

ABSTRACT

Title of dissertation: PRECURSORS IN GAMMA-RAY BURSTS
OBSERVED BY *FERMI*

Sylvia J. Zhu, Doctor of Philosophy, 2015

Dissertation directed by: Dr. Julie McEnery
Department of Physics

Gamma-ray bursts (GRBs) are some of the most energetic explosions in the universe. They come from the core collapses of massive stars and the mergers of compact objects, and are observed as bright flashes of gamma rays (prompt emission) followed by long-lived, fading emission (afterglow) across the electromagnetic spectrum. The instruments on the *Fermi* Gamma-ray Space Telescope provide excellent observations of GRBs across a large energy range. The Gamma-ray Burst Monitor (GBM, 8 keV to 40 MeV) is currently the most prolific detector of GRBs, and the Large Area Telescope (LAT, ~ 20 MeV to >300 GeV) has opened up the field of GRB observations to high-energy gamma rays.

In this thesis, I present studies on improving the LAT's capability to detect GRBs onboard in realtime, and analyses of both a single, extraordinary burst (the record-breaking GRB 130427A) and the population of GBM GRBs with precursors in their lightcurves. In a small fraction of GRBs, a dim peak appears before the much brighter peaks that are normally observed during the prompt emission. I explore whether the properties of GRBs with precursors suggests that precursors

have a distinct physical origin from the rest of the prompt emission, and discuss the implications for models of GRB precursor emission.

Precursors in gamma-ray bursts observed by *Fermi*

by

Sylvia Jiechen Zhu

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2015

Advisory Committee:

Professor Peter Shawhan, chair

Professor Julie McEnery, co-chair

Professor Jordan Goodman

Professor Kara Hoffman

Professor Cole Miller, Dean's Representative

© Copyright by
Sylvia Jiechen Zhu
2015

Acknowledgments

First, of course, I'd like to thank my advisor, Julie McEnery. You have been my mentor and friend through all these years. I have learned so many things from you — not only about science, but also on how to work in a group and bring people together — as well as reinforced some procrastination-related bad habits that we seem to share. You have taught me how to be a good scientist by example. I'm still not entirely sure you ever sleep.

I'd also like to thank Judy Racusin, who has been my mentor when Julie is busy keeping spacecrafts running. You have been my partner in space-themed cakes, and my go-to source for GRB knowledge. Your door was always open to me, and I really appreciate that.

Countless other people in the *Fermi* collaboration and especially the Goddard and NRL groups have helped me and inspired me during these years. I would particularly like to thank the grad students who have been with me for most of my own graduate career: David Green, Jamie Cohen, and Megan DeCesar. I don't know what I'll do when I have to go to conferences without you all.

Normally I'd thank the other physics graduate students but they've all left by now so they don't get to be acknowledged. Instead, I'd like to thank the astronomy graduate students, with whom I found a second home.

A few people who have brought me food during the dissertation writing process deserve to be especially acknowledged: ★Donna Perkins★ (are star symbols allowed in a dissertation?) for supplying me with an endless variety of snacks and

Jeremy Perkins for being the snack fairy, Nino Cucchiara, Rebekah Esmaili, Judy inadvertently, and Pepa Becerra.

And finally, I'd like to thank my parents and my sister for being there for me. We'll now have two doctors in the family, although only one of us will be qualified to give medical advice.

Table of Contents

List of Figures	vi
List of Abbreviations	xiv
1 Introduction	1
1.1 Gamma-ray burst observations	1
1.2 Thesis organization	2
2 A brief history of gamma-ray bursts	4
2.1 Discovery and early observations	4
2.2 The <i>CGRO</i> era of realtime detection	7
2.3 The era of afterglow discovery	15
3 The <i>Fermi</i> Gamma-ray Space Telescope	24
3.1 Large Area Telescope (LAT)	24
3.1.1 Design	27
3.1.2 LAT onboard trigger and event reconstruction	30
3.1.3 Ground event reconstruction	33
3.2 GBM	38
3.3 Observational strategies	41
3.4 GRB handling: Communication between the GBM and LAT	42
3.5 Ground GRB analysis with <i>Fermi</i>	45
4 The LAT onboard algorithm	48
4.1 Overview	49
4.2 LAT-only trigger	50
4.2.1 Tier 1	50
4.2.2 Tier 2	52
4.2.3 GBM-seeded window	53
4.2.4 Localization and Updates	54
4.3 Onboard algorithm studies	56
4.3.1 Improving sensitivity	56
4.3.2 Localization studies	60

5	GRB theory	65
5.1	The physical system	66
5.1.1	Progenitors	66
5.1.2	Central engine	67
5.2	Prompt emission	68
5.3	Afterglow emission	71
6	The Record-breaking GRB 130427A	74
6.1	An overview of high-energy emission from GRBs	76
6.2	LAT observations of GRB 130427A	80
6.2.1	Analysis	80
6.2.2	Interpretation	88
6.3	VERITAS observations of GRB 130427A	92
6.4	Multiwavelength observations	96
7	GBM observations of precursor emission in GRBs	101
7.1	Introduction	104
7.2	Brief overview of precursor models	109
7.3	Bayesian blocks	110
7.4	Systematic search of GBM precursors	113
7.5	Results	118
7.5.1	Type I	122
7.5.2	Type II	123
7.5.3	Type III	126
7.5.4	Comparison	127
7.6	Interpretation	138
7.6.1	Photospheric precursors	138
7.6.2	Shock breakout precursors	139
7.6.3	Fallback precursors	141
7.6.4	Multistage collapse precursors	141
7.6.5	NS-NS magnetic field interactions	142
7.6.6	NS crust cracking	142
7.7	Discussion	143
8	Conclusions and final thoughts	145
	Bibliography	147

List of Figures

2.1	The lightcurves of the one of the first GRBs ever detected, GRB 700822. The arrows point out features in the lightcurves that appeared in multiple detectors, and show the variability of the emission [1].	5
2.2	The lightcurve of GRB 830801B, a smoothly decaying peak preceded by a dim precursor pulse. [2]	6
2.3	The Compton Gamma Ray Observatory <i>CGRO</i> had four instruments, and covered a total energy range of 20 keV to 30 GeV. <i>CGRO</i> is the spiritual ancestor of <i>Fermi</i> , whose instrument designs were informed by BATSE and EGRET.	7
2.4	GRB lightcurves cover a wide range of properties. <i>CGRO</i> -BATSE observed over two thousand GRBs by the end of the mission; twelve lightcurves are shown here to illustrate some of the differences in duration, variability, etc. [Credit: J.T. Bonnell (NASA/GSFC)].	9
2.5	The Band function [17] was and is still often used as a non-thermal model of GRB spectra. In a few cases, such as GRB 990123, data from multiple detectors were combined to best constrain the parameters of the Band function [19].	10
2.6	In the <i>CGRO</i> era, there were already indications that GRBs could emit photons at both very high energies and very late times. For instance, EGRET observed an 18 GeV photon about an hour and a half after trigger for GRB 940217. (Note: the gap in the EGRET data, indicated by the dashed green line, is due to a combination of Earth occultation and telemetry limitations rather than the former alone, as <i>CGRO</i> was not in view of the satellites that provided contact with the ground during this time.) [22]	11
2.7	GRB fall into two rough groups based on their durations: long and short. About a quarter of observed GRBs have durations (T90s, described in the text) of less than approximately 2 seconds, while the remainder have durations of more than 2 seconds. The exact shape of the T90 distribution differs between different instruments, depending on factors such as the instrument's energy range [23].	12

2.8	GRBs are distributed isotropically across the sky. The figure on the left shows the distribution (in galactic coordinates) of the bursts in the first BATSE GRB catalog, published in 1994 [13], while the figure on the right shows the distribution of all the bursts detected by BATSE during its entire mission (available at http://gammaray.nsstc.nasa.gov/batse/gr)	
2.9	GRB 970228 was the burst with the first observed X-ray afterglow. <i>BeppoSAX</i> detected and localized the GRB, then slewed to observe it with its more sensitive NFIs starting eight hours after the burst began. The top plot shows the X-ray image of a small area around the GRB’s localization; there is a clear X-ray source that fades over the next few days. The bottom plot shows that the X-ray flux decreases smoothly over a few weeks’ time. [43]	17
2.10	GRB 970228 was also the burst with the first observed optical counterpart, which also faded in time. In addition, the optical lightcurve showed an excess or bump at late times, which was later interpreted as being evidence for an underlying supernova [44].	18
2.11	GRB 970508 was the GRB with the first redshift measurement and the first radio afterglow, and its optical and radio afterglows were bright enough that a complete spectrum across a large energy range could be constructed. The spectrum was well described by a power law with two breaks, which is consistent with the predictions of the fireball model [49].	19
2.12	Supernova 1998bw was found in the same region of the sky as GRB 980425, less than half an arcsecond apart. This marked the first clear association between GRBs and SNe. The archival image of the host galaxy is shown on the right, and the new observation on the left; the supernova is the bright new source marked by the red arrow [48].	19
2.13	The X-ray afterglows detected by <i>Swift</i> have a few common features. The lightcurves often follow a shallow-steep shallow-pattern, although there are sometimes deviations from this pattern. Afterglows can also have X-ray flares on top of the standard set of broken power laws, and for some — like GRB 050502B here — the X-ray flare releases a very large fraction of the total energy. Here, the BAT data points (extrapolated into the XRT energy range) are plotted with crosses and the XRT data points are plotted with circles [61].	21

2.14	A canonical X-ray afterglow lightcurve emerged thanks to <i>Swift</i>'s trove of observations. The lightcurves often feature rebrightenings or flares and multiple breaks, including jet breaks, from which a jet's opening angle can be calculated. The prompt emission (phase 0) is followed by a steep decay (phase I) as whatever process produced the prompt emission shuts off. The steep decay is sometimes followed by a shallow plateau (phase II) and x-ray flares (phase V) — suggesting renewed activity — or followed directly by a more shallow decay (stage III). The timing of the break in the lightcurve at the end of the shallow decay can be used to calculate the size of the jet which produced the GRB emission [59].	22
3.1	At energies above tens of MeVs, pair conversion has the largest interaction cross section. The different lines represent different target materials: silicon, germanium, iodine, bismuth, in order of both lowest to highest curves and smallest to largest atomic number. Tungsten is used in the LAT and has an atomic number (74) between iodine (53) and bismuth (83), so its cross-section curves fall between the top two. Original figure from [29].	25
3.2	A gamma ray converts into an electron-positron pair in the Tracker; the pair then deposits their energy into the Calorimeter. By reconstructing the energies and tracks of the charged particles, the energy and track of the original gamma ray can be determined.	27
3.3	A simplified diagram of a photon interacting with the Tracker. In this example, the photon pair-converts in the top tungsten layer. As the child particles propagate through the layers, they create “holes” in the SSDs; the x - y locations of these holes in the different layers are combined to determine the particle tracks.	29
3.4	Each Calorimeter module is made of 96 Cesium Iodide crystals in 8 layers. The layers contain 12 crystals each and are in a hodoscopic configuration. Every crystal has a set of photodiodes on each end to measure the energy deposited in the crystal. . .	29
3.5	The ACD is made of 89 plastic scintillator tiles that surround the Tracker and Calorimeter. The tiles overlap in one direction to minimize the amount of space left uncovered, and scintillating fiber ribbons cover the space in the other direction. Wavelength shifting fibers embedded in the tiles allow for uniform light collection; each of these is attached to two PMTs, for redundancy.	31
3.6	In the LAT-based coordinate system, an event's direction is described by the angle from the boresight (the z-direction), θ, and the azimuthal angle, ϕ. The x-axis is parallel to the solar panel arms of the spacecraft.	32

3.7	The acceptance of the LAT (the effective area integrated over the field of view) decreases at each stage of the event filtering process. Most of the events that are rejected at each stage are charged particles.	36
3.8	The event rates onboard is one or two orders of magnitude larger than the event rates for the standard event classes. . .	37
3.9	The twelve NaI detectors (labeled 0 through 11) and two BGO detectors (12 and 13) are positioned on opposite sides of the spacecraft. The NaIs are oriented so that the GBM observes the entire unocculted sky at all times.	39
4.1	The rate of events that pass the onboard filter is ~ 400 Hz, much larger than the loosest standard ground event class (transient). The LLE and Source class event rates are also plotted for comparison.	50
4.2	The offset between the onboard localization and the actual GRB position for different values of Tier 1 threshold, for the first six LAT-detected GRBs that could have been localized onboard. The different circles represent different Tier 1 thresholds, from 80 (the largest circle) to 100 (the smallest circle).	57
4.3	The onboard localization error decreases either if the number of photons increases or the clusteredness increases. This can be seen directly from the way the error is calculated (Eqn 4.10). .	61
4.4	The distribution of errors for the localizations that were calculated using 5 events, in the 'suspect' messages (i.e., the first message that is sent out in the set of 8). There is a peak at around 60 arcmin due to the sets of events in which one event was within a degree of the input position, and a similar one at around 44 arcmin ($60/\sqrt{2}$) due to the sets in which two events were within a degree of the input position. We also fit a Gaussian to the tail of the distribution to calculate the probability for very small errors.	62
5.1	Within the fireball model of GRBs, the prompt emission is caused by internal shocks within the jet while the afterglow emission is caused by the jet's later interactions with circumburst material. In the Poynting flux-dominated model, the overall GRB structure is the same, but the prompt emission instead comes from the reconnection of magnetic field lines within the jet. (Credit: NASA's Goddard Space Flight Center)	68
5.2	The canonical x-ray afterglow lightcurve is made of a multiply broken power law. The bolded segments are seen in most x-ray afterglows, while the dotted segments (including the flares marked as section V) sometimes occur as well. [76]	72

6.1	GRB 130427A was extraordinarily bright to observers (left two panels) but intrinsically ordinary (right two panels). Figure from [80].	75
6.2	LAT-detected prompt emission is usually delayed with respect to the GBM-detected emission, as seen here in GRB 090926A. The LAT emission (bottom three panels) does not begin until a few seconds after the GBM trigger (the vertical red line) [83].	77
6.3	A time-resolved spectral analysis of a bright LAT GRB often requires an extra power law component to model the high-energy emission, as seen here in GRB 090926A. The four time intervals here correspond to the marked time intervals in Figure 6.2 [83].	78
6.4	LAT-detected GRBs tend to be very energetic and/or very close. [77]	79
6.5	LAT-detected GRBs often come from jets with bulk Lorentz factors of a few hundred. [77]	79
6.6	The prompt emission of GRB 130427A was recording breakingly bright in the GBM energy range (top three panels). The LLE >10 MeV emission (fourth panel) shows similarities with the GBM emission. The LAT >100 MeV emission, however, brightens <i>after</i> the lower energy emission had started to fade.	81
6.7	During the brightest part of GRB 130427A, the GBM TTE data were saturated due to the high counts flux. The buffer filled up faster than it could be read out. The binned data types (CTIME and CSPEC) were unaffected by saturation, although all data types suffered from pulse pileup.	82
6.8	During the first two time bins (up to $T_0 + 11.5$ s), all of the GBM + LLE + LAT emission can be modeled by a single component, while a separate power law component is needed to fit the LAT emission after $T_0 + 11.5$ s. We did not use GBM data during the second time interval due to pulse pileup.	85
6.9	The LAT detected emission from GRB 130427A for almost 20 hours, an order of magnitude longer than the previous record of 1.5 hours. In the bottom panel, the larger circles mark the highest energy photons in each time interval.	87
6.10	The high-energy late-time photons detected by the LAT are difficult to explain using a standard synchrotron forward shock model. The red curves indicate a circumburst medium with a uniform density, while the blue curves indicate a wind-like ($\propto R^{-2}$) density. The two dotted lines mark the scenario in which the acceleration is extremely fast. The solid and dashed lines correspond to scenarios with $\Gamma_0 = 1000$ for an adiabatic and radiative blastwave, respectively. The blue dot-dashed and double dot-dashed lines represent an adiabatic blastwave with $\Gamma_0 = 500$ and $\Gamma_0 = 2000$	90

6.11	The VERITAS non-detection suggests a cutoff at around $\gtrsim 100$ GeV if the LAT observations are due to SSC emission. Alternatively, if the LAT observations are due to synchrotron emission, the VERITAS upper limits can put constraints on the magnetic fields around the blastwave front. The LAT $1\text{-}\sigma$ error contour is extrapolated in time from the power law fit to the LAT data at late times. The VERITAS upper limits were calculated assuming a cutoff in the Klein-Nishina regime at different energies.	94
6.12	The multiwavelength afterglow as measured of GRB 130427A, as measured by <i>Swift</i> and other instruments, was mostly well-modeled by a single synchrotron component, although this model underpredicted the GeV emission. See [85] for details and time interval definitions.	97
6.13	The bright optical flash is more consistent with reverse shock emission than forward shock emission. The late-time optical afterglow detected by RAPTOR decays similarly to the late-time LAT-detected afterglow, suggesting a common origin to both [75].	98
6.14	The Kouveliotou et al. study found that the multiwavelength observations of GRB 130427A at around 1.5 and 5 days after trigger are well-modeled by a single spectral component. See [99] for details.	99
6.15	When GRB 130427A occurred, HAWC was not operating at full sensitivity (solid black lines). If HAWC had had its full capabilities (green lines), it might have observed the first VHE photons ever detected from a GRB [100].	100
7.1	GRB 720427 was the first published GRB with a precursor. [104]	102
7.2	The properties that often are used to define precursors (dim and followed by a well-defined quiescent period) depend on the signal-to-noise ratio of the GRB. <i>Left:</i> The presence or absence of a quiescent period is background-dependent, and studies that require well-defined quiescent periods exclude some GRBs with precursor-like first peaks. The example shown here is GRB 130427.324, a particularly bright GRB. <i>Right:</i> Some GRBs (such as GRB 090529.564 here) have initial peaks that technically satisfy the definition of a precursor but are not much dimmer than the dominant emission. With a lower signal to noise, the flux ratio between the precursor and dominant emission becomes smaller, the quiescent period becomes more well-defined, and the first peak becomes more “precursor-like”.	107

7.3	The precursors were divided into three subtypes based on the relative fluxes and the existence of the quiescent periods. These plots show examples of the three types; the precursor is marked in red and the dominant emission in blue. The examples here are GRBs 100116.897, 081221.681, and 120530.121.	116
7.4	The low-energy index distributions for the Comptonized function and Band function (and smoothly broken power law) are generally consistent with each other, while the power law index distribution is more negative and represents softer spectra. In general, the Comptonized and Band α measure the same thing (the true low-energy index) whereas the power law γ is affected by the spectral turnover around E_{peak} while not having enough statistics to constrain E_{peak} . This plot is from the second GBM spectral catalog [114].	119
7.5	Precursors were found in GRBs across the whole range of T90 and peak flux. The values for GRBs before GRB 12071146 were obtained from the second set of GBM catalogs [114] [115] and are plotted as filled circles, while the values for all other GRBs were taken from the online GBM burst catalog and are plotted as open circles. The approximate division between short and long GRBs (roughly 2 seconds) is marked by the vertical dashed line.	120
7.6	The dominant emission and quiescent period durations for Type I GRBs span a similar range, while precursor durations are an order of magnitude smaller, on average.	123
7.7	We found mild correlations between T_P, T_Q, and T_D, but these correlations vanished when I normalized the durations by T_{tot}. The examples plotted here are the T_D vs T_P (top) and T_D/T_{tot} vs T_P/T_{tot} (bottom) for bursts with Type I precursors.	125
7.8	The Type II precursors have a wide range of “precursor-ness”. Some precursors were very clear peaks and almost resembled Type Is (left), while others were more “shoulder-like” and consistent with being the start of a slow-rising peak.	126
7.9	The E_{peak} and index values for the precursor and first dominant emission peak are plotted for the three precursor subtypes. In these plots, filled circles correspond to short GRBs and open circles to long GRBs. There are a few Type I precursors with particularly hard indices, although these have large error bars. For all three subtypes, the E_{peak} -index values for the precursor and first dominant emission peak span similar ranges, with a few exceptions.	128

7.10	The α (or γ) and E_{peak} distributions are plotted here for all the GRBs with precursors in our sample. In both plots, the darker histograms represent the emission episodes that were well fit by a Comptonized or Band function; in the plot on the left, the lighter histogram represents the emission episodes that were best fit by a power law. The hatched histograms in both plots correspond to the short GRBs.	130
7.11	For both α (left) and E_{peak} (right), we divided the value for the precursor by the value for the corresponding dominant emission. The short GRBs are marked by the hatched histograms.	133
7.12	For the most part, the D1 energy flux is larger than the P energy flux, for all three types. However, the relative difference between D1 and P energy flux is smaller on average for the Type III precursors than Type I or II.	137
7.13	Photospheric precursors would have particularly hard values of α, and their corresponding quiescent periods would be very short. The approximate parameter space they would cover is marked by the shaded region (assuming internal shocks). In this plot, spectra that were best fit by Comptonized or Band function are marked with filled circles, and spectra that were best fit by a power law with open circles.	139

List of Abbreviations

Γ	bulk Lorentz factor
ACD	Anti-Coincidence Detector
ARR	Autonomous Repoint Request
BA	Burst Advocate
BACODINE	BATSE Coordinates Distribution Network
BAT	Burst Alert Telescope
BATSE	Burst And Transient Search Experiment
BGO	Bismuth Germanate
BH	black hole
CAL	Calorimeter
<i>CGRO</i>	Compton Gamma-Ray Observatory
COMPTEL	Imaging Compton Telescope
CPF	charged particles in the field of view
C-STAT	Castor statistic
CsI	Cesium Iodide
DEC	declination
DRM	Detector Response Matrix
E_{iso}	isotropic energy release
EBL	Extragalactic Background Light
EGRET	Energetic Gamma Ray Experiment Telescope
FOV	field of view
FRED	fast rise, exponential decay
GBM	Gamma-ray Burst Monitor
GCN	Gamma-ray Circulars Network
GRB	Gamma-ray burst
GRBM	Gamma-Ray Burst Monitor
HAWC	High Altitude Water Cherenkov Gamma-ray Observatory
<i>HETE-2</i>	High Energy Transient Explorer
HR	hardness ratio
IR	infrared
ITS	Immediate Trigger Signal
IRF	Instrument Response Function
LAD	Large Area Detector
LAT	Large Area Telescope
LLE	LAT Low Energy
MET	Mission Elapsed Time
MIP	Minimum Ionizing Particle
NaI	Sodium Iodide
NFI	Narrow Field Instrument
NS	neutron star

NuSTAR	Nuclear Spectroscopic Telescope Array
OSSE	Oriented Scintillation Spectrometer Experiment
PHA	pulse height analysis
PMT	photomultiplier tube
POWL	power law
PSF	Point Spread Function
RA	right ascension
RAPTOR	Rapid Telescopes for Optical Response
RSP	Response
SAA	South Atlantic Anomaly
SBPL	smoothly broken power law
SD	Spectroscopy Detector
SGR	Soft Gamma Repeater
SN	supernova
SSC	synchrotron self-Compton
SSD	silicon strip detector
TASC	Total Absorption Shower Counter
TDRSS	Tracking and Data Relay Satellites System
TKR	Tracker
ToO	Target of Opportunity
TS	Test Statistic
TTE	Time-Tagged Events
UHE	Ultra High Energy (>100 TeV)
UVOT	Ultra-Violet / Optical Telescope
VERITAS	Very Energetic Radiation Imaging Telescope Array System
VHE	Very High Energy (>100 GeV)
WFC	Wide Field Camera
XRT	X-Ray Telescope

Chapter 1: Introduction

1.1 Gamma-ray burst observations

Gamma-ray bursts (GRBs) were discovered by accident in the late '60s. Since then, they have been observed over a wide range of energies with a large variety of instruments, yet many aspects of these violent explosions remain mysterious. They are observed as bright and highly variable flashes of gamma rays (prompt emission) followed by long-lived, fading emission across the electromagnetic spectrum. The Compton Gamma-ray observatory (*CGRO*) uncovered the startling variety of prompt emission lightcurves and space-based missions such as *BeppoSAX*, *HETE-2*, and *Swift* fully explored the X-ray afterglow emission, all accompanied by ground-based observatories at longer wavelengths.

Their observations all pointed to GRBs being produced by violent events: the core collapses of massive stars and the binary mergers of compact objects. Both scenarios lead to the formation of a compact central engine which launches a jet. The short-lived, highly variable prompt emission comes from interactions within the jet, while the longer and more smoothly fading afterglow comes from the jet's interactions with the environment. There are still many unanswered questions about GRBs, such as the structure and composition of the jet, the nature of the central engine, and the properties of the progenitor systems themselves.

By the time the *Fermi* Gamma-Ray Space Telescope was launched in 2008, earlier missions had observed GRBs well at lower energies, but had only scratched

the surface of the high-energy (>100 MeV) emission. *Fermi* has two instruments and covers a large energy range. The Large Area Telescope (LAT, ~ 20 MeV to >300 GeV) detects and tracks high-energy gamma rays from many sources, including GRBs. The Gamma-ray Burst Monitor (GBM, 8 keV to 40 MeV) detects photons from flaring sources in general and GRBs in particular.

The LAT and GBM work in tandem to observe GRBs. Both can independently detect and localize GRBs onboard and in realtime, and the LAT can utilize information from the GBM to seed a search for GRB emission. Data from both detectors is often used to construct and characterize the spectrum of a GRB, and different features are present in the different energy ranges. For instance, the GBM lightcurves sometimes have precursors (dim peaks that occur before the brightest part of the prompt emission) whose origins are still debated, and the LAT sometimes detects emission for hours after the GBM emission has faded, as in the case of GRB 130427A.

1.2 Thesis organization

In this work, I present studies of GRB observations by the *Fermi*-LAT and GBM. I begin with a historical overview of GRBs in Chapter 2, as told by the instruments that observed them: The initial discovery by the *Vela* satellites, the numerous observations of the prompt emission by *CGRO*, and the thorough observations of the afterglow emission by *BeppoSAX*, *HETE-2*, and *Swift*. This sets the scene for the field of GRB research at the time of *Fermi*'s launch.

In Chapter 3, I discuss the instrumentation and performance of the LAT and GBM. I detail the event reconstruction in the LAT and GBM, and the procedures for GRB analysis for both instruments. Chapter 4 focuses on the LAT onboard GRB detection algorithm, and describes studies we performed to improve the onboard sensitivity and localization capabilities.

I then focus on the science of GRBs. I discuss an overview of GRB theory in Chapter 5, focusing on the models that are most commonly used. In Chapter 6, I present observations of the record-breaking GRB 130427A by the LAT, and analysis of the LAT data both by itself and in conjunction with observations at energies >100 GeV by the Very Energetic Radiation Imaging Telescope Array System (VERITAS). In Chapter 7, I discuss a population study of GRBs with precursors observed by the *Fermi*-GBM, and the implications for GRB models. Finally, I summarize these results in Chapter 8.

Chapter 2: A brief history of gamma-ray bursts

Gamma-ray bursts (GRBs) were first detected in the late '60s. Since then, our understanding of them has grown in a sort of — to borrow a phrase from evolutionary biology — punctuated equilibrium, as each generation of telescopes with improved capabilities uncovers new and often surprising characteristics.

2.1 Discovery and early observations

GRBs were discovered by accident during the Cold War. The *Vela* satellites (named for the Spanish word “velar,” meaning “to watch”) were a series of US military satellites that were launched starting in 1963 to monitor the atmosphere and outer space for nuclear activity. The satellites had scintillator detectors that were sensitive to photons with energies between 0.2 to 1.5 MeV. Multiple satellites could roughly triangulate a source using the slightly different arrival times.

In 1969, the four operating *Vela* satellites observed flashes of high-energy photons coming from space. The flashes lasted for seconds to tens of seconds, and each was detected by multiple satellites almost simultaneously (see, e.g., Figure 2.1). They did not come from the Sun, and they were not associated with any known supernovae [1].

Following the official publication of the discovery in 1973, many other high-energy missions — both satellite and balloon-borne — also searched for and found gamma-ray bursts. They confirmed that GRBs were transient phenomena with a diverse set of lightcurves; some were rapidly varying (Figure 2.1) while others were




Figure 2.1: **The lightcurves of the one of the first GRBs ever detected, GRB 700822.** The arrows point out features in the lightcurves that appeared in multiple detectors, and show the variability of the emission [1].

smooth, and sometimes preceded by a dim precursor pulse (Figure 2.2) [2].

The energy spectra of GRBs were observed to be nonthermal; i.e., the spectra were consistent with radiative processes from energetic charged particles (e.g., [3] [4]). A few satellites also claimed evidence of absorption and emission lines in GRB spectra on top of the nonthermal continuum radiation [5] [6]; however, sometimes different instruments disagreed on the existence of the lines in the same GRB (e.g., [7]).

When examined as a population, GRBs were consistent with being distributed isotropically across the sky [8]. It was unclear whether they were also homogeneously




Figure 2.2: The lightcurve of GRB 830801B, a smoothly decaying peak preceded by a dim precursor pulse. [2]

distributed in their distance from the Earth, as instrumental effects obscured the true distribution, although there was some indication that they were indeed homogeneously distributed out to the detectors' sensitivity limits. Since there was no evidence of large-scale structure, if GRBs came from local objects, then they had to be either extremely close by or located in an extended halo so that the finite height of the galaxy did not affect their observed distribution [9].

The typical observed energy release was measured to be around 10^{-5} erg cm^{-2} , which corresponds to a luminosity of 10^{42} erg if the GRB progenitor is in our galaxy [10] but a problematically large 10^{53} erg s^{-1} for an object at a redshift of $z \sim 1$ [11] (equivalent to the entire rest energy of the Sun being emitted within seconds). At the same time, GRB lightcurves were often highly variable with fluctuations on sub-second timescales [12], which suggested that the emission source was a small, compact object. Neutron stars would successfully account for both the large energies involved and the rapid variations; if the neutron stars were strongly magnetized, they would also explain the spectral lines that were interpreted as cyclotron lines [8]. Many scenarios were suggested — starquakes, accretion, comet/asteroid impacts — but nothing was settled.

By the end of the second decade, the local neutron star model seemed to




Figure 2.3: The Compton Gamma Ray Observatory *CGRO* had four instruments, and covered a total energy range of 20 keV to 30 GeV. *CGRO* is the spiritual ancestor of *Fermi*, whose instrument designs were informed by BATSE and EGRET.

plausibly explain all of the observed properties of GRBs. However, the launch of the *Compton Gamma-ray Observatory* (*CGRO*) with its multiple gamma-ray detectors changed the GRB paradigm.

2.2 The *CGRO* era of realtime detection

The Compton Gamma-Ray Observatory *CGRO* was launched in 1991 with four instruments: the Burst And Transient Source Experiment (BATSE; 25 keV to 2 MeV), the Oriented Scintillation Spectrometer Experiment (OSSE; 50 keV to 10

MeV), the Imaging Compton Telescope (COMPTEL, 80 keV to 30 MeV), and the Energetic Gamma Ray Experiment Telescope (EGRET, 20 MeV to 30 GeV). (See Figure 2.3.)

BATSE was made up of eight detectors, each placed at a corner of the *CGRO* spacecraft so that BATSE had a near all-sky field of view. Each detector was made of two NaI scintillation detectors with complementary purposes: a Large Area Detector (LAD) for detection and localization and a Spectroscopy Detector (SD) for higher resolution energy measurements [13]. BATSE detected GRBs at an unprecedented rate of ~ 1 per day, and by the end of the mission in 2001, it had detected 2145 GRBs. A sample of BATSE lightcurves is shown in Figure 2.4, and illustrates the wide variety in lightcurve shapes.

EGRET was built to detect and localize individual gamma rays between 30 MeV and 30 GeV. It was made of spark chamber modules interleaved with tantalum foil. Photons interacted with the foil and pair produced, and the resultant charged particles were tracked by the spark chambers. Two sets of plastic scintillator tiles sat in the middle and bottom of the spark chamber systems to both trigger the spark chamber when the tiles detected charged particles and determine the particles' general detection (upward or downward). The energies were collected and measured by a Total Absorption Shower Counter (TASC), which could also function as an independent photon and charged particle detector. This was all surrounded by an Anti-Coincidence Counter, a scintillation detector which would detect the charged particle background and reject signals that were temporally coincident with these particles [14].

