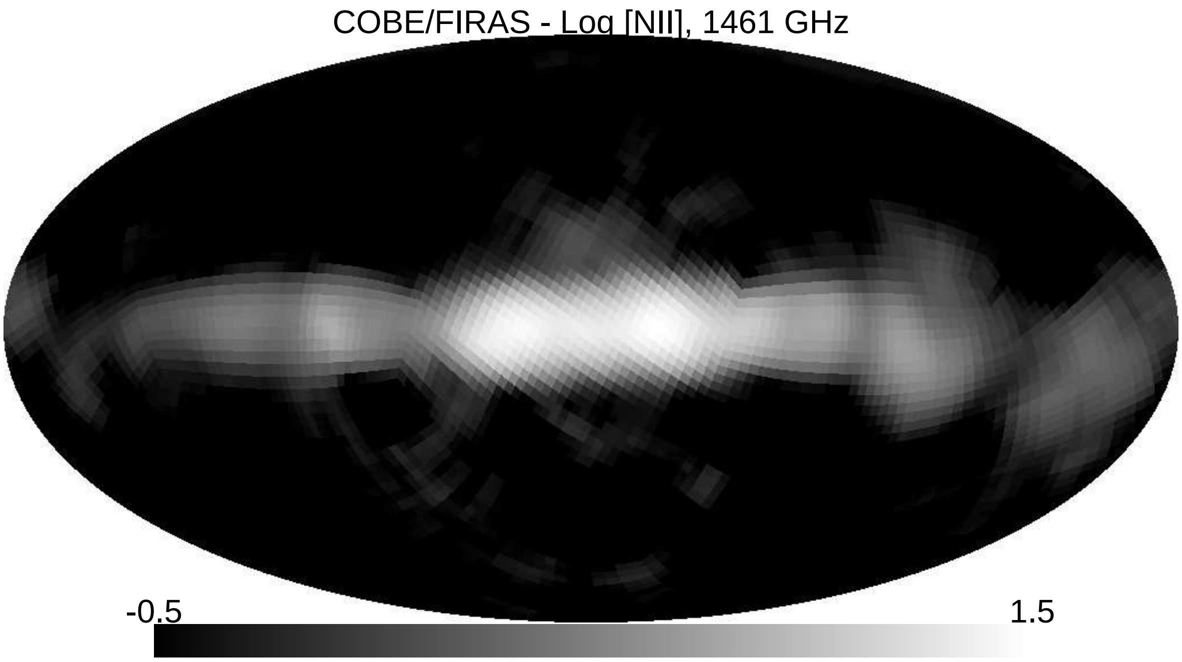

Fig. 8.9.— [NII] 1461 GHz flux map, logarithmic stretch — ($\mathrm{nWm^{-2}sr^{-1}}$). Another strong emission line associated with warm ionized as well as neutral gas.

9. Reference Data sets

The *FIRAS* data processing software has been designed so that the values of many parameters that affect how the programs operate are read from reference data sets. This design has allowed the modification of the values contained in the reference data sets without modifying data processing programs. A number of the parameters which comprise these reference data sets are contained in project data sets previously discussed. Several additional reference data sets may be of use in understanding how the *FIRAS* data have been processed. These data sets, which are in both ASCII and VAX binary file formats, are discussed below and are part of the *FIRAS* project data set release.

9.1. ASCII Format Reference Data sets

Glitch Rate Correction Parameters. Glitches add substantially to the noise in the data. Consequently, high glitch rate data are deweighted with respect to low glitch rate data. As described in Section 4.4, a linear least squares fit to the average variance and glitch rate of each calibrated sky spectrum generated a set of corrections for the numbers of IFGS in the *FIRAS* spectra. These corrections are shown in Table 4.8 and are listed in the reference data set FEX_GLTCHCOR.TXT given in Appendix I.

Vibration Correction Frequency Offset Indices. The calibration of the *FIRAS* spectra includes corrections for coherent vibrations (Sections 5.3.2 and 7.9.1). These corrections are applied at channel- and scan mode-specific offsets in frequency from the vibrations (Section 5.6). These offset indices are listed in the reference data set FEX_VIBCORRL.TXT given in Appendix I.

Actual Values of Commanded Instrument Gains. The coaddition program normalized the interferograms by the preamplifier gain before averaging them together, as discussed in Section 4.4. The values of the commanded gains for each channel are listed in the reference data set FEX_CMDGAIN.TXT given in Appendix I.

