

1997 1 1
 032 107

THE ANGULAR POWER SPECTRUM OF BATSE 3B GAMMA-RAY BURSTS

MAX TEGMARK

Max-Planck-Institut für Physik, Föhringer Ring 6, D-80805 München, Germany and Max-Planck-Institut für Astrophysik, Karl-Schwarzschild Strasse 1, D-85740 Garching, Germany; max@mppmu.mpg.de

DIETER H. HARTMANN

Department of Physics and Astronomy, Clemson University, Clemson, SC 29634; hartmann@grb.phys.clemson.edu

MICHAEL S. BRIGGS

Department of Physics, University of Alabama, Huntsville, AL 35899; briggs@gibson.msfc.nasa.gov

AND

CHARLES A. MEEGAN

NASA/Marshall Space Flight Center, Huntsville, AL 35812; meegan@ssl.msfc.nasa.gov

Received 1995 October 3; accepted 1996 March 14

ABSTRACT

We compute the angular power spectrum C_l from the BATSE 3B catalog of 1122 gamma-ray bursts and find no evidence for clustering on any scale. These constraints bridge the entire range from small scales (which probe source clustering and burst repetition) to the largest scales (which constrain possible anisotropies from the Galactic halo or from nearby cosmological large-scale structures). We develop an analysis technique that takes the angular position errors into account. For specific clustering or repetition models, strong upper limits can be obtained down to scales $l \sim 30$, corresponding to a couple of degrees on the sky.

The minimum-variance burst weighting that we employ is visualized graphically as an all-sky map in which each burst is smeared out by an amount corresponding to its position uncertainty. We also present separate bandpass-filtered sky maps for the quadrupole term and for the multipole ranges $l = 3-10$ and $l = 11-30$, so that the fluctuations on different angular scales can be inspected separately for visual features such as localized “hot spots” or structures aligned with the Galactic plane. These filtered maps reveal no apparent deviations from isotropy.

Subject headings: gamma rays: bursts — methods: statistical

1. INTRODUCTION

The BATSE experiment has now observed more than 1100 gamma-ray bursts. The observed angular distribution is isotropic, while the brightness distribution of bursts shows a reduced number of faint events. These observations favor a cosmological burst origin. The “great debate” on the distance scale of cosmic gamma-ray bursts (GRBs) (Fishman 1995; Lamb 1995; Paczyński 1995) considered two alternatives; cosmological bursts or events that occur in an extended Galactic halo (EGH). The old paradigm of nearby Galactic neutron stars with a Population I distribution perished due to the combined observations of an isotropic angular distribution of GRBs along with reduced source counts at the faint end of the apparent flux distribution (Meegan et al. 1992; Briggs et al. 1996). The absence of even a weak “Milky Way” band in the GRB distribution has indeed made it hard to retain the hypothesis that local neutron stars provide the underlying source population. Some recent reviews of these and related issues are given by Briggs (1995), Fishman & Meegan (1995), and Hartmann (1995).

Although no dominant anisotropies on the sky were found in the apparent sky distribution of gamma-ray bursts, even small effects might contain valuable information about the underlying sources. The detection of a small excess of events in special directions, such as nearby stars or the Andromeda galaxy, could be a unique signature of stellar or Galactic halo models, respectively. For example, a small asymmetry with respect to the Galactic plane might suggest a local disk origin (Hartmann, Greiner, & Briggs 1995).

Clustering of bursts beyond that expected from random alignments might be evidence of actual clustering of the sources or of repeated emission from some sources. Observation of repetition would seriously call into question the viability of those cosmological burst models that invoke unique events, such as mergers of neutron star binaries. On the other hand, a detection of the small anisotropy induced by the Earth motion relative to the cosmic microwave background (CMB) (Maoz 1994; Scharf, Jahoda, & Boldt 1995; but see Brainerd 1996) would constitute a convincing proof of the cosmological origin hypothesis. These various anisotropies manifest themselves on different angular scales and with different magnitudes. Galactic features would be expected to cause large-scale distortions, while burst repetition would show its effects on the scale that is typical for BATSE source localizations (in excess of $1^\circ 6$ for the 3B catalog). In addition, the instrument does not sample the sky uniformly so that we expect some distortions due to the nonuniform exposure map of BATSE.

How should we analyze the angular distribution of GRBs? Since the basic null hypothesis of isotropy states that burst directions are distributed randomly on the sky (which is the impression derived from visual inspection of GRB catalogs), we seek tests that can efficiently find small deviations. First we search for excess of sources toward some direction or a concentration toward some plane in the sky, i.e., we seek a dipole of quadrupole moment. It is perhaps preferable to search for such large-scale anisotropies in an unbiased way by not making reference to any particular coordinate frame (Hartmann & Epstein 1989; Briggs 1993). On the other hand, such

coordinate-free methods are not necessarily the most efficient ones. If a particular anisotropy is expected, then the tests should take this information into account to optimize the search efficiency. Paczyński (1990) introduced studies of the $\cos \theta$ and $\sin^2 b$ statistics, where b is the Galactic latitude of the GRB and θ is the angle between the GRB direction and the vector pointing to the Galactic center. It is now common practice to apply both the coordinate-free and the Galactic methods to the GRB distribution (Briggs 1993; Briggs et al. 1996). These dipole and quadrupole measures were sufficient to characterize the large-scale angular properties of GRBs when sample sizes were a few hundred bursts or less. However, the BATSE experiment has now observed so many bursts that an extension of these moment methods to higher orders is now useful. In this work, we use spherical harmonic analysis (SHA) to represent and interpret the angular distribution of GRBs.

It can be shown (Horack et al. 1993; Briggs et al. 1996) that the statistical estimates of low-order multipoles are not very sensitive to the angular smearing induced by statistical and systematic localization uncertainties. This is not the case for higher order multipoles, which probe the angular density field on smaller scales. We shall address this question very carefully in this work. Small angular scales may reveal important information about the nature of the GRB sources, and location accuracy is crucial. If associated with galaxies, we expect clustering on very small scales (Hartmann & Blumenthal 1989; Lamb & Quashnock 1993). If bursts repeat, we expect clustering at $\theta = 0$ (Quashnock & Lamb 1983a). Both effects are diluted by localization uncertainties (the point-spread function), and apparent power is transferred from small (or zero) angular scales to a scale given by the detector response. A traditional tool for the analysis of source clustering is the angular two-point correlation function, which was applied first to GRBs by Hartman & Blumenthal (1989). The severe reduction in correlation strength by positional smearing was demonstrated by Hartmann, Linder, & Blumenthal (1991). The two-point correlation function is closely related to the power spectrum (e.g., Peebles 1980) (in the ideal world with no measurement errors or shot noise, one would be found to be the spherical Fourier transform of the other). However, the correlation function and the power spectrum complement each other well, since they are affected by noise in quite different ways. This makes it worthwhile to estimate both from the data, just as has become the practice with galaxy surveys. Another method relevant to the study of clustering properties is the nearest neighbor (NN) method (e.g., Scott & Tout 1989). This method, applied first to GRBs by Quashnock & Lamb (1993a), probes only angles near the scale defined by the mean angular pair separation, which decreases with increasing sample size. We do not consider NN methods in this work.

