THE ASTROPHYSICAL JOURNAL, 458:L53–L55, 1996 February 20


# PRODUCING MEGAPIXEL COSMIC MICROWAVE BACKGROUND MAPS FROM DIFFERENTIAL RADIOMETER DATA

E. L. WRIGHT,[1] G. HINSHAW,[2] AND C. L. BENNETT[3]
*Received 1995 October 2; accepted 1995 December 7*

## ABSTRACT

A major goal of cosmology is to obtain sensitive, high-resolution maps of the cosmic microwave background anisotropy. Such maps, as would be produced by the recently proposed *Microwave Anisotropy Probe* (*MAP*), will contain a wealth of primary information about conditions in the early universe. To mitigate systematic effects when observing the microwave background, it is desirable for the raw data to be collected in differential form: as a set of temperature differences between points in the sky. However, the production of large (megapixel) maps from a set of temperature differences is a potentially severe computational challenge. We present a new technique for producing maps from differential radiometer data that has a computational cost that grows in the slowest possible way with increasing angular resolution and number of map pixels. The required CPU time is proportional to the number of differential data points, and the required random-access memory is proportional to the number of map pixels. We test our technique, and demonstrate its feasibility, by simulating 1 yr of a spaceborne anisotropy mission.

*Subject headings:* cosmic microwave background — methods: data analysis

## 1. INTRODUCTION

Experiments to measure the anisotropy of the cosmic microwave background radiation (CMB) need to measure signals $\Delta T$ of $O(10\ \mu\mathrm{K})$ in the presence of the large isotropic background with $T_0 = 2.73$ K and system noise that is usually $O(100$ K). The difficulty inherent in stabilizing the sensitivity of an instrument to much better than 1 part per million has led all CMB anisotropy experiments to use differential radiometers, which chop rapidly between the spot whose temperature is to be measured and one or more reference sources. The reference sources are usually other parts of the sky, although occasionally stabilized loads have been used. The resulting measurements of $\Delta T$ can be analyzed directly in terms of model predictions, but for high-resolution experiments covering a large fraction of the sky, the number of independent data points can make analysis extremely cumbersome. It is far more desirable to reduce the data by producing a *map* of the microwave temperature. For a given angular resolution, a map provides the most complete form of data reduction possible without loss of information, which in turn permits a full range of statistical tests to be performed on the data. The problem of converting a large set of differential sky observations into a map becomes increasingly challenging as the angular resolution, number of map pixels, and data-sampling rate all increase. The methods employed by the *COBE* Differential Microwave Radiometer (DMR) cannot be directly adapted to the high-resolution case because of limitations imposed by computer RAM. We present a new technique for producing maps that is equivalent to the *COBE* sparse-matrix technique and that has a computational cost that grows in the slowest possible way with increasing angular resolution and number of map pixels: the required CPU time is proportional to the number of differential data points, and the required RAM is proportional to the number of map pixels. We demonstrate the feasibility of this technique by simulating a 1 yr differential mission that produces sky maps with 1,572,864 square pixels of size $0\overset{\circ}{.}16$.

## 2. THE DMR EXPERIENCE

The only full-sky CMB maps produced to date use the differential data from the *COBE* DMR. The DMR data analysis required the construction of sky maps using the $6 \times 10^7$ differences that were collected each year in each of the six channels. Since there are only $\sim 10^3$ beam areas on the sky, the system that relates differences to sky temperature is highly overdetermined. The *COBE* team chose to analyze the data using 6144 pixels to cover the sky. These all have approximately the same area $\theta_{\mathrm{pix}}^2$ and are arranged in a square grid of $32 \times 32$ pixels on each of the six faces of a cube. Within each pixel it was assumed that the temperature was constant, so the basic problem can be represented as a least-squares system with $24 \times 10^7$ equations (after 4 yr of data taking) in 6144 variables. It is also desirable to extend this system to additionally account for systematic signals in the data due, for example, to the magnetic sensitivity of the ferrite switches used in the DMR instrument (Kogut et al. 1992). After calibration and baseline subtraction, one can represent the predicted output of the DMR instrument, in mK, as

