THE ASTROPHYSICAL JOURNAL, 473:310–321, 1996 December 10

BATSE OBSERVATIONS OF GAMMA-RAY BURST SPECTRA. III. LOW-ENERGY BEHAVIOR OF TIME-AVERAGED SPECTRA

R. D. PREECE, M. S. BRIGGS, G. N. PENDLETON, AND W. S. PACIESAS
Department of Physics, University of Alabama at Huntsville, Huntsville, AL 35899
J. L. MATTESON, D. L. BAND, AND R. T. SKELTON
Center for Astrophysics and Space Sciences, Code 0111, University of California at San Diego, La Jolla, CA 92093-0111
AND
C. A. MEEGAN
NASA/Marshall Space Flight Center, ES84, Huntsville, AL 35812
Received 1996 January 16; accepted 1996 June 21

ABSTRACT

We analyze time-averaged spectra from 86 bright gamma-ray bursts from the first 5 years of the Burst And Transient Source Experiment (BATSE) on board the *Compton Gamma Ray Observatory* to determine whether the lowest energy data are consistent with a standard spectra form fit to the data at all energies. The BATSE Spectroscopy Detectors have the capability to observe photons as low at 5 keV. Using the gamma-ray burst locations obtained with the BATSE Large Area Detectors, the Spectroscopy Detectors' low-energy response can be modeled accurately. This, together with a postlaunch calibration of the lowest energy Spectroscopy Detector discriminator channel, which can lie in the range 5–20 keV, allows spectral deconvolution over a broad energy range, ~5 keV to 2 MeV. The additional coverage allows us to search for evidence of excess emission, or for a deficit, below 20 keV. While no burst has a significant ($\geq 3\ \sigma$) deficit relative to a standard spectra model, we find that 12 bursts have excess low-energy emission, ranging between 1.2 and 5.8 times the model flux, that exceeds 5 σ in significance. This is evidence for an additional low-energy spectral component in at least some bursts, or for deviations from the power-law spectral form typically used to model gamma-ray bursts at energies below 100 keV.

Subject headings: gamma rays: bursts — methods: data analysis — X-rays: general

1. INTRODUCTION

Gamma-ray bursts (GRBs) are a transient astrophysical phenomenon in which the emission is confined exclusively to high energies (hard X-ray and gamma-ray bands; energies >2 keV). Moreover, the deficit of X-rays compared with gamma rays from these events is a serious constraint on theoretical models, such that it has been described as the "X-ray paucity" problem (see Harding 1991 and Fishman & Meegan 1995 for reviews of theory and observation, respectively). The general statement of X-ray paucity, derived from years of observations of GRBs by many instruments, is that the ratio of X-ray fluence to gamma-ray fluence is roughly $L_X/L_\gamma \approx 0.02$, where the energy marking the break between the two energy bands is 10 keV. GRBs are distributed isotropically on the sky, but inhomogeneously in space, in that there are far fewer weak bursts than expected for a constant density distribution in Euclidean space (Meegan et al. 1992). So far, no convincing counterpart has been observed at any wavelength. The origin of GRBs, whether in an extended Galactic halo or at cosmological distances, remains controversial. Previous analyses of BATSE data have characterized GRB spectra and temporal properties above ~20 keV (Schaefer et al. 1992; Band et al. 1993; Pendleton et al. 1994; Ford et al. 1995). In this study, we present a calibration of the Spectroscopy Detector (SD) discriminators, which allows us to extend the BATSE spectral coverage down to ~5 keV for a large numer of GRBs for the first time. We shall investigate whether there is an excess or deficit of X-ray emission relative to a canonical spectral form fit to time-averaged gamma-ray spectra.

There have been observations made with several instruments of X-ray emission simultaneous with GRBs. Most detectors that have observed bursts at low energies are constructed with a low-Z window, for which the transparency is a strong function of the path length through the material and, thus, the angle to the source. Experiments that do not have positional information therefore cannot model accurately their low-energy response, which limits severely the usefulness of the low-energy data. Trombka et al. (1974) reported on a burst (GRB 720427) in the *Apollo 16* gamma-ray spectrometer ($E > 67$ keV) that also had significant counts in the 2.8–7.9 keV range in an X-ray detector on the same spacecraft. Exact positional information for this event was not available, with large errors in the relative timing between the two spacecraft that observed the event (*Apollo 16* and *Vela 6A*), making analysis of the low-energy data difficult. Wheaton et al. (1973) published several spectra for GRB 720514 from the UCSD solar X-ray telescope on the *OSO 7* satellite and found hard-to-soft evolution in the 11–100 keV range. They also reported significant flux from one burst in the 7–10 keV range. Several bursts were reported with emission in the 3–10 keV range from the XMON experiment on the Air Force's *P78-1* satellite (Laros et al. 1984) simultaneous with gamma-ray observations on other spacecraft. An X-ray afterglow in GRB 790307 observed by XMON, which persisted for several tens of seconds after the end of the gamma-ray emission, was interpreted by Laros et al. as a cooling curve. In other cases as well, the X-ray and gamma-ray time histories presented by Laros et al. differed from each other. In their published spectra for two intervals of GRB 790307, the X-ray points lie above fitted model

curves by $\approx 2\ \sigma$, which the authors did not consider significant. Yoshida et al. (1989) reported on X-ray emission in 24 GRBs observed by *Ginga*. Spectral analyses of events observed by *Ginga* were limited to those for which locations could be determined, requiring the event to be observed by at least two other spacecraft. They observed hard-to-soft spectral evolution over a 100 s interval covering the bright portion of GRB 870303, fitting several spectra over 1.5–370 keV with thermal photon models, including a blackbody spectrum during the decay phase. However, other models for the low-energy spectra could not be ruled out due to the poor statistics. In addition, Strohmayer et al. (1996) find no evidence in the 2–10 keV spectra of 22 bursts observed by *Ginga* for deviations from the properties of spectra above 20 keV, although deconvolution of the low-energy spectra is plagued by an uncertain burst angle in most of their sample. In particular, the spectra continue smoothly from high to low energy and can be described well by concave-down spectra models. Soft X-ray tails were found in eight *Ginga* events by Murakami et al. (1991), as well as one precursor. The WATCH-*Granat* survey by Castro-Tirado (1994) reported X-ray flux in the energy range 6–15 keV, which was extended in time compared with the observed gamma-ray flux, for roughly 10% out of 70 bursts observed. They obtained spectral fits consistent with thermal emission in six cases of precursors and four X-ray tails. Observations of X-ray emission in GRB precursors or tails without accompanying gamma-ray emission belie the rule of X-ray paucity and raise the possibility of excess X-ray emission *during* a burst, which is the main concern in this paper.

Prior to BATSE, the theory of X-ray emission from GRBs was driven by the then-popular model that the sources of GRBs were nearby (Galactic disk population) strongly magnetized neutron stars (Harding 1991). Crustquakes were considered by Blaes et al. (1989) to explain X-ray emission for up to 150 s after a burst. Precursors in the X-ray band were predicted for the reprocessed pulsar model of Ruderman et al. (1988). More recent models have had to contend with the BATSE observations of GRB isotropy and inhomogeneity (Meegan et al. 1992, 1996; Briggs 1995), which are consistent with the sources of GRBs at cosmological distances or in an extended Galactic halo, but not with a nearby Galactic disk origin for the majority of the bursts. Low-energy precursors may occur in the electron-positron pair fireball model of Meszáros, Laguna, & Rees (1993) when the source becomes optically thin to X-rays. A universal low-energy power law of $N_\nu \propto \nu^{-2/3}$ arises via synchrotron emission by a shock-heated, nonthermal distribution of particles in blast-wave or synchrotron shock-emission models (Katz 1994; Tavani 1995). This prediction is consistent with the empirical fact of X-ray paucity but may require a separate component to account for excess X-ray emission, either simultaneous with the gamma-ray flux or as observed in soft precursors and tails. On the other hand, the Compton attenuation model of Brainerd (1994) predicts an X-ray turnup as a common feature of GRB spectra. In the context of the Compton attenuation model, the presence or absence of excess X-ray emission in any individual burst depends upon several parameters of the source, such as the opacity of surrounding material.

No single instrument has reported detailed spectral analyses of the low-energy behavior of a large number of GRBs. The problems involved in this task are manifold: First, bursts must be well located to account correctly for the instrument's low-energy response, as discussed above. In addition, the instrument must cover a broad energy range with good energy resolution, and it must have enough collecting area to obtain good statistics. Finally, an instrument should cover a fair fraction of the sky at all times, to accumulate a large burst sample in a reasonable time period. The unique capabilities of the BATSE SDs allow us to perform this type of analysis. We have analyzed the time-averaged spectra of 86 bursts and measured the agreement of the data from ~ 5 to 20 keV with a model obtained by a fit to up to 2 MeV.

The primary analysis method employed in this paper entails a joint spectral fit to SD high energy-resolution burst data (SHERB) and low-energy SD discriminator data (DISCSP1) covering the brightest portion of a burst and drawn from the single detector aligned most closely with the source. Details of the instrument, as well as a full description of each data type, are covered in the next section. A brief discussion of the energy calibration of DISCSP1 is given in § 3. Several tests of the channel-to-energy conversion algorithm derived from the energy calibration are mentioned in § 4. We discuss our analysis techniques and observations in § 5 and conclude with a discussion of the results and some of their implications in § 6.