The remainder of this paper is organized as follows. In § 2, we generalize the standard techniques of power spectrum estimation to properly take into account the location errors and the sky exposure of the BATSE catalog. In § 3 we apply this to the 3B data set, and in § 4 we discuss the results.

2. METHOD

In this section, we derive the power spectrum estimation technique that is employed in our analysis. The first sub-

section reviews the statistics of point processes on a sphere. This is standard material and has been discussed frequently in the literature in connection with the problem of estimating the angular power spectrum of point sources such as galaxies or quasars (Peebles 1973; Hauser & Peebles 1973; Peebles 1980); see Tegmark (1995) for a recent review. The extra twist, which makes the analysis of the BATSE data more challenging, is the presence of position errors. Since some bursts are more accurately localized than others, the question of how best to weigh the data is somewhat subtle; this is the topic of the second subsection. After that, we present the explicit expressions for computing the power spectrum estimates from a data set, including a simple beam function model.

2.1. Point Processes on a Sphere

We model the gamma ray burst distribution as a two-dimensional stochastic point process $n(\hat{r}) = \sum_i \delta(\hat{r} - \hat{r}_i)$, which is a Poisson process with intensity (average point density per steradian) $\lambda(\hat{r})$. Here δ denotes the Dirac delta function on the surface of the unit sphere, and the unit vectors \hat{r}_i correspond to the positions of the various bursts. If we had detected a nearly infinite number of bursts, then the function $\lambda(\hat{r})$ would be known with great accuracy, and the only source of errors when computing its power spectrum would be cosmic variance. Since in practice we have only a finite number of bursts (in our case 1122), our estimates of λ itself will be inexact, leading to the additional complication known as shot noise.

A Poisson process satisfies (see, e.g., Appendix A of Feldman, Kaiser, & Peacock 1994)

$$\langle n(\hat{r}) \rangle_p = \lambda(\hat{r}), \quad (1)$$

$$\langle n(\hat{r})n(\hat{r}') \rangle_p = \lambda(\hat{r})\lambda(\hat{r}') + \delta(\hat{r} - \hat{r}')\lambda(\hat{r}). \quad (2)$$

Here λ is itself a random field, $\lambda(\hat{r}) = \bar{n}(\hat{r})[1 + \Delta(\hat{r})]$, where the underlying density fluctuations Δ are modeled as a Gaussian random field. The function \bar{n} , which we will refer to as the *exposure function*, is thus the number of bursts per steradian expected a priori, not the number density actually observed. In other words, $\bar{n}(\hat{r})$ is proportional to the exposure time in the sky direction \hat{r} . As customary, we assume that the expectation value $\langle \Delta(\hat{r}) \rangle_g = 0$ and that the statistical properties of the field Δ are isotropic, which means that if we expand it in spherical harmonics¹ as

$$\Delta(\hat{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^l a_{lm} Y_{lm}(\hat{r}), \quad (4)$$

then

$$\langle a_{lm} a_{l'm'} \rangle_g = \delta_{ll'} \delta_{mm'} C_l, \quad (5)$$

where the coefficients C_l are known as the *angular power spectrum*. There are thus two separate random steps involved in generating n : first the generation of the smooth field Δ , then the Poissonian distribution of points. To make

¹ Since all our fields are real valued, we will find it convenient to use the *real-valued* versions of the spherical harmonics throughout. These are identical to the conventional spherical harmonics Y_{lm} as defined in, for instance, Abramowitz & Stegun (1965), except that the complex exponentials $e^{im\phi}$ are replaced by $2^{1/2} \sin(m\phi)$, 1 and $2^{1/2} \cos(m\phi)$ for $m < 0$, $m = 0$ and $m > 0$, respectively. With this convention, the standard identities involving spherical harmonics remain unchanged except that no complex conjugation is needed. For instance, the orthogonality relation becomes simply

$$\int Y_{lm}(\hat{r}) Y_{l'm'}(\hat{r}) d\Omega = \delta_{ll'} \delta_{mm'}. \quad (3)$$

this distinction clear, we use the notation $\langle \rangle_p$ for expectation values involving only the Poisson step (as in eqs. [1] and [2]) and write $\langle \rangle_g$ for expectation values involving only the Gaussian field Δ (as in eq. [5]). When we write simply $\langle \rangle$ without a subscript, we mean the expectation value involving both steps. For instance, $\langle n(\hat{r}) \rangle = \langle \langle n(\hat{r}) \rangle_p \rangle_g = \langle \lambda(\hat{r}) \rangle_g = \bar{n}(\hat{r})$.

Given the field $n(\hat{r})$, we wish to estimate the coefficients a_{lm} . We denote our estimates \tilde{a}_{lm} , and for reasons that will soon become clear, we define them as

$$\tilde{a}_{lm} \equiv \int Y_{lm}(\hat{r}) \frac{n(\hat{r})}{\bar{n}(\hat{r})} d\Omega - \delta_{l0} \delta_{m0} \sqrt{4\pi}. \quad (6)$$

We now compute the statistical properties of these estimates. By substituting equation (1) into equation (6), we obtain

$$\langle \tilde{a}_{lm} \rangle = \int Y_{lm}(\hat{r}) d\Omega - \delta_{l0} \delta_{m0} \sqrt{4\pi} = 0, \quad (7)$$

i.e., the expectation values vanish. Since the expectation values of the true coefficients a_{lm} vanish as well, this means that our estimates are unbiased. Notice that we chose to include the second term in equation (6) simply to cancel the bias arising from the monopole term $l = m = 0$. Using the expressions above, we find that the correlation between two multipole estimates is

$$\begin{aligned} \langle \tilde{a}_{lm} \tilde{a}_{l'm'} \rangle &= \iint Y_{lm}(\hat{r}) Y_{l'm'}(\hat{r}') \\ &\times \left[\langle \Delta(\hat{r}) \Delta(\hat{r}') \rangle_g + \frac{1}{\bar{n}(\hat{r})} \delta(\hat{r}, \hat{r}') \right] d\Omega d\Omega'. \end{aligned} \quad (8)$$

Substituting equation (4) into this, and using the spherical harmonic orthogonality relation (3) and equation (5), this reduces to

$$\langle \tilde{a}_{lm} \tilde{a}_{l'm'} \rangle = \delta_{ll'} \delta_{mm'} C_l + \int \frac{Y_{lm}(\hat{r}) Y_{l'm'}(\hat{r})}{\bar{n}(\hat{r})} d\Omega. \quad (9)$$

If \bar{n} is merely a constant, i.e., if the exposure time is the same for all parts of the sky, then the orthogonality relation will reduce the second term to simply $\delta_{ll'} \delta_{mm'} / \bar{n}$, and the various estimates \tilde{a}_{lm} will all be uncorrelated. Since the true exposure function \bar{n} for the BATSE 3B data set varies somewhat across the sky, a slight correlation will result.