$$S(t) = V(t)\mathbf{X}^{\mathrm{T}}, \tag{1}$$

where $t$ is the observation time. $V(t)$ is a vector of length $N_{\mathrm{pix}} + 3$, given by

$$V(t) = [0, \ldots, 0, +1, 0, \ldots, 0, -1, 0, \ldots, B_x(t), B_y(t), B_z(t)], \tag{2}$$

with the $+1$ in the pixel $p(t)$ that contains the plus-horn line of sight at time $t$, the $-1$ in the pixel $m(t)$ that contains the

[1] Department of Physics and Astronomy, UCLA, Los Angeles, CA 90095-1562; wright@astro.ucla.edu.
[2] Hughes STX Corporation, Laboratory for Astronomy and Solar Physics, Code 685, NASA/GSFC, Greenbelt, MD 20771.
[3] Laboratory for Astronomy and Solar Physics, Code 685, NASA/GSFC, Greenbelt, MD 20771.

minus-horn line of sight, and $B_i(t)$ the $i$th component of the magnetic field in spacecraft-fixed coordinates. The parameter vector $\boldsymbol{X}$ is given by

$$\boldsymbol{X} = [T_0, \ldots, T_{6143}, M_x, M_y, M_z], \tag{3}$$

where $T_i$ is the temperature of the $i$th pixel in mK and $M_i$ are the susceptibilities of the ferrite switches to external magnetic fields in mK G$^{-1}$. For a constant noise per observation, the least-squares solution for the map is (Lineweaver et al. 1994; Wright et al. 1994a, b)

$$\mathbf{A} = \sum_t V(t)^{\mathrm{T}} V(t),$$

$$\boldsymbol{B} = \sum_t S(t) V(t)^{\mathrm{T}},$$

$$\boldsymbol{X} = \lim_{\varepsilon \to 0^+} (\mathbf{A} + \varepsilon \mathbf{I})^{-1} \boldsymbol{B}. \tag{4}$$

Note that $\mathbf{A}$ is *sparse*, *symmetric*, and *singular*. The main diagonal of $\mathbf{A}$ has elements $A_{ii} = N_i^{\mathrm{obs}}$, the number of times each pixel was observed. The off-diagonal elements of $\mathbf{A}$ indicate the number of times a given pixel pair $ij$ was observed, either $i = p(t)$ and $j = m(t)$ or vice versa. Since the angle between the DMR horns was fixed at 60°, only pixel pairs separated by 60° ± $\theta_{\mathrm{pix}}$ could ever be observed. Thus, while $\mathbf{A}$ has $38 \times 10^6$ elements, all but $18 \times 10^5$ of these elements are zero, and since $A_{ij} = A_{ji}$, only $9 \times 10^5$ elements need to be kept. The limit $\varepsilon \to 0^+$ regulates the singularity of $\mathbf{A}$ and gives a map with zero mean.

The matrix $\mathbf{A}$ is sparse, but inverting $\mathbf{A} + \varepsilon\mathbf{I}$ does not preserve this sparseness. Thus, an iterative solution to the system $(\mathbf{A} + \varepsilon\mathbf{I})\boldsymbol{X} = \boldsymbol{B}$ is found by using algorithms that only require one to be able to find the product of $\mathbf{A}$ and a vector $\boldsymbol{Y}$. If we let $\mathbf{D}$ be a diagonal matrix whose diagonal elements are $N_i^{\mathrm{obs}}$, then a very simple iterative scheme for computing $\boldsymbol{X}$ is

$$\boldsymbol{X}^{(n+1)} = \boldsymbol{X}^{(n)} + \mathbf{D}^{-1}[\boldsymbol{B} - \mathbf{A}\boldsymbol{X}^{(n)}]. \tag{5}$$

Since differential observations do not constrain the mean of the map, the mean should be adjusted to zero after each iteration. With the DMR scan pattern, this simple iteration scheme converges quite rapidly. The computational cost of this algorithm for both RAM and CPU time is proportional to the number of nonzero off-diagonal elements in $\mathbf{A}$.