Thanks to BATSE's observations, GRBs could now be analyzed more thoroughly and systematically. GRB lightcurves came in a wide range and could be anywhere from smooth and single-peaked or highly variable (Figure 2.4). The individual peaks themselves were often found to be asymmetric, with a fast rise and a

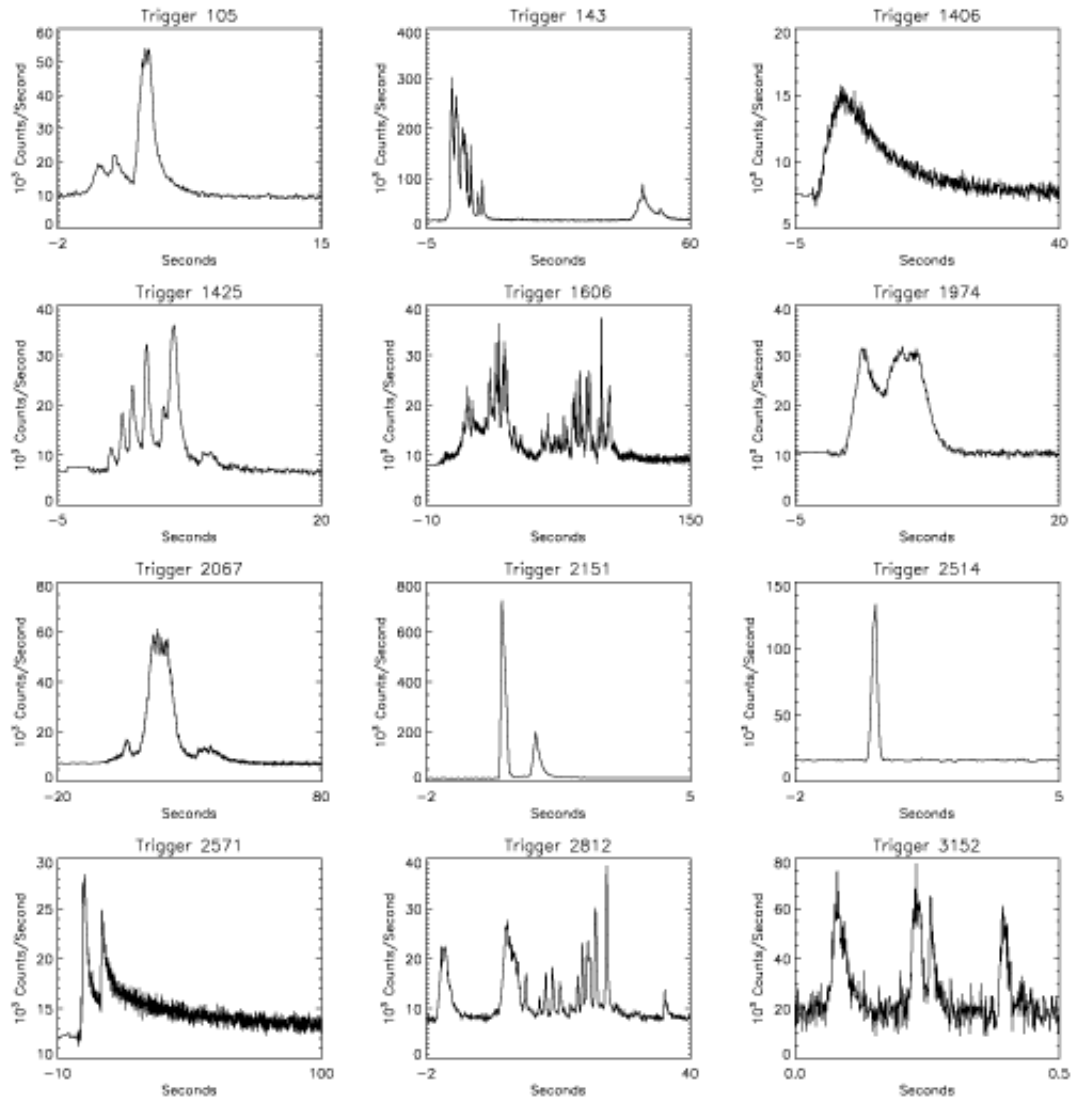


Figure 2.4: **GRB lightcurves cover a wide range of properties.** *CGRO-BATSE* observed over two thousand GRBs by the end of the mission; twelve lightcurves are shown here to illustrate some of the differences in duration, variability, etc. [Credit: J.T. Bonnell (NASA/GSFC)].




Figure 2.5: **The Band function [17] was and is still often used as a non-thermal model of GRB spectra.** In a few cases, such as GRB 990123, data from multiple detectors were combined to best constrain the parameters of the Band function [19].

slower, exponential decay (a.k.a. “FRED-like”), and they tended to be shorter at higher energies.

The SDs had the sensitivity to detect the spectral lines that had previously been claimed by some missions, but did not find any; previous reports of these lines were instead considered to have been caused by instrumental effects [15]. BATSE confirmed that GRB spectra are nonthermal, and the phenomenological Band function (a smoothly broken power law with a particular curvature) was put forth as a good general description of GRB spectra [17]. In a few cases, multiple *CGRO* detectors simultaneously observed a GRB, and all the data were combined to form a full picture of the prompt emission spectrum (see, e.g., Figure 2.5).

At the same time, deviations from the Band function were known from the




Figure 2.6: **In the CGRO era, there were already indications that GRBs could emit photons at both very high energies and very late times.** For instance, EGRET observed an 18 GeV photon about an hour and a half after trigger for GRB 940217. (Note: the gap in the EGRET data, indicated by the dashed green line, is due to a combination of Earth occultation and telemetry limitations rather than the former alone, as CGRO was not in view of the satellites that provided contact with the ground during this time.) [22]

beginning. In particular, EGRET observed individual high-energy gamma-rays (> 30 MeV) from four GRBs (GRBs 910503, 930131, 940217, and 940301 [18]), necessitating an additional spectral component in some time intervals to model the high-energy emission. In the case of GRB 940217, EGRET detected an 18 GeV photon an hour and half after trigger [22], so that whatever mechanism caused the high-energy photons could be very long-lived. In addition, the TASC on EGRET, being a scintillation detector, could detect lower energy events (1-200 MeV) independent of any spark chamber activity or scintillator tile trigger (although it could not veto charged particles and so could only be used for spectral analysis), and detected emission from a total of 18 GRBs (including GRB 940217) [18].

The large size of the BATSE sample led to the discovery that GRBs could be classified into two different subsets, short and long, based on their durations. GRB durations were (and still are) often reported as T90, the time between when 5% and 95% of the total counts in the GRB are observed. As Figure 2.7 shows, about 3/4 of GRBs have T90s of longer than a few seconds and are called long




Figure 2.7: **GRB fall into two rough groups based on their durations: long and short.** About a quarter of observed GRBs have durations (T_{90} s, described in the text) of less than approximately 2 seconds, while the remainder have durations of more than 2 seconds. The exact shape of the T_{90} distribution differs between different instruments, depending on factors such as the instrument’s energy range [23].

GRBs, while the remainder have shorter durations and are known as short GRBs. Various comparisons of their properties confirmed that these were indeed two different populations rather than a single underlying population separated into two by instrumental effects; for instance, short bursts tend to have harder spectra than do long bursts [23].

Thanks to the completeness of the BATSE sample and the detection of dim bursts in particular, it became more evident that GRBs could not be Galactic objects as there was no evidence for their being distributed in a disk [13]. GRBs had to be either located in a Galactic halo at a few to tens of kiloparsecs (see, e.g., [24]) or at cosmological distances. However, there was no direct evidence in favor of either, and the GRB community was split on this question. A public debate to discuss this took place in 1995 at the Smithsonian Museum of Natural History, with both a proponent for the Galactic halo scenario and a proponent for the cosmological scenario. On the one hand: high-velocity neutron stars were known to exist in the Galaxy, and a

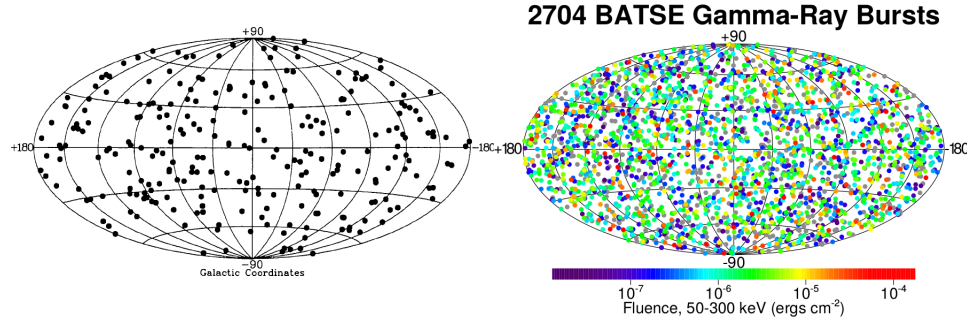


Figure 2.8: **GRBs are distributed isotropically across the sky.** The figure on the left shows the distribution (in galactic coordinates) of the bursts in the first BATSE GRB catalog, published in 1994 [13], while the figure on the right shows the distribution of all the bursts detected by BATSE during its entire mission (available at <http://gammaray.nsstc.nasa.gov/batse/grb/skymap/>).

population of these could form a halo; and, it was already known that neutron stars could emit energetic explosions in the form of Soft Gamma Repeaters (SGRs), so it was conceivable they could also be the sources of GRBs [25]. On the other hand: the proposed size of the GRB halo was far larger than any known Galactic halo; it would be difficult to produce a halo population without some measurable clustering near the Galactic center (which was not seen); and, a cosmological population would trivially explain the spatial distribution [26].

A point that seemed to be in favor of the Galactic halo scenario was the fact that a cosmological origin would require that GRBs have incredibly large energy releases (up to 10^{54} ergs if the energy were emitted isotropically) in very short periods of time. At the same time, the rapid variability seen in some lightcurves (Figure 2.4) required extremely compact source regions. This combination of high luminosities in very small spaces would mean that the photon densities in the source regions would be prohibitively large, and that no photons would be able to escape the region and actually be detected. However, this could be circumvented if the emission regions of GRBs were relativistic, which would simultaneously mean a larger emission region and lower photon energies in the rest frame; thus, the emission region would no longer be optically thick. This concept was first proposed in 1978 [27] and started

to gain traction with the observations of CGRO.

There was indirect support for both the Galactic halo and cosmological scenarios, and neither was entirely implausible, but there was still no direct evidence in favor of either one. Many instruments undertook searches for GRB emission at other wavelengths to try to uncover more information about the systems or objects they came from. Long-lived fading emission had been predicted as a natural consequence of the energy released during the prompt emission, as the outflow that causes the prompt emission would eventually interact with any material surrounding the GRB [58]. If this afterglow emission could be detected, it could say something about the nature or environment of GRBs.

The CGRO tape recorders had failed early in the mission, so all data detected by CGRO was instead downlinked near-continuously via the NASA Tracking and Data Relay Satellites System (TDRSS)¹. This meant that, so long as CGRO was in sight of TDRSS, the BATSE GRB data were available on the ground immediately after a trigger and localizations could be calculated on the ground in near-realtime. These localizations were then rapidly distributed to the community using the BATSE COordinates DIstribution NETwork (BACODINE [28]; later expanded and rebranded as the Gamma-Ray Coordinates Network, GCN), so that the community could observe the BATSE error circles for emission at other wavelengths. Unfortunately, none of these follow-up observations were successful, but the concept of immediate localizations and a rapid communication system would prove essential to the GRB community in the following generations.

Follow-up observations of BATSE GRBs had not been successful, largely due the size of the BATSE error circles (at least a few degrees) compared with the typical field-of-view for a longer wavelength telescope (arcminutes). The question of whether GRBs were Galactic or cosmological objects would not be settled until the

¹<http://tdrs.gsfc.nasa.gov/>

next generation. The launch of observatories that were capable of rapid slewing and greatly improved localizations finally allowed for GRBs’ distances to be determined.

2.3 The era of afterglow discovery

CGRO’s observations of GRBs gave the community a clearer picture of prompt emission but only hints of any long-lived emission. Follow-up observations of BATSE bursts were unsuccessful, primarily due to the size of the error circle. The next generation of GRB observatories were designed to have better localization capabilities, which led to the first detection of long-lived fading afterglow emission from GRBs, the first detection of emission from GRBs at other wavelengths (known as “counterparts”), and the first measurement of a redshift.

The *BeppoSAX* mission lasted from 1996 to 2002, overlapping with *CGRO* for a few years. It had multiple instruments, including a set of scintillators to detect GRBs (Gamma-Ray Burst Monitor, or GRBM) and a set of two coded mask cameras with large fields of view for subsequent localization (Wide Field Cameras, or WFCs), as well as four instruments at lower energies for longer observations (collectively called the Narrow Field Instruments, or NFIs). And, perhaps most importantly, *BeppoSAX* could be maneuvered to observe a particular part of the sky — say, a part of the sky in which a GRB was recently detected — with the sensitive NFIs for a long period of time [40]. These all led to the first detection of afterglow emission from a GRB, with GRB 970228.

The prompt emission of GRB 970228 was detected by the GRBM. It was also within the WFCs’ fields of view, and was localized to within an arcminute precision. *BeppoSAX* maneuvered to place the GRB within the NFIs’ field of view eight hours later, which was joined by other X-ray telescopes over the next couple of weeks. There was a clear X-ray source that had not previously been present, and it faded steadily as time passed (Figure 2.9).

Observations taken by optical observatories of this burst over the next couple of days also showed a new, rapidly fading optical source [41]. Curiously, the optical afterglow did not seem to simply decay as a power law, but instead flattened or rebrightened after about ten days (Figure 2.10). In addition, after the optical afterglow had faded sufficiently, a faint host galaxy at a redshift of 0.695 could be seen as well [42]; however, this was only observed a few years later, and so GRB 970228 was not the first GRB for which a redshift was measured.

That honor instead belongs to GRB 970508, also detected by *BeppoSAX*. The NFIs observations began even earlier this time (5.7 hours after trigger), and once again the X-ray and optical afterglows were clearly observed. This time, a radio afterglow was observed as well, and the unprecedented observations across such a large energy range allowed for the construction of a multiwavelength afterglow spectrum (Figure 2.11). In fact, its optical afterglow was sufficiently bright that measurements revealed redshifted emission lines in the afterglow itself, which solidly placed the GRB at a distance of $z = 0.835$ [45]. (This measurement was later confirmed by spectral measurements of the host galaxy [47].) For the first time since the start of the field, the question of the origin of GRBs had been answered with direct evidence.

The discoveries kept rolling in. In 1998, a supernova (SN 1998bw) was discovered less than half an arcsecond away from the position of GRB 980425 (Figure 2.12). This was a supernova of type Ic, a type of core collapse supernovae. With this in mind, the rebrightening in the optical afterglow lightcurve of GRB 970228 and others was interpreted as the emission from an underlying supernova [44]. Thus, GRBs must be closely connected with the deaths of massive stars.

All the previously mentioned GRBs in this section were long GRBs; to be precise, one should say that *long* GRBs must be closely connected with the deaths of massive stars. Much of this discovery was thanks to *BeppoSAX*'s accurate localiza-

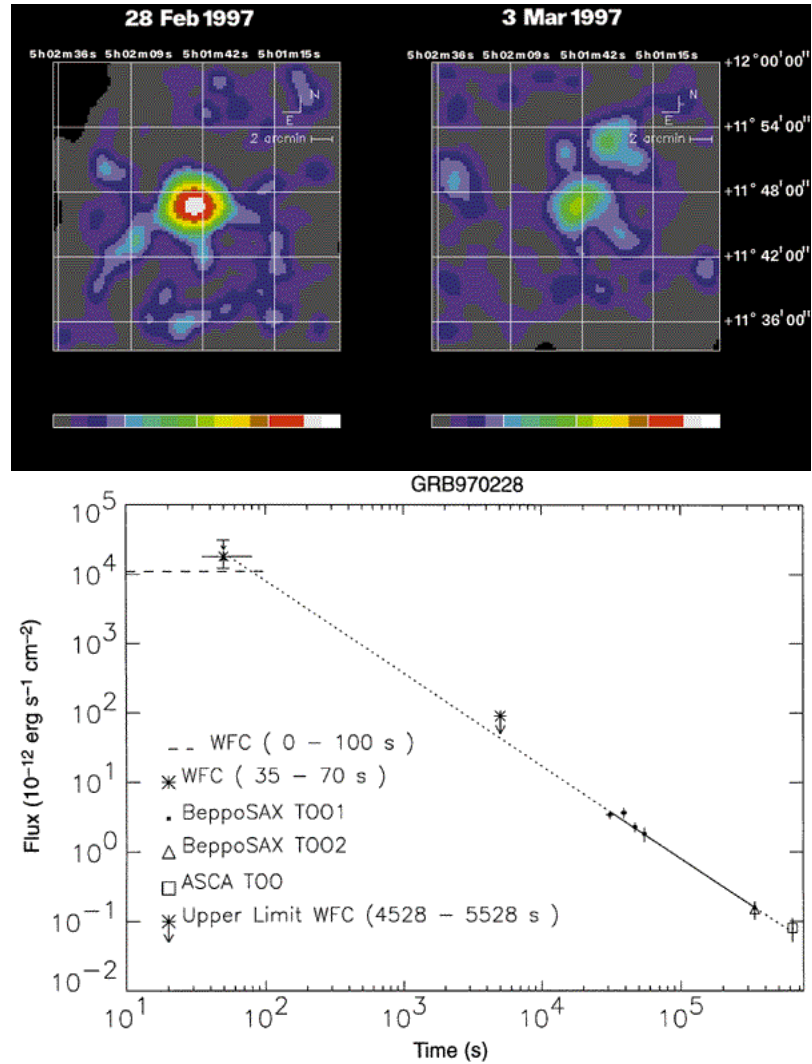


Figure 2.9: **GRB 970228 was the burst with the first observed X-ray afterglow.** *BeppoSAX* detected and localized the GRB, then slewed to observe it with its more sensitive NFIs starting eight hours after the burst began. The top plot shows the X-ray image of a small area around the GRB’s localization; there is a clear X-ray source that fades over the next few days. The bottom plot shows that the X-ray flux decreases smoothly over a few weeks’ time. [43]




Figure 2.10: **GRB 970228 was also the burst with the first observed optical counterpart, which also faded in time.** In addition, the optical lightcurve showed an excess or bump at late times, which was later interpreted as being evidence for an underlying supernova [44].

tions. The first optical counterpart of a *short* GRB, on the other hand, was detected thanks to the High Energy Transient Explorer (*HETE-2*, as the rocket that carried the original *HETE* suffered a failure soon after launch), which was launched in 2000 as the first GRB-focused mission. It had gamma-ray instruments for GRB detection, X-ray imagers for localization, and soft X-ray cameras to increase the energy range coverage. It was also the first satellite that could localize GRBs on board in realtime, and it transmitted this information to the ground within seconds, greatly increasing the speed with which other telescopes could observe a newly detected GRB [51].

HETE-2's rapid and precise localizations led to the detection of the first optical afterglow of a short GRB in GRB 050709. (The first X-ray afterglow was detected by *Swift*, discussed later.) The optical afterglow of this short GRB faded similarly to the optical afterglows of long GRBs; however, there was no evidence for a supernova




Figure 2.11: **GRB 970508** was the GRB with the first redshift measurement and the first radio afterglow, and its optical and radio afterglows were bright enough that a complete spectrum across a large energy range could be constructed. The spectrum was well described by a power law with two breaks, which is consistent with the predictions of the fireball model [49].




Figure 2.12: **Supernova 1998bw** was found in the same region of the sky as **GRB 980425**, less than half an arcsecond apart. This marked the first clear association between GRBs and SNe. The archival image of the host galaxy is shown on the right, and the new observation on the left; the supernova is the bright new source marked by the red arrow [48].

bump despite its relative proximity ($z = 0.16$). The GRB sat in a subluminous part of its host galaxy [52], whereas long GRBs (being related to core collapse supernovae) tended to be found in bright star-forming regions. This fact, combined with the energetics and short time scale, suggested that short GRBs came from the mergers of compact objects, such as black holes or neutron stars.

During this time, at the higher end of the electromagnetic spectrum, observations of GRBs by Very High Energy (VHE, >100 GeV) instruments had also been performed but had (and have) so far been unsuccessful. The Whipple Telescope [53], Milagro [54], the High Energy Stereoscopic System (HESS) [55], and the Very Energetic Radiation Imaging Telescope Array System (VERITAS) [56] have only been able to provide upper limits for VHE emission. Earlier, during the *CGRO* era, Milagro did observe an excess coincident with a GRB; unfortunately, the excess had only a borderline significance [57], and the question of VHE emission from GRBs remains for the current generation of VHE instruments.

We move now from telling a tale of individual GRBs to discussing the overall contributions from a single mission. As the knowledge of GRBs increased, the field had progressed to the point where single, game-changing GRBs had become rare².

Because *HETE-2* was a GRB-focused mission, its design choices included a large energy range and rapid localization. *Swift* (launched in 2004) expands and improves upon these abilities. The instruments on *Swift* are: the Burst Alert Telescope (BAT, 15 to 150 keV), for detecting GRBs; the X-Ray Telescope (XRT, 0.2 to 10 keV), for improved localization and long-duration observations of X-ray afterglows; and the Ultra-Violet / Optical Telescope (UVOT, 170 to 600 nm wavelengths), coaligned with the XRT for accompanying UV and optical afterglow observations. The BAT localizes GRBs with arcminute precision, and like *BeppoSAX* before it, *Swift* can slew to place a GRB in the XRT field of view. But unlike *BeppoSAX*,

²although not entirely gone; see the discussion of GRB 130427A in Chapter 5




Figure 2.13: **The X-ray afterglows detected by *Swift* have a few common features.** The lightcurves often follow a shallow-steep-shallow-pattern, although there are sometimes deviations from this pattern. Afterglows can also have X-ray flares on top of the standard set of broken power laws, and for some — like GRB 050502B here — the X-ray flare releases a very large fraction of the total energy. Here, the BAT data points (extrapolated into the XRT energy range) are plotted with crosses and the XRT data points are plotted with circles [61].

which could only slew with a manual command, *Swift* slews automatically as soon as the BAT detects a GRB, and the XRT can begin observing within a minute of the burst trigger.

Thanks to this unprecedented agility, *Swift* is often able to observe the start of the afterglow. For some lucky GRBs, *Swift* has unbroken observations from the start of the prompt emission to the late-time afterglow, by combining the observations of the BAT and XRT. It detects an average of two GRBs per week, which translates to a total of almost 1000 GRB afterglows since launch. On top of this, the XRT also often observes the afterglows of GRBs detected by other instruments.

A standard afterglow picture emerged from all of these observations. Figure 2.13 shows a few examples of *Swift*-detected afterglows and illustrates the features that are often observed. Figure 2.14 shows the canonical X-ray afterglow light




Figure 2.14: **A canonical X-ray afterglow lightcurve emerged thanks to *Swift*'s trove of observations.** The lightcurves often feature rebrightenings or flares and multiple breaks, including jet breaks, from which a jet's opening angle can be calculated. The prompt emission (phase 0) is followed by a steep decay (phase I) as whatever process produced the prompt emission shuts off. The steep decay is sometimes followed by a shallow plateau (phase II) and x-ray flares (phase V) — suggesting renewed activity — or followed directly by a more shallow decay (stage III). The timing of the break in the lightcurve at the end of the shallow decay can be used to calculate the size of the jet which produced the GRB emission [59].

curve with the standard features seen in many GRB afterglows. The *Swift* detections of the early afterglow emission showed that the afterglow is often complicated by rebrightenings or flares and multiple breaks in the light curve. In particular, many X-ray afterglows exhibit breaks at late times that are consistent with being due to a jet break: As the jet slows and spreads, when the relativistic beaming angle of the jet becomes larger than the geometric opening angle, the total observed flux drops suddenly [59]. Measurements of jet breaks indicate that GRB jets tend to have opening angles of 5 to 10° [60].

Thanks to many different instruments over the last few decades, GRBs were well-observed across a wide range of energies. However, except for a few hints of high-energy emission detected by EGRET, the behavior of GRBs at gamma ray energies and higher was still unclear. The *Fermi* Gamma-ray Space Telescope was

launched in 2008 and finally explored this behavior, answering some questions and introducing new ones.

Chapter 3: The *Fermi* Gamma-ray Space Telescope

Fermi was launched on June 11, 2008 as the spiritual successor to *CGRO*. The two instruments on *Fermi* are the Large Area Telescope (LAT) and the Gamma-ray Burst Monitor (GBM). The LAT is a pair conversion telescope that detects and tracks gamma rays between tens of MeV and hundreds of GeV. The GBM is a set of fourteen scintillators that detect photons from 8 keV to 40 MeV. Together, the LAT and GBM provide high-energy observations across almost eight orders of magnitude in energy.

The design of the *Fermi* instruments was influenced by and builds on the designs of predecessor experiments. The *Fermi* instrument teams took advantage of new technologies (such as silicon strip detectors for the LAT-Tracker), used increased onboard processing capability to implement sophisticated trigger and data filter algorithms, and made design choices to address known limitations of previous instruments.

3.1 Large Area Telescope (LAT)

For photons with energies of MeVs to GeVs, pair conversion has the largest interaction cross-section (Figure 3.1). The specific interaction is $p + \gamma \rightarrow p + e^+ + e^-$: For the purposes of a pair conversion telescope, an incoming gamma ray interacts with a proton in the nucleus of some target material and produces an electron-positron pair [30]. The interaction cross-section depends on the atomic number Z of the target material, among other things, so materials with large atomic numbers

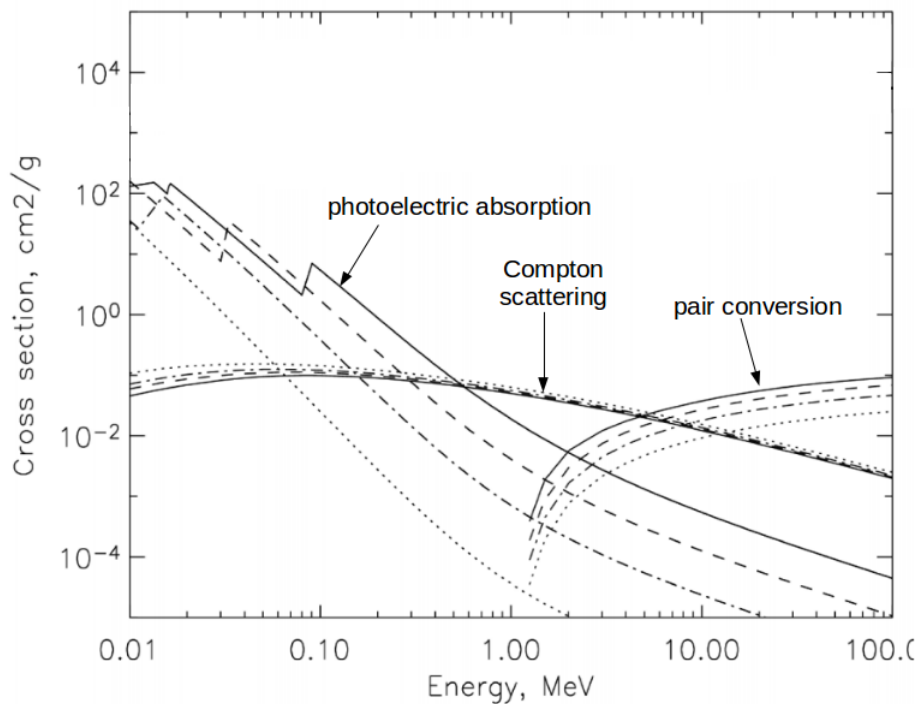


Figure 3.1: **At energies above tens of MeVs, pair conversion has the largest interaction cross section.** The different lines represent different target materials: silicon, germanium, iodine, bismuth, in order of both lowest to highest curves and smallest to largest atomic number. Tungsten is used in the LAT and has an atomic number (74) between iodine (53) and bismuth (83), so its cross-section curves fall between the top two. Original figure from [29].

are generally used [29].

After the photon pair converts into an electron-positron pair, the charged particles propagate through a tracker, which detects the electric signals left by the charged particles' interactions with materials within the tracker. By recreating the charged particles' paths from these signals, the trajectory of the original photon can be determined. The particles then encounter a heavy calorimeter, where they interact with the dense material and create a shower of particles. The calorimeter absorbs the energies of these secondary particles, and provides a measure of the energy of the original photon. Since the direction and energy measurements rely on charged particles, the tracker and calorimeter are surrounded by an anti-coincidence detector. This identifies charged particles and provides a way to determine whether a signal is due to a neutral (gamma ray) or charged (background) particle.

The Large Area Telescope (LAT), the primary instrument on *Fermi*, is a pair conversion telescope. The Tracker is made of 16 towers in a 4×4 configuration; each tower has alternating layers of tungsten foil to interact with the photons and silicon strip detectors (SSDs) to detect the charged particles that are produced. The Tracker towers sit atop Calorimeter towers made of Cesium Iodide crystals in a segmented design. The Tracker and Calorimeter are surrounded by a segmented Anti-Coincidence Detector (ACD) to reject the charged particle background, as well as a micrometeoroid shield and thermal blanket to passively protect the instrument (Figure 3.2).

The relevant references for this section are Atwood et al. (2009) [31] for an overview of the LAT; Atwood et al. (2007) [32] for details on the Tracker; Grove et al. (2010) [33] for details on the Calorimeter; and Moiseev et al. (2007) [34] for details on the ACD.




Figure 3.2: **A gamma ray converts into an electron-positron pair in the Tracker; the pair then deposits their energy into the Calorimeter.** By reconstructing the energies and tracks of the charged particles, the energy and track of the original gamma ray can be determined.

3.1.1 Design

The LAT has been described as a particle physics experiment that was shot into space. Unlike particle detectors on the ground, space-based particle detectors cannot be serviced (unless they're Hubble) or restocked. They must be efficient and light yet sturdy and able to withstand extreme conditions. The LAT was designed to minimize power requirements and have no consumables; its structure has minimum dead space, and makes maximum use of its elements. It has built-in redundancy, so that it can still function if part of it fails or is damaged.

Gamma rays interact with the tungsten foil in the Tracker and pair produce. The tungsten layers are thin (0.035 radiation lengths) in the top 12 out of the 18 layers of the detector (“front”) to minimize multiple scattering events of the electrons and positrons. In the four remaining layers of tungsten (“back”), the foil is thick (0.18 radiation lengths) to increase the chance of interactions with high-energy

photons. The bottommost two layers of the Tracker do not have any tungsten foil, since the Tracker requires three adjacent hits (i.e., signals in three adjacent SSD layers) to record a trigger, so a conversion event in the bottom two layers would be pointless.

Once an electron-positron pair is created, these charged particles propagate through the detector and their trajectories are tracked by the SSDs. The SSDs are arranged in x - y pairs (known as a hodoscopic configuration) in the layers of the Tracker (Figure 3.3). In the layers with tungsten foil, the SSDs are placed as close to the bottom of the foil as possible, to measure the position of the particle before the effects of scattering manifest. The SSDs act as both the triggering mechanism and the charged particle tracker. Thus, the Tracker is self-triggering, and the number of necessary parts is reduced. Each SSD layer can be read out from either end, and the readout systems for each tower are independent to provide robustness and redundancy. The Tracker layers are supported by a carbon composite structure, which was designed to be light, strong, minimally interacting, and shielding.

Each of the 16 Tracker towers sits on top of a corresponding 8-layer Calorimeter module. Each module has 96 CsI(Tl) crystals, and each of the 8 layers is perpendicular to its neighbors, mimicking the hodoscopic configuration of the SSDs. Every crystal has a large (sensitive to energies from 2 MeV to 1.6 GeV) and small area (100 MeV - 70 GeV) photodiode on each end, and a comparison of the readout from the two ends can determine where the particle interaction occurred along the crystal length. All of this combines to give the Calorimeter the capability to image shower profiles, on top of the standard energy deposition measurements. Thus, even if only a portion of the shower is contained in the Calorimeter, the shower image can provide an energy estimate (Figure 3.4).

As *Fermi* is a low Earth orbit, it is inundated by a large flux of charged particles and must be able to distinguish between external charged particles (the




Figure 3.3: **A simplified diagram of a photon interacting with the Tracker.** In this example, the photon pair-converts in the top tungsten layer. As the child particles propagate through the layers, they create “holes” in the SSDs; the x - y locations of these holes in the different layers are combined to determine the particle tracks.




Figure 3.4: **Each Calorimeter module is made of 96 Cesium Iodide crystals in 8 layers. The layers contain 12 crystals each and are in a hodoscopic configuration.** Every crystal has a set of photodiodes on each end to measure the energy deposited in the crystal.

background) and the charged particles produced within the LAT from gamma rays (the signal); at the energies that the LAT observes, the background rate is larger than the rate of gamma rays by 3 to 5 orders of magnitude. The Anti-Coincidence Detector (ACD) is a set of 89 plastic scintillator tiles that surround the Tracker and Calorimeter. The tiles overlap in one direction, to minimize the amount of dead area (i.e., areas that are not covered by an ACD tile through which charged particles can enter the LAT undetected), and the gaps in the other direction are covered with flexible scintillating fiber ribbons (Figure 3.5).

