

THE COSMIC MICROWAVE BACKGROUND SPECTRUM FROM THE FULL *COBE*¹ FIRAS DATA SET

D. J. FIXSEN,² E. S. CHENG,³ J. M. GALES,² J. C. MATHER,³ R. A. SHAFER,³ AND E. L. WRIGHT⁴

Received 1996 January 19; accepted 1996 July 11

ABSTRACT

We have refined the analysis of the data from the FIRAS (Far-InfraRed Absolute Spectrophotometer) on board the *COBE* (*COsmic Background Explorer*). The FIRAS measures the difference between the cosmic microwave background and a precise blackbody spectrum. We find new, tighter upper limits on general deviations from a blackbody spectrum. The rms deviations are less than 50 parts per million of the peak of the cosmic microwave background radiation. For the Comptonization and chemical potential, we find $|y| < 15 \times 10^{-6}$ and $|\mu| < 9 \times 10^{-5}$ (95% confidence level [CL]). There are also refinements in the absolute temperature, 2.728 ± 0.004 K (95% CL), the dipole direction, $(\ell, b) = (264^\circ 14 \pm 0.30, 48^\circ 26 \pm 0.30)$ (95% CL), and the amplitude, 3.372 ± 0.014 mK (95% CL). All of these results agree with our previous publications.

Subject headings: cosmic microwave background — cosmology: observations

1. INTRODUCTION

The FIRAS (Far Infrared Absolute Spectrophotometer) instrument aboard the *COBE* (*COsmic Background Explorer*) satellite (Boggess et al. 1992 and references therein; Bennett et al. 1992a) was designed to measure the spectrum of the cosmic microwave background radiation (CMBR). In the simple, hot big bang model, the spectrum has a blackbody form, but it could be distorted by energy release after a redshift $z \sim 3 \times 10^6$ (Peebles 1971; Sunyaev & Zeldovich 1980). After the annihilation of positrons and the decoupling of neutrinos until $z \sim 3 \times 10^4$, the CMBR was the dominant energy field. The number of photons exceeds the number of baryons by a factor $\sim 10^9$, so excellent sensitivity is required to detect even large radiant energy releases.

Spectral distortion limits from the FIRAS were presented by Mather et al. (1994). These were based on ~ 40 days of high Galactic latitude data from a single detector and scan mode. The dipole spectrum was also determined (Fixsen et al. 1994a) from this single detector. Wright et al. (1994) found that the spectrum and dipole are consistent with a simple big bang model.

All of the previous work in this frequency range ($2\text{--}21$ cm^{-1}) was based on part of the data from a single detector. This paper uses all of the low-frequency data from the 10 month mission, except for the first month when settings were frequently readjusted to reach the optimum condition, approximately doubling the statistical weight of the results. There have been several improvements in the calibration that we note, but the basic calibration process remains the same.

¹ The National Aeronautics and Space Administration/Goddard Space Flight Center (NASA/GSFC) is responsible for the design, development, and operation of the *Cosmic Background Explorer* (*COBE*). Scientific guidance is provided by the *COBE* Science Working Group. GSFC is also responsible for the development of the analysis software and for the production of the mission data sets.

² Applied Research Corporation, Code 685.3, NASA/GSFC, Greenbelt, MD 20771.

³ NASA/Goddard Space Flight Center, Code 685, Greenbelt, MD 20771.

⁴ UCLA, Physics and Astronomy Departments, P.O. Box 951562, Los Angeles, CA 90095-1562.

The instrument was recalibrated using the method described in Fixsen et al. (1994b). The calibration was applied to the sky data, producing spectra from 2 to 21 cm^{-1} (5000–480 μm). Five important improvements were made to the calibration: (1) a bias in some pixels is corrected; (2) insufficiently sampled pixels “borrow” data from neighboring pixels in order to determine a template for deglitching (removing effects of cosmic-ray hits on the detector); (3) data with a large number of glitches (cosmic-ray hits) are downweighted relative to data with few glitches; (4) we use 320 points in the spectrum rather than 256; and (5) the data are “destriped” after the calibration. Each of these processes is described in more detail below.

2. INSTRUMENT DESIGN, OPERATION, AND DATA

The FIRAS is a polarizing Michelson interferometer, described by Mather et al. (1990) and Mather, Fixsen, & Shafer (1993). It measures the spectral difference between a 7° patch of sky and an internal blackbody. The symmetric FIRAS optics are differential, with two input and two output ports. One input port receives emission from the sky, defined by a nonimaging concentrator. The other input port receives emission from an internal reference calibrator (emissivity ≈ 0.98) with an associated concentrator. Each of the two output beams (arbitrarily labeled “left” and “right”) is split by a dichroic filter into low- and high-frequency beams, separated at 20 cm^{-1} , feeding four silicon composite bolometer detectors operated simultaneously.

An external blackbody calibrator provides the critical absolute calibration. The other elements are operated in the same manner while taking sky data or calibration data, so the *only* requirement on them is that they be repeatable. Although we verify that the emissivities of the concentrators and the internal calibrator are reasonable, the fundamental comparison is between the sky and the external calibrator. During calibration, the sky aperture is completely filled by the external calibrator with an emissivity greater than 0.99997, calculated and measured. The external calibrator is isothermal to better than 1 mK at 2.7 K according to calculation. During calibration, the sky acts as a backdrop to the external calibrator, so residual transmission is still nearly 2.73 K radiation. Most of the radiation reflected by the external calibrator was originally emitted by it and then

returned by reflection within the concentrator, contributing to the effective emissivity. The next largest source of incident radiation toward the calibrator is the emission from the concentrator, but this is also at 2.7 K and thus introduces negligible error. The only important sources of radiation on the calibrator not at 2.7 K are the detectors at 1.65 K and the general instrument at ~ 1.5 K. The measured and calculated reflectance of the calibrator for this radiation is less than 3×10^{-5} at 1 cm^{-1} and diminishes rapidly with frequency. Thus, the spectrum uncertainty of the FIRAS instrument due to the external calibrator is approximately 10 parts per million. The full story of the external calibrator will be covered in a future paper.

The temperatures of the two calibrators and the two associated concentrators are controllable from 2 to 25 K. Redundant thermometers measure the temperatures of these four temperature-controlled elements and other infrared emitters such as the moving mirrors, the mechanical structure, and the detector housings. When observing the sky, the spectrometer is operated with its output nearly nulled, by adjusting the internal calibrator temperature. This reduces sensitivity to gain errors and instrument drifts.

The path difference x between the two arms of the interferometer is varied by scanning a pair of mirrors at a constant velocity. This motion produces a modulated power $P(x)$ at the output ports. $P(x)$ is approximately the cosine transform of the net power:

$$P(x) \approx \frac{1}{2} \int_0^{\infty} dv \sum_{\text{inputs}} E_i(v) S_i(v) \cos(2\pi vx), \quad (1)$$

where E_i is the effective instrument étendue, or light-gathering power, for each emitting element, S_i is the spectral intensity of each emitter, and v is the frequency in cm^{-1} . The principal inputs are the sky (or external calibrator) and the internal calibrator, which are 180° out of phase and thus enter into equation (1) with opposite signs. Each of the outputs of the interferometer also receives emission from other components of the instrument.

The mirrors are scanned at one of two velocities (“slow” and “fast”) and over one of two lengths (“short” and “long”). An insignificant amount of data was taken in the long slow mode, so there are 12 data sets (four detectors at each of three scan modes). The long data (up to 5.9 cm excursion) were truncated to match the short data (with maximum 1.2 cm excursion) for this analysis. In retrospect, perhaps all of the data should have been taken in the short slow mode, but at the launch of COBE, the detailed nature of the CMBR was not as clear as it is now. The use of two velocities allows clear separation of some instrument effects (e.g., vibration) from optical or sky effects.

Although the high-frequency ($> 20 \text{ cm}^{-1}$) radiation was effectively excluded from the low-frequency detectors, the filter reflected low-frequency radiation into the high-frequency detectors. We do not use the six high-frequency data sets since the calibrations for these are more susceptible to systematic error. The detectors do not have equal noise. Thus, even though we include five additional data sets, we only double the statistical weight of the “left low, short slow” data set used in the previous publications.