## 3. SCALING FOR SMALLER BEAMS

The number of off-diagonal elements can be calculated by considering the number of pixels in the reference ring of a given pixel. The radius of the reference ring is the chop angle $\theta_{\mathrm{chop}}$. The width of the reference ring is about twice the pixel size, $2\theta_{\mathrm{pix}}$. The solid angle of the reference ring is then $4\pi\theta_{\mathrm{pix}}$ sin $\theta_{\mathrm{chop}}$, and the number of pixels in the reference ring is $4\pi$ sin $\theta_{\mathrm{chop}}/\theta_{\mathrm{pix}}$. Since the total number of pixels is $N_{\mathrm{pix}} = 4\pi/\theta_{\mathrm{pix}}^2$, the number of nonzero off-diagonal elements in the upper triangle of $\mathbf{A}$ is $N_{\mathrm{off}} \approx 8\pi^2$ sin $\theta_{\mathrm{chop}}/\theta_{\mathrm{pix}}^3$. Even though this scaling is vastly superior to the $O(N_{\mathrm{pix}}^3) \propto \theta_{\mathrm{pix}}^{-6}$ operations required by direct inversion, this iterative approach using an explicitly stored sparse matrix will still be very difficult to apply to the proposed new generation of anisotropy experiments, *COBRAS/SAMBA*, *FIRE*, *PSI*, and *MAP*, which have beam sizes 10–60 times smaller than the 7° beam of the *COBE* DMR. Thus the RAM and CPU time required to invert a differential map by this technique would increase by a factor of $10^3$ to $2 \times 10^5$ over the time required for the *COBE* DMR maps. Because of this, several groups have abandoned the wide-angle differential radiometer design that worked so well in the *COBE* DMR and are using either total-power radiometers or chopping against a load, both of which are prone to serious systematic effects. However, we have found a different technique to execute the operations in equation (5) that reduces the computational load for small beams by a substantial factor.

## 4. THE TIME-ORDERED ALTERNATIVE

Many iterative methods for solving $\mathbf{A}\boldsymbol{X} = \boldsymbol{B}$ only require the ability to multiply a vector times $\mathbf{A}$. This multiplication can be performed by a sum over the elements of $\mathbf{A}$, or it can be performed by scanning the data in time order and processing each differential pixel pair as it occurs. The value of an off-diagonal element in $\mathbf{A}$ is just the number of times a given pixel pair occurs, and if the typical off-diagonal element is large, then using the sparse matrix is much faster than scanning the time-ordered data. However, the typical value of an off-diagonal element scales like $N_t/N_{\mathrm{off}}$, where $N_t$ is the total number of data samples taken. The proposed new experiments are all at least 10 times more sensitive per sample than the *COBE* DMR, so $N_t$ will be only somewhat larger than the *COBE* DMR data set while $N_{\mathrm{off}}$ will be at least $10^3$ times larger than it was for DMR. Thus, for maps with a large number of pixels, the typical magnitude of an off-diagonal element of the matrix will be quite small—ultimately declining to unity for very small pixels. In this limit the number of nonzero off-diagonal elements equals the number of time-ordered data points. For the *COBE* DMR maps, the time-ordered approach, which scales like $N_t$, was slower than scanning the sparse matrix, which scales like $N_{\mathrm{off}}$, since $N_t \gg N_{\mathrm{off}}$, but for the smaller pixels required by the new generation of experiments the time-ordered approach requires no more operations during the iterative solution for the map and avoids the work of constructing the sparse matrix. In addition, the RAM required by the time-ordered technique scales like $N_{\mathrm{pix}}$ while storing the sparse matrix scales like $N_{\mathrm{off}}$. Since new experiments with large maps will have data rates only somewhat larger than *COBE*'s, and since the speed of computers has grown by orders of magnitude, producing large all-sky maps from differential data is easily feasible for future experiments.

