Technische Universität München
TUM School of Natural Sciences

Gamma detection with CALIFA

Philipp Klenze

Vollständiger Abdruck der von der TUM School of Natural Sciences der Technischen Universität München zur Erlangung eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz: apl. Prof. Dr. Norbert Kaiser

Prüfende der Dissertation:

1. Prof. Dr. Laura Fabbietti
2. Prof. Dr. Lothar Oberauer

Die Dissertation wurde am 10.06.2024 bei der Technischen Universität München eingereicht und durch die TUM School of Natural Sciences am 23.07.2024 angenommen.

Gamma detection with CALIFA

PhD thesis

Philipp Klenze

Commit e6a6a1f
2024-09-25 16:50:30 +0200
Log: Changes to official title page, fixed date.

Abstract

The Facility for Antiproton and Ion Research (FAIR) currently under construction in Darmstadt will be the most powerful large-scale facility for heavy ion research in Germany. The main components of FAIR in the start version will be the SIS100 Synchrotron, the Super Fragment Separator (SuperFRS) and the NUSTAR (NUclear Structure, Astrophysics and Reactions) experiment installations. With a circumference of 1100 meters, SIS100 will allow for much higher beam intensities or energies compared to the current SIS18.

Primary beam intensities of 10^{11} particles per second allow to produce the most exotic fully stripped heavy exotic ions, precisely separated in a three-stage analysis of magnetic rigidity and energy loss in the large phase-space acceptance of the SuperFRS. As part of the NUSTAR collaboration, the R^3B experiment will be the important installation located at the focal plane of the high-energy branch after SuperFRS at FAIR. The R^3B (**R**eactions with **R**adioactive **R**elativistic **B**eams) setup is designed for kinematically complete reaction studies with exotic nuclei serving a broad physics program ranging from the investigation of the dynamics and structure of nuclei far off stability, studies of astrophysical reactions, investigation of collective properties in the most neutron rich species, to fundamental questions like the Equation Of State (EOS) of asymmetric nuclear matter at different densities.

The highly segmented CALIFA calorimeter – which has been continuously extended during the last years – is one of the key components in all these experiments. CALIFA is designed to determine the energy of γ rays in the ion frame of reference, it is highly segmented to measure the θ angle of the photons which is essential for the DOPPLER correction.

CALIFA currently consists of more than 1600 (out of an eventual total of 2540) CsI(Tl) scintillating crystals which will detect both γ rays and light charged particles emitted in an angular range of $7^\circ < \theta < 143^\circ$.

One focus of this work was the development of a new trigger system to get an improved time correlation between the individual crystal hits and to allow for an external reaction trigger especially important for high rate experiments.

Optimizing the detection of γ rays in CALIFA – especially at high energies ($E_{\gamma,Lab} \approx 5\,\mathrm{MeV}$) where e^+e^- pair production is the dominant interaction process – is the second key topic investigated here. Such γ rays generally deposit energy in multiple crystals which frequently form clusters.

Extensive GEANT4 simulation studies with a simplified version of the geometry were used to inform the parameters required for clustering (add-back) and give a deeper understanding of the cluster formation in CALIFA. For γ rays in the region of 3–10 MeV, a clustering radius of ten centimeters offers a good compromise between efficiency and rejection of random coincidences. In that energy region, the position resolution (σ) in θ direction is better than 7 mm for 98% of the events.

In a dedicated experiment, an Americium Beryllium (AmBe) reaction source which emits a γ ray at $E_\gamma = 4.4\,\mathrm{MeV}$ and neutrons was used to extend previous tests with calibration sources and verify the efficiency and background supression of the clustering scheme. Careful background subtraction showed an excess in γ,γ coincidences which correspond to γ cascade of ^{56}Fe which was excited from the neutrons.

As a byproduct of this work, a method to analyze crystal qualities with a single source measurement has been developed. This method is sensitive to distinguish between different imperfections like light transport, light collection or sensor problems in each detector element. All these major contributions had been already been implemented and used during the R^3B experiments in 2023 and 2024.

Contents

1 Physics motivation

Since prehistoric times, humans have the urge to understand how the world we inhabit came to be. While few if any creation myths credited the stars with being instrumental in the construction of everyday matter, 20th century cosmology, astrophysics and nuclear physics tells us that stars create most of the matter we consist of both during their billions of years of their lives and during the days of their violent ends.

From the discovery of the neutron by Chadwick in 1932 [1], our understanding of the atomic nucleus has progressed to a very detailed level. Modern calculations like large scale shell model or ab initio calculations can predict the nuclear wave functions and energies of individual states on a sub 100 keV level now. With the use of numerous rare isotope beam facilities like RIBF, FRIB, GSI-FAIR, ISOLDE and many more a huge number of very short-lived nuclei have been produced, identified and characterized by studying their masses, decay modes and level structure.

All known nuclei – namely those which were experimentally identified – are listed in the chart of nuclei, here shown in figure 1.