The FIRAS points $\approx 94^\circ$ from the Sun and away from the Earth as the COBE orbits the Earth with a period of 103 minutes. The orbit plane is maintained approximately perpendicular to the direction to the Sun. This scan pattern concentrates the observing time near the ecliptic poles. We

divide the sky into 6144 equal area pixels according to a quadrilateralized spherical cube representation (O’Neill & Laubscher 1976). Each pixel has ~ 0.2 of the beam area.

3. DATA REDUCTION

The data are sorted, “deglitched,” averaged, calibrated, averaged again, and “destriped” by a series of programs. At each step, the variance is calculated, and the statistics of the results are checked. The final result is a map of 6144 pixels (76 are unobserved) with a spectrum at each pixel.

3.1. Sorting

A total of 1,411,829 good (free from transmission errors), 512-point sky interferograms were received. Approximately 20% of the sky data is rejected for various conditions (e.g., sky horn temperature, high noise, bolometer temperature too high, etc.). After sorting the 12 data sets, 11 mission periods, 6114 pixels, and several instrument states, the remaining 1,138,149 interferograms are sorted into 272,525 groups with an average of about four interferograms per group or co-add. There are an average of 187 interferograms per pixel corresponding to an integration time of ~ 30 minutes.

In the first part of the mission, the calibrator was commanded into the horn approximately 2 days each month. However, for the last 6 weeks of the mission, the calibrator was commanded into the horn approximately 3 days each week. Thus, approximately 10% of the data set is calibration data.

There are 149,260 calibration interferograms. Approximately 10% of the calibration interferograms are rejected. The fraction is lower for the calibration data because it was primarily collected near the end of the mission when we had more experience operating FIRAS. The remaining 134,739 interferograms are sorted into 3028 co-adds with an average of 44 interferograms per co-add. The calibration co-adds are larger because they are not divided into pixels. The maximum size co-add is arbitrarily set at 100.

3.2. Deglitching

A robust estimate of each co-add average is formed at each point in the interferogram, to aid in recognizing glitches. This average is subtracted from each interferogram in order to yield a “noise” interferogram. The individual interferograms are “deglitched” by using a modified “CLEAN” algorithm (Isaacman, Read, & Barnes 1992). The glitches are recognized by their shape, which is due to the bolometer and amplifier response to a delta function. Glitches are probably due to cosmic particle hits on the detectors, so they are not correlated to the position of the mirrors. Thus, they add only noise. In the raw data, most of the noise is in the glitches, so deglitching is critical to the noise.

We average the center half of the data in order to form a robust template. In the case of $4n$ interferograms, the top n points and the bottom n points are discarded, and the remaining $2n$ points are averaged. In the case of $4n + 1$ or $4n - 1$ interferograms, the top n points and bottom n points are discarded, leaving $2n + 1$ or $2n - 1$ to be averaged. While this is not exactly the center half, it is close, and there is no bias. In the case of $4n + 2$ interferograms, our previous analysis ignored the bottom $n + 1$ points and the top n , leaving $2n + 1$ points to be averaged.

Clearly this biases the template, but a simple bias in the interferogram domain leads to only a DC term in the spectral domain, and the DC term (zero frequency) is not used. Since both sky and calibration observations are broadband features, the central peak of the interferogram has most of the signal, and hence it has some real signal. This extra signal causes the “CLEAN” algorithm to infer extra “glitches” at the peak only. In the unbiased case, this regresses the data toward the mean, which is a robust average of the co-add, so this is not a problem. In the case where the template is biased, this induces a small deviation in the interferogram, which becomes a small-amplitude wide-band distortion in the spectrum. This is not enough by itself to cause a problem, because the calibration data use the same procedure and thus suffer the same distortion. However, the calibration data are collected in large groups, while the sky co-adds form small groups, since the sky data are scattered over many pixels. This bias appeared most pronounced in the sky data that used only six interferograms.

To fix this bias, we changed the template for the case with $4n + 2$ interferograms in order to discard the bottom n points and the top n points. The top point and bottom point of the remaining $2n + 2$ points are included in the average with the other $2n$ points, but with only half the weight as a better approximation to the average of the center half.

For co-adds with fewer than eight interferograms, we included interferograms from neighboring pixels (up to eight total) in order to form the robust estimate for deglitching only. When there was a sufficient number of interferograms (eight or more), we did not use neighboring pixels. After the data were deglitched, the interferograms from the original pixel were averaged like the other co-adds. This borrowing of neighboring interferograms allows us to deglitch data where there are only a few interferograms in a co-add. This allows us to add an additional 15% of the data and closes many of the gaps that are otherwise unobserved. We checked this procedure to verify that it did not introduce biases. The noise in a single interferogram is high, so it is possible to use data from a neighboring pixel to approximate the original data. The neighboring pixels are sampled by the beam, which is almost 3 pixels across, and there is little gradient in most of the sky, so the actual signal difference is small. We selected the closest interferograms (in angle) from neighboring pixels from the same time period and at the same Galactic latitude for $|b| < 20^\circ$.

The same result is obtained when we exclude the data using neighboring pixels. We also compared directly the spectra with and without using neighbors in the deglitching. Finally, we compared the spectra with and without using neighbors in the high-frequency Galactic plane data. This is a more sensitive test, because the Galaxy exhibits high gradients and sharper spectral features. In all cases, the differences were consistent with the expected random noise.

The other features of the deglitching remain as in the previous work (see Isaacman et al. 1992). As in the previous calibration, in some cases a second template was subtracted, to allow for gradients across a pixel, and the “CLEAN” procedure was iterated.

3.3. Variance Estimation

Each interferogram is assigned a statistical weight before averaging. We estimate the variance as a function of glitch rate (or glitch number, since the duration of all interferograms is the same for a given scan mode). This is modeled as

$V = a + bg$, where V is the variance, g is the glitch rate, and a and b are parameters fitted for each detector and scan mode. From the deglitching, there are estimates of the variance for each of the 272,525 co-adds. These variance estimates are fitted to $V = a + bg$ for each detector and scan mode. The weight is then modeled as $1/(a + bg)$. Although the co-add estimated variances, v , are not biased, a weight, which is the inverse of the variance, is biased high by the presence of fluctuations; $\langle w_{\text{est}} \rangle = \langle 1/v \rangle > 1/\langle v \rangle = \langle w_{\text{true}} \rangle$. This is particularly important (25% effect) for the small co-adds of the sky, but it is also significant for the larger co-adds of the calibration data. To minimize this bias, we use an entire data set to fit for a and b rather than the variance of a single co-add to estimate variance and weight. Typically, 25 glitches halves the weight, and there are typically 15 glitches in each interferogram. The weighted average interferogram is calculated and converted to a spectrum by apodization, zero padding, and Fourier transformation. The total weight of the co-add is carried along as a weight of the spectrum.

The weighted average of the covariance at each frequency was estimated for each detector and scan mode using the dispersion of these weighted spectra. This variance was used in the following calculations in order to compute the statistical weights. When fitting in the frequency domain, we need to include the correlation effects. We use the inverse of the covariance matrix in order to weight the data. The covariance between frequencies is well approximated by $C_{\nu\nu'} = \sigma_\nu \sigma_{\nu'} Q(|\nu - \nu'|)$. The expression $|\nu - \nu'|$ can take on the values (0, 0.4538, ..., 19.06) for which Q takes on the values (1.000, 0.176, -0.203, 0.145, 0.077, -0.005, -0.022, 0.032, 0.053, 0.025, -0.003, 0.007, 0.029, 0.029, 0.003, -0.002, 0.016, 0.020, 0.011, 0.002, 0.007, 0.011, 0.009, 0.003, -0.004, -0.001, 0.003, 0.003, -0.001, -0.003, 0.000, 0.003, 0.009, 0.015, 0.008, 0.003, -0.002, 0.000, -0.006, -0.006, 0.000, 0.002, 0.008).

The covariance matrix could also be calculated analytically by propagating delta functions through the apodization, Fourier transform, and calibration process. These generally agree, but there are parameters of the electronic filters, effects of self apodization, and effects of the deglitching that complicate the calculation, so we use the estimated covariance matrix.