The iterative scheme in equation (5) can be implemented in time order by using the following scheme:

$$X_i^{(n+1)} = \frac{\sum_t \{\delta_{i,p(t)}[X_{m(t)}^{(n)} + S(t)] + \delta_{i,m(t)}[X_{p(t)}^{(n)} - S(t)]\}}{\sum_t [\delta_{i,p(t)} + \delta_{i,m(t)}]}. \tag{6}$$

The denominator in equation (6) is just the diagonal of $\mathbf{A}$ and does not need to be computed during each iteration. For each data point, this time-ordered technique requires obtaining the time-ordered data $S(t)$, $p(t)$, and $m(t)$ from disk or tape, two fetches, four adds, and two stores. It goes very much faster than the actual data taking.

In effect, this scheme evaluates the temperature in each map pixel by averaging all the differential observations of that pixel, after correcting each for an estimate of the signal in the reference beam, obtained from the previous iteration. We have tested this technique by simulating 1 yr of space-mission observations using a DEC 3000/600 Alpha workstation. The input map contains 1,572,864 pixels and includes both simu-

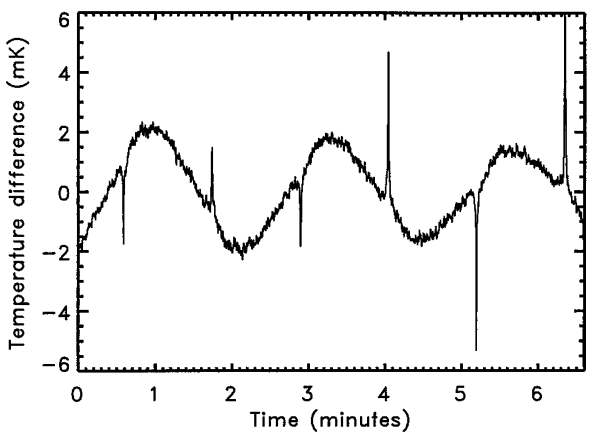


FIG. 1.—Raw differential data from simulated space-mission observations of the sky shown in Fig. 2. The dominant signal is the CMB dipole, modulated by the motion of the spacecraft. Galactic plane crossings are seen as alternating spikes while the small-amplitude fluctuations are from a cold dark matter CMB simulation.

lated CMB anisotropy, including the dipole, and a Galactic foreground model. We "observed" this map using a differential radiometer with 0°.5 FWHM beams separated by $\theta_{\rm chop} = 135°$ on a spinning spacecraft. Figure 1 shows a sample of the time-ordered data from this simulation, using a data rate of 20 observations $s^{-1}$. We processed 1 yr worth of differential observations, including instrument noise of 500 $\mu$K $Hz^{-1/2}$, and solved for the map using equation (6). The iterations were started with a pure dipole map for $\boldsymbol{X}^{(0)}$. Figure 2 (Plate L7) shows the input map and the iterations $\boldsymbol{X}^{(0)}$, $\boldsymbol{X}^{(1)}$, and $\boldsymbol{X}^{(20)}$; the input map and latter two iterations, minus the dipole, are shown with an expanded temperature scale in Figure 3 (Plate L8). Note the large artifacts in the first iteration due to the reference beam's crossing the Galactic plane, which is not modeled in $\boldsymbol{X}^{(0)}$. These artifacts are quickly removed in subsequent iterations, and by $N = 20$ there are no artifacts remaining in the map above the level of the instrument noise. Each iteration requires ~8 hr of CPU time on the existing Alpha workstation. The required time is proportional to the data rate and is virtually independent of the number of map pixels: increasing the number of pixels by a factor of 4 slows processing by only 2%. Significantly, 64% of the CPU time is spent in the subroutine that simulates the spacecraft's attitude. Thus, if applied to real data with stored attitude information, the processing would be even faster.