Mirror Transport Mechanism Sampling Rate. As discussed in Section 2, the Mirror Transport Mechanism samples the data at a fixed clockrate. This clockrate (681.43 Hz for on-orbit data, 673.29 Hz for preflight ground data) is listed in the reference data set FEX_SAMPRATE.TXT given in Appendix I.

Mirror Transport Mechanism Scan Times. As discussed in Section 2, the Mirror Transport Mechanism takes an average of about 40 seconds to collect the data averaged by the instrument microprocessors into single interferograms. The time required for a single sweep of the MTM in each scan mode is listed in the reference data set FEX_MTMSWEEP.TXT given in Appendix I.

Coaddition Consistency Check Thresholds. As discussed in Section 4.4, the coaddition program verified that the members of the ensembles of coaddable interferograms were in a consistent instrument state and had a consistent shape. The threshold values and tolerances used for these consistency checks are listed in the reference data set FEX_CTH.TXT given in Appendix I.

The following instrument state parameters were evaluated for the instrument state consistency check:

1. channel, scan mode, adds per group, number of MTM sweeps
2. attitude
3. bolometer voltage
4. GRT temperatures, spatial gradients, and temporal gradients

The following quantities determined whether or not secondary template formation and subtraction occurred:

1. amplitude and signal-to-noise ratio of primary template
2. amplitude and signal-to-noise ratio of secondary template.

Deglitching occurred after primary and secondary template formation and subtraction took place.

The following quantities were evaluated after deglitching for the shape consistency check:

1. interferogram noise
2. maximum number of outlier points

Minimum Number of IFGs. As discussed in Section 4.4, the coaddition program only creates templates for ensembles of interferograms containing at least three members. The minimum number of IFGs required for template formation is listed in the reference data set FEX_MINCOADD.TXT given in Appendix I.

GRT Weights for Coadded IFGs. The *FIRAS* has 24 GRTs and 8 calibration resistors mounted throughout its structure to measure the temperatures of various components of the instrument. Readings from the 24 GRTs are combined to yield ten temperatures that are used in calibrating the data. The relative weights used to combine the GRT readings for the calibration of coadded interferograms are listed in the reference data set FEX_GRTCOAWT.TXT given in Appendix I.

GRTs that are broken, are non-existent (not in the flight instrument or acting as a placeholder in the list of GRT properties), or give spurious readings have their weights set to zero, i.e. the XCAL "B" GRT and the collimator "B" GRT. GRTs that should be combined unequally are weighted accordingly. The calibration resistors also have their weights set to zero.

The Sky horn B GRT shows temperature spikes due to charged particle hits, so its weight is set to zero. The calibration program computes a lower χ^2 when the XCAL temperature is combined from the cone GRTs alone, so the XCAL A (tip) GRT has a weight of zero. The calibration program also computed a lower χ^2 when the structure temperature is read from the collimator mirror alone, so the Mirror A and Mirror B GRTs have their weights set to zero. Finally, the calibration program computes a lower χ^2 when the ICAL temperature is combined as 90% cone and 10% tip, so the ICAL A GRT has a weight of 0.1 and the ICAL B GRT has a weight of 0.9. All of the other GRTs are combined with equal weights.

GRT Weights for Raw IFGs. As discussed in Section 4.3 temperatures of various components of the *FIRAS* are used to sort raw IFGs into coaddable ensembles and to check the consistency of the members of these ensembles. The relative weights used to combine the GRT readings for the sorting and consistency checking of the raw IFGs are listed in the reference data set FEX_GRTRAWWT.TXT given in Appendix I. These weights are similar to but numerically diferent to the Coadd GRT weights, described above.

The XCAL B and Collimator B GRTs have zero weight. In general, the remaining weights are set to one since at this stage the instrument "sides" are not combined. The XCAL A and S5 GRTs are set to 0.5 for the XCAL combination, and the A side Mirror and Collimator GRTs are equally weighted for the A side structure measurement. The calibration resistors also have their weights set to zero.

GRT Low/High Current Transition Temperatures. The *FIRAS* GRTs have two temperature-dependent operating regimes. For the low-temperature regime, "low current" readings of the GRTs yield correct resistances and temperatures, while for the high-temperature regime, "high current" readings yield correct values. There are temperature transition regions for each GRT between these two regimes. The transition region midpoint and half-width for each GRT are listed in the reference data set FEX_GRTRRANS.TXT given in Appendix I.

