ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DOTTORATO DI RICERCA IN ASTROFISICA

Ciclo XXXVIII

Settore concorsuale: 02/C1 - ASTRONOMIA, ASTROFISICA, FISICA DELLA TERRA E DEI PIANETI

Settore Scientifico Disciplinare: FIS/05 - ASTRONOMIA E ASTROFISICA

Enhancing the simulation accuracy of the COSI Anticoincidence System: calibration, benchmarking, and detector effects prediction

Presentata da: Alex Ciabattoni

Coordinatore Dottorato
Prof. Andrea Miglio

Supervisore
Dr.ssa Valentina Fioretti

Co-supervisore
Prof. Cristian Vignali

Esame finale anno 2026

Abstract

The Compton Spectrometer and Imager (COSI) is a NASA Small Explorer satellite mission scheduled for launch in 2027. Operating as a Compton telescope, COSI will address the observational gap in the soft gamma-ray band by providing imaging, spectroscopy, and polarization in the $0.2-5$ MeV energy range thanks to its array of germanium detectors (GeDs). The GeDs are surrounded on lateral sides and underneath by 22 active panels of BGO scintillation detectors, acting as an Anticoincidence System (ACS). When an interaction with ionizing radiation occurs, the BGO panels produce light in the optical range, which is collected and converted into an electric signal by silicon photomultipliers (SiPMs) coupled to each BGO crystal. The ACS has the primary goal of shielding the GeDs from background radiation, but can also act as a stand-alone detector contributing to the detection and localization of gamma-ray transient sources, such as Gamma-Ray Bursts (GRBs). An on-board trigger system monitors the ACS ratemeters, sending prompt notifications and data telemetry to the ground for fast identification and communication of burst signals, with a required latency of 60 minutes.

Monte Carlo simulations using the Geant4 toolkit are widely used to predict the performance of high-energy space missions and drive the instrumentation design. Alongside simulations, calibration measurements using radioactive sources are commonly employed to refine the simulation parameters and guarantee that the simulations accurately reflect real data. The performance of the COSI ACS is closely linked to the scintillation processes occurring within the BGO panels and the electronics of the SiPM readout. The average path length that optical photons must travel to reach the SiPMs, and their corresponding probability of being detected before absorption in the BGO, depends on the location of the gamma-ray interaction within the crystal. Statistical fluctuations in the number of detected photons, their spatial dispersion across the BGO surface, and the SiPM electronic noise collectively determine the ACS energy resolution. Moreover, if the detection probability of the optical photons varies within the BGO, the energy threshold (defined by the minimum detectable number of optical photons) is also expected to be spatially non-uniform. These characteristics in the ACS performance are difficult to fully capture through laboratory measurements alone.

Although Geant4 was originally designed for high-energy experiments, it has been extended to simulate low-energy processes in the optical regime, including scintillation physics. However, the complexity of the optical interactions and the large number of the produced scintillation photons make it computationally prohibitive to include scintillation in every COSI simulation. To overcome these limitations, we propose the use of dedicated Geant4 optical simulations, properly verified and benchmarked against ACS calibration measurements, to encode the effects of scintillation physics into a correction matrix. This approach enables position-dependent corrections to the ACS energy response, bringing simulations into closer agreement with real measurements. Thanks to the flexibility of simulations, we can compute the correction with sufficiently high spatial resolution (~ 1 cm^2) and investigate energy ranges that are otherwise difficult to access through calibration measurements. Developing the correction matrix requires extensive optical simulations with sufficient statistics. In this perspective, we developed optimized Geant4 simulations that exploit High-Performance Computing (HPC) techniques, parallelizing execution across multiple computing resources to achieve the necessary efficiency.

First, we enhanced and expanded the simulation framework BoGEMMS-HPC, based on the Geant4 toolkit, to enable efficient parallel execution in HPC environments. Using up to 112 cores on pre-exascale Tier-0 supercomputer Leonardo, we achieved a speed-up of about 70 compared to serial execution. We then verified the Geant4-simulated transmission, reflection, and absorption probability of optical photons against theoretical models and reproduced an experimental setup involving a block of CsI(Tl) coupled with a photomultiplier tube to investigate the optical parameters and achieve the best agreement with laboratory measurements. After this verification of the Geant4 optical simulation, we performed calibration measurements of a lateral ACS wall to characterize its energy resolution and spatial response. The measured energy resolution is approximately 60% at 60 keV, decreases to $\sim 16\%$ at 511 keV, and plateaus at $\sim 8\%$ around 1836 keV. We used these data to benchmark the BoGEMMS-HPC simulations, which successfully reproduced the experimental results. Using the HPC-optimized and validated simulations, we performed extensive optical simulations on Leonardo to encode the true detector response on a correction matrix, which applies position-dependent corrections on the measured energy based on the optical processes. Since the ACS consists of BGO panels with different dimensions, we computed a dedicated correction for each panel, as the transport of optical photons also depends on the geometry of the crystal. Finally, we tested these corrections on GRB simulations. Specifically, we simulated a catalog of GRBs based on Fermi-GBM observations and analyzed their light curves to compute their significance. The correction matrix was then applied to these simulations to explore its impact on the GRB light curves, energy spectra, and the integrated counts of the individual ACS segments.

The development of a dedicated correction matrix, now integrated into the standard COSI data analysis pipeline, enables ACS simulations that closely reproduce real measurements while maintaining computational efficiency. As the ACS will contribute to the localization of GRBs by analyzing the relative count rates among different BGO panels, accurate modeling of each panel performance is essential for making reliable predictions of the ACS scientific contribution. Looking ahead, key perspectives of this work include refining the simulation benchmarking with future calibration campaigns in the final pre-launch phase, performing a robust assessment of the ACS response at different operating temperatures through dedicated experiments, and adapting the correction matrix accordingly. Moreover, the use of scintillation detectors coupled with SiPMs is emerging as a promising technique for future high-energy observations, making the validation of Geant4 optical physics carried out in this work valuable for such developments.

Contents

List of Figures

List of Tables

Chapter 1

The Anticoincidence System in the era of transient astronomy

In recent years, time-domain astronomy has emerged as a rapidly growing and highly promising field, playing a central role in modern astrophysics. The ability to detect and analyze transient and rapidly evolving phenomena has opened new windows into the dynamic Universe and provided unique insights into the underlying physical processes governing these extreme events. From an observational perspective, the Anticoincidence System (ACS), i.e., scintillating detectors surrounding the primary instruments to shield and reject the background, has played a key role in enabling gamma-ray space missions to detect and localize transient events, particularly Gamma-Ray Bursts (GRBs). These systems, typically composed of multiple scintillation detectors surrounding the primary instruments, are primarily designed to reduce background noise and improve the overall sensitivity of the mission. However, owing to their large sky coverage, they have also become effective tools for rapidly detecting transient phenomena across wide areas of the sky (typically from tens of keV to a few MeVs), enabling prompt communication with other space missions and facilitating rapid follow-up observations. The ACS will be a key instrumental component in an upcoming NASA Small Explorer mission, the Compton Spectrometer and Imager (COSI). Besides background rejection, the COSI ACS will act as a stand-alone gamma-ray transient detector, actively contributing to the detection and localization of short Gamma-Ray Bursts.

This chapter is organized as follows. Section 1.1 gives a brief description of Gamma-Ray Bursts and other typical gamma-ray transient sources. Section 1.2 moves to the observational domain, presenting the main interaction mechanisms of high-energy photons with matter. Section 1.3 focuses on scintillation detectors, listing their classification and operation mechanisms. Section 1.4 introduces the concept of Anticoincidence System and discusses how the ACS has played a key role as a gamma-ray transient monitor in past and current space missions. Finally, Section 1.5 introduces COSI, along with its instrumentation and scientific goals, with a special focus on its Anticoincidence System.

1.1 Transient gamma-ray astronomy

The Universe abounds with astronomical sources with variable characteristics and emission. These objects, which are commonly known as *transients*, have been widely observed throughout the whole electromagnetic spectrum, depending on the underlying physical process producing the variable emission. Transient events arise from a variety of astrophysical scenarios, such as neutron star mergers, interacting binary systems, supernova explosions, and stellar mass ejections. Among these, the most energetic events produce radiation peaking in the gamma-ray domain. Examples of these high-energy transients include Gamma-Ray Bursts, Active Galactic Nuclei (AGN), microquasars, solar flares, magnetars, and novae (Prokhorov and Moraghan 2023; Ramaty and Mandzhavidze 1998; Hurley 2011). In the following, we concentrate on Gamma-Ray Bursts and briefly outline the other classes of sources.

1.1.1 Gamma-Ray Bursts

GRBs are intense extragalactic transient bursts of gamma radiation. The initial emission (*prompt*) is usually characterized by a bright and impulsive burst peaking in the MeV band. Their origin is commonly associated with either the core-collapse of massive stars or mergers of binary compact objects. Their occurrence is of a few events per day with an isotropic distribution and isotropic equivalent energies of $\sim 10^{49} - 10^{55}$ erg (Luongo and Muccino 2021).

The standard scenario for the GRB formation is the *fireball* model. According to this scenario (Piran 1999), the central engine of GRBs is represented by a newborn black hole. The prompt emission is likely caused by the conversion of kinetic or magnetic energy into radiation via internal shocks in its jets, where particles are accelerated and reach ultra-relativistic velocities. Following the prompt emission, an *afterglow* emission is associated with external shocks between the jets and the surrounding medium.

GRBs are commonly classified based on the duration of their prompt gamma-ray emission. Conventionally, the burst duration is taken as the time interval over which 90% (from 5% to 95%) of the total background-subtracted counts are measured. This quantity is referred to as t_{90}. The first systematic study of the t_{90} distribution was performed using the CGRO-BATSE catalog (Kouveliotou et al. 1993). The resulting distribution (Figure 1.1) revealed a clear bimodality, with a division at approximately 2 seconds. This separation defines two main classes of GRBs: short GRBs ($t_{90} < 2$ s), generally associated with the merger of compact objects such as neutron stars, and long GRBs ($t_{90} > 2$ s), typically linked to the collapse of massive stars in supernova explosions (Berger 2014).

GRB light curves display a wide variety of temporal profiles. Some exhibit irregular, spiky, and erratic structures, while others appear smoother, with only one or a few distinct components. In certain cases, the emission period is clearly identifiable, whereas in others isolating the burst is more challenging. Representative GRB light curves are shown in Figure 1.2.

Within the t_{90} duration, the observed GRB spectrum is non-thermal and commonly best fitted by a phenomenological model called *Band* function (Band et al. 1993), which represents a smoothly joined broken power-law with the following functional form:

$$\frac{dN}{dE} = N_0 \times \begin{cases} \left(\frac{E}{100\ \mathrm{keV}}\right)^{\alpha} \exp\left[-\frac{(\alpha+2)E}{E_{peak}}\right], & E \leq \frac{\alpha-\beta}{\alpha+2} E_{peak} \\ \left(\frac{E}{100\ \mathrm{keV}}\right)^{\beta} \exp(\beta-\alpha) \left[\left(\frac{\alpha-\beta}{\alpha+2}\right)\left(\frac{E_{peak}}{100\ \mathrm{keV}}\right)\right]^{\alpha-\beta}, & E > \frac{\alpha-\beta}{\alpha+2} E_{peak} \end{cases} \tag{1.1}$$

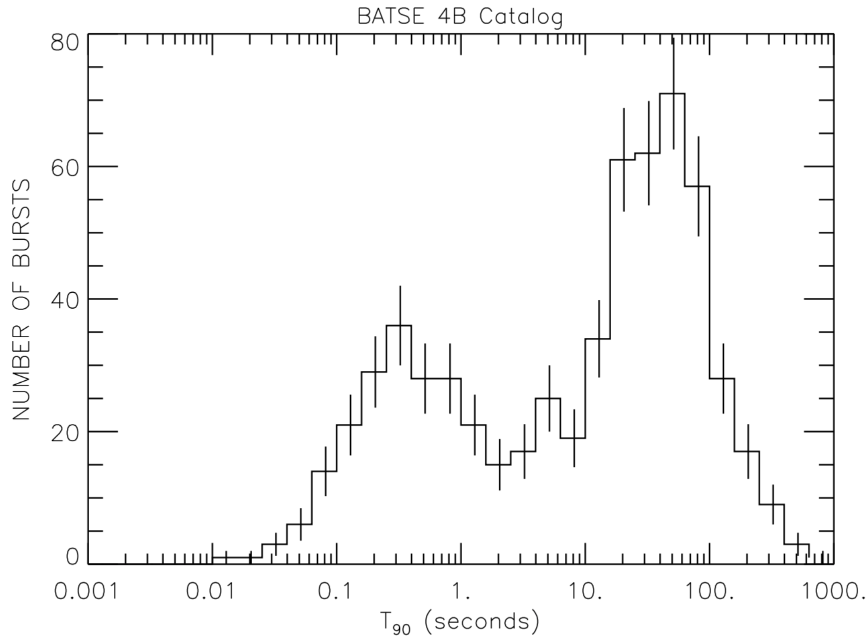

Figure 1.1 – Durations (t_{90}) of the 4B Catalog Gamma-Ray Bursts recorded with BATSE. Credit: https://gammaray.msfc.nasa.gov/batse/grb/duration/.

There are three independent parameters: a low-energy spectral index (α), a high-energy spectral index (β), and the peak of the energy spectrum (E_{peak}). In some cases, other models have been proposed to fit the observed spectrum, such as a power-law plus a thermal component (Ryde and Pe'er 2009).

In recent years, short GRBs have played a key role in multi-messenger astronomy, as they are associated with the electromagnetic counterparts of gravitational waves from binary neutron star mergers. In 2017, the short GRB 170817A was detected in temporal and spatial coincidence with the gravitational-wave event GW170817, both originating from a binary neutron star merger (Abbott et al. 2017). This event immediately triggered a massive follow-up observational campaign at all wavelengths. Today, many space missions are designed to efficiently detect and localize GRBs and to rapidly share alerts with other observatories, enabling fast multi-wavelength follow-up observations in coordination with gravitational-wave interferometers. In this context, a particularly useful instrument employed in space missions such as INTEGRAL (Winkler et al. 2003a) and AGILE (Tavani et al. 2009) is the ACS. Composed of scintillator detectors, the ACS is used to detect and localize GRBs over large portions of the sky while also minimizing the background rate in the primary instruments. The functioning and role of the ACS, along with its scientific impact on transient astronomy, are described in Section 1.4.

1.1.2 Other relevant transient gamma-ray sources

AGN AGN exhibit variable emission across the entire electromagnetic spectrum, with timescales ranging from minutes to decades. Observations from Fermi-LAT have provided significant progress in the study of AGN gamma-ray variability. Notable examples include the detection of quasi-periodic oscillations (QPOs) in some blazars (e.g., Ackermann et al. 2015) and fast gamma-ray flares on timescales of only a few minutes (e.g., Meyer et al. 2019). By comparing the variability properties of different AGNs, it is possible to probe the underlying

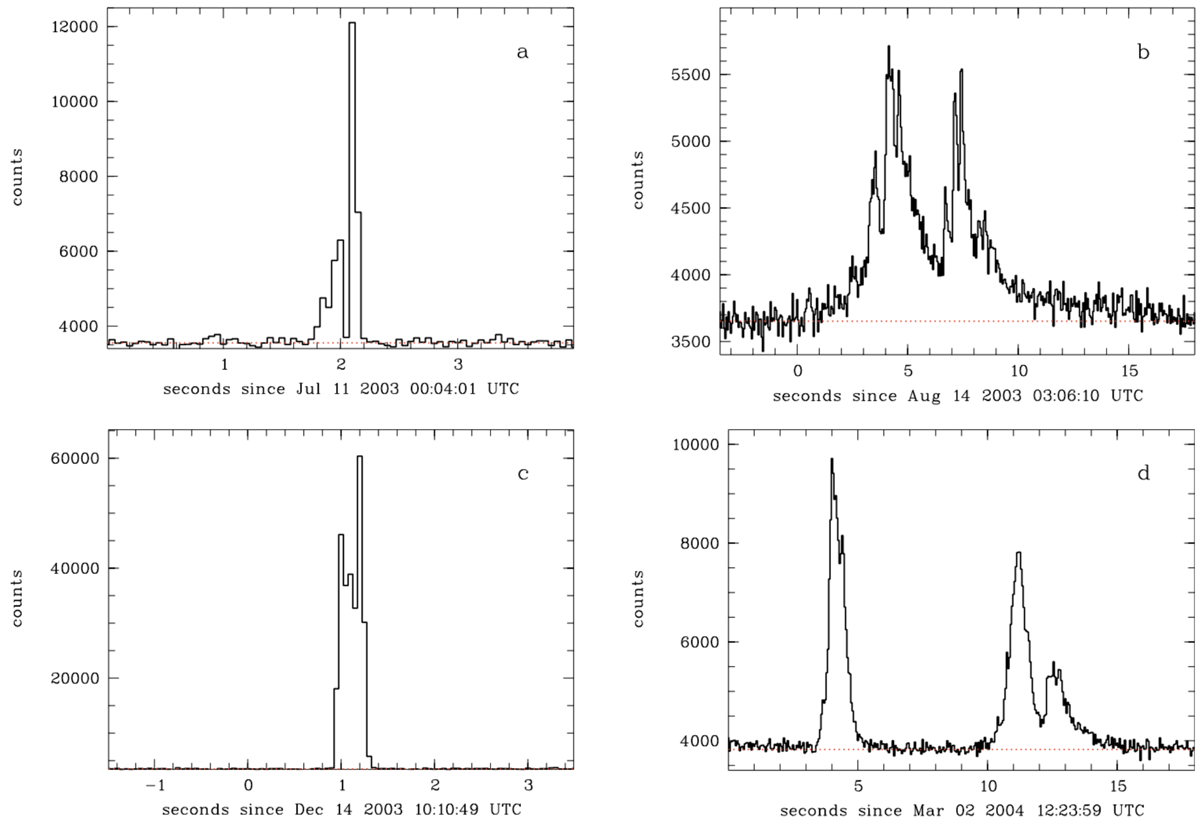

Figure 1.2 – Examples of GRB light curves as detected by INTEGRAL SPI/ACS (Rau et al. 2005).

physical mechanisms driving these events. In particular, blazars display strong variability due to Doppler boosting, which shortens the observed timescales (Bose et al. 2022).

Microquasars Microquasars are binary stellar systems consisting of a stellar-mass black hole gravitationally bound to a companion star. The black hole accretes matter from the companion, giving rise to intense radiation and the production of highly energetic particles. The first microquasar detected with Fermi-LAT was Cygnus X-3, which exhibited gamma-ray flares in 2008 and 2009 (Atwood et al. 2009). These events were also independently observed with AGILE (Tavani et al. 2009). The gamma-ray flares was associated with soft X-ray states and revealed a modulation of the gamma-ray flux with a period of 4.7917 ± 0.0011 hours, consistent with the orbital motion of the binary system (Prokhorov and Moraghan 2023). Another prominent microquasar, Cygnus X-1, has instead been detected at GeV energies during hard X-ray states (Zanin et al. 2016).

Solar flares Solar flares are sources of high-energy gamma rays that originate from the acceleration of ions during magnetic reconnection in the Sun's atmosphere. The solar flare gamma-ray emission can reach up to 100 MeV and exhibits both lines (Chupp et al. 1982) and continuum (Vilmer et al. 2003). A sample of 45 solar flares is currently available thanks to Fermi-LAT observations (Ajello et al. 2021). In more than 15 years of operation, the AGILE anticoincidence system detected 5003 transient events compatible with a solar origin in the $80 - 200$ keV.

Novae Classical novae occur when white dwarfs in binary systems accumulate material from a companion star, triggering a thermonuclear explosion. These eruptions expel high-velocity material enriched with newly formed isotopes. Among them, ^{13}N and ^{18}F decay rapidly (on timescales of minutes to hours) producing positrons that annihilate and generate a prompt gamma-ray signal, including a 0.511 MeV line and a continuum below this energy. Novae explosions usually happen in cycles, typically lasting $\sim 10^{4-5}$ yr. Fermi-LAT detected gamma-ray emission from the nova V407 Cyg by Fermi-LAT in 2010 (Abdo et al. 2010).

Soft gamma repeaters Soft gamma repeaters (SGRs) are sporadic sources of bursts of X-rays and gamma rays (Kouveliotou et al. 1998). They commonly have a short duration (about 100 ms) and may have long periods of inactivity. Spectra are often modeled either as a blackbody with temperature $\sim$ 10 keV or as Bremsstrahlung with temperature $\sim$ 25 keV, with isotropic energies up to $\sim 10^{40}$ erg (Hurley 2011). More rarely, SGRs may manifest as giant flares, with energies exceeding 10^{46} erg. To date, about 16 SGRs have been identified, 12 of which are confirmed[1]. The leading model associates SGRs with magnetars, i.e. highly magnetized neutron stars, although alternative scenarios, such as neutron stars surrounded by fossil disks, have also been proposed (Zhang et al. 2000).

Terrestrial gamma-ray flashes Terrestrial gamma-ray flashes (TGFs) are millisecond time-scale bursts of gamma rays produced in Earth's atmosphere. Discovered in 1992 by CGRO-BATSE (Fishman et al. 1994), they are associated with terrestrial lightning and thunderstorm activity (Inan et al. 1996). The underlying physical mechanism is the Relativistic Runaway Electron Avalanche (RREA) process. According to this process, a seed electron in a thundercloud is accelerated to relativistic speeds, producing secondary electrons. The collision of these secondary electrons with air molecules leads to the production of gamma rays via Bremsstrahlung. BATSE detected 76 TGFs during its 9 years of operation. Additional detections of TGFs were possible thanks to the Reuven Ramaty High Energy Solar Spectroscopic Imager (RHESSI) (Smith et al. 2005), AGILE (Marisaldi et al. 2010), Fermi (Briggs et al. 2010), and BeppoSAX (Ursi et al. 2017). In particular, AGILE detected more than 2000 events in 8 years of activity (Ursi et al. 2019), while the Fermi catalog contains 4144 TGFs.

1.2 Observing the High-Energy Universe

High-energy space-born astronomical instruments rely on two primary detection techniques: direct photo-detection or indirect reconstruction. In the first case, the photon is fully absorbed in the detector, and its energy is detected through the generation of a proportional electric signal. This method is particularly effective in the X-ray (up to a few hundred keV), where photoelectric absorption is the dominant interaction mechanism. At higher energies, however, efficient photo-absorption would require detectors of impractically large mass, while Compton scattering and pair production become increasingly dominant. In this regime, gamma-ray space telescopes must rely on indirect reconstruction techniques, where the photon properties are inferred from the kinematics of the secondary particles produced in the interaction.

[1] https://www.physics.mcgill.ca/ pulsar/magnetar/main.html

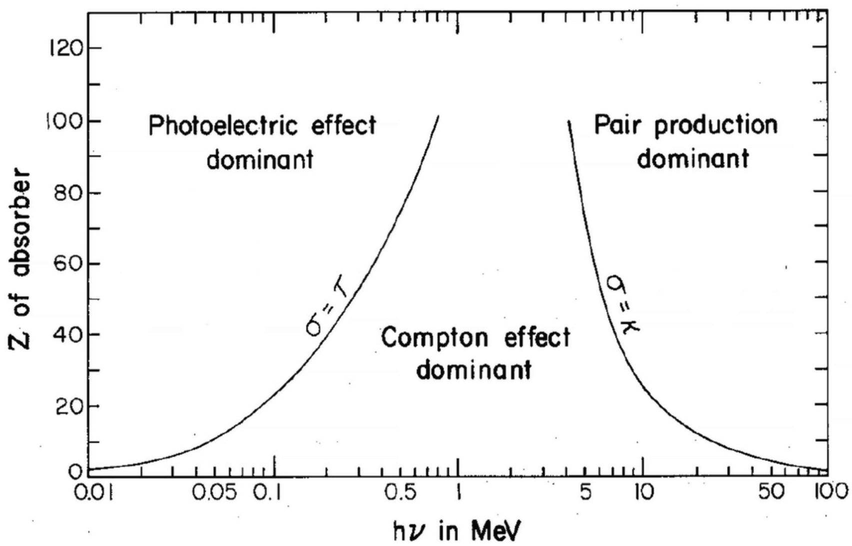

Figure 1.3 – Interplay between photoelectric absorption, Compton scattering, and pair production as a function of the detector material atomic number Z and energy of the incoming photon (Kierans et al. 2022).

1.2.1 Interaction of radiation with matter

From X-rays to gamma-rays, the main interaction processes between photons and matter are

- **photoelectric absorption** (up to a few hundred keV);
- **Compton scattering** (from ~ 100 keV to ~ 30 MeV);
- **pair production** (above a few MeV).

Their cross-section (e.g., the probability of occurrence) depends not only on the energy, but also on the properties of the material. Figure 1.3 illustrates the relative contributions of photoelectric absorption, Compton scattering, and pair production as functions of photon energy and the atomic number Z of the target material.

Photoelectric absorption

During a photoelectric absorption, the photon is absorbed by an atom and transfers all its energy to one of its bound electrons. This electron is then released with a kinetic energy E_k equal to the difference between the photon energy E_γ and the electron binding energy E_b:

$$E_k = E_\gamma - E_b \ . \tag{1.2}$$

The incoming photon must therefore have an energy larger than the binding energy to remove the electron from the atom. The ejected electron (also called photo-electron) usually ionizes surrounding atoms, progressively losing energy and ending up with a final deposit close to the photoelectric interaction. Most of the time, photo-electrons come from the innermost shell, like K or L shells. Once the electron is ejected, an electron from the outer shells moves to fill the vacancy. The transition of the electron from an outer to an inner shell can result in the emission of an X-ray fluorescence photon or an electron (Auger electron).

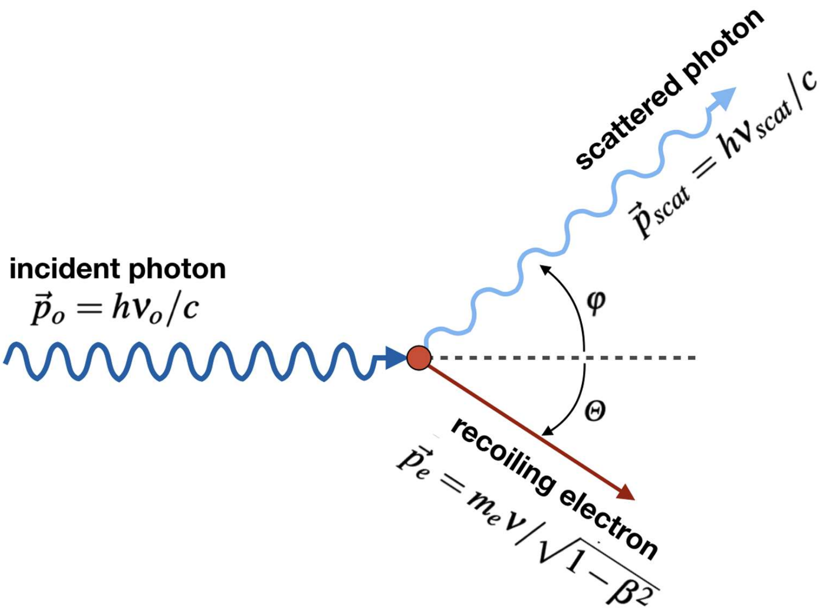

Figure 1.4 – Schematic diagram of a Compton scattering interaction between a photon and an electron (Kierans et al. 2022).

Compton scattering

Compton scattering is named after Arthur H. Compton, who discovered the scattering of X-ray and gamma-ray photons off of electrons in 1923 (Compton 1923), receiving the Nobel Prize in Physics in 1927. In the experiment, Compton irradiated the graphite with a beam of X-ray photons and found a scattered radiation component shifted to longer wavelengths (λ_{scat}) with respect to the original wavelength (λ_0), depending on the angle offset (φ). This result was not compatible with standard Thomson scattering of electromagnetic waves on free/bound electrons. Compton used the law of mechanics to explain the scattering of a quantum of radiation (photon) off an electron. The Compton scattering is shown schematically in Figure 1.4: an incoming photon with energy E_0 and momentum $\vec{p}_0$ scatters off an electron and transfers energy E_e and momentum $\vec{p}_e$ to the electron. The photon is scattered with a new energy E_{scat} and $\vec{p}_{scat}$. The scattering angle φ is called *Compton angle*. The conservation of energy and momentum gives

$$E_0 = E_{scat} + E_e \ , \tag{1.3}$$

$$\vec{p}_0 = \vec{p}_{scat} + \vec{p}_e \ . \tag{1.4}$$

Knowing that momentum $\vec{p}$ and frequency ν of a photon are related by $|\vec{p}| = h\nu/c$ (with h being the Planck constant and c the speed of light), and the electron momentum is given by $|\vec{p}_e| = \gamma\, m_e v$ (with m_e being the electron rest mass, v_e the electron velocity and $\gamma = 1/\sqrt{1 - v_e^2/c^2}$ the relativistic Lorentz factor), we end up with the well-known Compton equation:

$$\lambda_{scat} - \lambda_0 = \frac{h}{m_e c}(1 - \cos\varphi) \ . \tag{1.5}$$

Equation (1.5) can be solved for φ:

$$\cos\varphi = 1 - m_e c^2 \left(\frac{1}{E_{scat}} - \frac{1}{E_0} \right) \tag{1.6}$$

where wavelengths have been converted into energies.

The full quantum-mechanical treatment of the Compton scattering of unpolarized photons on free electrons was derived by Oscar Klein and Yoshio Nishina in 1928. Their study led to the Klein-Nishina formula for the differential cross-section of Compton scattering (Klein and Nishina 1929):

$$\frac{d\sigma}{d\Omega} = \frac{r_0^2}{2} \left(\frac{E_{scat}}{E_0} \right)^2 \left(\frac{E_{scat}}{E_0} + \frac{E_0}{E_{scat}} - \sin^2\varphi \right) \ , \tag{1.7}$$

where $r_0 = 2.818 \times 10^{-15}$ m is the classical electron radius.

Pair production

When a gamma-ray with an energy larger than twice the rest mass of the electron (1022 keV) is in a Coulomb field of a nucleus, it can be converted to an electron-positron pair. The kinetic energy of the pair corresponds to the excess energy of the photon above the 1022 keV threshold. A small, undetectable fraction of energy is also transferred to the nucleus. If the photon energy is more than four times the electron rest energy, pair production can also occur in the field of an electron. By measuring the directions and energies of the produced electron and positron with a tracker and absorber system, the properties of the original photon can be reconstructed. Notable instruments employing this technique include Fermi-LAT and AGILE.

1.3 Scintillation detectors

Scintillation is the physical process by which a material, known as a scintillator, emits ultraviolet or visible light when excited by high-energy photons (X-rays or gamma rays) or energetic particles (such as electrons, alpha particles, or neutrons). The use of scintillators for radiation detection is one of the oldest and most established techniques, as they offer an optimal balance of stopping power, fast response, cost-effectiveness, and geometrical flexibility. Photosensitive sensors, such as photomultiplier tubes (PMTs) and silicon photomultipliers (SiPMs), are coupled to the scintillator to collect the emitted optical photons and convert them into an electrical signal. The amplitude of this signal is proportional to the energy deposited by the original radiation. An ideal scintillator exhibits a linear conversion of deposited energy into light, emits scintillation photons promptly, and is optically transparent to ensure efficient light collection. In practice, scintillation detectors rarely fulfill all these ideal criteria simultaneously, but they are designed to approximate them as closely as possible.

A key parameter that characterizes scintillators is the *light yield* (or *light output*), which quantifies the number of optical photons produced per unit of deposited energy. It is typically expressed in units of photons/MeV or photons/keV and generally depends on the material and the temperature. Another critical property is the *decay time*, which describes how rapidly the scintillator emits light following excitation by ionizing radiation. Many scintillators exhibit

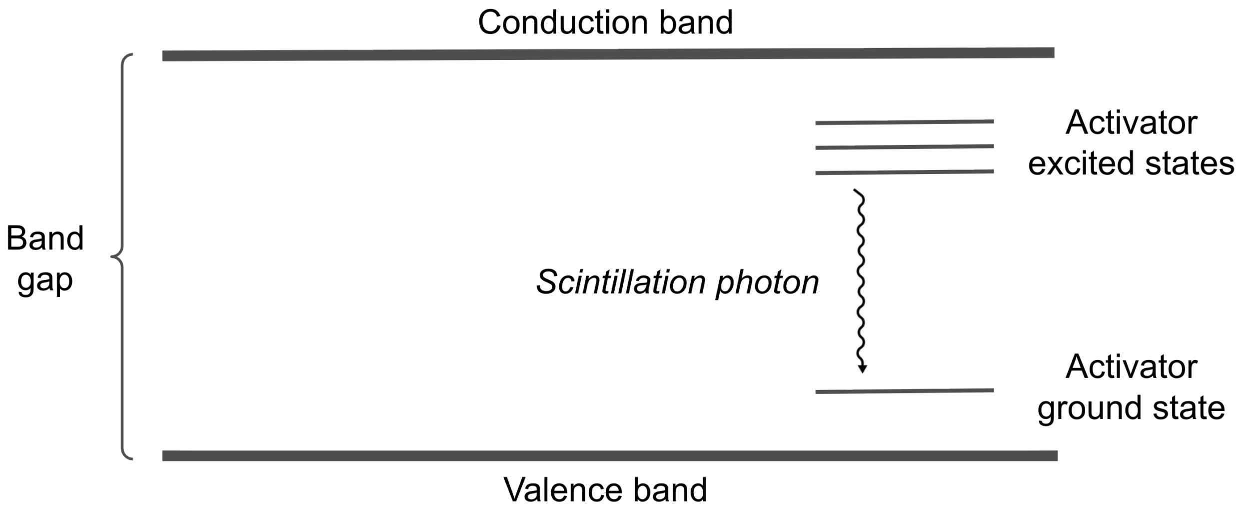

Figure 1.5 – Energy band structure for activated inorganic scintillators. Image adapted from Knoll (2010).

multiple decay components (such as fast and slow components), each with different time constants that influence the detector's timing performance.

There are two main classes of scintillators: *inorganic* and *organic* scintillators. In the following sections, we outline the operating principles of both classes. A comprehensive review of scintillation physics, which serves as the main reference for the following discussion, can be found in Knoll (2010).

1.3.1 Inorganic scintillators

In inorganic materials, electrons can occupy only discrete bands of energy. In the lower band, called the *valence band*, electrons are bound by the potential of the crystal atoms, while in the *conduction band* electrons have sufficient energy to move freely throughout the crystal. Between the valence and conduction bands, there is an intermediate band (the *forbidden band*) where electrons cannot exist in normal conditions. When a valence electron absorbs enough energy, it can transition to the conduction band, leaving behind a hole in the valence band. In a pure crystal, de-excitation from the conduction band to the valence band is inefficient and typically does not result in the emission of visible photons, as the energy gap is often too large (a few eV). To address this, small amounts of impurities, called *activators*, are introduced into the crystal. These impurities modify the energy band structure and create intermediate energy states within the forbidden gap. In this way, excited electrons can de-excite back to the valence band via these intermediate states, leading to the emission of photons in the optical range (Figure 1.5). The de-excitation sites are also called *luminescence centers*, and their energy structure determines the emission spectrum of the scintillator.

Typical half-lives of the luminescence centers are on the order of 50-500 ns, and electrons usually de-excite within this time range. The decay time of the luminescence centers defines the time characteristics of the scintillation light. An important consequence of the generation of these intermediate centers is that it allows the crystal to be transparent to the optical light. In fact, in a pure crystal, the emitted photons from electron de-excitation would have exactly the energy needed to excite other electrons, resulting in immediate self-absorption. Given that the electrons de-excite through intermediate steps, the emitted photons have less energy with respect to the gap between the valence and conduction band. The emission spectrum is therefore shifted to longer wavelengths and is not affected by self-absorption.

The most commonly used inorganic scintillators in a wide range of applications include NaI(Tl), CsI(Tl), LiI(Eu), BGO, $CdWO_4$, BaF_2, among others. In this section, we focus in particular on CsI(Tl) and BGO, as they are the scintillator materials selected for the anticoincidence systems of the COSI balloon mission and the COSI-SMEX mission (see Section 1.5), respectively.

CsI(Tl) Cesium iodide is one of the most widely used scintillation materials, typically doped with either thallium (CsI(Tl)) or sodium (CsI(Na)). CsI(Tl), in particular, offers a high absolute light yield of approximately 65 photons/keV at room temperature (Valentine et al. 1993a). One of its distinguishing characteristics is the relatively long scintillation decay time compared to other common scintillators. The light emission consists of two main components, with decay times and relative intensities of 0.68 μs (64%) and 3.34 μs (36%), respectively (Valentine et al. 1993b). CsI(Tl) has the interesting property of having a variable decay time for different exciting particles. This enables efficient pulse shape discrimination to differentiate among several types of radiation.

BGO Another widely used scintillation material is $Bi_4Ge_3O_{12}$, commonly abbreviated as BGO. A peculiar characteristic of BGO is that it does not require the presence of activator elements to enhance the scintillation. An advantageous feature of BGO is its high density (7.13 g/cm^3) and large atomic number ($Z = 83$) of the bismuth component. These properties allow the BGO to be particularly efficient in photoelectric absorption of gamma-rays. The BGO light yield is quite low with respect to other materials, ranging from 8 to 10 photons/keV. Moreover, it has a relatively high refractive index (~ 2.15) if compared to other scintillating crystals. This property can reduce the efficiency of light collection when the crystal is coupled with photosensitive devices, such as photomultiplier tubes or silicon photomultipliers, which typically use optical windows with lower refractive indices. Large differences in refractive index increase the probability that optical photons are reflected rather than absorbed, thus reducing the light collection efficiency. The principal decay time for the BGO is 300 ns. A faster scintillation component is also present, with a decay time of about 60 ns, contributing around 10% of the total light output. The BGO emission spectrum ranges from about 330 nm to 620 nm and has a peak at 480 nm (Akchurin et al. 2011).

1.3.2 Organic scintillators

In organic materials, the emission of optical light results from energy transitions within individual molecules. Most organic scintillators are based on molecules with specific symmetry properties that give rise to characteristic energy levels for the electrons, in what is called π-electron structure (Figure 1.6). The i-th singlet state (total spin 0) and triplet state (total spin 1) are denoted as S_i and T_i, respectively. Each of these electronic states is further split into a set of closely spaced vibrational levels, typically indicated as S_{01}, S_{02}, and so on. At room temperature, nearly all molecules reside in the ground vibrational state S_{00}, since the thermal energy is insufficient to populate the higher states. However, when the molecule of a scintillator absorbs the kinetic energy of an incoming photon or charged particle, it is excited to one of the higher electronic states. These excited molecules rapidly de-excite

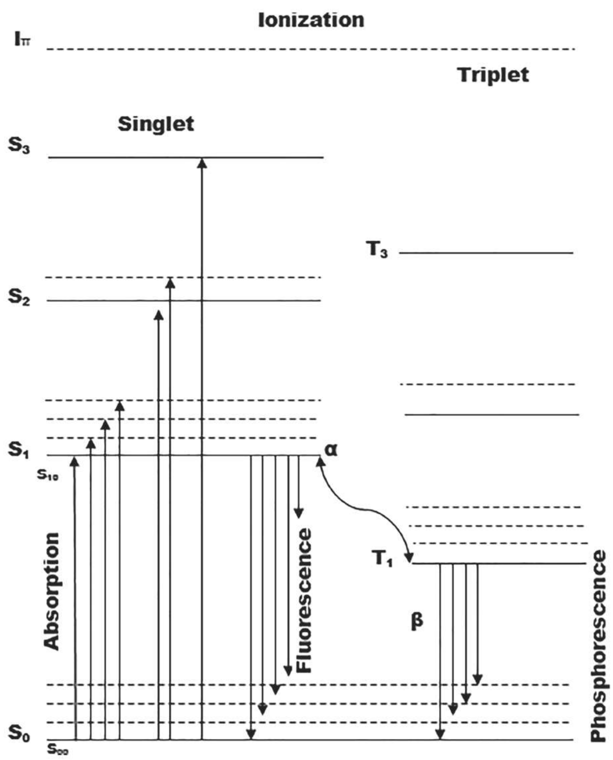

Figure 1.6 – Energy levels of an organic molecule with π-electronic structure (Gupta and Mao 2020).

to the S_{10} state. The primary scintillation light (also called *prompt fluorescence*) is then emitted as the molecule returns from S_{10} to one of the vibrational levels of the ground state S_0. The characteristic decay time of the prompt fluorescence is usually of the order of a few nanoseconds. Some excited singlet states can undergo intersystem crossing, transitioning into triplet states. The lifetime of T_1 is typically much longer than S_1 (about 10^{-3} s). As a result, the transition between T_1 and S_0 gives rise to a delayed light emission, known as *phosphorescence*. Since T_1 lies at a lower energy than S_1, the phosphorescence spectrum is shifted towards longer wavelengths compared to the prompt fluorescence. However, some molecules may be thermally excited back from T_1 to S_1 and subsequently decay through normal fluorescence, resulting in delayed fluorescence, which is indeed observed in certain organic scintillators.

There are three main categories of organic scintillators: pure organic crystals, liquid solutions, and plastic scintillators. Popular examples of pure organic crystals are Anthracene ($C_{14}H_{10}$) and Stilbene ($C_{14}H_{12}$). Liquid scintillators are obtained by dissolving an organic scintillator in specific solvents. A key advantage of liquid scintillators is their reduced susceptibility to radiation damage, as they lack a rigid crystal lattice. Plastic scintillators are produced by dissolving the scintillating compound (called the fluor) into a solvent that is then polymerized (called the base), forming a solid solution. Plastic scintillators can be easily manufactured in a variety of shapes and sizes, making them versatile and cost-effective. Due to their low production cost, they are often the preferred choice when large detector volumes are required. Some of the most popular bases include polyvinyltoluene (PVT) and polystyrene (PS).

1.3.3 Scintillation detectors in gamma-ray space missions

Scintillation detectors have been widely employed in gamma-ray space missions, serving different purposes. They act as active shields to reduce background (i.e., as anticoincidence systems, discussed in Section 1.4), as stand-alone detection devices, or as calorimeters complementary to trackers. For instance, COMPTEL was composed of two detector arrays: an upper liquid scintillator (NE 213A) and a lower NaI(Tl) array. Gamma-rays first underwent Compton scattering in the upper detector and were subsequently absorbed in the NaI crystals. The Fermi-GBM comprises 12 NaI detectors and 2 BGO crystals. The NaI detectors are sensitive from a few keV up to about 1 MeV, providing both burst triggers and source localization, while the BGO detectors extend the coverage from 150 keV to 30 MeV. In AGILE, the Mini-Calorimeter (MCAL) was composed of 30 CsI bars and was designed as a non-imaging scintillation detector operating from 350 keV to 100 MeV.

1.3.4 Photosensitive devices

In order to produce an actual measurement, the scintillation light must be collected and converted into an electric signal. This is typically achieved by coupling the scintillator to a photosensitive device capable of generating an output signal proportional to the number of collected optical photons. PMTs have been among the most commonly used devices in space missions, such as COMPTEL, INTEGRAL, Fermi, and AGILE. SiPMs, however, are increasingly adopted and are planned for future missions, including COSI.

Photomultiplier Tubes (PMTs)

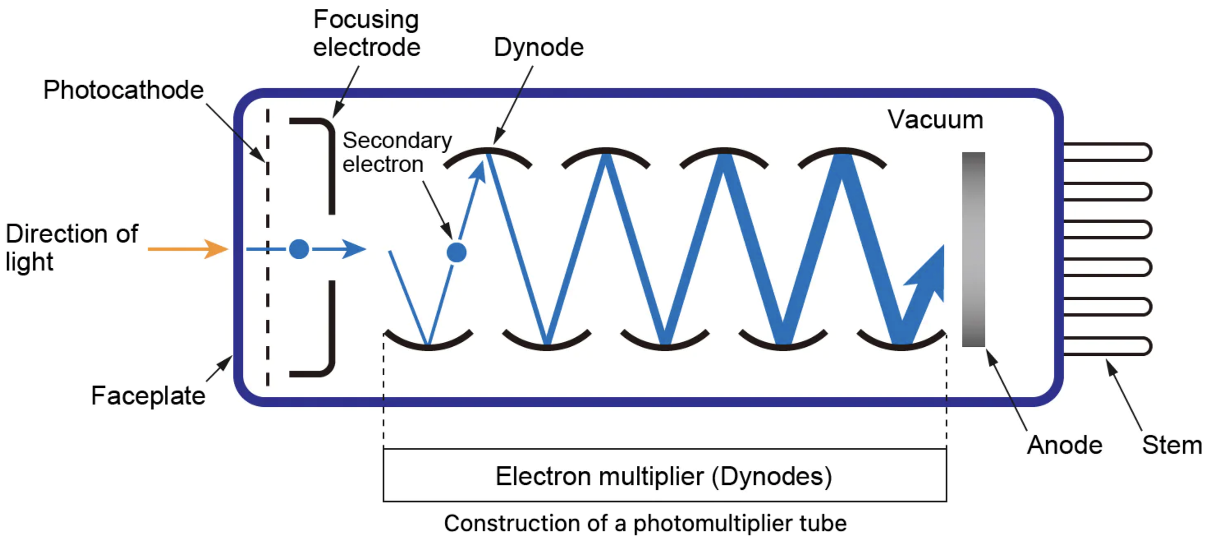

Figure 1.7 – Schematic illustration of the functioning of a photomultiplier tube. Credit: Matsusada Precision webpage.

PMTs are vacuum tubes composed of a photocathode, a series of dynodes, and an anode (Figure 1.7). Optical photons emitted by the scintillator strike the photocathode, ejecting electrons via the photoelectric effect. A focusing electrode directs these primary electrons towards the first dynode. Each dynode is maintained at a progressively higher potential (by

$\sim$ 100 Volts), causing the electrons to be accelerated and multiplied through the emission of secondary electrons. This process is repeated across the series of dynodes, resulting in an exponential amplification of electrons, approximately $N_e^{N_d}$, where N_e is the average number of secondary electrons for each incoming electron, and N_d is the number of dynodes. Finally, the anode collects the multiplied electrons, producing a detectable current pulse.

The capability of a photocathode to produce electrons is known as *quantum efficiency* (QE). The quantum efficiency is defined as the ratio of the number of electrons ejected by the photocathode to the number of incident photons, and is usually expressed as a percentage. It strongly depends on the wavelength of the incident light and its peak value is usually less than $\sim$ 35%. PMT photocathodes are commonly made of alkali antimonides, e.g. bialkali (SbKCs) and multialkali ($SbNa_2KCs$). The photocathode material is chosen to maximize the sensitivity to a specific region of the optical light spectrum, depending on the scintillation emission of the scintillator. Another key quantity is the *gain*, or current amplification, defined as the ratio of the anode current to the photocathode current. It is typically of the order of 10^5-10^7 and depends on the power of supplied voltage.

Silicon photomultipliers (SiPMs)

The SiPMs are single-photon-sensitive devices aimed at collecting the optical light generated inside the scintillator and the consequent generation of a proportional electric signal (Figure 1.8). They are composed of thousands of single-photon avalanche diodes (SPADs) operating in Geiger mode, that is, when the applied voltage is above the breakdown voltage[2]. In this regime, when a SPAD is hit by a photon, a large electric signal is generated from internal avalanche multiplication, which depends on the gain of the SiPM. The produced avalanche is eventually quenched[3] and, after a characteristic recovery time, the SPAD is restored and ready for the detection of another photon. Each SPAD detects photons independently, and the final output is given by the sum of the photocurrents from each SPAD. Another key quantity related to the SiPM performance is the *photon detection efficiency* (PDE), that is, the probability that the SiPM produces a signal in response to an incident photon. It is a function of the overvoltage (defined as the difference between the applied voltage and the breakdown voltage), and the wavelength of the incident light. It includes the geometrical fill factor (i.e., the fraction of the total area occupied by the microcells), the quantum efficiency (i.e., the probability that an incident photon generates an electron-hole pair), and the probability of Geiger discharge (i.e., the probability that the generated electron/hole triggers the avalanche process).

Compared with PMTs, the SiPMs offer substantial reductions in mass, volume, power consumption, and cost, and they can be packed into arrays. Although PMTs represent a well-established technology, SiPMs are a promising next-generation alternative, particularly suited for applications requiring compactness and low power consumption. On the other hand, SiPMs are often exposed to high levels of ionizing and non-ionizing radiation, leading to a deterioration of their performance over time (Merzi et al. 2024). Consequently, the selection of PMTs or SiPMs is commonly determined by the trade-off between the performance benefits and limitations of the two technologies.

[2] The breakdown voltage is the minimum voltage at which avalanche multiplication of carriers is possible.

[3] When an avalanche starts, the current passing through a quenching resistor causes a voltage drop below the breakdown voltage, consequently stopping the avalanche.

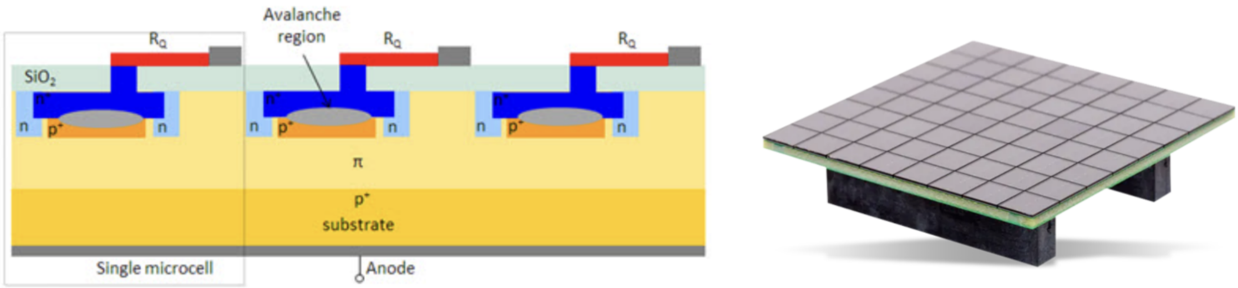

Figure 1.8 – On the left, typical structure of a SiPM (from Hamamatsu webpage). On the right, an array of SiPMs (from Onsemi datasheet).

1.4 Anticoincidence System (ACS)

For point sources, the energy-integrated flux F_s (defined as the number of counts per unit time and area) can be expressed as follows (Zoglauer 2006):

$$F_s = n_\sigma \frac{\sqrt{N_s + N_b}}{A_{eff}\,T} \,, \tag{1.8}$$

where N_s and N_b are the counts of the source and background, respectively, A_{eff} is the effective area of the telescope, T is the observation time, and n_σ is the significance of the detection in standard deviations. The minimum detectable flux of an instrument defines its sensitivity. From Equation (1.8), it is evident that the minimum detectable flux can be reduced (thereby improving sensitivity) by increasing either the effective area or the integration time. Equivalently, an increase in the background rate leads to a worsening of sensitivity. Therefore, minimizing the background flux in the detector is essential to ensure optimal telescope performance. A common technique is the use of an Anticoincidence System, consisting of scintillating detectors that surround the primary instruments. The ACS acts as a veto detector: whenever an interaction is recorded in one of its panels, a veto signal is raised, and any coincident event in the main detector is rejected. One of the earliest implementations was proposed by Kenneth John Frost in 1962, with an active CsI(Tl) scintillation shield around the gamma-ray detector (Frost and Rothe 1962). Modern ACSs are typically built from scintillating materials, most commonly CsI, NaI, BGO, or plastic scintillators. Scintillators offer several benefits, as they provide fast and efficient detection of charged particles. Additionally, they can be made into panels or segmented shapes to effectively surround the main detector. Organic scintillators are commonly preferred when a fast response, low-mass, and large-area coverage are required, and timing performance is prioritized over energy resolution. Inorganic scintillators, on the other hand, are favored when good energy resolution and compact geometries are needed (Knoll 2010). The scintillation panels are typically coupled to photosensitive devices that collect the emitted optical photons and generate electrical pulses proportional to the deposited energy. The rejection of background events is particularly useful for Compton telescopes, which rely on Compton scattering to reconstruct incoming photons. In such instruments, only photons that are fully absorbed in the primary detector can be effectively reconstructed. If a photon escapes and deposits energy in an ACS panel, the event can be discarded, thus improving reconstruction fidelity. This approach is especially relevant for the upcoming space mission COSI, which will be presented in Section 1.5.

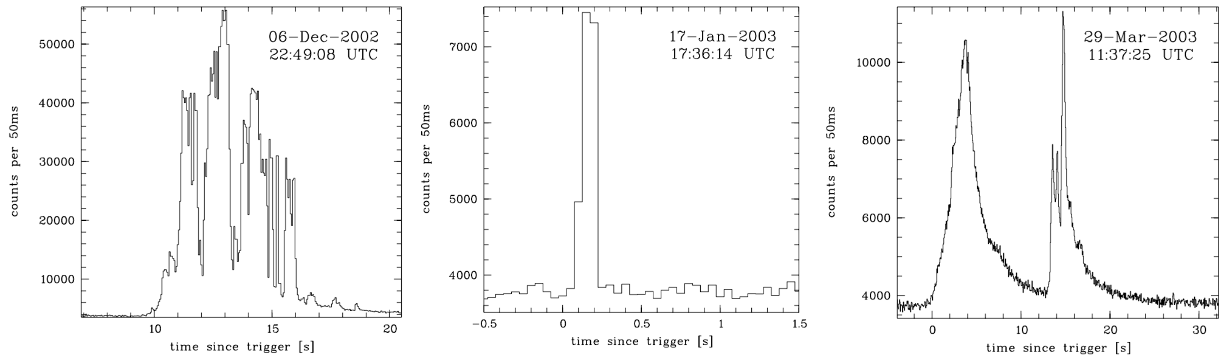

Figure 1.9 – Light curves for three typical GRBs detected by SPI-ACS with 50 ms time resolution (von Kienlin et al. 2003).

1.4.1 The ACS as gamma-ray transient monitor

Anticoincidence systems, traditionally designed to suppress background events in the primary instruments for high-energy observations, have also demonstrated their scientific utility beyond background rejection. Thanks to photosensitive devices like PMTs and SiPMs, the scintillation light produced within the ACS modules can be detected and converted into an electric signal, leaving a clear footprint of the passage of an energetic photon or charged particle. By recording and analyzing signals from these systems, it is possible to detect and characterize transient astrophysical phenomena such as GRBs, solar flares, and magnetar outbursts. In several space missions, anticoincidence detectors have thus operated as complementary instruments, allowing for a faster detection of transient events and contributing to constraining their localization. As illustrative examples, we describe the use of anticoincidence systems in three major astrophysical space missions, INTEGRAL, AGILE, and Fermi, highlighting their design and scientific contributions.

The INTEGRAL SPI-ACS

The INTErnational Gamma-Ray Astrophysics Laboratory (INTEGRAL, Winkler et al. 2003) is a retired space telescope designed for the detection of gamma rays up to 8 MeV. It was equipped with two instruments, IBIS (the imager) and SPI (the spectrometer), along with an anticoincidence system (SPI-ACS) for background rejection. The SPI-ACS consisted of 91 BGO crystals positioned on the lateral and bottom sides of the spectrometer. Each of the 90 BGO crystals was coupled with two PMTs, whose signals are read out by 90 front-end electronic (FEE) boxes, while the 91st BGO crystal was read out by a single PMT and FEE. Besides providing veto signals for photons and charged particles coming from outside the FoV, it offered excellent sensitivity for the detection of GRBs, with a threshold of $\sim$ 75 keV. The count rates of the 91 FEEs were sampled every 50 ms. The SPI-ACS count rates were monitored simultaneously at 8 different timescales (from 0.05 s to 5 s). An alert was generated if at least two of the monitor programs registered a count rate excess above a predefined significance threshold, adjustable for each timescale. The SPI-ACS made a significant contribution to GRB science, with 388 GRB detections in the first 26.5 months of operation (Rau et al. 2005). Figure 1.9 shows three examples of GRB light curves measured by the SPI-ACS.