2. BATSE SPECTROSCOPY DETECTORS

The SDs are a component of the BATSE instrument on board the *Compton Gamma Ray Observatory* (Fishman et al. 1989). They consist of eight cylindrical modules of NaI(Tl), 127 cm^2 on each face and 7.2 cm thick, each mounted adjacent to a BATSE Large Area Detector (LAD), but offset in angle by 18°.5. With the faces of the LADs parallel to the faces of a regular octahedron, the SDs are canted toward the center plane of the spacecraft. The offset is compensated by the more isotropic response of the SDs above 20 keV in comparison to the LADs, ensuring full sky coverage. The NaI crystal in each SD is optically coupled to a single photomultiplier tube (PMT), which has adjustable high voltage to set the desired energy range. Most of the face and sides of the SDs are covered by 0.13 cm of aluminum, which is essentially opaque to photons below 20 keV. However, centered on the face of each detector is a circular beryllium window, 7.62 cm in diameter, which allows some response to lower energy photons. Within this aperture, the only materials between the detector crystal and a cosmic source are light-reflective paper, the 0.127 cm thick Be window, and the spacecraft thermal blanket. These transmit, for example, 1% at 5 keV and 32% at 8 keV, on-axis. The transmission decreases rapidly with angle off-axis due to the increased path length; for this reason, an accurate burst location is essential for determining flux from detector counts. Our most recent in-flight calibrations show that it is necessary to include the reflective paper in the detector response function to ensure the accuracy of observations of solar flares and other sources.

Since details of detector physics enter into this paper, it is worthwhile to review how the BATSE SDs work and define some useful terms at the same time. Counts in the detector are produced by the conversion of the energy of gamma rays (and energetic charged particles) in the NaI crystal to visible light, which is collected by the PMT and converted into a voltage pulse by a preamplifier. The pulse is amplified and shaped before it is sent onto two different paths for

analysis by the on-board electronics. On one branch, a pulse-height analyzer (PHA), triggered by a commandable lower-level discriminator (LLD), produces a high-resolution, 2752 channel count spectrum. The high-resolution data are compressed to 256 channels (Spectroscopy High Energy Resolution = SHER), starting at a fixed offset channel, or pedestal, and spaced quasi-logarithmically in energy. The energy coverage of the PHA spectrum is therefore determined by three factors: the PMT gain (via the high-voltage setting), the commandable LLD setting, and the fixed PHA pedestal setting. The gain scales the entire spectrum up or down in energy (high gain gives low energies); a nominal, or $1\times$, gain has an upper energy of 10 MeV and a typical LLD threshold of ~ 50 keV. Most of the data reported here were taken in $4\times$ and $8\times$ gains. The latter has an upper energy of 2 MeV and a typical LLD threshold of ~ 10 keV.

On the other branch, the pulse passes through the four-channel fast discriminator (SFAST1–4) electronics and is accumulated into four integral data channels (SPectroscopy DISCriminator = DISCSP1–4). The energy calibration of the thresholds of the lowest two discriminators is crucial to the present work. The same discriminator, SFAST2, sets the energy threshold of DISCSP2 and the LLD. The electrical thresholds of the first (and lowest) discriminator (SFAST1) is half that of SFAST2. However, the energy threshold of SFAST1 will not be half of SFAST2 because of the nonlinear light output of NaI, as described in § 3.

The data collected from each of these branches are also grouped into bursts and nonburst data types, with the burst data accumulation initiated by the on-board trigger. Typically, background SHER data are accumulated over ~ 300 s but sometimes as short as 32 s. Higher time resolution burst data (SHER Burst = SHERB) are accumulated after a trigger, but only for the four detectors whose adjoining LADs have the largest increases in counts at the time of the trigger. The durations of individual SHERB spectral accumulations are integer multiples of 64 ms based upon the count rate in the LADs, but they are no shorter than 128 ms. The nonburst four-channel SFAST data (DISCSP) are collected every 2.048 s at all times, serving as a background data type that aids burst analysis. In addition, burst four-channel discriminator DISCSP data are accumulated simultaneously with the SHERB data, so the two data sets cover exactly the same time intervals. The first integral data channel (DISCSP1) accumulates counts above the SFAST1 discriminator, and so on. We will be dealing with the differential counts in each channel for our analysis; henceforth, DISCSP1 will stand for counts accumulated between SFAST1 and SFAST2, etc.

For each PHA or discriminator channel, the observed counts may be due to photons with energies ranging from slightly below the channel's lower energy to infinite energy; thus, the channel edges are more appropriately termed "energy-loss thresholds." Photons may be downshifted in energy due to scattering of gamma rays outside the detector's NaI crystal or may suffer partial deposition of their energy in the detector crystal. This is a property of the detector and its surroundings, and it varies as a function of angle between the source and the face of the detector, energy, detector geometry, and geometry of likely scattering sources, such as the Earth's atmosphere and components of the spacecraft, etc. These effects are modeled by the detector response matrix (DRM; Pendleton et al. 1995), which takes into account the efficiency of the detector to observe photons at all energies (the response is thus a lower triangular matrix, mixing each detector channel with all the input photon energy bins for energies greater than the output channel).

The attribution of counts to photons in the DISCSP1 energy ranges is dependent on the counts in that channel due to higher energy photons, and self-consistency is required between the measured and model-predicted counts at high energies and low energies. Due to the effects of the iodine K-escape feature at 33.2 keV, the spectral shape from 38 to 53 keV has the most influence on counts recorded in the 5–20 keV range that encompasses DISCSP1. For example, in 3B 911118 (burst names from the BATSE 3B Catalog [Meegan et al. 1996] are used throughout), a burst characterized by a fairly hard -0.58 spectra index for the photon flux up to ≈ 120 keV, 43% of the 5–10 keV counts in DISCSP1 came from higher energy photons. This is a worst-case example: most bursts are softer, with typically a -1 power-law index in the 20–60 keV range. In such a case, high-energy photons contribute only $\approx 20\%$ of the counts in DISCSP1. The true photon flux in DISCSP1 is therefore model dependent; it is determined relative to a spectral model fit jointly to all the data, which takes into account the detector response.

3. ENERGY CALIBRATION OF DISCSP1

The energy-loss thresholds of the PHA spectra data were determined by calibrating the SDs in the lab against sources with spectral lines at known energies (Band et al. 1992). We have had to determine the DISCSP thresholds after the spacecraft was launched into orbit. The PMT gain, LLD position, and PHA pedestal together determine the energy threshold of SFAST2, which, as designed, corresponds to the LDD channel number in the SHER data. To determine the correspondence between LDD setting and SHER channel number, the LLD for each detector was commanded to several levels in orbit. Background SHER count spectra obtained for each setting were examined to determine the channel number of the LLD. In practice, the LLD cutoff is spread over several channels due to variations in the threshold voltage at which the discriminator triggers. Going upward from the lowest channels, the LLD is manifested as a smoothly rising curve, similar to the hyperbolic tangent function (Fig. 1). We define the LLD as the channel in which the counts have diminished to half the counts for a flat spectrum. Since most count spectra are not flat, the slope of the count spectra at the LLD will affect the determination of the LLD by one to two channels. We made a linear fit for each detector through the points obtained by plotting the commanded LLD setting versus channel number. The slopes from these fits agree well with laboratory data; however, the intercepts are systematically different, as expected: They corresponded to the PHA pedestal settings, which were changed before launch (which explains why the ground measurements cannot be used directly). We estimate that this relationship is accurate to within two channels ($=1$ keV, for $8\times$ gain), which is well below the energy resolution of the SDs.

To calculate the energy edges of DISCSP1, we first use the linear fits to the in-orbit measurements to predict the SHER channel number of the LLD. We find good agreement between the predicted SHER LLD channel, for several detectors at different gains, and the actual position

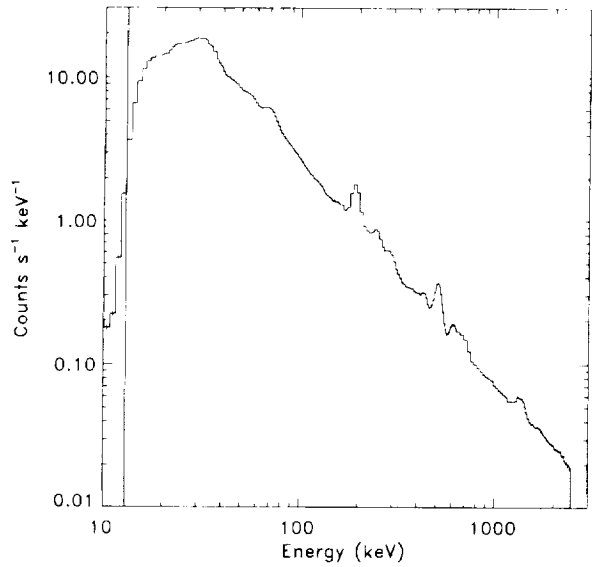

FIG. 1.—Typical SHER background counts spectrum from 3B 940329 accumulated from −3356 to +2946 s, relative to the trigger, in SD 4 at 4× gain. At the low end, the calculated position of the LLD is indicated by the vertical line. For clarity, several of the lowest and highest channels have been omitted by the selection of the horizontal axis. Several prominent features also stand out with good statistics in this long accumulation: the Na K-edge, visible as a change in slope between 30 and 40 keV; the 511 keV electron-positron annihilation line; as well as several activation lines, including the brightest at ~200 keV.

of the LLD in the corresponding LLD spectra, to within the stated accuracy of two channels (e.g., Fig. 1). The channel-to-energy conversion algorithm of Band et al. (1992) relates a SHER channel number to its lower energy threshold, given the gain setting of the detector. The energy threshold of the LLD SHER channel is the same as that of SFAST2, the upper edge of the differential DISCSP1 channel. The lower edge of DISCSP1 is derived from the upper: the voltage at which SFAST1 triggers is designed to be one-half that of SFAST2. Due to interference between SFAST1 and SFAST2, the actual voltage ratio, as measured in the lab, is slightly different than 0.5. For typical operation of the detector in orbit, this introduces a 0.4 keV systematic uncertainty in the determination of the SFAST1 thresholds, which is less than our 1 keV error estimate for a detector at 8× gain. Also, one-half the light output of the NaI crystal is close to, but not exactly equal to, half the energy-loss threshold. This is because the light output versus photon energy for NaI is nonlinear. The nonlinearity for the SDs was measured before launch (Band et al. 1992). We interpolate the prelaunch measurements to determine the lower energy-loss threshold. The lowest energy datum is at 8.0 keV; an extrapolation of the light output curve is used at lower energies, which introduces only a 2.5% error at 6 keV.