3.4. Calibration

The Fixsen et al. (1994b) calibration procedure was used. This includes ~ 3000 parameters for each detector and scan mode. Some of the parameters are nonlinear, and the process involves considerable computation. For a more detailed description, see Fixsen et al. (1994b). Here we will only address the changes in that procedure. Rather than weighting according to the variance of each spectrum at each frequency, we used the mission average for the noise spectrum shape (dependence on frequency) for each type of spectrum (particular setting of temperatures) and the glitch weighting to scale the noise of each co-add.

The Nyquist frequency (set by the sampling rate) is 145.212 cm^{-1} . Our use of a 320-point spectrum leaves 43 significant elements in the $2\text{--}21 \text{ cm}^{-1}$ frequency range. The excursion of the mirrors is not symmetric with about 1.2 cm in one direction and 0.5 cm in the other. Thus, there is information down to 0.4 cm^{-1} rather than only to 0.56 cm^{-1} . By zero padding the short end of the interferogram out to a total of 640 points, we can use more of the informa-

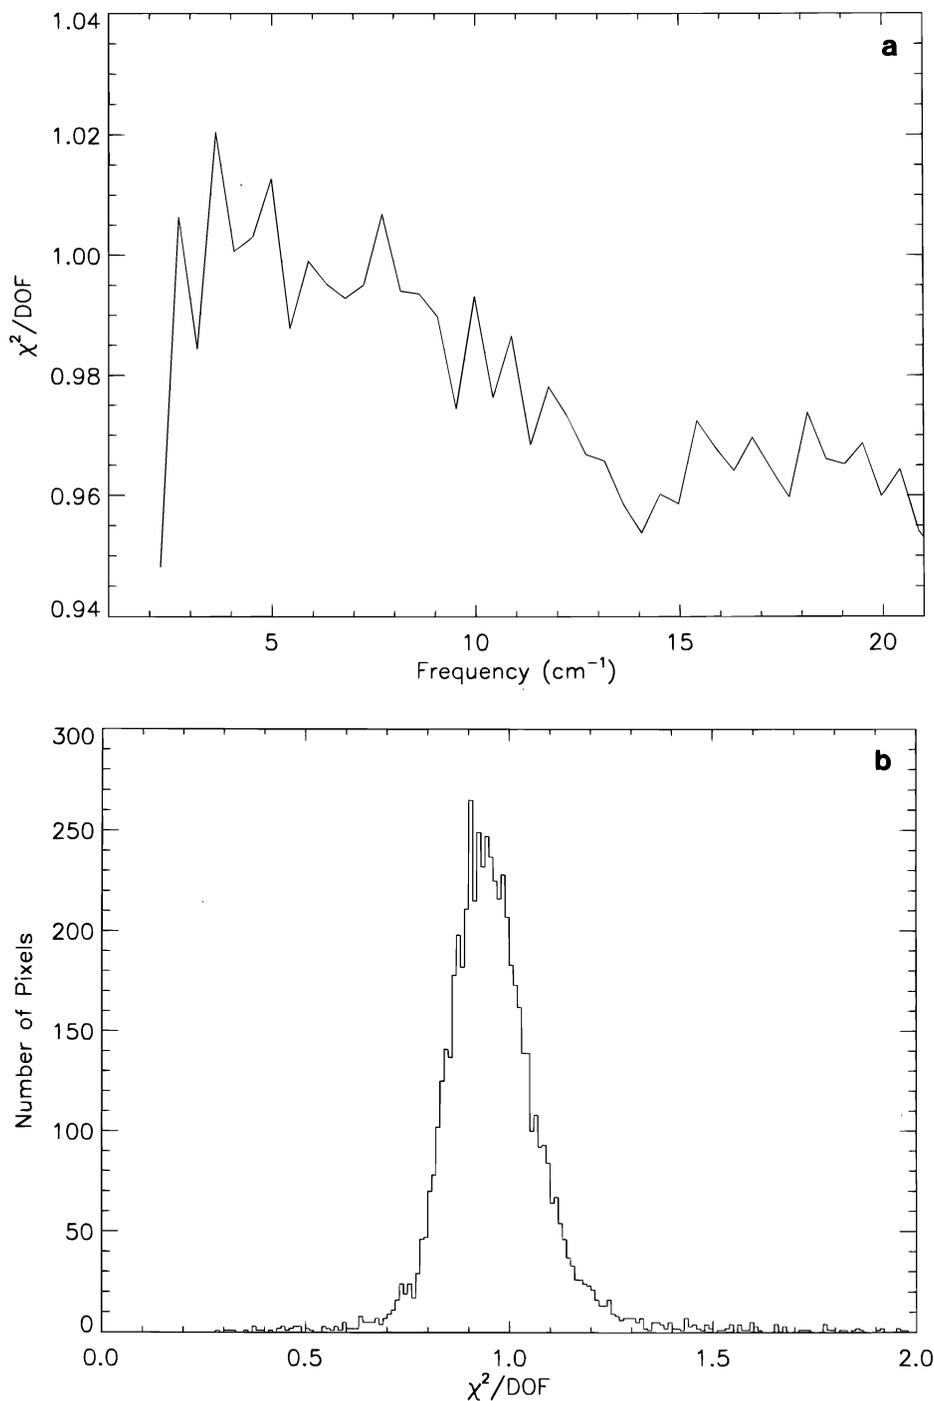


FIG. 1.—The χ^2/dof in the destriping process: (a) χ^2/dof as a function of frequency for low-frequency channels; (b) χ^2 distribution for pixels for low-frequency data.

tion from the Fourier transform that otherwise winds up in the imaginary part of the spectrum. The resulting real elements are not independent, but our analyses include correlation effects.

The truncated long and short interferograms are calibrated together. There are eight data sets or “channels” (four detectors \times two speeds). The data taken with the external calibrator in the input aperture were used to construct a calibration model for each channel. Each model includes seven parameters to describe the bolometer, two parameters for the JFET (junction field effect transistor) amplifier, six parameters for the spectrometer, and five

parameters for the unwanted vibrations. The model is nonlinear in these parameters. The temperature of each of the four “controllables” and an overall time shift for each co-add are also parameters of the fit, with constraints consistent with the thermometer readout noise (temperature dependent). The model is almost linear in these parameters. A complex number at each frequency parameterizes the emission from each of the seven sources: internal and external calibrators, sky and reference concentrators, mirrors, structure, and bolometer housing, with the Kirchhoff constraint that they sum to zero. The model is linear in these parameters. The complex number for the external calibrator

is identified as the transfer function, since the external calibrator is black and fills the beam.

The χ^2 for the calibration for each channel is given in Table 1. Excess χ^2 indicates an imperfect model of the bolometer and vibration, and other possible errors. The imperfect bolometer model is also indicated by preflight data on the bolometers alone. The $\chi^2/\text{degrees of freedom}$ (dof) for the cold, null, skylike co-adds are much closer to unity. This is expected because the hot calibrations stress the model much more than the near-null sky conditions.

The fitted calibration models are then applied to both the sky and the calibration spectra, resulting in sky co-adds and calibration co-adds in absolute flux units. After calibration, the sky spectra are adjusted to the solar system barycenter in order to account for the Earth's orbital velocity around the Sun, which would otherwise be readily detectable as an additional variable dipole of amplitude 0.3 mK. Although this further correlates the data, the correlation effect is negligible since the Earth's velocity around the Sun is $\sim 10^{-4}c$.

3.5. Destriping

After the barycenter-corrected co-adds were produced, a linear fit was made to the sky and calibration co-adds from all four low-frequency channels at each frequency. The fit includes 14 time-dependent templates, four temperature

TABLE 1
CALIBRATION χ^2 AND NUMBER OF DEGREES OF FREEDOM FOR EACH CHANNEL

PARAMETER	LOW FREQUENCY		HIGH FREQUENCY	
	Left	Right	Left	Right
Short slow	68143	53795	451509	558325
dof	34856	35276	205650	200216
Fast	20785	17198	228365	229018
dof	13688	14360	96552	95298

templates, and a parameter for the sky emission at each pixel and frequency. Thus, 18 stripe parameters and 6068 fluxes are used to fit $\sim 100,000$ co-adds separately at each frequency. By simultaneously fitting all of the data sets with a single "best" sky and a photometric model, we get the maximum amount of cross checking on each data set. By simultaneously fitting the calibration data, we effectively recalibrate the data, get the best sky signal, and get a covariance matrix to measure the correlations we have introduced by this process. The χ^2 for this fitting is shown in Figure 1.