The AGILE anticoincidence system

AGILE (Tavani et al. 2009) is a retired ASI light satellite aimed at gamma-ray observations from 30 MeV to 50 GeV. The AGILE scientific instrument was composed of three detecting systems: a silicon tracker, a scintillation minicalorimeter, and a coded mask (Super-AGILE). To reduce the background rate, the three instruments were shielded on the top and lateral sides by an anticoincidence system. The ACS was composed of 13 modules, 3 on each lateral side and 1 on the top. The modules consisted of BC-404 plastic scintillators, coupled with PMTs for light collection. The ACS was sensitive to hard X-ray photons in the 50-200 keV range and continuously recorded the count rate of each panel as ratemeter data, with a time resolution of 1.024 s. Alongside the tracker and the minicalorimeter, the ACS has played a key role in the prompt detection of GRBs and other transient events, and dedicated machine learning techniques have been developed for rapid signal identification in ACS data (Parmiggiani et al. 2024).

Fermi GBM

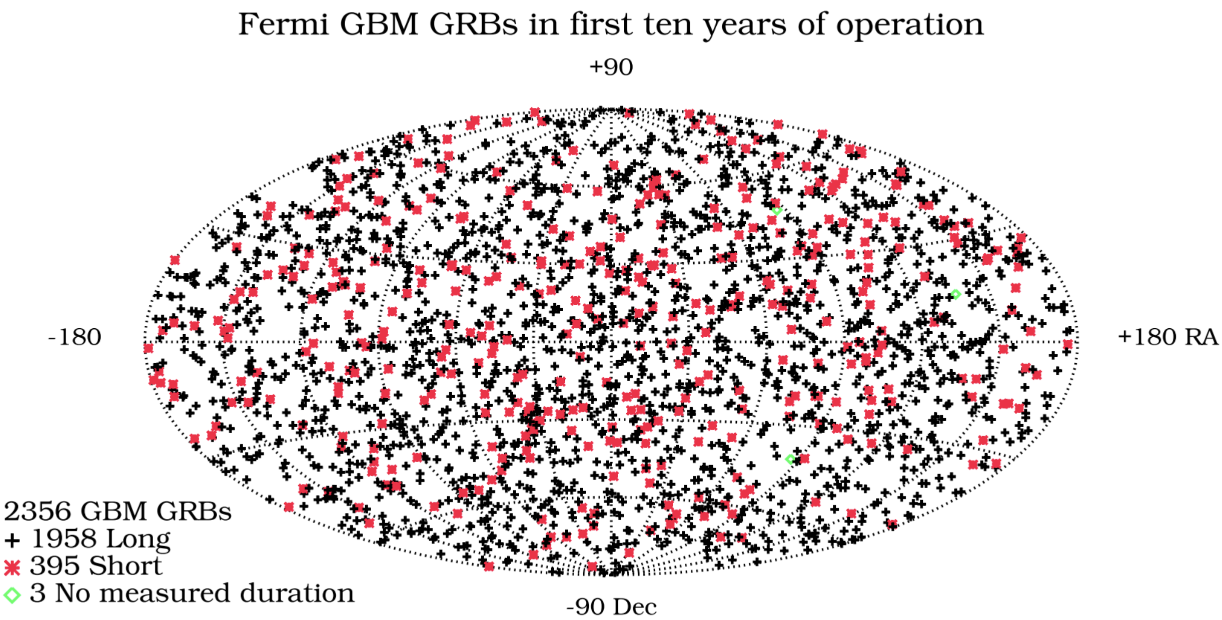

Figure 1.10 – Sky distribution of detected GRBs from GBM in celestial coordinates (von Kienlin et al. 2020a).

The Fermi Gamma-Ray Space Telescope (FGST) is a major NASA mission launched in 2008 and aimed at observations of gamma rays. The main instrument, the Large Area Telescope (LAT, Atwood et al. 2009), is designed to observe gamma rays in the energy range from about 20 MeV to 300 GeV. To reduce the charged-particle background, the LAT is equipped with an Anticoincidence Detector (ACD), a segmented plastic scintillator shield that surrounds the tracker and helps discriminate between gamma rays and charged cosmic rays. Complementing the LAT, the Gamma-ray Burst Monitor (GBM) (Meegan et al. 2009) extends coverage to transient sources, detecting X-rays and gamma rays between 8 keV and 40 MeV. The GBM includes 12 NaI scintillation detectors, which cover the lower part of the energy range up to $\sim$ 1 MeV, and 2 BGO scintillation detectors, which cover the energy range of $\sim$ 150 keV to $\sim$ 40 MeV. Although not designed to serve as an anticoincidence system for background rejection, the GBM is a clear example of how scintillation detectors can actively contribute to the detection and localization of transient events. Coupled with PMTs for light collection, the GBM detectors produce three types of data: Time-Tagged Events (TTE), which record the time and energy (in 128 channels) of each photon; CTIME, consisting of 8-channel spectra every 0.256 s; and CSPEC, providing 128-channel spectra every 4.096 s, with a finer time resolution of 1.024 s immediately following a trigger. Since its launch, GBM has detected about 240 GRBs per year, releasing catalogs on a regular basis with the main characteristics of triggered bursts (Gruber et al. 2014; von Kienlin et al. 2014; Narayana Bhat et al. 2016; von Kienlin et al. 2020a). The last released catalog contains 2356 GRBs detected in 10 years of operation, from 2008 July 12 to 2018 July 11 (Figure 1.10).

1.5 The Compton Spectrometer and Imager (COSI)

1.5.1 The MeV gap

Besides impressive scientific advances in high-energy astronomy, the soft gamma-ray band ($\sim$ 0.2-10 MeV) still remains largely unexplored. This region of the electromagnetic spectrum is often referred to as *MeV gap*, to underline the observational gap in this particular energy band. COMPTEL performed the first all-sky survey in the 0.75-30 MeV energy band and currently represents the instrument achieving the best sensitivity at these energies. Nevertheless, its sensitivity was several orders of magnitude worse than that of other instruments in neighboring bands (Figure 1.11). The main problems in the soft gamma-ray domain are due to physical, technological, and environmental challenges. In fact, the MeV band is the realm of nuclear gamma-ray lines, which inevitably leads to high levels of instrumental background from Cosmic Rays interactions and deactivation of irradiated materials in space. Moreover, unlike optical and X-ray photons, gamma-rays in the MeV range cannot be focused using lenses or mirrors due to their high penetration power and relatively low cross section. The interaction of photons with matter has the lowest cross-section in the MeV band and is dominated by Compton scattering, which is more difficult to reconstruct than photo-absorption (relevant at lower energies) and pair production (relevant at higher energies). This portion of the electromagnetic spectrum, however, is particularly important for gaining insight into many astrophysical objects and for developing a more comprehensive understanding of high-energy astrophysics. Many sources in the Universe, such as Gamma-Ray Bursts, blazars, and pulsars, exhibit peak emissivity at photon energies between 0.2 MeV and 100 MeV. Additionally, this energy band is characterized by nuclear lines from radioactive elements (important to understand the chemical evolution of our Galaxy), electron-positron annihilation line at 0.511 MeV, and emission from accreting black holes (BHs). It is therefore crucial to improve our observational capabilities and deepen our understanding in this energy band, as it can reveal important information about many high-energy astrophysical phenomena and the physical processes that shape the Universe.

1.5.2 COSI

The Compton Spectrometer and Imager (COSI, Tomsick et al. 2024) is an upcoming NASA Small Explorer (SMEX) satellite mission that aims to address the observational gap in the soft gamma-ray band. Scheduled for launch in 2027, it is a wide-field gamma-ray telescope aimed

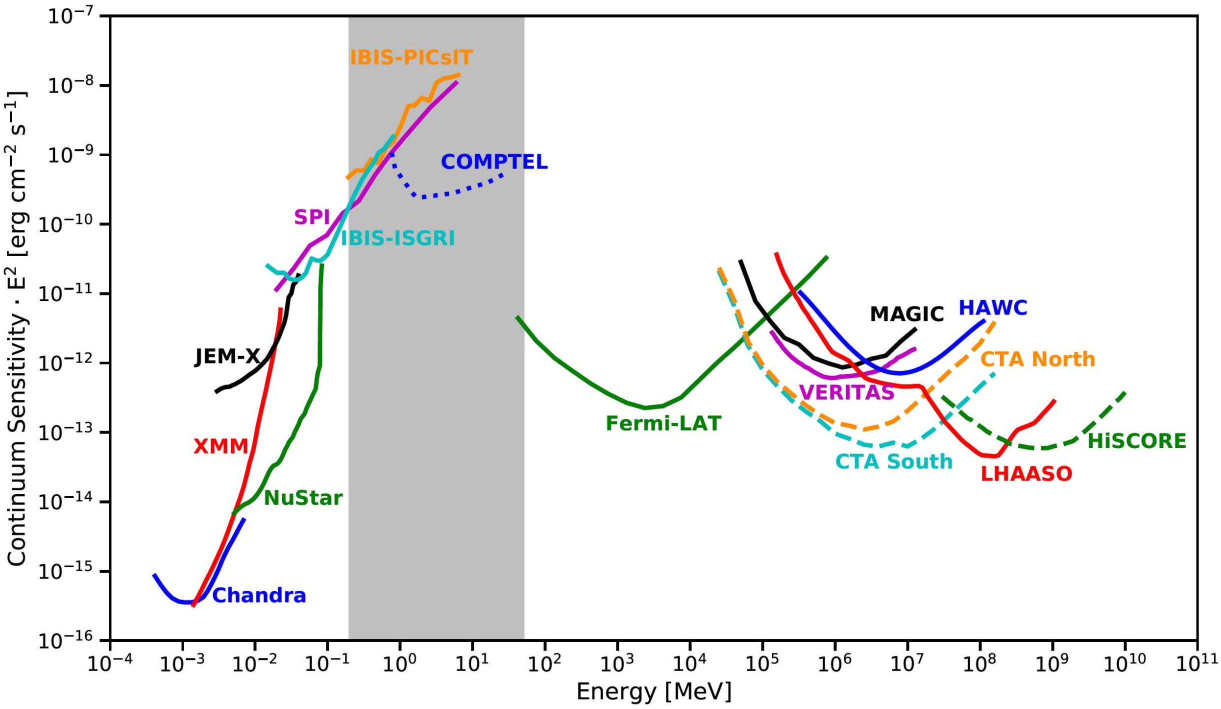

Figure 1.11 – Differential sensitivity for current and past X-ray and gamma-ray missions (Lucchetta et al. 2022).

at surveying the entire sky at 0.2-5 MeV and will operate in a low-Earth near-equatorial orbit. It is designed to perform imaging, spectroscopy, and polarization of astrophysical gamma-ray sources. In 2016, the COSI instruments experienced a 46-days long balloon flight (Kierans et al. 2016), allowing for the acquisition of the very first scientific data and testing their calibration and general performance.

COSI operates as a Compton telescope, exploiting multiple Compton scattering events within the instrument before the final photo-absorption. The instrument consists of 16 crossed-strip cryogenic germanium detectors (GeDs). Each crystal measures 8 cm×8 cm×1.5 cm and has 64 aluminum strip electrodes placed orthogonally on the anode and cathode. This configuration allows the GeDs to have accurate energy and 3D position sensitivity. At each Compton scatter, numerous electron-hole pairs are produced from ionization by the recoil electron. A high voltage is applied across the detectors, so that the electrons and holes travel in the opposite direction and produce a signal. The number of electrons/holes, and therefore the produced pulse height, is proportional to the energy of the recoil electron, that is, the energy deposited by the gamma photon. For position measurement, electrons and holes generated by the interaction are collected by orthogonal strip electrodes located directly above and below the interaction region. The intersection of the activated strips determines the X-Y position of the event. The Z-position is inferred from the difference in arrival times of the charge carriers at their respective electrodes: if the interaction occurs closer to the anode, the electrons have a shorter distance to travel and will reach the electrode earlier than the holes reaching the cathode. The detectors are cooled to 80-90 K using a mechanical cryocooler and placed in an aluminum cryostat. The GeDs are surrounded by an Anticoincidence System, which is discussed in details in Section 1.6.

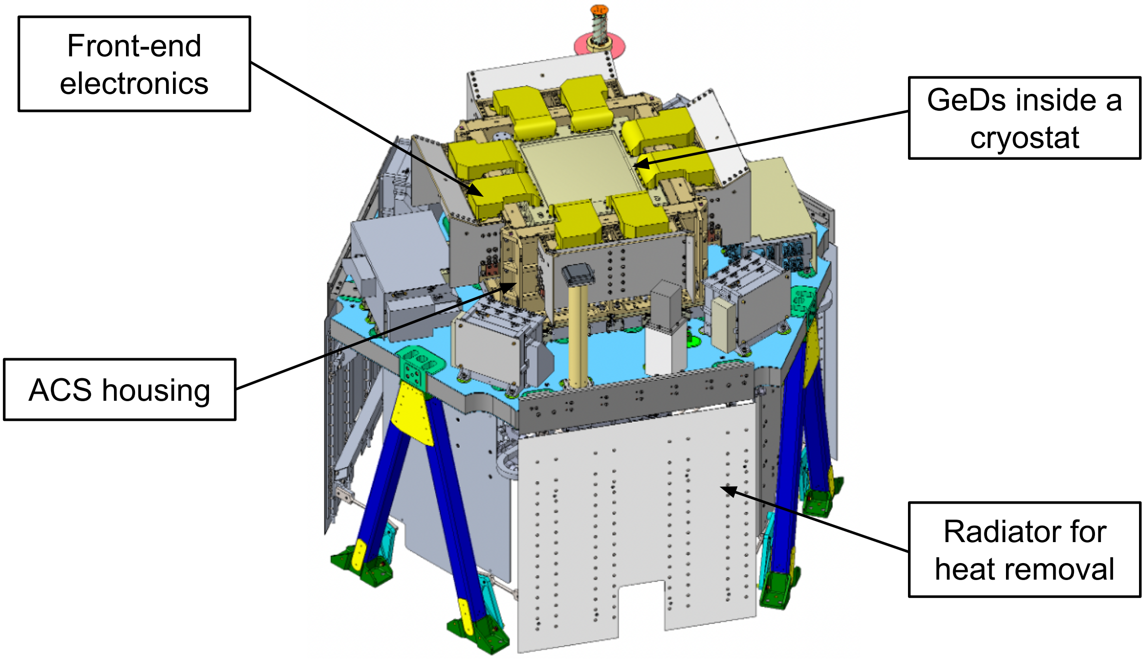

Figure 1.12 – COSI payload design, which shows the cryostat enclosing the GeDs, the ACS housing, the front-end electronics with ASIC readout, and the radiators for heat removal.

1.5.3 COSI key science goals

The main science goals for COSI can be summarized in the following points: (*i*) uncovering the origin of Galactic positrons; (*ii*) revealing the dynamics of nucleosynthesis in the Galaxy; (*iii*) gaining insights into extreme environments through polarization measurements of compact objects; (*iv*) actively contributing to multi-messenger astrophysics.

The specific requirements for spectral resolution, angular resolution, and line sensitivity necessary for meeting these goals are summarized in Table 1.1. The required narrow-line and continuum sensitivities compared to current and previous missions are shown in Figure 1.13.

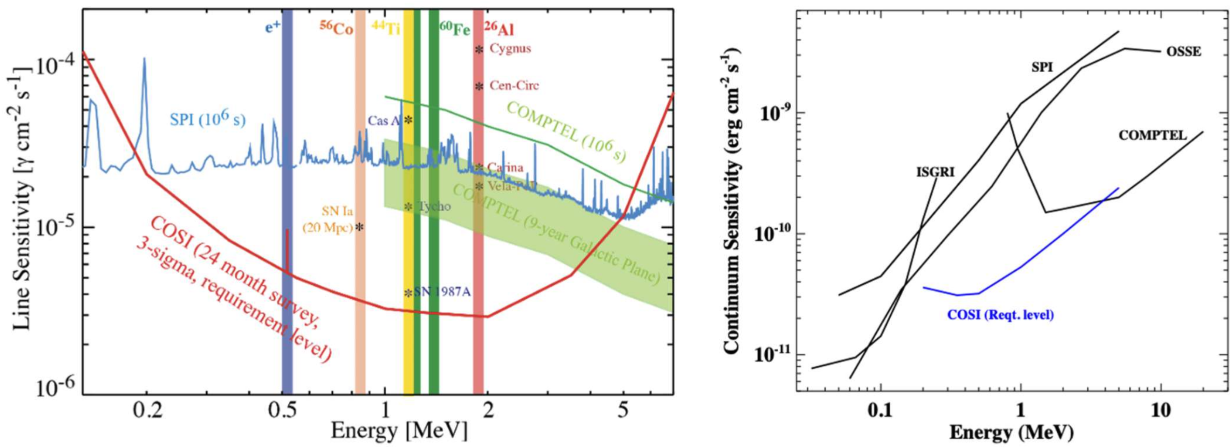

Figure 1.13 – Narrow-line (left) and continuum (right) sensitivities at requirement level compared to current and previous space missions (Tomsick et al. 2024).

Table 1.1 – COSI requirements (Tomsick et al. 2024)

Parameter	**Requirement**
Sky Coverage	25%-sky instantaneous field of view (all-sky every day in survey mode)
Spectral Resolution (FWHM)	6.0 keV at 0.511 MeV and 9.0 keV at 1.157 MeV
Angular Resolution (FWHM)	4.1° at 0.511 keV and 2.1° at 1.809 MeV
Line sensitivity[a]	1.2×10^{-5} photons cm^{-2} s^{-1} at 0.511 MeV 3.0×10^{-6} photons cm^{-2} s^{-1} at 1.157 MeV 3.0×10^{-6} photons cm^{-2} s^{-1} at 1.1173 MeV 3.0×10^{-6} photons cm^{-2} s^{-1} at 1.333 MeV 3.0×10^{-6} photons cm^{-2} s^{-1} at 1.809 MeV
Flux limit for polarization[b]	1.4×10^{-10} erg cm^{-2} s^{-1} (0.2-0.5 MeV)
Reporting short GRB detections	< 1 hour reporting time $\sim 1°$ localizations (accuracy depends on GRB fluence) 100 ms absolute time accuracy

[a] 3σ narrow line point source sensitivity in 2-years of survey observations.
[b] For 50% minimum detectable polarization in 2-years of survey observations.

Origin of Galactic positrons

The production of positrons and their subsequent annihilation in the Galactic interstellar medium (ISM) remains a cornerstone topic in gamma-ray astronomy. Although the 511 keV annihilation line from the inner Galaxy was discovered over five decades ago (Johnson et al. 1972), the precise origin of the positrons responsible for this emission is still not fully understood. Several satellite and balloon-borne missions have detected the 511 keV signal, but the most significant contribution came from INTEGRAL/SPI, which provided a detailed characterization of both the spectral and spatial properties of the emission (Bouchet et al. 2010). The SPI map of the 511 keV signal (left panel of Figure 1.14) shows a bright emission from the Galaxy bulge and a fainter component emanating from the disk. However, the origin of this emission remains unknown. A primary challenge lies in identifying the sources of Galactic positrons. Several candidates have been proposed, including nucleosynthesis products (e.g., ^{26}Al, ^{44}Ti, ^{56}Co), individual sources such as low-mass X-ray binaries, hypernovae, pulsar winds, microquasars, neutron star mergers, and even exotic scenarios involving dark matter annihilation or primordial black holes. However, none of these candidates can fully account for the observed morphology of the 511 keV emission. Additionally, the initial energy of the positrons at the moment of their creation, as well as their propagation distances within the Galaxy, remains poorly understood. The upcoming COSI mission will map the 511 keV line with unprecedented sensitivity (1.2×10^{-5} photons cm^{-2} at 0.511 MeV) and spectral resolution (6 keV at 0.511 MeV), offering crucial insights into the origin of Galactic positrons (right panel of Figure 1.14). It aims to potentially identify individual sources, study the annihilation mechanisms, and constrain the positron injection energies.

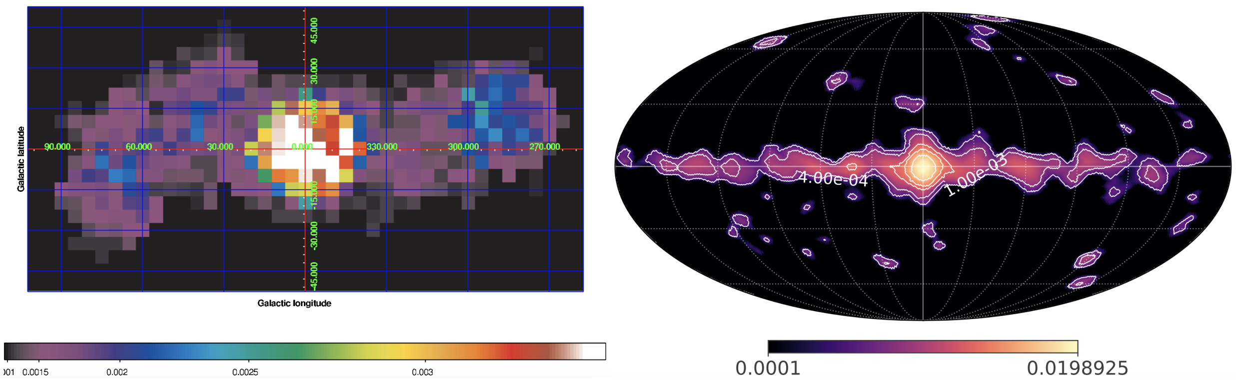

Figure 1.14 – On the left, SPI map of positron annihilation (511 keV) from the Galaxy (Bouchet et al. 2010). On the right, reconstructed 0.511 MeV image with COSI (Yoneda et al. 2025). Pixel size is $5° \times 5°$ for both images.

Nucleosynthesis in the Galaxy

Nucleosynthesis models predict that ^{26}Al and ^{60}Fe are produced throughout the lifetimes of massive stars and during core-collapse supernova (CCSN) explosions. ^{26}Al has a half-life of 7.2×10^5 years and decays to ^{26}Mg, emitting a characteristic 1.809 MeV gamma-ray photon. Mapping the 1.809 MeV line provides a powerful tool for tracing recent CCSN activity and massive star winds in the Milky Way. Similarly, ^{60}Fe decays to ^{60}Co with a half-life of 2.6×10^6 years; ^{60}Co then decays rapidly to ^{60}Ni, emitting two gamma-ray lines at 1.173 and 1.333 MeV. While ^{26}Al is released throughout all stages in the evolution of massive stars, ^{60}Fe is produced only by core-collapse supernovae. Therefore, comparing the relative fluxes of the two emission lines provides valuable insights into the evolutionary stages of different portions of the Galaxy. COMPTEL mapped the ^{26}Al emission in the Galactic plane with 3.8° FWHM angular resolution (Diehl et al. 1995). SPI measurements provide limits on the ^{60}Fe/^{26}Al ratio (Bouchet et al. 2015), but only averaging over all Galaxy, without providing information on its spatial distribution. COSI will map the ^{26}Al and ^{60}Fe Galaxy emission with improved resolution and sensitivity, enabling a more in-depth understanding of the stellar evolution in our Galaxy.

Extreme environments

Polarization measurements represent a powerful tool to investigate emission mechanisms and source geometries (e.g., magnetic field, accretion disk, and jet) of extreme environments, like GRBs, black holes (BHs), and active galactic nuclei (AGN). Thanks to its high polarization sensitivity, COSI will measure the polarization of a large number of GRBs with a minimum detectable polarization of less than 50%, shedding light on the physical processes behind their production. Many BH systems show a non-thermal emission which extends into the COSI bandpass. Polarization measurements of the gamma emission from Galactic BHs will enable a better understanding of the physics in such extreme environments. COSI will also utilize polarization measurements to investigate jet composition and emission mechanisms in AGN. In the case of blazars, COSI aims to resolve the long-standing debate over whether their high-energy emission is of hadronic or leptonic origin. A high degree of polarization ($\sim 60\%$) would strongly support a hadronic scenario (Zhang and Böttcher 2013).

Multi-messenger astronomy

The coordinated observation and interpretation of multiple signals received from the same astronomical event marked a major breakthrough in modern astronomy. Astrophysical events often involve the interplay between complex physical processes, with the consequent emission of different "messengers": electromagnetic radiation, gravitational waves (GWs), neutrinos, and cosmic rays. One of the most significant developments in the past decade has been the simultaneous detection of gravitational waves (GWs) and gamma-ray bursts (GRBs) from neutron star mergers (Abbott et al. 2017). Prompt detection and accurate localization of GRBs from these events are crucial for enabling rapid follow-up observations by other instruments, thereby offering deeper insights into the underlying physics of these extreme astrophysical phenomena. COSI will provide $\sim$ degree localizations for GRBs within its instantaneous field of view of $\sim 25\%$ of the sky. Moreover, thanks to its full coverage of the unocculted sky, the COSI ACS will enhance the detection rate of GRBs and enable rapid reporting times within 60 minutes (see Section 1.6).

1.6 The COSI Anticoincidence System

The COSI GeDs are surrounded on four lateral sides and underneath by active ACS shields, defining an instantaneous field of view (FoV) for the GeDs of $\sim$25% of the sky. In the 2016 balloon configuration (Kierans et al. 2016), the ACS was built from CsI panels, each coupled to two PMTs for light collection. In the NASA SMEX configuration (COSI-SMEX), shown in Figure 1.15, the ACS consists of 22 BGO crystals: three on each lateral side and ten positioned beneath the GeDs (Figure 1.16). The crystal sizes are reported in Table 1.2. BGO was preferred over CsI mainly for its high density and non-hygroscopic nature: an equivalent shielding with CsI would be achieved with crystals 1.3 cm thicker than the BGO design of the COSI balloon (Patel et al. 2025). Instead of PMTs, each BGO module is coupled with a 3×3 array of Onsemi J-Series 6x6 mm SiPMs, positioned on the BGO smaller edges. When a photon or particle releases energy in a BGO crystal, it produces scintillation photons that are detected by the SiPMs, with a threshold of $\sim$ 80 keV. The 3×3 SiPM configuration was found to provide the best compromise between minimizing electronic noise and maximizing signal output, compared to configurations with a larger number of SiPMs (Patel et al. 2025). The BGO crystals are wrapped with reflective layers to reduce optical photon losses and enhance the signal collected by the SiPMs. The SiPM outputs are combined into six groups, each forming an independent shield module connected to a dedicated ASIC (Application-Specific Integrated Circuit) for signal processing. After a hit in a BGO crystal is recorded, the corresponding ASIC has an expected dead time of about 1.5 μs[4], during which it cannot register subsequent events.

The ACS serves both active and passive attenuation of background events in the GeDs. Photons entering the instrument from outside the GeD field of view are effectively detected by the BGO shields, due to their high stopping power. Moreover, photons that are not fully absorbed in the GeDs cannot be reliably reconstructed. These photons would likely be absorbed in one of the surrounding BGO crystals, causing the ACS to raise a veto to reject such problematic events. The ACS also plays a key role in the detection and localization of GRBs and other transients outside the FoV of the GeDs, acting as a stand-alone gamma-

[4] This value was provided by Parshad Patel (NRL, private communication).

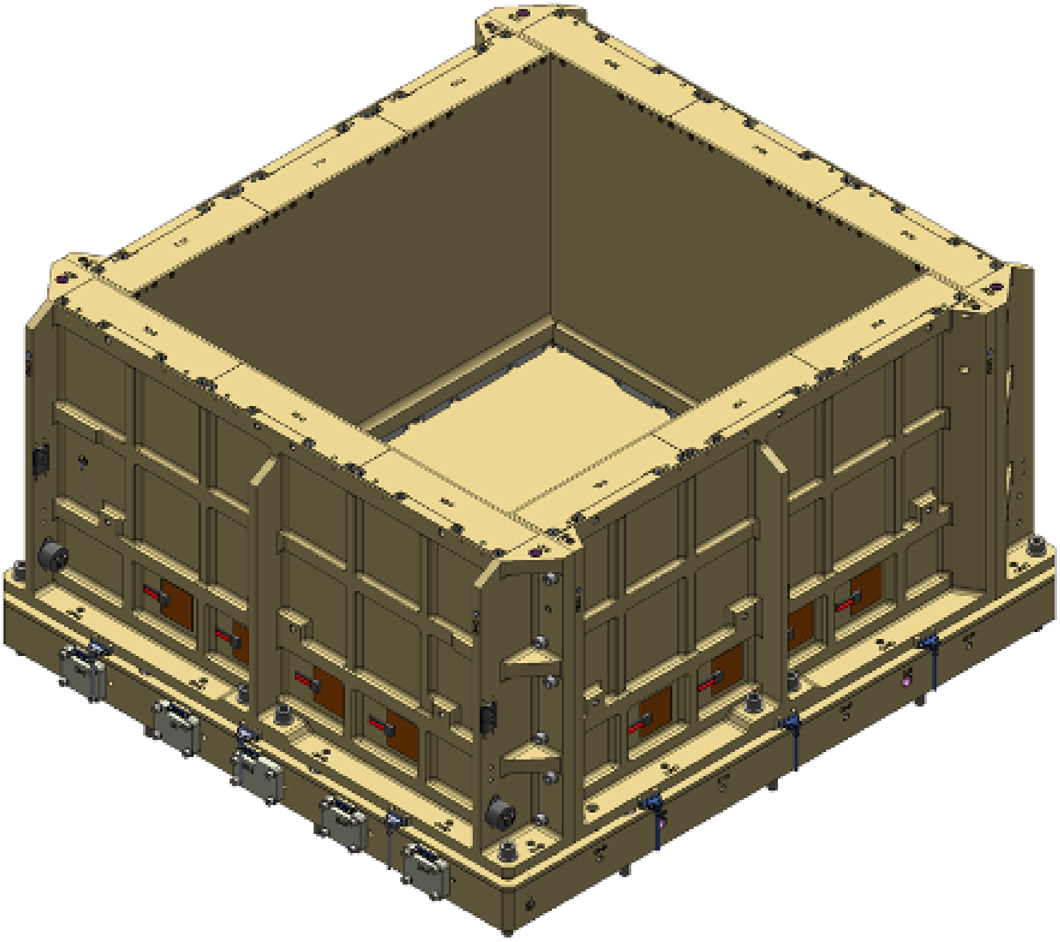

Figure 1.15 – COSI ACS design. The ACS is composed of four lateral walls, each containing 3 BGO crystals, and one bottom wall, containing 10 BGO crystals. Each wall is enclosed in an aluminum housing. Image provided by the COSI team.

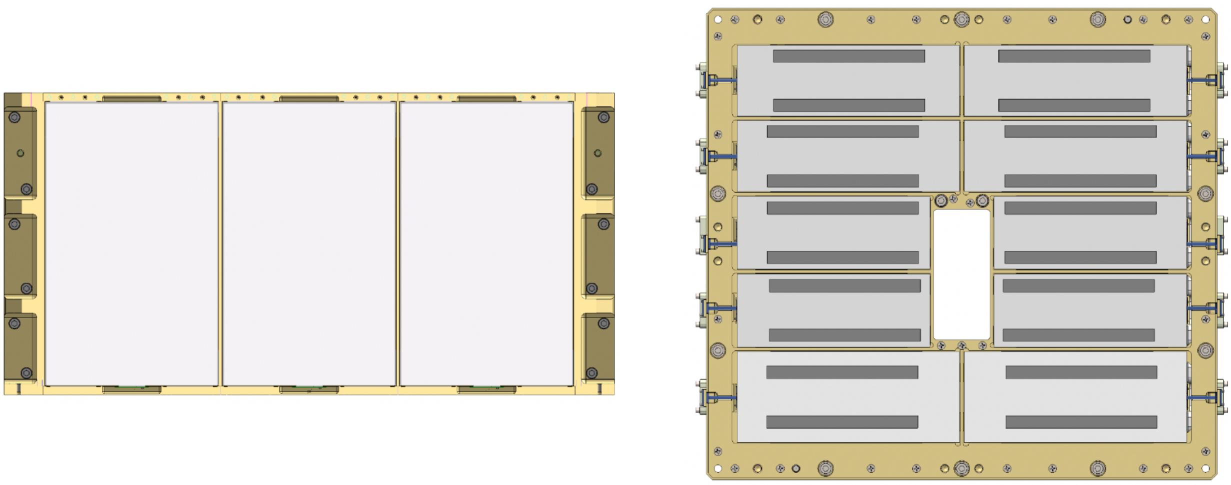

Figure 1.16 – Lateral ACS wall with 3 BGO crystals (left) and bottom ACS wall with 10 BGO crystals (right). The opening in the bottom wall is designed to enable mechanical integration with the spacecraft instrument deck. Image provided by the COSI team.

Table 1.2 – BGO crystal sizes (Patel et al. 2025). The 22 crystals are grouped into five classes according to their size: two classes corresponding to the lateral X and Y sides, and three classes corresponding to the bottom (Z) side. All quantities are specified with a tolerance of ± 0.2 mm.

	X [cm]	**Y [cm]**	**Z [cm]**
BGO X-crystal	19.4	11.8	2.30
BGO Y-crystal	19.4	11.0	2.30
BGO Z-crystal A	19.8	6.25	2.40
BGO Z-crystal B	17.1	6.45	2.40
BGO Z-crystal C	19.8	8.05	2.40

ray transient monitor. The COSI ACS data will be collected into 6 different independent modules, corresponding to each ACS ASIC. Each of them will provide two energy channels, below and above 2 MeV, for a total of 12 channels of BGO data. An on-board trigger system monitors the ACS ratemeters and, when a statistical significance threshold is exceeded, the data is downlinked via the Tracking and Data Relay Satellite System (TDRSS, Teles et al. 1995). By analyzing the relative count rates among different panels, the ACS can provide localization of GRBs outside the GeD FoV, or provide a quick localization before the GeD signal is reconstructed, allowing for quick follow-up observations by other telescopes. The downlinked ACS data passes through the fast transient pipeline to obtain a first source type identification, significance, and localization of the detected source, and a GCN notice is generated within 60 minutes after the initial trigger. Subsequently, a comprehensive transient pipeline performs a reanalysis of the transient with all data downlinked to the ground stations during the orbital passages and generates a GCN within a few hours of the trigger.

1.7 Enhancing the ACS simulation accuracy: from calibrations to GRB detection

Accurately predicting the ACS performance, both as a veto detector and as a transient monitor, is essential for assessing its scientific contribution (e.g., GRB detection and localization) and the overall mission capabilities. Monte Carlo simulations based on the Geant4 toolkit are the standard approach for this purpose. To ensure reliable and accurate results, modeling the COSI ACS performance requires accounting for several effects intrinsically linked to the scintillation processes within the BGO crystals. In this context, dedicated simulations including the optical physics provide a powerful and flexible tool to evaluate these effects in detail. The Geant4 toolkit, the Geant4-based framework used for COSI (MEGAlib), and the key parameters defining the ACS performance are discussed in Chapter 2.

The explicit simulation of optical physics is particularly challenging for two main reasons. First, it requires a large computational effort due to the vast number of optical photons produced during the scintillation process (for BGO, approximately $8-10$ photons per keV, as discussed in Section 1.3.1). Second, it involves numerous parameters that govern the physical interaction of the scintillation light with the surrounding materials. To address the first issue, optimized simulations exploiting parallel execution on High-Performance Computing (HPC) infrastructures are essential to make the problem computationally feasible. This aspect is the focus of Chapter 3, where we introduce the simulation framework BoGEMMS-HPC and present its performance results.

The second challenge, namely the complexity of configuring the optical simulation parameters, must be addressed by directly comparing the simulations with both theoretical models and experimental data, in order to validate their capability to reproduce realistic results and to constrain the relevant parameters. The verification of the optical physics simulation and the simulation benchmarking using calibration measurements of a lateral ACS wall are presented in Chapter 4 and Chapter 5, respectively.

Once validated, the simulations can be used to derive a correction matrix, which encodes the true ACS response and provides energy corrections based on the true deposited energy and the interaction position within the BGO crystals. The development of this correction matrix is the subject of Chapter 6.

The application of the correction matrix to realistic GRB simulations, while exploring the detection capability of the COSI ACS as a gamma-ray transient monitor, is finally discussed in Chapter 7.

Chapter 2

COSI ACS simulations and calibrations

The ever-increasing complexity of the technologies and instruments involved in astrophysical observations requires simulation tools with high predictive accuracy. In this perspective, Geant4 is the standard tool of choice for many past, current, and future X-ray and gamma-ray missions, such as Chandra, XMM-Newton, AGILE, Fermi, NewAthena, COSI, NewASTROGAM, and THESEUS. Achieving reliable simulation results demands precise treatment of particle interactions with all spacecraft components, along with accurate modeling of the design and physical properties of the materials. For COSI, the Detector Effects Engine introduces realistic detector performance, enabling simulations to reflect real measurements. Developing a dedicated Detector Effects Engine for the COSI ACS enhances the predictive power of its performance both as a veto detector and as a gamma-ray transient monitor. The ACS performance, characterized by parameters such as quantum detection efficiency, energy resolution, and others, is intrinsically determined by the scintillation processes in the BGO crystals and the electronic noise of the SiPM readout system. Furthermore, these parameters are expected to vary with the interaction position inside the crystal, as the average path length of scintillation photons to the SiPMs changes accordingly. In this context, the Geant4 capability to simulate scintillation processes can be exploited to predict the ACS performance over a wide energy range and with spatial detail that would be challenging to obtain through laboratory measurements alone. Nevertheless, laboratory calibration measurements remain essential to fine-tune (i.e., *benchmark*) the simulation, ensuring that the simulated performance is representative of the real ACS behavior.

This chapter is organized as follows. Section 2.1 provides an overview of the Geant4 toolkit. Section 2.2 presents MEGAlib, the Geant4-based simulation framework adopted as the standard for COSI. Section 2.3 describes the Geant4 simulations performed for COSI, with particular attention to the ACS. Section 2.4 introduces the concept of the Detector Effects Engine (DEE) for COSI, focusing on the ACS and its performance parameters. Section 2.5 discusses how the DEE for the COSI ACS can be modeled through Geant4 simulations. Finally, Section 2.6 outlines the COSI data analysis pipeline and the main software tools employed.

2.1 Geant4

Geant4 (Agostinelli et al. 2003; Allison et al. 2006; Allison et al. 2016) is an object-oriented Monte Carlo simulation toolkit for the simulation of particle transport through matter. It is an open-source library, written in C++ (with available Python bindings), that the user can use to develop a stand-alone program. Originally developed at CERN for the simulation of high-energy experiments, it is widely used for astrophysical applications, such as sensitivity evaluation, design optimization, instrumental background estimation, and detector response modeling. The key domains of the simulation framework include geometry and materials, particle interaction with matter, tracking, hits, event management, and visualization. Each domain is associated with a set of classes that the user can handle and customize. The toolkit allows the user to build a three-dimensional virtual model with several components of different shapes and materials, where particles and their secondaries are tracked along their paths. The user can define sensitive volumes that record information about the passage of particles and their energy deposits. The particles injected in the simulation are called primary particles, which also define an *event*. Within an event, primary particles can produce secondary particles through processes such as pair production or the ejection of a bound electron from an atom. These secondary particles are then tracked alongside the primaries. The user can simulate multiple independent events. Geant4 offers a comprehensive suite of physics processes to model particle interactions, which can be implemented and customized by the user. Besides high-energy physics, Geant4 handles lower-energy processes involving optical photons such as scintillation and Cherenkov radiation.

2.2 MEGAlib

MEGAlib (Medium-Energy Gamma-ray Astronomy library; Zoglauer et al. 2008) is an object-oriented C++ software library designed for simulating and analyzing Compton and pair telescopes. It provides a comprehensive set of tools for a complete data analysis pipeline, from simulation to high-level analysis:

- **Geomega** for the construction of the virtual telescope;
- **Cosima** for the Geant4-based simulation;
- **Revan** for the reconstruction of the events;
- **Mimrec** for the high-level data analysis (event selections, image reconstruction, polarization analysis, etc.).

2.2.1 Geomega

Geomega (Geometry for MEGAlib) is a tool that allows the user to construct and visualize the geometry and the detector. The description of all materials, volumes, and detectors is specified through plain text files using keywords. The `.geo.setup` file is the central input for Geomega, while `.geo` and `.det` files are used to describe subvolumes and detector characteristics, respectively, and are included in the main file. The user can position, rotate, enlarge, and shrink a set of volumes (e.g. boxes, spheres, tubes, trapezoids, polygons), specify their materials and place them within mother volumes. In the `.det` files, the user can set typical gamma-ray detectors such as scintillators and strip/pixel detectors (i.e., with and without depth resolution) and specify their energy resolutions, position resolutions, trigger thresholds, noise thresholds, number of strips/pixels, etc. Moreover, Geomega enables the specification of customized trigger and veto criteria, such as the number of hits in the same detectors or the number of detectors required to raise a trigger, and the type of hit if above the threshold (trigger or veto). Geomega allows for the selection between Gaussian and Lorentzian noise for the energy resolution, the implementation of overflows (energy deposits exceeding a certain value are accumulated at a specified energy), and discretization (summation of the energy deposits in spatial pixels).

2.2.2 Cosima

Cosima (Cosmic Simulator for MEGAlib) is a Geant4-based simulator. The simulation input is a plain text file (`.source`), where the user specifies the desired configuration via keywords. These include options such as the output mode (e.g., storing all or only part of the information), the choice of physics lists, source geometry (point-like or extended, isotropic or user-defined emission, etc.), spectral shape (monoenergetic, power law, broken power law, user-defined, etc.), particle type, stopping criterion (simulation time or number of triggers/events), and several other parameters. Cosima reads the Geomega files and converts the Geomega geometries into Geant4 volumes. The output is written in a `.sim` file, an ASCII file containing hits in sensitive detectors and, if specified by the user in the `.source` file, information about the ideal interactions in either active or passive volumes. It is also possible to retrieve from the output file the ideal energy deposits (i.e., without any noise), and apply the corrections in a post-simulation step (e.g., with Nuclearizer).

2.2.3 Revan

Revan (Real Event Analyzer) performs the event reconstruction. It analyzes the hits from the `.sim` file and identifies the interaction processes (photo-absorption, Compton scattering, pair production). For Compton events, the algorithm is composed of 4 steps: *clustering* (which combines adjacent hits into one single hit), *tracking* (which searches for pair events, Compton electron tracks, showers, and muons), *Compton sequence reconstruction* (which determines the sequence of Compton interactions), and *decay detection* (which searches for beta-decay structures and rejects them). The output is written to an ASCII file (with extension `.tra`) which lists the parameters (positions, directions, energy) of the interactions.

2.2.4 Mimrec

Mimrec (MEGAlib image reconstruction) performs the final high-level data analysis step. Its task is to reconstruct the properties of the source from the interaction processes (e.g., Compton scattering) listed in the `.tra` file. It allows the user to perform event selections on all relevant parameters, including event type, energies, scattering angles, interaction distances, number of interactions, and several more. Moreover, it is equipped with several diagnostic tools to visualize the data, such as energy spectra, ARM distributions, polarization, and many others. The principal (and most challenging) task of Mimrec is, however, the image reconstruction. For Compton scattering events, the origin of the photon can only be restricted to a cone, which complicates the process of image reconstruction. The approach is based on a list-mode likelihood image reconstruction algorithm, i.e., analyzing data event-by-event and searching for the image that maximizes the likelihood of observing the recorded events. The image reconstruction can be performed in various coordinate systems (spherical, Galactic, 2D or 3D Cartesian coordinates).

2.3 MEGAlib simulations of the COSI ACS

As discussed in Section 2.1, Geant4 simulations are a key tool in high-energy astrophysics missions, as they enable accurate performance predictions, guide instrument design, and support the development of detector response files. For COSI, most of the Geant4 simulation efforts using MEGAlib have so far focused on the GeDs, including, for instance, benchmarking simulations for the balloon design (Sleator et al. 2019b), background characterization (Gallego et al. 2025), and scientific studies such as mapping Galactic positron annihilation (Siegert et al. 2020). In addition, numerous simulations have been carried out for the yearly released data challenges (Spectrometer and Collaboration 2025), providing data for imaging, GRB localization and spectral fitting, background estimation, polarization analysis, detector response files, and others. On the other hand, dedicated simulations of the COSI ACS are now being carried out for multiple purposes. A primary goal is to estimate the background rate recorded by the ACS, which is essential for evaluating its effectiveness as a veto detector and, when simulated in combination with the GeDs, for assessing the overall sensitivity of the mission. Figure 2.1 illustrates the ACS background obtained from a 24-hour simulation performed by the COSI team. The plot shows the counts over time recorded by each of the six ACS panels, where each gap corresponds to a passage through the South Atlantic Anomaly. In addition, the ACS can be simulated as a stand-alone detector to investigate its capability to detect and localize GRBs. Figure 2.2 presents an example of simulated GRB light curves as measured by the six ACS modules. Dedicated ACS simulations were also carried out to quantify the GRB localization performance (Parmiggiani et al. 2025).

2.4 Detector Effects Engine

Given the importance and variety of simulation applications, it is essential to have simulations that include realistic imperfections and distinctive characteristics of the detectors. This task is carried out by the Detector Effects Engine (DEE), a dedicated pre-processing stage that introduces realistic detector effects, such as energy resolution and thresholds, and converts physical quantities into detector-level observables (e.g., ADC units, strip numbers, and detector IDs). This step is performed by the software Nuclearizer (Zoglauer et al. 2021), a calibration tool specifically developed for COSI. A prototype of DEE has been developed for the COSI balloon GeDs (Sleator et al. 2019a). It applies realistic effects such as charge sharing, charge loss, crosstalk, threshold, dead strips, and dead time. Physical quantities are then converted into detector units: (x, y, z) positions are mapped to detector and strip IDs, while energies are converted into pulse heights.

Unlike the GeDs, no dedicated DEE module had been developed for the ACS prior to this

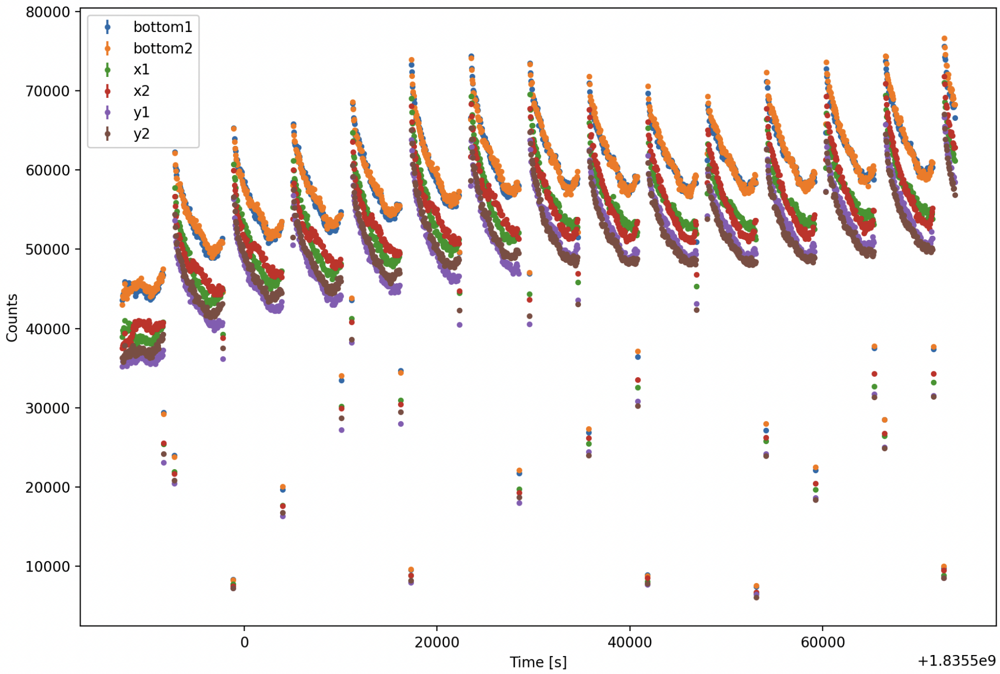

Figure 2.1 – Simulated 24-hour light curve of the background as measured by the 6 panels of the COSI ACS. The periodic variation in the count rate is due to the satellites passages through the South Atlantic Anomaly (the actual SAA passages are removed from the simulated data). The simulation was conducted by Savitri Gallego of the COSI team.

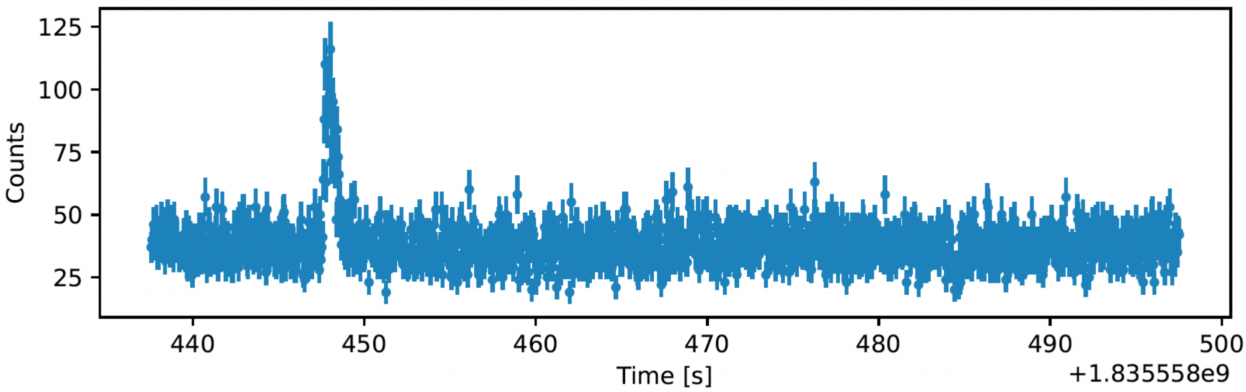

Figure 2.2 – Simulated GRB light curve as seen by a lateral panel of the COSI ACS.

work. To realistically simulate the ACS, several key parameters must be considered: the quantum detection efficiency, the effective area, the optical light collection efficiency, the energy resolution, and the energy threshold. Another important factor is the dead time, defined as the interval following an interaction in a BGO crystal during which subsequent events cannot be recorded. As for GeDs, after modeling these parameters to enhance the realism of the simulations the DEE must also apply an *inverted calibration*, namely a conversion from simulated physical quantities into detector units. In the case of the ACS, this corresponds to mapping the energy deposited in the BGO crystals to ADC channels (i.e., pulse heights). The treatment of dead time and inverted calibration, however, is beyond the scope of this PhD project.

2.4.1 ACS performance parameters

Quantum detection efficiency The *quantum detection efficiency* is defined as the fraction of detected photons out of the total incident radiation. It can be a *total* efficiency, if all detected photons are considered regardless of their energy released in the material, or a *photopeak* efficiency, if only photons depositing their full energy are considered as detected. The latter is particularly important if the aim of the detector is to measure the energy of the incoming radiation. The detection efficiency can be either an energy-integrated quantity or an energy-dependent value. For the latter case, the photopeak efficiency ϵ at a certain energy E is defined as the ratio of detected photons N_{det} with energy deposit E over the number of incident photons N_{in} with initial energy E:

$$\epsilon(E) = \frac{N_{det}(E)}{N_{in}(E)} \ . \tag{2.1}$$

A related quantity is the *effective area* A_{eff}. It is defined as the geometrical area A_{geom} exposed to the incoming radiation, multiplied by the detection efficiency ϵ:

$$A_{eff} = \epsilon \times A_{geom} \ . \tag{2.2}$$

The effective area represents the area that would be required for an ideal detector (i.e., with 100% efficiency) to produce the same count rate as the real detector under the same incident flux. The quantum detection efficiency and effective area are already correctly taken into account in the MEGAlib simulation, as they mostly depend on the electromagnetic interaction between photons and the BGO crystal.

Optical light collection efficiency The *optical light collection efficiency* is defined as the fraction of the optical photons detected out of the total generated photons. This property depends on several parameters, such as the BGO absorption capability, the reflectivity of the reflective layers wrapping the BGO crystal, and the configuration of the readout system. It may depend on the generation position of the scintillation light within the crystal. In fact, the average path length that optical photons must travel to reach the SiPMs, and their corresponding probability of being detected, varies within the BGO volume (see Figure 2.3). As a result, the same energy deposit may produce, in average, more or less detected optical photons, depending on the interaction position within the BGO crystal. This process

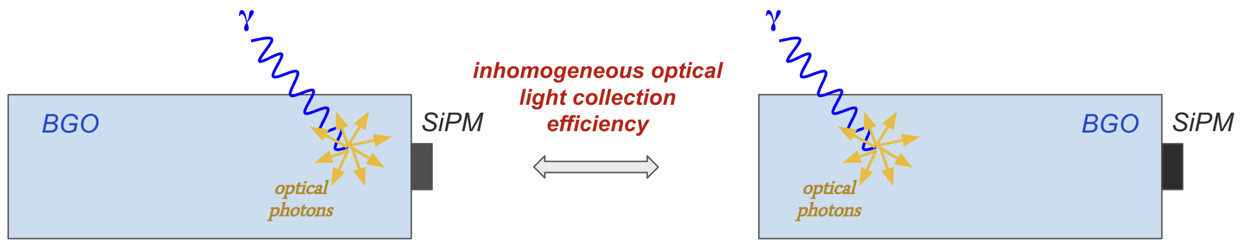

Figure 2.3 – Schematic illustration of a high-energy photon (γ) hitting in different positions a BGO crystal coupled to SiPMs. As the distance from the SiPMs changes, the optical light collection efficiency may be not the same for the two situations.

has an impact on the determination of the energy stored in the crystal. During the energy calibration procedure, the BGO crystal is fully irradiated by radioactive sources to define the relation between deposited energy and the SiPM electronic pulse (see Section 5.1.3). Spatial non-uniformity in the optical light collection means that, when calibrating from pulse to energy using the averaged relation, the same true energy deposited in different positions results in a different measured energy. This bias, present when calibrating real on-ground or in-flight ACS data, should also be implemented in its simulations.

Energy resolution The *energy resolution* defines the ability of a scintillator to distinguish between photons of different energies. When photons emitted by a source interact with the detector, their original energies are reconstructed by the instrument, resulting in a measured energy distribution. For a source emitting mono-energetic photons, the ideal response would be a delta function centered at the true energy. In reality, due to the intrinsic limitations of the detector, the measured distribution exhibits a finite width. This distribution, known as the energy response function, reflects the detector energy resolution and quantifies the uncertainty in reconstructing the photon energies. Defining the full width at half maximum (FWHM) of the response function as the width of the distribution at half of its maximum, the energy resolution R is defined as

$$R = \frac{\Delta E}{E_0} = \frac{\text{FWHM}}{E_0} \,, \tag{2.3}$$

where E_0 is the peak centroid of the distribution. The energy resolution is typically expressed as a percentage. Commonly, it is determined in laboratory using radioactive sources emitting mono-energetic lines. The typical spectrum measured by a scintillator, for a particular radioactive source (^{137}Cs) is shown in Figure 2.4. The spectrum exhibits a photopeak (where photons are fully absorbed in the detector) and a continuum tail at lower energy (where photons deposit part of their initial energy through scattering in the detector or surrounding material). The FWHM of the photopeak, normalized to the line energy of the radioactive isotope, represents the energy resolution at that particular energy. The energy resolution of a scintillator depends on three factors: *i)* the electronic noise of the SiPM readout, *ii)* the Poissonian statistical fluctuations in the number of generated and detected optical photons, and *iii)* the overall inhomogeneity of the optical light collection efficiency across the BGO crystal. These factors introduce non-trivial energy dependence for the energy resolution, which will be discussed in more detail in Chapter 5. MEGAlib applies a Gaussian constant noise, without taking into account any of these effects.