4. TESTING

We have tested our calibration of the energy edges of the DISCSP data in several ways. First, the overall accuracy was assessed by comparing spectral fits to data obtained by several detectors simultaneously. If one of the detectors is at a low gain setting, while another is at a high gain, the DISCSP1 energy interval for the low-gain detector can overlap the SHERB data of the detector at high gain. The raw counts cannot be compared directly due to the different detector-to-source geometries, so a joint spectral fit of the observed spectra from the two detectors for a well-located (bright) GRB is necessary. In the cases studied, the count rate of the DISCSP1 bin of the low-gain detector lay within 2 σ or 4% of the predicted model rate. Figure 2 gives an example of how well this verification procedure works for 3B 940703, using SHERB data from an 8× gain detector (SD 5) and DISCSP1 data from a 0.4× gain detector (SD 1). The figure shows both sets of data on one plot by converting the data count rates to photon fluxes in the usual forward-folding model-dependent manner. Both the error bars and the offset of the DISCSP1 model rate from the data are invisibly small on the scale of the plot, due to the large number of counts in the data. The good statistics ensures that this technique is very sensitive. The DISCSP1 data from SD 1 cover approximately the energies 35–70 keV. The edges are determined to within 10% of the mean energy at this gain; that is, shifting them by 5 keV changes the count rate by 1 σ (for a detector at 4× gain, this means a 1 σ uncertainty would be 0.5 keV, well within our 1 keV error estimate). Leaving out this single data point from the joint fit decreases χ^2 by ≈ 2, while removing one degree of freedom.

As a second test, we analyze sources that have well-defined spectral behavior in the low energies and see how well the DISCSP1 data agree with the spectral model fit at higher energies. Solar flares are a comon BATSE burst trigger (with a well-known source location!) throughout the mission, and many are quite well fit by the sum of power-law and optically thin thermal bremsstrahlung (OTTB) spectra (Schwartz 1995). An example joint fit using this two-component model is illustrated in Figures 3*a* and 3*b*, which show count and photon spectra, respectively, from a C7.3 solar flare that occurred at 1715 UT on 1992 June 26. The

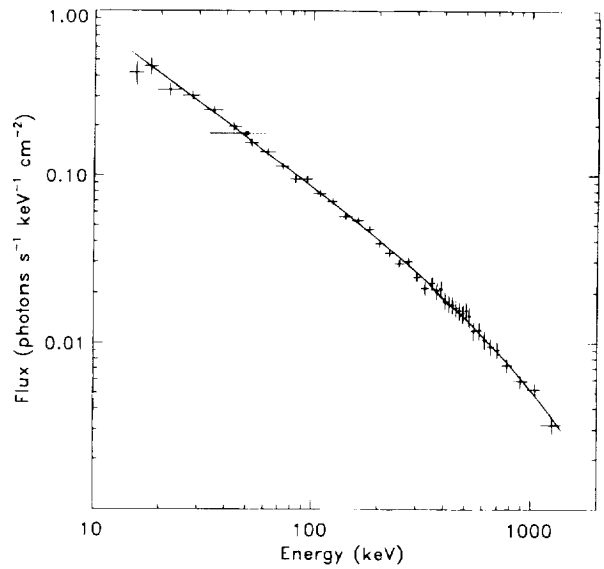

FIG. 2.—Deconvolved photon spectrum for 3B 940703, accumulated from 30.6 to 33.5 s after the trigger. SHERB data (*crosses*) are shown for SD 5, which is at high gain (8×), and compared with DISCSP1 data (*horizontal line with dot*) for SD 1, which is at low gain (0.4×). A continuous curve, representing the fitted GRB photon model, overlays the data. The DISCSP1 channel lies 0.25 σ from the model, which has $\chi^2 = 172.5$ for 213 degrees of freedom. Note that the errors on the DISCSP1 rate and the horizontal line indicating the model rate are invisible on this scale, also that the placement of the dot is not indicative of the weighted average of the model rate over the energy interval; rather, it is positioned arbitrarily at the average of the energy thresholds of the interval.

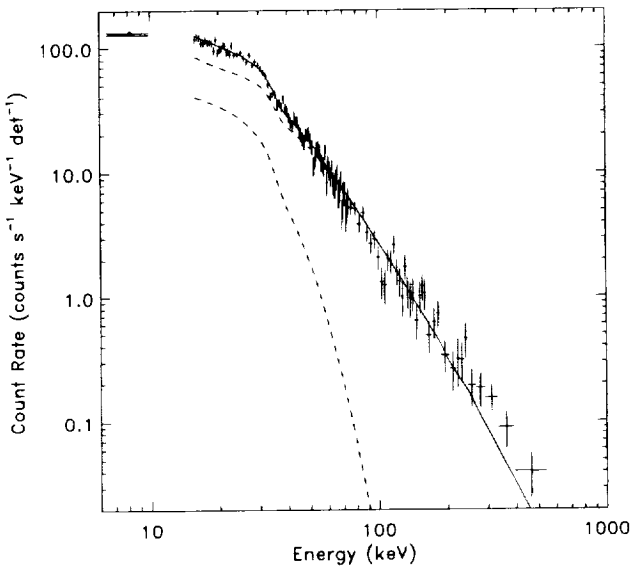

FIG. 3a

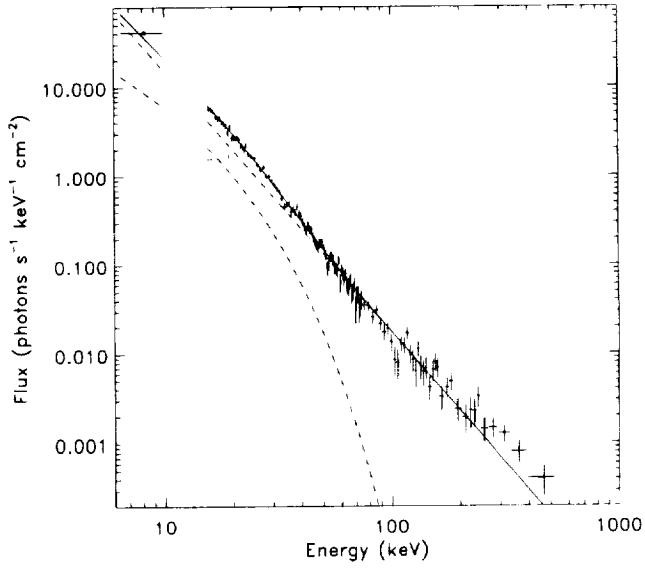

FIG. 3b

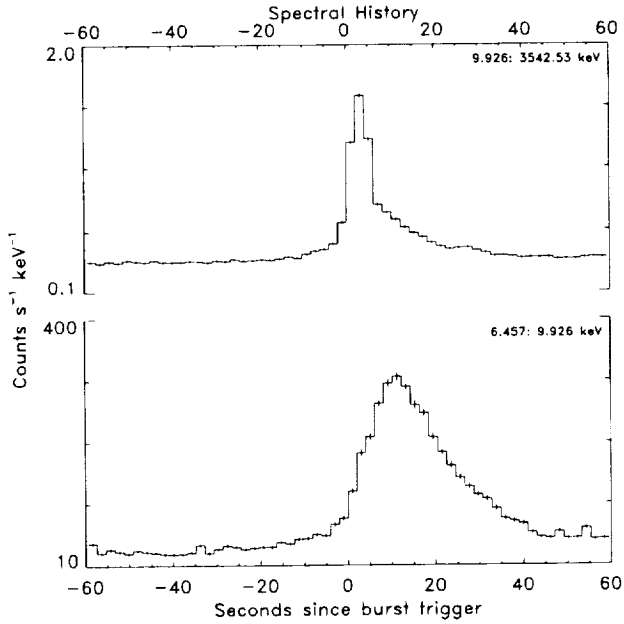

FIG. 3c

FIG. 3.—(*a*) Joint fit to counts accumulated in DISCSP1 and SHERB in SD 0 during the first 6.14 s after the trigger for the 1992 June 26 solar flare at 1715 UT. The DISCSP1 data point is indicated by a dot. The two separate components of the fit, indicated by dashed lines, are a power law of index −2.9 and OTTB with $kT = 9.4$ keV. The total model is indicated by the continuous curve for the SHERB data and by a horizontal line for DISCSP1. (*b*) Photon spectrum for the same fit as in (*a*), obtained by forward-folding the count spectrum through the detector response. (*c*) Count rate history for DISCSP1 (*bottom*) and DISCSP2 (*top*), showing the different temporal behavior of the two components.

two spectral components have distinct time histories, as shown in Figure 3*c*, with the thermal component having a peak that is later than, and broader than, the higher energy emission. The spectrum in Figures 3*a* and 3*b* is dominated by OTTB emission with a fitted temperature of 14 keV. The total model is shown as a continuous line, while the separate components are indicated by dashed lines. The DISCSP1 energy channel covers approximately 5–10 keV and is consistent with the extrapolation of the thermal component fit to the higher energy data. This is a sensitive test because the small errors constrain the fit tightly. A 1 keV shift in the edges changes the DISCSP1 counts by 1 σ; thus, the determination of the edges is good to within 13% of the mean energy. The 1992 June 26 solar flare was unusually hard; the DISCSP1 rates for most flares are considerably higher, relative to the SHERB data, than in this example. For these cases, the joint fit of an OTTB component may be driven strongly by the single, high DISCSP1 rates. The six other flares we examined all resulted in good fits over their entire energy ranges; half of these were well fit by a single power-law component model, and none exhibited an anomalous residual in DISCSP1, which would indicate a failure of the energy calibration.