Of course, we are also interested in estimating the angular power spectrum C_l . Defining the quantities

$$\tilde{C}_{lm} \equiv \tilde{a}_{lm}^2 - b_{lm}, \quad (10)$$

we find that they are unbiased power estimates (in the sense that $\langle \tilde{C}_{lm} \rangle = C_l$) if we choose our *bias correction* to be

$$b_{lm} \equiv \int \frac{Y_{lm}^2(\hat{r})}{\bar{n}(\hat{r})} d\Omega. \quad (11)$$

If \bar{n} is constant, then the bias correction becomes simply $b_{lm} = 1/\bar{n}$, independent of l and m .

It should now be clear why we divided by \bar{n} in equation (6). If we had not divided by the exposure function, then our power estimator \tilde{C}_{lm} would not have measured only what we wanted it to, i.e., C_l . Rather, $\langle \tilde{C}_{lm} \rangle$ would also have received contributions from other multipoles $C_{l'}$, with $l \neq l'$. The quantities \tilde{C}_{lm} are thus good estimates of C_l for each m -value separately. To reduce error bars, we estimate the

power by averaging the \tilde{C}_{lm} :

$$\tilde{C}_l \equiv \frac{1}{2l+1} \sum_{m=-l}^l \tilde{C}_{lm}. \quad (12)$$

Defining b to be the average of the bias corrections b_{lm} , we find that b is in fact independent of l : by substituting the spherical harmonic addition theorem (16) into equation (11) and using the fact that $P_l(1) = 1$, we obtain

$$b \equiv \frac{1}{2l+1} \sum_{m=-l}^l b_{lm} = \frac{1}{4\pi} \int \frac{d\Omega}{\bar{n}(\hat{r})}, \quad (13)$$

i.e., b is just the spherical average of $1/\bar{n}$. This means that the coefficients b_{lm} , which would be slightly cumbersome to compute numerically, need never be computed at all, since the power estimate \tilde{C}_l is simply the average of the squared \tilde{a}_{lm} -coefficients minus b .

2.2. The Effect of Position Errors

The discussion in the previous section applies to any population of point sources on the celestial sphere, not merely gamma-ray bursts. However, analyzing the BATSE catalog involves an extra complication that is absent in, for instance, galaxy and quasar catalogs: position errors.

Let us study first the simple case in which the position errors are the same for all bursts in the catalog. If the true direction to a burst is \hat{r} , then we model the apparent direction \hat{r}' as a random variable whose probability distribution depends only on the angle between \hat{r} and \hat{r}' . Thus, we can write the probability distribution as $B(\hat{r} \cdot \hat{r}')$ for some function B that we will refer to as the *beam function*.

Above, we characterized the distribution of the true burst positions as a Poisson process with intensity $\lambda(\hat{r})$, where λ was in turn a Gaussian random field. From now on, we will let the density $n(\hat{r}) = \sum_i \delta(\hat{r} - \hat{r}_i)$ refer not to the true burst positions but to the apparent positions. It is easy to show that this n will also be a Poisson process, but with a different intensity function λ . Specifically, the apparent intensity is the true one convolved with the beam function, i.e.,

$$\lambda_{\text{app}}(\hat{r}) = (B * \lambda_{\text{true}})(\hat{r}) \equiv \int B(\hat{r} \cdot \hat{r}') \lambda_{\text{true}}(\hat{r}') d\Omega'. \quad (14)$$

Thus, the effect of the position errors is to smooth out sharp features in the expected burst density, which as we will see limits our ability to measure fluctuations on scales below the beamwidth. Let us expand the beam function in Legendre polynomials as

$$B(\hat{r} \cdot \hat{r}') = \sum_{l=0}^{\infty} \left(\frac{2l+1}{4\pi} \right) B_l P_l(\hat{r} \cdot \hat{r}'). \quad (15)$$

By using the spherical harmonic addition theorem,

$$\sum_{m=-l}^l Y_{lm}(\hat{r}) Y_{lm}(\hat{r}') = \left(\frac{2l+1}{4\pi} \right) P_l(\hat{r} \cdot \hat{r}'), \quad (16)$$

together with the orthogonality relation (3), we can thus write the beam function as

$$B(\hat{r} \cdot \hat{r}') = \sum_{l=0}^{\infty} \sum_{m=-l}^l B_l Y_{lm}(\hat{r}) Y_{lm}(\hat{r}'). \quad (17)$$

Applying the beam convolution to equation (4) and using the orthogonality relation, we thus obtain the spherical

version of the convolution theorem:

$$(B * \Delta)(\hat{r}) = \sum_{l=0}^{\infty} \sum_{m=-l}^l a_{lm} B_l Y_{lm}(\hat{r}). \quad (18)$$

In other words, convolution with B corresponds simply to multiplying the multipole coefficient a_{lm} by B_l .

Repeating the analysis of the previous section including position errors (replacing λ by $B * \lambda$), the case in which \bar{n} is constant² thus yields the simple result

$$\langle \tilde{C}_{lm} \rangle = B_l^2 C_l. \quad (19)$$

In practice, some sources are localized more accurately than others, and we clearly want to make use of this fact to make the most of the data. Suppose that the total population, with number density \bar{n} , consists of a number of subpopulations with number densities \bar{n}_i (so that $\sum \bar{n}_i = \bar{n}$), and that all bursts in the i th subpopulation are equally accurately localized, as specified by a beam function B_i . (Since their shape depends only on the sky exposure, all the functions \bar{n}_i are identical apart from their normalization, so $\bar{n}_i \propto \bar{n}$.) Estimating the power spectrum can now be split into two steps:

1. Estimate a_{lm} separately from each population, as above, and call the results \tilde{a}_{lm}^i ;
2. Combine these estimates into one by some weighted averaging,

$$\tilde{a}_{lm} = \sum W_{li} a_{lm}^i. \quad (20)$$

Obviously we want the weights W_{li} to be larger for those populations i that are better localized. Let us now determine which weighting scheme is optimal. The generalization of equation (19) to multiple populations is readily found to be

$$\langle \tilde{C}_{lm} \rangle = \langle \tilde{a}_{lm}^2 \rangle - b_l = \left(\sum_i B_{li} W_{li} \right)^2 C_l, \quad (21)$$

where the bias correction is

$$b_l \equiv \sum_i \frac{W_{li}^2}{\bar{n}_i}. \quad (22)$$

How should we choose our weights W_{li} ? First of all, to make \tilde{C}_{lm} an unbiased estimate of C_l , clearly we want to normalize the weights so that the expression in parentheses in equation (21) equals unity, i.e., so that $\sum_i B_{li} W_{li} = 1$. Second, we want the error bars on our estimate to be as small as possible, i.e., we want to minimize the variance of

² When \bar{n} is not constant, the $B * (\bar{n}\Delta)$ term in addition gives rise to a weak mode coupling between the different multipoles. As discussed below, the \bar{n} of the BATSE 3B data set is basically constant except for small dipole and quadrupole corrections. This means that a_{lm} will pick up small contributions from $a_{l'm}$, where $|l' - l| \leq 2$, which is completely irrelevant for this analysis. The reason is that it is merely a second-order effect: we are investigating whether there is any signal at all apart from the shot noise, and this coupling effect would only alter the *relative* level of the signal by a few percent. Thus, the only instance in which the anisotropy of \bar{n} must be taken into account is when computing the noise bias with eq. (11), since an error of a few percent in the (much larger) shot noise contribution could be of the same order as the weak signal we are trying to detect.