9.2. VAX Binary Format Reference Data sets

Digital Transient Response Functions. The onboard digital filters add a transient signal to the samples at the beginning of each interferogram. As discussed in Section 4.4, the coaddition program corrects for this signal by fitting a digital transient response function to the first 128 points of the IFG and subtracting that fitted response from the IFG. The functional form of the digital transient response function is the Z-transform of the lowpass digital filter contained within the electronics transfer function (Section 5.3.1).

The *FIRAS* reference data set FEX_DTRF.DAT is a VAX binary format file which contains eight digital transient response function records, each of which is a floating-point array 128 samples long. These functions are dependent on the frequency (high or low) and the scan mode (short slow, short fast, long slow, or long fast) of the data. The Record Definition Language file for this data set is given in Appendix H.

Glitch Profiles. As discussed in Section 4.4 charged particles incident on the *FIRAS* bolometers give rise to "glitches" in the interferograms. The charged particles deposit a spike of energy in the detectors, giving rise to a delta function response on the part of the detectors. Consequently, the glitch spectrum is the product of the instrument electronics transfer function and the detector transfer function:

$$GS(\omega) \ = \ \frac{Z(\omega)}{1 \ + \ i\omega\tau} \tag{51}$$

where ω is the audio frequency of the instrument, $Z(\omega)$ is the electronics transfer function and τ is the detector time constant. The audio frequency ω in rad/s is related to the optical frequency ν in cm^{-1} by the MTM scan speed v in cm/s:

$$\omega \ = \ v\,\nu \tag{52}$$

The glitch profile in the time domain is the inverse Fourier transform of the glitch spectrum:

$$GP(x) \;=\; \Omega^{-1} \int_{-\infty}^{+\infty} d\nu \, e^{-2\pi i\nu x} \, \frac{Z(v\nu)}{1 \,+\, iv\nu\tau} \tag{53}$$

where Ω is the normalization of the transform. As discussed in Section 4.4, the coaddition program uses the glitch profiles in the time domain to "deglitch" the interferograms.

The *FIRAS* reference data set FEX_GLTCHPRO.DAT is a VAX binary format file which contains 312 glitch profile records in the time domain, each of which is a floating-point array 512 samples long. The first 510 points of the record store the actual glitch profile; point 511 is index of glitch peak position; and point 512 is offset from the glitch peak position to the "actual" peak position obtained by parabolic interpolation. The glitch profiles have been normalized so that they have zero integral and peak values of unity. The glitch profiles are dependent on the MTM scan speed, on the channel, and on the position of the glitch within the raw data stream (the data stream prior to on-orbit averaging by the instrument microprocessors, as discussed in Section 2). Consequently there are 26 profiles for each channel/scan speed combination. The Record Definition Language file for this data set is given in Appendix H.

Baseline Basis Vectors. Internal defocussing of the instrument gives rise to non-zero baselines in the interferograms at the ends of the Mirror Transport Mechanism sweeps. As discussed in Section 4.4, the coaddition program corrects for this baseline by fitting a fourth-order polynomial baseline to the coadded IFG and subtracting that fitted baseline from the IFG. The basis vectors for the polynomial baseline are the first five Legendre polynomials on a 512-point scale running between -255/256 and 1.0.

The *FIRAS* reference data set FEX_BASIS.DAT is a VAX binary format file which contains the five basis vectors, each of which is a double-precision array 512 samples long. The Record Definition Language file for this data set is given in Appendix H.

10. Reading *FIRAS* Data

Several *FIRAS*-specific FORTRAN and *IDL* programs which aid the user in reading the *FIRAS* FITS files and the *FIRAS* native VAX binary format files have been created and are available.

A complete *IDL* package of *COBE* analysis software, *UIDL* , is available through the *COBE* home page

`http://www.gsfc.nasa.gov/astro/cobe/cobe_home.html`

The *IDL* library contains the complete set of *COBE* Guest Investigator analysis tools. A FORTRAN library is also available which contains only data I/O and coordinate conversion routines.

Appendix J contains listings of the FORTRAN programs and informational headers for the IDL programs mentioned below.

The primary and binary extension headers used to create the FITS files are included in this document as Appendix G.

10.1. Reading *FIRAS* FITS Files

The *FIRAS* project data sets are stored in FITS Binary Table format, which is easily accessible by an appropriate FITS reader, such as the Fortran FITSIO package. The Fortran program FIRAS_READ uses FITSIO calls to read in the pixel number, spectrum, and spectrum sigma fields. The fields are stored in pixel list order, *i.e.* not as a rasterized sky map, so additional calls to routines that convert pixel numbers into raster or sky coordinates would be necessary to display the data as an image.