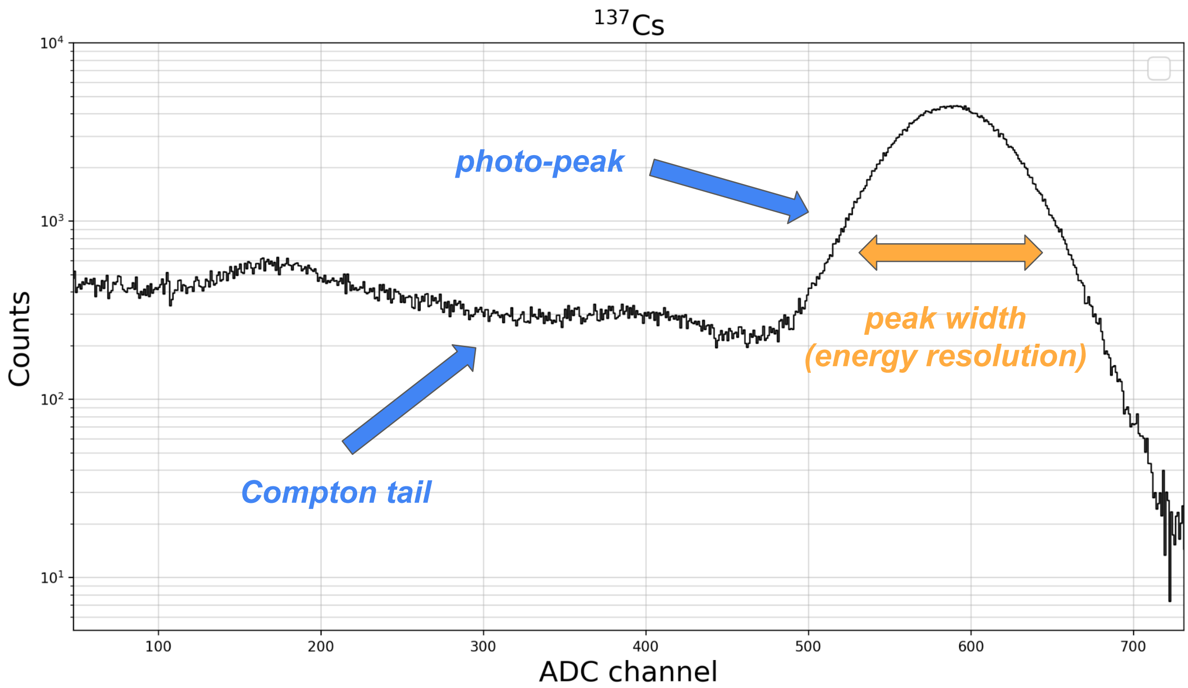

Figure 2.4 – Background-subtracted ^{137}Cs spectrum measured by a COSI ACS BGO detector. The x-axis is expressed in ADC channels, representing the pulse height proportional to the deposited energy.

Energy threshold Another key parameter is the *energy threshold*, defined as the minimum energy detectable by the ACS. This is related to the minimum electric pulse (or, equivalently, the minimum number of optical photons) that can be detected. The mission requirement is 80 keV. The threshold depends on the crystal size, BGO properties, and SiPM electronic noise. Moreover, if the optical light collection efficiency varies within the BGO volume, the energy threshold is also expected to be spatially non-uniform.

2.5 Modeling the ACS performance through Geant4 simulations

All the parameters presented in Section 2.4 must be modeled and properly incorporated into the DEE for the ACS to enhance the simulation accuracy and ensure fidelity to real measurements. Some of these parameters, such as the quantum detection efficiency and effective area, are already included in the MEGAlib simulation, whereas others need to be determined separately. Moreover, the energy resolution can actually influence the detection efficiency (and consequently the effective area), as it may change the number of photons falling below and above the energy threshold. Typically, all these effects are constrained through dedicated laboratory measurements of the instrument. However, obtaining them with high spatial resolution and across wide energy ranges is particularly challenging and hardly achievable with laboratory measurements alone.

The strategy adopted in this work is to exploit the flexibility of Geant4 simulations and its capability to simulate the scintillation process to evaluate the ACS performance. In fact, with Geant4 we can explore the ACS response with high spatial resolution and probe energy ranges that are otherwise difficult to access experimentally. While including optical processes increases the realism of the simulations, it also introduces significant challenges.

The first issue is the higher computational cost. Scintillation simulations require tracking thousands of optical photons: for BGO, about 8000 photons are produced per MeV. As an example, simulating 500 photons at 511 keV impinging on a BGO crystal takes roughly one hour on a single CPU core when optical physics is enabled, whereas the same simulation without scintillation completes in less than a minute. Efficient simulation strategies are therefore mandatory, including the use of large-scale computational resources to enable parallel executions. In this work, we will rely on BoGEMMS-HPC, a Geant4-based framework optimized for high-performance simulations on large infrastructures (see Chapter 3).

The second issue is the increased complexity of the simulations. Including scintillation in the simulation requires specifying many parameters that govern the interaction of optical photons with surrounding materials, such as the refractive index, absorption length (mean free path before absorption), surface roughness, reflectivity, reflection model, and others. Some of these parameters can be known from the literature or manufacturer information; others require benchmarking through comparison with laboratory measurements. This procedure of *simulation benchmarking*, which is presented in Chapter 5, provides a fundamental step to ensure that the simulation results are representative of the real detector behavior while constraining parameters that are otherwise unknown. The validated simulation framework can subsequently be used to model the ACS performance parameters in a dedicated correction matrix to be included in the DEE, with the ultimate goal of enhancing the realism of the simulation results.

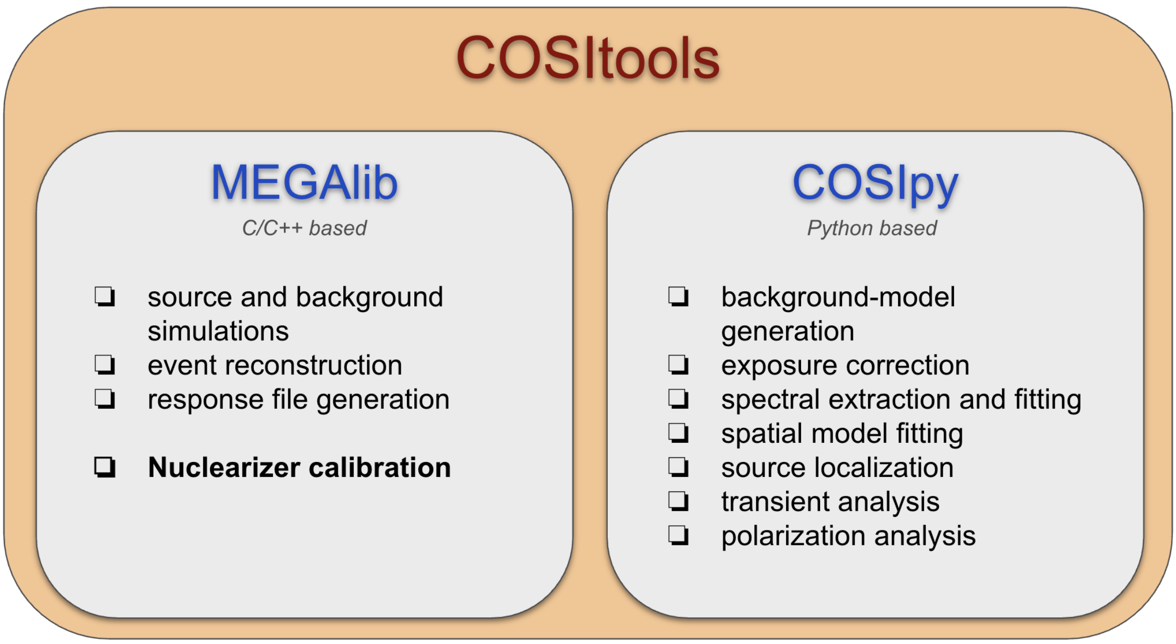

Figure 2.5 – Set of tools and software composing COSItools.

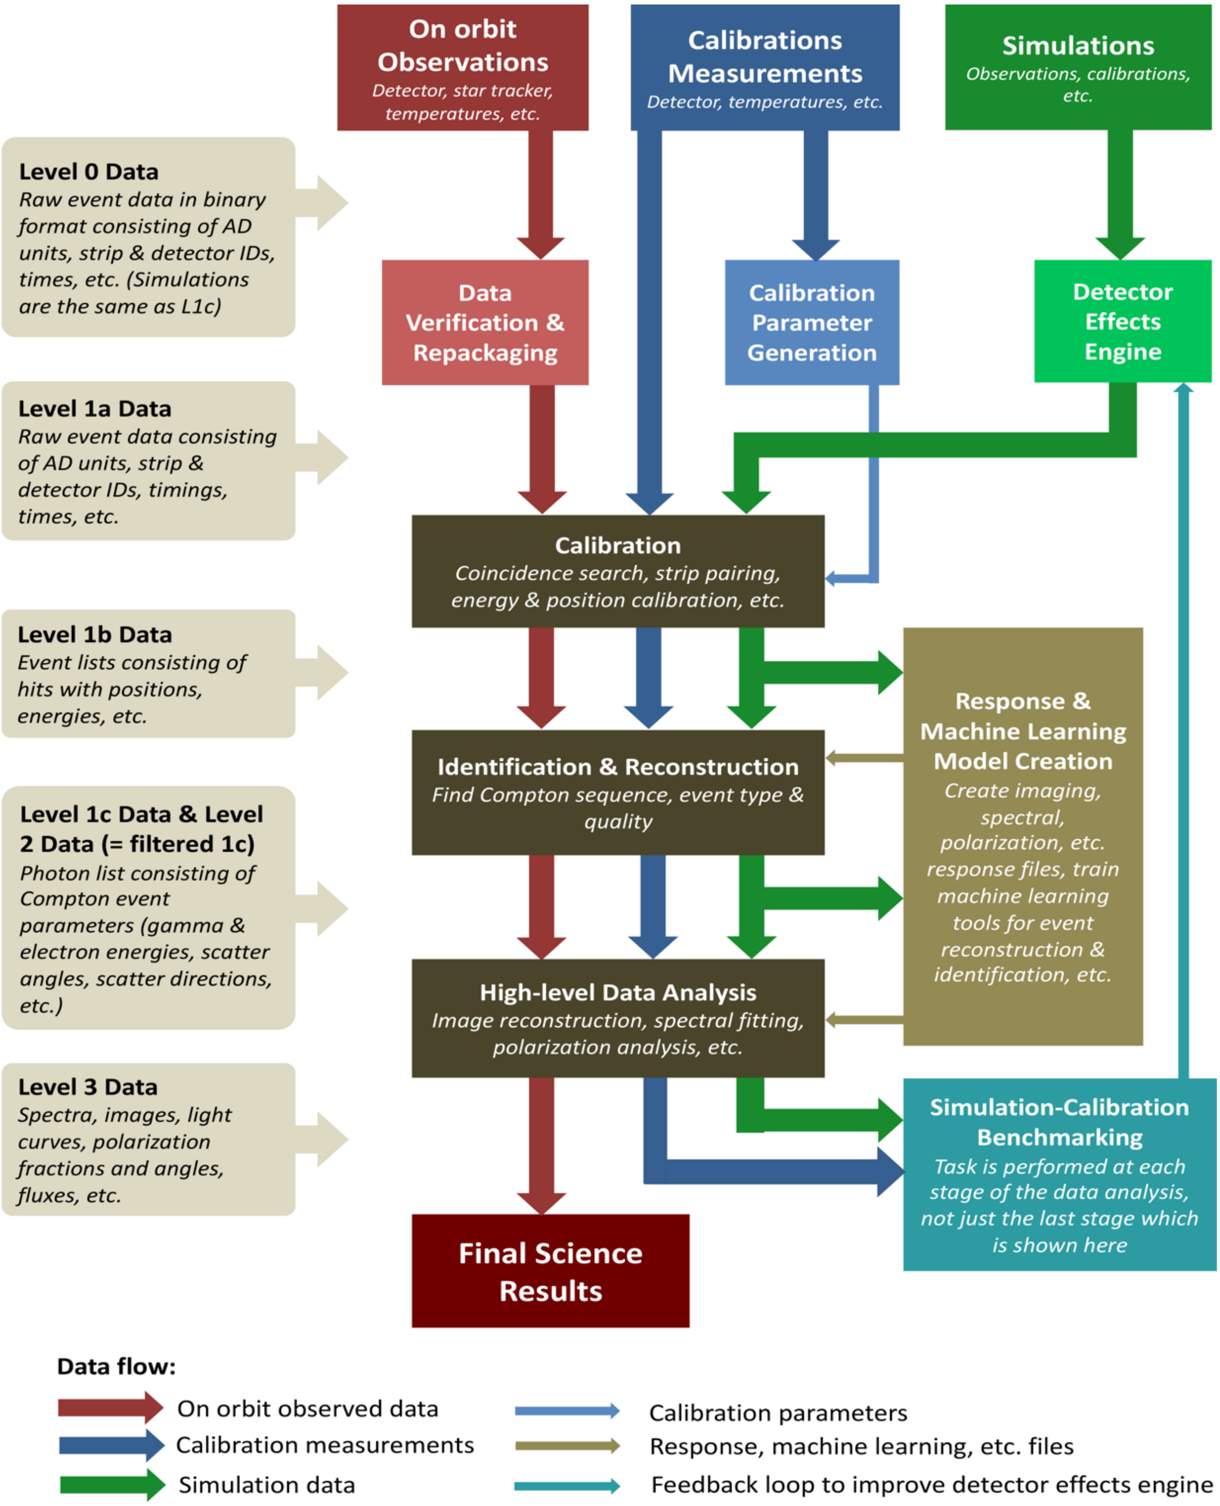

Figure 2.6 – Diagram of the COSI data flow and data products from on orbit observations, calibration measurements, and simulations to the final science results (Zoglauer et al. 2021).

2.6 The COSI data analysis pipeline

The path from row detector data to final science products (energy spectra, all-sky images, light curves, etc.) requires multiple processing steps. The COSI pipeline utilizes a set of software components which are collected under the name *COSItools*. The main software components constituting COSItools are MEGAlib (for simulations and event reconstruction), Nuclearizer (for calibrations and DEE), and COSIpy (Martinez-Castellanos et al. (2023), for high-level data analysis), as illustrated in Figure 2.5.

The COSI data analysis pipeline is shown in Figure 2.6 (Zoglauer et al. 2021). The COSI pipeline follows the workflow for three types of data: on orbit observations, on ground observations, and simulations. As discussed in Section 2.4, simulations must pass through the DEE to match real data. After this preliminary step, the three data types pass through the same analysis pipeline. At Level 1a, the data contain detector-specific quantities such as ADC units, strip IDs, and detector IDs. The first major step is calibration, which converts these detector units into physical quantities (e.g., from ADC channels to deposited energy). The output of this step, Level 1b, consists of hits with assigned energies, positions, and related parameters. The next stage is event reconstruction. For GeD data, this includes reconstructing the Compton interaction sequence, which requires determining the positions and energies of the hits as well as the Compton scatter angle. The reconstructed output is defined as Level 2 data. Finally, the data enter the high-level analysis stage, which includes image reconstruction, spectral fitting, and polarization studies. The resulting Level 3 products consist of the final science data: energy spectra, light curves, images, polarization fractions, and related observables.

Chapter 3

Geant4 simulations on HPC environments

The increasing complexity of the technologies used to observe the high-energy universe requires modern, large-scale simulations and presents new challenges in the management of ever-larger datasets, as well as higher demands for storage and computational resources. Next-generation High-Performance Computing (HPC) techniques and efficient I/O management are therefore essential for Geant4-based simulations. A clear example of this is the simulation of the COSI Anticoincidence System: tracking thousands of optical photons produced within the scintillation crystals is unfeasible using standard techniques. However, the inclusion of the optical physics is essential to produce realistic simulations and accurately modeling the ACS performance. During my PhD, we addressed these issues by implementing multi-threading and multi-process capabilities on BoGEMMS-HPC (Bulgarelli et al. 2012; Fioretti et al. 2014; Fioretti 2025), a Geant4-based simulation framework tailored for HPC environments. To evaluate the performance gains, we performed dedicated tests on local and HPC architectures, such as the Leonardo supercomputer. This work was carried out within the framework of the newly established Italian Research Center on High-Performance Computing, Big Data, and Quantum Computing (ICSC), which funds my PhD project. The optimization and acceleration of the code represent a fundamental step that enables extensive COSI ACS simulations and supports the work presented in Chapter 5 and Chapter 6. The code is currently released in a public repository and is based on Geant4 v11.1.

This chapter is organized as follows. Section 3.1 provides an overview of High-Performance Computing, the main parallelization schemes, and the current status of parallel Geant4 simulations. Section 3.2 introduces BoGEMMS-HPC and its main features. Section 3.3 describes the HPC capabilities implemented in BoGEMMS-HPC during my PhD, focusing on multi-threading and multi-process execution. Finally, Section 3.4 presents the results of performance tests conducted on the Tier-0 Leonardo supercomputer and a local INAF OAS computing node.

3.1 High-Performance Computing

HPC refers to the use of supercomputers or large computing clusters to address complex and computationally intensive problems. Supercomputers are equipped with numerous processors (even thousands), each containing multiple cores, enabling programs to be executed in parallel and significantly reducing computation time. The applications of HPC strategies are numerous, ranging from climate and weather modeling, medical imaging, artificial intelligence, finance, particle physics, and astrophysics. In the context of astrophysics, the progress in observational techniques, the advancement in complexity of the scientific targets, and the increase of sophisticated instrumental technologies, often operating in radiation-intensive environments, require efficient and accurate simulations to provide reliable predictions and drive the instrumental design. Implementing HPC strategies becomes therefore essential to achieve these tasks.

Regarding parallel simulations, there is a fundamental distinction between *thread-level* parallelism and *process-level* parallelism. The difference between processes and threads is the following:

- **processes** are individual instances of executing programs, each with its own dedicated execution context, memory, and other system resources;
- **threads** represent parallel execution flows within a single process, sharing that process's memory and resources.

In thread-level parallelism, all threads share the same address space, making communication essentially equivalent to direct memory access. In contrast, process-level parallelism requires explicit interprocess communication (IPC) mechanisms to enable data exchange between processes. A common approach is message passing, where processes explicitly send and receive messages. The standardized library specification for implementing message passing in HPC architectures is the Message Passing Interface (MPI), and the execution processes are also called MPI tasks. There are several widely used implementations of MPI, such as OpenMPI, MPICH, and vendor-optimized solutions like Intel MPI and IBM Spectrum MPI. One advantage of using MPI over multi-threading is that parallelism can be extended to multiple independent computers (such as those in large-scale clusters), with message passing taking place over a network connecting the computers. While this inevitably provides a wider opportunity to exploit larger resources, it can also introduce overheads if communication between processes is inefficient (if, for example, communication is too frequent or the data exchanged is too large).

Starting from the 10.0 release, Geant4 is equipped with a dedicated library for multi-threading (MT) parallelism based on the POSIX standard. The parallelization is at *event* level, i.e. the events are distributed among the parallel instances. In Geant4, an event is the simulation of all interactions generated by a primary particle and it is associated to a unique ID. A Geant4 event includes the secondary particles that can be produced by the interaction of the primary particles with the environment, such as optical photons generated in a scintillator, an electron-positron pair after a gamma photon conversion, and so on. Geant4 MT is implemented using a master-worker model: a central *master* thread manages the overall simulation flow, while multiple *worker* threads are responsible for simulating individual events in parallel (Figure 3.1). To optimize memory usage in multi-threaded applications, worker threads share memory-intensive components such as geometry and physics tables. Meanwhile, each thread maintains private instances of objects that must remain thread-local, including sensitive detectors and hit collections. Currently, the default installation of Geant4 does not include a library for process-level parallelism. However, an external library (called G4MPI) that uses OpenMPI has been developed and released as part of the Geant4

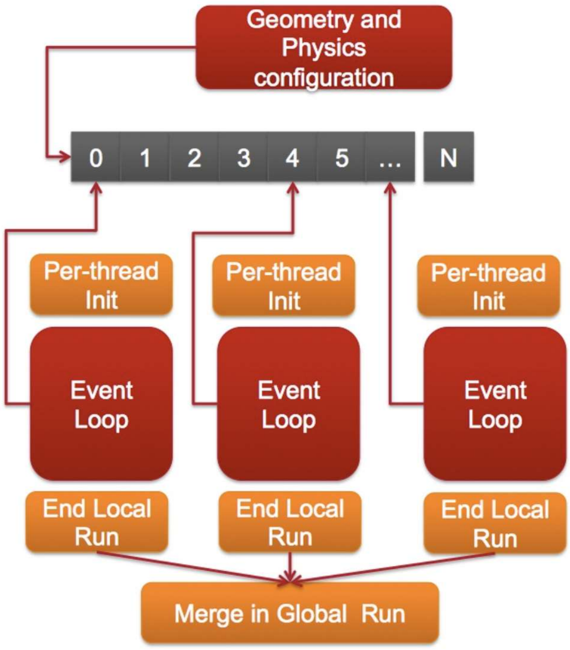

Figure 3.1 – Schema of the master-worker model employed in Geant4. The geometry and physics configuration is handled by a master thread, while the events are distributed among worker threads; the results are then merged and returned to the master thread. Credit: Geant4 documentation.

example applications. The G4MPI library (Dotti et al. 2016) adopts a master/worker model, in which the master process distributes commands and input to the worker processes, which in turn return the results to the master. It is currently the only readily usable MPI interface for the Geant4 framework.

3.2 BoGEMMS-HPC

In this section, we provide an overview of the BoGEMMS-HPC design and capabilities, whose details can be found in Bulgarelli et al. (2012), Fioretti et al. (2014), and Fioretti (2025).

BoGEMMS-HPC is a multi-purpose Geant4-based framework based on BoGEMMS (Bologna Geant4 Multi-Mission Simulator), a software project started at the INAF/OAS institute in 2010 (Bulgarelli et al. 2012), with the aim of building an astronomy-oriented multi-application simulation framework. While Geant4 provides all the core libraries needed for the simulation, it leaves to the user the task of configuring the whole setup. This includes choosing appropriate physics models and specifying which physical quantities (e.g., energy deposition, particle direction, interaction type) should be recorded. Typically, these choices must be hard-coded into the simulation source files, requiring the development of a dedicated simulation for each new experiment. BoGEMMS-HPC addresses this limitation by allowing users to fully configure the simulation geometry, physics models, and output formats (see Section 3.2.1) via a configuration file. Moreover, the user can activate or deactivate multi-threading execution (see Section 3.3), specifying the number of threads. The geometry and physics classes are selected by setting the corresponding keywords to the names of the classes. Geometry classes support reading CAD and GDML files. The user can also define custom keywords directly in the source code and reference them in the configuration file, enabling flexible run-time selection of specific features.

The source particle distribution (energy, direction, and position) can be specified using a macro file with the Geant4 General Particle Source (GPS) tool. GPS is part of the Geant4 toolkit and allows for the specification of the spectral, spatial, and angular distribution of the primary particles. In addition, the user can simulate from a list of particles provided in ASCII, FITS, or ROOT format. Thanks to its great flexibility, BoGEMMS has been employed in numerous space missions, such as the soft proton scattering in XMM-Newton and Athena (Fioretti et al. 2024), and non-X-ray background in the New Hard X-ray Mission (Fioretti et al. 2012).

3.2.1 Output data formats

BoGEMMS-HPC records the entrance and exit of the particles (and their secondaries) in the volumes that the user defines as sensitive. During the simulation, hits are collected in the activated sensitive volumes. When a particle enters (or is generated within) the volume, the energy deposit is accumulated along its path until it exits the volume, is absorbed, or undergoes a conversion. If secondary particles are generated within the volume, their energy deposits are also tracked and recorded throughout their trajectories. The output data format is defined as XYZ. It is an event list, each row showing information about the particles interacting with the sensitive volumes. The user can then select in the configuration file which information is stored in the output file.

There are three data output formats available: FITS files, SQLite, and MongoDB databases.

FITS A FITS file (Flexible Image Transport System [1]) is a standard digital file format used in astronomy for storing, transmitting, and processing scientific data. BoGEMMS-HPC uses the CFITSIO interface library to open FITS files, write data, and close the files.

SQLite Structured Query Language (SQL) is a specialized programming language designed for managing and manipulating data in relational database management systems (RDBMS). SQLite[2] is a lightweight, fast, and reliable SQL database engine implemented as a C library. It is self-contained, requires minimal setup, and represents the most used database engine. A notable application example is the Sloan Digital Sky Survey (SDSS), one of the largest multi-spectral imaging and spectroscopic redshift surveys with exceptional data volume (York et al. 2000).

MongoDB MongoDB[3] is a non-relational (NoSQL) database that stores data in a flexible, non-tabular format. In the BoGEMMS-HPC framework, the MongoDB C++ driver enables the simulation output to be saved in BSON (Binary JSON), a JSON-like format. Unlike traditional databases that use tables and rows, MongoDB organizes data into collections of documents, where each document is a set of key-value pairs. In this context, each row in the BoGEMMS-HPC FITS output, representing the entry and exit of a particle in a sensitive volume, is stored as an individual document within a collection. A MongoDB database can hold multiple collections. MongoDB supports two modes of operation for writing data: it

[1] https://heasarc.gsfc.nasa.gov/docs/heasarc/fits.html
[2] https://www.sqlite.org/
[3] https://www.mongodb.com/

can store documents on a remote server (allowing concurrent access from multiple threads or processes), or it can write them locally on disk, when remote access is not available.

3.3 HPC features

A relevant portion of my PhD project has been devoted to implementing and testing multi-threading and multi-process parallelism in BoGEMMS-HPC. We implemented the Geant4 built-in multi-threading library for multi-threading. For multi-processing, we use the Geant4 G4MPI library.

3.3.1 Multi-threading (MT)

As mentioned in Section 3.2, the user can enable MT and select the number of threads directly through dedicated keywords in the configuration file. The behavior of threads during output writing depends on the selected data format. When using FITS files or SQLite databases, concurrent writing is not supported: only one thread can write at a time, while the others are placed on hold. This was done by including dedicated locks that prevent other threads from writing while one thread is writing. In contrast, MongoDB offers a more scalable solution for parallel output. It enables simultaneous access from multiple threads through the use of a connection pool (a cache of pre-established, reusable database connections) which allows all threads to share the same client instance efficiently.

When MT is enabled, the master thread generates random number seeds for each event, ensuring that every event is simulated independently.

3.3.2 G4MPI

Beyond multi-threading, BoGEMMS-HPC has been extended to support multi-process parallelization using the G4MPI library. The user can enable multi-process parallelization by executing the simulation with the command `mpiexec -n <task number>`. The `<task number>` defines the number of processes, referred to as MPI tasks, used for the parallel simulation. Each MPI task is an independent instance of BoGEMMS-HPC, with its own memory space and access to input and output files. If the user specifies N events in the macro file, each MPI task will simulate N events, resulting in a total of $N\times$`<task number>` events being processed across all tasks. Unlike multi-threading, each MPI process writes to its own separate output file. This eliminates any concurrency issues during output writing, regardless of the selected data format. Each output file is uniquely identified by including the label *task* followed by the task number in its filename, indicating which MPI task generated it. If MT is also enabled, the events within a task are distributed among the threads within the MPI process. The scheme of the multi-threading and multi-tasking implementation design is illustrated in Figure 3.2.

To ensure that each task runs independently and generates statistically uncorrelated results, the MPI task ID is used to initialize the seed of the corresponding Monte Carlo simulation.

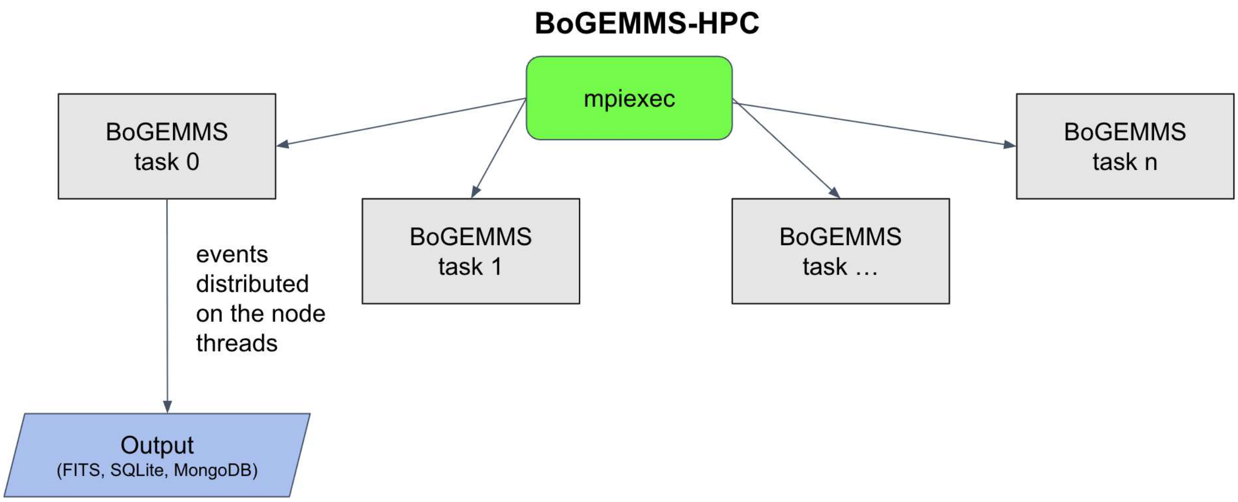

Figure 3.2 – Schema of G4MPI and MT implementation in BoGEMMS-HPC (Fioretti and Ciabattoni 2025). With `mpiexec`, the user executes indipendent BoGEMMS-HPC applications (tasks); within a task, the events can be further distributed among the threads. Each task writes the output to an independent file, which can be FITS, SQLite, and MongoDB.

3.4 Performance tests on HPC environments

The performance of BoGEMMS-HPC has been evaluated under both parallelization schemes (MT and MPI) and for all supported output data formats (FITS, SQLite, and MongoDB). The tests were performed on *Leonardo*, the pre-exascale Tier-0 supercomputer of the European High-Performance Computing Joint Undertaking (EuroHPC JU), hosted by CINECA at the Bologna DAMA-Technopole in Italy. Through the ICSC Resource Allocation Committee (RAC), we accessed the X2140 CPU-only blades, based on Intel Sapphire Rapids processors, within the Data Centric General Purpose (DCGP) partition. We used the Slurm[4] workload manager to run the simulation on one of the available nodes. When using G4MPI, MPI is invoked within a single node, and the tasks are distributed across the different cores of that node. Additional tests were executed on a node from the local OAS cluster (*node05*).

The test case selected for the performance evaluation is the simulation of the light output of a CsI crystal coupled with a PMT, an application used to verify the generation of optical photons in Geant4 and discussed in Chapter 4 (Ciabattoni et al. 2024). The simulation involves the irradiation of 122 keV photons onto a CsI panel (Figure 3.3). The deposited energy in the crystal is converted into optical photons, which then propagate through the CsI until they are either absorbed or detected by the PMT.

For each parallelization scheme, we iteratively ran the same simulation while progressively increasing the number of threads or tasks. We recorded the wall time of each run to assess the performance gain and scalability of the code with respect to computational resources. In the following plots, we present both the wall time in seconds (i.e., the actual time that passes from the start to the end of the simulation) and the corresponding speed-up, defined as the ratio between the wall time for a single thread/task and that for n threads/tasks. Both quantities are shown as a function of the number of threads/tasks. For reference, the ideal (linear) speed-up is also plotted for comparison.

[4] https://slurm.schedmd.com

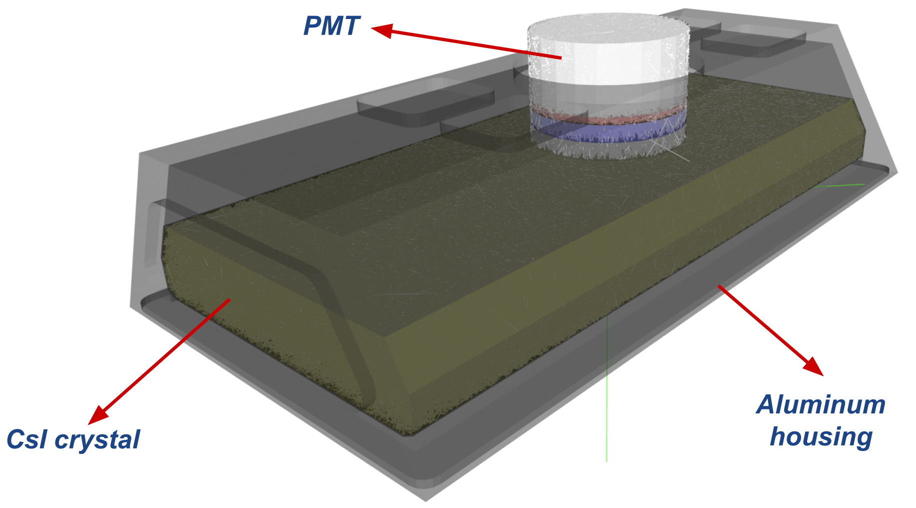

Figure 3.3 – Snapshot of the simulation used for the performance test. A CsI crystal, enclosed in a aluminum casing, is irradiated by a 122 keV photon (the green ray). The energy deposited in the crystal leads to the generation of thousands of optical photons (white rays) that are reflected within the CsI. Part of the optical photons enters the PMT window and are absorbed.

3.4.1 Test on Leonardo

The tests were conducted on a 112-core node of Leonardo, using the DCGP partition. In this setup, MongoDB was evaluated by writing output data to local files instead of a remote server, as remote server access could not be configured on the Leonardo nodes. Figure 3.4 shows the recorded wall times and speed-ups for both MT and MPI parallelization schemes, across all output data formats. In all cases, MPI outperforms MT, with the performance gap widening as the number of cores increases. With MPI, we achieve a speed-up of approximately 70 using 112 cores, compared to the serial case. The scaling remains nearly ideal up to about 10 cores; after that, it starts to deviate, likely due to memory management overhead as the node becomes fully occupied, along with the increased overhead from the initialization and finalization of MPI tasks. In contrast, MT reaches a maximum speed-up of about 15 at 112 cores. Both parallelization schemes begin to plateau at around 100 tasks/threads, although the trend suggests that larger speed-ups could be achieved with more than 112 cores. The better performance of MPI with respect to MT is primarily due to a limitation in the MT approach: threads cannot write simultaneously to the same file, which leads to contention and bottlenecks as more cores are involved. Conversely, MPI processes generate separate output files, avoiding write contention and enabling more efficient parallel execution. The three output data formats show comparable wall times and speed-ups as the number of cores increases.

3.4.2 Test on local machine

The optimal use of MongoDB involves writing output data to a remote server, allowing parallel threads to access the database without local file system contention. However, due to limitations in the Leonardo environment, it was not possible to configure remote server access from the compute nodes. For this reason, we conducted additional performance tests

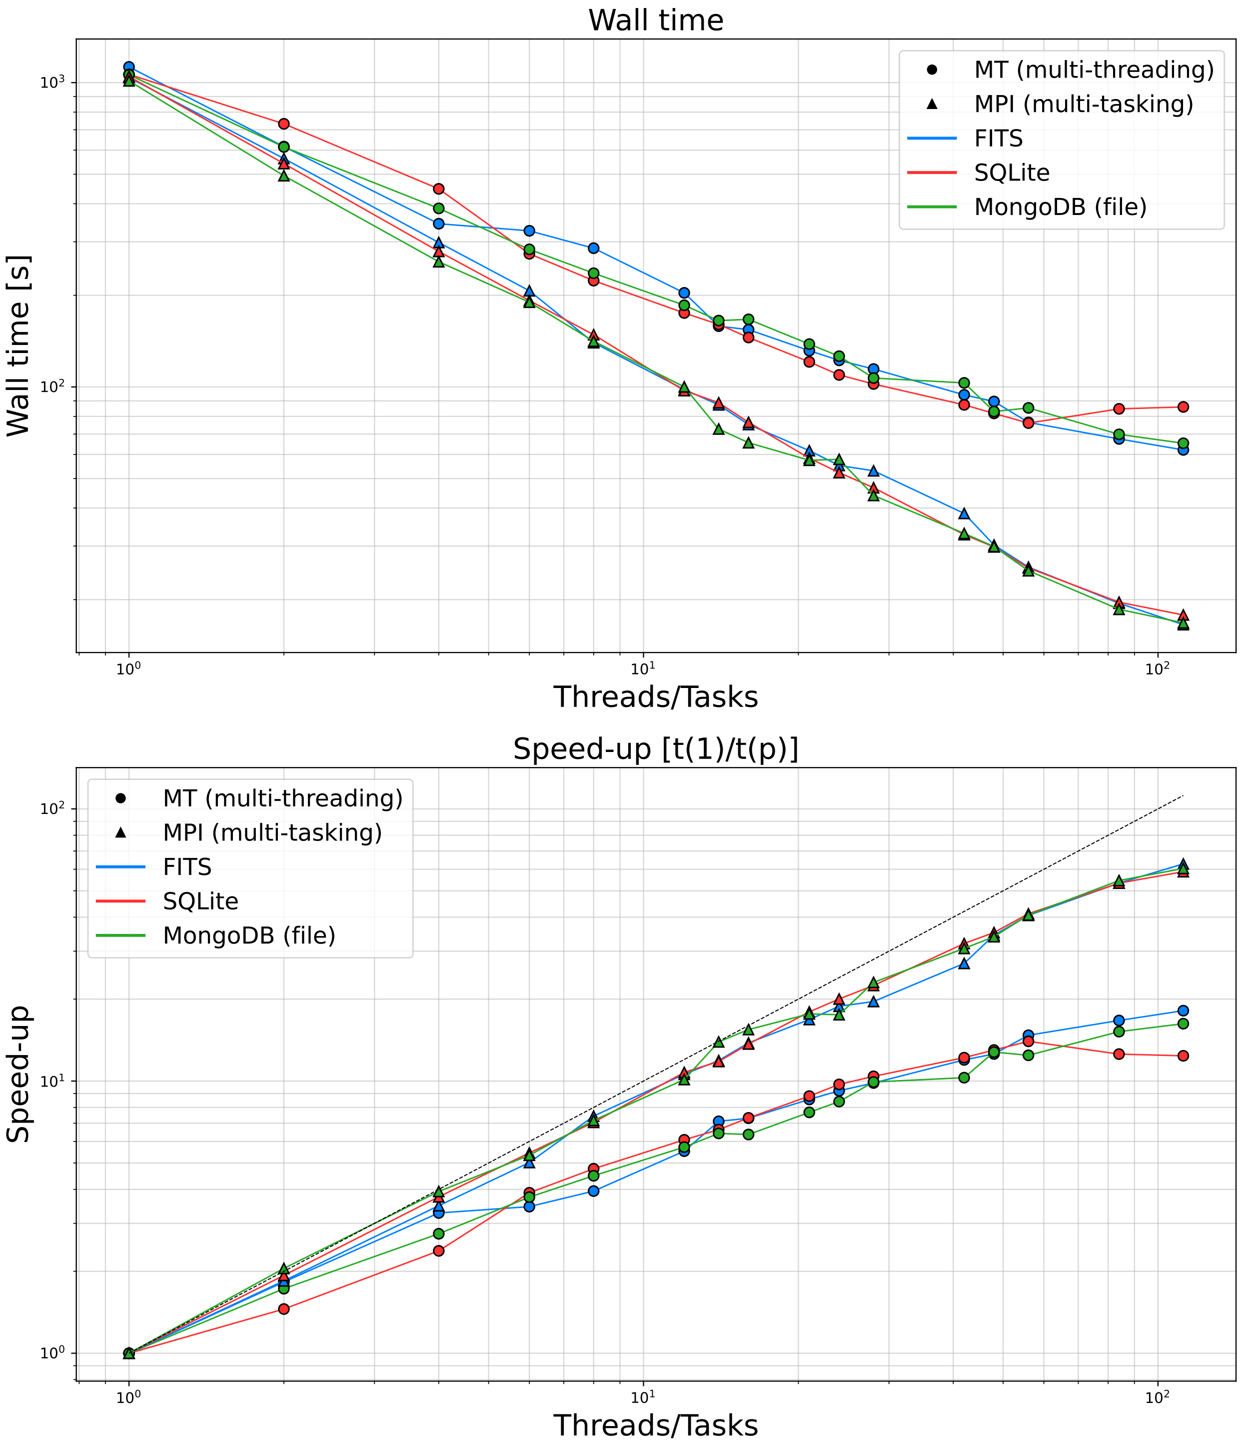

Figure 3.4 – Wall time (top) and speed-up (bottom) of BoGEMMS-HPC simulations as a function of the number of threads/processes, as measured in Leonardo.

on a smaller local cluster at INAF OAS. The test node at OAS is equipped with two sockets, each hosting 10 physical cores, for a total of 20 cores. In this setup, we were able to configure MongoDB with remote server access. The corresponding results are shown in Figure 3.5. In terms of wall time, MongoDB with server access is the slowest among the three output data formats for both parallelization schemes. In the serial case, MongoDB is approximately 50% slower with respect to FITS and SQLite. However, this gap progressively narrows under MT as the number of cores increases. With 20 cores under MT, MongoDB is only about 13% slower, whereas the performance gap with MPI remains nearly constant. The improved scalability of MongoDB with MT is primarily due to the reduced thread-level write contention when accessing a remote server. In contrast, FITS and SQLite suffer from increased contention with more threads, which progressively degrades their performance. As a result, MongoDB achieves better speed-up compared to the other formats beyond approximately 10 cores. At 20 cores, MongoDB reaches a speed-up of about 10, while FITS and SQLite plateau around 8.

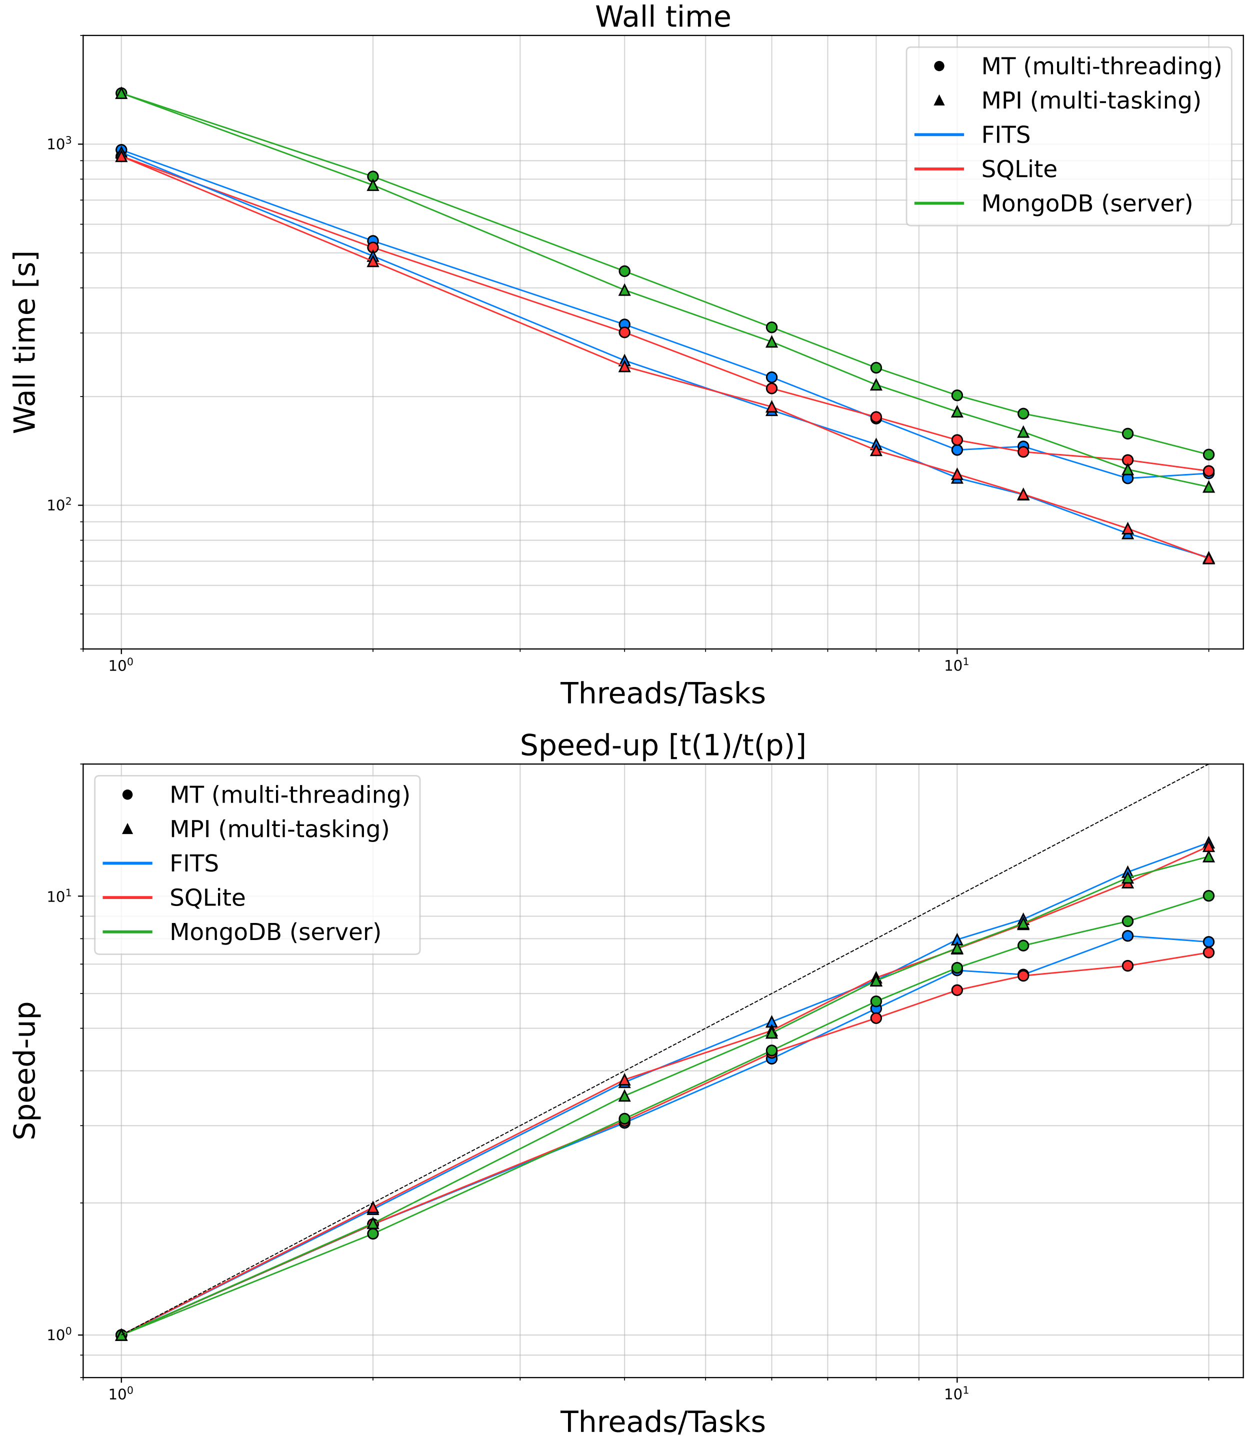

Figure 3.5 – Wall time (top) and speed-up (bottom) of BoGEMMS-HPC simulations as a function of the number of threads/processes, as measured in OAS.

3.4.3 Memory consumption

Among the three output data formats, MongoDB results in the highest storage usage for the same number of simulated events. On average, FITS and SQLite require approximately 67% and 69% less storage, respectively, compared to MongoDB. In addition, BoGEMMS-HPC provides the option to enable FITS file compression at the end of the simulation. Although this significantly reduces the required storage space, it introduces an overhead in execution time, affecting the overall performance. The simulations are about 40% slower when using MPI, regardless of the number of cores used. With MT instead, this gap progressively increases with the number of threads, ranging from 40% for serial simulations to about 320% with 20 threads. This result stems from the fact that, in the MPI case, each task compresses its own output file independently, effectively sharing the compression overhead across all processes. In contrast, under MT, the compression is executed serially by the master thread after the simulation has completed. As the number of threads increases, the workload per thread decreases, but the compression time remains constant, thus becoming a relatively larger fraction of the total runtime. Nevertheless, FITS file compression offers a significant advantage in terms of storage, reducing the occupied disk space by approximately 67%.

Chapter 4

Verification of Geant4 optical simulations

Although Geant4 was originally designed for high-energy processes, it has been extended to simulate low-energy physics in the optical range, including scintillation. Geant4 treats optical light as individual photons, which obey standard optical interactions such as reflection and refraction. We can use this Geant4 capability to explicitly simulate the scintillation processes within the COSI ACS crystal and accurately model its performance. As a preliminary yet essential step, we performed a verification of the Geant4 optical simulation using both analytical models and experimental measurements, while exploring the large set of configurable optical parameters. We verified the simulated transmission and reflection probabilities of optical photons against the theoretical expectations and replicated in the simulation an experimental setup involving a PMT coupled with a block of CsI to investigate the optical parameters and reproduce the laboratory measurements. We used the framework BoGEMMS-HPC (presented in Chapter 3) to perform the Geant4 simulations. The results of this work have been published in Ciabattoni et al. 2024.

This chapter is organized as follows. Section 4.1 presents the Geant4 optical physics library and the main parameters. Section 4.2 shows the verification of the simulated transmission and reflection probabilities against theoretical models. Section 4.3 discusses the comparison between the Geant4 simulation of the CsI-PMT system and laboratory measurements, focusing on the impact of the optical parameters on the results.

4.1 The Geant4 optical physics library

The Geant4 optical physics library (Geant4-collaboration 2023) enables the generation and tracking of radiation in the optical band. The transport of optical photons through matter is affected by two processes: absorption and boundary effects. The absorption capability of a medium is defined through the `ABSLEN` parameter and represents the mean free path of an optical photon in that medium. The boundary effects are regulated by a set of parameters that characterize the optical properties of the surface and the reflection and refraction probabilities for the optical photons. In Geant4, two models for the boundary effects can be selected: the `UNIFIED` model (Levin and Moisan 1996) and the `GLISUR` model. There are also look-up tables, such as LBNL (Janecek and Moses 2010) and LUTDAVIS (Roncali and Cherry 2013), which are tailored to specific surface types and derived from laboratory measurements. In this verification work (and also for the simulation benchmarking in Chapter 5), we adopted the `UNIFIED` model, as it extends the original `GLISUR` model (developed for Geant3) and offers a more flexible and customizable description of surface properties than LUT-based approaches. When a scintillator is wrapped with a reflective material, optical photons are either reflected or absorbed. In Geant4, this interface is referred to as a *painted* surface. There are two main characteristics which discriminate different types of painted surfaces: the finish (*polished* or *ground*) and the presence or absence of a gap between volume and paint (*frontpainted* and *backpainted*). Hence, for dielectric-dielectric interfaces we have four different combinations: `PolishFrontPainted`, `GroundFrontPainted`, `PolishBackPainted` and `GroundBackPainted`. The level of roughness is parametrized by the parameter σ_α. A rough surface is modeled as consisting of microfacets, each of which has an angle (α) between its normal and the average surface normal. The alpha angle is sampled from a Gaussian distribution, with the mean given by the average surface normal and a standard deviation of σ_α. Hence, a surface with higher σ_α is a rougher surface. Another key property is the reflective type. There are four types of reflection that an optical photon can undergo: specular spike (SS), specular lobe (SL), backscattering (BS) and Lambertian (L) reflection (see Figure 4.1). In the verification work against laboratory measurements (Section 4.3), we will test SS reflection (i.e. the incidence angle is equal to the reflection angle), and L reflection (i.e. the reflection angle follows a Lambert cosine distribution). The main properties of the four types of surfaces are summarized as follows:

- **PolishedFrontPainted (PFP)** The optical photon can only be reflected or absorbed. The probability of being reflected is set by the parameter Reflectivity; if reflection occurs, then it is specular spike reflection.

- **GroundFrontPainted (GFP)** The optical photon can only be reflected or absorbed. The probability of being reflected is set by the parameter Reflectivity; if reflection occurs, then it is Lambertian reflection.

- **PolishedBackPainted (PBP)** There is an air gap between the crystal and the paint, the latter being a perfectly smooth mirror. The user must set σ_α, which is associated to the crystal-gap interface, and a refractive index, which is associated to the gap and then used to apply the Snells law between the crystal and the gap. If reflection occurs, it can be SS, SL, BS or L according to their assigned probabilities. If refraction occurs, the optical photon enters the gap and interacts with the paint. It can then be either SS reflected or absorbed according to the parameter Reflectivity as in the previous cases. If reflected, the photon interacts again with the gap-crystal interface and Snells law is again applied. At this point, the photon can be refracted into the crystal, or reflected back and interact again with the paint.

- **GroundBackPainted (GBP)** As for PBP, there is an air gap between the crystal and the paint, the latter being now a ground mirror instead of a polished one. The user must set σ_α, which is associated to the crystal-gap interface, and a refractive index, which is associated to the gap and then used to apply the Snells law between the crystal and the gap. The reflection can be SS, SL, BS or L according to their assigned probabilities. If refraction occurs, the optical photon enters the gap and can then be either L reflected or absorbed from the paint according to the parameter Reflectivity. If reflected, the photon interacts again with the gap-crystal interface and Snells law is again applied. In summary, it is the same as PBP, but the finish of the mirror is ground and the reflection type is Lambertian.

The complete scheme of all optical surfaces within the Geant4 `UNIFIED` model is shown in Figure 4.2 (from the Geant4 manual). The dielectric-dielectric surfaces `Ground` and `Polished` are appropriate for transparent interfaces (e.g., CsI-air). In this case, the absorption probability is determined by 1 minus the reflection coefficient (set by Reflectivity). If not absorbed, the photons are transmitted or reflected according to the refractive indices of the two media and the Snell's law (see Eq. (4.2) of Section 4.2.2).

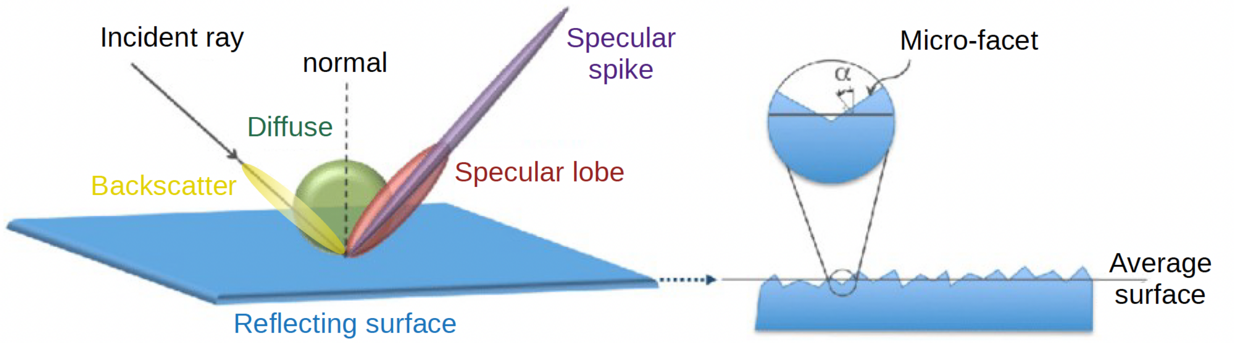

Figure 4.1 – Light reflection models in the Geant4 `UNIFIED` model: specular spike, specular lobe, backscattering, and Lambertian (marked as *diffuse* in the figure). The ground surface is shown as composed of micro facets, with α being the angle between the micro facet normal and the average surface normal. Figure from Bismark (2019).