\tilde{C}_l . In the approximation that \tilde{a}_{lm} is Gaussian,³ we have simply $V(\tilde{C}_l) = 2\langle \tilde{a}_{lm}^2 \rangle^2$, so that we minimize the variance by minimizing the expectation value $\langle \tilde{a}_{lm}^2 \rangle = C_l + b_l$. Since C_l is independent of our weights, we thus wish to choose W_{li} so as to minimize the bias correction b_l , given the above-mentioned normalization constraint $\sum_i B_{li} W_{li} = 1$. This constrained optimization problem is solved readily by the method of Lagrange multipliers, and the solution is

$$W_{li} = b_l \bar{n}_i B_{li}, \quad (23)$$

where the minimal bias correction is

$$b_l = \left(\sum_i \bar{n}_i B_{li}^2 \right)^{-1}. \quad (24)$$

In summary, we have found our best multipole estimate to be

$$\tilde{a}_{lm} \equiv \frac{Nb_l}{4\pi} \sum_i B_{li} \int Y_{lm}(\hat{r}) \frac{n_i(\hat{r})}{\bar{n}(\hat{r})} d\Omega. \quad (25)$$

(For brevity, we omit the trivial monopole correction eq. [7] here and below, since we are never interested in computing the monopole anyway.)

2.3. Power Spectrum Estimation in Practice

For any given data set, the density field n_i is just a sum of delta functions, one for each burst, so equation (25) reduces to

$$\tilde{a}_{lm} = \frac{Nb_l}{4\pi} \sum_i B_{li} \sum_{j=1}^{N_i} \frac{Y_{lm}(\hat{r}_j)}{\bar{n}(\hat{r}_j)}, \quad (26)$$

where N_i denotes the number of bursts in the i th subpopulation. We can simplify this expression further by a mere change of notation. We let the index k refer to sums over the entire burst sample ($k = 1, \dots, N$), and from here on, we simply let B_{lk} denote the beam factor corresponding to the subpopulation that the k th burst belongs to. Then

$$b_l^{-1} = \sum_i \bar{n}_i B_{li}^2 = \frac{1}{4\pi} \sum_i N_i B_{li}^2 = \frac{1}{4\pi} \sum_k B_{lk}^2 = \frac{N_l^{\text{eff}}}{4\pi}, \quad (27)$$

where we have defined the *effective number of bursts* at a given multipole as $N_l^{\text{eff}} = \sum_{k=1}^N B_{lk}^2$. With this same convention, replacing the double sum over subpopulations and their members by a single sum over all bursts, equation (25) simplifies to

$$\tilde{a}_{lm} = \frac{N}{N_l^{\text{eff}}} \sum_{k=1}^N B_{lk} \frac{Y_{lm}(\hat{r}_k)}{\bar{n}(\hat{r}_k)}, \quad (28)$$

Thus, we have eliminated the need to keep track of

³ The Gaussian approximation is good when the number of bursts is large, by the central limit theorem. It should be emphasized that even under circumstances in which this approximation is poor, our \tilde{C}_l will be a good estimate of the power spectrum: it will simply have slightly larger error bars than it would with optimal weighting.

subpopulations altogether,⁴ and expressed our multipole estimates directly in terms of the observed quantities.

Repeating the analysis for an arbitrary exposure function \bar{n} , equation (27) becomes generalized to

$$b_l = \frac{N}{N_l^{\text{eff}}} \frac{1}{4\pi} \int \frac{d\Omega}{\bar{n}(\hat{r})}. \quad (29)$$

We estimate the C_l by averaging over m -values as before, i.e.,

$$\tilde{C}_l \equiv \left(\frac{1}{2l+1} \sum_{m=-l}^l \tilde{a}_{lm}^2 \right) - b_l. \quad (30)$$

In the above-mentioned Gaussian approximation, the \tilde{C}_{lm} of equation (10) are almost independent with variance $V(\tilde{C}_{lm}) = V(\tilde{a}_{lm}^2) = 2(\tilde{a}_{lm}^2)^2$, since the b_{lm} are mere constants. Hence, the 1σ error bar is

$$\Delta\tilde{C}_l \approx \frac{1}{2l+1} \left[\sum_{m=-l}^l V(\tilde{C}_{lm}) \right]^{1/2} \approx \left(\frac{2}{2l+1} \right)^{1/2} (C_l + b_l). \quad (31)$$

Thus, as l increases, the error bars will typically decrease first due to the growing number of independent m -modes and then gradually start increasing again around the scale corresponding to the position errors as N_l^{eff} eventually approaches zero, making b_l explode.

2.4. The Beam Function

We model the BATSE beam function as a Fisher function (Fisher, Lewis, & Embleton 1987):

$$B_k(\hat{r} \cdot \hat{r}') = \frac{\exp(\sigma_k^{-2} \hat{r} \cdot \hat{r}')}{4\pi\sigma_k^2 \sinh(\sigma_k^{-2})} \quad (32)$$

characterized by a location error σ_k . This is often considered by mathematicians to be the spherical version of the Gaussian distribution, and it reduces to

$$B^k(\cos \theta) \approx \exp\left(-\frac{1}{2} \frac{\theta^2}{\sigma_k^2}\right) / 2\pi\sigma_k^2, \quad (33)$$

when $\sigma_k \ll 1$ radian $\approx 60^\circ$. The Fisher function has the advantage that it is correctly normalized (its integral over the sphere is unity) for arbitrarily large angles σ_k , which is not the case for the plane Gaussian of equation (33). It should be emphasized that although the BATSE location error distribution has usually been modeled as a Gaussian distribution, it is currently not well enough known that one particular distribution is preferred over another, so the choice is merely one of convenience.

⁴ The Gaussian assumption that we used for computing error bars was strictly valid only when $N_l \gg 1$ for each subpopulation. However, since the BATSE 3B distribution of position errors forms a smooth continuum, we expect the error bars derived from the Gaussian approximation to remain accurate anyway, as long as $N_l^{\text{eff}} \gg 1$, and this is indeed confirmed numerically by Monte Carlo simulations. We generated 1000 mock BATSE 3B catalogs with no clustering and analyzed them with the same software as the real data, extracting the multipoles $l \leq 40$. To within the Monte Carlo errors (a relative error of order $1000^{-1/2} \approx 3\%$), the actual error bars were identical to those expected analytically when making the Gaussian approximation.

In the limit $\sigma_k \ll 1$ (valid for all bursts in the sample as shown in Fig. 2), we have to a good approximation that

$$B_{lk} \approx e^{-(1/2)\sigma_k^2 l(l+1)}. \quad (34)$$

The position uncertainties $\Delta\theta$ quoted in the BATSE 3B catalog are defined as the radius of the 1σ circle, i.e., of the circle that contains $\text{erf}[2^{-1/2}] \approx 68\%$ of the probability. Thus, in the limit $\sigma_k \ll 1$, the conversion between $\Delta\theta$ and σ is

$$\frac{\sigma}{\Delta\theta} = \left\{ -2 \ln \left[1 - \text{erf} \left(\frac{1}{\sqrt{2}} \right) \right] \right\}^{-1/2} \approx 0.66. \quad (35)$$

Note that the values of $\Delta\theta$ quoted in the BATSE 3B catalog do not include the systematic error contribution of 1.6° , which is to be added to the quoted values in quadrature. This yields the distribution shown in Figure 2.