The templates consist of two global-decaying exponentials with time constants of 61 and 153 days, and step functions for three time periods for each of the four channels. The exponentials were included because an exponential-like decay was seen in the data, both in the sky data and in the calibration data (Fixsen et al. 1994b). The cause of this is not known, but it is seen in all of the detectors and scan modes. The time-period templates were unity during a certain period and zero at other times. Two of the periods were when the concentrator temperatures were raised to 6 and 4 K, respectively. The third period covers the rest of the time when the concentrator temperatures were 2.73 K. The transition between these periods was very short and not used. The temperature of the dihedral mirrors was also used as a template because parts of the instrument had temperature fluctuations of ~ 2 K, correlated with the dihedral mirror temperature. We use a separate model for each channel because the higher temperatures may have different effects on the different detectors. The calibration corrected for these, but there is uncertainty in the correction. We reduce the residual uncertainty by using the more abundant sky data to determine the corrections.

4. THE FIRAS TEMPERATURE SCALE

Mather et al. (1994) gave the absolute temperature of the CMBR as 2.726 ± 0.010 K, with a conservative systematic

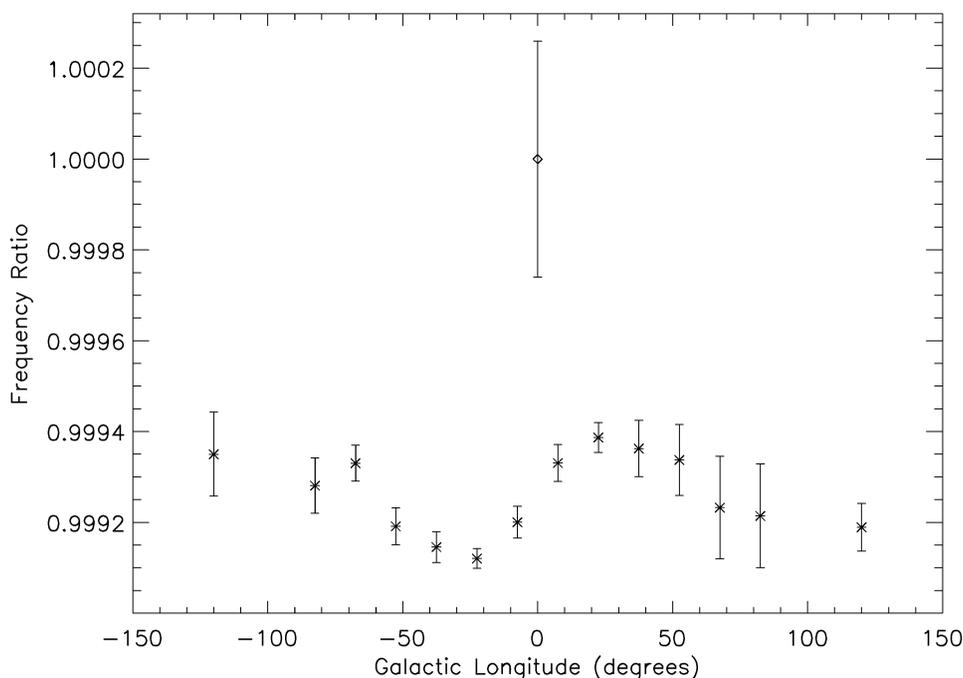


FIG. 2.—The frequency of the C^+ line as a function of Galactic longitude. Galactic rotation is clearly evident. The diamond shows the low-frequency measurement, which was used to set the frequency scale.

uncertainty estimate. They noted a discrepancy between the thermometers and the color temperature scale. We have accounted for most of the discrepancy by using the low-frequency lines to calibrate the frequency scale rather than the high-frequency scale. Since the uncertainty in going from the FIRAS temperature scale to the CMBR temperature scale is only a few μK , we discuss them together. Only the absolute temperature of the external calibrator is an issue since the other parts of the instrument only need to be stable. There are three ways of determining the CMBR temperature from the FIRAS data set:

1. We use the preflight calibration of the external calibrator thermometers, which should be good to the nominal 1 mK accuracy of the calibration specification. The task is quite simple since there exist substantial calibration data with the external calibrator at 2.73 K and other temperatures near the sky conditions. This gives a CMBR temperature of 2.730 ± 0.001 K, with the error entirely dominated by the absolute thermometry calibration error on the external calibrator. The main difficulty with this is that there is evidence that the internal calibrator thermometry has errors of ~ 3 mK and that the thermometer at the tip of the external calibrator differs from the other two by 3 mK (at 2.73 K). There is some evidence from testing the alternate calibrators on the ground that both of these problems may be due to insufficient thermal contact for the thermometers embedded in epoxy at the tip, and that the two thermometers on the external calibrator fixed to the copper base ring should not have this problem. These two were used in estimating the external calibrator temperature, and the tip temperature is not used.

2. We use the calibration data to determine the temperature scale. This effectively uses the "color" to determine the temperature and so depends on the frequency determination. We have improved the frequency determination of the FIRAS (see Fig. 2). While Mather et al. (1994) use the C^+ (63 cm^{-1}) and N^+ (49 cm^{-1}) lines, we note the possibility that the high-frequency calibration and the low-frequency calibration need not agree. A discrepancy can arise because the beam in the interferometer has divergence (5.5 HWHM on average) that, because of the cosine effect, implies a 0.5% frequency shift. However, the beam profile of the high-frequency light may be more strongly attenuated than the beam profile of the low-frequency data. To improve the frequency scale, we used a combination of CO and C^+ lines (7.69 , 11.53 , 15.38 , and 16.42 cm^{-1}) to make our low-frequency determination to $\pm 0.03\%$. The difference between the high-frequency and the low-frequency scales implies a difference of 1° in average beam divergence. The 0.03% frequency uncertainty implies a temperature uncertainty of 0.82 mK. There is an additional 0.2 mK error

in determining the color temperature once the frequency scale is set, but since this adds in quadrature, it is negligible. The result of this analysis is that the CMBR temperature is 2.7255 ± 0.0009 K.

3. We can also use the CMBR itself. If we assume the dipole is a result of a Doppler shift, the shape of the differential spectrum should be dB_ν/dT , where $B_\nu(T)$ is the Planck function. The best-fit temperature to the dipole spectrum (see below) is a CMBR temperature of 2.717 ± 0.007 K. This result depends on a differential spectrum, so offsets of the instrument calibration do not matter. Since the measurements are taken 50 minutes apart (on opposite sides of the orbit), long-term drifts (of the sort we see and correct for) also do not matter. This measurement is not completely independent because it still depends on the frequency calibration. The uncertainty is dominated by the uncertainty in fitting the Galaxy radiation that modulates the dipole signal, which is only 0.1% of the CMBR signal in the Rayleigh-Jeans region.

These three methods give answers that are within 3σ . Apparently there are additional systematic errors in one or more of these methods. We conclude that the absolute temperature of the CMBR is 2.728 ± 0.004 K (95% confidence level [CL]), entirely dominated by the systematic errors. While this is not a true statistical uncertainty, it is a useful summary of the uncertainty in the result.

5. ANALYSIS OF SKY DATA

The calibrated destriped sky spectra were then fitted to four spatial templates. We use the FIRAS data to determine the spectra of the four components. We model the data $S(\nu; \ell, b)$, where ℓ, b are Galactic coordinates and ν is frequency, as follows:

$$S(\nu; \ell, b) = I_0(\nu) + D(\ell, b)d(\nu) + G_1(\ell, b)g_1(\nu) + G_2(\ell, b)g_2(\nu). \quad (2)$$

The monopole is represented by the spectrum $I_0(\nu)$, the dipole variation is represented by the spatial distribution $D(\ell, b)$ and the spectrum $d(\nu)$, and the Galactic emission is represented by one or two spatial distributions $G_k(\ell, b)$ and corresponding spectra $g_k(\nu)$. The fit is made independently at each frequency; only the spatial variation is assumed. Each of the resulting spectra will be discussed and further modeled below. In all cases, the uncertainties due to this template fit are propagated to the following analysis.