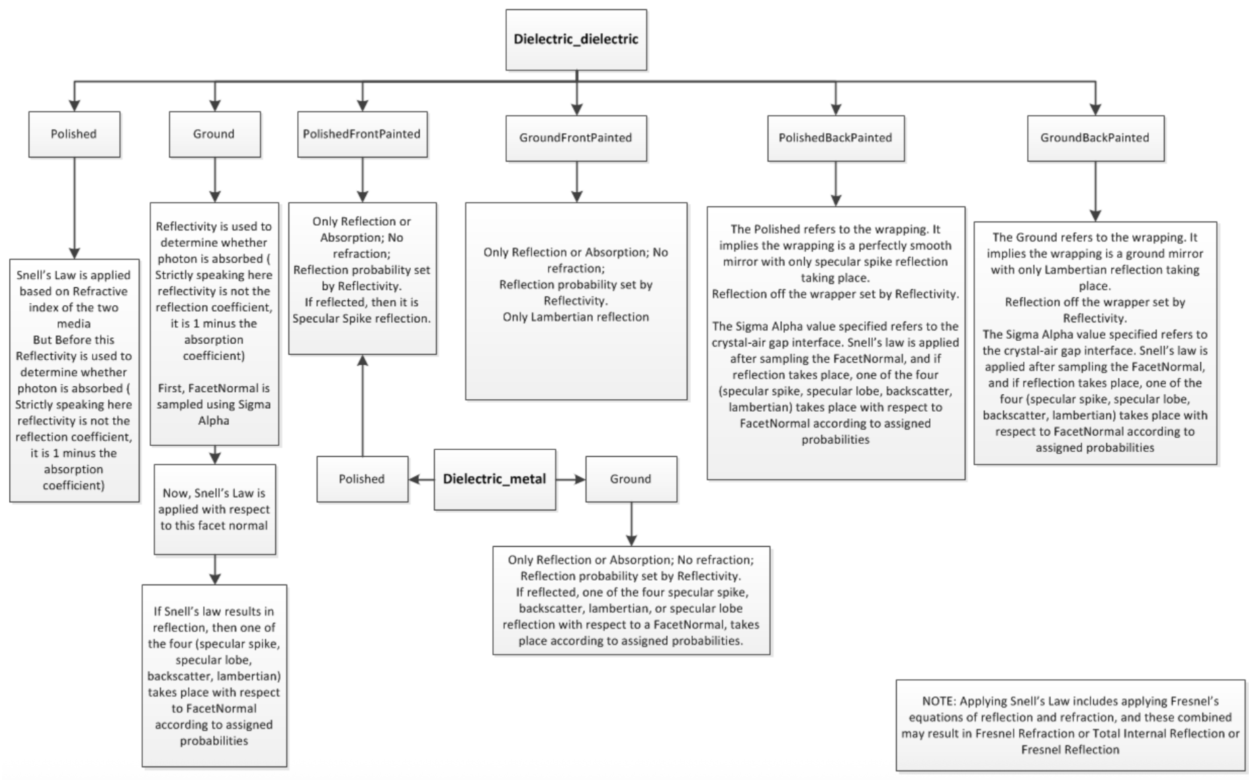

Figure 4.2 – Diagram of the `UNIFIED` model for the Geant4 optical surfaces (figure from the Geant4 manual).

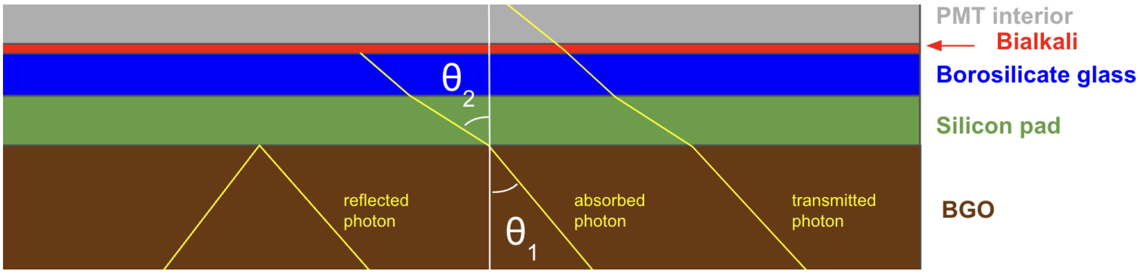

Figure 4.3 – Schematic view of the propagation of optical photons through the BGO-PMT interface.

4.2 Verification with analytical models

With BoGEMMS-HPC we built a toy model that approximates the interface between a BGO-based scintillator and a PMT[1], and studied the transmission, reflection, and absorption probabilities as functions of the angle of incidence (Alvarez Franco 2022).

4.2.1 Geant4 simulation setup

The mass model, shown schematically in Figure 4.3, consists of a five-layered slab composed as follows:

1. the first layer is the BGO, with a thickness of 2 cm;
2. the second layer used as optical coupler between BGO and PMT is the silicon pad (SiPad), polymer made of siloxane ($R_2SiOSiR_2$) and with a thickness of 5 mm;
3. the third layer is the borosilicate glass, made of SiO_2 (80 %), B_2O_3 (13%), Na_2O (4%) and Al_2O_3 (3%); it is used as a window on the PMT and it has a thickness of 5 mm;
4. the fourth layer is the bialkali photocathode, which is made of potassium, cesium and antimony (KCsSb) and is where a fraction of the optical photons is absorbed and generates photoelectrons; it has a thickness of 20 nm;
5. the fifth layer is the PMT interior, which is simulated as an empty slab; its thickness is set to 1 cm.

The lateral dimension of the slab was set to be quite large (1 m) to prevent possible boundary effects in the simulation. The simulated optical properties of all the materials were specified by their refractive index, which is listed in Table 4.1 as a function of the wavelength. As can be noted, only the bialkali photocathode has a non-vanishing imaginary refractive index, which means that it is the only absorbing material in the simulation. In the Geant4 simulation, this absorbing property was taken into account not through the use of a complex refractive index, but rather by setting an (energy-dependent) absorption length (d) to the bialkali material, which is related to the imaginary refractive index (k) in the following way:

$$k = \frac{\lambda}{4\pi\, d} \tag{4.1}$$

[1]The PMT was the photosensitive device originally selected for the COSI ACS, before transitioning to SiPMs.

Wavelength (mm)	BGO	SiPad	Glass	Bialkali	PMT
380	2.244	1.42	1.514	$1.92 + 1.69\,i$	1
395	2.224	1.42	1.514	$2.18 + 1.69\,i$	1
410	2.208	1.42	1.515	$3.38 + 1.71\,i$	1
425	2.193	1.42	1.516	$2.61 + 1.53\,i$	1
440	2.181	1.42	1.516	$2.70 + 1.50\,i$	1
455	2.169	1.42	1.517	$2.87 + 1.44\,i$	1
470	2.159	1.42	1.517	$3.00 + 1.34\,i$	1
485	2.151	1.42	1.518	$3.00 + 1.11\,i$	1
500	2.143	1.42	1.519	$3.00 + 1.06\,i$	1
515	2.135	1.42	1.520	$3.09 + 1.05\,i$	1
530	2.129	1.42	1.520	$3.26 + 0.86\,i$	1
545	2.123	1.42	1.521	$3.20 + 0.63\,i$	1
560	2.118	1.42	1.522	$3.12 + 0.53\,i$	1
575	2.113	1.42	1.524	$3.06 + 0.46\,i$	1
590	2.110	1.42	1.525	$3.01 + 0.42\,i$	1
605	2.104	1.42	1.526	$2.98 + 0.38\,i$	1
620	2.100	1.42	1.528	$2.96 + 0.37\,i$	1
635	2.097	1.42	1.530	$2.95 + 0.35\,i$	1
650	2.094	1.42	1.532	$2.95 + 0.34\,i$	1
665	2.091	1.42	1.534	$2.95 + 0.34\,i$	1
680	2.088	1.42	1.515	$2.96 + 0.33\,i$	1

Table 4.1 – Refractive indices of all the materials for different photon wavelengths (Alvarez Franco 2022).

where λ is the photon wavelength in vacuum. Physically, the absorption length represents the mean free path for the photon before being absorbed in the medium.

4.2.2 The analytical model

The fundamental equation which relates the incident angle θ_1 and the refractive angle θ_2 (with respect to the vertical line perpendicular to the incidence surface) of the light passing through two different media is Snell's law:

$$\sin\theta_1 \cdot n_1^* = \sin\theta_2 \cdot n_2^* \tag{4.2}$$

where n_1^* and n_2^* are the complex refractive indices of the two materials. From Eq. (4.2) it is possible to express the cosine of θ_2 as follows:

$$\cos\theta_2 = \sqrt{1 - \left(\frac{n_1^*}{n_2^*} \cdot \sin\theta_1\right)^2} \; . \tag{4.3}$$

We note that when $\sin\theta_1 > (n_2^*/n_1^*)$ the cosine of θ_2 is purely imaginary, i.e. no refraction is possible and total internal reflection takes place. The smallest incidence angle for which we have total internal reflection is called *critical angle* (θ_{crit}), that means that for $\theta_1 > \theta_{\text{crit}}$ reflection always occurs.

The transmission and reflection probabilities are calculated through the amplitude coefficients r and t, which are polarization-dependent and specified by the following Fresnels equations:

$$\begin{aligned} r_s &= \frac{n_1^* \cdot \cos\theta_1 - n_2^* \cdot \cos\theta_2}{n_1^* \cdot \cos\theta_1 + n_2^* \cdot \cos\theta_2} \\ r_p &= \frac{n_2^* \cdot \cos\theta_1 - n_1^* \cdot \cos\theta_2}{n_1^* \cdot \cos\theta_2 + n_2^* \cdot \cos\theta_1} \\ t_s &= 1 + r_s \\ t_p &= \frac{n_1^*}{n_2^*} \cdot (1 + r_p) \end{aligned} \tag{4.4}$$

where the subscripts p and s refer to parallel and perpendicular polarization with respect to the plane of incidence, respectively.

In a system of three materials, photons which are generated in the material 1 and are transmitted to the material 2 can undergo multiple internal reflections before being transmitted to the material 3. Then, the amplitude coefficients for such three-layered configuration (T and R) are

$$\begin{aligned} T_{(s,p)}(\lambda,\theta_1) &= \frac{t_{(s,p)12} \cdot t_{(s,p)23} \cdot e^{i\delta_2 \cos^2\theta_2}}{1 + r_{(s,p)12} \cdot r_{(s,p)23} \cdot e^{i(2\delta_2 - \delta_3)}} \ , \\ R_{(s,p)}(\lambda,\theta_1) &= r_{(s,p)12} + \frac{t_{(s,p)12} \cdot t_{(s,p)21} \cdot r_{(s,p)23} \cdot e^{i(2\delta_2 - \delta_1)}}{1 + r_{(s,p)12} \cdot r_{(s,p)23} \cdot e^{i(2\delta_2 - \delta_1)}} \ , \end{aligned} \tag{4.5}$$

where

$$\begin{aligned} \delta_1 &= 4\pi\frac{h_2}{\lambda}n_1^* \frac{\sqrt{1-\cos^2\theta_1} \cdot \sqrt{1-\cos^2\theta_2}}{\cos\theta_2} \ , \\ \delta_2 &= \frac{2\pi h_2 n_2^*}{\lambda\cos\theta_2} \ , \\ \delta_3 &= 4\pi\frac{h_2}{\lambda}n_3^* \frac{\sqrt{1-\cos^2\theta_3} \cdot \sqrt{1-\cos^2\theta_2}}{\cos\theta_2} \ , \end{aligned} \tag{4.6}$$

are the phase differences; h_2 refers to the thickness of the intermediate layer and λ is the photon wavelength (both of them must have the same measurement unit). Such phases take into account the wave nature of light and the fact that it can take paths with different lengths before reaching the material 3, depending on how many reflections occur inside the material 2; this translates in differences among the phases and hence interference effects.

Finally, the total reflection and transmission probability for a three-layered slab, equally weighting on both polarizations, is

$$\begin{aligned} T(\lambda,\theta_1) &= 0.5 \cdot \left(T_s \cdot \bar{T}_s + T_p \cdot \bar{T}_p\right) \cdot \frac{n_3 \cdot \mathbb{R}(\cos\theta_3)}{n_1 \cdot \mathbb{R}(\cos\theta_1)} \ , \\ R(\lambda,\theta_1) &= 0.5 \cdot \left(R_s \cdot \bar{R}_s + R_p \cdot \bar{R}_p\right) \ . \end{aligned} \tag{4.7}$$

where $\mathbb{R}(x)$ is the real part of x. For a system of five materials, it is possible to consider it as composed of two three-layered blocks, with the material 3 of the first one being also the material 1 for the second one (so for our case, BGO-SiPad-glass is the first block, glass-bialkali-PMT is the second one). In this perspective, the total probability for a photon to reach the final layer is that it is both transmitted in the first and second block, while the probability of being reflected is that it is reflected in the first block, or it is transmitted in the first block and then reflected in the second one. Hence

$$\begin{aligned} T(\lambda,\theta_1) &= T_1(\lambda,\theta_1)\cdot T_2(\lambda,\theta_1')\ , \\ R(\lambda,\theta_1) &= R_1(\lambda,\theta_1) + T_1(\lambda,\theta_1)\cdot R_2(\lambda,\theta_1')\ . \end{aligned} \tag{4.8}$$

where the incident angle of the second block (θ_1') is actually the transmission angle from the first one. Since all probabilities must sum up to 1, the absorption is then just

$$A(\lambda,\theta_1) = 1 - T(\lambda,\theta_1) - R(\lambda,\theta_1)\ . \tag{4.9}$$

Once a photon reaches the photocathode, it might be absorbed and generate a photoelectron. The probability that such an event occurs is then

$$P_{phe}(\lambda,\theta_1) = T(\lambda,0°)\cdot\epsilon_Q(\lambda,0°)\cdot\frac{A(\lambda,\theta_1)}{A(\lambda,0°)}\ . \tag{4.10}$$

The first term takes into account the probability that the photon is transmitted from the first block and hence actually reaches the photocathode; the second term is the quantum efficiency (Hamamatsu 2007) at $\theta_1 = 0°$ for the bialkali photocathode; the third term (the ratio of the two absorption terms) encodes the angular dependence which is indeed absent in the quantum efficiency.

4.2.3 Comparison with analytical expectations

For each simulation, we generated in the BGO 1000 optical photons at 500 nm (2.48 eV) with a fixed angle of incidence with respect to the BGO-SiPad interface, varying the latter from 0° to 90° at a step of 0.1° and obtaining each time the number of transmitted, reflected, and absorbed photons. The transmitted (T) photons are those that reach the PMT interior, the reflected (R) ones are the photons that are reflected back into the BGO layer. The absorbed (A) photons are the ones absorbed in the photocathode and the only ones that can generate photoelectrons (and, therefore, a signal in the PMT). For both analytical and simulated probabilities, we have $T + R + A = 1$.

The left panel of Figure 4.4 shows the transmission, reflection and absorption probabilities obtained by the simulation and the analytical model as a function of the incidence angle, while the probability of emission of photoelectrons is instead plotted in the right panel Figure 4.4. The Geant4 simulation reproduces the behavior of the theoretically predicted probabilities, presenting however some deviation from the analytical model. As can be noticed, the transmission probability drops at an incidence angle of $\sim 27.3°$. In fact, when θ_1 is larger than such an angle, the optical photons will always arrive at the bialkali-PMT boundary with an angle larger than the critical value of the surface and, therefore, they are reflected. The absorption probability is nearly constant at $\sim 40\%$ for $\theta_1 \lesssim 27.3°$, then increases up to $\sim 60\%$ and finally drops to 0% at $\sim 41°$, that corresponds to the critical angle of the BGO-SiPad interface. For θ_1 larger than this angle, only reflection takes place.

The deviation between Geant4 and the analytical prediction is probably caused by having a very thin layer (the bialkali photocathode) that has a thickness of the same order of the

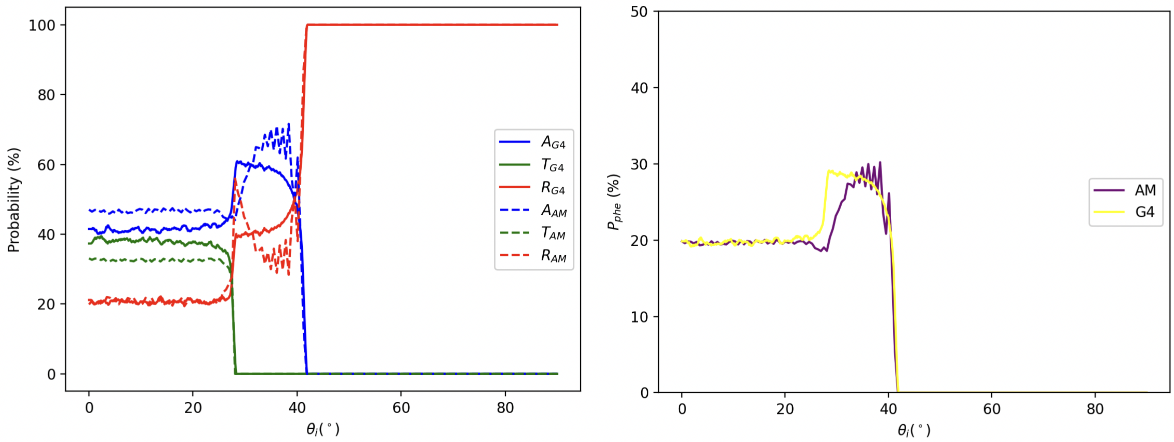

Figure 4.4 – On left panel, transmission, reflection, and absorption probabilities as functions of the incidence angle for the PMT toy model from both the simulation (G4) and the analytical model (AM) (left). On right panel, probability for the emission of a photoelectron as a function of the incidence angle, from both the simulation (G4) and the analytical model (AM).

photon wavelength. In fact, in this condition, the wave nature of optical photons becomes relevant, and interference effects and frustrated transmission (i.e. transmission through a thin layer) may not be completely captured by the Geant4 simulation, leading to the observed discrepancy. We further investigated this issue by reproducing the same simulation, first with a larger thickness of the photocathode (20 mm) and a proportionally larger absorption length (otherwise, the optical photons would always be absorbed in such a thick layer). Secondly, we removed the absorbing property of the photocathode, that is, without assigning to it an absorption length (keeping its thickness at 20 nm). The results are shown in Figure 4.5, where it can be seen that, in the former case, the simulation roughly represents the average behavior of the analytical curve, whereas in the latter case no relevant deviations are observed. The fluctuations in the analytical model are due to the large interference phases in Eq. (4.5), since we increased the parameter h_2. An important remark is that Geant4 is meant to deal mostly with high energy photons, whose wave nature can be safely neglected, but in situations where optical photons interact with very thin layers, the particle description of light may not be a satisfactory approximation. However, despite the lack of such physics in the simulation, the transmission, reflection, and absorption probabilities are still quite close to their theoretically expected values, at every incidence angle.

As the PMT response is proportional to the total number of absorbed photons, regardless of their incidence angle, we compute the three probabilities by integrating them over the incidence angle, for both the simulation and the analytical model. The results are shown in Table 4.2. Geant4 is capable of reproducing the analytical model with an overall accuracy of $< 15\%$. In particular, the discrepancy for the number of absorbed photons, particularly important for predicting the number of photoelectrons emitted from the photocathode, is $\sim 8\%$.

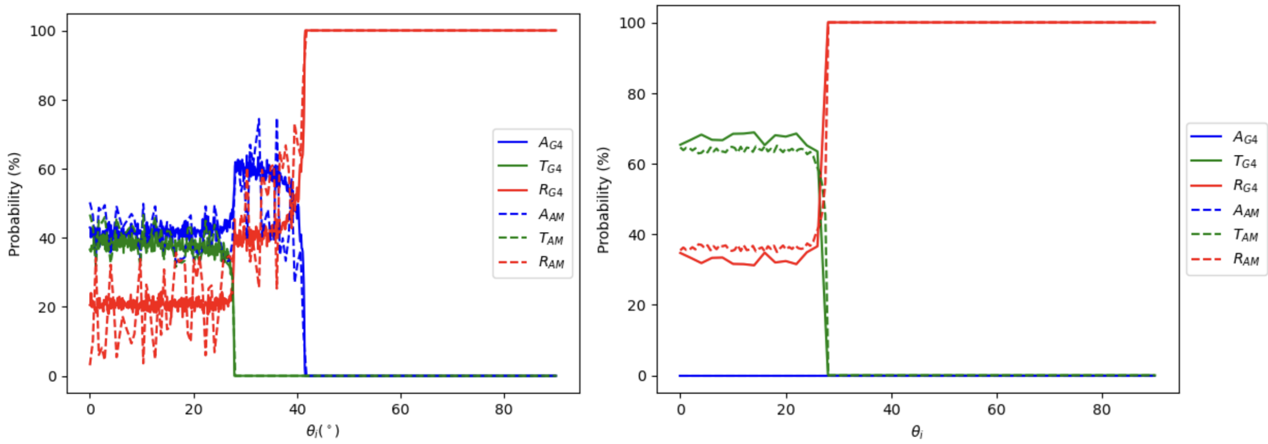

Figure 4.5 – On the left, transmission, reflection, and absorption probabilities for the PMT toy model as functions of the incidence angle for a bialkali thickness and absorption length both 100 times larger than the standard ones; on the right, the same quantities with no absorption in the photocathode.

	Analytical model	Geant4 simulation
Transmission	10.0%	11.6%
Reflection	66.7%	67.1%
Absorption	23.2%	21.4%

Table 4.2 – Integrated transmission, reflection, and absorption probabilities for both the analytical model and the Geant4 simulation.

4.3 Verification with laboratory measurements

We simulated the interaction of X-ray photons with a block of CsI(Tl) coupled with a PMT to estimate the PMT relative response for different beam positions and then compared it with the laboratory measurements described in Alvarez Franco (2022).

4.3.1 The CLAIRE experiment

The experimental setup consists of a trapezoidal block of CsI(Tl), with a length of 198 mm, height of 20 mm, and widths of 93 mm (for the larger side on the top) and 70 mm (for the smaller side on the bottom), placed inside an aluminum case. The block of CsI is wrapped with 2 layers of Teflon and 2 layers of aluminum foil, whose role is to reflect and keep scintillation photons inside the scintillator. An open window with a diameter of 55 mm in the aluminum case is designed to allow the optical coupling of a PMT with the CsI block, through a silicon pad. Finally, a collimated radioactive source (^{57}Co) was placed in front of the bottom side of the aluminum case and 40 measurements of the 122 keV decay line were taken. Each time, the radiation hits the block in a different position, whose distribution and indices are shown in Figure 4.6, and the optical scintillation photons are consequently produced inside the CsI crystal; finally, the PMT responses (proportional to the number of scintillation photons that reach the PMT) were recorded for each position.

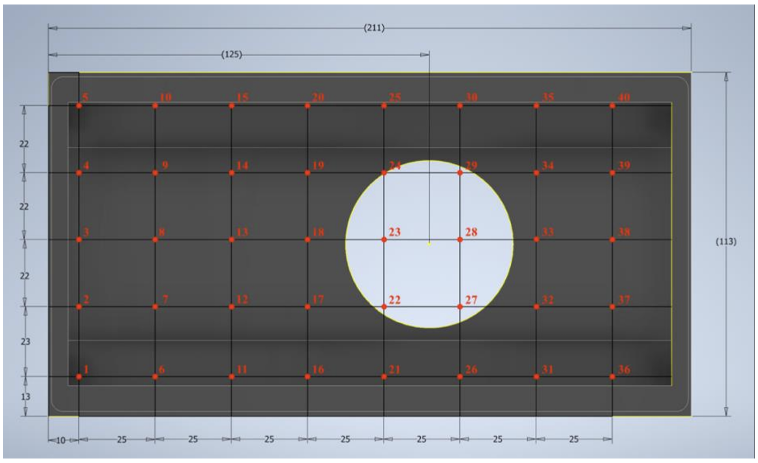

Figure 4.6 – Interaction positions of the 122 keV photons in the bottom face of the CsI crystal and corresponding labels (Alvarez Franco 2022). The white circle corresponds to the PMT window.

4.3.2 Geant4 simulation setup

The Geant4 simulations were conducted with BoGEMMS-HPC. Concerning the simulated mass model, the geometries of the CsI block and the aluminum case were imported using the CADMesh library (Poole et al. 2012), whereas all the other components were constructed as Geant4 standard volumes. The crystal-PMT interface was simulated as follows: a SiPad (with a thickness of 5 mm) optically couples the crystal to the borosilicate glass (with the same thickness) within the open window; above the glass, there is the photocathode (with a thickness of 20 nm) and finally the PMT (with a height of 2 mm). All the layers are cylinders with a diameter of 46 mm.

We simulated 1000 X-ray photons of 122 keV hitting the CsI block in each position as shown in Figure 4.7. Then, for each run, we considered the number of optical photons that were absorbed in the photocathode and normalized it to the sum of all the 40 measurements. We will refer to this quantity as PMT relative response. Although we are not including the PMT electronic readout (the PMT is simulated just as an empty volume in BoGEMMS-HPC), we assume that the actual response of the PMT (which would be expressed in Volt) is related to the number of absorbed optical photons in the photocathode through a proportionality factor, which becomes irrelevant when dealing with the relative response (i.e. normalizing to the sum of all the measurements).

The number of photons absorbed by the photocathode in the simulation is assumed to follow a Poisson distribution, i.e. if we count N photons being absorbed in the photocathode, the associated error is $\sqrt{N}$. In order to best reproduce the experiment, only events in which the X-ray photon entirely deposits its energy in scintillation photons had to be taken into account. To do so, for each beam position we collected the total energy deposit in the photocathode for each input X-ray. Only the events whose total energy deposit lies within ± 2 standard deviations from the mean were kept, all the others being associated with partial conversions of the primary X-ray photon. Having 1000 input events ensured enough statistics to perform this kind of analysis.

As discussed in Section 4.1, Geant4 allows the setting of various parameters to specify the optical properties of the surfaces, e.g. the roughness, and how optical photons interact with them. These parameters influence the simulation outcome and its comparison with

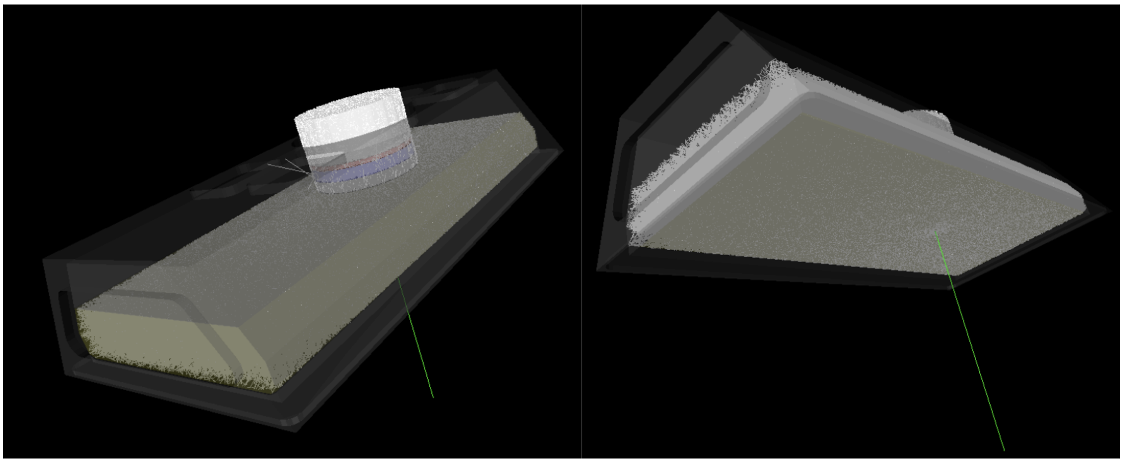

Figure 4.7 – Snapshots of the Geant4 simulation: the green line is the gamma photon which hits the aluminum case, enters the CsI and produces scintillation photons, whose trajectories are so crowded that they almost fill the crystal.

experimental data, hence their proper definition and understanding become crucial for our purposes. We then performed several comparisons for different combinations of those parameters, in particular those corresponding to the CsI - Teflon boundary, and tried to select the setup which best reproduces the data. In addition, we investigated the effect of changing the CsI absorption length.

4.3.3 Comparison with experimental measurements

We present here the results of the simulations for the relative response of the PMT for each position, along with the experimental data for comparison. We perform the comparison for several combinations of the optical parameters to find the set-up that best matches the data. For each of the following cases, we set the parameter Reflectivity to 0.99 (Janecek 2012a) and the absorption length to 2 m for the CsI crystal (Alvarez Franco 2022).

We first compare the results for ground painted surfaces between CsI and Teflon; first, we study the effect of having different levels of roughness, which is parametrized by σ_α. The left panel of Figure 4.8 shows the relative responses for each position with a GBP surface, 100% specular spike (SS) reflection and three different values of σ_α (2°, 12°, 40°), together with the experimental data. The results show that having a higher roughness corresponds to better agreement with the data. In particular, better agreement is achieved for $\sigma_\alpha = 40°$, where the deviation from the experiment is, on average, below ∼5%. For $\sigma_\alpha = 12°$ the deviations increase up to ∼ 10%, while for $\sigma_\alpha = 2°$ they lie near 10% on average with peaks of ∼ 15% for beams 23 and 28 (i.e., the ones closer to the PMT). The Lambertian (L) reflection probability can be set for `GroundBackPainted` surfaces, whereas for `GroundFrontPainted` only Lambertian reflection can occur. The right panel of Figure 4.8 shows the relative responses for a fixed $\sigma_\alpha = 40°$, for both GBP and GFP, with 100% Lambertian reflection; for comparison, also the previous case with 100% SS reflection with $\sigma_\alpha = 40°$ is shown, along with the experimental data. The result with SS reflection represents the best agreement with the data; changing reflection to the Lambertian type slightly worsens the comparison for GBP, with larger deviations of a few %, while selecting a GFP surface results in a larger discrepancy (∼ 15% deviation far from the PMT and ∼ 25% nearby).

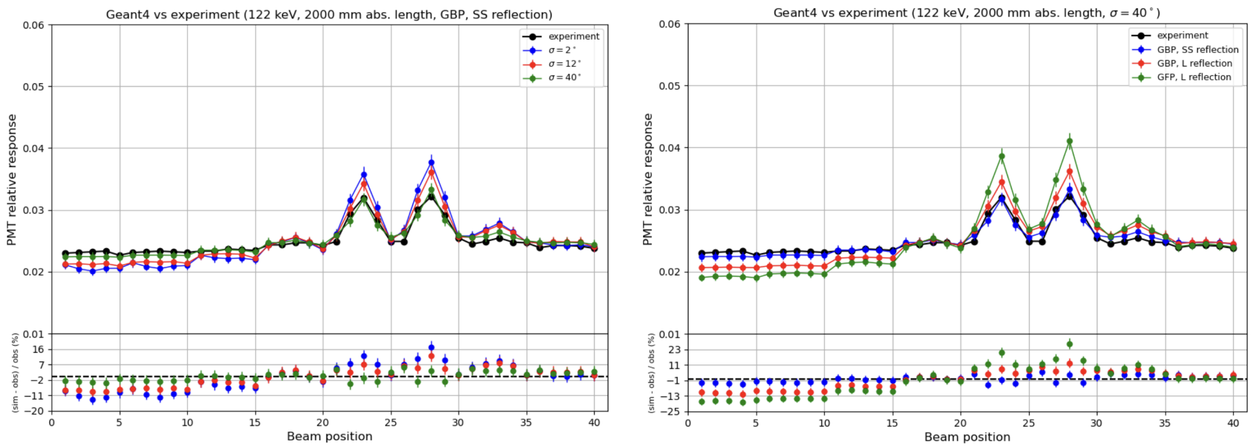

Figure 4.8 – Left panel: PMT relative responses from the simulation for each beam position, for a GBP surface, 100% SS reflection and different values of σ_α (see the legend). Right panel: PMT relative responses from the simulation for each beam position, for $\sigma_\alpha = 40°$ with a different type of surfaces and reflection (see the legend). For both panels, the black points refer to experimental data and the bottom panels show the residuals for the simulations with respect to experimental data.

In the left panel of Figure 4.9, we show the results for a PBP surface for two different levels of roughness, a very low one ($\sigma_\alpha = 1.3°$ as applied in Janecek and Moses 2010) to best mimic the polished nature of the surface, and a higher one ($\sigma_\alpha = 40°$). Note that the roughness here is actually associated with the CsI-gap interface and not with the wrapping, which is instead perfectly smooth. Increasing the roughness does not significantly impact the distribution, contrary to the ground surfaces, and both cases show a very good agreement with the data (deviations below $\sim 2\%$ far from the PMT and $\sim 5\%$ nearby). The right panel of Figure 4.9 shows the case with PBP and 100% Lambertian reflection. As for the PBP with SS reflection, different levels of roughness provide similar results within the errors. However, a higher roughness results in better agreement with the experiment for beam positions far from the PMT, with respect to the case with smaller σ_α. Still, the Lambertian reflection generally worsens the comparison with the data with respect to having 100% SS reflection, also for polished painted surfaces, since the deviations from the experiment are roughly 10% either far from or near the PMT.

In the left panel of Figure 4.10, we compare the relative responses with fixed $\sigma_\alpha = 40°$ and PBP surfaces, having SS reflection in one case and L reflection in the other case. We also include the case with PFP surface with Lambertian reflection (which, in this case, is the only option). As for ground surfaces, SS reflection provides the best results, with discrepancies lower by a few % than those of the other cases. It is also interesting to compare ground and polished painted surfaces, in particular the cases with the best agreement with the data for both categories: (polished and ground) back painted surfaces, with SS reflection and with $\sigma_\alpha = 1.3°$ and $\sigma_\alpha = 40°$ respectively. As shown in the right panel of Figure 4.10, the two results are similar and in agreement, within the statistical error, with the experimental data.

So far, for all cases we kept the absorption length of the CsI crystal fixed to the value of 2 m. In the following, we test a different choice. In the left panel of Figure 4.11, we compare the results for GBP surface, SS reflection, and $\sigma_\alpha = 40°$ with two different absorption lengths: 2

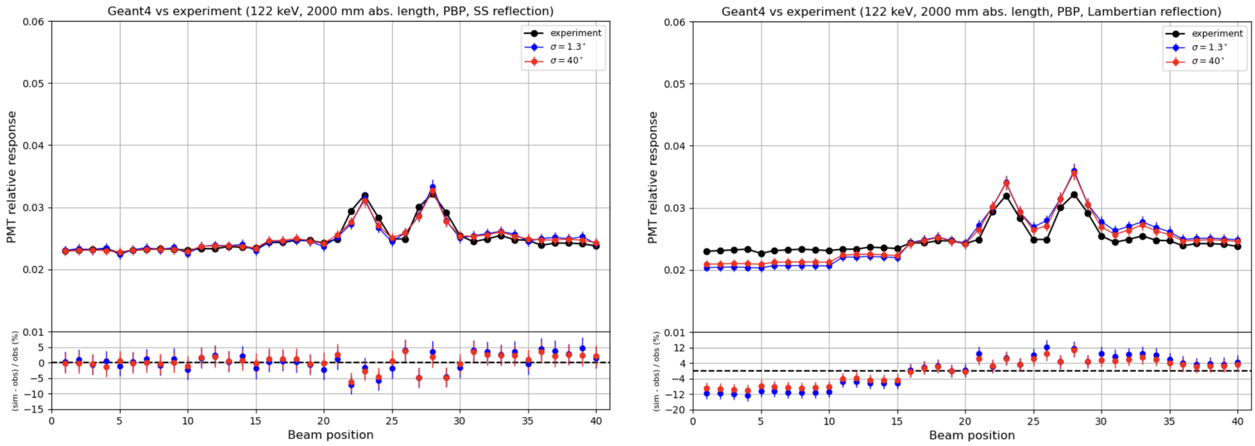

Figure 4.9 – Left panel: PMT relative responses from the simulation for each beam position, for a PBP surface, 100% SS reflection and different values of σ_α (see the legend). Right panel: PMT relative responses from the simulation for each beam position, for a PBP surface, 100% L reflection and different values of σ_α (see the legend). For both panels, the black points refer to experimental data and the bottom panels show the residuals for the simulations with respect to experimental data.

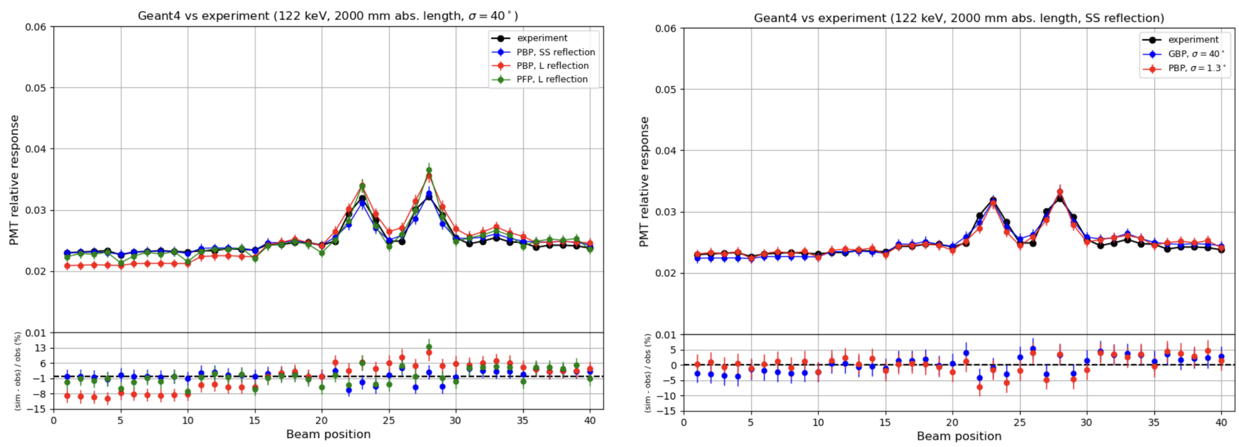

Figure 4.10 – Left panel: PMT relative responses from the simulation for each beam position, for $\sigma_\alpha = 40°$ with different type of surfaces and reflection (see the legend). Right panel: PMT relative responses from the simulation for each beam position, for 100% SS reflection, $\sigma_\alpha = 40°$ and for GBP, PBP surfaces (see the legend). For both panels, the black points refer to experimental data and the bottom panels show the residuals for the simulations with respect to experimental data.

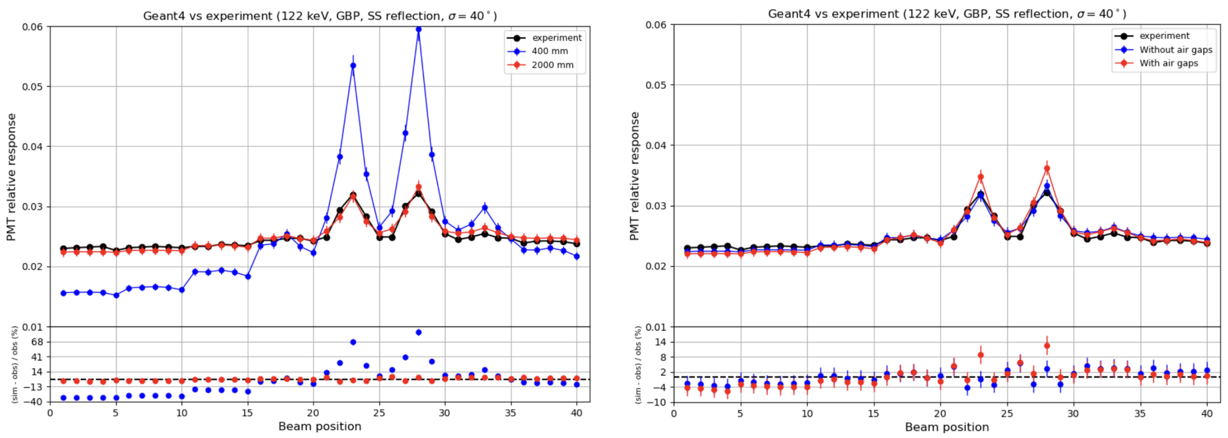

Figure 4.11 – Left panel: PMT relative responses from the simulation for each beam position, for a GBP surface, 100% SS reflection, $\sigma_\alpha = 40°$ and for absorption lengths of 400 mm, 2000 mm (see the legend). Right panel: PMT relative responses from the simulation for each beam position, for a GBP surface, 100% SS reflection, $\sigma_\alpha = 40°$ with and without air gaps above and below the SiPad (see the legend). For both panels, the black points refer to experimental data and the bottom panels show the residuals for the simulations with respect to experimental data.

m (the one used so far) and 40 cm. We can immediately conclude that a shorter absorption length dramatically worsens the agreement with the experiment (the discrepancy is almost 40% far from the PMT and even $\sim 70\%$ for positions 23 and 28), thus ensuring that selecting an absorption length of 2 m represents a good choice. Finally, we investigate the possibility of imperfect coupling between the CsI and PMT by introducing two air gaps above and below the SiPad (both with a thickness of 0.01 mm). We show here, in the right panel of Figure 4.11, the results for GBP, SS reflection, and $\sigma_\alpha = 40°$ in one case with air gaps, and in another case without them. Having the air gaps results in worse agreement with the experimental data since the PMT relative response is overestimated for positions in front of the PMT (23 and 28) and slightly suppressed for further positions. A possible explanation is that optical photons originating from positions farther from the PMT have a higher probability of entering the SiPad at larger incidence angles compared to photons generated near the PMT, and are therefore more likely to be reflected back at one of the two air gaps.

4.3.4 Verification results

We have compared the simulated and experimental relative responses of a CsI crystal coupled to a PMT, for several sets of Geant4 optical parameters. Among all the parameter combinations that we tested, two of them gave the most satisfactory results: the `GroundBackPainted` and the `PolishedBackPainted` surfaces, both with 100% specular spike reflection and with $\sigma_\alpha = 40°$ and $\sigma_\alpha = 1.3°$, respectively. However, the case with $\sigma_\alpha = 1.3°$ represents a more realistic scenario than $\sigma_\alpha = 40°$. A surface with $\sigma_\alpha = 40°$ is extremely rough, and the suspect is that the GBP case matches the experiment only if one artificially tunes the roughness to such high levels. The summary of the accuracy achieved with each optical model is reported in Table 4.3.

Table 4.3 – Maximum discrepancies between simulation and experimental results for each Geant4 optical configuration.

Optical surface	Refl. type	Roughness	CsI abs. length	Max. discrepancy
GBP	SS	$\sigma_\alpha = 2°,\ 40°$	2 m	15%, 5%
GBP	SS	$\sigma_\alpha = 40°$	0.4 m	70%
GBP (with air gaps)[a]	SS	$\sigma_\alpha = 40°$	2 m	12%
GBP	L	$\sigma_\alpha = 40°$	2 m	12%
GFP	L	$\sigma_\alpha = 40°$	2 m	25%
PBP	SS	$\sigma_\alpha = 1.3°,\ 40°$	2 m	5%
PBP	L	$\sigma_\alpha = 1.3°,\ 40°$	2 m	12%
PFP	L	$\sigma_\alpha = 40°$	2 m	14%

[a] Model including air gaps between the optical interfaces.

Chapter 5

Benchmarking of Geant4 simulations for the COSI ACS

Calibration campaigns are essential for accurately assessing detector performance and for benchmarking simulations. Validating simulations against laboratory data ensures that key ACS characteristics, such as energy resolution and threshold, are accurately modeled. Benchmarked simulations can then be used to tune and develop a dedicated Detector Effects Engine that implements realistic effects in simulated data, making it consistent with actual measurements (see Section 5.1). The simulation of the intricate scintillation within the COSI ACS is the perfect use case to apply the new HPC features of the BoGEMMS-HPC simulation framework outlined in Chapter 3. In this chapter, we present the calibration campaign that I performed within the framework of my Marco Polo project at the Space Sciences Laboratory (SSL) in Berkeley, California, USA. We used these calibration measurements (along with additional data taken by the COSI team at the Naval Research Laboratory, NRL) to define the ACS instrument parameters, including the BGO absorption length and electronic noise. We then built a dedicated BoGEMMS-HPC simulation of the calibration experimental apparatus, including the optical physics to model scintillation light transmission and signal readout. This work has been published in Ciabattoni et al. (2025) and builds upon the verification of the Geant4 optical physics, carried out through theoretical models and experimental data with PMTs described in Chapter 4.

Once benchmarked against laboratory measurements, the simulations are used to encode the optical physics effects on the ACS performance into a correction matrix. This correction will be applied to the ACS MEGAlib simulations using the DEE module in Nuclearizer, in order to obtain a realistic model of the ACS both as a vetoing and stand-alone detector. The development of the correction matrix is discussed in detail in Chapter 6.

This chapter is organized as follows. Section 5.1 presents the NRL and SSL calibration campaigns conducted for a lateral ACS wall and the results on its energy resolution and spatial response. Section 5.2 shows the Geant4 simulation of the experimental setup. Section 5.3 discusses the simulation benchmarking and shows the results of the comparison between simulations and laboratory measurements.

Figure 5.1 – Wrapped BGO crystal with coupled SiPM array on the left (picture taken at NRL) and the EM X-wall with the three BGO crystals on the right (picture taken at SSL).

5.1 Calibration measurements of the EM X-wall

We used two calibration campaigns for the COSI ACS, one conducted at the Naval Research Laboratory (NRL, Washington DC) and one at the Space Sciences Laboratory (SSL, Berkeley). The measurements involve the Engineering Model (EM) X-wall, i.e. a lateral ACS panel consisting of 3 BGO crystals and read-out each by SiPM arrays (Figure 5.1). Each BGO crystal consists of a $198 \times 118 \times 23$ mm scintillator block, and a 3×3 array of SiPMs (onsemi J-series 60035) is attached to the 118×23 mm face with the use of a silicone optical coupler. The BGO is wrapped with two layers of VM2000, one layer of Tetratex and one layer of Teflon PTFE. The face where the SiPMs are attached is instead covered only with two layers of Tetratex. These reflective layers prevent scintillation light from escaping the BGO and maximize the SiPM response. The measurements at NRL were taken at room temperature using three radioactive sources: ^{241}Am, ^{22}Na, and ^{137}Cs. The ^{241}Am and ^{137}Cs sources were measured with and without collimation, while ^{22}Na was used only without collimation. At SSL, we used a larger number of radioactive sources spanning a broader range of energy, and only measurements without collimation were taken. The full list of radionuclides used, along with their line energies and the corresponding calibration campaigns, is shown in Table 5.1. The dataset of each calibration campaign consists of a measured ADC channel spectrum for each source and configuration (collimated or uncollimated), and one or more background spectra. The ADC channels are proportional to the charge collected by the SiPMs coupled with the BGO and hence to the energy deposit in the crystal. Each spectrum exhibits one or more photopeaks where the photons are fully absorbed in the BGO, and a continuum tail at lower energy which arises from the Compton scattering of the source photons within the crystal or the surrounding environment. We present here the measurements for the central BGO crystal.

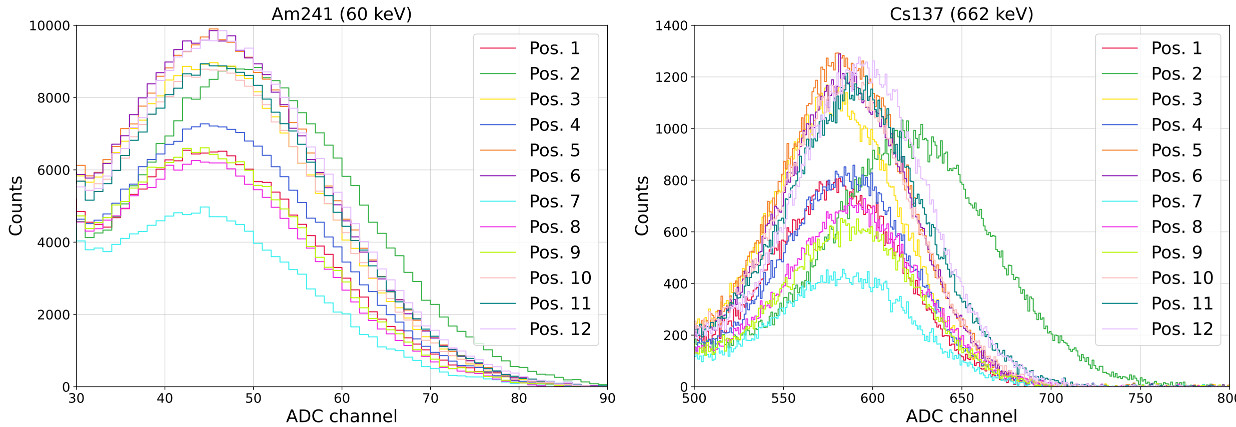

Figure 5.2 – Background-subtracted spectra for ^{241}Am (left) and ^{137}Cs (right) for each collimated position (see the positions in Figure 5.7).

5.1.1 NRL campaign

The measurements at NRL were taken with a multi-channel analyzer and an amplifier. The X-wall was placed on a metal table and covered with a black cloth to prevent light leakage. The uncollimated spectra were taken with ^{241}Am and ^{137}Cs placed on top of a foam block that sits on an aluminum frame, with the sources at ∼32 cm above the central BGO crystal. The ^{22}Na source was instead positioned directly on top of the aluminum housing and centered on the crystal. The collimated measurements were performed using a lead cylinder with a 16 mm diameter hole and a 3D printed aligner in order to place the collimator at 12 different positions to scan the BGO surface. The background spectrum was taken only once for both collimated and uncollimated measurements and subsequently subtracted from the source spectra. Figure 5.2 shows the background-subtracted spectra for ^{241}Am and ^{137}Cs for each collimated position.

5.1.2 SSL campaign

At SSL, we conducted additional calibration measurements for the same EM X-wall, but using a larger number of radioactive sources which are listed in Table 5.1. With respect to the NRL experiment, we had the same lowest energy line (60 keV from ^{241}Am) which is just below the required threshold for the ACS (80 keV), but we extended the measurements up to 1836 keV from the ^{88}Y emission. The spectra were acquired using a Picoscope[1].

The X-wall was placed on a wooden table and anti-static mat and covered with a black cloth. A stand was used to hold the sources at ∼60 cm from the central BGO crystal, without collimation. We measured the background spectra three times, i.e. before, in between, and after the measurements with the sources, in order to check for possible time variations. We subtracted from the source spectra the average of the three background spectra, and included their difference as a systematic uncertainty in the estimated error. An example of the source spectrum, along with the background, and a background-subtracted spectrum is shown in Figure 5.3, for the ^{22}Na. In Appendix A we show all the spectra measured at NRL and SSL.

[1]Picoscope is a digital oscilloscope used for high-speed signal acquisition and analysis.

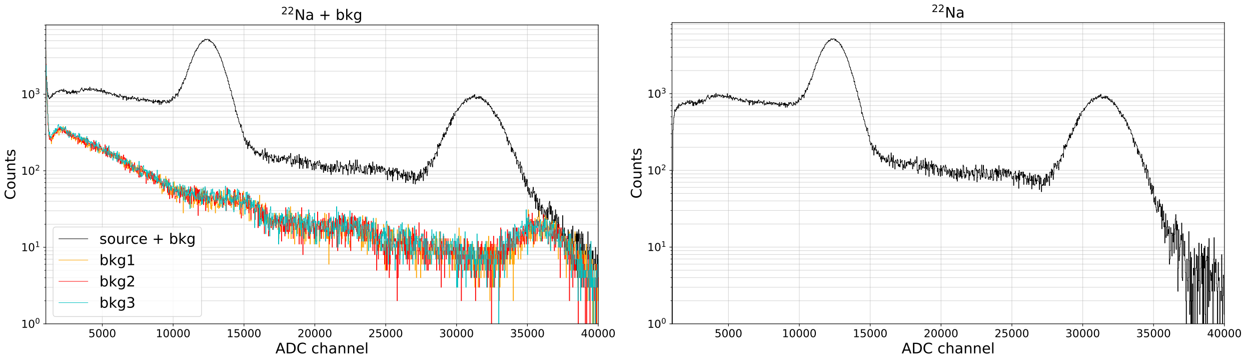

Figure 5.3 – Source + background spectra (on the left) and background-subtracted spectra (on the right) for ^{22}Na, measured at SSL. In the caption of the left plot, 'bkg1', 'bkg2' and 'bkg3' refer to the background measured before, in between and after the data acquisition with the sources.

Table 5.1 – Radioactive sources used for COSI ACS calibration measurements, including their most relevant line energies (neglecting lines below 60 keV), source activity, the campaigns in which they were used, and the configuration (uncollimated or collimated).

Radionuclide	**Line energies [keV]**	**Activity [μCi]**[a]	**NRL**	**SSL**
^{241}Am	60	98.03 (103)	Yes (uncoll. & coll.)	Yes (uncoll.)
^{109}Cd	88	1000.0	No	Yes (uncoll.)
^{57}Co	122, 136	7.65	No	Yes (uncoll.)
^{133}Ba	81, 276, 302, 356, 383	0.05289	No	Yes (uncoll.)
^{22}Na	511, 1275	11.83 (0.7)	Yes (uncoll.)	Yes (uncoll.)
^{137}Cs	662	75.19 (9)	Yes (uncoll. & coll.)	Yes (uncoll.)
^{60}Co	1173, 1333	55.69	No	Yes (uncoll.)
^{88}Y	898, 1836	52.7	No	Yes (uncoll.)

[a] The source activities reported here are based on the most recent measurements available at the time of the experiments. Values refer to the eight SSL sources; values in parentheses correspond to the three NRL sources.

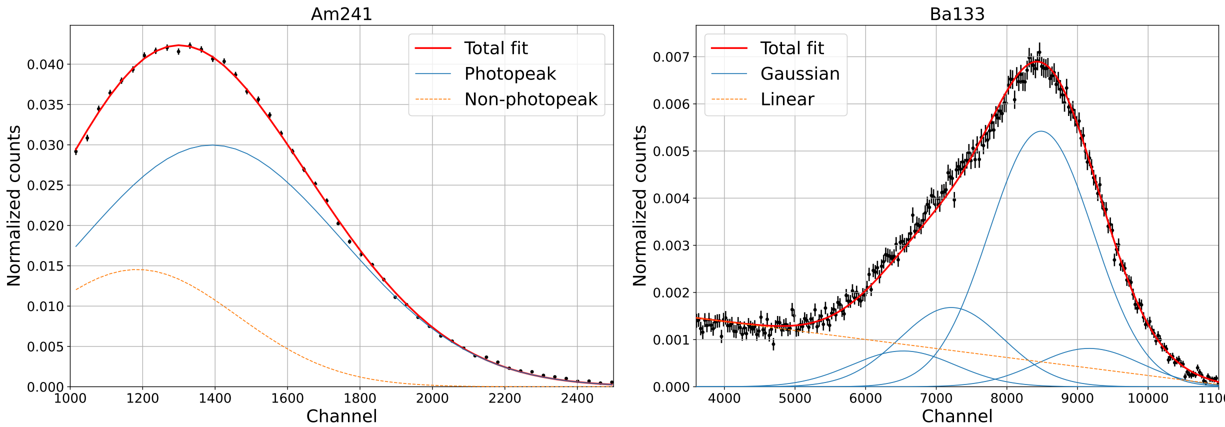

Figure 5.4 – Fit of the photopeaks for ^{241}Am (left) and ^{133}Ba (right) measured at SSL. For ^{241}Am we use an extra Gaussian for the non-photopeak contribution, which is not fully resolved and therefore contaminates the photopeak. ^{133}Ba has 4 overlapping photopeaks (276, 302, 356, 383 keV) and we consider the dominant one at 356 keV for the energy calibration. The linear model takes into account the continuum component of the spectra.

5.1.3 Energy calibration

We performed an energy calibration to convert the ADC channels into energies and obtain for each source the corresponding energy spectrum. The conversion from channel to energy was achieved as follows. We fitted the photopeak with a Gaussian for each source to get the centroid. To account for the non-photopeak component, we used a linear model. For low-energy sources (like ^{241}Am, ^{109}Cd and ^{57}Co), where the Compton tail has a larger overlap with the photopeak, instead of a linear model we found a Gaussian better describes the non-photopeak component. For ^{133}Ba and ^{60}Co, we used multiple Gaussian models to account for their partially overlapping photopeaks. In Figure 5.4 we show the cases (from the SSL dataset) of ^{241}Am, for which we used two Gaussian functions, and ^{133}Ba, for which we used four Gaussian distributions to model the 276, 302, 356, and 383 keV overlapping photopeaks, plus a linear model to account for the continuum. Once we performed the fitting for each source, we associated with each centroid obtained from the fit the corresponding line energy of the source. Finally, we fitted with a linear model ($C = m \cdot E + q$, with C being the channel and E the energy) the centroids and the corresponding energies to find the channel-energy relation, which is shown in Figure 5.5 for both NRL and SSL datasets. The best-fit parameters are: $m = 0.904 \pm 0.001$, $q = -10.1 \pm 0.8$ for NRL, $m = 24.524 \pm 0.002$, $q = -86.2 \pm 1.3$ for SSL.