The *UIDL* routines DATAIN and FIRASMOD also read *FIRAS* FITS files: DATAIN reads the *FIRAS* spectra, number of observations, and pixel numbers from one face of the sky cube and forms a weighted average into a single spectrum; FIRASMOD reads *FIRAS* FITS files that are not pixelized.

10.2. Reading *FIRAS* Native VAX Binary Format Files

All data in the *FIRAS* native VAX binary format files were written by VAX FORTRAN software which made use of Record Definition Language (RDL) files, which are source (text) files that define data structures in a language resembling the VAX FORTRAN Structure statement. The files are listed in Appendix H.

The Record Definition Language was developed in order to place the *FIRAS* data structures into the VAX Common Data Dictionary (CDD), so that they could be accessed using COBEtrieve, a set of database access routines built on the FORTRAN-callable forms of VAX Datatrieve routines.

In order to read a *FIRAS* native VAX binary format file, the user must first find the corresponding RDL file and translate it into a programming language that supports data structures. The translated source text must then be included into the declaration section of a program which invokes the appropriate READ statement, making the *FIRAS* data available to the program. A sample program READ_FSS, which reads *FIRAS* Sky Coadd Index Records, is given in Appendix J in both VAX FORTRAN and C versions.

10.3. Using Skycube Data

COBE skymap data are pixelized using a variation of a quadrilateralized spherical projection that was originally developed for Earth science databases (Chan and O'Neill 1975). In this "*COBE* Quadrilateralized Spherical Cube" (*CSC*), the sky is projected onto a cube (called a skycube), with faces numbered 0 through 5. Face 0 contains the North Ecliptic Pole and face 5 contains the South Ecliptic Pole. The Ecliptic plane spans the remaining faces. The reference frame for the cube is defined by axes normal to faces 1 and 0. In Geocentric Ecliptic coordinates, these directions correspond to the vernal equinox and the North Ecliptic Pole.

For the *FIRAS* instrument, each face of the *CSC* projection is divided into $32 \times 32 = 1024$ squares (res 6), yielding 6144 pixels across the sky. The projection is adjusted to form equal area pixels having a solid angle of $4\pi/6144$ sr or 6.7 square degrees. Because the 7° beamwidth of the sky horn is greater than the separation between pixels (2.6° average) this binning oversamples the sky and the pixels are not independent.

All *FIRAS* skymaps are in the *CSC* projection in Geocentric Ecliptic J2000.0 coordinates. The best photometric integrity is maintained by using the data in the *CSC* projection. All of the skymaps contain the *CSC* pixel number; the spectral skymaps also contain the ecliptic, Galactic, and equatorial coordinates of the average pointing direction of the IFGs that were combined into the spectra. The discrepancy between the average pointing position and the *CSC* pixel centers is generally small (up to a few tenths of a degree). However, the average pointing coordinates should be used instead of the *CSC* pixel centers whenever possible.

The *CSC* projection for *FIRAS* data are shown in Figures 10.1, 10.2, and 10.3. Documentation and references on the *CSC* may be found in the accompanying FITS file DMR_FIRAS_SKYMAP_INFO.FITS.