The ACD detects charged particles, so to zeroth order, any signal in the Tracker or Calorimeter that is accompanied by an ACD signal is a background charged particle. However, when signal particles encounter the dense calorimeter, they can produce secondary particles that propagate back upwards through the Tracker, known as “backsplash.” These secondary particles would also create hits in the ACD, which would cause it to erroneously veto these events (“self-veto”). This had caused problems for EGRET, which lost a great deal of effective area at higher energies because of the high frequency of vetos caused by backsplash from energetic gamma rays. The LAT ACD is segmented to address this issue; an ACD hit that accompanies an event only causes a veto if the hit is in a tile close to where the incident event entered the LAT (which would indicate that the incident event was a charged particle).

3.1.2 LAT onboard trigger and event reconstruction

On average, the charged particle rate experienced by the LAT is several orders of magnitude larger than either the photon rate or the maximum LAT downlink rate. Events must undergo multiple stages of filtering in order to make it to the downlink stage: hardware trigger request, hardware trigger accept, and onboard gamma-ray filter. The main reference for this section is Ackermann et al. (2012) [35].




Figure 3.5: **The ACD is made of 89 plastic scintillator tiles that surround the Tracker and Calorimeter.** The tiles overlap in one direction to minimize the amount of space left uncovered, and scintillating fiber ribbons cover the space in the other direction. Wavelength shifting fibers embedded in the tiles allow for uniform light collection; each of these is attached to two PMTs, for redundancy.

Each event detection begins with a trigger request, or trigger primitive. One type of trigger primitive requires hits to three adjacent layers of SSDs in the Tracker (also known as the TKR trigger primitive), where a “hit” means the SSDs record a signal above a threshold. Two other types of trigger primitives involve the Calorimeter (`CAL_LO` / `CAL_HI`) and require signals above 100 MeV / 1 GeV to have been read out by any Calorimeter crystal. A fourth type of trigger primitive is the ROI (essentially a veto), which requires a signal in specified groups of ACD tiles relative to the particular Tracker tower that registered a trigger, to help determine if an event is a charged particle. There are also four other types of trigger primitives that are not crucial to the subjects discussed in this work; please refer to Ackermann et al. (2012) [35] for more details.

Certain combinations of trigger primitives are required in order for these to become upgraded to a single trigger accept; these combinations are known as the trigger engines. For instance, a TKR (track detection) without an ROI marks an event as a potential gamma ray, as does a TKR in combination with a `CAL_LO` or `CAL_HI`. The event then passes into the trigger accept stage during which the data




Figure 3.6: **In the LAT-based coordinate system, an event’s direction is described by the angle from the boresight (the z-direction), θ , and the azimuthal angle, ϕ .** The x-axis is parallel to the solar panel arms of the spacecraft.

is read out, and then enters the filtering stage. The relevant filter for the purpose of this work is the GAMMA filter, which selects for gamma rays using a series of hierarchical tests.

The GAMMA filter is a series of fifteen simple tests or vetos with increasing complexity. In order to be computationally efficient, the filter starts with the vetos that take the least computational time. The goal is to reject events as early as possible, and an event that fails at any step is immediately rejected. Roughly speaking, the steps are as follows: First, events with certain patterns of ACD tile hits in conjunction with energy above a certain threshold are vetoed. The Calorimeter information is then used with the ACD information to veto events that are consistent with being low energy charged particles. The Tracker hits are compared to the ACD hits to veto events for which the Tracker and ACD hits are spatially consistent. The Tracker hits are then analyzed to determine the two dimensional projections of the track (in the x - z and y - z planes; Figure 3.6), and these are used with the Calorimeter measurement and ACD hits to veto events whose projected tracks are spatially consistent with the ACD hits.

3.1.3 Ground event reconstruction

When events reach the ground, their associated Tracker, Calorimeter, and ACD signals are analyzed and reconstructed to provide the direction, energy, and quality of the event. The Tracker is used to identify candidate tracks and parametrize the topology of the hits. The Calorimeter is used to measure the amount of energy deposited and the shape of the deposit, to estimate the energy of the incoming event. The ACD hits are used on a tile by tile basis to help reject charged particles, based on the distance between the hit tiles and the reconstructed direction of the event.

Tracks are reconstructed from the Tracker hits using a Calorimeter-seeded pattern recognition and a blind search pattern recognition method. The Calorimeter-seeded pattern recognition uses clusters of Tracker hits as candidate tracks, and includes only the tracks that point toward the centroid of the Calorimeter energy deposition. (The blind search pattern recognition does not have this requirement.) Additional clusters are added to the track if they are deemed likely to have been produced by the same event, and the track reconstruction is updated after each addition. Candidate tracks that point back to the Calorimeter centroid and that are longer and straighter are weighted more favorably. A classification tree analysis is used to calculate the probability P_{core} that the reconstructed track falls within the core of the point-spread function (PSF); i.e., how well the track is reconstructed.

The Calorimeter reconstruction uses three different methods of energy estimation, which were tuned using simulations to be sensitive to different energy ranges and incidence angles. The best track that was found from the Tracker reconstruction is used to correct for any missing energy due to energy leaking out of the sides or back of the Calorimeter or in inert material. A classification tree analysis is then used to select the best energy estimate out of the three and calculate the probability P_E that the reconstructed energy is in the core of the energy dispersion.

The ACD hits are used to classify events as photons or charged particles. Events for which the best track points back to an area of ACD tile hits or regions of lower ACD sensitivity (e.g., corners and gaps, although these have miniscule influences within the overall event rate) are filtered out as background. Next, events for which there are signs that the direction was not accurately reconstructed (e.g., the event passes mostly through the gaps between the Tracker towers) are filtered out. If the total energy deposited in the ACD is too large to be due to backscatter when compared to the amount of energy deposition measured by the Calorimeter, the event is considered a charged particle and filtered out (this requirement is sometimes loosened during activity such as solar flares that deposit large amounts of energy in the ACD). A final classification tree analysis is performed on the remaining events to calculate P_{CPF} (“charged particles in the field of view”).

Further filtering on charged particles is performed by considering the topology of the Tracker and Calorimeter signals. The topology of the Tracker signal is put through a classification tree analysis to recognize charged particles that might have been missed with the previous analysis as well as calculate the probability P_{TKR} that the event is a gamma ray. Similar analyses are performed on the Calorimeter signal to give P_{CAL} . Finally, all of the probabilities are combined into P_{all} , the overall probability that an event is a gamma ray.

The events must pass a few minimal requirements for the standard event classes. Trivially, the event must have a reconstructed track, and the track must point into the Calorimeter where they must pass through four radiation lengths, and deposit 5 MeV of energy. (Events that do not pass these stages are often used in non-standard event types.)

Roughly speaking, the event classes have different amounts of background rejection; the loosest standard event class (“Transient”) is much less rigorous about excluding charged particles than is the most stringent standard event class (“Ultr-

aclean”). Having multiple event classes with various levels of background rejection allows for the optimization of different kinds of analyses. For GRB analysis with standard event classes, the Transient event class is used for the prompt emission (practically speaking, less than 100 seconds after a GRB begins) and the Source event class is used for the temporally extended emission.

In order to perform analyses of the standard event classes, it is necessary to know how well photons are reconstructed. This information is contained within a set of Instrument Response Functions (IRFs): the point-spread function (PSF), the energy dispersion, and the effective area (A_{eff}). The PSF is the probability density for a reconstructed photon direction; the energy dispersion is the probability density for a reconstructed photon energy; and the A_{eff} is the product of the collecting area and the efficiency for detection. All three of these depend on the energy, direction, and event class of a photon.

The event reconstruction and IRFs require detailed knowledge of how photons interact with and are detected by the LAT. Prior to launch, all of this information was based on simulations, and the event analysis that corresponds with this is known as PASS 6. This was used until August of 2011, at which point the LAT team released PASS 7 to the public. PASS 7 improved on PASS 6 by using on-orbit data, to address effects that were not known prior to launch. In November of 2013, PASS 7 was updated with improved instrument calibrations and a new set of IRFs, and was released as PASS 7 **reprocessed**, or P7REP. The newest iteration of the data analysis, PASS 8, was released in June of 2015, and is a complete reprocessing of all LAT data and has improved reconstruction accuracy, a greater energy range, and a larger acceptance. The LAT standard analysis in this thesis uses only PASS 7.

The PASS 7 Transient event class (“P7TRANSIENT”) has the following cuts: 1) $P_E > 0$ and $P_{\text{core}} > 0$; 2) final reconstructed energy must be greater than 10 MeV; 3) final reconstructed energy must not be more than five times the energy deposited




Figure 3.7: The acceptance of the LAT (the effective area integrated over the field of view) decreases at each stage of the event filtering process. Most of the events that are rejected at each stage are charged particles.

in the Calorimeter; 4) $P_{\text{CPF}} > 0.1$; and 5) $P_{\text{all}} > 0.2$. Thus, P7TRANSIENT events only have minimal requirements on being well-constructed.

The P7SOURCE events are a subset of the P7TRANSIENT events; the requirements are along similar lines but stricter. The requirements of the P7SOURCE event class are: 1) The event must not have been flagged as a charged particle in either the Calorimeter or Tracker topological analysis, and $P_{\text{E}} > 0.1$ or 0.3 (depending on the energy estimate method); 2) the event must not have been flagged as a Minimum Ionizing Particle (MIP, an energetic charged particle that deposits a minimum amount of energy in a material); 3) the event must have a good agreement between the Tracker and Calorimeter reconstructed directions; 4) the event must pass a tighter and energy-dependent cut on P_{core} ; and 5) the event must pass a tighter and energy-dependent cut on P_{all} . (See [35] for specific formulas.)

In addition to the standard event classes, the GRB group also uses the LAT Low Energy (LLE) event selection. This non-standard event selection has an extremely loose cut on background, and is therefore not suitable for event-by-event




Figure 3.8: **The event rates onboard is one or two orders of magnitude larger than the event rates for the standard event classes.**

analysis. However, this is fine for short-duration, high-signal events such as solar flares or GRB prompt emission. It has an energy range of 30 MeV to 100 MeV, thereby connecting the GBM energy range with the LAT standard energy range. At these energies, the uncertainties on the event reconstruction are too large for these low-energy events to be used in standard LAT analysis. The idea of the LLE is to ignore the LAT's imaging capabilities and use it instead as a rate detector. The LLE definition has extremely loose cuts on background events, after including a cut on radius around a source position.

The LLE event class has the following requirements: 1) the event must pass the GAMMA filter; 2) the event must have a reconstructed track; 3) the event must have either triggered the Tracker or deposited at least 1 GeV in the Calorimeter; 4) the event must have an energy reconstruction using the method that is best for low energies; 5) the event came in at less than 90° degrees from the boresight (that is, the event did not move upward through the LAT), $\theta < 90^\circ$; and 6) the event direction is within the PSF of the source position (such as a GRB).

An event's time of arrival is recorded in Mission Elapsed Time (MET), which

are the number of seconds since 1 January 2001, UTC (MET = 0).

3.2 GBM

The Gamma-ray Burst Monitor (GBM) observes a variety of sources in hard x-rays, but — as one might guess from its name — its primary task is to detect GRBs. It is a collection of 14 scintillation detectors and is the second instrument on *Fermi*. There are 12 Sodium Iodide (NaI) detectors positioned on two opposite sides of the spacecraft (six on each side) and pointed in different directions; each set of six NaIs is accompanied by a Bismuth Germanate (BGO) detector as well. The GBM observes the entire unocculted sky at all times. There is some overlap in the energy ranges of the NaIs (8 keV to 1 MeV) and the BGOs (200 keV to 40 MeV), and between the BGOs and the LAT. The primary reference for this section is Meegan et al. (2009) [36].

Each NaI detector is made of a 5-inch diameter circular crystal protected by an aluminum casing and glass window. They are pointed in different directions (Figure 3.9), and are used to localize a burst by comparing the strength of the signals observed by the different detectors. A PMT is attached to the back of each NaI within the casing to convert the scintillation light produced by a photon detection into an electrical signal.

Each BGO detector is made of a 5-inch diameter and 5-inch long cylindrical crystal; the circular edges are the detecting surfaces and are protected by glass windows and a carbon fiber-reinforced plastic. A PMT is attached to either end of the crystal.

The signal from each PMT is analyzed separately. A larger signal indicates a larger energy deposit; if the signal is above a certain threshold, the peak height is detected and the signal is marked as an event. The peak height is then dropped into one of either 8 or 128 pre-determined bins, known as Pulse Height Analysis (PHA)




Figure 3.9: The twelve NaI detectors (labeled 0 through 11) and two BGO detectors (12 and 13) are positioned on opposite sides of the spacecraft. The NaIs are oriented so that the GBM observes the entire unocculted sky at all times.

bins, based on onboard lookup tables.

The GBM records data in three different data types: CTIME, CSPEC, and TTE. CTIME data has 8 PHA bins and a minimum temporal resolution of 64 ms; the default time binning is 256 ms, but drops to 64 ms for a short period of time after a GRB detection (described later). CSPEC data has 128 PHA bins, and has a default time binning of 4.096 s but drops to the minimum of 1.024 s after a GRB trigger. Like CSPEC, the TTE data (which stands for “Time-Tagged Events”) has 128 PHA bins but encodes the time-of-arrival for every photon rather than binning. Prior to 26 November 2012, the TTE data were stored in a buffer, and ~ 30 s of TTE data before trigger were made available for every GBM detection. Since then, Continuous TTE (CTTE) data have been available during normal spacecraft operations.

The GBM triggers on GRBs onboard, in realtime. The GBM flight software searches for increases above threshold over the average background rate in the count rates of least two NaIs. There are multiple triggering algorithms in place, covering a variety of energy and time ranges, so that the GBM is sensitive to different types of transient phenomena (e.g., soft gamma repeaters versus GRBs, or short vs long GRBs). When a GRB is detected, the TTE data (starting 30 seconds before the trigger) is sent to the spacecraft until ~ 300 seconds after trigger; at the same time, the CTIME and CSPEC temporal resolutions drop to 64 ms and 1.024 s, respectively, for ~ 600 s.

After a GRB detection, the GBM localizes the GRB onboard by comparing the relative rates in the NaIs to lookup tables. These tables report the expected relative rates in different NaIs for a set of positions on the sky on a $\sim 5^\circ$ grid. (However, this lookup table assumes that the spacecraft is pointed at the zenith which was an acceptable approximation for the 35° rocking angle at launch, whereas the rocking angle has been 50° as of September of 2009; see next subsection.) More precise

localizations are performed on the ground, both automated (using a lookup table with a finer grid) and manual (using a full analysis with detector responses).

The Detector Response Matrices (DRMs) are the GBM instrument response functions and are required for any analysis of GBM data. The DRMs account for factors that affect how the detectors measure a signal, such as the detector's angle to the source and the spacecraft's orientation with respect to the Earth. They are generated using Monte Carlo simulations and provided as RSP files.

3.3 Observational strategies

Fermi is in a low-Earth orbit at $\sim 26^\circ$ inclination and with a 96 minute orbital period. Most of the time, *Fermi* operates in a survey mode, alternating between rocking 50° below the orbit plane for one orbit and 50° above for the next. This results in complete coverage of the sky every two orbits (~ 3 hours). Up until September of 2009, the rocking angle was 35° , but was changed to reduce the batteries' exposure to thermal radiation from the Earth.

Fermi can also be pointed at a particular part of the sky for an extended period of time. Targets of Opportunity, or ToOs, are often used for flaring sources (such as Active Galactic Nuclei or bright solar flares) and are manually scheduled. In contrast, Autonomous Repoint Requests, or ARR, automatically occur when the GBM triggers on a GRB that is more likely to be detected by the LAT.

During an ARR, the spacecraft slews within minutes to place the GRB position (that was calculated by the GBM onboard) near the center of the LAT field of view. This is maintained for a maximum of 2.5 hours; if the target position is obscured by the Earth within this time, the spacecraft points away from it and returns to observe the source as soon as it can. Initially, ARRs lasted for 5 hours, but this was changed to 2.5 hours when it was realized that the LAT-detected emission from GRBs usually lasts for only an hour or so. (The striking exception is GRB 130427A, from which

the LAT detected emission lasting almost a day; see discussion in Chapter 7.)

Initially, the spacecraft was programmed to slew due to an ARR immediately after trigger. This was detrimental to GBM analysis, since the background during a slew is very difficult to model. This also proved detrimental to the LAT, as we discovered that the LAT onboard GRB detection algorithm was disabled during ARRs and prevented the LAT from triggering on a GRB onboard and providing a rapid localization. A 30-second delay was instituted in August of 2011 to address both these issues.

During normal spacecraft operations, the LAT and GBM are constantly taking data except when the spacecraft is within the South Atlantic Anomaly (SAA). This is an area over the southern Atlantic ocean with a particularly large charged particle flux. Since charged particles are background events for both the *Fermi* instruments, the LAT does not record data and the GBM does not trigger (but still records data) when the spacecraft passes through the SAA. The LAT and GBM boundaries for the SAA are independent and slightly different.

3.4 GRB handling: Communication between the GBM and LAT

The GBM is a collection of scintillation detectors with no imaging capabilities; it is adept at detecting GRBs but can only roughly localize them. The LAT can localize GRBs with greater precision, but since GRB emission spectra tend to peak at 100s of keV, the signal in the LAT energy range is much smaller than the signal in the GBM energy range. The LAT therefore uses information provided by the GBM to aid in the LAT's search for GRB emission.

When the GBM detects a GRB onboard, it sends an Immediate Trigger Signal (ITS) to the LAT ~ 2 ms after trigger to inform the LAT of detection and prepare the LAT to receive information. Within the next few minutes, the GBM sends a series of messages (telecommands) to the LAT at predetermined times ($\sim 2, 5, 10, 20,$

30, 60, 90, and 150 seconds after trigger). Each telecommand contains housekeeping data plus the following information related to the GRB detection:

- Trigger time in MET, up to tenths of a second
- Message number in sequence
- Right Ascension (RA) of localization, in arcminutes
- Declination (DEC) of localization, in arcminutes
- Error of localization, in arcminutes
- Which localization algorithm was used
- The two most likely event classifications with their probabilities (e.g., “GRB”, “local particles”)
- Which detection algorithm was triggered

Each time it receives a localization telecommand, the LAT runs its onboard GRB detection algorithm around the GBM detection time and localization and searches for an excess that could be due to a GRB. (This is discussed further in Chapter 4.) After these 8 localization telecommands, the GBM sends a final closeout telecommand.

While the GBM is sending these messages, it simultaneously and continuously checks its onboard detection against a set of criteria to see if the detection merits the spacecraft’s repointing. It uses the peak counts flux and energy fluence (both measured in the 50 to 300 keV energy band), as well as hardness ratios. The fluence is the integral flux over a period of time; here, it refers to the total energy since the trigger time. The hardness ratios are the ratios of peak flux or fluence between a higher energy bin and a lower energy bin. The two energy ranges are (300 keV to 30 MeV) and (150 keV to 300 keV). All of the calculated values (peak flux, fluence,

Within LAT field of view	
peak flux	$10 \text{ ph cm}^{-2} \text{ s}^{-1}$
fluence	$1.5 \times 10^{-5} \text{ erg cm}^{-2}$
peak flux and HR	$2 \text{ ph cm}^{-2} \text{ s}^{-1}$ $\beta > -2$
fluence and HR	$3 \times 10^{-6} \text{ erg cm}^{-2}$ $\beta > -2$
Outside LAT field of view	
peak flux	$40 \text{ ph cm}^{-2} \text{ s}^{-1}$
fluence	$1 \times 10^{-4} \text{ erg cm}^{-2}$
peak flux and HR	$10 \text{ ph cm}^{-2} \text{ s}^{-1}$ $\beta > -1.8$
fluence and HR	$2 \times 10^{-5} \text{ erg cm}^{-2}$ $\beta > -1.8$

Table 3.1: **The thresholds for an ARR are different based on whether the GBM localizes the GRB to within the LAT field of view (defined here as 60°).** If the burst properties pass any of the relevant thresholds, the GBM sends a request to repoint the spacecraft. β refers to the high-energy power law index of the GRB spectrum, assuming the spectrum can be modeled as a smoothly broken power law.

and hardness ratios) are continuously updated as the GRB emission continues and evolves.

As a first step, the GBM flight software checks that the detection is likely to be a GRB ($>70\%$ according to its flight software). If this is true, it checks the burst properties against a series of thresholds. The values of these thresholds depend on whether the GRB is within the LAT field of view (defined here as 60°) or outside.

Table 3.1 lists the possible ARR criteria. If a GRB passes any of these, the GBM sends a request to repoint the spacecraft. β refers to the high-energy index of the Band function (a smoothly broken power law with a particular curvature), which is often used to model the spectrum of GRBs. $\beta > -2$ or $\beta > -1.8$ indicates a GRB with a particularly hard spectrum. The thresholds were all set based on

pre-flight simulations of GBM detections, and requiring a certain number of ARRs per year.

The single-branched criteria (either high peak flux or high fluence by itself) will be passed by bright bursts regardless of hardness ratio; the other two (requiring an appropriate hardness ratio as well) will be passed by bursts that are particularly hard but not as bright. Both are more likely to have emission in the LAT energy range than the average GRB.

3.5 Ground GRB analysis with *Fermi*

For ground GRB analysis, the CSPEC and TTE data types are most commonly used for GBM data, as these have 128 PHA bins and are more appropriate for characterizing GRB spectra. The Transient event and Source event classes are used for the LAT analysis (for analysis of times earlier than and later than 100 seconds after trigger, roughly speaking), as well as the non-standard LLE data type.

Both the LAT and GBM use a maximum likelihood analysis as their standard analyses. In this context, the likelihood is the probability of obtaining an observation based on a model. A maximum likelihood analysis is performed by varying the parameters of the model until the set of parameters that maximizes the likelihood is found.

The primary analysis performed with GBM data (besides localization studies) is spectral analysis. The background before and after the GRB is fit with a polynomial (up to 4 dimensions). A Detector Response Matrix (DRM) is convolved with a model to produce an expected photon spectrum, which is compared to the observed counts. The parameters of the chosen model are varied to find the one that minimizes a fit statistic. A commonly used fit statistic is χ^2 ; however, this is appropriate only for Gaussian data, whereas the GBM (and LAT) data are Poisson distributed. For fitting, the GBM team primarily uses the Castor statistic (C-STAT), which mod-

ifies the Cash statistic [37] for Poisson statistics so that the C-STAT approaches the χ^2 distribution at very large counts. Unlike χ^2 , C-STAT cannot be used to judge a model’s goodness of fit, and should only be used to judge two nested models’ *relative* goodnesses of fit.

For standard LAT analysis, a model with both spatial and spectral parts is convolved with the IRFs (the PSF, energy dispersion, and A_{eff}) and compared to the observed signal. The model parameters are varied to find the set of parameters that maximize the likelihood, either binned or unbinned. For a binned likelihood analysis, the events are binned into counts maps, and the number of counts in each bin is described by a Poisson distribution. (For unbinned analysis, the bins are treated as being infinitely small so that each bin contains either one or no events.) The likelihood is then the product across all bins of the Poisson probability in each bin of obtaining the observed counts based on the model. The likelihood by itself cannot be used to judge a model’s goodness of fit; instead, two models’ likelihoods are compared in a likelihood ratio to determine how much better one model fits than the other. The models are compared using the test statistic, $\text{TS} = -2(\ln(L_0) - \ln(L_1))$, where L_0 and L_1 are the likelihoods of two different models. (Roughly speaking, TS is approximately the square of the significance σ ; this is not quite accurate but gives a sense of whether a particular TS value is interesting.) For instance, when searching for the presence of a signal, L_0 would be the likelihood of the model where there is no signal (i.e., all photons are due to the background) and L_1 the likelihood of the model where a signal exists in some prespecified form. For background modeling, the LAT team uses either a set of background templates (for the Source event class and up) or the Background Estimator (for the Transient event class). The Background Estimator uses previous spacecraft data to estimate the background at a given geomagnetic position and spacecraft pointing direction, as both of these influence the background observed by the LAT [38].

For GRB analysis, LLE is used to bridge the gap between the standard LAT events (over 100 MeV) and the GBM (up to 40 MeV) when analyzing spectra. Analysis of LLE data requires a Detector Response Matrix (DRM), similar to the DRMs required by the GBM. These require large simulations to generate and are only made for GRBs and solar flares¹. The LLE background is estimated, and LLE analysis proceeds along the same lines as GBM analysis.

Once LAT data arrives on the ground, a few different automated processes run through the data to search for transient or flaring events on different time scales. The automated process most relevant to GRBs is the internal Burst Advocate (BA) Tool. The BA Tool performs a standard likelihood analysis to search for significant signals within 100 seconds of a GRB trigger (either from the GBM or from *Swift*) using the Transient event class, and within 1000 seconds using the Source event class. It also searches for an increase in the count rate of LLE data. If any of these searches finds a signal above a threshold (roughly speaking, $> 5\sigma$ for LLE and $TS > 25$ for the standard event classes), the LAT GRB group is notified. The group then performs more fine-tuned analyses and writes a GCN circular to notify the GRB community of the detection. As of December of 2013, the group also sends out a LAT Offline Position Notice, a standardized machine-readable notice with position information for robotic telescopes and other interested individuals.

¹<http://heasarc.gsfc.nasa.gov/W3Browse/fermi/fermille.html>

Chapter 4: The LAT onboard algorithm

GRBs fade rapidly, so rapid communication between observatories is essential for successful joint observations. The *Fermi*-LAT and GBM are both able to localize GRBs onboard and in realtime. This enables the announcement of detections and localizations to be distributed to the broader GRB community in less than a minute rather than the ~ 12 hours for ground localizations. In this chapter, I describe the work I performed to optimize the algorithms that detect and localize GRBs onboard the LAT.

A typical optical observatory has a field of view of only a few arcmin, and *Swift*-XRT has a field of view of about a third of a degree.¹ The GBM excels at detecting GRBs, but can only crudely localize them. The typical GRB statistical uncertainty is a few degrees, and systematic uncertainties increase these error regions by at least a few more degrees. In contrast, when the LAT detects a GRB onboard, it can localize the burst to about half a degree or less, which is a much more manageable size for follow-up observations.

In this chapter, I use the term “event” instead of “photon” or “gamma ray.” The onboard background rejection and classification are not nearly as robust as the methods used on the ground, and we cannot be sure that the particles that go through the onboard algorithm are photons.

¹*Swift*-XRT now has the ability to conduct tiled observations of larger areas; see <https://www.swift.psu.edu/secure/toop/tiling.php> for more information.

4.1 Overview

The LAT onboard GRB detection algorithm is subject to the limited resources onboard the spacecraft. Because of this constraint, every calculation involved in the onboard algorithm is computationally simpler and more efficient than the analysis on the ground. The onboard localization is therefore not as precise as the ground localization, but it is transmitted to the GRB community in seconds rather than hours. As the history of GRB research shows, this makes all the difference in a field where the objects fade rapidly.

The onboard algorithm has three stages: event selection, track reconstruction, and cluster finding. In the event selection stage, the onboard algorithm currently uses the onboard filter that selects events for gamma ray analysis; this is a loose selection with an event rate of ~ 400 Hz (Figure 4.1). The events that pass the onboard filter then undergo a track reconstruction, to roughly determine the directions these events came from.

The cluster finding stage of the LAT onboard GRB detection algorithm takes a two-tier approach. In Tier 1, the algorithm searches through a list of events for clusters of events in space and time; in Tier 2, the algorithm conducts a deep search around the clusters that passed a significance threshold in Tier 1. If a cluster passes the Tier 2 thresholds, the algorithm localizes the cluster to a position on the sky, and reports a detection to the *Fermi* GRB team and the wider GRB community. The two-tier design reduces the CPU requirements, since the thresholds can be chosen so that most Tier 1 clusters do not make it onto the more resource-intensive Tier 2 stage. A cluster is said to trigger the onboard algorithm if it successfully passes both stages.

The LAT can simultaneously run up to 32 versions of the algorithm in independent “windows.” Currently, three windows are operational; two of these are




Figure 4.1: **The rate of events that pass the onboard filter is ~ 400 Hz**, much larger than the loosest standard ground event class (transient). The LLE and Source class event rates are also plotted for comparison.

LAT-only blind searches, while the third takes GBM detections as input. In the following sections, I describe the LAT-only and GBM- seeded triggers, and the studies we performed to improve the onboard GRB detection capabilities of the LAT.

4.2 LAT-only trigger

4.2.1 Tier 1

The onboard algorithm maintains a running list of the energies, directions, and arrival times of all the events that pass the onboard filter. When a new event is added, it is considered as a seed event, and the algorithm searches for clusters around this seed.

Consider a seed event that was detected at time T_{seed} (the seed time) with a reconstructed direction $(\alpha_{seed}, \delta_{seed})$ (the seed location, in RA and dec). The algorithm takes the list of N_{WS} events that arrived *before* the seed event (N_{WS} is

the window size; currently, 40 and 80), and considers only the set of events in this list that are within T seconds and θ_1 degrees of $(\alpha_{seed}, \delta_{seed})$. It then calculates a set of temporal and spatial test statistics for this cluster to quantify the “clusteredness.” (Despite their names, these are not technically test statistics, as they are not directly related to hypothesis testing.)

The temporal test statistic TS_T is defined as

$$TS_T = \sum_{i=1}^{N_{WS}} \left| \log_{10} \left(1 - \exp^{-r_t \Delta T_i} \right) \right|, \quad (4.1)$$

where

$$r_t = \frac{\pi \theta_1^2}{2\pi (1 - \cos(\theta_m))} \times \frac{1}{\delta t}. \quad (4.2)$$

ΔT_i is the difference in seconds between two consecutive events’ arrival times; θ_1 is the maximum acceptable spatial distance between a cluster event and the seed event in Tier 1; θ_m is approximately the LAT field of view (currently set to 115°); and δt^{-1} is the typical onboard background rate in Hz (currently, 250 Hz). Thus, r_t is a normalization factor.

The spatial test statistic TS_S is defined as

$$TS_S = \sum_{i=1}^{N_{WS}} \left| \log_{10} \left(\frac{1 - \cos(\theta_i)}{1 - \cos(\theta_m)} \right) \right|, \quad (4.3)$$

where θ_i is the angular separation (in degrees) between the seed location and the reconstructed direction of the i th event.

The algorithm then combines $TS_{tot} \equiv c \cdot TS_T + TS_S$. In this calculation, c is a multiplicative factor (the “contrast”) that can be used to weight the spatial and temporal TSs differently. c is set to unity based on prelaunch simulations (i.e., TS_S and TS_T contributed roughly equally to TS_{tot}), and there is currently no plan to

change this.

If TS_{tot} does not exceed some preset Tier 1 threshold for a given cluster of events, then the algorithm uses the next event in the list as the seed event and restarts the Tier 1 calculations around the new seed event. On the other hand, if TS_{tot} is greater than the threshold, then the seed time and location are passed along to Tier 2.

4.2.2 Tier 2

In Tier 2, the algorithm uses the same formulas for TS_T and TS_S as before, although most of the parameters can be — and generally are — different from their Tier 1 equivalents (e.g., θ_2 is not necessarily the same as θ_1). The main difference is that Tier 2 maintains a running cumulative TS_{tot} instead of calculating TS_{tot} for a finite number of events. With each new event, TS_{tot} is recalculated and checked against the Tier 2 threshold (the contrast c is the same for both Tier 1 and Tier 2). The Tier 2 equivalent of the window size N_{WS} is the list size N_{LS} (currently, 200 events); if the current list of Tier 2 events has the maximum number of events N_{LS} , then the chronologically first event is dropped from the list when a new one is added.

Tier 1 searches backwards in time from the seed time, so that the seed event is the *latest* event in a cluster. To account for this, the Tier 2 search begins a preset number of events N_{HD} before the seed event (N_{HD} is called the history depth and is currently set to 100 for the LAT-only windows). The Tier 2 search continues until either elapsed time reaches the timeout T_{TO} (currently, 2 seconds) or TS_{tot} reaches the Tier 2 threshold. In the first case, the algorithm returns to the Tier 1 stage; in the second case, the algorithm proceeds to localize the presumed GRB.

The LAT-only windows were set at launch to have reasonable thresholds for Tier 1 but prohibitively high thresholds for Tier 2. These were set so that the

team could study the algorithm performance and CPU usage without risking false detections. They have not been changed, and the LAT team currently relies on the GBM-seeded window for onboard triggers.