The stable-ish nuclei which are involved in the creation of heavier elements within the lifetime of a star are already well-characterized. However, a significant fraction of metals[1] is created in processes which happen during extreme astrophysical conditions such as supernova explosions and involve nuclei which have not been created on Earth – yet.

[1] In the context of astrophysics, metals are any elements heavier than helium.

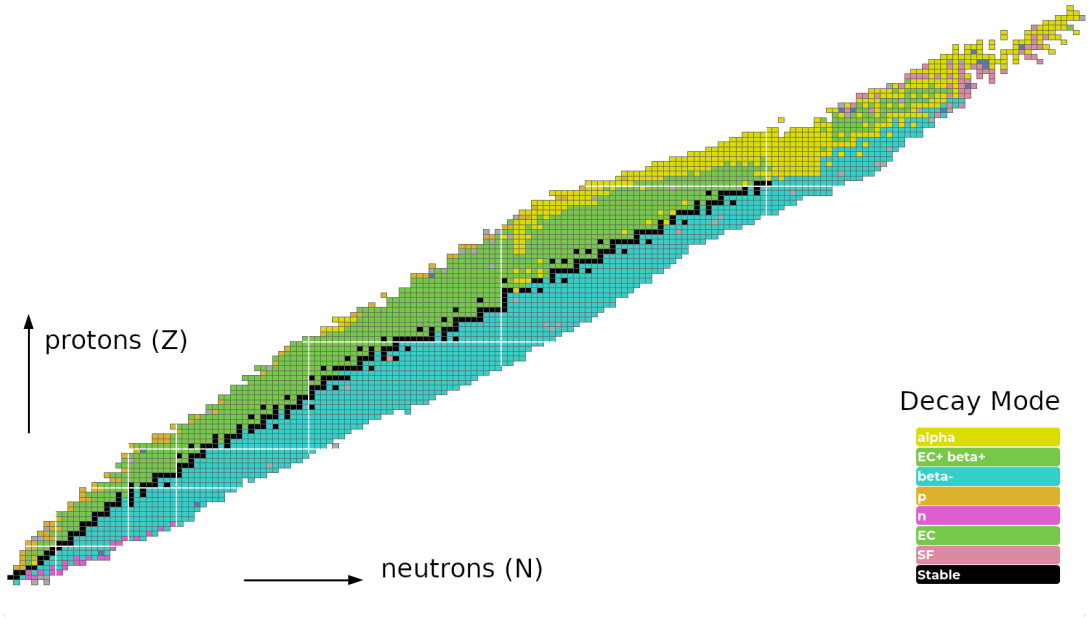

Figure 1: From [2], the chart of the known nuclides. The black squares indicate stable nuclei while the colored squares visualize the main decay mode of the isotope. Nuclides with white lines underneath or to the left of them have a magic number of protons or neutrons.

All known nuclides are visualized in Figure 1 with the color indicating their primary decay mode. Stable nuclei are indicated in black. Basically all matter on Earth is made just from these few isotopes. They form the base of the *valley of stability*.

The isotopes not located near the bottom of the valley of stability can often be transformed into their isobars through weak decays. The binding energy difference between neighboring isobar nuclides scales roughly linearly with the distance to the point of beta stability. Theory gives us *Sargent's rule*, which states that with regard to β decay, the inverse life-time τ^{-1} is proportional to ΔE^5. This arises from integrating the state density of the electron (or positron) and antineutrino (or neutrino) over the energy.

Very far from stability, the binding energy of the least bound nucleon will eventually be negative. That nucleon will still have to tunnel through the centrifugal barrier (and, in the case of protons, the Coulomb barrier) of the nucleus, so even such states have a finite half-life. The points where this happens are called the proton and neutron drip lines.

Nucleosynthesis describes the creation of new nuclei from other nuclei and nucleons. **Big bang nucleosynthesis** started about 20 seconds after big bang when the system had cooled down sufficiently so that the first composite nuclei were able to form without being immediately disintegrated by the radiation field. During this process, a free neutron would fuse with a proton to form deuterium, and two deuterium nuclei would fuse to form ^{4}He. In total, about a quarter of the baryon mass was thus fused into helium, with the remainder remaining free protons[2] – which is mostly still the present state.

Much later, some of that matter gravitationally aggregated into large celestial bodies. High pressure and temperature conditions within the cores of such objects lead to **stellar nucleosynthesis** – exothermic reactions in which lighter elements are fused into heavier ones. This process generally starts from hydrogen which is fused to ^{4}He by processes such as the **p–p chain**. If the mass of the star is high enough, stellar nucleosynthesis may continue all the way until ^{62}Ni, which has the highest binding energy per nucleon. The heaviest isotope typically produced in large quantities is the slightly lighter ^{56}Fe.