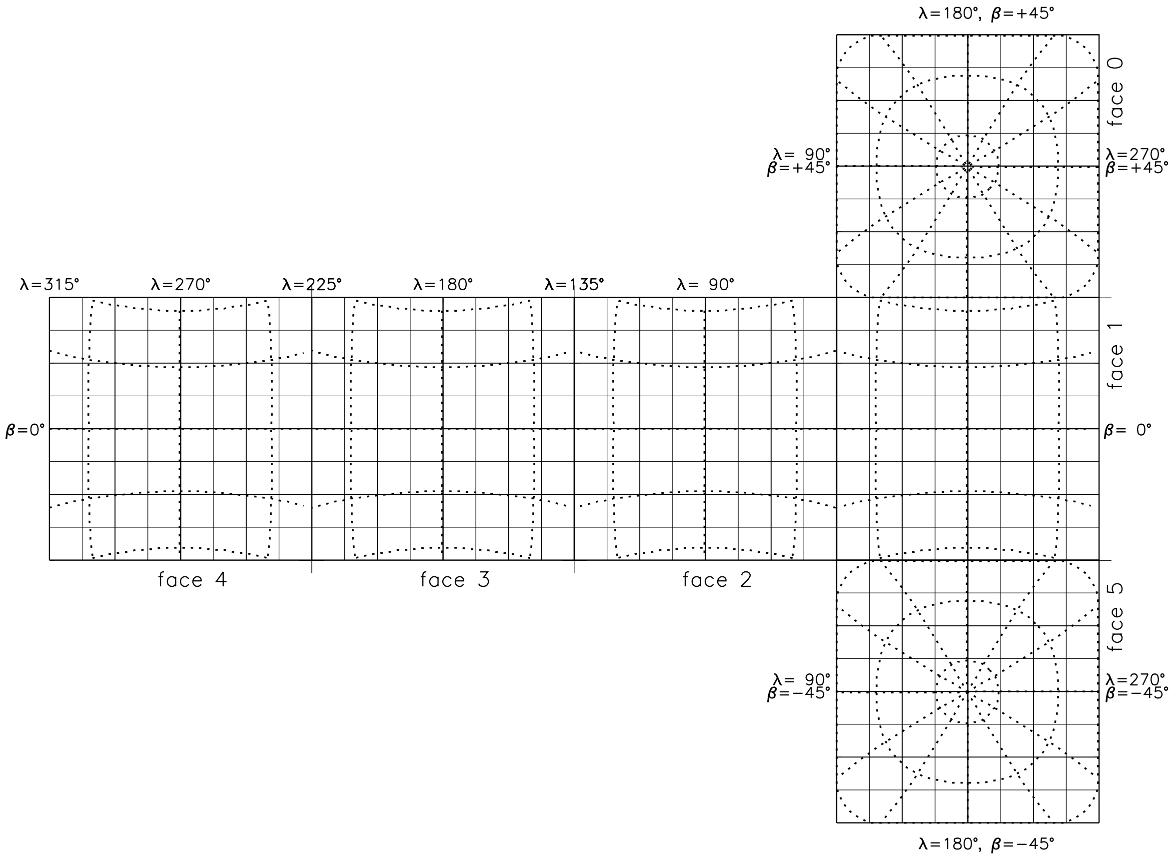

Fig. 10.1.— Unfolded, skyward-looking cube with ecliptic coordinate overlay — Dotted lines show ecliptic coordinate grid, with Δ longitude $= 30^\circ$, Δ latitude $= 20^\circ$, and coordinates indicated around cube face edges.