4.2.3 GBM-seeded window

When the GBM detects a GRB onboard, the GBM trigger time and flight localization are sent to the LAT onboard algorithm and used as a Tier 2 seed. The GBM-seeded window runs in parallel with the LAT-only windows, and is equivalent to the LAT-only Tier 2 stage. The same calculations are performed, but with an independent set of parameter values.

Prior to the launch of *Fermi*, there were hints from the EGRET observations that the high-energy (GeV) emission in GRBs was delayed with respect to the lower-energy (keV-MeV) emission. To account for this potential delay, the GBM trigger information is sent to the Tier 2 stage of the LAT onboard algorithm at 2, 5, 10, 20, 30, 60, and 150 seconds after the GBM trigger. Thus, the LAT has eight attempts to trigger on a GBM detection.

An especially bright or hard GBM detection can trigger an ARR as soon as the GRB is detected. For the first few years, onboard detections were disabled during ARRs. We found this out the hard way when the LAT onboard algorithm failed to trigger on the GRB 110731A, despite this burst being bright enough to reach the threshold. To bypass this for future GRBs, a 30-second delay was introduced after ARR messages to give the onboard algorithm time to trigger. This also benefits the GBM team, since the changing GBM background is difficult to model during the rapid spacecraft slewing of an ARR.

For the first two years, the GBM-seeded search was configured to always send information to the ground following a GBM trigger. This meant that the onboard algorithm was not given the opportunity to trigger on the later GBM messages. The

	parameter	name	Win 0	Win 1	GBM
general	θ_m ($^\circ$)	LAT FOV	115.0	115.0	n/a
	δt^{-1} (Hz)	background rate	250.0	250.0	n/a
Tier 1	θ_1 ($^\circ$)	angular cut	10.0	10.0	n/a
	N_{WS}	window size	40	80	n/a
	c_1	contrast	1.0	1.0	n/a
	$TS_{tot,1}$	threshold	90.0	90.0	n/a
Tier 2	θ_2 ($^\circ$)	angular cut	10.0	10.0	10.0
	N_{LS}	list size	200	200	200
	N_{HD}	history depth	100	100	500
	c_2	contrast	1.0	1.0	1.0
	t_{TO} (s)	timeout	1.0	1.0	1.0
	$TS_{tot,2}$	threshold	1000.0	1000.0	150.0

Table 4.1: **At the moment, the GBM-seeded window (“GBM”) has a reasonable threshold for detecting GRBs onboard, while two LAT-only windows (“Win 0” and “Win 1”) have prohibitively high thresholds;** however, this is currently being changed. The parameters included here are the ones that are most relevant for onboard algorithm optimization studies and are not an exhaustive list.

threshold for the GBM-seeded search was changed in May of 2010 to address this; I describe this later in the chapter.

4.2.4 Localization and Updates

When the algorithm successfully triggers on a cluster, it begins iteratively recalculating the cluster localization. It starts with the events that were included in the Tier 2 TS_{tot} calculation and continues to add events within a radius θ_{loc} of the current localization. The calculation for each event is weighted by the distance between that event and seed event (for the very first iteration) or the most recent localization (for all other iterations); events that are closer to the localization are

weighted more heavily. The localization formulas are:

$$X = \sum_{i=1}^N w_i^2 \text{dircos}(x_i) \quad (4.4)$$

$$Y = \sum_{i=1}^N w_i^2 \text{dircos}(y_i) \quad (4.5)$$

$$Z = \sum_{i=1}^N w_i^2 \text{dircos}(z_i) \quad (4.6)$$

$$w_i = \frac{1}{\theta_i} \quad (4.7)$$

$$\alpha_{loc} = \arctan\left(\frac{\sqrt{X^2 + Y^2}}{Z}\right) \quad (4.8)$$

$$\delta_{loc} = \arctan\left(\frac{X}{Y}\right) \quad (4.9)$$

$$err_{loc} = \frac{1}{\sqrt{\sum_i w_i^2}}, \quad (4.10)$$

where $(\alpha_{loc}, \delta_{loc})$ is the current localization, err_{loc} is the uncertainty assigned to the localization, and $\text{dircos}(x_i)$, $\text{dircos}(y_i)$, and $\text{dircos}(z_i)$ are the direction cosines of the i th event in the x, y, and z spacecraft directions. θ_i is the distance between the i th event and the current localization; it is set to a minimum of 1° to avoid giving undue weight to events that are coincidentally close to the current localization. (Based on prelaunch simulations, there is a range of values for the minimum of θ_i that produce reasonable localizations, and 1 degree was chosen from this range.) During the first iteration, w_i is set to 1, so the first localization is always the average of all the events.

The algorithm remains in this stage until a predetermined amount of time has passed (currently set to 10 minutes) and recalculates localizations as appropriate. It sends a series of automated messages to the ground at predetermined times (six total), each with the latest localization.

4.3 Onboard algorithm studies

4.3.1 Improving sensitivity

To optimize the LAT's onboard GRB detection capabilities, we sought to simultaneously maximize the sensitivity to real GRBs while minimizing the rate of false triggers. In addition, not only did we require that the algorithm be able to trigger on the GRB, but that it also be able to localize the GRB accurately enough for follow-up observations. Thus, in studies to improve the algorithm sensitivity, we are necessarily checking the algorithm localization abilities at the same time.

We first studied the effect of changing various algorithm parameters of the LAT-only windows. The parameters we focused on were the Tier 1 and Tier 2 thresholds, as these are most directly related to whether or not a cluster successfully triggers the algorithm.

We began by changing the Tier 1 threshold while keeping the Tier 2 threshold low so that it did not affect the results (since any cluster that passed Tier 1 would automatically pass Tier 2). Figure 4.2 shows the offsets between onboard localizations and the actual GRB positions across different values of Tier 1 threshold, for the first six LAT-detected GRBs that could theoretically have been detected onboard. Two of these are short (GRBs 081024B and 090510), and GRB 090510 was bright enough to bypass the exceedingly high Tier 2 threshold for the LAT-only algorithm; the rest are long. The different circle sizes represent different Tier 1 thresholds, with 80 corresponding to the largest circle and 100 corresponding to the smallest. For diagnostic purposes, we changed the time of the closeout message from 600 s to 60 s after trigger; all other parameters are unchanged from their defaults.

In general, the offsets decreased with later update messages, which is consistent with the fact that later messages include more burst photons. Prior to launch, it was determined that the second message (first update) generally had the most accurate

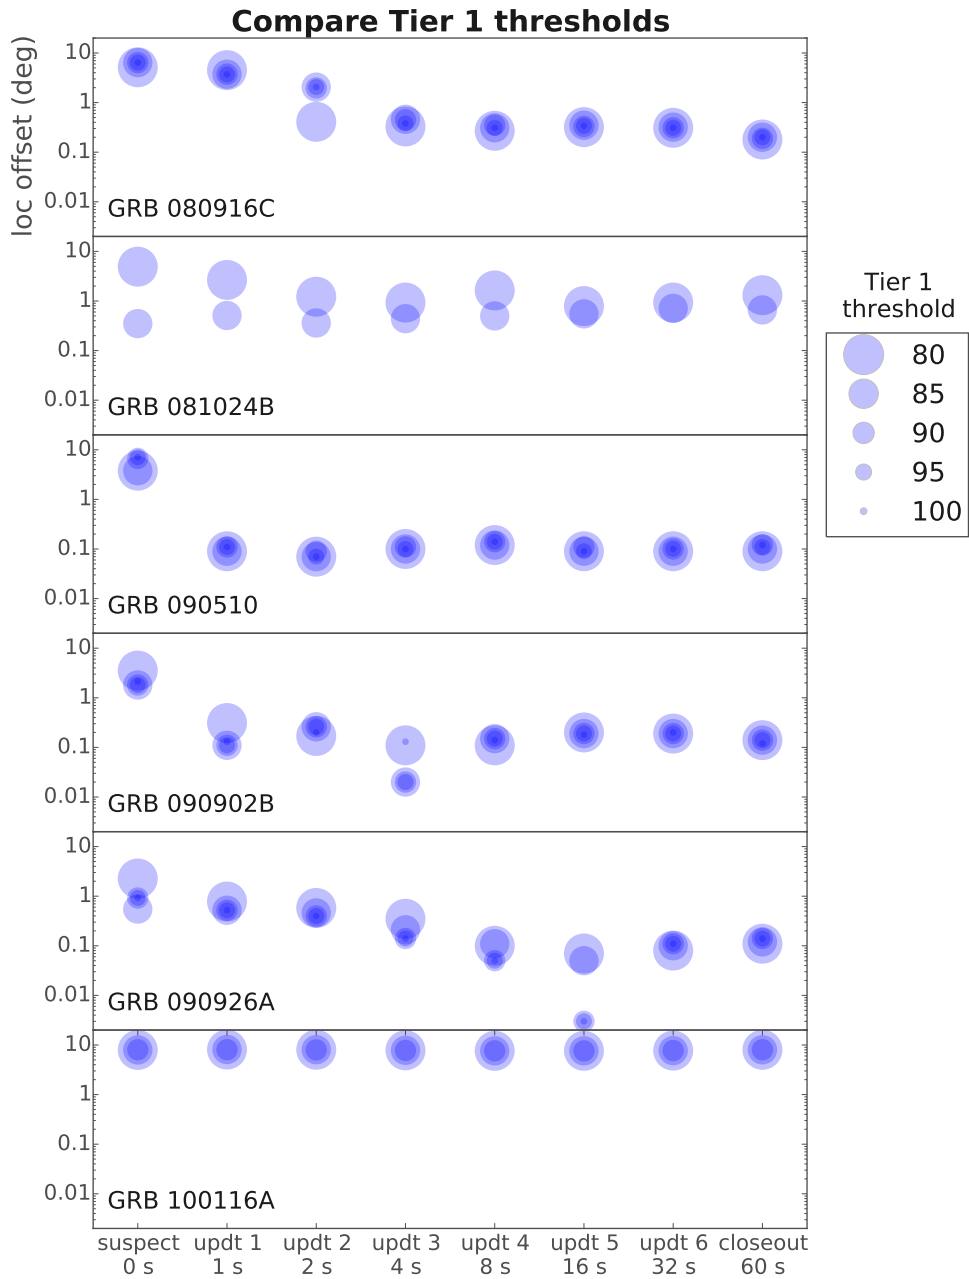


Figure 4.2: **The offset between the onboard localization and the actual GRB position for different values of Tier 1 threshold, for the first six LAT-detected GRBs that could have been localized onboard. The different circles represent different Tier 1 thresholds, from 80 (the largest circle) to 100 (the smallest circle).**

localization, or at least a sufficiently accurate localization. It turns out that this is not necessarily true (see, e.g., GRB 090926A), likely because the prompt GeV emission is often delayed with respect to the keV-MeV emission. I discuss further this issue of which message to use in the next subsection.

We did not find a systematic difference in localization accuracy between the different Tier 1 threshold settings, although there were a few cases for which the lowest or highest Tier 1 thresholds did not localize as accurately as the intermediate ones. We also did not find a systematic effect when we changed the Tier 2 threshold for a few choices of Tier 1 threshold. Therefore, for any reasonable combination of Tier 1 and Tier 2 thresholds, the only important distinguishing factor is the false trigger rate. We found that with a Tier 1 threshold of 80 to 100, any Tier 2 thresholds $\gtrsim 150$ had negligible false trigger rates during survey mode. However, there were still false triggers due to the Earth limb during pointed observations. This would not be a problem during ARRs since the onboard algorithm is disabled during these periods, but it would have to be accounted for during TOOs and other nonstandard observation modes.

We performed similar studies to improve the GBM-seeded window. When *Fermi* was launched, the threshold for the GBM-seeded window (the equivalent of the LAT-only Tier 2 threshold) was set to zero onboard, with a threshold of 120 applied on the ground before sending out notices to the public. This meant the GBM-seeded window would trigger as soon as the GBM detected a burst, but since LAT prompt emission is often delayed with respect to GBM emission (as previously discussed), there was very little chance of a trigger passing the ground threshold of 120. At the same time, the other windows were effectively disabled for the 10 minutes the algorithm remained in the elevated mode, further hampering the LAT onboard GRB detection capabilities.

We generated a set of 80,000 inputs (positions and times) across 8 different

data runs and ran them through the GBM-seeded window with different thresholds. We found that a threshold of 150 can be expected to have an acceptable false trigger rate of < 0.33 triggers per year. With this threshold, the algorithm successfully triggered on 5 of the 6 GRBs that were discussed in the “LAT-only” section (see, e.g., Figure 4.2). The exception was GRB 100116A; for this GRB, the high-energy emission was brightest over a minute after the GBM trigger, falling between the GBM message at 60 s and the one at 150 s (§ 4.2.3).

This threshold was put into effect in May of 2010 [39]. To date, it has successfully triggered on GRB 131108A (and had the potential to trigger on GRB 110731A), and has produced zero false triggers.

Out of the six bursts that were detectable onboard, four were detectable up to higher Tier 1 thresholds ($\gtrsim 120$) and with sufficiently accurate localizations. These four (GRBs 080916C, 090510, 090902B, and 090926A) make up the LAT “Fantastic Four,” and are four of the brightest GRBs detected by the LAT with especially bright prompt emission [77]. The majority of LAT GRBs are not as high flux as these and would not be easily detectable using these thresholds; in order to detect more GRBs onboard, we would have to lower the thresholds to a level that would lead to large false trigger rates. To mitigate this, however, we could selectively lower the thresholds for the LAT-only windows for a short period of time after a GBM detection. This would only make us vulnerable to false detections for a short period of time.

We revisited all of the 63 LAT-detected bursts up to October of 2013 and ran them through the onboard algorithm with lower thresholds. Out of these 63, six were bright enough to be detectable onboard with the more stringent Tier 1 threshold of > 90 , like we previously found; these six are made of the four found in the earlier study, plus GRBs 110721A and 110731A, and were particularly well-localized as well (offsets of 0.1 to 0.4°). Four other GRBs (081006, 081024B, 100116A, and 130427A)

were detectable with lower thresholds (Tier 1 threshold = 50 to 80). In general, a Tier 1 threshold of 50 did not include enough burst photons and led to some particularly large offsets while some dimmer bursts were lost with thresholds of 70 or 80, so a Tier 1 threshold of 60 is a good compromise in this regard.

We then tested these thresholds on a set of 520 runs from 14 September 2009 to 21 October 2009, excluding any runs that contained GRBs and/or Earth limb emission. We found that if the thresholds are dropped for 30 seconds after a GBM trigger, a Tier 1 threshold of 60 with a Tier 2 threshold of >100 would yield <0.1 false triggers per year.

We can further reduce the false trigger rate by prohibiting the LAT from triggering onboard when the spacecraft rocking angle is larger than a particular amount. *Fermi* normally rocks at $\pm 50^\circ$, but pointed observations can increase the rocking angle, which (to first order) would also increase the size of the Earth limb in the field of view. We are currently in the process of implementing this rocking angle cut on the ground (before dispersing the information to the public) to increase the LAT's ability to detect GRBs onboard.

4.3.2 Localization studies

There is no single message out of the eight that consistently provides the most accurate localization for a GRB (see, e.g., Figure 4.2), since this is often affected by the shape of the high-energy lightcurve. Besides the localization, each message also reports the number of photons used in the localization and the error (calculated according to Equation 4.10). We explored the possibility of using these pieces of information to determine *a priori* the most accurate localization.

It is important to note that the error calculation is necessarily simplistic due to the available computation resources onboard. The error doesn't have a direct physical meaning (e.g., it is not a confidence interval) but it can still give a sense

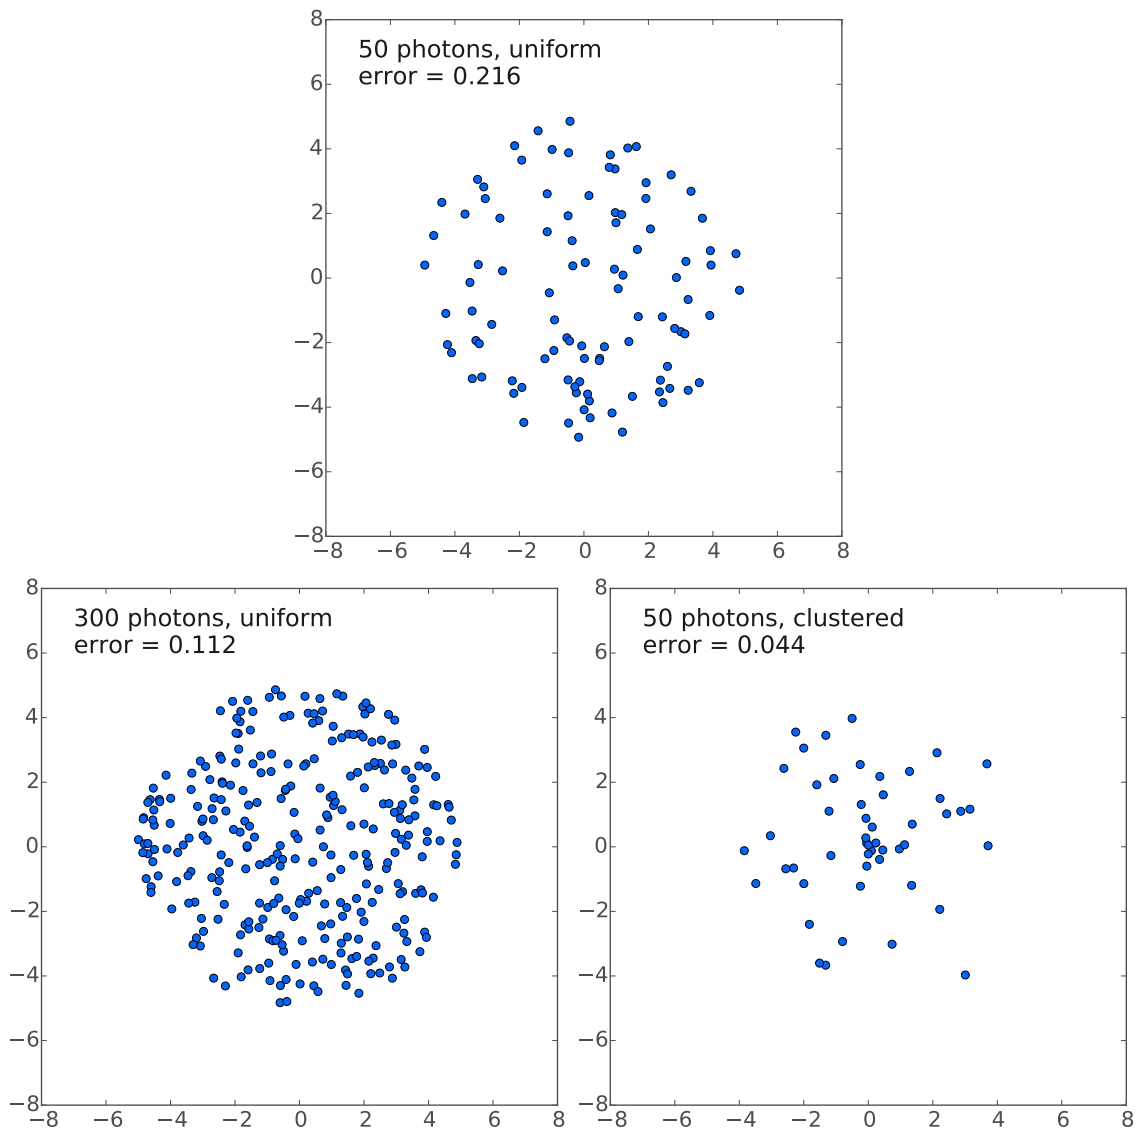


Figure 4.3: **The onboard localization error decreases either if the number of photons increases or the clusteredness increases.** This can be seen directly from the way the error is calculated (Eqn 4.10).

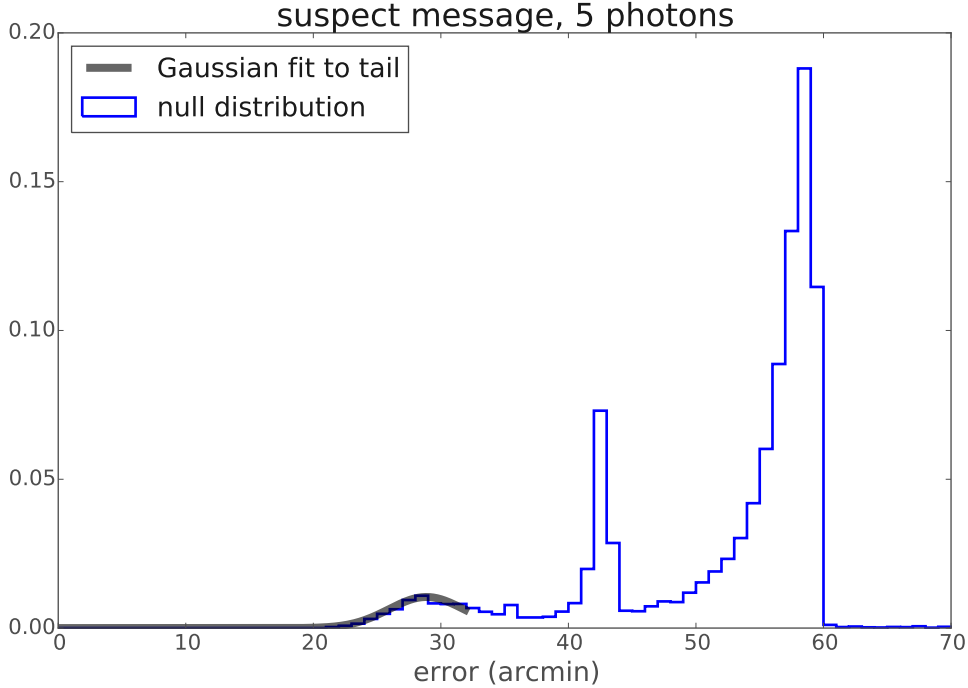


Figure 4.4: **The distribution of errors for the localizations that were calculated using 5 events, in the 'suspect' messages** (i.e., the first message that is sent out in the set of 8). There is a peak at around 60 arcmin due to the sets of events in which one event was within a degree of the input position, and a similar one at around 44 arcmin ($60/\sqrt{2}$) due to the sets in which two events were within a degree of the input position. We also fit a Gaussian to the tail of the distribution to calculate the probability for very small errors.

of how clustered a group of events is. Figure 4.3 shows that the error depends on both the number of events and the clusteredness of the photons; increasing either property decreases the error. Therefore, it is theoretically possible to use the number of events and the error to get a rough sense of how clustered a group of events is and, presumably, how accurately it represents the GRB's true position.

In order to get a sense of the error distribution for events that are purely background (i.e., that contain no signal from a GRB), we generated a set of about 600,000 sets of uniformly distributed seed times and positions. We ran these through the GBM-seeded trigger, and set the threshold to be zero so that the algorithm successfully triggered on each seed.

Figure 4.4 shows a sample distribution of errors; in particular, it corresponds

to the localizations that were calculated using 5 events and reported in the 'suspect' messages (i.e., the first message that is sent out of the set of 8). The algorithm sets a minimum θ of 1 degree or 60 arcminutes, so any event with $\theta < 1$ from the seed or input position gets reassigned to $\theta = 1$. The peak at around 60 arcminutes represents the clusters with 1 event within a degree of the input position, and the peak at around $60/\sqrt{2}$ arcminutes represents the clusters with 2 events within a degree of the input position. In this way, a localization with a particular number of events and error can be compared against the corresponding null distribution to get a sense of how unlikely this cluster of events could have come to background. We created null distributions for clusters of 3 to 599 photons, for the first seven messages (the closeout message arrives 600 seconds after trigger and is not suitable for localization purposes).

Generally, groups of events from real GRBs will be more tightly clustered and have errors that were too small to appear in our null distributions. To account for this, for each distribution, we fit a Gaussian to the leftmost feature of the distribution to approximate the shape of the distribution at small errors.

This feature has been implemented, and all GCNs of LAT onboard detections include a "LOC_QUALITY" value between 0 and 1 (essentially, one minus the probability of obtaining this error from the null distribution). Larger values of LOC_QUALITY indicate more tightly clustered events, and their corresponding localizations should in theory be more accurate.

We have been able to test this with GRB 131108A, a bright GBM burst that successfully triggered the LAT onboard algorithm. The localizations from the first three onboard GCN notices were over a degree away from the actual position (as later determined by *Swift*-XRT) while the later notices were only a quarter degree away. The LOC_QUALITY parameter reported in the first three notices were <1 , while the parameter in the later notices were all exactly 1.0. Because of this information,

we chose a position from the second set of localizations to communicate to *Swift*, which performed a tiled observation with a radius of $\sim 0.5^\circ$ and successfully detected the burst.

Chapter 5: GRB theory

Broadly speaking, a GRB occurs when a particular type of system undergoes a violent event and releases an ultrarelativistic jet. Interactions within the jet give rise to the bright and rapidly varying prompt emission, while the jet's subsequent interactions with the external medium cause the long-lived, fading afterglow emission. Due to both the energetics (up to 10^{51} or 10^{52} erg after accounting for beaming) and the short timescales (often less than tens of seconds and with subsecond variability), GRBs must be closely related to black holes or neutron stars, as these are the only objects that are compact enough to simultaneously account for such a massive energy release over a short time scale.

However, the details of GRB modeling are in flux and can change from year to year and conference to conference. A model will often naturally explain a particular observation or set of observations, but require cycles and epicycles to explain the rest. The difficulty lies in the extreme nature of the GRB central engine, which is opaque to electromagnetic observations. In addition, GRB observables can vary a lot from burst to burst, and GRBs have so far defeated all attempts to interpret them as standard candles.

Despite all this, a grand unified theory of GRBs is starting to settle, although many questions remain unanswered. In this chapter, I present a basic overview of GRB theory. There are many alternative models of GRB formation and emission, and I will not attempt to do them justice. Instead, I will only present the most commonly discussed models.

5.1 The physical system

5.1.1 Progenitors

Long GRBs have been observed in association with core-collapse supernovae, and are often found in the star-forming regions of galaxies. Because of this, the progenitors of long GRBs are well established to be very massive stars at the ends of their lifetimes. The associated supernovae are of Type Ib/c: the core collapses of massive stars that have shed their outer hydrogen (and sometimes helium) layers. The progenitors are thought to be Wolf-Rayet stars with initial masses of $\sim 40M_{\odot}$ (where M_{\odot} is the mass of the Sun). Wolf-Rayet stars start at $20M_{\odot}$, but the larger mass is necessary to account for the large energy releases of long GRBs (10^{50} - 10^{51} erg, after accounting for beaming), although confounding factors make it difficult to predict the progenitor mass exactly. The progenitor stars must be rapidly rotating (and possibly highly magnetized as well) to explain why the emission from a long GRB is variable and sustained instead of an instantaneous release of energy; without the rapid rotation, the core would just collapse and would not be able to sustain a jet. With this in mind, the progenitors of the most energetic GRBs must be low metallicity objects, to allow for both the large mass and rapid rotation [62].

Short GRBs are sometimes found a good distance from their host galaxies, consistent with having been given a “kick” at some point. This piece of circumstantial evidence, combined with their large energetics and short timescales, suggests that they come from the mergers of compact object binaries, such as NS-NS or NS-BH systems [63]. Unlike long GRBs, there have been no “smoking gun” observations of associated systems to directly support the binary coalescence scenario. However, recently, evidence for a rebrightening in the near-IR lightcurve of short GRB 130603B has been interpreted as a “kilonova” (also known as a “macronova” or “r-process supernova”), a transient event that is expected to occur in a binary

merger [64] [65].

5.1.2 Central engine

In both the collapsing star and binary merger scenarios, the central engine must be a compact object to explain the rapid variability found in both short and long GRBs. Crucially, GRB emission is not instantaneous but rather sustained, which requires some sort of quasistable central engine; the single, catastrophic act of collapsing or merging cannot explain all the features of GRB lightcurves.

The central engine is usually assumed to be a black hole powered by an accretion disk. The gravitational energy of the accretion disk is extracted and converted into a jet, whose axis is aligned with the rotation axis of the black hole and the original system [62]. (If the central engine initially emits a weak jet, material from the jet could fall back and accrete back onto the black hole, powering a second, stronger jet [66].) Alternatively, instead of being powered by the accretion of matter, the central engine could be powered by the accretion of magnetic flux (e.g., [67]). The nature of the accreted material is closely tied with the nature of the jet: the initial bulk Lorentz factor of the jet (i.e., how fast the jet moves as a whole), the duration of the GRB, whether the jet is matter- or Poynting flux-dominated, what the microphysical parameters within the jet are.

The central engine could also be a newly formed millisecond magnetar, a rapidly rotating and highly magnetized neutron star ($B \gtrsim 10^{14}$ G). In this scenario, the jet would be powered by the spindown of the magnetar. The magnetar is assumed to be hyper/supramassive and held up by its rapid rotation and large magnetic field, but would eventually collapse as it spins down (e.g., [68]).

In a short GRB, the objects in the progenitor system determine what the central engine is. If the progenitor binary contains a black hole, the central engine will necessarily be a black hole as well. However, if the progenitor system is a NS-NS




Figure 5.1: **Within the fireball model of GRBs, the prompt emission is caused by internal shocks within the jet while the afterglow emission is caused by the jet’s later interactions with circumburst material.** In the Poynting flux-dominated model, the overall GRB structure is the same, but the prompt emission instead comes from the reconnection of magnetic field lines within the jet. (Credit: NASA’s Goddard Space Flight Center)

binary, then the central engine could be a hypermassive magnetar depending on factors such as the masses and equations of state. For long GRBs, both central engine systems can theoretically form from the progenitor, again depending on factors such as mass, rotation, and magnetic field.

5.2 Prompt emission

Once the jet is formed, interactions within the jet accelerate charged particles, which then emit the photons that are observed as prompt emission. When discussing the jet, the main question is whether the energy of the jet is primarily carried by the matter or the magnetic fields. Both are expected to contribute, but the relative amount of energy carried by the one versus the other determines how the charged particles are accelerated. Most of the information in this section was gathered from Piran (1999) [69] (for the fireball model) and Lyutikov (2006) [70] (for the magnetic

field-dominated jet model).

In the fireball model, the bulk of the energy is carried by the matter in the jet (Figure 5.1). The “fireball” in the name refers to the fact that the total energy is much greater than the rest energy, that there is a very large amount of radiation in a very small space. Within the fireball model, the jet is mostly made of radiation and electron-positron pairs. There are likely a small number of baryons as well (known as “baryon loading”). The amount of baryons in the jet affects its evolution; the more baryonic matter present, the more energy is converted into the kinetic energy of the baryons. With too much baryon loading, the jet would never achieve the bulk Lorentz factors that are required based on observations ($\Gamma > 100$). Therefore, while the energy density in a jet is very high, the matter density must be very low.

The jet is optically thick when it is first launched; that is, the density is high enough that any photons created within the jet interact with other photons to produce electron-positron pairs before they can escape. As the jet propagates, it also slows and cools, until it becomes transparent to these photons. This transition surface is the photosphere, and it is expected to contribute to the prompt emission, or even precede it and manifest as a precursor peak in the lightcurve [71]. (In some models of GRB emission, it is the *primary* source of the prompt emission; see, e.g., [72] [73].)

Some of the energy in the bulk motion of the jet must be used to accelerate charged particles. Presumably, the central engine does not emit the jet in a perfectly uniform stream; instead, there are areas of high and low density along the length of the jet, which can be thought of as shells. These shells all travel with slightly different Lorentz factors, and if a faster shell is emitted after a slower shell, these two shells will eventually interact in the form of an internal shock. When the shells meet, they form an area of very high pressure, and any magnetic field lines in this area become highly condensed, thus creating very strong local magnetic fields. Every

time an electron or positron in the jet crosses this area of strong magnetic fields, it gets scattered and gains energy from this interaction. If it crosses multiple times, it gets boosted each time. This process is known as first-order Fermi acceleration, and results in a population of energetic electrons.

However, these shocks cannot excite the electrons to arbitrarily high energies. The same magnetic fields that accelerate the electrons can cause the electrons to lose energy via synchrotron emission. This occurs when an electron spirals around a magnetic field line, causing it to emit photons as it is radially accelerated. In this context, the Band function is often interpreted as an approximation to the synchrotron spectrum.

As we have seen, in the fireball model, magnetic fields play an important part but they are only required to be local (i.e., much smaller than the size of the jet). In contrast, the models of a Poynting flux-dominated jet begin with magnetic fields that are global and large-scale; for these jets, more energy is carried by the Poynting flux than the kinetic energy of the matter. Within this model, the particles are accelerated when magnetic field lines are split and later reconnect, depositing energy into the surrounding plasma and accelerating the electrons indirectly. The primary photon emission mechanism is still assumed to be synchrotron emission.