5.1.4 Experimental energy resolution

From the Gaussian fitting of the photopeaks we obtained, for each source, the full width at half maximum (FWHM). The energy resolution R at a particular energy E is defined as

$$R(E) = \frac{\text{FWHM}}{E} \ , \tag{5.1}$$

In Table 5.2, we list the energy resolution for each measured emission line and both calibration campaigns. We measured an energy resolution of approximately 55-60% at 60 keV, ∼16%

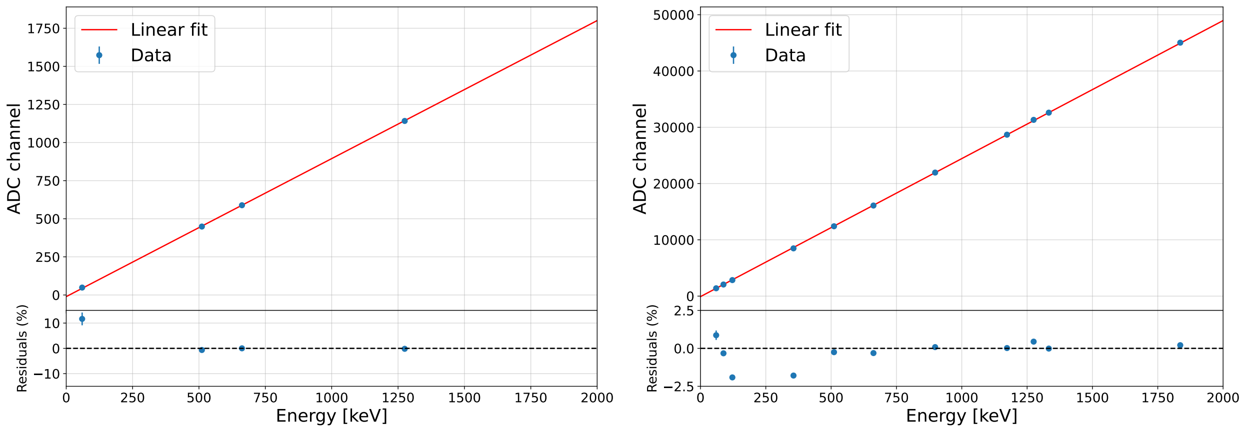

Figure 5.5 – Channel-energy relation for the NRL (left) and SSL (right) datasets and linear fit. The error bars are not visible because they are smaller than the size of the data points.

Table 5.2 – Experimental energy resolution (R) for each emission line and both calibration campaigns.

Emission line (keV)	R (%) at NRL	R (%) at SSL
60	54.5 ± 3.4	60.1 ± 1.8
88	/	44.9 ± 1.6
122	/	37.7 ± 1.1
356	/	19.7 ± 0.4
511	15.8 ± 0.3	16.3 ± 0.1
662	14.0 ± 0.1	14.3 ± 0.1
898	/	12.3 ± 0.04
1173	/	10.8 ± 0.04
1275	10.3 ± 0.3	10.0 ± 0.1
1333	/	10.1 ± 0.1
1836	/	8.5 ± 0.02

at 511 keV and ~8.5% at 1836 keV. Besides the different readout system, the NRL and SSL resolutions are consistent within the 2σ confidence level.

The dependence of the resolution on the energy can be estimated by simple considerations about the photoelectron statistics. In fact, the energy corresponding to the peak pulse is proportional to the number of photoelectrons producing the signal. If the average number of photoelectrons is N, their statistical fluctuations would go as $\sim \sqrt{N}$. If we assume that the width of the photopeak is only due to this statistical uncertainty, we have $R = \mathrm{FWHM}/E \propto 2.355\,\sqrt{N}/N = 2.355\,N^{-1/2}$. However, in a real detector, there are other effects that can change the energy resolution, like the non-uniform optical light collection efficiency across the detector and the electronic noise of the SiPMs. To take into account all these effects, we follow the parametrization given in Bissaldi et al. (2009):

$$R = \frac{\mathrm{FWHM}}{E} = \frac{\sqrt{a^2 + b^2\,E + c^2\,E^2}}{E}\ . \tag{5.2}$$

In the above equation, the parameter a describes the limiting electronic resolution. The parameter b accounts for the statistical nature of the light production, attenuation, photon-electron conversion, and electron amplification. The parameter c represents the position

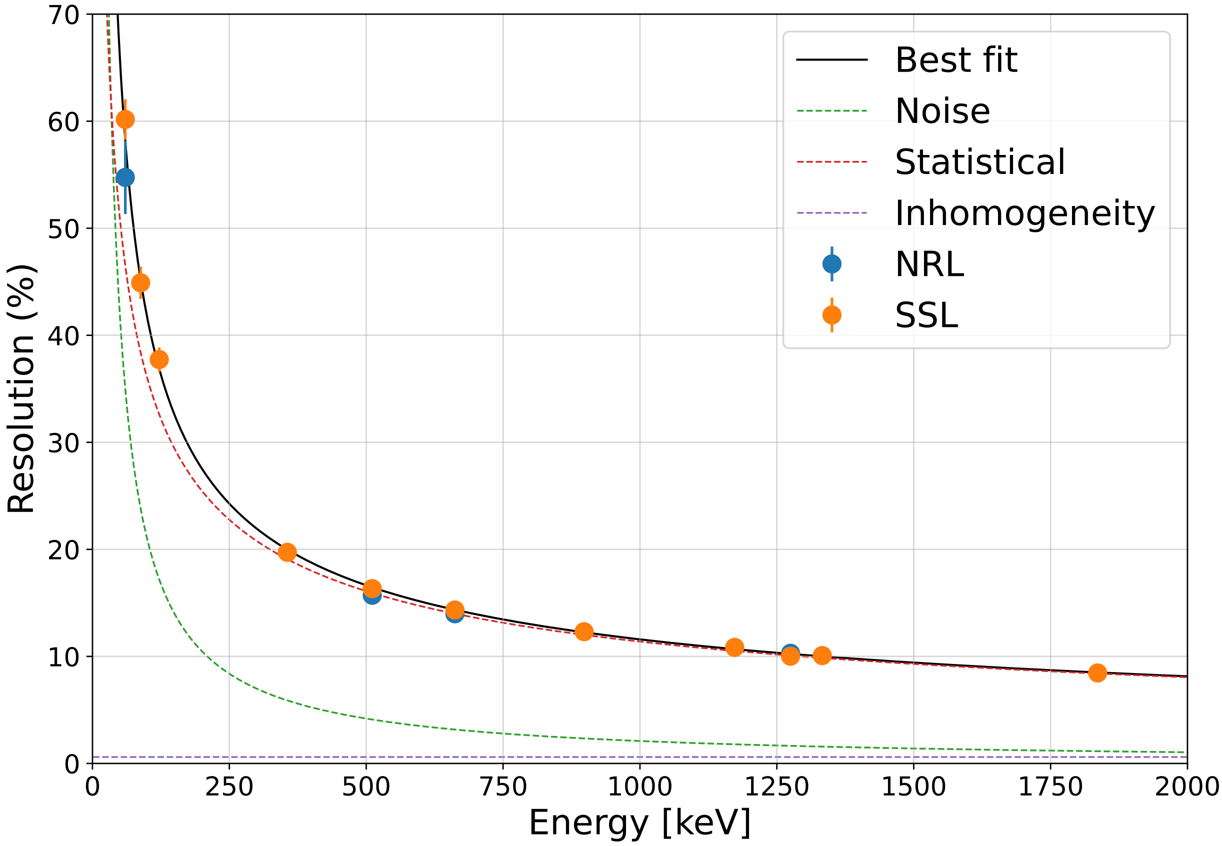

Figure 5.6 – Experimental BGO energy resolution as a function of the energy. The best-fit model (reported in Eq. (5.2)) along with the individual contributions to the resolution are also shown. The best-fit parameters are: $a = 20.98 \pm 1.19$, $b = 3.60 \pm 0.01$, $c = 0.007 \pm 0.004$, with a reduced $\chi^2_{\mathrm{red}} = 3.2$.

dependence of the light transmission from the scintillator to the SiPM. The evaluated energy resolution for each emission line and calibration campaign, along with the best-fit model following Eq. (5.2), is shown in Figure 5.6. We also underline in the same plot the individual components contributing to the resolution. The best-fit parameters are: $a = 20.98 \pm 1.19$, $b = 3.60 \pm 0.01$, $c = 0.007 \pm 0.004$, with a reduced $\chi^2_{\mathrm{red}} = \chi^2/dof = 38.4/12 = 3.2$. The energy resolution is primarily determined by stochastic fluctuations in the generation and detection of optical photons across all energy levels. However, at lower energies, the electronic noise increases significantly and becomes comparable to the statistical fluctuations, thereby contributing notably to the overall resolution. The inhomogeneity component, which acts as an energy-independent factor, has a negligible contribution to the overall energy resolution.

5.1.5 BGO response spatial distribution

The measurements with collimated sources conducted at NRL were used to scan the response across the BGO surface. In fact, different interaction positions of the high-energy photons within the crystal imply a different generation location for the optical photons, possibly affecting their chance of being detected. Having an enhanced/suppressed number of detected optical photons would produce a shift of the photopeak towards higher/lower channels (which indeed are proportional to the number of detected optical photons by the SiPMs), as shown in Figure 5.2. This shift introduces an additional uncertainty in the determination of the

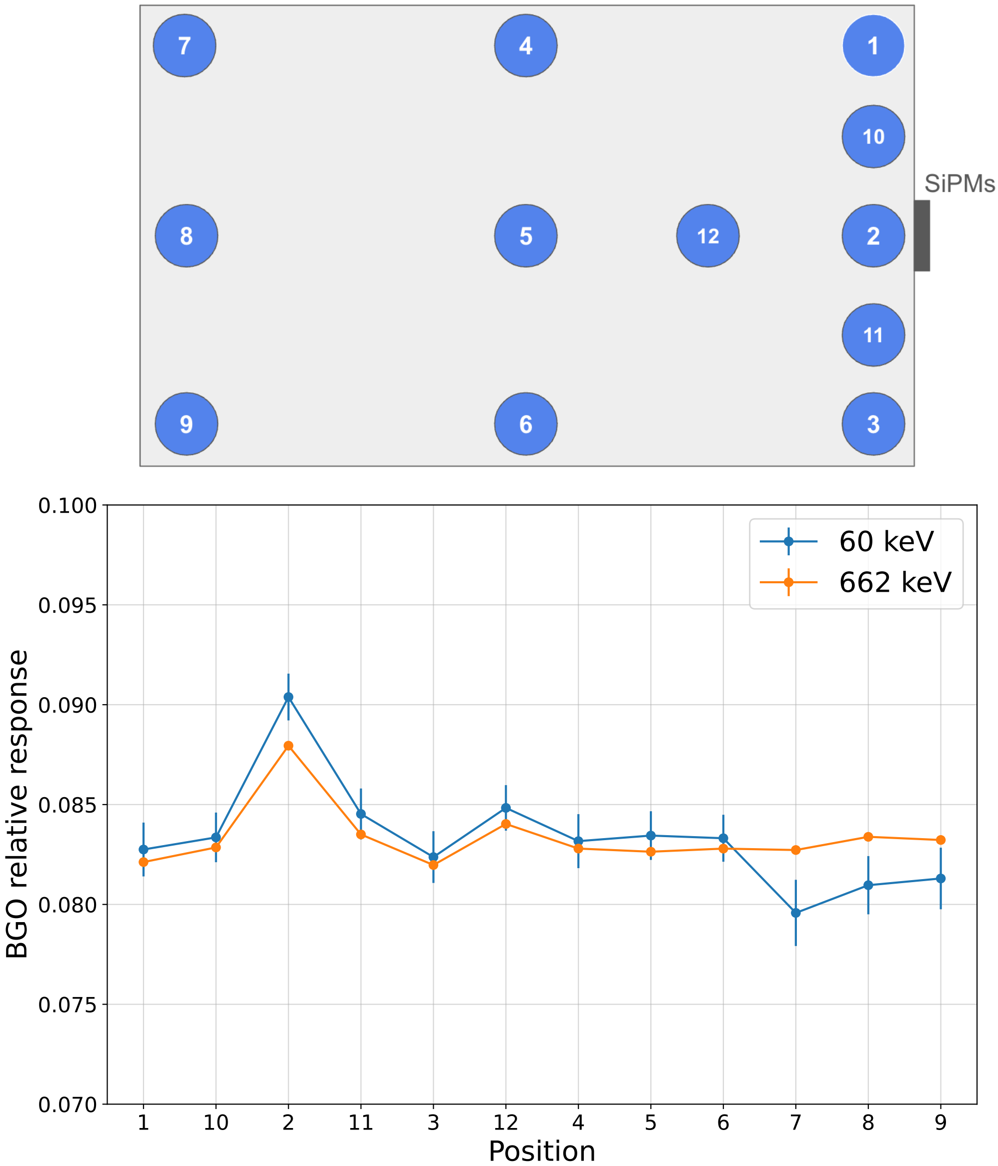

Figure 5.7 – On the top: positions of the collimated sources on top of the 198 × 118 mm BGO surface. On the bottom: experimental BGO relative response, measured at NRL, as a function of the collimated positions, for both ^{241}Am and ^{137}Cs sources.

energy deposit in the BGO. To assert this effect, we fitted the 60 keV (from ^{241}Am) and 662 keV (from ^{137}Cs) photopeaks for each collimated position and extract the centroid. Since we are interested in the relative response across the BGO surface, we normalized the ADC channels measured at each position by the sum of the channels from all positions. In Figure 5.7 we show the positions of the collimated sources and the corresponding measured relative responses, for both 60 keV and 662 keV lines. The result shows an overall homogeneous response across the BGO surface for both energies. We observe an enhancement of ∼8% with respect to the average response for the position in front of the SiPMs, and a slight suppression of ∼4% far from the SiPMs for 60 keV.

5.2 BoGEMMS-HPC simulation of the EM X-wall

We used the BoGEMMS-HPC simulation framework (see Chapter 3) to reproduce the calibration measurements and produce the X-wall correction matrix. We opted for multi-process parallelization using G4MPI to run the simulations, as this proved to be the most efficient way to speed up execution (see Chapter 3). Regarding the output, we used the FITS data format for which we already have a robust analysis framework.

5.2.1 The Geant4 optical physics setup

The generation and transport of the optical photons are regulated through a set of simulation parameters. The `SCINTILLATIONYIELD` parameter specifies the mean number of optical photons generated per unit of energy deposited in the material. We set the BGO light yield to 8.2/keV (Holl et al. 1988). The `SCINTILLATIONTIMECONSTANT1` parameter defines the decay time of the BGO pulse shape and is set to 300 ns (Moszyski et al. 1981). The `SCINTILLATIONCOMPONENT1` is used to define the spectral shape of the generated optical photons, and we set it according to the emission spectrum shown in Akchurin et al. (2011).

The absorption length of the BGO crystals used in the experiments was not specified, and several values are found in the literature (Roncali et al. 2019; Akchurin et al. 2011). We therefore used the benchmarking to constrain its value (see Section 5.3.3). The BGO crystal has a polished surface and is coated with a reflective material, where optical photons are either reflected or absorbed. The distribution of the reflected photons, from specular to diffusive Lambertian, depends on the BGO and the coating surface finish. In our simulations, the BGO-VM2000 reflective surface is set to `PolishBackPainted`, i.e. a polished coating with the presence of an air gap, with 100% specular reflection and reflectivity of ∼99% @ 480 nm (Janecek 2012b). The BGO-Tetratex interface in the SiPM face is set to `GroundBackPainted`, i.e. a ground coating with an air gap, with 100% diffusive reflection and reflectivity of ∼94% @ 480 nm (Janecek 2012b). Finally, we assign a wavelength-dependent refractive index to the BGO (2.15 @ 480 nm, Williams et al. (1996)) and to the optical pad (1.4118 @ 589 nm, info from the manufacturer).

The SiPMs were simulated as empty, fully absorbing boxes, so that every optical photon that reaches the SiPMs is automatically absorbed. The actual number of optical photons detected was then obtained by applying the SiPM energy-dependent PDE[2] for an overvoltage of 2.5 V (see Figure 5.8). The final output of the simulation is a histogram of the numbers of detected optical photons for each radioactive source. Note that in this simulation setup, there are no thin layers between the scintillation crystal and the photosensitive device, unlike in Section 4.2 for the analytical verification. Here, the number of detected photons is obtained directly from the PDE; therefore, the issue of frustrated transmission does not occur.

[2]The PDE also depends on the incident angle of the optical photons (Nakarmi et al. 2020). In our simulation, we applied the manufacturer-provided PDE uniformly, due to the lack of angular dependence data.

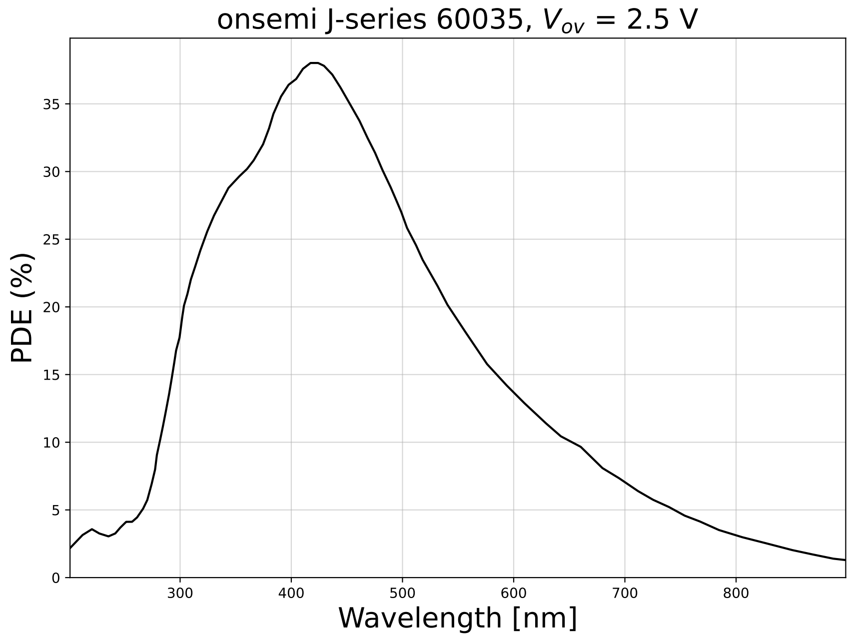

Figure 5.8 – Wavelength-dependent photon detection efficiency for onsemi J-series 60035 SiPM with an overvoltage of 2.5 V (reproduced from the onsemi datasheet).

5.2.2 Simulated mass models

We reproduced in BoGEMMS-HPC the calibration campaigns at NRL and SSL by modeling the corresponding experimental setups as accurately as possible. We built the BGO mass model as a $198 \times 118 \times 23$ mm block coated with two layers of VM2000, one layer of Tetratex and one layer of Teflon on the faces without SiPMs, and with two layers of Tetratex on the face with the SiPMs attached. The $18 \times 18 \times 1$ mm optical pad was placed at the center of the 118×23 mm face, with the Tetratex layers having a cavity of the same size to allow the pad to be in direct contact with the BGO. The 3×3 array of SiPMs was coupled with the optical pad.

We simulated the mass model of the aluminum housing using its CAD model, whose file was provided by NRL. For the NRL experiment, we reproduced the metal table and the plastic piece below the X-wall. The aluminum frame and the foam block were included for the uncollimated ^{241}Am and ^{137}Cs. For the collimated measurements, we reproduced the collimator as a lead cylinder with a central hole of 16 mm diameter. In the SSL simulation we included the wooden table, the anti-static mat, the stand holding the sources and a lead block that was in proximity to the X-wall.

For all the experiments, we used the Geant4 `G4RadioactiveDecay` physics list to simulate the radioactive sources. The sources were modeled as 5 mm disks, with a thickness of 1 mm for NRL and 3.18 mm for SSL, and enclosed in a 25.4×6.35 mm plastic disk. Moreover, ^{22}Na for NRL and ^{133}Ba, ^{60}Co, ^{88}Y for SSL were further encased in plastic and placed on top of a 1 mm thick foam piece, and we included these materials in the mass model.

Table 5.3 – List of simulation parameters used for the benchmarking.

Parameter	Value/Range	Fixed/Free	References
BGO light yield	8.2/keV	Fixed	Holl et al. (1988)
BGO decay time	300 ns	Fixed	Moszyski et al. (1981)
BGO emission range	330-620 nm (peak @ 480 nm)	Fixed	Akchurin et al. (2011)
BGO refractive index	2.15 @ 480 nm	Fixed	Williams et al. (1996)
Optical pad refr. index	1.4118 @ 589 nm	Fixed	Manufacturer
VM2000 reflectivity	99% @ 480 nm	Fixed	Janecek (2012b)
VM2000 reflection type	Specular	Fixed	Janecek (2012b)
Tetratex reflectivity	94% @ 480 nm	Fixed	Janecek (2012b)
Tetratex reflection type	Diffusive	Fixed	Janecek (2012b)
PDE	38% @ 420 nm	Fixed	onsemi datasheet
BGO absorption length	5 m	Free	/
Electronic noise	21 keV	Free	/

5.3 Simulation benchmarking for the EM X-wall

The results presented in this section are based on simulations using a BGO absorption length of 5 m, selected as the optimal trade-off between agreement with collimated and uncollimated source measurements. Increasing the absorption length enhances the spatial homogeneity of the BGO response; however, it also decreases the FWHM of the photopeaks, hence worsening the comparison with the experimental energy resolution. This decrease occurs because a larger number of detected optical photons, thanks to the increased absorption length, reduces statistical fluctuations, thereby improving the simulated energy resolution. The chosen value of 5 m provides a balance that maintains acceptable the benchmarking using both collimated and uncollimated sources, with an uncertainty of the order of ±1 m (see Section 5.3.3). We carried out this trade-off with the goal of limiting the discrepancy with the experiment to approximately 20%, which is the target accuracy provided by the COSI team.

Another key parameter of the benchmarking is the electronic noise, which contributes to the experimental energy resolution, as discussed in Section 5.1.4. While the Geant4 simulation models the statistical nature of optical photon production and detection, as well as light transport inhomogeneities within the crystal, it does not account for the electronics of the SiPMs. We therefore incorporated the electronic noise in the simulation by adding a Gaussian smearing to the simulated energies, with a FWHM corresponding to the experimental value of $a \simeq 21$ keV.

Table 5.3 summarizes the key parameters relevant to the benchmarking, including their values or ranges, whether they were treated as fixed or free in the analysis, and the corresponding references.

5.3.1 Relative response

We used the collimated simulations to compare the response spatial distribution with the experimental one, as done in Section 4.3. Following the same procedure conducted for the experimental data (Section 5.1.5), we quantified the response of the ACS for each position of the collimated source. We fitted with a Gaussian the simulated photopeaks for both 60 and

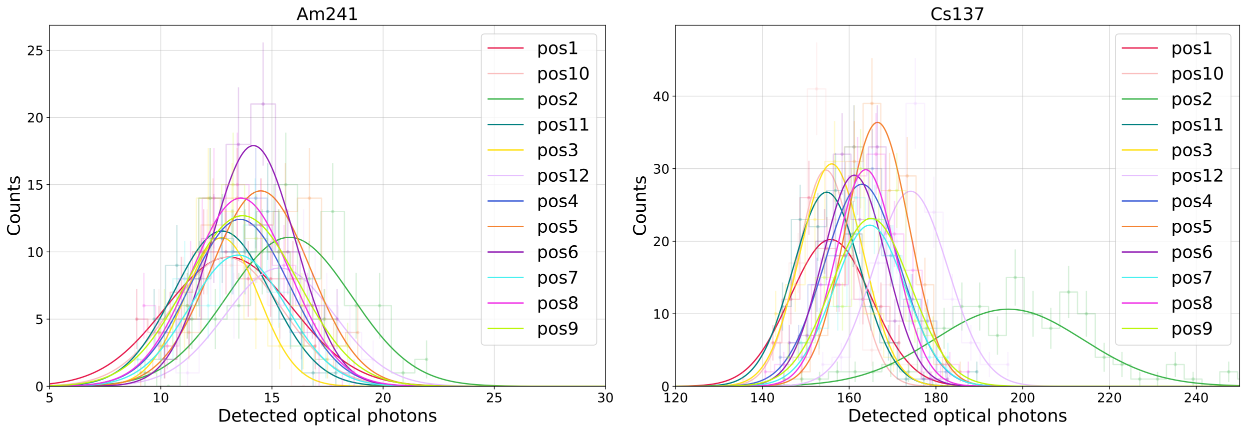

Figure 5.9 – Simulated photopeaks for each position and corresponding Gaussian fit for 60 keV (left) and 662 keV (right).

662 keV (Figure 5.9) and assigned to each position the number of detected optical photons, that defines the ACS response. We then normalized such value by the sum of the response of all the positions to get the relative response. This normalization also allowed for a direct comparison between the experimental and simulated relative responses, as ADC channels and $N_{\rm det}$ are linearly proportional.

The uncertainty in the positioning of the collimated sources, estimated to be ∼1 mm, represents a source of systematic error in the response estimation. We quantified its impact by repeating the simulations with the source positions shifted by 1 mm in four orthogonal directions and evaluating the change in response for each configuration. The maximum change in the response due to the positioning error was taken into account as a systematic error and summed with the statistical one. The same systematic error is also included in the measured response of Figure 5.7.

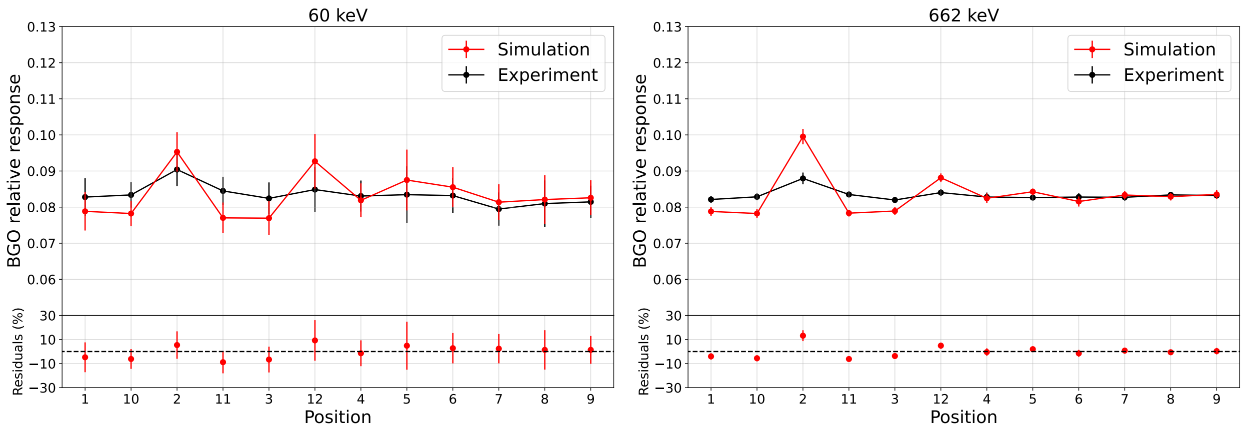

Figure 5.10 – Comparison between experimental and simulated relative response for each collimated position, for ^{241}Am (left) and ^{137}Cs (right).

In Figure 5.10 we show the simulated relative response for each position, for both ^{241}Am and ^{137}Cs, compared with the collimated measurements obtained at NRL. The error bars represent the total (statistical and systematic) error for both simulated and experimental data points. All the residuals shown in the following plots are defined as $(S - L)/L$ (%), where S and L are simulated and laboratory data, respectively. Overall, the simulation reproduces the homogeneity of the experimental response with a maximum discrepancy of

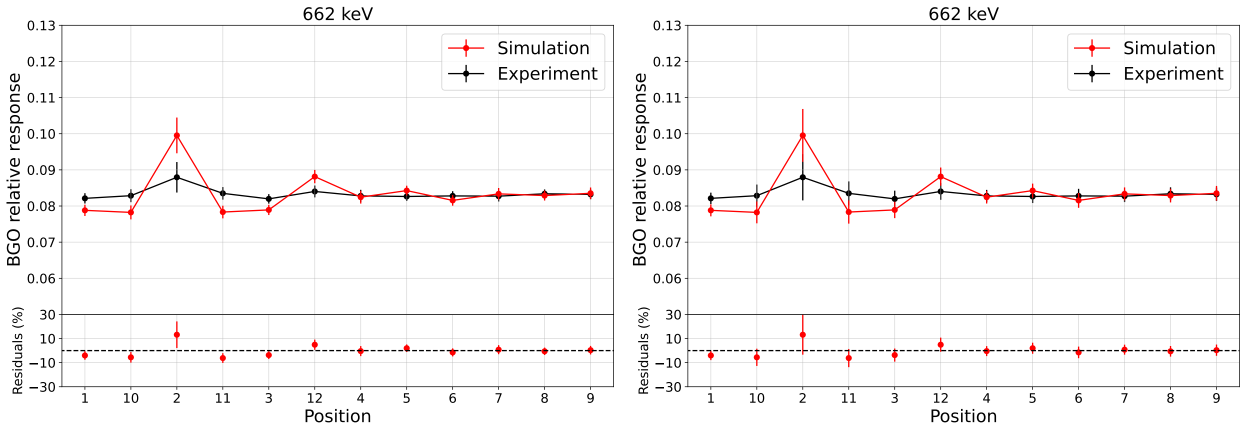

Figure 5.11 – Comparison between the experimental and simulated relative response, using a systematic uncertainty on the positions of 5 mm (left) and 8 mm (right).

∼10%. The largest discrepancy is observed near the SiPMs, where the simulation is predicting a larger inhomogeneity compared with the laboratory measurements. Many factors may contribute to this discrepancy. Increasing the BGO absorption length to 6 m reduces the discrepancy between simulation and experiment by approximately ∼15%. Another potential source of discrepancy is the uncertainty in the positioning of the collimated sources. Although a systematic uncertainty of 1 mm was already included, this may underestimate the true positional uncertainty. To investigate this, we considered a systematic offset of 8 mm, corresponding to the radius of the collimator hole. In Figure 5.11, we show the comparison between the experimental and simulated relative response, using a systematic uncertainty on the positions of 5 mm and 8 mm. The response at the position directly in front of the SiPMs is particularly sensitive to small positional shifts. Assuming a systematic uncertainty of 8 mm, the simulated results become consistent with the experimental measurements within the associated uncertainties. Nevertheless, with the current setup, the maximum discrepancy between the experiment and the simulated response spatial distribution remains well below 20%, consistently meeting our target.

5.3.2 Spectral comparison

We used the uncollimated simulations to compare the energy resolution and energy spectra with the experimental values. The simulated spectra were first converted to energy and then compared with the calibrated spectra obtained in the laboratory with uncollimated sources. The conversion from optical photons to energy was obtained following the same procedure as for the experimental energy calibration, described in Sec. 5.1.3. The resulting photon-energy relation is shown in Figure 5.12. The best-fit parameters are: $m = 0.2522 \pm 0.0002$, $q = -1.13 \pm 0.05$ for NRL, $m = 0.2506 \pm 0.0002$, $q = -1.36 \pm 0.05$ for SSL.

After applying the energy calibration for the simulations, we computed the simulated energy resolution for each energy line, fitted the data with the model described in Eq. (5.1), and compared it to the experimental values.

In Figure 5.13, we show the comparison between the experimental energy resolution and the simulated results, with and without the inclusion of electronic noise. In the absence of electronic noise, the best-fit parameters are: $a = 0.63 \pm 0.58$, $b = 2.67 \pm 0.02$, and $c = 0.06 \pm 0.01$. The value of a is consistent with zero, confirming the lack of electronic noise in the simu-

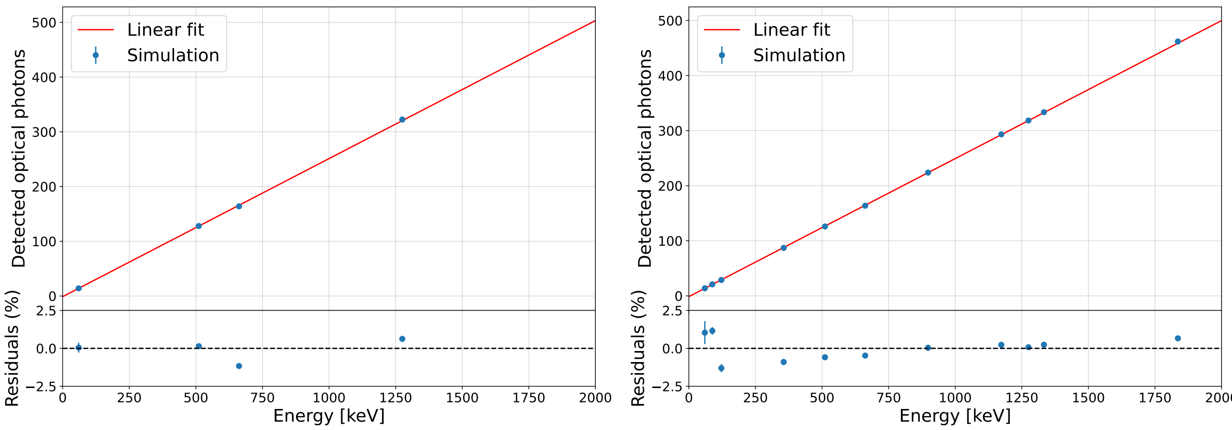

Figure 5.12 – Photon-energy relation for the simulated NRL (left) and SSL (right) datasets and linear fit.

lation, as expected. With the inclusion of the electronic noise, the new best-fit parameters are: $a = 20.77 \pm 0.69$, $b = 2.63 \pm 0.06$ and $c = 0.06 \pm 0.01$. While b and c remain largely unchanged, the value of a increases significantly, becoming consistent with its experimental counterpart. Incorporating the experimental electronic noise into the simulation notably improves the agreement, particularly at low energies, resulting in a maximum discrepancy of 20% across the investigated energy range. The maximum discrepancy arises in the range 100 keV $\lesssim E \lesssim$ 500 keV, where the simulation predicts a better energy resolution compared with the experimental results. Although we accounted for electronic noise in the simulation, some components contributing to the statistical term may still be missing, such as gain fluctuations arising from statistical variations in breakdown voltage and temperature (Dorosz et al. 2013). This effect would introduce additional broadening in the measured photopeaks, leading to a degradation of the experimental energy resolution. Unlike electronic noise, this additional source of statistical uncertainty cannot be isolated and extracted from the fit of the experimental energy resolution.

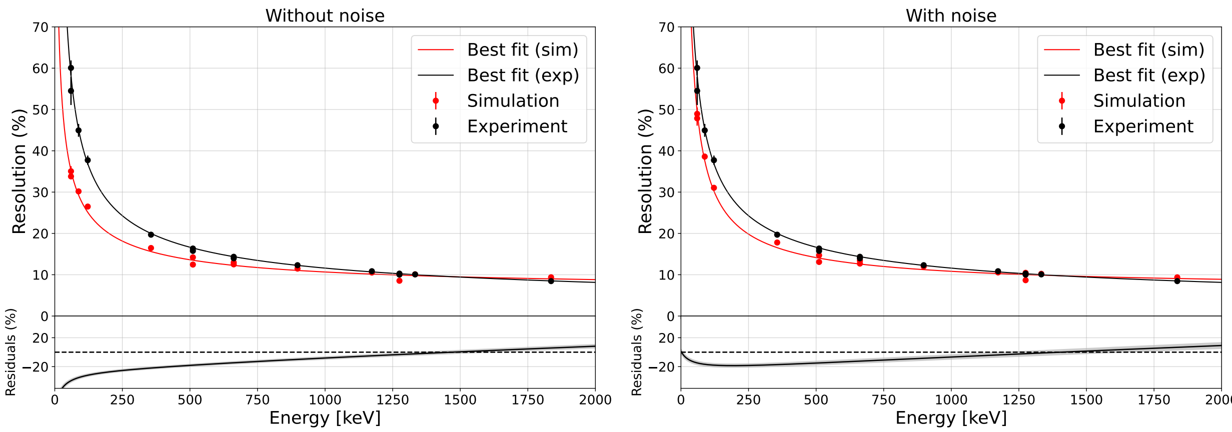

Figure 5.13 – Experimental and simulated energy resolution for each source energy and the corresponding best-fit model, without (left) and with (right) the application of the electronic noise to the simulation.

Once the electronic noise is applied to the energy resolution, we compared the experimental and simulated spectra. We normalized both spectra by the total number of counts in the photopeak, within one FWHM from the peak. The error in the simulated count rate is the sum of the statistical fluctuations propagated with the uncertainty in the photon-to-energy conversion and the electronic noise. For the experimental counts, we propagated the uncertainty in the channel-to-energy conversion and the systematic in the background subtraction. On average, the systematic uncertainty contributes to the total error by ∼25% for the simulation, and by ∼45% for the experiment. In Figure 5.14, we show the comparison between the SSL experimental and simulated ^{137}Cs spectrum.

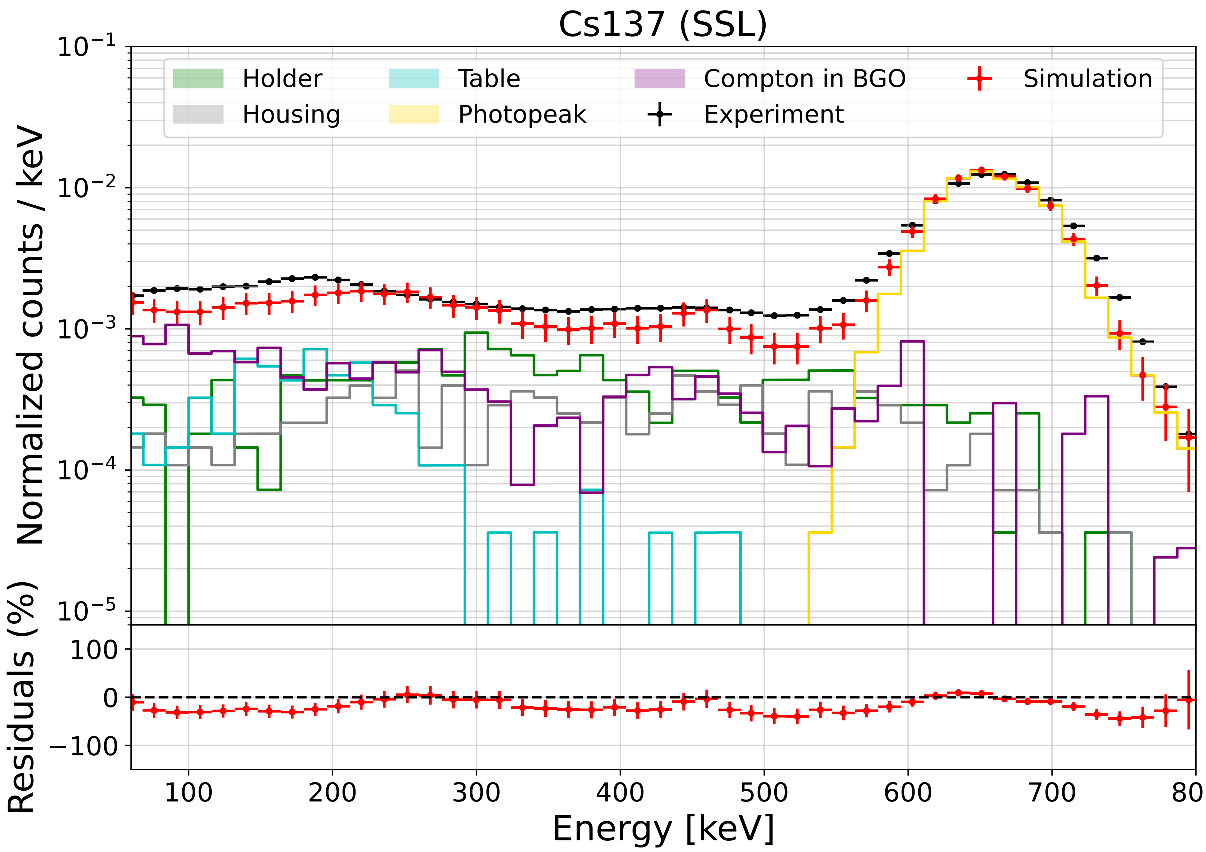

Figure 5.14 – SSL experimental and simulated ^{137}Cs spectrum, along with the scattering components (i.e. those photons which scattered first on that component before entering the BGO) contributing to the simulated spectrum.

The simulation provides an opportunity to investigate the individual components contributing to the spectrum: the photopeak arises from photons depositing their entire energy in the BGO, while the low-energy tail results from Compton scattering of the primary photons within the BGO or the surrounding environment. At $E \gtrsim 300$ keV, the source holder is the dominant contribution, followed by the Compton in the BGO and the aluminum housing. At $E \sim 200$ keV, an important contribution arises from back-scattered photons from the table, while the lowest energy region is dominated by Compton scattering in the BGO. This continuum region of the spectrum is therefore highly dependent on the specific experimental setup, which can vary between calibration campaigns and with unknown density and composition of the surrounding materials, increasing the discrepancy with the simulation. For example, the table used in the SSL calibration is made of wood, a material with a wide range of densities (Ishak et al. 2015), introducing an inherent source of uncertainty. Additionally, the simulation does not account for all objects present in the laboratory, focusing instead on those immediately surrounding the BGO detectors. In Figure 5.15 and Figure 5.16, we present the comparison for all sources and calibration campaigns. The error bars in the simulated spectra are generally larger than those in the experimental ones, as the laboratory

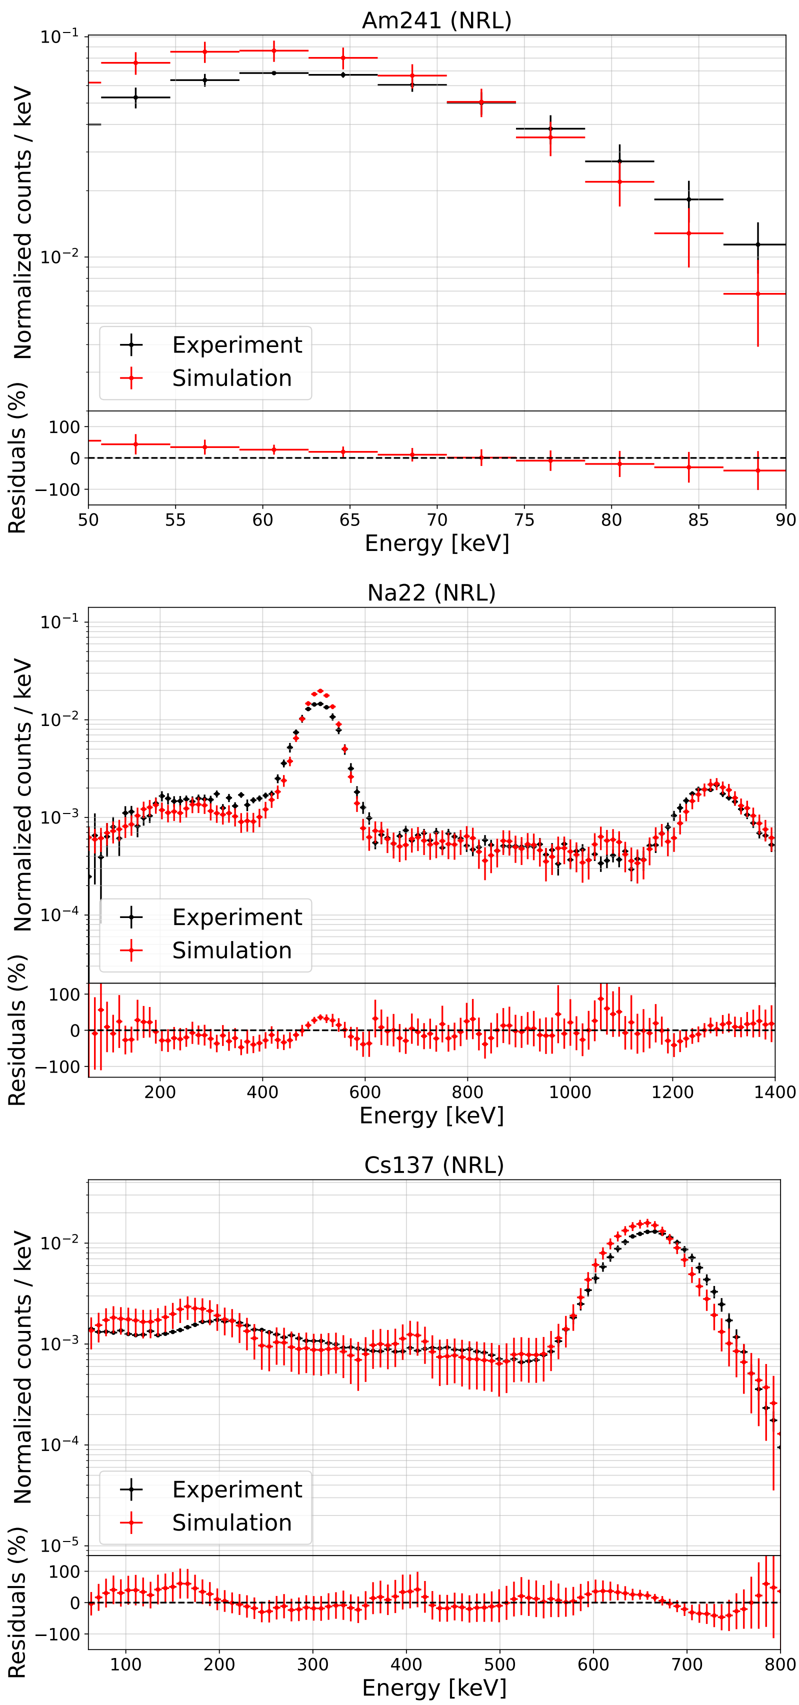

Figure 5.15 – Comparison between NRL experimental (black) and simulated (red) energy spectra for each radioactive source, namely ^{241}Am, ^{22}Na, and ^{137}Cs.

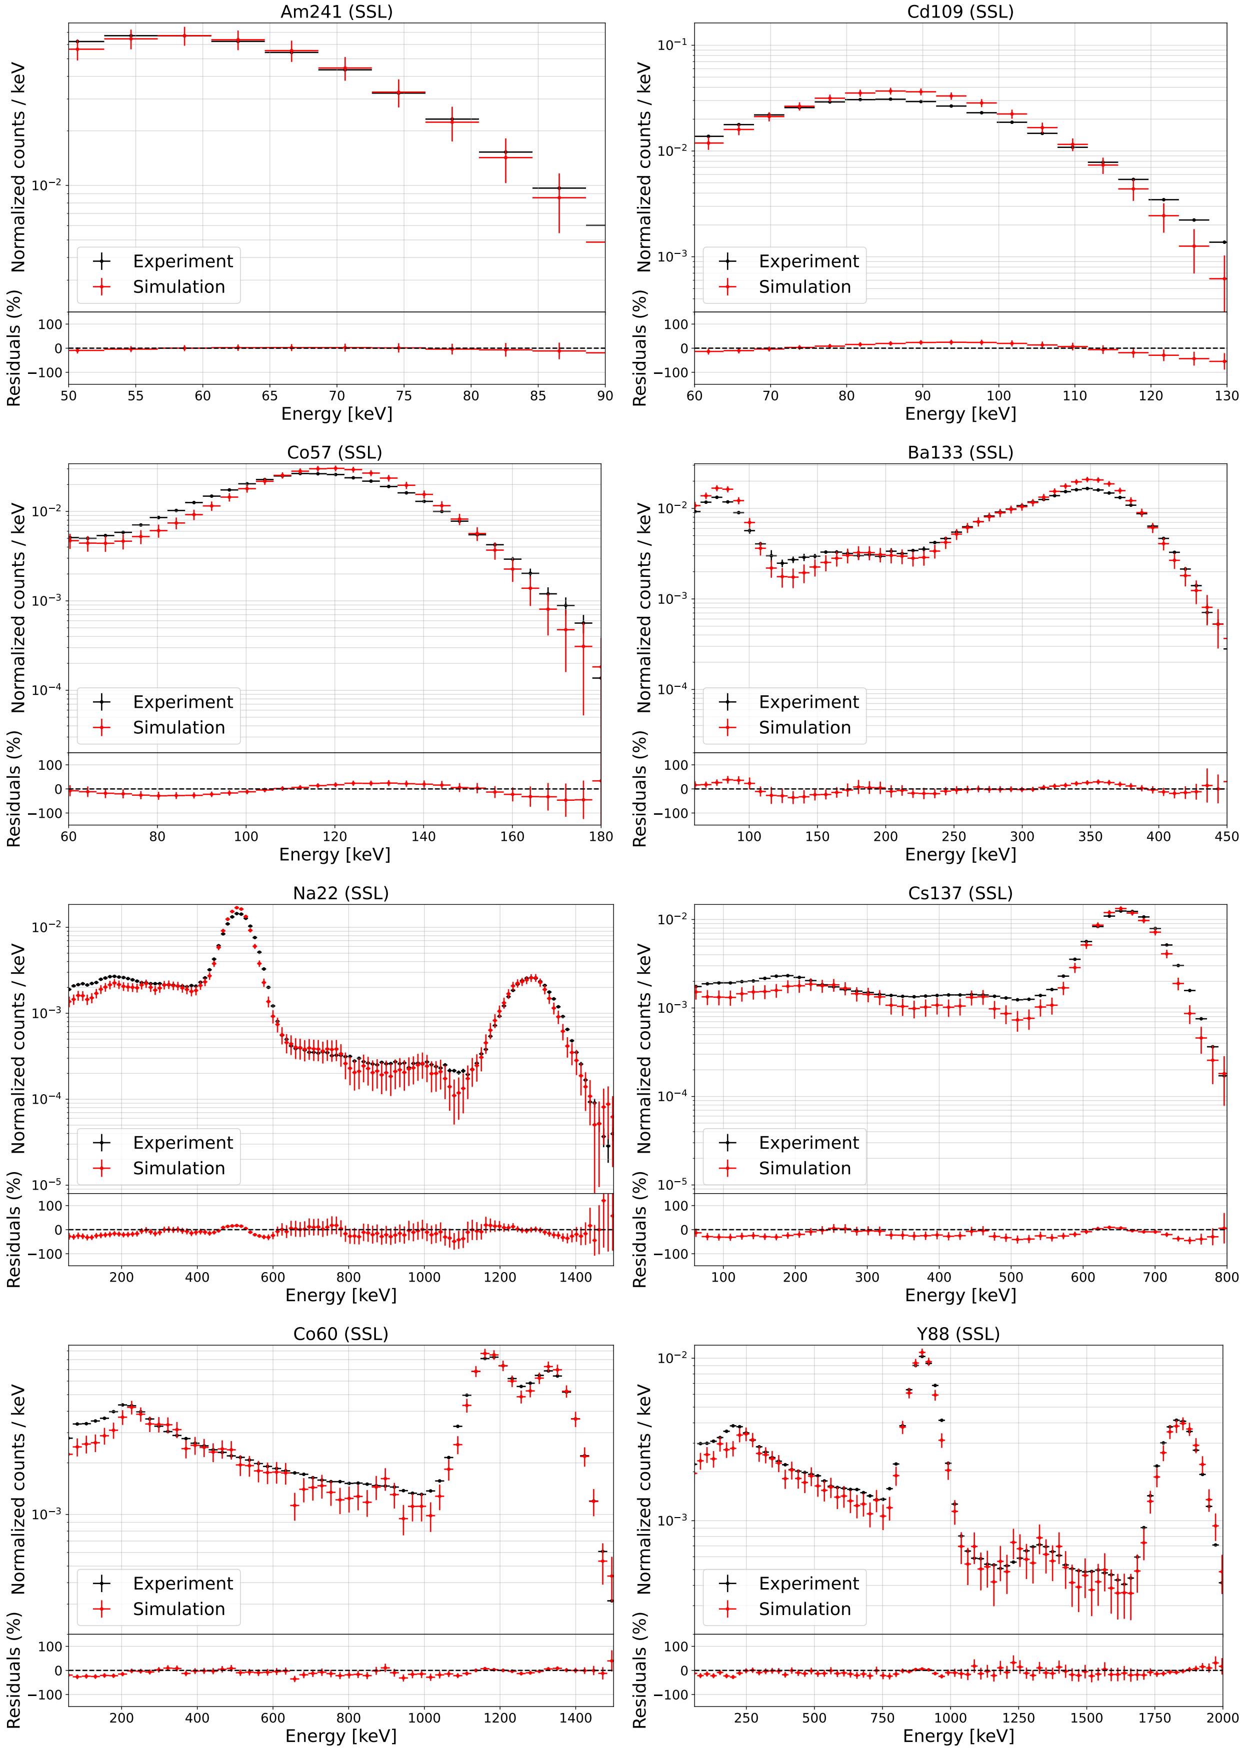

Figure 5.16 – Comparison between SSL experimental (black) and simulated (red) energy spectra for each radioactive source, namely ^{241}Am, ^{109}Cd, ^{57}Co, ^{133}Ba, ^{22}Na, ^{137}Cs, ^{60}Co, and ^{88}Y.

measurements have significantly higher counts per bin, resulting in lower statistical fluctuations. Also, the statistics vary among the sources due to differences in source activity and their distances from the BGO. The average discrepancy between simulation and experiment is ∼12% in reproducing the photopeaks and ∼18% in reproducing the Compton tails. In particular, all photopeaks are reproduced within the 2σ confidence level, except for the 511 keV line from the ^{22}Na source. In this case, the experimental photopeaks from the NRL and SSL measurements are broader than the simulated ones, with the discrepancy reaching up to 3σ. A likely explanation for this behavior is the Doppler broadening of 511 keV photons, an intrinsic effect arising from electron-positron annihilation which additionally degrades the experimental energy resolution (do Nascimento et al. 2005). Although this broadening is more visible in high-resolution detectors such as germanium detectors, it might still contribute to the experimental BGO resolution and increase the observed discrepancy with the simulation. Regarding the continuum regions, a larger discrepancy is observed in the low-energy portion of the spectra for SSL high-energy sources. In particular, the simulated counts are systematically lower than the laboratory ones. As previously discussed, this portion of the spectrum is strongly influenced by the experimental environment, suggesting that some components may be absent in the simulation. In contrast, the low-energy comparison shows better agreement with the NRL measurements.

5.3.3 Selection of the BGO absorption length

In Figure 5.17, we show the simulated relative response and energy resolution for three different values of the BGO absorption length (4, 5, 6 m), compared with the experimental measurements. Our results show that while absorption lengths greater than 6 m would lead to a discrepancy larger than 20% in the energy resolution, absorption lengths below 4 m would result in a discrepancy exceeding 20% in the relative response. This trade-off allowed us to select the BGO absorption length to 5 m, with an uncertainty of the order of ±1 m.

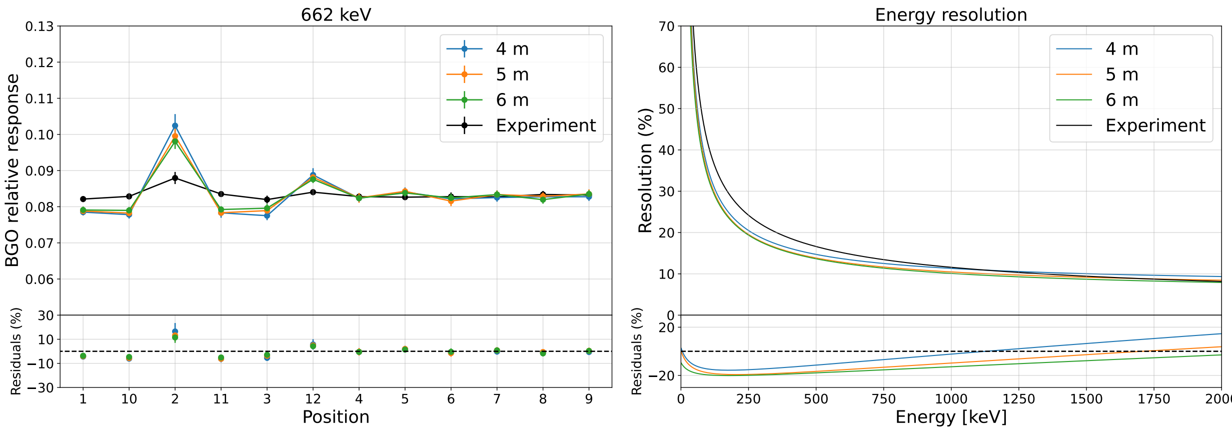

Figure 5.17 – On the left, simulated BGO relative response at 662 keV for different BGO absorption lengths (4, 5, 6 m), compared with the experimental one. On the right, simulated best-fit model of the ACS energy resolution for different BGO absorption lengths (4, 5, 6 m), compared with the experimental best-fit model.