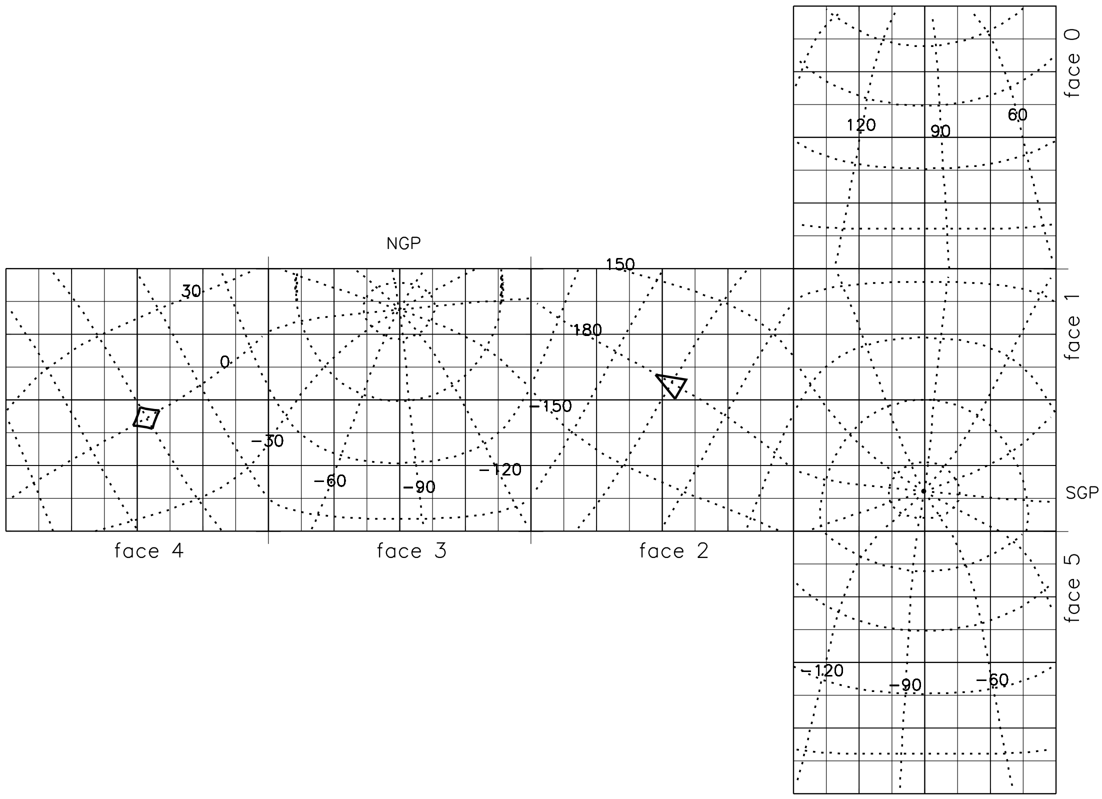

Fig. 10.2.— Unfolded, skyward-looking cube with Galactic coordinate overlay — Dotted lines show Galactic coordinate grid, with Δ longitude $= 30^\circ$, Δ latitude $= 20^\circ$. North and south Galactic poles are indicated beyond cube faces. $\diamond$ marks Galactic center and $\triangle$ marks anticenter.

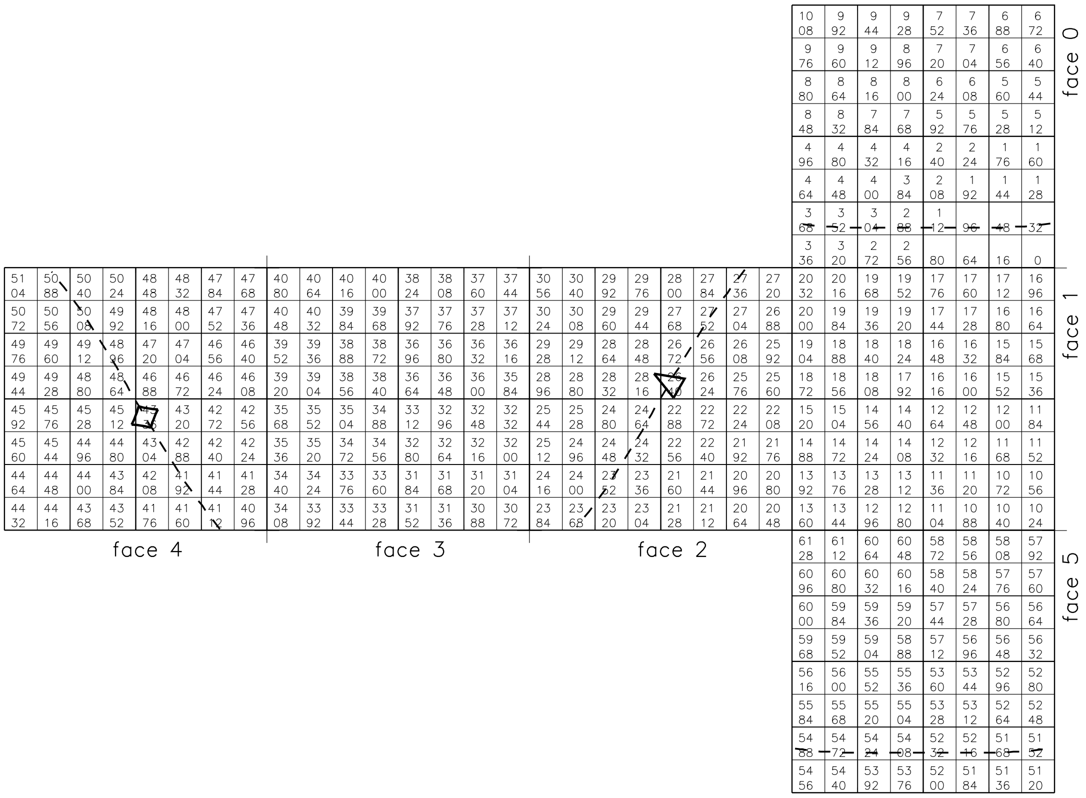

Fig. 10.3.— Unfolded *FIRAS* skycube with pixel numbering — The number in each box is the number in the lower right corner of that portion of the cube face; top row gives thousands, hundreds and bottom gives tens, units. Each face has 32 by 32 pixels. Dashed line shows Galactic plane. $\diamond$ marks Galactic center and $\triangle$ marks anticenter.