In either scenario, we are left with a population of hard x-ray / low energy gamma ray synchrotron photons as well as a population of electrons. If a lower energy synchrotron photon Compton scatters off of a higher energy electron, the photon is inverse Compton scattered into the high-energy gamma ray regime, in a process known as synchrotron self-Compton.

If the jet is sufficiently dense, high-energy emission involving hadrons is also possible. The same Fermi process that accelerated the electrons can also accelerate protons, which can scatter with the jet photons and lead to an electromagnetic cascade that produces high-energy photons. These cascades are also expected to

produce neutrinos [74].

5.3 Afterglow emission

As the jet propagates, it encounters the material that surrounded the progenitor system and/or material that was ejected in the earlier part of the GRB. At this point, for both the fireball and Poynting flux-dominated models, much of the energy is in the shock wave at the head of the jet. As the jet head collides with the circumburst medium, it launches external shocks: a forward shock downstream of the jet head, and a reverse shock back into the jet. (The reverse shock still moves forward, so to speak, but is slower than the jet itself, so within the frame of the jet it propagates backward.) Synchrotron emission is assumed to be the dominant emission mechanism for the afterglow as well, and has been successfully used to model the broadband afterglow emission from radio to X-ray observations. The main reference for this section is Zhang (2007) [59].

When the reverse shock crosses the head of the jet, it heats up the material and energizes the electrons. (The rest of the jet is not empty of material, of course, but the bulk of it is at the head.) This causes the electrons to emit a single flash of photons as they cool. The reverse shock is relatively weak, so the electrons are not accelerated as much as they would have been in an internal or forward shock; the photons they emit are at optical energies and lower. (This is a simplified scenario; the actual processes near the reverse shock can be much more complicated, as in [75].)

If the jet has a sufficiently high bulk Lorentz factor, the forward shock can emit high-energy gamma rays (hundreds of GeVs). As it slows, the typical energy of the synchrotron photons decreases, and the forward shock component to the afterglow is emitted in successively lower energy bands. Sources of afterglow emission are often simultaneous; for instance, the GeV emission from the forward shock can begin while the prompt emission from the internal interactions is still occurring [77].




Figure 5.2: **The canonical x-ray afterglow lightcurve is made of a multiply broken power law.** The bolded segments are seen in most x-ray afterglows, while the dotted segments (including the flares marked as section V) sometimes occur as well. [76]

The forward shock propagates through and heats up the cold circumburst material. As it does so, it produces a “canonical” x-ray afterglow lightcurve, as shown in Figure 5.2. In this figure, the bolded segments of the lightcurve are most commonly observed, while the dotted segments are only sometimes present. The lightcurve decays as a multiply broken power law, sometimes with flares or rebrightenings. The afterglow does not solely consist of x-ray emission, but the x-ray afterglow is particularly well-studied and can reveal crucial information about the GRB system.

After the prompt emission (Phase 0) ends, it is usually followed by a period of very steep decay of slope (Phase I). This is likely due to the curvature effect: The central engine activity that caused the prompt emission has ceased, but the emission from higher latitudes of the jet takes a while to reach the observer. Phase I is sometimes followed by a plateau phase (Phase II) which could be due to some underlying sub-dominant activity of the central engine, and sometimes accompanied

by x-ray flares (Phase V). After this time, the central engine has completely shut off, and the emission during this time is solely caused by the jet's interactions with the ambient medium (Phase III). Finally, in a fraction of afterglows, this is followed by a final Phase IV with a steeper decay. The break between Phases III and IV is interpreted as a jet break; i.e., as the jet slows and spreads, the dynamics of the jet are such that the observer sees the emission as decaying more rapidly than before. Initially, the jet moves relativistically and the emission is relativistically beamed within an angle smaller than the geometric opening angle; the observer only sees the flux contained within the relativistic beaming angle. As the jet slows, the relativistic beaming angle increases, but while it is smaller than the jet opening angle, the observer still sees the same flux. However, when the beaming angle becomes larger than the opening angle, the total observed light is spread out over a larger surface, and the observed flux drops. At around the same time this is occurring, the jet begins to laterally spread, which also causes the observed flux to drop. These effects combine to cause the jet break between Phases III and IV, and the timing of the jet break reveals the jet opening angle [60].

Chapter 6: The Record-breaking GRB 130427A

Ever since the EGRET era, GRBs have been known to emit high-energy photons. The LAT was programmed with features such as onboard triggering and ARRs to better explore this high-energy behavior of GRBs.

The LAT has detected emission from $\sim 6\%$ of GBM GRBs. Studies of LAT GRBs have indicated that they tend to be inherently more energetic than the average GRB, that they are in some way “extraordinary.” However, one of the brightest GRBs detected by the LAT is also one of the most ordinary: GRB 130427A, the “nearby ordinary monster” [85]. This GRB broke many observational records; at the same time, its inherent physical properties were somewhat average (Figure 6.1). As it turned out, it was merely an ordinary GRB that happened to be very close; a “monster in our own backyard.” (It certainly caused many members of the GRB community to wax poetic.)

GRB 130427A triggered both the GBM and *Swift*-BAT. It was by far the brightest GRB ever detected by the GBM, and was detectable by the LAT for a record-breaking ~ 20 hours. It was also well-observed in X-ray and optical photons by *Swift*, at optical wavebands by RAPTOR, and in hard X rays at late times by *NuSTAR*, as well as many other instruments across the electromagnetic spectrum.

In this chapter, I discuss the observations of GRB 130427A and their implications, with a focus on the high-energy emission. I begin with an overview of the LAT observations of GRBs to provide context. I then discuss the LAT observations of GRB 130427A, and the implications of the long-lived high-energy emission. I

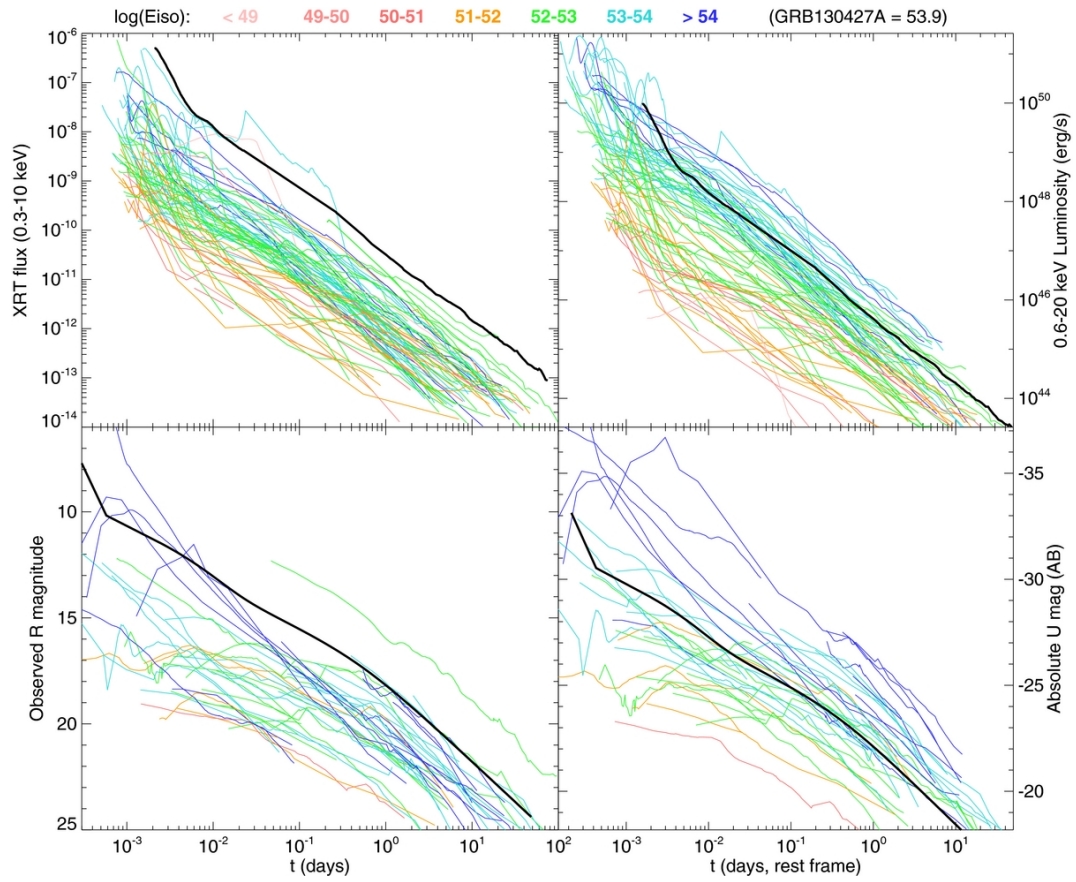


Figure 6.1: GRB 130427A was extraordinarily bright to observers (left two panels) but intrinsically ordinary (right two panels). Figure from [80].

also discuss the VERITAS (>100 GeV) upper limits for this GRB, and summarize observations from other instruments. The section on the LAT observations is drawn from Ackermann et al. (2014) [81], and the section on the VERITAS observations from Aliu et al. (2014) [82].

6.1 An overview of high-energy emission from GRBs

The LAT has detected high-energy emission from just over 100 GRBs so far (103 as of early September 2015). During this time, the GBM has detected about 1700 GRBs, and half of these have been in the LAT field of view, so that the LAT detects high-energy emission from roughly 1 in 8 or 9 GBM GRBs.

Out of the $\gtrsim 100$ LAT detections, roughly half were detected by the LAT solely as temporally extended emission. Some of these were not initially within the LAT field of view, which precluded the possibility of detecting prompt emission, but others genuinely were not detected as prompt emission despite being well observed. Out of the bursts that were detected as prompt emission, about half of these were were only detected with the LLE data type and not the standard analyses, usually because the burst was outside of the LAT standard field of view at the time of trigger.

When the LAT detects the prompt emission, the onset of the LAT emission is usually delayed with respect to the GBM emission (e.g., Figure 6.2). For most of these bursts, the LAT emission can only be modeled as a separate power-law component rather than an extension of, for instance, the high-energy power law of the Band function; in some cases, the power law even requires a high-energy cutoff (Figure 6.3). They tend to be very energetic and/or very close (Figure 6.4); in either case, their observed flux tends to be high. They often have bulk Lorentz factors of a few hundred (Figure 6.5), so that they come from very energetic jets.

While the keV-MeV emission is usually modeled as synchrotron emission from

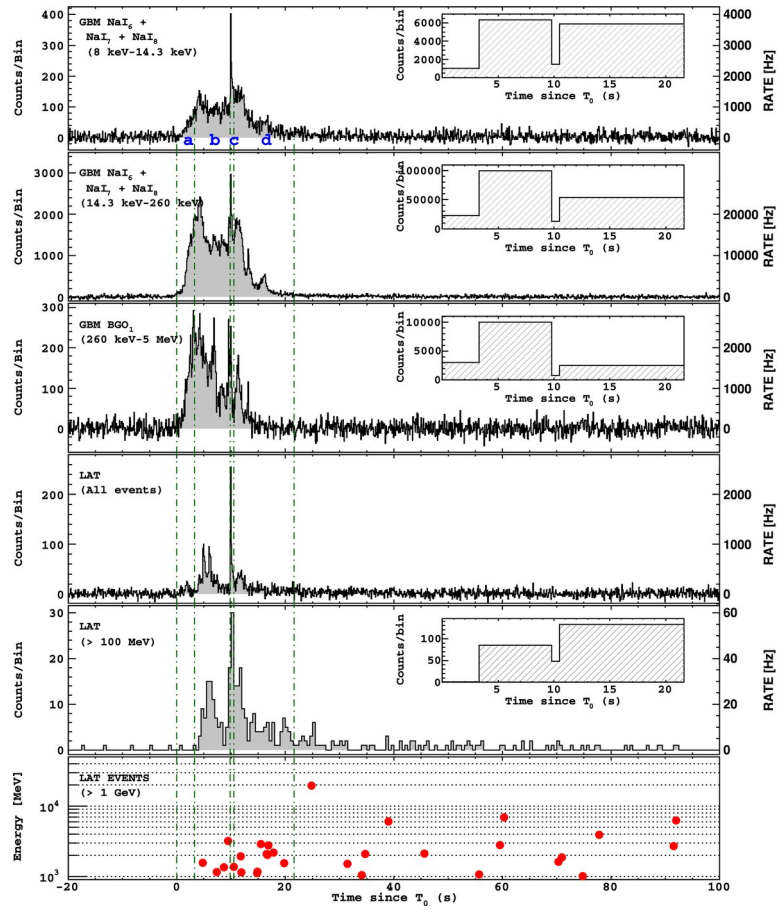


Figure 6.2: LAT-detected prompt emission is usually delayed with respect to the GBM-detected emission, as seen here in GRB 090926A. The LAT emission (bottom three panels) does not begin until a few seconds after the GBM trigger (the vertical red line) [83].




Figure 6.3: A time-resolved spectral analysis of a bright LAT GRB often requires an extra power law component to model the high-energy emission, as seen here in GRB 090926A. The four time intervals here correspond to the marked time intervals in Figure 6.2 [83].




Figure 6.4: LAT-detected GRBs tend to be very energetic and/or very close. [77]




Figure 6.5: LAT-detected GRBs often come from jets with bulk Lorentz factors of a few hundred. [77]

internal shocks (or an equivalent intra-jet mechanism), the mechanisms for >100 MeV emission are still unclear. The LAT-detected emission could be an extension of the synchrotron emission, or could come from inverse Compton emission. The LAT sometimes observes emission for over an hour, suggesting an external shock origin.

For further details, please see the first LAT GRB catalog [77].

6.2 LAT observations of GRB 130427A

6.2.1 Analysis

At 07:47:06.42 UTC on 27 April 2013 (T_0), while *Fermi* was in its regular survey mode, the GBM triggered on GRB 130427A. This burst was sufficiently hard and intense to initiate an ARR, as it had both a large peak flux and hardness ratio. At the time of the GRB trigger, the GRB was 47.3° from the LAT boresight. The ARR brought the burst to within the LAT field of view (FOV). It remained in the FOV for 715 s until it became occulted by the Earth, and re-emerged from Earth occultation at $T_0 + 3135$ s. After this, *Fermi* returned to its normal survey mode. The burst remained detectable for about 19 hours after T_0 .

The temporal profile of the emission from GRB 130427A varies strongly with energy from 10 keV to ~ 100 GeV (Figure 6.6). The GBM light curves consist of an initial precursor-like peak with a duration of a few seconds, a much brighter multi-peaked emission episode lasting ~ 10 s, and a dim, broad peak at $T_0 + 120$ s, which fades to an undetectable level after ~ 300 s. (This later peak is also seen in the *Swift* light curve [85].) The GBM peak between ~ 4.5 s and ~ 11.5 s was the brightest emission ever detected by the GBM; in fact, it was so bright that it caused problems for the instrument. At its brightest, the unprecedentedly high number of counts filled the TTE buffer before it could be read out, so the TTE data at the peak is saturated; this is evident by comparing the TTE and CTIME lightcurves (Figure 6.7). The

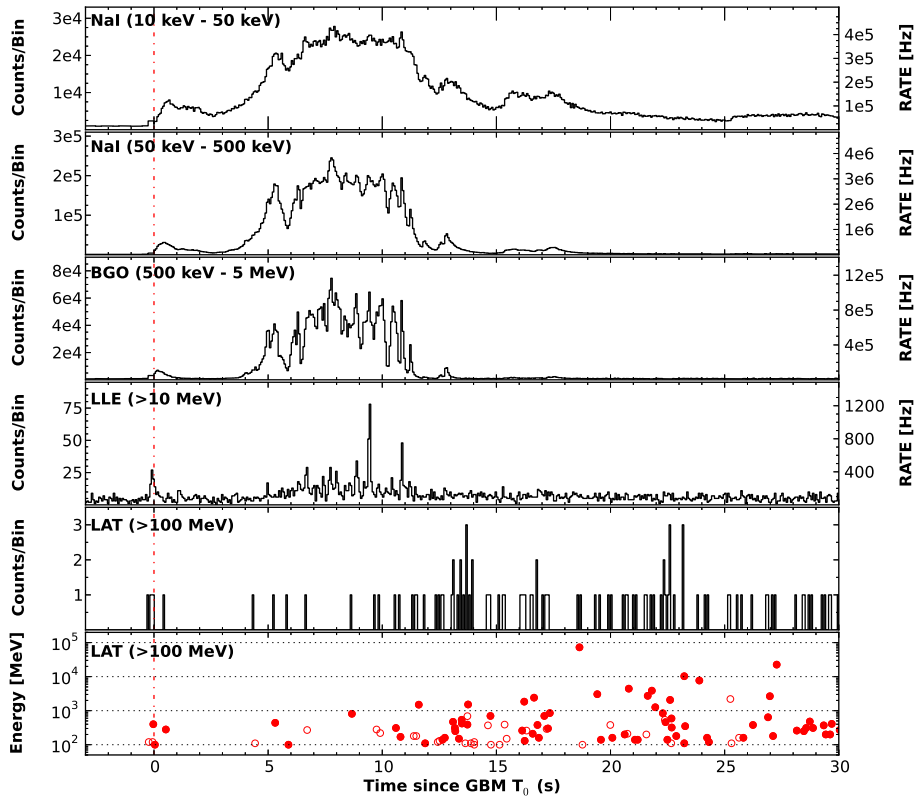


Figure 6.6: The prompt emission of GRB 130427A was recording breakingly bright in the GBM energy range (top three panels). The LLE >10 MeV emission (fourth panel) shows similarities with the GBM emission. The LAT >100 MeV emission, however, brightens *after* the lower energy emission had started to fade.




Figure 6.7: **During the brightest part of GRB 130427A, the GBM TTE data were saturated due to the high counts flux.** The buffer filled up faster than it could be read out. The binned data types (CTIME and CSPEC) were unaffected by saturation, although all data types suffered from pulse pileup.

binned CTIME and CSPEC data types were unaffected by saturation. However, all three data types were affected by pulse pileup: The photon rate was so high that the detectors did not have sufficient time to recover between detections, and multiple low energy photons were read as a single, higher energy photon [86]. Pulse pileup both decreases the count rate and shifts the spectrum to higher energies. This rendered the GBM data between ~ 4.5 s and ~ 11.5 s unusable, and this time interval was not included in any analysis.

The triggering pulse observed in the LLE (> 10 MeV) light curve is more sharply peaked than the GBM-detected emission at T_0 . The LLE light curve between $T_0 + 4$ s and $T_0 + 12$ s exhibits a multi-peaked structure. Some of these peaks have counterparts in the GBM energy range, although the emission episodes are not perfectly correlated; e.g., the sharp spike in the LLE light curve at $T_0 + 9.5$ s is not very bright in the GBM light curves.

The LAT-detected emission does not appear to be temporally correlated with either the LLE or GBM emission beyond the initial spike at T_0 . Photons with ener-

gies >1 GeV were observed ~ 10 s after T_0 , after the brightest GBM had faded. This is consistent with a delayed onset of the high-energy emission, although this particular case is different from other bright LAT GRBs in that there are no simultaneous peaks in the GBM and LAT (> 100 MeV) (compare with Figure 6.2). The ARR slew started at $T_0 + 33$ s, so the lack of LAT emission before 10 s is not due to the changing LAT exposure but rather reflects the true evolution of the GRB emission.

In order to understand this evolution, we performed a time-resolved joint spectral analysis on the *Fermi* data. We used the TTE data between 8 keV and 900 keV from NaI detectors 6, 9, and 10 (which had the smallest source angles and therefore the best exposures to the source), and the corresponding BGO detector 1 (which is positioned on the same side of the spacecraft as NaI detectors 6, 9, and 10). For the NaI detectors, we excluded the data from the energy bins near the K-edge [87]. We also used the LLE data above 30 MeV and LAT Pass 7 *Transient* class events above 100 MeV.

The spectral fits were performed using both RMFIT (version 4.3BA) and XSPEC (version 12.8), with consistent results. I report here the fits obtained with RMFIT by minimizing C-STAT. The models we used were a single power law (“POWL”) and a smoothly broken power law (“SBPL”). The POWL has the form

$$f_{\text{POWL}}(E) = A \left(\frac{E}{E_{\text{piv}}} \right)^\gamma, \quad (6.1)$$

where A is the normalization (in photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$) and E_{piv} is the pivot energy (used to normalize the model to the energy range in question; fixed at 100 keV). In our analyses, the POWL was only used as an extra component.

The SBPL model is parameterized as

$$f_{\text{SBPL}}(E) = A \left(\frac{E}{E_{\text{piv}}} \right)^b 10^{a-a_{\text{piv}}}, \text{ where} \quad (6.2)$$

$$a = m\Delta \ln \left(\frac{e^q + e^{-q}}{2} \right), \quad a_{\text{piv}} = m\Delta \ln \left(\frac{e^{q_{\text{piv}}} + e^{-q_{\text{piv}}}}{2} \right), \quad (6.3)$$

$$q = \frac{\log(E/E_b)}{\Delta}, \quad q_{\text{piv}} = \frac{\log(E_{\text{piv}}/E_b)}{\Delta}, \quad (6.4)$$

$$m = \frac{\lambda_2 - \lambda_1}{2}, \quad b = \frac{\lambda_2 + \lambda_1}{2}, \quad (6.5)$$

where E_b is the break energy (in keV), γ_1 and γ_2 are the low- and high-energy indices, respectively, and Δ is the break scale (in decades of energy) [87]. The Band function, a particular type of SBPL, is often used to model GRB spectra; however our time intervals were large and presumably contained a lot of spectral evolution, so we used the more general SBPL instead to avoid putting artificial constraints on our spectral fitting. We chose to use large time intervals because we were primarily concerned with the behavior of the GeV emission relative to the keV-MeV emission rather than the intricacies of the keV-MeV emission itself, and the LAT-detected emission was not bright enough to be divided into many small time bins.

We divided the prompt emission into three time bins: The initial peak ($T_0 - 0.1$ s to $T_0 + 4.5$ s), the brightest part ($T_0 + 4.5$ s to $T_0 + 11.5$ s), and the decaying part ($T_0 + 11.5$ s to $T_0 + 33.0$ s). We used the full data set (NaI + BGO + LLE + LAT) for the first and third time intervals but only used the LLE + LAT for the second time interval due to the pulse pileup issues mentioned previously.

During the initial peak, there are only a few LAT-detected photons, and the emission is well fit by the SBPL. For the brightest part, there are no photons with energies >1 GeV in this time interval, and the spectrum >30 MeV is well described by a single power law without a break. This slope of this power law is similar to the high-energy slope in the first time interval, and suggests that all the emission in this time interval could be described by a single component (Figure 6.8).




Figure 6.8: During the first two time bins (up to $T_0 + 11.5$ s), all of the GBM + LLE + LAT emission can be modeled by a single component, while a separate power law component is needed to fit the LAT emission after $T_0 + 11.5$ s. We did not use GBM data during the second time interval due to pulse pileup.

However, the spectrum of the decaying part (after $T_0 + 11.5$) is not well fit by a SBPL, and requires an extra POWL component to model the high-energy emission. Many photons with energies greater than 1 GeV were detected in this last time interval, including a 73 GeV photon at $T_0 + 19$ s. Unlike other bright LAT bursts, the LAT-detected emission from GRB 130427A appears to be temporally distinct from the GBM-detected emission, suggesting that the GeV and keV-MeV photons come from different emission regions or mechanisms.

The LAT detected photons from this GRB well after the GBM emission had faded below background. To characterize the temporally extended emission, we performed the standard unbinned maximum likelihood analysis on the LAT >100 MeV data. We modeled the LAT photon spectrum as a power law with a spectral index α (i.e., the spectrum $N(E) \propto E^\alpha$). We found that the LAT >100 MeV spectrum is well described by a power law at all times, but with a varying spectral index. For analysis of time intervals before the first occultation ($T_0 + 710$ s), we used *Transient* class events within 12° of the burst position. This was reduced to 8° for the last interval before the Earth occultation in order to reduce the background contamination from the Earth limb. After the first occultation, we used *Source* class events within 10° , which reflects the better PSF of the *Source* class.

In order to characterize the emission, we modeled the GRB as a point source at the burst position. For analysis with the *Transient* class events, we used the Background Estimator to model the background. For analysis with the *Source* class events, we used the publically available Isotropic and Galactic templates; the former includes the contributions from the residual charged particle backgrounds and the time-averaged celestial γ -ray emission, and the latter models the contribution from the Milky Way. We also included all the known LAT sources within the region of interest [88].

In all the time intervals we used, we found that the LAT >100 MeV spectrum

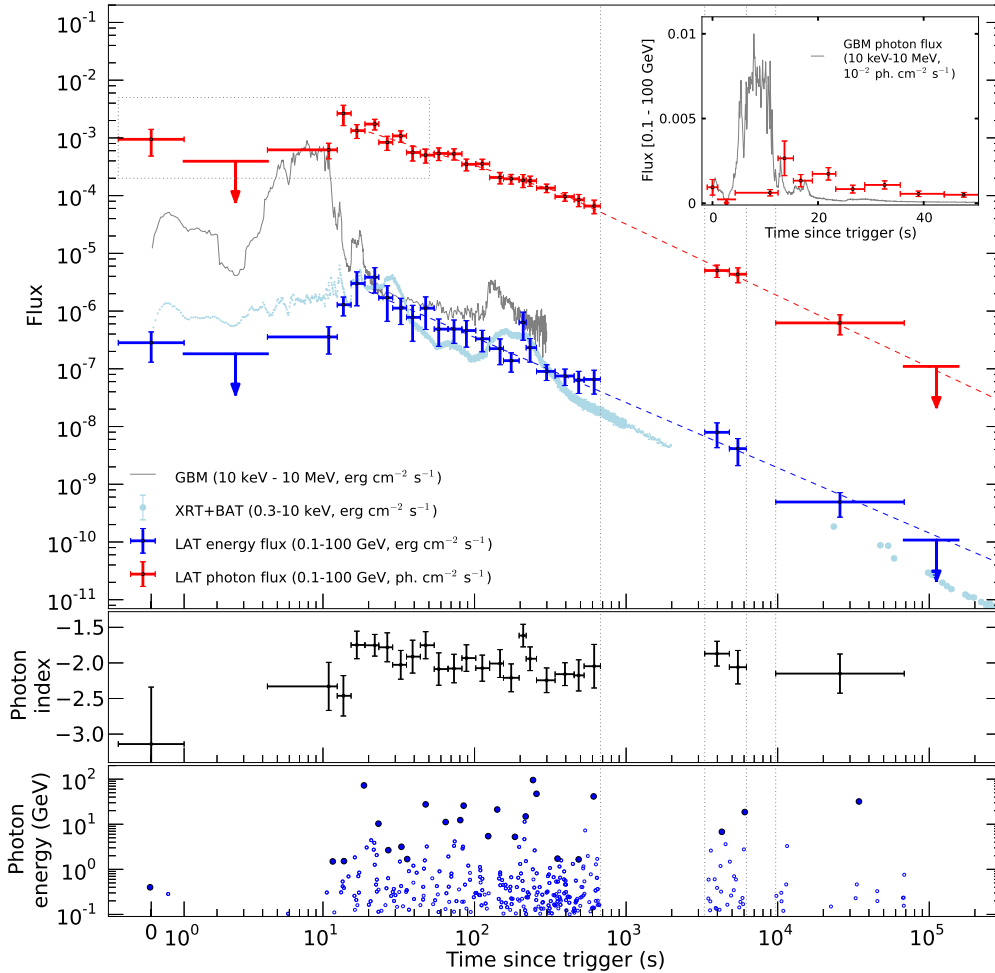


Figure 6.9: **The LAT detected emission from GRB 130427A for almost 20 hours, an order of magnitude longer than the previous record of 1.5 hours.** In the bottom panel, the larger circles mark the highest energy photons in each time interval.

is well described by a power law at all times, but with a varying spectral index. Figure 6.9 shows the behavior of the temporally extended LAT emission. During the earlier part of the GRB (the first couple of hundred seconds), the spectral index fluctuates between $\alpha \sim -2.5$ and $\alpha \sim -1.7$. The index then settles to $\alpha \sim -2$, consistent with the indices of other LAT bursts [77]. During the time intervals with the hardest spectra, the LAT recorded the highest energy photons.

The temporally extended photon flux light curve is better fit by a broken power law than a single power law. We found a break after a few hundred seconds, with the temporal index steepening from -0.85 ± 0.08 to -1.135 ± 0.08 ($\chi^2/\text{dof} = 36/19$ for a single power law, 16/17 for a broken power law). The energy light curve is well fit by a single power law with a temporal index of -1.17 ± 0.06 , consistent with other LAT bursts [77].

6.2.2 Interpretation

GRB 130427A is among the brightest LAT bursts. The 10 keV - 20 MeV fluence measured by the GBM in the 400 seconds following T_0 is $\sim 4.2 \times 10^{-3}$ erg cm $^{-2}$. We estimate the systematic error to be less than 20%, primarily due to the pulse pileup issue and uncertainties in the detector calibrations. (Compared to this, the statistical uncertainty is negligible.) The >100 MeV fluence measured by the LAT in the 100 ks following T_0 is $(7 \pm 1) \times 10^{-4}$ erg cm $^{-2}$. The total LAT fluence is therefore $\approx 20\%$ of the GBM fluence, similar to other bright LAT GRBs [77] [89]. Combining the GBM and LAT fluences gives a total 10 keV - 100 GeV fluence of $\approx 4.9 \times 10^{-3}$ erg cm $^{-2}$; at a redshift of $z = 0.34$, this translates to an isotropic energy release of $E_{\text{iso}} = 1.40 \times 10^{54}$ erg. This value is slightly less than but similar to the values for other bright LAT events, including GRBs 080916C, 090902B, and 090926A. The true energy release is only a fraction of this, since E_{iso} does not take beaming into account.

The emission region must be transparent against absorption by photon-photon pair production, which has a significant effect at energies >100 MeV. This requires the jet to be relativistic. The 73 GeV photon at $T_0 + 19$ s provides the most stringent limit on the bulk Lorentz factor γ . We performed a minimum variability analysis by using a Maximum Overlap Discrete Wavelet Transform, similar to what is described in [90]. During the brightest part of the burst, we found a minimum variability timescale of 0.04 ± 0.01 s in the LLE data (i.e., the shortest-duration features in the light curve have widths of 0.08 ± 0.02 s), which sets a constraint on the size of the emission region. If we assume that the MeV and GeV emissions around the time of the 73 GeV photon at $T_0 + 19$ s come from the same region, this implies a minimum bulk Lorentz factor of $\Gamma_{min} = 455^{+16}_{-13}$. (See [81] for calculation details.) However, this entire calculation relies on the assumption that the MeV and GeV emission are cospatial, whereas the joint spectral analysis suggested that the MeV and GeV emission come from different sources, so this calculation should be taken with a grain of salt. In any case, we note that it is consistent with values of Γ_{min} that were calculated for other bright LAT GRBs [77].

From the spectral analysis (Figure 6.8), we see that the LAT emission becomes harder and more intense after the GBM-detected emission has faded, and is better modeled as a distinct spectral component. These suggest that the GeV emission is produced later than the keV-MeV emission and in a different region. If the keV-MeV emission is the standard prompt emission and comes from interactions within the outflow (e.g., internal shocks), then the GeV emission comes from the external shock that occurs after when the outflow interacts with the circumburst medium. The reverse shock cannot heat material enough to emit photons at GeV energies, so the forward shock must be the source of the LAT-detected photons.

The circumburst medium could have a uniform density n_0 or fall off as $n(r) \propto r^{-2}$ like what a wind would produce; realistically, the density profile roughly falls




Figure 6.10: **The high-energy late-time photons detected by the LAT are difficult to explain using a standard synchrotron forward shock model.** The red curves indicate a circumburst medium with a uniform density, while the blue curves indicate a wind-like ($\propto R^{-2}$) density. The two dotted lines mark the scenario in which the acceleration is extremely fast. The solid and dashed lines correspond to scenarios with $\Gamma_0 = 1000$ for an adiabatic and radiative blastwave, respectively. The blue dot-dashed and double dot-dashed lines represent an adiabatic blastwave with $\Gamma_0 = 500$ and $\Gamma_0 = 2000$.

somewhere in between these two possibilities (ignoring small-scale inhomogeneities). The uniform medium would slow down the jet more rapidly than would the wind-like medium, limiting the energy of the photons produced by the jet.