To make this separation, the functions $D(\ell, b)$ for the dipole and $G_k(\ell, b)$ for the Galactic emission must be specified. The dipole is $D(\ell, b) = \cos(\theta)$, where θ is the angle between the observation and the maximum of the dipole, $(\ell, b) = (264^\circ 26', +48^\circ 22')$ (Bennett et al. 1996).

TABLE 2
RESULTS OF FITS

Parameter	Fit Value	Statistical Uncertainty	Systematic Uncertainty	Final Uncertainty
Galaxy temperature.....	13.3	0.6	0.8	1.0 K
Dipole amplitude	3.372	0.004	0.006	0.007 mK
Dipole temperature	2717	3	6	7 mK
Galactic latitude.....	48.26	0.11	0.10	$0^\circ 15'$
Galactic longitude	264.14	0.14	0.06	$0^\circ 15'$
CMBR temperature.....	2.728	0.00001	0.002	0.002 K
y parameter	-1	6	4	7×10^{-6}
μ parameter.....	-1	4	1	4×10^{-5}

For $G_k(\ell, b)$, we tried several combinations of predictors. Five templates for $G(\ell, b)$ are as follows:

1. A plane-parallel, $\csc |b|$ distribution: This is clearly an oversimplification of the Galaxy, but our results are independent of the Galactic template used. This provides strong evidence that our results are not contaminated by any inadequacy of the model of the Galactic emission.

2. The spatial distribution of the total power received in the high-frequency FIRAS channel above 25 cm^{-1} : This is used under the assumption that the high-frequency ($25 < \nu < 80 \text{ cm}^{-1}$) radiation is well correlated to the low-frequency ($2 < \nu < 21 \text{ cm}^{-1}$) Galactic radiation. The high-frequency channel is centered on 41 cm^{-1} with a 44 cm^{-1} effective bandwidth. Since the high- and low-frequency data are taken simultaneously through the same concentrator, there is little mismatch in sky coverage and no need for beam convolution.

3. COBE Diffuse Infrared Background Experiment (DIRBE) $240 \mu\text{m}$ map, convolved to FIRAS resolution: This has the advantage of being totally independent and having low noise, but it could suffer from beam convolution errors. The DIRBE resolution is $\sim 0.7^\circ$, and the FIRAS resolution is $\sim 7^\circ$.

4. COBE DIRBE $140 \mu\text{m}$ map, convolved to FIRAS resolution: This, in combination with the $240 \mu\text{m}$, allows adjustments for varying dust temperatures.

5. COBE DIRBE $100 \mu\text{m}$ map, convolved to FIRAS resolution: This has lower noise than the $140 \mu\text{m}$ data, but it is farther afield in frequency and may be contaminated by other effects (zodiacal emission).

For Galaxy templates two through five, we use the normalization $\langle G(\ell, b) \rangle_{|b| > 60^\circ} = 1.074$, the natural normalization of the $\csc |b|$ model. The only effect of the normalization, of course, is to rescale the Galaxy spectrum.

With this normalization, the resulting Galaxy spectrum is the sky brightness near the Galactic poles.

The principal fit for this paper uses the DIRBE 240 and DIRBE 140 combination for $G_k(\ell, b)$ and a 5° Galactic cut. There is a high χ^2 “tail” (see Fig. 1), so after the fit was made, pixels with a $\chi^2/\text{dof} > 2$ were removed. The fit was repeated with the remaining pixels, and again pixels with a $\chi^2/\text{dof} > 2$ were removed. This process was repeated until the pixels removed formed a stable set. This removes 58 pixels.

To investigate the sensitivity to these parameters, many other fits were made. We also tried Galactic cuts of 0° , 10° , 15° , 20° , 30° , and 40° . We tried fits to subsets of the data. We also tried each of the five Galactic templates individually and the DIRBE 240 and DIRBE 100 combination, the DIRBE 240 and high-frequency FIRAS combination, and the high-frequency FIRAS and DIRBE 140 combination. While one spatial Galactic template gives a reasonable result, a proper combination of two templates greatly reduces the dependence on the Galactic cut. We also varied the χ^2/dof cutoff from 1.3 to 3. For other Galactic models or other Galactic cutoffs, of course, different pixels were removed. In particular, for single component Galactic models, many pixels at low Galactic latitudes were removed.

These other fits form a background “cloud” of solutions that test the result for systematic effects. We use the rms deviation of a solution amongst these fits to estimate the size of the systematic uncertainties.

The data are weighted according to the channel variance, including glitch effects (see § 3.3). The spectra that are determined by the fit are the linear coefficients of the three or four spatial components discussed above, made independently at each frequency. The result can be viewed as three or four spectra, one for each spatial component. It is important to note that this fit makes no a priori assumptions

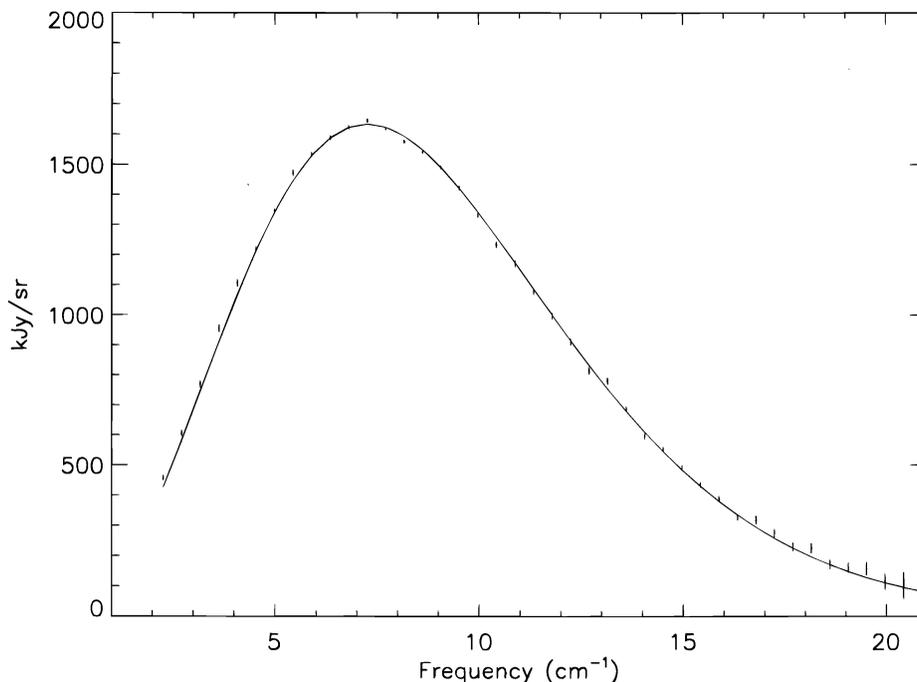


FIG. 3.—Dipole spectrum and fit to dB/dT . The vertical lines indicate plus and minus 1σ uncertainties. The peak of uniform CMBR is approximately 400 MJy sr^{-1} .

about the forms of these spectra; only their spatial distribution is posited.

5.1. Galactic Spectrum

To get a single Galactic spectrum, we add the Galactic spectra if there are more than one. The result is the average spectrum normalized at the Galactic poles. The Galactic spectrum can be approximately fitted by functions of the form $\nu^n B_\nu(T_{\text{dust}})$. Wright et al. (1991) examined the FIRAS Galaxy spectrum and found that it was fitted by $n = 1.65$ and $T_{\text{dust}} = 23.3$ K. The Galactic spectrum derived here, which is determined by high Galactic latitude data, produces a lower temperature, $T_{\text{dust}} = 13 \pm 1$ K, with n fixed at 2 and an emissivity of $\sim 10^{-4}$ at 60 cm^{-1} . With only the low-frequency data used here, we cannot fit for both n and the temperature. About half of the temperature difference is due to the change in index from 1.65 to 2, and half from the change to the higher latitudes. Even when emphasizing the high Galactic latitude regions, the CO and C⁺ lines are still evident.

When using the DIRBE 240 and DIRBE 140 templates to fit the Galactic radiation, we note that the DIRBE 140 template produces a spectrum that is negative for all frequencies. This is an indication of the Galactic temperature variation. Where the temperature is high, the DIRBE 140 is more intense relative to the DIRBE 240, but this is precisely where the DIRBE 240 is more intense relative to the low frequencies under scrutiny, and hence the negative DIRBE 140 spectrum makes a “color” correction to the DIRBE 240 estimate. The strong correlation between DIRBE 240 and DIRBE 140 ensures that the sum is positive everywhere.