This technique can be readily extended to include observations of microwave polarization by accounting for the orientation of the observing horns with respect to the sky. We have tested this extension with similar mission simulations and find suitable convergence on maps of the Stokes parameters $I$, $Q$, and $U$. The computational cost of including polarization is less than a factor of 2, relative to the pure intensity case.

## 5. CONCLUSION

We have demonstrated a computational method that allows for the production of megapixel CMB anisotropy maps from large differential radiometer data sets. Thus, the proposed new anisotropy experiments, which would produce maps with hundreds of times more pixels than *COBE* DMRs, can still use the time-tested wide-separation differencing technique successfully employed by the DMR.

We gratefully acknowledge NASA's Office of Space Sciences for the support provided to study the *Microwave Anisotropy Probe* (*MAP*) under the New Mission Concepts program, and G. H. has been partly supported by the Long Term Space Astrophysics (LTSA) program.

## REFERENCES

Kogut, A., et al. 1992, ApJ, 401, 1
Lineweaver, C. H., et al. 1994, ApJ, 436, 452
Wright, E. L., Smoot, G. F., Bennett, C. L., & Lubin, P. M. 1994a, ApJ, 436, 443
Wright, E. L., Smoot, G. F., Kogut, A., Hinshaw, G., Tenorio, L., Lineweaver, C., Bennett, C. L., & Lubin, P. M. 1994b, ApJ, 420, 1

