UNIVERSITÀ DI TRENTO

PhD SST
Space Science and Technology

Università degli Studi di Palermo

Doctorate of National Interest

on

Space Science and Technology

Temporal Analysis of High-Energy Transient Sources

An Adaptive Binning Approach in Time-Domain Astronomy

The Author
Wladimiro Leone

Supervisor
Prof. Tiziana di Salvo

Co-Supervisor
Prof. Giovanni Marsella
Co-Supervisor
Prof. Luciano Burderi

January, 2026

Acknowledgements

This thesis and related research have been conducted during and with the support of the Italian national inter-university PhD programme in Space Science and Technology.

I would like to express my sincere gratitude to all the people who have contributed, in different ways, to the completion of this PhD journey.

First and foremost, I am deeply thankful to *Prof. Tiziana Di Salvo* and *Prof. Luciano Burderi* for their invaluable guidance, constant support, and trust throughout these years. Their expertise, scientific insight, and passion have been an endless source of motivation and inspiration.

I am also grateful to my co-supervisor, *Prof. Giovanni Marsella*, for his insightful comments and for the discussions that contributed to refining several aspects of this research.

My heartfelt thanks go to my colleagues and friends from the University of Palermo — *Alessio Anitra*, *Francesco Barra*, *Simona Caserta*, *Rosario Iaria*, *Claudia Maraventano*, and *Carlotta Miceli* — from the IASF/Palermo — *Melania Del Santo*, *Ciro Pinto*, and *Fabio Pintore* — from the University of Cagliari — *Alessandro Riggio* and *Andrea Sanna* — and from the University of Ferrara — *Cristiano Guidorzi* and *Romain Maccary* — for their technical and scientific advice, as well as for the many fruitful discussions that have enriched and shaped this work.

I am also very grateful to all the collaborators of the HERMES-SP and SpIRIT missions, with a special mention to *Fabrizio Fiore* and *Riccardo Campana*, for their continuous support, insightful discussions, and guidance during the development, calibration, and data analysis phases of the instruments. I would also like to warmly thank the Masaryk University team in Brno, *Norbert Werner*, *Michaela Durisková*, and *Jakub Řípa*, for their hospitality and collaboration during my research stay. Likewise, I am grateful to *Michele Trenti* at the University of Melbourne for his valuable guidance and support during my time in Australia.

To my friends, thank you for being a constant source of encouragement and joy, for reminding me that beyond the code, data, and deadlines, there is always time for laughter, music, and good company.

Finally, my deepest gratitude goes to my family. Your unconditional love, patience, and support have been the foundation upon which this achievement stands. This milestone is, above all, dedicated to you.

«Praeterita non revocantur, futura incerta sunt; praesentia sola nostra sunt.»

“The past cannot be recalled, the future is uncertain; only the present is ours.”
— Seneca, *De Brevitate Vitae*

Abstract

Context and motivation. In the era of multi–messenger astrophysics, the landmark joint detection GRB170817A/GW170817 established the first, and so far only, confirmed association between a short Gamma Ray Burst and a binary neutron star merger. This event demonstrated that precise, near, simultaneous localization of electromagnetic and gravitational wave signals is essential for unambiguous source identification. Maximizing the scientific return of future joint detections requires continuous high-energy monitoring and rapid, accurate localization capabilities, naturally enabled by distributed, constellation-style missions that triangulate events through inter-satellite time-of-arrival delay measurements.

Nanosatellites and the low–count challenge. Nanosatellites drastically reduce barriers to space science through small mass, rideshare compatibility, and COTS–based integration. However, their modest effective areas (typically $\lesssim 100\ \mathrm{cm}^2$) constrain time-domain analyses to a low-count Poisson-limited regime. In such conditions, fixed–bin light curves may either smooth out fast variability (if bins are too coarse) or become dominated by shot noise (if too fine), biasing traditional timing estimators.

Methods: Estimation of delays. This thesis develops a timing toolkit designed for such constraints. First, light curves are constructed using an *adaptive binning* approach that stabilizes the per-bin count statistics without sacrificing temporal resolution, enabling the recovery of variability across a wide range of timescales. Second, I introduce a simulation-free cross-correlation estimator that operates directly on photon time-of-arrival lists. By minimizing the Poisson imprint intrinsic to flux randomization techniques, which repeatedly simulate from noisy templates, this method yields stable, unbiased lag estimates with near nominal coverage. Across photon count regimes, the *Modified Double Pool* method improves both accuracy and computational efficiency, with only the expected $\sqrt{2}$ precision penalty due to halving-reshuffle process.

Constellations for localization: performance and FoV. Localization accuracy is quantified for realistic three-satellite constellations, exploring platform sizes (3U–27U) and baselines ranging from low–Earth separations to Earth–L1 distances. Using non–imaging, wide field detectors, such configurations achieve quasi–hemispherical instantaneous coverage, a major advantage over

the ~ 1.4 sr field of view and Earth occultation of low Earth orbit instruments such as *Swift*/BAT. A triplet with $A_{\rm eff} \sim 360$ cm^2 per node and baselines of $\sim 2.5 \times 10^6$ km can reach *Swift*/BAT class localizations, while smaller nodes still achieve arcminute accuracy under favorable geometry and count rate. These results should enable rapid, targeted follow-up by partner facilities with minimal tiling, reducing the risk of missing the earliest and most informative transient phases.

SpIRIT in–orbit operations. This thesis also encompasses the in–orbit activities of the 6U *SpIRIT* satellite, focusing on the onboard *HERMES* payload. These include commissioning and first-light operations, execution of the observing program, and an end-to-end data reduction pipeline producing time-tagged event lists, light curves, and preliminary spectra. Operational effects affecting background, calibration, timing, pointing, and data continuity are identified and mitigated through dedicated procedures. Bright X–ray targets are analyzed, including a tentative *HERMES* detection of the Sco X–1 spectrum, with detailed interval selection, background modeling, response application, and consistency checks.

On–orbit validation of microsecond timing. High precision timing is demonstrated using the 6U *SpIRIT* CubeSat equipped with a *HERMES* detector. During a single 730 s on–axis observation of the Crab pulsar, the canonical double-peaked profile was resolved at $> 5\sigma$, showcasing microsecond-level performance previously associated only with flagship observatories.

Fundamental physics: limits on Lorentz invariance violation. The timing framework is further applied to tests of Lorentz Invariance Violation (LIV) through energy-dependent photon propagation in GRBs. By measuring inter-band lags between MeV–GeV light curves and a keV reference band, and statistically disentangling intrinsic source delays from propagation-induced terms, I obtain, at 95% C.L., first-order lower limits of $E_{\rm QG} > 1.25 \times 10^{16}$ GeV (subluminal) and $E_{\rm QG} > 2.49 \times 10^{16}$ GeV (superluminal).

Faint source timing: ULX hard lags. Beyond Gamma Ray Bursts, the toolkit proves effective for extremely faint sources such as Ultraluminous X–ray sources (ULXs). Statistically significant *hard lags* are detected both over full *XMM–Newton* exposures and within windows of $\sim$10 ks, supporting a scenario of inward propagation of accretion rate fluctuations in super Eddington, geometrically thick flows. The measured lags, of order 1 ks, correspond to characteristic distances of ~ 0.001 AU from the compact object, where the soft X-Ray emission mechanisms remain under investigation.

Contents

List of Figures

List of Tables

1

Introduction

1.1.0 Story Pills

In the 1960s, during the Cold War, the United States launched the Vela satellites to monitor compliance with the Partial Nuclear Test Ban Treaty by detecting potential atmospheric nuclear detonations through γ-ray detection. On July 2, 1967, the Vela satellites recorded a brief and intense burst of gamma radiation that did not correspond to the known signatures of terrestrial nuclear events. These signals, detected intermittently and originating from various directions in space, could not be attributed to solar or geophysical phenomena. Following extensive analysis, the findings revealed the existence of a previously unknown class of high-energy astrophysical transients, later termed Gamma-Ray Bursts (GRBs).

The unexpected nature of GRBs spurred a decades-long observational effort to understand their origins and mechanisms. Beyond GRBs, the advent of X-Ray and gamma-ray astronomy also unveiled a rich zoo of high-energy sources, including accreting or isolated neutron stars, magnetars, black hole binaries, and active galactic nuclei. These discoveries transformed our understanding of the high-energy sky, revealing extreme environments where matter and radiation interact under the most intense gravitational, magnetic, and relativistic conditions.

A significant breakthrough came with the Burst and Transient Source Experiment (BATSE) aboard the Compton Gamma Ray Observatory (CGRO) in the 1990s, which provided the first compelling evidence that GRBs are isotropically distributed across the sky, strongly suggesting a cosmological origin. This realization marked a paradigm shift, firmly placing GRBs among the most powerful extragalactic phenomena.

The definitive breakthrough in understanding the origin of Gamma-Ray Bursts (GRBs) was achieved with the Italian–Dutch satellite BeppoSAX, whose unique instrumental configuration combined wide-field gamma-ray detection with narrow-field X-Ray imaging. In 1997, BeppoSAX enabled, for the first time, arcminute-level localizations of GRBs within a few hours of the prompt emission, leading to the discovery of their X-Ray afterglows. This capability marked a paradigm shift: rapid and accurate localizations made it possible to trigger systematic multi-wavelength follow-up observations, resulting in the identification of optical and radio counterparts and, crucially, the measurement of redshifts for GRB host galaxies. These observations provided unambiguous proof that GRBs originate at cosmological distances, with isotropic-equivalent energies far exceeding those of any Galactic phenomenon.

A milestone in the study of short GRBs was reached on August 17, 2017, with the joint detection of the gravitational wave (GW) signal GW170817 and the short GRB 170817A, observed by the LIGO/Virgo interferometers and the Fermi/INTEGRAL satellites. This event provided the first direct evidence that at least some short GRBs originate from binary neutron star (BNS) mergers and marked the beginning of multi-messenger astrophysics, in which GWs and electromagnetic radiation jointly inform our understanding of cosmic events. Beyond validating theoretical models, it highlighted the challenge of capturing such rare and rapidly evolving phenomena and underscored the importance of coordinated, multi-wavelength monitoring strategies.

Astrophysical transients such as GRBs are inherently unpredictable in both time and location. Their brief duration and fast-fading emission require instruments capable of continuous, wide-field monitoring and low-latency response. Joint detections depend on the integration of multiple GW detectors and γ-ray and X-Ray satellites, each with its own field of view, sensitivity, and cadence. The X-Ray afterglow often bridges the gap between the prompt γ-ray emission and longer-wavelength signatures, offering a more accurate localization and key diagnostics on the energetics, geometry, and environment of the source, but it occurs after the catastrophic merger that generates both the GW and electromagnetic signals.

However, current X-Ray facilities face a trade-off between sensitivity and sky coverage. High-resolution instruments such as Chandra and XMM-Newton offer excellent sensitivity but narrow fields of view, limiting their use in prompt transient discovery. Conversely, wide-field monitors like Swift-BAT or MAXI provide broader coverage, but with lower spatial resolution or reduced temporal sampling. These limitations constrain our ability to systematically detect and study joint events in real time. Looking ahead, the next generation of wide-field, high-sensitivity X-Ray missions will be critical to unlock the full potential of multi-messenger astrophysics, enabling rapid discovery and characterization of transient phenomena across the Universe. In particular, emerging missions based on the concept of distributed architecture offer pivotal improvements in the localization accuracy of variable sources. Their coordinated design enhances temporal analysis and lag estimation across wide baselines, thereby expanding our ability to reconstruct the spatial and dynamical properties of high-energy transients.

1.2.0 The Advent of the Nanosatellite Era in Astrophysics

The distribution in Figure 1.1 shows how, over the past decade, nanosatellite launches have grown exponentially, driven by faster development cycles, markedly lower costs, and reduced mission risk. Standardized form factors (e.g., CubeSats) enable rideshare opportunities that cut access-to-space expenses, while the low mass eases mechanical and environmental constraints, mitigating launch risk. Compared with bespoke, large satellites, nanosatellites are simpler and quicker to design, integrate, and qualify, supporting rapid design–test–flight iteration and timely

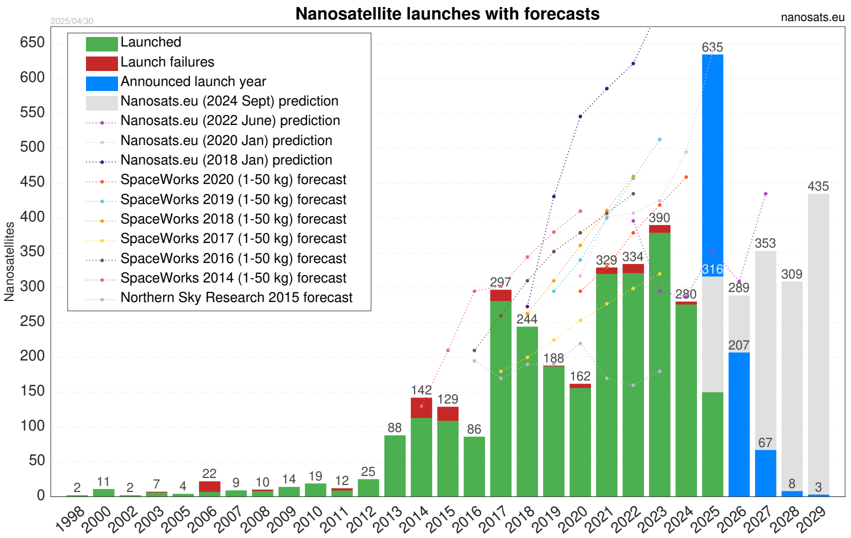

Figure 1.1: Annual nanosatellite launches and forecasts (1998–2029). Bars show successful launches and launch failures; lighter bars indicate missions with an announced launch year. Overlaid curves trace forecasts from Nanosats.eu (2018–2024), SpaceWorks (2014–2020), and NSR (2015), illustrating the post-2013 surge and the sustained cadence expected through 2029. Data/forecasts: Nanosats.eu (30 Apr 2025).

scientific return.

A key enabler is the use of Commercial Off-The-Shelf (COTS) components. Leveraging COTS reduces procurement time and cost, streamlines payload integration, and makes space-based astrophysics accessible even to single research institutes. Unlike custom platforms that demand extensive one-off manufacturing, COTS-based nanosatellites can be assembled rapidly and affordably while still achieving science-grade performance. Moreover, once an efficient payload and spacecraft design has been developed and flight-demonstrated, assembly becomes effectively plug-and-play: the integration and test chain can be replicated with minimal customization, enabling fleets or iterative upgrades and drastically shortening time-to-science.

In this context, nanosatellites are especially well-suited to time-domain astronomy, where sky coverage, rapid cadence, and technological iteration are decisive, and where distributed constellations can deliver prompt detections, cross-checks, and triangulation of transient sources. To illustrate their effectiveness, if one CubeSat in a constellation fails, it can often be replaced in well under a year and re-flight at modest cost (cf. Falcon Users Guide), thanks to short build/integration cycles and the high cadence of low-cost rideshares to polar Sun-synchronous orbits (SSO). SpaceX's SmallSat Rideshare Program lists \$325,000 for up to 50 kg to SSO, plus \$6,500 per additional kg (SpaceX). Dedicated SSO rideshares fly multiple times per year (e.g., Transporter–14 on 14 March 2025; mission page, coverage). Launch risk is correspondingly low given Falcon 9's demonstrated reliability (∼99% mission success; NASASpaceflight).

1.3.0 The Low-Statistics Problem in CubeSat Timing

While nanosatellites offer compelling advantages, those same attributes also constrain their scientific return. Because of the small volume factor ($\lesssim 1\mathrm{m}^3$), instrument effective areas are typically $\lesssim$ 100 cm^2, yielding low photon counts and background-dominated regimes in many observations. In time-domain studies, these challenges limit standard fixed-bin-size methods devised for large-area instruments: coarse binning smears fast variability and time lags, whereas fine binning inflates noise and biases inference under non-Gaussian (Poisson) statistics.

This thesis tackles the limitations of time-domain astronomy for variable high-energy sources by developing physics-informed, Poisson-aware timing techniques tailored to the low-count regime. Although motivated by nanosatellite observations, the methods apply broadly wherever low photon statistics impede robust inference (e.g., faint sources) and integrate seamlessly with event-level pipelines and future distributed-constellation datasets. The approach centers on adaptive binning for temporal analysis, which captures intrinsic variability in X-Ray and γ-ray signals, increases sensitivity to weak or short-lived features, improves time-lag measurements, and yields a more faithful reconstruction of transient evolution. Applied to real observations, these techniques provide new insight into the temporal behavior of extreme phenomena and deliver more robust tools for analysis and localization in the time domain.

2

Gamma Ray Burst

Summary

This chapter frames the research question within high-energy, time-domain astrophysics by moving from context to consequence. It begins with a concise overview of the stellar life cycle to show how different endpoints naturally produce explosive transients, then traces how the rise of X-Ray and γ-ray astronomy transformed that picture, culminating in the discovery and interpretation of GRBs and a new, dynamic view of the high-energy sky. From there, the focus shifts to the modern multi-messenger era, where rapid, coordinated observations across the electromagnetic spectrum, together with gravitational-wave and neutrino detections, enable prompt localization, secure host association, and physically grounded modeling. Within this framework, GRBs emerge as natural laboratories for central-engine physics (collapsars and compact-object mergers), jet geometry and energetics, radiation mechanisms, and environment coupling; their fast variability encodes properties of both the engine and the circumburst medium. The chapter thus sets an evolutionary backdrop and then builds toward a quantitative, multi-messenger synthesis that motivates the methods and results developed in the remainder of the thesis.

2.1.0 The Path to Collapse: From Equilibrium to GRB Progenitors

Stars are self-gravitating, quasi-hydrostatic plasma spheres powered by nuclear burning. Their long-term evolution is governed by the competition between gravity and pressure support, while their *endpoints* are set by what happens after nuclear fuel is exhausted and thermal pressure can no longer oppose gravitational collapse. At that stage, the (final) core mass plays the central role in determining the fate of the dying star.

The core contraction of low- and intermediate-mass progenitors, whose star mass is $\lesssim 8M_\odot$, is arrested by electron degeneracy pressure. This implies a maximum stable mass, the Chandrasekhar limit, $M_{\rm Ch} \simeq 1.4\,M_\odot$ (Chandrasekhar 1931), and the formation of a white dwarf (WD), an Earth-sized remnant ($R \sim 10^4$ km, decreasing with mass) composed of ionized nuclei (typically C/O, or O/Ne/Mg at higher masses) embedded in a sea of degenerate electrons.

The Chandrasekhar limit follows from equating hydrostatic support by electron degeneracy to the characteristic gravitational stress and can be derived from elementary considerations. For a zero–temperature Fermi gas with electron number density n_e, the Fermi momentum is $p_F = \hbar(3\pi^2 n_e)^{1/3}$ and the equation of state is

$$P_{\rm deg}^{\rm NR} = \frac{\hbar^2}{5m_e}(3\pi^2)^{2/3} n_e^{5/3} \quad \text{(nonrelativistic)}, \qquad P_{\rm deg}^{\rm UR} = \frac{\hbar c}{4}(3\pi^2)^{1/3} n_e^{4/3} \quad \text{(ultrarelativistic)}. \tag{2.1}$$

Balancing $P_{\rm deg}$ against the gravitational stress $P_g \sim GM^2/R^4$ (with M and R the WD mass and radius) yields the familiar scaling $R \propto M^{-1/3}$ in the nonrelativistic regime: increasing mass compresses the star, raising the electron Fermi momentum. Once the electrons become relativistic, the softer scaling $P_{\rm deg} \propto n_e^{4/3}$ can no longer sustain arbitrarily large M, and a maximum mass emerges. Again, accounting for the electron mean molecular weight μ_e, one finds $M_{\rm Ch} \simeq 5.83\,\mu_e^{-2}\,{\rm M}_\odot \simeq 1.4\,{\rm M}_\odot$ for $\mu_e \approx 2$.

Thus, as $M \to M_{\rm Ch}$ the equilibrium radius tends to zero and the configuration becomes unstable to further collapse, marking the transition beyond which electron degeneracy alone cannot support the star.

Stars whose core mass exceeds this limit can give rise to some of the most catastrophic events in the Universe, including GRBs.

2.2.0 Gamma Ray Burst discovery

The 1963 Nuclear Test Ban Treaty, a key milestone in the easing of Cold War tensions, was established to limit the spread of nuclear weapons by prohibiting tests in both the atmosphere and outer space. As part of the treaty's enforcement, the U.S. Air Force initiated the launch of a series of surveillance satellites in October 1963, known as the "Vela" program—named after the Spanish verb velar, meaning "to watch". These satellites were designed to detect the brief and intense gamma-ray flashes expected from unauthorized nuclear detonations. Although it remains

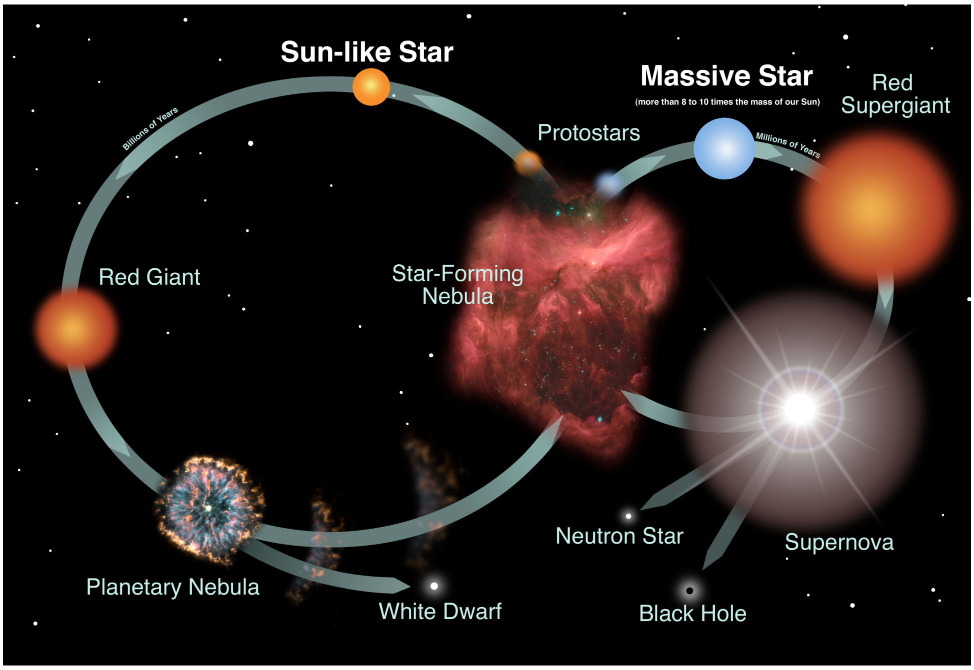

Figure 2.1: Star lifecycle as sketched by NASA (https://imagine.gsfc.nasa.gov/star_life)

uncertain whether the Vela satellites ever recorded any illegal nuclear tests, their deployment led to a groundbreaking scientific discovery: a new category of cosmic transients known as GRBs. The first GRB was observed on July 2, 1967, by the Vela 3 and 4 satellites as a sudden and bright burst of gamma radiation, illustrated in Figure 2.2.

Despite the Vela system's limited localization accuracy—relying on detection time differences between satellites and Earth occultation—it was soon suspected that the origin of these bursts was extraterrestrial. The characteristics of the signals, particularly their short duration, spectral features, and variability, clearly did not align with what would be expected from nuclear detonations. It was not until 1973 that R. W. Klebesadel and colleagues (Klebesadel et al. 1973) formally reported the discovery of this and 15 additional similar events to the scientific community. This delay was due to the complexity of the data analysis and the necessity for further observations to support a cosmic or galaxy origin hypothesis. These unexpected bursts of gamma rays immediately captured the attention of the astronomical community. During the two decades that followed the announcement, the true nature of Gamma-Ray Bursts remained elusive, despite the deployment of several satellites and the publication of hundreds of theoretical models. GRBs were observed exclusively in the gamma-ray band and typically lasted only a few tens of seconds. The absence of reliable distance measurements and the limited sample size made it impossible to determine the total energy output or spatial distribution of these events, fueling an ongoing debate as to whether their origin was within the Milky Way or beyond.

A major advancement in the understanding of Gamma-Ray Bursts came with the launch of the Compton Gamma-Ray Observatory (CGRO) in 1991. Among its key instruments were the BATSE and the Energetic Gamma-Ray Experiment Telescope (EGRET). Together, they enabled the detection of more than 2,700 GRBs and provided extensive spectral coverage of gamma-ray emissions, spanning from approximately 10 keV to several GeV. One of the most striking findings

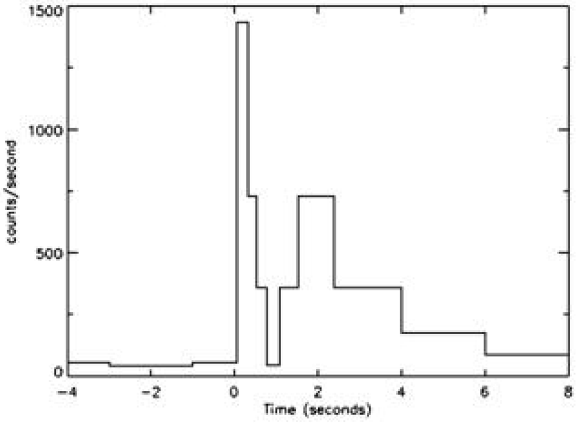

Figure 2.2: Light curve of the first Gamma-Ray Burst recorded by the Vela 3 and 4 military satellites. The identification of these short and intense γ-ray flashes, originating from beyond Earth, was not disclosed to the public until seven years later (Klebesadel et al. 1973). Image credit: R. Klebesadel, I. Strong R. Olson

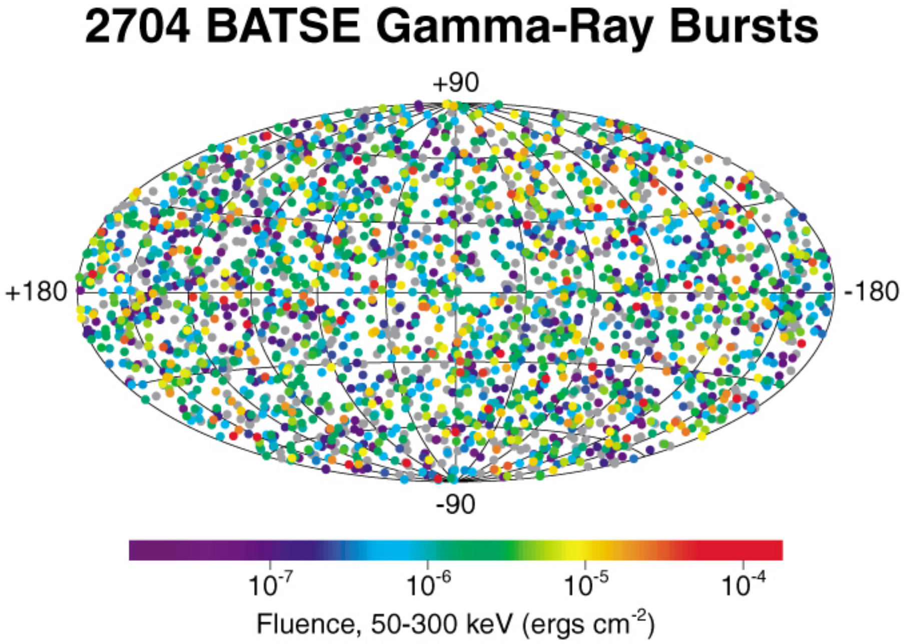

Figure 2.3: Sky distribution of 2,704 Gamma-Ray Bursts recorded by BATSE over a 9-year period. Each point is color-coded according to the measured fluence—that is, the flux integrated over the duration of each individual burst. The data reveal a uniform, isotropic distribution across the sky, consistent across all fluence ranges. Image source: https://heasarc.gsfc.nasa.gov/docs/cgro/batse/

was the isotropic distribution of GRBs across the sky (see Figure 2.3), suggesting no strong preference for any particular direction in space. Moreover, the cumulative fluence distribution of these events was found to deviate from the -3/2 power-law slope expected for a homogeneous and uniformly distributed local population. These observations offered compelling evidence in favor of an extragalactic origin for GRBs. Nevertheless, the true nature of GRBs continued to defy explanation, prompting a central question in high-energy astrophysics: which astrophysical environments give rise to these powerful and transient phenomena? This uncertainty culminated in the renowned "Great Debate" of 1995, which featured Bohdan Paczyński, advocating for a

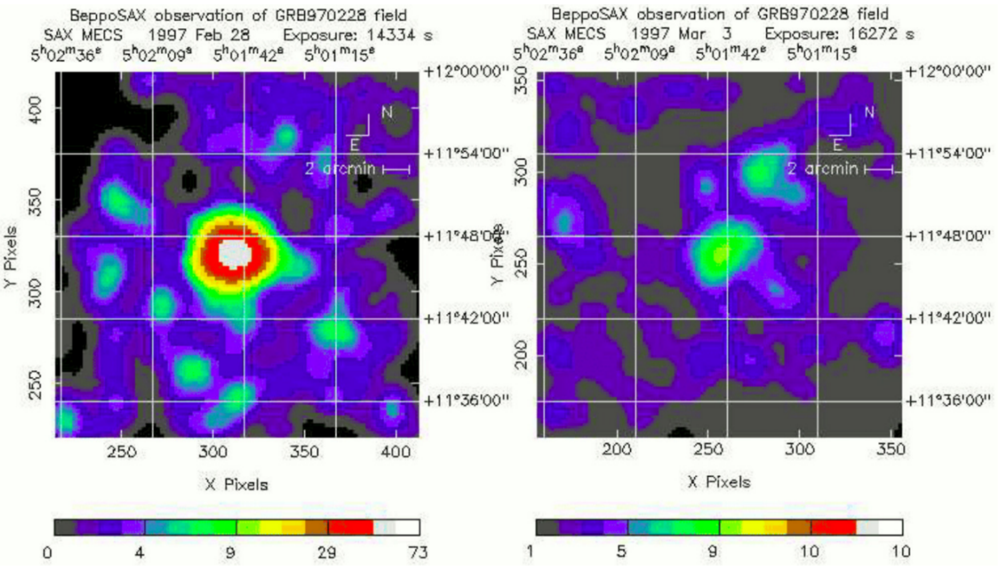

Figure 2.4: BeppoSAX localization of GRB 970228. The image on the left captures X-Ray emission within the γ-ray detection region roughly 8 hours after the initial burst, while the right image shows the same region observed approximately 3 days later. A fading X-Ray source was distinctly observed at the GRB site, representing the first confirmed identification of an afterglow in the X-Ray band. Credit to Costa et al. (1997).

cosmological origin (Paczynski 1995), and Donald Lamb, supporting a Galactic interpretation (Lamb 1995). If GRBs originated from extragalactic distances, their extreme luminosity and rapid variability would suggest the existence of a longer-lasting emission, known as the afterglow (literally after the glow), observable at lower energies, from X-Rays to radio wavelengths, as treated in Paczynski & Rhoads (1993). At the time, however, no such multiwavelength counterparts were identified. This was largely due to the large positional uncertainties associated with BATSE detections, which produced error regions containing numerous potential sources, thereby hindering effective follow-up observations.

At 02:58 UT on February 28, 1997, the Wide Field Cameras (WFC) aboard BeppoSAX successfully localized GRB 970228 (Costa et al. 1997) within a narrow error region of just 3 arcminutes. This precise localization enabled the subsequent discovery of the event X-Ray afterglow a few hours later, when the satellite Narrow Field Instruments (NFI) were directed toward the same region of the sky (see Figure 2.4). Shortly thereafter, van Paradijs et al. (1997) identified the optical afterglow using the 4.2-meter William Herschel Telescope in La Palma, approximately 1.9 days after the burst. Although no spectroscopic data were obtained at the time, and thus its redshift remained undetermined, this event marked the beginning of afterglow studies.

A major step forward occurred just months later, on 8 May 1997, when BeppoSAX detected GRB 970508 at 21:42 UT. This burst was followed by the detection of a radio afterglow by the Very Large Array (VLA) (Frail et al. 1997), along with its optical counterpart. The Metzger et al. (1997) Spectroscopic observations with the Keck Telescope in Figure 2.5 revealed absorption features corresponding to a redshift of z = 0.835, offering the first direct confirmation that GRBs occur at cosmological distances.

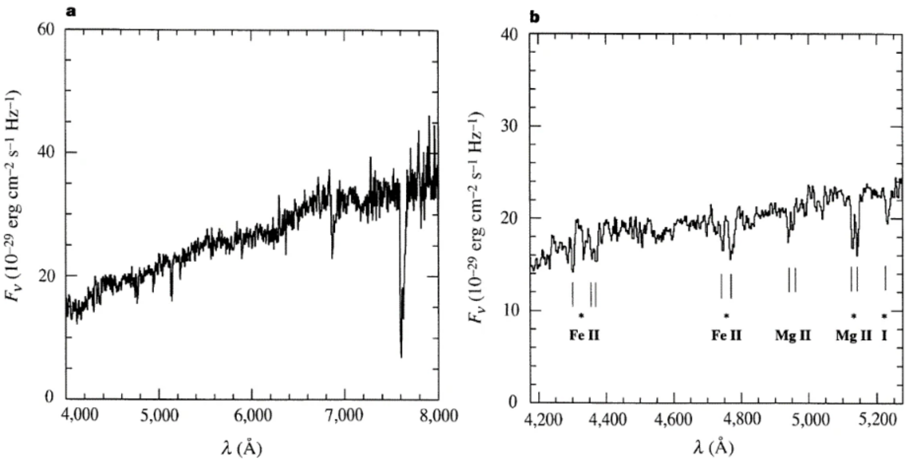

Figure 2.5: (a) Full optical spectrum of GRB 970508 take with the Keck Telescope; (b) Zoomed-in view of a selected wavelength range, highlighting prominent absorption lines with their identifications. Features marked with an asterisk correspond to an absorption system at redshift $z = 0.835$, while the remaining lines are associated with a system at $z = 0.767$. The spectrum has been smoothed using a three-pixel boxcar filter. Several additional weak features (not shown here) have also been tentatively attributed to the $z = 0.767$ system. F_ν denotes the flux density, and λ represents the wavelength in Å. Credit to (Metzger et al. 1997)

In the decades following their initial discovery, Gamma-Ray Bursts have remained a topic of intense and growing interest within the astrophysical community. Numerous X-Ray and γ-ray observatories have been launched to investigate the remarkable properties of these cosmic phenomena. Notable missions include KONUS/WIND (1994), the High Energy Transient Explorer (HETE-2, 2000), the INTErnational Gamma-Ray Astrophysics Laboratory (INTEGRAL, 2002), the Neil Gehrels Swift Observatory (Swift, 2004), the Astro-Rivelatore Gamma a Immagini Leggero (AGILE, 2007), and the Fermi Gamma-ray Space Telescope, launched in 2008.

2.3.0 The compactness problem

Early works by Ruderman (1975) and Schmidt (1978) argued that, if GRBs were at cosmological distances, the observed fluences imply enormous isotropic energies and very compact emission regions (from millisecond variability), which together make the source highly opaque to $\gamma\gamma \to e^+ + e^-$. Yet GRB spectra are non–thermal and extend to high energies, implying the source must be optically thin. This tension is the *compactness problem*.

For a burst with fluence F at distance D, the isotropic–equivalent energy is

$$E_{\rm iso} \simeq 4\pi D^2 F,$$

and variability on a timescale δt limits the source size to

$$R \lesssim c\,\delta t.$$

The spectra contain high–energy photons that can annihilate with softer photons,

$$\gamma + \gamma \to e^+ + e^-, \qquad E_1 E_2 (1 - \cos\theta) \geq 2(m_e c^2)^2,$$

so the average optical depth to pair production, $\tau_{\gamma\gamma}$, is (very roughly) large for the compact, luminous non-relativistic source.

Relativistic resolution. If the emitting region moves toward the observer with Lorentz factor $\Gamma \gg 1$, several effects drastically reduce $\tau_{\gamma\gamma}$: (i) observed photon energies are blueshifted, so comoving energies are lower by $\sim \Gamma^{-1}$, pushing many pairs below threshold; (ii) the photon field is beamed into small angles ($\theta \sim 1/\Gamma$), reducing $(1 - \cos\theta)$; (iii) the variability constraint relaxes to a larger emission radius,

$$R \lesssim \frac{\Gamma^2 c\, \delta t}{1+z},$$

which lowers the photon density in the comoving frame.

Combining these factors yields a strong Γ–dependence of the opacity (schematically $\tau_{\gamma\gamma} \propto \Gamma^{-a}$ with $a \gtrsim 4$ for typical spectra). For standard GRB parameters and high–energy spectral index $\alpha \sim 2$, one finds that

$$\tau_{\gamma\gamma} < 1 \quad \text{for} \quad \Gamma \gtrsim 10^2,$$

which allows the escape of high–energy γ–rays and resolves the compactness problem. Empirically inferred Γ values from afterglow and prompt–emission studies are indeed of this order or larger, fully consistent with a cosmological origin for GRBs (Kumar & Zhang 2015; Piran 1999; Lithwick & Sari 2001).

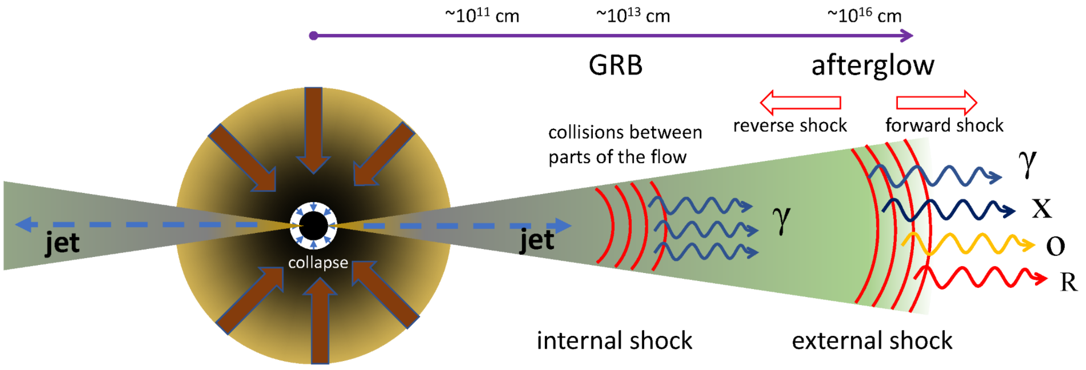

Figure 2.6: This artistic illustration depicts a representative scenario for the Fireball model depicting the various emission components in a GRB (Dado & Dar 2019).

2.4.0 The fireball model

By assuming the relativistic effects discussed so far, it is now useful to summarize the key elements of the generic GRB model that have historically emerged. The compactness problem demands ultra-relativistic motion, implying that the outflow must have a Lorentz factor of at least $\Gamma \gtrsim 100$. Aiming to constrain the total energy budget, it is natural to attribute the observed gamma-ray emission to the deceleration of a relativistic energy flow, occurring at a stage when the emitting particles are still moving at relativistic speeds.

This leads to the original fireball model in which a GRB is understood as a multi-stage phenomenon consisting of (Piran 1999):

1. The presence of a compact object, most commonly a black hole, surrounded by an accretion disk, forms the so-called *accretion-powered central engine*. This engine accretes infalling/fallback matter and can launch two highly relativistic, collimated outflows (jets/fireballs). Their composition may include a baryonic component and/or be magnetically dominated, with photons and pairs produced within the dissipative regions of the flow. A *magnetar-powered* engine is an alternative: a rapidly rotating, highly magnetized neutron star (a millisecond magnetar), where the outflow is powered by the extraction of rotational energy (spin-down) mediated by intense magnetic fields (Bucciantini et al. 2009).

2. Internal relativistic shocks occur when multiple fireballs, ejected at different velocities by the central engine, collide. These shocks convert part of the kinetic energy into thermal energy, accelerating particles and giving rise to the prompt gamma-ray emission observed over short timescales.

3. As the fireball expands and interacts with the circumburst medium, it decelerates, producing an external relativistic shock. This interaction leads to the long-lasting afterglow emission, which is observed at longer wavelengths, including X-Rays, optical, and radio.

An artistic representation that illustrates the fireball framework of Gamma-Ray Bursts is shown in Figure 2.6.

While this framework has played a crucial historical role in shaping our understanding of GRB prompt emission, it is no longer regarded as the most compelling model in light of more recent theoretical and observational developments. In particular, the intrinsically low radiative efficiency of internal shocks and their difficulty in reproducing key spectral properties, such as the observed low-energy photon indices and the presence of quasi-thermal components, have motivated the investigation of alternative scenarios (e.g. Kumar & Zhang 2015). Photospheric emission models have therefore gained increasing support, as they naturally predict high efficiencies and can accommodate both thermal and non-thermal spectral components through sub-photospheric dissipation processes (Pe'er et al. 2006; Rees & Meszaros 2005; Beloborodov 2010). In parallel, models invoking weakly magnetized or hybrid jets have emerged as viable frameworks, in which magnetic fields play a non-negligible role in the jet dynamics and energy dissipation. A notable example is the Internal-Collision-induced Magnetic Reconnection and Turbulence (ICMART) model, where internal collisions trigger magnetic reconnection events that efficiently power the prompt emission (Zhang & Yan 2010). These advances suggest that GRB outflows are likely more complex than originally envisioned, and that the standard hydrodynamic fireball/internal-shock model should be viewed as one possible scenario among several, rather than the prevailing paradigm.

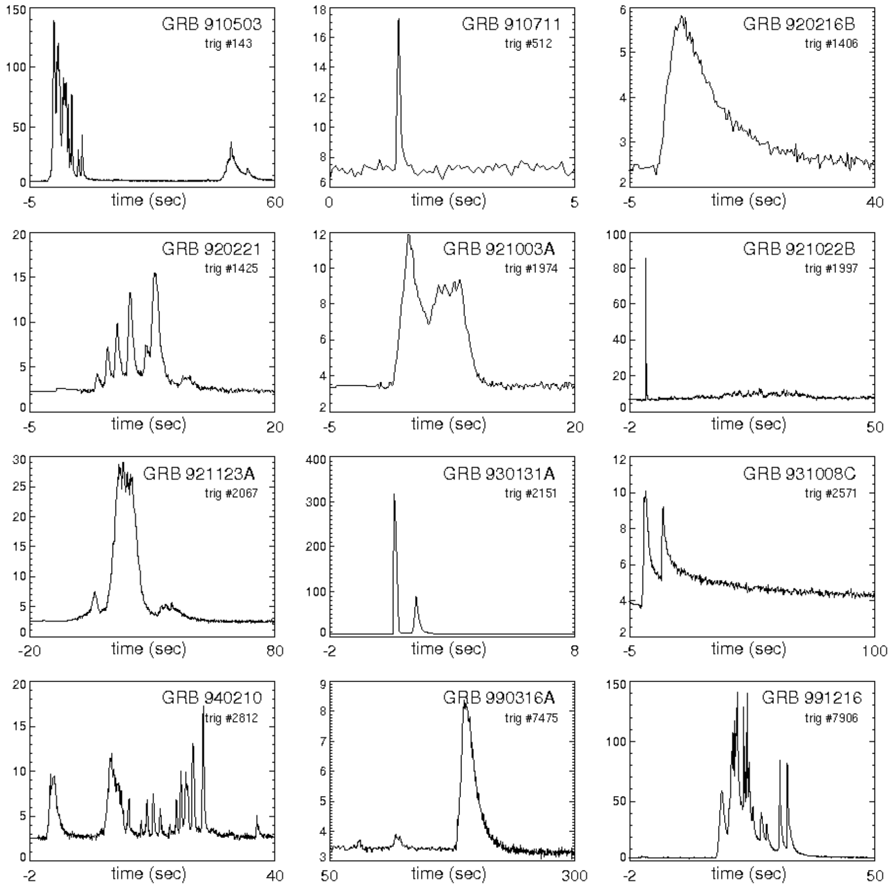

Figure 2.7: Light curves of 12 bright Gamma-Ray Bursts observed by BATSE, part of NASA Compton Gamma-Ray Observatory. Data sourced from the public BATSE archive: `http://gammaray.msfc.nasa.gov/batse/grb/catalog/`

2.5.0 GRB prompt phase

The prompt phase is the initial and most energetic stage of a GRB, characterized by an intense flash of gamma radiation that lasts from a few milliseconds to several minutes. In the historical framework, the emission arises from the *internal shock scenario* (Rees & Meszaros 1992), in which the central engine injects energy in an unsteady way along the jet axis, producing multiple relativistic shells (fireballs) composed of baryons, electrons, and high-energy photons.

The high internal pressure inside each shell drives rapid acceleration, and internal energy is converted into bulk kinetic energy. The ejecta expand and achieve highly relativistic speeds, with Lorentz factors $\Gamma \sim 10^2$, consistent with the compactness constraints discussed earlier.

Following Zhang & Mészàros (2004), the evolution begins with a hot, radiation–dominated plasma launched from the compact central engine at an initial radius $R_0 \gtrsim 10^7$ cm. Photons and electron–positron pairs are initially in thermal equilibrium. As energy injection continues, the fireball expands adiabatically, and the bulk Lorentz factor Γ increases approximately linearly with radius until it saturates at a terminal value Γ_0 (Panaitescu et al. 1997).

Under these conditions, the temperature can reach $T \sim 10^{10}$ K, which implies an abundant pair population in equilibrium with photons. This motivated the *hot fireball* model proposed by Paczynski (1986) and Goodman (1986), in which a sudden energy release within a compact

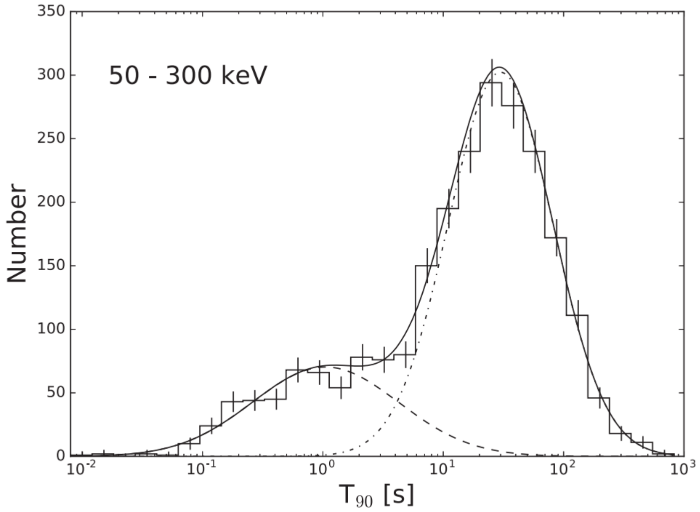

Figure 2.8: T90 distributions and for the events given in the 4th Fermi/GBM catalogue (von Kienlin et al. 2020).

region creates an optically thick mixture of photons and leptons. As the fireball expands, the temperature drops, and photons can escape. A purely radiative fireball would produce a thermal spectrum, which does not match the non-thermal prompt emission observed in GRBs.

If the outflow carries a small but non-negligible baryon load, part of the internal energy is converted into baryonic bulk kinetic energy (Shemi & Piran 1990). To reach the ultra-relativistic velocities inferred for GRB jets, the baryon loading must remain low; however, the acceleration mechanism and subsequent jet dynamics depend sensitively on the jet magnetization. In weakly magnetised, baryon-dominated outflows, bulk acceleration is primarily driven by thermal pressure, and kinetic energy can be dissipated through internal shocks, re-energizing electrons that radiate non-thermal emission via synchrotron and inverse Compton processes (Rees & Meszaros 1994; Daigne & Mochkovitch 2000). In contrast, in magnetised outflows, acceleration may proceed through the gradual conversion of magnetic (Poynting) flux into kinetic energy, leading to different dynamical profiles and dissipation regimes (Drenkhahn & Spruit 2002; Komissarov et al. 2009). In such cases, energy release is expected to occur predominantly through magnetic reconnection, rather than purely hydrodynamic shocks, and the resulting prompt emission may reflect a combination of magnetic dissipation and radiative processes (Zhang & Yan 2010).

Individual pulses in the prompt phase typically show approximately constant durations over time (Ramirez-Ruiz & Fenimore 2000), which suggests localized and quasi-stationary emitting regions. The corresponding emission radius R_γ is constrained, according to the emission model (e.g., internal shocks vs photosphere scenario vs magnetic dissipation scenarios), to $R_\gamma \sim 10^{12}$ to 10^{16} cm (see Figure 2.6), based on observed variability timescales and the requirement that the region is optically thin to allow photon escape.

2.5.1 Temporal properties

The light curves associated with the prompt emission of GRBs exhibit a high degree of complexity and variability, with substantial differences observed from one burst to another. Only

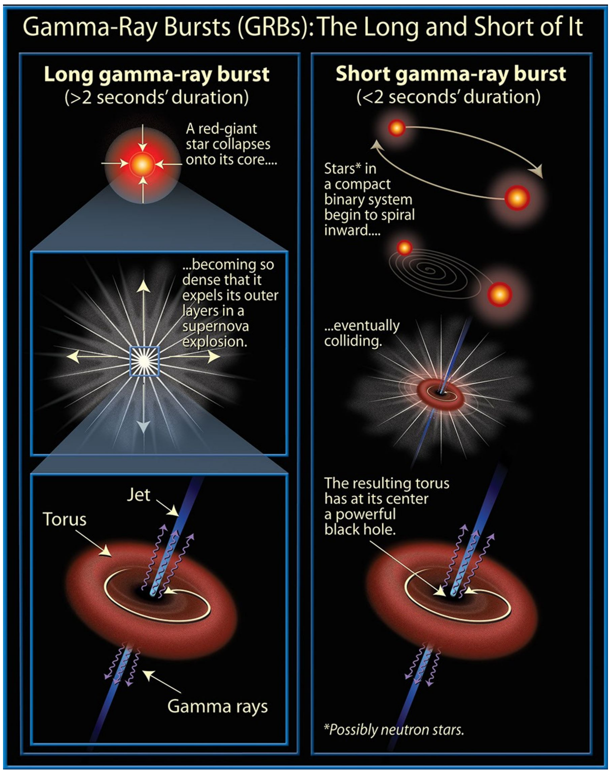

Figure 2.9: Schematic illustration of the two favored progenitor scenarios for gamma–ray bursts: long GRBs (left), from the collapse of massive stripped stars, and short GRBs (right), from compact object mergers. Credit: NASA and A. Feild (STScI). Source: https://hubblesite.org/contents/media/images/GRB_progenitor.

a small number of events display similar temporal profiles. The Figure 2.7 clearly highlights the intrinsic complexity and pronounced lack of uniformity characteristic of this phenomenon.

They are typically characterized by rapid variability, with timescales ranging from a few milliseconds up to several tens or even hundreds of seconds as treated in Camisasca et al. (2023). To quantify the duration of a burst, the parameter T_{90} has traditionally been used. It represents the time interval during which the cumulative flux rises from 5% to 95% of the total detected flux, measured relative to the background level—thus encompassing 90% of the burst's total emission.

The T90 distribution shown in Figure 2.8 is well described by a bimodal fit, with an overlap occurring around ~2 seconds. This empirical feature forms the basis for the historical classification of Gamma-Ray Bursts into short GRBs (T90 < 2 s) and long GRBs (T_{90} > 2 s).

As will be discussed throughout this thesis, the distinction between long and short GRBs extends well beyond the duration of the prompt emission. It also reflects fundamental differences in their astrophysical environments, progenitor systems, spectral properties, and other key characteristics.

2.5.2 Gamma Ray Burst progenitors

Exceeding the Chandrasekhar mass in the collapsing core is a *necessary* but not *sufficient* condition for a gamma–ray burst. The vast majority of core collapses produce ordinary core collapse supernovae without a successful relativistic jet. Long GRBs occur only when a rapidly

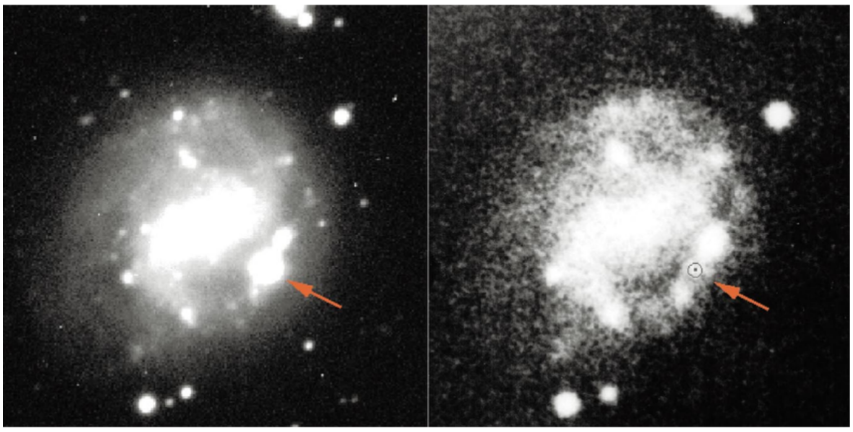

Figure 2.10: Side-by-side images documenting the discovery of SN 1998bw associated with the long GRB 980425. *Left:* discovery frame obtained with ESO's New Technology Telescope on 1998 May 1. *Right:* archival image from the UK Schmidt Telescope taken in 1976, where the supernova is absent. Adapted from Galama et al. (1998).

rotating, hydrogen/helium stripped massive progenitor (large C/O core) forms a central engine (BH + accretion disk or a millisecond magnetar) and the jet breaks out of the stellar envelope; even after beaming corrections, such events constitute $\ll 1\%$ of the CCSN rate (MacFadyen & Woosley 1999; Woosley & Bloom 2006; Woosley & Heger 2006; Guetta & Della Valle 2007; Wanderman & Piran 2010; Langer 2012). Short GRBs, in contrast, arise predominantly from compact object mergers rather than core collapse, as established by the binary neutron star event GW170817/GRB 170817A (Abbott et al. 2017).

The millisecond variability of the prompt emission implies, by causality, an emitting region only a few $\times 10^7$–10^8 cm across, naturally pointing to engines powered by compact remnants. In both the collapsar (long GRB) and merger (short GRB) channels, a stellar mass black hole (BH) surrounded by a transient accretion flow can drive a narrowly collimated, relativistic jet. Observational demographics reinforce this dichotomy: long GRBs favor actively star-forming hosts, whereas short bursts occur in a broad range of environments from ellipticals to irregulars (e.g. Berger 2014).

Long GRB

The *collapsar* paradigm—core collapse of a massive, stripped star forming a central engine and relativistic jet, was articulated by Woosley (1993), who identified Wolf–Rayet–like progenitors with masses $\gtrsim$ 10–15 $M_\odot$ as promising candidates. The first compelling empirical support arrived with the nearby event GRB 980425/SN 1998bw (Galama et al. 1998), where the supernova was directly seen in the host (cf. Figure 2.10: bright post–discovery source absent in the pre–explosion image). In the years that followed, additional GRB–SN associations were established: the supernova component typically emerges ~2–3 weeks after the prompt emission, producing a late–time "bump" in the optical afterglow light curve and broad spectral features characteristic of broad–lined Type Ic explosions.

The resolved hosts of long GRBs are predominantly irregular, strongly star–forming galaxies,

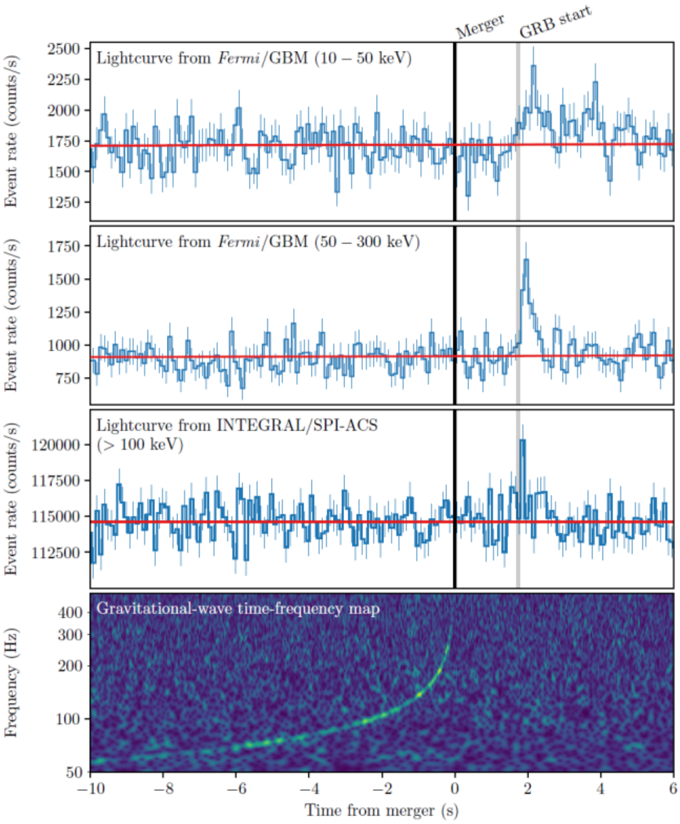

Figure 2.11: Joint detection of GW170817 and GRB 170817A. Panels 1–2: *Fermi*/GBM light curves in the 10–50 keV and 50–300 keV bands. Panel 3: *INTEGRAL* SPI–ACS count rate above 100 keV. Panel 4: time–frequency spectrogram of GW170817 obtained by coherently combining data from LIGO–Hanford and LIGO–Livingston. Adapted from the LIGO–Virgo Collaboration (https://www.ligo.org/science/Publication-GW170817GRB/).

and GRB sites concentrate in their brightest star–forming regions (Fruchter et al. 2006). Taken together, these photometric, spectroscopic, and environmental diagnostics firmly link many long GRBs to the deaths of massive stars and support the collapsar interpretation (Cano et al. 2017; Hjorth & Bloom 2011; Thoene & Fynbo 2007; Kulkarni et al. 1998).

Short GRB

Regarding short GRBs, the long–standing hypothesis is that they originate from the catastrophic coalescence of either BNS or NS–BH systems was first proposed by Eichler et al. (1989). Several population clues support this picture: short bursts are often found in galaxies with relatively low star–formation activity and at large projected offsets from their hosts, as expected for NS binaries that merge after long gravitational–wave–driven inspirals and natal–kick propagation (Fong et al. 2010). A decisive test came with GW170817/GRB 170817A: a short GRB detected ~1.7 s after the binary–neutron–star gravitational–wave signal on 2017 August 17 (Abbott et al. 2017), whose time–frequency map and γ–ray light curve are shown in Figure 2.11. Within hours, follow–up observations pinpointed the host galaxy NGC 4993 and revealed a bright optical/near–infrared counterpart (e.g., Coulter et al. 2017; Drout et al. 2017; Chornock et al. 2017); X–ray and radio emission emerged days later and were monitored for years (Troja et al.

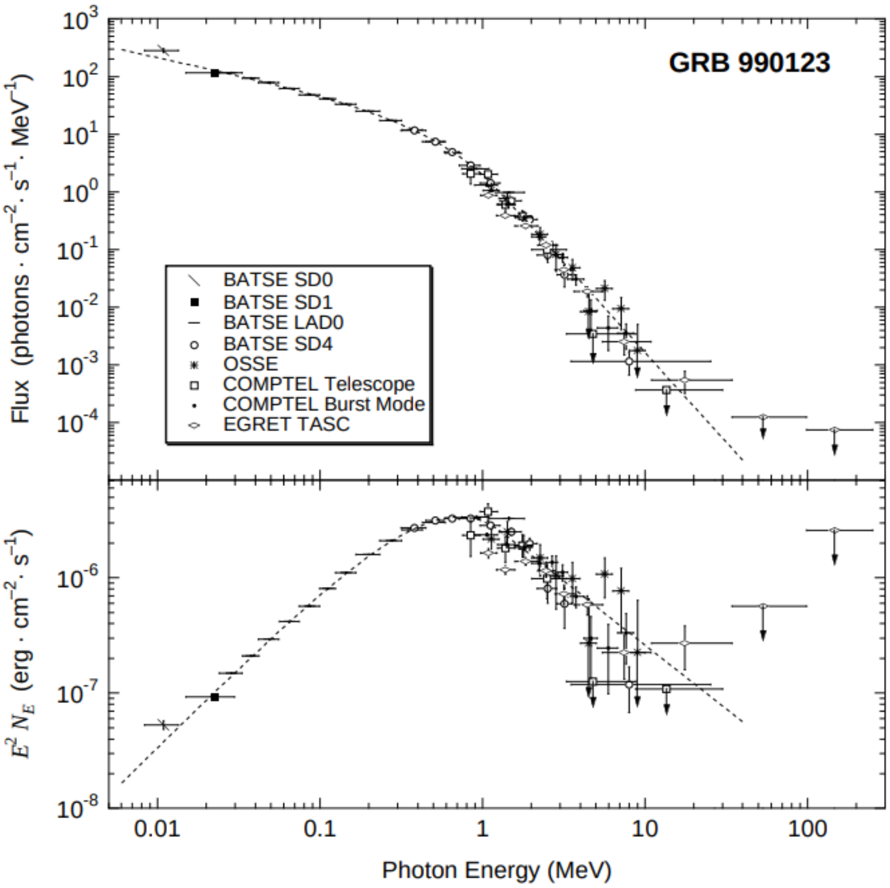

Figure 2.12: Deconvolved spectra of GRB 990123 from the CGRO detectors, fitted with a Band function. It is presented both as photon flux $N(E)$ and in spectral energy distribution units $E^2N(E) = \nu F_\nu$ (Briggs et al. 1999).

2017; Haggard et al. 2017; Alexander et al. 2018; D'Avanzo et al. 2018; Mooley et al. 2018; Hajela et al. 2019). This multi–messenger/multi-band data set established that (i) the merger ejected $\sim 0.03\,M_\odot$ of material at $v \sim 0.2\,c$, powering the first securely identified kilonova via the radioactive decay of r–process nuclei (Pian et al. 2017), and (ii) a relativistic, angularly structured jet was launched and viewed $\sim 20°$ off–axis, producing apparent super luminal motion in very long baseline interferometry, evidence for a narrowly collimated outflow (Mooley et al. 2018; Ghirlanda et al. 2019).