3. RESULTS

We have used the improved BATSE positions of the 3B catalog (Meegan et al. 1995b, 1995c) to expand the angular distribution of GRBs in terms of spherical harmonics. The 3B catalog contains 1122 bursts with known best-fit positions (shown in Fig. 1a) and their statistical uncertainties. In addition to statistical shifts, we must also include (in quadrature) a 1.6° systematic uncertainty. This value is significantly lower than the 4° of earlier catalogs, and it allows us to extend spherical harmonic analysis to $l \sim 50$ – 60 before localization uncertainties completely wash out any possible intrinsic angular power in the GRB sky map. The distribution of the actual statistical errors is shown in Figure 2.

Because the sky exposure of BATSE is not uniform (Fishman et al. 1994; Meegan et al. 1995b), artificial moments are induced (e.g., Briggs et al. 1996). The BATSE experiment does not exclude any area of the sky, but due to blocking by the Earth and detector gaps during passages of the South Atlantic Anomaly (SAA), some positions on the sky have a reduced probability for burst detection. The associated exposure map is thus best described as a semi-transparent mask. While the exposure corrections are not as severe as those encountered in galaxy surveys, it should and can be included in the analysis. We shall discuss the effect of uneven sampling in the next section.

3.1. The Exposure Function

Because of problems due to the loss of the spacecraft tape recorders, the absolute efficiency has not been determined since the release of the 1B data set. However, the shape of the exposure function \bar{n} is essentially independent of time, and since the shape is all that matters for the present analysis, we employ the 1B estimate (Fishman et al. 1994). This function \bar{n} depends on declination only and is independent of right ascension. This means that in equatorial coordinates, the multipole coefficients \bar{n}_{lm} vanish except when $m = 0$. The dominant deviation from uniformity is a quadrupole ($\bar{n}_{20}/\bar{n}_{00} \approx 8.8\%$) depletion of bursts near the equator due to the shadowing of the sky by the Earth. The second largest anisotropy is a dipole moment ($\bar{n}_{10}/\bar{n}_{00} \approx 4.5\%$) toward the Earth's north pole, due to the South Atlantic Anomaly, which requires disabling triggers. Compared to the shot noise, the higher multipoles ($l \geq 3$) are




FIG. 1.—The BATSE 3B data set and the smoothed burst map. The measured locations of the BATSE 3B sample of 1122 gamma-ray bursts are shown in Hammer-Aitoff projection in galactic coordinates (*top*) and with each burst smeared out by an amount corresponding to the uncertainty in its position (*bottom*).

negligible ($a_{10}/a_{00} \lesssim 1\%$), but for completeness, they have nonetheless been included in our analysis.

3.2. The Power Spectrum

The power spectrum \tilde{C}_l extracted from the BATSE 3B data set is shown in Figure 3, and as can be seen, there is no evidence of deviations from isotropy on any angular scale. What is plotted is, of course, the difference between two positive quantities, the power in the data minus the bias correction, according to equation (30), which is why some (unphysical) negative estimates occur. Thus, if the gamma-ray bursts are in fact completely uncorrelated, we would expect the points in Figure 3 to be scattered symmetrically around zero, with roughly equal numbers above and below the horizontal axis, and about 68% within the shaded region. Since all power is by definition positive, the presence of any type of correlation would shift the distribution upward, leading to a positive excess.

In Figure 3, we have divided the power spectrum by 4π to make the interpretation of the numbers simpler. A mono-

pole $C_0/4\pi = 0.0001$ would simply correspond to a fluctuation of $0.0001^{1/2} = 1\%$ in the average burst density. Likewise, $(C_l/4\pi)^{1/2}$ can be interpreted as the density fluctuation on the angular scale $\theta \approx 60^\circ/l$.

Let us comment briefly on this factor of 60° and the correspondence between l and θ . From equation (34), we see that roughly speaking, a burst probes the multipole l only if the factor B_{lk} is of order unity, i.e., if $\sigma_k \lesssim 1$. Here σ_k is measured in radians, so since 1 radian is $180^\circ/\pi \approx 57^\circ$, this means that only bursts with a location error $\sigma \lesssim 60^\circ/l$ are sensitive to the multipole l .

3.3. The Error Bars

The size of the error bars (the height of the shaded region) in Figure 3 is readily understood from equation (31). For $l = 0$, we have $N_l^{\text{eff}} = N = 1122$, so apart from the factor of $2^{1/2}$, the shot noise gives just the familiar Poisson variance $1/N$. As l increases, the $(2l + 1)$ denominator reduces the error bars, since many independent modes are being averaged. However, as l increases beyond the scale correspond-




FIG. 2.—A histogram of the position errors $\Delta\theta$ is shown for the BATSE 3B sample of 1122 gamma-ray bursts. The $1^\circ6$ systematic errors are included here, added to the statistical errors in quadrature.

ing to the typical location errors, the sharp drop in N_l^{eff} causes the error bars to increase dramatically. Thus, we cannot place strong constraints for $l \geq 60$ simply because there are no bursts that are localized to better than $1^\circ6$




FIG. 3.—The shot-noise-corrected angular power spectrum. Filled squares show the multipoles estimated from the BATSE 3B data set with minimum-variance burst weighting and shot noise removed (this is why unphysical negative values occur). The shaded region shows the 1σ shot noise error bars, so if there is no clustering whatsoever, about 68% of the squares would be expected to fall within this region, distributed symmetrically around zero. Any type of clustering would drive the points upward, leading to more points above zero than below. The double-shaded region shows that the error bars would be if there were no position errors.




FIG. 4.—The factor by which fluctuations are suppressed by the effect of position errors, N_l^{eff}/N , is plotted as a function of multipole l . Our method corrects for the smearing by dividing by this suppression factor, which is the reason that the error bars in Fig. 3 explode for large l . The suppression factor for the real data (shaded) is compared with the hypothetical situation in which all bursts have the same position errors $\Delta\theta$, taken to be $1^\circ6$ (solid line), 4° (long-dashed line), and 10° (short-dashed line).

($l \sim 35$), so $l \geq 60$ would be more than “ 2σ ” out in the Gaussian tail of B_l , causing the shot noise to explode.⁵

This effect is the reason that the actual error bars become so much larger than the “ideal world” error bars (double-hatched) that would result if there were no position errors. This is also illustrated in Figure 4, where N_l^{eff} is plotted as a function of l . For $l = 30$, for instance, we are effectively only making use of about 25% of all bursts, the remainder being too poorly localized to contribute much information about the power on this small a scale. Conversely, Figure 3 shows also that for $l \leq 5$, the location errors have little impact on the error bars, confirming the results of Horack et al. (1993) and Briggs et al. (1996) for dipole and quadrupole moments.