GRB emission in the GBM energy range is well described by synchrotron radiation (§ 5.2), but the emission mechanism in the LAT energy range is uncertain. If the LAT-detected photons are also produced by synchrotron emission, the temporal behavior of the maximum possible synchrotron photon energy can be calculated by equating the radiation and acceleration timescales (assuming that Fermi acceleration is the mechanism for the acceleration). When we did so, we found that the highest energy photons could not easily be explained by synchrotron emission, even

with extreme assumptions.

Figure 6.10 shows the highest-energy photons detected by the LAT plotted against curves of maximum synchrotron photon energy. The parameters we varied were the circumburst medium density (either uniform or wind-like), the bulk Lorentz factor when the blastwave first interacts with the circumburst medium Γ_0 , and the type of blastwave, either adiabatic (energy losses caused by electron cooling are negligible) or radiative (energy losses must be taken into account). We also explore the scenario of extremely fast acceleration (the acceleration timescale is shorter than the Larmor timescale for a gyration by a factor of 2π). In the figure, the red/blue curves correspond to a uniform/wind-like circumburst medium. The solid and dashed lines correspond to scenarios with $\Gamma_0 = 1000$ for an adiabatic and radiative blastwave, respectively. The dot-dashed and double dot-dashed lines (both blue) represent the adiabatic case with a lower ($\Gamma_0 = 500$) and higher ($\Gamma_0 = 2000$) value of Γ_0 . The two dotted lines show the scenario of extremely fast acceleration; even these have trouble explaining some of the highest energy photons. As expected, the maximum synchrotron photon energy is lower in a uniform than a wind-like medium scenario, and the radiative blastwave curves drop more rapidly than the adiabatic curves. (These calculations all assume that the acceleration mechanism is Fermi acceleration. It is always possible to invoke a process faster than Fermi acceleration, such as magnetic reconnection (e.g., [91]). We did not explore this scenario.)

Synchrotron self-Compton (SSC) emission is often invoked as a possible source of the GeV emission. The ingredients for SSC emission are all present: a population of target photons and a population of energetic electrons, all within a small space. SSC emission is expected to peak at TeV and higher energies during the prompt phase (although no GRB, including 130427A, has been detected at TeV energies; see next section) and then pass through the lower energy bands. It would cause the

GeV light curve to flatten and the LAT spectrum to harden when it passes through the LAT energy range [92] [93] [94]. Such a feature can be seen in the light curves of GRBs 090902B and 090926A at 15 s - 30 s after T_0 [77], but no such hardening or plateau can be found in the LAT light curve in 130427A. (Extreme parameters might still allow an SSC interpretation; see, eg., [95].) There is very little variation in the photon or energy flux light curves (Figure 6.9), and none that are indicative of an SSC flattening or hardening. The LAT spectrum does harden at around 250 s, but this behavior lags behind the brightening in the GBM and *Swift* light curves. Overall, while SSC emission is not ruled out, it is also not indicated.

Inverse Compton emission is also possible with a different, non-synchrotron population of target photons. However, any source of target photons is difficult to maintain over the long period of time during which the LAT emission is observed.

The late-time high-energy emission could also be caused by electromagnetic cascades. These can occur when the blastwave is opaque to ultra high energy photons (UHE, >100 TeV), and would produce a photon index ~ -2 , consistent with the photon index of the late-time LAT observations. An electromagnetic cascade induced by ultra-relativistic hadrons would be accompanied by a neutrino flux [74], but no neutrinos were detected by IceCube [96].

Because GRB 130427A was both moderately energetic and extremely close, the high-energy observations were thorough enough that they challenged the standard afterglow models. This could suggest that other mechanisms are at work, or that more sophisticated/complicated models are required. In any case, the observations of this burst at other wavelengths can help reveal other properties of this GRB.

6.3 VERITAS observations of GRB 130427A

Observations of very high energy photons (VHE, >100 GeV) can help constrain GRB emission models. Unfortunately, there have been no clear detections of GRBs

at VHE energies. GRB 130427A was an excellent candidate for a VHE detection, given its closeness and its brightness in the LAT energy band. The Very Energetic Radiation Imaging Telescope Array System (VERITAS), a VHE observatory, began searching for emission as soon as it was able, but did not detect any. However, the non-detection itself is interesting, and sets limits on the model of SSC emission from this GRB.

VERITAS is a collecting of four imaging atmospheric Cherenkov telescopes. It detects photons with energies between ~ 100 GeV and several tens of GeV, overlapping with the LAT. Since full array operations began in 2007, VERITAS has followed up on more than 100 GRBs to search for VHE emission [56]. At T_0 , GRB 130427A was at a favorable elevation of 52° for VERITAS. Unfortunately, at the time, the bright moonlight conditions (97% full and $\sim 30^\circ$ above the horizon) precluded observations. Since the GRB was extraordinarily bright in the LAT energy range, however, VERITAS observations were initiated the following night, 28 April 2013, at 03:32:35 UTC, ~ 71 ks after T_0 . (In general, VERITAS follow-up observations of GRBs are limited to three hours after trigger.) Observations continued on the following two nights, lasting ~ 2 and ~ 2.5 hours, respectively (see [82] for details). During the first VERITAS observation (71.0 to 75.0 ks), the GRB was in the LAT FOV from 72.1 to 73.4 ks and from 73.5 to 74.9 ks; the last >1 GeV photon detected by the LAT was observed at 68.4 ks.

The VERITAS data were analyzed using a standard VERITAS software package. Absorption by the Extragalactic Background Light (EBL) greatly modifies the emission at VHE energies, so the GRB was modeled as a soft-spectrum ($\frac{dN}{dE} \propto E^{-3.5}$), weak (5% Crab Nebula flux) point source. We found no evidence for >100 GeV emission from this GRB in any single-night observation or in all the nights combined. We calculated the VERITAS upper limits by extrapolating the LAT spectrum (a power law with a photon index of ~ -2 with no intrinsic cutoff).




Figure 6.11: **The VERITAS non-detection suggests a cutoff at around $\gtrsim 100$ GeV if the LAT observations are due to SSC emission.** Alternatively, if the LAT observations are due to synchrotron emission, the VERITAS upper limits can put constraints on the magnetic fields around the blastwave front. The LAT 1- σ error contour is extrapolated in time from the power law fit to the LAT data at late times. The VERITAS upper limits were calculated assuming a cutoff in the Klein-Nishina regime at different energies.

We analyzed the LAT data using the unbinned maximum likelihood method as described in §6.2. The last LAT time interval with a significant detection was between 10 ks and 70 ks after T_0 . The spectrum during this time is well-fit by a power law with a spectral index of -2.2 ± 0.2 . (We also fit the emission with a power law plus an exponential cutoff and with a broken power law to search for spectral curvature, but neither of these models was statistically preferred.)

The first VERITAS observation and the last significant LAT detection are not simultaneous. However, the late-time emission (>200 s) measured by the LAT shows no deviation from a power-law behavior in either time or energy (Figure 6.6). We extrapolated the LAT data and fit from the last time interval (10 ks to 70 ks) to the first VERITAS observing interval using the photon flux relation $dN/dt \approx t^{-1.35 \pm 0.08}$ (reported in the previous section) to create the joint VERITAS-LAT spectral energy distribution (SED).

The joint VERITAS-LAT SED is shown in Figure 6.11. The LAT $1-\sigma$ error contour to the fit is shown in gray. The VERITAS upper limits (the three boxes) are compatible with an extrapolation of the LAT measurement, but they disfavor a scenario in which there is an enhanced VHE component. The SSC models plotted here are taken from the slow-cooling scenario described in [93] and fit to the LAT emission.

If the high-energy observations are caused by synchrotron emission, then the VERITAS upper limits can place constraints on the relative magnetic field strengths around the blastwave front. As discussed in the previous section, there is a maximum possible synchrotron photon energy. If we assume that the VERITAS non-detection is due to this cutoff, then a conservative estimate is that the cutoff occurs around ~ 100 GeV. With some assumptions, it can be shown that the magnetic field strength immediately behind the shock front is $\gtrsim 200$ times the magnetic field strength of the shocked material in front of the shock front.

If we instead interpret the high-energy observations as inverse Compton emission, the VERITAS non-detection implies the existence of a break in the spectrum. The most likely source of this break is the Klein-Nishina cutoff: At high enough energies, the electrons can lose most of their energy in a single scattering event rather than multiple small-loss events, so that the photon spectrum drops sharply above a certain energy [97]. We explored a few spectral breaks around 100 GeV, and found that the VERITAS upper limits are incompatible with a break above ~ 120 GeV (Figure 6.11). Thus, within the inverse Compton scenario, the most plausible interpretation for the VHE non-detection is that there is a Klein-Nishina cutoff before the VHE range.

6.4 Multiwavelength observations

GRB 130427A was so energetic in the GBM energy range that the first peak was bright enough for a detailed time-resolved analysis. The Preece et al. study analyzed the first peak of the GBM emission within the context of the physically-motivated synchrotron shock model. They found that synchrotron emission could qualitatively explain the properties of the first pulse, but that quantitatively, the temporal and spectral behaviors were difficult to reconcile [84].

The Maselli et al. study interpreted the multiwavelength afterglow in the context of the standard afterglow emission model: forward shocks, synchrotron emission. They found that a single synchrotron spectrum modeled the optical and X-ray afterglows well but underestimated the GeV emission (Figure 6.12). They concluded that GRB 130427A was a perfectly “ordinary” GRB (as opposed to the subluminal ones that are usually observed at low redshifts), but that it was the closest “ordinary” GRB ever observed [85].

Three independent RAPTOR (Rapid Telescopes for Optical Response) telescopes observed an incredibly bright optical flash from this GRB, the second bright-




Figure 6.12: The multiwavelength afterglow as measured of GRB 130427A, as measured by *Swift* and other instruments, was mostly well-modeled by a single synchrotron component, although this model underpredicted the GeV emission. See [85] for details and time interval definitions.




Figure 6.13: **The bright optical flash is more consistent with reverse shock emission than forward shock emission.** The late-time optical afterglow detected by RAPTOR decays similarly to the late-time LAT-detected afterglow, suggesting a common origin to both [75].

est ever observed (aside from the “naked-eye burst” [98]). The optical flash reached 7th magnitude between about 10 and 20 seconds after T_0 . The Vestrand et al. study interpreted the optical flash at the beginning as being caused by the reverse shock, due to its brightness and rapid decay. However, they found that the optical emission after $T_0 + 100$ s decayed more slowly and in fact was similar to the decay of the LAT-detected emission, possibly suggesting a common reverse shock origin for the LAT and RAPTOR emission (Figure 6.13). They also found that the optical afterglow became bluer at late times, which is consistent with the forward shock component becoming dominant after the bright reverse shock component had faded [75].

NuSTAR (Nuclear Spectroscopic Telescope ARray) observed the late-time emission of this GRB in the hard X-ray regime, the first time a GRB afterglow

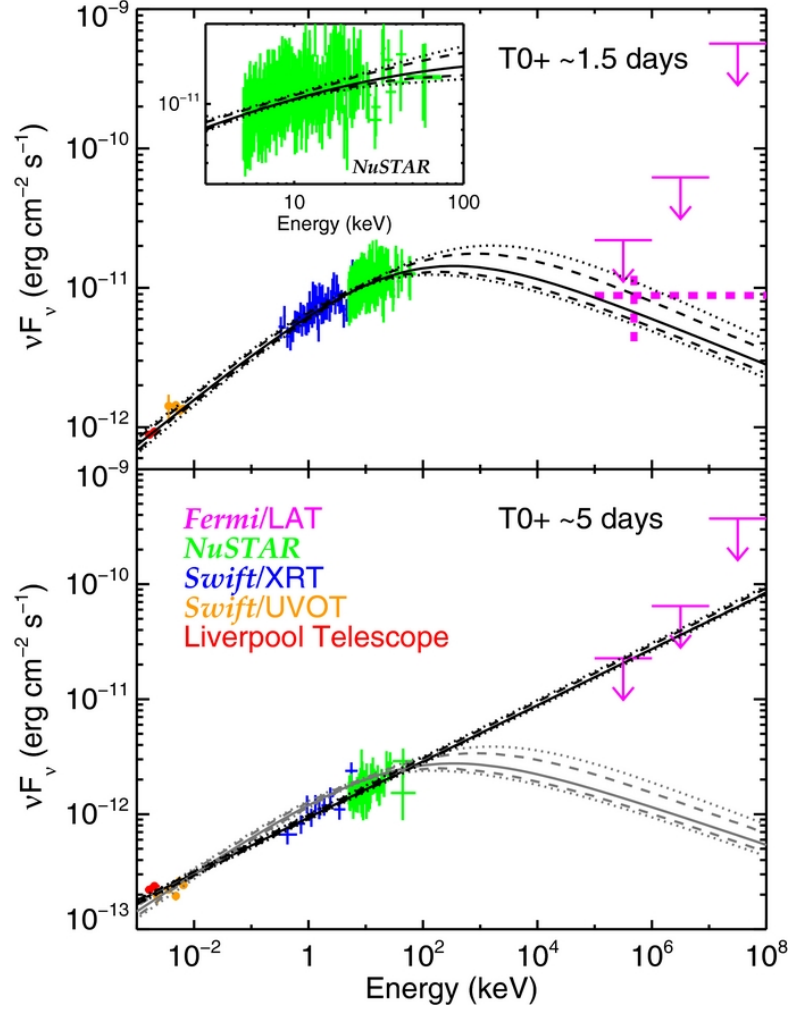


Figure 6.14: The Kouveliotou et al. study found that the multiwavelength observations of GRB 130427A at around 1.5 and 5 days after trigger are well-modeled by a single spectral component. See [99] for details.

had ever been detected at these energies. The Kouveliotou et al. study found that the joint optical, *Swift*-XRT and UVOT, and *Fermi*-LAT observations at ~ 1.5 and ~ 5 days is well modeled by a single component, which they interpreted as synchrotron emission (Figure 6.14). They calculated that the density of the circumburst medium falls as $E^{-1.4}$, between the uniform and wind-like density scenarios [99].

HAWC (the High Altitude Water Cherenkov Gamma-ray Observatory), with its unprecedentedly large field of view and excellent sensitivity and angular resolution, was well-placed to observe VHE emission from this GRB. Unfortunately, at the




Figure 6.15: **When GRB 130427A occurred, HAWC was not operating at full sensitivity** (solid black lines). If HAWC had had its full capabilities (green lines), it might have observed the first VHE photons ever detected from a GRB [100].

time the GRB occurred, HAWC was operating only 10% of its detectors and only with the scaler data acquisition system (which searches for statistical excesses in the photon rate) rather than the full shower reconstruction system. (The construction has since been completed.) HAWC was not able to detect >100 GeV emission from this GRB; however, it likely would have had the full array been operational (Figure 6.15) [100].

Chapter 7: GBM observations of precursor emission in GRBs

Precursors have been present in the field since the beginning. In 1974, a single year after the original discovery publication, observers using high-energy spectrometers on *Apollo 16* reported seeing a dim peak at the start of the lightcurve of GRB 720427 (Figure 7.1). They called it a “probable precursor,” on the grounds that it was $\sim 3\sigma$ above the average background level but presumably much dimmer than the other peaks. They pointed out that “this event [was] not initiated by its most explosive phase,” which is puzzling; if GRBs are caused by violent events, then to zeroth order, one would expect the start of the GRB to be dominated by larger energy releases.

As we have seen in the previous chapters, however, GRB emission is not a simple matter of a violent event causing a bright peak. GRB lightcurves are highly variable and structured, and the emission comes from complicated interactions within the jet. If the prompt emission is caused by a later, faster shell colliding with an earlier and slower one, precursors could simply occur when the Lorentz factors are very similar. On the other hand, nothing requires precursors to come from the same sort of physical jet processes that cause the brighter prompt emission peaks (the “dominant emission”), or from the jets at all. The first step to figuring out whether precursors are inherently different is to determine whether they are observationally different. This cannot be done by looking at single examples (although single bursts can certainly be illuminating).

In this work, we explored whether precursors come from a different source or




Figure 7.1: GRB 720427 was the first published GRB with a precursor. [104]

physical process than the dominant emission that follows them. We systematically searched for and characterized precursors in GRBs detected by the *Fermi*-GBM. The GBM is currently the most prolific detector of GRBs, and observes approximately one GRB every two or three days. Its wide energy range (8 keV to 40 MeV) and large set of available data (over 1700 GRBs) make it an excellent candidate for a population study.

In order to be as model-agnostic as possible, we began with a broad definition of “precursor” and subdivided the sample based on characteristics such as the delay time (the “quiescent period”) and the relative brightness of the precursor and dominant emission. We compared the different subtypes of precursors with each other and against the larger sample of all GBM GRBs to determine whether there were any systematic differences in spectral or temporal properties, which could indicate a unique physical origin for precursors.

I begin this chapter by discussing some of the work that has already been performed on precursors, as well as the most common theoretical models for precursor emission. I summarize the Bayesian blocks algorithm, which is the main tool we used to both detect and characterize precursors. I then discuss the specifics of the analysis — the subtype definitions, the models we used to fit the GRB spectra — and the results we obtained when we compared distributions of temporal and spectral parameters. Finally, I discuss what these results mean in the context of the precursor models. A table of the relevant data in this study is included in the appendix.

Note: In this chapter, instead of using the “date + letter” format that GRBs are officially named with, I am using the format of “date + fraction-of-day”.

7.1 Introduction

Previous population studies have found that precursors occur in a few to 20% of GRBs, depending on the instrument, the study’s definition of a precursor, and the method of finding precursors. The studies generally required that the precursor be dimmer than the dominant emission, and precede the dominant emission by a well-defined quiescent period during which the emission returned to a level comparable to background. Summaries of these studies are below:

- **Koshut et al. 1995** looked at the first 1000 BATSE GRBs using the DISCLA data from the LADs (20 keV to ~ 120 MeV in four energy bins with a temporal resolution of 1.024 s). In addition to the basic requirements (that precursors be dimmer and quiescent periods exist), they also required that the quiescent period durations be longer than the dominant emission durations. They searched the sample by eye, defined durations using signal-to-noise ratios, and quantified spectral properties using hardness ratios (i.e., the ratio of either the total counts or count rates between two different energy channels). They found precursors in only 3% of these GRBs. They concluded that there was no strong evidence that the dominant emission properties depended on the precursor properties, but also no strong evidence that the dominant emission and precursor are inherently different [105].
- **Lazzati 2005** limited their search to bright, long BATSE GRBs. They required that the precursors be *untriggered* — that is, that precursors be too dim to have tripped any detection algorithm — as they worked under the assumption that precursors should be weak, as some models predicted. They used a wavelet transform analysis to search for precursors, fit Gaussian profiles to the precursors to characterize their temporal properties, and approximated their spectral properties (when possible) using three or four low-resolution chan-

nels. They found precursors in as many as 20% of these GRBs; the precursors tended to be softer than the dominant emission, and there was a mild correlation between the precursor duration and both the dominant emission duration and the variability timescale. They noted that some quiescent periods were very long, up to hundreds of seconds [106].

- **Burlon et al. 2008** studied the first 105 *Swift*-BAT GRBs with redshift measurements, using only the basic requirements. They do not mention how they found the precursors (presumably they searched by eye); they used T90 measurements on both the precursor and dominant emission to characterize temporal properties, and performed standard BAT spectral analysis while using signal-to-noise ratios to define time intervals. They found precursors in 14% of their sample, and concluded that precursors were indistinguishable from the dominant emission due to their similar characteristics [107].
- **Burlon et al. 2009** performed a similar, time-resolved study of 2121 BATSE GRBs with fine time-resolved data (64-ms bins), using the same general requirements on precursors as Burlon et al. (2008) (but without requiring redshift measurements). They used T90 measurements to characterize temporal properties, and obtained spectral characteristics from a catalog [109]. They found that 12.6% of their sample had precursors. As in the Burlon et al. (2008) study, this study found no indication that precursors and dominant emission are distinct, regardless of the quiescent period duration [108].
- **Troja et al. 2010** focused on 49 short GRBs detected by *Swift*, including 11 that were characterized as “short bursts with extended emission” [110]. They used the basic definition and searched for precursors with a wavelet transform analysis, and used hardness ratios to characterize spectral properties. They found precursors in 8-10% of their sample, and concluded that precursors

tended to be softer than the dominant emission but that there was no discernable difference between the sets of GRBs with and without precursors [111].

- **Hu et al. 2014** searched through the first 613 BAT GRBs for multiple temporal features, including precursors, for which they did *not* require a well-defined quiescent period. Unlike previous studies, they used a Bayesian blocks algorithm (discussed in §7.3) with multiple lightcurve binnings. They characterized the temporal properties using the Bayesian blocks bins, and performed standard spectral analyses of BAT GRBs. They found 8% of their sample had precursors, but only one of 57 short GRBs did. They found that the spectra of precursors and dominant emission episodes were consistent with each other [112].
- **Charisi et al. 2015** searched through long GRBs from BATSE, *Swift*-BAT, and *Fermi*-GBM (a total of 2710 GRBs) using the basic precursor definition. They used a method developed for gravitational wave analysis to search for clusters of emission in time and energy. They found precursors in $\sim 10\%$ of their sample, and determined that the percentages were consistent between the three instruments. They found no correlations between precursor and main emission temporal properties (they did not study spectral properties) [113].

Despite their differences in instrument, precursor definition, and detection technique, most of the studies found no evidence that precursors are distinct from the dominant emission in their measured properties. This suggests that precursors and dominant emission episodes come from the same source, and that precursors are not physically unique. On the other hand, many of the studies found long quiescent periods (>100 seconds, in bursts where the dominant emission duration was $\ll 100$ seconds) that were difficult to explain using standard models of GRB emission.




Figure 7.2: **The properties that often are used to define precursors (dim and followed by a well-defined quiescent period) depend on the signal-to-noise ratio of the GRB.** *Left:* The presence or absence of a quiescent period is background-dependent, and studies that require well-defined quiescent periods exclude some GRBs with precursor-like first peaks. The example shown here is GRB 130427.324, a particularly bright GRB. *Right:* Some GRBs (such as GRB 090529.564 here) have initial peaks that technically satisfy the definition of a precursor but are not much dimmer than the dominant emission. With a lower signal to noise, the flux ratio between the precursor and dominant emission becomes smaller, the quiescent period becomes more well-defined, and the first peak becomes more “precursor-like”.

We selected our criteria for defining precursors based on these studies’ definitions, but with slight adjustments. Previous definitions of precursors did not fully take into account instrumental effects, such as the background-dominated nature of BATSE and GBM. For instance, quiescent periods are defined as intervals during which the rate returns to background, which strongly depends on the signal-to-noise ratio of a burst; the first peak of GRB 130427.324 (130427A) was much dimmer than the rest of the prompt emission, but because this GRB was so extraordinarily bright, the emission never completely returned to background levels in between (Figure 7.2, left). At the same time, some GRBs have initial peaks that technically satisfy the definition of a precursor but are not much dimmer than the dominant emission; these seem to follow the letter of the rule but not the spirit (compare Figure 7.2, right, with Figure 7.1). But if the signal-to-noise ratio were lower for these GRBs, the flux ratio between the initial and later peaks would be smaller and the quiescent period more defined, and the first peak becomes more “precursor-like.”

With this in mind, we began with a broad definition: That precursors are emission episodes that are dimmer than the rest of the emission (the “dominant emission”), and that precede the dominant emission. We did not require the existence of a well-quiescent period between the precursor and dominant emission. We then split the sample of precursors into three subsets, based on the relevant peak fluxes and the presence or absence of a quiescent period. We used the Bayesian blocks algorithm to both find precursors and characterize the temporal properties, and fit the spectra to models to characterize the spectral properties. We searched for correlations between the precursor and dominant emission properties within bursts and performed statistical tests to determine whether the precursor and dominant emission properties were consistent with each other. We also compared the properties across the different subtypes, and between the subtypes and the general population of GRBs.

7.2 Brief overview of precursor models

In this section, we begin with the assumption that precursors are genuinely “different” from the dominant emission; i.e., that whatever causes the precursor emission is distinct from what causes the dominant emission. If the dominant emission is emitted after the jet has become optically thin, then the precursor must be emitted before this point occurs during GRB formation. In this section, I start with the models for which the precursor emission site is closest to the dominant emission region and work backwards, which is roughly equivalent to beginning with the models that predict the shortest quiescent times. A more detailed discussion of these models is presented in §7.6.

- **Photospheric precursors:** The jet is optically thick when it is initially emitted, and becomes optically thin as it propagates and slows. When this transition occurs, the thermal photons that had previously been trapped within the jet are released within a short time as photospheric emission. If precursors are due to photospheric emission, then they are expected to occur very shortly before the dominant emission, and to have thermal spectra [71].
- **Shock breakout precursors:** For long GRBs, the central engine is surrounded by the stellar envelope of the progenitor star, and the jet must penetrate this envelope in order to be observable. The jet’s interactions with the envelope heat the material in front of the jet, and when this material breaks through, it could release thermal emission in the form of a shock breakout precursor. In addition, energetic particles within the jet can interact with these thermal photons and release a second, upscattered precursor. For short GRBs, a similar scenario could occur if the central engine releases a dense wind before emitting a jet; this requires that the central engine be a magnetar rather than a black hole [119].

- **Fallback precursors:** In this scenario, the first jet that the central engine emits penetrates the stellar envelope (or dense wind) but is too weak to propagate through; the material falls back and is accreted by the central engine, powering a second, stronger jet. The first jet produces the precursor while the later jet produces the dominant emission [66].
- **Multistage collapse precursors:** Depending on the parameters of the progenitor system, the core collapse or binary merger could produce a temporarily stable intermediate object before collapsing into the final black hole. The precursor comes from the initial collapse, and the dominant emission from the final collapse [123].
- **NS-NS magnetic field interactions:** For short GRBs, if the two progenitors are highly magnetized neutron stars (NS), the magnetospheres can interact when they come in close enough contact [124].
- **NS crust cracking:** For short GRBs, if at least one of the two merger objects is a NS, it could be tidally affected before the merger causing the crust to crack [125].

Some models (e.g., photospheric precursors) predict thermal emission, while others predict nonthermal emission. Some can only explain very short quiescent periods, while others can accommodate much longer ones. By quantifying the spectral and temporal properties of the GRBs with precursors, we can exclude some of these models of precursor emission.

7.3 Bayesian blocks

We used the Bayesian blocks algorithm developed by Jeff Scargle to both find precursors and define durations within GRBs. This algorithm determines the opti-

mal piecewise-constant representation of a light curve, within a Bayesian framework. That is, it searches for significant jumps in the light curve (called “change points”) and connects consecutive change points by flat lines, so that the lightcurve is approximated by a series of blocks of different heights. It makes no assumptions about pulse shapes or backgrounds. The relevant references for this section are Scargle (1998) [101] (for the mathematical underpinning) and Scargle et al. 2013 [102] (for the details of the algorithm itself).

Within the Bayesian approach to statistics, one starts from an observation and asks whether a particular model makes sense given the observation and any prior knowledge one has. If one were to flip a coin a hundred times and get “heads” 90 of those times, how sure can one be that the coin is fair and unweighted? The frequentist approach might be to flip a fair coin a large number of times and build a distribution of how often “heads” comes up (or, alternatively, to look up “binomial distribution” on Wikipedia), then see where the 90/100 result falls on this distribution. The Bayesian approach, on the other hand, would be to directly determine the confidence in saying that the coin is fair, based on the fact that heads came up 90 out of 100 times instead of the ~ 50 times one would normally expect, and the fact that one would generally expect coins to be fair.

Given a particular model M and a set of observations D , (one version of) Bayes’s theorem gives the probability of M being true given that D is observed as

$$P(M|D) = \frac{P(D|M)P(M)}{P(D)}. \quad (7.1)$$

Here, $P(D|M)$ is the probability of getting D given the model M (in the coin flip example, it is the probability of getting “heads” when using a fair coin); $P(M)$ is the prior probability of the model M (how often does one expect a coin to be completely fair?); and $P(D)$ is the probability of getting D across all models (how

often would “heads” come up 90 out of 100 times when considering coins with all possible weightings?).

$P(D)$ is a normalizing factor, and is the same for all models. In the coin flip example, $P(D)$ can be difficult to calculate. However, if we want to compare two models M_1 and M_2 to see which one is more likely to be true, $P(D)$ would drop out, and we are left with:

$$\frac{P(M_1|D)}{P(M_2|D)} = \frac{P(D|M_1)P(M_1)}{P(D|M_2)P(M_2)} \quad (7.2)$$

The algorithm works in the following way: It begins with one element of data (for instance, the very first event in the case of TTE data, or the number of events in the very first time bin in the case of binned data) and adds new elements one at a time. The first element can trivially be represented by a single block (i.e., a period of time over which the photon rate is constant). After each new addition, the algorithm decides whether the new element is a continuation of the previous block or signals the start of the next block, using Eqn. 7.2. In this way, the algorithm steps through the data one element at a time and builds up a piecewise-constant representation of the data.

The inputs to the Bayesian blocks algorithm are the data itself (either TTE or binned) and the prior ncp_prior . For TTE data, ncp_prior is explicitly related to the choice of a false positive rate, and is a way to adjust the complexity of the resultant Bayesian blocks lightcurves. Scargle et al. 2013 performed simulations on BATSE TTE data to find the value of ncp_prior which minimizes the number of false change points, and found that

$$ncp_prior = 4 - 73.53 p_0 N^{-0.478}, \quad (7.3)$$

where p_0 is the user-selected false positive rate and N is the total number of events

in the data. I found that choices of *ncp_prior* between ~ 5 and 10 did not yield substantially different light curves, and used *ncp_prior* = 6 for TTE data.

For binned data, *ncp_prior* does not have a straightforward connection to p_0 . Scargle et al. performed simulations based on a false positive rate of 0.05 and determined the optimal choices for *ncp_prior* based on, essentially, the average event rate in the lightcurve.

For further details, please see Scargle et al. 2013.

7.4 Systematic search of GBM precursors

We searched through all the bursts observed by the GBM up to the end of 2013, for a total of 1275 GRBs.¹ We used the data from the NaI detectors, and included only the two or three detectors with the best statistics / brightest signals; these are the detectors with the smallest viewing angles to the burst. By using multiple detectors for each burst, we had a check against spurious fluctuations which would show up in only one detector. We did not include more than three detectors; the detectors we used had the brightest signals, and if we couldn't find a significant precursor in them, we were unlikely to find one in the detectors with dimmer signals. We also checked the detectors we used against the ones selected in the GBM catalog² (if available) or the GBM internal team wiki (if not) to check that the detectors we chose had well-modeled responses (e.g., were not blocked by other parts of the spacecraft in a particular burst-spacecraft geometry).

We ran the raw NaI lightcurves through the Bayesian blocks algorithm, both individually and summed. We used both CTIME and TTE data for each burst; TTE has better time resolution (crucial for short GRBs) but for most of the GRBs

¹The latest GBM catalog encompasses the first four years and stops in the middle of 2012; we originally planned on matching the dates of this set, but — full disclosure — we wanted to include GRB 130427A in this sample.

²<http://heasarc.gsfc.nasa.gov/W3Browse/fermi/fermigbrst.html>

in the sample, it is only available from tens of seconds before a burst to hundreds of seconds after a burst. (Continuous TTE was implemented in December of 2012.) In addition, the Bayesian blocks algorithm treats each TTE event as a separate bin, so running it on TTE data for longer durations is computationally impractical. We used CTIME data between $(T_0 - 150, T_0 + 300)$ s (where T_0 is the GBM trigger time) and TTE data between $(T_0 - 20, T_0 + 20)$ s. For both data types, we extended the time interval on a case-by-case basis if there was any indication that a particular burst had emission beyond this time interval.

We inspected the Bayesian blocks light curves by eye to search for GRBs with precursors. In order to pass inspection, a precursor had to appear in both the summed lightcurves and at least two of the individual lightcurves of either data type. We then divided the precursors into three subtypes (illustrated in Figure 7.3):

- **Type I:** The precursor was much dimmer than the dominant emission and preceded it by a well-defined quiescent period. In practice, this meant that the background-subtracted precursor peak flux was less than approximately a third of the background-subtracted dominant emission peak flux, and that the Bayesian blocks lightcurve flux returned to approximately background levels in between. For this subtype, the precursor, dominant emission, and quiescent period durations are all well-defined.
- **Type II:** The precursor was much dimmer than the dominant emission but the quiescent period was not well-defined; i.e., the dominant emission began before the precursor had fully faded. We defined the precursor duration as the time between the start of the precursor and the start of the dominant emission, so that the precursor durations for the Type II GRBs are lower limits and not physically meaningful.
- **Type III:** The precursor was dimmer than the dominant emission but not

by much; i.e., the background-subtracted precursor peak flux was more than a third of but less than the background-subtracted dominant emission peak flux. The quiescent period was *not* required. In practice, any GRB for which the first peak was slightly dimmer than the brightest of the rest of the peaks was included in this category.