The Galactic fit to $\csc |b|$ has more radiation at the Galactic poles than the fit to the other Galactic models, indicating that the Sun is in a local “hole” in the Galaxy. This local “hole” is approximately $\frac{2}{3}$ as dense as a typical Galactic location.

The dust can be modeled with other forms as well. Most of the information on the Galactic regions is at higher frequencies. A full discussion of the Galaxy requires the high-frequency data (see Bennett et al. 1992b, 1994 and Reach et al. 1995); here the Galactic radiation is treated as a contaminant of the CMBR.

5.2. Dipole Spectrum

We fitted the dipole spectrum to the derivative of a Planck function, $T_{\text{amp}} dB_\nu(T_{\text{cmb}})/dT$, where we fix $T_{\text{cmb}} = 2.728$ K and allow T_{amp} to vary. The best-fit value of the dipole amplitude is $T_{\text{amp}} = 3.369 \pm 0.004$ mK. However, the value is dependent on the Galactic cutoff and the Galaxy model. In our “cloud” of solutions, we find an rms variation of 0.006 mK. Adding this in quadrature gives an uncertainty of 0.007 mK. We also make a 0.003 mK correction for the finite beam size, which is approximately 7° FWHM. The result is summarized in Table 2. This amplitude of the dipole is in agreement with $3.353 \pm .024$ mK measured by the COBE DMR (Bennett et al. 1996), derived from observations at 1.05, 1.77, and 3.00 cm^{-1} . The DMR beam is also approximately 7° , but the details of the beam patterns are different. We do not include possible effects of the octopole and higher harmonics, although, because of our nonuniform sky coverage, these might perturb our result. Bennett et al. 1996 includes these effects.

TABLE 3

DIPOLE SPECTRUM AND RESIDUALS (INTENSITIES IN kJy sr^{-1})

Frequency (cm^{-1})	Measurement	Residual (kJy sr^{-1})	Uncertainty (1σ)
2.27	458	33	7
2.72	605	23	9
3.18	770	23	12
3.63	956	45	11
4.08	1106	38	11
4.54	1218	5	10
4.99	1346	5	9
5.45	1473	26	9
5.90	1534	5	8
6.35	1588	1	7
6.81	1623	2	6
7.26	1645	12	6
7.71	1619	-3	5
8.17	1576	-17	5
8.62	1541	-7	5
9.08	1490	1	6
9.53	1423	2	7
9.98	1331	-12	8
10.44	1234	-26	9
10.89	1171	-3	11
11.34	1077	-10	11
11.80	996	-4	11
12.25	910	-5	11
12.71	814	-18	11
13.16	780	27	11
13.61	687	8	10
14.07	595	-14	10
14.52	551	7	9
14.97	492	7	9
15.43	435	5	10
15.88	388	8	10
16.34	326	-9	11
16.79	318	24	13
17.24	275	18	14
17.70	230	5	15
18.15	224	28	16
18.61	172	2	16
19.06	162	14	17
19.51	157	30	20
19.97	113	3	27
20.42	102	7	43
20.87	161	79	76
21.33	45	-25	139

The dipole spectrum and the fitted Planck derivative are shown in Figure 3 and in Table 3. The weighted deviations, $[\sum_\nu (\Delta_\nu/\sigma_\nu)^2/\sum_\nu (1/\sigma_\nu^2)]^{1/2}$, from the blackbody derivative dipole spectrum are 14 kJy sr^{-1} rms, smaller than the deviations of the monopole from its fitted blackbody, primarily because the dipole is less susceptible to systematic errors.

By choosing the monopole temperature as the point to evaluate dB_ν/dT , we have forced the dipole temperature to be that of the monopole. If we fitted both T_{amp} and T_{cmb} , we obtain $T_{\text{cmb}} = 2.717 \pm 0.003$ K (1σ statistical uncertainty). However, the result is sensitive to the Galactic cut. We add a systematic uncertainty of 0.006 to account for this effect. This results in a final estimation of $T_{\text{cmb}} = 2.717 \pm 0.007$ K (1σ combined uncertainty). The agreement of the dipole temperature with the monopole temperature is consistent with the conventional Doppler interpretation of the dipole. It is also a confirmation that the overall temperature scale and calibration for the FIRAS are consistent.

We can also find the dipole direction from the FIRAS data with a different fit. Here we reverse the assumptions of the principal fit. We assume three spectra, a CMBR temperature and two Galactic dust optical depths, one for each

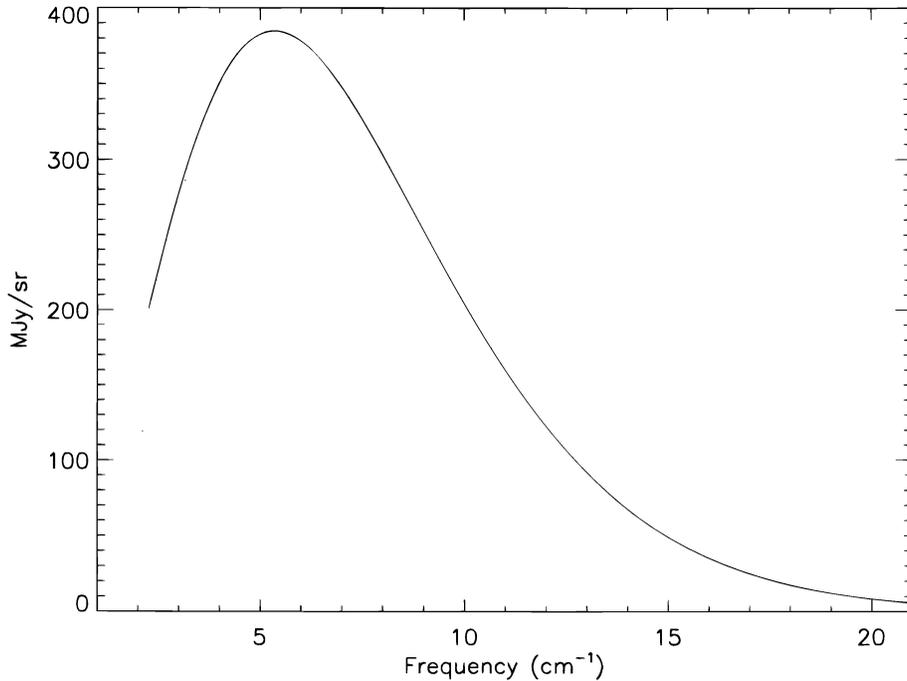


FIG. 4.—Uniform spectrum and fit to Planck blackbody (T). Uncertainties are a small fraction of the line thickness.

of the spectra derived from the DIRBE templates in the principal fit, $g_k(\nu)$, and fit for their intensities at each pixel. The CMBR temperature assumes a Planck spectrum. These spectra were chosen to match the shape of the spectrum in the data. This yields maps of the CMBR temperature and dust intensities. A monopole plus three dipole components are then fitted to the resulting temperature map.

The vector sum of the dipole coefficients points in the direction $(\ell, b) = (264^\circ.14 \pm 0.15, 48^\circ.26 \pm 0.15)$, consistent with the direction from the DMR results. Data for $|b| < 10^\circ$ were excluded from the dipole fit because of the potential inaccuracy of the model of the Galaxy. The direction is particularly sensitive to the Galaxy because it is almost orthogonal to the direction of the Galactic center. Galactic variations in spectral shape as a function of longitude couple into the angle directly, while for the fixed angle case, they come in as second-order terms for the same reason.