Note that N_l^{eff} in Figure 4 is far from being Gaussian: for small l , it falls off roughly as a Gaussian with $\Delta\theta = 10^\circ$, but for larger l , the tail falls off much more slowly, since most of the contribution is coming from the best localized bursts. It should also be noted that since the C_l -coefficients are rotationally invariant quantities, Figure 3 would look identical

⁵ If the true location errors should turn out to be larger than those we have assumed here, then the error bars would thus explode at lower l -values. For instance, Graziani & Lamb (1995) analyzed the distribution of 3B locations in comparison to the positions derived from the IPN³ network and conclude that the systematic error of the 3B data should be about 4° instead of the advertised $1^\circ6$. In addition, there may be correlations in the data that suggest a brightness dependence for the systematic error, instead of the constant value suggested by Meegan et al. (1995c). The studies by Graziani & Lamb (1995) did not take systematic effects in the IPN localization method into account and also do not incorporate the fact that some IPN locations are based upon the earlier BATSE 2B locations and thus may be biased against the 3B locations. Although we tend to agree more with the error budget prescribed by the BATSE Team, this is an unsettled question, and it should be borne in mind that whereas our multipole estimates for $l \leq 15$ (corresponding to $\theta \geq 4^\circ$) are virtually independent on this controversy, the estimates for $l \geq 30$ should be taken with a grain of salt if the reader favors a 4° systematic error.

if Galactic rather than equatorial coordinates had been used when generating it.

3.4. Is the Exposure Function Correct?

If the estimate of \bar{n} (Fishman et al. 1994) were incorrect, this could introduce artificial signals into our power spectrum. Because of the azimuthal symmetry, this would affect only those coefficients \tilde{a}_{lm} that have $m = 0$. These are plotted in Figure 5. Thus, if the bursts are uncorrelated and the \bar{n} estimate is correct, the points should scatter symmetrically around zero with about 68% of them in the shaded region, which appears to be the case. Figure 5 thus provides reassuring evidence that \bar{n} has been correctly modeled. To indicate the sensitivity of this analysis, the figure shows also the dipole and quadrupole that would be expected if we had failed to correct for the above-mentioned Earth-shadow quadrupole and the South Atlantic Anomaly. Since the quadrupole correction was about 9%, this shows that uncertainties in the modeling of the higher multipoles of \bar{n} , which are typically at least an order of magnitude smaller, will not be important compared to the $(N_l^{\text{eff}})^{-1/2}$ errors caused by shot noise.

3.5. The Minimum-Variance-Weight Burst Map

Using equation (17) and the orthogonality relation (3), we can rewrite equation (28) as

$$\tilde{a}_{lm} = \frac{N}{N_l^{\text{eff}}} \int Y_{lm}(\hat{r}) x(\hat{r}) d\Omega, \quad (36)$$

where we have defined x , the *smoothed burst map*, as

$$x(\hat{r}) \equiv \sum_{k=1}^N \frac{B^k(\hat{r} \cdot \hat{r}_k)}{\bar{n}(\hat{r})}. \quad (37)$$

Thus, we see that the minimum-variance method we derived above has a very simple interpretation: apart from the




FIG. 5.—The multipole coefficients with $m = 0$. The multipole coefficients a_{l0} in equatorial coordinates, corresponding to fluctuations independent of right ascension, are shown (filled squares) together with the 1σ region expected from shot noise alone. For any isotropic fluctuations, the distribution should be symmetric around zero. The triangle and the star show the effect that the South Atlantic Anomaly and Earth shadowing would have if they were not taken into account in \bar{n} .

overall weighting factor N/N_l^{eff} , our optimal estimates of the multipoles a_{lm} were just the spherical harmonic coefficients of a map in which *each burst is smeared out by its own beam function*, and corrected for the uneven sky exposure. This map is shown in Figure 1b and Figure 6 (Plate 1) (upper left). A comparison of this map with that using earlier BATSE data (Hartmann et al. 1994) shows the tremendous improvements due to the reduction of systematic position uncertainties from 4° to 1.6° and the increase in sample size.

It is quite useful for inspecting the data set visually, since in a sense it displays only the information that is really present in the data and not more. It does not mislead the eye by exaggerating the accuracy to which the burst locations are known, enabling those bursts that are well localized to visually stand out against the background.

3.6. Bandpass-Filtered Maps

Although the angular power spectrum C_l provides a useful measure of the amount of clustering on different angular scales, it should be borne in mind that it does not contain any information about the relative *phases* of the different multipoles a_{lm} . The same can be said about the correlation function, a useful statistical quantity that has been estimated elsewhere (Meegan et al. 1995b, 1995c; Blumenthal 1995). The loss of phase information means that although the power spectrum may tell us that there is extra power on some scale, it does not tell us anything about where in the sky this power is coming from; we may, for instance, be interested in knowing if there are any signals localized near the Large Magellanic Cloud or aligned with the Galactic plane. Fortunately, this type of information (which can be seen as complementary to that provided by the power spectrum) is easy to extract with the formalism developed above. We define $x_l(\hat{r})$, the *multipole map* corresponding to multipole l , as the sky map

$$x_l(\hat{r}) \equiv \sum_{m=-l}^l \tilde{a}_{lm} Y_{lm}(\hat{r}), \quad (38)$$

where the estimated spherical harmonic coefficients \tilde{a}_{lm} are those defined by equation (28). Similarly, we define the *bandpass-filtered map* corresponding to a multipole range $l_1 \leq l \leq l_2$ as the sum of the multipole maps for the different l values in the range. Figure 6 shows the filtered maps corresponding to $l = 2$ (the quadrupole), $l = 3-10$, and $l = 11-30$, respectively, and the reader is encouraged to scrutinize these images in search of any features that are spatially localized or aligned with the Galactic plane, both of which would provide evidence against isotropy. The quadrupole, for instance, is neither aligned with the Galactic plane nor with the equator of Earth, and as is seen in Figure 3, its amplitude is of the order that is expected from mere shot noise fluctuations.

Using the orthogonality relation, we see that apart from the shot noise correction and a proportionality constant, our multipole estimate \tilde{C}_l is just the integral of the square of the corresponding multipole map, $\int x_l^2 d\Omega$. It is in this sense that the filtered maps allow us to see where the power (the fluctuations) is coming from. Also, apart from normalization issues (for instance, the density modulation in \bar{n} is eliminated in the filtered maps), the smoothed burst map in Figure 1b is just an average of all the multipole maps, weighted by inverse noise level. Thus, we can think of the filtered maps roughly as a decomposition of the smoothed

burst map into its different frequency components, into its contributions from different angular scales.

4. DISCUSSION

Much of the current debate on the origin of GRBs rests upon a careful analysis of their angular and brightness distribution. Without established counterparts or other burst properties that could be used to estimate distances, we do not even know their distance scale, which in turn leaves burst energetics undetermined. Building models is a challenge under such conditions. One of the most important pieces of information that we can obtain is the angular distribution of GRBs. Deviations from isotropy on some angular scale for some or all bursts could provide crucial hints to the distance scale. The lack of large anisotropies makes it very hard to retain traditional models of neutron stars in the Galactic disk. But even models that invoke a very extended halo do predict small anisotropies that should emerge eventually from the data. And while cosmological models generically result in isotropic distributions, they too may have tell-tale deviations. We may consider the small deviations due to the Earth's motion with respect to the CMB, or the granularity due to local superstructures in the cosmic mass distribution. In addition, the well-known angular correlations of many cosmological objects or clustering that would result from burst recurrences would lead to some deviations from isotropy. The distribution of burst positions on the sky could be the primary source of information leading to an understanding of the burster distance scale, and perhaps their nature as well.