2.5.3 Spectral properties

The prompt emission spectrum of GRBs exhibits a non-thermal profile and is generally described by two power-law segments that are smoothly connected at a characteristic break energy. When the photon statistics are sufficiently robust, this spectral shape is commonly modeled using an empirical function known as the Band function (Band et al. 1993). This model consists of two smoothly joined power laws and is widely used to fit GRB prompt emission data. The photon number spectrum of the Band function is given by:

$$N(E) = A \begin{cases} \left(\frac{E}{100\,\mathrm{keV}}\right)^{\alpha} \exp\left(-\frac{E}{E_0}\right), & E < (\alpha-\beta)E_0 \\ \left(\frac{(\alpha-\beta)E_0}{100\,\mathrm{keV}}\right)^{\alpha-\beta} \exp(\beta-\alpha) \left(\frac{E}{100\,\mathrm{keV}}\right)^{\beta} & E \geq (\alpha-\beta)E_0 \end{cases}. \tag{2.2}$$

Here, $N(E)\,dE$ denotes the photon flux (in units of photons cm^{-2} s^{-1} keV^{-1}) within the energy interval dE. The parameter A is the normalization constant at 100 keV, also expressed in photons cm^{-2} s^{-1} keV^{-1}. The indices α and β represent the low-energy and high-energy power-law slopes, respectively, while E_0 indicates the break energy, measured in keV.

The spectral energy distribution (SED) is defined as $E^2N(E)$, which corresponds to the commonly used νF_ν representation. In this form, the prompt emission spectrum typically exhibits a pronounced peak, allowing for a straightforward identification of the peak energy.

The energy at which the νF_ν spectrum reaches its maximum is referred to as the peak energy, denoted by E_{peak}. In the context of the Band function, E_{peak} is related to the break energy E_0 by the expression:

$$E_{\mathrm{peak}} = (2+\alpha)\,E_0. \tag{2.3}$$

See Figure 2.12 for an example of the Band function fitted to GRB 990123 data. Based on BATSE observations, the typical spectral indices are approximately $\alpha \sim -1$ and $\beta \sim -2.25$, a result later confirmed by several observatories, (Guidorzi et al. 2010; Goldstein et al. 2012; Kaneko et al. 2006; Tsvetkova et al. 2021). The peak energy E_{peak} commonly centers around 250 keV, although it can span a broad range, from a few keV up to several MeV (Nava et al. 2011).

Nevertheless, the Band function is purely empirical and lacks a direct physical interpretation. Observations have also revealed the presence of additional spectral components that contribute to the overall complexity of GRB prompt emission. In particular, the fireball model predicts the existence of a thermal component, which is expected to arise at the photosphere, where the relativistic outflow becomes optically thin (Ryde 2005; Ryde et al. 2011).

The prompt emission is typically attributed to the dissipation of kinetic energy within a relativistic jet. Accounting for the widely accepted framework of the hot fireball model (Goodman 1986; Paczynski 1986), a compact and highly energetic region produces a photon-lepton-baryon plasma that expands relativistically.

Assuming a prompt isotropic-equivalent luminosity $L \sim 10^{52}$ erg s^{-1} and a variability timescale $\delta t \sim 10^{-2}$ s, the *initial* fireball temperature in the *comoving* frame can be estimated by treating the radiation field as approximately thermal at the engine scale:

$$T' \sim \left(\frac{L}{4\pi(c\delta t)^2\sigma_{\mathrm{B}}}\right)^{1/4} \sim 2\times10^{10}\ \mathrm{K}, \tag{2.4}$$

where σ_{B} is the Stefan–Boltzmann constant. Here $R_0 \sim c\,\delta t$ denotes the characteristic *initial* size of the energy release region set by the central engine variability, *before* significant bulk acceleration. The larger effective scale $\sim \Gamma^2 c\,\delta t$ that appears in compactness and transparency arguments pertains instead to the *later*, ultra-relativistic stage of the outflow, when relativistic light-travel-time and beaming effects control the observed variability. At $T' \sim 10^{10}$ K, photons are tightly coupled to leptons and baryons and remain close to thermal equilibrium. As the plasma expands adiabatically, radiation pressure transfers internal energy to the baryons, accelerating the flow and progressively converting radiation energy into bulk kinetic energy.

Photons decouple from the plasma at the photospheric radius, defined as the distance where the optical depth to Thomson scattering becomes unity ($\tau = 1$). The optical depth is approximated by:

$$\tau \sim \frac{n_p \sigma_T R}{2\Gamma}, \tag{2.5}$$

where n_p is the comoving proton number density, σ_T is the Thomson cross-section, R is the radius, and Γ is the Lorentz factor.

Assuming a mass outflow rate $\dot{M}$, the co-moving proton density is:

$$n'_p = \frac{\dot{M}}{4\pi R^2 m_p c \Gamma}. \tag{2.6}$$

Introducing the dimensionless baryon loading parameter:

$$\eta = \frac{L}{\dot{M} c^2}, \tag{2.7}$$

I derive the photospheric radius as:

$$R_{\rm ph} = \frac{L \sigma_T}{8\pi m_p c^3 \eta \Gamma^2}. \tag{2.8}$$

For typical GRB parameters ($L \sim 10^{52}$ erg/s, $\Gamma \sim 100$, $\eta \sim 100$), I obtain $R_{\rm ph} \sim 6 \times 10^{12}$ cm.

Detecting a thermal spectral component in the prompt emission would allow estimation of η, providing insights into the initial energy composition of the jet. However, clear observational evidence for such a component remains limited.

The resulting population of relativistic electrons is then expected to emit radiation primarily via synchrotron emission and inverse Compton scattering.

2.6.0 Radiative processes

The observed spectrum of the prompt emission in GRBs is significantly broader than a blackbody spectrum and typically features a power-law tail at high energies. This empirical evidence suggests that non-thermal radiative mechanisms are the dominant contributors to the emission process. Both fireball and magnetically dominated jet models anticipate the presence of highly energetic particles and strong magnetic fields. Among various radiation mechanisms, synchrotron radiation, the emission from relativistic electrons in magnetic fields, is widely considered a principal candidate for explaining GRB prompt emission. This section outlines the fundamentals of synchrotron and inverse Compton radiation, with a more detailed treatment available in Rybicki & Lightman (1986).

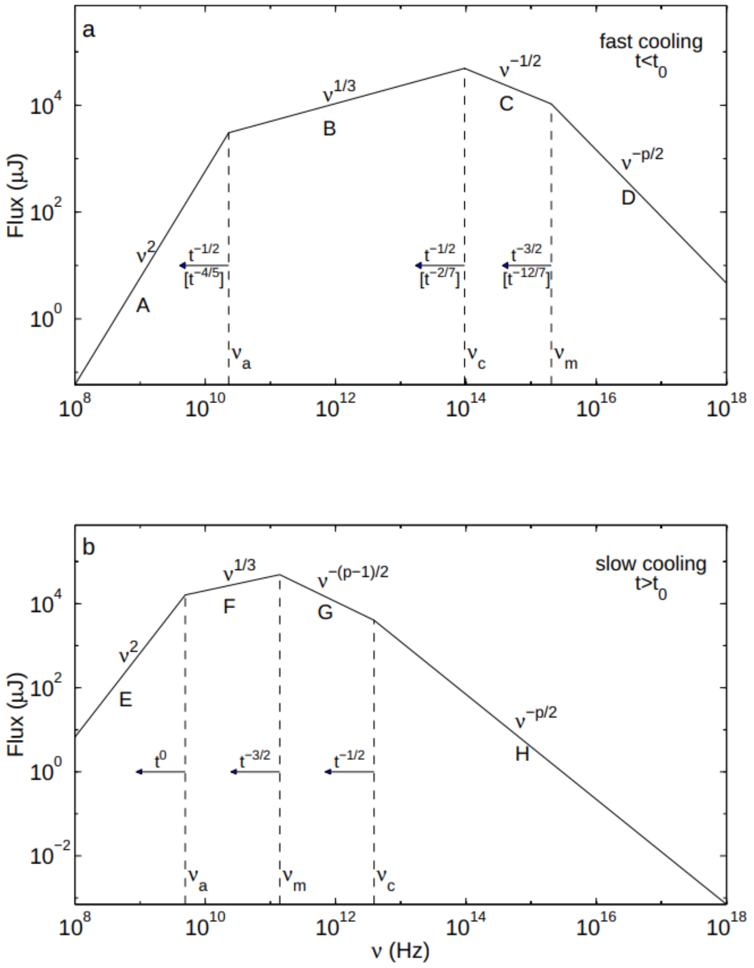

Figure 2.13: The synchrotron spectrum produced by a relativistic shock (Piran 1999), where the accelerated electrons follow a power-law energy distribution, exhibits characteristic differences depending on the cooling regime. During the fast-cooling phase, which typically occurs at early times ($t < t_0$), the spectrum can be divided into four regions (A, B, C, and D) (Sari et al. 1998), with synchrotron self-absorption becoming relevant below the characteristic frequency ν_a. As time progresses, the break frequencies, namely the injection frequency ν_m, the cooling frequency ν_c, and the self-absorption frequency ν_a, shift to lower values.

Following Sari et al. (1998), I define t_0 as the transition time from fast to slow cooling, i.e., the epoch when the cooling and injection break frequencies intersect, $\nu_c(t_0) = \nu_m(t_0)$.

2.6.1 Synchrotron Spectrum from Single Electrons and Populations

Synchrotron radiation, produced by ultra–relativistic electrons spiraling around magnetic field lines, represents one of the fundamental emission mechanisms in high–energy astrophysics, and is widely invoked to explain the observed spectra of sources such as GRBs. The motion of a charged particle in a magnetic field causes an acceleration perpendicular to its velocity due to the Lorentz force, which in turn results in the emission of radiation.

The total power radiated by a single relativistic electron is given by

$$P_{\rm syn} = \frac{4}{3}\sigma_T c \gamma^2 \beta^2 U_B, \tag{2.9}$$

where σ_T is the Thomson cross–section, $\beta = v/c$, and $U_B = B^2/(8\pi)$ is the magnetic energy density. The emitted power peaks around a characteristic frequency

$$\nu_s = \frac{\gamma^2 eB}{2\pi m_e c}, \tag{2.10}$$

which depends quadratically on the electron Lorentz factor. The single–electron spectrum rises as $F_\nu \propto \nu^{1/3}$ at low frequencies and declines exponentially beyond the critical value $\nu_c \approx 0.29\,\nu_s$.

In realistic astrophysical environments, however, the emission arises from a population of electrons with a broad range of energies. Considering the GRB scenario, the SED derives from several diffusive shock accelerations and is well described by a power law of the form

$$\frac{dN}{d\gamma} \propto \gamma^{-p}, \quad \gamma > \gamma_m, \tag{2.11}$$

where γ_m denotes the minimum Lorentz factor and p is the spectral index of the accelerated electrons. The resulting synchrotron spectrum depends critically on the cooling regime of the shocked electrons.

In the *slow–cooling regime*, where the radiative losses occur on timescales longer than the dynamical evolution of the system, the lowest–energy electrons retain most of their energy. The emergent flux density follows

$$F_\nu \propto \begin{cases} \nu^{1/3}, & \nu < \nu_m, \\ \nu^{-(p-1)/2}, & \nu > \nu_m, \end{cases}$$

with a spectral break near the frequency ν_m associated with γ_m.

Conversely, in the *fast–cooling regime*, electrons lose energy efficiently through synchrotron emission. The corresponding cooling frequency can be expressed as

$$\nu_c \sim \frac{36\pi e m_e c}{\sigma_T^2 B^3 t^2}, \tag{2.12}$$

and the resulting spectrum takes the form

$$F_\nu \propto \begin{cases} \nu^{1/3}, & \nu < \nu_c, \\ \nu^{-1/2}, & \nu_c < \nu < \nu_m, \\ \nu^{-p/2}, & \nu > \nu_m. \end{cases}$$

At the lowest frequencies, the emitted radiation becomes self–absorbed due to synchrotron self–absorption, which limits the escape of photons from the emitting region. Below the self–absorption frequency $\nu_{\rm sa}$, the spectrum follows a much steeper dependence, $F_\nu \propto \nu^2$ or $F_\nu \propto \nu^{5/2}$, typically occurring in the radio domain.

A schematic representation of the synchrotron spectrum in both cooling regimes, including the characteristic break frequencies $\nu_{\rm sa}$, ν_c, and ν_m, is shown in Figure 2.13.

It must be noted that the temporal scalings in Figure 2.13 are derived under the assumption that the blast wave propagates into a *homogeneous interstellar medium* (ISM), for which ν_m and ν_c both decrease with time in the adiabatic regime (Sari et al. 1998). In the later slow-cooling phase ($t > t_0$), the system is therefore expected to follow an adiabatic evolution while retaining the same segmented spectral structure. It is worth noting, however, that in a stratified circumburst environment, such as a stellar-wind density profile ($n \propto r^{-2}$), the temporal behavior of the break frequencies is modified: in particular, the cooling frequency ν_c is expected to *increase* with time rather than decrease (Chevalier & Li 2000).

2.6.2 Inverse Compton Radiation

Low-energy photons can gain energy via scattering off relativistic electrons in a process known as inverse Compton (IC) scattering. If the photon energy in the electron rest frame is small compared to $m_e c^2$, the scattering remains elastic and the average energy of the upscattered photon is $\nu_{\rm IC} \sim \gamma_e^2 \nu$. The total IC power for an electron in a photon field with energy density $u_{\rm ph}$ is:

$$P_{\rm IC} \sim \sigma_T u_{\rm ph} \gamma_e^2 c. \tag{2.13}$$

A particularly relevant case is Synchrotron Self-Compton (SSC), in which synchrotron photons serve as the seed photons for IC scattering. The corresponding photon energy density is:

$$u_{\rm ph} = \frac{\Delta R \sigma_T n_e B^2 \bar{\Gamma}^2}{6\pi}, \tag{2.14}$$

where ΔR is the emission region size, n_e is the electron density, and $\bar{\Gamma}$ is the average Lorentz factor.

The relative strength of the IC process compared to synchrotron emission is quantified by the Compton Y-parameter:

$$Y = \frac{P_{\rm IC}}{P_{\rm syn}} \sim \tau_e \bar{\gamma}^2, \tag{2.15}$$

where τ_e is the Thomson optical depth. In the Klein-Nishina regime ($h\nu\gamma_e \gtrsim m_e c^2$), scattering becomes inelastic and the cross-section decreases, leading to spectral distortions and limiting photon energy gain.

2.7.0 Afterglow phase

From an observational perspective, the GRB afterglow appears as a gradually fading emission that typically follows the prompt phase and can persist for hours to several months. Since its first

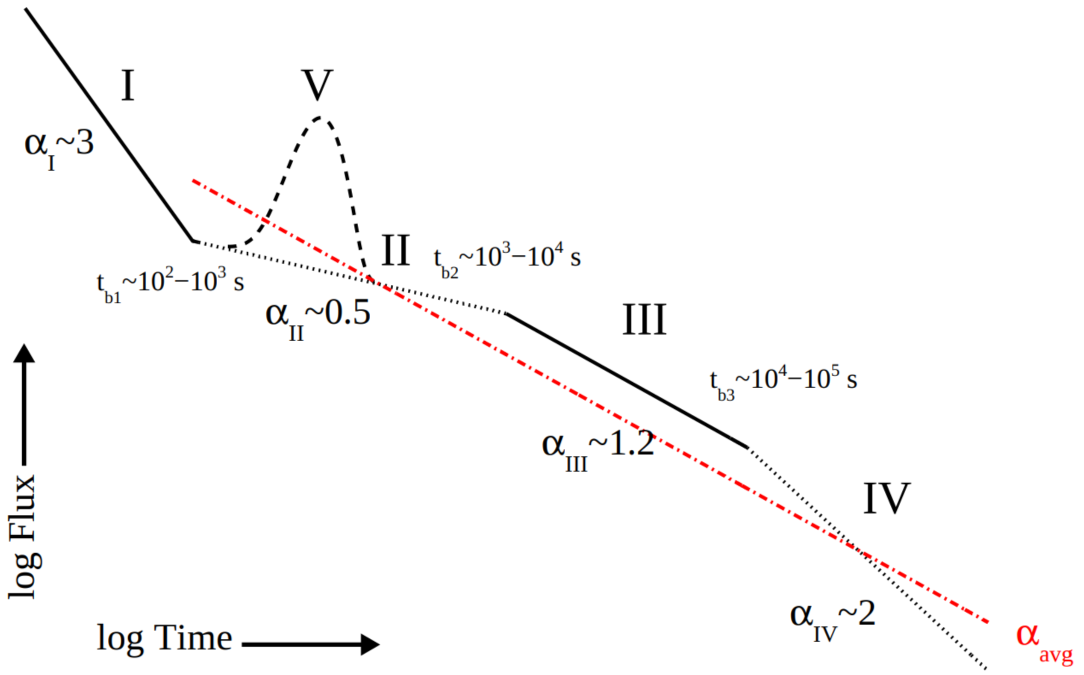

Figure 2.14: The canonical morphology of GRB X-Ray afterglows, as outlined by Zhang et al. (2006) and Nousek et al. (2006), is typically characterized by a sequence of temporal segments. One possible interpretation is the following: (I) a steep initial decline attributed to high-latitude (curvature) emission from the tail of the prompt phase; (II) a plateau phase likely driven by sustained energy injection into the forward shock; (III) a standard afterglow decay governed by the external forward shock; (IV) a post–jet break steepening indicating the collimated nature of the outflow; and (V) flares, which are observed in approximately 50% of GRBs and are thought to result from late-time internal activity. The average decay indices and break times exhibit significant variation across the GRB population (Racusin et al. 2009; Evans et al. 2009; Margutti et al. 2012).

detection by the *BeppoSAX* satellite in 1997 (Costa et al. 1997), afterglows have been extensively studied through multi-wavelength follow-up observations. This emission spans a broad range of the electromagnetic spectrum and shows an overall evolution toward lower characteristic frequencies with time.

2.7.1 Temporal Properties of GRB Afterglows

Prior to the launch of *Swift*, GRB afterglows were already routinely detected and systematically studied in X-Rays, most notably by *BeppoSAX* through its Narrow Field Instruments, which enabled the construction of a statistically significant catalog of X-Ray afterglows (De Pasquale et al. 2006).

With the advent of *Swift*, early and densely sampled multi-wavelength afterglow observations became possible, revealing a richer temporal behavior than previously accessible. Prior to *Swift*, optical light curves, typically obtained at relatively late times, often exhibited a simple power-law decay in flux, $F_\nu \propto t^{-\alpha}$, with a temporal decay index $\alpha \sim 1$. Compared to the highly variable prompt emission, the afterglow phase generally showed a smoother temporal evolution, as sketched in Figure 2.14.

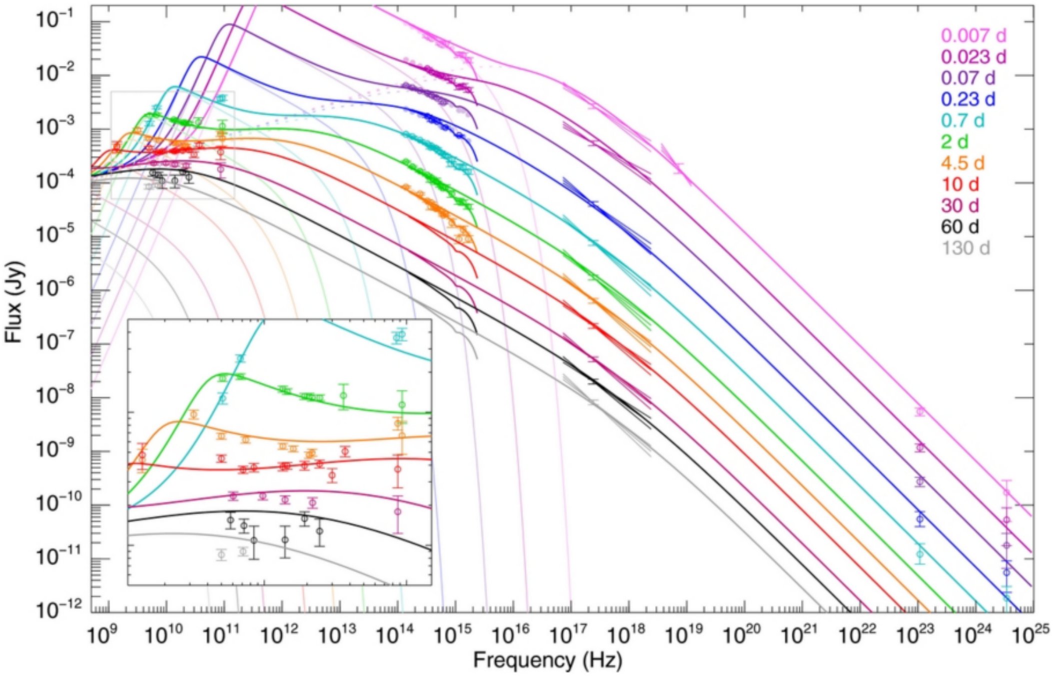

Figure 2.15: Broadband spectral energy distribution (SED) of the afterglow of GRB 130427A, compiled from observations spanning frequencies from the radio band up to 100 GeV. The various colors correspond to different epochs of observation. The solid lines represent the best-fit synchrotron emission models based on the standard afterglow theory (see Figure 2.13), which accurately reproduce the observed spectra across 18 orders of magnitude in frequency and four orders of magnitude in time. The inset in the lower left corner presents a zoomed view of the radio SED for times $t > 0.7$ days. Figure adapted from Perley et al. (2014).

The advent of the *Swift* mission revolutionized our understanding of afterglow emission, particularly through data collected by the X-Ray Telescope (XRT) onboard the satellite (Gehrels et al. 2009). Due to its rapid slewing capability, *Swift* can reorient toward a newly detected GRB within approximately 60-100 seconds from the trigger, allowing precise localization (within a few arcseconds) and extended monitoring of the X-Ray afterglow in the 0.3–10 keV energy range.

These early X-Ray observations revealed that the temporal behavior of GRB afterglows is far more complex than previously thought. Light curves often begin with a steep decay phase (decay index $\alpha < -2$), which is commonly attributed to high-latitude emission or the tail end of the prompt phase. This is frequently followed by a plateau phase with a much shallower decay ($\alpha \gtrsim -0.5$), sometimes superimposed with flaring activity whose physical origin remains under investigation (Zhang et al. 2006; Kann et al. 2010).

Subsequently, many GRBs transition into a normal decay segment with a slope around $\alpha \sim -1$, consistent with the optical afterglow behavior documented before the *Swift* era. For particularly long-lasting and luminous afterglows, a late-time steepening in the light curve is occasionally observed. This achromatic break, often occurring several days post-burst, is interpreted as a jet break, indicating the collimated nature of the outflow.

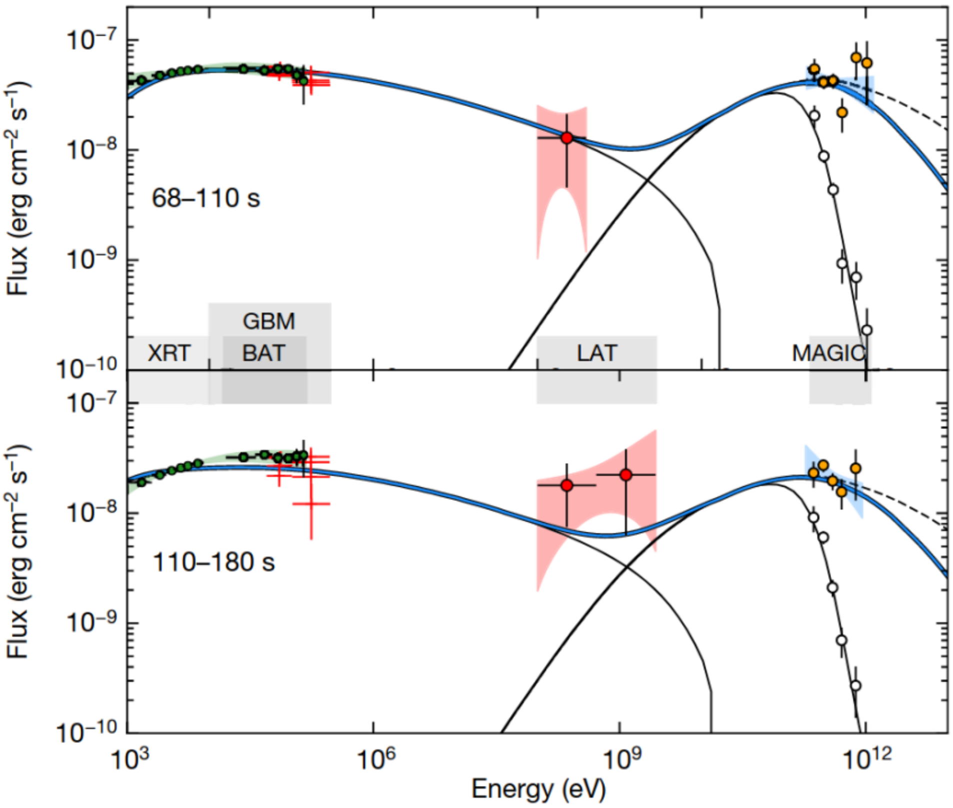

Figure 2.16: First high-significance detection of very high energy (VHE) photons from a GRB afterglow, observed by the MAGIC telescopes following the prompt phase of GRB 190114C. The data clearly reveal a second spectral peak, interpreted as the Synchrotron Self-Compton (SSC) component. The broadband spectral energy distribution is modeled over two time intervals (68–110 s and 110–180 s), with the synchrotron and SSC components indicated by thin solid curves. The thick blue line shows the combined model across more than nine decades in energy. Empty circles represent the observed MAGIC data, uncorrected for attenuation by the Extragalactic Background Light (EBL). Adapted from Acciari et al. (2019b).

2.7.2 Spectral Properties

GRB afterglows have been extensively observed across a broad range of frequencies, from GeV gamma-rays to the radio band. Thanks to simultaneous multi-wavelength campaigns, it has been possible to construct broadband spectral SED of several events. These SEDs are generally consistent with a single, non-thermal emission component, typically exhibiting multiple power-law segments connected by spectral breaks. As the afterglow evolves, these spectral breaks shift progressively toward lower frequencies.

The synchrotron emission mechanism is widely regarded as the dominant process responsible for the afterglow radiation. Indeed, synchrotron models are generally successful in reproducing the observed SEDs across different energy bands. A notable example is provided by GRB 130427A in Figure 2.15, whose afterglow spectrum has been modeled over 18 orders of magnitude in frequency and 4 orders of magnitude in time, offering a compelling match to synchrotron predictions (Perley et al. 2014).

2.7.3 Very High energy emission

Recent advancements in ground-based gamma-ray astronomy have allowed for the first robust detection of very high energy (VHE; $\gtrsim 100$ GeV) photons from a GRB afterglow. This breakthrough came with GRB 190114C, observed by the Major Atmospheric Gamma Imaging Cherenkov (MAGIC) telescopes (Acciari et al. 2019a). The detection marked a pivotal moment, confirming long-standing theoretical predictions that GRBs can emit in the TeV range.

The observation was particularly significant given the challenges in detecting such emissions.

This detection revealed a distinct high-energy component in the GRB afterglow spectrum (see Figure 2.16). Multi-wavelength modeling strongly supports the SSC mechanism as the dominant process responsible for the VHE photons (Acciari et al. 2019b). The energy released in the IC component was found to be comparable to that in the synchrotron component, providing key insights into particle acceleration and radiative efficiency.

Following this milestone, additional GRBs have been detected at VHE energies, including GRB 180720B, GRB 190829A, GRB 201015A, GRB 201216C (Abdalla et al. 2019; Blanch et al. 2020; Abe et al. 2023) and GRB 221009A, whose TeV afterglow rise and decay was clearly detected by LHAASO (Cao et al. 2023). Notably, all these detections pertain to the afterglow phase and originate from long-duration GRBs.

2.8.0 Multi-messenger astrophysics

The joint observation of GW170817 and GRB 170817A inaugurated the era of multi–messenger studies GRBs, demonstrating how GWs and electromagnetic signals together can decode the engines of the most extreme transients. GRB progenitors are natural factories of additional messengers, GWs, high–energy neutrinos, and ultra–high–energy cosmic rays, so coincident detections provide incisive diagnostics. In particular, the GW signal identifies the progenitor channel: compact–object mergers produce a characteristic chirp during inspiral, enabling a definitive association with short GRBs, whereas core–collapse events are expected to yield more complex and model–dependent GW emission (Flanagan & Hughes 1998; Bartos et al. 2013; Bethe 1990; Logue et al. 2012; Powell et al. 2016; Andersson et al. 2010). Post–merger GW features can, in principle, discriminate between prompt black–hole formation (ring-down) and a long–lived, supermassive neutron star (extended emission). For long GRBs, a GW counterpart would uniquely probe collapse dynamics hidden from EM view, although predicted strains remain highly uncertain because they sensitively depend on the supernova mechanism.

GW170817 occurred at a distance of ~ 40 Mpc and remains the only, benchmark joint detection. As the sensitivity and duty cycle of the LIGO–Virgo–KAGRA network improve, BNS mergers will be detectable to hundreds of megaparsecs, greatly expanding the pool of GRB–GW associations and enabling systematic searches for EM counterparts to BH–NS coalescences, systems that can launch a relativistic jet under favorable conditions (Abbott et al. 2020; Etienne et al. 2012; Ruiz et al. 2018). Such events will clarify how often mergers produce successful jets and delineate differences between NS–NS and BH–NS progenitors.

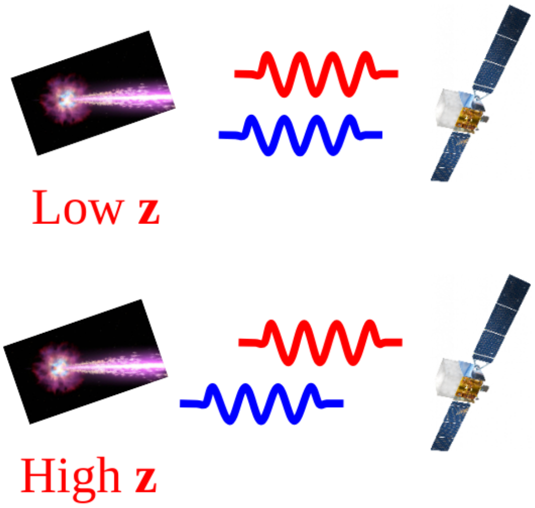

Figure 2.17: Cartoon of the time–of–flight test for Lorentz invariance with GRBs. Red (high–energy) and blue (low–energy) photons are emitted simultaneously at the source and detected by a space telescope (right). For very nearby burst (low z) any energy–dependent delay is negligible, whereas for a distant burst (high z) a putative LIV–induced dispersion can accumulate along the path, producing a measurable lag. Colours are symbolic; not to scale.

Beyond source physics, GW–GRB coincidences deliver cosmological leverage. Standard-siren distances from the GW signal, combined with a GRB–enabled redshift (via host identification), yield an independent measurement of the Hubble constant, as showcased by GW170817(Abbott et al. 2017; Guidorzi et al. 2017). With larger samples and next–generation facilities such as the Einstein Telescope, standard–siren cosmology can reach percent–level precision, helping arbitrate the present tension between early– and late–Universe determinations of H_0 (Punturo et al. 2010; Maggiore et al. 2020).

2.9.0 Quantum gravity with the GRB prompt emission

Lorentz invariance (LI) is a cornerstone of modern physics and underpins Einstein's theory of relativity, asserting that the speed of light in a vacuum is constant and independent of the inertial reference frame. Nevertheless, several theoretical approaches to Quantum Gravity (QG) predict a possible violation of this symmetry at energies approaching the Planck scale, $E_{\rm Pl} \sim 10^{19}$ GeV (e.g. Amelino-Camelia et al. 1998; Amelino-Camelia 2005; Kováčik et al. 2024; Pan et al. 2024). In particular, modified dispersion relations may introduce a weak energy dependence on photon propagation speed, an effect that can accumulate over cosmological distances.

In QG scenarios where Lorentz invariance may be violated at energies approaching $E_{\rm Pl}$ (Amelino-Camelia et al. 2011), the photon group velocity acquires an energy dependence, leading to minute but potentially observable deviations from the standard speed of light (Jacob & Piran 2008). This behavior can be described by the phenomenological dispersion relation (see e.g.

Tsvetkova et al. 2023, and references therein):

$$\frac{v_{\rm ph}}{c} - 1 = \xi \left(\frac{E_{\rm obs}}{E_{\rm QG}}\right)^n , \tag{2.16}$$

where ξ is ± 1 for the superluminal or subluminal scenario, respectively, $E_{\rm obs}$ is the observed photon energy, and $E_{\rm QG} = \zeta\, m_{\rm Pl} c^2$ is the effective QG scale, with $m_{\rm Pl} = 2.176 \times 10^{-5}$ g and ζ expresses the strength of the QG energy in units of the Planck energy ($E_{Pl} = m_{\rm Pl} c^2 = \sqrt{\hbar c^5/G} \simeq 1.22 \times 10^{19}$ GeV). The integer n denotes the leading-order term in the expansion, typically $n = 1$ or $n = 2$.

The resulting arrival-time delay $\Delta t_{\rm QG}$ between two photons with energy difference $\Delta {\rm E}_{\rm obs}$, emitted simultaneously from a source at comoving cosmological distance $D_{\rm trav}$, is (e.g. Burderi et al. 2021; Jacob & Piran 2008):

$$\Delta t_{\rm QG} = \xi \left(\frac{D_{\rm trav}}{c}\right) \left(\frac{\Delta E_{\rm obs}}{E_{\rm QG}}\right)^n . \tag{2.17}$$

Although small, a QG-induced dispersion becomes measurable when photon energies are high and path-lengths are cosmological, as for Gamma-ray bursts (GRBs).

To estimate the magnitude of LIV delays, let us adopt in Equation 2.17 the $\frac{D_{\rm trav}}{c} \approx 10$ billion years $\approx 3.5 \times 10^{17}$ s (corresponding to a redshift z=2 in a standard ΛCDM model), $\Delta E_{\rm obs} \sim 1$ GeV, $\zeta \sim 1$ that implies ${\rm E}_{\rm QG} \sim 10^{19}$ GeV. This means that $\Delta t_{QG} \sim 3 \times 10^{-2}$ s, for $n = 1$, and $\sim 3 \times 10^{-21}$ s, for $n = 2$. Therefore, reasonably, a direct measure of the Quantum Gravity induced delays can be attempted for n=1; I adopt n=1 hereafter in this thesis.

GRBs are among the most promising astrophysical sources for testing such violations. Their extreme energetics, large redshifts, and millisecond variability make them highly sensitive probes of *in-vacuo* dispersion effects (Amelino-Camelia et al. 1998; Vardanyan et al. 2023). In a Lorentz-invariance-violation (LIV) scenario, photons of different energies emitted simultaneously may arrive at Earth at slightly different times, producing measurable lags (see e.g. Burderi et al. 2021, and references therein).

Most limits on LIV have come from one-off analyses of a few extraordinarily bright GRBs, as in Vasileiou et al. (2013). Classic examples include GRB 221009A, which yields $E_{\rm QG} \gtrsim 5.9\, E_{\rm Pl}$ thanks to TeV-energy photons (Piran & Ofengeim 2024), or GRB 080916C, which pushes the limit to $E_{\rm QG} > 1.2\, E_{\rm Pl}$
(Abdo et al. 2009). In these studies, the time history of the keV prompt emission is cross-matched with the arrival times of individual GeV–TeV photons. Any residual delay that grows with energy is interpreted as a potential propagation effect and translated into a limit on LIV, with the strong hypothesis of negligible intrinsic delays between photons emitted at very different (by orders of magnitude) energies.

This "spectro-temporal" strategy is highly sensitive, but it relies on two fragile pillars: (i) it requires exceptionally luminous bursts that yield many high-energy photons, and (ii) it must assume a specific, and often uncertain, model for the intrinsic emission timeline of each GRB. Attempts to reduce such biases with a single burst include GRB 160625B, where cross-correlating forty energy bands (between 14 and 2000 keV) with a 10–12 keV reference energy band yields

$E_{\mathrm{QG}} > 5 \times 10^{15}$ GeV (Wei et al. 2017); yet the result remains sample-limited.

To sidestep those limitations, recent work has moved toward model-independent, population-based tests that fold intrinsic variability into the statistical treatment of the data. A model-independent analysis of 37 GRBs at redshifts $0.12 \leq z \leq 2.0$ gives $E_{\mathrm{QG}} > 2.6 \times 10^{15}$ GeV (Pan et al. 2024), while a statistical test based on the irregularity, skewness, and kurtosis of 8 *Fermi*/LAT GRBs sets $E_{\mathrm{QG}} \gtrsim 2.4 \times 10^{17}$ GeV (Ellis et al. 2019).

The central difficulty—whether for single-burst or ensemble analyses—remains the same: disentangling delays in the arrival time of photons of different energies due to intrinsic emission mechanisms from those induced by a dispersion law for photons of different energies expected in some QG theories. Authors in (Bernardini et al. 2017) restrict their analyses to *short* GRBs, arguing that the intrinsic spectral lags in these bursts are negligible compared with those observed in the long class and can be safely ignored. The article sets a 2σ lower limit of $E_{QG} > 1.48 \times 10^{16}$ GeV.

3

Compact objects

Summary

This chapter examines compact remnants above the Chandrasekhar threshold, emphasizing *neutron stars*, their pulsar manifestations, and *ULXs*. It outlines the key microphysics (degeneracy, magnetism, rotation), connects spin-powered emission to pulsar behavior, and explains how super-Eddington accretion and winds energize ULXs.

3.1.0 Neutron Stars

Soon after the neutron's discovery in 1932 (Chadwick 1932), theorists began to consider whether a self–gravitating object made predominantly of neutrons could exist in stable equilibrium. In 1934, Baade & Zwicky (1934) introduced the neutron–star concept, emphasizing their extraordinary density and proposing a supernova origin. A few years later, Oppenheimer & Volkoff (1939) constructed the first relativistic model, treating the star as cold neutron matter at extreme densities ($\rho \simeq 5 \times 10^{15}\,\mathrm{g\,cm^{-3}}$) and inferring a limiting mass of $\sim 0.7\,M_\odot$ with a radius $R \simeq 9.6\,\mathrm{km}$. For more than three decades, however, NSs were largely set aside observationally, owing to the belief that their small size would render them unobservable with optical telescopes at astrophysical distances.

Interest revived in the 1960s with the dawn of high–energy astronomy: the first extra–solar X–ray source was identified in 1962 (Giacconi et al. 1962b), prompting a surge of theoretical work on compact objects. Soon after, on 28 November 1967, Jocelyn Bell and Anthony Hewish detected a periodically pulsing radio source with $P = 1.3373$ s (Hewish et al. 1968). Initially nicknamed "LGM–1" (for "Little Green Men"), it was quickly recognized as the first radio pulsar (PSR B1919+21). The discovery led to the Nobel Prize for Hewish; Bell's crucial role, though acknowledged, was not similarly recognized.

As outlined in the Equation 2.1, increasing density drives the electron Fermi energy upward. Once the total electron energy exceeds the neutron–proton mass–energy difference ($\Delta E \simeq 1.29\,\mathrm{MeV}$), inverse β decay ($p + e^- \rightarrow n + \nu_e$) becomes energetically allowed, reducing the electron fraction

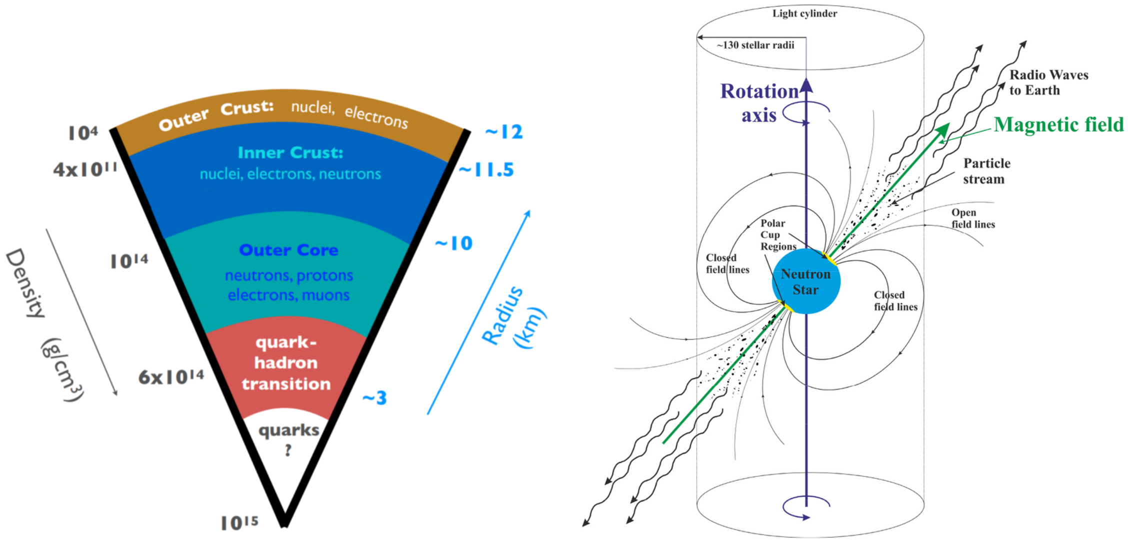

Figure 3.1: **Left panel**: Representative stratification of a neutron star's interior (Longair 2011). **Right panel**: Schematic magnetic–dipole model of a pulsar (Błaszkiewicz et al. 2016).

and initiating *neutronisation*.

In white–dwarf–like matter dominated by C/O/Fe nuclei, this process loads nuclei with excess neutrons; beyond a threshold, they fragment and release free neutrons into the relativistic, degenerate electron gas.

The onset of free–neutron leakage, *neutron drip*, occurs at $\rho \sim 4 \times 10^{14}\,\mathrm{kg\,m^{-3}}$ (Longair 2011). Progressively, the innermost regions convert almost entirely to neutrons. Because neutrons are ~ 2000 times more massive than electrons, their typical velocities are nonrelativistic at comparable number densities, yielding a degeneracy pressure with the steeper scaling $P_D \propto n^{5/3}$; this enhanced stiffness balances gravity and produces a new equilibrium configuration: the **neutron star**.

3.1.1 How Neutron stars exist

Free neutrons decay with a mean lifetime of ≈ 14.8 minutes. In an NS, however, β decay is Pauli–blocked: the decay electron must occupy a state above the electron Fermi surface. If the electron Fermi energy already exceeds the kinetic energy available to the emitted electron (from the Q-value of $n \to p + e^- + \bar{\nu}_e$), the process is energetically disfavored. This blocking becomes effective above a critical density $\rho \simeq 1.2 \times 10^{10}\,\mathrm{kg\,m^{-3}}$ (Longair 2011), stabilizing bulk neutron matter in the star.

Internal structure and equation of state. NSs are among the densest known objects, with core densities exceeding those in atomic nuclei, inaccessible to terrestrial experiments. As a result, a variety of theoretical equations of state (EoS) have been developed to describe ultra–dense matter. The left panel of Figure 3.1 sketches a standard stratification:

- *Outer layer:* a thin surface/atmospheric region above $\rho \lesssim 10^9\,\mathrm{kg\,m^{-3}}$, often idealized as highly conducting iron–group matter arranged in ordered structures with effectively negligible resistivity.

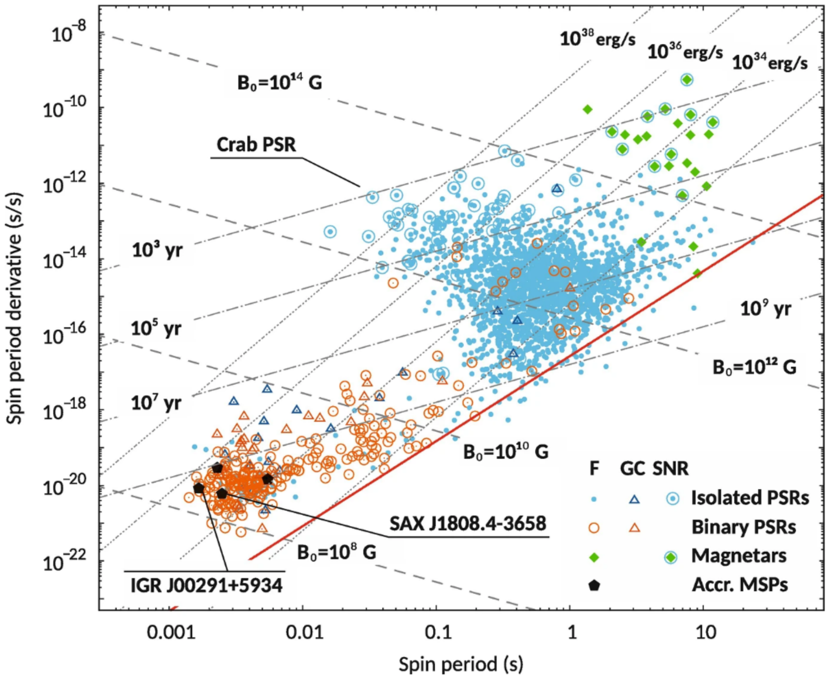

Figure 3.2: $P - \dot{P}$ diagram of rotation–powered pulsars showing distinct sub–classes. Dashed, dotted, and dot–dashed curves mark constant surface dipole field B, characteristic age τ_c, and spin–down power $\dot{E}$, respectively; the red solid curve is the putative death line (Di Salvo & Sanna 2022).

- *Outer crust:* from $\rho \sim 10^9$ to $4.3 \times 10^{14}\,\mathrm{kg\,m^{-3}}$, a Coulomb lattice of heavy nuclei embedded in a relativistic, degenerate electron gas.

- *Inner crust:* from $\rho \sim 4.3 \times 10^{14}$ to $2 \times 10^{17}\,\mathrm{kg\,m^{-3}}$, neutron–rich nuclei coexist with a gas of free degenerate neutrons and relativistic electrons; with depth, nuclei dissolve into a neutron liquid.

- *Central core:* beyond $\rho \gtrsim 3{\times}10^{18}\,\mathrm{kg\,m^{-3}}$, neutrons overlap geometrically and non–nucleonic degrees of freedom may appear; possibilities include solid or superfluid neutron phases, hyperons, meson condensates, or deconfined quark matter.

While the crustal layers are comparatively well constrained (Baym et al. 1971), the innermost composition remains a matter of active debate and is encoded in the dense–matter EoS. Broadly, *stiff* EoS (pressure rising rapidly with density) produce stars that are harder to compress, hence larger radii at a given mass and potentially higher maximum masses; *soft* EoS (more compressible matter) yield more compact stars with lower $M_{\max}$ (Lattimer & Prakash 2016). This dichotomy imprints directly on the mass–radius relation, making precise astrophysical constraints on M and R, together with complementary probes, central to discriminating among competing EoS models (Degenaar & Suleimanov 2018).

3.2.0 Pulsars

The first neutron star identified by Jocelyn Bell (Hewish et al. 1968) turned out to be a *radio pulsar*: a rapidly rotating, highly magnetized NS whose misaligned magnetic and spin axes produce a beamed lighthouse–like signal. Magnetic flux freezing during core collapse implies $B \propto R^{-2}$, so compressing a $\sim 10^6$ km progenitor to $R \sim 10$ km can amplify modest stellar fields to $B \sim 10^{12}$ G; in some objects, *magnetars*, even stronger fields $B \sim 10^{14}$–10^{15} G are inferred (Kaspi & Beloborodov 2017; Turolla et al. 2015). Angular–momentum conservation ($J = I\Omega$) simultaneously drives the spin to short periods.

In vacuum dipole braking, a star with polar field B_p, radius R, and obliquity α loses rotational energy at a rate

$$\dot{E}_{\rm dip} = -\frac{B_p^2 R^6 \Omega^4}{6c^3} \sin^2\alpha, \tag{3.1}$$

implying $\dot{\Omega} \propto -\Omega^n$ with $n \simeq 3$ for a pure dipole (Ostriker & Gunn 1969; Jackson 1998). The associated *magneto–dipole* wave has frequency $\sim \Omega$ (Hz), far below radio bands; the *observed* radio through γ–ray emission instead arises in the magnetosphere, where unipolar induction generates a charge density $\rho_{\rm GJ} \simeq -\mathbf{\Omega}\cdot\mathbf{B}/(2\pi c)$ (Goldreich & Julian 1969). Particle acceleration in gaps near the polar caps/slot gap/outer gap triggers pair cascades; curvature, synchrotron, and inverse–Compton radiation from these pairs produce the broadband beamed emission (Ruderman & Sutherland 1975; Daugherty & Harding 1996a; Cheng et al. 1986; Lyutikov 2013; Daugherty & Harding 1996b). Pulses are seen whenever the rotating beam intersects our line of sight.

The diversity of pulsars is conveniently organized in the P–$\dot{P}$ plane (Figure 3.2). Young, rotation–powered pulsars (often still embedded in supernova remnants) cluster at $P \sim 0.03$–1 s with large $\dot{P}$, strong fields $B \simeq 3.2\times10^{19}\sqrt{P\dot{P}}$ G, and spin–down ages $\tau_c = P/(2\dot{P}) \sim 10^3$–$10^5$ yr. *Magnetars* occupy the upper–left region with $B \gtrsim 10^{14}$ G and comparatively large $\dot{P}$ (Kaspi & Beloborodov 2017). *Rotating Radio Transients* (RRATs) show sporadic radio bursts and sit among ordinary pulsars in P–$\dot{P}$ but with extreme intermittency (McLaughlin et al. 2006). The lower–left corner hosts *millisecond pulsars* (MSPs), with $P \lesssim 10$ ms and $B \sim 10^8$–10^9 G.

Their demographics, low orbital eccentricities and a high binary fraction, support the *recycling* scenario: prolonged Roche–lobe accretion in low–mass X–ray binaries spins up the neutron star to millisecond periods; when the accretion rate drops, the magnetosphere clears, coherent radio emission turns on, and the already fast rotator is observed as a radio MSP (Alpar et al. 1982; Backer et al. 1982; Bhattacharya & van den Heuvel 1991).

As pulsars age, their spin slows and pair production becomes less efficient, driving them toward the *death line* where coherent radio emission can no longer be sustained; beyond this boundary they join the so–called pulsar graveyard (Chen & Ruderman 1993).

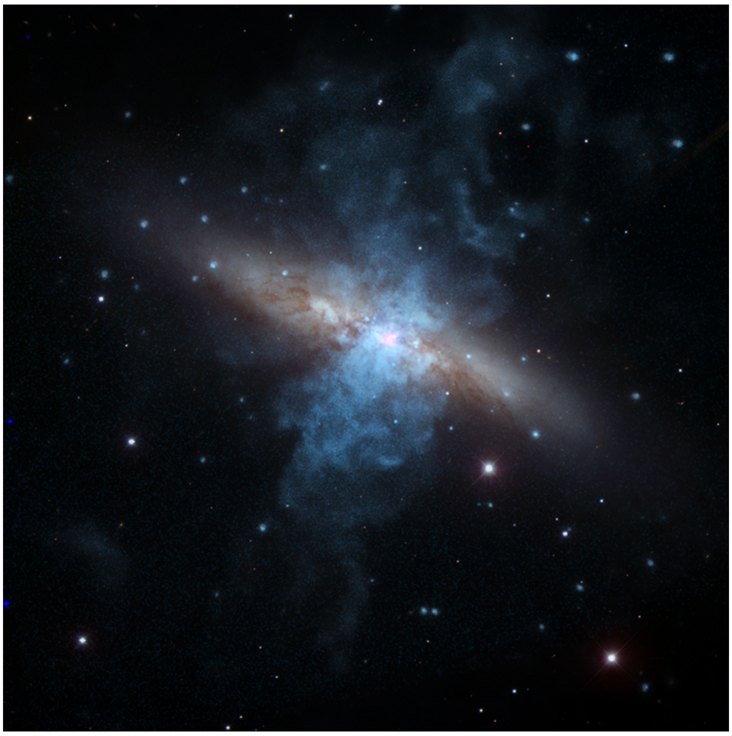

Figure 3.3: Ultraluminous X-Ray *pulsar* M82 X–2 in the starburst galaxy M82. The composite shows X-Rays from *NuSTAR* (purple) and *Chandra* (blue) overlaid on optical data from the NOAO 2.1 m telescope (gold). Thanks to its angular resolution, *Chandra* isolates M82 X–2 from the neighboring ULX M82 X–1, while *NuSTAR* discovered coherent 1.37 s pulsations establishing a neutron-star accretor (PULX) (Bachetti & et al. 2014). **Credit:** X-Ray: NASA/CXC/Univ. of Toulouse/M. Bachetti et al.; optical: NOAO/AURA/NSF.

3.3.0 Ultraluminous X-ray Sources

ULXs are off nuclear X–ray emitters in external galaxies whose *isotropic equivalent* luminosities in the 0.3–10 keV band exceed $L_X \gtrsim 10^{39}$ erg s^{-1}, apparently above the Eddington output of a standard stellar mass accretor:

$$\mathrm{L_{edd}} = 1.3 \times 10^{38} \left(\frac{\mathrm{M}}{\mathrm{M_{\odot}}} \right) \mathrm{erg\ s^{-1}}. \tag{3.2}$$

Early surveys already revealed a non negligible population (Fabbiano 1989), and modern compilations list ~ few × 10^3 candidates, with a bright tail up to $L_X \sim 10^{41}$–10^{42} erg s^{-1} (Walton & et al. 2022). Their displacement from galactic nuclei distinguishes them from accreting supermassive black holes.

From the outset, two broad explanatory tracks emerged. One invoked unusually massive accretors: Colbert & Mushotzky (1999) argued that the most luminous systems could host *intermediate mass black holes* (IMBHs; $M \sim 10^2$–10^4 $M_{\odot}$), a term soon adopted in the literature (Taniguchi et al. 2000), with proposed formation channels via runaway mergers of compact objects or stars in dense clusters (Miller & Hamilton 2002; Portegies Zwart & McMillan 2002). The alternative kept *stellar mass* accretors but in atypical radiative states: ideas included intrinsically super Eddington emission from inhomogeneous, radiation pressure dominated discs (Begelman

2002), relativistic beaming from microquasar jets (Körding et al. 2002), or reduced electron scattering opacity in strongly magnetized neutron star columns (Canuto et al. 1971; Mushtukov et al. 2015). However, these mechanisms generally struggle beyond $L_{\mathrm{sph}} \sim 10^{40}$–$3 \times 10^{40}$ erg s^{-1} and cannot alone account for the most extreme ULXs (Narayan et al. 2017).

A unifying picture arose from binary evolution arguments: episodes of supercritical mass transfer drive powerful, radiation-pressure winds that collimate the escaping X-Rays into a funnel, boosting the apparent luminosity for sightlines near the disc axis. This framework, rooted in super Eddington disc physics (Shakura & Sunyaev 1973) and high mass transfer phases, identified in X–ray binary evolution (King & Ritter 1999; Podsiadlowski et al. 2002), naturally predicts outflows with characteristic speeds $\sim 0.1c$ (King & Pounds 2003, as a wind diagnostic) and an association with vigorously star forming environments, as observed in systems like the Antennae and the Cartwheel (Fabbiano et al. 2001; Gao et al. 2003).

A decisive development was the discovery of coherent pulsations from M82 X–2 (Bachetti & et al. 2014), proving that *neutron stars* can power ULX–level luminosities. This result removed the need to invoke IMBHs as a *generic* explanation and sharpened tests of disc wind collimation versus strong field NS models; for complementary viewpoints see Fabrika et al. (2021) and citetMushtukovTsygankov2022. Looking ahead, establishing whether ULXs form a single supercritical accretion class with geometry-dependent observables, or comprise multiple subpopulations (including a minority of genuine IMBHs), requires uniform, multi-epoch spectroscopy and timing across large, well-defined samples (Walton & et al. 2022). This endeavour has been enabled by *Chandra* and *XMM–Newton*, whose imaging and spectroscopic capabilities in the 0.3–10 keV band revolutionized the field (Weisskopf & et al. 2000; Jansen & et al. 2001).

4

Data and instrumentation

Summary

This chapter briefly introduces the facilities used in this thesis for timing analysis: the HERMES nano-satellite constellation, the *Fermi* instruments (GBM and LAT), *Insight–HXMT*, and *XMM–Newton*. For each, I summarize the energy coverage, timing resolution, and data products relevant to burst detection, time–tagging, and light-curve analysis.

4.1.0 The High Energy Rapid Modular Ensemble of Satellites

HERMES–TP/SP (High Energy Rapid Modular Ensemble of Satellites—Technological and Scientific Pathfinder) is a constellation of six 3U nano-satellites equipped with compact yet innovative X–ray/soft γ–ray detectors that combine a broad energy passband with excellent timing accuracy. The system is optimized for the continuous monitoring of high–energy cosmic transients, most notably gamma–ray bursts and the electromagnetic counterparts of gravitational–wave events, and for the prompt determination of their sky positions. The projects are funded by the Italian Ministry of University and Research (MUR) and the Italian Space Agency (ASI), with additional support from the European Union's Horizon 2020 Research and Innovation Programme (Grant Agreement No. 821896). Conceived as an in–orbit demonstration with initial activities planned from 2022, *HERMES–TP/SP* adopts an intrinsically modular architecture that can be naturally scaled into a distributed, sensitive all–sky monitor for high–energy transients.

An HERMES detector is on board the Australian nanosatellite SpIRIT, launched on December 2023, and it is currently active and operational (Trenti et al. 2024). The full HERMES-TP/SP constellation was launched on 15 March 2025, and the nanosatellites are presently in the commissioning phase.

Since the primary goal of HERMES is high-energy sky monitoring, the detector has no imaging capability. The pointing is typically kept radial, i.e., the satellite zenith points toward the Earth's center at all times.

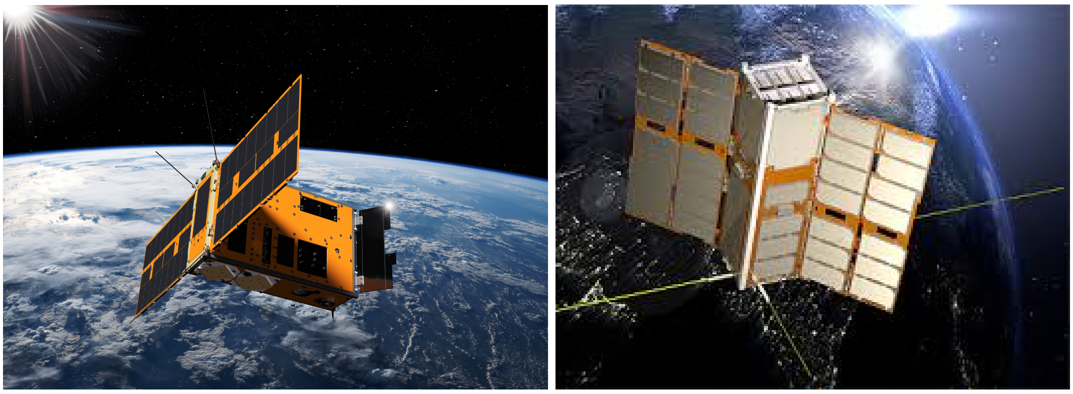

Figure 4.1: Right: The *HERMES–TP/SP* 3U nanosatellites hosting the siswich X–ray/soft γ–ray detector (launched in March 2025). Left: The *SpIRIT* 6U CubeSat carrying a HERMES detector (launched in December 2023). Together they provide complementary baselines and sub–μs timing for inter–satellite time–delay localization of high–energy transients.

4.1.1 HERMES scientific goals

The *HERMES–TP/SP* pathfinder pursues five tightly coupled objectives Fiore et al. (2020):

1. *Miniaturized instrumentation for breakthrough science.* Design and qualify compact X–ray/soft γ–ray detectors capable of capturing GRB and X–ray transient signals within the strict mass, volume and power envelopes of 3U nano-satellites, while preserving the sensitivity and performance required for high–impact science.

2. *Preparation for high–value data production.* Demonstrate that accurate sky positions can be recovered from time–difference–of–arrival (TDOA) measurements across multiple spacecraft. The in–orbit demonstration targets localization uncertainties of a few degrees, comparable to or better than *Fermi*/GBM, with arcminute–class localizations achievable in a future, larger *HERMES Full Constellation* (HFC) observing the same transient with tens to hundreds of units. The sub–μs timing intrinsic to the payload also enables temporal studies of prompt emission down to $O(10^{-7})$ s.

3. *COTS in challenging space missions.* Reduce cost and time–to–orbit by leveraging miniaturized platforms and commercial off–the–shelf (COTS) components, while establishing a production life–cycle (screening, parts derating, radiation/thermal testing) that raises COTS reliability without eroding the schedule and budget advantages typical of CubeSats.

4. *Space 4.0 contributions.* Identify and standardize novel approaches to the manufacture, assembly, and verification of miniaturized subsystems, fostering supply–chain evolution and lowering barriers for agile small and medium enterprises to participate in science–grade space hardware.