5. ANALYSIS AND OBSERVATIONS

The main objective of this work is to determine the relationship between the low-energy ($\approx$5–20 keV) photon flux and the spectrum at higher energies. In particular, do

any GRBs in the ensemble exhibit significant excesses or deficits in the low-energy regime? Bursts with large deviations may indicate either that extra curvature is required for a truly acceptable spectral model, that an additional low-energy spectral component exists, or possibly that there is some absorption at low energies. X-ray paucity can be understood within the current state of spectral analysis by the observation that bursts generally emit most of their power at a peak energy that lies between 40 keV and 2 MeV. That is, the concave-downward curvature inherent in most fitted spectra naturally satisfies X-ray paucity (Band et al. 1993; Ford et al. 1995). Here we wish to compare the observed DISCSP1 data to the rate predicted from a model fitted to both the DISCSP1 and the higher-energy SHERB data.

We have selected our burst sample from the entire set of available bursts, from the beginning of the BATSE mission up until the end of 1995, with the requirement that the total fluence (as determined by fits to the LAD data when coverage of the entire burst was available) be greater than 2×10^{-5} ergs cm^{-2}. In addition, we include bursts with a signal-to-noise ratio in the 25–35 keV band of the SDs of greater than 7.5. For each burst trigger, there are data available from the four detectors whose associated LADs had the highest count rates at the time of the trigger. We use the calculated burst location to determine which of the four selected detectors have angles to the source of less than 60°. The transmission of the Be window at 5 keV (at 60°: 40% of the 0° transmission, down to $\approx$0% at 85°) as well as the reflective paper (at 60°: $\approx$13% of the 0° transmission) decreases more rapidly than cos θ, which describes more closely the angular response above 10 keV. We require also that a detector have an 8 × or 4 × gain setting, high enough that the mean of the DISCSP1 energy range lies below 20 keV. The gain settings are changed with each spacecraft pointing to satisfy several ongoing science objectives. At any time, at least four out of the eight SDs are at high (8 ×) gain, while one or two are at a medium (4 ×) gain, which is still adequate for our purposes. Eighty-six bursts survived the selection criteria to make up our data set.

The characteristics of the DISCSP1 background are such that a systematic noise component dominates at very low energies. Most of this noise is due to either locally produced low-energy trapped particles or long-lived phosphorescence from energetic single cosmic rays. The latter has the distinctive time profile of a spike-like onset folllowed by a very fast decay, usually less than 3 s. The LLD eliminates most of this noise in the SHERB data, but DISCSP1 is susceptible to it. If we fit a polynomial trend through the DISCSP1 background as a function of time, the residuals of the fit exceed the deviations expected from Poisson counting statistics. We multiply the background model variance by max [1, $(\chi^2_\nu)^{1/2}$], where the χ^2 of the background model is divided by the number of degrees of freedom ν, rescaling the background fit residuals into a Gaussian distribution with the correct width. The rescaled background variances are added in quadrature to the variance of the observed rates to obtain the uncertainties of the background-subtracted data. This works well when the distribution of cosmic-ray spikes during the burst interval are of the same magnitude as those in the selected background intervals, since the spikes are expected to raise the count rate somewhat overall. However, the presence of individual large spikes in the selected burst intervals will not be modeled correctly by our procedure. Spikes that do not also show up in the higher energy data (DISCSP2) are excluded from the selected burst interval. Using this criterion, we have thrown out suspicious fluctuations in the time histories of $\approx$3% of the bursts.

For each burst, the interval with the largest signal-to-noise ratio in the energies between 25–35 keV was selected, which sometimes included several peaks. We sum all count spectra in the selection, finally dividing by the total live time to obtain a single, time-averaged spectrum. The same time intervals were selected for both the SHERB and DISCSP1 data, whether or not there was any apparent emission in the DISCSP1 data. We generally select SHERB energy channels starting directly above the LLD to use in fitting; however, a local nonlinearity in the low-energy SHERB channels, called SLED (spectroscopy low-energy distortion; Band et al. 1992), sometimes affects the selection of useful data. Although it has been corrected largely through recalibration of the SHER energy-loss thresholds (Band et al. 1992), residual nonlinearities due to the SLED persist in roughly one-third of the bursts. For those cases, we do not choose SHERB channels between the LLD channel and the lowest channel unaffected by the SLED, which creates a gap about 10 keV wide above the DISCSP1 upper edges at 8 × gain. We also do not include the highest channel of the SHERB data type, which is an integral channel with no well-defined upper edge.

The combined data set of the time-averaged SHERB spectrum and count rate in DISCSP1 for each burst is fit jointly by the standard forward-folding technique. This procedure convolves a chosen photon spectral model with the DRM by straightforward matrix multiplication, thus generating a trail count spectrum that is compared to the actual count spectrum using the χ^2 statistic. This test statistic is minimized by variation of the photon model spectral parameters until a best fit is obtained. We use the Levenberg-Marquardt nonlinear fitting algorithm (Press et al. 1992, p. 675), modified to make use of the mathematically correct model variances rather than the data variances (Ford et al. 1995). Even when an optimized spectral model has an acceptable χ^2 value, the solution is not unique, since other (untried) spectral models may also have acceptable χ^2 values.

In order to have a clear basis for comparison between fits to different bursts, we start out with the "GRB" spectral form of Band et al. (1993). The GRB spectral form has been successfully fitted to SHERB spectra throughout the time histories of a wide variety of bursts (Ford et al. 1995). It has the form of two power-law segments, E^α and E^β, smoothly connected, with continuous derivatives. We have chosen a parameterization of the model that fits the energy of the peak (E_p) of the spectrum in νF_ν; the low- and high-energy power-law indices, α and β, respectively; and the amplitude. E_p is actually defined only when the condition $\beta \leq -2$ holds; otherwise, the parameter value is identical to the "break" energy at which the two power-law components join. In two limiting cases, either for extremely soft spectra, or for spectra that approximate a single power law (the spectral break is outside the BATSE energy band), β becomes unconstrained in the fit, and we can drop the high-energy power law. In the latter case, even though the BATSE spectra do not require a high-energy power law, spectral fits to data (if available) from higher energy detectors such as COMPTEL and EGRET can show that a higher energy power law is actually present and that the

true break energy is close to or above the BATSE upper energy limit (e.g., Hurley et al. 1994; Hanlon et al. 1994). The restricted form of the GRB model without a high-energy power-law segment is similar to unsaturated inverse-Comptonized thermal emission (Rybicki & Lightman 1979, p. 221), which consists of a low-energy power law attenuated by an exponential cut-off. A large fitted value for E_p in this case would indicate that the spectral break is too high in the BATSE energy range to allow β to be determined reliably. We have used this exponential form in spectral fits for 20 bursts in our sample, nine of which had values of E_p greater than ~ 700 keV.

We present a typical example burst, GRB 941020, in Figure 4. The burst count rate time histories, shown in Figure 4*a*, demonstrate that there is some flux present in DISCSP1 (*bottom*), which has a similar profile to that of DISCSP2 (*top*). In this case, DISCSP1 covers the energies 6.4–9.9 keV and has the same time resolution as the SHERB data (no background subtraction has been performed in the figure). Note that the two histories are not identical, since the evolution of the second peak at around 55–60 s is significantly different in the two energy ranges. For the joint fit, DISCSP1 and SHERB spectra from 15.7 to 70.3 s are summed together to obtain an average spectrum. The joint, background-subtracted, time-averaged count spectrum is shown in Figure 4*b*, along with the fitted model, which is indicated by a continuous curve. For the single DISCSP1 point, indicated by the filled circle, a horizontal line indicates the model rate over the entire bin. Some of the higher energy SHERB channels have been binned together to increase the readability of the figure; the individual unbinned channels were used in the actual fit. A value of