1996ApJ...458L..53W

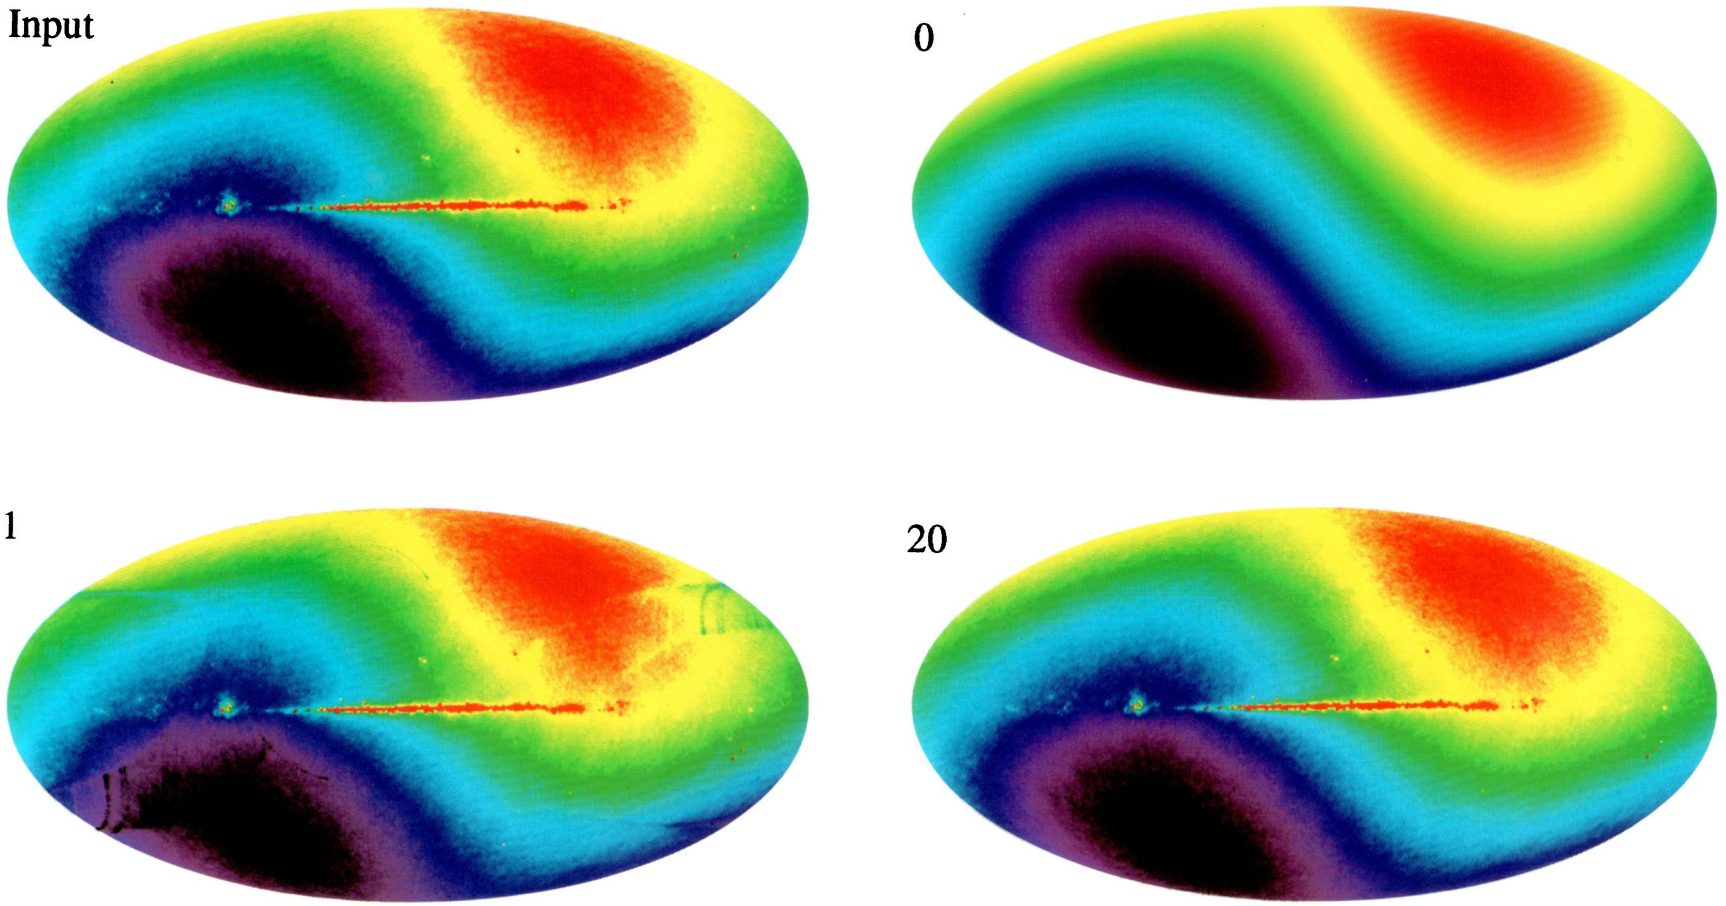


FIG. 2.—*Top left*: Full-sky map used as input for the mission simulation. The map includes simulated CMB anisotropy, the CMB dipole, and a model Galactic signal. The Mollweide equal-area projection is used. *Top right*: Pure dipole signal used for the zeroth iteration. *Bottom left*: Recovered map after one iteration of eq. (6). The Galactic plane signal appears coherently, though with echoes that are ~10% of the plane signal. This efficient reduction of plane echoes after only one iteration requires a scan pattern that successfully connects a given sky pixel to many other pixels, both on and off the Galactic plane. *Bottom right*: Recovered map after 20 iterations, by which time no significant artifacts remain.

WRIGHT, HINSHAW, & BENNETT (see 458, L55)