5. *Scaling distributed architectures.* Quantify how constellation design choices impact timeliness and localization precision for unpredictable transients: number of satellites versus ground stations, sparing policies, heterogeneous versus homogeneous orbital

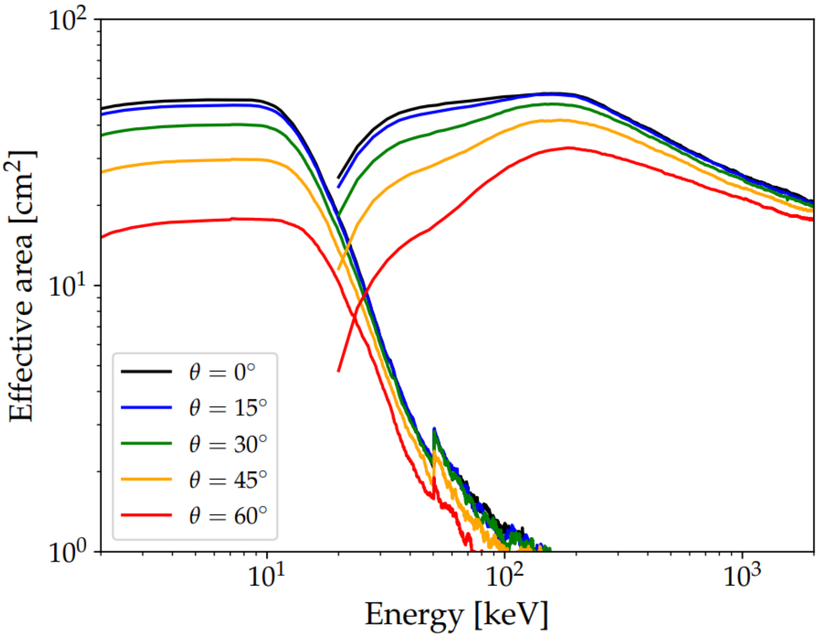

Figure 4.2: Effective area of the HERMES detector as a function of photon energy for several off–axis angles (see legend). Curves clustered on the left correspond to the X–mode response (direct SDD detection), whereas those on the right trace the S–mode response (Campana et al. 2020).

geometries, constellation coordination (science/relay roles), subsystem performance and lifetime versus replacement, and overall reliability, cost, scalability, and operational readiness for large–scale distributed science missions.

High–level scientific requirements (Pathfinder). The pathfinder shall (i) detect GRBs with peak flux $\gtrsim$ 0.5–1 ph s^{-1} cm^{-2} in the 50–300 keV band; and (ii) during a nominal two–year mission, detect $\gtrsim$ 40 long and $\gtrsim$ 8 short GRBs *simultaneously in at least three units*, with per–unit efficiency $\gtrsim$ 40–50%, so that TDOA analysis can deliver a sky position. Meeting (i) implies, per detector, an effective collecting area in 50–300 keV of $\gtrsim$ 50 cm^2 and a background $\lesssim$ 1.5 counts cm^{-2} s^{-1} (Campana et al. 2020; Fuschino et al. 2019; Evangelista et al. 2020a). Meeting (ii) requires a wide FoV (order few sr) and, operationally, the capability to co–align the instantaneous FoVs of at least three detectors; attitude knowledge should be better than a small fraction of a single–detector FoV (several degrees). A dedicated mission analysis confirms the feasibility of these conditions Colagrossi et al. (2020).

Localization performance and scaling. The key science goal is to prove that precise localizations are achievable with miniaturized instruments on nano-satellites by exploiting inter–satellite TDOAs over baselines of $\sim 10^3$ km. In the timing–dominated limit, the 1σ angular uncertainty scales approximately as

$$\sigma_\theta \;\propto\; \frac{c\,\sigma_t}{B}$$

where σ_t is the uncertainty on the inter–spacecraft delay estimate and B the projected baseline between the spacecraft. For the pathfinder mini–constellation (three to four units rather than the $\sim$ 25–30 envisaged for HFC, and smaller per–unit area), achieving the HFC target of $\sim 15'$

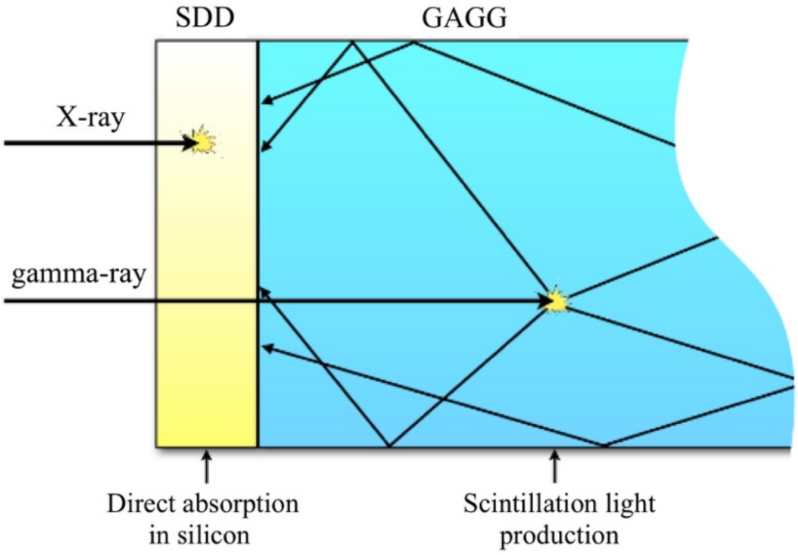

Figure 4.3: In the “siswich” architecture, a single Silicon Drift Detector (SDD) provides dual readout: it operates as a photodiode to measure the scintillation light generated in the crystal by hard X/γ–ray interactions, and it simultaneously serves as a direct sensor for soft X–rays absorbed in the silicon.

requires demonstrating $\sim 2.5^\circ$ localizations. This, in turn, demands delay constraints at the level of $\sigma_t \sim 1$ ms or better between two GRB light curves, which are attainable with $\gtrsim 50$ cm^2 per detector for bright bursts ($\gtrsim$ 10 ph s^{-1} cm^{-2}) that yield $\gtrsim$ 1 count every few milliseconds Sanna et al. (2020); Fuschino et al. (2019).

4.1.2 HERMES Payload

To meet the required broad bandpass with a compact payload, *HERMES* adopts an integrated, “double detection” (siswitch–like) architecture in which soft X–rays are recorded directly by Silicon Drift Detectors (SDDs) in the X-Mode, while hard X–/γ–rays are measured via scintillation in GAGG:Ce crystals and subsequent optical readout with the same SDDs as shown in Figure 4.3 in the S-Mode. The two channels partially overlap near $\sim$ 20 keV, enabling cross–calibration and robust event typing (see Figure 4.2).

The SDDs, implemented as $\lesssim$ 1 cm^2 tiles with very low anode capacitance, deliver room–temperature noise of only a few tens of electrons rms and thus afford low thresholds down to a few keV; with a 450 μm silicon thickness they retain useful efficiency up to $\sim$ 20–30 keV. Although GAGG has seen limited flight heritage so far, radiation–environment studies indicate its suitability for space use (Yoneyama et al. 2018; Gatti & Rehak 1984), a conclusion supported by dedicated irradiation tests (e.g. Trento Protontherapy Centre, January 2019)(Dilillo et al. 2020). Relative to photomultiplier tubes, silicon readout is favored for its much higher photon–detection efficiency ($\sim$ 90% versus $\sim$ 20%), smaller volume, and lower power; compared with SiPMs, SDDs additionally provide direct sensitivity to low–energy X–rays, enabling a truly broadband instrument, and they have demonstrated superior tolerance to radiation–induced degradation. In operation, incident hard photons are converted to visible light in GAGG and the SDDs act as

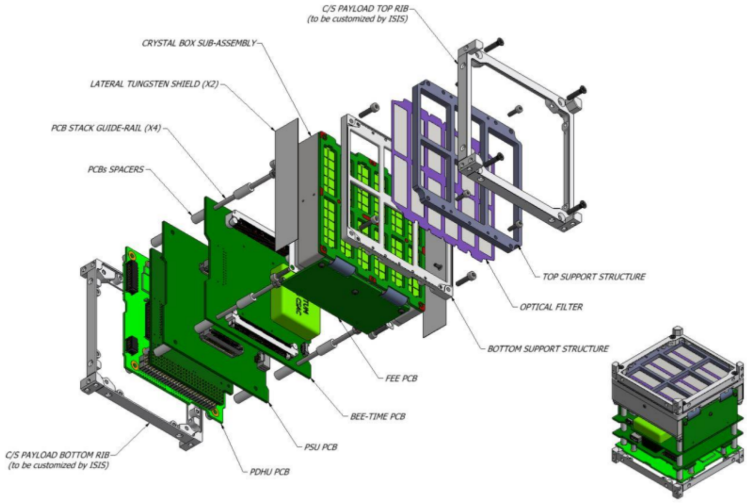

Figure 4.4: Extensive view of the HERMES payload (Evangelista et al. 2020a).

photodiodes that integrate the scintillation signal, whereas soft X–rays are absorbed directly in silicon. Event discrimination exploits a segmented geometry in which each crystal is optically coupled to two SDDs: single–anode hits are associated with direct X–ray absorption, while near–simultaneous signals on both SDDs tag scintillation events. GAGG:Ce is selected for its sub–μs decay time and high light yield (~ 56 ph/keV), which lowers the energy threshold compared to denser but dimmer scintillators such as BGO (~ 8 ph/keV).

To exploit the SDDs' low noise, in the front–end is implemented an application–specific integrated circuit. The baseline comprises two ASICs: *LYRA–FE*, a charge–sensitive preamplifier co–located with the SDD to minimize input capacitance and series noise; and *LYRA–BE*, which implements the full spectroscopic chain (including ADCs) and a fast discriminator for precise event timing. The analog front–end draws heritage from the VEGA mixed–signal, low–noise multichannel ASIC developed by Politecnico di Milano and Università di Pavia for SDD readout. A back–end electronics (BEE) board hosts the secondary analog stage, digitization, and control logic, and, critically, the on–board time (OBT) management. The BEE FPGA maintains two counters: one disciplined by the GPS pulse–per–second (PPS) and the other driven by the 10 MHz output of a miniaturized atomic clock. With GPS lock, the payload PPS is phase–aligned each second to the spacecraft bus PPS to continuously correct the OBT; during GPS outages (e.g. Earth occultation, fewer than four tracked satellites) the atomic clock supplies both 1 PPS and 10 MHz in free–running mode. This scheme provides stable, sub–μs, event–by–event time stamping, which underpins high–fidelity temporal analyses and inter–satellite time–difference–of–arrival localization.

The *HERMES* detector is segmented into four identical, independent quadrants (Q0–Q3). Each quadrant has its own front–end chain and housekeeping/rate counters, so data products provide per–quadrant count rates in the three energy bands. This layout adds redundancy and enables cross–checks for systematics or partial occultations, as well as basic localization/diagnostics from inter–quadrant count–rate differences (see Figure 4.4).

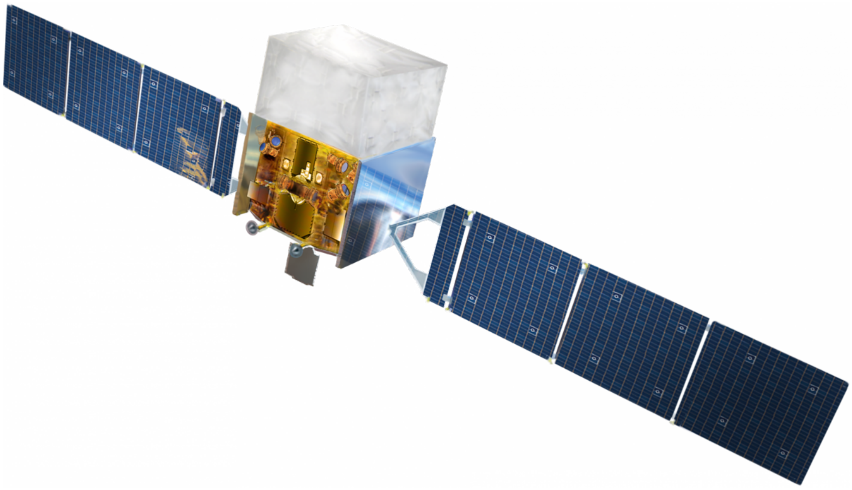

Figure 4.5: The *Fermi* Gamma-ray Space Telescope: the LAT in the upper tower and the GBM mounted on the spacecraft bus below, together surveying the high–energy sky from ~8 keV to > 300 GeV.

4.1.3 HERMESdas data reduction

The HERMES data reduction pipeline comprises two stages (Puccetti et al. 2024). The Raw→Level 0 segment (developed at INAF) converts spacecraft telemetry into OGIP-compliant FITS files, performing packet decommutation and translating electronic housekeeping into physical SI units. The Level 0→Level 2 processing (developed at SSDC) is carried out by HERMESDAS within the NASA HEASoft environment, reusing multi-mission infrastructure (makefiles and core libraries such as FITSIO and PIL/FTOOLS) and delivering calibrated, cleaned science products. Inputs and outputs use FITS following NASA-OGIP standards, ensuring immediate interoperability with standard High-Energy analysis tools, XSELECT for event filtering and data-product extraction, XSPEC for spectral modeling, and XRONOS for timing analysis. Instrument calibration is managed via the HEASARC CALDB system, which provides the appropriate response and ancillary files.

4.2.0 Fermi observatory

The *Fermi Gamma-ray Space Telescope* in Figure 4.5 (formerly *GLAST*), launched on 11 June 2008 into low–Earth orbit, is a flagship observatory for high–energy astrophysics that surveys the non–thermal Universe from ~ 8 keV to > 300 GeV through two complementary instruments. The *Large Area Telescope* (LAT) is a wide–field, pair–conversion γ–ray telescope operating from ~ 20 MeV to beyond 300 GeV, optimized for continuous all–sky monitoring with uniform exposure, source detection, spectroscopy, and variability studies on time scales from minutes to years (Atwood et al. 2009). The *Gamma-ray Burst Monitor* (GBM) monitors the hard X-Ray and soft γ-ray sky. Operating predominantly in survey mode, *Fermi* has delivered transformative results on γ–ray pulsars, active galactic nuclei, diffuse emission, and, critically for multi–messenger astrophysics, the prompt and afterglow emission of GRBs, including the landmark short burst GRB 170817A temporally associated with the binary neutron–star merger

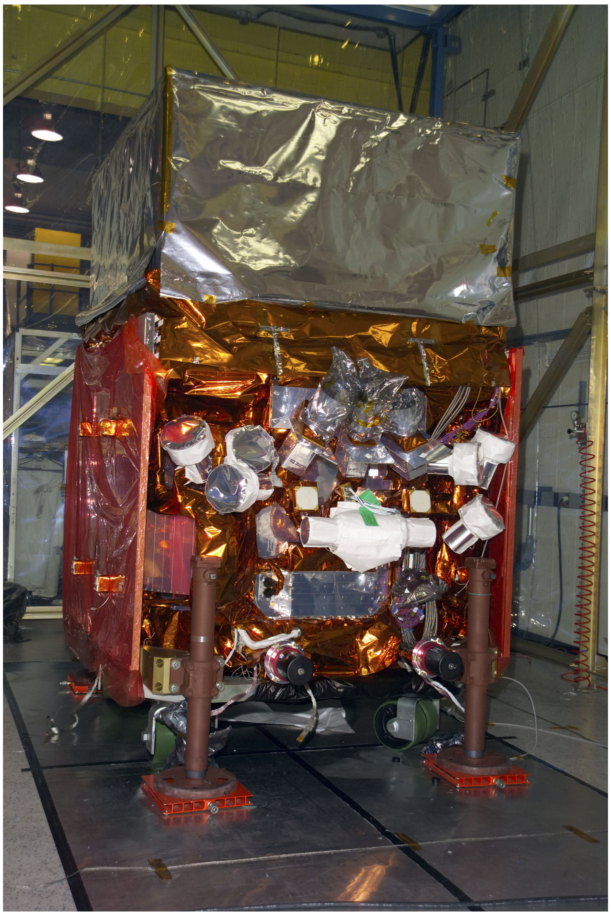

Figure 4.6: *Fermi*/GBM flight hardware during integration. The twelve NaI(Tl) scintillation detectors (white conical housings) are mounted around the spacecraft at different viewing angles to provide nearly all–sky coverage and enable relative–count–rate localization of transients.

GW170817 (Goldstein et al. 2017; Abbott et al. 2017), as well as observations supporting neutrino+EM studies of blazars (e.g. IceCube Collaboration et al. 2018; Aartsen et al. 2018). Comprehensive performance advances (e.g. LAT Pass 8) and decade–long catalogs further underscore *Fermi* role as the reference facility for time–domain and population studies in the high–energy sky (Ajello et al. 2019; Abdollahi et al. 2020).

4.2.1 The *Fermi* Gamma-ray Burst Monitor

The *Fermi* GBM is the mission's wide–field, non–imaging transient monitor that extends *Fermi*'s sensitivity from the LAT band down to hard X–rays, providing near all–sky coverage and prompt detection/localization of high–energy transients (Meegan et al. 2009b). GBM comprises 14 scintillation detectors: twelve NaI(Tl) modules (each 12.7 cm in diameter by 1.27 cm thick) sensitive to ~8 keV–1 MeV, and two BGO modules (each 12.7 × 12.7 cm) covering ~200 keV–40 MeV. The NaI units are arranged around the spacecraft at different viewing angles to enable relative–rate localization over the full unocculted sky (as shown in Figure 4.6), while the BGOs provide high–energy coverage and cross–calibration with the LAT (Meegan et al. 2009b). GBM's on–board triggering operates primarily in 50–300 keV on multiple timescales (from tens of μs to seconds) and computes an initial sky position from the detector response

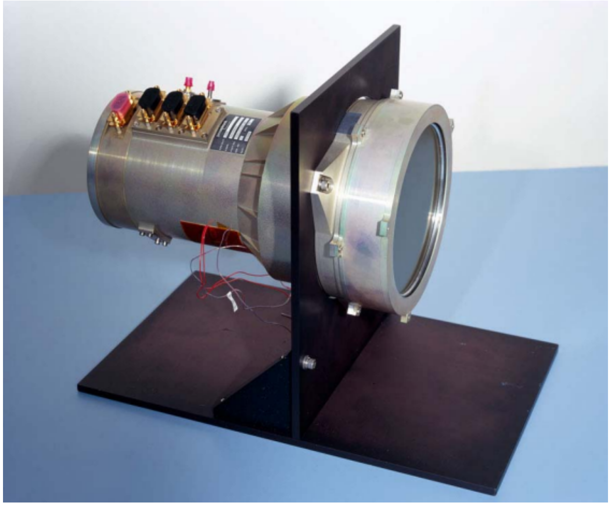

Figure 4.7: NaI(Tl) detector flight unit comprising a circular NaI(Tl) crystal (12.7 cm diameter, 1.27 cm thickness; 5 in × 0.5 in) optically coupled to a photomultiplier tube (PMT).

model; refined ground localizations reach a few degrees for bright events once systematic terms are included (Meegan et al. 2009b). The instrument has delivered cornerstone results for time–domain and multi–messenger astrophysics, including the prompt γ–ray counterpart of GW170817 (GRB 170817A) detected by GBM (Goldstein et al. 2017; Abbott et al. 2017).

Building on the legacy of BATSE aboard the CGRO, which pioneered all–sky monitoring and timing analyses of accreting pulsars, the Fermi/GBM continues this heritage with modern instrumentation and data products.

NaI detector

The twelve NaI(Tl) modules provide the lower–energy response of GBM (~8 keV–1 MeV) and furnish the directional information for GRB localization via relative–rate analysis. Each crystal is a circular disk of diameter 12.7 cm (5 in) and thickness 1.27 cm (0.5 in), enclosed in a hermetic, light–tight aluminum housing to protect the hygroscopic NaI. A 0.6 cm thick N–BK7 glass plate serves as the optical interface to the photomultiplier tube (PMT), bonded to the enclosure with a white Araldite adhesive. The entrance window toward the sky is a 0.2 mm beryllium foil to retain low–energy transmission; a 0.7 mm silicone layer, required for mechanical robustness, ultimately sets the practical low–energy threshold near 8 keV. To enhance light collection, the crystal is wrapped with diffuse reflectors (Tetratex on the front face and PTFE/Teflon around the circumference). A flight unit is shown in Figure 4.7.

The NaI axes are deliberately distributed over the spacecraft to enable near all–sky coverage and to derive burst positions from the pattern of relative count rates, a technique established by *Konus* and by BATSE. The pointing of each detector is specified by a zenith angle measured from the spacecraft $+Z$ axis (nominally aligned with the LAT boresight) and an azimuth measured clockwise from the $+X$ (Sun–facing) side. The complete set of NaI orientation angles in spacecraft coordinates is reported in Meegan et al. (2009b); their physical placement on the spacecraft is illustrated in Figure 4.6.

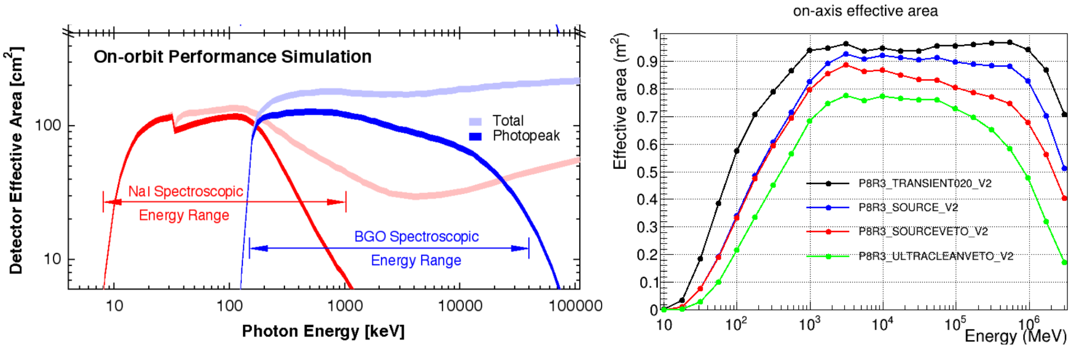

Figure 4.8: Left: *Fermi*/GBM on–axis effective area as a function of energy for the NaI (8 keV–1 MeV) and BGO (0.2–40 MeV) detectors Meegan et al. (2009b). Right: *Fermi*/LAT effective area versus energy (performance curves; credit: https://www.to.infn.it/test_LATperf.

4.2.2 The Fermi Large Area Telescope

The LAT is an imaging wide field, high-energy γ–ray instrument covering from below $\sim$20 MeV to beyond 300 GeV. Built by an international consortium of space agencies, high–energy physics institutes, and universities in France, Italy, Japan, Sweden, and the United States, the LAT is a pair–conversion telescope composed of a precision tracker and a deep segmented calorimeter arranged as a 4×4 array of identical modules. A segmented anticoincidence detector envelopes the tracker to suppress charged–particle backgrounds, and a programmable trigger/DAQ orchestrates event selection and readout. Each tracker module comprises a vertical stack of 18 x, y tracking planes, with per–tray layers of single–sided silicon strip detectors interleaved with high–Z tungsten converter foils. Each calorimeter module contains 96 CsI(Tl) logs in an 8–layer hodoscopic geometry, providing a total depth of 8.6 radiation lengths and furnishing both longitudinal and transverse shower information; the depth and segmentation underpin the LAT's high–energy reach and its background rejection power. The tracker aspect ratio (height/width $\simeq$ 0.4) yields a large instantaneous field of view of $\sim$ 2.4 sr and ensures that most conversion pairs initiated in the tracker are sampled by the calorimeter for energy reconstruction.

Pre–flight performance expectations and the principal science drivers were as follows: (i) rapid identification and dissemination of high–energy GRBs and other transients, with continuous monitoring of variable sources; (ii) an all–sky survey delivering a catalog of several thousand γ–ray emitters; (iii) spectral measurements from $\sim$20 MeV to $>$ 50 GeV for hundreds of sources; (iv) point–source localizations at the sub–arcminute to few–arcminute level ($\sim$0.3–2') for bright objects; (v) mapping and spectroscopy of extended structures (e.g., supernova remnants, molecular clouds, nearby galaxies); (vi) measurements of the isotropic diffuse γ–ray background up to TeV energies; and (vii) exploration of discovery space relevant to particle dark matter.

The Fermi LAT detector

The Fermi LAT detector principle is sketched in Figure 4.9. In the LAT energy range, photon interactions with matter are dominated by pair conversion, which the instrument exploits to distinguish γ rays from the vastly more numerous charged cosmic rays and to reconstruct incident

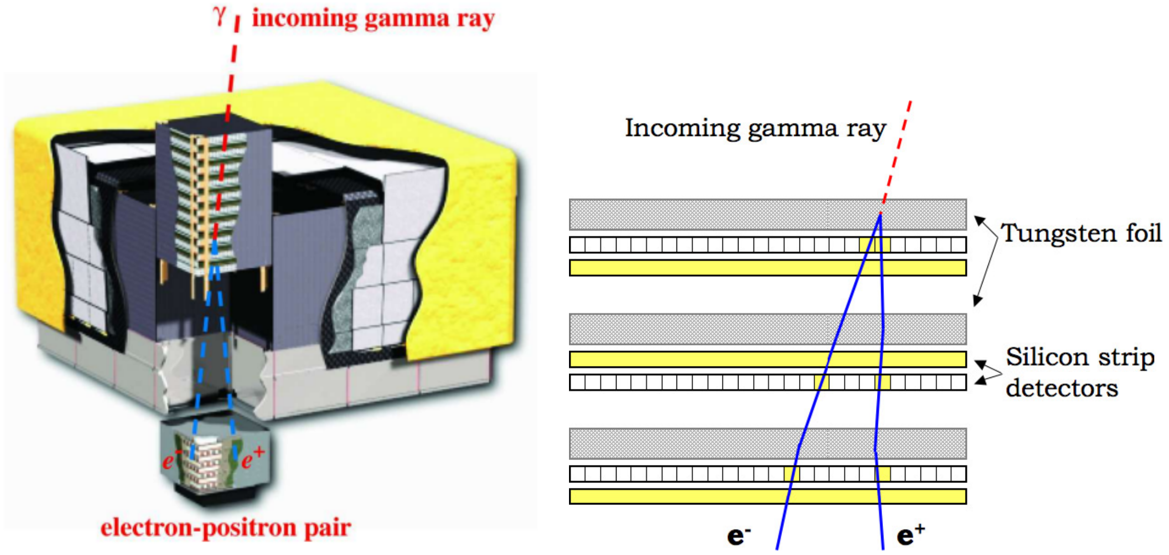

Figure 4.9: Left:Schematic diagram of the *F*ermi/LAT. Overall dimensions which are 1.8 × 1.8 × 0.72 m; the instrument draws ∼ 650 W of power and has a mass of 2,789 kg. Right: Zoomed view of the Fermi/LAT calorimeter. The incoming gamma photon interacts with the Tungsten foil; the pair production electrons are read out by the silicon strip detectors, allowing for reconstruction of energy and direction of the photon.

directions from the trajectories of the resulting e^+e^- pairs. Incident radiation first encounters the segmented anticoincidence detector (ACD), which tags charged particles; surviving candidates traverse thin high-Z tungsten foils in the tracker where conversions occur, and the emergent tracks are measured with 18 x–y layers of single–sided silicon strip detectors. The shower energy is then sampled in a hodoscopic CsI(Tl) calorimeter that provides both longitudinal and transverse information. A canonical LAT γ–ray signature thus comprises: (1) no ACD tile firing consistent with a charged primary, (2) two or more tracks with a common vertex in the tracker, and (3) an electromagnetic shower in the calorimeter. The baseline instrument is modular, a 4×4 array of towers, each $\sim 40 \times 40$ cm^2, with a tracker, calorimeter (total depth $\simeq 8.6\,X_0$ at normal incidence), tower-level trigger/DAQ, and a programmable, instrument–wide trigger that orchestrates a self-triggered readout. This architecture, together with finely segmented subsystems (Si strips, CsI logs, and tiled ACD with wavelength–shifting fiber readout), underpins both background rejection and angular/energy reconstruction performance (Atwood et al. 2009; Moiseev et al. 2007).

The LAT design was matured through detailed Monte Carlo studies and validated empirically with a staged program of accelerator beam tests of engineering models and a high–altitude balloon flight of a full tower–equivalent payload (the BFEM), which exercised end–to–end triggering, readout, reconstruction, and background rejection in a realistic environment (do Couto e Silva et al. 2001; Thompson et al. 2002). Calorimeter element performance and readout dynamic range were characterized with dedicated CsI(Tl) detector studies (Bederede et al. 2004). On orbit, hardware triggering and on–board software filters suppress the raw ∼kHz trigger to a manageable downlink while preserving high efficiency for celestial γ rays, with the ACD providing the first–line veto against charged particles and the tracker–calorimeter topology furnishing powerful residual background discrimination (Atwood et al. 2009; Moiseev et al. 2007).

4.2.3 Fermi data

Within this thesis, I analyze GBM data extracted in the 10 – 900 keV energy range of the NaI detectors, and γ-ray data with E>10 MeV from Fermi/LAT data.

I retrieve NaI data from the GBM burst table[1] by downloading cleaned TTE data of the chosen detector. Data are selected by inserting the GRBs trigger name that is recovered from the NASA Gamma-ray Coordinates Network (GCN) [2].

Fermi LAT Low-Energy Events (LLE) are obtained by the LLE search table [3] by considering the GRB trigger number.

The Data Query interface [4] is used to achieve LAT event files (photon files). The coordinates and observation time windows are obtained from the GCN circulars [5]. A circular mask is considered to increase the signal-to-noise ratio (S/N), selecting photons within a small radius centered around the GRB position.

4.3.0 Insight–HXMT (Hard X–ray Modulation Telescope)

Insight–HXMT is China's first X–ray astronomy satellite, launched on 2017 June 15 into a LEO (altitude ~550 km, inclination 43°). The payload comprises three co–aligned telescopes that collectively span from ~1 keV up to ~3 MeV, enabling strictly simultaneous broad–band timing and spectroscopy in pointed, scanning, and high–energy (GRB) monitoring modes (Zhang et al. 2020). An all–sky γ–ray capability is provided in the high–energy low–gain configuration, with microsecond–scale event time–tagging; the ground segment distributes OGIP/FITS data products and a dedicated HXMTDAS analysis suite integrated with HEASoft Zhang et al. (2020). Operationally, the mission alternates between pointed observations (hours to weeks), small–area scans (e.g., Galactic plane monitoring) supported by direct–demodulation imaging, and a GRB mode that repurposes the active anticoincidence as a near all–sky monitor (Zhang et al. 2020; Guan et al. 2020). Calibration relies on in–flight standards (Crab pulsar/nebula), on–board lines, and dedicated background tracers (blocked apertures), with cross–checks against contemporaneous ephemerides delivering absolute timing accuracies at the few×10 μs level across the payload (Tuo et al. 2022; Li et al. 2020).

4.3.1 Detector suite and key characteristics

Medium Energy (ME). ME covers ~5–30 keV using 1728 Si–PIN pixels arranged in three boxes (nine units total, 32 pixels per module), delivering a geometric area of ~952 cm^2 Zhang et al.

---

[1] see https://heasarc.gsfc.nasa.gov/db-perl/W3Browse/GBB_burst_table

[2] see `https://gcn.nasa.gov/circulars`

[3] see https://heasarc.gsfc.nasa.gov/cgi-bin/W3Browse/LLE_search_table

[4] heasarc.fermi.gsfc.nasa.gov.LAT

[5] `https://gcn.nasa.gov/circulars`

Figure 4.10: Schematic of *Insight–HXMT* and its three co–aligned telescopes (LE/ME/HE).

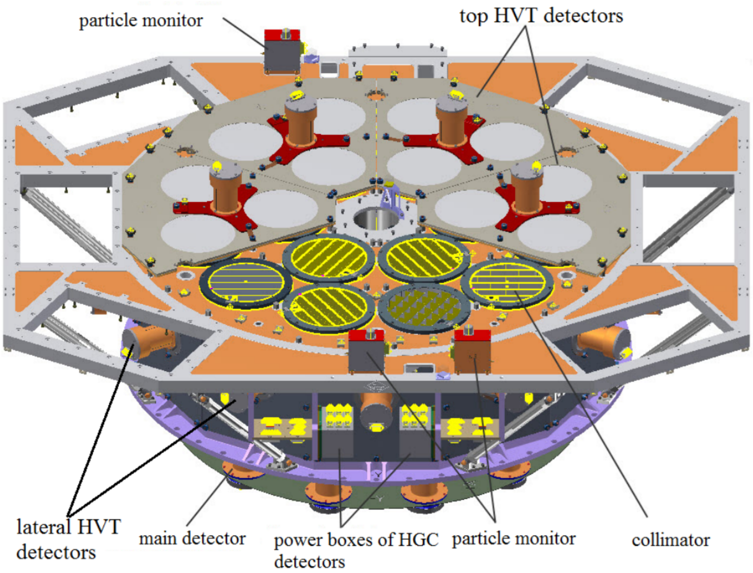

Figure 4.11: Schematic of *Insight–HXMT* payload.

(2020). Slat collimators define 1° × 4° and 4° × 4° FOVs, with additional blocked apertures for background modeling (Zhang et al. 2020). The native time resolution is ~280 μs; absolute timing shows ~10 μs offsets and ~8.6 μs systematics relative to NICER on the Crab (Tuo et al. 2022). The ME electronics do not carry a temperature–compensated oscillator; dedicated corrections mitigate thermal/long–term drifts (Tuo et al. 2022).

Low Energy (LE). LE operates over ~1–15 keV with swept–charge devices (CCD236; three boxes, two modules each; total area ~384 cm^2), offering multiple FOV options (~1.6° × 6°, 4° × 6°, narrow strips) that enable flexible throughput/confusion trade–offs (Zhang et al. 2020). Continuous readout affords ~1 ms native time resolution; a characteristic readout–phase delay is calibrated in orbit (yielding systematics ~16 μs after correction) and must be accounted for in high–precision timing analyses (Tuo et al. 2022; Zhou et al. 2021).

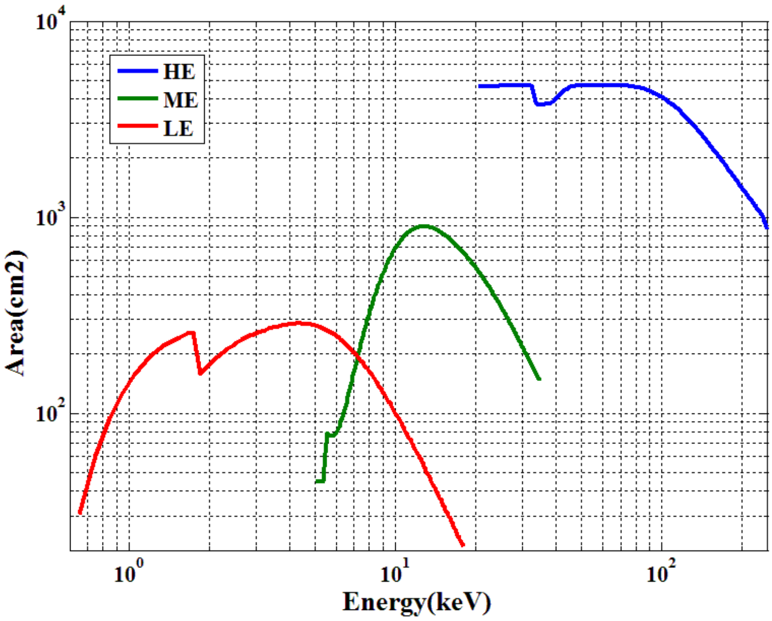

Figure 4.12: The effective areas of the three Insight-HXMT detectors.

4.3.2 Insight–HXMT/HE: architecture and subsystems

The High Energy (HE) payload on *Insight–HXMT* is a collimated hard X–ray instrument optimized for simultaneous spectroscopy, fast timing, and coarse imaging in the ~20–250 keV band Liu et al. (2019). The camera is built from eighteen identical detector units as shown in Figure 4.11. Each unit couples a cylindrical NaI(Tl) primary scintillator to a thicker CsI(Na) crystal in a phoswich configuration (diameter 19 cm; NaI thickness 3.5 mm; CsI 40 mm), viewed by a common photomultiplier and read out with pulse–shape discrimination (PSD) to separate NaI events, used for science, from energy deposits in CsI, which are rejected as background (typically from the rear hemisphere) (Liu et al. 2019). Summed over the array, the total geometric collecting area is ~5100 cm^2 (see Figure 4.12). Active thermal control maintains the detector environment at 18 ± 2°C to stabilize gain and light yield Liu et al. (2019).

Collimators and field of view. Each phoswich is paired with a slat collimator that defines the field of view (FOV) and provides the spatial modulation needed for direct–demodulation imaging during scan observations (Zhang et al. 2020). Fifteen collimators have FOV $\simeq 1.1^\circ \times 5.7^\circ$, two have wider $5.7^\circ \times 5.7^\circ$ FOVs for enhanced background tracking, and one narrow FOV unit is fully blocked by a 2 mm Ta lid to monitor the local instrumental background Liu et al. (2019). The eighteen FOVs are grouped in three orientation sets separated by 60°, yielding the required modulation pattern in scanning mode (Liu et al. 2019). The slat grids are realized in a high–Z absorber; the aperture dimensions (e.g. 6×30×300 mm width×length×height for the $1.1^\circ \times 5.7^\circ$ design) set the angular response Liu et al. (2019).

Particle environment monitoring and protection. To mitigate radiation–induced backgrounds and protect the photomultipliers, HE employs both passive shielding and a set of three particle

Figure 4.13: The ESA *XMM–Newton* observatory in flight configuration (optics module, telescope tube, service module, and focal–plane instruments).

monitors (PMs) that measure the local flux of electrons and protons Liu et al. (2019). Each PM uses a small plastic scintillator read by a PMT; the effective low–energy thresholds are ∼1 MeV (electrons) and ∼20 MeV (protons) (Liu et al. 2019). When the measured flux exceeds a preset limit, e.g. during South Atlantic Anomaly passages, the instrument autonomously lowers or disables the high voltages on the main and anticoincidence detectors to prevent damage and suppress spurious triggers Liu et al. (2019).

Active anticoincidence. Charged–particle rejection is further enhanced by an active anticoincidence detector surrounding the front and sides of the phoswich array. The ACD is segmented into 6 top tiles and 12 side tiles; veto signals are recorded alongside each science event (18–bit flags), enabling offline optimization of anticoincidence selections and background modeling Liu et al. (2019). In routine operations, CsI–based PSD plus ACD vetoes significantly reduce the charged–particle and rear–incident γ backgrounds (Liu et al. 2019).

Electronics, timing, and data interface. All front–end analog signals are processed by the HE power–control and data–handling electronics, which also manage high voltages, calibration control, and telecommand/telemetry via 1553B and LVDS links to the spacecraft bus Liu et al. (2019). For timing, the HE electronics ingest the spacecraft time reference and a 5 MHz clock to ensure stable time–tagging of events; in–orbit cross–calibration with the Crab pulsar shows native microsecond–scale resolution and absolute time alignment at the few×10 μs level once systematic offsets are accounted for (Tuo et al. 2022; Li et al. 2020).

Operations, calibration, and software. Data products follow HEASARC/OGIP conventions for spectra, event lists and timing files, enabling analysis with standard multi–mission tools (HXMTDAS within HEASoft; Zhao 2020; Arnaud et al. 1992; George et al. 1994). Pointed and scanning observations are reconstructed via direct–demodulation imaging (Guan et al. 2020); GRB analyses employ GEANT4–based response and localization tools for the HE/CsI low–gain mode (Zhang et al. 2020). These capabilities make *Insight–HXMT* particularly effective for broad–band spectral–timing studies of compact objects and high–energy transients in coordination with other facilities (Zhang et al. 2020).

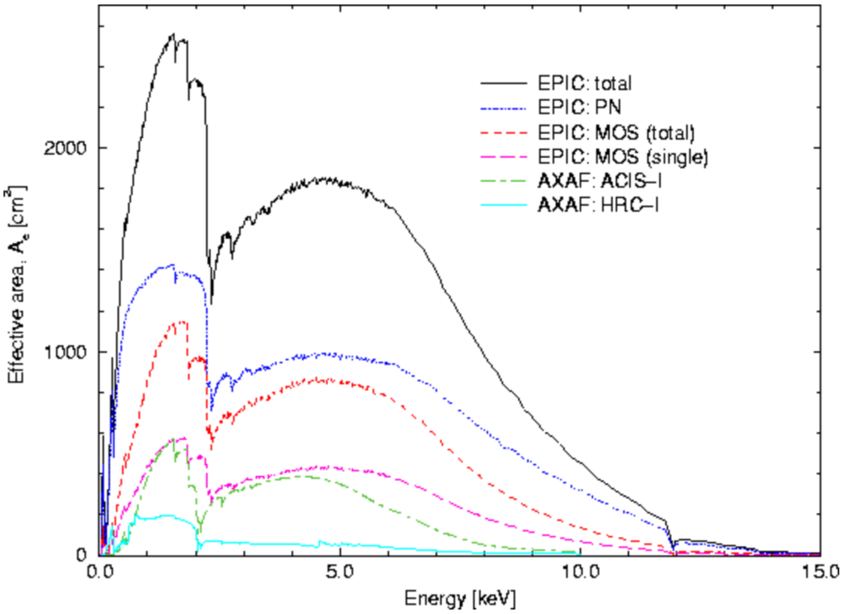

Figure 4.14: Effective-area curves for XMM–Newton in different EPIC configurations—EPIC total, pn only, two MOS, and single MOS—compared on a linear scale with AXAF/Chandra ACIS–I and HRC–I, credit: https://www.mssl.ucl.ac.uk/XMM_Aeff

4.3.3 Insight data

I retrieve Insight-HXMT HE data from the HXMT data archive[6]. Data reduction is operated by following the the HXMT Data Reduction Guide[7]. For the timing analysis of HE data, I apply a timedel=0.00008 (time difference between two adjacent events to discarding spike events). Photon energy is derived by:

$$E = \mathrm{PI} \cdot \frac{370}{256} + 15\ \ (\mathrm{keV}), \tag{4.1}$$

where PI is the value of Pulse Invariant (PI) calculated from the raw values of Pulse Height Amplitude (PHA) of each event.

4.4.0 XMM–Newton

The *X–ray Multi–Mirror Mission XMM–Newton* is ESA’s large X–ray observatory, launched on 10 December 1999 (overall length ∼10 m; launch mass ∼4 t). The spacecraft comprises four principal elements (see Figure 4.13): the *Mirror Support Platform* carrying three Wolter I mirror modules and the co–aligned Optical Monitor (OM), the carbon–fiber *telescope tube* that maintains alignment and focal length, the *Focal Plane Platform* (FPP) with the EPIC imaging cameras (pn, MOS1/2) and the two Reflection Grating Spectrometers (RGS1/2), and the *Service Module* (SVM) housing support subsystems (Jansen & et al. 2001; Mason et al. 2001; Strüder et al. 2001; den Herder et al. 2001). The XMM–Newton optics maximize collecting power over a

---

[6] see `http://archive.hxmt.cn/grb`

[7] see `http://hxmten.ihep.ac.cn/SoftDoc.jhtml`

broad band with specific attention to Fe–K energies around 7 keV. Each telescope is a Wolter I system of 58 nested, gold–coated shells (paraboloid + hyperboloid) to sustain high reflectivity up to several keV (Aschenbach et al. 2000; Jansen & et al. 2001). Performance is characterized by image quality, effective area, and stray–light control. On–axis image quality is described by a nearly energy–independent PSF from ~0.1–6 keV, with modest degradation at higher energies; the FWHM is ~6' (HEW ~15') (Gabriel et al. 2006; Aschenbach et al. 2000). As illustrated in Figure 4.14 EPIC–pn attains a larger effective area than the two MOS cameras because two mirror modules host the Reflection Grating Assemblies, which diffract a significant fraction of the beam to the RGS and thereby reduce the MOS throughput (den Herder et al. 2001; Strüder et al. 2001). Overall, the mirrors are most efficient between ~0.1 and 10 keV, peaking near ~1.5 keV, yet they retain substantial area at 2–7 keV (Jansen & et al. 2001; Gabriel et al. 2006). Stray light from off–axis directions is mitigated by X–ray baffles with concentric annular apertures that block singly–reflected paths (Jansen & et al. 2001).

Key capabilities relevant here are the simultaneous operation of all science instruments, the very large collecting area for a focusing X–ray facility (~1.5×10^3 cm^2 at 1.5 keV for EPIC–pn), arcsecond imaging (FWHM ~6'), moderate EPIC spectral resolution (FWHM ~70–80 eV near 1 keV) with high–resolution dispersive spectroscopy from RGS in the 0.35–2.5 keV band, and long uninterrupted visibility windows (up to ~40 h) enabled by the highly elliptical 48 h orbit (Strüder et al. 2001; den Herder et al. 2001; Mason et al. 2001; Jansen & et al. 2001; Gabriel et al. 2006).

4.4.1 European Photon Imaging Camera (EPIC)

The EPIC instrument on *XMM–Newton* comprises two MOS CCD cameras and one pn–CCD camera that operate simultaneously, delivering sensitive imaging and spectroscopy over a $\sim 30'$ field of view in the ~0.2–12 keV band (Strüder et al. 2001; Jansen & et al. 2001). The angular resolution is set by the mirror PSF (FWHM ~6', HEW ~15') and is essentially common to all three focal planes (Aschenbach et al. 2000; Gabriel et al. 2006).

Cameras and focal planes. The pn camera is built on a single silicon wafer integrating 12 back–illuminated, fully–depleted pn–CCDs with fast parallel readout, yielding high quantum efficiency and very short frame times (Strüder et al. 2001). Each MOS camera employs seven front–illuminated EEV CCD–22 devices arranged in a central+hexagonal pattern; two of the three telescopes host reflection gratings so that a substantial fraction of the beam is diffracted to the RGS, reducing the flux that reaches the corresponding MOS focal planes (to roughly ~40–50%) (den Herder et al. 2001; Gabriel et al. 2006). The different detector technologies and optical paths explain the larger effective area achieved by EPIC–pn relative to an individual MOS camera (see also Figure 4.15) (Strüder et al. 2001).

Observing modes and pile–up control. EPIC offers a set of standard modes tailored to source brightness and science goals: *imaging* (full frame, large window, small window), *timing* (one

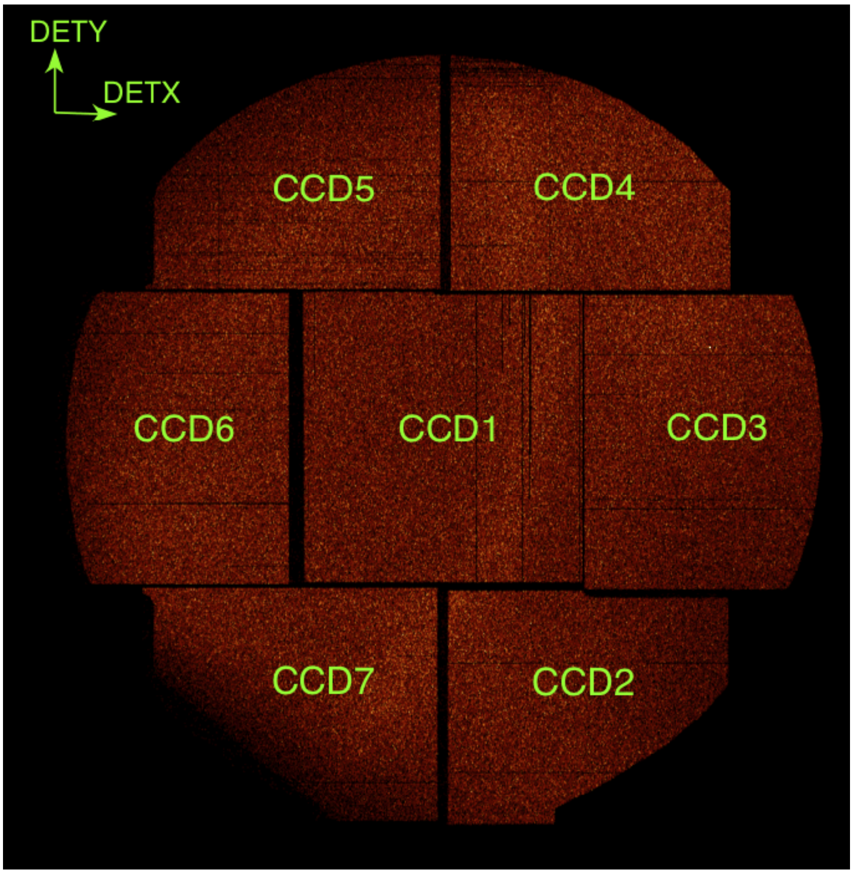

Figure 4.15: **Left:** EPIC–MOS focal plane layout (seven CCDs; each chip covers $\sim 10.9' \times 10.9'$); gaps are inactive border regions as detailed in the *XMM–Newton Users Handbook*. **Right:** schematic of a Reflection Grating Assembly (RGA).

spatial dimension preserved to accelerate readout), and, for EPIC–pn only, the *burst* mode that trades duty cycle for extremely short frame times (Gabriel et al. 2006; Strüder et al. 2001). At high count rates CCDs suffer *pile–up*, when multiple photons are registered within one frame as a single event with summed energy, biasing both the spectrum and the inferred flux. Windowed/timing modes and pn burst mode mitigate pile-up by reducing the effective integration time and readout path, thereby preserving spectral fidelity for bright targets (Gabriel et al. 2006).

5

Precision Timing for Transient sources: Adaptive bin-size approach to time-domain astronomy

Summary

This chapter presents a novel methodological framework for the timing analysis of transient astrophysical events, addressing several open problems in modern high-energy astrophysics. It introduces the mathematical and physical tools required to estimate time delays and their uncertainties between independent photon arrival time lists, with a focus on Cross-Correlation Function–based techniques. Since high-energy detectors record individual photon arrival times through a stochastic process rather than continuous intensities, photon lists are converted into light curves that approximate continuous rate functions, enabling the application of cross-correlation methods. Owing to Poisson statistics, repeated observations of the same signal yield different realizations, causing the measured CCF to fluctuate around the true delay. A Monte Carlo approach is therefore developed to provide a robust estimate of the delay and its associated uncertainty induced by counting noise. Gamma-ray bursts are used as benchmark cases, as their rapid variability and typically low photon statistics make them ideal testbeds. The proposed techniques offer a significant computational advantage and are particularly well-suited for low-count data, overcoming key limitations of traditional fixed bin-size analyses. The work presented in this chapter was published under the title "**Time domain astrophysics with transient sources. Delay estimate via Cross Correlation Function techniques**" in Astronomy & Astrophysics (`https://doi.org/10.1051/0004-6361/202453054`).

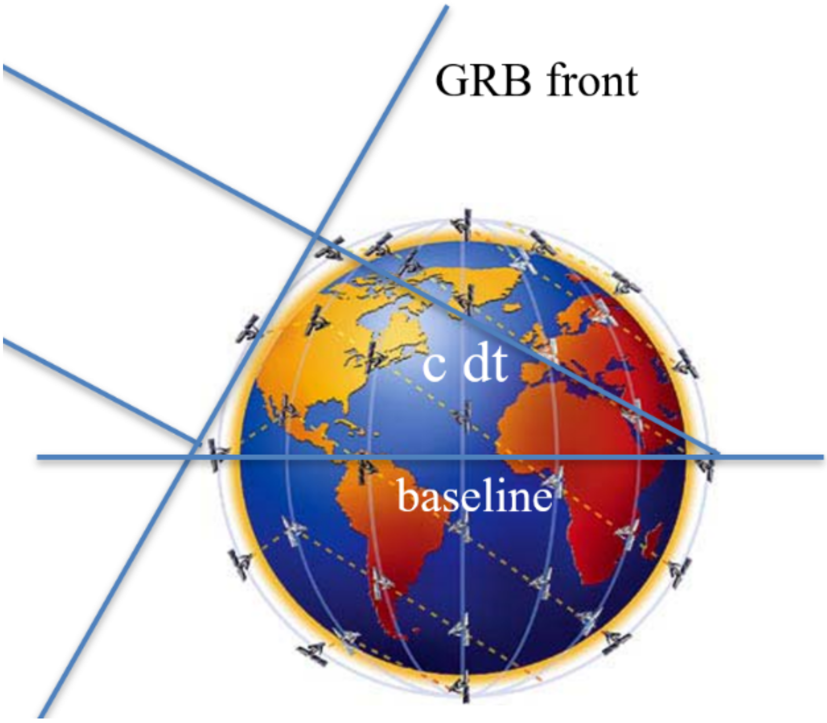

Figure 5.1: Time–difference–of–arrival (TDoA) localization with a satellite constellation. A plane GRB wavefront (oblique lines) sweeps across two spacecraft separated by the baseline $\mathbf{b}$ (horizontal). The measured delay Δt satisfies $c\,\Delta t = \mathbf{b}\cdot\hat{\mathbf{n}}$, where $\hat{\mathbf{n}}$ is the source direction. A single baseline yields a great–circle annulus on the sky; intersecting annuli from multiple baselines (satellites around Earth) produce a compact localization region.

5.1.0 The role of time delays in astronomy

The delay estimate plays a pivotal role in several fields of modern astrophysics. I roughly characterize two types of delays: “spectral“ lags and “temporal“ delays.

Spectral lags may be present when observing a source in different energy bands. Several factors can lead to the formation of delays between light curves obtained by two detectors under such conditions. In the case of GRB, emission mechanisms can drive this effect, spanning a range from a second to even tens of seconds (Giuliani et al. 2008; Frontera et al. 2000; Tsvetkova et al. 2023). Some quantum gravity theories predict that spectral lags depend on a dispersion law for light in vacuo (Amelino-Camelia et al. 1998; Piran 2004). Delays can also be estimated between the continuum and ionized line emission (e.g., Mg II line) light curves of bright sources such as Active Galactic Nuclei (AGN). That allows for probing the AGN geometry and the spatial extent of the accretion disk using reverberation mapping techniques as in Zajaček et al. (2019).

Temporal delays are also crucial for triangulating the position of transient events, which is the purpose of the HERMES-SP mission. Such delays arise between detectors located at different positions in space while observing the same event, as depicted in Figure 5.1. The accurate and rapid localization of the events is key to a rapid and effective follow-up of the source by another in-orbit or ground-based instrument along several energy bands. This article covers the physical and mathematical tools that enable the estimation of this type of delay between two

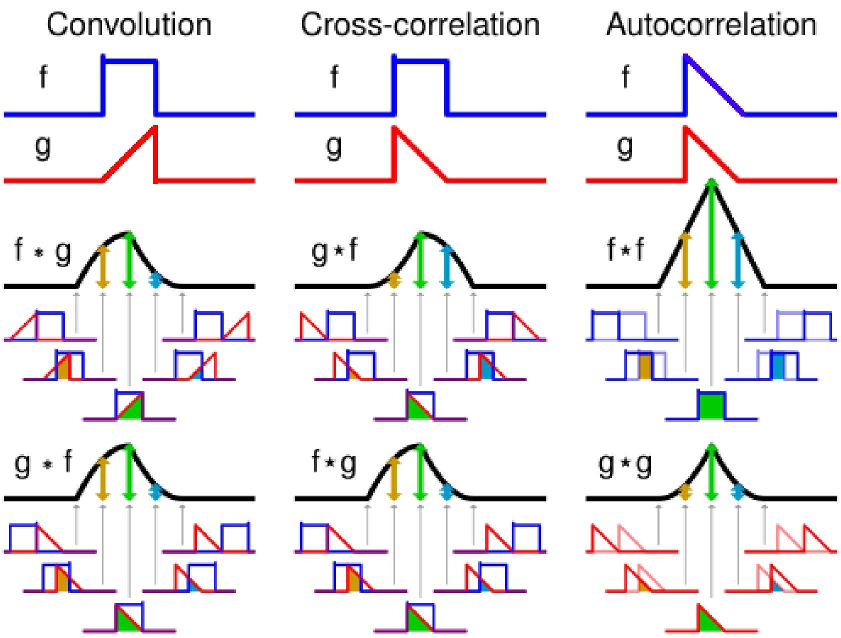

Figure 5.2: Schematic comparison of *convolution*, *cross–correlation*, and *autocorrelation*. Top rows: input signals f (blue) and g (red). Middle and bottom rows: illustrative outcomes for the two possible orderings ($f * g$ vs. $g * f$, $g \star f$ vs. $f \star g$), highlighting that convolution is not commutative for causal signals with finite support, whereas correlation measures similarity as a function of lag. Autocorrelation (right column) peaks at zero lag and quantifies the characteristic width/periodicity of a single signal.

times-of-arrival (ToA) lists.

5.2.0 Probing time with electromagnetic waves

Electromagnetic waves transport, at the speed of light, an amount of energy per square centimeter per second (flux ϕ) along the propagation direction. Plane waves in vacuo are associated with an energy flow whose intensity equals the modulus of the pointing vector $\vec{s}$:

$$|\vec{\mathrm{S}}| = \phi = \frac{\mathrm{E}^2}{4\pi}\mathrm{c}. \tag{5.1}$$

A transitory phenomenon is, by definition, characterized by a variable flux during the occurrence of a transient source.

One can consider a scenario in which a series of theoretically identical detectors can be positioned on the wavefront defined in Equation 5.1. Each detector measures the same intensity at the same time.

On the other hand, if the detectors are displaced arbitrarily in space each detector measures the same intensity at a delayed time, τ, which equals the scalar product of the line of propagation and the vector connecting the positions of the detectors, divided by the speed of light.

By measuring τ, one can deduce the projected distances along the line of propagation and therefore, determine the direction of the wave. At least three detectors are required to determine the direction of propagation, based on geometrical considerations. This can be intuitively understood by considering that three points are sufficient to define a plane in space and, consequently, its perpendicular direction (Sanna et al. 2020). This method, known as the temporal triangulation technique (e.g. Hurley et al. (2013)) relies on the experimental determination of time delays between signals observed by different detectors.

Delays can be obtained by cross-correlating two light curves-defined as the product of ϕ and the instrument's effective area projected along the line of sight-obtained from detector photon lists. To perform the cross-correlation function (CCF), a continuous function $f(t)$ must be derived from each ToA list. Once two continuous functions, f(t), are obtained for a couple of detectors (1 and 2) the delay can be computed using the CCF as depicted in Figure 5.2:

$$\mathrm{CCF}_{1,2}(\tau) = \int_{-\infty}^{+\infty} \mathrm{f}_1(\mathrm{t})\mathrm{f}_2(\mathrm{t}+\tau)\mathrm{dt}. \tag{5.2}$$

The value of τ at which $CCF_{1,2}(\tau)$ reaches its maximum, is the expected delay between the two light curves (MIT 2008).

It is important to note that the detector does not directly measure the intensity of the observed signal, making the derivation of the light curve from a ToA list a non-trivial task.

5.2.1 Statistics of times of arrival

When using a counting device (detector) the energy is recorded discretely, as a list of ToA of photons (quanta of energy). If the wave is monochromatic (single-frequency ν) each energy grain transports the same amount of energy $\mathrm{E} = \mathrm{h}\nu$. In the case of multi-frequency electromagnetic spectra, the same argument can be applied to the “average quanta“:

$$< \mathrm{h}\nu >= \frac{\int_{\nu_{\min}}^{\nu_{\max}} \mathrm{h}\nu\, \mathrm{f}(\nu)\, \mathrm{d}\nu}{\int_{\nu_{\min}}^{\nu_{\max}} \mathrm{f}(\nu)\mathrm{d}\nu}. \tag{5.3}$$

While photons can be detected, the variation of ϕ over time is not directly measurable. Instead, detectors measure the ToA of photons associated with a given rate r(t), where r(t) represents the continuous rate at which photons are detected by the detector. In the ideal case of 100% detection efficiency, the clear relation between ϕ and r is:

$$\mathrm{r}(\mathrm{t}) = \begin{cases} \frac{\phi(\mathrm{t})}{\mathrm{h}\nu}, & \text{multi-frequency spectra} \\ \frac{\phi(\mathrm{t})}{\langle \mathrm{h}\nu \rangle}, & \text{mono-frequency spectra} \end{cases} \tag{5.4}$$

Following Theorem 5.2 in Park (2018) and section A.1, for a given photon arrival rate r(t), the normalized Poisson probability function, associated with the detection of N photons within a time interval Δt, is the Erlang distribution:

$$\mathrm{Q}_{\mathrm{N},\Delta\mathrm{t}}(\mathrm{r}) = \Delta\mathrm{t}\frac{(\mathrm{r}\,\Delta\mathrm{t})^{\mathrm{N}}\mathrm{e}^{-\mathrm{r}\Delta\mathrm{t}}}{\mathrm{N}!}, \tag{5.5}$$

where Δt is small enough r(t) doesn't vary significantly over Δt.

Since the detection of N photons depends on a specific rate chosen among all possible rates, I determine the corresponding confidence interval for the rate at a given confidence level (CL), in accordance with the condition described in subsection A.1.1. As illustrated in Figure 5.3, the same average rate (1 c/s) corresponds to a broader or narrower confidence interval depending on the number of observed counts. These two cases highlight how statistical regimes, defined by low or high counts, influence the accuracy of an otherwise identical rate measurement.

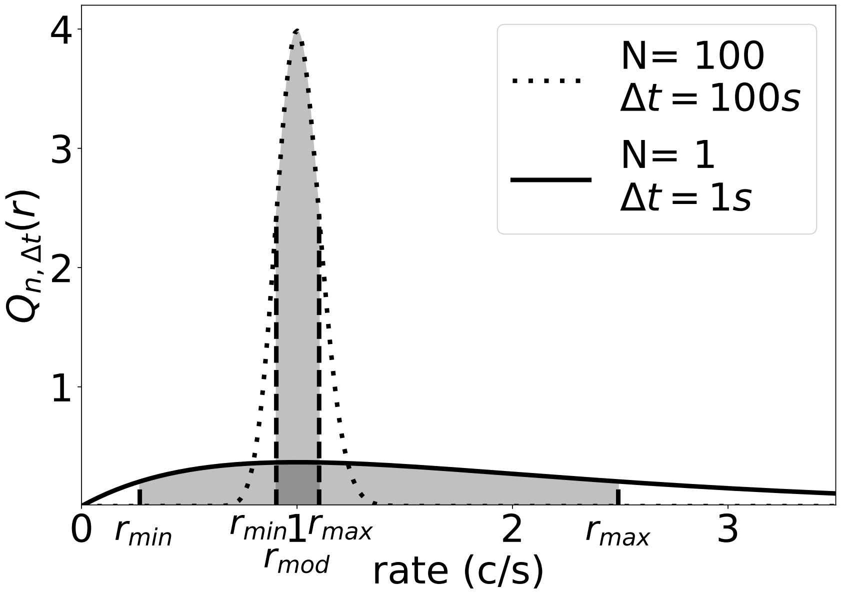

Figure 5.3: Normalized Poisson probability function $Q_{N,\Delta\mathrm{t}}$ shown as a function of the rate r for N=1 (solid line) and n=100 (dotted line). The gray areas indicate a confidence level (CL) of 0.68 corresponding to a 1 σ CL of a Gaussian distribution.