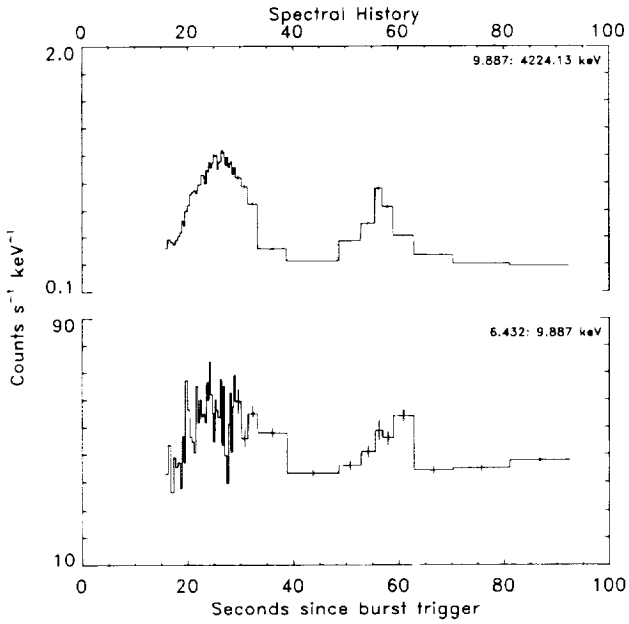

FIG. 4*a*

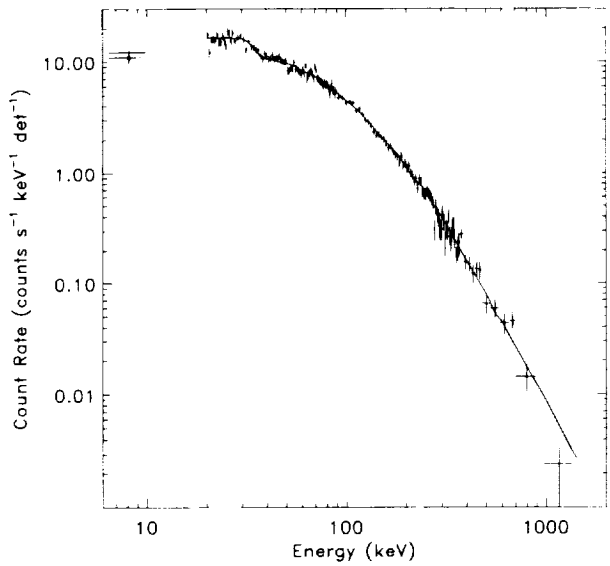

FIG. 4*b*

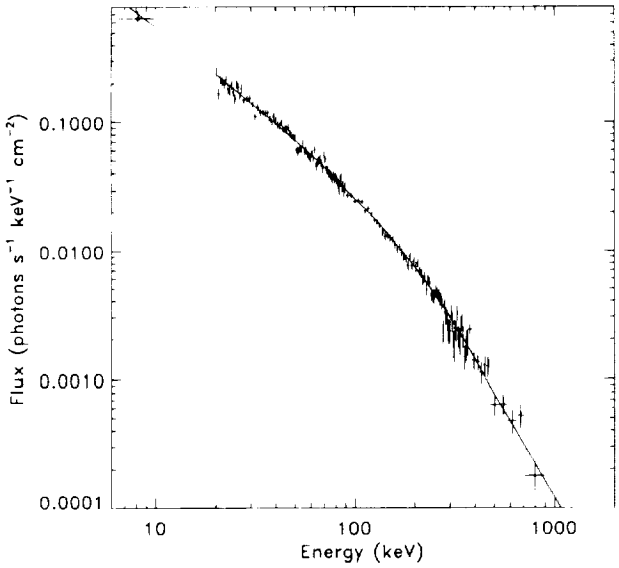

FIG. 4*c*

FIG. 4.—Example of a burst with good agreement between the low-energy data point and the fitted model. (*a*) Count rate history for GRB 941020, for DISCSP1 and DISCSP2 in SD 5. (*b*) GRB model joint fit to the interval 15.7–70.3 s of GRB 941020, in SD 5, with fitted parameters $E_p = 191 \pm 7$ keV, $\alpha = -1.16 \pm 0.03$, $\beta = -2.7 \pm 0.1$, and $\chi^2 = 221.3/201$ degrees of freedom. (*c*) Photon spectrum for the same fit as in (*b*).

$\chi^2 = 221.3$ for 201 degrees of freedom was obtained in a joint fit of the data, using the GRB spectral form. Figure 4*c* plots the photon model on top of the flux in each channel, which is inferred from the count data in the usual way by calculating the detector's model-dependent efficiency to detect photons at that energy, and dividing the derived efficiency into the actual counts accumulated. The good statistics of the lower energy SHERB channels constrain the fitted value of the low-energy power-law index quite well ($\alpha = -1.16 \pm 0.03$). In addition, $E_p = 191 \pm 7$ keV is high enough in energy to ensure that the low-energy power law can be well determined. As seen in Figure 4*c*, the fitted model curve is nearly a pure power law below ≈ 30 keV. The model lies 0.8 σ above the DISCSP1 rate, which is 90% of the model rate; there is clearly no evidence for any significant X-ray excess or deficit in this example.

We need a quantitative measure of the deviation of the DISCSP1 count rate from the model rate at DISCSP1 energies for each time-averaged burst spectrum. The difference between the model and background-subtracted source counts for DISCSP1 can be expressed in units of σ by dividing by the error derived from the variance of the model, a quantity we call the sigma residual. We use the model variance in the calculation of the parameters of the spectral model in the fit, since in that case we are assuming that the best-fit model is correct; in calculating the sigma residuals, we are testing the hypothesis that the model is correct and that a deviation in the data represents a fluctuation around the model rate, considered as a mean. We also express the magnitude of the deviation in each burst in physical terms as a percentage, dividing the observed counts in DISCSP1 by the fitted model. Thus, bursts with an excess or deficit will have a DISCSP1 count rate more or less than 100% of (or 1.0 times) the model rate. The joint fit is tightly constrained by the SHERB data at higher energies, so the extension of the model down to lower energies is very much like an extrapolation. In some cases, the DISCSP1 rate may be significantly higher or lower than the model rate, significantly increasing χ^2 overall.

The results of our survey are summarized in Table 1, which contains the burst name, DISCSP1 channel edges, sigma residual of the DISCSP1 count rate, fractional deviation of DISCSP1, and the total fluence of the interval chosen. We do not show the parameters of the spectral fits; they are similar to those of Band et al. (1993), except for those cases in which the model is pulled severely by a low-energy deviation. We have indicated in Table 1 which bursts were also observed by WATCH (Castro-Tirado 1994). Although it is difficult to compare the two catalogs directly, in most of the cases in which Castro-Tirado has indicated an X-ray excess, we also have an excess. Also indicated in Table 1 are those bursts for which the fitted value for E_p is less than 45 keV. In these four bursts, the low-energy power-law index is not well constrained by the data, as discussed below.

Figure 5*a* is a plot of deviations from the fit for each burst in the sample from Table 1, defined as the observed count rate in DISCSP1 divided by the model, as a function of the average of the two DISCSP1 energy thresholds. The points cluster around a deviation of 1.0, which indicates agreement between the data and the fitted model. The vertical bars indicates the 1 σ errors of the deviation for each burst. The deviations at energies below 8 keV are consistently greater than 1, although the significances of these points are typically less than 2 σ, with the largest at 4.4 σ. The consistent excess at these energies may be an indication that our determination of the SFAST1 threshold may suffer from a systematic uncertainty at energies below 8 keV. However, as discussed in § 3, these uncertainties are only 2.5% at a mean energy of 6 keV. Indeed, applying a worst-case systematic error of 10% all across the board would not change the significances for most of the observed rates. As discussed above, the systematic error in the energy calibration due to the uncertainty in the position of the LLD is actually 4%.

Figure 5*b* shows the significances illustrated in Figure 5*a* more clearly by plotting sigma residuals for the same bursts. Notice that there are both positive and negative sigma residuals. Clearly, the majority of the bursts in our sample have a low-energy flux consistent with the overall fitted spectrum. The most frequent values for the residuals are

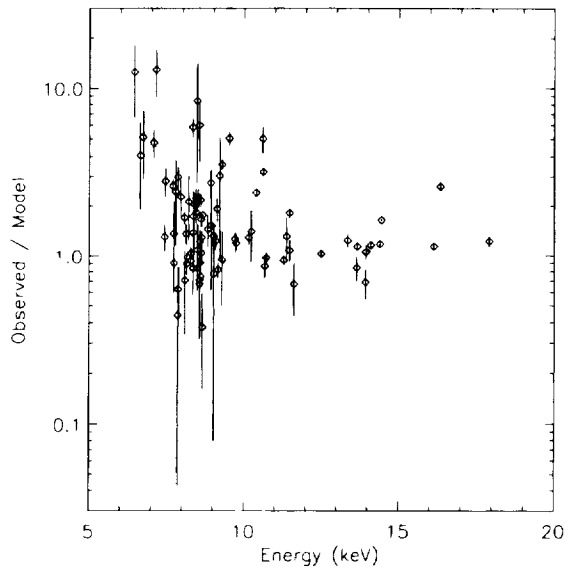

FIG. 5*a*

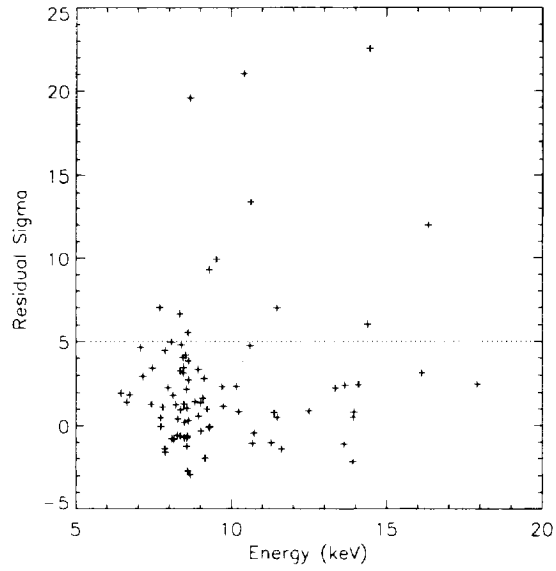

FIG. 5*b*

FIG. 5.—(*a*) Observed counts divided by modeled counts from the fitted DISCSP1 data for 86 bursts observed by BATSE, vs. the average of the DISCSP1 energy edges. Error bars are indicated by the vertical lines. (*b*) Sigma residuals of fitted DISCSP1 data for 86 BATSE bursts vs. the DISCSP1 average energy.