6. THE COSMIC SPECTRUM

The monopole spectrum from equation (2) (see Fig. 4) is well fitted by a Planck blackbody spectrum, and deviations are small, consistent with the earlier FIRAS results within their larger uncertainties (Shafer et al. 1991; Cheng 1992; Mather et al. 1994). To determine or constrain any deviations from a blackbody, let us consider a generic cosmological model $S_c(\nu; p)$, where p is some cosmic parameter quantifying the deviation from a blackbody, such as the Kompaneets y parameter for Comptonized spectra or the dimensionless chemical potential μ for a Bose-Einstein photon distribution. Because the deviation is small, a linear fit

$$I_0(\nu) = B_\nu(T_0) + \Delta T \frac{\partial B_\nu}{\partial T} + G_0 g(\nu) + p \frac{\partial S_c}{\partial p} \quad (3)$$

can be performed on the unknown parameters p , G_0 , and ΔT . The first two terms are the Planck blackbody spectrum, with the temperature $T_0 + \Delta T$. It is important to have the second term in order to properly estimate the uncertainty since the ΔT is strongly correlated with the resulting p (95% in the case of the Bose-Einstein distortion). The third term allows for Galactic contamination to remain in the monopole spectrum. The final term is the modeled deviation. We fit either the Kompaneets parameter or the chemical potential, but the two are too similar to fit simultaneously. The uncertainties are propagated from the template fits, and the correlation between the $g(\nu)$ and $\partial S_c / \partial p$ increases the uncertainty of G_0 and p .

6.1. Galactic Contamination

Most of the Galactic emission has been removed, but there is a small residual contamination. We use either the $\sum_k g_k(\nu)$ derived from the all-sky data set (see § 4) or the $\nu^2 B_\nu(T)$, and we fit a temperature and an emissivity. The $\nu^2 B_\nu(T)$ model with a temperature of 9 K produces a lower χ^2 , and we use this model for the analysis in Table 4. The problem is to fit the emission missed by the DIRBE maps, not the total Galactic dust. One possible interpretation of this fit is a 9 K Galactic halo or a cosmic background (uniform component) with a spectrum similar to the Galactic spectrum. One must be cautious, however, since subtle variations in dust temperature or emissivity can produce similar effects.

The determination of $g(\nu)$ is dominated by low Galactic latitude emission, and there is some variation from this form at higher latitudes. We vary the Galactic latitude cutoff used in deriving $I_0(\nu)$ in equation (3) in order to test the effect of variations in the Galactic spectrum from $g(\nu)$. Variations greater than the statistical uncertainty in any derived parameters, such as the cosmological term p , would most likely be due to an inadequacy in our Galactic model.

TABLE 4
RESIDUAL OF MONOPOLE SPECTRUM
(INTENSITIES IN kJy sr^{-1})

Frequency (cm^{-1})	Residual (kJy sr^{-1})	Uncertainty (1σ)	Galaxy (at Pole)
2.27	5	14	4
2.72	9	19	3
3.18	15	25	-1
3.63	4	23	-1
4.08	19	22	3
4.54	-30	21	6
4.99	-30	18	8
5.45	-10	18	8
5.90	32	16	10
6.35	4	14	10
6.81	-2	13	12
7.26	13	12	20
7.71	-22	11	25
8.17	8	10	30
8.62	8	11	36
9.08	-21	12	41
9.53	9	14	46
9.98	12	16	57
10.44	11	18	65
10.89	-29	22	73
11.34	-46	22	93
11.80	58	23	98
12.25	6	23	105
12.71	-6	23	121
13.16	6	22	135
13.61	-17	21	147
14.07	6	20	160
14.52	26	19	178
14.97	-12	19	199
15.43	-19	19	221
15.88	8	21	227
16.34	7	23	250
16.79	14	26	275
17.24	-33	28	295
17.70	6	30	312
18.15	26	32	336
18.61	-26	33	363
19.06	-6	35	405
19.51	8	41	421
19.97	26	55	435
20.42	57	88	477
20.87	-116	155	519
21.33	-432	282	573

Indeed, if the χ^2 cutoff is raised and the $\text{csc } |b|$ model is used, such effects can be seen.

When p is fixed at 0, i.e., there is only a Galactic deviation fit, the weighted rms deviation in the frequency range 2–21 cm^{-1} is 50 parts per million (ppm) of the peak brightness. The weighted rms is defined as $[\sum_v (\Delta_v/\sigma_v)^2/\sum_v (1/\sigma_v^2)]^{1/2}$, where Δ_v is the difference of the model and the data at spectrum point v , and σ_v is the uncertainty of the spectrum point v . The formal χ^2 per dof = 46/40 = 1.15. The residuals, with the errors, are shown in Figure 5 and listed in Table 4. The best-fit temperature is 2.728 K, which is determined by our choice of frequency and thermometry scales (see § 4). When comparing this with values from other experiments, the systematic thermometry uncertainty of 2 mK (1σ) should be included. The detector noise part of the temperature uncertainty is only 10 μK and is entirely negligible.

6.2. Bose-Einstein Distortion

There are two likely forms for distortions of the primeval spectrum produced during the plasma epoch before recom-

bination, and we compute independently a limit for each. Energy release or conversion in the redshift range $10^5 < z < 3 \times 10^6$ produces a Bose-Einstein distribution, where the Planck law is modified by a dimensionless chemical potential μ (Zeldovich & Sunyaev 1970):

$$S_\mu(\nu; T, \mu) = \frac{2hc^2\nu^3}{e^{x+\mu} - 1}, \quad (4)$$

where $x = hc\nu/kT$, and ν is measured in cm^{-1} . The linearized deviation of S_μ from a blackbody is the derivative of equation (4) with respect to μ :

$$\frac{\partial S_\mu}{\partial \mu} = \frac{-T_0}{x} \frac{\partial B_\nu}{\partial T}. \quad (5)$$

The current FIRAS result is $\mu = -1 \pm 4 \times 10^{-5}$, or a 95% CL upper limit of $|\mu| < 9 \times 10^{-5}$. This result and those following are summarized in Table 2, including the effects of systematic uncertainties and the variation of the result according to the Galactic model used in the fit.

6.3. Compton Distortion

Energy release at later times, $z < 10^5$, produces a Comptonized spectrum, a mixture of blackbodies at a range of temperatures. In the case of nonrelativistic electron temperatures, this spectrum is described by the Kompaneets (1957) equation, parameterized by the value of y (Zeldovich & Sunyaev 1969):

$$y = \int \frac{k(T_e - T_\gamma)}{m_e c^2} d\tau_e, \quad (6)$$

where T_e , T_γ , and τ_e are the electron temperature, the CMBR photon temperature, and the optical depth to electron Compton scattering, respectively. The distortion will be of the form (Zeldovich & Sunyaev 1969)

$$\frac{\partial S_y}{\partial y} = T_0 \left[x \coth \left(\frac{x}{2} \right) \right] - 4 \frac{\partial B_\nu}{\partial T}. \quad (7)$$

The results are $y = -1 \pm 6 \times 10^{-6}$. There is some dependence on the Galactic cut and on the Galactic model. This indicates either a residual Galactic contamination or a nearly uniform source with a spectrum similar to the Galactic spectrum. We include a systematic uncertainty of 4×10^{-6} to account for these effects for a final 95% CL upper limit of 15×10^{-6} .

The limit for $|\mu|$ is about $\frac{1}{3}$ as large as that given in Mather et al. (1994), and the limit on $|y|$ is about $\frac{1}{2}$ as large. The $|y|$ is more strongly coupled to the residual Galactic contamination. These new values place stringent constraints on theories of the early universe and the development of cosmic structure.

7. INTERPRETATION

The cosmological interpretation of the limits on y and μ are developed by Wright et al. (1994). We summarize their conclusions here briefly. Since the FIRAS spectrum with the Galactic dust emission removed is close to a blackbody, strict limits can be set on energy release after the big bang. For energy release during the Bose-Einstein period, the limit is $\Delta U/U = 0.71\mu$, where U is the energy of the cosmic background and ΔU is the energy converted from other forms. For energy release after that but before decoupling,