5.3.4 Benchmarking accuracy

We have replicated with BoGEMMS-HPC the experimental setups for the EM X-wall described in Section 5.1 and compared the simulated spatial response distribution (using the collimated sources), the energy resolution, and the energy spectra (using uncollimated sources) with the corresponding laboratory measurements. The BGO spatial response was reproduced with a maximum discrepancy of $\sim 10\%$ in regions close to the SiPMs. Factors contributing to this discrepancy likely include the uncertainty in the BGO absorption length, systematic errors in the positioning of the collimated sources, and intrinsic limitations of the Geant4 optical simulation. The simulated energy resolution matched the experimental one within $\leq 20\%$, with the largest deviations observed in the range $100\,\mathrm{keV} \lesssim E \lesssim 500\,\mathrm{keV}$. In this case, the implementation of the electronic noise, absent in the initial simulation and extracted directly from the experimental data, was crucial in mitigating the discrepancy. Without including electronic noise, the simulated energy resolution would deviate by more than 20% for $E \lesssim 250\,\mathrm{keV}$. We also compared the simulated and experimental energy spectra for each radioactive source. The simulations reproduced the photopeaks and Compton tails with average discrepancies of 12% and 18%, respectively. The larger discrepancy observed in the continuum region of the spectra likely arises from the stronger influence of the experimental environment, which would require a detailed modeling of all components surrounding the EM X-wall in the laboratory setup. Overall, the simulation reproduced the experimental data within 20%, consistent with the target accuracy defined by the team. The benchmarking process also allowed us to constrain two key parameters that were otherwise unknown: the BGO absorption length (5 m) and the electronic noise (21 keV).

Chapter 6

Correction matrix for the ACS

Accurate ACS simulations require a detailed treatment of the production, tracking, and detection of scintillation photons within the BGO crystals. The production and transport of optical photons to the SiPMs introduce biases in the reconstruction of the energy deposited in the crystal, as well as spatial inhomogeneities in the BGO response. Such effects are difficult to capture through laboratory measurements alone, which typically provide only average parameters. In this context, well-benchmarked simulations offer a powerful tool to investigate the impact of optical physics on the ACS performance with much greater detail and accuracy. Since including the optical processes in the ACS MEGAlib simulations would be largely time-consuming, we used the benchmarked BoGEMMS-HPC simulator (see Chapter 5) to develop a correction matrix that encodes the energy correction as a function of the interaction position within the BGO crystal. The correction matrix is then applied to the MEGAlib simulations using Nuclearizer. Each crystal has its own correction file, as differences in size also impact their performance.

This chapter is organized as follows. Section 6.1 discusses how the optical physics affects the ACS performance. Section 6.2 shows the development of the correction matrix for the ACS EM X-wall (the one used for the calibration measurements and benchmarking presented in Chapter 5). Section 6.3 presents the simulations conducted for bottom ACS crystals, highlighting the performance differences due to their smaller size. Section 6.4 explores the impact of different temperatures on the ACS performance. Finally, Section 6.5 shows the development of correction files for the entire ACS using the latest spacecraft mass model.

6.1 Impact of optical physics on the ACS performance

In Chapter 5, we employed dedicated calibration campaigns for a lateral wall of the ACS to benchmark the simulations and ensure their reliability, particularly concerning the ACS energy resolution and response spatial homogeneity. In MEGAlib, ionizing radiation (photons and particles) releases a certain amount of energy that is recorded in the simulation output without detector effects, such as a position-dependent Gaussian noise. We used the BoGEMMS-HPC simulation of the full detection chain up to the SiPM signal to model the energy correction to the MEGAlib energy deposition, which we call *correction matrix*. These corrections are implemented in a dedicated DEE module, bringing the ACS simulations closer to real-world measurements. As illustrated in Figure 6.1, the energy deposit is first converted into scintillation optical photons, a fraction of which is eventually detected by the SiPMs to produce the signal. Using the channel-energy relation obtained from the energy calibration, the signal is finally converted back into the measured energy. This entire process is intrinsically affected by the details of the scintillation process in the BGO, i.e. the generation of optical photons and their transport through the crystal until they are either absorbed or detected by the SiPMs.

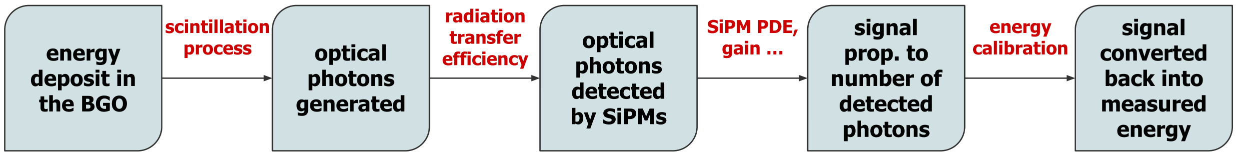

Figure 6.1 – Path from energy deposit in BGO to the actual energy measurement.

As discussed in Section 2.4, accurately simulating the ACS requires accounting for several key parameters, particularly the optical light collection efficiency and the energy resolution. The light collection efficiency depends on several factors, including the BGO absorption length, the reflectivity of the layers wrapping the BGO, and the SiPM PDE. Moreover, the scintillation light collection efficiency varies throughout the crystal, as photons generated at different locations have different probabilities of being detected, depending on their proximity to the SiPMs. This non-uniformity also affects another fundamental performance parameter, the energy threshold, which tends to be higher where the light collection efficiency is lower, and vice versa. The energy resolution, as described in Eq. (5.2), is determined by the SiPM electronic noise, the statistical fluctuations in the number of optical photons, and the spatial inhomogeneity of the BGO light collection efficiency. It can also exhibit a positional dependence and may indirectly influence the number of detected photons and, consequently, the overall quantum detection efficiency.

6.2 Correction matrix for the EM X-wall

We present here the characterization of the optical light collection efficiency, energy resolution and their positional dependence for a BGO X-crystal with dimension $198 \times 118 \times 23$ mm. We used the same simulation parameters set by the benchmarking (see Chapter 5): a BGO absorption length of 5 m and an electronic noise of 21 keV.

6.2.1 Optical light collection and energy resolution map

To achieve sufficient statistics for the production of the correction matrix, we needed to simulate at least ~ 30 photons per spatial and energy bin, leading to a total of about 50000 high-energy photons. Considering that the BGO produces 8.2 optical photons per keV, this corresponds to more than 4000 photons for a 511 keV gamma-ray and more than 10^8 optical photons to generate the correction matrix, translating into ~ 2 yr of simulation time on a single machine. It is therefore clear that such simulations require large-scale HPC infrastructure. Under the ICSC Resource Allocation Committee (RAC), we were granted access to 100000 core-hours on the X2140 CPU-only blades of the Leonardo supercomputer, using the DCGP partition. The simulations were carried out using the multi-tasking capabilities of BoGEMMS-HPC, which provides a speed-up factor of ~ 70 compared to serial execution

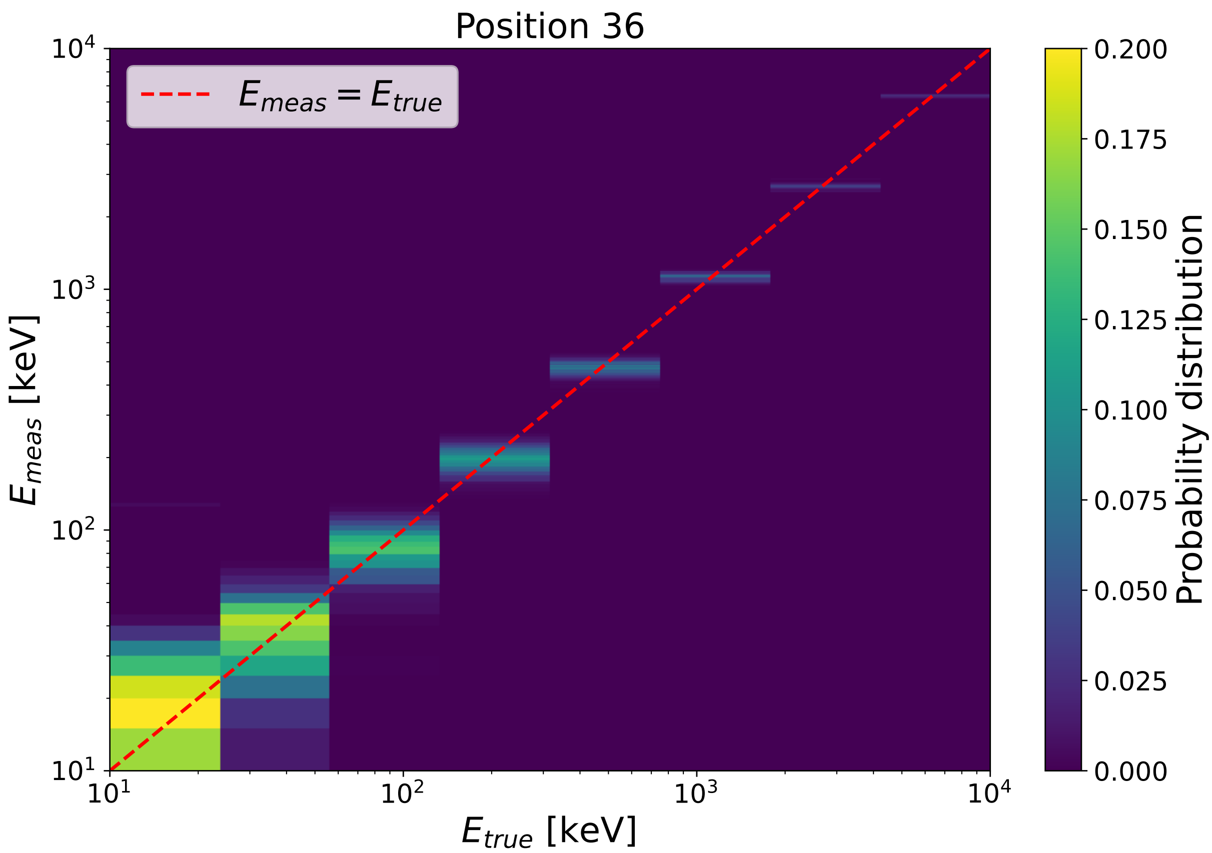

Figure 6.2 – Response matrix (E_{meas} against E_{true}) for the BGO X-crystal for a particular spatial bin.

(see Chapter 3). Thanks to parallel execution and the large computational resources, the simulations were completed in about two weeks.

We simulated photons from 10 keV to 10 MeV hitting perpendicularly the 198×118 mm face of the BGO and, for each photon, we recorded the energy deposit (E_{true}), the 2-dimensional position of the deposit (x and y) and the corresponding number of optical photons detected by the SiPMs[1]. We then used the energy calibration obtained from the SSL simulations (left plot of Figure 5.12) to convert the number of detected optical photons into measured energies (E_{meas}). At the end of this process, we obtained $\{E_{\text{true}}, x, y, E_{\text{meas}}\}$. We then performed a spatial x-y binning of the 198×118 mm BGO surface into 20×10 bins. Within each spatial bin, E_{true} is further divided into 8 logarithmic energy bins spanning from 10 keV to 10 MeV. For each bin, we analyzed the distribution of the differences between E_{meas} and E_{true}, shifted to the corresponding energy bin center, and fit them with a Gaussian model. In Figure 6.2 we show the measured energy distribution against the true energy deposit for a particular spatial bin. In Figure 6.3, we show, for the same spatial bin, the same distributions with the corresponding Gaussian fit.

The fit of the distributions provides, for each spatial and energy bin, two key parameters: the centroids ($\bar{E}_{\text{meas}}$) and the FWHMs. These two quantities are related to two fundamental

[1]For each high-energy photon interacting with the crystal, multiple energy deposits can occur at different locations within the crystal, and each energy deposit generates its own set of optical photons. In this case, we treat each energy deposit independently, each of which is associated with its own x-y position and number of detected optical photons.

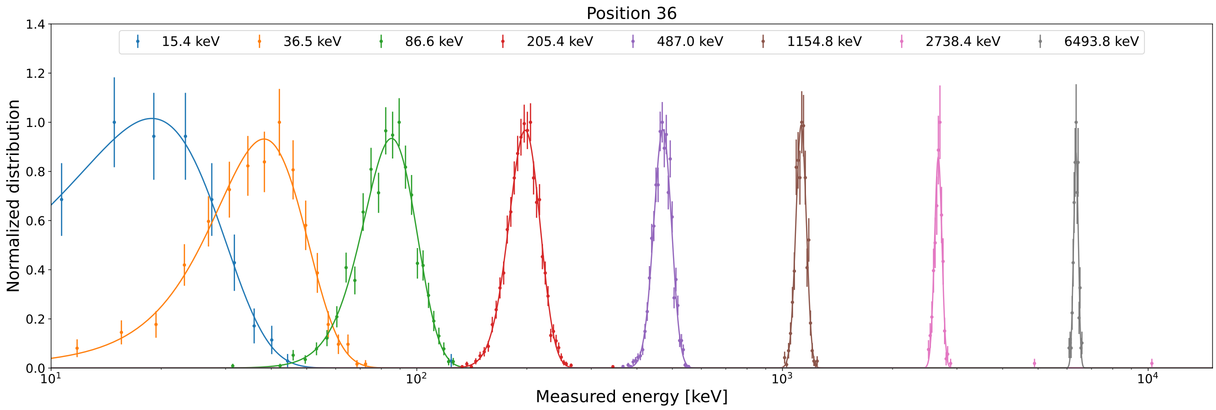

Figure 6.3 – Simulated normalized distributions of the deviations of E_{meas} from E_{true} for each energy bin, for a particular spatial bin taken as example.

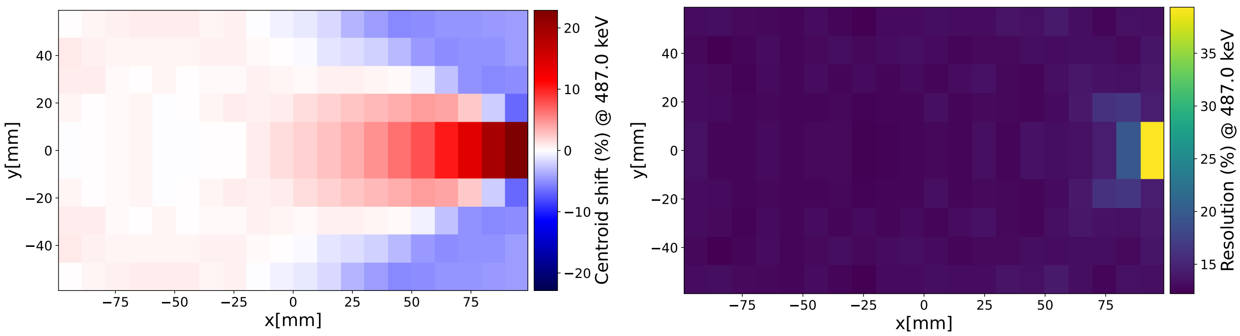

Figure 6.4 – Centroid shift map (left) and energy resolution map (right) at 487 keV.

properties of the BGO performance: the optical light collection efficiency and the energy resolution, respectively. The relationship between FWHM and energy resolution is straightforward: the FWHM quantifies the precision of the true energy deposit estimation. Regarding the centroids, if for a particular spatial and energy bin we have an enhanced light collection efficiency, the mean measured energy E_{meas} would systematically be larger than the true energy E_{true} and the centroid of the distribution would be shifted towards higher energies. This arises from using a single average calibration to convert the number of optical photons into energy, which fails to account for the position-dependent variation in the relationship between detected optical photons and energy. As discussed in Section 2.4, this introduces a bias in the energy determination that must be accounted for in simulations to reflect the real ACS performance. In Figure 6.4, we show the relative centroid shift (with respect to the bin center) and the energy resolution map for the energy bin at 487 keV (the closest one to 511 keV). The SiPMs are centered on the right side of the map. The centroid shift distribution is affected by the SiPMs position: it is higher for the central bins in front of the SiPMs, reaching a maximum enhancement of $\sim$22%, suppressed by $\sim$5% in the corners adjacent to the SiPMs, and nearly uniform in the region distant from the SiPMs. The energy resolution is instead highly uniform, except for the two bins right in front of the SiPMs. In these bins, the detector response exhibits significant variability depending on the exact location of the energy deposition. Due to the close proximity between the interaction point and the SiPMs, even small displacements within a single bin can lead to relatively large variations in the light transmission.

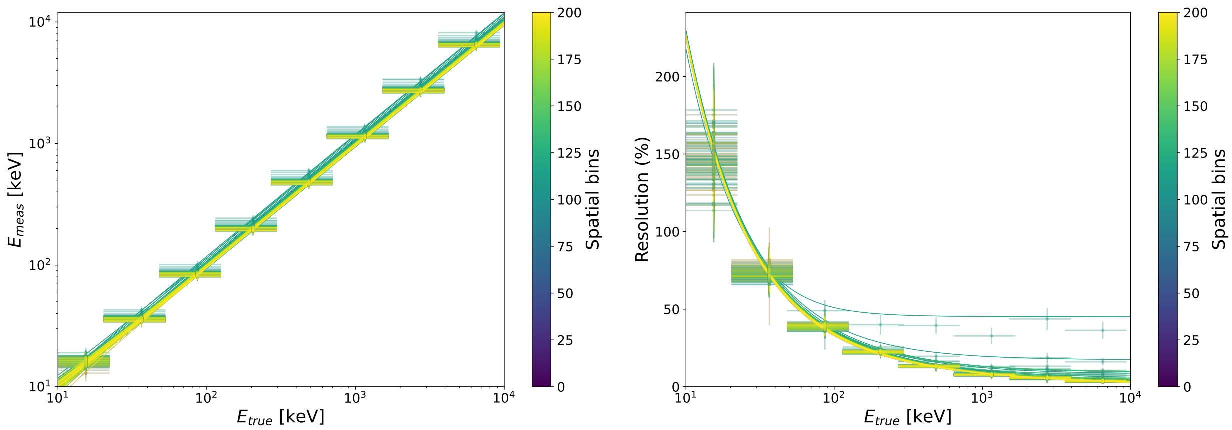

Figure 6.5 – Centroids (left) and energy resolutions (right) as a function of the energy, color coded with the corresponding spatial bin, and best-fit models for each distribution. The outlier curve for the energy resolution, visible in the right panel, corresponds to the spatial bins in front of the SiPMs, which have a large value of c_{xy} due to the strong variability of the detection probability within the bin.

In Figure 6.5 we show the centroids and energy resolutions as a function of the energy bins, color coded with the x-y bin. We fit the two relations with the following models:

$$\bar{E}_{\rm meas} = m_{xy} * E_{\rm true} + q_{xy} \ , \tag{6.1}$$

$$\text{FWHM} = \sqrt{a_{xy}^2 + b_{xy}^2\, E_{\rm true} + c_{xy}^2\, E_{\rm true}^2} \ , \tag{6.2}$$

with the second equation having the same description as Eq. (5.2). The best-fit provides, for each x-y bin, a set of five parameters $\{m_{xy}, q_{xy}, a_{xy}, b_{xy}, c_{xy}\}$. These parameters allow us to take as input the true energy deposit $E_{\rm true}$ and its x-y position, and give as output a new measured energy $E_{\rm meas}$ drawn from a Gaussian distribution with centroid $\bar{E}_{\rm meas}$ and FWHM as given by Eqs. 6.1 and 6.2.

This correction also has an impact on the overall gamma-ray quantum detection efficiency of the BGO. The quantum detection efficiency is defined as the fraction of detected high energy photons (i.e. with $E_{\rm meas} > 80$ keV) out of the total incident photons. We simulate a flat energetic spectrum of photons with an initial energy < 300 keV hitting the BGO and calculate the detection efficiency for each x-y bin. The detection efficiency map is shown in the left plot of Figure 6.6, while on the right plot we show the relative difference between the evaluated efficiency and the same efficiency evaluated without the correction, to underline the impact of the optical physics. We observe a similar distribution obtained for the centroid shift in Figure 6.4: the efficiency is nearly uniform far from the SiPMs, larger when approaching the SiPMs, and suppressed in the corner adjacent to the SiPMs. The effect of implementing the optical physics correction is to enhance the efficiency by at most ∼8% close to the SiPMs, and suppress it by at most ∼2% in the adjacent corners.

6.3 Bottom ACS simulation

We used the benchmarked simulations (Chapter 5) to investigate the performance of the BGO bottom crystals (Z-crystals), in particular to estimate their energy threshold and resolution. The main difference between the bottom BGO crystals and the side panels is the size. The

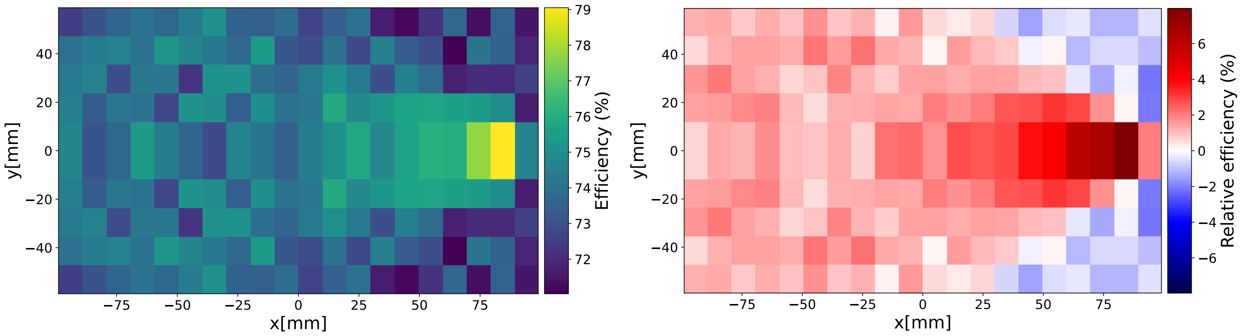

Figure 6.6 – Quantum detection efficiency map (left) and relative quantum detection efficiency map (right).

bottom EM ACS contains 10 BGO crystals of varying dimensions: four measuring 198 × 62.5 × 24 mm (Z-crystal A), four measuring 171 × 64.5 × 24 mm (Z-crystal B), and two measuring 198 × 80.5 × 24 mm (Z-crystal C). All of them are smaller than the lateral ACS crystals.

The crystal size has a potential impact on the overall optical light collection efficiency of the SiPMs and, consequently, on the energy threshold. In fact, in a smaller crystal the active detection area corresponding to the SiPMs occupies a larger fraction of the total BGO surface, enhancing the probability for the optical photons to reach the SiPMs before being absorbed. Hence, if for the same energy deposit the average number of detected optical photons is larger, the energy threshold would be accordingly shifted to a lower energy. To quantify this effect, we simulate 3 parallel monochromatic beams at 60, 88, and 122 keV, hitting orthogonally the largest surface of each type of Z-crystal, and one of the X-wall crystals, that has a dimension of 194 × 118 × 23 mm. Then, for each energy and crystal, we collect the number of detected optical photons in the SiPMs by fitting with a Gaussian the photopeaks. The photopeaks, illustrated in Figure 6.7, show indeed that the smaller crystals have a larger number of detected optical photons, caused by an enhanced optical light collection efficiency of the SiPMs.

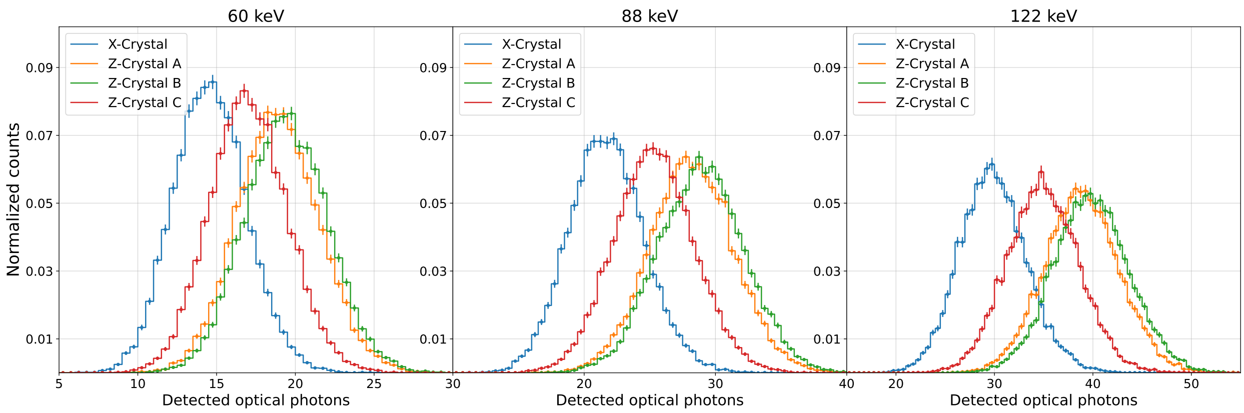

Figure 6.7 – Simulated photopeak distributions detected by the BGO X-crystal and Z-crystals, for 60 keV (left), 88 keV (middle) and 122 keV (right) source photons, respectively.

To estimate the impact on the energy threshold, we perform an energy calibration by taking the peak values from the Gaussian fitting of the photopeaks and associating to each of them

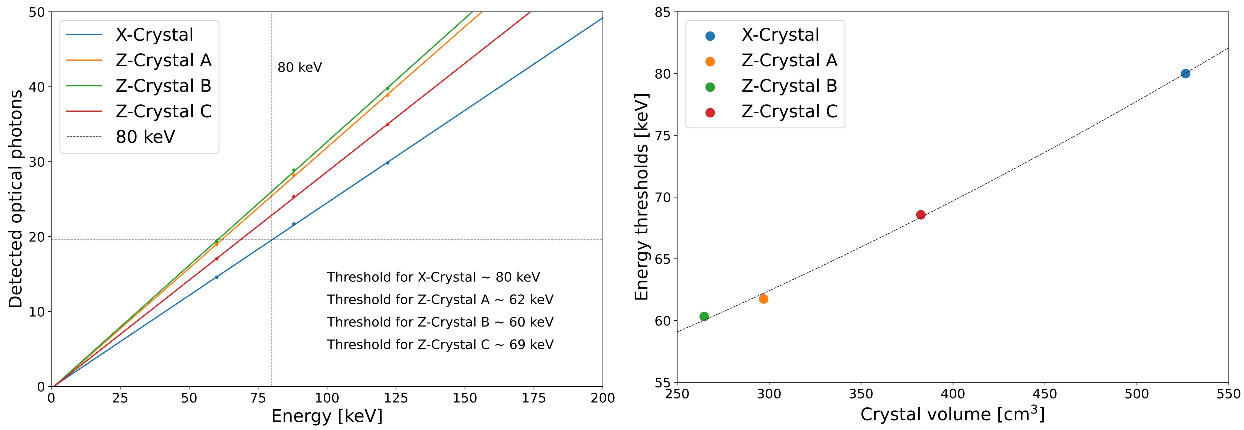

Figure 6.8 – On the left, energy calibrations for the lateral and bottom crystals, with highlighted the optical photons emitted at the 80 keV threshold in the lateral crystal. On the right, energy threshold as a function of the crystal volume.

the corresponding source energy, in a similar fashion as described in Section 5.1.3. From the linear fit of the data points, we obtain the photon-energy relation for each crystal, as shown in the left plot of Figure 6.8. Smaller crystals exhibit steeper slopes, as a result of an enhanced optical light collection efficiency. This means that, for smaller crystals, the same number of detected optical photons corresponds to a lower energy or, equivalently, the same amount of energy deposit produces a larger signal. From the photon-energy relation of the X-crystal, we derive the number of optical photons corresponding to 80 keV, which is the goal for the ACS threshold. We then use the other photon-energy relations to convert back this number into energy for each Z-crystal. We find an energy of ∼62 keV for the Z-crystal A, ∼60 keV for the Z-crystal B and ∼69 keV for the Z-crystal C. These are the energy thresholds expected for the Z-crystals, if we assume an energy threshold of 80 keV for the X-crystal. In terms of relative variation, we have a reduction of ∼22%, ∼25% and ∼14% for the Z-crystal A, Z-crystal B and Z-crystal C, respectively.

In the right plot of Figure 6.8 we show the energy thresholds as a function of the crystal volume. The trend demonstrates that the energy threshold progressively decreases with the crystal volume, as we indeed expected. We fit the relation with the following empirical model:

$$\frac{\Delta E_t}{E_t} = a_0 + a_1 V + a_2 V^2 \ , \tag{6.3}$$

where E_t is the energy threshold, $\Delta E_t/E_t$ is the relative variation with respect to a reference value, V is the crystal volume, and a_i are the model parameters. Using as a reference threshold the requirement of 80 keV, we obtain the following best-fit parameters: $a_0 = -0.43 \pm 0.08$ keV, $a_1 = (6 \pm 3) \times 10^{-3}$ keV·cm^3, $a_2 = (5 \pm 2) \times 10^{-7}$ keV·cm^6. This relation allows for a quick estimate of the energy threshold of a BGO crystal with arbitrary dimensions, assuming the same readout system.

Having a larger number of detected optical photons for smaller crystals has an impact also on the energy resolution. In fact, the FWHM increases with the energy, as described in Eq. (5.1), and is affected by the electronic noise (constant term), statistical fluctuations of the detected optical photons ($\propto \sqrt{E}$) and response inhomogeneity ($\propto E$). We can assume the electronic noise to be the same for all the crystals, as it does not depend on the BGO geometry. However, for the same energy we have more detected optical photons for smaller

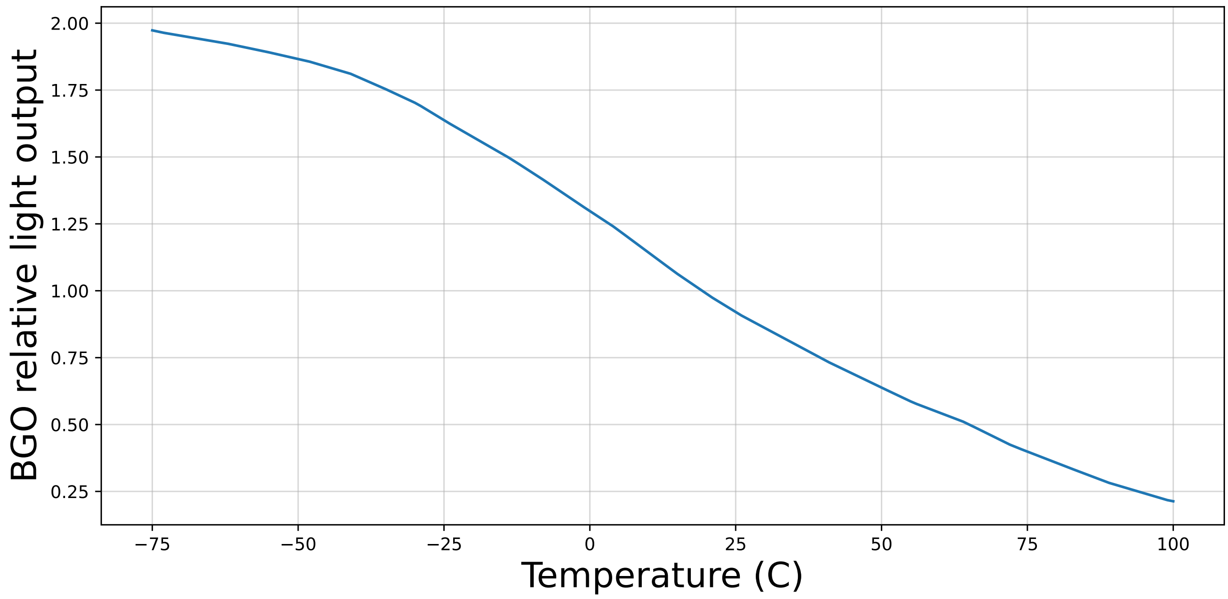

Figure 6.9 – Relative light output of BGO as a function of temperature, normalized to the value at room temperature. Data from Knoll (1979).

crystals, resulting in lower statistical fluctuations. Moreover, for a smaller BGO we expect a smaller contribution in the resolution from the response inhomogeneity with respect to a larger crystal, where the increased size introduces larger variation in the optical photon collection. After converting the photopeaks in energy using the corresponding photon-energy relation for each crystal, we fit the obtained distributions with a Gaussian to compute the FWHM and, consequently, the energy resolution. Among the three investigated energies, we obtain an average improvement in the energy resolution of $\sim$8% for Z-crystal A and Z-crystal B, and of $\sim$5% for Z-crystal C.

6.4 Temperature effects on the ACS performance

The light yield of scintillators commonly depends on temperature. Typically, it decreases with increasing temperature. For BGO, the variation of the light yield with respect to the temperature is shown in Figure 6.9. Different light outputs inevitably impact the BGO response: at higher temperatures, a lower number of optical photons per energy deposit leads to larger statistical fluctuations and, consequently, poorer energy resolution. Furthermore, a decrease in the number of optical photons would increase the energy threshold, resulting in a reduction in detection efficiency. Overall, we expect the COSI ACS to perform worse at higher temperatures. Another effect of temperature on the ACS is that SiPM electronic noise increases with temperature, mainly due to an increase in SiPM dark current.

The calibration measurements, simulation benchmarking, and development of the correction matrix presented in Chapter 5 and in this chapter were all performed at room temperature. We explored the effect of different operating temperatures on the ACS performance, as it might vary during flight. We performed additional simulations of the BGO response, varying its light yield according to Figure 6.9. We simulated a ^{137}Cs source on top of the EM X-wall, replicating the experimental setup described in Section 5.1. We tested two different temperatures: -40°C and 45°C. For each temperature, we computed the energy resolution,

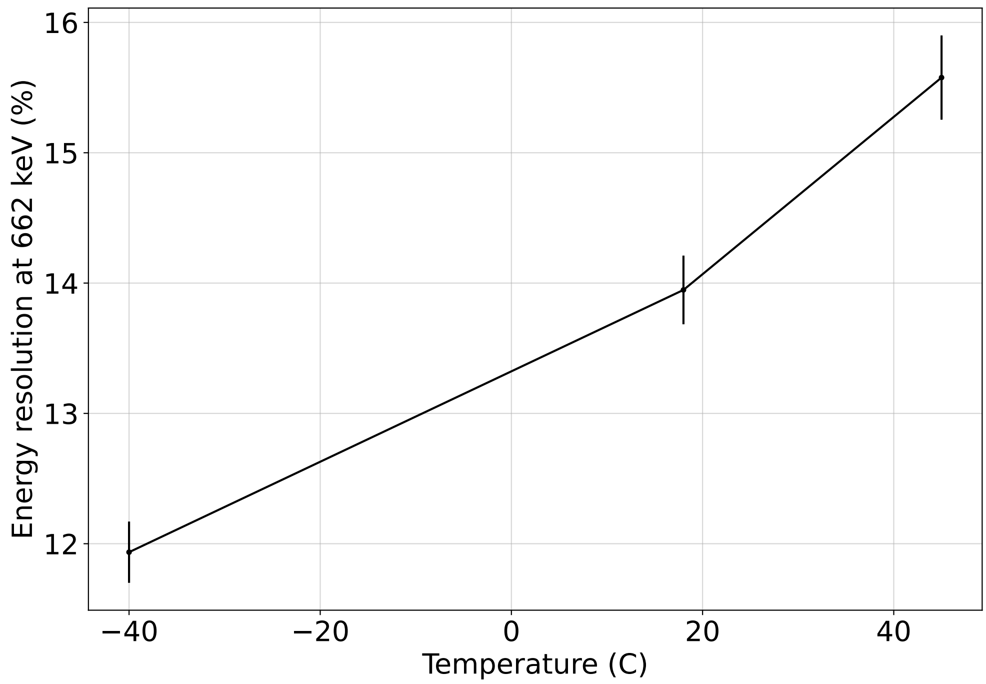

Figure 6.10 – BGO energy resolution at 662 keV for three different temperatures: -40 °C, 18 °C, and 45 °C.

following the same procedure as in Chapter 5.

The result is shown in Figure 6.10, where we also added the energy resolution obtained at room temperature to better display the trend. The energy resolution is solely due to changes in the BGO light yield with temperature and, consequently, changes in the statistical fluctuations in the number of optical photons. As expected, the energy resolution deteriorates as the temperature increases, varying almost linearly from approximately 12% at -40 °C to almost 16% at 45 °C.

A higher temperature also leads to a higher threshold, as the same energy deposit results in fewer optical photons and, consequently, a weaker signal from the SiPMs. We replicated the work carried out in Section 6.3 on the energy threshold, but instead of varying the crystal volume, we changed the temperature (i.e., the BGO light yield). The result is shown in Figure 6.11: the energy-photon relationship is presented for 18 °C and 45 °C. The different slopes of the relationship are due to the different light yields and, therefore, the different numbers of scintillation photons produced per unit of energy. Using the same reasoning as in Section 6.3, we can calculate the expected energy threshold for 45 °C. We obtained an increase of about 44% compared to the threshold at room temperature. Assuming a value of 80 keV at room temperature, the expected threshold at 45 °C is approximately 115 keV.

This work has been carried out to investigate how the BGO response varies with operating temperature. The variation in energy resolution is relatively small if compared to the overall benchmarking uncertainty of 20%. It should be noted, however, that additional factors may contribute to the temperature dependence of the BGO performance, including changes in electronic noise and variations in the SiPM gain. Dedicated calibration measurements across different temperatures are planned to quantify these effects. Only once these measurements are available will it be possible to update the simulation framework and, if necessary, adapt the correction matrix accordingly.

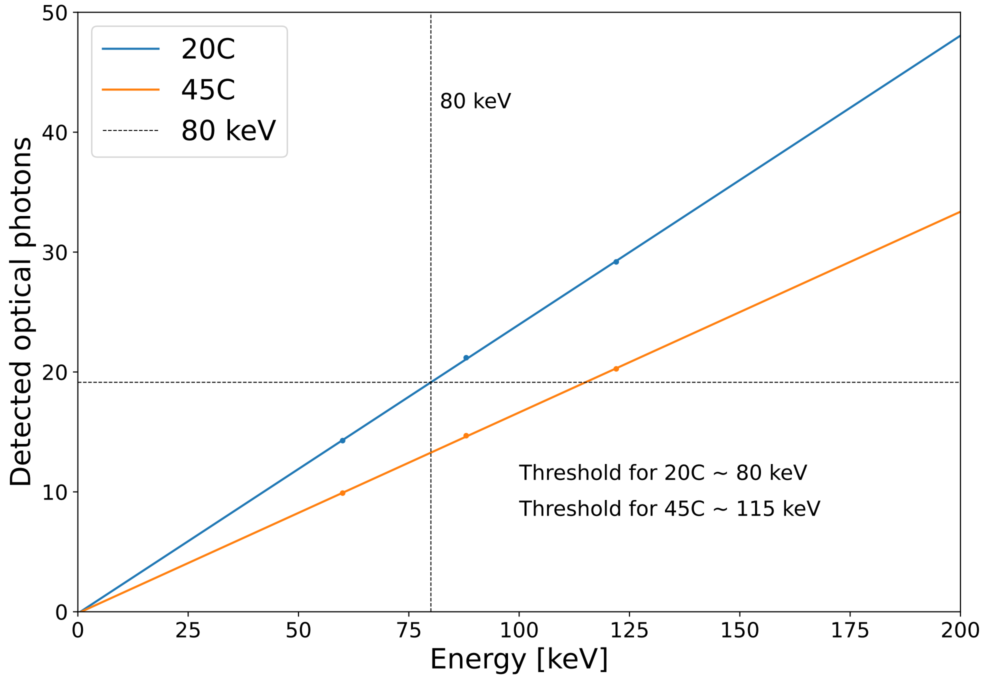

Figure 6.11 – Energy calibrations for the EM X-wall for two temperatures (18 °C and 45 °C.), with highlighted the optical photons emitted at the 80 keV threshold.

6.5 Correction matrix for the DC3 mass model

Performing realistic MEGAlib simulations of GRBs and testing their detection with the COSI ACS requires modeling the entire spacecraft and the full ACS, including every BGO crystal. To account for this, we developed an ad-hoc correction matrix tailored to the full ACS geometry, which will be applied through a dedicated Python script to evaluate its impact on the GRB simulations (see Chapter 7). The most recent spacecraft mass model available is the one used for the Data Challenge 3 (DC3, Spectrometer and Collaboration 2025). The size of the BGO crystals in the DC3 mass model differs slightly from that in the EM version and corresponds to the dimensions given in Table 1.2. Moreover, each BGO detector is divided into spatial bins (voxels), with each of them acting as an independent detector that records its own energy. The adopted spatial binning is the following: $19 \times 12 \times 2$ for the BGO X-crystal, $19 \times 11 \times 2$ for the BGO Y-crystal, $20 \times 8 \times 3$ for the BGO Z-crystal A, $17 \times 6 \times 3$ for the BGO Z-crystal B, and $20 \times 6 \times 3$ for the BGO Z-crystal C.

We replicated the analysis shown in Section 6.2 for the DC3 design. We ran simulations for each BGO panel as in Section 6.2, using a node of the local OAS cluster. Due to the limitation in computational resources, the energies of the simulated photons range from 10 keV to 1 MeV. The spatial binning matches the voxelization of the BGO detectors in the DC3 mass model, while the energy binning is limited to 4 logarithmic bins. In Figure 6.12 and Figure 6.13, we present the centroid shift and resolution maps for each BGO crystal, for a particular energy bin (562.3 keV). Although all crystals display similar spatial patterns in both centroid and resolution, the lateral panels exhibit a broader spread of values compared to the bottom panels. This effect is most evident in the centroid shift, where the lateral BGO crystals reach a maximum shift of about $\sim 25\%$, while the bottom crystals do not exceed $\sim 15\%$. These results indicate that smaller crystals provide a more homogeneous response, whereas wider crystals are more affected by spatial variations.

We then fitted the centroids and energy resolutions as a function of energy using Equations (6.1) and (6.2), obtaining for each BGO crystal the corresponding fit parameters m_{xy}, q_{xy}, a_{xy}, b_{xy}, c_{xy}. Table 6.1 reports the energy resolutions of the BGO crystals at selected energies, chosen to match the emission lines employed for the benchmarking (Chapter 5). The listed values correspond to averages over all spatial bins, with the uncertainties reflecting the standard deviation of the respective distributions. As anticipated from the analysis in Section 6.3, the bottom crystals exhibit systematically better (lower) energy resolutions compared to the lateral crystals. This behavior can be attributed to the higher optical light collection efficiency, which reduces statistical fluctuations, and to the higher homogeneity in the detector response due to their increased compactness. Moreover, the bottom crystals also exhibit a systematically lower dispersion of the resolutions across spatial bins compared to the lateral ones, reflecting the differences in compactness.

We finally encoded the correction matrix in a single ASCII file (see Figure 6.14), which contains the following information: BGO crystal name, voxel IDs, x-y position of the voxel, and the best-fit parameters $\{m_{xy}, q_{xy}, a_{xy}, b_{xy}, c_{xy}\}$. Although the BGO crystals are also spatially binned along the z direction in the MEGAlib mass model, the correction was computed only in the x and y directions. This choice is motivated by the fact that bins along the z axis are equally distant from the SiPMs and, therefore, we do not expect significant variation in the response. Consequently, the best-fit parameters were replicated uniformly along the z axis.

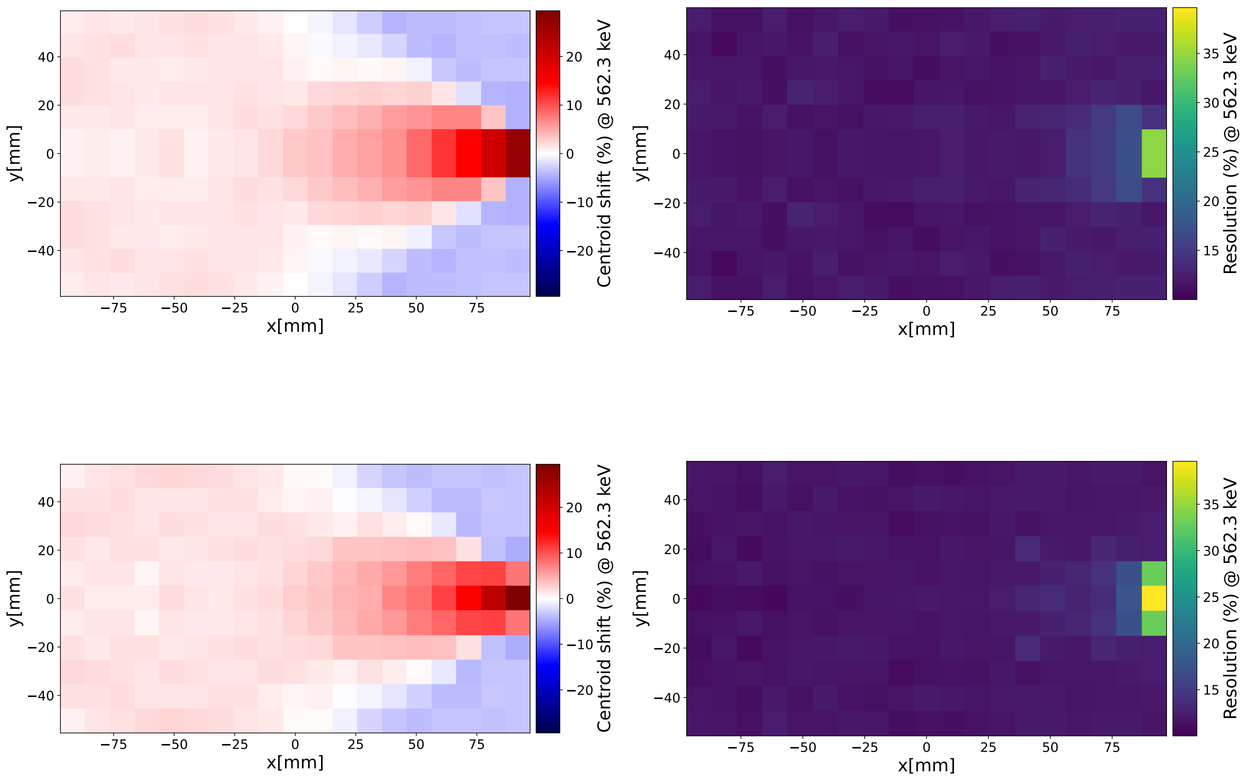

Figure 6.12 – Centroid shift map (left panels) and energy resolution map (right panels) at 562.3 keV for the BGO X-crystal (top) and BGO Y-crystal (bottom). The SiPMs are on the right corner.

Table 6.1 – Energy resolution (R) at different energies for each BGO crystal, as derived from the correction matrix. The reported values are obtained by averaging the energy resolutions over all spatial bins, and the quoted errors correspond to the standard deviation of the respective distributions.

Energy (keV)	$\mathbf{R_X}$ **(%)**	$\mathbf{R_Y}$ **(%)**	$\mathbf{R_{ZA}}$ **(%)**	$\mathbf{R_{ZB}}$ **(%)**	$\mathbf{R_{ZC}}$ **(%)**
60	51.99 ± 4.13	52.09 ± 3.64	49.82 ± 2.65	48.06 ± 1.85	50.29 ± 2.78
88	39.13 ± 3.22	39.06 ± 2.91	37.44 ± 2.17	36.04 ± 1.48	37.72 ± 2.48
122	31.19 ± 2.99	31.05 ± 2.85	29.79 ± 2.16	28.63 ± 1.50	29.96 ± 2.49
356	16.28 ± 3.17	16.26 ± 3.17	15.39 ± 2.63	14.75 ± 2.12	15.38 ± 2.76
511	13.45 ± 3.29	13.53 ± 3.30	12.63 ± 2.78	12.11 ± 2.33	12.59 ± 2.84
662	11.83 ± 3.38	11.98 ± 3.41	11.04 ± 2.87	10.58 ± 2.47	10.99 ± 2.90
898	10.24 ± 3.50	10.47 ± 3.56	9.48 ± 2.98	9.09 ± 2.62	9.41 ± 2.97
1173	9.08 ± 3.61	9.39 ± 3.70	8.33 ± 3.06	8.00 ± 2.73	8.26 ± 3.03
1275	8.75 ± 3.65	9.08 ± 3.74	8.01 ± 3.09	7.69 ± 2.77	7.94 ± 3.05
1333	8.59 ± 3.67	8.93 ± 3.77	7.85 ± 3.11	7.53 ± 2.79	7.77 ± 3.06
1836	7.51 ± 3.80	7.94 ± 3.94	6.79 ± 3.20	6.52 ± 2.91	6.70 ± 3.13

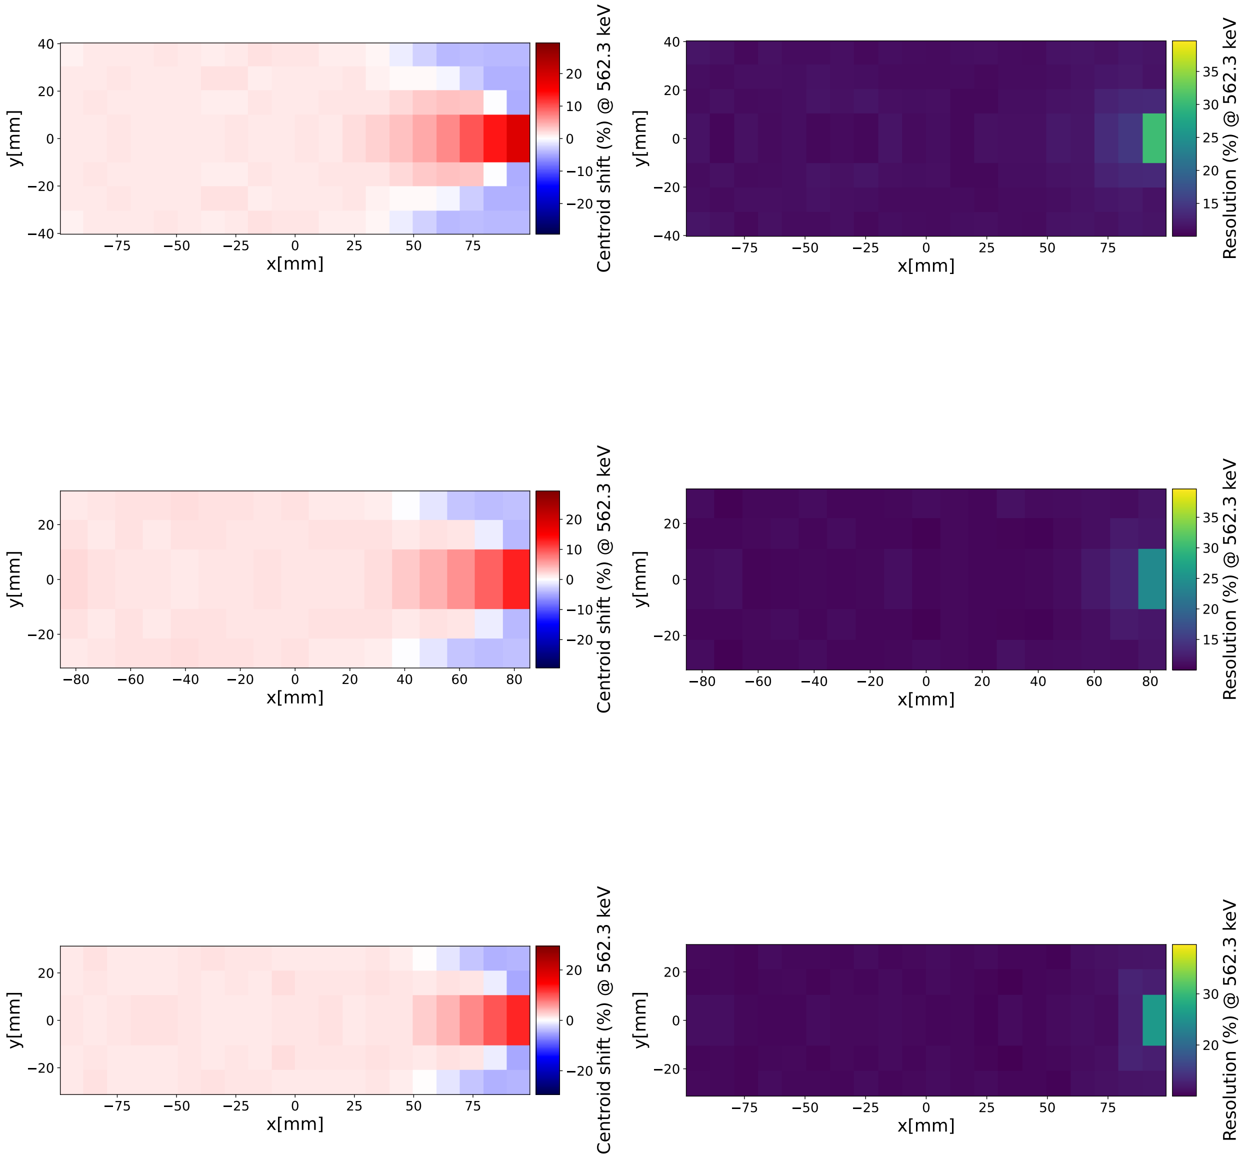

Figure 6.13 – Centroid shift map (left panels) and energy resolution map (right panels) at 562.3 keV for the BGO Z-crystal A (top), BGO Z-crystal B (middle), and BGO Z-crystal C (bottom). The SiPMs are on the right corner.

```
# DetectorName voxel_X voxel_Y voxel_Z x[mm] y[mm] m q a b c
BGO_X0_0          0          0         18     -91.89     -54.08       1.00       2.17       27.47       2.19       0.07
BGO_X0_0          0          0         17     -81.68     -54.08       1.00       2.43       16.38       2.72       0.00
BGO_X0_0          0          0         16     -71.47     -54.08       0.99       3.87       19.87       2.78       0.00
BGO_X0_0          0          0         15     -61.26     -54.08       1.06       5.15       37.47       0.00       0.11
BGO_X0_0          0          0         14     -51.05     -54.08       1.00       3.13       23.31       2.61       0.00
BGO_X0_0          0          0         13     -40.84     -54.08       0.99       4.25       21.20       2.78       0.00
BGO_X0_0          0          0         12     -30.63     -54.08       1.00       1.58       24.37       2.13       0.08
BGO_X0_0          0          0         11     -20.42     -54.08       1.00       2.07       22.98       2.86       0.00
BGO_X0_0          0          0         10     -10.21     -54.08       0.97       4.57       18.89       2.88       0.00
BGO_X0_0          0          0          9       0.00     -54.08       1.00       1.55       27.46       2.48       0.00
```

Figure 6.14 – Portion of the ACS correction file.

Chapter 7

GRB detection with the COSI ACS

The COSI ACS will operate as a stand-alone detector for gamma-ray transients, with GRBs as its primary targets. The dedicated analysis software for COSI is currently under development by the team. As part of this PhD work, in collaboration with the COSI GRB-focused science team, I investigated the GRB detection capability of the COSI ACS using an existing tool, bc-tools (Martinez-Castellanos 2023), originally developed for the BurstCube mission (Smith 2019). The bc-tools library has also been employed to investigate the COSI ACS localization capability (Parmiggiani et al. 2025). Furthermore, the study of GRB detection with the ACS provided an opportunity to explore the impact of the correction matrix (discussed in Chapter 6) on the ACS detection capability.