As is standard in statistical analysis, I can define the mode (the most probable value), average, and median of the distribution. The mode is given by $\frac{\mathrm{n}}{\Delta\mathrm{t}}$ (see section A.2) and differs from the median and the mean (defined in section A.2). I note that, for the case $n \longrightarrow +\infty$, the mode, mean, and median converge to the same value.

In general, r_{min} and r_{max} depend on the chosen CL, and I evaluate them numerically using Equation A.6 and Equation A.7, as in section A.3.

Figure 5.4 shows the relative upper and lower errors associated with several CLs. The gray lines in Figure 5.4 correspond to the $1 \pm \frac{1}{\sqrt{\mathrm{N}}}$ symmetric Gaussian relative error, at the 1 sigma CL. It is clear that, for $N \gtrsim 10$, the lower relative limit, $\epsilon^{-}(N,\ 1\sigma\ CL)$, and the upper relative limit, $\epsilon^{+}(N,\ 1\sigma\ CL)$, approximately equal the values of the Gaussian relative symmetric error (see, e.g., Table 5.1), confirming the discussion presented above.

5.2.2 Construction of a light curve

If N photons are detected in a given Δt, the mode of the Poisson distribution, $\mathrm{r} = \mathrm{N}/\Delta\mathrm{t}$, represents the most probable rate value within that interval. This value corresponds to an average rate associated with a confidence interval that depends on the chosen CL and the number of photons N, as illustrated in Figure 5.4.

A statistically uniform representation of a light curve can be achieved by fixing the number

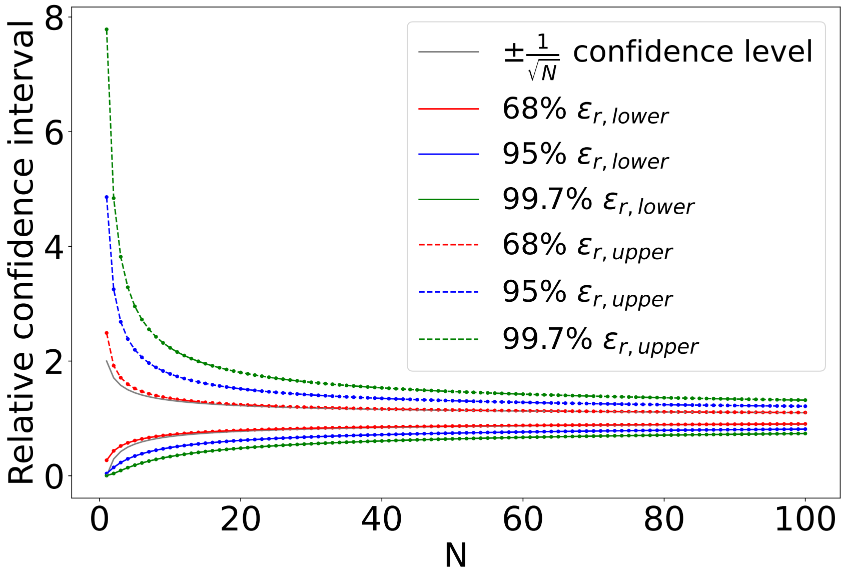

Figure 5.4: Relative errors (for different CLs) as a function of the observed number of events. The upper and lower limits are expressed in units of the mode $r_{mod} = N/\Delta t$ (see section A.2 for a detailed computation). The 1,2 and 3 σ confidence levels are respectively associated with 68%, 95% and 99.7%d. The gray lines represent the $\pm 1/\sqrt{N}$ confidence level as a function of the observed number of events.

of photons, N, used to compute each rate point. Examining a ToA list, N photons are detected during Δt_i:

$$\Delta t_i = t_{i\cdot N-1} - t_{(i-1)\cdot N}, \tag{5.6}$$

and the corresponding rate is:

$$r_i = \frac{N}{\Delta t_i}. \tag{5.7}$$

The time t_i associated with each rate point r_i is the barycenter of the ToA in the time interval $[t_{i\cdot N-1}, t_{(i-1)\cdot N}]$, i.e.,

$$t_i = \frac{1}{N} \sum_{k=(i-1)\cdot N}^{i\cdot N-1} t_k. \tag{5.8}$$

The relative confidence interval is $[\epsilon^-(N, CL),\ \epsilon^+(N, CL)]$, as shown in Figure 5.4 and section A.3, and the absolute error confidence interval is

$$\begin{aligned} r_i^-(N, \Delta t) &= \epsilon^-(N, CL) \cdot r_i, \\ r_i^+(N, \Delta t) &= \epsilon^+(N, CL) \cdot r_i. \end{aligned} \tag{5.9}$$

Once N is fixed, this method guarantees the same relative accuracy for each estimated rate point.

I note that the rate r_i depends on the Δt_i required to collect N photons. By increasing the number of photons N, I obtain smaller confidence intervals (see Figure 5.4). For N >> 1, the relative uncertainties converge as $\epsilon^- \approx \epsilon^+ \approx \frac{1}{\sqrt{N}}$.Increasing the accuracy of the rate measurement requires increasing N, which consequently reduces the temporal resolution.

However, the number of detected photons increases with the detector's effective area (A_{eff}).

Therefore, for a fixed N, an increase in A_{eff} allows smaller time scales to be explored with the same precision.

The method described above allows for a continuous function to be obtained by linearly connecting the rate points $\mathrm{r_i(t_i)}$. The light curve obtained in this way is a continuous function of the generic variable t.

Table 5.1: Poisson vs. Gaussian 1σ relative confidence intervals

N	$\epsilon^{-}(N,\ 1\sigma\ \mathrm{CL})$	$\epsilon^{+}(N,\ 1\sigma\ \mathrm{CL})$	$1-\frac{1}{\sqrt{N}}$	$1+\frac{1}{\sqrt{N}}$
1	0.29	2.49	0	2
5	0.61	1.52	0.55	1.45
10	0.72	1.35	0.68	1.32
50	0.86	1.15	0.86	1.14
100	0.90	1.10	0.90	1.10

Comparison between the Poisson confidence interval [$\epsilon^{-}(N,1\sigma)$, $\epsilon^{+}(N,1\sigma)$] and the Gaussian approximation $1 \pm 1/\sqrt{N}$ as a function of the number of counts N.

5.2.3 Light curve variability

The variability of the light curve obtained using the method described above arises from three different phenomena:

1. The intrinsic variability of the unknown light curve, which represents the genuine variability of the source.

2. The variability induced by the Poissonian detection process. As shown in subsection 5.2.2 and Table 5.1, for N >> 1, the relative weight of this variability scales as the inverse of the square root of the total number N of photons used to build the light curve. In particular, Table 5.1 illustrates that for N = 1 the Gaussian and the $[\epsilon^{-}(\mathrm{N,CL})\ ,\epsilon^{+}(\mathrm{N,CL})]$ relative confidence intervals differ. For N=1, the asymptotic formula $\epsilon^{\pm}_{\mathrm{a}}(\mathrm{N}) = 1 \pm \frac{1}{\sqrt{\mathrm{N}}}$ overestimates the upper limit of the CL and underestimates the lower limit due to the intrinsic skewness of the Poisson distribution.

3. The spurious variability introduced by the linear interpolation between rate points to generate a continuous light curve from a ToA list (see subsection 5.2.2). This variability is independent of the chosen N. Notably, linear interpolation between rate points introduces the minimum possible spurious variability, as it minimizes the required distance to connect consecutive rate points.

Therefore, a trade-off arises. On the one hand, I want to keep N (the number of photons used to build each light-curve point) as small as possible to exploit the detector's minimal temporal resolution for observing the shortest intrinsic variability of the light curve. On the other hand, a larger N is required to minimize the variability induced by Poisson fluctuations. For instance, N=10 represents a good compromise between achieving an approximately symmetric confidence interval, with reasonably small relative errors (around 30%), and exploring finer temporal resolutions.

5.2.4 Cross-correlation Function

Once the rate $r(t)$ is obtained from a ToA list, the CCF can be performed between two rates, $r_1(t)$ and $r_2(t)$, defined for the same time interval $t_1 < t < t_2$:

$$\mathrm{CCF}_{1,2}(\theta) = \int_{t_1}^{t_2} r_1(t) r_2(t+\theta) dt \tag{5.10}$$

The best-fit maximum of $\mathrm{CCF}_{1,2}(\theta)$ corresponds to the best estimate of the delay τ between the two light curves:

$$\max\{\mathrm{CCF}_{1,2}(\theta)\}_{\text{for } t_1<\theta<t_2} = \mathrm{CCF}_{1,2}(\tau) \tag{5.11}$$

Since CCFs are almost symmetric functions (as shown in Figure A.1), I adopt a Gaussian profile to estimate this parameter.

5.3.0 Comparing fixed and adaptive binning for CCF

Several studies employ fixed bin-size light curves to estimate time lags using the CCF (Sanna et al. 2020) or the discrete correlation function (Castignani et al. 2014). However, adaptive binning became particularly advantageous in low-count regimes. While the techniques presented here yield results practically indistinguishable from fixed bin-size methods when applied to high-count-rate signals (e.g., $>10^3$ c/s), their advantages become evident when dealing with sparser data.

Figure 5.5 illustrates how fixed bin-size techniques exhibit clear limitations in low-count scenarios, as commonly encountered in high-energy astrophysics.

Suppose I observe a signal with an average count rate of 10 cts/s, featuring a sharp spike lasting 0.1 s with a peak of 100 cts/s. As shown in Figure 5.5, fixed bin-size light curves may inadequately represent the source's temporal variability. Using this theoretical signal, I simulated ToA lists following the procedure detailed in the next section. The resulting ToA lists were simulated under the same conditions and resemble realistic observations from an X-Ray detector. If a fixed bin size of 1s was chosen (much larger than the spike duration), the rebinning process effectively smooths out the intrinsic variability, rendering the spike invisible. Conversely, selecting a bin size of 0.05 s, comparable to the spike duration, allowed the spike to be captured. However, this came at the cost of introducing significant statistical noise: many bins contained only 0 or 1 photon, leading to substantial spurious variability unrelated to the original signal. In such cases, detecting a single photon corresponded to a 100% uncertainty in the measured rate. These issues directly impact the reliability of CCF-based delay measurements.

To further demonstrate this, I injected a 1-second delay into the theoretical signal and simulated the corresponding ToA lists. In order to estimate the expected delay, a Gaussian fit

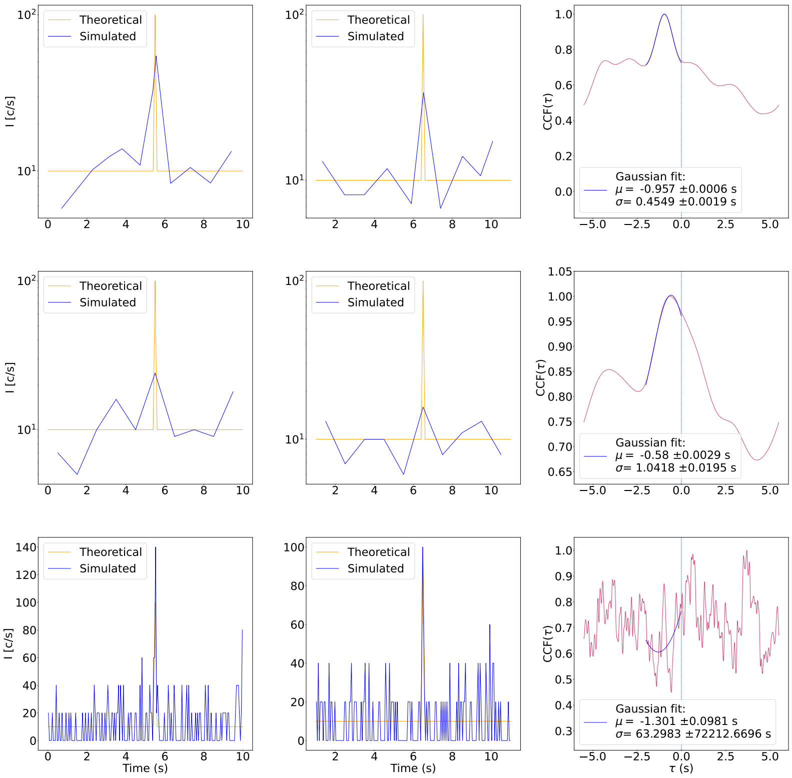

Figure 5.5: Top row: Light curves (left and central panels) and corresponding CCF (right panel) obtained using the adaptive rebinning method (10 photons per bin). Middle and bottom rows: Same configuration, but using fixed bin sizes of 1 s and 0.05 s, respectively. In each row, the left panel shows in blue the simulated signal obtained by rebinning the ToA list generated by simulating the theoretical profile shown in orange. The center panel displays the same signal, delayed by 1 s before rebinning. The right panel presents the CCF between the two simulated light curves shown in the left and central panels of the corresponding row.

was applied to each CCF in Figure 5.5, centered at -1 s and spanning a width of 1 s. As clearly shown, only the adaptively binned light curves recovered the expected delay.

5.4.0 Treatment of errors

Due to the probabilistic nature of the process, when two identical detectors observe the same GRB, the obtained rates, r_1 and r_2, differ due to Poisson fluctuations. Consequently, the delay estimate τ generally differs from the expected value $\tau = 0$s when cross-correlating light curves of the same event as observed by two detectors positioned side by side.

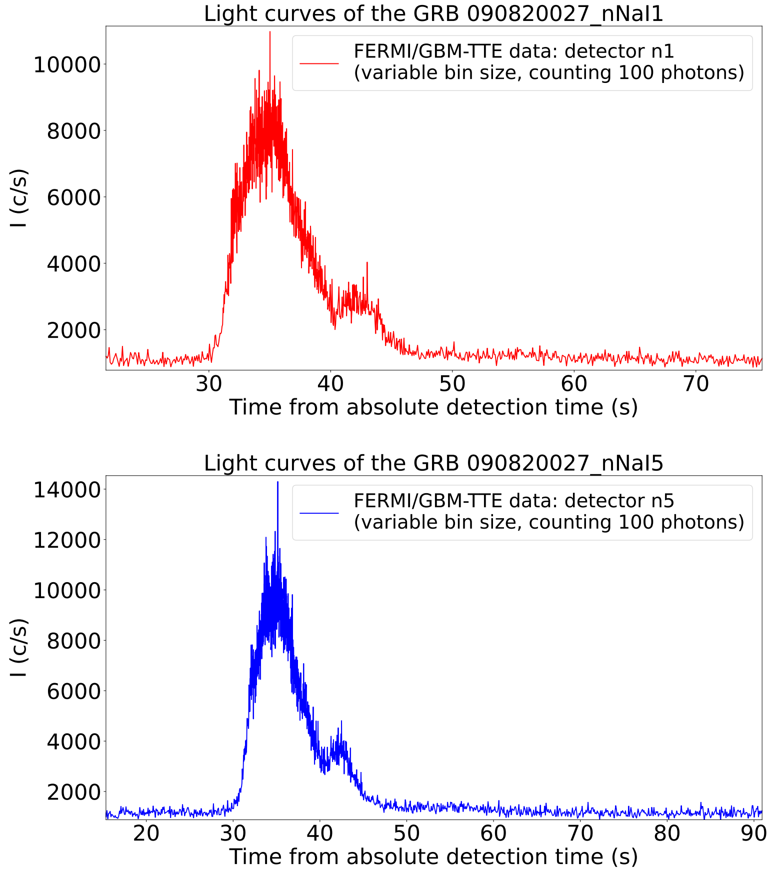

Figure 5.6: Light curves of GRB 090820 obtained by counting N=100 photons per bin. The upper panel shows the n1 detector light curve and the lower panel displays the n5 detector light curve.

For example, I considered the ToA lists associated with GRB 0908207 (see Figure 5.6) as observed by two Fermi-GBM detectors (NaI detector 1 and NaI detector 5: 10 keV - 900 keV). These detectors are physically separated by a maximum distance of 5 m, corresponding to the diagonal of the almost cubic shape of the Fermi satellite, which implies a maximum theoretical delay of $\tau_{\text{th}} = 15$ns (Bissaldi et al. 2009; Meegan et al. 2009a). I also note that the two detectors had different pointing directions at the time of the burst. Consequently, the observed photon count and the respective rate varied depending on the off-axis angle relative to the source.

I estimated the delay between the two detector rate curves, obtained with N = 10 and a sampling resolution of 1μs. The CCF in Figure 5.7 was computed by using the procedure described above, and the CCF upper region was fitted with a Gaussian profile over a 1-second baseline. The lag estimate $\tau_{\rm exp} = (-3.5 \pm 0.068) \times 10^{-2}$s corresponds to the Gaussian best-fit parameter μ and its associated error, as shown in Figure 5.7. It is important to note that the delay estimation result, based on the procedures discussed in the next section, is independent of the Gaussian profile's width.

Taken at face value, this result would imply that some unknown systematic effect has biased the measurement. The single CCF formally yields a significant lag with minimal uncertainty, as it inherently captures the particular statistical fluctuation in the pair of detector measurements. This small uncertainty, however, is purely mathematical and pertains only to the statistical variation specific to that individual realization. Repeating the measurement under identical conditions would result in a different lag estimate due to the random plot showing the CCF fluctuations.

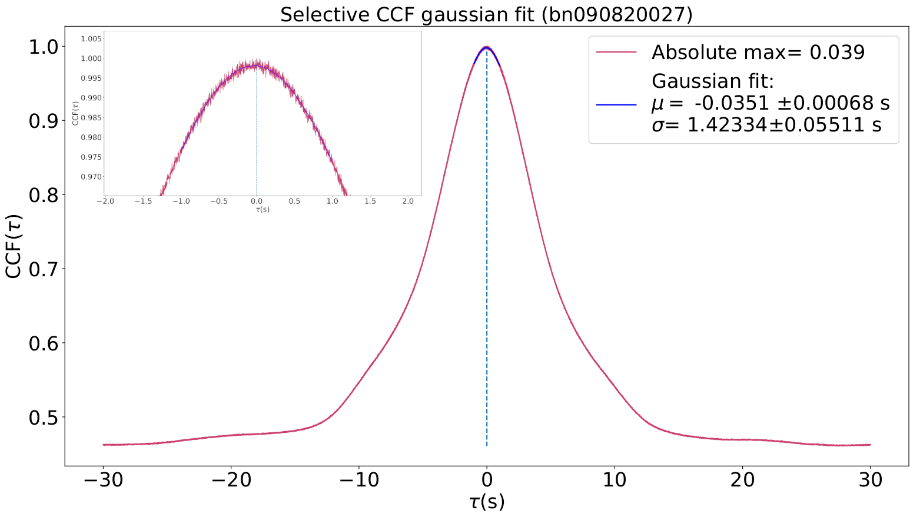

Figure 5.7: Upper panel: Cross-correlation function between light curves from the n1 and n5 detector ToA lists, using a variable bin size of 10 photons per bin and a 1μs resolution. Lower panel: Zoom on the Gaussian fit centroid fluctuation relative to the vertical green dashed line, which indicates the null theoretical delays.

Therefore, a Monte Carlo (MC) simulation approach is necessary to accurately estimate the overall uncertainty, incorporating the full range of possible statistical fluctuations (Zhang et al. 1999). The MC distribution of delays was centered around the best experimental estimate μ and the associated error is the standard deviation of the distribution σ.

5.4.1 The methods

Standard MC methods are based on simulating light curves by the 'flux-randomization' process, as described by Peterson et al. (1998). I explored two alternative methods for delay estimation: the Double Pool (DP) method revisits the concept of flux randomization, while the Modified Double Pool (MDP) method departs from this approach entirely, providing an experimental delay estimate without relying on simulations.

I intentionally retained the background, particularly when it is comparable to the GRB signal, since background fluctuations significantly impact the observed variability. Background subtraction would artificially enhance statistical fluctuations, potentially causing random variations to be mistaken for genuine source variability. Thus, preserving the background allowed us to evaluate variability under realistic observational conditions, avoiding the attribution of false significance to statistically insignificant features. The cross-correlation functions (CCFs) were computed over the T_{90} and the background data intervals of $1.5 \cdot \mathrm{T}_{90}$ before and after the T_{90} interval. These intervals ensure that the resulting CCFs exhibit the characteristic “wings,“ thus enabling a correct interpretation of any potential physical delay.

5.4.2 Double Pool Method: Simulation of Light curves

I propose an alternative method that is conceptually consistent with the real detection process of a detector. This method is based on the generalization of the inversion method in Klein & Roberts (1984) for variable light curves. Instead of using flux randomization, the proposed method generates of a simulated ToA list from a given rate curve r(t) defined over an interval $\mathrm{t}_1 < \mathrm{t} < \mathrm{t}_2$:

$$\int_{\mathrm{T\_SIM[N-1]}}^{\mathrm{T\_SIM[N]}} \mathrm{r(t')dt'} = -\ln\{1 - \mathrm{RND}(0,1)\}, \tag{5.12}$$

where RND(0,1) denotes a value drawn from a uniform distribution between 0 and 1. The ToA T_SIM[N] is recursively simulated, starting from the previous time $\mathrm{T\_SIM[N-1]}$. In the first step, $\mathrm{T\_SIM[N-1]}$ is t_1. This approach emulates the detector measurement process by applying a Poisson arrival process to the rate of the observed signal.

In the case of a constant rate $\mathrm{r(t)} = \lambda$ the integral in Equation 5.12 becomes:

$$\lambda \cdot (\mathrm{T\_SIM[N]} - \mathrm{T\_SIM[N-1]}) = -\ln\{1 - \mathrm{RND}(0,1)\}, \tag{5.13}$$

and each simulated time $T\_SIM[N]$ is:

$$\mathrm{T\_SIM[N]} = \mathrm{T\_SIM[N-1]} - \frac{\ln\{1 - \mathrm{RND}(0,1)\}}{\lambda}. \tag{5.14}$$

As shown in the sketch of Figure 5.8, the integral in Equation 5.12 corresponds to the area of the trapezoid between $\mathrm{T\_SIM[N-1]}$ and T_SIM[N], under the given rate function. According to these considerations, the Poisson arrival process T_SIM[N] can be solved analytically as in section A.4.

The net result is the generation of a ToA list, produced by a counter subject to Poissonian (quantum) fluctuations that observe the rate r(t).

Double Pool method

The DP method exploits this simulation technique to perform the required MC analysis between two detector ToA lists. In principle, one can use a light curve derived from a ToA list

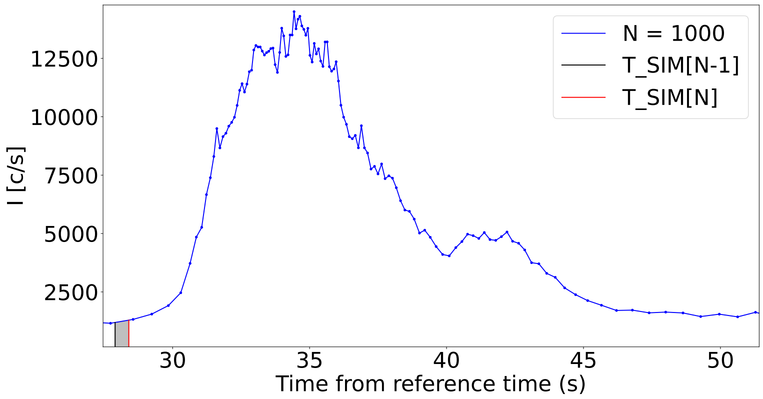

Figure 5.8: Luminosity curve illustrating the trapezoidal integral in Equation 5.12. The black point indicates a generic ToA that was just simulated in the previous step, and the red point represents the ToA to be simulated in the current step.

as a "theoretical" template to generate a large number of simulated light curves. However, as discussed in subsection 5.2.3, the "theoretical" template is affected by Poisson fluctuations that depend on the chosen value of N. This implies that, when using as a template the light curve derived by a particular ToA list as a template, Poisson fluctuations cannot be distinguished from the genuine variability of the source.

This issue can be mitigated in several ways. When estimating the CCF between two detectors, an effective approach is to use the two ToA lists to build two separate templates. Each of these templates is affected by the Poisson fluctuations discussed above. However, since these fluctuations are randomly distributed around the true rate value, any correlation between the Poisson-induced variability between the two curves is effectively removed. For a fixed, reasonable value of N (e.g., N=10) I used each of these two templates to generate several simulated ToA lists (pool of ToA lists) using the inversion method described above. I therefore performed CCF using the method described in subsection 5.2.4 between a pair of ToA lists, each extracted only once from each of the two pools. In this way, Poisson fluctuation variability imprinted on each template did not correlate between the light curves and, therefore, did not bias the CCF result.

By generating a large number of ToA pairs, each belonging to one pool, and cross-correlating the corresponding light curves, I obtained a distribution of delays that was approximately Gaussian due to the central limit theorem. The mean of this distribution represents the expected delay, while the sigma indicates the uncertainty.

5.4.3 Modified Double Pool method

The MDP method allows one to obtain the required delay distribution while providing an exceptional computational time gain. No simulations are required to obtain this distribution. Consider a list of ToA events obtained by a detector. This list can be divided into two independent lists of ToA by generating a random number RND(0,1) for each ToA. Each ToA belongs to one of the two lists depending on the outcome of the RND(0,1). In particular, if RND(0,1)<0.5, the ToA belongs to the first list; otherwise, it belongs to the second list. Since the spatial position of the photons on the detector area is randomly distributed with a flat distribution over the entire detection area, this splitting procedure produces two ToA lists as if obtained by two identical detectors observing the same GRB in the same spatial position, each with an effective area that is half the original one. Cross-correlating the two light curves derived from these ToA lists yields a temporal delay that fluctuates around the expected null value. These fluctuations are purely statistical in origin.

By repeating the splitting procedure with different random realizations, two new ToA lists are obtained. I note that this second couple is not fully statistically independent of the first one, as the original ToA list remains unchanged. However, this statistical dependence is weak because each point in the light curve is constructed using a large number of photons ($N \sim 10$), and it does not significantly affect the computation of the sigma of the distribution, as demonstrated by numerical computation (see subsection 5.5.2).

By averaging each rate point over N photons, each resulting light curve within the same pool represents a distinct Poissonian realization, with each rate value and associated time being approximately statistically independent of any other realization.

By recursively applying this method, I obtain a pool of nearly statistically independent ToA lists from the detector. I explicitly note that the splitting is necessary only to obtain statistically independent ToA lists, and the fact that this method produces a pair at each step is irrelevant to the statistical independence of the final sample of ToA lists in the pool.

Now consider a second detector observing the source. The procedure described above can be applied to obtain a pool of nearly statistically independent ToA lists.

I can now cross-correlate light curves, each extracted from one of the two pools as depicted in Figure 5.9. The average value of this distribution represents the expected delay and the sigma represents the associated uncertainty.

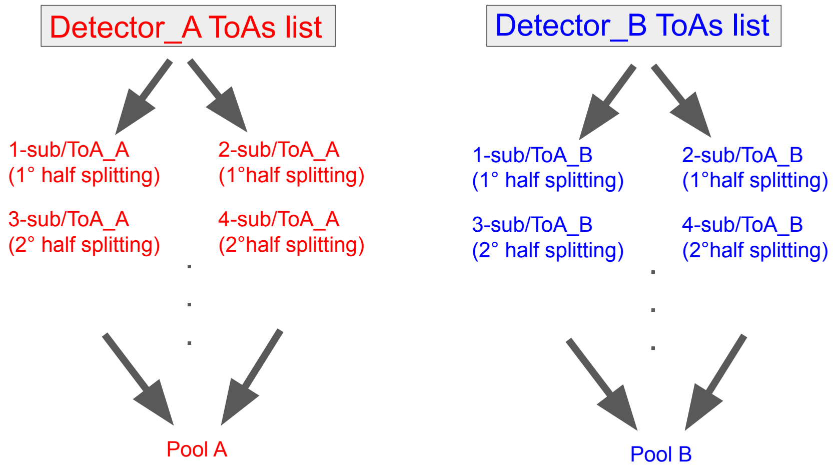

Figure 5.9: Scheme of the MDP splitting procedure.

5.5.0 DP and MDP testing

Gamma Ray Bursts (GRBs) are optimal candidates for testing the capability of the developed tools. I aimed to obtain experimental delays that are consistent with the expected true delays. The GRB sample was randomly selected from bursts observed by Fermi-GBM, considering a broad range of fluence values. As shown in Figure A.1, the discussed CCF techniques ensure a Gaussian profile under all flux and background conditions.

To this end, the test allowed us to discriminate between the proposed procedures, defining the most effective model. During the test, the Gaussian fit guess parameters were fixed for both methods. The number of photons per bin, used to construct the light curves, was set to N=10.

5.5.1 MDP and DP methods comparison

To demonstrate the effectiveness of the two methods, I considered a representative sample of 20 GRBs observed by one Fermi-GBM detector. For each GRB, two independent ToA lists were generated by randomly splitting the considered Fermi-GBM detector data. This approximately corresponded to having two ToA lists with nearly half the effective area of the GBM detector (e.g., as observed by two HERMES detectors). I applied the MDP and DP methods to each pair of ToA lists. The expected theoretical temporal delay between the two ToA lists must be zero (see the discussion in section 5.4).

Figure 5.12 compares the results of the two MC methods applied to the considered sample. Both methods accurately estimate the delays statistically compatible with the true zero delay, considering the standard deviation of each distribution. The residual histograms are compatible with zero.

As evident from Figure 5.12 and Figure A.1, the precision of the estimated delays decreases as the number of photons associated with the source diminishes relative to the background.

The Gaussian fit of the residual Dp distribution indicates that the DP method is less accurate in ensuring compatibility with the true delay. This discrepancy may stem from the simulation process used in the MC procedure. Specifically, the ToA lists from the two starting detectors were employed to define the initial templates for the DP method, which were then fixed during the MC simulations. Since the templates remained fixed throughout the MC simulations, any injected Poisson variability may propagate through all the simulated light curves. This may result in an MC distribution influenced by the Poissonian variability of the initially generated templates. In contrast, the MDP method uses reshuffling of the ToA, which ensures that no privileged light curves are considered.

5.5.2 Demonstrating the independence of the MDP method

Each MDP step is statistically independent, even though the split ToA lists are always derived from the same set of events. Due to the random nature of the halving procedure, each generated light curve represents a specific Poisson realization of the true signal light curve. As a result, each delay constitutes an independent estimate, forming a delay distribution with the correct associated error, as shown in Figure 5.12.

To demonstrate this, I used data from the Insight-Hard X-Ray Modulation Telescope (Insight-HXMT) instrument (Zhang et al. 2018, 2020), specifically focusing on GRB 180113C in the 1 keV - 600 keV energy band. With an effective area of approximately 5000 cm^2, I randomly split the initial ToA list into 200 independent sub-lists. These truly correspond to the lists obtained by 200 detectors observing the same GRB under identical conditions (effective area, detector response, attitude, off-axis angle) and spatial location. I injected a 1s delay into 100 ToA and applied CCF techniques to estimate delays between the delayed and not delayed groups. The resulting distribution of 100 values is shown in Figure 5.10, with an average of $\mu = -0.95$s and an associated error of $\sigma = 0.295s$.

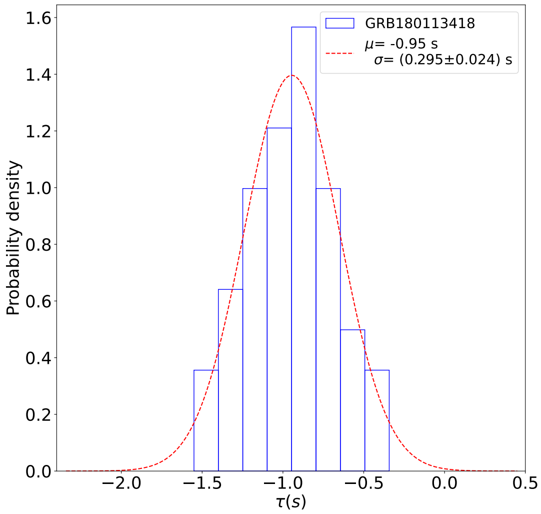

Figure 5.10: Delay distribution is obtained by cross-correlating 100 pairs of ToA lists, derived from the random division of the GRB 180113418 event file. A 1-second delay is injected into one of the ToA lists in each pair.

This experiment remains conceptual, as 200 identical detectors observing the same source are not available. Typically, I estimate the delay between two instruments, so for this analysis, I randomly selected two lists from the sample of 200. I again injected a 1s delay in one of the two lists, and the MDP method in Figure 5.9 was applied to estimate the delay.

The distribution in Figure 5.11 is centered at $\mu = -1.1$s with an associated error of $\sigma = 0.299s$. This result demonstrates the efficiency of the MDP method in estimating the actual existing delays and their associated errors. Moreover, the standard deviation is approximately the same as in the previous case. This indicates that the MDP estimates are accurate even though the generated lists are not strictly statistically independent due to the reshuffling of the same ToA.

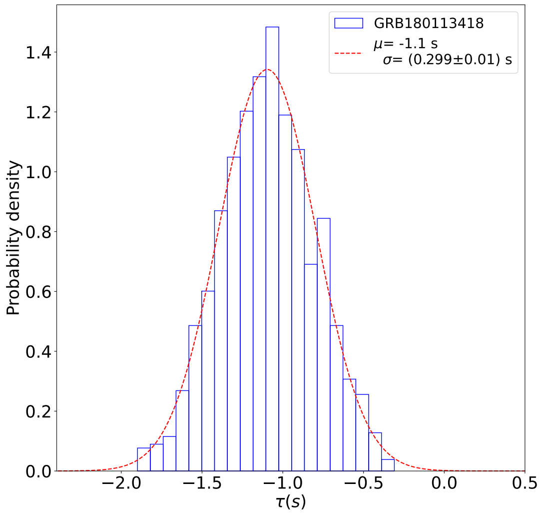

Figure 5.11: Delay distribution obtained by applying the MDP method to two of the 200 ToA lists derived from the random division of the GRB 180113418 event file. A 1-second delay is introduced into one ToA list. The MDP procedure is carried out by randomly splitting the initial ToA lists 500 times, resulting in two pools of 1000 light curves each.

5.6.0 Conclusion

Deriving the light curve associated with the observation of a cosmic source is a nontrivial task that requires careful handling of detector data. This step is critical in timing astronomy, where delays between ToA lists from different detectors are estimated via cross-correlation since this tool is defined on continuous functions. The proposed variable bin size method facilitates the construction of "averaged“ light curves from ToA data, sampling the observed electromagnetic signal with uniform statistical accuracy. This approach allows sampling the light curve at a finer temporal resolution when the intensity is higher. By applying linear interpolation between the rate points of the sampled curves, it is possible to obtain a rate function, enabling the estimation of delays via cross-correlation on continuous functions. It is important to note that linear interpolation introduces minimal variability between consecutive rate points.

The MDP Monte Carlo procedure enables the generation of a distribution of delays, where the mean is the experimental value of the delay and the standard deviation is the associated experimental error.

The results confirm the ability of the MDP method to deliver reliable scientific results, providing a significant increase in both accuracy and computational efficiency. I conclude that the MDP method reduces the impact of intrinsic Poisson fluctuations in input templates, preventing their amplification in Monte Carlo simulations. Despite a $\sqrt{2}$ loss in precision due to halved statistics, it ensures unbiased delay estimates.

Furthermore, the developed techniques demonstrate crucial effectiveness in low-statistics regimes, where traditional methods may struggle to yield consistent results. These techniques remain effective regardless of the transient signal luminosity, although the precision of the

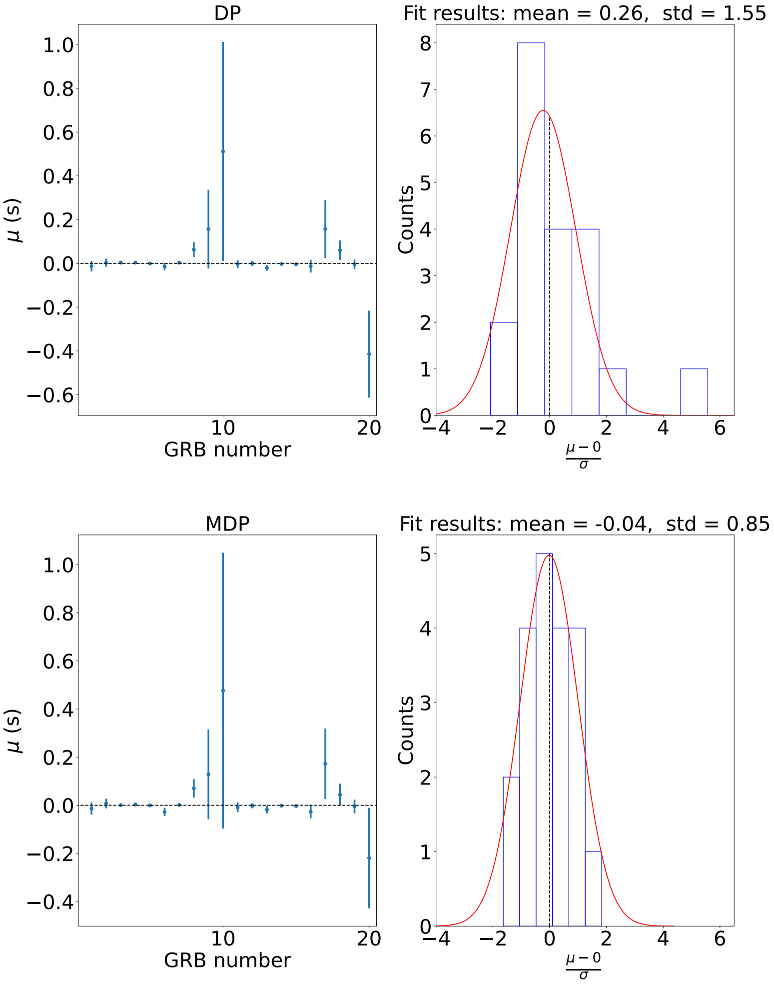

Figure 5.12: Comparison of DP and MDP methods. Left panels: Experimental delays estimated via MC procedures with associated errors shown as vertical error bars. Right panels: Residual distribution for each method in units of sigma.

estimated lag improves with the increasing number of source-associated photons.

6

Transient sources localization via triangulation method

Summary

In this chapter, I extend the temporal-analysis framework to source localization. In the GRBs realm, precise localization is essential for multi-messenger and multi-wavelength studies: they enable secure associations with candidate gravitational-wave events and drive rapid afterglow follow-up.

Thanks to distributed satellite architectures, sky position can be inferred by temporal triangulation, measuring inter-spacecraft arrival-time delays for the same transient.

I quantify time-lag precision as a function of burst count rate using a sample of 230 GRBs from multiple instruments. From these results, I derive the expected localization accuracy versus count rate for representative orbital configurations and instrument effective areas.

Authorship and context. The work presented in this chapter was carried out by the candidate, Wladimiro Leone, and the related article is still in preparation.

6.1.0 Introduction

Realizing the full scientific return from joint GW–GRB events requires arcminute-scale localizations delivered in (near) real time (Pang et al. 2023; Mo et al. 2023). Continuous, wide-field monitoring in the X- and γ-ray bands, coupled with rapid localization pipelines, increases the likelihood of secure GW–GRB associations and enables rapid, multi-wavelength follow-up of the explosion and its environment (Fong et al. 2022). Nanosatellite constellations distributed across space are well suited to this role, as exemplified by the InterPlanetary Network (IPN)[1]. Several mission concepts adopt this architecture, including *HERMES* (Ghirlanda et al. 2024), *ALBATROS* (Burderi et al. 2022), *GRID* (Wen et al. 2021), and *CAMELOT* (Werner et al. 2018). In a triangulating array, the first–order uncertainty on the burst position angle can be approximated as

$$\sigma_{\rm PA} \simeq \frac{\sqrt{\sigma_{\rm Lag}^2 + \sigma_{\rm tpos}^2 + \sigma_{\rm time}^2}}{\langle {\rm Baseline} \rangle \sqrt{n-1-2}}, \tag{6.1}$$

where $\sigma_{\rm Lag}$ is the statistical error on the pairwise cross–correlation delay, $\sigma_{\rm tpos} = \sigma_{\rm pos}/c$ accounts for spacecraft ephemeris errors, and $\sigma_{\rm time}$ is the absolute timekeeping uncertainty of each instrument (Sanna et al. 2020). Here $n-1$ is the number of independent baselines, and $\langle$Baseline$\rangle$ denotes their mean projection along the line of sight; the -2 term reflects the loss of degrees of freedom due to fitted nuisance parameters. Baseline lengths can be increased through orbital design (Ghirlanda et al. 2024), while $\sigma_{\rm tpos}$ depends on onboard navigation (e.g., GPS/GNSS). In low–Earth orbit the inter–satellite light–travel time is $\lesssim 10\,$ms, so $\sigma_{\rm Lag}$ rapidly becomes the dominant contribution for faint events. Mapping $\sigma_{\rm Lag}$ versus count rate is therefore essential for designing and forecasting the performance of future or present triangulation missions.

6.2.0 GRB Sample

To investigate how the estimated lag precision depends on the burst rate, I analyzed a sample of 230 GRBs observed by multiple instruments. For each burst, I extracted light curves across the full calibrated energy range of the corresponding instrument; the relevant bandpasses and data products are summarized in Table 6.1.

Instrument	**Energy Range (keV)**	**Data Type**
Insight/HXMT	20 - 250	HE-Evt data
Fermi/GBM	10 - 900	TTE data
RHESSI	20 - 15 000	1μs event list
INTEGRAL	10 - 900	event list (61 μs)

Table 6.1: Energy range and data type for each dataset used for the analysis. TTE data stands for Time Tagged Event data observed by NaI Fermi/GBM detectors.

The T_{90} duration and the start time of the T_{90} interval are taken, when available, from the fourth

---

[1] `https://heasarc.gsfc.nasa.gov/docs/heasarc/missions/ipn.html`

Fermi/GBM GRB catalog (von Kienlin et al. 2020). For GRBs not observed by GBM, I compiled the same quantities from mission-specific resources: IBAS burst properties from the IBAS results page[2]; *RHESSI* GRB entries from the GSFC archive[3] (as listed in the RHESSI catalog; Lin et al. 2002; Ripa et al. 2009); and *Insight*-HXMT parameters from the catalog of Song et al. (2022). The final sample comprises 50 short and 180 long GRBs. Data retrieval and preprocessing were as follows:

- *Fermi*/GBM (single detector; blue). Products were downloaded from the official burst table[4]. For each event, I selected cleaned TTE data from the NaI detector with the highest peak count rate.
- *RHESSI* (purple). Event lists (native 1 μs time tags) were reconstructed from Level–0 telemetry[5]; for each burst is retained the detector with the highest counts.
- *INTEGRAL*/IBIS/ISGRI (green). Data were processed with OSA v11.2 (Courvoisier et al. 2003), including energy calibration, good–time selection, dead–time correction, and barycentric conversion to the Solar System barycenter. Pixel Illumination Fraction (PIF) weighting was not applied because the background was estimated from the event-derived light curve. ISGRI time tags have an uncertainty of 61 μs (Kuiper et al. 2003). For each GRB, I used the science window containing the burst.
- *Fermi*/GBM (merged; magenta). To emulate an instrument with three times the effective area of a single *Fermi*/GBM detector, merged data products were generated from the three brightest NaI detectors, selected for their similar peak count rates within the same dataset.
- *Fermi*/GBM (halved; orange). To emulate an instrument with half the effective area, a thinned dataset was created by applying Bernoulli sampling ($p = 0.5$) to the photon list from the single brightest NaI detector.
- *Insight*/HXMT HE-Evt (red). Products were retrieved from the official repository[6]. I mitigated instrumental spikes by requiring successive times of arrival to be separated by at least 80 μs[7] before analysis.

For each data set, I converted the FITS files to lightweight ASCII event lists, applied energy selection, and extracted ToA data in the T_{90}.

6.3.0 The method

Each detector records a Single Photon (SP) list for the GRB, whose ToA array represents a Poisson realization of the transient signal.

---

[2] `http://ibas.iasf-milano.inaf.it/IBAS_Results.html`

[3] `https://hesperia.gsfc.nasa.gov/hessidata/`

[4] `https://heasarc.gsfc.nasa.gov/GBM_table`

[5] see https://hesperia.gsfc.nasa.gov/rhessi/software/software-overview/software-overview/index.html

[6] See `https://hxmt.nssdc.ac.cn/data.html`; DOI: `10.12176/02.03.001.2019-03-01`

[7] `http://hxmten.ihep.ac.cn/SoftDoc/501.jhtml`

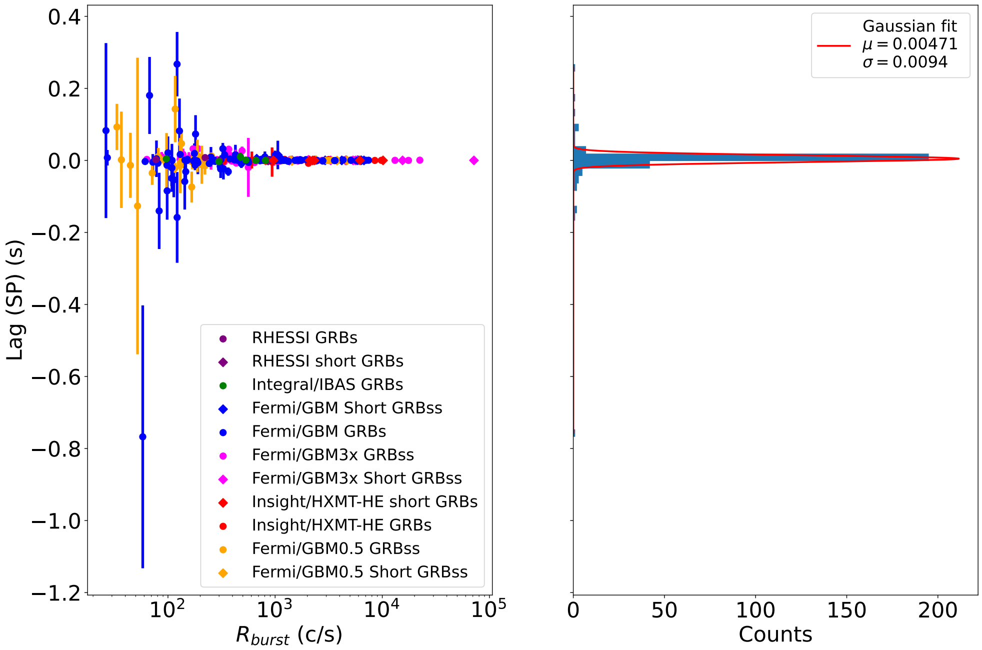

Figure 6.1: The central panel shows the centroids of the lag distributions as a function of the burst rate, R_{burst}. Short GRBs ($T_{90} < 2$, s) are marked with square symbols. The right-hand panel displays the distribution of single-pair (SP) lag measurements about zero.

Using the MDP procedure outlined in chapter 5, an ensemble of N_{pair} quasi-independent ToA series is generated from the SP list. Computing the CCF for independent pairs, drawn from this ensemble, yields a distribution of measured delays.

Because all the N_{pair} originate from the same underlying event (with no intrinsic interband offset), the centroid of the SP lag distribution is expected to be consistent with zero, as shown in Figure 6.1.

The reported uncertainties are the standard error of the mean of that distribution,

$$\sigma_{\bar{x}} = \frac{\sigma_{\text{Lag,SP}}}{\sqrt{N_{\text{pair}}}},$$

where $\sigma_{\text{Lag,SP}}$ is the standard deviation of the SP lag distribution.

In this context, $\sigma_{\text{Lag,SP}}$ represents the statistical limit on the achievable delay precision for a pair of identical detectors observing the same transient under equivalent conditions. The MDP procedure effectively mimics independent realizations of the same Poisson process, providing an empirical estimate of the intrinsic timing accuracy.

This framework enables us to quantify how the timing precision scales with the burst rate. By evaluating the background-subtracted burst rates R_{burst} (computed as described in section C.1) and $\sigma_{\text{Lag,SP}}$ for each event, a clear correlation with R_{burst} is observed.

The trend in Figure 6.2 shows that the precision scales approximately inversely with burst rate; a

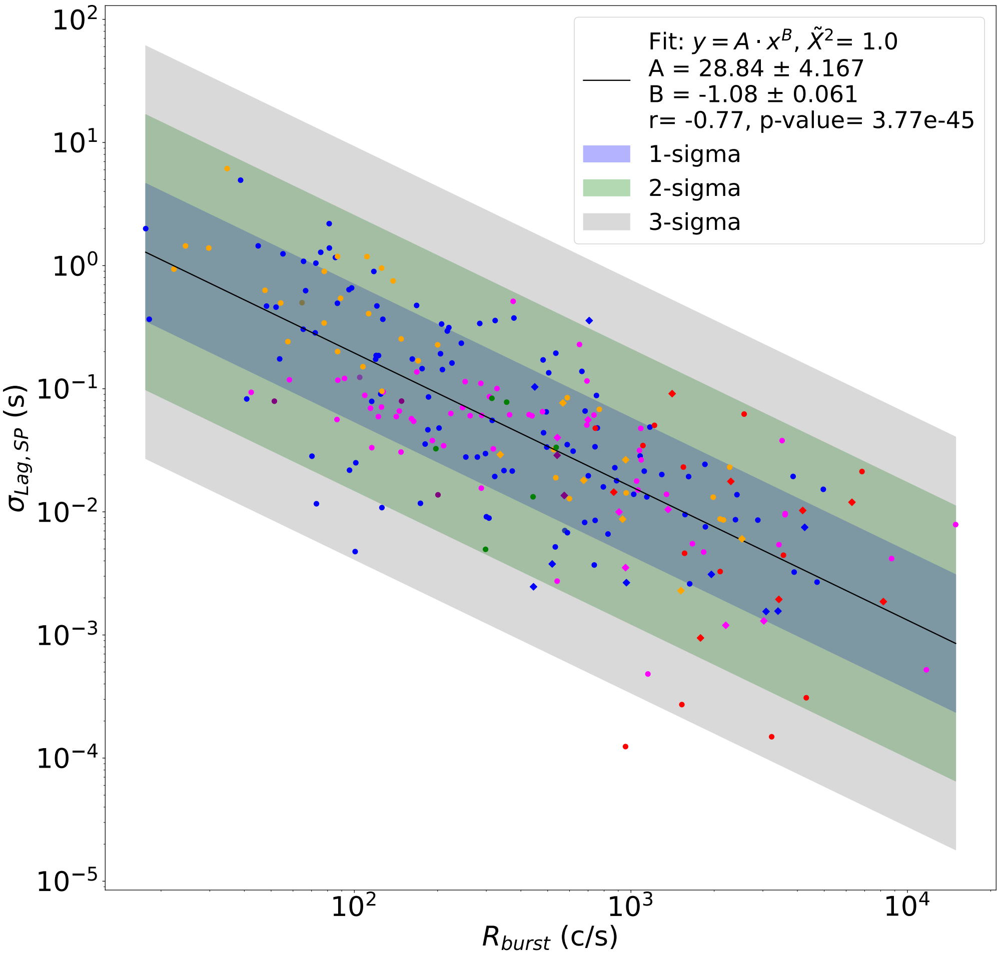

Figure 6.2: Lag precision $\sigma_{Lag,SP}$ as a function of the R_{burst}. Power law parameters are estimated by fitting a linear regression to the data on a log-log scale. The parameters and the n-σ associated region are estimated via the post-fit sigma technique (68% C.L.).

power–law fit (in log–log space) yields

$$\sigma_{\mathrm{lag,SP}} \;\propto\; R_{\mathrm{burst}}^{\alpha}, \qquad \alpha \;=\; -1.08 \pm 0.06,$$

consistent with an $\sim R_{\mathrm{burst}}^{-1}$ dependence.
The Pearson correlation coefficient is $r = -0.77$, with a p-value $p = 3.77 \times 10^{-45}$, indicating a highly significant anti-correlation that is extremely unlikely to arise by chance.

6.4.0 Discussion

Localization accuracy in timing-based triangulation arises from simple geometrical considerations. An uncertainty in the measured time delay uncertainty σ_t corresponds to an uncertainty in the photon path length $\sigma_\ell = c\,\sigma_t$. For a pair of detectors separated by a projected baseline $\langle B \rangle$, this

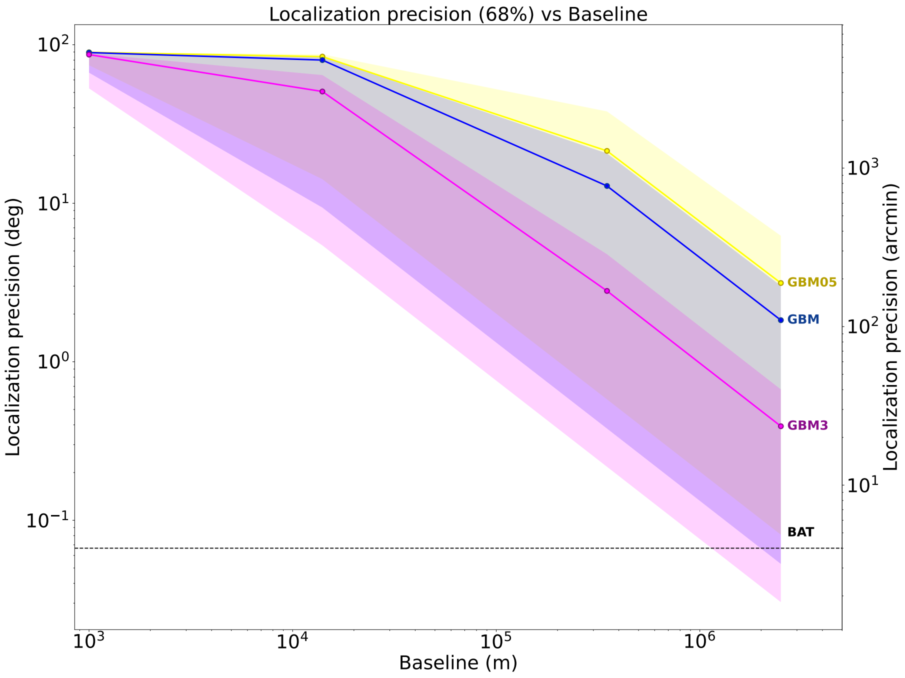

Figure 6.3: Localization precision (68% C.L.) versus inter-satellite average baseline. Curves show three nanosatellite configurations, whose effective areas are respectively GBM05=60 cm^2, GBM=120 cm^2, GBM3=360 cm^2; the horizontal black band marks the typical Swift/BAT on-board localization (4’). Left axis in degrees; right axis in arcminutes.

path-length uncertainty subtends an angular error σ_{PA} on the sky, defined by:

$$\tan(\sigma_{\mathrm{PA}}) = \frac{c\,\sigma_t}{\langle B\rangle}. \tag{6.2}$$

The position accuracy can therefore be written as:

$$\sigma_{\mathrm{PA}} \simeq \frac{180}{\pi}\arctan\left(\frac{c\,\sigma_t}{\langle B\rangle}\right), \tag{6.3}$$

$$\approx \frac{180}{\pi}\frac{c\,\sigma_t}{\langle B\rangle} \qquad \text{for } c\,\sigma_t \ll \langle B\rangle, \tag{6.4}$$

where the small-angle approximation applies in the regime relevant for practical triangulation.

In particular, the σ_t can be expressed as:

$$\sigma_t = \sqrt{\sigma^2_{\mathrm{Lag}} + \sigma^2_{\mathrm{tpos}} + \sigma^2_{\mathrm{time}}} = \sqrt{\sigma^2_{\mathrm{Lag,SP}} + \sigma^2_{\mathrm{tpos}} + \sigma^2_{\mathrm{time}}}.$$

When the lag term dominates the budget (*i.e.*, $\sigma_{\mathrm{Lag,SP}} \gg \sigma_{\mathrm{time}}$, $\sigma_{\mathrm{Lag,SP}} \gg \sigma_{\mathrm{tpos}}$), I obtain, by using the empirical scaling $\sigma_{\mathrm{Lag,SP}} \propto R_{\mathrm{burst}}^{-1.08}$ from Figure 6.2, the final equation:

$$\sigma_{\mathrm{PA}} \simeq \frac{180}{\pi}\,arctan\left(\frac{c\,\sigma_{\mathrm{Lag,SP}}}{\langle B\rangle}\right), \tag{6.5}$$

$$\approx \frac{180}{\pi}\frac{c\,\sigma_{\mathrm{Lag,SP}}}{\langle B\rangle} \qquad \text{for } c\,\sigma_t \ll \langle B\rangle. \tag{6.6}$$

Combining Figure 6.2 with Equation 6.6, I derive 68% confidence regions for the localization accuracy of each instrument: for the burst sample, I compute the mean lag uncertainty $\langle\sigma_{\mathrm{Lag,SP}}\rangle$ and its 68% bounds, then convert these to σ_{PA} via Equation 6.6.

Figure 6.3 illustrates the potential of distributed architectures. While *Swift*/BAT covers ~ 1.4 sr and is periodically Earth–occulted in low–Earth orbit, a triplet of 12U satellites, each with $A_{\mathrm{eff}} = 360\ \mathrm{cm}^2$, deployed on long baselines (e.g., a cartwheel around L1) could achieve comparable localization accuracy with near–hemispheric instantaneous coverage using non–imaging, wide–field detectors.

The estimates are based on triplets of satellites distributed according to the proposed average baselines. Clearly, the accuracy further improves, according to Equation 6.1, by adding more satellites to the constellation.

I note, however, that deploying and operating a CubeSat constellation at L1/L2 would generally increase mission cost and complexity (e.g., launch, communications, and operations) compared to low-Earth orbit, potentially reducing the cost advantage of CubeSat-class spacecraft. A practical mitigation is that such small platforms could be flown as secondary payloads (“ride-share”/“piggyback”) on launchers already targeting L1/L2, thereby amortizing launch costs while retaining the benefits of a distributed, wide-field architecture.

6.5.0 Conclusions

I established a practical link between timing accuracy and localization for GRB triangulation. Empirically, the lag uncertainty scales nearly inversely with burst brightness, $\sigma_{\mathrm{Lag,SP}} \propto 1/R_{\mathrm{burst}}$; propagated through the baseline geometry, this relation directly determines the achievable sky-position precision.

These results provide design guidance for next-generation, multi-satellite GRB facilities targeting rapid, arcminute-scale localizations for multi-messenger astronomy: (i) maximize effective area and timing fidelity, and (ii) pursue extended baselines via distributed architectures (e.g., Earth–Moon or cartwheel constellations).

As an example, a three-node 12U constellation with $A_{\mathrm{eff}} \approx 360\ \mathrm{cm}^2$ per satellite, operating from L1, can achieve *Swift*/BAT-class localizations in favorable cases while monitoring nearly half the sky instantaneously. An additional advantage of these techniques is that the improvement scales with the square root of the number of satellites added, ensuring that localization precision can be improved arbitrarily as more spacecraft are launched into orbit.

7

SpIRIT payload operational activities

Summary

This chapter documents the in–orbit activities of the 6U *SpIRIT* satellite with a focus on the *HERMES* payload hosted on board. It begins with commissioning and first–light operations, then presents the observing program and data reduction workflow used to produce time–tagged event lists, light curves, and preliminary spectra. In flight, the payload encountered a number of routine operational effects that impact background, calibration, timing, pointing, and data continuity. I show in the following the solutions to the presented issues. The chapter culminates with the analysis of bright X–ray targets, including a tentative HERMES detection of the spectrum of Sco X–1, outlining the selection of intervals, background modeling, response application, and consistency checks, and summarizes how the lessons learned informed subsequent operations and analysis choices.

Authorship and context. The work presented in this chapter was carried out by the candidate, Wladimiro Leone, during a research stay at the University of Melbourne (Australia) and formed part of the FM1/*SpIRIT* Status Report submitted to the Italian Space Agency (ASI).

7.1.0 Space Industry – Responsive – Intelligent – Thermal

The *Space Industry – Responsive – Intelligent – Thermal* (SpIRIT) nanosatellite in Figure 4.1 is a 6U Australian CubeSat, successfully launched on December 1st, 2023, with ~ 11.5 kg mass and linear dimensions of approximately 30 × 20 × 10 (Trenti et al. 2024). Continued SpIRIT operations showcase the durable, in-orbit performance of Australian-built hardware. The mission has been running since December 2023 and is expected to conclude only with its natural de-orbit around mid-2027. SpIRIT employs an Australian-designed platform—the Apogee satellite bus from Inovor Technologies—and carries a single primary science payload for advanced gamma- and X-Ray remote sensing: the HERMES instrument (see section 4.1), developed with support from the Italian Space Agency (HERMES Technological Pathfinder) and the EU's Horizon 2020 program (HERMES Scientific Pathfinder).