GRBs were allowed to fall into multiple categories and have multiple precursors. The basic definition used by previous studies encompasses the Type I and some of the Type III precursors. We separated out the Type I precursors in keeping with the spirit of the Lazzati study, which focused only on dim, weak precursors.

All of these definitions are heavily background-dependent. If the background level for a burst with a Type II precursor were raised slightly relative to the signal, the Type II could become a Type I. If the background level for a Type I or II were raised significantly relative to the signal, the precursor could disappear altogether. If the background level for a Type III were raised enough, it could turn into a Type I.

A large fraction of GBM GRBs do not have measured redshifts. The most common way to obtain a redshift requires a *Swift*-BAT detection, since the optical telescopes that can measure the redshifted spectral lines require much better localizations than the GBM itself can provide. Only a third of GBM GRBs are co-detected by *Swift*-BAT, and the GBM error regions are generally too large to be followed up by *Swift*-XRT. Recently, some GBM GRBs have also been successfully localized by the Intermediate Palomar Transient Factory, or iPTF (e.g., [103]), but the majority still come from *Swift*. Unfortunately, this means that we do not have redshift information for most of the bursts in the sample, which — among other things — makes it difficult to directly compare durations between bursts. Instead, when possible, we compared *relative* durations by scaling a redshift-dependent duration by another redshift-dependent duration.

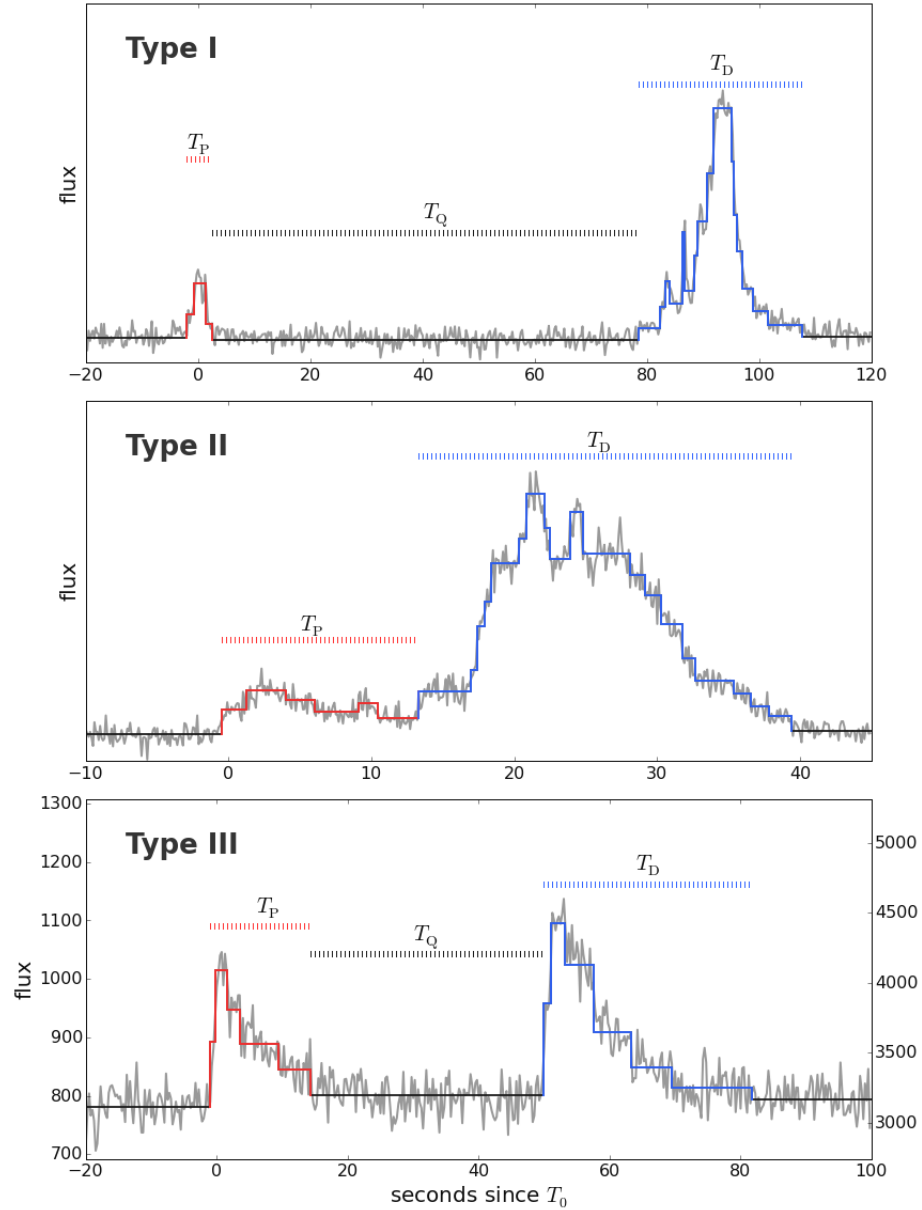


Figure 7.3: The precursors were divided into three subtypes based on the relative fluxes and the existence of the quiescent periods. These plots show examples of the three types; the precursor is marked in red and the dominant emission in blue. The examples here are GRBs 100116.897, 081221.681, and 120530.121.

We binned the data according to the Bayesian blocks time bins and fit the spectra in these time bins with standard nonthermal spectral models. When the dominant emission had multiple peaks, we only modeled the spectrum of the *first* peak to avoid smearing out any spectral evolution and biasing the fits. We used the standard techniques recommended by the GBM team, with the official software RMFIT version 4.3.2, and the standard energy selections: ~ 10 keV to ~ 900 keV for the NaIs (excluding the K-edge, ~ 25 keV to ~ 40 keV), and ~ 300 keV to ~ 38 MeV for the BGOs.

The spectral models we used were the Power Law (POWL), “Comptonized” model (COMP), and Band function (BAND). These are all standard non-thermal functions that are often used to model GRB spectra. They are empirical rather than physically motivated, and are not meant to be strictly interpreted in the context of any particular physical model.

The power law has two degrees of freedom — an amplitude A and a spectral index γ — and is defined as:

$$f_{\text{POWL}}(E) = A \left(\frac{E}{E_{\text{piv}}} \right)^\gamma \quad (7.4)$$

where the pivot energy $E_{\text{piv}} = 100$ keV normalizes the model to the energy range in question. The Comptonized function is a power law with an exponential high-energy cutoff; it has three degrees of freedom — amplitude A , spectral index α , the peak of the νF_ν spectrum E_{peak} — and is defined as:

$$f_{\text{COMP}}(E) = A \left(\frac{E}{E_{\text{piv}}} \right)^\alpha \exp \left[-\frac{(\alpha + 2)E}{E_{\text{peak}}} \right] \quad (7.5)$$

where E_{piv} is defined as before. The Band function is a smoothly broken power law whose curvature depends on the low- and high-energy spectral indices; it has four degrees of freedom — A , low-energy index α , high-energy index β , and E_{peak} — and

is defined as:

$$f_{\text{BAND}}(E) = A \begin{cases} \left(\frac{E}{E_{\text{piv}}}\right)^{\alpha} \exp\left[-\frac{(\alpha+2)E}{E_{\text{peak}}}\right], & E < \frac{(\alpha-\beta)E_{\text{peak}}}{\alpha+2} \\ \left(\frac{E}{E_{\text{piv}}}\right)^{\beta} \exp(\beta - \alpha) \left[\frac{(\alpha-\beta)E_{\text{peak}}}{E_{\text{piv}}(\alpha+2)}\right]^{\alpha-\beta}, & E \geq \frac{(\alpha-\beta)E_{\text{peak}}}{\alpha+2}. \end{cases} \quad (7.6)$$

In general, if a time interval is best fit by a power law or Comptonized function, this does not mean that the spectrum is *intrinsically* a power law or power law with an exponential cutoff. Rather, it means that there were not enough statistics to constrain the behavior at high energies. The Comptonized function is the Band function with $\beta \rightarrow -\infty$, and the power law is the Comptonized function with $E_{\text{peak}} \rightarrow \infty$. Roughly speaking, the Comptonized function and Band function α are equivalent and can be compared. However, the power law γ often measures something different; it is mostly used for dim emission episodes, so in many cases, rather than being a true measure of the low-energy index, it is affected by the turnover around what would have been E_{peak} and is softer than it would have otherwise been had the emission been brighter. This effect can be seen in Figure 7.4; the Comptonized and Band α distributions are consistent with each other, whereas the power law γ distribution is much softer.

For further details on spectral fitting of GBM bursts, see [114].

7.5 Results

Out of the 1275 bursts detected by the GBM up until the end of 2013, we found precursors in 202 of these, or 16%. However, the definitions of Type I and II precursors effectively exclude dimmer bursts, since the precursors (if they exist) of these bursts would be too dim to be detectable. This is essentially a cut on the peak *observed* photon flux, but the effect of the *fitted* photon flux can be seen in Figure 7.5. We defined a “high-flux” group whose energy fluxes are larger than the




Figure 7.4: The low-energy index distributions for the Comptonized function and Band function (and smoothly broken power law) are generally consistent with each other, while the power law index distribution is more negative and represents softer spectra. In general, the Comptonized and Band α measure the same thing (the true low-energy index) whereas the power law γ is affected by the spectral turnover around E_{peak} while not having enough statistics to constrain E_{peak} . This plot is from the second GBM spectral catalog [114].




Figure 7.5: **Precursors were found in GRBs across the whole range of T90 and peak flux.** The values for GRBs before GRB 12071146 were obtained from the second set of GBM catalogs [114] [115] and are plotted as filled circles, while the values for all other GRBs were taken from the online GBM burst catalog and are plotted as open circles. The approximate division between short and long GRBs (roughly 2 seconds) is marked by the vertical dashed line.

subtype	definition	# found
Type I	precursor peak flux <1/3 dominant emission peak flux; well-defined quiescent period	65 (5.1% / 6.5%)
Type II	precursor peak flux <1/3 dominant emission peak flux; no quiescent period	77 (6.0% / 7.7%)
Type III	precursor peak flux <dominant emission peak flux but >1/3 dominant emission peak flux; no requirement on quiescent period	60 (4.7% / 5.9%)

Table 7.1: **We found a total of 202 out of 1275 GRBs had precursors (16%).** These were split evenly between the three different subtypes. The peak fluxes in the definitions are the observed background-subtracted photon fluxes; the 1/3 ratio is approximate. The two numbers in parentheses are the percentage of GRBs with these precursor subtypes in the total sample (1275 GRBs) and the percentage in the “high-flux” sample (1006 GRBs).

dimmiest bursts for which we found a precursor; roughly speaking, this is the group of GRBs for which we *could* have found a precursor using our definition and detection technique. This is not a perfect approach, as the fitted photon flux depends on the spectral model used to fit the data; we would need to use simulations to more rigorously determine the observed flux limits of our search. With this caveat in mind, we are left with 1006 GRBs in the “high flux” group. Out of this group, we found that 20% had precursors.

Traditionally, the division between short and long GRB durations is 2 seconds [23]. This is not an absolute or exact definition, but our results do not substantially change if we move this division around by a second or two. Out of approximately 207 short GRBs in the full sample of 1275, we found that only 10 had precursors (5%), compared to 197 out of 1068 long GRBs (18%). In particular, we found that no short GRB had a Type II precursor, and that the GRBs with Type II precursors on average had longer T90s than the Type I or III precursors. This is possibly because the Type

II precursors were much harder to find in the shortened timescales since they would be “absorbed” into the dominant emission in the Bayesian Blocks lightcurve. We would have to do more studies to exclude the possibility that this has a real, physical basis; i.e., that short GRBs are genuinely less likely to have precursors. Since long and short GRBs have different progenitor systems, if a light curve feature appears in one type and not the other, this would suggest that precursors come from some aspect that is different in short and long GRBs (i.e., not the jet).

GRBs are known to undergo (sometimes rapid) spectral evolution, and the spectral analysis of long time intervals in GRBs can muddy true features and introduce spurious ones. However, a fully time-resolved analysis of all the GRBs in the sample is beyond the scope of this work. As a compromise, we modeled the spectra of individual peaks in a GRB rather than the entire duration. In particular, for each precursor, we compared the spectrum of the precursor to the spectrum of only the *first* peak of the dominant emission.

7.5.1 Type I

Out of the first 1275 GBM GRBs, we found that 65 GRBs had Type I precursors ($\sim 5\%$ of all GBM GRBs, and $\sim 6\%$ of high-flux GRBs). A few had multiple precursors, for a total of 86 Type I precursors.

We calculated the durations of the precursor (T_P), quiescent period (T_Q), and dominant emission (T_D) using the change points of the Bayesian blocks lightcurve (Figure 7.6). T_D and T_Q span similar ranges, while T_P is on average an order of magnitude smaller. In general, the precursors tended to be single peaks (we separated them into multiple precursors if there were multiple, well-defined peaks) while the dominant emission could be — and often was — multi-peaked, so it is not surprising that the T_D distribution has larger than values than the T_P distribution does.




Figure 7.6: **The dominant emission and quiescent period durations for Type I GRBs span a similar range, while precursor durations are an order of magnitude smaller, on average.**

We calculated the Pearson’s r (linearity) and Spearman’s ρ (monotonicity) correlation coefficients between the different durations. We did not find any strong correlations between T_P , T_Q , and T_D (Table 7.2). When we normalized the durations by the total duration ($T_{\text{tot}} = T_P + T_Q + T_D$), the signs of the correlation coefficients became negative, since T_P and T_D are anti-correlated with T_Q ; that is, for a given GRB, if the signal-to-noise ratio increases then T_P and T_D increase while T_Q decreases, and vice versa. This effect is particularly noticeable when comparing T_D/T_{tot} and T_Q/T_{tot} . (See Figure 7.7 for an example of durations plotted against each other.)

7.5.2 Type II

Out of the first 1275 GBM GRBs, we found that 77 GRBs ($\sim 6\%$ of all GRBs and $\sim 7\%$ of high-flux GRBs) had Type II precursors; including multiple precursors,

Type I	T_P	T_Q	T_D
T_P	—	$r = 0.24$ $\rho = 0.50$	$r = 0.17$ $\rho = 0.45$
T_Q	$r = 0.24$ $\rho = 0.45$	—	$r = 0.35$ $\rho = 0.65$
T_D	$r = 0.17$ $\rho = 0.45$	$r = 0.35$ $\rho = 0.65$	—

Type I	T_P/T_{tot}	T_Q/T_{tot}	T_D/T_{tot}
T_P/T_{tot}	—	$r = -0.20$ $\rho = -0.15$	$r = -0.42$ $\rho = -0.39$
T_Q/T_{tot}	$r = -0.20$ $\rho = -0.15$	—	$r = -0.80$ $\rho = -0.80$
T_D/T_{tot}	$r = -0.42$ $\rho = -0.39$	$r = -0.80$ $\rho = -0.80$	—

Table 7.2: **For GRBs with Type I precursors, we found no meaningful significant correlations between any duration measurements.** The top table lists correlation coefficients for the absolute values; the bottom table, the values normalized by the total durations. In these tables, r is Pearson’s r for linearity, and ρ is Spearman’s ρ for monotonicity. The negative correlation coefficients in the normalized table are due to the fact that higher signal-to-noise ratios mean T_P and T_D increase while T_Q decreases, and vice-versa. This effect is particularly noticeable when comparing T_Q/T_{tot} and T_D/T_{tot} .

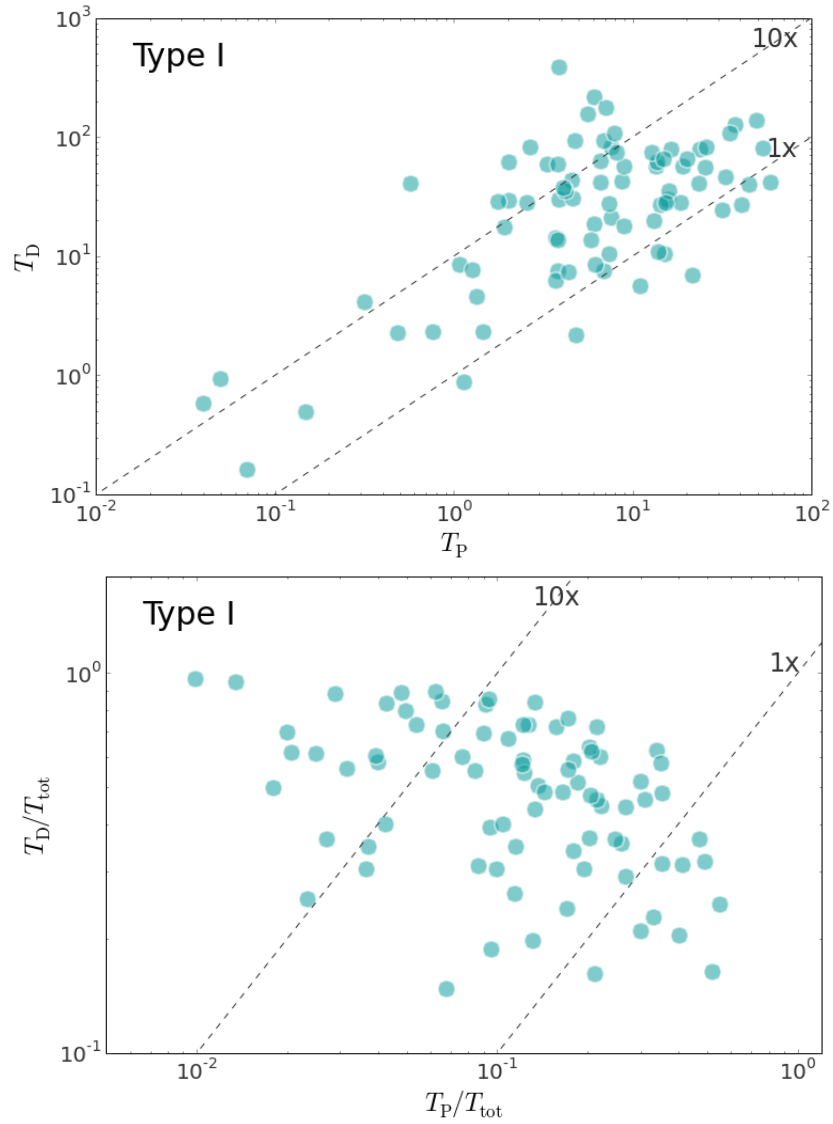


Figure 7.7: We found mild correlations between T_P , T_Q , and T_D , but these correlations vanished when I normalized the durations by T_{tot} . The examples plotted here are the T_D vs T_P (top) and T_D/T_{tot} vs T_P/T_{tot} (bottom) for bursts with Type I precursors.




Figure 7.8: **The Type II precursors have a wide range of “precursor-ness”.** Some precursors were very clear peaks and almost resembled Type Is (left), while others were more “shoulder-like” and consistent with being the start of a slow-rising peak.

there were a total of 95 Type II precursors. There were more bursts with multiple Type II precursors than with Type I precursors, and some had both Types I and II.

Within the Type II group, some precursors were clearly FRED-like before being overrun by the dominant emission, whereas others did not have a clear rise-and-fall shape and could merely be the slowly rising start of the dominant emission. At the same time, we also found “shoulders,” or long periods of low-level emission before the dominant emission that did not appear to be merely the start of a slowly rising peak. (See Figure 7.8 for examples of either extreme of the Type II group.) Overall, the Type II group is the least well-defined.

We did not compare the durations for Type II GRBs since, for this group, the dominant emission always starts before the precursor has ended.

7.5.3 Type III

Out of the first 1275 GBM GRBs, we found that 60 GRBs ($\sim 5\%$ of all GRBs and $\sim 6\%$ of high-flux GRBs) had Type III precursors, for a total of 67 Type III precursors after counting multiples. We only required that Type III precursors be distinct peaks that were dimmer than the emission that followed; since the flux

	α_P and α_{D1}	$E_{\text{peak,P}}$ and $E_{\text{peak,D1}}$
Type I	$r = 0.63$ $\rho = 0.66$ $p_{\text{KS}} = 0.3$ $p_{\text{AD}} = 0.3$	$r = 0.48$ $\rho = 0.61$ $p_{\text{KS}} = 0.02$ $p_{\text{AD}} = 0.03$
Type II	$r = 0.58$ $\rho = 0.56$ $p_{\text{KS}} = 0.7$ $p_{\text{AD}} = 0.5$	$r = 0.66$ $\rho = 0.74$ $p_{\text{KS}} = 0.2$ $p_{\text{AD}} = 0.1$
Type III	$r = 0.64$ $\rho = 0.56$ $p_{\text{KS}} = 0.90$ $p_{\text{AD}} = 0.93$	$r = 0.69$ $\rho = 0.75$ $p_{\text{KS}} = 0.2$ $p_{\text{AD}} = 0.2$

Table 7.3: **Within each subtype, we found that the precursor and first dominant emission peak spectral parameters are mildly correlated within bursts and that the distributions are consistent with each other among the bursts.** We calculated correlation coefficients (Pearson’s r and Spearman’s ρ) between α and E_{peak} . We also calculated the two-sample Kolmogorov-Smirnov (p_{KS}) and Anderson-Darling (p_{AD}) tests to determine the p -values for the two samples being drawn from the same population (small p -values indicate that the two samples are likely drawn from different populations). For these calculations, we used only the subset of GRBs for which both the precursor and first dominant emission peak were well-modeled by a Comptonized or Band function.

requirement is less stringent, the subset of GRBs with Type III precursors is dimmer on average than the subset with Type I or II precursors.

We did not compare the durations for Type III GRBs as only some of these GRBs had well-defined quiescent periods.

7.5.4 Comparison

We modeled the precursor(s) and the first dominant emission peak for each GRB in our sample, using the standard nonthermal models for GRB prompt emission described before (power law, Comptonized function, Band function). The values of α and E_{peak} are plotted in Figure 7.9 for the three types, with the precursor in

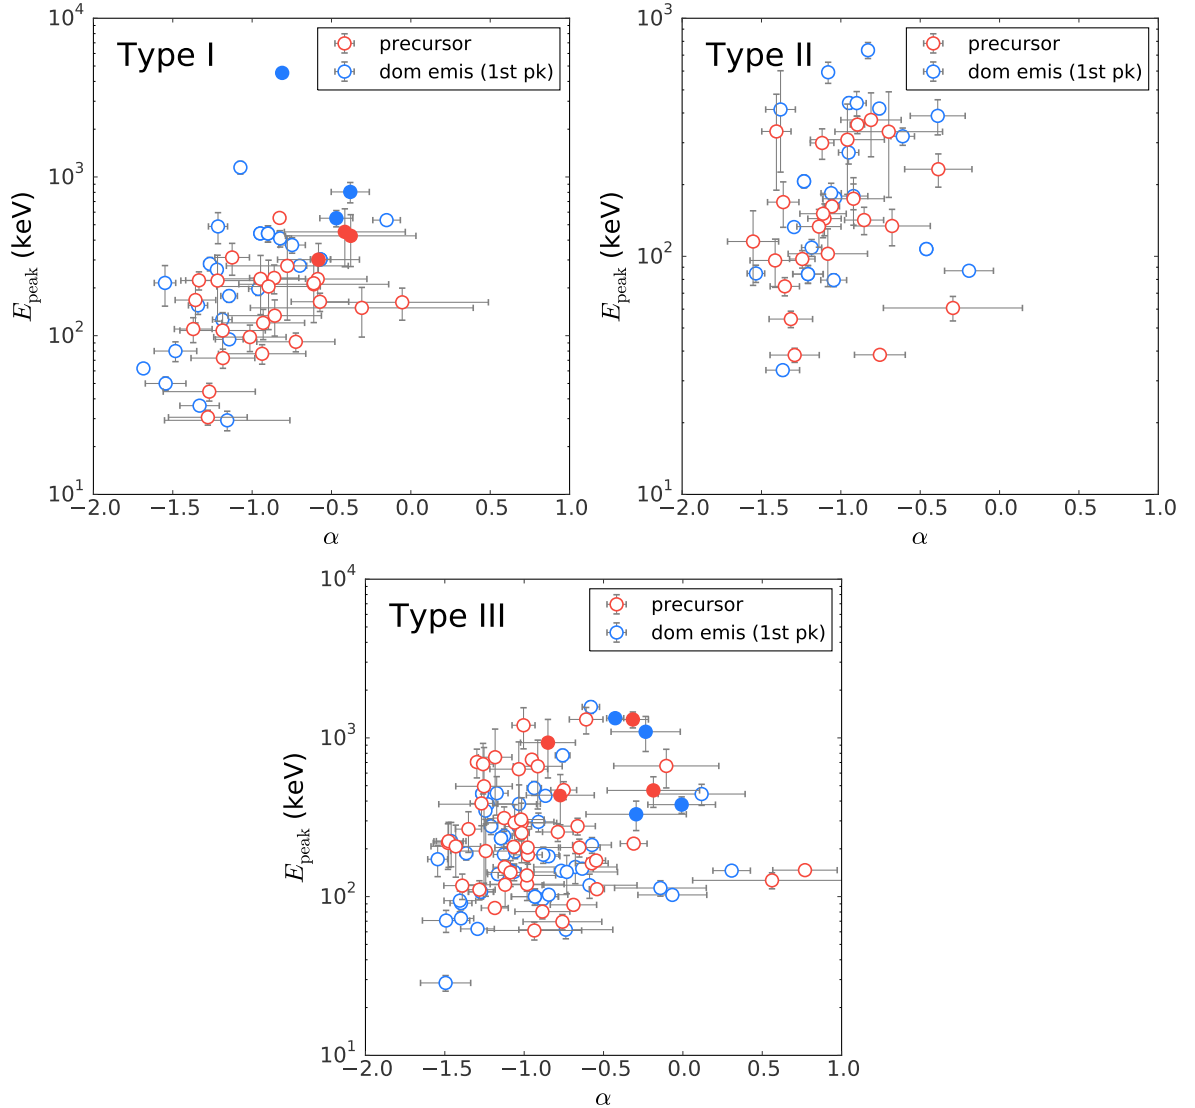


Figure 7.9: The E_{peak} and index values for the precursor and first dominant emission peak are plotted for the three precursor subtypes. In these plots, filled circles correspond to short GRBs and open circles to long GRBs. There are a few Type I precursors with particularly hard indices, although these have large error bars. For all three subtypes, the E_{peak} -index values for the precursor and first dominant emission peak span similar ranges, with a few exceptions.

red and first dominant emission peak in blue. Since these plots require E_{peak} measurements, they only include GRBs for which both the precursor and first dominant emission peak were well fit by either a Comptonized or Band function. The short GRBs ($T_{90} < 2$) are plotted with filled circles, long GRBs with open circles. The P and D1 subscripts correspond to “precursor” and “first dominant emission peak,” respectively.

We compared the α and E_{peak} distributions within each subtype. To determine whether the distributions are correlated (e.g., whether a harder precursor α also indicates a harder first dominant emission peak α within a burst), we used r and ρ as previously described. To determine whether the distributions were consistent with each other over the group of bursts within a subtype, we used the two-sample Kolmogorov-Smirnov (KS) and Anderson-Darling (AD) tests to calculate the approximate p -values for the samples being drawn from the same underlying population. (For the KS and AD tests, small p -values indicate that the samples are likely drawn from different populations.) The results are listed in Table 7.3. We found mild correlations between the precursor and first dominant emission peak α , as well as between the E_{peak} values. That is, bursts that have harder precursors also tended to have harder first dominant emission peaks. We also found no indication that either the α or E_{peak} distributions were inconsistent when comparing the precursor to the first dominant emission peak ($p > 0.01$).

We also compared the α_{P} , α_{D1} , $E_{\text{peak,P}}$, and $E_{\text{peak,D1}}$ distributions across the three different subtypes (Table 7.4). We found that the distributions were all consistent with each other. There is some indication that the E_{peak} distributions are different for Type II and III precursors ($p \gtrsim 0.001$), but this only has borderline significance, especially with the small sample sizes in question.

The distributions of α or γ for the precursors (red) and dominant emission first peaks (blue) are shown in Figure 7.10 on the left. The darker histograms

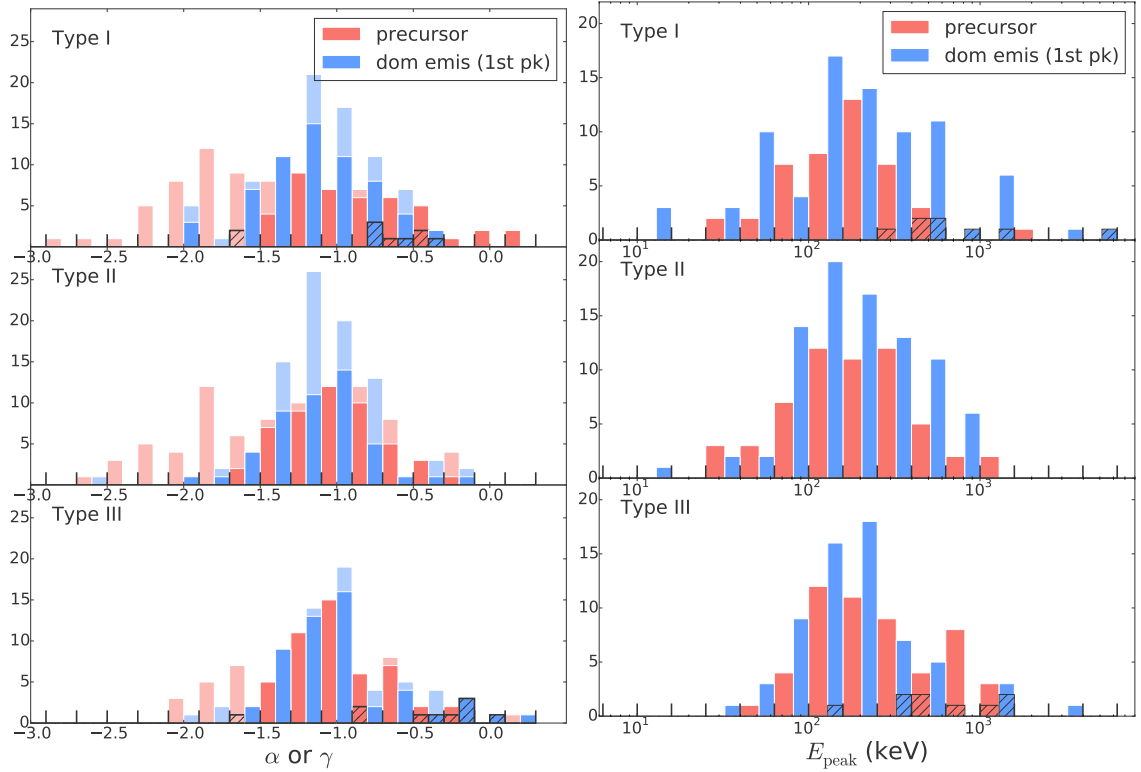


Figure 7.10: **The α (or γ) and E_{peak} distributions are plotted here for all the GRBs with precursors in our sample.** In both plots, the darker histograms represent the emission episodes that were well fit by a Comptonized or Band function; in the plot on the left, the lighter histogram represents the emission episodes that were best fit by a power law. The hatched histograms in both plots correspond to the short GRBs.

		Type I	Type II	Type III
α_P	Type I	—	$p_{KS} = 0.3$ $p_{AD} = 0.1$	$p_{KS} = 0.3$ $p_{AD} = 0.5$
	Type II	$p_{KS} = 0.3$ $p_{AD} = 0.1$	—	$p_{KS} = 0.6$ $p_{AD} = 0.5$
	Type III	$p_{KS} = 0.3$ $p_{AD} = 0.5$	$p_{KS} = 0.6$ $p_{AD} = 0.5$	—
α_{D1}	Type I	—	$p_{KS} = 1.0$ $p_{AD} = 1.0$	$p_{KS} = 0.9$ $p_{AD} = 0.5$
	Type II	$p_{KS} = 1.0$ $p_{AD} = 1.0$	— —	$p_{KS} = 0.7$ $p_{AD} = 0.7$
	Type III	$p_{KS} = 0.9$ $p_{AD} = 0.5$	$p_{KS} = 0.7$ $p_{AD} = 0.7$	—
$E_{\text{peak},P}$	Type I	—	$p_{KS} = 0.4$ $p_{AD} = 0.5$	$p_{KS} = 0.07$ $p_{AD} = 0.01$
	Type II	$p_{KS} = 0.4$ $p_{AD} = 0.5$	—	$p_{KS} = 0.007$ $p_{AD} = 0.002$
	Type III	$p_{KS} = 0.07$ $p_{AD} = 0.01$	$p_{KS} = 0.007$ $p_{AD} = 0.002$	—
$E_{\text{peak},D1}$	Type I	—	$p_{KS} = 0.3$ $p_{AD} = 0.2$	$p_{KS} = 0.2$ $p_{AD} = 0.2$
	Type II	$p_{KS} = 0.3$ $p_{AD} = 0.2$	—	$p_{KS} = 0.6$ $p_{AD} = 0.6$
	Type III	$p_{KS} = 0.2$ $p_{AD} = 0.2$	$p_{KS} = 0.6$ $p_{AD} = 0.6$	—

Table 7.4: We checked for differences in the spectral parameters between the subtypes and found that the parameter distributions were all consistent with each other, with the possible exception of $E_{\text{peak},P}$ between bursts with Type II and III precursors.

represent the emission episodes for which we were able to constrain E_{peak} ; that is, the episodes that were well-fit by a Comptonized or Band function. In essence, the α distributions are plotted in the dark histograms and the γ distributions in the light histograms. On average, the α distribution is harder than the γ distribution, since softer indices ($\lesssim -2$) usually indicate that we are probing the falling part of the νF_ν spectrum (i.e., the energies above E_{peak}); this is consistent with the distributions shown in Figure 7.4.