The crucial objective of our study is thus an advanced analysis of GRB positions. There are two significant steps in this field:

1. Providing accurate locations for all bursts;
2. Analyzing this position information with appropriate statistical tools.

The BATSE Team has made great progress in the first area, now providing location accuracies down to about 2° for many bursts and about 5° for the average burst (Meegan et al. 1995c). The reduction of systematic uncertainties is essential for studies of small-scale anisotropies, but it also contributes to better estimates of more global patterns that may be present in the data. The smearing of burst positions, unavoidable from the instrumental point of view, must be included in the data analysis. Additional features that must be accounted for are temporal and angular gaps in the observations. Here we do not consider possible structure in burst arrival times, but we study exclusively their arrival directions. The exposure function of BATSE must be and has been included in this work.

The remaining question is about selecting appropriate tools. This depends somewhat on the question we wish to address. Global anisotropies present in many Galactic burst models can be studied through low-order multipole expansion, e.g., dipole-quadrupole statistics, while clustering is generally approached with angular correlation functions or nearest neighbor distributions. Because of the larger database and the superior position accuracy of the 3B data studied here, we are actually able to bridge these two distinct approaches by extending dipole and quadrupole analysis of the angular distribution of GRBs to higher order multipoles. The technique is the well known spherical harmonic analysis (SHA), i.e., expansion of the burst distribu-

tion in terms of spherical harmonics, Y_{lm} . As discussed above, the angular scale probed by a given harmonic is approximately $60^\circ/l$, so that the expansion must be carried out to l -values in excess of 30 if we wish to reach the smearing scale of the current BATSE configuration.

4.1. Limits on Repetition

If some fraction of the observed GRBs are repeat events, the sky distribution should show angular concentrations on small scales (roughly given by the beam smearing of the instrument). Evidence for burst recurrence was found in the 1B data (Quashnock & Lamb 1993a), but subsequent 2B data did not confirm this result (e.g., Meegan et al. 1995a).

The 3B data set is greatly improved over the 2B data in its ability to test the repeater hypothesis for the following reasons:

1. The systematic position uncertainty has been reduced from 4° to 1.6° , and
2. In addition to the overall exposure time being increased by about a year, the post-2B portion of the 3B catalog has a greater fractional exposure (live time), which is important for repeater models in which the bursting phase of sources is less than the BATSE lifetime.

Burst recurrence is expected to generate excess correlations at $\theta = 0$, which corresponds to excess power at all multipoles.⁶ Our study does show some modes with deviations around the 2σ level, but this is by no means a significant excess of power because only about the expected number of points deviate by about 2σ and the points are generally scattered with 1σ of no power. The data are consistent with the hypothesis of no recurrences. The alternative hypothesis tested most frequently in the literature are repeater models in which a fraction f of all observed bursts can be labeled as repeaters that are observed to burst ν times each. Tegmark et al. (1996, hereafter T96) employed an SHA-based technique to test this two-parameter family of models against the BATSE 3B data, and they find that all models with $(\nu - 1)f \geq 0.05$ are ruled out at 99% confidence, as compared to the best previous 99% limit $(\nu - 1)f \geq 0.27$ (Meegan et al. 1995c). Thus, even a cluster of six events from a single source would have caused excess power above that present in the 3B catalog. In other words, the multipole information that our SHA extracts from the data, as plotted in Figure 3, translates into sharp quantitative limits on repetition.

It is conceivable that bursts repeat once or more often on a timescale of months and become dormant afterward for a much longer period. In that situation, accumulation of bursts into a growing sample would dilute the repeater

⁶ From the additional theorem (16), one obtains the well-known result that

$$\langle \Delta(\hat{r})\Delta(\hat{r}') \rangle = \sum_{l=0}^{\infty} \left(\frac{2l+1}{4\pi} \right) P_l(\hat{r} \cdot \hat{r}') C_l, \quad (39)$$

i.e., the C_l -values are basically the spherical Fourier coefficients of the angular correlation function. Correlations only at zero angular separation (before position errors are added in) correspond to the correlation function being a Dirac delta function. Just as the regular Fourier transform of $\delta(x)$ is a constant, the power spectrum corresponding to repeaters is C_l constant, independent of l .




FIG. 6.—All-sky maps of the 1122 gamma-ray bursts in the BATSE 3B data set, in Hammer-Aitoff projection. In the smoothed burst map (*top left*), each burst is smeared out by an amount corresponding to the uncertainty in its position. The other maps are bandpass filtered, showing the quadrupole ($l = 2$, *top right*), intermediate scales ($3 \leq l \leq 10$, *bottom left*), and small scales ($11 \leq l \leq 30$, *bottom right*).

TEGMARK et al. (see 468, 221)

signal. When the 3B set is divided into four sets of roughly equal number of bursts (not equal time), the correlation function shows some small-angle excess at the approximately $1-2\sigma$ level in all subsets (Meegan et al. 1995b; Blumenthal 1995). Adding these correlation functions together generates a noticeable, but still not highly significant, excess of burst pairs near about 5° . Our corresponding SHA analysis for the four subsets (Fig. 7) reveals this effect also, but it is evident from this figure that the significance of this increase is marginal at best. In other words, SHA yields results that are consistent with those obtained by correlation function analysis. This emphasizes the fact that the SHA method now bridges the range of power estimators previously employed in GRB studies.

4.2. Limits on Large-Scale Clustering

Angular power spectra also constrain burst models that trace the large-scale structure of the universe. If GRBs trace the galaxy distribution (as neutron star binary mergers would), we expect to find angular correlations similar to those observed for galaxies or clusters of galaxies (Hartmann & Blumenthal 1989; Lamb & Quashnock 1993). However, if BATSE samples to a redshift of order unity (assuming standard cosmology and standard candles for bursts), the sparse sampling of the galaxy density (specific rate of bursts inside or near galaxies is $\sim 10^{-6} \text{ yr}^{-1}$) and the imperfect angular resolution reduces the expected strength of the clustering signal. With increasing sample size, it will be possible to apply brightness selections, while retaining good angular resolution. So far, only the dipole and quadrupole term have been investigated as a function of apparent source brightness and interpreted in the context of

Galactic GRB models (Quashnock & Lamb 1993b; Gurevich et al. 1993, 1994) and cosmological GRB models (Hartmann, Briggs, & Mannheim 1995).

Figure 3 shows that the BATSE 3B data are consistent with the null hypothesis of no clustering at all. Translating this multipole information into quantitative limits of course requires a detailed specification of a clustering model. For the above-mentioned repeater case, this was done by defining a two-parameter family of repeater models and placing lower limits on the parameter values (T96). Although any type of clustering would produce excess power on some scale, addressing the issue of how high above the Poisson noise this excess would be, and thus the issue of whether such clustering is ruled out by BATSE or not, would require detailed modeling of the clustering pattern in question. Since SHA tightened the previous repeater limits by a factor of 5, it appears quite promising to carry out such modeling of other clustering scenarios as well.