11. References

Bennett, C.L., *et al.* 1994, "Morphology of the Interstellar Cooling Lines Detected by *COBE*", ApJ, **434,** 587.

Bevington, P.R. 1969, *Data Reduction and Error Analysis for the Physical Sciences*, (McGraw-Hill, New York), p. 134.

Chan, F.K. and I.M. O'Neill 1975, "Feasibility Study of a Quadrilateralized Spherical Cube Earth Data Base", *Computer Science Corporation EPRF Tech. Rep.*, p. 2.

Dwek, E., *et al.* 1997, "Detection and Characterization of Cold Interstellar Dust and Polycyclic Aromatic Hydrocarbon Emission, from *COBE* Observations", ApJ, **475,** 565.

Fixsen, D.J., *et al.* 1994a, "Cosmic Microwave Background Dipole Spectrum Measured by the *COBE FIRAS* Instrument", ApJ, **420,** 445.

Fixsen, D.J., *et al.* 1994b, "Calibration of the *COBE FIRAS* Instrument", ApJ, **420,** 457.

Fixsen, D.J., *et al.* 1996, "The Cosmic Microwave Background Spectrum from the Full *COBE FIRAS* Data Set", ApJ, **473,** 576.

Fixsen, D.J., *et al.* 1997a, "The Spectrum of the Cosmic Microwave Background Anisotropy from the Combined *COBE FIRAS* and *DMR* Observations", ApJ, **486,** 623.

Fixsen, D.J., *et al.* 1997b, "Comparison of the *COBE FIRAS* and *DIRBE* Calibrations", ApJ, **490,** December 1, 1997.

Isaacman, R.B., S.M. Read, and W.J. Barnes 1992, "Pattern Matching in *COBE* Spectral Analysis", in *Proc. Astronomical Data Analysis Software and Systems (ADASS) I*, Astron. Soc. Pacific Conference Series 25, ed. D.M. Worrall, C. Biemesderfer, and J. Barnes (ASP, San Francisco), p. 403.

Mather, J.C. 1982, "Bolometer Noise: Nonequilibrium Theory", *Applied Optics,* **21,** 1125.

Mather, J.C. 1984a, "Bolometers: Ultimate Sensitivity, Optimization, and Amplifier Coupling", *Applied Optics,* **23,** 584.

Mather, J.C. 1984b, "Electrical Self-Calibration of Nonideal Bolometers" , *Applied Optics,* **23,** 3181.

Mather, J.C., *et al.* 1990, "A Preliminary Measurement of the Cosmic Microwave Background Spectrum by the Cosmic Background Explorer (*COBE*) Satellite", ApJ, **354,** L37.

Mather, J.C., D.J. Fixsen, and R.A. Shafer 1993, "Design for the *COBE* Far Infrared Absolute Spectrophotometer (*FIRAS*)", in *Proc. SPIE*, vol. 2019, pp. 168-179, Conf. on Infrared Spaceborne Remote Sensing, in San Diego, CA, 11-16 July 1993 (SPIE: Bellingham, WA).

Mather, J.C., *et al.* 1994, "Measurement of the Cosmic Microwave Background Spectrum by the *COBE FIRAS* Instrument", ApJ, **420,** 439.

Petuchowski, S.J. and C.L. Bennett 1993, "Galactic Fine Structure Lines: Morphologies of the Warm Ionized Medium", ApJ, **405,** 591.

Polavarapu, P., *et al.* 1997, *Principles and Applications of Polarization Division Interferometry*, (John Wiley & Sons, Ltd., Sussex).

Reach, W.T., *et al.* 1995, "Far-Infrared Spectral Observations of the Galaxy by *COBE*", ApJ, **451,** 188.

Read, S.M., *et al.* 1992, "The *COBE FIRAS* Project Pipeline Software System", in *Proc. Astronomical Data Analysis Software and Systems (ADASS) I*, Astron. Soc. Pacific Conference Series 25, ed. D.M. Worrall, C. Biemesderfer, and J. Barnes (ASP, San Francisco), p. 524.

Shafer, R.A., *et al.* 1997, "The Far Infrared Background as Measured by *COBE FIRAS* I: Limits from Dark Sky Measurements", accepted by *The Astrophysical Journal*.