TABLE 1

RESIDUALS OF JOINT FITS

Burst Name (3B Catalog)[a]	DISCSP1 E Range (keV)	Sigma Residual[b]	(Obs. Counts) / Model	Fluence[c] (Photon cm^{-2})
3B910425[d]	8.6 – 14.2	0.77	1.30 ± 0.39	110.8
3B910503	10.2 – 18.7	22.54	1.64 ± 0.03	280.9
3B910717	9.8 – 17.5	–1.10	0.84 ± 0.14	34.8
3B910807	9.9 – 17.9	–2.17	0.69 ± 0.14	103.8
3B910809[e]	6.7 – 10.3	0.19	1.09 ± 0.46	105.4
3B910814C	6.7 – 10.3	1.31	8.41 ± 5.64	123.8
3B910814[d]	11.3 – 21.4	11.99	2.64 ± 0.14	142.9
3B910927[d]	8.7 – 14.3	0.49	1.08 ± 0.16	77.2
3B911016[d]	6.5 – 10.0	1.27	2.13 ± 0.89	34.9
3B911031	5.8 – 9.1	3.42	2.83 ± 0.53	100.3
3B911106	6.1 – 9.5	1.12	2.47 ± 1.31	102.2
3B911118	6.5 – 10.0	0.39	1.04 ± 0.10	375.7
3B911126	6.4 – 9.8	4.96	1.69 ± 0.14	243.1
3B911127	7.1 – 11.1	1.67	1.23 ± 0.14	129.8
3B911202[d]	8.0 – 12.8	21.03	2.43 ± 0.07	169.3
3B911209[d]	6.9 – 10.7	1.46	1.44 ± 0.30	79.7
3B920110	8.1 – 13.1	4.74	5.01 ± 0.85	67.9
3B920130	6.1 – 9.4	–0.05	0.90 ± 1.85	112.9
3B920210	5.0 – 7.9	1.98	12.47 ± 5.80	112.5
3B920221	6.8 – 10.4	5.55	2.18 ± 0.21	59.6
3B920513	6.7 – 10.4	2.20	5.99 ± 2.27	120.9
3B920517	6.8 – 10.5	19.56	1.75 ± 0.04	57.5
3B920525	5.2 – 8.2	1.87	5.11 ± 2.19	224.0
3B920617	7.0 – 10.9	3.34	2.76 ± 0.53	136.3
3B920622	6.6 – 10.2	4.79	2.03 ± 0.21	331.5
3B920627	6.2 – 9.5	–1.39	0.44 ± 0.40	91.9
3B920711	7.4 – 11.6	9.92	5.04 ± 0.41	689.8
3B920720[d]	6.6 – 10.1	6.64	5.83 ± 0.73	24.2
3B920723[d]	6.2 – 9.6	–1.63	0.63 ± 0.23	261.5
3B920902[d]	12.2 – 23.6	2.49	1.22 ± 0.09	110.6
3B921003	6.5 – 10.0	–0.57	0.93 ± 0.12	106.9
3B921009	6.4 – 9.8	–0.79	0.71 ± 0.37	247.9
3B921123	6.7 – 10.3	4.02	2.24 ± 0.31	329.1
3B921206	7.3 – 11.3	–0.12	0.95 ± 0.45	165.9
3B921207	7.0 – 10.3	–2.94	0.37 ± 0.21	290.9
3B921209	6.8 – 10.4	1.05	1.22 ± 0.21	58.3
3B930120	6.8 – 10.4	–2.74	0.75 ± 0.09	245.9
3B930127[e]	6.8 – 10.4	–0.62	0.90 ± 0.16	143.1
3B930201	9.8 – 17.6	2.45	1.14 ± 0.06	311.0
3B930405	7.3 – 11.3	9.28	3.56 ± 0.28	80.9
3B930406	6.6 – 10.1	–0.61	0.84 ± 0.27	42.7
3B930425	11.2 – 21.1	3.13	1.14 ± 0.04	87.3
3B930426	7.1 – 11.0	1.38	1.32 ± 0.23	108.3
3B930523	8.2 – 13.3	–0.44	0.97 ± 0.07	107.5
3B930612	7.1 – 11.0	–0.32	0.78 ± 0.70	208.8
3B930614[e]	8.6 – 14.0	–1.03	0.94 ± 0.06	56.9
3B930706	6.8 – 10.5	0.29	1.04 ± 0.14	50.2
3B930916	6.7 – 10.3	–0.73	0.84 ± 0.22	337.3
3B930922	9.2 – 15.8	0.88	1.03 ± 0.03	96.1
3B931024	7.6 – 11.9	1.17	1.19 ± 0.16	69.2
3B931026	7.9 – 12.6	0.83	1.40 ± 0.48	136.2
3B931031	6.8 – 10.5	3.83	1.65 ± 0.17	38.3
3B931103	6.8 – 10.4	–1.24	0.67 ± 0.26	112.2
3B931204	10.2 – 18.6	6.04	1.18 ± 0.03	216.6
3B940218	7.6 – 11.9	2.36	1.26 ± 0.11	120.9
3B940228	8.1 – 13.1	13.41	3.19 ± 0.16	109.9
3B940302	8.8 – 14.5	–1.42	0.67 ± 0.23	761.1
3B940329[e]	7.2 – 11.2	–1.98	0.82 ± 0.09	138.6
3B940330	7.2 – 11.2	0.99	3.03 ± 2.05	115.0
3B940414	8.2 – 13.2	–1.07	0.86 ± 0.13	154.2
3B940429	5.2 – 8.1	1.42	4.07 ± 2.16	153.7
3B940526B	5.6 – 8.7	2.96	12.87 ± 4.02	26.6
3B940526	6.3 – 9.6	2.29	2.29 ± 0.56	94.6
3B940703	6.2 – 9.5	4.45	2.98 ± 0.45	837.3
GRB940921	6.1 – 9.4	7.01	2.65 ± 0.24	81.8
GRB941008	6.4 – 9.9	1.84	1.36 ± 0.19	288.9
GRB941014	6.6 – 10.1	0.95	1.72 ± 0.76	176.5
GRB941017	6.7 – 10.4	4.17	2.25 ± 0.30	549.8
GRB941020	6.4 – 9.9	–0.81	0.90 ± 0.13	434.1
GRB950111	6.1 – 9.4	0.47	1.36 ± 0.76	193.1
GRB950208	7.2 – 11.1	2.82	1.91 ± 0.32	336.8
GRB950211	6.8 – 10.4	–0.74	0.71 ± 0.39	207.3
GRB950305	5.5 – 8.6	4.64	4.78 ± 0.81	40.6
GRB950325	6.6 – 10.1	3.26	1.38 ± 0.12	114.0
GRB950403_47987	5.8 – 9.1	1.29	1.30 ± 0.23	250.3
GRB950403_84826	6.8 – 10.5	2.74	1.29 ± 0.11	181.7
GRB950425	8.7 – 14.3	6.98	1.80 ± 0.11	257.4
GRB950513	9.6 – 17.1	2.26	1.24 ± 0.10	52.1
GRB950701_12758	9.9 – 18.0	0.50	1.05 ± 0.10	82.7
GRB950701_23737	9.9 – 18.0	0.81	1.06 ± 0.07	70.5
GRB950804	10.0 – 18.2	2.47	1.16 ± 0.06	71.3
GRB950818	6.7 – 10.3	3.44	2.09 ± 0.32	182.4
GRB951011	7.0 – 10.9	0.56	1.51 ± 0.91	60.2
GRB951102	6.7 – 10.3	3.14	1.82 ± 0.26	30.7
GRB951104	7.3 – 11.4	–0.06	0.94 ± 0.98	83.7
GRB951203	7.8 – 12.5	2.37	1.28 ± 0.12	148.6

[a] BATSE 3B Catalog names (Meegan et al. 1996) have been assigned only for bursts up to 1994 September 19; thereafter, we have given the traditional name: GRB*yymmdd*. Seconds of day is appended to distinguish between two events on the same day.

[b] Defined as (observed counts – model counts)/data error.

[c] Defined as the fitted photon spectrum integrated over the selected times and energies, which, depending upon gain, is usually 5 keV–2 MeV; these are thus different from the 3B Catalog values, given as an indication of relative intensity.

[d] Observed also by WATCH-*Granat* (Castro-Tirado 1994).

[e] Fitted value of E_p is below 45 keV.

within $\pm 2\ \sigma$; however, there is a significant excess of bursts with positive deviations above 5 σ (12 out of 86, or 14% of the sample). The mean energies for DISCSP1 of the bursts that have a large positive deviation range from 8 to 16 keV. The large number of significant excesses implies that, at the very least, some bursts require either a spectral form that turns up at low energies or an additional low-energy spectral component. We have verified for at least one burst (3B 920517) that an observed large excess is consistent in two detectors that had similar viewing angles and gains. There are no bursts with low-energy deficits greater than 3 σ, which is consistent with the statistics of our sample: given 86 bursts, we expect 0.14%, or 0.1 bursts with deviations greater than 3 σ, for a Gaussian probability distribution.

6. DISCUSSION

For the majority of the bursts sampled, we find good agreement with the GRB model at low energies (e.g., Fig. 4). Given the tremendous variety of burst time histories, it is remarkable that a single spectral form with two power-law segments suffices to fit the time-averaged spectra of the majority of bursts over a broad energy range. The present work expands upon the results of Band et al. (1993), extending fits down to lower energies, and thus obtaining better determinations of the fitted values of α. We also see strong evidence for the breakdown of this simple form at low energies for 14% of the bursts.