This chapter is organized as follows. Section 7.1 provides an overview of the operation of the COSI ACS as a gamma-ray transient monitor and describes the fast transient pipeline. Section 7.2 introduces the bc-tools library and the algorithm employed for burst detection. Section 7.3 presents the results of short GRB detections with the COSI ACS, based on simulations using the fourth Fermi-GBM catalog. Finally, Section 7.4 discusses the impact of the correction matrix on the GRB data as detected by the COSI ACS.

7.1 The COSI ACS as transient detector

COSI is expected to detect and localize more than 10 short GRBs in its first two years of operation, with a requirement on the localization accuracy of $< 2.5°$, time accuracy of < 100 ms, and reporting time via GCN of < 1 hr (Tomsick et al. 2019; Spectrometer and Collaboration 2025). The COSI ACS will contribute to the detection and localization of such transient events. Its full view of the sky makes it highly effective for detecting GRBs from all directions, especially outside the GeDs FoV. As discussed in Section 1.4, the COSI ACS is divided into six independent segments: four corresponding to the lateral walls and two to the bottom side. Each segment is read out by a dedicated ASIC, which can handle up to four BGO crystals. Consequently, two of the ten bottom crystals are grouped together with the lateral walls, as illustrated in Figure 7.1. Each ASIC provides two energy channels, below and above 2 MeV, resulting in a total of twelve ACS channels. The ACS energy threshold requirement is 80 keV, though the actual value is expected to vary across different BGO crystals (as discussed in Section 6.3) and will ultimately be determined with dedicated laboratory measurements.

The GeDs and ACS data are downlinked to the ground via the Malindi ground station at each orbital pass. Additionally, the onboard trigger system monitors the ACS ratemeters in each channel: if the count rate exceeds a statistical significance threshold in two or more channels, the ACS data are downlinked via TDRSS. The ACS data consist of light curves with 50 ms time bins. The burst search and analysis will be carried out through two complementary pipelines: the fast transient and the comprehensive pipelines. The fast transient pipeline will operate on ACS light curves downlinked through the TDRSS to provide rapid identification of burst candidates and to generate GCN notices within the required latency of 60 minutes, enabling prompt follow-up by other facilities. The comprehensive pipeline, on the other hand, will process the full data products downlinked to the ground stations, including GeDs data, ACS light curves, and occasionally ACS event-level data. This more detailed analysis will refine the results produced by the fast transient pipeline, updating the preliminary GCN circulars with improved localization, duration, and significance within a few hours. Moreover, the comprehensive pipeline will extend the search to sub-threshold signals that did not trigger the onboard burst detection but may still be visible in the GeDs.

The development of a fast transient pipeline for the detection and localization of GRBs is fundamental in providing results that can be promptly shared with other instruments, enabling rapid follow-ups and contributing to the search for electromagnetic counterparts of gravitational waves. The AGILE mission, whose operation ended in 2024, made use of a real-time analysis pipeline (Bulgarelli 2019; Parmiggiani et al. 2023) to analyze the data acquired by the onboard instruments and received at each contact with the ground station (every about 90 minutes). The alerts were shared with other facilities to enable a fast repointing and follow-up. Fermi and Swift (Gehrels et al. 2004) utilize a real-time analysis system on-board to detect and localize GRBs and allow the spacecraft to repoint (Connaughton et al. 2015; Cummings and Swift BAT Team 2002). The results are then sent to the ground via GCN notices. For COSI, the fast transient pipeline is currently under development by the COSI team. In this work, we focus on the case of data downlinked via TDRSS and analyzed by the fast transient pipeline, emulating realistic ACS observations and performing the analysis using the bc-tools software suite.

7.2 GRB time analysis

The bc-tools library is an open-source Python package developed to detect and localize GRB light curves and analyze their energy spectra. It has been employed for the BurstCube mission, which used CsI scintillators coupled to SiPMs to detect and localize GRBs. The library builds upon the functionality of the GBM Data Tools[1], originally developed for Fermi-GBM (see Section 1.4.1), and extends them to provide a detector-agnostic framework. The bc-tools library is therefore suitable for the analysis of other count-based GRB detectors. While the cosipy library is dedicated to the high level analysis of the COSI Compton data, the software used for time domain analysis, detection, localization, and classification of GRBs from the ACS is still under definition. The exploration of the bc-tools library in this work could open the possibility of integrating such algorithms into the COSI official analysis suite.

The bc-tools library utilizes the Bayesian blocks algorithm for the analysis of GRB light curves. The Bayesian blocks algorithm (Scargle 1998) represents a method for the analysis

[1] https://fermi.gsfc.nasa.gov/ssc/data/analysis/gbm/

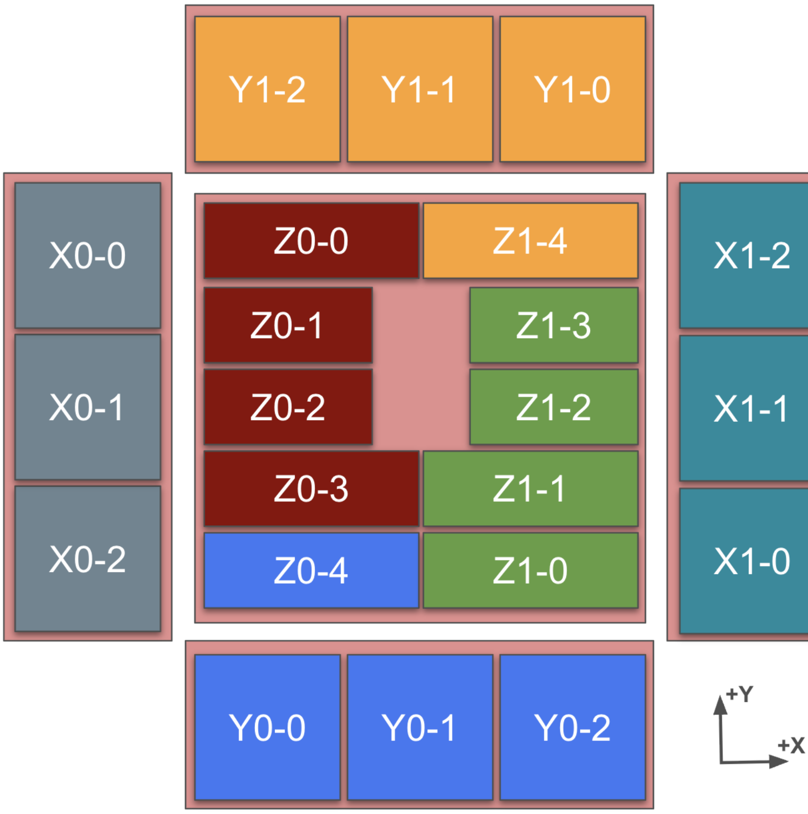

Figure 7.1 – Schematic illustration of the BGO crystals composing the COSI ACS. Each group of color corresponds to an independent ASIC readout system.

of time series data, in order to detect and characterize structures localized in time. Invented by Jeffrey D. Scargle, the idea of the algorithm is to segment the data, which consists of a histogram of counts over time, into variable-sized blocks. The algorithm operates in a Bayesian framework and utilizes posterior probabilities:

$$P(M|D) \propto P(D|M)P(M) \ , \tag{7.1}$$

where $P(D|M)$ is the likelihood for a model M given a certain set of data D, and $P(M)$ represents the prior probability for the same model. In this case, the model corresponds to a particular partition of the data into blocks, and the algorithm seeks the partition that maximizes a fitness function derived from the posterior probability. The Bayesian blocks algorithm is particularly effective for identifying change points in time series such as light curves, highlighting statistically significant variations in the data. Time intervals with no appreciable variability are typically represented as a single, extended bin, while regions characterized by rapid variations are represented with a finer binning, thus adapting the resolution of the representation to the signal property.

The algorithm employed by bc-tools proceeds as follows:

- The background is fitted over a selected region (by default, the entire light curve) using a polynomial function. The polynomial order increases at each iteration (the first iteration corresponds to simply taking the mean).
- An effective light curve is created, i.e. a light curve with the same binning and counts, but where each bin is weighted by the corresponding background counts from the model

obtained in the previous step. This procedure removes any possible time dependence of the background.

- The Bayesian blocks algorithm is applied to the effective light curve, identifying peaks and defining the signal range (from the leftmost to the rightmost changing point).
- The procedure is repeated; this time, the background is fitted only in the region outside the signal range.

The algorithm stops once the background fit no longer changes. The algorithm provides as output the light curve bins in the resulting Bayesian blocks and the signal range. The signal range is defined as the time interval spanning from the first to the last detected change point in the binned light curve.

7.3 Short GRB population analysis with the COSI ACS

We used bc-tools to analyze GRBs as they would be observed by the COSI ACS. Specifically, we constructed mock GRBs and emulated the downlinked ACS data through the TDRSS. For this purpose, we simulated GRBs with MEGAlib using the DC3 mass model.

7.3.1 Fourth Fermi-GBM catalog

To ensure a representative and realistic GRB sample, we considered the fourth Fermi-GBM catalog (von Kienlin et al. 2020b) of GRBs. This catalog covers 10 years of operation, from 2008 July 12 to 2018 July 11. During this period, GBM identified 2356 GRBs. This catalog provides detailed spectral and temporal analysis for each burst, including the duration, the flux, the fluence, and the spectral parameters obtained by fitting different models.

In our analysis, we considered only short GRBs, selecting from the catalog all bursts with $t_{90} < 2$ s for a total of 541 GRBs. The obtained catalog contains three spectral models: Power Law, Band, and Comptonized. They take the following forms:

$$\frac{dN}{dE} = N_0 \left(\frac{E}{E_0}\right)^{\gamma} , \tag{7.2}$$

$$\frac{dN}{dE} = N_0 \times \begin{cases} \left(\frac{E}{100\ \mathrm{keV}}\right)^{\alpha} \exp\left[-\frac{(\alpha+2)E}{E_{peak}}\right], & E \leq \frac{\alpha-\beta}{\alpha+2} E_{peak} \\ \left(\frac{E}{100\ \mathrm{keV}}\right)^{\beta} \exp(\beta-\alpha) \left[\left(\frac{\alpha-\beta}{\alpha+2}\right)\left(\frac{E_{peak}}{100\ \mathrm{keV}}\right)\right]^{\alpha-\beta}, & E > \frac{\alpha-\beta}{\alpha+2} E_{peak} \end{cases} \tag{7.3}$$

$$\frac{dN}{dE} = N_0 \left(\frac{E}{100\ \mathrm{keV}}\right)^{\alpha} \exp\left[-\frac{(\alpha+2)E}{E_{peak}}\right] , \tag{7.4}$$

respectively.

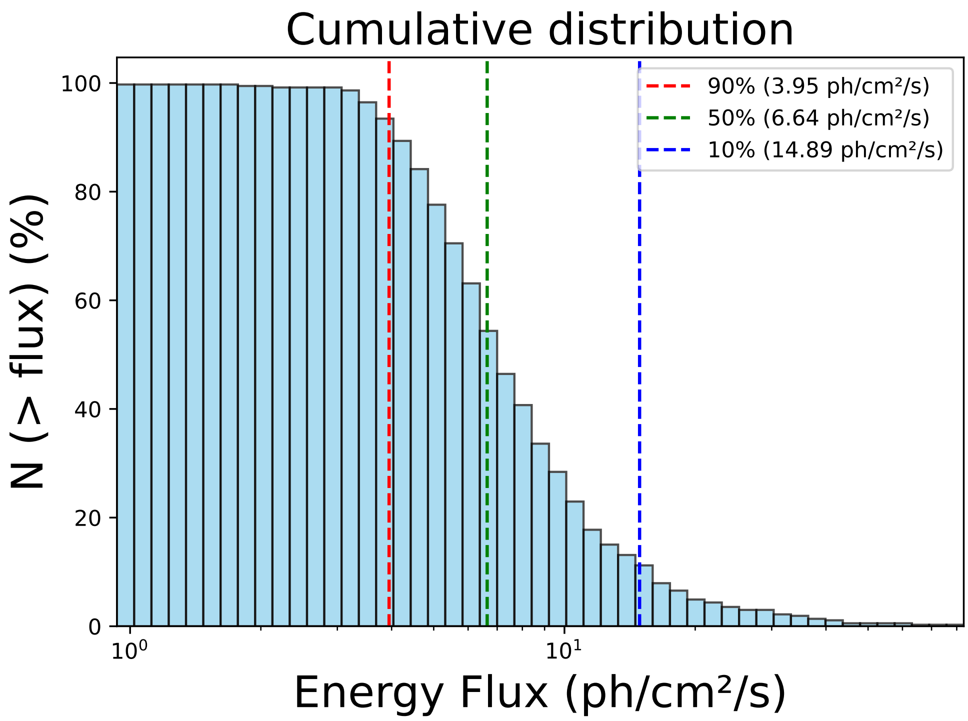

Figure 7.2 – Cumulative distribution of the GRB fluxes from the Fermi-GBM catalog ($t_{90} < 2$ s).

For each burst, we extracted the reported fluxes and spectral model. Since the fluxes were reported for the 10 keV - 1 MeV energy range, we rescaled them to the COSI ACS energy range (10 keV - 10 MeV). The rescaled COSI flux (F_{COSI}) was obtained using the spectral model $f(E)$ of the GRB and integrating it over the COSI energy range:

$$F_{\mathrm{COSI}} = F_{\mathrm{GBM}} \times \frac{\int_{10\ \mathrm{keV}}^{10\ \mathrm{MeV}} f(E)\, dE}{\int_{10\ \mathrm{keV}}^{1\ \mathrm{MeV}} f(E)\, dE} \ , \tag{7.5}$$

with F_{GBM} being the GBM flux. The rescaled cumulative flux distribution (in ph/cm^2/s) of this GRB sample is shown in Figure 7.2.

7.3.2 GRB simulations with MEGAlib

The simulations were conducted with MEGAlib. The GRBs were simulated from random directions[2]. For each GRB, we assigned a random light curve from a set of templates extracted from GBM data (work done by G. Panebianco), which defines the temporal profile of the injected photons. The flux sets the correct normalization of the temporal profile so that the average flux over the burst duration corresponds to the specified input value. The energy spectrum is assumed to be constant during the entire duration of the burst. MEGAlib applies an energy threshold of 80 keV and a constant Gaussian noise with $\sigma = 10$ keV to the energies recorded by the BGO crystals.

[2]In this work, we are neglecting the Earth occultation.

7.3.3 Time analysis of the COSI ACS ratemeters

Once the GRBs are simulated, we used the Python package PyROOT[3] to extract the hits in the ACS and aggregate them in the 6 ASIC segments, according to Figure 7.1. We restrict our analysis to the low-energy channel ($E < 2$ MeV), as it will be the one mostly adopted for transient detection. To construct a realistic light curve, we performed the following calculations. The expected data rate for the ACS low-energy channel is approximately 6 kbps, assuming 20 bins per second with a bin width of 50 ms for the six segments. Given an alert size limit of roughly 500 kbits, the total ACS exposure that can be downlinked per alert is estimated to be about 80 seconds. Since the GeD data will likely be included in the same data package, the total duration of the ACS data will probably be shorter. To be conservative, we approximated the light curve duration as 60 seconds. For the background, we considered the 24 hours simulation conducted by the COSI team. The resulting background as observed by the ACS is shown in Figure 7.3, for each of the 6 segments. The gaps represent the passages in the South Atlantic Anomaly (SAA), an area where Earth's inner Van Allen radiation belt comes closer to the Earth's surface. This leads to an increased flux of energetic particles in this region, exposing the satellites to higher levels of radiation. For this reason, data are usually not taken when passing in the SAA. For each GRB, we randomly extracted 60 seconds of the ACS background, excluding the first few orbits where the background had not yet stabilized[4]. For each ACS segment, we combined GRB and background events to form a single light curve and binned the summed data with 50 ms time bins. We assumed 10 seconds of background before the start of the GRB. In the following, we show the results using a representative case, GRB 120410585, which lies close to the median of the cumulative flux distribution (6.64 photons/cm^2/s). Its light curve, combined with the background, as detected by the six ACS segments is shown in Figure 7.4.

After constructing the light curve for each GRB, we used bc-tools to apply the Bayesian blocks algorithm and identify the signal. We summed the light curves from all panels to obtain a single light curve. Figure 7.5 shows the light curve binned using the Bayesian blocks algorithm for the example GRB. The gray data points represent original binning of the summed light curve (0.05 s), the dotted red line is the fitted background, and the blue line represents the Bayesian blocks. Two green vertical lines define the signal range, i.e. the time interval between the first and the last changing point in the Bayesian blocks. In total, the algorithm successfully recognized about 97% of the total 541 GRBs. For the remaining 3%, the algorithm failed to identify a significant variation in the data due to the limited flux of the GRB, thus recognizing the light curve as background only.

As a final step, we evaluated the significance of each GRB. The significance was calculated following the method of Li & Ma (Li and Ma 1983), which provides a statistically rigorous estimate of the detection significance for counting experiments with background. It compares the number of observed counts in the signal region to those in a background region, accounting for their relative exposure times. According to this method, the significance (S) can be obtained using the following formula:

$$S = \sqrt{2}\left\{N_{on}\ln\left[\frac{1+\alpha}{\alpha}\left(\frac{N_{on}}{N_{on}+N_{off}}\right)\right] + N_{off}\ln\left[(1+\alpha)\left(\frac{N_{off}}{N_{on}+N_{off}}\right)\right]\right\}^{1/2} , \quad (7.6)$$

[3]https://root.cern/manual/python/

[4]A fraction of the background arises from activation of spacecraft materials, producing radioactive isotopes that subsequently decay over time.

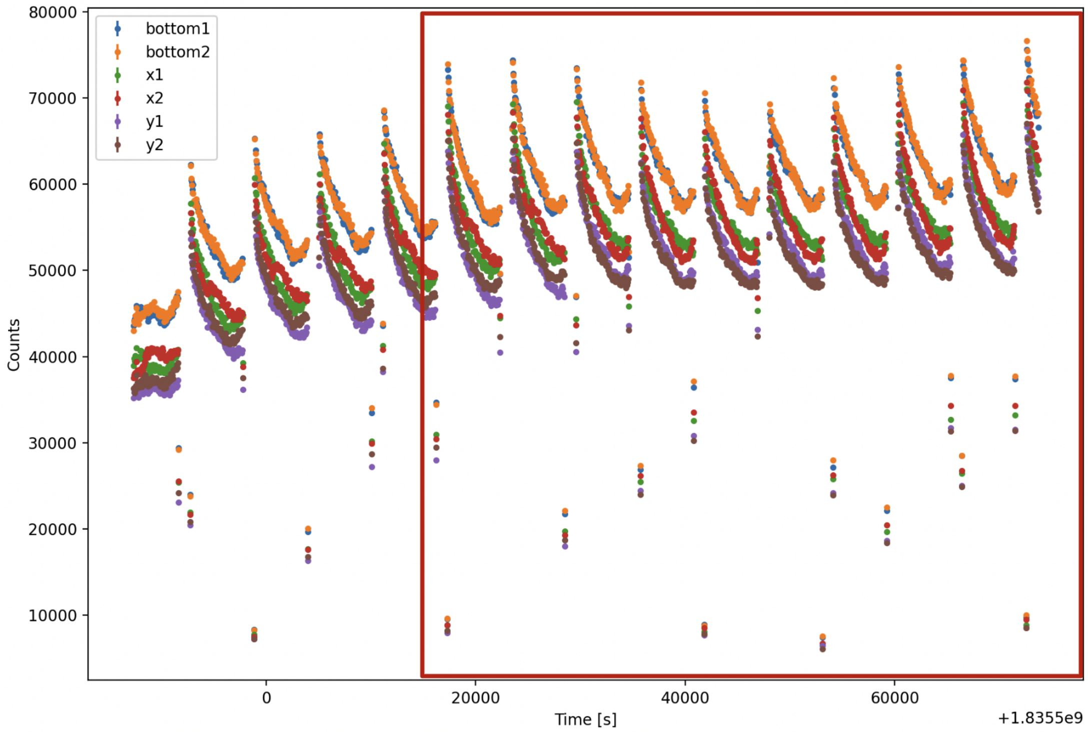

Figure 7.3 – Simulated 24-hour light curve of the background as measured by the 6 panels of the COSI ACS. The red box highlights the time interval from which a 1-minute background segment was extracted. The simulation was conducted by Savitri Gallego of the COSI team.

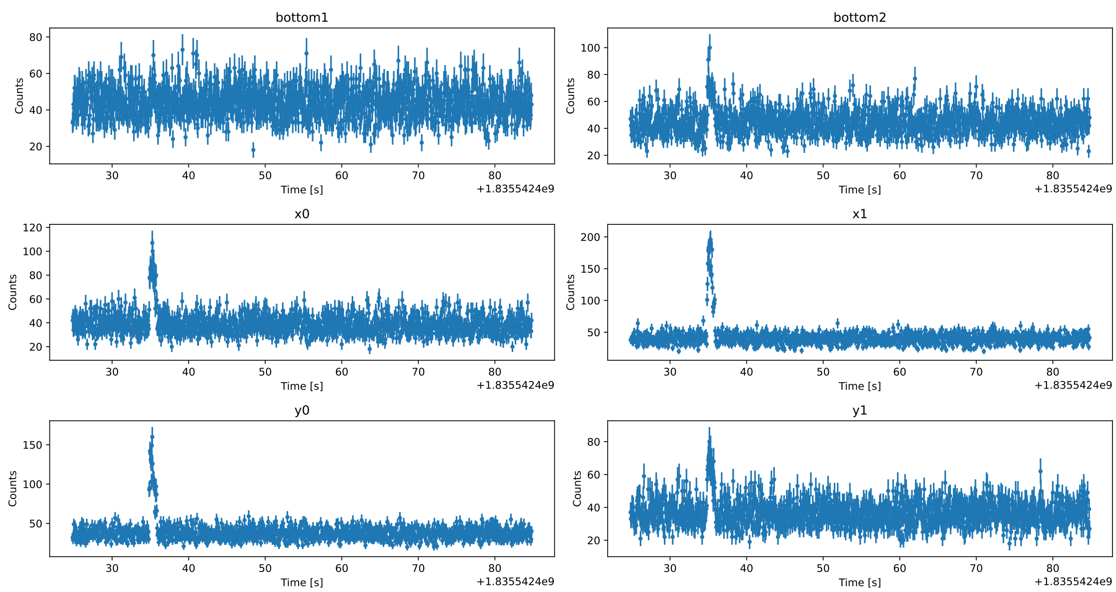

Figure 7.4 – Simulated light curve of GRB 120410585 detected by the 6 segments of the COSI ACS.

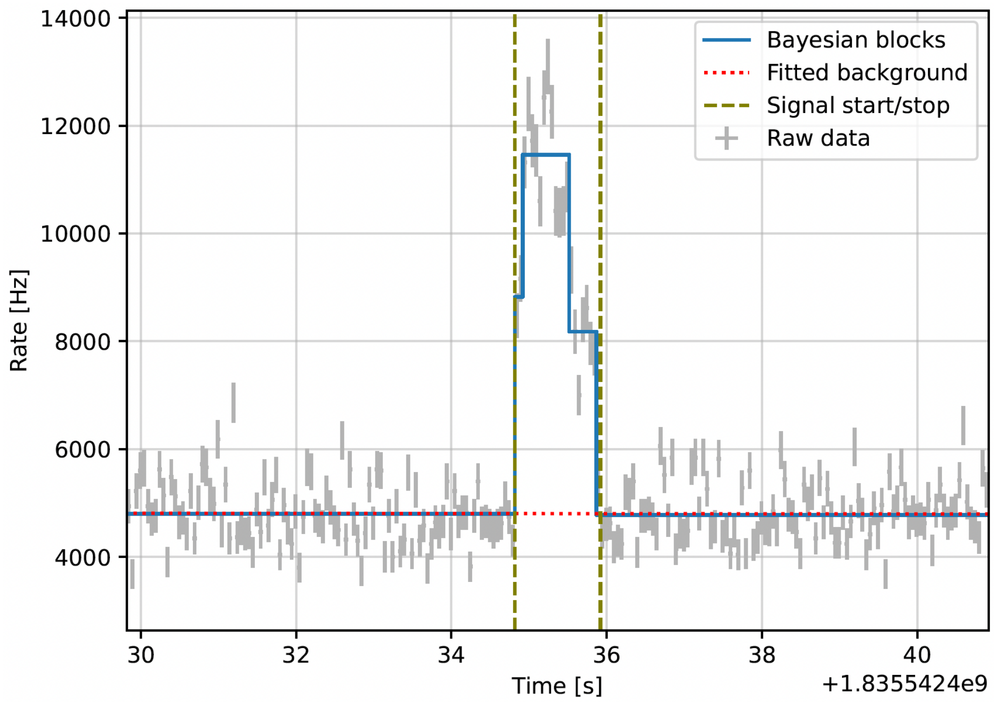

Figure 7.5 – Simulated COSI ACS light curve of GRB 120410585, with the Bayesian blocks representation, background fit, and signal range indicated.

where N_{on} is the total number of counts detected within the signal range (with duration t_{on}), N_{off} is the total number of counts within a time range with no signal (with duration t_{off}), and $\alpha = t_{on}/t_{off}$. As the time range with no signal, we selected the portion before the start of the signal (hence, $t_{off} = 10$ s). As the signal range, we considered the one obtained with the Bayesian blocks algorithm. Figure 7.6 shows the significance as a function of the flux, for each GRB. Among all simulated GRBs, about 95% achieve a significance above 5σ. The large fraction of detected GRBs indicates a strong detection capability of the COSI ACS. The 86% of GRBs have significance between 10σ and 100σ, with an average value of 48σ. Using the Pearson coefficient to quantify the correlation between significance and flux, we obtained a moderate correlation of about 0.58. The deviation from an ideal correlation between flux and detection significance arises because each GRB is associated with a randomly selected background interval, leading to inhomogeneous background rates across the sample. Figure 7.7 shows the all-sky map of the simulated GRBs, color-coded with their significance. No dependence of the significance on the GRB direction is observed.

7.4 Impact of the correction matrix

In Section 7.3 we investigated the capability of the COSI ACS to detect representative short GRBs, using the bc-tools library and a constant Gaussian noise of $\sigma = 10$ keV. However, the impact of the detector effects on the ACS simulation is an energy correction that is energy and position dependent, encoded in a correction matrix (Chapter 6). By applying this correction to the MEGAlib simulations of the ACS, we can explore the impact of the ACS DEE on the GRB detection and localization capability.

We developed a dedicated Python script that applies the correction matrix to the simulated data. Specifically, it reads from the simulation output (the `.sim` file from Cosima) the true

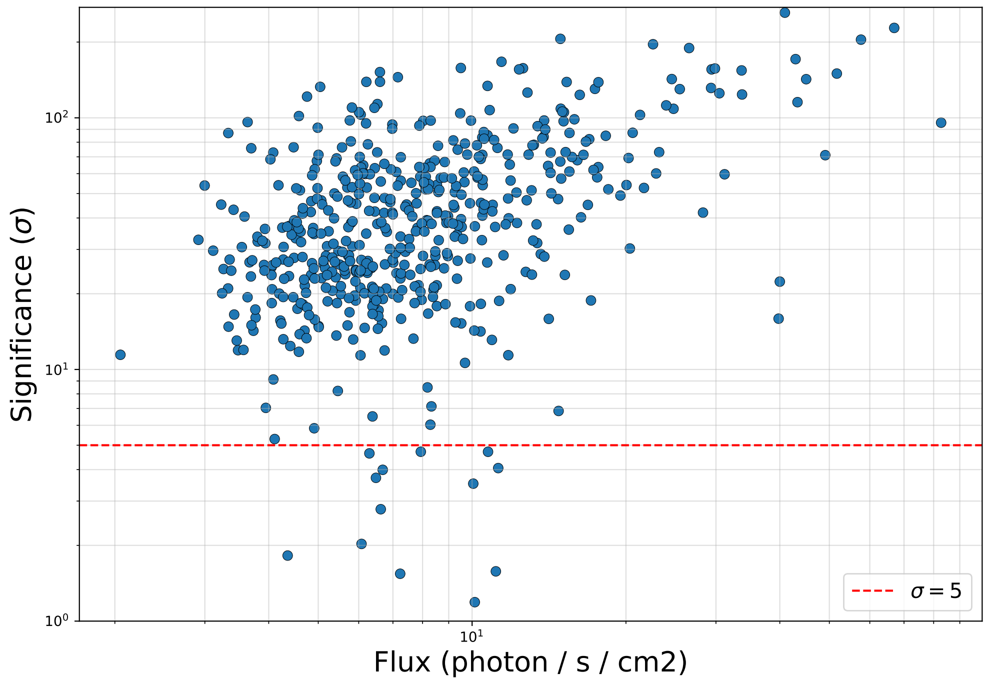

Figure 7.6 – Detection significance of the simulated GRBs from the Fermi-GBM catalog, shown versus their flux.

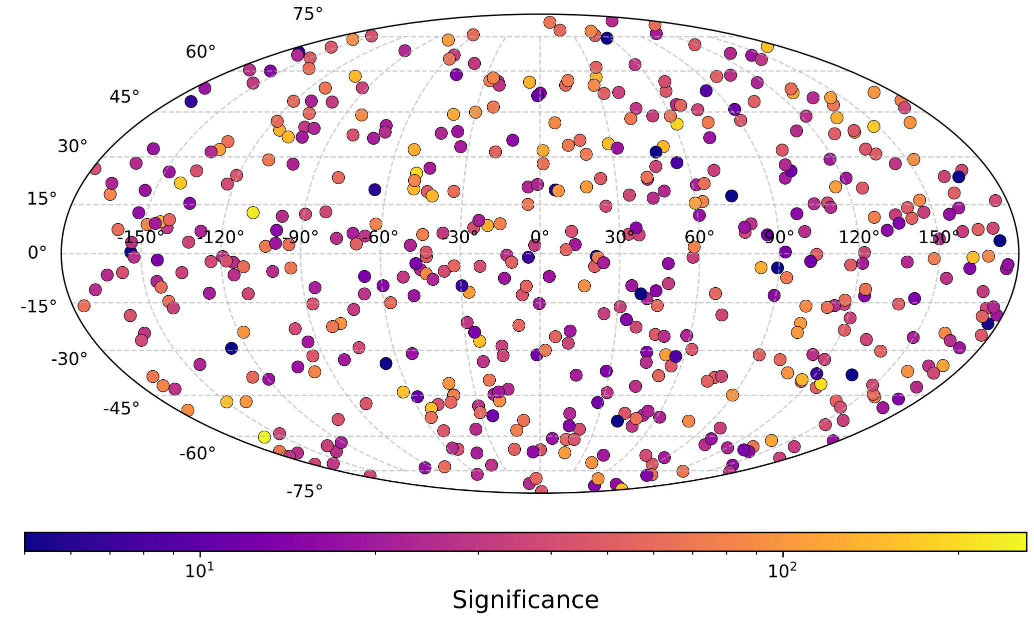

Figure 7.7 – All-sky map of the simulated GRBs, color-coded with their detection significance by the COSI ACS.

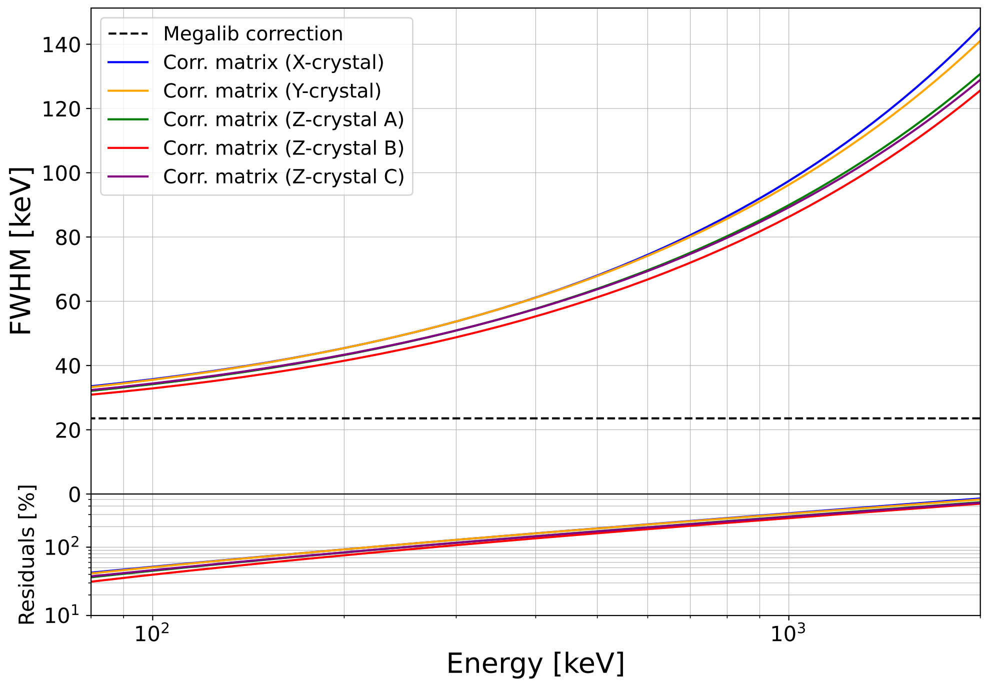

Figure 7.8 – The top panel displays the FWHMs (in keV) of the Gaussian noise applied by MEGAlib (constant per BGO crystal) compared to the FWHM values adopted in the correction matrix (crystal-dependent). Each correction matrix value is obtained by averaging the FWHMs measured in every voxel of the corresponding BGO crystal. The bottom panel shows the residuals (in %) between the correction-matrix values and the MEGAlib values.

energy deposits in the BGO crystals (i.e., without any noise) and the voxel where the deposit occurred. It then identifies the fit parameters of the corresponding voxel from the correction file and computes the centroid and FWHM of the Gaussian distribution. It finally draws a new corrected energy from the distribution and applies the energy threshold of 80 keV. We compare the results obtained using the correction matrix with those using the standard MEGAlib correction. The constant Gaussian noise of $\sigma = 10$ keV applied by MEGAlib corresponds to FWHM $\simeq 2.355\,\sigma = 23.55$ keV. Figure 7.8 shows the FWHM of the Gaussian noise for both the standard MEGAlib correction and the correction matrix, the latter being different for each BGO crystal. Since the correction matrix is not only energy dependent but also position dependent, we report the average value over the entire crystal surface. The deviation between the noise applied with the correction matrix and the MEGAlib value increases with energy: it is $\sim 35\%$ at 80 keV and rises to $\sim 500\%$ at 2 MeV.

7.4.1 Light curves and energy spectra

We investigate the impact of the correction matrix on GRB light curves and energy spectra. Although spectral information is not expected to be available through TDRSS, but only occasionally from standard ground-station contacts, we nonetheless include the energy spectra in our analysis for completeness. The energy spectra were constructed by extracting the events within the signal range identified by the Bayesian blocks algorithm with bc-tools. For both the light curves and the spectra, we calculated the relative variations[5] between the

[5] The relative variations are computed as $R(\%) = (C_{corr} - C_{mega}) \times 100/C_{mega}$, where C_{corr} and C_{mega} are the counts with the correction matrix and with the MEGAlib correction, respectively.

counts obtained with the correction matrix and those obtained with the standard MEGAlib correction, for each temporal and energy bin. Figure 7.9 shows the distributions of these variations collected from all GRBs, for both light curves and energy spectra. For the light curves, the variations are mostly within ±10%, while for the energy spectra they can reach up to ±40%. The smaller variations in the light curves likely arise from the integration over the full energy range, which reduces the impact of energy-dependent corrections. Both distributions are slightly skewed towards positive variations (about 70% and 54% of light curve and energy spectra variations are positive), indicating that the correction matrix tends to enhance the detected counts. This effect arises likely because the correction matrix modifies the overall detection efficiency of the ACS. As shown in Figure 6.6, the regions of the BGO crystals with enhanced detection efficiency occupy a larger fraction of the surface and are therefore more likely to be hit by photons, resulting in an overall increase in the detection efficiency.

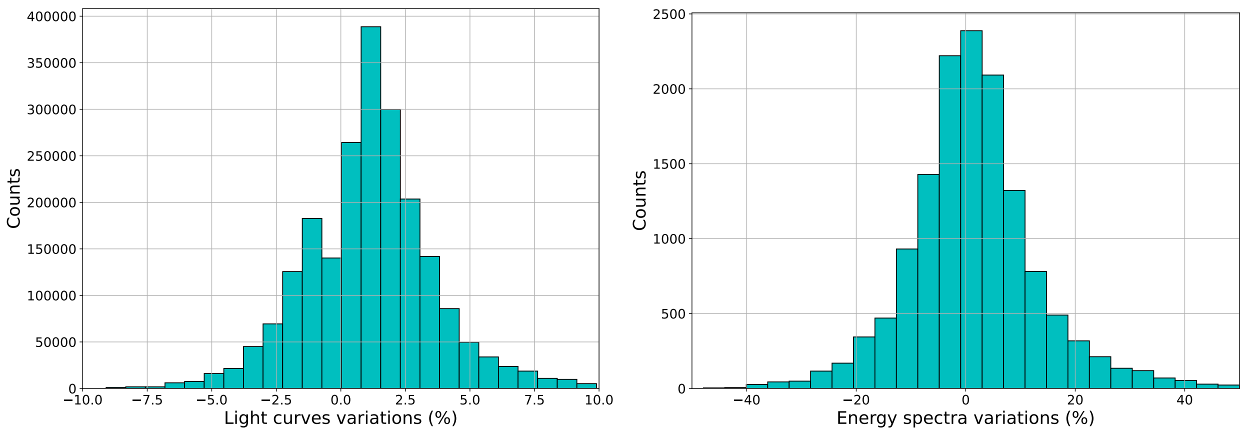

Figure 7.9 – Distributions of the relative variations between the counts with correction matrix and those with the standard MEGAlib correction, computed from the light curves (left) and energy spectra (right) at each temporal and energy bin.

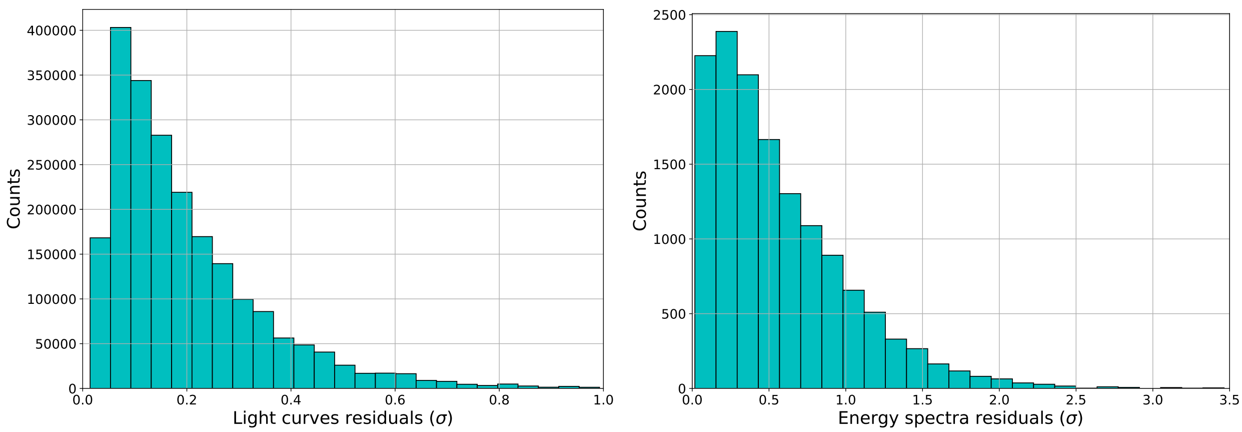

Figure 7.10 – Distributions of the variations in units of standard deviations between the counts with correction matrix and those with the standard MEGAlib correction, computed from the light curves (left) and energy spectra (right) at each temporal and energy bin.

To assess the significance of these variations, we also computed them in units of standard deviation (σ)[6]. Figure 7.10 shows the distribution of these residuals, aggregated over all GRBs,

[6] $R(\sigma) = |C_{corr} - C_{mega}| / \sqrt{\sigma^2_{C,corr} + \sigma^2_{C,mega}}$

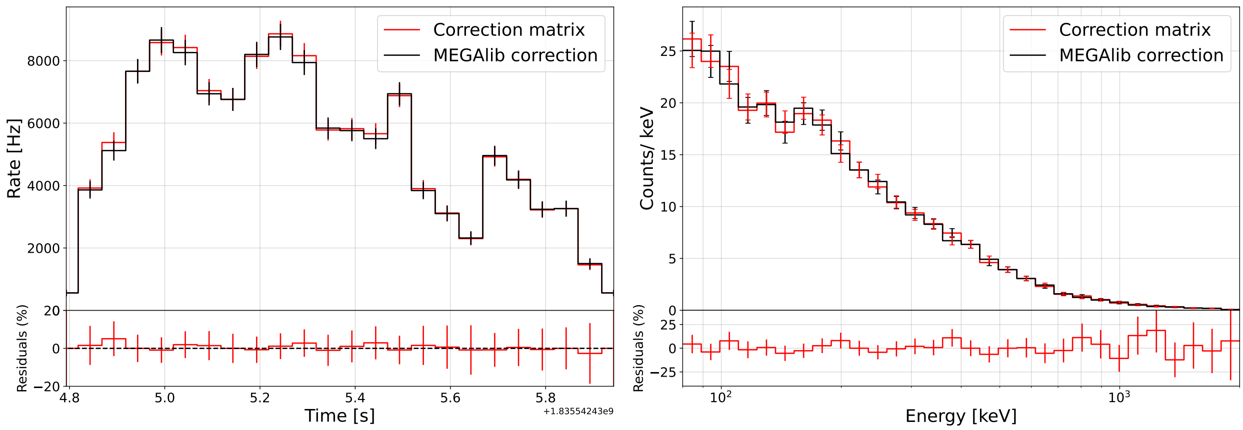

Figure 7.11 – Light curve (left) and energy spectrum (right) of the example GRB with and without correction matrix.

for both light curves and energy spectra. For light curves, the deviations are generally not significant, being mostly contained within 1σ. For energy spectra, about 17% of the residuals are larger than 1σ, reaching a maximum value of $\sim 3\sigma$. These deviations may become relevant for particularly bright GRBs, where even small variations can produce noticeable effects. Figure 7.11 illustrates the light curve and the energy spectrum of the reference GRB, comparing the standard MEGAlib correction with the correction matrix.

The fact that the differences introduced by the correction matrix are relatively small can be attributed to several factors. First, as mentioned earlier, the light curves are composed of energy-integrated counts, thereby reducing the visibility of any energy-dependent corrections. In the case of the spectra, we verified that even when comparing the corrected spectra (with the correction matrix) with those obtained without any correction, the magnitude of the variations remains largely the same. This suggests that the applied noise level is not sufficient to produce statistically significant effects, either for the MEGAlib correction or for our correction matrix. This is particularly true for the smooth, approximately power-law-shaped energy spectra that do not exhibit sharp spectral features. Contrary to the MEGAlib correction, our matrix applies a correction that is also position-dependent. This dependence, however, does not introduce significant variations either. In Chapter 6, a spatial pattern in the BGO response was observed, arising from the arrangement of the SiPMs. Nonetheless, the absolute amplitude of these spatial variations remains small: for the quantum detection efficiency, we measured deviations within $\pm 7\%$ (Figure 6.6). This limited variation is mostly due to the high efficiency of the SiPMs and the large absorption length of the BGO crystals, which together produce a large spatial uniformity in the detector response.

We also investigated whether the impact of the correction matrix depends on the GRB spectral hardness. Following a similar definition used in von Kienlin et al. (2020a) but adapted for the COSI ACS energy range, we define the hardness (H) of a GRB as

$$H = \frac{F(100 - 10000 \text{ keV})}{F(10 - 100 \text{ keV})} , \tag{7.7}$$

where $F(10 - 100$ keV) and $F(100 - 10000$ keV) are the integrated fluxes over the indicated energy range, using the GRB spectral model. Using this definition, we find the hardness to range from ~ 0.1 to ~ 10, with a median value of approximately $H \sim 1$. Figure 7.12 shows the variation in units of σ (averaged over the time bins for the light curves and energy bins

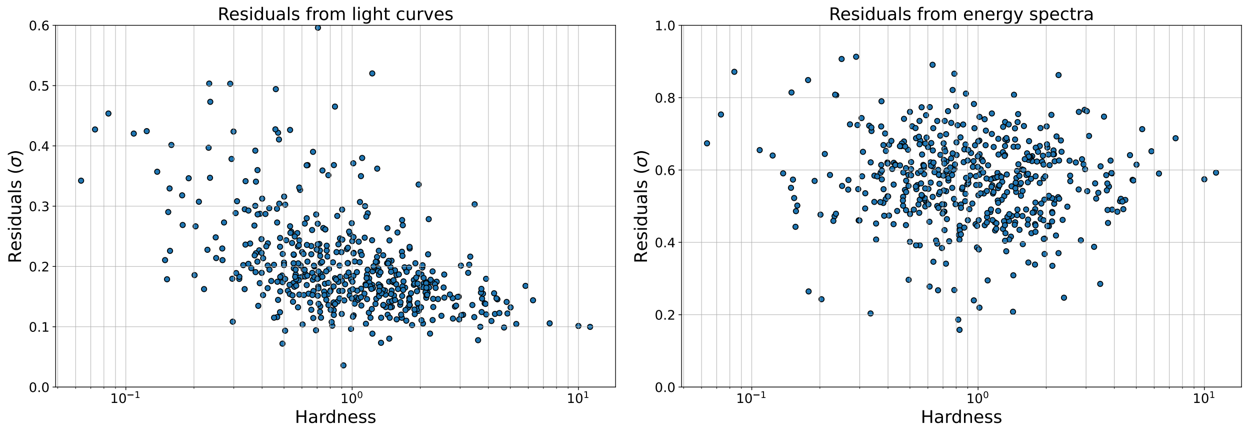

Figure 7.12 – Mean residuals in the GRB light curves (left) and energy spectra (right) as a function of the GRB hardness.

for the spectra) introduced by the correction matrix, as a function of the flux and H. The computed Pearson correlation coefficients are all very close to zero ($|r| < 0.1$), indicating that the average residuals, both from the spectra and from the light curves, show no significant correlation with the hardness.

7.4.2 Relative counts among ACS segments

The GRB localization from ACS data relies on the analysis of the relative counts among the different panels. In fact, depending on the GRB arrival direction, the individual ACS segments are more or less illuminated, leading to variations in their count rates (Parmiggiani et al. 2025). We investigated whether the correction matrix affects the integrated counts of the individual ACS segments. We computed the relative variations between the MEGAlib correction and the correction matrix in the integrated counts (over the signal range) for each ACS segment. Figure 7.13 illustrates the distributions of such variations. The correction matrix impacts the integrated counts within $\pm 2\%$. We also explored whether these variations depend on the GRB arrival directions. In Figure 7.14, we present HealPix-binned sky maps[7] showing the relative variations in the integrated counts (over the signal range) for each ACS segment. Each map pixel corresponds to the average value of all simulated directions falling within that HealPix bin. Although the statistical fluctuations are generally larger ($\sim 5\%$), the sky distributions reveal a spatial pattern, with the positive variations concentrated in the areas exposed to the ACS segment.

We then computed the ratios of the integrated counts between pairs of ACS segments, with the MEGAlib correction and the correction matrix, and evaluated the relative variations induced by the correction. The corresponding HealPix maps, for a subset of possible combinations, are shown in Figure 7.15. As for the rates for the individual ACS segments, these variations exhibit a spatial dependence. For example, the ratio $Y1/Z0$ is systematically enhanced in the north-western portion of the map (where the Y1 panel is directly exposed), whereas it is suppressed in the south-eastern region. These variations can reach up to 2-3%, although the statistical errors are estimated to be larger ($\sim 6\%$).

The origin of these asymmetries lies in the fact that the ACS panels are not uniformly exposed

[7] HealPix (Hierarchical Equal Area isoLatitude Pixelation) is a pixelation scheme of a spherical surface in which each pixel covers the same surface area as every other pixel.

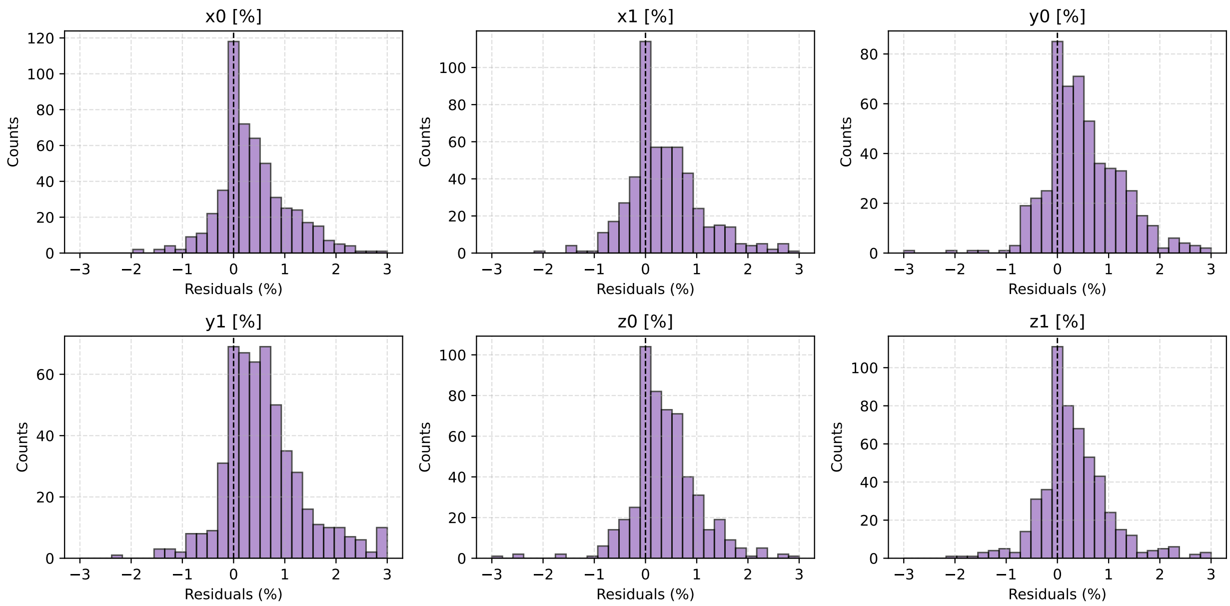

Figure 7.13 – Distribution of the relative variations between the integrated counts (over the signal range) obtained with the MEGAlib correction and those obtained with the correction matrix, for each ACS segment.

to incoming radiation, and their exposure changes depending on the GRB incident direction. In Figure 7.16, we show the voxel hit distribution map for the X0 panel, obtained by summing all counts within the signal range from all simulated GRBs. Specifically, we computed the maps separately for GRBs arriving from the northern and southern hemispheres[8]. In both cases, the panel is not uniformly illuminated: for instance, the outer borders are, on average, more exposed than the central regions. Moreover, the two distributions are not identical: for GRBs coming from the north, the ACS panel is more illuminated along the northern edge, whereas for those coming from the south, the hit distribution is more concentrated in the lower part. In this configuration, the SiPMs are mounted on the lower side of each BGO crystal. As shown in Figure 6.6, the regions of the BGO crystals close to the SiPMs include several spatial bins whose detection efficiency is suppressed by $\sim 2\%$ with respect to the average, except for a few bins directly in front of the SiPMs. On the other hand, the regions far away from the SiPMs exhibit an overall enhanced efficiency of $\sim$ 2-3%. Therefore, the asymmetries observed in Figure 7.14 are likely due to these variations in the interaction distribution, as the correction is applied as a function of the interaction position within the BGO crystals. As discussed in Section 7.4.1, despite the presence of spatial variations, the BGO response remains largely uniform, owing to the excellent SiPM performance and the large absorption length of the BGO. Consequently, these variations are small and fall within the range of statistical fluctuations.

[8]We emphasize that, in this analysis, the Earth occultation effect has been neglected. Therefore, in a real scenario a portion of the southern hemisphere is actually not visible.

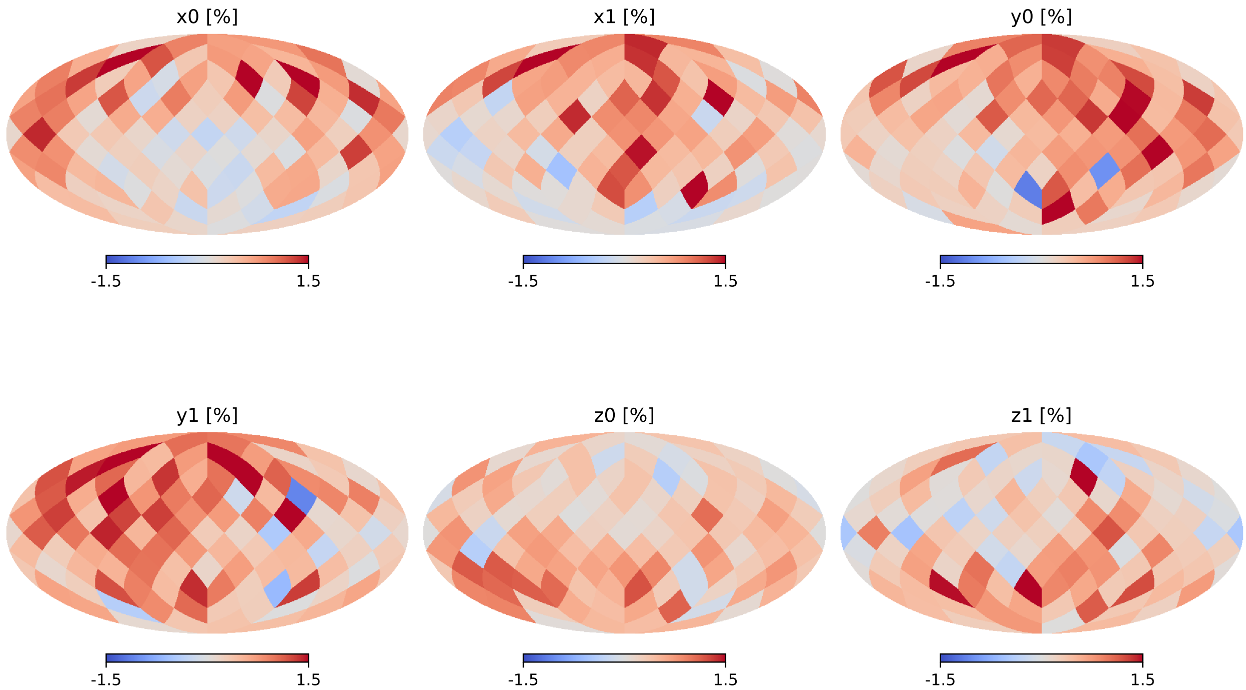

Figure 7.14 – HEALPIX maps of the relative variations in the integrated counts induced by the correction matrix, for each ACS segment.

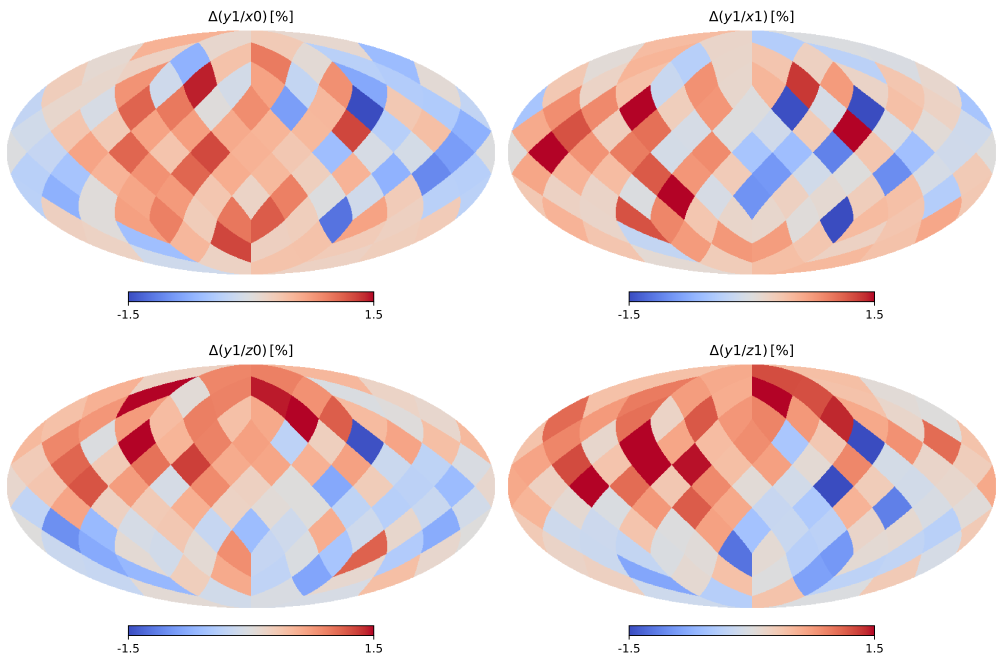

Figure 7.15 – HEALPIX maps of the relative variations in the ratio of counts between pairs of ACS segments, for a few possible combinations.