If a model predicts clustering that is anisotropic (for instance, an excess in the supergalactic plane), then it is likely that even stronger constraints can be obtained by making use of this spatial template information, since our averaging over m -values can "dilute" an anisotropic power contribution by up to a factor of $(2l+1)$. In other words, if one knows exactly what kind of clustering one is looking for, one should adopt a data analysis technique that incorporates this information. In contrast, since the angular power spectrum involves no preferred directions in the sky, SHA is a useful general-purpose diagnostic tool which will detect clustering on any scale (if it is strong enough to stand out over the Poisson noise), without requiring any prior knowledge as to what kind of clustering one is looking for.

4.3. Outlook

We conclude that multipole expansion of the projected distribution of GRBs does not show evidence for clustering on any angular scale. This argues against the recurrence of a substantial fraction of burst sources (at 99% confidence, not more than 5% of the BATSE 3B bursts can be due to repeating sources [T96]) and against any source population with intrinsically strong anisotropies resulting from an intrinsically special position of the observer. The remarkable degree of isotropy of GRBs constrains severely any burst model that invokes traditional geometric features of the Milky Way (disk, bulge, or halo). If one wishes to retain the Galactic origin hypothesis by introducing very extended halo distributions, it seems that these populations cannot contribute significantly to the dynamics of the Galaxy (those that do are all known to be highly anisotropic) but must constitute a trace population. Whether high-velocity neutron stars injected into the Galactic halo can indeed provide the necessary isotropy remains to be determined, and model builders should verify that the proposed spatial distributions indeed generate essentially zero angular power on all scales, as our analysis suggests. The currently fashionable paradigm of cosmological bursts now passes this test, but eventually some deviations from isotropy are expected, and spherical harmonic analysis is a tool well suited to detect such deviations.

The authors would like to thank Ted Bunn for help with Hammer-Aitoff projections, and G. R. Blumenthal and an anonymous referee for valuable comments on the manuscript. This work has been partially supported by European




FIG. 7.—Angular power spectrum when grouping into quarters. Filled squares in the bottom plot show the multipoles estimated as in Fig. 3, except that the data have been split into four sequential quarters according to when the burst occurred. Each square shows the average of the four estimates that the multipole, and the corresponding error bars are seen to be twice as large as in Fig. 3, shown above for comparison. The slight apparent excess of power in the bottom figure is consistent with the correlation function analysis of Meegan et al. (1995b, 1995c), which finds weak clustering when the data is time binned into four quarters.

Union contract CHRX-CT93-1020, Deutsche Forschungsgemeinschaft grant SFB-375 and NASA grant NAGW 5-1578. D. H. H. is grateful to J. Trümper for hospitality and support of a summer visit to the MPE, where this study

was initiated, and to the staff and scientific members of the ITP and UCSB for their hospitality and support during the final stages of this project. Work carried out at the ITP was supported in part through NSF grant PHY 94-07194.

REFERENCES

- Abramowitz, M. & Stegun, I. A., ed. 1965, *Handbook of Mathematical Functions* (New York: Dover)
- Blumenthal, G. R., Hartmann, D. H., & Linder, E. V. 1995, *AIP Conf. Proc.* 307, *Gamma-Ray Bursts*, ed. G. J. Fishman, J. J. Brainerd, & K. Hurley (New York: AIP), 117
- Brainerd, J. J. 1996, *ApJ*, 456, 449
- Briggs, M. S. 1993, *ApJ*, 407, 126
- . 1995, *Ap&SS*, 231, 3
- Briggs, M. S., et al. 1996, *ApJ*, 459, 40
- Feldman, H. A., Kaiser, N., & Peacock, J. A. 1994, *ApJ*, 426, 23
- Fisher, N. I., Lewis, T., & Embleton, B. J. J. 1987, *Statistical Analysis of Spherical Data* (Cambridge: Cambridge Univ. Press)
- Fishman, G. J. 1995, *PASP*, 107, 1145
- Fishman, G. J., & Meegan, C. A. 1995, *ARA&A*, 33, 415
- Fishman, G. J., et al. 1994, *ApJS*, 92, 229
- Graziani, C., & Lamb, D. Q. 1995, *Proc. Third Huntsville Workshop on Gamma-Ray Bursts*, ed. C. Kouveliotou, M. S. Briggs, & G. J. Fishman (New York: AIP), in press
- Gurevich, A. V., Zharkov, G. F., Zybin, K. P., & Ptitsyn, M. O. 1993 *Phys. Lett. A*, 181, 289
- . 1994, *Phys. Lett. A*, 192, 305
- Hartmann, D. H. 1995, *A&A Rev.*, 6, 225
- Hartmann, D. H., & Blumenthal, G. R. 1986, *ApJ*, 342, 521
- Hartmann, D. H., Briggs, M. S., & Mannheim, K. 1995, *ApJ*, submitted
- Hartmann, D. H., & Epstein, R. I. 1989, *ApJ*, 349, 960
- Hartmann, D. H., Greiner, J., & Briggs, M. S. 1995, *A&A*, submitted
- Hartmann, D. H., Linder, E., & Blumenthal, G. R. 1991, *ApJ*, 367, 186
- Hartmann, D. H., et al. 1994, *ApJS*, 90, 893
- Hauser, M. G., & Peebles, J. P. E. 1973, *ApJ*, 185, 757
- Horack, J. M., Meegan, C. A., Fishman, G. J., Wilson, R. B., Paciesas, W. S., Emslie, A. G., Pendleton, G. N., & Brock, M. N. 1993, *ApJ*, 413, 293
- Lamb, D. Q. 1995, *PASPP*, 107, 1152
- Lamb, D. Q., & Quashnock, J. 1993, *ApJ*, 415, L1
- Maoz, E. 1994, *ApJ*, 428, 454
- Meegan, C. A., et al. 1992, *Nature*, 355, 143
- . 1995a, *ApJ*, 446, L15
- . 1995b, electronic catalog: use username gronews at gronews.gsfc.nasa.gov
- . 1995c, *ApJ*, submitted
- Paczynski, B. 1990, *ApJ*, 348, 485
- . 1995, *PASP*, 107, 1167
- Peebles, P. J. E. 1973, *ApJ*, 185, 413
- . 1980, *The Large Scale Structure of the Universe* (Princeton: Princeton Univ. Press)
- Quashnock, J., & Lamb, D. Q. 1993a, *MNRAS*, 265, L59
- . 1993b, *MNRAS*, 265, L45
- Scharf, C. A., Jahoda, K., & Boldt, E. 1995, *ApJ*, 454, 573
- Scott, D., & Tout, C. A. 1989, *MNRAS*, 241, 109
- Tegmark, M. 1995, *ApJ*, 455, 429
- Tegmark, M., Hartmann, D. H., Briggs, M. S., Hakkila, J., & Meegan, C. A. 1996, *ApJ*, 466, 757 (T96)