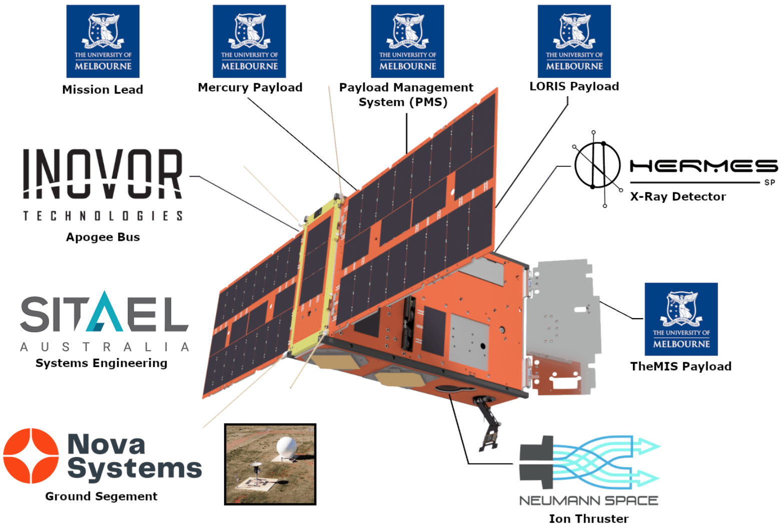

Figure 7.1: SpIRIT spacecraft, payload suite, and partners. The Australian-built platform uses Inovor Technologies' Apogee bus and hosts the HERMES X-Ray payload together with University of Melbourne payloads—Mercury, LORIS, TheMIS, and the Payload Management System (PMS). Systems engineering by SITAEL; ground segment by Nova Systems; electric propulsion by Neumann Space. Credit to https://spirit.research.unimelb.edu.au/.

It has also showcased a suite of innovative technologies as shown in Figure 7.1.

7.2.0 SpIRIT data acquisition

The most recent observations with the *HERMES* payload aboard the *SpIRIT* nanosatellite were carried out on *2025-04-01*. Up to that date, I obtained, in total, 11 *science* observations producing *scientific rate meters* (SRA) data products and 11 observations executed in *Calibration* (CAL) mode. When the payload is in *Observation* mode, HERMES records both SRA and CAL telemetry. However, constraints imposed by the host platform severely limit the number and size of acquisitions; approximately half of the executed observations were configured to downlink SRA data only, which are substantially lighter in terms of onboard storage and telemetry volume. The dominant bottleneck is the VHF downlink: with ~3 ground–station passes per day and a carrier at 401.8950 MHz, the daily telemetry budget is of order 100–200 kB, which sets the practical ceiling for data return per observing day.

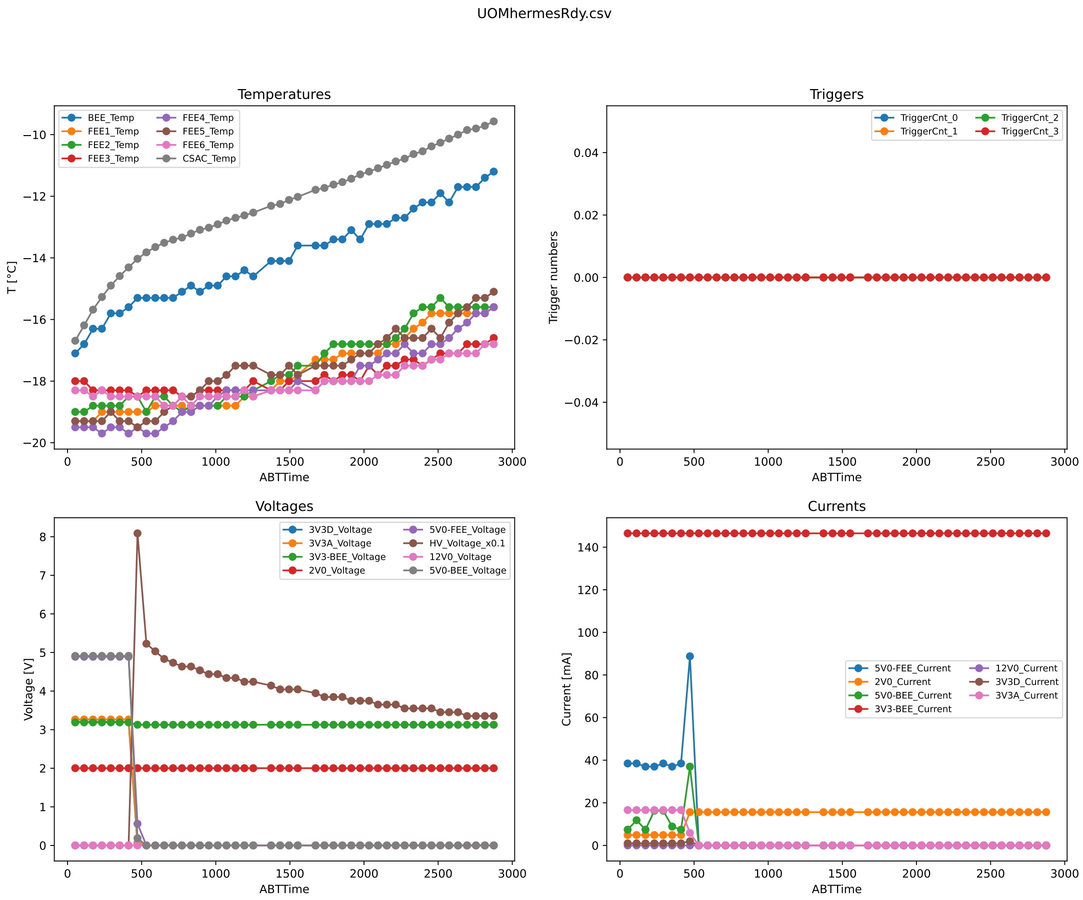

Figure 7.2: The four panels show an example of HERMES housekeeping. These indicate that the thermal and electrical subsystems are in steady state, suitable for SRA acquisition sessions without penalties from payload instabilities. Temperatures lie within a narrow range (10–20 °C on the axis), with nearly parallel trends and no abrupt transients. The 3V3D, 3V3A, 3V3-BEE, 2V0, 5V0-FEE, 5V0-BEE, and 12V0 supplies are flat and free of steps; the high-voltage trace (HV_Voltage, scaled ×0.1) is likewise stable, as desired for detector gain. Since these are housekeeping (HK) from SRA data, there is no direct information on event triggers, as the payload is not operating in photon-counting (“photon”) mode.

The telemetry is downlinked as raw binary files (`.bin`) and converted on the ground into Level 0 `FITS` products by a converter developed within the HERMES collaboration. The Level–0 step performs packet decommutation and integrity checks, orders packets by on–board time, and writes the corresponding headers/metadata, without applying any scientific calibration.

Depending on the downlinked data product, once the Level–0 `FITS` file is available, the scientific extraction proceeds to Level 1 for SRA and to Level 2 for CAL. The telemetry also includes payload housekeeping (e.g., temperatures, high–voltage set/measure, bias, currents). An example is shown in Figure 7.2.

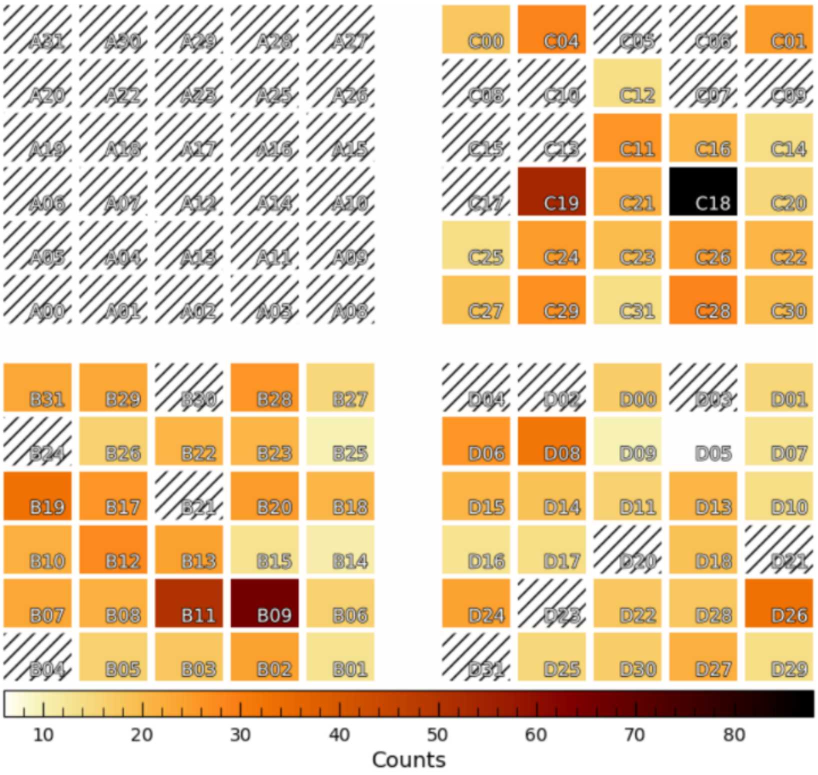

Figure 7.3: MESCAL count map. The grid tracks the physical layout of readout channels on the detector plane, while the color scale encodes the total events recorded per channel. Channels shown with dashed outlines were temporarily disabled (e.g., for excessive noise). Only three quadrants of four are working, granting 70/120 operating channels as treated in Baroni et al. (2024).

I note that, during transport of the payload to Australia, the SpIRIT launch site, one of the four quadrants of the HERMES payload, was damaged and has remained non-operational to date. The total effective area of the detector decreases according to this.

Because only a subset of channels are operational, the effective area is reduced to about 58% of the nominal value as shown in Figure 7.3. Moreover, because the current trigger is set to 5 keV (photons below this threshold are not recorded), the SDDs, whose role is to read the scintillation produced in the crystals by higher energy photons, cannot measure scintillation events induced by incident photons with $E < 90\,\mathrm{keV}$, since their light yield corresponds to $< 5\,\mathrm{keV}$ at the SDD.

7.3.0 SRA data products

The SRA product contains per–quadrant count rates integrated in 100 ms bins and sorted into three predefined energy channels. In the current flight firmware on the HERMES PDHU, the channel boundaries are hard–coded as:

$$\text{low: } 5\text{–}10\ \text{keV}, \qquad \text{mid: } 10\text{–}100\ \text{keV}, \qquad \text{high: } > 100\ \text{keV}.$$

Typical observations last ~10 minutes and accumulate integrated counts in the three predefined energy bands. Because the counts are integrated, it is not possible to perform event–by–event filtering to suppress spurious spikes caused by *retriggering* on active channels.

I define *retrigger* as the condition in which a single physical interaction produces multiple discriminator firings on the same channel. This usually occurs when the front–end is momentarily saturated, or the pulse/baseline has not fully recovered (e.g., due to long scintillation tails or analog ringing), so that the signal re–crosses the threshold one or more times within a short time window. As a result, one event is recorded as several close (in time) counts, artificially inflating the rate in the affected bin. Diagnostic maps, as in Figure 7.3, can reveal pixels suffering from saturation or retrigger issues and show where they sit on the detector plane and which S–mode channel (i.e., which scintillator crystal read out by the SDD photodiode) they are coupled to. In integrated SRA products, this manifests as narrow, impulsive count excesses that cannot be removed a posteriori without the availability of time–tagged events.

Due to the discussed issue, the light curve from SRA data products appears as in Figure 7.4.

7.3.1 Spike filtering of SRA data

The goal of the filter is to suppress instrumental artifacts (spikes or “holes” caused by retrigger/instabilities) while preserving the astrophysical variability recorded in the three active quadrants and their energy channels. The method exploits (i) redundancy across quadrants and (ii) statistical compatibility tests based on Poisson confidence limits at a user–selected confidence level.

Inputs and normalization. Let $q_B(t_i)$, $q_C(t_i)$, and $q_D(t_i)$ denote the per–bin counts (with $\Delta t = 0.1$ s) in the three active quadrants. To place the three time series on a common scale, each quadrant $X \in \{B, C, D\}$ is normalized by a quadrant–specific factor norm_X:

$$q_X^{\mathrm{norm}}(t_i) = \frac{q_X(t_i)}{\mathrm{norm}_X}\,. \tag{7.1}$$

The normalization factors are defined from the total counts accumulated by each quadrant,

$$q_X^{\mathrm{all}} \equiv \sum_i q_X(t_i), \qquad q_3^{\mathrm{all}} \equiv q_B^{\mathrm{all}} + q_C^{\mathrm{all}} + q_D^{\mathrm{all}}, \tag{7.2}$$

as

$$\mathrm{norm}_X \equiv \frac{q_X^{\mathrm{all}}}{q_3^{\mathrm{all}}}\,. \tag{7.3}$$

With this choice, inter–quadrant throughput/exposure differences are removed and the traces become directly comparable across quadrants. In particular, each normalized quadrant has identical integrated counts, $\sum_i q_X^{\mathrm{norm}}(t_i) = q_3^{\mathrm{all}}$, so that residual differences are purely temporal.

For $N_\sigma \in \{1, 2, 3, 4, 5\}$ I tabulate Poisson *relative* bounds $\left[\epsilon^-(n, N_\sigma),\, \epsilon^+(n, N_\sigma)\right]$ from Leone et al. (2025a) for integer counts $n = 0, \ldots, 100$, and I adopt the large-n Gaussian limit $1 \mp N_\sigma/\sqrt{n}$ for $n > 100$. Importantly, these bounds are applied to the *raw* (un-normalized) counts in each time bin, which are Poisson-distributed by construction. Any subsequent conversion to rates or re-normalized quantities (e.g., division by the bin width, effective area, or an overall light-curve normalization) is performed *afterwards* as a deterministic scaling, so that the confidence bounds are propagated by the same multiplicative factor. For the zero–count case I adopt the one–sided upper limit at the chosen CL.

Local baseline. At each time t_i, I estimate a local baseline $q_{\mathrm{mean}}(t_i)$ as the mean, over a sliding window W of width 2×25 bins (shortened near the edges), of the three normalized series concatenated:

$$q_{\mathrm{mean}}(t_i) = \left\langle q_B^{\mathrm{norm}},\ q_C^{\mathrm{norm}},\ q_D^{\mathrm{norm}} \right\rangle_{t\in[i-W,\,i+W]}.$$

Pairwise compatibility tests. For each pair $(X, Y) \in \{(B, C), (B, D), (C, D)\}$ let $n_X = q_X^{\mathrm{norm}}(t_i)$ and $r_X = n_X/\Delta t$ (similarly for Y) be the normalized counts and rates per quadrants, respectively. I define the *multiplicative* Poisson N_σ confidence bounds for $n > 0$ as

$$I_X = \Big[r_X\, \epsilon^-(n_X, N_\sigma),\ r_X\, \epsilon^+(n_X, N_\sigma) \Big], \qquad I_Y = \Big[r_Y\, \epsilon^-(n_Y, N_\sigma),\ r_Y\, \epsilon^+(n_Y, N_\sigma) \Big].$$

For the zero–count case I adopt a one–sided upper limit, so if $n_X = 0$ then $I_X = [\,0,\ r_0^{\mathrm{UL}}\,]$ (and analogously for Y).

The pair (X, Y) is deemed *compatible* when the two intervals overlap:

$$I_X \cap I_Y \neq \varnothing.$$

Equivalently, ordering the two rates so that $r_{\max} = \max(r_X, r_Y)$ with count $n_{\max}$ and $r_{\min} = \min(r_X, r_Y)$ with count $n_{\min}$,

$$\text{compatible} \iff r_{\max}\, \epsilon^-\big(n_{\max}, N_\sigma\big) \ \le\ r_{\min}\, \epsilon^+\big(n_{\min}, N_\sigma\big).$$

If $n_{\min} = 0$ in the right–hand side, use r_0^{UL} as the upper bound for that term.

Anomaly identification rules. From the three pairwise results (BC, BD, CD) at t_i:

- **Single outlier (spike or hole).** If the two pairs that *both* involve the same quadrant are *incompatible* and the remaining pair is *compatible* (e.g. BC = F, BD = F, CD = T), the shared quadrant ($X \in \{B, C, D\}$) is flagged as anomalous. Classify it as a *spike* if $q_X^{\mathrm{norm}}(t_i) > 0$ (after renormalization/rounding) and as a *hole* if $q_X^{\mathrm{norm}}(t_i) = 0$.
- **Triple incompatibility.** All three pairs are *incompatible* (i.e., BC = F, BD = F, CD = F).
- **One non–zero channel (two holes).** Exactly one among ($q_B^{\mathrm{norm}}, q_C^{\mathrm{norm}}, q_D^{\mathrm{norm}}$) is non–zero, compare it to q_{mean} (and the zero–count upper limit).
- **Three zeros.** If all three are zero, test compatibility against q_{mean} and the zero–count upper limit.

Cleaning strategy.

- *Single outlier:* replace

  $$q_X^{\text{norm}}(t_i) \;\leftarrow\; \frac{q_Y^{\text{norm}}(t_i) + q_Z^{\text{norm}}(t_i)}{2},$$

  where $\{Y, Z\} = \{B, C, D\} \setminus \{X\}$. Log the event (time, channel, before/after) and whether it was a spike or a hole.

- *Triple incompatibility:* assume a simultaneous disturbance and restore all three channels toward the local baseline (short temporal average):

  $$q_B^{\text{norm}}(t_i),\; q_C^{\text{norm}}(t_i),\; q_D^{\text{norm}}(t_i) \;\leftarrow\; q_{\text{mean}}(t_i),$$

  with $q_{\text{mean}}(t_i)$ computed over short windows W' (typically ±5 bins centrally; 5 forward/backward bins at edges).

- *One non–zero channel (two holes):* if the non–zero channel is *incompatible* with $q_{\text{mean}}(t_i)$, pull it to the baseline,

  $$q_X^{\text{norm}}(t_i) \;\leftarrow\; q_{\text{mean}}(t_i).$$

  also set $q_Y^{\text{norm}}, q_Z^{\text{norm}} \leftarrow q_{\text{mean}}(t_i)$.

  The heal is operated if the two–hole filling flag (`hf2`) is enabled, otherwise the pattern is only flagged.

- *Three zeros:* if *incompatible* with the baseline and the three–hole filling flag (`hf3`) is enabled, set

  $$q_B^{\text{norm}}(t_i),\; q_C^{\text{norm}}(t_i),\; q_D^{\text{norm}}(t_i) \;\leftarrow\; q_{\text{mean}}(t_i);$$

  otherwise only flag the event.

Renormalization and outputs. Finally, de–normalize counts and form rates are defined as

$$q_X^{\text{healed}}(t_i) = q_X^{\text{norm}}(t_i) \times \text{norm}_X, \qquad r_X^{\text{healed}}(t_i) = \frac{q_X^{\text{healed}}(t_i)}{\Delta t},$$

yielding three “cleaned” rate curves suitable for timing analyses (e.g. CCF and triangulation).

Control parameters. Key knobs are: confidence level N_σ (typically 3–5 σ), sliding–windows half–width $W \simeq 25$ bins and $W' \simeq 5$ and the behavior for two/three holes (hf2/hf3 filling flags). Larger N_σ makes the filter more conservative; shorter windows reduce temporal smearing at the cost of a noisier baseline.

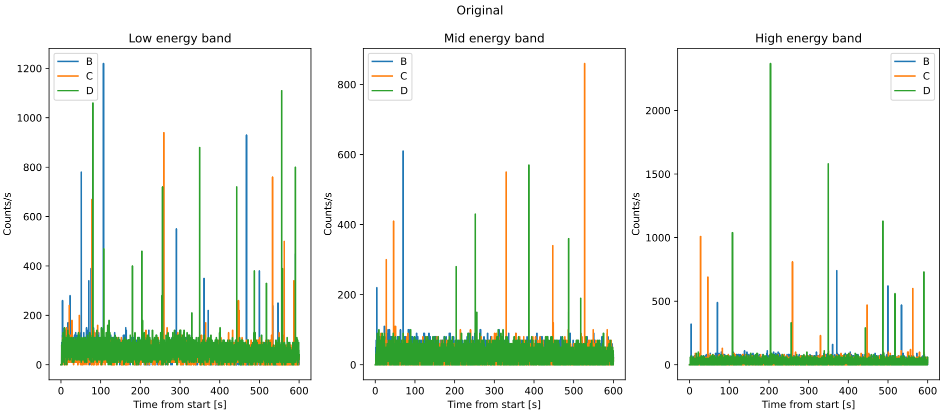

Figure 7.4: Non-filtered rate curve (i.e., counts s^{-1}) of the 10 minutes 06.06.2024 SRA observation.

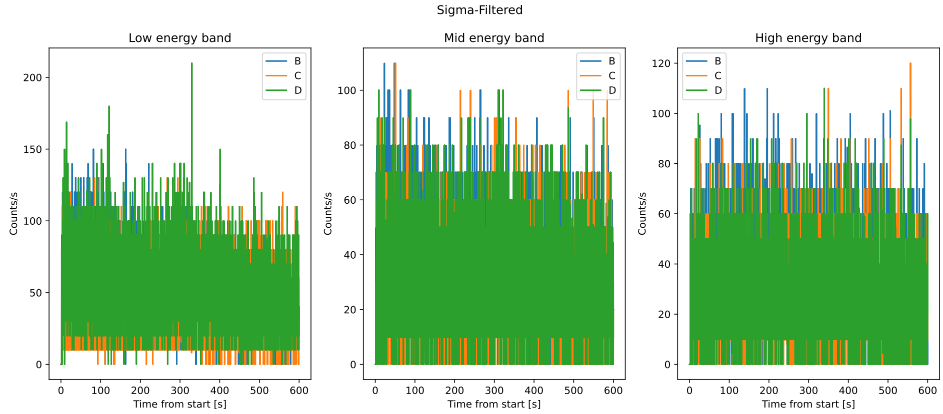

Figure 7.5: Filtered rate curve (i.e., counts s^{-1}) of the 10 minutes 06.06.2024 SRA observation. The flag hf2 and hf3 are enabled.

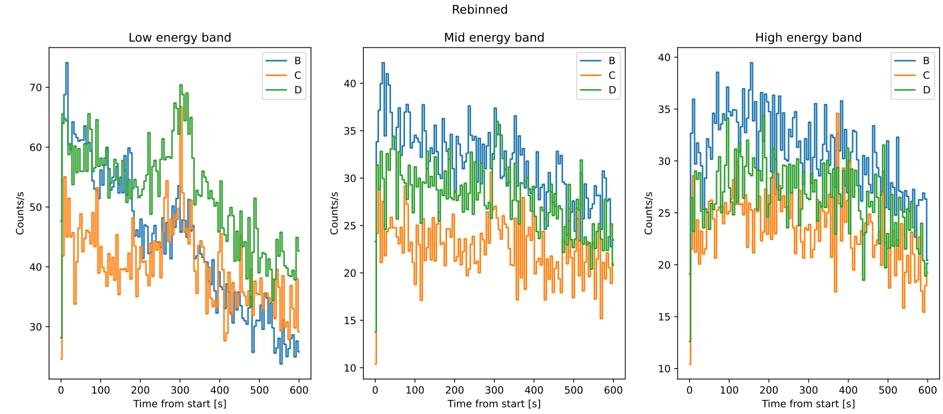

Figure 7.6: Filtered SRA counts per quadrant (B, C, D) for the three energy bands (low, mid, high), rebinned by a factor of 25 (bin width 2.5 s).

Notes and limitations. Because the filter is driven by inter–quadrant coherence, a *real* phenomenon present in only one quadrant and absent in the others (e.g. a genuinely energy–selective glitch) could be attenuated. To mitigate this, replacement occurs only when two channels agree, and the outlier is statistically incompatible with *both* at the chosen N_σ. The local baseline is used only in degenerate cases (2–3 holes or triple incompatibility) and over short windows to minimize impact on intrinsic variability.

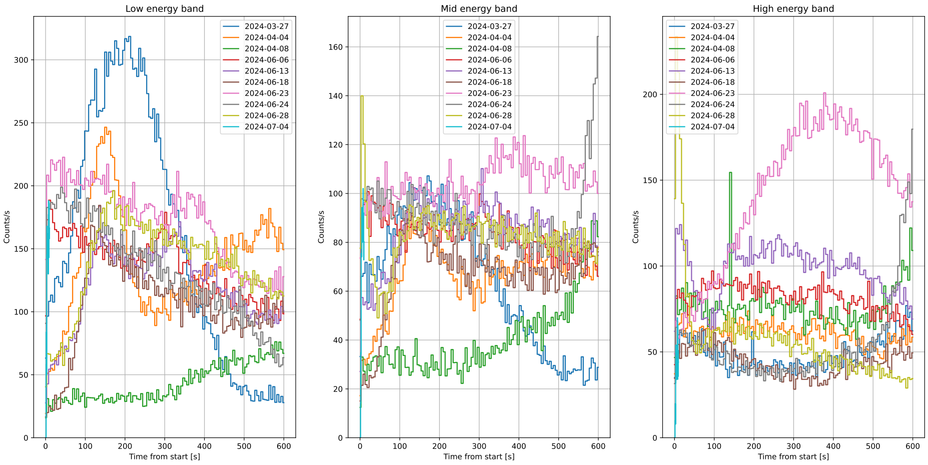

Figure 7.7: All SRA observations: count–rate light curves in the low, mid, and high energy bands (top to bottom). Each panel spans 0–600 s from the start, rebinned by a factor of 25; dates label the corresponding epochs.

7.3.2 Rebinning of filtered SRA counts

The filtered SRA counts (see, e.g., Figure 7.5) are rebinned to emphasize the temporal trend over the observation. Let $q = (q_0, \ldots, q_{N-1})$ be the one–dimensional array of *filtered counts per native bin* (time step $\Delta t = 0.1$ s), and let $f \in \mathbb{N}$ be the rebinning factor. The routine

$$\texttt{rebin}(q, f) = \left(\sum_{k=0}^{f-1} q_{j+k} \right)_{j=0}^{M-1}, \qquad M = \frac{N}{f},$$

is implemented as `q.reshape((q.shape[0]//f, f)).sum(axis=1)`: it groups consecutive, non–overlapping blocks of length f and *sums* their counts, yielding a series of length $M = N/f$ defined on bins of width $f\,\Delta t$. This is the appropriate operation for Poisson counts, since both counts and variances add across independent bins.

7.3.3 SRA light curves

The filtered SRA light curves, in Figure 7.5, show smooth temporal trends once quadrant glitches are removed. The corresponding per–epoch, rebinned and integrated overview over the three quadrants is reported in Figure 7.7.

By considering the detector trigger (as treated in section 7.2) and the typical soft spectra of the Cosmic X-Ray Background (CXB), the low–energy band carries the largest count level, the mid–energy band lies intermediate, and the high–energy band is the faintest (Gruber et al. 1999a). This monotonic ordering is stable across the observation and across epochs, indicating a consistent response and effective area after quadrant healing.

Naturally, even the nominal energy threshold is currently set at 5 keV, each channel exhibits a different effective response, resulting in a significant spread (still not quantified) in the global

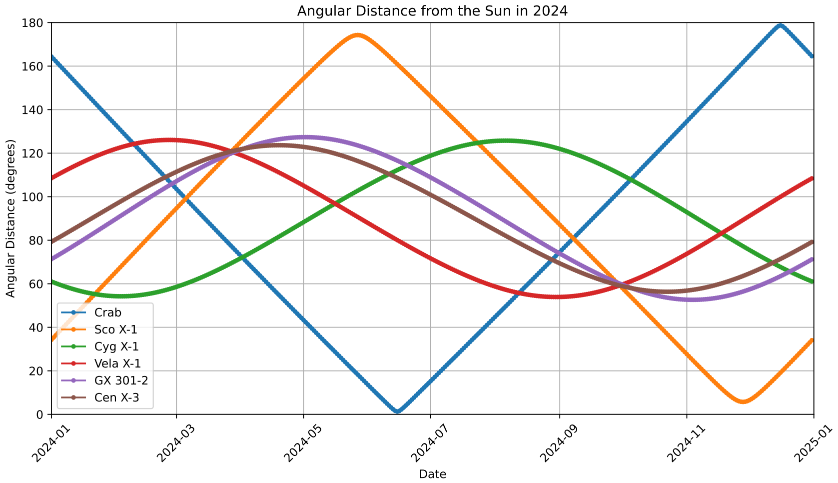

Figure 7.8: Angular separation from the Sun during 2024 for selected bright X-Ray sources (Sco X-1, Crab, Cyg X-1, Vela X-1, GX 301–2, Cen X-3). The ordinate is elongation (degrees) and the abscissa is calendar date; minima indicate intervals of poor visibility due to Sun constraints.

threshold. As a result, counts are still recorded in the nominal 10–100 keV band from two sources: (i) the instrument x (the SDDs), whose intrinsic energy range extends up to $\sim$ 30 keV and thus populates 10–30 keV; and (ii) a residual component from instrument s, i.e. scintillation light produced in the crystals and read out by the SDDs, whenever the corresponding SDD pulse exceeds the 5 keV trigger threshold (which occurs for incident photons with $E \gtrsim$ 90 keV). This point will be further clarified in the next section.

In most light curves, a monotonically decreasing trend is observed. This may indicate an increase in the angular separation between the satellite zenith and a particularly bright source such as Sco X-1.

Since the pointing is generally kept radial during these observations, as the satellite moves along its orbit a source will gradually leave the payload field of view, or at least drift off-axis.

Moreover, due to occasional issues with the orbital platform, the satellite may enter a free–spinning state, causing the Earth's limb to re-enter the field of view and leading to count variations across multiple energy bands as the 2024-03-27 in Figure 7.7.

7.3.4 Sco-X1

Sco X-1 is the brightest persistent extrasolar X-Ray source and the first to be discovered. It is a low-mass X-Ray binary hosting a neutron star and belongs to the "Z-source" class, identified by its characteristic track in X-Ray color–color diagrams (Giacconi et al. 1962a; Hasinger et al. 1989). A trigonometric VLBI parallax places it at a distance of ~2.8 kpc (Bradshaw et al. 1999), and it exhibits transient, mildly relativistic radio jets/lobes on AU scales (Fomalont et al. 2001).

During June, Sco X-1 lies at (or near) its maximum solar elongation and is therefore easily observable (see Figure 7.8). In all June observations, as anticipated, the count rate shows a clear decreasing trend. Hereafter, I hypothesize that this trend is driven by the increase, during each observation, of the off-axis angle between the source and the satellite pointing.

To validate this hypothesis, I adopted the spectral parameters of Sco X-1 reported by Mazzola

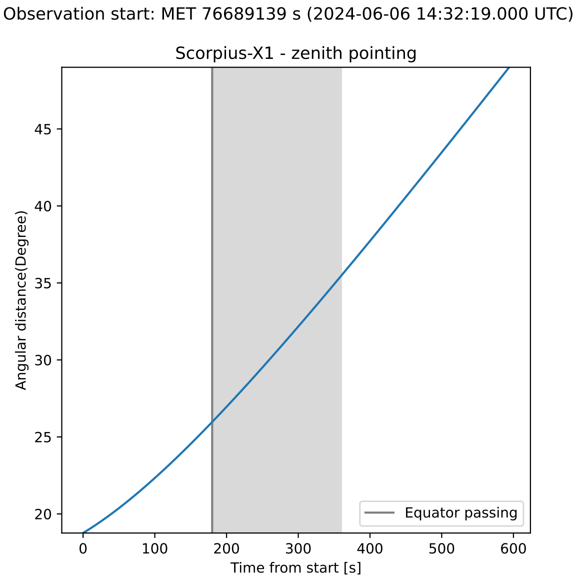

Figure 7.9: Angular distance between Sco-X1 and payload pointing as a function of Time. In gray is highlighted the equator passing area and the tumbling phase of the satellite.

et al. (2021) and generated HERMES-predicted spectra with XSPEC using `fakeit` [1]. I assumed the absorbed thermal–plus–Compton model `tbabs*(diskbb+comptt)` and folded it through the HERMES response (RMF/ARF), matching the observation's exposure.

Since I do not have precise estimates of the HERMES background, I used the `cutoffpl` model and adopted the Gruber parameters from Gruber et al. (1999a):

$$F_{CXB}(E) \; = 7.877\, E^{-1.29} \exp\left(-\frac{E}{41.13\ \mathrm{keV}}\right) \ \mathrm{ph\ keV^{-1}\ cm^{-2}\ sec^{-1}\ sr^{-1}}. \tag{7.4}$$

To obtain the correct normalization, the spectral intensity (given per steradian) in Equation 7.4 must be integrated over the instrument's effective field of view so that the steradian dependence is removed.

These steps provide an estimate of both the background contribution and that of Sco X-1 as if observed on-axis. A further limitation arises from the lack of precise knowledge of the pointing during the observations: although it is nominally kept zenithal, it is affected by errors due to possible satellite tumbling. To estimate the angular separation between the source and the pointing direction, assuming zenithal pointing, I use the SGP4 orbital propagator (Vallado et al. 2006) to compute the SpIRIT position as a function of time. The zenithal pointing assumption is valid for a non-tumbling spacecraft. When tumbling occurs, a gradual pointing drift leads to measurable deviations in the reconstructed light curves, including features such as the "bump" visible in Figure 7.10.

From the spacecraft position, the pointing follows directly: for zenithal (radial) pointing, it coincides with the right ascension and declination of the geocentric radial direction at that epoch (i.e., the RA/Dec of the spacecraft position vector).

It is therefore possible to estimate the time–dependent angular separation between the pointing

---

[1] see `https://heasarc.gsfc.nasa.gov/docs/software/lheasoft/download.html`

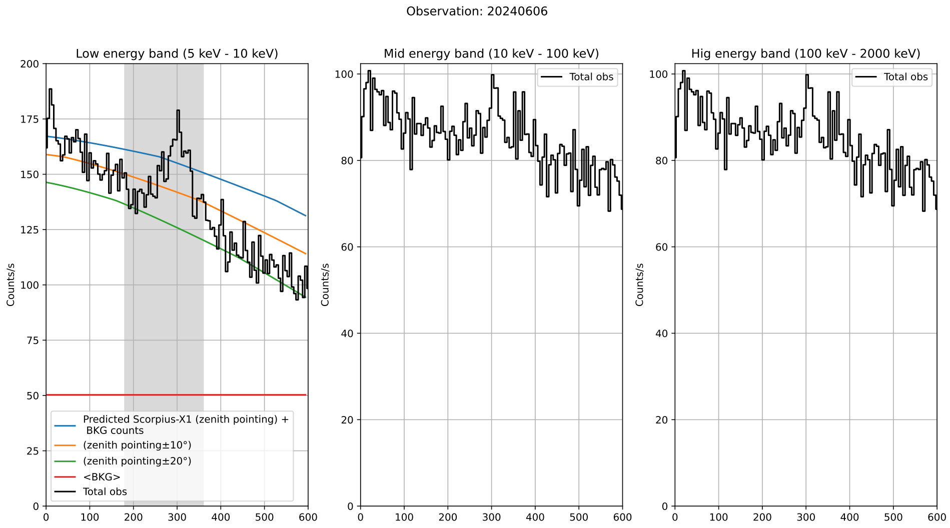

Figure 7.10: Filtered and rebinned SpIRIT/HERMES SRA light curves from the 2024–06–06 observation in three energy bands, *Low* (5–10 keV), *Mid* (10–100 keV), and *High* (100–2000 keV). In the Low band, the measured rate is overplotted with simple Sco X–1 predictions that include the time–dependent off–axis angle in the condition of: nominal zenith pointing (source+background) and pointing offsets of ±10° and ±20°. The Mid and High panels show the total observed rate only. Time is referenced to the start of the SRA window.

direction and the source (see Figure 7.9), $\theta(t)$, and to rescale the on–axis source counts by convolving the Sco-X1 simulated spectra with the different off-axis angle effective area in Figure 4.2, rescaled by active channels factor.

The example in Figure 7.10 illustrates the method described in this section. The measured count rates are consistent with a boresight offset of roughly 10°–20° from zenith pointing. Moreover, the smooth, monotonic decline is well reproduced once the time–dependent effective area $A_{\text{eff}}(\theta)$ is applied, i.e., when the off–axis angle $\theta(t)$ derived from the orbital propagator is folded into the model.

The measured count rates are lower than the on–axis expectation; the actual line of sight must be farther from the source than the assumed zenithal pointing; in other words, the pointing must be offset away from the source.

However, this agreement indicates that the propagator reliably reconstructs the spacecraft position/attitude and that the observed rate is chiefly modulated by the expected variation with off–axis angle.

It is also possible to discern a $\sim$ 100 s bump in the count rate at $\approx$ 250 s from the start of the observation. After the spacecraft's equator crossing, as indicated in the Figure 7.9, the attitude appears to have tumbled due to interactions between the onboard magneto-torques and the changing geomagnetic field, causing a transient pointing excursion—likely toward bright source or sources (i.e., the Galactic Center or Sco-X1). This would naturally produce the observed increase in counts, consistent with some bright high-energy emission.

Concluding, the proposed technique is nevertheless effective in explaining all the June observations in Figure 7.14, in which Sco X–1 falls within the HERMES field of view. The

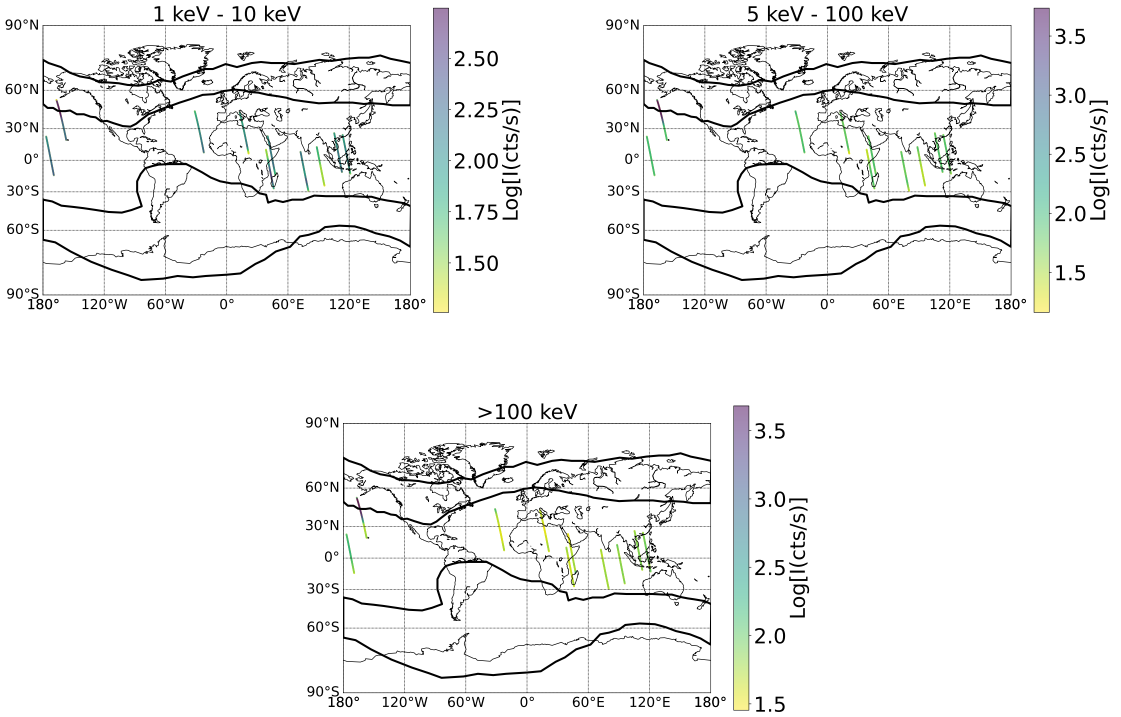

Figure 7.11: Global background count–rate maps in three energy bands: low (< 10 keV), medium (10–100 keV), and high (> 100 keV). Colors indicate the logarithm of the rate (cts s^{-1}) on a geographic latitude–longitude grid. Most variations reflect spacecraft pointing jitter, while black contours mark radiation critical regions with elevated charged particle flux.

discrepancy between the count rates expected for zenith pointing and those actually measured indicates that a real pointing offset is present in these observations. Unfortunately, platform limitations prevent us from determining this offset with high accuracy.

7.3.5 Background Map

Given the 10 acquired observations in SRA mode, I produced the background map in Figure 7.11, by considering the orbital position of the satellite.

During transits through high–radiation zones, most notably the South Atlantic Anomaly (SAA) and the polar belts (PB), the instrument is normally switched off. That said, the extent and boundaries of these regions drift with time, so static masks can be imprecise; occasionally the payload stays on for a few seconds to minutes inside the SAA/PB, adding spurious fluctuations to the recorded count rates as shown Figure 7.11.

7.4.0 CAL data

In CAL mode the data are genuine event by event lists: each photon is time–tagged with nominal resolutions of 320 ns in X–mode and 216 ns in S–mode, and the average energy resolution is ~ 15% at 30 keV.

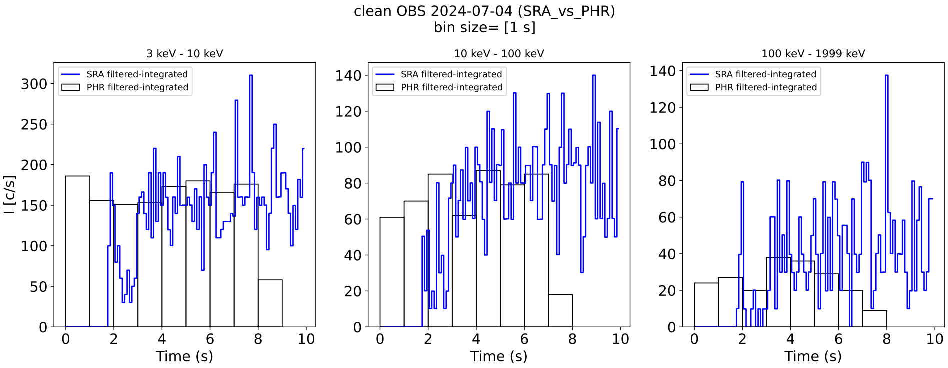

Figure 7.12: Comparison between SRA (rate–meter) and CAL/PHR (photon–by–photon, integrated) light curves for the 2024–07–04 SpIRIT/HERMES observation, shown in three energy bands (3–10 keV, 10–100keV, and 100–2000keV). CAL data are rebinned to 1s, while ratemeters are kept to their original resolution (0.1s).

I have at my disposal CAL data from 11 observations: 4 with durations of ~ 10 s, 3 of ~ 40 s, and the remaining ones lasting several hundreds of seconds.

7.4.1 2024.07.04 observation

The only observation that includes both SRA and CAL data was obtained on 2024-07-04. This dataset enables a direct cross-check of the SRA acquisition and processing by comparison with the photon-by-photon CAL events.

The close agreement between SRA and PHR products in Figure 7.12 validates the SRA acquisition chain and the developed filter, for this dataset; residual differences are compatible with counting statistics and minor dead–time/threshold effects.
However, a gap is visible in the *SRA* stream around t≈3s, and the first ~ 2s of the *SRA* time series contain no recorded counts. By contrast, the *CAL* (photon–by–photon) data are steady; this observation is instead truncated at ~ 8s due to an interruption in the acquisition pipeline.

7.4.2 Sco-X1 spectra with SpIRIT/HERMES

As part of the HERMES-PF commissioning and performance verification activities, I observed Sco X-1 on the 2024.07.04. Due to its brightness and variability, Sco X-1 is often used as a benchmark target for testing and calibrating X-Ray instruments. The detection and analysis of Sco X-1 with HERMES-PF provided a valuable opportunity to assess the instrument spectral response and overall sensitivity in a real astrophysical context.

I extract spectra from the July observation by producing level 2 data with the HERMESdas software (see subsection 4.1.3).

I subtract to the July observed spectra the background model as estimated in the previous section. As shown in Figure 7.13, Sco X-1 spectrum is dominated by the background above approximately 15 keV. I can fit the background subtracted spectra with a Compton model as expected for the source (Mazzola et al. 2021). The best fit parameter and associated errors are

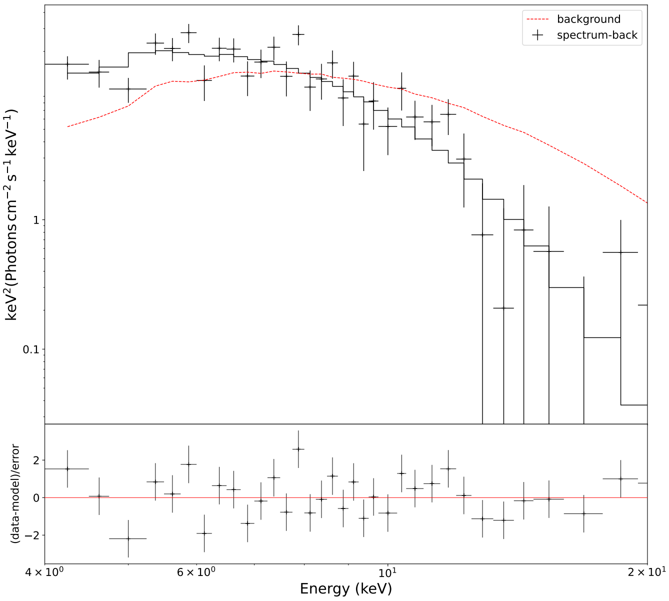

Figure 7.13: 2025.07.04 observation spectra reduced with the HERMESdas software. In black, the Sco-X1 background-subtracted spectra, while in red, the simulated background via fakeit.

compatible with the ones in the literature. I note that I only fix the seed photon temperature, and all the other parameters are estimated via fitting the observed spectra. I was also able to estimate the cut-off temperature $T_{cut-off}$ by fitting the spectra with a cut-off power law model, finding a $T_{cut-off}$ that is compatible with the literature (Mazzola et al. 2021). I finally derive the luminosity of the source in the 1 – 30 keV energy range, and I obtain a value of $L_{SpIRIT} = 2.6080 \cdot 10^{38}$ergs^{-1}.

Model	Component	Sco X-1
nthComp	Gamma	$2.15^{+0.63}_{-0.66}$
	kT$_e$(keV)	$2.22^{+2.1}_{-0.6}$
	kT$_{bb}$(keV)	0.33 (frozen)
	inp$_{type}$(0/1)	0 (frozen)
	Redshift	0 (frozen)
	norm	52^{+57}_{-37}
	χ^2/dof	40.4/31

To achieve an independent measure of the off-axis angle between source and S/C pointing and verify the correctness of the estimated luminosity, I retrieve the light curve between 2-20 keV from the MAXI database around July 4, 2024 (between MJD 60492 and 60408).

The flux extraction in the 2-20 keV band, estimates approximately $2.8 \cdot 10^{-7}$ erg s^{-1}/cm^{-2}, which, for a distance of 2.8 kpc, should correspond to a luminosity of $L_{MAXI} \sim 3.5 \cdot 10^{38}$ erg s^{-1}.

The correction $\frac{\mathrm{L_{SpIRIT}}}{\mathrm{L_{MAXI}}} \sim 0.75$ should be compatible with a source angle relative to the pointing of about 40 degrees, in line with S/C pointing.

Unfortunately, the 05/07/2024 is the only available event file associated with a Sco-X1 in FoV observation.

7.5.0 Summary of SpIRIT/HERMES Commissioning status as of 10 July 2024

At the time of writing (01 September 2025), the commissioning of the HERMES payload on board *SpIRIT* can be summarized as follows.

1. Data acquisitions completed so far include eleven photon–by–photon runs and eleven ~600 s observations for which only the standard rate meters were downlinked, owing to UHF telemetry limitations. I will discuss the remaining photon-by-photon observations in the next chapter.

2. A stable, continuous temperature regulation is not yet in place. All acquisitions have been performed with payload temperatures typically between −10°C and +10°C. As a consequence, higher payload temperatures may lead to an increased occurrence of retriggering spikes.

3. Spacecraft pointing/attitude information is not yet available.

4. From the datasets in hand I infer that:

    (a) the payload operates nominally;

    (b) housekeeping (voltages and currents) remains within expected ranges;

    (c) due to radiation damage and the relatively elevated operating temperatures, the detector noise is high, producing retrigger chains that contaminate the SRA files;

    (d) the average count rates are in line with pre–launch simulations of the in–orbit (equatorial) background.

Open items under investigation.

1. Further quantifying the origin of the observed light–curve modulations, the role of off–axis sources, and observed spike/holes in data.

2. Quantify the spread of the lower energy threshold in SDD.

3. Assessing possible anisotropic particle components (e.g., near–equatorial soft protons (Petrov et al. 2009).

Progress on these topics requires reasonably accurate attitude knowledge during the measurements and the availability of simultaneous housekeeping.

A robust in–orbit characterization of the payload will also require regular, stable access to photon–by–photon data. To this end, a 10 s observation with simultaneous SRA and PHR (photon–by–photon) downlink was executed on 04 July 2024. This dataset enables a direct comparison between the ratemeters and the underlying event stream.

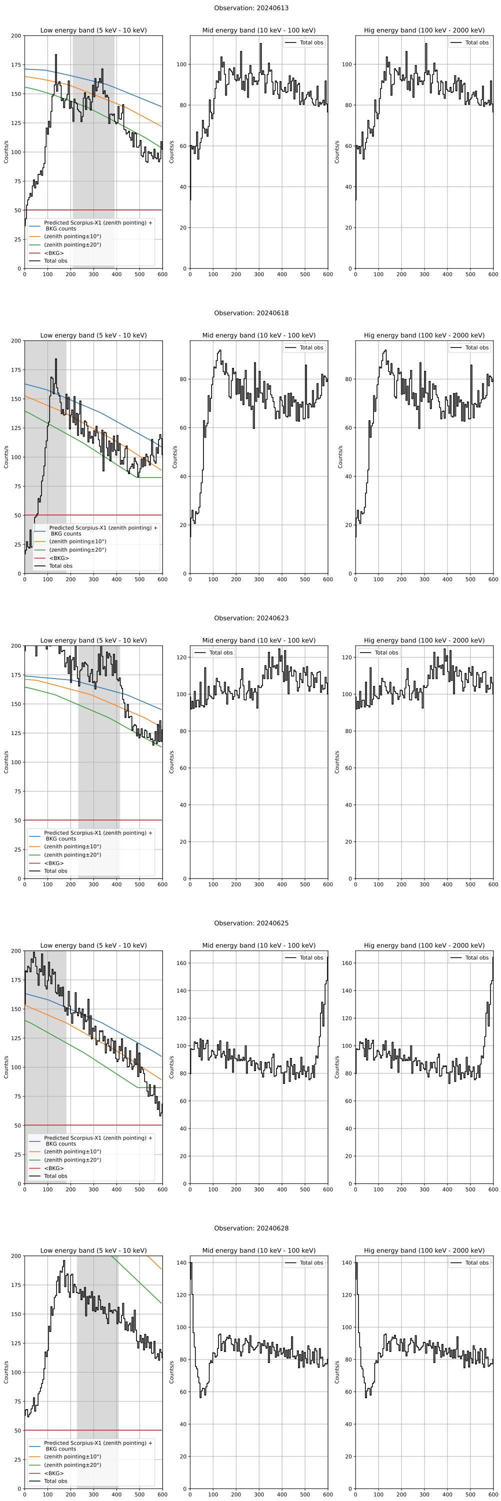

Figure 7.14: Filtered and rebinned SpIRIT/HERMES SRA light curves obtained by June observations.

8

CRAB absolute timing with SpIRIT/HERMES

Summary

This chapter summarizes the SpIRIT CubeSat demonstration of high–precision X–ray timing on the Crab pulsar using the *HERMES* hybrid SDD+GAGG detector. During a single on–axis exposure of $T_{\rm exp} = 730$ s, the 6U *SpIRIT* mission resolved the canonical double–peaked pulse profile with a $\sim 5\sigma$ significance in the 3–11.5 keV band, after adopting $n_{\rm bin} = 11$ phase bins and shifting bin edges to maximize the variance of fractional counts across phase. A key platform limitation was the absence of an operational GPS receiver during the observation, hence no definitive onboard orbit solution; the analysis addressed this by propagating the latest TLE with SGP4 and applying Solar–System barycentric corrections (via `barycorr`) using ICRS source coordinates, preserving absolute phase alignment to within a few microseconds over the ~12 min span. An *efsearch* around the Jodrell Bank ephemeris with up to 0.1 μs resolution yielded a best period $P_{\rm best} \simeq 33.8394$ ms, consistent with expectations and within the Fourier resolution limit $\Delta P \simeq P^2/T_{\rm exp} \approx 1.6\,\mu$s; the recovered profile is coherent, within uncertainties, with contemporaneous high–energy templates (e.g., *NICER*/*NuSTAR*) after amplitude rescaling and phase registration. A simple exposure–time model, validated through Monte Carlo simulations of Crab–like profiles, predicts that detectors with *HERMES*–class effective area and background achieve $\gtrsim 3\sigma$ pulsation detections in the majority of 700–800 s pointings, with significance peaking in the 3–11.5 keV interval where source SED and instrument response combine most favorably. These results show that microsecond–class timing, once the domain of flagship observatories, can be robustly executed from nanosatellites, opening practical avenues for time–domain high–energy astrophysics (e.g., pulsars, magnetars, short–timescale transients) and enabling rapid localization via inter–satellite triangulation in the future *HERMES*/*SpIRIT*–like constellations.

Authorship and context. The work presented in this chapter is submitted to the international journal APJL with the title "CubeSats Reach the Microsecond X-Ray Domain: Crab Pulsar Timing with SpIRIT/HERMES".

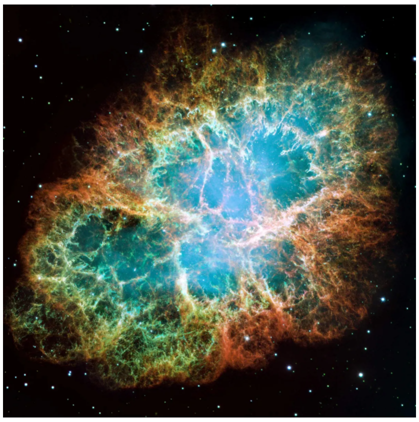

Figure 8.1: CRAB wind Nebula view as seen by the Hubble space telescope, credit to https://science.nasa.gov/

8.1.0 The Crab Pulsar

I present an on-orbit timing analysis with the *HERMES* detector aboard the *SpIRIT* nanosatellite. I resolve the Crab pulsar's X-Ray pulse profile and show that a low-cost, CubeSat-class platform can deliver timing performance previously associated with larger, dedicated observatories—an important step toward precision time-domain astrophysics with nanosatellites.

The Crab is a composite system consisting of a young pulsar (PSR B0531+21) and a surrounding synchrotron nebula. It is one of the most widely used calibration sources in high-energy astrophysics. Both the Pulsar and the Nebula emit across the entire electromagnetic spectrum. Owing to its high and stable flux, especially in the X-Ray and γ-ray bands, the Crab is considered a "standard candle" for instrument verification (Kirsch et al. 2005). Furthermore, the Crab pulsation, characterized by their well-known double-peaked profile, is exploited to evaluate the timing accuracy of any instrument (Rots et al. 2004).

It should be noticed that an independent CubeSat mission, the NinjaSat instrument carrying X-Ray gas detectors sensitive in the 2–50 keV band, has also carried on successful Crab pulsar timing observations (Tamagawa et al. 2025), highlighting the broader emergence of nanosatellites as capable instruments for time-domain astronomy. This significance was achieved with an exposure of 11 ks and thanks to a collimator that narrows the field of view (FoV) to $\approx 2.1^\circ$, thereby substantially reducing the non-pulsed contribution from the cosmic X-Ray background.

8.2.0 Crab absolute timing

Aiming to perform energy and timing calibration of the detector and demonstrating its timing capabilities, as a part of the commissioning activities the Crab Nebula was targeted in CAL mode. An extensive treatment of the chosen exposure time is reported in Chapter B.

The Crab pulsar light curve is folded into $n_{\text{bins}} = 15$ phase channels using the *SpIRIT* exposure of $T_{\text{exp}} = 730$ s performed on 2025-04-01 at 15:33:45 UTC. During the observation, spacecraft pointing, temperature, and state were stable. I restrict the analysis to *X-mode* (SDD) events

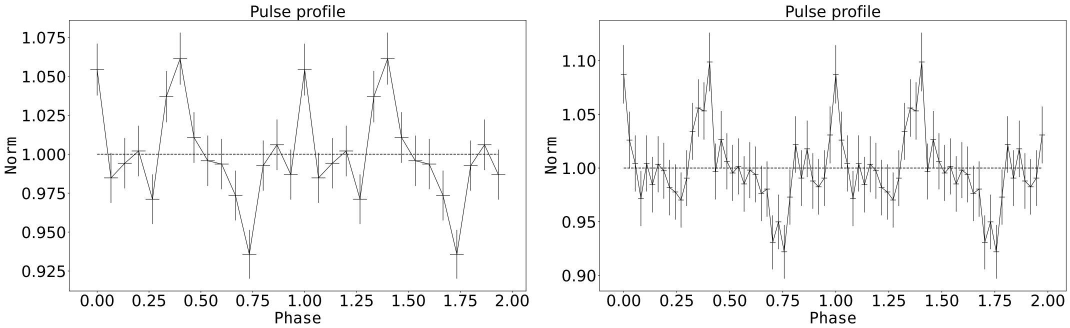

Figure 8.2: Left: SpIRIT/HERMES CRAB pulse profile in the 3–11.5 keV energy band at JB ephemerid with 15 folding channels (~5σ pulse profile significance). Right: SpIRIT/HERMES CRAB pulse profile in the 3–11.5 keV energy band at JB ephemerid with 37 folding channels (~3.3σ pulse profile significance).

only. Event energies are reconstructed with the standard gain/offset calibration by using the HERMESDAS reduction pipeline (Puccetti et al. 2024), and I retain photons in the 3–22 keV X-mode band. S-mode, from $\approx$ 30 keV to $\approx$ 2 MeV (Campana et al. 2020), photon statistics is not sufficient to produce a significant pulse profile. I note that the payload temperature was maintained under control during this observation.

8.2.1 Barycentric Correction and Ephemeris Folding

Pulse phases were computed by correcting the photon times of arrival with respect to the ephemeris reported in the Jodrell Bank (JB) monthly Crab database[1] after extrapolating the database solution closest to the observation epoch (2025-03-17, MJD 60751). Other solutions for previous or following epochs were also tested with similar results.

Since the GPS on board was not active, the spacecraft orbital file for the *SpIRIT* nanosatellite was not yet available at the time of the observation. The spacecraft trajectory was reconstructed by propagating the closest-in-time two-line-element (TLE) orbital ephemeris set with the SGP4 algorithm (Vallado et al. 2006). Due to the limited GPS time discipline of the HERMES payload, the spacecraft position reconstruction is affected by an uncertainty of approximately ±20 km. Photon arrival times were subsequently converted to the Solar System barycenter with the HEASOFT task `barycorr`, adopting the pulsar International Celestial Reference System (ICRS) coordinates and the propagated spacecraft ephemeris.

The Crab pulsar coordinates were taken from the Gaia DR3 catalog[2] (Gaia Collaboration et al. 2023). The Gaia DR3 coordinates were propagated from the reference epoch J2016.0 to the observation date t by applying the proper motions in right ascension and declination as:

$$\begin{aligned} \alpha(t) &= \alpha_0 + \mu_{\alpha*}\,(t - t_0), \\ \delta(t) &= \delta_0 + \mu_\delta\,(t - t_0) \end{aligned} \tag{8.1}$$

where $\mu_{\alpha*}$ and μ_δ are expressed in mas yr^{-1} and reported in the Gaia DR3 catalog.

---

[1] https://www.jb.man.ac.uk/pulsar/crab/all.gro

[2] https://gea.esac.esa.int/archive/

Starting from the JB solution for the P_0 and $\dot{P}$ quoted in the catalog at t_0 = MJD 60751, I propagate the spin period to the epoch of the *SpIRIT* observation t:

$$P(t) \; = \; P_0 \; + \; \dot{P}\,(t - t_0). \tag{8.2}$$

Hereafter, I adopt the resulting $P(t) \approx 33.84$ ms as the folding period.

8.2.2 Energy resolved profile

I select data in several energy bands, aiming to assess the impact of the noise contribution on the signal significance, by keeping constant folding parameters.

Consistent with the Crab pulsar spectral energy distribution peaking at 3–5 keV (Brandt et al. 2003), I found that the pulse-profile significance is maximized when a low-energy cut is set to 3 keV. As previously discussed, although the instruments effective area decreases between 3 and 5 keV and the nominal hardware threshold is ~5 keV, channel-to-channel non-uniform response yields a spread in the *effective* thresholds, leaving non-negligible acceptance down to ~3 keV.

The profile significance is evaluated by testing the null hypotheses that a flat signal is in the data (i.e., the mean of the profile) and evaluating the χ^2:

$$\chi^2 = \sum_{i=1}^{N} \frac{(I_i - \langle I \rangle)^2}{\sigma_i^2}, \tag{8.3}$$

where I_i is the intensity of the profile and σ_i the associated uncertainties.

The 3–11.5 keV energy range was identified as the one showing the most significant pulse profile; the obtained profile is shown in Figure 8.2, and the significance is $\sim 5\sigma$ (i.e., $\chi^2 = 57.5$ with dof = 14). This is in agreement with the distribution of the expected χ^2 value from the Monte Carlo simulations in subsection B.1.2. The phase separation between the two peaks is compatible with the typical Crab profile ($\Delta\phi = 0.4$). The apparent dip at $\phi \approx 0.75$ depends on the chosen number of phase bins and is consistent with the canonical off-pulse phase of the Crab profile (as seen with Jain & Paul (2011) or in Figure B.3). With the current choice of $n_{\text{bins}} = 15$, the off-pulse interval spans two adjacent bins, which enhances the contrast and thus the statistical significance of the dip. Using a different binning (e.g., increasing n_{bins}) redistributes the off-pulse counts and reduces the effect as shown in Figure 8.2. In this sense, the dip is primarily a binning-driven effect rather than an independent feature.