SIMULATION INPUT

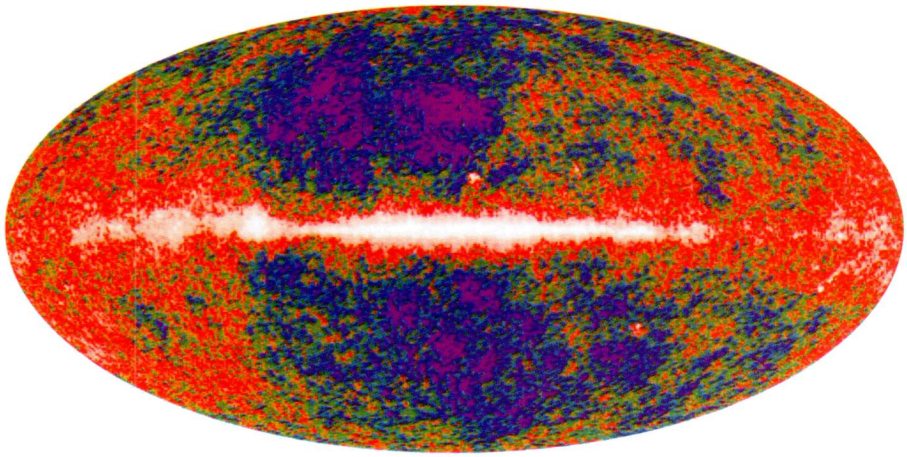

ITERATION 1

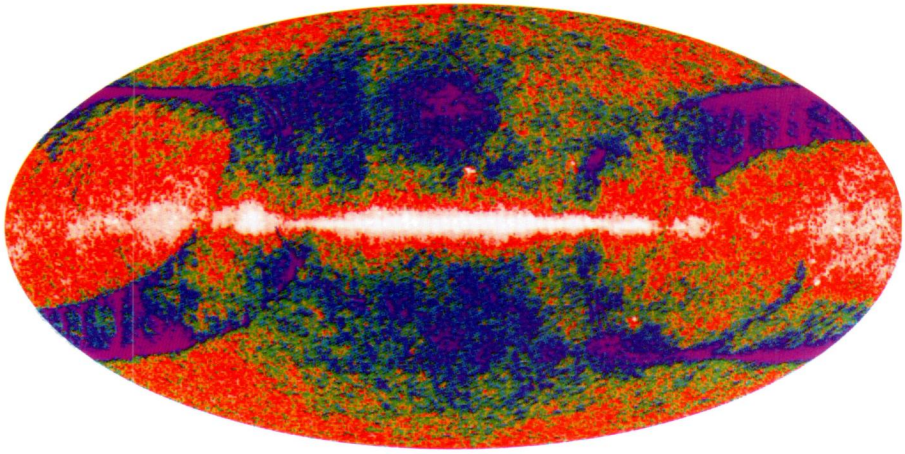

ITERATION 20

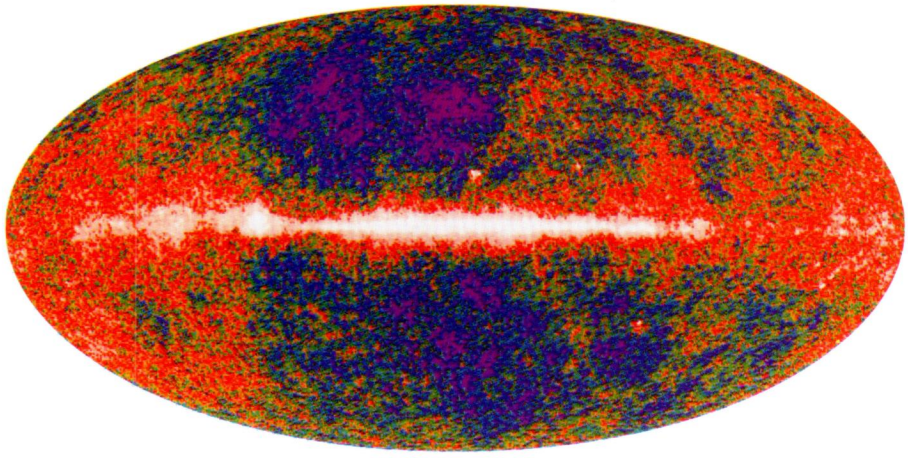

FIG. 3.—Same as Fig. 2, but with the dipole removed and the temperature scale expanded. The zeroth iteration is omitted from this figure.

WRIGHT, HINSHAW, & BENNETT (see 458, L55)