We show in Figure 6 the relationship between the fitted values of α and E_p, which agrees well qualitatively with Figure 5 of Band et al. (1993). The most probable value for α is near -1.0, while the bounds of the distribution are at

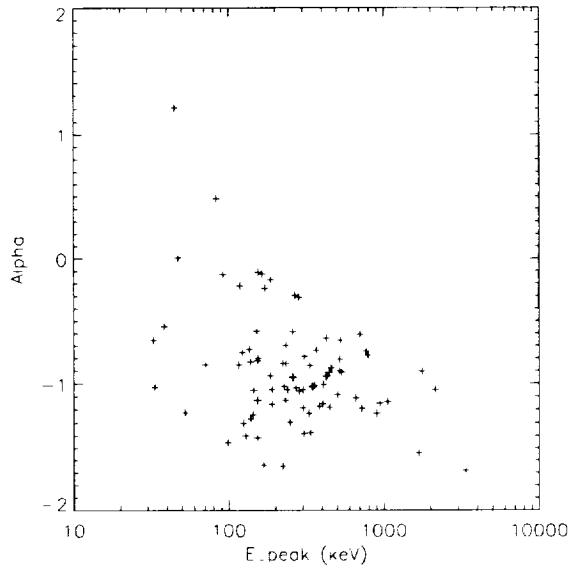

FIG. 6.—Correspondence between two parameters of the GRB model, α and E_p, for the joint fits used in producing Fig. 5. Errors on the parameters are not shown.

-1.7 and 1.2. Quantitatively, we find good agreement between their fitted values for α and ours, except in those cases in which a difference is expected, i.e., bursts with an excess at low energies as well as bursts with a value for E_p below 45 keV. In each case in which a low-energy excess has influenced the fitted value of α, the difference is consistent with the residual sigma for the burst reported in Table 1. We note that $\alpha > -2$ naturally satisfies X-ray paucity, with increasingly less power per decade toward lower energies. As discussed in Band et al. (1993) and Mallozzi et al. (1995) and also seen here in Figure 6, the fitted values of E_p have a broad distribution, covering nearly the entire available energy range. However, the average value for α clearly changes with energy: below $E_p \approx 100$ keV, the average value for α is ≈ -0.3, while it is lower (≈ -1.0) for higher values. The entire upper right half of the plot is unpopulated.

The most significant deviations from the assumed model occur as low-energy excesses, as seen in Figure 5*b*. The effect is observed at all mean energies from 8 to 16 keV; no particular energy appears to be a threshold for this behavior. Although the largest excesses seen in Figure 5*a* tend to occur at lower energies, the most significant excesses, greater than 5 σ (Fig. 5*b*), occur at all energies, From Table 1, the observed versus model counts ratios for these bursts range between 1.2 and 5.8. The largest factor observed, 13 times the model (for 3B 940526B), is significant at 3 σ. We note that whereas the survey of Castro-Tirado (1994) found that 10% of the bursts had 6–15 keV X-ray emission that was extended in time beyond the gamma-ray outburst, here we have found that 14% have X-ray excesses in their spectra *simultaneous* with the burst. This raises the possibility that hard-to-soft spectral evolution within a burst (Ford et al. 1995) might account for low-energy excesses in the time-averaged spectra. However, this is not evident in the time history of 3B 920517 (Figure 7*a*), which has one of the most significant low-energy excesses. In contrast, there is some evidence for excess soft emission in the second peak in the time history of 3B 941020 (Fig. 4*a*), where no excess was found in the time-averaged spectrum. A time-resolved spectral analysis of the bursts in which a low-energy excess is present should determine better how well the low- and high-energy emission track each other, and this will be the topic of a later paper in this series.

If we fit an additional component, such as optically thin thermal bremsstrahlung (OTTB), to the data for bursts that show an excess, we can obtain quite good fits (see Fig. 7). However, the extra component of the model is not unique (almost any component that is steeper than a -2 power law will do), especially since the parameters are largely constrained only by the single DISCSP1 channel. In most cases, there is no corresponding enhancement of counts in the low-energy bins of the SHERB data to help constrain the fit; which is in contrast to fits to solar flare data, as seen in Figure 3, where the SHERB data constrain the thermal component of the model quite well. The Compton attenuation model of Brainerd (1994) also results in good fits for most cases in which there is a low-energy enhancement, since it predicts a concave upward low-energy tail to the data. To account for X-ray paucity, this model includes absorption at the site of the GRB source.

In some bursts, the fitted value for E_p is so low in energy that it falls near the lowest energy of the SHERB data, which by itself can be fitted by a single power law. In addition, the DISCSP1 data point falls below the extrapolation of this power law down to energies covered by DISCSP1. In such cases, the GRB model obtains a good fit by allowing the low-energy power-law segment to compensate for the deviation of the DISCSP1 data. An example of this effect is shown in Figure 8 for a spectrum accumulated in the first 96.4 s of 3B 930127. The fitted value of E_p is only 38 keV, so the low-energy power law is not very well constrained ($\alpha = -0.5 \pm 0.4$), especially taking into account the curvature of the GRB spectral form below E_p. The low-energy power law is essentially determined by two points DISCSP1 and the lowest energy resolution element (group of channels totaling one resolution width) of the SHERB data. This is backed up by the fact that the observed and model DISCSP1 rates agree very well. Their ratio from Table 1 is 0.9 ± 0.16, or roughly $-0.6\ \sigma$. Without including the DISCSP1 datum in the joint fit, the constraint on the low-energy power-law index is poorer: $\alpha = -1.0 \pm 0.6$, and the data can just as well be fit by a single power law ($\chi^2 = 221.7$ for 219 degrees of freedom, without DISCSP1 vs. $\chi^2 = 304.6/220$, with DISCSP1). A single power-law spectrum with an index of roughly -2 extrapolated down to the X-ray (or lower!) energy band emits too much power in X-rays to be consistent with the observation of X-ray paucity in GRBs, and indeed we do not observe this. Thus we expect the spectrum from such a burst to exhibit curvature at low energies and interpret a deficit of counts in DISCSP1 as evidence for this curvature. In all such cases we have kept the GRB spectral form as the model, and we indicate in our results, presented in Table 1, the four cases in which the fitted value of E_p falls below 45 keV.

While the BATSE SDs are not optimized for detecting low-level X-ray emission in GRBs, it is worth noting that we have observed low-energy excesses in a similar fraction of bursts as reported by *Ginga* and WATCH, which detected X-ray precursors and tails. Indeed, where we have observed bursts in common with WATCH, we generally agree on which bursts have excess X-ray emission. Where fits were made to the X-ray precursor and tail spectra obtained by *Ginga*, they were consistent (but not conclusively) with blackbody emission. Blackbody emission equivalent to the Eddington luminosity from a compact solar mass object, such as a neutron star, will have a characteristic energy of about 1 keV. For a typical bright (10^{-5} ergs s^{-1} cm^{-2}) GRB observed by BATSE, such an object would have to be closer than ~ 100 pc. However, the blackbody interpretation is far from required by the data.

If GRB sources are at cosmological distances relative to us, what can we say about the observation of low-energy excesses? First, many cosmological scenarios for GRBs involve a series of shocks arising between relativistic matter and interstellar material at the source (Meszáros et al. 1993; Katz 1994; Tavani 1995). If several bunches of beamed particles exist, then it is likely that there will be several components to the emission, possibly at different energies. The simplest example of this would be an observation of the remnant of the original thermal pair fireball: either shocked particles from the expanding fireball are reaching the ISM or possibly the fireball is emitting strongly enough to be observed by itself, after adiabatic cooling. Timing would be crucial in this case for both components to be observed simultaneously. However, precursor and tail X-ray emission could be a result, as well as a low-energy spectral component that has a time history distinct from the bulk of the