Wright, E.L., *et al.* 1990, "*COBE* Attitude Refinement", in *AIP Conference Proceedings*, vol. 222, pp. 179-182, Conf. After the First Three Minutes, in Greenbelt, MD, 1990 (American Institute of Physics).

Wright, E.L., *et al.* 1991, "Preliminary Spectral Observations of the Galaxy with a 7° Beam by the Cosmic Background Explorer (*COBE*)", ApJ, **381,** 200.

Wright, E.L., *et al.* 1994, "Interpretation of the *COBE FIRAS* CMBR Spectrum", ApJ, **420,** 450.

12. Appendices

The Appendices to the *FIRAS* Explanatory Supplement are as follows:

A. Full text of Fixsen, D.J., *et al.* 1996, "The Cosmic Microwave Background Spectrum from the Full *COBE FIRAS* Data Set", ApJ 473, 576.

B. Full text of Mather, J.C., D.J. Fixsen, and R.A. Shafer 1993, "Design for the *COBE* Far Infrared Absolute Spectrophotometer (*FIRAS*)", in *Proc. SPIE*, vol. 2019, pp. 168-179, Conf. on Infrared Spaceborne Remote Sensing, in San Diego, CA, 11-16 July 1993 (SPIE: Bellingham, WA).

C. Full text of Fixsen, D.J., *et al.* 1994b, "Calibration of the *COBE FIRAS* Instrument", ApJ 420, 457.

D. Full text of Bennett, C.L., *et al.* 1994, "Morphology of the Interstellar Cooling Lines Detected by *COBE*", ApJ 434, 587.

E. Full text of Read, S.M., *et al.* 1992, "The *COBE FIRAS* Project Pipeline Software System", in *Proc. Astronomical Data Analysis Software and Systems (ADASS) I*, Astron. Soc. Pacific Conference Series 25, ed. D.M. Worrall, C. Biemesderfer, and J. Barnes (ASP, San Francisco), p. 524.

F. Full text of Isaacman, R.B., S.M. Read, and W.J. Barnes 1992, "Pattern Matching in *COBE* Spectral Analysis", in *Proc. Astronomical Data Analysis Software and Systems (ADASS) I*, Astron. Soc. Pacific Conference Series 25, ed. D.M. Worrall, C. Biemesderfer, and J. Barnes (ASP, San Francisco), p. 403.

G. Concatenated primary and extension headers for *FIRAS* FITS format files:

Coadded Calibration Interferograms

Coadded Sky Interferograms

Calibrated Calibration Spectra

Calibrated Sky Spectra

Differential Calibration Spectra

Differential Sky Spectra

Destriped Sky Spectra

Coadd-Based Zodiacal Light Model

Pixel-Based Zodiacal Light Model

DIRBE Gradient Kernel

Physical Stripes

Orthogonal Stripes

C Vector

A Vector

Covariance Matrix

Chi-Squared

Calibration Model Solution

Calibration Errors

Combined Calibration Errors

Model Selection Errors

Low Frequency Line Profiles

High Frequency Line Profiles

Low Frequency Line Map

High Frequency Line Map

Dust Spectrum Map

CMBR Temperature Map

H. Record Definition Language Files for *FIRAS* Native VAX Binary Format Files:

Sky and Calibration Data Index Records

Time-Ordered Interferogram Records

Engineering Data Records

Housekeeping and Ancillary Data Records

Engineering Mode Timing Records

Reference Datasets

I. *FIRAS* ASCII Format Reference Datasets:

FEX_GLTCHCOR.TXT: Glitch rate correction parameters

FEX_VIBCORRL.TXT: Vibration correction frequency offset indices

FEX_CMDGAIN.TXT: Actual values of commanded instrument gains

FEX_SAMPRATE.TXT: Mirror transport mechanism sampling rate

FEX_MTMSWEEP.TXT: Mirror transport mechanism scan times

FEX_CTH.TXT: Coaddition consistency check thresholds

FEX_MINCOADD.TXT: Minimum number of IFGs

FEX_GRTCOAWT.TXT: GRT weights for coadded IFGs

FEX_GRTRAWWT.TXT: GRT weights for raw IFGs

FEX_GRTTRANS.TXT: GRT low/high current transition temperatures

J. Listings for the *FIRAS* FITS and Native VAX Binary Format File Reading Programs:

DATAIN: Header for the *UIDL* program

FIRASMOD: Header for the *UIDL* program

FIRAS_READ: Listing of the FORTRAN program

READ_FSS: Listings of the FORTRAN and C programs