For the subset of emission episodes for which we were able to measure E_{peak} , the distributions of E_{peak} are plotted in the right-hand plot of Figure 7.10. Very few emission episodes had enough statistics above E_{peak} to be well-fit by a Band function and constrain the high-energy power law index β , so we did not visually distinguish them in the set of plots.

In both sets of panels in Figure 7.10, short GRBs are represented with hatched histograms. These tend to have a harder α and higher E_{peak} than average in both their precursor and dominant emission episodes. This is consistent with the tendency for short GRBs to be harder on average than long GRBs [23].

GRB spectra are constantly evolving with time. In order to quantify how the spectrum changes between the precursor and the dominant emission, we took the ratios of α and E_{peak} between the precursor and first dominant emission peak for each GRB. These are plotted in Figure 7.11 and listed in Table 7.5. The α ratio distributions (left) are similar between the three types, and they all peak around unity; i.e., α does not change much in going from P to D1. On the other hand, the E_{peak} ratio distribution (right) for the Type III bursts looks different from the Type I and II distributions ($p \lesssim 1 \times 10^{-4}$). In particular, the E_{peak} ratio distribution is skewed to values less than 1 for Types I and II (indicating that E_{peak} increases in going from P to D1), whereas the distribution is more symmetric around 1 (with some outliers) or skewed to values greater than 1 for Type III bursts.

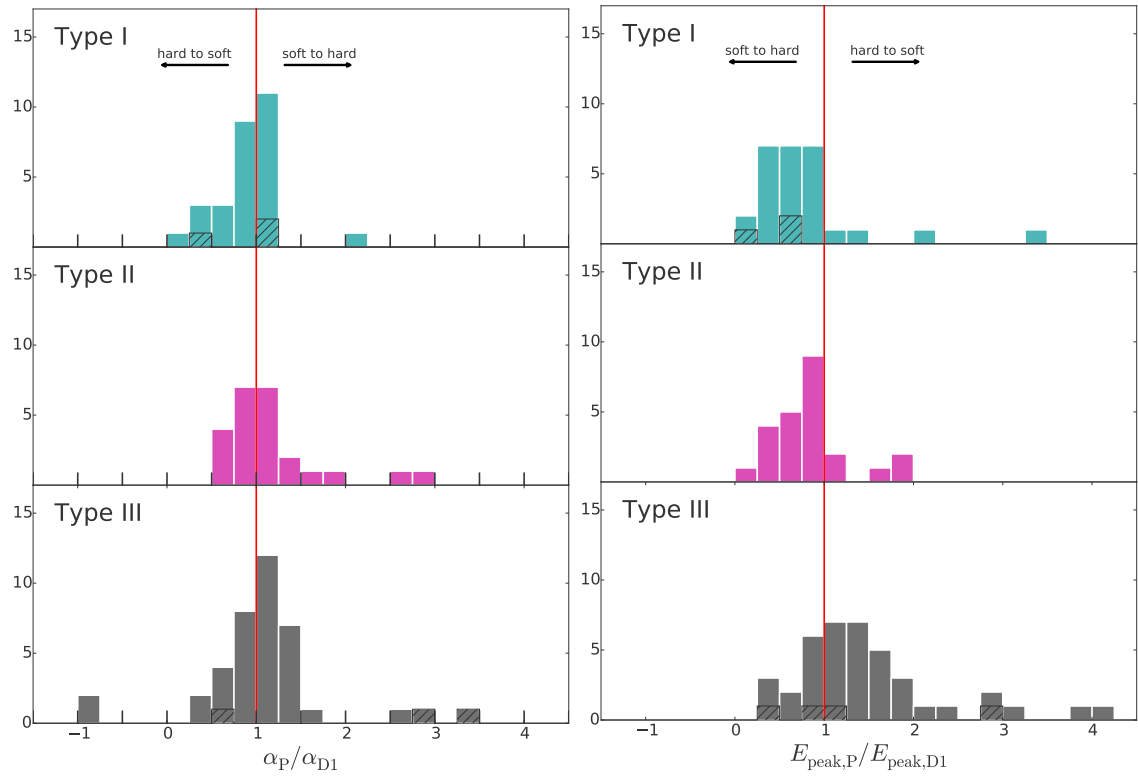


Figure 7.11: For both α (left) and E_{peak} (right), we divided the value for the precursor by the value for the corresponding dominant emission. The short GRBs are marked by the hatched histograms.

		Type I	Type II	Type III
α_P/α_{D1}	Type I	—	$p_{KS} = 0.4$ $p_{AD} = 0.2$	$p_{KS} = 0.1$ $p_{AD} = 0.1$
	Type II	$p_{KS} = 0.4$ $p_{AD} = 0.2$	—	$p_{KS} = 0.7$ $p_{AD} = 0.6$
	Type III	$p_{KS} = 0.1$ $p_{AD} = 0.1$	$p_{KS} = 0.7$ $p_{AD} = 0.6$	—
$E_{\text{peak,P}}/E_{\text{peak,D1}}$	Type I	—	$p_{KS} = 0.2$ $p_{AD} = 0.2$	$p_{KS} < 10^{-5}$ $p_{AD} < 10^{-5}$
	Type II	$p_{KS} = 0.2$ $p_{AD} = 0.2$	—	$p_{KS} < 10^{-5}$ $p_{AD} = 2 \times 10^{-4}$
	Type III	$p_{KS} < 10^{-5}$ $p_{AD} < 10^{-5}$	$p_{KS} < 10^{-5}$ $p_{AD} = 2 \times 10^{-4}$	—

Table 7.5: When we compared the spectral parameter ratios (α_P/α_{D1} and $E_{\text{peak,P}}/E_{\text{peak,D1}}$) across the different subtypes, we found indications that the GRBs with Type III precursors are inconsistent with the GRBs with Type I or II precursors. The p -values for the E_{peak} ratios are mostly $\lesssim 10^{-5}$, suggesting that the distributions are inconsistent with each other.

It is possible that the difference in the E_{peak} ratio distributions signifies that Type III precursors are inherently different from Types I and II. The Type I and II precursors have a relative flux requirement, while the Type IIIs do not; in essence, the Type III precursors look more similar to their dominant emission episodes than the Type I or II precursors. This is supported by the fact that, if we ignore the outliers in the E_{peak} ratio for the Type III bursts, the rest of the distribution is relatively symmetric around 1, suggesting that less spectral evolution occurs for bursts with Type III precursors than for Types I or II. Perhaps this means that Type III precursors are not special; that they are merely the first emission episode of the dominant emission and slightly dimmer by chance; that their physical origin is the same as the origin of the dominant emission. Then, if bursts with Type I and II precursors evolve differently from Type IIIs, this suggests that Type I and II precursors *do* have a distinct physical origin.

However, this difference is possibly a selection effect. The Amati, Ghirlanda, and Yonetoku relations all find a correlation between E_{peak} and some measure of the total energy release or luminosity (E_{iso} , E_{γ} , L_{iso}). Figure 7.12 compares the precursor and first dominant emission peak energy fluxes for the three subtypes. The distribution of the ratio of precursor to first dominant emission peak energy flux is different for Type IIIs than for Type Is or IIs (Table 7.6); this seems to be primarily driven by the differences in the precursor energy flux, with Type III precursors being more high-flux than Type I or IIs, likely due to the looser flux requirement for a Type III precursor. If the Type III precursors have larger energy fluxes, then by a loose interpretation of the Amati/Ghirlanda/Yonetoku relations, they could also have higher E_{peak} values on average, which would explain the E_{peak} ratio being more symmetric around 1 for bursts with Type III precursors. This is somewhat borne out by the comparisons of $E_{\text{peak,P}}$ across the different subtypes (Table 7.4). More work is necessary to disentangle this effect.

		Type I	Type II	Type III
P energy flux	Type I	—	$p_{\text{KS}} = 0.002$ $p_{\text{AD}} = 0.002$	$p_{\text{KS}} < 10^{-5}$ $p_{\text{AD}} < 10^{-5}$
	Type II	$p_{\text{KS}} = 0.002$ $p_{\text{AD}} = 0.002$	—	$p_{\text{KS}} = 0.002$ $p_{\text{AD}} = 0.005$
	Type III	$p_{\text{KS}} < 10^{-5}$ $p_{\text{AD}} < 10^{-5}$	$p_{\text{KS}} = 0.002$ $p_{\text{AD}} = 0.005$	—
D1 energy flux	Type I	—	$p_{\text{KS}} = 0.008$ $p_{\text{AD}} = 0.007$	$p_{\text{KS}} = 0.5$ $p_{\text{AD}} = 0.2$
	Type II	$p_{\text{KS}} = 0.008$ $p_{\text{AD}} = 0.007$	—	$p_{\text{KS}} = 0.005$ $p_{\text{AD}} = 0.003$
	Type III	$p_{\text{KS}} = 0.005$ $p_{\text{AD}} = 0.003$	$p_{\text{KS}} = 0.008$ $p_{\text{AD}} = 0.007$	—
ratio	Type I	—	$p_{\text{KS}} = 0.9$ $p_{\text{AD}} = 0.7$	$p_{\text{KS}} < 10^{-5}$ $p_{\text{AD}} < 10^{-5}$
	Type II	$p_{\text{KS}} = 0.9$ $p_{\text{AD}} = 0.7$	—	$p_{\text{KS}} < 10^{-5}$ $p_{\text{AD}} = 1 \times 10^{-4}$
	Type III	$p_{\text{KS}} < 10^{-5}$ $p_{\text{AD}} < 10^{-5}$	$p_{\text{KS}} < 10^{-5}$ $p_{\text{AD}} = 1 \times 10^{-4}$	—

Table 7.6: **When we compared the energy fluxes and energy flux ratios between the different subtypes, we found that the distributions for GRBs with Type III precursors were inconsistent with the distributions for GRBs with Type I or II precursors.** In other words, the Type IIIs seemed different from the Type I or IIs when it came to the precursor energy flux or the precursor to first dominant emission peak energy flux ratio.




Figure 7.12: **For the most part, the D1 energy flux is larger than the P energy flux, for all three types.** However, the relative difference between D1 and P energy flux is smaller on average for the Type III precursors than Type I or II.

7.6 Interpretation

7.6.1 Photospheric precursors

The jet is optically thick when it is first emitted from the central engine. As the jet propagates, it slows down and expands, and at some point becomes optically thin. Photospheric emission is naturally expected around this transition region, as the thermal photons that were previously contained are able to escape. The question is not whether photospheric emission is present, but whether it is bright enough to be detected as a separate component.

Precursors could be photospheric emission that is separated in time from the dominant emission. If this is true, then the precursor spectrum should be thermal or quasithermal and peak in X rays, while the dominant emission spectrum should be nonthermal and mostly in gamma rays. If the dominant emission comes from internal shocks, then any quiescent period would be very short, from a millisecond to a second before the dominant emission [71] [116].

Longer quiescent periods are possible if the jet is magnetically dominated. In this scenario, the dominant emission would be expected at $\sim 10^{16}$ cm rather than the $\sim 10^{13}$ cm in the internal shock scenario [117]. However, photospheric emission is expected to be especially dim in a Poynting flux-dominated jet. Larger magnetic fields allow for more delayed dominant emission (and therefore long quiescent periods), but also make it difficult to have any detectable photospheric emission in the first place [118].

A thermal spectrum can be approximated by a nonthermal spectrum with a particularly hard low-energy index. For instance, during an early part of the emission, the *Fermi* data for GRB 090902B were best fit by a Band function with $\alpha = 0.07$ rather than the more standard $\alpha \sim -1$ [128], and values of $\alpha > 0$ are often interpreted as thermal emission [129]. Based on temporal properties alone, the




Figure 7.13: **Photospheric precursors would have particularly hard values of α , and their corresponding quiescent periods would be very short.** The approximate parameter space they would cover is marked by the shaded region (assuming internal shocks). In this plot, spectra that were best fit by Comptonized or Band function are marked with filled circles, and spectra that were best fit by a power law with open circles.

Type II precursors are the most likely to be photospheric precursors; however, we did not find any Type II precursors with particularly hard low-energy indices ($\alpha > 0$) (Figure 7.9). We also did not find any Type I precursors that unambiguously met both criteria (Figure 7.13); there were no precursors that were particularly hard or followed by particularly short quiescent periods. This does not preclude the possibility of thermal precursors; possibly, their fluxes are too dim to have been detected using these methods.

7.6.2 Shock breakout precursors

Long GRBs occur when the core of a massive star collapses, leaving the outermost layers intact in the form of a stellar envelope. A jet emitted by the central

engine would eventually run into this material and would have to successfully penetrate it in order to be observable. As it does so, the interactions heat the material immediately ahead of the jet; when this heated material breaks through the stellar envelope, it releases thermal emission in the form of a shock breakout [119]. Shock breakout precursors are expected to be very weak, partly because they would be only weakly beamed and their energy would spread out over a larger surface. Assuming the dominant emission is caused by internal shocks, the thermal precursor would occur roughly ~ 1 -10 s before the dominant emission, and would peak in X rays [120]. There are indeed a few Type I precursors that fit this description (Figure 7.13), although the majority of these precursors are too soft.

Particles within the jet can then interact with the shock breakout emission and upscatter this to form another precursor, this time at gamma-ray energies. These would be expected only about a second before the dominant emission, and would emit a nonthermal spectrum at around hundreds of keV. However, this would be extremely difficult to disentangle from the dominant emission, which manifests as a nonthermal spectrum at around hundreds of keV as well. For all intents and purposes, these “upscattered shock breakout” precursors would look just like another dominant emission peak [120].

Shock breakout emission can occur in short GRBs if the merger event or what immediately follows it produces a wind. This wind would have to be baryon-loaded in order to form a sufficiently dense cloud. This is most naturally accomplished if the progenitor is a NS-NS system and forms an unstable hypermassive neutron star (HMNS). As the HMNS spins down, it can release a dense neutrino-driven wind [121]. Presumably, the precursor properties in this scenario would mirror those in the long GRB scenario, just on a shorter timescale.

Interestingly, one would expect shock breakout precursors to appear more often in long GRBs (the stellar envelope is always expected to be present) than in short

GRBs (the progenitor system must be a NS-NS system, and the central engine must be a magnetar rather than a black hole), and this is indeed the case (Figure 7.5).

7.6.3 Fallback precursors

Like the shock breakout scenario, the fallback precursor scenario more naturally applies to long than short GRBs. In this model, the central engine releases an initial weak jet. The jet manages to penetrate the stellar envelope, but the interactions slow it down so that some of the jet material falls back onto and is accreted by the central engine, powering a second, stronger jet. The typical timescale for enough material to fall back is around 100 s, much longer than in most other scenarios.

Since the precursor is caused by a jet (in particular, a weak jet that is separate from the jet that causes the dominant emission), the precursor emission should look like the standard nonthermal GRB emission, just weaker and at lower energies [66]. Based on the quiescent period timescale and the weakness of the jets, fallback precursors would manifest as Type I precursors.

7.6.4 Multistage collapse precursors

If the progenitor system had a sufficiently fast rotation, the core collapse or binary merger could produce a temporarily stable intermediate object before collapsing into the final black hole. This intermediate object is more generally known as a “spinar,” defined as a collapsing object that is temporarily stable by balancing centrifugal and gravitational forces [122].

The general spinar scenario can support a large range of precursor energies and quiescent period durations. Most notably, depending on the initial parameters, the quiescent period can be up to hundreds of seconds long [123]. Depending on the progenitor system, precursors followed by quiescent periods of less than 100 s should occur in $\sim 10\%$ of GRBs, which is roughly consistent with our results. However, this

model does not naturally explain the presence of multiple precursors.

7.6.5 NS-NS magnetic field interactions

In the binary merger scenario, the two objects are often assumed to be neutron stars (NS). The interactions between their magnetospheres can emit energy in the form of a thermal X-ray precursor that hardens as the objects get closer; larger initial magnetic fields would result in higher energy precursors [124]. These would probably be too weak to detect, however [125].

7.6.6 NS crust cracking

In the binary merger scenario of short GRBs, if at least one of the compact objects is a NS, then that NS will be tidally affected before the merger. This can shatter or crack the crust and release energy. If crust cracking takes it place, it would primarily occur when the crust is directly tidally deformed some milliseconds before the merging event [126] [127]. It could also manifest as a resonant effect a few seconds before, allowing for longer quiescent periods as in Troja et al. (2010) [125]. Both can be reasonably expected to occur, but it is unclear whether the energy release would be large enough to be detected as a precursor. The resonant crust cracking scenario could theoretically account for quiescent periods as long as a few tens of seconds before the dominant emission in short GRBs.

NS resonant crust cracking was invoked to explain some of the longer quiescent periods found by Troja et al. [111], which found precursors as long as 100 s before the trigger time in short GRBs detected by *Swift*. We did not find any such quiescent periods in our search; however, we used a different method from Troja et al. (Bayesian blocks rather than a wavelet transform analysis) and effectively set a higher discovery threshold (requiring the peak to appear in multiple detectors simultaneously). The longest quiescent period we found for a short GRB (defined

by the GBM T90) was 3.4 s, for a GRB whose dominant emission was 0.9 s long. (Other quiescent periods of short GRBs were all less than 0.5 s in duration, and shorter than their respective dominant emission durations.) If this a true short burst caused by a binary merger, then resonant crust cracking could explain its origin.

7.7 Discussion

We found some indication that a subset of precursors are statistically distinct from the rest of the population, but cannot exclude the possibility of systematic effects or observational biases causing this. If the precursors we found *are* caused by a different emission mechanism from the internal jet interactions that cause the dominant emission, then a few conclusions can be drawn. Most of the precursors had spectra with soft low-energy indices that did not suggest thermal emission, which disfavors the photospheric and shock breakout models. We found quiescent periods with durations longer than 100 s, which is most naturally explained by the fallback accretion or multistage collapse models.

Most of the previous precursor studies concluded that, as a whole, the properties of precursors and dominant emission episodes were similar; in particular, the studies that did not have extra requirements on the precursor or quiescent period duration (both Burlon et al. studies, the Troja et al. study, the Hu et al. study) all found that the precursor and dominant emission spectral properties were consistent with each other. Similarly, when we compared the precursor and dominant emission spectral properties, we found that they were consistent with being drawn from the same population; i.e., there were no statistically significant differences between the precursor and dominant emission.

It is possible, then, that the relevant quantity is not the relative flux of the first peak but rather the presence of the quiescent period. One particular study of quiescent periods in BATSE GRBs found that the post-quiescent emission tends

to be longer than the pre-quiescent emission (which is also a property we found (Figure 7.7)); and, based on some of the longer quiescent period durations, they also favored models that involved a change in the central engine behavior [130]. In our study, these “central engine” precursor models are the fallback accretion and multistage collapse models, but these were developed to account for the relatively low flux in the precursor / pre-quiescent emission. If, on the other hand, the relative flux is unimportant, then this opens up the possibility of other emission models (e.g., [67]).

Chapter 8: Conclusions and final thoughts

Many aspects of GRBs remain mysterious today, more than four decades after their initial discovery. Observations often led to the discovery of new information about GRBs. These observations were made possible by a series of great observatories: *CGRO*, *BeppoSAX*, *HETE-2*, *Swift*, many ground-based optical and radio observatories, and countless others.

Fermi continues this history of GRB observations that shed new light on their structures and physical mechanisms. Together, the *Fermi*-LAT and GBM provide observations over more than eight orders of magnitude in energy. Both the LAT and GBM can detect and localize GRBs onboard and in realtime, and communicates the localizations with the rest of the GRB community to facilitate follow-up observations. The LAT onboard GRB detection algorithm can provide localizations that are precise enough for *Swift* to observe in a few tiled observations, and improvements in both the sensitivity and localization capability of the onboard algorithm will soon be implemented.

The LAT has observed high-energy emission from over 100 GRBs so far, including some of the brightest and most energetic. Both the GBM and LAT observed the record-breaking GRB 130427A, a luminous burst located close enough that its temporal and spectral behavior could be thoroughly analyzed and characterized. The synchrotron shock model for afterglow emission cannot easily explain the high-energy photons at late times that were detected by the LAT, which suggests that the standard fireball model of GRBs is insufficient to model the afterglow.

The GBM has detected over 1700 GRBs to date, and such a large data set has made it a population study of GBM GRBs with precursors possible. Precursors appeared in roughly 15% of GBM GRBs, although there are no indications that these are caused by an emission source or region distinct from the interactions within the jet that generally cause the prompt emission. However, standard jet models of GRB prompt emission cannot easily explain the long quiescent periods that appear in some GRBs with precursors. Further work is needed to ascertain whether precursors are indicative of special behavior.

Open questions remain in the field of GRBs more than forty years after their discovery. Knowledge about GRBs is and has been best gained with multiple instruments working together. In the coming years, *Fermi* will continue to work with electromagnetic detectors such as *Swift* and HAWC, neutrino detectors such as IceCube, and gravitational detectors such as Advanced LIGO (the Laser Interferometer Gravitational-Wave Observatory) in a truly multimessenger era of GRB observation.

Bibliography

- [1] Klebesadel, R.W., Strong, I.B., & Olson, R.A. 1973, ApJ, 182, L85
- [2] Kuznetsov, A.V., et al. 1986, Pis'ma Astron. Zh., 12, 755
- [3] Cline, T.L., et al. 1973, ApJ, 185, L1
- [4] Mazets, E.P., et al. 1981, Ap&SS, 80, 3
- [5] Mazets, E.P., et al. 1981, Nature, 290, 378
- [6] Murakami, T., et al. 1988, Nature, 335, 234
- [7] Nolan, P.L., et al. 1984, Nature, 311, 360
- [8] Mazets, E.P. & Golenetskii, S. V. 1981, Ap&SS, 75, 47
- [9] Mazets, E.P. 1985 in ICRC conference proceedings, 9
- [10] Manchanda, R.K. & Ramsden, D. 1977, Nature, 266, 425
- [11] Paczyński 1986, ApJ, 308, L43
- [12] Vedrenne, G. 1981, Phil Trans R Soc Lond A, 301, 645
- [13] Fishman, G.J., et al. 1994, ApJSS, 92, 229
- [14] Kanbach, G., et al. 1989, Space Sci Rev, 49, 69
- [15] Palmer, D.M., et al. 1994, ApJ, 433, L77

- [16] Goldstein, A., et al. 2013, ApJSS, 208, 21
- [17] Band, D., et al. 1993, ApJ, 413, 281
- [18] Gonzalez Sanchez, M.M. 2005, PhD Thesis, University of Wisconsin-Madons
- [19] Briggs, M.S., et al. 1999, ApJ, 524, 82
- [20] Hurley, K., et al. 1994, Nature, 372, 652
- [21] Share, G.H., et al. 1994 in Proceedings of the Second Gamma-Ray Bursts Workshop, 283
- [22] Hurley, K., et al. 1994, Nature, 372, 652
- [23] Kouveliotou, C., et al. 1993, ApJ, 413, L101
- [24] Mao, S. & Paczyński, B. 1992, ApJ, 389, L13
- [25] Lamb, D.Q. 1995, PASP, 107, 1152
- [26] Paczyński, B. 1995, PASP, 107, 1167
- [27] Cavallo, G. & Reese, M. J. 1978, MNRAS, 183, 359
- [28] Barthelmy, S.D., et al. 1995, Ap&SS, 231, 235
- [29] Smith, D. 2010, in Observing Photons in Space, 21
- [30] Dermer, C. & Menon, G. 2009, High Energy Radiation From Black Holes
- [31] Atwood, W.B., et al. 2009, ApJ, 697, 1071
- [32] Atwood, W.B., et al. 2007, Astropart Phys, 28, 422
- [33] Grove, J.E., et al. 2010, in Space Telescopes and Instrumentation 2010: Ultraviolet to Gamma Ray
- [34] Moiseev, A.A., et al. 2007, Nucl Instrum Meth A, 583, 372
- [35] Ackermann, M., et al. 2012, ApJSS, 203, 4

- [36] Meegan, C., et al. 2009, ApJ, 702, 791
- [37] Cash, W. 1979, ApJ, 228, 939
- [38] Vasileiou, V. 2013, Astropart Phys, 48, 61
- [39] McEnery, J., et al. 2010, GCN Circ 10777
- [40] Boella, G., et al. 1997, A&AS, 122, 299
- [41] Galama, T., et al. 1997, Nature, 387, 479
- [42] Bloom, J.S., et al. 2001, ApJ, 554, 678
- [43] Costa, E., et al. 1997, Nature, 387, 783
- [44] Galama, T.J., et al. 2000, ApJ, 536, 286
- [45] Metzger, M.R. & Cohen, J.G. 1997, IAUC 6676
- [46] Metzger, M.R., et al. 1997, Nature, 387, 878
- [47] Bloom, J.S., et al. 1998, ApJ, 507, L25
- [48] Galama, T.J., et al. 1999, A&AS, 138, 465
- [49] Galama, T.J., et al. 1998, ApJ, 500, L97
- [50] Piro, L. et al., 1998, A&A, 331, L41
- [51] Kawai, N., et al. 2001, in Small Telescope Astronomy on Global Scales, 149
- [52] Hjorth, J., et al. 2005, Nature, 437, 859
- [53] Horan, D., et al. 2007, ApJ, 655, 396
- [54] Atkins, R., et al. 2004, ApJ, 604, L25
- [55] Aharonian, F., et al. 2009, A&A, 495, 505
- [56] Acciari, V.A., et al. 2011, ApJ, 743, 62

- [57] McEnery, J.E., et al. 1999, in GeV-TeV Astrophysics: Toward a Major Atmospheric Cherenkov Telescope VI, 1999
- [58] Mészáros, P. & Rees, M.J. 1993, ApJ, 405, 278
- [59] Zhang, B. 2007, Adv Space Res, 40, 1186
- [60] Racusin, J.L., et al. 2009, ApJ, 698, 43
- [61] O'Brien, P.T., et al. 2006, New Journal Phys, 8, 121
- [62] Kouveliotou, C., Wijers, R.A.M.J, Woosley, S.E., et al. 2013, Gamma-ray Bursts, Chapter 10
- [63] Nakar, E. 2007, Phys Reports, 442, 166
- [64] Tanvir, N.R., et al. 2013, Nature, 500, 547
- [65] Berger, E., et al. 2013, ApJ, 774, L23
- [66] Wang, X.-Y. & Mészáros, P. 2007, ApJ, 670, 1247
- [67] Tchekhovskoy, A. & Giannios, D. 2015, MNRAS, 447, 327
- [68] Lü, H.J., et al. 2015, ApJ, 805, 89
- [69] Piran, T. 1999, Phys Reports, 314, 575
- [70] Lyutikov, M. 2006, New Journal Phys, 8, 119
- [71] Lyutikov, M. & Usov, V. 2000, ApJ, 543, L129
- [72] Rees, M.J. & Mészáros, P. 2005, ApJ, 628, 847
- [73] Ryde, F., et al. 2011, MNRAS, 415, 3693
- [74] Dermer, C.D. & Atoyan, A. 2006, New Journal Phys, 8, 122
- [75] Vestrand, W.T., et al. 2014, Science, 343, 38
- [76] Zhang, B., et al. 2006, ApJ, 642, 354

- [77] Ackermann, M., et al. 2013, ApJSS, 209, 11
- [78] Burrows, D.N., et al. 2007, Phil Trans R Soc Lond A, 365, 1213
- [79] Li, L.-B., et al. 2015, Ap&SS, 359, 37
- [80] Perley, D.A., et al. 2014, ApJ, 781, 37
- [81] Ackermann, M., et al. 2014, Science, 343, 42
- [82] Aliu, E., et al. 2014, ApJ, 795, L3
- [83] Ackermann, M., et al. 2011, ApJ, 729, 114
- [84] Preece, R., et al. 2014, Science, 343, 51
- [85] Maselli, A., et al. 2014, Science, 343, 48
- [86] Chaplin, V., et al. 2013, NIMPA, 717, 21
- [87] Goldstein, A., et al. 2012, ApJSS, 199, 19
- [88] Nolan, P.L., et al. 2012, ApJSS, 199, 31
- [89] Cenko, S.B., et al. 2011, ApJ, 732, 29
- [90] MacLachlan, G.A., et al. 2013, MNRAS, 432, 857
- [91] Giannios, D., 2012 A&A, 480, 305
- [92] Dermer, C., et al. 2000, ApJ, 537, 785
- [93] Sari, R. & Esin, A. 2001, ApJ, 538, 787
- [94] Zhang, B. & Mészáros, P. 2001, ApJ, 559, 110
- [95] Liu, R.-Y., et al. 2013, ApJ, 773, L20
- [96] Blaufuss, E. 2013, GCN Circ 14520
- [97] Moderski, R., et al. 2005, MNRAS, 363, 954

- [98] Racusin, J.L., et al. 2008, *Nature*, 455, 183
- [99] Kouveliotou, C., et al. 2013, *ApJ*, 779, L1
- [100] Abeysekara, A.U., et al. 2015, *ApJ*, 800, 78
- [101] Scargle, J.D. 1998, *ApJ*, 504, 405
- [102] Scargle, J.D., et al. 2013, *ApJ*, 764, 167
- [103] Singer, L.P., et al. 2013, *pJ*, 776, L34
- [104] Metzger, A.E., et al. 1974, *ApJ*, 194, L19
- [105] Koshut, T.M., et al. 1995, *ApJ*, 452, 145
- [106] Lazzati, D. 2005, *MNRAS*, 357, 722
- [107] Burlon, D., et al. 2008, *ApJ*, 685, L19
- [108] Burlon, D., et al. 2009, *A&A*, 505, 569
- [109] Kaneko, Y., et al. 2006, *ApJSS*, 166, 298
- [110] Norris, J.P. & Bonnell, J.T. 2006, *ApJ*, 643, 266
- [111] Troja, E., et al. 2010, *ApJ*, 723, 1711
- [112] Hu, Y.-D., et al. 2014, *ApJ*, 789, 145
- [113] Charisi, M., et al. 2015, *MNRAS*, 448, 2624
- [114] Gruber, D., et al. 2014, *ApJSS*, 211, 12
- [115] von Kienlin, A., et al. 2014, *ApJSS*, 211, 13
- [116] Li, L.-X. 2007, *MNRAS*, 380, 621
- [117] Lyutikov, M. & Blandford, R.D. 2003, *arXiv:astro-ph/0312347*
- [118] Zhang, B. & Yan, H. 2011, *ApJ*, 726, 90

- [119] MacFadyen, A.I., et al. 2001, ApJ, 550, 410
- [120] Ramirez-Ruiz, E., et al. 2002, MNRAS, 331, 197
- [121] Murguia-Berthier, A., et al. 2014, ApJ, 788, L8
- [122] Lipunov, V.M. & Gorbovsky, E.S. 2008, MNRAS, 383, 1397
- [123] Lipunova, G.V., et al. 2009, MNRAS, 397, 1695
- [124] Hansen, B.M.S. & Lyutikov, M. 2001, MNRAS, 322, 695
- [125] Tsang, D., et al. 2012, PRL, 108, 011102
- [126] Kochanek, C.S. 1992, ApJ, 398, 234
- [127] Penner, A.J., et al. 2011, PRD, 84, 103006
- [128] Abdo, A.A., et al. 2009, ApJ, 706, L138
- [129] Hascoët, R., et al. 2013, A&A, 551, A124
- [130] Drago, A. & Pagliara, G. 2007, ApJ, 665, 1227
- [131] Bernardini, M.G., et al. 2013, ApJ, 775, 67