I note that the χ^2 of the pulse profile in Figure 8.2 attains its maximum when the initial phase is shifted by $\frac{6}{15}$ of the period. For other phase shifts, however, the χ^2 values deviate only within their expected statistical uncertainty, $\sigma_{\chi^2} = \sqrt{2\,\text{dof}}$. These results are nevertheless consistent with an uncertainty in the absolute phase arising from time-tagging inaccuracies, themselves a consequence of the absence of on-board GPS time discipline.

8.3.0 NuSTAR and NICER comparison

As a further check of the derived SpIRIT/HERMES pulse profile I use NICER and *NuSTAR* data to benchmark the timing performance of *SpIRIT* against both instruments and to validate the recovered pulse profile by adopting the NICER/*NuSTAR* profiles as reference templates.

8.3.1 NuSTAR and NICER data extraction

I extract data from the closest NuSTAR observation to the SpIRIT Crab Pulsar observation, which started on 2025-04-05 at 22:32:21.9340 UTC with an exposure of $T_{\mathrm{exp}} = 14.7$ ks. Data were retrieved from the NuSTAR target search website[3]. I considered filtered data and orbital file as provided by the archive and applied the NuSTAR CALDB clock correction to the event list.

I retrieved also NICER data from the NICER archive [4]. I consider the latest available NICER observation of the Crab pulsar, taken on 2025-02-28 18:55:18.00, with 1.4 ks exposure time, and downloaded cleaned data with the associated orbital file.

8.3.2 SpIRIT profile modeling

I derive the NuSTAR pulse profile I_{Nu} by considering the same number of channels as the SpIRIT profile in Figure 8.2. I analyze the SpIRIT profile by fitting it with the model profile:

$$I_{\mathrm{mp}} = A \times I_{\mathrm{Nu}}(x + \phi) + C, \tag{8.4}$$

where A is the scale factor of the model and C an additive constant. The two profiles are not in phase, so I consider a free parameter ϕ to apply a shift to the model in order to minimize the χ^2 that depends on the grade of similarity between the NuSTAR profile and the SpIRIT one. I_{mp} is the intensity of the model profile derived from the NuSTAR profile I_{Nu}. I apply 1000 positive phase shifts, from 0 to 1, to the I_{Nu} to model the SpIRIT profile. For each ϕ I evaluate the parameters A and C that minimize the χ^2, obtaining the best phase shift. The best-fit phase shift is $\phi = 0.667$, while the best-fit parameters are $A = 0.41 \pm 0.073$ and $C = 0.59 \pm 0.072$.

Considering the uncertainties associated with the SpIRIT observation and the best-fit parameters (A,C,ϕ), Figure 8.3 shows the reliability of the pulse profile. The reduced chi-square is $\chi^2_\nu = 1.33$ with $\nu = 15 - 2 = 13$ degrees of freedom (i.e., $\chi^2 = 17.29$), yielding a right-tail p-value of $p \approx 0.19$. Therefore, the fit is acceptable and would not be rejected at conventional significance levels (e.g., 10%, 5%, or 1%).

8.3.3 Epoch folding search

Further proof of the profile accuracy is obtained by performing an epoch folding search (`efsearch`) around the Jodrell Bank catalog expected Crab Pulsar ephemeris, as in Equation 8.2,

---

[3] https://heasarc.gsfc.nasa.gov/db-perl/NuSTAR_table

[4] https://heasarc.gsfc.nasa.gov/db-perl/NICER_table

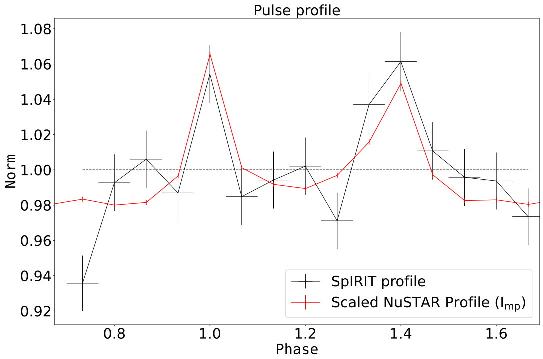

Figure 8.3: In black the Crab Pulse profile measured by *SpIRIT* (15 phase bins) with Poisson uncertainties, in the 3–11.5 keV band. The solid red curve is the best-fit NuSTAR modulated template, and associated uncertainties, $I_{\text{mp}}(\phi)$ from Equation 8.4, fitted *to* the *SpIRIT* and shown for best A, C, and ϕ parameters.

with a 1 μs resolution. This value is chosen given that this is the accuracy of the on-board time stamping. The efsearch best period is $P_{\text{best}} = 33.8394$ ms, which is 0.5 μs smaller with respect to the period derived from the radio ephemeris.

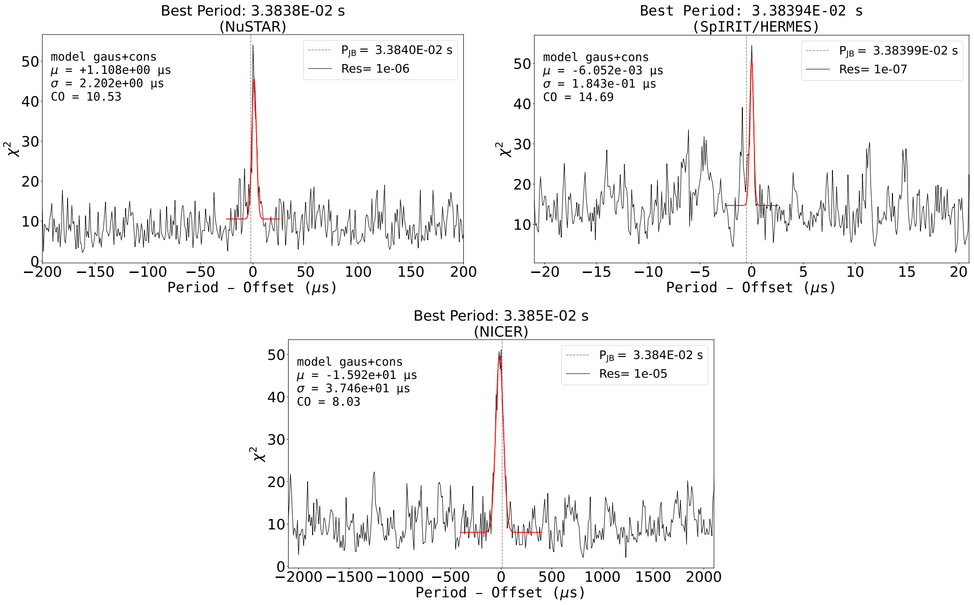

Figure 8.4: Epoch folding search around the expected JB catalog period at different search grid resolution, depending on the considered exposure time. I consider $n_{\text{bins}} = 15$, the 3–11.5 keV energy band and P, $\dot{P}$ of the epoch (MJD 60751). Panels show the Gaussian fit of the efsearch for each considered instrument.

The expected period is highlighted in Figure 8.4, whose difference from the best `efsearch` period is perfectly compatible with the first-order error on the period determination ΔP. By taking into account the natural width of a coherent periodogram of an epoch-folding peak imposed solely by the finite time baseline T_{exp}, from $\Delta\nu = 1/T_{\text{exp}}$, I derive $\Delta \text{P} = |\text{dP}/\text{d}\nu|\,\Delta\nu = \text{P}^2\Delta\nu$, and:

$$\Delta P = \frac{P^2}{T_{\text{exp}}} \approx 1.6\,\mu s. \tag{8.5}$$

However, by considering the Gaussian fit of the narrow efsearch peak, the SpIRIT best period is $\sim 2.5\sigma$ compatible with the JB ephemeris at the observation epoch (see Figure 8.4).

To compare the timing capabilities of *SpIRIT* with *NuSTAR* and *NICER*, I applied `efsearch` to *NuSTAR* and *NICER* data in the same energy band considered for the efsearch in Figure 8.4. I restrict the exposure so that the resulting detection significance matches that obtained with *SpIRIT* (see Figure 8.4) and take into account the same Crab ephemerids used for the *SpIRIT* data. I therefore decrease the exposure time for *NuSTAR* and *NICER* data. Due to that, the efsearch step is:

$$t_{\text{step}} = \frac{P^2}{T_{\text{exp}}\, n_{\text{bins}}}, \tag{8.6}$$

which sets the period resolution shown in Figure 8.4. As expected, thanks to the larger effective area of the two instruments, they require a smaller exposure time to reach the HERMES detection significance (*NuSTAR* $\approx$ 95 s of exposure time, while *NICER* $\approx$ 9 s).

8.4.0 Discussion and Conclusion

I performed sub-microsecond X-ray timing with the wide field of view HERMES payload onboard the 6U ~12 kg nanosatellite SpIRIT. A single on-axis exposure lasting $T_{\text{exp}} = 730$ s resolved the Crab pulsar characteristic double-peaked pulse profile with a statistical significance of 5.0σ in the $3-11.5$ keV band, with the primary and secondary peaks separated by a phase interval of ≈ 0.4, as expected.

To validate the correct operation of the payload, I performed an analytical and a Monte Carlo analysis presented in subsection B.1.2. I normalize an RXTE CRAB Pulsar pulse profile to the count rate (i.e., counts s^{-1}) of the HERMES payload by considering the background and the Nebula/Pulsar contributions. In a first step, I derive an analytical estimate of the minimum exposure needed for significant pulsation detection (see section B.2). This shows that an exposure of 730 s is sufficient to achieve a pulse-profile significance of at least 3σ.

I then quantify the detection probability at fixed exposure (730 s) by Monte Carlo simulation under Poisson statistics. Specifically, I generate realizations of the HERMES normalized pulse profile and, for each, compute the χ^2 with respect to the mean profile (exposure held fixed). The simulations show that, for a 730 s observation, the probability that χ^2 exceeds the threshold $\chi^2_{\text{th}} = 35.25$ (the 3σ criterion for a 15-bin profile) is $\simeq 82\%$. This means that χ^2 measured for the Crab pulse profile reported in the paper lies within the non-central χ^2 distribution (see subsection B.1.2).

Both analyses neglect instrumental systematics and are consistent with the significance obtained at the available exposure and conditions, indicating that the instrument performed correctly and that the expected Poisson statistics governs the results

Folding the photon time of arrivals at the Jodrell Bank ephemeris extrapolated to the observation epoch reproduced the Crab Pulsar profile, demonstrating that the spacecraft clock, the SGP4 orbit propagation, and the `barycorr` correction preserve the absolute phase to within a few microseconds over the 12-minute observation window. Indeed, an independent epoch-folding search recovered a best period of 33.8394 ms, only 0.5 μs shorter than the expected value of the catalog and entirely consistent with the Fourier resolution limit of ~ 1.5 μs achievable with the available *SpIRIT* exposure time. A template fit using the closest *NuSTAR* data to the *SpIRIT* observation confirms both pulse shape compatibility.

This result demonstrates that high-precision timing, previously the domain of large and resource-intensive observatories, can now be achieved by cost-effective small-scale space missions. A $\gtrsim 5\sigma$ detection in $\lesssim 15$ minutes of net exposure validates the performance of the compact *HERMES* payload and supports the feasibility of low-cost, short-development-cycle nanosatellite missions for high-energy astrophysics.

9

Quantum-Gravity Phenomenology

Summary

This chapter presents a technique for constraining an nth-order violation of Lorentz invariance by measuring energy-dependent arrival-time delays in a homogeneous sample of 25 gamma-ray bursts. All events were observed simultaneously by the *Fermi* Gamma-ray Burst Monitor and the *Fermi* Large Area Telescope and have reliable redshifts. The sample includes GRBs in the $0.34 < z < 4.35$ redshift range. For each burst, I compare the light curves of photons in several high-energy ($E > 1$ MeV) sub-bands with those in a fixed 10–100 keV reference band. Cross-correlation of the light curves yields 123 rest-frame lag measurements covering the 8–2200 MeV energy range. Each lag is interpreted as the sum of an intrinsic emission term and a propagation term produced by quantum-gravity dispersion; the statistical analysis provides a method to disentangle the intrinsic component from the quantum-gravity delays. This allowed us to put robust constraints on the quantum gravity energy, only assuming that the GRB mechanism is redshift-invariant. However, this assumption is not contradicted by the sample and is a far milder assumption than those adopted in single bright-GRB analyses.

Considering the first-order correction term in units of energy with respect to the Quantum Gravity energy, I obtain the following constraints: $E_{\mathrm{QG}} > 1.25 \times 10^{16}$ GeV (sub-luminal), $E_{\mathrm{QG}} > 2.49 \times 10^{16}$ GeV (super-luminal) at 2σ (95% confidence).

Although some recent constraints were closer to the Planck energy scale, the robustness of the method enables its application to increasingly large GRB samples. The limit improves roughly in proportion to the square root of the sample size, allowing sensitivity to approach the Planck scale as the statistics grow. With future datasets, such as those expected from THESEUS, which are anticipated to provide approximately 500 GRBs per year, I foresee at least an order-of-magnitude improvement within a few years of operation. The inclusion of additional bright, high-z bursts observed in a broad MeV–GeV energy range will further enhance this trend and tighten the resulting constraints until reaching the Planck Energy scale. *Authorship and context.* The work presented in this chapter is under revision (after major comments) to the international journal A&A with the title *"Constraints on Lorentz invariance violation from energy-resolved timing of gamma-ray bursts: A technique independent of the Gamma Ray Burst emission model."*

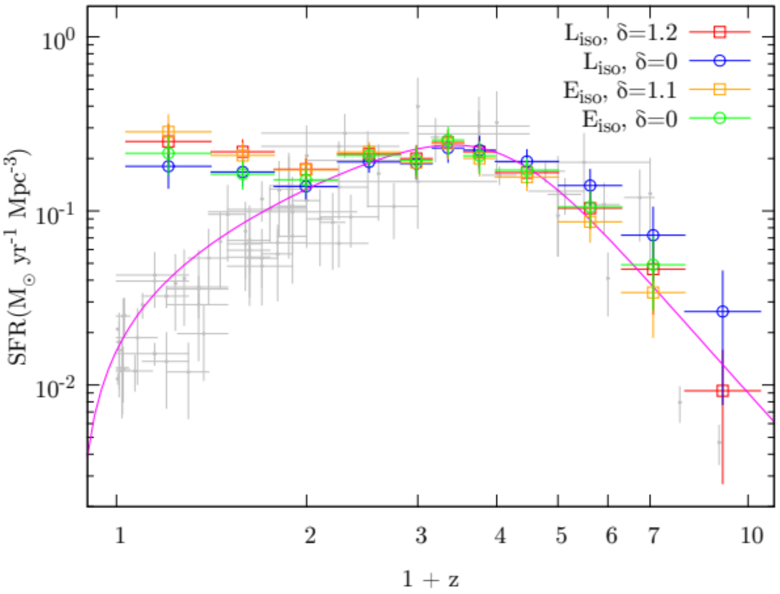

Figure 9.1: Cosmic GRB formation rate (GRBFR) from Tsvetkova et al. (2023) overlaid on literature star–formation–rate (SFR) measurements. Gray symbols: SFR compilations from Hopkins (2004), Bouwens et al. (2011), Hanish et al. (2006), and Thompson et al. (2006). The solid curve shows the SFR parametrization of Li (2008). For display, all four GRBFR determinations share the same normalization, and the GRBFR points are shifted to align with the SFR around $(1 - z) \sim 3.5$. Adapted from Fig. 5b of Tsvetkova et al. (2021). © AAS; reproduced with permission.

9.1.0 Novel methodologies for quantum gravity

The techniques presented in chapter 5 of this thesis have broad applications in the field of quantum gravity. As introduced in section 2.9, GRB prompt emission provides an ideal probe for studying quantum-gravity phenomenology, given its high-energy output from keV to TeV and the occurrence of GRBs across virtually the full range of cosmological redshifts, up to $z \approx 10$.

Building on the methodology of Tsvetkova et al. (2023) and Burderi et al. (2021), this paper presents the analysis conducted on a sample of 25 GRBs of known redshift observed by the Fermi satellite. I model the observed time lags in the sample as the sum of intrinsic and LIV-induced delays to disentangle the two effects. This may be possible since intrinsic delays should not depend on distance, while QG delays should be correlated with it.

As shown in Figure 9.1, the GRB occurrence rate remains roughly constant out to $z \sim 6$. This is a key advantage for LIV searches, since any propagation–induced time delay scales with cosmological distance and is therefore amplified at higher redshift.

In the following, I restrict the analysis to the more pronounced first-order term (i.e., I will consider $n = 1$), for the reasons outlined in section 2.9.

9.2.0 Data reduction

Aiming to explore possible LIV effects, I analyze a sample of 25 GRBs with known redshifts, $0.36 \lesssim z \lesssim 4.35$, detected by the Fermi Gamma-Ray Space Telescope. All the considered GRBs are listed in Table 9.3. For each GRB, I retrieve the Fermi Gamma-ray Burst Monitor (Meegan et al. 2009) (GBM) data from the Burst Catalog [1]. I consider the brightest among the twelve thallium-activated sodium iodide (NaI(Tl)) detectors and download the Time-Tagged Events (TTE) data.

I apply an energy slice to the GBM data by selecting the 10–100 keV energy band photons (I refer to it as HE_{keV}).

To include very high-energy photon data (E>1 MeV, hereafter HE_{MeV}), I take into account the Fermi Large Area Telescope (LAT) (Atwood et al. 2009) and its Data Query interface [2]. The coordinates and observation time windows are obtained from the NASA Gamma-ray Coordinates Network (GCN) circulars [3]. A circular mask is considered to increase the signal-to-noise ratio (S/N) to select photons within a small radius centered on the GRB position.

I consider as input a region of interest with a 60-degree radius around source coordinates and an energy range from 30 MeV to 10 GeV. I also include in the analysis LAT lower-energy (LLE) data from the Fermi LLE Events Catalog [4]. I download the LLE data of the corresponding GRBs file from the trigger catalog. In that case, I also apply a circular mask to increase the S/N.

The LLE and LAT T_{90} and $T_{90,start}$ are retrieved from the GBM burst table and used during the temporal analysis.

The energy binning for both LAT and LLE datasets ensures that photons are selected within energy intervals that satisfy the condition:

$$\nu = \frac{\Delta \mathrm{E}}{\bar{\mathrm{E}}} \lesssim 0.25, \tag{9.1}$$

where $\bar{E} = \sqrt{E_{\max} E_{\min}}$ is the average energy (geometric mean) of the selected energy band ΔE and $E_{\max}$ and $E_{\min}$ are, respectively, the upper and lower boundaries of the sub-band from which the $\mathrm{HE}_{\mathrm{MeV}}$ photons are extracted.

The criteria ensure that the spread of photon energies within an energy bin ΔE is sufficiently narrow such that potential QG-induced time delays remain negligible within the bin. The value of 0.25 was the lowest threshold that allowed us to keep the energy window as narrow as possible around the reference energy while still retaining sufficient statistics for the temporal analysis. I eventually consider broader bins (i.e., exceeding the 0.25 criterion) in the cases where photon statistics is low (see Table 9.3).

I note that, for each GRB, including more than one NaI detector dataset does not affect the precision of the estimated lag, given the poor statistics of Fermi/LAT data. Considering the bursts discussed, e.g., in Leone et al. (2025b), the typical lag precision achievable for bright

---

[1] heasarc.fermi.gsfc.nasa.gov.GBM

[2] heasarc.fermi.gsfc.nasa.gov.LAT

[3] `https://gcn.nasa.gov/circulars`

[4] heasarc.fermi.gsfc.nasa.gov.LLE

GRBs with GBM detectors is of order $\sigma_{\tau,\mathrm{GBM}} \sim 10^{-2}$ s. This improves as the square root of the number of contributing detectors N_{det} when two or more GBM units observe the burst under comparable statistics. By contrast, the typical lag precision for *Fermi*/LAT is $\sigma_{\tau,\mathrm{LAT}} \gtrsim 0.11$ s. For the purposes of this work, to first approximation, the total timing precision is given by the quadrature sum of these independent contributions:

$$\sigma_{\tau,\mathrm{tot}} \simeq \sqrt{\left(\frac{\sigma_{\tau,\mathrm{GBM}}}{\sqrt{N_{\mathrm{det}}}}\right)^2 + \sigma^2_{\tau,\mathrm{LAT}}},$$

therefore, for GBM–LAT lag measurements, the optimal choice is to use a single GBM detector: since $\sigma_{\tau,\mathrm{LAT}} \gtrsim 10 \times \sigma_{\tau,\mathrm{GBM}}$, the total lag error is dominated by the LAT term and gains negligibly from $N_{\mathrm{det}} > 1$. By contrast, for GBM–GBM lags the precision improves as $N_{\mathrm{det}}^{-1/2}$.

This choice is also motivated by computational efficiency: processing a larger GBM photon list would needlessly slow the CCF computation without changing the outcome.

9.3.0 The method

The core idea is to model the total observed delay, $\tau_{\mathrm{total,obs}}(E_{\mathrm{obs}}, z)$, measured between a fixed low-energy reference band (10–100 keV) and the observed energy $\mathrm{E_{obs}}$ above 1 MeV, as the sum of two separate contributions.

First one is the intrinsic lag due to the emission mechanism of the sources $\tau_{int,obs}(E_{obs}, z)$.

The effect of cosmological expansion affects both the energy and the observed duration of the intrinsic delay. Due to cosmological redshift, the energy in the GRB reference frame as E_{rf} is related to the E_{obs} through:

$$\mathrm{E_{rf}} = \mathrm{E_{obs}}\,(1 + \mathrm{z}). \tag{9.2}$$

Moreover, because of the cosmological time dilation, the GRB reference frame intrinsic delay is related to the observed one by:

$$\tau_{int,obs}(E_{obs}; z) \equiv (1+z)\tau_{int}\Big(E_{obs}(1+z); z\Big) = \tau(E_{rf})(1+z). \tag{9.3}$$

The second term is the observed propagation delay induced by LIV effects $\tau_{QG,obs}(E_{obs}, z)$ (Tsvetkova et al. 2023). Following Jacob & Piran (2008) this can be expressed as a function of the E_{rf} and redshift:

$$\begin{aligned}\tau_{QG}(E_{rf}, z) = \left(\frac{1}{H_0}\right)\left(\frac{1}{\xi}\right)\left(\frac{E_{rf}}{\zeta E_{pl}}\right)^n \left(\frac{1+n}{2}\right)\left(\frac{1}{1+z}\right)^n \\ \times \int_0^z \frac{(1+z')^n\,dz'}{\sqrt{\Omega_m(1+z')^3 + \Omega_\Lambda}}.\end{aligned} \tag{9.4}$$

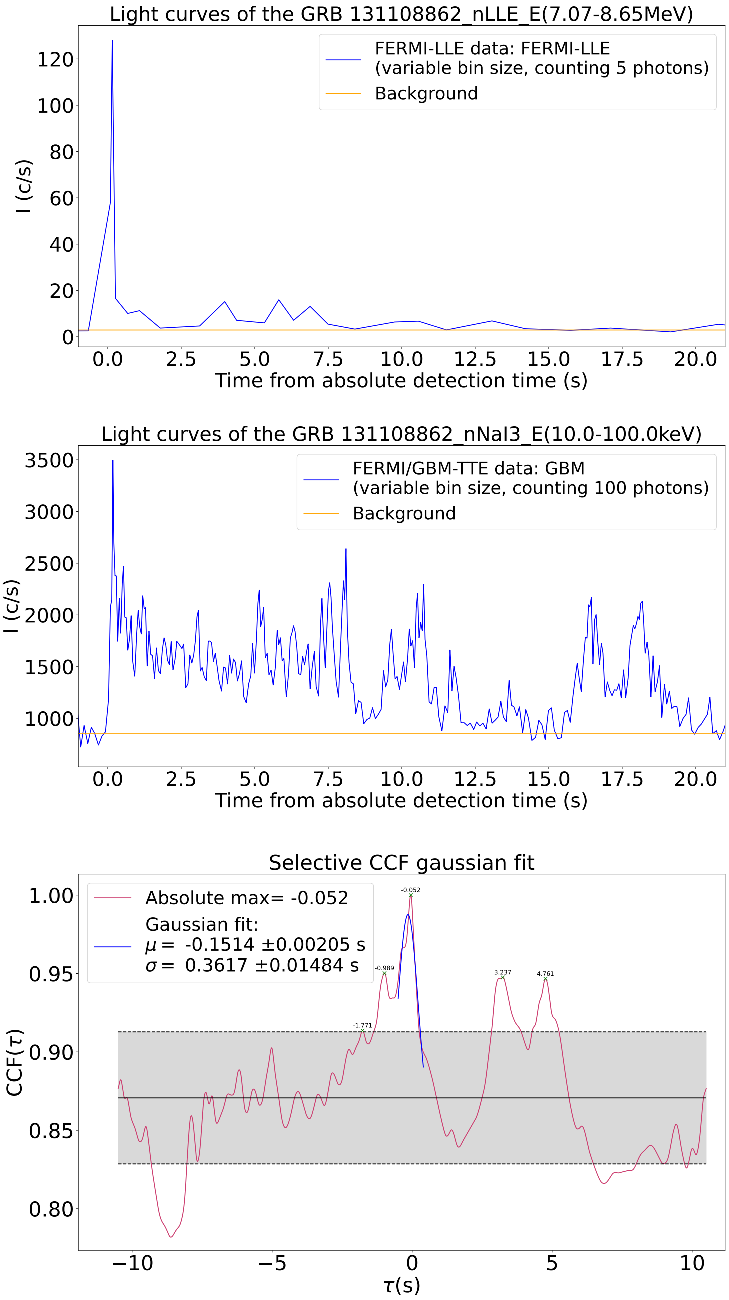

Figure 9.2: Top and central panel represent respectively HE_{keV} and HE_{MeV} light curves of GRB131108A. I report the Fermi/LLE energy band in MeV and the Fermi/GBM (NaI) energy band in keV. The bottom plot shows the Cross-Correlation Function between the two light curves above. A positive/negative delay means that HE_{MeV} photons lag/anticipate with respect to HE_{keV}photons.

The total delay is therefore:

$$\tau_{total}(E_{rf}, z) = \tau_{int}(E_{rf}) + \tau_{QG}(E_{rf}, z). \tag{9.5}$$

Since

$$\lim_{E_{obs}\to 0} E_{rf} = 0, \tag{9.6}$$

and

$$\lim_{(E_{rf}/\zeta E_{pl})\to 0} \tau_{QG}(E_{rf}, z) = 0, \tag{9.7}$$

it is possible to compute the QG delays in the MeV band by cross-correlating the light curve in the MeV band with the light curve in the keV band. I note that, according to Equation 9.7, QG–induced delays in the keV band are negligible compared to those in the MeV band.

9.3.1 CCF lags

The light curves of these HE$_{MeV}$ bands are then cross-correlated with the HE$_{keV}$ counterpart using the Cross-Correlation Function (CCF) method described in Leone et al. (2025b), to estimate the total lag as in Figure 9.2. The considered light curves, although they may differ in morphology and duration because of energy difference, dependent responses, and bandpass effects, are assumed to trace the same transient. Under this hypothesis, the measured CCF lag reflects an interband delay rather than unrelated variability.

The novel methods proposed in Leone et al. (2025b) are equivalent to standard CCF techniques as in Maraventano et al. (2025), but are pivotal in the case of low statistics (i.e., low counts, low signal-to-noise ratio).

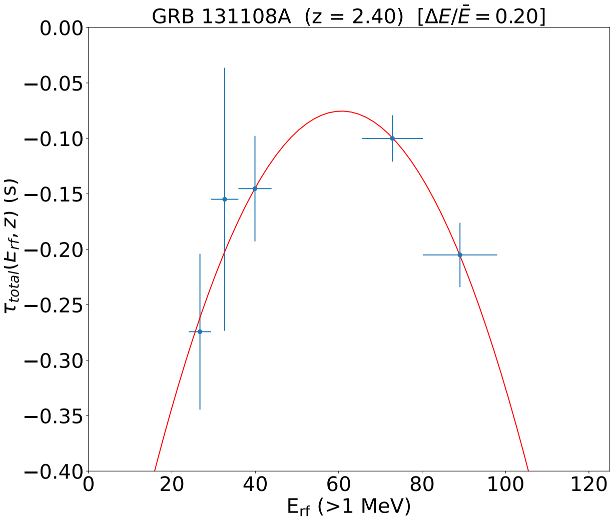

Figure 9.3: Total delays $\tau_{total}(E_{rf}, z)$ estimates as a function of energy E_{rf} at redshift z=2.40 for the GRB 131108A. An F-test indicates that upgrading from a constant to a quadratic model yields a chance-improvement probability of $p = 0.016$ (98.4% significance; $\sim 2.4\sigma$). Relative to a linear fit, the quadratic model gives $p = 7.9 \times 10^{-3}$ (99.2%; $\sim 2.7\sigma$), again favoring a quadratic model.

I evaluate the lag of HE$_{MeV}$ photons with respect to the HE$_{keV}$ photons as depicted in Figure 9.2.

The delay and associated uncertainty are estimated using the Monte-Carlo analysis (MPD method) of Leone et al. (2025b). The initial fit parameters are chosen once and kept fixed across all Monte Carlo realizations.

In particular, I fit the primary CCF peak with a Gaussian profile only if it satisfies the significance requirement

$$y_{\rm max} \;>\; y_{\rm sec} + 1\sigma, \tag{9.8}$$

where $y_{\rm max}$ is the height of the highest peak, $y_{\rm sec}$ is the height of the next-highest peak, and σ is the standard deviation of the CCF values over lag (e.g., the scatter of the CCF about its mean). If the condition is not fulfilled, the detected peaks are treated as a single, broader structure. In that case, I fit one Gaussian whose initial width is set equal to the separation between the leftmost and rightmost peaks of the CCF.

When a single, statistically significant peak is identified and thus interpreted as the time delay between the HE_{MeV} and HE_{keV} photon streams, I verify its consistency with the light–curve morphology by time–shifting one curve by the best–fit lag and confirming that the main structures align (see Figure 9.2 for an example).

The example in Figure 9.3 shows the delay estimates across the selected energy intervals for GRB 131108A. The lag trend follows a quadratic profile, which suggests the presence of a structured jet in the GRB, as proposed in Vyas et al. (2024). I note that the delays shown in Figure 9.3 are marginally consistent with those reported by Maraventano et al. (2025) (i.e., $\tau \approx -0.11$), albeit with the opposite sign as a consequence of the authors' sign convention.

I estimate delays as a function of the E_{rf} for the whole GRB sample. All the estimated total delays $\tau_{total}(E_{rf}, z)$ are shown in Figure 9.4 and include 123 rest-frame lag measurements covering the 8–2200 MeV energy range.

9.3.2 Disentangling Quantum Gravity delays from GRBs emission mechanism

To perform the analysis described in Section 3.3 of Tsvetkova et al. (2023), I divide the E_{rf} axis into adjacent energy channels such that:

(i) each energy channel of width $\Delta E = E_{max} - E_{min}$ is sufficiently wide to include data of several (≥3) GRBs;

(ii) it is sufficiently narrow to include light curves of almost the same energy, i.e., to have $\frac{\Delta E}{\bar{E}_{ch}} = \frac{\Delta E}{\sqrt{E_{max}\,E_{min}}} \leq 0.33$.

These energy channels and the number of GRBs within each channel are shown as a grey rectangle in Figure 9.4.

For each energy channel (that contains delays belonging to different GRBs and therefore of different redshift z), I have from Equation 9.4 (adopting n=1):

$$\tau_{QG} = f(\bar{E}_{ch})\, u(z), \tag{9.9}$$

where

$$f(\bar{E}_{ch}) = \left(\frac{1}{H_0}\right)\left(\frac{1}{\xi}\right)\left(\frac{\bar{E}_{ch}}{\zeta E_{pl}}\right), \tag{9.10}$$

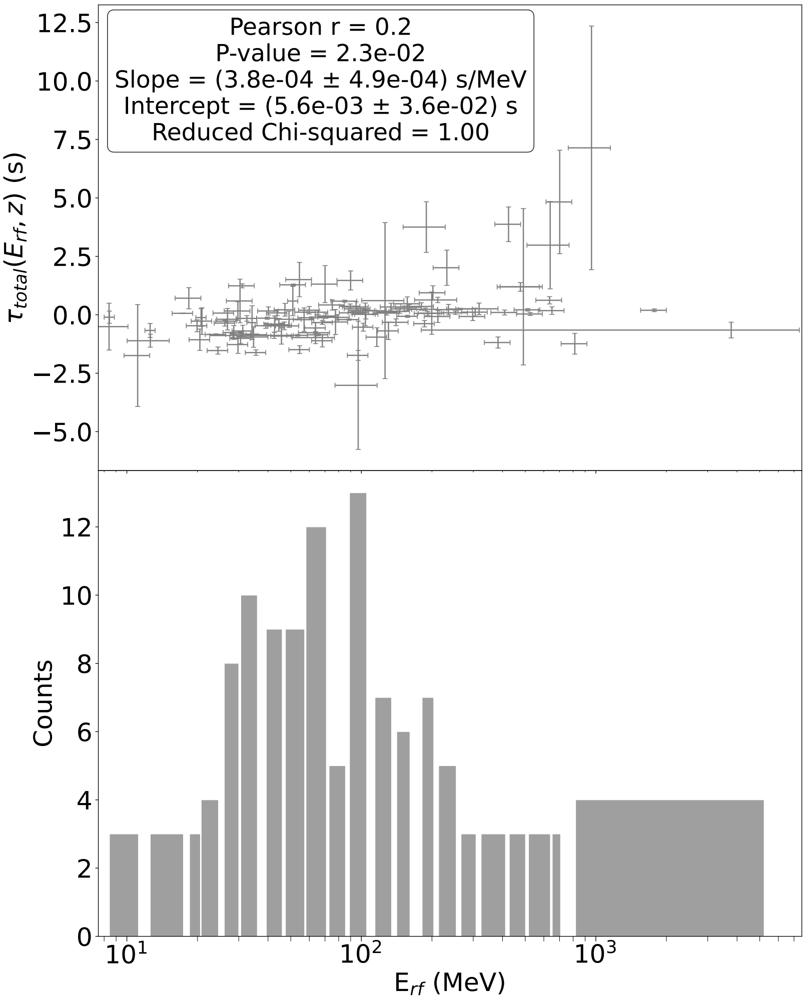

Figure 9.4: Upper Panel: Rest-frame total lags $\tau_{total}(E_{rf}, z)$ for the GRB sample as a function of energy. I report the linear model's best-fit parameters, achieved by considering the post-fit sigma and the correlation parameters. Lower Panel: Histogram of the E_{ref} at which delay data are evaluated.

and

$$u(z) = \left(\frac{1}{1+z}\right) \int_0^z \frac{(1+z')\,dz'}{\sqrt{\Omega_m (1+z')^3 + \Omega_\Lambda}} \tag{9.11}$$

is computed for each redshift adopting a flat ΛCDM cosmology ($\Omega_M = 0.301 \pm 0.045$ and $\Omega_\Lambda = 0.699 \pm 0.0402$). Redshift and cosmological parameter uncertainties are propagated into $u(z)$.

With these definitions Equation 9.5 reads

$$\tau_{total}(\bar{E}_{ch}, z) = \tau_{int}(\bar{E}_{ch}) + f(\bar{E}_{ch})\,u(z). \tag{9.12}$$

In the hypothesis that the central engine is the same for all, it is reasonable to assume that the

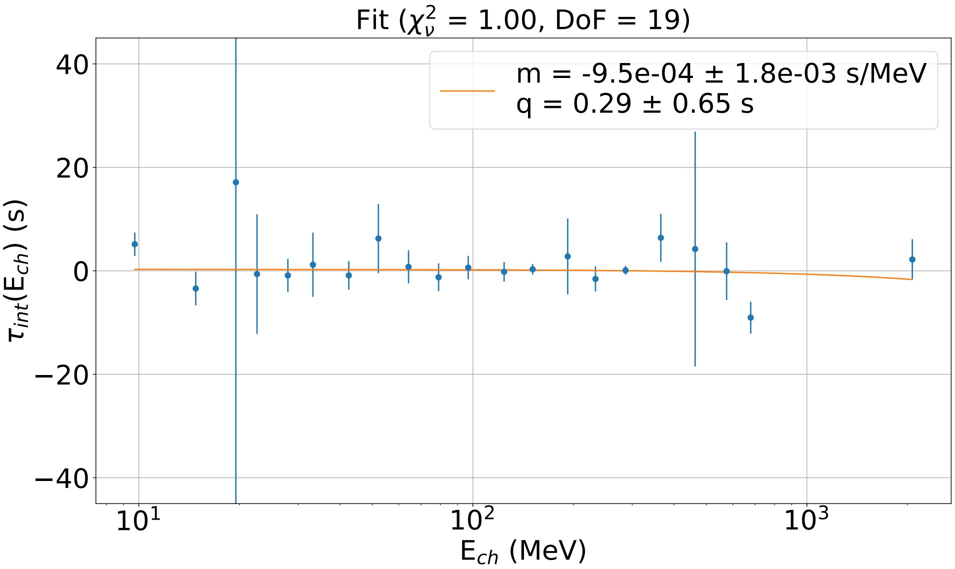

Figure 9.5: Best fit of the $\tau_{int}(E_{ch})$ values from Figure 9.7, taking into account post-fit uncertainties.

intrinsic delays are the same for all the GRBs in a given energy channel $\bar{E}_{ch}$.

Therefore I fitted the delays withing an energy channel $\bar{E}_{ch}$ with a linear function of $u(z)$, obtaining a value for $\tau_{int}(\bar{E}_{ch})$ and $f(\bar{E}_{ch})$ for each energy channel. Fit results are shown in Figure 9.7.

The central engine assumption is not contradicted by Figure 9.5, where the intrinsic delays, $\tau_{int}(\bar{E}_{ch})$, fitted with a linear model, yield slope and intercept consistent with zero. Thus, no statistically significant energy dependence is detected, indicating that keV–MeV emission mechanisms do not introduce systematic inter-band time offsets. These findings are consonant with Maraventano et al. (2025), which, for lags between LLE (30–100 MeV) and GBM (10–100 keV), reports a near-symmetric sign distribution, about 40% positive and 37% negative, consistent with no statistically significant energy-dependent offsets in the keV–MeV correlation.

Since I have successfully disentangled the two contributions, I focus only on the $f(\bar{E}_{ch})$ term, which is proportional to the energy of the channels $\bar{E}_{ch}$. The constant of proportionality contains the Quantum Gravity strength, in determining LIV delays, in the sub-luminal and super-luminal cases $\alpha = \pm\zeta^{-1}$ respectively:

$$f(\bar{E}_{ch}) = \alpha \, \frac{\bar{E}_{ch}}{H_0 \, E_{pl}}. \tag{9.13}$$

Therefore, it is possible to fit a direct proportionality relation to the $f(\bar{E}_{ch})$ obtained from the 21 energy channels, allowing for the computation of the LIV Quantum Gravity strength α and its associated error. This fit is shown in Figure 9.6.

The estimated value of α for n=1 is:

$$\alpha_{\text{fit}} = (2.05 \pm 3.08) \times 10^2. \tag{9.14}$$

with an associated error $\sigma_\alpha = 3.08 \times 10^2$ (68% c.l.).

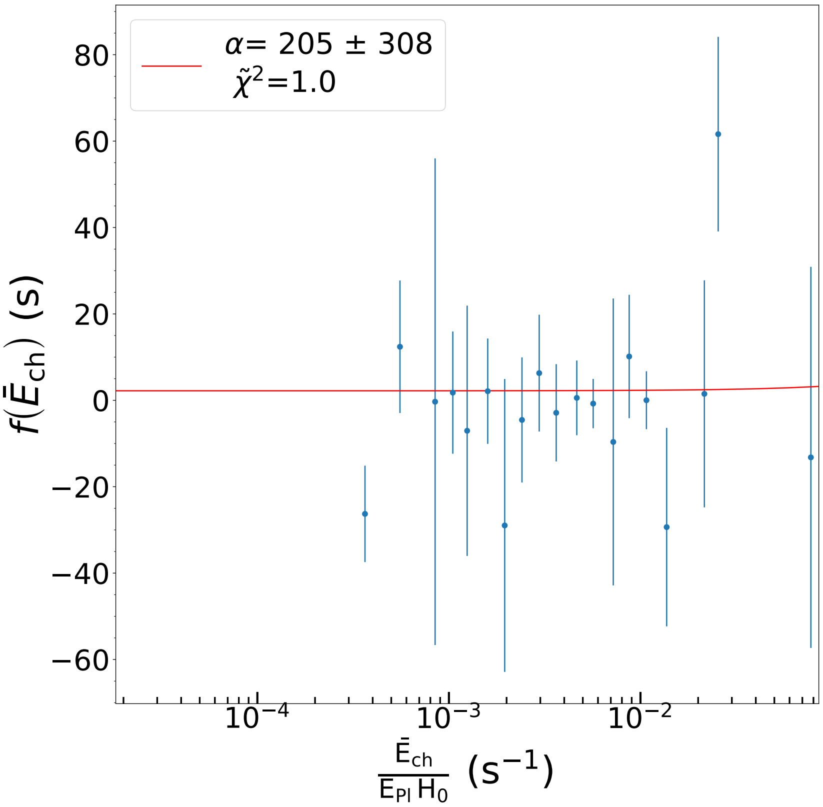

Figure 9.6: Best fit of the slope values from Figure 9.7, taking into account post-fit uncertainties.

9.4.0 Discussion

In this paper, I introduce a model-independent framework to search for LIV using energy-resolved light curves of GRBs with known redshift. The method exploits the possibility that LIV induces tiny, energy-dependent photon delays during cosmological propagation, while simultaneously accounting for intrinsic energy-dependent lags associated with GRB emission processes. I analyze a sample of 25 Fermi GRBs with known redshifts, extracting keV–MeV-GeV light curves from GBM, LLE, and LAT data and computing inter-band delays using a robust cross-correlation framework optimized for low-count statistics. The measured total lag is modeled as the sum of an intrinsic component—corrected by cosmological redshift and time dilation—and a potential LIV propagation term. By grouping measurements in narrow rest-frame energy channels and fitting the lag–redshift relation within each channel, the method cleanly separates intrinsic delays from propagation-induced ones, enabling a direct estimation of the LIV term without assumptions about the prompt-emission physics.

Although the current sample size is limited and photon statistics above 100 MeV remain modest, the method already yields meaningful constraints on first-order LIV. The absence of any detectable energy-dependent propagation delay allows us to place robust lower limits on the quantum gravity energy scale.

Interpreted within the LIV framework, the fact that I measure no residual energy-dependent lag—after subtracting the source-intrinsic contribution—implies that the best-fit slope (α) in Figure 9.6 is statistically consistent with zero. In other words, any LIV-induced propagation delay must be smaller than the sensitivity of the analysis. This non-detection directly translates into a lower limit on the QG energy scale. Since, in the linear LIV scenario, the slope (α) is inversely proportional to ($\mathrm{E_{QG}}$) (i.e., ($\zeta = \pm\alpha^{-1}$) for the subluminal and superluminal cases, respectively), the absence of a measurable trend means that (E_{QG}) must be larger than the value that would have produced a detectable non-zero slope in Figure 9.6. This interpretation is consistent with the timing accuracy shown in Figure D.1. At GeV energies—the regime in which LIV effects are expected to be largest—the attainable delay precision is at best on the order of sub-seconds. According to the first-order estimates in Equation 2.17, such timing uncertainties are insufficient to reveal a LIV signature at the current instrumental sensitivity, naturally explaining the absence of a significant detection in the dataset.

Adopting the 95% confidence level (2σ) around the measured value, I compute these constraints in the two Quantum Gravity regimes, sub-luminal ($\alpha > 0, \xi = 1$) and super-luminal ($\alpha < 0, \xi = -1$).

In the sub-luminal regime, high-energy photons lag low-energy photons. The slope in Figure 9.6 must be positive, therefore
$0 < \alpha < \alpha_{max} = \alpha_{\mathrm{fit}} + 2\sigma_\alpha = 821$ (95% C.L). Since $\mathrm{E_{QG}} = \mathrm{E_{Pl}}/\alpha$, the lower limit can be expressed as $\mathrm{E_{QG}} > \mathrm{E_{Pl}}/\alpha_{\mathrm{max}}$. Therefore $\mathrm{E_{QG}} > E_{Pl}/821 = 1.2 \times 10^{-3} E_{Pl} = 1.5 \times 10^{16}$ GeV.

In the super-luminal regime, high-energy photons precede low-energy photons. Therefore, the slope must be $\alpha_{\mathrm{min}} = \alpha - 2\sigma_\alpha = -411 < \alpha < 0$ (95% C.L.). In this case the lower limit can be expressed as $\mathrm{E_{QG}} > E_{Pl}/|\alpha_{min}| = 2.4 \times 10^{-3} E_{Pl} = 3 \times 10^{16}$ GeV.

Note that the method allows for disentangling source–intrinsic emission lags from propagation lags induced by general theories in first-order QG expansion, without requiring assumptions on the fireball-emission model. The only physical assumption in the analysis is that the GRB explosion mechanism remains unchanged with redshift (this is reasonable since the physical law that rules the fireball model should be independent of the age of the Universe). This is not the case in several papers discussing LIV delays with GRBs. For instance, Piran & Ofengeim (2024) analyzed GRB 221009A under the assumption that photons of different energies were emitted simultaneously, so that any observed delay could be attributed to propagation effects rather than to intrinsic emission processes. Similarly, Cao et al. (2024) assumed that intrinsic delays were negligible when analyzing the same event. Again, these kinds of assumptions were present in earlier works, suchas the hypothesis that a single 31 GeV photon was emitted at the same time as a peak of emission detected in the keV band Abdo et al. (2009). In summary, although previous studies have reported constraints that are closer to the Planck energy with respect to the constraint reported here, their reliance on specific modeling of the fireball emission implies that such constraints could be systematically biased.

Moreover, the analysis seems to indicate that there are no significant delays between the emission of photons in different energy bands for the whole GRBs sample (see the flat trend in Figure 9.5). This result does not rule out the possibility that a significant trend may be present in the delays of any individual GRB (e.g., the quadratic trend in Figure 9.3). Indeed, the method offers a robust statistical treatment to isolate the intrinsic delays due to the GRB emission mechanisms across a large sample of GRBs.

The result is consistent with previous studies based on the analysis of a population of GRBs. For instance, Bernardini et al. (2017) obtained a comparable lower limit by using a sample of 15 short GRBs, again under the assumption that short bursts exhibit negligible intrinsic spectral lags. I note explicitly that the model-agnostic approach allows, in principle, to use short and long GRBs in the same sample.

On the other hand, the analysis yields a constraint approximately five times tighter than that reported by Pan et al. (2024). Although the dataset is smaller, this is expected since the sample spans higher redshifts, up to $z = 4.35$ versus $z \leq 1.99$ in Pan et al. (2024) . More in detail, in their approach, the source–intrinsic delay term is explicitly modeled as a function of the energy and subtracted from the measured lags to isolate a putative quantum-gravity contribution. By contrast, I find no statistically significant energy dependence of intrinsic delays between the keV and GeV bands (see Figure 9.5), in agreement with Maraventano et al. (2025).

A more stringent constraint (by a factor of ten) is reported by Ellis et al. (2019), which analyzed a sample of 8 bright GRBs of known redshift detected by Fermi/LAT. The authors adopted a sophisticated statistical approach that combines multiple estimators (e.g., skewness and kurtosis). However, they make an a priori assumption about the shape of the prompt emission, i.e., a burst-like Gaussian intensity profile.

9.4.1 Future Perspective

Years	N_{GRBs}	E_{QG} [GeV] β=-1.5	E_{QG} [GeV] β=-2.6	E_{QG} [GeV] β=-3.5
1	500	$> 1.1 \times 10^{17}$	$> 7.9 \times 10^{16}$	$> 4.3 \times 10^{16}$
2	1000	$> 1.9 \times 10^{17}$	$> 1.1 \times 10^{17}$	$> 5.9 \times 10^{16}$
4	2000	$> 2.8 \times 10^{17}$	$> 1.5 \times 10^{17}$	$> 7.4 \times 10^{16}$
8	4000	$> 4.4 \times 10^{17}$	$> 2.6 \times 10^{17}$	$> 9.0 \times 10^{16}$

Table 9.1: First order E_{QG} constraint evaluated as a function of the nominal operation year of the THESEUS mission in sub-luminal regime. The $\mathrm{N_{GRBs}}$ column reports the number of generated GRBs with redshift association and $1-10$ MeV spectral coverage, while β is the high-energy spectral index, adopted to simulate the GRBs spectra.

The method disentangles intrinsic from propagation-induced lags without relying on a specific prompt-emission model. It can be applied directly to a bigger GRBs sample as obtained by next-generation missions such as *THESEUS*.

I generate the synthetic GRB population, detectable by the future THESEUS mission Bozzo et al. (2024), with a high-energy spectrum described by a single high-energy spectral index β in Equation D1. Moreover, I considered a Quantum Gravity effective energy equal to the Planck Energy (i.e., $\alpha = 1$). The detailed procedure is described in section D.1.

Table 9.1 summarizes the expected subluminal LIV constraints as a function of the mission operating years and the high-energy spectral index β. Specifically, the two extreme values of

β correspond to the lower and upper bounds of the high-energy spectral index reported for the Fermi/LLE data in Maraventano et al. (2025), while the central value is the arithmetic mean of all listed indices. Using the technique on the sample of GRBs detected by THESEUS in eight years, I expect to be able to confirm the constraint reported in Ellis et al. (2019), without making any assumption on the GRB emission model or shape.

I note that, even if the number of GRBs with measured redshift were halved with respect to the forecast adopted here, the techniques described in chapter 5 would readily allow us to increase by a factor of two the number of $\tau_{\text{total}}(E_{\text{rf}}, z)$ measurements obtainable from each GRB, thereby compensating for the reduced sample size and preserving the overall statistical power of the analysis.

Coupling a GeV–sensitive instrument similar to Fermi/LAT, but with an effective area greater by one order of magnitude, with an instrument like *THESEUS*, capable of measuring the redshfit of almost all the observed GRBs, would rapidly raise the lower limit on the QG energy, up to a fraction of the Planck Energy or to eventually detect the long sought fingerprint of the quantum nature of the space time. This is obtained by taking into account that the precision improves as the square root of the GRB number and the square root of the effective area. Since the spectra of GRBs rapidly fade above 1 GeV, I predict that the GeV band is the most promising band to search for these delays.

9.5.0 Conclusion

In the emerging field of quantum gravity, it is crucial to have a model-independent and robust technique that can confirm previous results, even when applied to a small sample. I analysed 25 *Fermi* GRBs with secure redshift measurements ($0.36 \lesssim z \lesssim 4.35$), measuring arrival-time lags between high-energy sub-bands ($E > 1$ MeV) and the 10–100 keV reference band. The total lag $\tau_{\text{total}}(E_{\text{rf}}, z)$ was modelled as the sum of a source–intrinsic term and a first-order LIV propagation term. The cosmology-aware, model-agnostic framework separates these contributions without assuming an emission template; the only physical assumption is that there is no evolution with redshift of the fireball emission mechanism, once I have corrected the energies for the redshift in order to recover the intrinsic value from the observed one. Moreover, I find no significant energy dependence of intrinsic delays, in the GRBs sample, in the keV–GeV band, in line with Maraventano et al. (2025). Applying this technique, I obtain lower limits $E_{\text{QG}} \gtrsim 10^{-3} E_{\text{Pl}}$ (2σ c.l.) in both sub- and super-luminal cases, in line with several previous multi-GRB studies.

Next-generation missions will both validate and improve these limits. In particular, *THESEUS*, with rapid repointing and prompt optical/NIR follow-up, will secure redshifts for all the GRBs in its field of view. With hundreds of events per year in the 1–10 MeV band, the constraint will improve in the near future. In the far future, coupling a “redshift-hunter” mission like *THESEUS* with a GeV-sensitive instrument would allow us to explore the quantum gravity nature of the space-time close to the Planck limit.

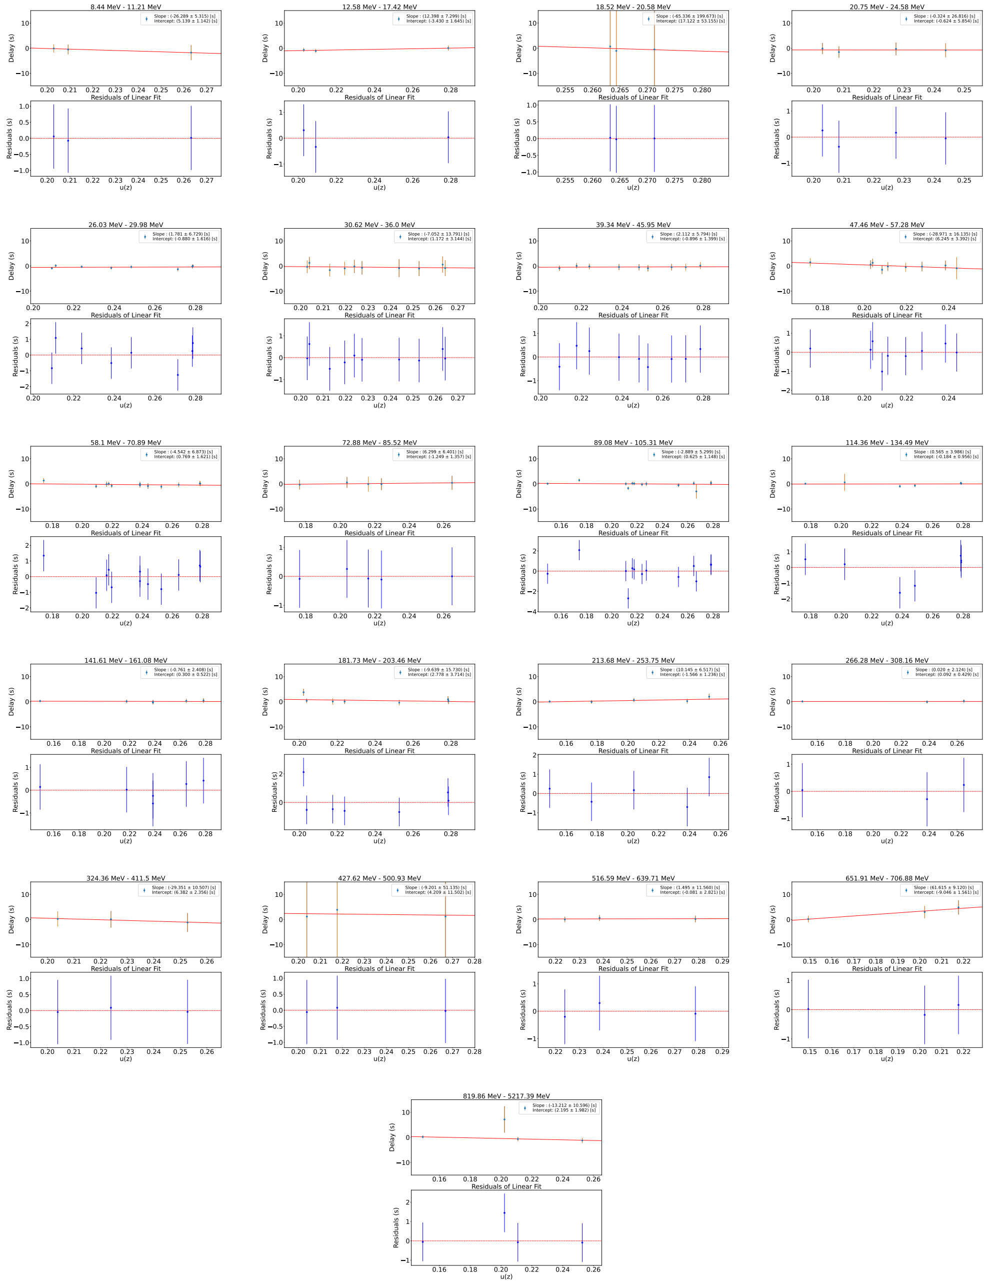

Figure 9.7: The linear trends obtained for the delays between the high-energy (HE_{keV}) and lower-energy (LE) light curves are shown for the individual energy bins.

Table 9.2: Analyzed GRBs with known redshift

GRB name	RA (°)	DEC (°)	Redshift	ν_{LLE}	ν_{LAT}
GRB 080916C	121.8	-61.3	4.35 ± 0.15	0.25	0.25
GRB 081024B	322.9	+21.2	2.56 ± 1.63	0.25	–
GRB 090323	190.69	+17.08	3.57	0.25	–
GRB 090328	0.87	-41.95	0.7354 ± 0.0003	0.2	–
GRB 090510	333.400	-26.767	0.903 ± 0.003	0.2	0.25
GRB 090902B	264.5	+26.5	1.822	0.2	0.25
GRB 090926A	353.56	-66.34	2.1062 ± 0.0004	0.2	0.25
GRB 100724B	124.16	+74.42	1.00	0.2	–
GRB 110721A	333.4	-39.0	1.94 ± 1.56	0.2	–
GRB 110731A	280.39	-28.53	2.83	0.25	0.4
GRB 130427A	173.139	+27.692	0.338 ± 0.002	0.1	0.4
GRB 130518A	355.809	+47.641	2.49	0.3	–
GRB 131108A	156.47	+9.90	2.4	0.2	0.25
GRB 131231A	10.585	-1.845	0.6439	0.25	–
GRB 140206B	315.26	-8.51	2.73	0.2	–
GRB 140619B	132.68	-9.66	2.67	0.25	–
GRB 141028A	322.70	-0.28	2.33	0.2	–
GRB 150403A	311.79	-62.76	1.91 ± 0.15	0.2	–
GRB 160509A	310.1	+76.0	1.17	0.2	0.25
GRB 170214A	256.33	-1.88	2.53	0.2	0.25
GRB 170405A	219.37	-25.23	3.510	0.25	–
GRB 180720B	0.58	-2.95	0.654	0.2	–
GRB 230812B	250.1	+46.2	0.3602 ± 0.0006	0.4	–
GRB 240825A	341.6	+5.9	0.659	–	0.25
GRB 241228B	122.2	+14.0	2.674	–	0.8

Complete list of GRBs with redshift, right ascension (RA), declination (DEC), and the parameter ν considered for the analysis with LLE or LAT data. When the ν is not reported, I do not consider the LLE or LAT GRB data.

Table 9.3: The f($\bar{\mathrm{E}}_{\mathrm{ch}}$) and $\tau_{int}(\bar{E}_{\mathrm{ch}})$ reported in Figure 9.7

$\bar{E}_{\mathrm{ch}}$ (MeV)	f($\bar{\mathrm{E}}_{\mathrm{ch}}$) (s)	$\tau_{int}(\bar{E}_{\mathrm{ch}})$ (s)
9.72	-26.3 ± 5.3	5.1 ± 1.1
14.8	12.4 ± 7.3	-3.4 ± 1.6
19.5	-65.3 ± 199.7	17.1 ± 53.1
22.6	0.3 ± 26.8	-0.6 ± 5.8
27.9	1.8 ± 6.7	-0.9 ± 1.6
33.2	-7.0 ± 13.8	1.2 ± 3.1
42.5	2.1 ± 6.9	-0.9 ± 1.4
52.1	-28.97 ± 16.1	6.2 ± 3.4
64.2	-4.54 ± 6.8	0.8 ± 1.6
78.9	6.3 ± 6.4	-1.2 ± 1.4
96.9	-2.9 ± 5.3	0.6 ± 1.1
124	0.6 ± 3.9	-0.2 ± 0.9
151	-0.8 ± 2.4	0.3 ± 0.5
192.3	-9.6 ± 15.7	2.8 ± 3.7
232.9	10.1 ± 6.5	-1.6 ± 1.2
286.4	0.02 ± 2.1	0.1 ± 0.4
365.3	-29.3 ± 10.5	6.4 ± 2.3
462.8	-9.2 ± 51.1	4.2 ± 11.5
574.9	1.5 ± 11.5	-0.1 ± 2.8
678.8	61.6 ± 9.1	-9.0 ± 1.6
2068.2	-13.2 ± 10.6	2.2 ± 2.0

Linear coefficient $f(\bar{E}_{ch})$ *of the delay slopes in Figure 9.7 and intrinsic delay* $\tau_{int}(\bar{E}_{\mathrm{ch}})$ *for the respective energy channel* E_{ch}*.*

10

Energy-Dependent Time-Lag Detection in ULXs

Summary

In this chapter, I present the temporal analysis of NGC 7456 ULX-1 with *XMM-Newton*. I build soft (0.3–1 keV) and hard (1–10 keV) light curves using fixed-count adaptive binning, and measure inter-band delays with the developed CCF techniques.

I measure, over two XMM observations (2018-05-19 and 2023-04-29), a hard lag of the order of 10^3s. Segmenting the data into 10-ks windows, I identify two significant lag episodes, -1987.5 ± 336.6 s and -982.4 ± 345.7 s. These are temporally coincident with intervals of rising flux, where the temporal flux derivative dF/dt becomes large and positive (steepening). This concurrence suggests that lag detection is variability-dependent, being strongest during phases of rapid brightening. The magnitude ($\sim 10^3$ s) and sign of the delay disfavor simple reverberation (which would imply nonphysical light-travel paths). This instead supports the inward propagation of accretion-rate fluctuations in a super-Eddington, geometrically thick flow: perturbations first emit soft X-Rays at larger radii and, after the measured delay, reach the inner corona, producing the hard emission.

Authorship and context. The work presented in this chapter was carried out by the candidate, Wladimiro Leone, and the related article is still in preparation.