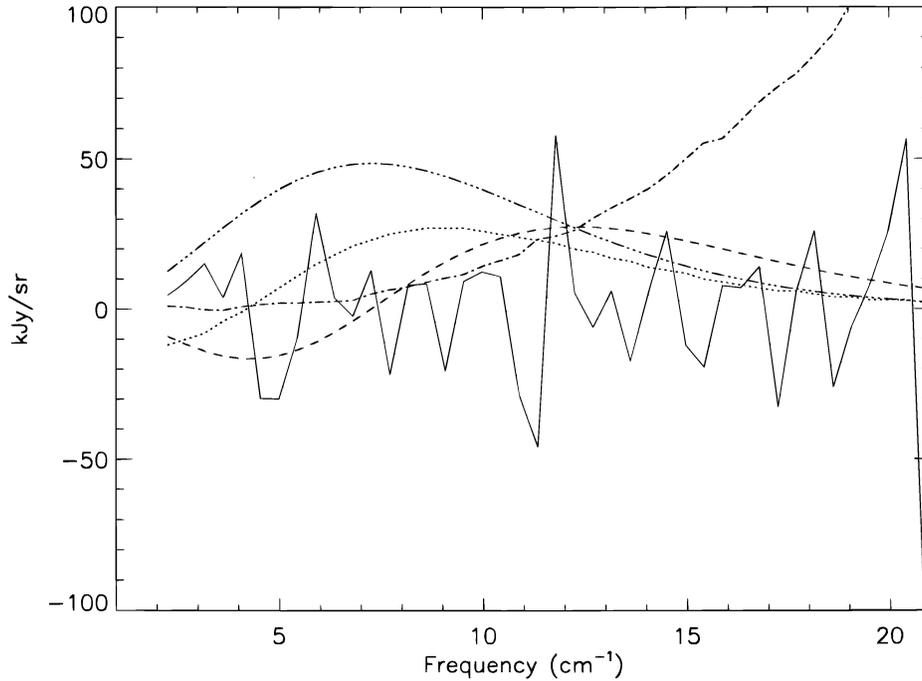


FIG. 5.—FIRAS-measured CMBR residuals, $I_0 - B_\nu(T_0) - \Delta T(dB/dT) - G_0 g(\nu)$ (solid line). Spectrum model components: the maximum allowed distortions (95% CL) $y = 15 \times 10^{-6}$ (dashed line) and $|\mu| = 9 \times 10^{-5}$ (dotted line), the Galaxy spectrum $g(\nu)$ scaled to $\frac{1}{4}$ the flux at the Galactic pole (dot-dashed line), and the effects of a 100 μ K temperature shift in T_0 , 0.0001 dB/dT , (Dot-dot-dot-dashed line).

the limit is $\Delta U/U = 4y$. The result can be summarized as a limit on $\Delta U/U$ as a function of redshift shown in Figure 6.

8. SUMMARY AND CONCLUSIONS

The FIRAS spectrum of the cosmic microwave background radiation agrees with a blackbody spectrum to high accuracy. The CMBR monopole and dipole spectra are the

result of fitting a model, including a dust map derived from the DIRBE data. The CMBR dipole has a spectrum consistent with its thermal origin and a Doppler shift. The dipole itself has a differential thermal spectrum the temperature of which, 2.717 ± 0.007 K, agrees with the monopole temperature. The Doppler shift implies that the Sun's peculiar velocity relative to the comoving frame is 371 ± 1 km s^{-1}

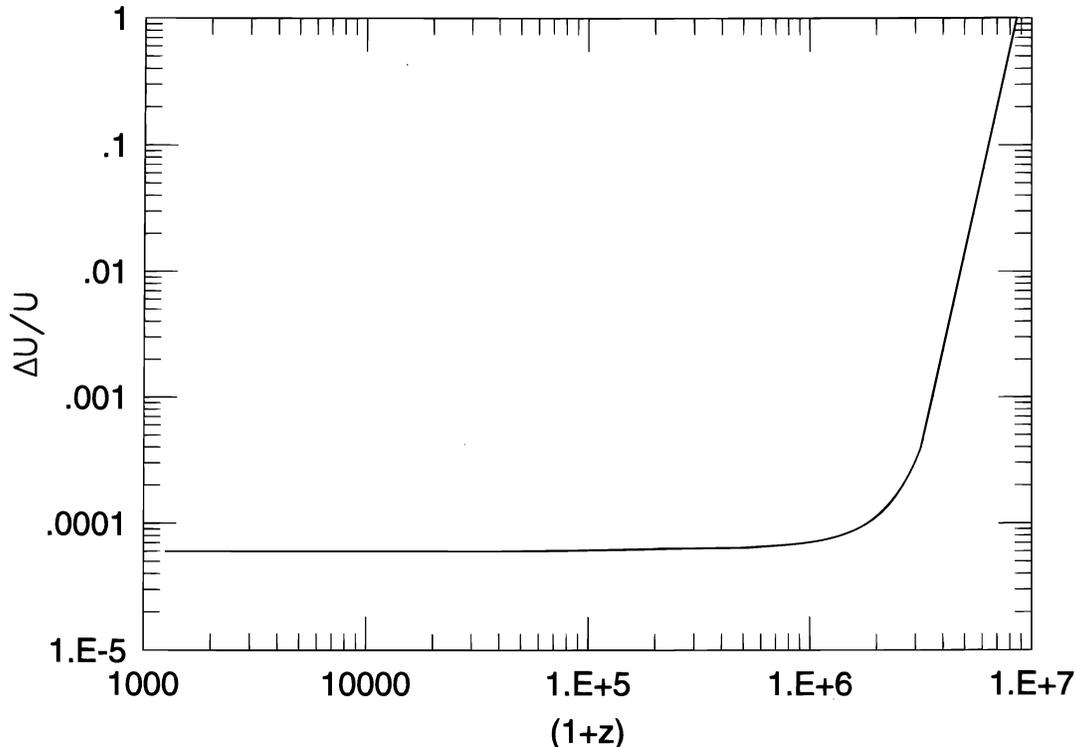


FIG. 6.—Derived limit on injected energy as a function of redshift

(95% CL) toward $(\ell, b) = (264^\circ 14 \pm 0.15, 48^\circ 26 \pm 0.15)$, in agreement with the microwave results from the DMR.

The CMBR temperature is 2.728 ± 0.004 K (95% CL), where the error is dominated by our estimate of the thermometry errors. The weighted rms deviation from the fit is 50 ppm of the peak brightness. The limit on $|y|$ is 15×10^{-6} , and the limit on $|\mu|$ is 9×10^{-5} (95% CL).

We thank the computer scientists and analysts who operated the *COBE*, processed the data, and solved problems in software that could no longer be solved with hardware. In particular, we thank S. Alexander, K. Jensen, S. Read, F. Shuman, and A. Trenholme. This work was supported by the Astrophysics Division of NASA Headquarters.

REFERENCES

- Bennett, C. L., et al. 1992a, in Proc. Third Teton Summer School, The Evolution of Galaxies and Their Environment, ed. H. A. Thronson & J. M. Shull (Dordrecht: Kluwer), 27
 ———. 1992b, ApJ, 396, L7
 ———. 1994, ApJ, 436, 423
 ———. 1996, ApJ, 464, L1
 Boggess, N. W., et al. 1992, ApJ, 397, 420
 Cheng, E. S. 1992, in ASP Conf. Ser. 51, Observational Cosmology, ed. G. Chincarini (San Francisco: ASP), 501
 Fixsen, D. J., et al. 1994a, ApJ, 420, 445
 ———. 1994b, ApJ, 420, 457
 Isaacman, R., Read, S., & Barnes, W. 1992, in ASP Conf. Ser. 25, Astronomical Data Analysis Software and Systems I, ed. D. M. Worall, C. Biemesderfer, & J. Barnes (San Francisco: ASP), 403
 Kompaneets, A. S. 1957, J. Exp. Theor. Phys., 4, 730
 Mather, J. C., et al. 1990, ApJ, 354, L37
 ———. 1994, ApJ, 420, 439
 Mather, J. C., Fixsen, D. J., & Shafer, R. A. 1993, Proc. SPIE, 2019, 168
 O'Neill, E. M., & Laubscher, R. E. 1976, Extended Studies of a Quadrilateralized Spherical Cube Earth Data Base (NEPRF Tech. Rep. 3-76) (NTIS Rep. AD-A026294) (Computer Sciences Corp. CSC/TR-76/6008)
 Peebles, P. J. E. 1971, Physical Cosmology (Princeton: Princeton Univ. Press)
 Reach, W. T., et al. 1995, ApJ, 451, 188
 Shafer, R. A., et al. 1991, Bull. Am. Phys. Soc., 36, 1398
 Sunyaev, R. A., & Zeldovich, Ya. B. 1980, ARA&A, 18, 537
 Wright, E. L., et al. 1991, ApJ, 381, 200
 ———. 1994, ApJ, 420, 450
 Zeldovich, Ya. B., & Sunyaev, R. A. 1969, Ap&SS, 4, 301
 ———. 1970, Ap&SS, 7, 20