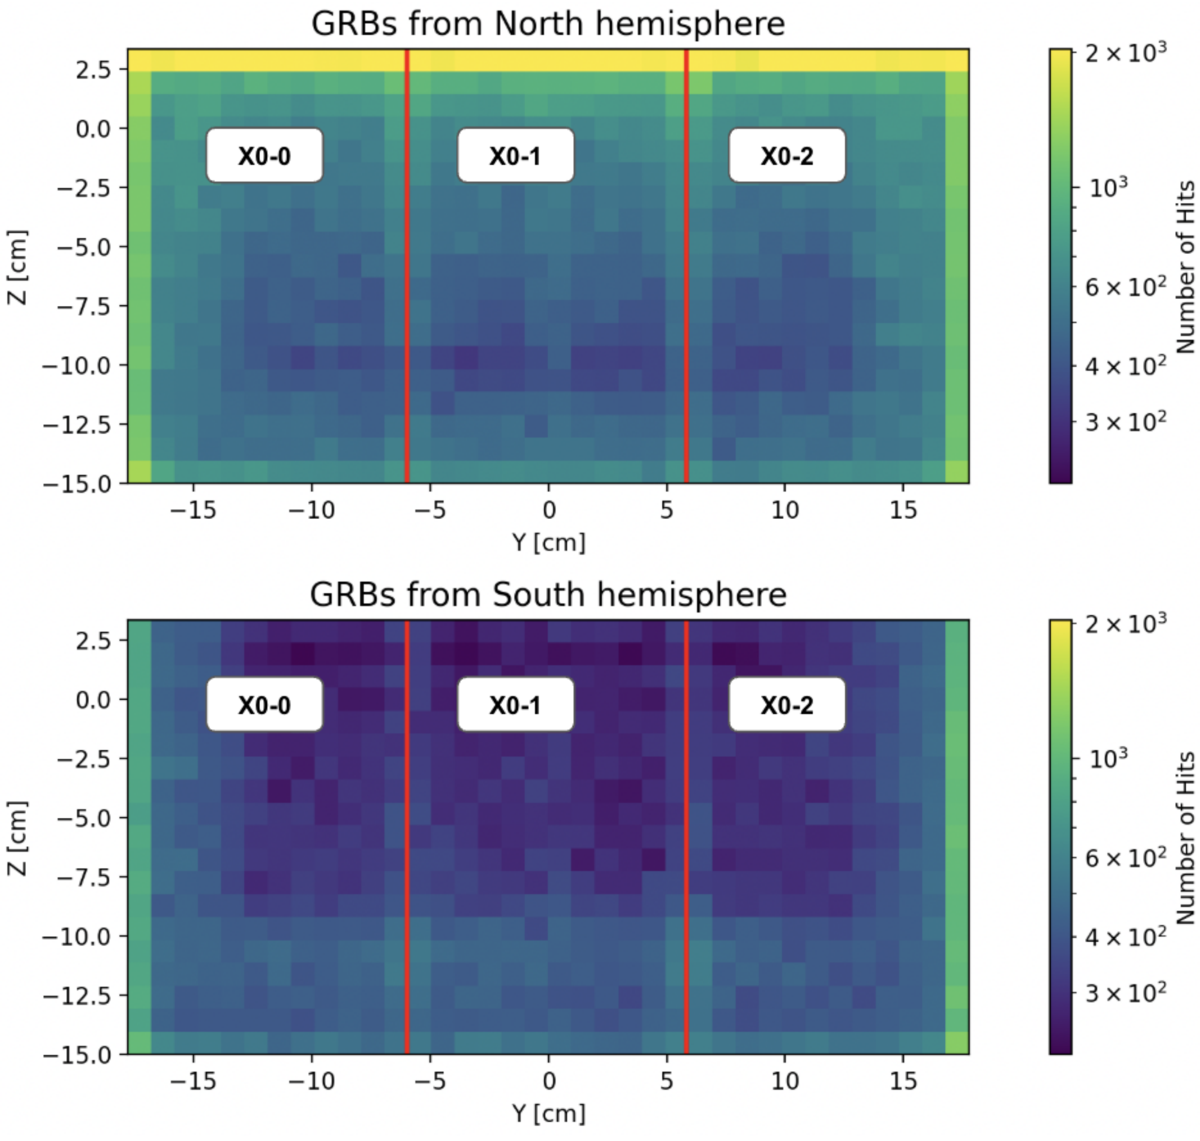

Figure 7.16 – Voxel hit distribution map for the X0 panel, obtained by summing all hits from all simulated GRBs from the northern (top panel) and southern (bottom panel) hemispheres.

Chapter 8

Conclusion and future perspectives

COSI represents a promising NASA SMEX mission that aims to address the observational gap in the 0.2-5 MeV energy band. Scheduled for launch in 2027, it will provide imaging, spectroscopy, and polarization capabilities through an array of 16 cross-strip GeDs. Surrounding the GeDs, the ACS will serve as a key instrument for the COSI mission. Comprising 22 BGO crystals, it will function not only as a veto detector for background suppression, but also as an active gamma-ray monitor for transient sources, with GRBs as its primary scientific targets. The ACS will contribute to the detection and localization of these events, thereby playing a fundamental role in multi-messenger astronomy. Ensuring realistic and accurate simulations of the ACS is therefore essential for predicting its performance both as a veto and as an active detector. The performance of the COSI ACS, especially its energy resolution and spatial response distribution, is intrinsically tied to the scintillation processes occurring within the BGO panels. Optical photons are produced whenever an interaction with ionizing radiation occurs and are collected by the SiPMs, which convert the scintillation light into measurable electrical signals.

In this work, we developed a correction matrix that adjusts the energy measured by the BGO crystals of the COSI ACS based on the interaction position and the true deposited energy. The correction matrix, to be included in the dedicated DEE for the ACS, accounts for the scintillation processes within the BGO crystals and the electronic noise from the SiPM readout, thereby enhancing the realism and accuracy of the ACS simulations. Characterizing these corrections required the development of optimized high-performance simulations, verification of the optical simulations through comparisons with theoretical models and experimental data, and the acquisition of laboratory measurements to enable accurate benchmarking of the simulations. The correction matrix was then tested on a representative population of simulated GRBs, assessing the capability of the COSI ACS to detect such transient events.

The most relevant results achieved during my PhD program and discussed in this thesis can be summarized as follows:

Development of HPC-optimized Geant4 simulations The simulation of the scintillation process usually requires tracking a large amount of optical photons. To achieve sufficient statistics for the correction matrix, more than 10^8 optical photons have to be generated. This led to the necessity of efficient simulations conducted in HPC environments, exploiting parallel execution in multiple resources. We used as baseline the simulation framework

BoGEMMS-HPC to enable efficient parallel execution in HPC environments. We implemented the Geant4 built-in multi-threading library and used the G4MPI library for multi-task execution. We evaluated the performance of the code under both parallelization schemes and for all supported output data formats (FITS, SQLite, MongoDB) on the Leonardo supercomputer. Using up to 112 cores, MPI parallelization provided a speed-up of about 70 with respect to serial execution and nearly 30 with respect to MT. This optimization allowed us to have an efficient simulation framework that addresses the computationally demanding optical simulations.

Verification of the Geant4 optical simulation To verify the Geant4 optical simulation, we compared the simulated transmission, reflection, and absorption probabilities of optical photons with theoretical expectations derived from Snell's and Fresnel's equations. In addition, we reproduced an experimental setup involving a PMT coupled to a CsI crystal to investigate the optical parameters and achieve the best agreement with laboratory measurements. Geant4 successfully reproduced both the analytical model and the experimental data, with maximum discrepancies of approximately 15% and 5%, respectively. The main source of discrepancy is attributed to the Geant4's limited capability to accurately simulate the transmission of optical photons through thin layers, where the wave nature of light becomes significant. For the subsequent simulations for the COSI ACS (with BGO and SiPMs), there are no such thin layers, thereby limiting the impact of such discrepancies in the benchmarking results.

Benchmarking of the COSI ACS simulations To ensure that the simulation accurately reproduces the ACS performance while constraining unknown parameters such as the BGO absorption length and the SiPM electronic noise, we used calibration measurements of the lateral ACS wall conducted with radioactive sources to benchmark the simulation. By combining the laboratory data acquired at NRL (by the COSI team) and at SSL (by myself), we first characterized the experimental ACS energy resolution (16% at 511 keV) and spatial response (variations contained within $\sim$ 10%). We reproduced the experimental setups with BoGEMMS-HPC and compared the resulting spectra, energy resolution, and spatial response with the experimental values. The BGO spatial response and energy resolution were reproduced with a maximum discrepancy of $\sim$ 10% and $\sim$ 20%, respectively. Regarding the energy spectra of the individual calibration sources, the simulations reproduced the photopeaks and Compton tails with average discrepancies of 12% and 18%, respectively. Overall, the simulation benchmarking achieved discrepancies within 20%, meeting the target accuracy defined by the COSI team, while constraining the BGO absorption length to 5 m and the electronic noise to 21 keV.

Development of the ACS correction matrix Once the simulation was benchmarked against laboratory measurements, we used it to develop the correction matrix that adjusts the energy measured by the ACS according to the interaction position and the actual energy deposited in the crystal. Using up to 100000 core-hours on the Leonardo supercomputer, we performed dedicated BoGEMMS-HPC simulations of the lateral ACS wall including the full optical physics processes. The BGO surface was divided into a 20×10 spatial grid, and for each spatial bin we characterized the ACS energy response as a Gaussian distribution, from which to sample the the corrected energy. The centroids of these distributions reveal a clear spatial pattern related to the SiPM positioning, with variations of up to $\pm$20%.

The FWHMs (i.e., the energy resolutions) exhibit greater uniformity, except in the regions directly facing the SiPMs, where the proximity enhances the positional sensitivity and thus increases the energy dispersion. Since simulations of the entire spacecraft require including the full ACS, we also carried out smaller-scale simulations on the local OAS computing node to develop a correction matrix tailored to each individual BGO crystal. Similar spatial patterns were observed across all crystals, though the smaller bottom one exhibited overall higher uniformity in both centroids and FWHMs due to their compact geometry.

GRB time analysis with the COSI ACS The correction matrix was finally tested on a representative population of short GRBs extracted from the fourth Fermi-GBM catalog. At the same time, we explored the capability of the COSI ACS as a stand-alone instrument for GRB detection. The GRBs were properly combined with background data to emulate realistic ACS ratemeters and analyzed with the bc-tools library, which employs the Bayesian Blocks algorithm for signal identification. Using the standard MEGAlib correction (constant Gaussian noise of $\sigma = 10$ keV), about 95% of the 541 GRBs were detected with a significance larger than 5σ. The resulting significances showed no dependence on the GRB arrival direction, while a moderate correlation with the GRB flux was found. We then compared these results with those obtained using the correction matrix. For the light curves, the relative variations are mostly within $\pm 10\%$, while for the energy spectra the variations can reach up to $\pm 40\%$. Both distributions have a larger fraction of positive variation (about 70% and 54% for the light curves and energy spectra, respectively), suggesting that the correction matrix tends to enhance the overall ACS quantum detection efficiency. The smaller variations observed in the light curves are due to the integration over the full energy range, which mitigates the effect of energy-dependent corrections. For light curves, the variations are within the Poisson uncertainty ($< 1\sigma$); for energy spectra, a small fraction ($\sim 17\%$) of the deviations lies in the range 1σ-3σ. We further investigated its impact on the integrated counts of the individual ACS segments, as these quantities are used for GRB localization. The variations induced by the correction matrix in both the integrated counts and the ratios of counts between pairs of ACS segments are within $\pm 2\%$, still inside the statistical uncertainty. Nevertheless, a dependence on the GRB arrival direction is observed, due to the position dependence of the correction matrix.

Final remarks and future perspectives

The DEE represents a fundamental component of the COSI simulation pipeline. The work I carried out during the PhD contributed directly to its development, meeting the required specifications and providing a concrete addition to a NASA mission. The DEE will be integrated into the fifth data challenge (DC5), making it publicly available and usable by the entire scientific community. Furthermore, all the codes developed during my PhD activity for the simulation benchmarking, analysis, and development of the correction matrix are collected into a dedicated GitLab repository, allowing other users to reproduce, verify, and extend this work.

The development of the correction matrix allowed us to quantitatively assess whether position- and energy-dependent effects are relevant, rather than assuming they are negligible a priori. One of the key outcomes of this work has been the assessment of the high homogeneity of the BGO crystals and the effectiveness of their coupling with the SiPMs. The large transparency of BGO to optical photons, combined with the strategic placement of the SiPMs, ensures an overall uniform ACS response and thus reduces the impact of the developed correction. Nevertheless, it is important to underline that the conditions under which the correction matrix was computed may differ from those in orbit. Variations in the operating temperature or in the ACS energy threshold could amplify the effects of the correction matrix. A proper understanding of the temperature effect on the ACS performance through future dedicated laboratory measurements is mandatory to accurately quantify the impact of such corrections.

Future calibration campaigns will lead to a potential refinement of the simulation benchmarking and a subsequent update of the correction matrix. It is also worth noting that, in this work, we focused primarily on GRBs, which are the main scientific targets of the ACS. However, the ACS may also detect other transient sources with different temporal or spectral properties, for which dedicated simulations are required. Broadening the analysis to include a wider range of scientific targets to test the correction matrix would be an interesting extension of this work.

The investigation of the correction impact on GRB simulations also provided an opportunity to assess the performance of the COSI ACS as a stand-alone transient detector. The analysis demonstrated its capability to efficiently detect short GRBs, confirming its key scientific contribution to the overall mission.

This work also required the development and optimization of a simulation framework capable of exploiting HPC technologies, including multi-threading and MPI multi-tasking. The resulting framework, BoGEMMS-HPC, is now publicly available for the scientific community, providing an HPC-compliant simulation environment that can be applied to a wide range of X-ray and gamma-ray detectors. Future developments of BoGEMMS-HPC could focus on enhancing its parallelization efficiency, improving its user accessibility, and integrating database-driven tools for online or streaming analysis.

Finally, this research provided an in-depth exploration and validation of Geant4 optical physics and the related simulation parameters. To date, only a limited number of studies have addressed this topic in detail. As scintillator panels coupled to SiPM are becoming increasingly common in next-generation space instruments, such as in compact CubeSat missions (e.g., the GRID mission, which employs GAGG:Ce scintillators read out by SiPM arrays), this work provides a valuable methodological framework that can be adapted and extended to future detectors using similar architectures.

Appendix A

NRL and SSL experimental spectra for the EM X-wall

We show in Figure A.1 and Figure A.2 the experimental spectra for each uncollimated source, with the corresponding background, for both NRL and SSL experiments.

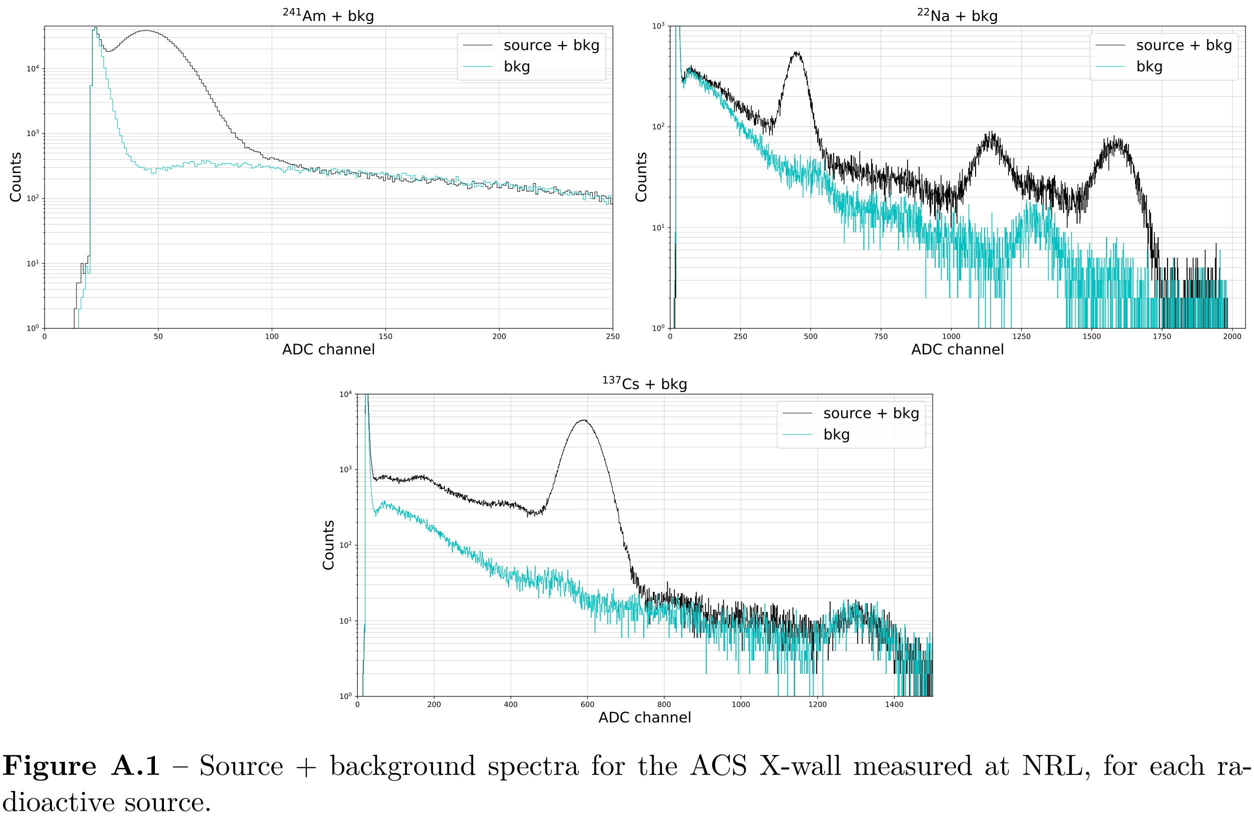

Figure A.1 – Source + background spectra for the ACS X-wall measured at NRL, for each radioactive source.

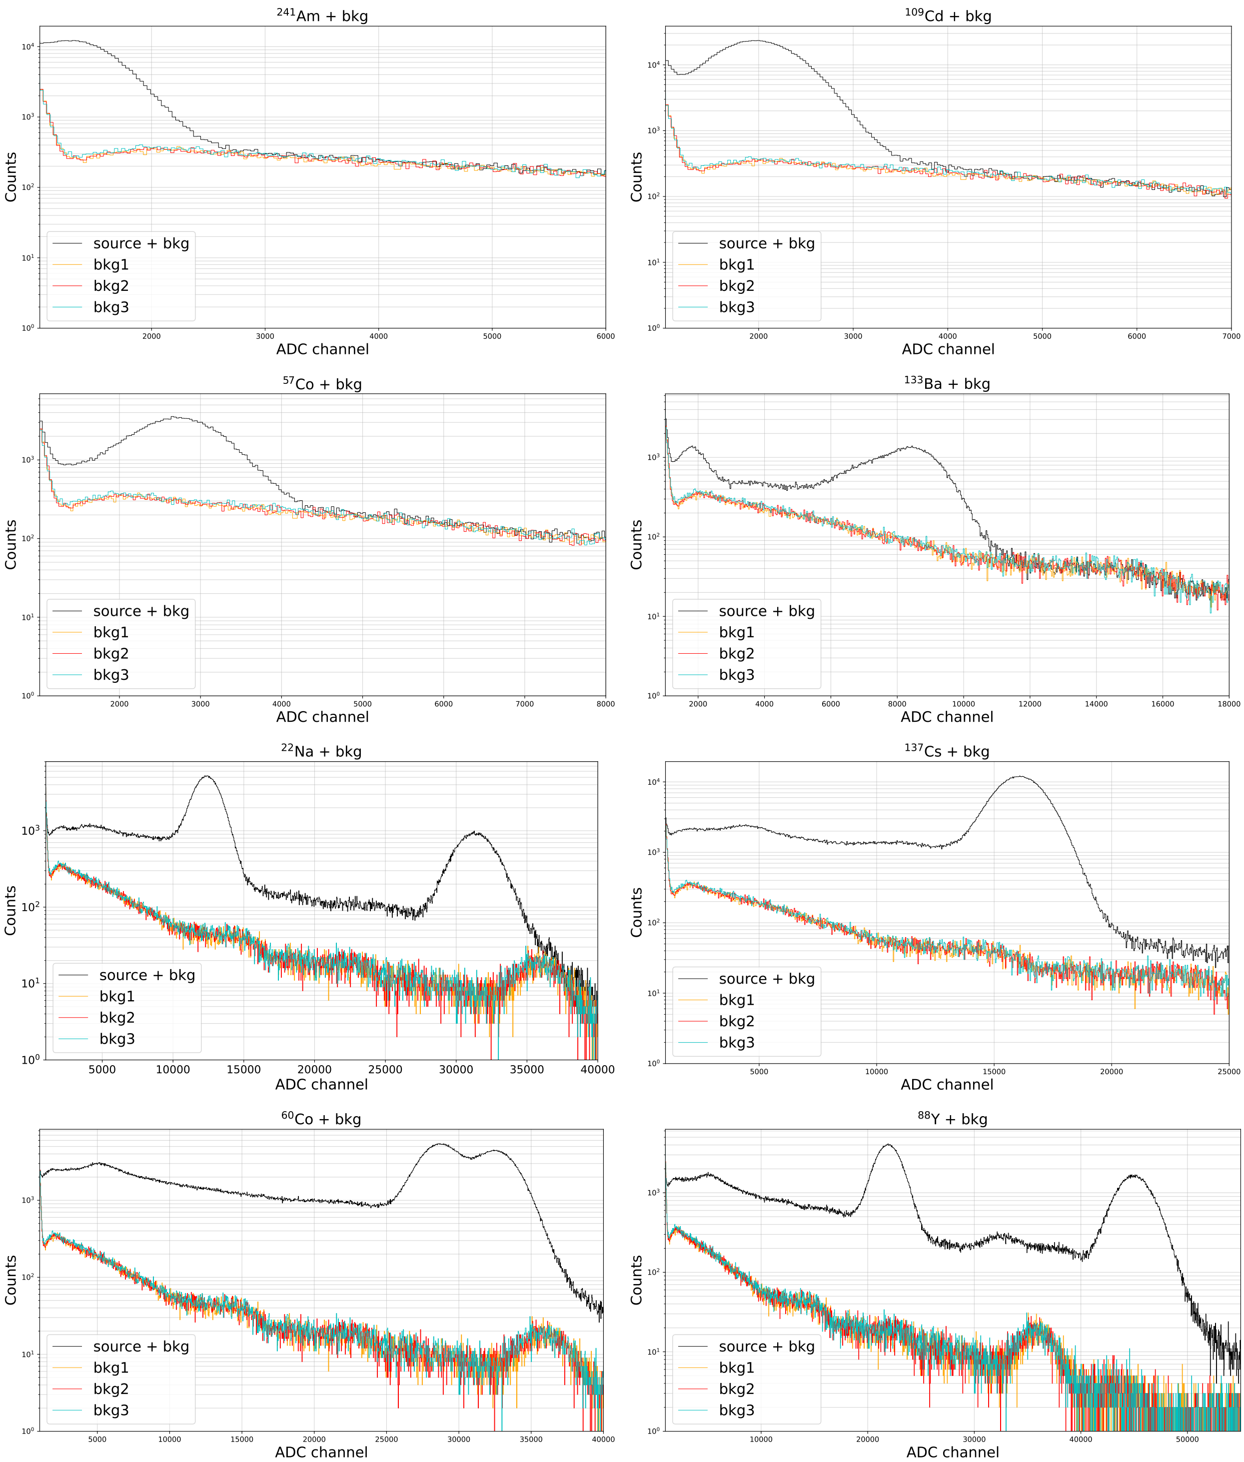

Figure A.2 – Source + background spectra for the ACS X-wall measured at SSL, for each radioactive source.

Acknowledgements

Supported by Italian Research Center on High Performance Computing Big Data and Quantum Computing (ICSC), project funded by European Union - NextGenerationEU - and National Recovery and Resilience Plan (NRRP) - Mission 4 Component 2 within the activities of Spoke 2 (Fundamental Research and Space Economy), (CN 00000013 - CUP C53C22000350006).

The Compton Spectrometer and Imager is a NASA Explorer project led by the University of California, Berkeley with funding from NASA under contract 80GSFC21C0059.

We acknowledge the use of lab space, equipment, and engineering support at UC Berkeley's Space Sciences Laboratory.

References

Abbott, B. P. et al. (Oct. 2017). 'Gravitational Waves and Gamma-Rays from a Binary Neutron Star Merger: GW170817 and GRB 170817A'. In: *The Astrophysical Journal Letters* 848.2, L13, p. L13. DOI: 10.3847/2041-8213/aa920c. arXiv: 1710.05834 [astro-ph.HE].

Abdo, A. A. et al. (Aug. 2010). 'Gamma-Ray Emission Concurrent with the Nova in the Symbiotic Binary V407 Cygni'. In: *Science* 329.5993, pp. 817–821. DOI: 10.1126/science.1192537. arXiv: 1008.3912 [astro-ph.HE].

Ackermann, M. et al. (Nov. 2015). 'Multiwavelength Evidence for Quasi-periodic Modulation in the Gamma-Ray Blazar PG 1553+113'. In: *The Astrophysical Journal Letters* 813.2, L41, p. L41. DOI: 10.1088/2041-8205/813/2/L41. arXiv: 1509.02063 [astro-ph.HE].

Agostinelli, S. et al. (July 2003). 'G EANT4a simulation toolkit'. In: *NIM A* 506.3, pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.

Ajello, M. et al. (Feb. 2021). 'First Fermi-LAT Solar Flare Catalog'. In: *The Astrophysical Journal Supplement Series* 252.2, 13, p. 13. DOI: 10.3847/1538-4365/abd32e. arXiv: 2101.10010 [astro-ph.HE].

Akchurin, N. et al. (June 2011). 'A comparison of BGO and BSO crystals used in the dual-readout mode'. In: *Nuclear Instruments and Methods in Physics Research A* 640.1, pp. 91–98. DOI: 10.1016/j.nima.2011.03.013.

Allison, J. et al. (Feb. 2006). 'Geant4 developments and applications'. In: *IEEE Trans. Nucl. Sci.* 53.1, pp. 270–278. DOI: 10.1109/TNS.2006.869826.

Allison, J. et al. (Nov. 2016). 'Recent developments in GEANT4'. In: *NIM A* 835, pp. 186–225. DOI: 10.1016/j.nima.2016.06.125.

Alvarez Franco, D. (2022). 'Performance of the photodetector for the active shield of the COSI space telescope'. In: *Luleå University of Technology (dissertation thesis).*

Atwood, W. B. et al. (June 2009). 'The Large Area Telescope on the Fermi Gamma-Ray Space Telescope Mission'. In: *Astrophysical Journal* 697.2, pp. 1071–1102. DOI: 10.1088/0004-637X/697/2/1071. arXiv: 0902.1089 [astro-ph.IM].

Band, D. et al. (Aug. 1993). 'BATSE Observations of Gamma-Ray Burst Spectra. I. Spectral Diversity'. In: *Astrophysical Journal* 413, p. 281. DOI: 10.1086/172995.

Berger, Edo (Aug. 2014). 'Short-Duration Gamma-Ray Bursts'. In: *ARAA* 52, pp. 43–105. DOI: 10.1146/annurev-astro-081913-035926. arXiv: 1311.2603 [astro-ph.HE].

Bismark, Alexander (Dec. 2019). 'Simulation and characterization of a LXe TPC for DARWIN RD'. PhD thesis. DOI: 10.6094/UNIFR/154696.

Bissaldi, E. et al. (May 2009). 'Ground-based calibration and characterization of the Fermi gamma-ray burst monitor detectors'. In: *Experimental Astronomy* 24.1-3, pp. 47–88. DOI: 10.1007/s10686-008-9135-4. arXiv: 0812.2908 [astro-ph].

Bose, D. et al. (Jan. 2022). 'Galactic and extragalactic sources of very high energy gamma rays'. In: *European Physical Journal Special Topics* 231.1, pp. 27–66. DOI: 10.1140/epjs/s11734-022-00434-8. arXiv: 2201.06789 [astro-ph.HE].

Bouchet, L. et al. (Sept. 2010). 'On the Morphology of the Electron-Positron Annihilation Emission as Seen by Spi/integral'. In: *Astrophysical Journal* 720.2, pp. 1772–1780. DOI: 10.1088/0004-637X/720/2/1772. arXiv: 1007.4753 [astro-ph.HE].

Bouchet, Laurent et al. (Mar. 2015). 'The Galactic ^{26}Al Emission Map as Revealed by INTEGRAL SPI'. In: *Astrophysical Journal* 801.2, 142, p. 142. DOI: 10.1088/0004-637X/801/2/142. arXiv: 1501.05247 [astro-ph.HE].

Briggs, M. S. et al. (July 2010). 'First results on terrestrial gamma ray flashes from the Fermi Gamma-ray Burst Monitor'. In: *Journal of Geophysical Research (Space Physics)* 115.A7, A07323, A07323. DOI: 10.1029/2009JA015242.

Bulgarelli, A. et al. (2012). 'BoGEMMS: the Bologna Geant4 multi-mission simulator'. In: *High Energy, Optical, and Infrared Detectors for Astronomy V*. Ed. by Andrew D. Holland and James W. Beletic. Vol. 8453. International Society for Optics and Photonics. SPIE, p. 845335. DOI: 10.1117/12.926065. URL: https://doi.org/10.1117/12.926065.

Bulgarelli, Andrea (Dec. 2019). 'The AGILE Gamma-Ray observatory: software and pipelines'. In: *Experimental Astronomy* 48.2-3, pp. 199–231. DOI: 10.1007/s10686-019-09644-w.

Chupp, E. L. et al. (Dec. 1982). 'A direct observation of solar neutrons following the 0118 UT flare on 1980 June 21.' In: *The Astrophysical Journal Letters* 263, pp. L95–L99. DOI: 10.1086/183931.

Ciabattoni, Alex et al. (Aug. 2024). 'Towards a response function for the COSI anticoincidence system: preliminary results from Geant4 simulations'. In: *Space Telescopes and Instrumentation 2024: Ultraviolet to Gamma Ray*. Ed. by Jan-Willem A. den Herder et al. Vol. 13093. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 130937Y, 130937Y. DOI: 10.1117/12.3020252. arXiv: 2409.12327 [astro-ph.IM].

Ciabattoni, Alex et al. (Aug. 2025). 'Benchmarking of Geant4 simulations for the COSI Anticoincidence System'. In: *Experimental Astronomy* 60.1, 9, p. 9. DOI: 10.1007/s10686-025-10019-7. arXiv: 2507.21275 [astro-ph.IM].

Compton, Arthur H. (May 1923). 'A Quantum Theory of the Scattering of X-rays by Light Elements'. In: *Phys. Rev.* 21 (5), pp. 483–502. DOI: 10.1103/PhysRev.21.483. URL: https://link.aps.org/doi/10.1103/PhysRev.21.483.

Connaughton, V. et al. (Feb. 2015). 'Localization of Gamma-Ray Bursts Using the Fermi Gamma-Ray Burst Monitor'. In: *The Astrophysical Journal Supplement Series* 216.2, 32, p. 32. DOI: 10.1088/0067-0049/216/2/32. arXiv: 1411.2685 [astro-ph.IM].

Cummings, J. R. and Swift BAT Team (Dec. 2002). 'The Swift Burst Alert Telescope (BAT)'. In: *American Astronomical Society Meeting Abstracts*. Vol. 201. American Astronomical Society Meeting Abstracts, 41.14, p. 41.14.

Diehl, R. et al. (June 1995). 'COMPTEL observations of Galactic 2̂6Âl emission.' In: *Astronomy and Astrophysics* 298, p. 445.

do Nascimento, E. et al. (2005). 'Doppler broadening of positron annihilation radiation: fitting the coincidence spectrum'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 538.1, pp. 723–730. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2004.09.013. URL: https://www.sciencedirect.com/science/article/pii/S0168900204021072.

Dorosz, P. et al. (2013). 'Silicon photomultiplier's gain stabilization by bias correction for compensation of the temperature fluctuations'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equip-*

ment 718. Proceedings of the 12th Pisa Meeting on Advanced Detectors, pp. 202–204. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2012.11.116. URL: https://www.sciencedirect.com/science/article/pii/S0168900212014635.

Dotti, Andrea et al. (May 2016). 'Extending Geant4 Parallelism with External Libraries (MPI, TBB) and Its Use on HPC Resources'. In: *arXiv e-prints*, arXiv:1605.01792. DOI: 10.48550/arXiv.1605.01792. arXiv: 1605.01792 [physics.comp-ph].

Fioretti, V. et al. (2012). 'The low Earth orbit radiation environment and its impact on the prompt background of hard x-ray focusing telescopes'. In: vol. 8453, 845331, p. 845331. DOI: 10.1117/12.926248.

Fioretti, V. and A. Ciabattoni (2025). *Pipeline for GEANT4 simulations in HPC environments, sub-project of Pipeline optimization for space and ground based experiments (PSGE)*. Tech. rep. MS7 KPI Report. ICSC. URL: https://drive.google.com/file/d/1p6zeKekNMb4mIQ72MBVgjgtPsqp4Dr9k/view.

Fioretti, V. et al. (Nov. 2024). 'Unveiling the origin of XMM-Newton soft proton flares: I. Design and validation of a response matrix for proton spectral analysis'. In: *Astronomy and Astrophysics* 691, A229, A229. DOI: 10.1051/0004-6361/202451231.

Fioretti, V. et al. (2025). 'BoGEMMS-HPC: a Geant4 simulation framework tailored for high performance computing applications'. In: in preparation.

Fioretti, Valentina et al. (2014). 'Monte Carlo simulations of gamma-ray space telescopes: a BoGEMMS multi-purpose application'. In: *Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray*. Ed. by Tadayuki Takahashi et al. Vol. 9144. International Society for Optics and Photonics. SPIE, 91443N. DOI: 10.1117/12.2056442. URL: https://doi.org/10.1117/12.2056442.

Fishman, G. J. et al. (Jan. 1994). *Discovery of intense gamma-ray flashes of atmospheric origin*. Draft Report National Aeronautics and Space Administration. Marshall Space Flight Center, Huntsville, AL.

Frost, K. J. and E. D. Rothe (1962). 'Detector for Low Energy Gamma-Ray Astronomy Experiment'. In: *IRE Transactions on Nuclear Science* 9.3, pp. 381–385. DOI: 10.1109/TNS2.1962.4316023.

Gallego, Savitri et al. (June 2025). 'Bottom-up Background Simulations of the 2016 COSI Balloon Flight'. In: *Astrophysical Journal* 986.2, 116, p. 116. DOI: 10.3847/1538-4357/add6a0. arXiv: 2503.02493 [astro-ph.IM].

Geant4-collaboration (2023). *Physics Reference Manual*.

Gehrels, N. et al. (Aug. 2004). 'The Swift Gamma-Ray Burst Mission'. In: *Astrophysical Journal* 611.2, pp. 1005–1020. DOI: 10.1086/422091. arXiv: astro-ph/0405233 [astro-ph].

Gruber, David et al. (Mar. 2014). 'The Fermi GBM Gamma-Ray Burst Spectral Catalog: Four Years of Data'. In: *The Astrophysical Journal Supplement Series* 211.1, 12, p. 12. DOI: 10.1088/0067-0049/211/1/12. arXiv: 1401.5069 [astro-ph.HE].

Gupta, Santosh and Yuanbing Mao (Sept. 2020). *Nano Scintillator-Book*.

Hamamatsu (2007). *Photomultipliers Tubes - Basics and Applications*. Hamamatsu Photonics K. K.

Holl, I. et al. (1988). 'A measurement of the light yield of common inorganic scintillators'. In: *IEEE Transactions on Nuclear Science* 35.1, pp. 105–109. DOI: 10.1109/23.12684.

Hurley, Kevin (2011). 'Soft gamma repeaters'. In: *Advances in Space Research* 47.8. Neutron Stars and Gamma Ray Bursts, pp. 1326–1331. ISSN: 0273-1177. DOI: https://doi.org/10.1016/j.asr.2010.03.001. URL: https://www.sciencedirect.com/science/article/pii/S0273117710001523.

Inan, Umran S. et al. (May 1996). 'On the association of terrestrial gamma-ray bursts with lightning and implications for sprites'. In: 23.9, pp. 1017–1020. DOI: 10.1029/96GL00746.

Ishak, Nurliyana et al. (Dec. 2015). 'Estimating single tree stem and branch biomass using terrestrial laser scanning'. In: *Jurnal Teknologi* 77. DOI: 10.11113/jt.v77.6860.

Janecek, Martin (2012a). 'Reflectivity Spectra for Commonly Used Reflectors'. In: *IEEE Transactions on Nuclear Science - IEEE TRANS NUCL SCI* 59, pp. 490–497. DOI: 10.1109/TNS.2012.2183385.

Janecek, Martin (June 2012b). 'Reflectivity Spectra for Commonly Used Reflectors'. In: *IEEE Transactions on Nuclear Science* 59.3, pp. 490–497. DOI: 10.1109/TNS.2012.2183385.

Janecek, Martin and William Moses (2010). 'Simulating Scintillator Light Collection Using Measured Optical Reflectance'. In: *Nuclear Science, IEEE Transactions on* 57, pp. 964–970. DOI: 10.1109/TNS.2010.2042731.

Johnson III, W. N. et al. (Feb. 1972). 'The Spectrum of Low-Energy Gamma Radiation from the Galactic-Center Region.' In: *The Astrophysical Journal Letters* 172, p. L1. DOI: 10.1086/180878.

Kierans, Carolyn et al. (2022). 'Compton Telescopes for Gamma-Ray Astrophysics'. In: *Handbook of X-ray and Gamma-ray Astrophysics*. Ed. by Cosimo Bambi and Andrea Santangelo. Singapore: Springer Nature Singapore, pp. 1–72. ISBN: 978-981-16-4544-0. DOI: 10.1007/978-981-16-4544-0_46-1. URL: https://doi.org/10.1007/978-981-16-4544-0_46-1.

Kierans, Carolyn A. et al. (2016). 'The 2016 Super Pressure Balloon flight of the Compton Spectrometer and Imager'. In: *PoS INTEGRAL2016*. Ed. by Ed van den Heuvel, p. 075. DOI: 10.22323/1.285.0075. arXiv: 1701.05558 [astro-ph.IM].

Klein, O. and T. Nishina (Nov. 1929). 'Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac'. In: *Zeitschrift fur Physik* 52.11-12, pp. 853–868. DOI: 10.1007/BF01366453.

Knoll, G. F. (1979). *Radiation detection and measurement.*

Knoll, Glenn F. (2010). *Radiation Detection and Measurement, 3rd ed.* 4th edition. New York: John Wiley and Sons. ISBN: 978-0-471-07338-3, 978-0-471-07338-3.

Kouveliotou, C. et al. (May 1998). 'An X-ray pulsar with a superstrong magnetic field in the soft γ-ray repeater SGR1806 - 20'. In: *Nature* 393.6682, pp. 235–237. DOI: 10.1038/30410.

Kouveliotou, Chryssa et al. (Aug. 1993). 'Identification of Two Classes of Gamma-Ray Bursts'. In: *The Astrophysical Journal Letters* 413, p. L101. DOI: 10.1086/186969.

Levin, A. and C. Moisan (1996). 'A more physical approach to model the surface treatment of scintillation counters and its implementation into DETECT'. In: *1996 IEEE Nuclear Science Symposium. Conference Record*. Vol. 2, 702–706 vol.2. DOI: 10.1109/NSSMIC.1996.591410.

Li, T. -P. and Y. -Q. Ma (Sept. 1983). 'Analysis methods for results in gamma-ray astronomy.' In: *Astrophysical Journal* 272, pp. 317–324. DOI: 10.1086/161295.

Lucchetta, Giulio et al. (Aug. 2022). 'Introducing the MeVCube concept: a CubeSat for MeV observations'. In: *Journal of Cosmology and Astroparticle Physics* 2022.08, p. 013. DOI: 10.1088/1475-7516/2022/08/013. URL: https://dx.doi.org/10.1088/1475-7516/2022/08/013.

Luongo, Orlando and Marco Muccino (Oct. 2021). 'A Roadmap to Gamma-Ray Bursts: New Developments and Applications to Cosmology'. In: *Galaxies* 9.4, 77, p. 77. DOI: 10.3390/galaxies9040077. arXiv: 2110.14408 [astro-ph.HE].

Marisaldi, M. et al. (Mar. 2010). 'Detection of terrestrial gamma ray flashes up to 40 MeV by the AGILE satellite'. In: *Journal of Geophysical Research (Space Physics)* 115.53, A00E13, A00E13. DOI: 10.1029/2009JA014502.

Martinez-Castellanos, Israel (Sept. 2023). 'The bc-tools library: an open-source software for BurstCube and other GRB satellites'. In: *AAS/High Energy Astrophysics Division*. Vol. 20. AAS/High Energy Astrophysics Division, 103.66, p. 103.66.

Martinez-Castellanos, Israel et al. (Aug. 2023). 'The cosipy library: COSI's high-level analysis software'. In: *arXiv e-prints*, arXiv:2308.11436, arXiv:2308.11436. DOI: 10.48550/arXiv.2308.11436. arXiv: 2308.11436 [astro-ph.HE].

Meegan, Charles et al. (Sept. 2009). 'The Fermi Gamma-ray Burst Monitor'. In: *Astrophysical Journal* 702.1, pp. 791–804. DOI: 10.1088/0004-637X/702/1/791. arXiv: 0908.0450 [astro-ph.IM].

Merzi, Stefano et al. (Aug. 2024). 'Radiation Damage on Silicon Photomultipliers from Ionizing and Non-Ionizing Radiation of Low-Earth Orbit Operations'. In: *Sensors* 24.15, 4990, p. 4990. DOI: 10.3390/s24154990. arXiv: 2405.07191 [physics.ins-det].

Meyer, Manuel et al. (May 2019). 'Characterizing the Gamma-Ray Variability of the Brightest Flat Spectrum Radio Quasars Observed with the Fermi LAT'. In: *Astrophysical Journal* 877.1, 39, p. 39. DOI: 10.3847/1538-4357/ab1651. arXiv: 1902.02291 [astro-ph.HE].

Moszyski, M. et al. (Sept. 1981). 'Timing properties of BGO scintillator'. In: *Nuclear Instruments and Methods in Physics Research* 188.2, pp. 403–409. DOI: 10.1016/0029-554X(81)90521-8.

Nakarmi, P. et al. (Jan. 2020). 'Reflectivity and PDE of VUV4 Hamamatsu SiPMs in liquid xenon'. In: *Journal of Instrumentation* 15.1, P01019. DOI: 10.1088/1748-0221/15/01/P01019. arXiv: 1910.06438 [physics.ins-det].

Narayana Bhat, P. et al. (Apr. 2016). 'The Third Fermi GBM Gamma-Ray Burst Catalog: The First Six Years'. In: *The Astrophysical Journal Supplement Series* 223.2, 28, p. 28. DOI: 10.3847/0067-0049/223/2/28. arXiv: 1603.07612 [astro-ph.HE].

Parmiggiani, N. et al. (2023). 'The AGILE real-time analysis software system to detect short-transient events in the multi-messenger era'. In: *Astronomy and Computing* 44, p. 100726. ISSN: 2213-1337. DOI: https://doi.org/10.1016/j.ascom.2023.100726. URL: https://www.sciencedirect.com/science/article/pii/S2213133723000410.

Parmiggiani, N. et al. (Sept. 2024). 'A New Deep Learning Model to Detect Gamma-Ray Bursts in the AGILE Anticoincidence System'. In: *Astrophysical Journal* 973.1, 63, p. 63. DOI: 10.3847/1538-4357/ad64cd.

Parmiggiani, N. et al. (2025). 'COSI Short Gamma-Ray Burst Localization Using BGO Shield Data'. In: submitted.

Patel, Parshad et al. (2025). 'Goals, Development, and Performance of the Compton Spectrometer and Imagers Anti-Coincidence Subsystem'. In: *PoS* ICRC2025, p. 799. DOI: 10.22323/1.501.0799.

Piran, Tsvi (1999). 'Gamma-ray bursts and the fireball model'. In: *Physics Reports* 314.6, pp. 575–667. ISSN: 0370-1573. DOI: https://doi.org/10.1016/S0370-1573(98)00127-6. URL: https://www.sciencedirect.com/science/article/pii/S0370157398001276.

Poole, C. M. et al. (2012). 'A CAD Interface for GEANT4'. In: *Australasian Physical & Engineering Science in Medicine*. DOI: 10.1007/s13246-012-0159-8. URL: http://www.springerlink.com/content/u563877422284578.

Prokhorov, D. A. and A. Moraghan (Feb. 2023). 'An update on Fermi-LAT transients in the Galactic plane, including strong activity of Cygnus X-3 in mid-2020'. In: *Monthly Notices* *of the Royal Astronomical Society* 519.2, pp. 2680–2689. DOI: 10.1093/mnras/stac3453. arXiv: 2209.12461 [astro-ph.HE].

Ramaty, Reuven and Natalie Mandzhavidze (Oct. 1998). 'Solar Flares: Gamma Rays'. In: *arXiv e-prints*, astro-ph/9810089, astro–ph/9810089. DOI: 10.48550/arXiv.astro-ph/9810089. arXiv: astro-ph/9810089 [astro-ph].

Rau, A. et al. (Aug. 2005). 'The 1st INTEGRAL SPI-ACS gamma-ray burst catalogue'. In: *Astronomy and Astrophysics* 438.3, pp. 1175–1183. DOI: 10.1051/0004-6361:20053159. arXiv: astro-ph/0504357 [astro-ph].

Roncali, Emilie and Simon Cherry (Mar. 2013). 'Simulation of light transport in scintillators based on 3D characterization of crystal surfaces'. In: *Physics in medicine and biology* 58, pp. 2185–2198. DOI: 10.1088/0031-9155/58/7/2185.

Roncali, Emilie et al. (Apr. 2019). 'Cerenkov light transport in scintillation crystals explained: realistic simulation with GATE'. In: *Biomedical Physics & Engineering Express* 5.3, p. 035033. DOI: 10.1088/2057-1976/ab0f93. URL: https://dx.doi.org/10.1088/2057-1976/ab0f93.

Ryde, Felix and Asaf Pe'er (Sept. 2009). 'Quasi-blackbody Component and Radiative Efficiency of the Prompt Emission of Gamma-ray Bursts'. In: *Astrophysical Journal* 702.2, pp. 1211–1229. DOI: 10.1088/0004-637X/702/2/1211. arXiv: 0811.4135 [astro-ph].

Scargle, Jeffrey D. (Sept. 1998). 'Studies in Astronomical Time Series Analysis. V. Bayesian Blocks, a New Method to Analyze Structure in Photon Counting Data'. In: *Astrophysical Journal* 504.1, pp. 405–418. DOI: 10.1086/306064. arXiv: astro-ph/9711233 [astro-ph].

Siegert, Thomas et al. (July 2020). 'Imaging the 511 keV Positron Annihilation Sky with COSI'. In: *Astrophysical Journal* 897.1, 45, p. 45. DOI: 10.3847/1538-4357/ab9607. arXiv: 2005.10950 [astro-ph.HE].

Sleator, Clio et al. (Mar. 2019a). 'Benchmarking simulations of the Compton Spectrometer and Imager with calibrations'. In: *AAS/High Energy Astrophysics Division*. Vol. 17. AAS/High Energy Astrophysics Division, 109.63, p. 109.63.

Sleator, Clio et al. (Jan. 2019b). 'Benchmarking simulations of the Compton Spectrometer and Imager with calibrations'. In: *American Astronomical Society Meeting Abstracts #233*. Vol. 233. American Astronomical Society Meeting Abstracts, 146.22, p. 146.22.

Smith, David M. et al. (Feb. 2005). 'Terrestrial Gamma-Ray Flashes Observed up to 20 MeV'. In: *Science* 307.5712, pp. 1085–1088. DOI: 10.1126/science.1107466.

Smith, J. (July 2019). 'BurstCube: Mission Concept, Performance, and Status'. In: *36th International Cosmic Ray Conference (ICRC2019)*. Vol. 36. International Cosmic Ray Conference, 604, p. 604. DOI: 10.48550/arXiv.1907.11069. arXiv: 1907.11069 [astro-ph.IM].

Spectrometer, The Compton and Imager (COSI) Collaboration (Apr. 2025). *COSI Data Challenges*. Version 3.1. Zenodo. DOI: 10.5281/zenodo.15126188. URL: https://doi.org/10.5281/zenodo.15126188.

Tavani, M. et al. (Aug. 2009). 'The AGILE Mission'. In: *Astronomy and Astrophysics* 502.3, pp. 995–1013. DOI: 10.1051/0004-6361/200810527. arXiv: 0807.4254 [astro-ph].

Teles, J. et al. (Aug. 1995). 'Overview of TDRSS'. In: *Advances in Space Research* 16.12, pp. 67–76.

Tomsick, J. et al. (Sept. 2024). 'The Compton Spectrometer and Imager'. In: *38th International Cosmic Ray Conference*, 745, p. 745. DOI: 10.48550/arXiv.2308.12362. arXiv: 2308.12362 [astro-ph.HE].

Tomsick, John A. et al. (2019). *Compton Spectrometer and Imager (COSI): Concept Study Report.* Tech. rep.

Ursi, A. et al. (Apr. 2017). 'Terrestrial gamma-ray flashes in the BeppoSAX data archive'. In: *Journal of Atmospheric and Solar-Terrestrial Physics* 156, pp. 50–56. DOI: 10.1016/j.jastp.2017.02.014.

Ursi, Alessandro et al. (Dec. 2019). 'Detection of terrestrial gamma-ray flashes with the AGILE/MCAL'. In: *Rendiconti Lincei. Scienze Fisiche e Naturali* 30, pp. 265–269. DOI: 10.1007/s12210-019-00762-3. arXiv: 2108.11108 [physics.space-ph].

Valentine, John D. et al. (1993a). 'Temperature dependence of CsI (T1) absolute scintillation yield'. In: *IEEE Transactions on Nuclear Science* 40.4, pp. 1267–1274. DOI: 10.1109/TNS.1993.8526779.

Valentine, John D. et al. (1993b). 'Temperature dependence of CsI(Tl) gamma-ray excited scintillation characteristics'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 325.1, pp. 147–157. ISSN: 0168-9002. DOI: 10.1016/0168-9002(93)91015-F. URL: https://www.sciencedirect.com/science/article/pii/016890029391015F.

Vilmer, N. et al. (Apr. 2003). 'On the solar origin of interplanetary disturbances observed in the vicinity of the Earth'. In: *Annales Geophysicae* 21.4, pp. 847–862. DOI: 10.5194/angeo-21-847-2003.

von Kienlin, A. et al. (Nov. 2003). 'INTEGRAL Spectrometer SPI's GRB detection capabilities. GRBs detected inside SPI's FoV and with the anticoincidence system ACS'. In: *Astronomy and Astrophysics* 411, pp. L299–L305. DOI: 10.1051/0004-6361:20031231. arXiv: astro-ph/0308346 [astro-ph].

von Kienlin, A. et al. (Apr. 2020a). 'The Fourth Fermi-GBM Gamma-Ray Burst Catalog: A Decade of Data'. In: *Astrophysical Journal* 893.1, 46, p. 46. DOI: 10.3847/1538-4357/ab7a18. arXiv: 2002.11460 [astro-ph.HE].

von Kienlin, A. et al. (Apr. 2020b). 'The Fourth Fermi-GBM Gamma-Ray Burst Catalog: A Decade of Data'. In: *Astrophysical Journal* 893.1, 46, p. 46. DOI: 10.3847/1538-4357/ab7a18. arXiv: 2002.11460 [astro-ph.HE].

von Kienlin, Andreas et al. (Mar. 2014). 'The Second Fermi GBM Gamma-Ray Burst Catalog: The First Four Years'. In: *The Astrophysical Journal Supplement Series* 211.1, 13, p. 13. DOI: 10.1088/0067-0049/211/1/13. arXiv: 1401.5080 [astro-ph.HE].

Williams, P. A. et al. (July 1996). 'Optical, thermo-optic, electro-optic, and photoelastic properties of bismuth germanate ($Bi_4Ge_3O_{12}$)'. In: 35.19, pp. 3562–3569. DOI: 10.1364/AO.35.003562.

Winkler, C. et al. (Nov. 2003a). 'The INTEGRAL mission'. In: *Astronomy and Astrophysics* 411, pp. L1–L6. DOI: 10.1051/0004-6361:20031288.

Winkler, C. et al. (Nov. 2003b). 'The INTEGRAL mission'. In: *Astronomy and Astrophysics* 411, pp. L1–L6. DOI: 10.1051/0004-6361:20031288.

Yoneda, H. et al. (May 2025). 'Enhancing Compton telescope imaging with maximum a posteriori estimation: A modified RichardsonLucy algorithm for the Compton Spectrometer and Imager'. In: *Astronomy and Astrophysics* 697, A117, A117. DOI: 10.1051/0004-6361/202453528. arXiv: 2504.02468 [astro-ph.IM].

York, Donald G. et al. (Sept. 2000). 'The Sloan Digital Sky Survey: Technical Summary'. In: *Astronomical Journal* 120.3, pp. 1579–1587. DOI: 10.1086/301513. arXiv: astro-ph/0006396 [astro-ph].

Zanin, R. et al. (Nov. 2016). 'Gamma rays detected from Cygnus X-1 with likely jet origin'. In: *Astronomy and Astrophysics* 596, A55, A55. DOI: 10.1051/0004-6361/201628917. arXiv: 1605.05914 [astro-ph.HE].

Zhang, Bing et al. (Dec. 2000). 'Nature and Nurture: a Model for Soft Gamma-Ray Repeaters'. In: *The Astrophysical Journal Letters* 545.2, pp. L127–L130. DOI: 10.1086/317889. arXiv: astro-ph/0010225 [astro-ph].

Zhang, H. and M. Böttcher (Sept. 2013). 'X-Ray and Gamma-Ray Polarization in Leptonic and Hadronic Jet Models of Blazars'. In: *Astrophysical Journal* 774.1, 18, p. 18. DOI: 10.1088/0004-637X/774/1/18. arXiv: 1307.4187 [astro-ph.HE].

Zoglauer, Andreas et al. (July 2008). 'MEGAlib: simulation and data analysis for low-to-medium-energy gamma-ray telescopes'. In: *Space Telescopes and Instrumentation 2008: Ultraviolet to Gamma Ray*. Ed. by Martin J. L. Turner and Kathryn A. Flanagan. Vol. 7011. Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, 70113F, 70113F. DOI: 10.1117/12.789537.

Zoglauer, Andreas et al. (Feb. 2021). 'COSI: From Calibrations and Observations to All-sky Images'. In: *arXiv e-prints*, arXiv:2102.13158, arXiv:2102.13158. DOI: 10.48550/arXiv.2102.13158. arXiv: 2102.13158 [astro-ph.IM].

Zoglauer, Andreas Christian (Jan. 2006). 'First light for the next generation of Compton and pair telescopes : Development of new techniques for the data analysis of combined Compton and pair telescopes and their application to the MEGA prototype'. PhD thesis. -.