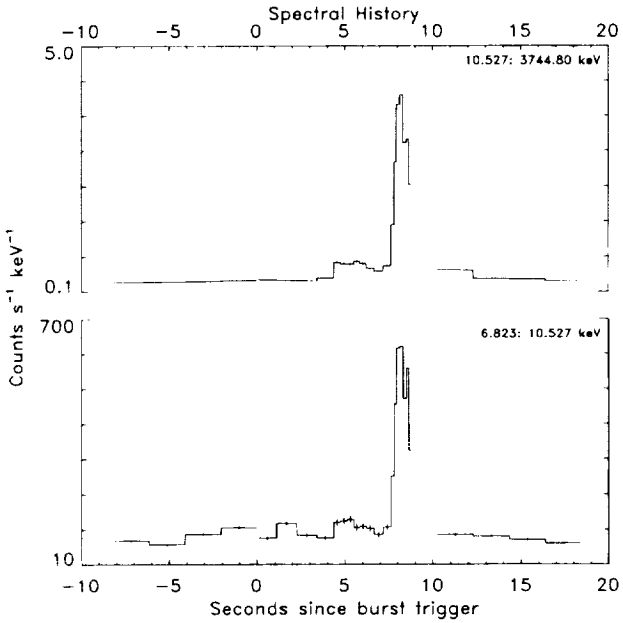

FIG. 7a

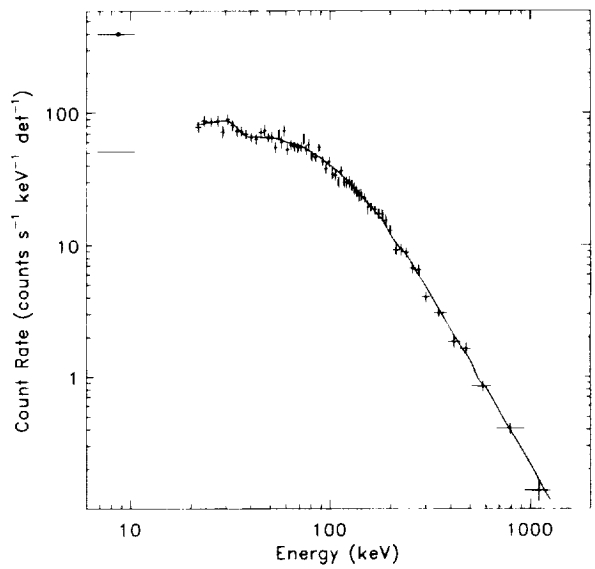

FIG. 7b

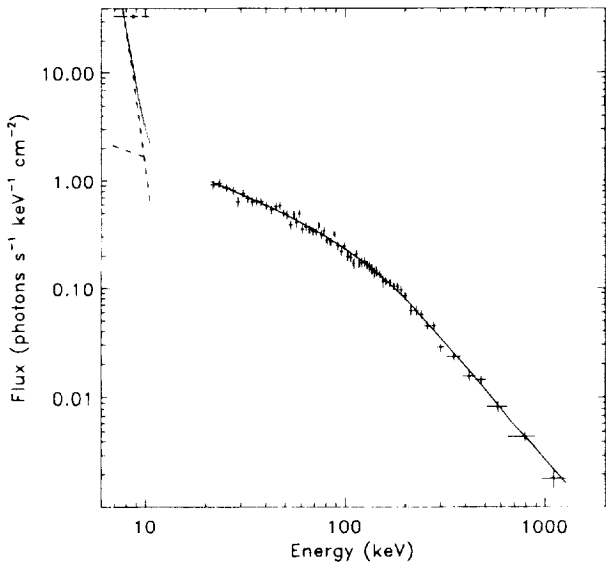

FIG. 7c

FIG. 7.—(*a*) Count rate history for 3B 920517, in DISCSP1 and DISCSP2, SD 7. (*b*) GRB model plus OTTB fit to spectra accumulated from 7.6 to 8.8 s after the trigger for 3B 920517, with fitted parameters $E_p = 216 \pm 17$ keV, $\alpha = -0.61 \pm 0.08$, $\beta = -2.10 \pm 0.06$, $kT = 1.5 \pm 9$ keV, and $\chi^2 = 189.0/195$ degrees of freedom. Note that the temperature is not well constrained by the fit. The GRB model component is shown as a solid continuous line, except for DISCSP1 (data point indicated by a dot), where it is a horizontal line. The total model differs from the GRB model only at the DISCSP1 point, where it agrees identically with the data rate. (*c*) Photon spectrum for the same fit as in (*b*), showing the two separate components as dashed lines and their sum as a solid continuous line.

gamma-ray flux. Brainerd's Compton attenuation model, discussed above, is another possible source of excess low-energy emission consistent with a cosmological source scenario. This model has specific predictions that can be tested, one of the most important of which is the simultaneity of the low- and high-energy spectral regimes, which are not derived from separate emission components. A future paper will describe the temporal evolution of the low-energy excesses observed by BATSE and will address the issue of simultaneity between these two regimes during bursts. It is hoped that our observation of excess emission at low energies in some GRBs will contribute to the development of burst models in the future.

7. SUMMARY AND CONCLUSIONS

This paper presents spectral analyses of 86 GRBs, including low-energy (5–20 keV) emission, comparing the low-energy count rate to the model rate predicted by a joint fit of a canonical spectral form to BATSE SD data. We have extended the spectral coverage of the BATSE SDs to below 20 keV for the first time by an in-flight calibration of the SD discriminators. We quantify the deviation of the observed low-energy discriminator count rate both in terms of the ratio between the observed and modeled rate, as well as by the significance of the residual between the observed and model. We have shown that, for 74 of the GRBs surveyed, a model fitted to data covering all the available energies (~ 5

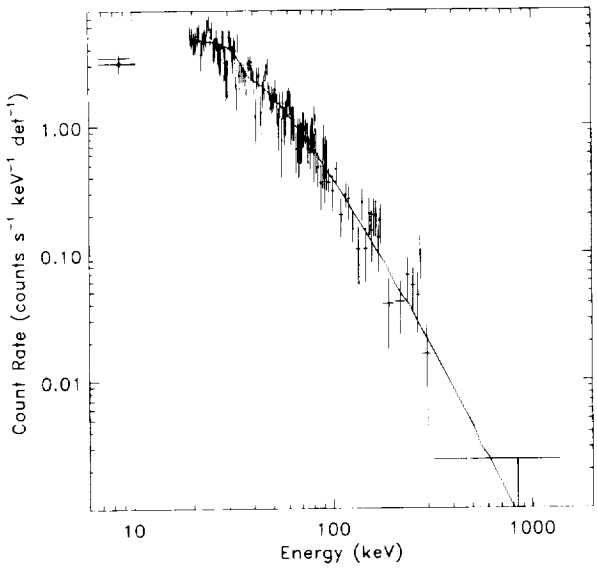

FIG. 8.—Example of burst with $E_p < 45$ keV. GRB model joint fit to the first 96.4 s of 3B 930127, in SD 2, with fitted parameters $E_p = 38 + 3$ keV, $\alpha = -0.5 \pm 0.4$, $\beta = -2.51 \pm 0.08$, and $\chi^2 = 197.3/218$ degrees of freedom. Note that α is not very well constrained by the fit. The DISCSP1 datum is indicated with a dot, while the model rate at that energy is indicated by the horizontal line.

keV–2 MeV) has no significant deviation in the low-energy data, as shown in Figure 5*a*, with a typical spectrum shown in Figure 4. The model used is the "GRB" spectral form, consisting of two smoothly joined power-law segments (Band et al. 1992), where the joining energy has been parameterized in terms of the energy of the peak of the spectrum (E_p) in νF_ν. There is a special category of four bursts in which the fitted value of E_p is so low that there is insufficient SHERB data below E_p to constrain the low-energy power-law fit, as presented in Figure 8. Agreement between the model and the observed count rate is guaranteed in this case, so this class is a subset of the bursts showing no significant deviations. All 12 of the cases in which the model deviates from the low-energy data by greater than 5 σ are due to low-energy excesses; one example is shown in Figure 7. These are evidence that an additional spectral component or an upturn in the spectrum is sometimes required to fully characterize GRBs, but there is not enough information from the single low-energy datum to constrain the shape of any such component. The observed deviations indicate that the GRB spectral form is not sufficient to fit every burst-averaged spectrum below 20 keV. Finally, no burst showing a significant deficit relative to the model is observed. The GRB spectral form incorporates concave-downward spectral curvature for energies lower than E_p, and thus it is able to account for even the weakest low-energy emission we observed. The low-energy power-law index never becomes steeper than -1.7, except for the 12 cases with excesses, for values of E_p ranging from 30 keV to 3 MeV. An excess thermal blackbody component (which cannot be proven in our data) is hard to reconcile with anything other than a Galactic source; however, a low-energy excess in GRB spectra is predicted in at least one theoretical model (Brainerd 1994) that assumes the sources to be at cosmological distances.

We are grateful for discussions with Richard Schwartz, which resulted in ideas on how to refine the calibration of the DISCSP1 energy edges. The anonymous referee's comments helped us clarify the exposition. Also appreciated are the efforts of Robert Mallozzi (UAH) to improve the user interface of our spectral analysis software (WINGSPAN). The software is publicly available from the *Compton Gamma Ray Observatory* Science Support Center at Goddard Space Flight Center, which also archives data produced by BATSE. BATSE work at UCSD is supported under NASA contract NAS 8-36081.

REFERENCES

Band, D. L., et al. 1992, Exp. Astron., 2, 307
———. 1993, ApJ, 413, 281
Blaes, O., et al. 1989, ApJ, 343, 839
Brainerd, J. J. 1994, ApJ, 428, 21
Briggs, M. S. 1995, Ap&SS, 231(1), 3
Castro-Tirado, A. 1994, Ph.D. thesis, Univ. Copenhagen
Fishman, G. J., & Meegan, C. A. 1995, ARA&A, 33, 415
Fishman, G. J., et al. 1989, in Proc. GRO Science Workshop, ed. W. N. Johnson (Washington: NRL), 2-39
Ford, L. A., et al. 1995, ApJ, 439, 307
Hanlon, L. O., et al. 1994, A&A, 285, 161
Harding, A. K. 1991, Phys. Rep., 206, 327
Hurley, K., et al. 1994, Nature, 372, 652
Katz, J. I. 1994, ApJ, 432, L107
Laros, J. G., Evans, W. D., Fenimore, E. E., Klebesadel, R. W., Shulman, S., & Fritz, G. 1984, ApJ, 286, 681
Mallozzi, R. S., et al. 1995, ApJ, 454, 597
Meegan, C. A., et al. 1992, Nature, 355, 143
Meegan, C. A., et al. 1996, ApJS, 106, 65 (BATSE 3B Catalog)
Meszáros, P., Laguna, P., & Rees, M. J. 1993, ApJ, 415, 181
Murakami, T., et al. 1991, Nature, 350, 592
Pendleton, G., et al. 1994, ApJ, 431, 416
———. 1995, Nucl. Instrum. Methods, 364, 567
Press, W., Teukolsky, S., Vetterling, W., & Flannery, B. 1992, Numerical Recipes in FORTRAN (2d ed.; New York: Cambridge Univ. Press)
Ruderman, M., et al. 1988, ApJ, 335, 306
Rybicki, G. B., & Lightman, A. P. 1979, Radiative Processes in Astrophysics (New York: Wiley)
Schaefer, B. E., et al. 1992, ApJS, 92, 285
Schwartz, R. 1995, private communication
Strohmayer, T., Fenimore, E. E., Murakami, T., & Yoshida, A. 1996, submitted
Tavani, M. 1995, Ap&SS, 231(1), 181
Trombka, J. I., et al. 1974, ApJ, 194, L27
Wheaton, W. A., et al. 1973, ApJ, 185, L57
Yoshida, A., et al. 1989, PASJ, 41, 509