10.1.0 NGC 7456 ULX-1

NGC 7456 is a spiral galaxy at a distance of ~15.7 Mpc (Tully et al. 2016). It is the host of a numerous X-Ray population, with at least 4 ULXs (Walton et al. 2011; Pintore et al. 2020) plus a likely ULX transient (ULX-5; Pintore et al. 2020). Two ULXs (ULX-1 and ULX-2) are persistent sources and the brightest of the galaxy, with average X-Ray luminosity well above 5×10^{39}erg s^{-1} (see section 3.3 for further details) while the other two (ULX-3 and ULX-4) are transient objects with flux variations of orders of magnitude. The source was observed three times in the past by XMM and SWIFT/XRT. In particular, during a 2018 XMM observation, the source ULX-1 was caught during an extreme short-term flux variability of one order of magnitude on timescales of minutes. This is one of the highest short-term variability ever observed in a ULX (fractional variability $> 40\%$ on timescales of 500s; Sutton et al. 2013; Middleton et al. 2015a; Pintore et al. 2020), however, not coupled with any significantly marked spectral variability. ULX-1 can be classified as a soft-ultraluminous ULX (Sutton et al. 2013) since its spectra peak around 1 keV. I extend the temporal analysis in chapter 5 by exploiting a never reported in the literature XMM observation taken in 2023, where the source is found at a higher flux but with less marked short-term temporal variability.

10.2.0 Data reduction

I analyzed all available observations of XMM [, 2], i.e., ObsID: 0303560701 (PI: D. Rosa-Gonzalez, XMM1 hereafter), 0824450401 (PI: G. Israel, XMM2 hereafter), and 0901330101 (PI: F. Pintore, XMM3 hereafter), which were taken in 2005, 2018 and 2023, respectively (Table 10.1).

Table 10.1: Observations used in this work.

Instr.	Obs.ID	Start	Stop	Exp.^{a}
		[YYYY-MM-DD hh:mm:ss (UTC)]		[ks]
XMM	0303560701	2005-05-06T20:05	2005-05-06T22:25	2.9
XMM	0824450401	2018-05-18T11:24	2018-05-19T12:10	63
XMM	0901330101	2023-04-28T23:50	2023-04-30T08:49	94

^{a} Net EPIC-pn exposure after background screening.

Data are reduced with SAS v21.0.0 and adopting recent calibrations. Data are selected from both EPIC-pn and EPIC-MOS using the task EPPROC and EMPROC. Events were filtered with standard criteria, i.e., FLAG=0 and PATTERN≤4 and PATTERN≤12 for pn and MOS, respectively. Good times are obtained by removing all the time intervals with a high background. In particular, lightcurves and spectra were extracted in the following time intervals: [231797254:231800156]s, [643029762:643092614]s, and [799126347:799220865]s from the XMM reference time, resulting in a net exposure time of 2.9 ks, 63 ks, and 94 ks for XMM1, XMM2, and XMM3, respectively.

ULX and background events were selected from circular regions of 30”and 45”, the latter being taken from a free-of-source region close to ULX-1 on the same CCD. The photon times of arrival (ToAs) were reported to the barycenter of the Solar System using the task BARYCEN.

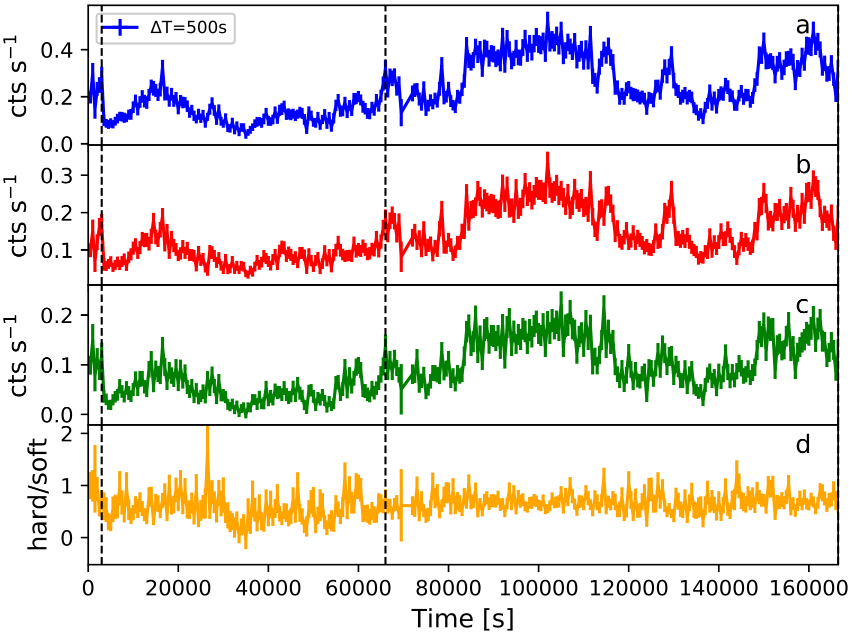

Figure 10.1: XMM light curves of ULX-1. The four panels display, respectively, the full band (0.3–10 keV; blue), the soft band (0.3–1 keV; red), the hard band (1–10 keV; green), and the corresponding hardness ratio (yellow). The time binning used in each panel is indicated in the labels. Dashed vertical lines mark data gaps between separate *XMM-Newton* pointings; the time axis (in seconds) concatenates the good-time intervals and therefore does not represent the actual wall–clock timeline of the observations.

10.3.0 Temporal Analysis

10.3.1 Long-term analysis

The long-term light-curve of the ULXs in NGC 7456 covers almost 20 years, although the coverage is very scanty and many gaps are present along the time.

ULX–1 remains persistently active, exhibiting ~5-fold variability in the 0.3–10 keV band, with various flux dips. The soft (0.3–1 keV) and hard (1–10 keV) bands vary in step, indicating co–co-evolution of the emitting components and corroborating the behavior seen in the XMM observations in Figure 10.1.

10.3.2 Short–term variability

ULX–1: All the XMM observations give strong indications that ULX-1 is highly variable on short timescales.

To quantify the source variability, I considered the average flux temporal derivative, $\Delta F/\Delta t$, estimated as the linear slope of the count–rate versus time in successive, non-overlapping 5 ks segments of data.

For each 5 ks segment, I bin ToA with a fixed bin-size of 250 s, which implies that each of

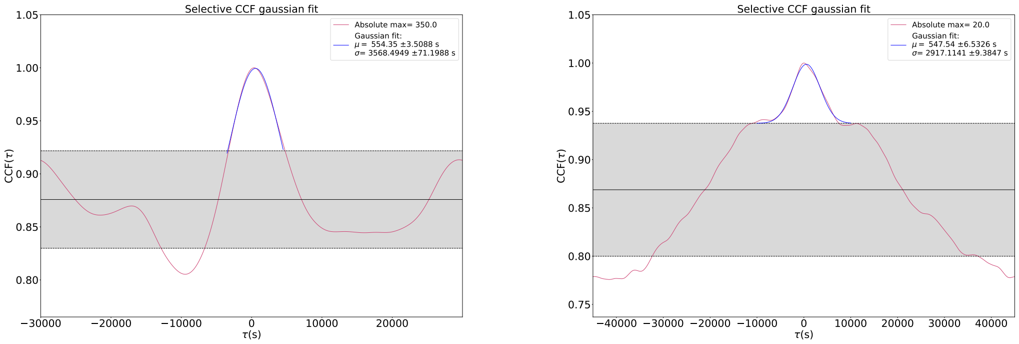

Figure 10.2: CCFs between the 1–10 keV and 0.3–1 keV bands computed with the MDP procedure (n = 100 photons per bin). The displayed profiles are single Poisson realizations within the MDP procedure. Grey shaded areas indicate the ±68% confidence region around the mean CCF. **Left:** The CCF is evaluated over the first ∼63 ks from the start of the XMM2 observation; a Gaussian fitted to the entire CCF profile is overplotted, and the corresponding best–fit parameters are quoted. **Right:** Same CCF setup (MDP with n = 100), computed over the first ∼94 ks from the start of the XMM3 observation. Here, the Gaussian is fitted only to the CCF peak region; the best–fit values are reported.

the bins has $\gtrsim 25$ photons. Under these conditions, the Gaussian statistic is adequate, and the per–bin rate uncertainty is

$$\sigma_{\text{rate}} \simeq \frac{\sqrt{N}}{\Delta t} = \frac{\text{rate}}{\sqrt{N}},$$

where N is the number of counts in each $\Delta t = 250\,\text{s}$ bin.

10.3.3 Time–lag estimation

To estimate the time lag between hard X-Ray (1 keV - 10 keV) and soft X-Ray (0.3 keV - 1 keV) photons I employ the CCF formalism that is reported in chapter 5.

I adopted a fixed counts per bin of $n = 100$ photons for both observations, which sets the relative Poisson uncertainty of each rate point to $1/\sqrt{n} = 10\%$. The total lags for the two datasets are then evaluated over the full duration of each observation (XMM2 and XMM3).

The resulting CCFs are shown for the two observations in Figure 10.2. By the sign convention, when cross-correlating the hard-photon light curve with the soft-photon one, a negative/positive lag means soft-photons lag/anticipate the hard-photons.

Lags and relative uncertainties are obtained with the MDP procedure in chapter 5.

Fitting a Gaussian either to the peak region alone or to the entire CCF (as in Fig. 10.2) yields consistent lag estimates.

The total lags and associated uncertainties (68 % C.L.) are respectively: $\tau_{tot,XMM2}$ = (1230 ± 463) s for the first 62 ks (XMM2) observation, while $\tau_{tot,XMM3}$ = (1010 ± 499) s in the second 94 ks (XMM3), both reported in Figure 10.4.

As an independent check, I reproduced the analysis with the HEASoft task `crosscor`[1], reaching consistent results. I also measured the magnitude-squared coherence via the Welch

---

[1] See `https://heasarc.gsfc.nasa.gov/docs/software/lheasoft/help/crosscor.html`.

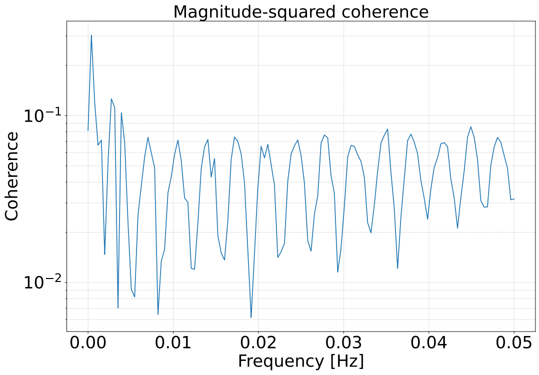

Figure 10.3: Magnitude-squared coherence (Welch averaged periodogram) between the 0.3–1.0 keV and 1.0–10.0 keV light curves, using 100 photons per bin, for the first 10 ks of XMM2 observation. The 30% coherence peak is at f=4×10^{-4} Hz.

averaged periodogram (Welch 1967), between the hard- and soft-X-Ray light curves, as in Figure 10.3, and found it to be highest at low Fourier frequencies ($\lesssim$mHz) corresponding to the computed lag time scales. Taken together, the independent replication and the low-frequency coherence support the robustness of the measured delay and indicate that it reflects a genuine physical process rather than a statistical artifact.

To validate the total evaluated lag, I investigated the time dependence of the lag by dividing each observation into 10-kS subsequent windows.

For each 10-ks time window, lag computation is automated as follows: if the CCF displays a single statistically significant peak relative to the rest of the profile, I fit a Gaussian to that peak alone; if no unique significant peak is present, to avoid arbitrary peak selection, I fit a single Gaussian to the entire CCF.

Results of the analysis are shown in Figure 10.4.

10.4.0 Discussion

The two observation show a total *hard* lag of the order of $t_{\rm lag} \sim 10^3$ s. Such a long lag is *not* compatible with a pure light–travel echo: a reverberation delay of 10^3 s would require an extra path length $c\,t \simeq 3\times10^{13}$ cm (≈ 2 AU), far exceeding the characteristic scales of ULX inner flows and at odds with the usual expectation that reverberation produces *soft* lags at high Fourier frequencies (Uttley et al. 2014).

I propose a possible interpretation: the inward propagation of an accretion–rate fluctuation in a *viscous* timescale in a super–Eddington, geometrically thick flow. In an α–disc the viscous time at radius R is (Shakura & Sunyaev 1973; Pringle 1981):

$$t_{\rm visc}(R) \;\simeq\; \frac{1}{\alpha}\left(\frac{R}{H}\right)^2 \,\Omega^{-1}, \tag{10.1}$$

where H is the vertical scale height and $\Omega(R)$ the orbital frequency. For super–Eddington ULXs one expects $H/R \sim 1$ and $\alpha \sim 0.1$, hence $t_{\rm visc} \approx 10\,\Omega^{-1}$. Identifying the observed delay with the

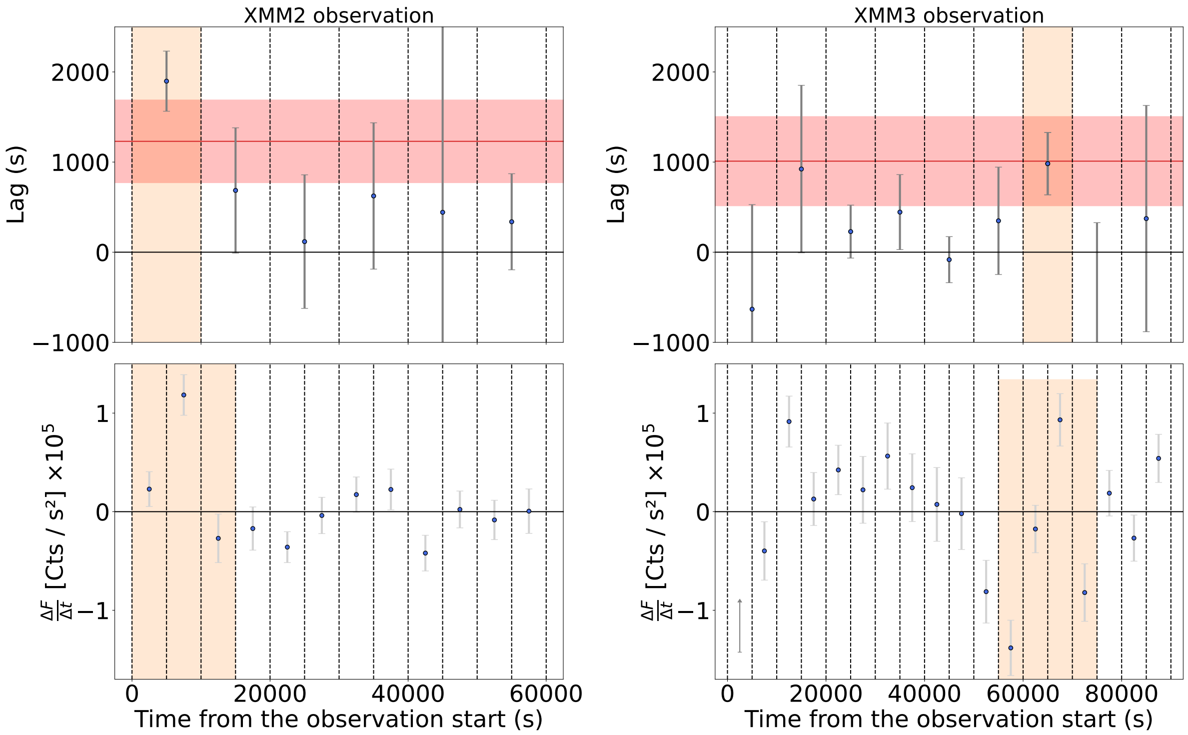

Figure 10.4: **Top:** Time–resolved lags measured in 10 ks windows for the two *XMM-Newton* observations (XMM2 on the left, XMM3 on the right). The solid red line, and its associated uncertainties (highlighted as red shaded band) marks the lag estimated over the full exposure ($\tau_{\mathrm{tot},XMM2}$ and $\tau_{\mathrm{tot},XMM3}$ in the respective panels); the orange shaded band shows the corresponding 1σ interval from the MDP procedure. Positive lags indicate that the soft band leads the hard band (hard lag). **Bottom:** Evolution of the $\Delta F/\Delta t$ (i.e., slope of the count–rate versus time) in adjacent 5 ks windows for the 0.3 - 1 keV energy band. Error bars are at 1σ confidence level.

propagation time, $t_{\mathrm{visc}} \approx t_{\mathrm{lag}} \sim 10^3$ s, gives $\Omega \sim 10^{-2}$ s^{-1} at the relevant radius. Using Keplerian scaling,

$$R \simeq \left(\frac{GM}{\Omega^2}\right)^{1/3}, \tag{10.2}$$

I obtain $R \approx 8.2 \times 10^{-4}$ AU for a 1.4 $M_\odot$ neutron star and $R \approx 1.6 \times 10^{-3}$ AU for a 10 $M_\odot$ black hole, values consistent with plausible disc/inner–wind scales in ULXs (e.g. Middleton et al. 2015a).

The techniques introduced in chapter 5 allow us to track the lag as a function of time by estimating it in successive 10 ks windows. However, I was only able to detect two significant hard lags in the reported time ranges, Obs. XMM2 (0 ks - 10 ks): −1987.5 ± 336.6 s and Obs. XMM3 (60 ks - 70 ks): −982.4 ± 345.7 s (68% c.l.).

Interestingly, the measured lag time windows coincide with marked temporal variations of the Flux $\frac{\Delta F}{\Delta t}$ for the two light curves, as shown in Figure 10.2. This is valid both for the soft and hard X-Ray bands. In addition, within this time window, the coherence profile peaks at the frequency corresponding to the lag timescale (e.g., f$\approx 4 \times 10^{-4}$ Hz $\rightarrow \tau \approx 2500$ s, see Figure 10.3).
In this scenario, fluctuations originate at larger, cooler radii, first enhancing the soft emission, and then propagate inward to the inner Comptonizing corona and produce the hard X-Ray emission. The finite propagation time naturally yields interband lags that increase with the steepness of the

flux rise (large positive dF/dt). When the flux is nearly constant, the driving perturbations and the inward response are weak, and the lag becomes small or undetectable.

This inward propagation does not exclude lags seen in several ULXs soft lags at high frequencies, where light travel (reverberation) effects dominate, and hard lags at low frequencies where viscous propagation times are longer (e.g. De Marco et al. 2013; Hernández-García et al. 2015).

However, this should be viewed as one plausible interpretation; the concentration of lags during flux-rise intervals may simply reflect the greater temporal variability present in those segments. It remains under investigation whether X-Ray emission can originate at the radii discussed here, given that a standard (non-super-Eddington) accretion disk is typically too cool to produce soft X-Ray emission.

10.5.0 Conclusion

In the two XMM oservation (XMM2-XMM3) I measure a hard lag of order 10^3 s. A delay this long is difficult to reconcile with a pure light–travel echo, which would require path lengths of $\sim c\,t \simeq 3\times10^{13}$ cm (i.e. ~ 2 AU), unrealistically large for a ULX environment. A more natural interpretation is that the lag traces the inward propagation, from a radius R, of an accretion–rate fluctuations on the *viscous* timescale of a super–Eddington, geometrically thick flow. In an α–disc with $H/R \sim 1$ and $\alpha \sim 0.1$, the viscous time is $t_{\rm visc} \simeq \frac{1}{\alpha}\left(\frac{R}{H}\right)^2\Omega^{-1} \approx 10\,\Omega^{-1}$. Identifying the observed lag with this timescale ($t_{\rm lag} \approx t_{\rm visc} \sim 10^3$ s) implies $\Omega \sim 10^{-2}\,{\rm s}^{-1}$. Assuming Keplerian rotation around a neutron star, $\Omega^2 = GM/R^3$ then gives a characteristic radius for the accretion–rate fluctuations of $R \sim \left(\frac{GM}{\Omega^2}\right)^{1/3} \sim 10^{10}$ cm, which is consistent with inner disc/wind scales expected in super–Eddington ULXs.

Time–resolved measurements in 10 ks windows detect two statistically significant hard lags that occur precisely during intervals when the light curves show the steepest gradients. This correlation supports a "rate–dependent" picture: rapid rises (strong variability) may reflect perturbations seeded at larger, cooler radii (traced by the soft band) that propagate inward and, after a delay, modulate the inner Comptonizing region emitting the hard X–rays. In flatter intervals, the driving fluctuations are weak or stationary, and the measured lag correspondingly diminishes.

Overall, the data favor propagating–fluctuation models as the origin of the low–frequency hard lags, but it is still under investigation what process could have heated the disc at the inferred characteristic radii, causing it to emit in the soft X-Rays. The proposed mechanism does not exclude reverberation or reprocessing at higher frequencies. Methodologically, tracking the lag in sliding windows (CCF with MDP uncertainties) is effective in the low–count regime of ULXs and reveals when and how the lag grows. Longer, uninterrupted observations and coordinated high–energy coverage will enable lag–frequency–energy mapping and tighter constraints on the coupling between the soft thermal and hard Comptonizing components in super–Eddington accretion flows.

11

Conclusions and Outlook

Context and motivation In the emerging landscape of multi-messenger astrophysics, time-domain astronomy has become central to tackling long-standing questions. The landmark joint detection of GRB 170817A and GW170817 provided the first confirmed association between a short GRB and a binary neutron star merger, showing that such mergers can power the GRB electromagnetic emission. Precise, near-simultaneous localization of both the gravitational wave and electromagnetic signals is therefore essential to secure source associations. Continuous monitoring of the high-energy sky will maximize the chances of future joint detections, an objective well served by constellation-style, distributed satellite missions.

With multiple detectors distributed at different locations in space, one can maximize coverage of the transient sky and localize events by triangulation, exploiting their intrinsic variability.

Nanosatellites and the low-count challenge The use of nanosatellites can effectively meet this need, thanks to the ease and low cost of payload assembly and the reduced launch expenses enabled by the spacecraft's low mass. Conversely, their small effective areas (typically $\lesssim 100\,\mathrm{cm}^2$) force operations in a low statistics regime. In low count conditions (i.e., $\lesssim 10^3$ cts s^{-1}, fewer than a few counts per ms), standard timing techniques tend to degrade as the observed rate decreases. Fixed-width binning, in particular, imposes an unfavorable trade-off: bins that are too wide wash out rapid variability and inter-band lags, whereas overly narrow bins render the light curves shot noise-dominated, break Gaussian error assumptions, and skew lag estimators.

Methods: Estimation of delays The time series analysis techniques developed in Chapter 5 are essential for low-statistic timing, where modest effective areas (typically $\sim 50\,\mathrm{cm}^2$) lead to low photon counts. By adaptively varying the bin width to stabilize counting statistics, they preserve temporal resolution and yield more reliable lag estimates and uncertainties than fixed bin methods. Because the binning responds to the instantaneous rate, the approach is especially effective at detecting short timescale variability in low-count regimes, precisely where fixed-size binning tends to wash out features or become noise-dominated.

Across all count regimes, lag estimates and their uncertainties remain stable thanks to the simulation-free *Modified Double Pool* (MDP) estimator. The results show that MDP delivers science-grade inferences while substantially improving both accuracy and computational efficiency. In particular, it suppresses realization-specific Poisson structure present in input templates, preventing its artificial amplification across Monte Carlo trials. Although each split halves the available counts, incurring the expected $\sqrt{2}$ loss in precision, the resulting lag estimates remain unbiased.

Constellations for localization: performance and FoV These techniques provide a robust route for the future of transient localization via inter–detector time delays (i.e, triangulation method). The results in Chapter 6 validate the use of CubeSats in strategically designed orbital configurations to achieve precise localization and continuous monitoring. Notably, a constellation of three satellites, each with an effective area of 360 cm^2, placed near a Lagrange point and separated by baselines of $\sim 2.5 \times 10^6$ km can reach *Swift*/BAT–class localization accuracy in some cases. However, bringing a dedicated nano-satellite constellation to L2 would largely wash out the intrinsic low-cost advantage of CubeSat missions; this limitation can be mitigated only if such spacecraft are accommodated as secondary payloads on larger missions already targeting L2, thereby sharing launch and transfer costs. The key advantage of such a configuration is the quasi–hemispherical instantaneous field of view achievable with non–imaging, wide–field detectors. This can set a major improvement over the ~ 1.4 sr FoV of *Swift*/BAT, which is also periodically Earth–occulted in low–Earth orbit. It should also be noted that the localization precision improves with the square root of the number of satellites added, allowing the accuracy to be increased arbitrarily as more spacecraft are placed in orbit.

These capabilities underscore the strong scientific potential for multi–messenger and multi–wavelength astrophysics. Arcminute localizations enable rapid, targeted follow–up: facilities can promptly repoint to the inferred sky position without extensive tiling or blind searches, reducing the risk of missing the earliest, and most informative, phases of events such as GRB afterglows.

SpIRIT: On-orbit validation of microsecond timing In parallel, the emerging nanosatellite paradigm is opening promising avenues for astrophysics, especially for high time resolution studies. Chapter 7 and Chapter 8 of this thesis detail the timing performance of a 1U detector ($10 \times 10 \times 10$ cm^3) flown on the 6U SpIRIT nanosatellite, demonstrating, for the first time, that an ultra compact, two in one instrument with broadband coverage from 2 keV to 2 MeV can deliver competitive, science grade timing results. Notably, the limited exposure time of only 730 s makes the 5 σ detection of the Crab pulse profile even more remarkable, demonstrating the outstanding capabilities of the innovative technology hosted by the HERMES payload.

Fundamental physics: limit on Lorentz invariance violation The techniques of this thesis enable stringent tests of LIV as contemplated in some quantum gravity frameworks (e.g., loop quantum gravity). Gamma ray bursts detected at very high energies by *Fermi*/LAT ($E \gtrsim 10$ MeV) typically yield low count rates ($\lesssim 10$ $\mathrm{cts\,s^{-1}}$), making adaptive, Poisson aware timing essential. By measuring inter-band lags between MeV–GeV light curves and a keV reference band, and exploiting the long cosmological propagation baselines, I constrain in Chapter 9 the phenomenological dispersion expected from quantum gravity. Each lag is interpreted as the sum of two components: an intrinsic emission delay and a propagation term induced by potential LIV effects within quantum gravity frameworks. This statistical analysis disentangles these contributions, enabling a clean measurement of the cosmological propagation term and allowing us to place robust constraints on the quantum-gravity energy scale. The only physical assumption made is that the GRB emission mechanism does not evolve with redshift, an

assumption supported by the shown work and far milder than those adopted in single bright-GRB studies.
At the 2σ (95% confidence) level, I obtain the following first-order (linear) lower limits:

$$E_{\mathrm{QG},1} > 1.25 \times 10^{16}\ \text{GeV (subluminal)}, \qquad E_{\mathrm{QG},1} > 2.49 \times 10^{16}\ \text{GeV (superluminal)}.$$

Although some recent analyses have yielded limits closer to the Planck scale, the strength of this approach lies in its robustness and scalability to larger GRB samples in the future. Since the sensitivity improves approximately as $N^{-1/2}$ with the number of bursts, future missions such as *THESEUS*, expected to detect and localize hundreds of GRBs per year with high redshift completeness, will enable at least an order-of-magnitude improvement within a few years of operation. Extending the analysis to the full sample of GRBs with contemporaneous keV and MeV–GeV coverage *and* secure redshifts (e.g., *Swift*/BAT+XRT, *Fermi*/GBM+LAT, *Konus–Wind*, *INTEGRAL*/SPI–ACS, *AGILE*/MCAL, *CALET*/CAL) would further tighten the constraint, progressively approaching the regime of quantum gravity at the Planck energy scale.

Faint source timing: ULX hard lags Among the many additional applications, I have successfully detected *significant hard lags* in extremely faint sources such as ULXs, not only over full *XMM-Newton* exposures but also within shorter segments (e.g., 10 ks). As treated in Chapter 10, the emission and propagation mechanisms that could lead such delays are still under investigation, even though they may be linked to the viscous propagation time of accretion fluctuation. Systematically verifying and interpreting this phenomenology, and extending the analysis to the broader ULX population, could clarify such processes in super Eddington accretion discs.

Final remarks All the discussed results demonstrate the versatility and broad applicability of the timing techniques developed in this work. Although validated in the X–/γ–ray regime, the methodology naturally extends to other wavelength bands, and even to different research domains, whenever precise analysis of time-of-arrival series is required. The framework delivers decisive advantages in low-count (Poisson-limited) regimes, where standard approaches often break down, enabling reliable inference when it matters most.

Bibliography

Aartsen, M. G., Ackermann, M., Adams, J., & et al. 2018, Science, 361, 147

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2020, Living Reviews in Relativity, 23

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, Nature, 551, 85–88

Abbott, B. P., Abbott, R., Abbott, T. D., et al. 2017, The Astrophysical Journal Letters, 848, L13

Abbott, B. P., Abbott, R., Abbott, T. D., & et al. 2017, The Astrophysical Journal Letters, 848, L12

Abdalla, H., Adam, R., Aharonian, F., et al. 2019, Nature, 575, 464–467

Abdo, A. A., Ackermann, M., Ajello, M., et al. 2009, Nature, 462, 331

Abdollahi, S., Acero, F., Ackermann, M., & et al. 2020, The Astrophysical Journal Supplement Series, 247, 33

Abe, H., Abe, S., Acciari, V. A., et al. 2023, Monthly Notices of the Royal Astronomical Society, 527, 5856–5867

Acciari, V. A., Ansoldi, S., Antonelli, L. A., et al. 2019a, Nature, 575, 455–458

Acciari, V. A., Ansoldi, S., Antonelli, L. A., et al. 2019b, Nature, 575, 459–463

Ajello, M., Arimoto, M., Axelsson, M., & et al. 2019, The Astrophysical Journal, 878, 52

Alexander, K. D., Margutti, R., Blanchard, P. K., et al. 2018, The Astrophysical Journal Letters, 863, L18

Alpar, M. A., Cheng, A. F., Ruderman, M. A., & Shaham, J. 1982, Nature, 300, 728

Amati, L., O'Brien, P., Götz, D., et al. 2021, Experimental Astronomy, 52, 183–218

Amelino-Camelia, G. 2005, Introduction to Quantum-Gravity Phenomenology (Springer Berlin Heidelberg), 59–100

Amelino-Camelia, G., Ellis, J., Mavromatos, N. E., Nanopoulos, D. V., & Sarkar, S. 1998, Nature, 393, 763

Amelino-Camelia, G., Freidel, L., Kowalski-Glikman, J., & Smolin, L. 2011, Physical Review D, 84

Andersson, N., Ferrari, V., Jones, D. I., et al. 2010, General Relativity and Gravitation, 43, 409–436

Arnaud, K. A. et al. 1992, The OGIP Standard PHA File Format, HEASARC/OGIP Memo 92-007

Aschenbach, B., Briel, U. G., Haberl, F., et al. 2000, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 4012, X-Ray Optics, Instruments, and Missions III, ed. J. E. Truemper & B. Aschenbach, 731–739

Atwood, W. B., Abdo, A. A., Ackermann, M., et al. 2009, The Astrophysical Journal, 697, 1071

Atwood, W. B., Abdo, A. A., Ackermann, M., & et al. 2009, The Astrophysical Journal, 697, 1071

Baade, W. & Zwicky, F. 1934, Physical Review, 46, 76

Bachetti, M. & et al. 2014, Nature, 514, 202

Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, Nature, 300, 615

Band, D., Matteson, J., Ford, L., et al. 1993, The Astrophysical Journal, 413, 281

Baroni, G., Campana, R., Evangelista, Y., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13093, Space Telescopes and Instrumentation 2024: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, S. Nikzad, & K. Nakazawa, 130935O

Bartos, I., Brady, P., & Márka, S. 2013, Classical and Quantum Gravity, 30, 123001

Baym, G., Pethick, C., & Sutherland, P. 1971, The Astrophysical Journal, 170, 299

Bederede, D., Bougamont, E., Bourgeois, P., et al. 2004, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 518, 15

Begelman, M. C. 2002, The Astrophysical Journal Letters, 568, L97

Beloborodov, A. M. 2010, Monthly Notices of the Royal Astronomical Society, 407, 1033–1047

Berger, E. 2014, Annual Review of Astronomy and Astrophysics, 52, 43

Bernardini, M. G., Ghirlanda, G., Campana, S., et al. 2017, Astronomy & Astrophysics, 607, A121

Bethe, H. A. 1990, Reviews of Modern Physics, 62, 801

Bhattacharya, D. & van den Heuvel, E. P. J. 1991, Physics Reports, 203, 1

Bissaldi, E., von Kienlin, A., Lichti, G., et al. 2009, EXP ASTRON, 24, 47–88

Blanch, O., Gaug, M., Noda, K., et al. 2020, GRB Coordinates Network, 28659, 1

Bouwens, R. J., Illingworth, G. D., Oesch, P. A., et al. 2011, The Astrophysical Journal, 737, 90

Bozzo, E., Amati, L., O'Brien, P., Goetz, D., & Santangelo, A. 2024, Proc. SPIE Int. Soc. Opt. Eng., 13093, 130932B

Bradshaw, C. F., Fomalont, E. B., & Geldzahler, B. J. 1999, The Astrophysical Journal Letters, 512, L121

Brandt, S., Budtz-Jørgensen, C., Lund, N., et al. 2003, Astronomy and Astrophysics, 411, L243

Briggs, M. S., Band, D. L., Kippen, R. M., et al. 1999, The Astrophysical Journal, 524, 82–91

Bucciantini, N., Quataert, E., Metzger, B. D., et al. 2009, Monthly Notices of the Royal Astronomical Society, 396, 2038–2050

Burderi, L., Sanna, A., Di Salvo, T., Fiore, F., & Riggio, A. 2022, in 44th COSPAR Scientific Assembly. Held 16-24 July, Vol. 44, 3009

Burderi, L., Sanna, A., Di Salvo, T., et al. 2021, Experimental Astronomy, 52, 439

Błaszkiewicz, L., Lewandowski, W., Krankowski, A., et al. 2016, Acta Geophysica, 64

Camisasca, A. E., Guidorzi, C., Amati, L., et al. 2023, Astronomy & Astrophysics, 671, A112

Campana, R., Fuschino, F., Evangelista, Y., Dilillo, G., & Fiore, F. 2020, in Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, K. Nakazawa, & S. Nikzad (SPIE)

Cano, Z., Wang, S.-Q., Dai, Z.-G., & Wu, X.-F. 2017, Advances in Astronomy, 2017, 1–41

Canuto, V., Lodenquai, J., & Ruderman, M. 1971, Physical Review D, 3, 2303

Cao, Z., Aharonian, F., An, Q., et al. 2023, Science, 380, 1390–1396

Cao, Z., Aharonian, F., Axikegu, et al. 2024, Physical Review Letters, 133

Castignani, G., Guetta, D., Pian, E., et al. 2014, A&A, 565, A60

Chadwick, J. 1932, Proceedings of the Royal Society of London Series A, 136, 692

Chandrasekhar, S. 1931, The Astrophysical Journal, 74, 81

Chen, K. & Ruderman, M. 1993, The Astrophysical Journal, 408, 179

Cheng, K. S., Ho, C., & Ruderman, M. 1986, The Astrophysical Journal, 300, 500

Chevalier, R. A. & Li, Z. 2000, The Astrophysical Journal, 536, 195–212

Chornock, R., Berger, E., Kasen, D., et al. 2017, The Astrophysical Journal Letters, 848, L19

Colagrossi, A., Prinetto, J., Silvestrini, S., & Lavagna, M. 2020, Journal of Astronomical Telescopes, Instruments, and Systems, 6

Colbert, E. J. M. & Mushotzky, R. F. 1999, The Astrophysical Journal, 519, 89

Costa, E., Frontera, F., Heise, J., et al. 1997, Nature, 387, 783

Coulter, D. A., Foley, R. J., Kilpatrick, C. D., et al. 2017, Science, 358, 1556

Courvoisier, T. J. L., Walter, R., Beckmann, V., et al. 2003, Astronomy and Astrophysics, 411, L53

Dado, S. & Dar, A. 2019, Critical Tests Of Leading Gamma Ray Burst Theories

Daigne, F. & Mochkovitch, R. 2000, Astronomy and Astrophysics, 358, 1157

Daugherty, J. K. & Harding, A. K. 1996a, The Astrophysical Journal, 458, 278

Daugherty, J. K. & Harding, A. K. 1996b, The Astrophysical Journal, 458, 278

De Marco, B., Ponti, G., Miniutti, G., et al. 2013, Monthly Notices of the Royal Astronomical Society, 436, 3782–3791

De Pasquale, M., Piro, L., Gendre, B., et al. 2006, Astronomy &; Astrophysics, 455, 813–824

Degenaar, N. & Suleimanov, V. F. 2018, arXiv e-prints, arXiv:1806.02833

Dekking, F. M., Kraaikamp, C., Lopuhaä, H. P., & Meester, L. E. 2005, The Poisson process (London: Springer London), 167–179

den Herder, J. W., Brinkman, A. C., Kahn, S. M., et al. 2001, Astronomy and Astrophysics, 365, L7

Di Salvo, T. & Sanna, A. 2022, in Astrophysics and Space Science Library, Vol. 465, Astrophysics and Space Science Library, ed. S. Bhattacharyya, A. Papitto, & D. Bhattacharya, 87–124

Dilillo, G., Campana, R., Zampa, N., et al. 2020, in Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, K. Nakazawa, & S. Nikzad (SPIE), 185

do Couto e Silva, E., Anthony, P., Arnold, R., et al. 2001, Nuclear Instruments and Methods in Physics Research A, 474, 19

Drenkhahn, G. & Spruit, H. C. 2002, Astronomy and Astrophysics, 391, 1141

Drout, M. R., Piro, A. L., Shappee, B. J., et al. 2017, Science, 358, 1570–1574

D'Avanzo, P., Campana, S., Salafia, O. S., et al. 2018, Astronomy & Astrophysics, 613, L1

Eichler, D., Livio, M., Piran, T., & Schramm, D. N. 1989, Nature, 340, 126

Ellis, J., Konoplich, R., Mavromatos, N. E., et al. 2019, Physical Review D, 99

Etienne, Z. B., Liu, Y. T., Paschalidis, V., & Shapiro, S. L. 2012, Physical Review D, 85

Evangelista, Y., Fiore, F., Fuschino, F., et al. 2020a, in Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, K. Nakazawa, & S. Nikzad (SPIE)

Evangelista, Y., Fiore, F., Fuschino, F., et al. 2020b, in Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, K. Nakazawa, & S. Nikzad (SPIE)

Evans, P. A., Beardmore, A. P., Page, K. L., et al. 2009, Monthly Notices of the Royal Astronomical Society, 397, 1177

Fabbiano, G. 1989, Annual Review of Astronomy and Astrophysics, 27, 87

Fabbiano, G., Zezas, A., & Murray, S. S. 2001, The Astrophysical Journal, 554, 1035

Fabrika, S., Ueda, Y., Vinokurov, A., Sholukhova, O., & Shidatsu, M. 2021, Universe, 7, 293

Fiore, F., Burderi, L., Lavagna, M., et al. 2020, in Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, K. Nakazawa, & S. Nikzad (SPIE)

Flanagan, É. É. & Hughes, S. A. 1998, Physical Review D, 57, 4535

Fomalont, E. B., Geldzahler, B. J., & Bradshaw, C. F. 2001, The Astrophysical Journal, 558, 283–301

Fong, W., Berger, E., & Fox, D. B. 2010, The Astrophysical Journal, 708, 9

Fong, W.-f., Nugent, A. E., Dong, Y., et al. 2022, The Astrophysical Journal, 940

Frail, D. A., Kulkarni, S. R., Nicastro, L., Feroci, M., & Taylor, G. B. 1997, Nature, 389, 261

Frontera, F., Amati, L., Costa, E., et al. 2000, APJ Supplement Series, 127, 59–78

Fruchter, A. S., Levan, A. J., Strolger, L., et al. 2006, Nature, 441, 463

Fuschino, F., Campana, R., Labanti, C., et al. 2019, Nuclear Instruments and Methods in Physics Research A, 936, 199

Gabriel, C., Guainazzi, M., & Metcalfe, L. 2006, Astronomy and Astrophysics, Supplement, 305, 315

Gaia Collaboration, Vallenari, A., Brown, A. G. A., et al. 2023, Astronomy and Astrophysics, 674, A1

Galama, T. J., Vreeswijk, P. M., van Paradijs, J., et al. 1998, Nature, 395, 670–672

Gao, Y., Wang, Q. D., Appleton, P. N., & Lucas, R. A. 2003, The Astrophysical Journal Letters, 596, L171

Gatti, E. & Rehak, P. 1984, Nuclear Instruments and Methods in Physics Research, 225, 608

Gehrels, N., Ramirez-Ruiz, E., & Fox, D. 2009, Annual Review of Astronomy and Astrophysics, 47, 567–617

George, I. M., Rots, A., & Mukai, K. 1994, Guidelines for Defining FITS Formats for Event Lists, HEASARC/OGIP Memo 94-003

Ghirlanda, G., Nava, L., Salafia, O., et al. 2024, Astronomy & Astrophysics, 689, A175

Ghirlanda, G., Salafia, O. S., Paragi, Z., et al. 2019, Science, 363, 968–971

Giacconi, R., Gursky, H., Paolini, F. R., & Rossi, B. 1962a, Physical Review Letters, 9, 439

Giacconi, R., Gursky, H., Paolini, F. R., & Rossi, B. B. 1962b, Phys. Rev. Lett., 9, 439

Giuliani, A., Mereghetti, S., Fornari, F., et al. 2008, A&A, 491, L25–L28

Goldreich, P. & Julian, W. H. 1969, The Astrophysical Journal, 157, 869

Goldstein, A., Burgess, J. M., Preece, R. D., et al. 2012, The Astrophysical Journal Supplement Series, 199, 19

Goldstein, A., Veres, P., Burns, E., & et al. 2017, The Astrophysical Journal Letters, 848, L14

Goodman, J. 1986, The Astrophysical Journal Letters, 308, L47

Gruber, D. E., Matteson, J. L., Peterson, L. E., & Jung, G. V. 1999a, The Astrophysical Journal, 520, 124

Gruber, D. E., Matteson, J. L., Peterson, L. E., & Jung, G. V. 1999b, The Astrophysical Journal, 520, 124

Guan, J., Lu, F.-J., Zhang, S., et al. 2020, Journal of High Energy Astrophysics, 26, 11

Guetta, D. & Della Valle, M. 2007, The Astrophysical Journal Letters, 657, L73

Guidorzi, C., Lacapra, M., Frontera, F., et al. 2010, Astronomy &; Astrophysics, 526, A49

Guidorzi, C., Margutti, R., Brout, D., et al. 2017, The Astrophysical Journal Letters, 851, L36

Haggard, D., Nynka, M., Ruan, J. J., et al. 2017, The Astrophysical Journal Letters, 848, L25

Hajela, A., Margutti, R., Alexander, K. D., et al. 2019, The Astrophysical Journal Letters, 886, L17

Hanish, D. J., Meurer, G. R., Ferguson, H. C., et al. 2006, The Astrophysical Journal, 649, 150–162

Hasinger, G., Priedhorsky, W. C., & Middleditch, J. 1989, The Astrophysical Journal, 337, 843

Hernández-García, L., Vaughan, S., Roberts, T. P., & Middleton, M. 2015, Monthly Notices of the Royal Astronomical Society, 453, 2878–2885

Hewish, A., Bell, S. J., Pilkington, J. D. H., Scott, P. F., & Collins, R. A. 1968, Nature, 217, 709

Hjorth, J. & Bloom, J. S. 2011, The Gamma-Ray Burst - Supernova Connection

Hopkins, A. M. 2004, The Astrophysical Journal, 615, 209

Hurley, K., Pal'shin, V. D., Aptekar, R. L., et al. 2013, The Astrophysical Journal Supplement, 207, 39

IceCube Collaboration, Fermi-LAT, MAGIC, et al. 2018, Science, 361, eaat1378

Jackson, J. D. 1998, Classical Electrodynamics, 3rd edn. (New York: Wiley)

Jacob, U. & Piran, T. 2008, Journal of Cosmology and Astroparticle Physics, 2008, 031

Jain, C. & Paul, B. 2011, Research in Astronomy and Astrophysics, 11, 1134–1142

Jansen, F. & et al. 2001, Astronomy & Astrophysics, 365, L1

Kaneko, Y., Preece, R. D., Briggs, M. S., et al. 2006, The Astrophysical Journal Supplement Series, 166, 298–340

Kann, D. A., Klose, S., Zhang, B., et al. 2010, The Astrophysical Journal, 720, 1513

Kaspi, V. M. & Beloborodov, A. M. 2017, Annual Review of Astronomy and Astrophysics, 55, 261–301

King, A. R. & Pounds, K. A. 2003, Monthly Notices of the Royal Astronomical Society, 345, 657

King, A. R. & Ritter, H. 1999, Monthly Notices of the Royal Astronomical Society, 309, 253

Kingman, J. F. C. & Taylor, S. J. 1966, Characteristic functions (Cambridge University Press), 314–334

Kirsch, M. G., Briel, U. G., Burrows, D., et al. 2005, in UV, X-Ray, and Gamma-Ray Space Instrumentation for Astronomy XIV, ed. O. H. W. Siegmund (SPIE)

Klebesadel, R. W., Strong, I. B., & Olson, R. A. 1973, The Astrophysical Journal Letters, 182, L85

Klein, R. W. & Roberts, S. D. 1984, Simulation, 43, 193

Komissarov, S. S., Vlahakis, N., Königl, A., & Barkov, M. V. 2009, Monthly Notices of the Royal Astronomical Society, 394, 1182–1212

Körding, E., Falcke, H., & Markoff, S. 2002, in Proceedings of the 6th European VLBI Network Symposium

Kováčik, S., Ďuríšková, M., & Rusnák, P. 2024, Phenomenology of the dispersion law in three-dimensional quantum space

Kuiper, L., Hermsen, W., Walter, R., & Foschini, L. 2003, Astronomy and Astrophysics, 411, L31

Kulkarni, S. R., Frail, D. A., Wieringa, M. H., et al. 1998, Nature, 395, 663

Kumar, P. & Zhang, B. 2015, Physics Reports, 561, 1–109

Lamb, D. Q. 1995, Publications of the Astronomical Society of the Pacific, 107, 1152

Langer, N. 2012, Annual Review of Astronomy and Astrophysics, 50, 107

Lattimer, J. M. & Prakash, M. 2016, Physics Reports, 621, 127

Leone, W., Burderi, L., di Salvo, T., et al. 2025a, Time domain astrophysics with transient sources. Delay estimate via Cross Correlation Function techniques

Leone, W., Burderi, L., di Salvo, T., et al. 2025b, Astronomy &; Astrophysics, 701, A50

Li, L.-X. 2008, Monthly Notices of the Royal Astronomical Society, 388, 1487–1500

Li, X., Song, L., Li, X., Tan, Y., & Yang, Y. 2020, Journal of Astronomical Instrumentation, reports ~50 μs absolute timing accuracy early in mission

Lin, R. P., Dennis, B. R., Hurford, G. J., et al. 2002, Solar Physics, 210, 3

Lithwick, Y. & Sari, R. 2001, The Astrophysical Journal, 555, 540–545

Liu, C. Z., Zhang, Y. F., Li, X. F., et al. 2019, The High Energy X-ray telescope (HE) onboard the Insight-HXMT astronomy satellite

Logue, J., Ott, C. D., Heng, I. S., Kalmus, P., & Scargill, J. 2012, Physical Review D, 86, 044023

Longair, M. S. 2011, High Energy Astrophysics, 3rd edn. (Cambridge, UK: Cambridge University Press)

Lyutikov, M. 2013, Monthly Notices of the Royal Astronomical Society, 431, 2580–2589

MacFadyen, A. I. & Woosley, S. E. 1999, The Astrophysical Journal, 524, 262–289

Maggiore, M., Broeck, C. V. D., Bartolo, N., et al. 2020, Journal of Cosmology and Astroparticle Physics, 2020, 050–050

Maraventano, C., Ghirlanda, G., Nava, L., et al. 2025, Astronomy & Astrophysics, 697, A161

Margutti, R., Zaninoni, E., Bernardini, M. G., et al. 2012, Monthly Notices of the Royal Astronomical Society, 428, 729–742

Mason, K. O., Breeveld, A., Much, R., et al. 2001, Astronomy and Astrophysics, 365, L36

Mazzola, S. M., Iaria, R., Di Salvo, T., et al. 2021, Astronomy and Astrophysics, 654, A102

McLaughlin, M. A., Lyne, A. G., Lorimer, D. R., et al. 2006, Nature, 439, 817

Meegan, C., Lichti, G., Bhat, P. N., et al. 2009, The Astrophysical Journal, 702, 791

Meegan, C., Lichti, G., Bhat, P. N., et al. 2009a, The Astrophysical Journal, 702, 791

Meegan, C., Lichti, G., Bhat, P. N., & et al. 2009b, The Astrophysical Journal, 702, 791

Metzger, M. R., Djorgovski, S. G., Kulkarni, S. R., et al. 1997, Nature, 387, 878

Middleton, M. J., Heil, L., Pintore, F., Walton, D. J., & Roberts, T. P. 2015a, Monthly Notices of the Royal Astronomical Society, 447, 3243

Miller, M. C. & Hamilton, D. P. 2002, Monthly Notices of the Royal Astronomical Society, 330, 232

MIT. 2008, Signal Processing - Continuous and Discrete (Department of Mechanical Engineering)

Mo, G., Jayaraman, R., Frostig, D., et al. 2023, Multi-messenger astrophysics in the gravitational-wave era

Moiseev, A. A., Hartman, R. C., Ormes, J. F., et al. 2007, Astroparticle Physics, 27, 339

Mooley, K. P., Deller, A. T., Gottlieb, O., et al. 2018, Nature, 561, 355–359

Mushtukov, A. A., Suleimanov, V. F., Tsygankov, S. S., & Poutanen, J. 2015, Monthly Notices of the Royal Astronomical Society, 454, 2539

Narayan, R., Sądowski, A., & Soria, R. 2017, Monthly Notices of the Royal Astronomical Society, 469, 2997

Nava, L., Ghirlanda, G., Ghisellini, G., & Celotti, A. 2011, Astronomy and Astrophysics, 530, A21

Nousek, J. A., Kouveliotou, C., Grupe, D., et al. 2006, The Astrophysical Journal, 642, 389

Oppenheimer, J. R. & Volkoff, G. M. 1939, Physical Review, 55, 374

Ostriker, J. P. & Gunn, J. E. 1969, The Astrophysical Journal, 157, 1395

Paczynski, B. 1986, The Astrophysical Journal Letters, 308, L43

Paczynski, B. 1995, Publ. Astron. Soc. Pac., 107, 1167

Paczynski, B. & Rhoads, J. E. 1993, The Astrophysical Journal Letters, 418, L5

Pan, Y., Tian, J., Cao, S., Jiang, Q.-Q., & Qian, W.-L. 2024, Model-independent constraints on Lorentz Invariance Violation with update observations of Gamma-Ray Bursts

Panaitescu, A., Wen, L., Laguna, P., & Mészáros, P. 1997, The Astrophysical Journal, 482, 942

Pang, P. T. H., Dietrich, T., Coughlin, M. W., et al. 2023, Nature Communications, 14

Park, K. 2018, Fundamentals of Probability and Stochastic Processes with Applications to Communications

Pelassa, V., Preece, R., Piron, F., Omodei, N., & Guiriec, S. 2010, The LAT Low-Energy technique for Fermi Gamma-Ray Bursts spectral analysis

Perley, D. A., Cenko, S. B., Corsi, A., et al. 2014, The Astrophysical Journal, 781, 37

Pescalli, A., Ghirlanda, G., Salvaterra, R., et al. 2016, Astronomy & Astrophysics, 587, A40

Peterson, B. M., Wanders, I., Horne, K., et al. 1998, Publications of the Astronomical Society of the Pacific, 110, 660

Petrov, A. N., Grigoryan, O. R., & Kuznetsov, N. V. 2009, Advances in Space Research, 43, 654

Pe'er, A., Meszaros, P., & Rees, M. J. 2006, The Astrophysical Journal, 642, 995–1003

Pian, E., D'Avanzo, P., Benetti, S., et al. 2017, Nature, 551, 67–70

Pintore, F., Marelli, M., Salvaterra, R., Israel, G. L., & et al. 2020, The Astrophysical Journal, 890, 166

Piran, T. 1999, Physics Reports, 314, 575–667

Piran, T. 2004, Gamma-Ray Bursts as Probes for Quantum Gravity (Springer-Verlag), 351–362

Piran, T. & Ofengeim, D. D. 2024, Phys. Rev. D, 109, L081501

Pishro-Nik, H. 2014, Introduction to Probability, Statistics, and Random Processes (Kappa Research, LLC)

Podsiadlowski, P., Rappaport, S., & Pfahl, E. 2002, The Astrophysical Journal, 565, 1107

Portegies Zwart, S. F. & McMillan, S. L. W. 2002, The Astrophysical Journal, 576, 899

Powell, J., Gossan, S., Logue, J., & Heng, I. S. 2016, Physical Review D, 94, 123012

Pringle, J. E. 1981, Annual Review of Astronomy and Astrophysics, 19, 137

Puccetti, S., Perri, M., Campana, R., et al. 2024, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 13098, Observatory Operations: Strategies, Processes, and Systems X, ed. C. R. Benn, A. Chrysostomou, & L. J. Storrie-Lombardi, 130980O

Punturo, M., Abernathy, M., Acernese, F., et al. 2010, Classical and Quantum Gravity, 27, 194002

Racusin, J. L., Liang, E. W., Burrows, D. N., et al. 2009, The Astrophysical Journal, 698, 43–74

Ramirez-Ruiz, E. & Fenimore, E. E. 2000, The Astrophysical Journal, 539, 712

Rees, M. J. & Meszaros, P. 1992, Monthly Notices of the Royal Astronomical Society, 258, 41

Rees, M. J. & Meszaros, P. 1994, The Astrophysical Journal Letters, 430, L93

Rees, M. J. & Meszaros, P. 2005, The Astrophysical Journal, 628, 847–852

Ripa, J., Meszaros, A., Wigger, C., et al. 2009, VizieR Online Data Catalog: GRB observed with RHESSI satellite (Ripa+ 2009), VizieR On-line Data Catalog: J/A+A/498/399. Originally published in: A&A 498, 399 (2009), accessed via CDS VizieR service

Rots, A. H., Jahoda, K., & Lyne, A. G. 2004, The Astrophysical Journal, 605, L129–L132

Ruderman, M. 1975, in Seventh Texas Symposium on Relativistic Astrophysics, ed. P. G. Bergman, E. J. Fenyves, & L. Motz, Vol. 262, 164–180

Ruderman, M. A. & Sutherland, P. G. 1975, The Astrophysical Journal, 196, 51

Ruiz, M., Shapiro, S. L., & Tsokaros, A. 2018, Physical Review D, 97

Rybicki, G. B. & Lightman, A. P. 1986, Radiative Processes in Astrophysics (Wiley)

Ryde, F. 2005, The Astrophysical Journal, 625, L95–L98

Ryde, F., Pe'er, A., Nymark, T., et al. 2011, Monthly Notices of the Royal Astronomical Society, 415, 3693–3705

Sanna, A., Burderi, L., Di Salvo, T., et al. 2020, in Space Telescopes and Instrumentation 2020: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, K. Nakazawa, & S. Nikzad (SPIE)

Sari, R., Piran, T., & Narayan, R. 1998, The Astrophysical Journal Letters, 497, L17

Schmidt, W. K. H. 1978, Nature, 271, 525

Shakura, N. I. & Sunyaev, R. A. 1973, Astronomy & Astrophysics, 24, 337

Shemi, A. & Piran, T. 1990, The Astrophysical Journal Letters, 365, L55

Song, X.-Y., Xiong, S.-L., Zhang, S.-N., et al. 2022, The Astrophysical Journal Supplement, 259, 46

Strüder, L., Briel, U., Dennerl, K., et al. 2001, Astronomy and Astrophysics, 365, L18

Sutton, A. D., Roberts, T. P., & Middleton, M. J. 2013, Monthly Notices of the Royal Astronomical Society, 435, 1758

Swank, J. H. 1999, Nuclear Physics B Proceedings Supplements, 69, 12

Tamagawa, T., Enoto, T., Kitaguchi, T., et al. 2025, Publications of the Astronomical Society of Japan, 77, 466–479

Taniguchi, Y., Shioya, Y., Tsuru, T. G., & Ikeuchi, S. 2000, Publications of the Astronomical Society of Japan, 52, 533

Thoene, C. C. & Fynbo, J. P. U. 2007, The spatially resolved host of GRB 060505 and implications for the nature of the progenitor

Thompson, D., Godfrey, G., Williams, S., et al. 2002, IEEE Transactions on Nuclear Science, 49, 1898–1903

Thompson, R. I., Eisenstein, D., Fan, X., et al. 2006, The Astrophysical Journal, 647, 787

Trenti, M., del Castillo, M. O., Mearns, R., et al. 2024, SpIRIT Mission: In-Orbit Results and Technology Demonstrations

Troja, E., Piro, L., van Eerten, H., et al. 2017, Nature, 551

Tsvetkova, A., Burderi, L., Riggio, A., Sanna, A., & Di Salvo, T. 2023, Universe, 9, 359

Tsvetkova, A., Frederiks, D., Svinkin, D., et al. 2021, The Astrophysical Journal, 908, 83

Tully, R. B., Courtois, H. M., & Sorce, J. G. 2016, The Astronomical Journal, 152, 50

Tuo, Y., Li, X., Ge, M., et al. 2022, The Astrophysical Journal Supplement Series, 259, 14

Turolla, R., Zane, S., & Watts, A. L. 2015, Reports on Progress in Physics, 78, 116901

Uttley, P., Cackett, E. M., Fabian, A. C., Kara, E., & Wilkins, D. R. 2014, Astronomy and Astrophysics Review, 22, 72

Vallado, D. A., Crawford, P., Hujsak, R., & Kelso, T. S. 2006, in AIAA/AAS Astrodynamics Specialist Conference and Exhibit (Keystone, Colorado: American Institute of Aeronautics and Astronautics)

van Paradijs, J., Groot, P. J., Galama, T., et al. 1997, Nature, 386, 686

Vardanyan, V., Takhistov, V., Ata, M., & Murase, K. 2023, Physical Review D, 108

Vasileiou, V., Jacholkowska, A., Piron, F., et al. 2013, Physical Review D, 87

von Kienlin, A., Meegan, C. A., Paciesas, W. S., et al. 2020, The Astrophysical Journal, 893, 46

Vyas, M. K., Pe'er, A., & Iyyani, S. 2024, Unified theory of negative and positive spectral lags in GRB prompt phase due to shear Comptonization from a structured jet

Walton, D. J. & et al. 2022, Monthly Notices of the Royal Astronomical Society, 509, 1587

Walton, D. J., Roberts, T. P., Mateos, S., & Heard, V. 2011, Monthly Notices of the Royal Astronomical Society, 416, 1844

Wanderman, D. & Piran, T. 2010, Monthly Notices of the Royal Astronomical Society, 406, 1944

Wei, J.-J., Zhang, B.-B., Shao, L., Wu, X.-F., & Mészáros, P. 2017, The Astrophysical Journal Letters, 834, L13

Weisskopf, M. C. & et al. 2000, Proceedings of SPIE, 4012, 2

Weisskopf, M. C., Tennant, A. F., Yakovlev, D. G., et al. 2011, The Astrophysical Journal, 743, 139

Welch, P. 1967, IEEE Transactions on Audio and Electroacoustics, 15, 70

Wen, J.-X., Zheng, X.-T., Yu, J.-D., et al. 2021, Nuclear Science and Techniques, 32

Werner, N., Řípa, J., Pál, A., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10699, Space Telescopes and Instrumentation 2018: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, S. Nikzad, & K. Nakazawa, 106992P

Woosley, S. & Bloom, J. 2006, Annual Review of Astronomy and Astrophysics, 44, 507–556

Woosley, S. E. 1993, The Astrophysical Journal, 405, 273

Woosley, S. E. & Heger, A. 2006, The Astrophysical Journal, 637, 914

Yoneyama, M., Kataoka, J., Arimoto, M., et al. 2018, Journal of Instrumentation, 13, P02023

Zajaček, M., Czerny, B., Martínez-Aldama, M. L., & Karas, V. 2019, Astronomische Nachrichten, 340, 577

Zhang, B., Fan, Y. Z., Dyks, J., et al. 2006, The Astrophysical Journal, 642, 354

Zhang, B. & Mészàros, P. 2004, International Journal of Modern Physics A, 19, 2385–2472

Zhang, B. & Yan, H. 2010, The Astrophysical Journal, 726, 90

Zhang, S., Zhang, S. N., Lu, F. J., et al. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10699, Space Telescopes and Instrumentation 2018: Ultraviolet to Gamma Ray, ed. J.-W. A. den Herder, S. Nikzad, & K. Nakazawa, 106991U

Zhang, S.-N., Li, T., Lu, F., et al. 2020, Science China Physics, Mechanics & Astronomy, 63

Zhang, S.-N., Li, T., Lu, F., et al. 2020, Science China Physics, Mechanics, and Astronomy, 63, 249502

Zhang, Y. H., Celotti, A., Treves, A., et al. 1999, APJ, 527, 719–732

Zhao, H. 2020, in Astronomical Society of the Pacific Conference Series, Vol. 527, Astronomical Data Analysis Software and Systems XXIX, ed. R. Pizzo, E. R. Deul, J. D. Mol, J. de Plaa, & H. Verkouter, 469

Zhou, D.-K., Zheng, S.-J., Song, L.-M., et al. 2021, Research in Astronomy and Astrophysics, 21, 005

Appendix

A.1.0 normalized Poisson Probability Function

The detection process is a probabilistic process in which the infinitesimal probability of detecting a photon within an infinitesimal time interval *dt* is:

$$\mathrm{dP(t) = r(t)dt.} \tag{A.1}$$

The probability to detect *N* photons in a time Δt for a given rate r(t) is (Kingman & Taylor 1966; Pishro-Nik 2014; Dekking et al. 2005):

$$\mathrm{P}(\mu, \mathrm{N}) = \frac{\mu^{\mathrm{N}} \mathrm{e}^{-\mu}}{\mathrm{N}!} = \mathrm{P}_{\mathrm{N},\Delta \mathrm{t}}(\mathrm{r}), \tag{A.2}$$

where $\mu(\mathrm{t}) = \mathrm{r(t)}\Delta\mathrm{t}$ and I assumed that Δt is small with respect to the time scale on which $\mu(\mathrm{t})$ varies. Therefore:

$$\mathrm{P}_{\mathrm{N},\Delta \mathrm{t}}(\mathrm{r}) = \frac{(\mathrm{r}\,\Delta\mathrm{t})^{\mathrm{N}} \mathrm{e}^{-\mathrm{r}\Delta\mathrm{t}}}{\mathrm{N}!}. \tag{A.3}$$

Equation A.3 represents the probability of detecting *N* photons in a time interval Δt given a rate r(t). Apart from the overall normalization factor, Equation A.3 can be interpreted as the probability that *N* photons, detected in a time interval Δt, derive from a particular rate r(t). Since the detection of *N* photons must depend on a rate, among all the possible rates, the normalization factor A is obtained by integrating Equation A.3 in rate between 0 and ∞:

$$\mathrm{A} \int_0^{+\infty} \mathrm{P}_{\mathrm{N},\Delta \mathrm{t}}(\mathrm{r}) \mathrm{dr} = 1, \tag{A.4}$$

which gives $A = \frac{1}{\Delta \mathrm{t}}$.

Therefore, given that *N* photons are detected within a time interval Δt, the probability distribution of the rate *r* is:

$$\mathrm{Q}_{\mathrm{N},\Delta \mathrm{t}}(\mathrm{r}) = \Delta\mathrm{t} \frac{(\mathrm{r}\,\Delta\mathrm{t})^{\mathrm{N}} \mathrm{e}^{-\mathrm{r}\Delta\mathrm{t}}}{\mathrm{N}!}. \tag{A.5}$$

A.1.1 Statistical Confidence Level

To evaluate the confidence level (CL) for the rate (as in any statistical treatment), I must integrate Equation 5.5 between two points of equal probability, one below and one above the (unique) maximum of the distribution. Notably, in the case of $N = 0$, the function simplifies to $e^{-r\Delta t}$, which is monotonically decreasing, allowing me to determine an upper limit.

By this definition, the CL corresponds to the area under the normalized probability distribution $Q_{N,\Delta t}(r)$, enclosed between the upper and lower bounds of the rate confidence interval:

$$\mathrm{CL} = \int_{r_{min}}^{r_{max}} Q_{N,\Delta t}(r)dr, \tag{A.6}$$

with the constraint:

$$Q_{N,\Delta t}(r_{min}) = Q_{N,\Delta t}(r_{max}). \tag{A.7}$$

A.2.0 Poisson characteristic values

A.2.1 Mode

The Poisson mode is the rate value where the $Q(r\,\Delta t; N)$ is maximum, $\frac{\partial Q_{N;\Delta t}(r)}{\partial r}|_{r=r_{mode}} = 0$:

$$\frac{\Delta t^2}{N!} e^{-r_{mod}\,\Delta t}\left[N(r_{mod}\,\Delta t)^{N-1} - (r_{mod}\,\Delta t)^N\right] = 0 \tag{B1}$$

that leads to $N(r_{mod}\,\Delta t)^{N-1} = (r_{mod}\,\Delta t)^N$, so the mode value is

$$r_{mod}(N, \Delta t) = \frac{N}{\Delta t}$$

A.2.2 Median

The Poisson median divides the area under the Poisson distribution into two identical parts ($x = r\Delta t$):

$$\int_0^{r_{med}} Q(r\,\Delta t; N)dr = \frac{\Delta t}{N!}\int_0^{r_{med}} x^N\, e^{-x}dx = \frac{1}{2} \tag{B2}$$

The integral $\int (xk)^N\, e^{-xk}\, dk = -\frac{e^{-xk}}{k} \sum_{l=0}^{N} \frac{N!}{l!}(xk)^l$

So the median value can be numerically solved by imposing:

$$\left[-\,e^{-r\,\Delta t}\sum_{l=0}^{N}\frac{(r\,\Delta t)^l}{l!}\right]_0^{r_{med}} = \frac{1}{2} \tag{B3}$$

A.2.3 Mean

The Poisson mean can be evaluated by considering the expectation value formula:

$$\mathrm{r_{mean}} = \frac{\int_0^\infty \mathrm{r\, Q_{N;\Delta t}(r) dr}}{\int_0^\infty \mathrm{Q_{N;\Delta t}(r) dr}} = \frac{\mathrm{\Delta t^{N+1}}}{\mathrm{N!}} \int_0^\infty \mathrm{r^{N+1} e^{-r\,\Delta t} dr} \tag{B4}$$

but $\int_0^\infty x^{\mathrm{N+1}} \mathrm{e^{-k\,x}dx} = \frac{\mathrm{(n+1)!}}{\mathrm{k^{N+2}}}$, therefore:

$$\mathrm{r_{mean}(N, \Delta t)} = \frac{\mathrm{\Delta t^{N+1}}}{\mathrm{N!}} \frac{\mathrm{(N+1)!}}{\mathrm{\Delta t^{N+2}}} = \frac{\mathrm{N+1}}{\mathrm{\Delta t}} \tag{B5}$$

A.3.0 CL analytical solution

By substituting the expression of $\mathrm{Q_{N,\,\Delta t}\,(r)}$ in Equation 5.5, the the confidence level condition in Equation A.6 is:

$$\frac{\mathrm{\Delta t^{N+1}}}{\mathrm{N!}} \int_{\mathrm{r_{min}}}^{\mathrm{r_{max}}} \mathrm{(r\Delta t)^n\, e^{-r\,\Delta t} dr = CL} \tag{C1}$$

but $\int \mathrm{(xk)^N\, e^{-xk} dk} = -\frac{\mathrm{e^{-xk}}}{\mathrm{k}} \sum_{\mathrm{l=0}}^{\mathrm{N}} \frac{\mathrm{N!}}{\mathrm{l!}} \mathrm{(xk)^l}$, so I can write the previous equation as:

$$\left[\mathrm{e^{-r\,\Delta t}} \sum_{\mathrm{l=0}}^{\mathrm{N}} \frac{\mathrm{(r\,\Delta t)^l}}{\mathrm{l!}}\right]_{\mathrm{r_{max}}}^{\mathrm{r_{min}}} = \mathrm{CL} \tag{C2}$$

By considering the probability condition in Equation A.7 I obtain:

$$\frac{\mathrm{\Delta t}}{\mathrm{N!}} \mathrm{(r_{min}\,\Delta t)^N\, e^{-r_{min}\,\Delta t}} = \frac{\mathrm{\Delta t}}{\mathrm{N!}} \mathrm{(r_{max}\,\Delta t)^N\, e^{-r_{max}\,\Delta t}} \tag{C3}$$

therefore:

$$\left(\frac{\mathrm{r_{max}}}{\mathrm{r_{min}}}\right)^N = \mathrm{e^{(r_{max} - r_{min})\Delta t}} \tag{C4}$$

A.3.1 Numerical solution

This 2-equation and 2-variable system can be numerically solved. By defining $x = (r_{max} - r_{min}) \cdot \Delta\mathrm{t}$ and considering the condition above.

$$\mathrm{F_2(x, y(x,n), z(x,n), n) = exp(-y)} \left[\sum_{\mathrm{l=0}}^{\mathrm{N}} \frac{\mathrm{y^l}}{\mathrm{l!}}\right] - \mathrm{exp(yz)} \left[\sum_{\mathrm{l=0}}^{\mathrm{N}} \frac{\mathrm{(yz)^l}}{\mathrm{l!}}\right] - \mathrm{CL} \tag{C5}$$

where $\mathrm{y(x,n)}$ is:

$$\mathrm{y(x,n)} = \frac{\mathrm{x}}{\mathrm{exp(\frac{x}{n}) - 1}} \tag{C6}$$

and $\mathrm{z(x,n)}$ is:

$$\mathrm{z(x,n)} = \mathrm{exp}\left(\frac{\mathrm{x}}{\mathrm{n}}\right) \tag{C7}$$

By imposing $F_2(x, y(x, n), z(x, n), n) = 0$, I can find a numerical solution by solving this equation in the variable x. This yields, as a result, a x that depends on the number of events considered. The x_n corresponds to the $x_{N,CL}$ that satisfies the confidence level conditions for a certain number of counts in a given Δt. At this point, I can express the relative confidence interval $[\epsilon^{-}(N, CL)\ ,\ \epsilon^{+}(N, CL)]$ as

$$\epsilon^{-}(N, CL) = \epsilon^{-}(x_{N,cl}) = \frac{\frac{x_{N,CL}}{N}}{\exp\frac{x_{N,cl}}{N} - 1} \tag{C8}$$

$$\epsilon^{+}(N, CL) = \epsilon^{+}(x_{N,cl}) = \frac{\frac{x_{N,CL}}{N} \cdot \exp\frac{x_{N,CL}}{n}}{\exp\frac{x_{N,CL}}{n} - 1} \tag{C9}$$

Therefore the absolute confidence interval $[r^{-}(N, CL, \Delta t)\ ,\ r^{+}(N, CL, \Delta t)]$

$$r^{-}(N, CL, \Delta t) = \epsilon^{-}(N, CL) \cdot r_{mod}(N, \Delta t) \tag{C10}$$

$$r^{+}(N, CL, \Delta t) = \epsilon^{+}(N, CL) \cdot r_{mod}(N, \Delta t) \tag{C11}$$

A.4.0 Generalized inversion method analytical solution

The inversion method integral in Equation 5.12 can be considered as a trapezoidal integral when the rate curve r(t) is a continuous piecewise linear function. By looking at Figure 5.8 the integral can be rewritten as:

$$\begin{aligned}&\frac{r(T\_SIM[N]) + r(T\_SIM[N-1])}{2} \cdot \\ &\cdot (T\_SIM[N] - T\_SIM[N-1]) = -\ln\{1 - RND(0, 1)\}\end{aligned} \tag{D1}$$

In the most general case, $T\_SIM[N-1]$ is between two rate points of r(t) as in Figure 5.8. The rate $r(T\_SIM[N-1])$ can be linearly extrapolated by considering the intensities of the two rate points r_1 and r_2 that are before and after $T\_SIM[N-1]$), as well as their respective associated times t_1 and t_2:

$$\begin{gathered}m = \frac{r_2 - r_1}{t_2 - t_1} \\ r(T\_SIM[N-1]) = r_1 + m \cdot (T\_SIM[N-1] - t_1)\end{gathered} \tag{D2}$$

The same procedure can be performed for the unknown T_SIM[N] from the T_SIM[N − 1] where $r(T\_SIM[N-1])$ is known. Let us define T_SIM[N − 1] as as $\bar{t}$ and T_SIM[N] as x:

$$\begin{aligned} \mathrm{r(x)} &= \mathrm{r_1 + m \cdot (x - t_1)} \\ \frac{\mathrm{r(\bar{t}) + r(x)}}{2} \cdot \mathrm{(x - \bar{t})} &= -\ln\{1 - \mathrm{RND}(0,1)\} \\ \frac{\mathrm{r(\bar{t}) + r_1 + mx - mt_1}}{2} \cdot \mathrm{(x - \bar{t})} &= -\ln\{1 - \mathrm{RND}(0,1)\} \end{aligned} \tag{D3}$$

The equation can then be rewritten by rearranging the random terms with ZETA $\equiv -\ln\{1 - \mathrm{RND}(0,1)\}$:

$$\mathrm{mx^2 + x(r(\bar{t}) + r_1 - mt_1 - m\bar{t}) + mt_1\bar{t} - \bar{t}r(\bar{t}) - \bar{t}r_1 - 2ZETA = 0} \tag{D4}$$

The only possible solution is therefore when $\mathrm{x} > \bar{\mathrm{t}}$:

$$\begin{aligned} B &= \mathrm{r(\bar{t}) + r_1 - mt_1 - m\bar{t}} \\ C &= \mathrm{mt_1\bar{t} - \bar{t}r(\bar{t}) - \bar{t}r_1 - 2ZETA} \\ \mathrm{x} &= \frac{\mathrm{-B + \sqrt{B^2 - 4mC}}}{\mathrm{2m}} \end{aligned} \tag{D5}$$

A.5.0 CCF examples

Figure A.1 from pages 12 to 16 shows the GRB considered in the comparison in Figure 5.12 between DP and MDP methods.

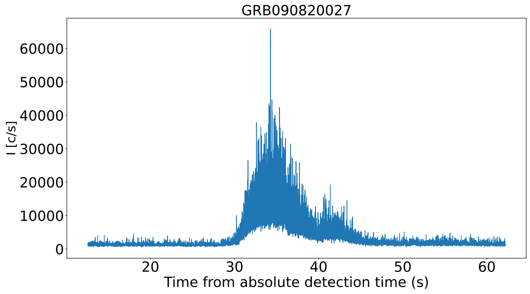

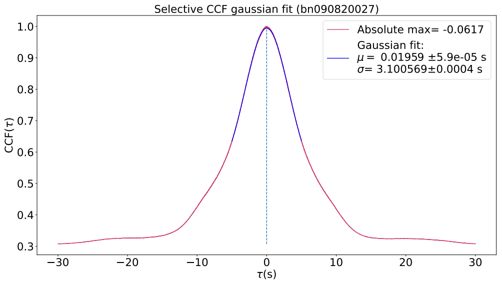

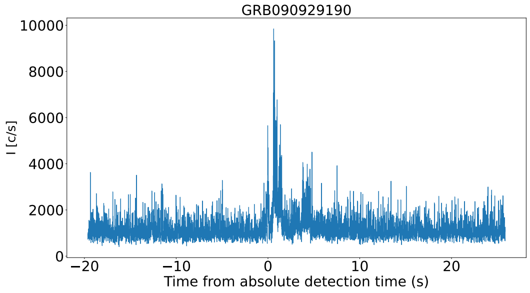

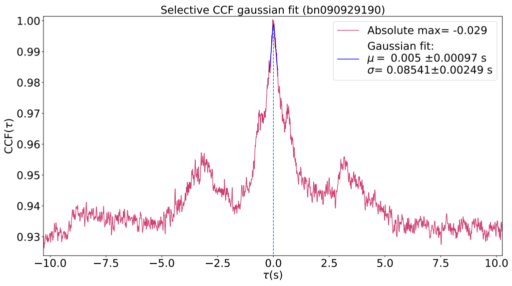

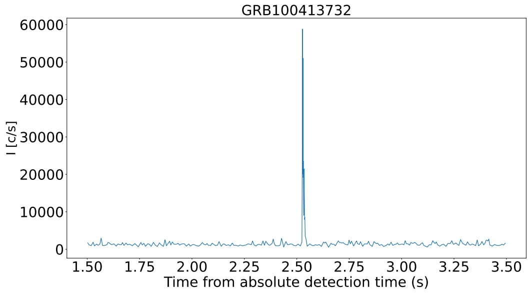

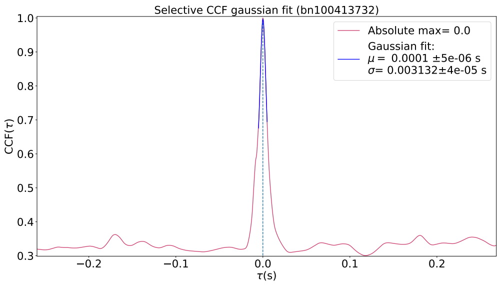

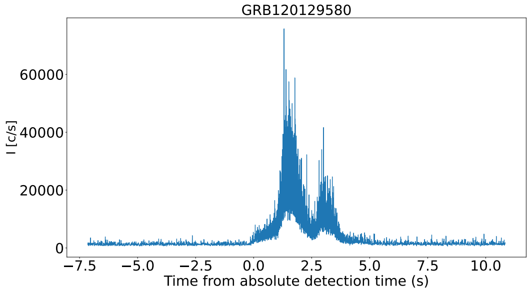

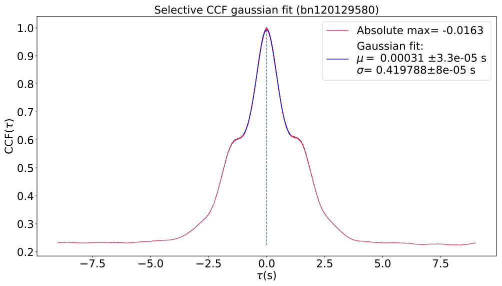

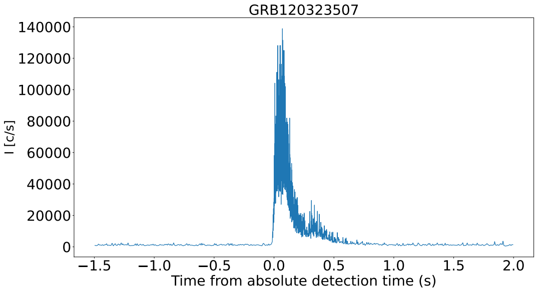
GRB120323507
I [c/s]
140000
120000
100000
80000
60000
40000
20000
0
−1.5
−1.0
−0.5
0.0
0.5
1.0
1.5
2.0
Time from absolute detection time (s)

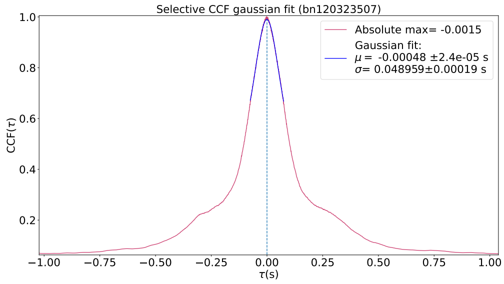
Selective CCF gaussian fit (bn120323507)
Absolute max= -0.0015
Gaussian fit:
μ = -0.00048 ±2.4e-05 s
σ= 0.048959±0.00019 s
CCF(τ)
1.0
0.8
0.6
0.4
0.2
−1.00
−0.75
−0.50
−0.25
0.00
0.25
0.50
0.75
1.00
τ(s)

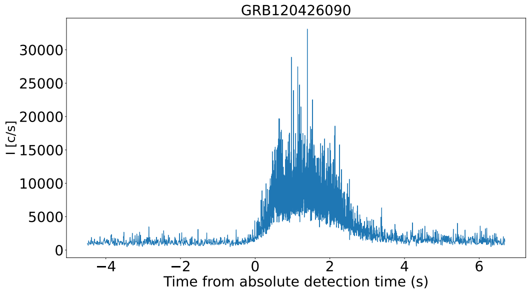
GRB120426090
I [c/s]
30000
25000
20000
15000
10000
5000
0
−4
−2
0
2
4
6
Time from absolute detection time (s)

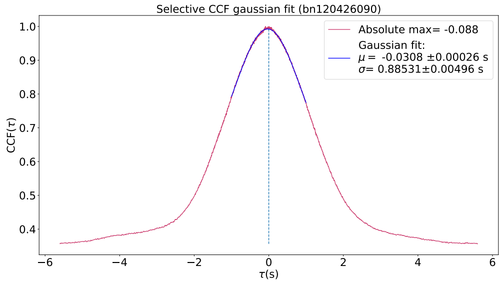
Selective CCF gaussian fit (bn120426090)
Absolute max= -0.088
Gaussian fit:
μ = -0.0308 ±0.00026 s
σ= 0.88531±0.00496 s
CCF(τ)
1.0
0.9
0.8
0.7
0.6
0.5
0.4
−6
−4
−2
0
2
4
6
τ(s)

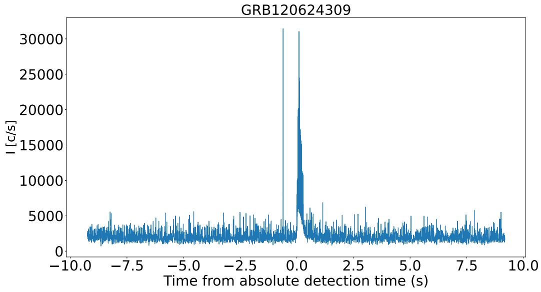
GRB120624309
I [c/s]
30000
25000
20000
15000
10000
5000
0
−10.0
−7.5
−5.0
−2.5
0.0
2.5
5.0
7.5
10.0
Time from absolute detection time (s)

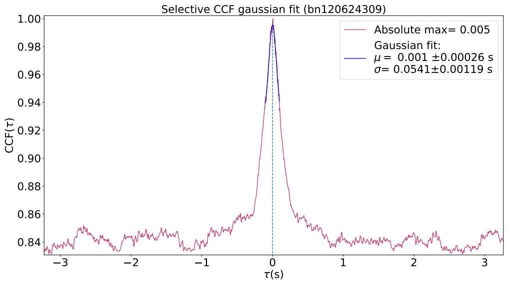
Selective CCF gaussian fit (bn120624309)
Absolute max= 0.005
Gaussian fit:
μ = 0.001 ±0.00026 s
σ= 0.0541±0.00119 s
CCF(τ)
1.00
0.98
0.96
0.94
0.92
0.90
0.88
0.86
0.84
−3
−2
−1
0
1
2
3
τ(s)

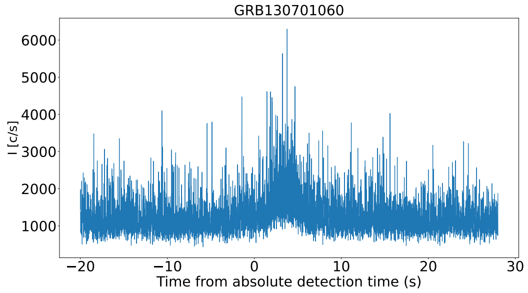
GRB130701060
I [c/s]
6000
5000
4000
3000
2000
1000
−20
−10
0
10
20
30
Time from absolute detection time (s)

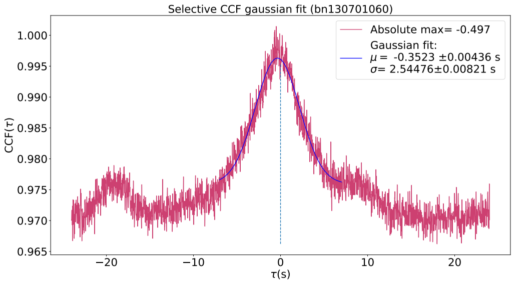
Selective CCF gaussian fit (bn130701060)
Absolute max= -0.497
Gaussian fit:
μ = -0.3523 ±0.00436 s
σ= 2.54476±0.00821 s
CCF(τ)
1.000
0.995
0.990
0.985
0.980
0.975
0.970
0.965
−20
−10
0
10
20
τ(s)

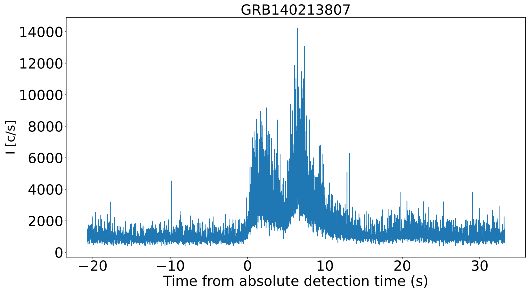
GRB140213807
I [c/s]
14000
12000
10000
8000
6000
4000
2000
0
−20
−10
0
10
20
30
Time from absolute detection time (s)

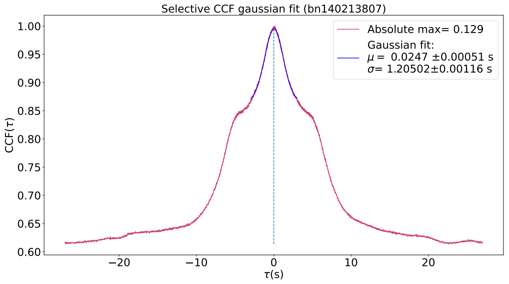
Selective CCF gaussian fit (bn140213807)
Absolute max= 0.129
Gaussian fit:
μ = 0.0247 ±0.00051 s
σ= 1.20502±0.00116 s
CCF(τ)
1.00
0.95
0.90
0.85
0.80
0.75
0.70
0.65
0.60
−20
−10
0
10
20
τ(s)

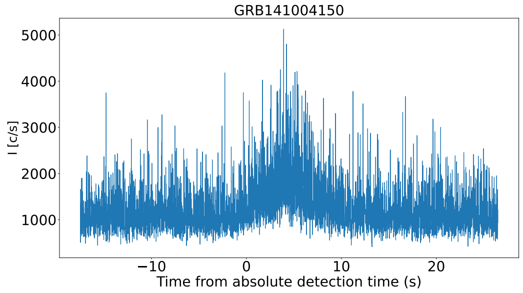
GRB141004150
I [c/s]
5000
4000
3000
2000
1000
−10
0
10
20
Time from absolute detection time (s)

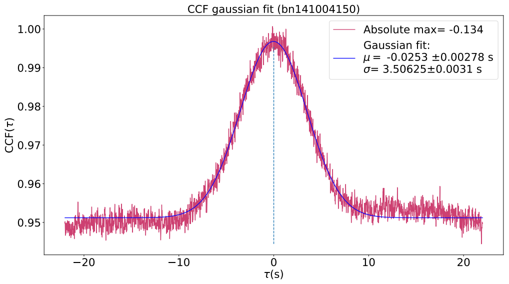
CCF gaussian fit (bn141004150)
Absolute max= -0.134
Gaussian fit:
μ = -0.0253 ±0.00278 s
σ= 3.50625±0.0031 s
CCF(τ)
1.00
0.99
0.98
0.97
0.96
0.95
−20
−10
0
10
20
τ(s)

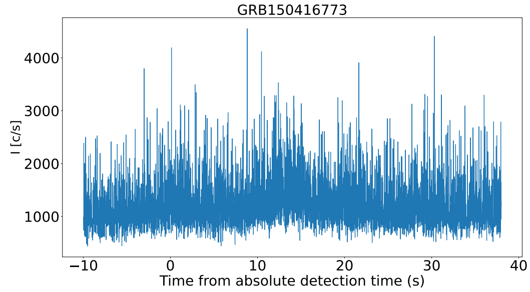
GRB150416773
I [c/s]
1000 2000 3000 4000
−10 0 10 20 30 40
Time from absolute detection time (s)

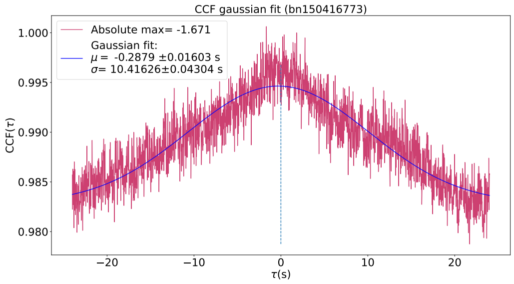
CCF gaussian fit (bn150416773)
Absolute max= -1.671
Gaussian fit:
μ = -0.2879 ±0.01603 s
σ= 10.41626±0.04304 s
CCF(τ)
0.980 0.985 0.990 0.995 1.000
−20 −10 0 10 20
τ(s)

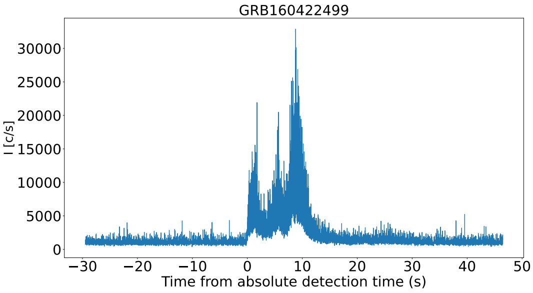
GRB160422499
I [c/s]
0 5000 10000 15000 20000 25000 30000
−30 −20 −10 0 10 20 30 40 50
Time from absolute detection time (s)

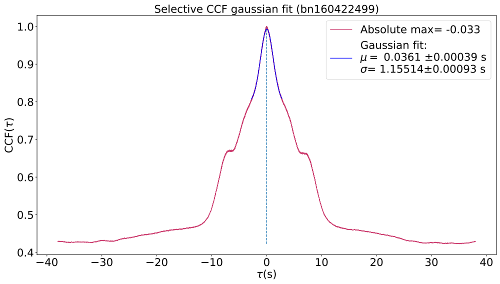
Selective CCF gaussian fit (bn160422499)
Absolute max= -0.033
Gaussian fit:
μ = 0.0361 ±0.00039 s
σ= 1.15514±0.00093 s
CCF(τ)
0.4 0.5 0.6 0.7 0.8 0.9 1.0
−40 −30 −20 −10 0 10 20 30 40
τ(s)

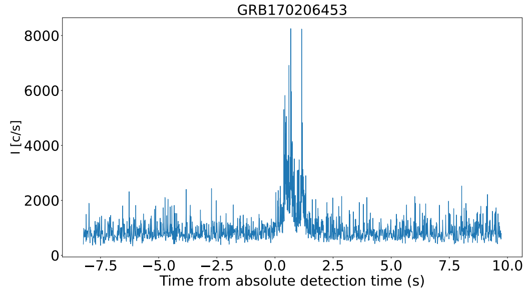
GRB170206453
I [c/s]
0 2000 4000 6000 8000
−7.5 −5.0 −2.5 0.0 2.5 5.0 7.5 10.0
Time from absolute detection time (s)

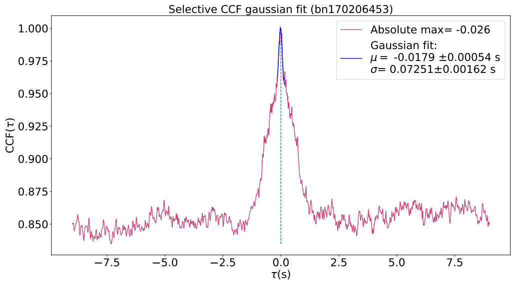
Selective CCF gaussian fit (bn170206453)
Absolute max= -0.026
Gaussian fit:
μ = -0.0179 ±0.00054 s
σ= 0.07251±0.00162 s
CCF(τ)
0.850 0.875 0.900 0.925 0.950 0.975 1.000
−7.5 −5.0 −2.5 0.0 2.5 5.0 7.5
τ(s)

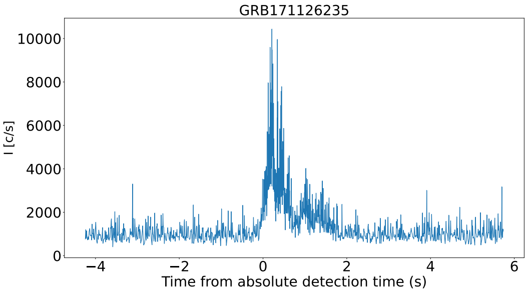
GRB171126235
I [c/s]
0 2000 4000 6000 8000 10000
−4 −2 0 2 4 6
Time from absolute detection time (s)

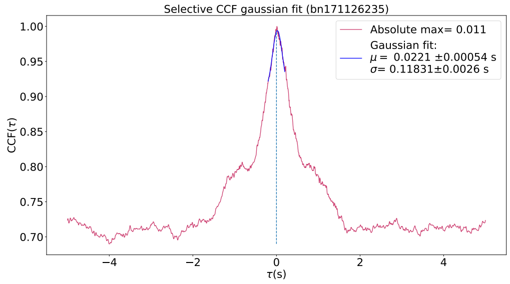
Selective CCF gaussian fit (bn171126235)
Absolute max= 0.011
Gaussian fit:
μ = 0.0221 ±0.00054 s
σ= 0.11831±0.0026 s
CCF(τ)
0.70 0.75 0.80 0.85 0.90 0.95 1.00
−4 −2 0 2 4
τ(s)

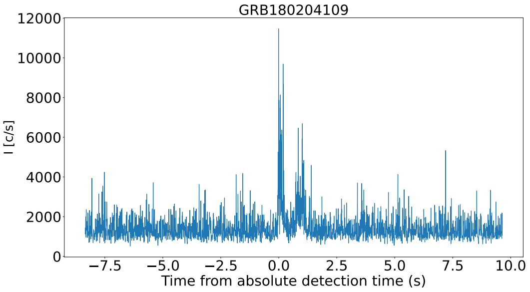
GRB180204109
I [c/s]
0 2000 4000 6000 8000 10000 12000
−7.5 −5.0 −2.5 0.0 2.5 5.0 7.5 10.0
Time from absolute detection time (s)

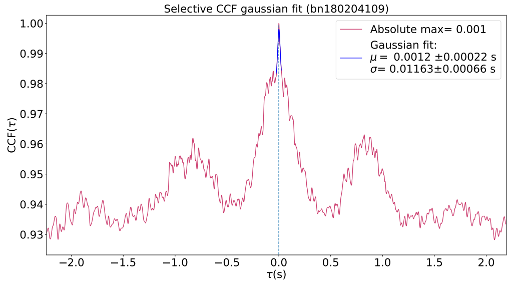
Selective CCF gaussian fit (bn180204109)
Absolute max= 0.001
Gaussian fit:
μ = 0.0012 ±0.00022 s
σ= 0.01163±0.00066 s
CCF(τ)
0.93 0.94 0.95 0.96 0.97 0.98 0.99 1.00
−2.0 −1.5 −1.0 −0.5 0.0 0.5 1.0 1.5 2.0
τ(s)

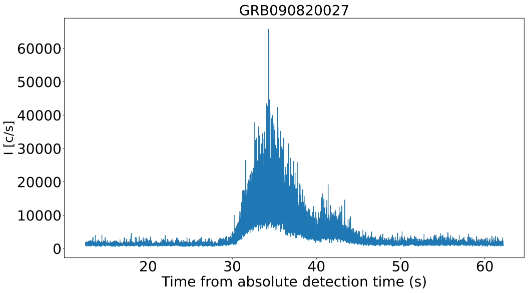
GRB090820027
I [c/s]
0 10000 20000 30000 40000 50000 60000
20 30 40 50 60
Time from absolute detection time (s)

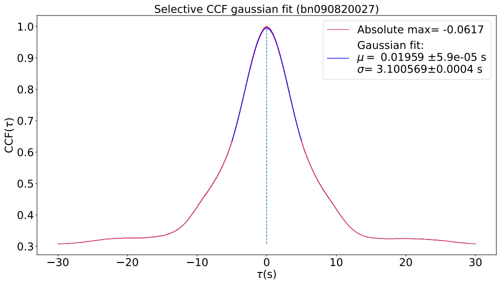
Selective CCF gaussian fit (bn090820027)
Absolute max= -0.0617
Gaussian fit:
μ = 0.01959 ±5.9e-05 s
σ= 3.100569±0.0004 s
CCF(τ)
0.3 0.4 0.5 0.6 0.7 0.8 0.9 1.0
−30 −20 −10 0 10 20 30
τ(s)

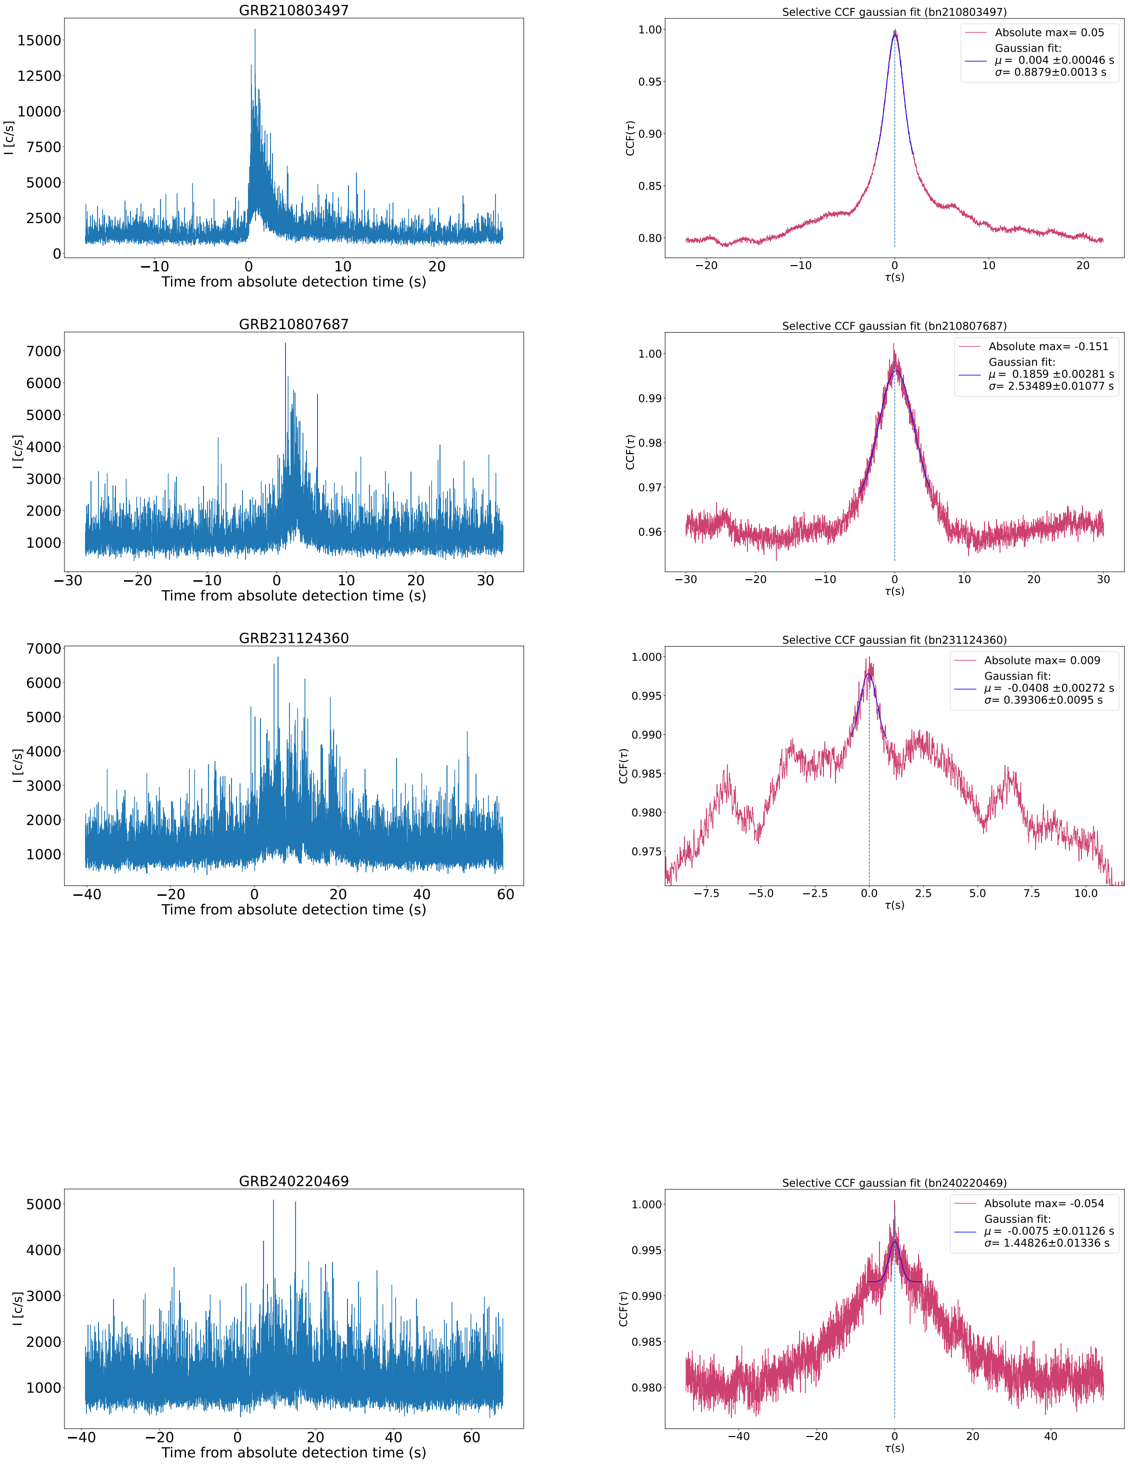

Figure A.1: Right panel: Example CCFs performed between ToA lists (see subsection 5.2.4) obtained via the MDP method (see subsection 5.4.3). Gaussian fit parameters are highlighted in each plot and fixed for both the MDP and the DP methods testing (see section 5.5). ToA lists are retrieved from GRB data as observed by the brightest Fermi-GBM detector monitoring the bursts. Left panel: Light curves from the brightest detector, computed using an adaptive bin size of 10 photons per bin.

B

Appendix

B.1.0 Crab Pulsar Exposure with HERMES: Analytical Estimate with MC Validation

In general, determining the exposure time required to detect a neutron star pulsed profile is not a straightforward task. The calculations presented in this appendix show, both analytically and through simulations, how to assess an $n\sigma$-level detection as a function of a given exposure time, $T_{\rm exp}$.

I adopt as a template the folded *Rossi X-Ray Timing Explorer* (RXTE, Swank 1999) pulse profile of the Crab pulsar reported by Jain & Paul (2011). Assuming an on–axis pointing, I compute the exposure time required for the two peaks to reach the desired statistical significance.

By considering the detector response files, I simulated the count spectrum expected for *HERMES* with the `fakeit` task in XSPEC, modeling it as the superposition of three components: (i) the Crab Nebula, whose spectral parameters were taken from Kirsch et al. (2005); (ii) the pulsed emission of the Crab pulsar, described by the model in Weisskopf et al. (2011); and (iii) the diffuse X-ray background.

The Cosmic X-Ray Background (CXB) is derived from the diffused CXB model proposed in Gruber et al. (1999b). I assume a low energy HERMES/SpIRIT FoV of 1.57 steradians as in Ghirlanda et al. (2024), for the 3–11.5 keV energy band.

In the current configuration of the HERMES instrument onboard SpIRIT, the nominal energy lower cut threshold is set at a level equivalent to about 5 keV. However, the precise energy threshold for each one of the active channels may vary, due to the channel-to-channel process variations and the working principle of the discrimination circuit Evangelista et al. (2020b). This leads to a significant threshold spread, and a certain sensitivity also in the 3–5 keV band. Accordingly, throughout I adopted a lower-energy cut of 3 keV.

An appropriate set of HERMES Ancillary Response Files (*arf*) and Redistribution Matrix Files (*rmf*) (Campana et al. 2020) has been built to take this effect into account, and was used to evaluate the average count rate in the 3–11.5 keV energy range for the Crab Nebula+Pulsar ($C_{\rm Crab}$) and the CXB ($C_{\rm BKG}$).

B.1.1 Analytical determination of the Crab exposure time

Because of platform constraints, the dominant practical limitation is the on–axis exposure that can be accumulated overall. In this section I show the analytical treatment to derive a relation linking the expected χ^2 of the Crab pulse profile to the observing time T_{exp}.

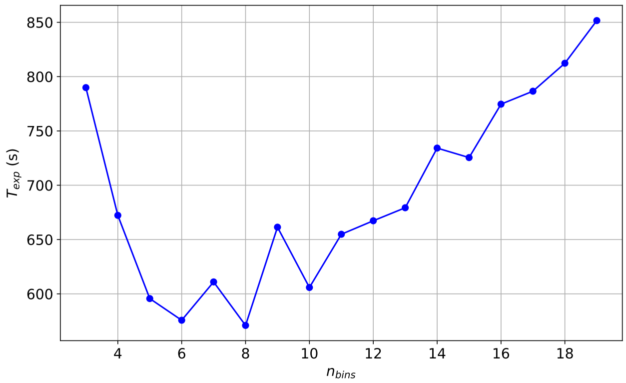

Figure B.1: Required exposure time to obtain a χ^2 threshold for a 3σ level detection as function of n_{bins}. Although analytically derived, the observed scatter in Figure B.1 arises because the solution is applied to an observed profile with finite counts and is therefore subject to statistical (Poisson) fluctuations.

section B.2 shows the analytical solution to the problem and, by solving the Equation D9 for T_{exp}, I find that the χ^2 value can be expressed as:

$$\chi^2_{\mathrm{th}} = T_{\mathrm{exp}} \cdot C_{\mathrm{Crab}} \cdot \sum_{i=1}^{n_{\mathrm{bins}}} \frac{(\mathrm{FRAC}_i - \Delta\phi)^2}{\mathrm{FRAC}_i + \Delta\phi \cdot \left(\frac{C_{\mathrm{BKG}}}{C_{\mathrm{Crab}}}\right)}, \tag{D1}$$

where $FRAC_i$ is the per–bin pulsed fraction that maximizes the pulse significance, as defined in section B.3 and Equation D10, and $\Delta\phi$ is the dimension of the phase bin.

I now assume, as null hypotheses, a flat signal in the data (H_0 corresponds to no Crab pulsation detection). The probability that, for a given threshold $\chi^2_{\mathrm{th,dof}}$, the observed χ^2 is greater, comes from a well-known function $\mathrm{P}(\chi^2 \geq \chi^2_{\mathrm{th,dof}})$, for a certain number of degrees of freedom dof.

Equation D1 can be inverted to derive the T_{exp} necessary to get the desired $\chi^2_{\mathrm{th,dof}}$ for a given number of degrees of n_{bins}. I compute the $\chi^2_{\mathrm{th,dof}}$ to reach a 3σ significance level for each bin, and obtain the trend shown in Figure B.1.

B.1.2 The Monte Carlo validation

The previous discussion is associated with the χ^2 threshold and so with the probability $\mathrm{P}(\chi^2 \geq \chi^2_{\mathrm{th,dof}})$. This establishes the required exposure time to reach a minimum significance limit.

In order to verify the accuracy of the previous solution and the expected χ^2 distribution as a function of a given T_{exp}, I perform a simulation-based analysis starting from a theoretical profile with $n_{\mathrm{bins}} = 15$ as shown in Figure B.3.

I adopt a fixed pulse shape template and denote its amplitude by I_p (in counts) as the sum of the CRAB and background contributions:

$$I_p = I_{\text{ratio}} \cdot T_{\text{exp}} \cdot C_{\text{Crab}} + \frac{1}{n_{\text{bins}}} \cdot T_{\text{exp}} \cdot C_{\text{BKG}}, \tag{D2}$$

where I_{ratio} is the ratio of the counts in every bin of the RXTE profile over the total counts as shown in Figure B.3. When using the Equation D2, I extract the value of I_p with a Poisson distribution, in order to achieve a series of Poisson realizations of the Crab pulse profile at a given exposure time.

According to the observation exposure, in Equation D2, I adopt $T_{\text{exp}} = 730$ s and for each Poisson realization I evaluate the χ^2 with respect to the average of the pulse profile.

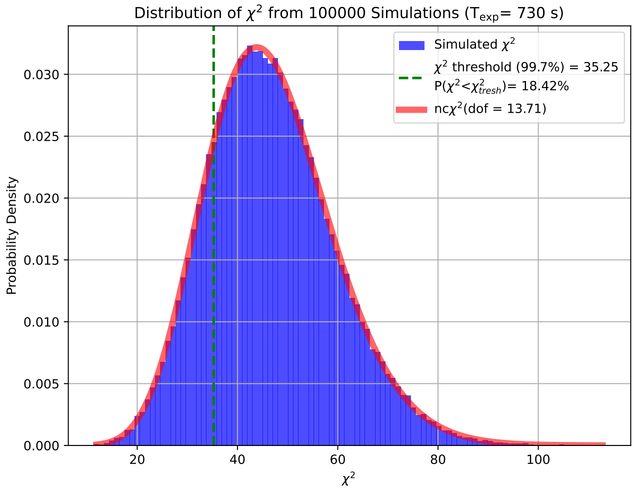

Figure B.2: Monte Carlo distribution of the χ^2 statistic for simulated Crab-like pulse profiles with $n_{\text{bins}} = 15$. The histogram is fitted with the probability density in Equation D3. The vertical green line marks the 3σ detection threshold: for dof = 14 I adopt the threshold $\chi^2 = 35.25$, i.e. the 99.7% upper-tail critical value of the central χ^2 distribution.

B.2.0 The analytical solution

The Monte Carlo distribution is shown in Figure B.2 is fitted with a non-central χ^2 (ncχ^2) distribution:

$$f(x, dof, \lambda) = \frac{1}{2} \exp\left(-\frac{\lambda + x}{2}\right) \left(\frac{x}{\lambda}\right)^{(dof-2)/4} L_{(dof-2)/2}\left(\sqrt{\lambda x}\right), \tag{D3}$$

where λ is the non-centrality parameter. $L_{(k-2)/2}$ denotes the modified Bessel function of first order of degree and x the χ^2. The parameters from the ncχ^2 fit match expectations, and the dof are in line with the theoretical value $n_{\text{bins}} - 1 = 15 - 1 = 15$.

Therefore, I expect that in ≈ 81.58% of cases, a 730 s on-axis HERMES observation grants at least a 3σ detection.

The total detected count rate (in counts/s) is given by the sum of the Crab emission (nebula + pulsar) C_{Crab} and the background in the HERMES FoV C_{BKG}:

$$C_{\mathrm{tot}} = C_{\mathrm{Crab}} + C_{\mathrm{BKG}} \tag{D4}$$

If I divide the pulse profile into n_{bins} bins, as in Figure B.3, the total rate in each phase bin ϕ_i is:

$$C_{\mathrm{tot}}(\phi_i) = \Delta\phi \cdot C_{\mathrm{BKG}} + \mathrm{FRAC}_i \cdot C_{\mathrm{CRAB}} \quad \text{(counts/s)}, \tag{D5}$$

where $\Delta\phi$ is the bin width and FRAC_i the count ratio evaluated for each bin as in Figure B.3.

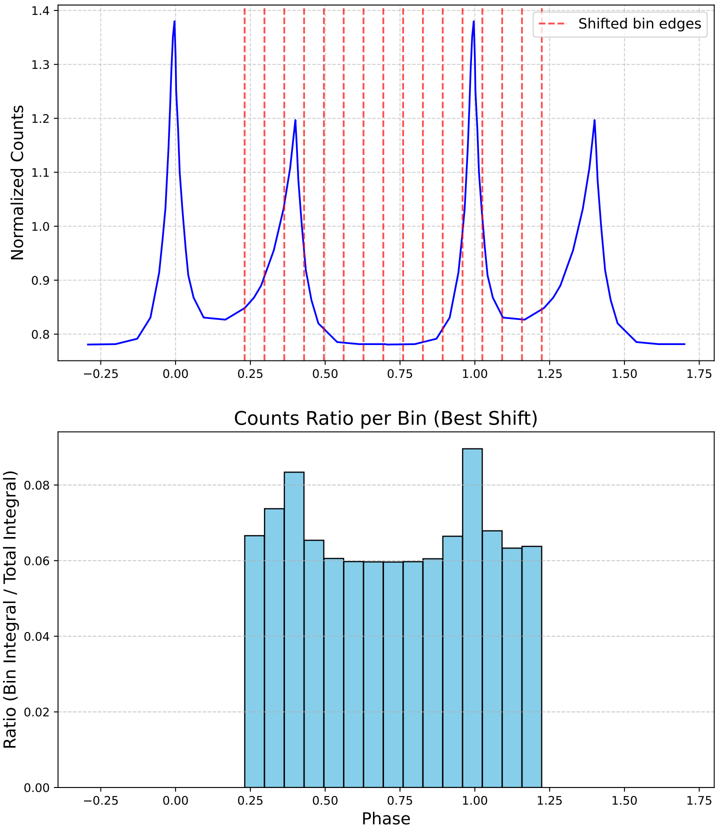

Figure B.3: Upper plot: RXTE Crab Pulsar pulse profile is rebinned to obtain a 15-channel profile, for the given best shift. Lower plot: Normalized count ratio as evaluated in the red bins, with respect to total counts of the profile ($FRAC_i$). In this case, the on-pulse peaks are the 3rd and the 12th; the others are considered as off-pulse bins.

For an observation time T_{exp}, the expected counts in each phase bin are:

$$N_{\mathrm{tot}}(\phi_i) = C_{\mathrm{tot}}(\phi_i) \cdot T_{\mathrm{exp}}, \tag{D6}$$

with an associated Poissonian error of:

$$\sigma_{N_i} = \sqrt{N_{\rm tot}(\phi_i)} \tag{D7}$$

The average total value of the expected counts is computed as:

$$\overline{N_{\rm tot}} = \sum_{i=1}^{n_{\rm bins}} \frac{N_{\rm tot}(\phi_i)}{n_{\rm bins}} \times \Delta\phi, \tag{D8}$$

where $\Delta\phi_{BINS}$ is the peak dimension in unit of phase. I can now compute the theoretical chi-square value as:

$$\chi^2_{\rm th} = \sum_{i=1}^{n_{\rm bins}} \frac{\left(N_{\rm tot}(\phi_i) - \overline{N_{\rm tot}}\right)^2}{N_{\rm tot}(\phi_i)}. \tag{D9}$$

B.3.0 Pulsation significance maximization

To maximize the pulsation significance, I slide the edges of phase bins across the folded profile as in Figure B.3.

For every trial position I compute, in each energy channel, the ratio

$$\mathrm{FRAC}_i \;=\; \frac{N_i}{N_{\rm on} + N_{\rm off}}, \qquad i = 1, \ldots, n_{\rm bins}, \tag{D10}$$

i.e. the ratio of counts falling in each bin window N_i to the total counts on–pulse ($N_{\rm on}$) plus off–pulse ($N_{\rm off}$), as defined in Figure B.3.

The optimal phase shift is found by maximizing the variance

$$\sigma^2 \;=\; \frac{1}{n_{\rm bins}} \sum_{i=1}^{n_{\rm bins}} (1 - \mathrm{FRAC}_i)^2, \tag{D11}$$

which reaches its maximum when the binning is such that two bins are centered on the main emission peaks of the pulse profile.

Applying this optimization for different numbers of channels, in the 3–11.5 keV energy band, I produce a lookup table that records the phase boundaries and the corresponding FRAC_i for each profile represented with a given bin number.

C

Appendix

C.1.0 BKG treating

During the burst analysis, the average background (BKG) was estimated using a third-degree polynomial regression model based on the rate evaluations before and after the burst temporal interval that is reported in the HXMT catalog. The background regression model is exploited to calculate the predictive function, allowing the estimation of background values within the burst interval. Subsequently, the background integral (in the T_{90}) is computed to estimate the BKG counts and subtracted from the total event count during the T_{90}.

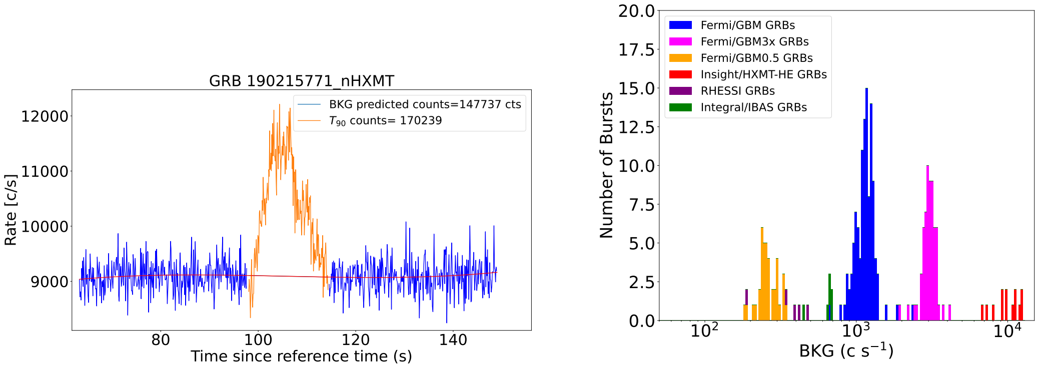

Figure C.1: Left: Example of GRB background and rate extraction. The red line is the background predicted rate during and not the burst region (in light blue). Right: Background estimates as obtained for the considered instruments.

The background-subtracted average count rate R_{burst} was estimated by dividing the number of remaining events, after subtracting the background integral, by the burst duration T_{90}:

$$\mathrm{R_{burst}} = \frac{\mathrm{C_{T_{90}}} - \mathrm{C_{BKG}}}{\mathrm{T_{90}}}, \qquad \text{(D1)}$$

where $C_{T_{90}}$ are the total counts during the T_{90} and C_{BKG} the predicted counts in this window.

The uncertainty on R_{burst} is calculated by assuming the Poissonian statistic of the counts (i.e., $\sigma_{C_{BKG}} = \sqrt{C_{BKG}}$) and propagating the T_{90} error on the total error.

D

Appendix

D.1.0 THESEUS forecast

To quantify how the constraints would scale with an enlarged GRB sample, I generate a synthetic population of GRBs, as observed by the future THESEUS mission. Each GRB provides three measurable delays, in three $\bar{E}_{\rm ch}$ channel, relative to the 10–100 keV reference band. Each GRB redshift is drawn from the cumulative distribution of GRBs with known redshift, as reported in Bozzo et al. (2024), by considering the $z \in [0.11, 4.35]$ range. Delays are generated, with the Equation 2.17, by taking into account a Quantum Gravity effective energy equal to the Planck Energy (i.e., $\alpha = 1$).

To forecast the THESEUS timing accuracy (i.e., the precision associated with the total observed delay $\sigma_{\tau_{total}}$), I rescale the observed Fermi/LAT-LLE trend in Figure D.1, letting the $\sigma_{\tau_{total}}$ improve as the square root of the detected photons. Two factors lead to this effect: (i) the ratio between the Fermi/LLE-LAT (Pelassa et al. 2010) and the THESEUS effective area (Amati et al. 2021), and (ii) the spectral energy distribution due to the high-energy power-law index β:

$$F(E, \beta) \propto E^{\beta}. \tag{D1}$$

I did not apply any additional cosmological correction to the delay accuracy in the *THESEUS* synthetic sample since the spread in the Fermi/LLE-LAT uncertainties already incorporates all processes that modulate the detected photon counts, including the GRB luminosity–redshift evolution (see, e.g., Pescalli et al. (2016)). Therefore, I let the *THESEUS* accuracies in Figure D.1 to lie within the 99.7% confidence interval obtained by rescaling the Fermi dispersion, according to spectral flux modulation in Equation D1. Since the synthetic GRB population is constructed to span the same redshift range as the Fermi dataset, redshift–driven variations in timing accuracy, encoded in the observed dispersion, are automatically propagated into the *THESEUS* forecast.

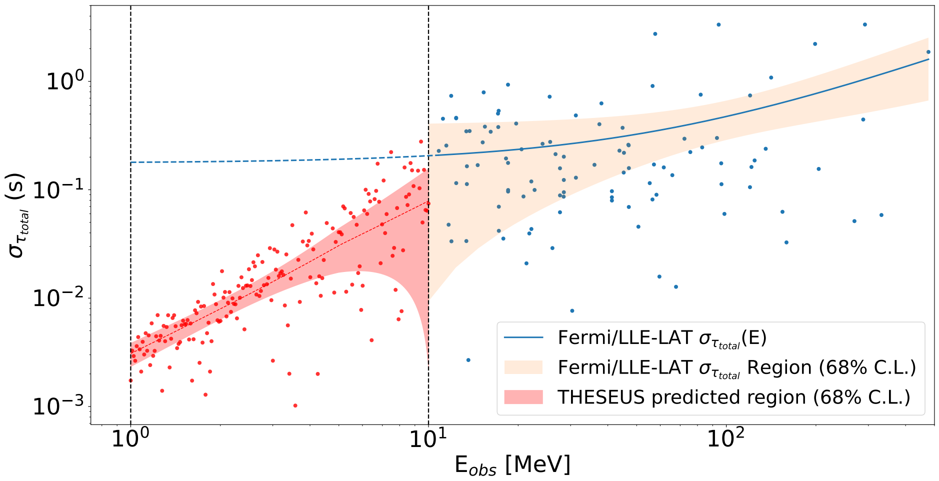

Figure D.1: Accuracy of the total observed delay ($\sigma_{\tau_{total}}$) as a function of the observed energy (E_{obs}), shown in log–log scale. The blue line represents the best-fit linear regression of the delay accuracy derived from the Fermi/LLE–LAT dataset. The shaded region indicates the 68% confidence region obtained from the covariance matrix of the regression parameters. The red band shows the predicted *THESEUS*/XGIS precision in the 1–10 MeV range (delimited by the vertical black lines), scaled from the observed uncertainties. A representative sample of 200 synthetic delay accuracies within the *THESEUS* prediction range is also displayed.

Acronym	Meaning
GRB	Gamma-Ray Burst
VHE	Very High Energy
NS	Neutron Star
BNS	Binary Neutron Star
BH	Black Hole
RRAT	Rotating Radio Transient
MSP	Milli-Second Pulsar
SED	Spectral Energy Distribution
IC	Inverse Compton
SSC	Synchrotron Self Compton
XMM-Newton	X-ray Multi-Mirror Mission - Newton
BATSE	Burst and Transient Source Experiment
CGRO	Compton Gamma Ray Observatory
EGRET	Energetic Gamma-Ray Experiment Telescope
INTEGRAL	INTErnational Gamma-Ray Astrophysics Laboratory
Swift	Neil Gehrels Swift Observatory
Swift/BAT	Swift/Burst Alert Telescope
AGILE	Astro-Rivelatore Gamma a Immagini Leggero
MAGIC	Major Atmospheric Gamma Imaging Cherenkov
HXMT	Hard X-ray Modulation Telescope
GBM	Gamma-ray Burst Monitor
LAT	Large Area Telescope
HERMES -TP/SP	High Energy Rapid Modular Ensemble of Satellites
Technological and Scientific Pathfinder	
HFC	Hermes Full Constellation
WFC	Wide Field Camera
NFI	Narrow Field Instruments
XRT	X-Ray Telescope
SDD	Silicon Drift Detectors
CCD	Charge-Coupled Device
ACD	Anti Coincidence Detector
SP	Single Photons
BKG	Background
PSF	Point Spread Function
FoV	Field of View
ToA	Time of Arrival
TDOA	Time–Difference–of–Arrival
COTS	Commercial Off–The–Shelf
SAA	South Atlantic Anomaly

PB	Polar Belt
EM	Electro-Magnetic
GW	Gravitational Wave
LI	Lorentz Invariance
LIV	Lorentz Invariance Violation
EoS	Equation of State
C.L.	Confidence Level
ASI	Italian Space Agency
MUR	Ministry of University and Research
LEO	Low Earth Orbit