In later generations of stars, preexisting elements such as iron can slowly capture additional neutrons (emitted for example during the fusion of rare ^{13}C with helium). As this happens slowly – so that beta-unstable nuclei can generally decay before another neutron is captured – this is called the **s-process**. This can go on until ^{210}Po is formed, which will decay by emitting an *alpha*, reducing its mass number by four, long before it can hope to capture more neutrons. Thus the quasi-stable ^{209}Bi is the heaviest element produced here.

In very massive stars ($M > 10M_\odot$) the pressure within the core can eventually not withstand the gravitational pull, leading to a **core collapse supernova** as the degeneracy pressure of the electrons is overcome, forcing them to merge with the protons to form neutrons in the creation of a neutron star[3]. This is a highly exothermic process, and the resulting shock wave is an environment of very high pressure, temperature and neutron density. In such an environment, nuclei can absorb neutrons much faster than the beta decay converts them to protons, reaching close to the **neutron dripline**, where the neutron capture is in equilibrium with the neutron evaporation due to the high temperatures. As nucleosynthesis happens rapidly, this is called

[2]Give or take minute amounts of ^{2}H, ^{3}He and lithium.

[3]The neutron star may then collapse into a black hole.

the **r-process**. After eventual expansion, cooling and beta decay, the resulting (quasi-) stable nuclei also include the elements up to uranium and thorium.

While the above explains most of the abundances, it is also observed that stable proton-rich nuclei called **p-nuclei** such as ^{92}Mo exist in nature. These can not be formed by the beta decay of neutron rich nuclei, so some other mechanism has to produce them. Photodisintegration (**γ-process**) of heavier nuclei is one such process. Proton capture processes are also thought to play a role in forming these isotopes. While neutrons are easily absorbed by nuclei, protons carry an electrical charge and thus have to pass through the Coulomb barrier of the nucleus, so very high temperatures are required.[4] The rapid proton capture process (**rp-process**) is thought to take place on the surface of neutron stars which are parts of binary star systems and cause X-ray bursts. The pathway of the rp-process is thought to move along the **proton dripline**.

The nucleosynthesis processes which take place near either of the driplines involve nuclei which have not been discovered yet. With the advent of new high intensity accelerator facilities such as FAIR, these key regions of the nuclide chart will come within experimental reach.

[4]The Coulomb barrier is also why creating fusion reactors is hard, and why stars last for billions of years.

2 The FAIR facility and the R^{3}B experiment

The Facility for Antiproton and Ion Research (FAIR) – which is currently under construction near Darmstadt, Germany – will offer unprecedented opportunities to access ions far from the valley of stability.

One of the key questions FAIR is to study is the production of heavy elements in the universe. The modelling of astrophysical scenarios such as supernovae depends on the nuclear properties of the relevant isotopes both for the evolution of the process and the expected isotope yields. By determining the properties of the nuclei on the pathways, FAIR will offer valuable insight to astrophysics.

While the study of ions near both the neutron and proton dripline will allow new insights into both the r-process and rp-process pathways, the nuclear properties of these ions are also interest for other fields. Exotic nuclei can serve as a touchstone to test different nuclear models to test our current understanding of the nuclear many body system at the most extreme neutron to proton ratios.

FAIR is producing exotic ions by having a high intensity primary beam (such as ^{238}U) hitting a production target. This will produce secondaries via projectile fragmentation or fission. As can be estimated from figure 3, most of the reaction products will be well known isotopes. The **SuperFRS** (Super fragment separator) will impose a tiny moveable window of the production chart, forwarding just a handful of isotopes to the experiments at a time.

Depending on their half-life, different experiments can be done with these ions. They can be decelerated and stored in storage rings or even ion traps for mass measurements and decay spectroscopy, or the exotic beam can directly be used for in-flight reaction studies by placing a target in its path. The target will induce nuclear reactions in some of the ions, which can be used to probe the nuclear properties. Due to the fact that the role of target and beam are exchanged, compared to the traditional setup, this is called **inverse kinematics**.

The R^{3}B experiment is a key experiment which will offer insight into the properties of very exotic nuclei through reactions. Both the reaction channels and the detector setup they necessitate will be discussed.

2.1 Ion production and selection

The production of exotic ions will be accomplished by accelerating a stable ions up to ^{238}U in the SIS100 accelerator and hitting a rather thick production target made from light nuclei like beryllium or carbon at the entrance of the SuperFRS. Two processes are dominant to produce exotic ions from stable beams: fission and fragmentation. Fission naturally requires a fissile beam such as ^{238}U and can either be induced by exciting the nucleus using a virtual photon from a high-Z target material such as lead or inelastic scattering processes on a light target. SuperFRS will use light targets such as Li [3]. During the fission, the nucleus splits into two daughter nuclei. As the fissile material has a much higher neutron to proton ratio than the lighter products, the fission products are generally neutron-rich. Splitting fissile nuclei is an exothermic process releasing about 170 MeV as kinetic energy. This means that the daughter nuclei will cover a large momentum space with regard to the primary beam, so a high angular acceptance of the fragment separator is beneficial.

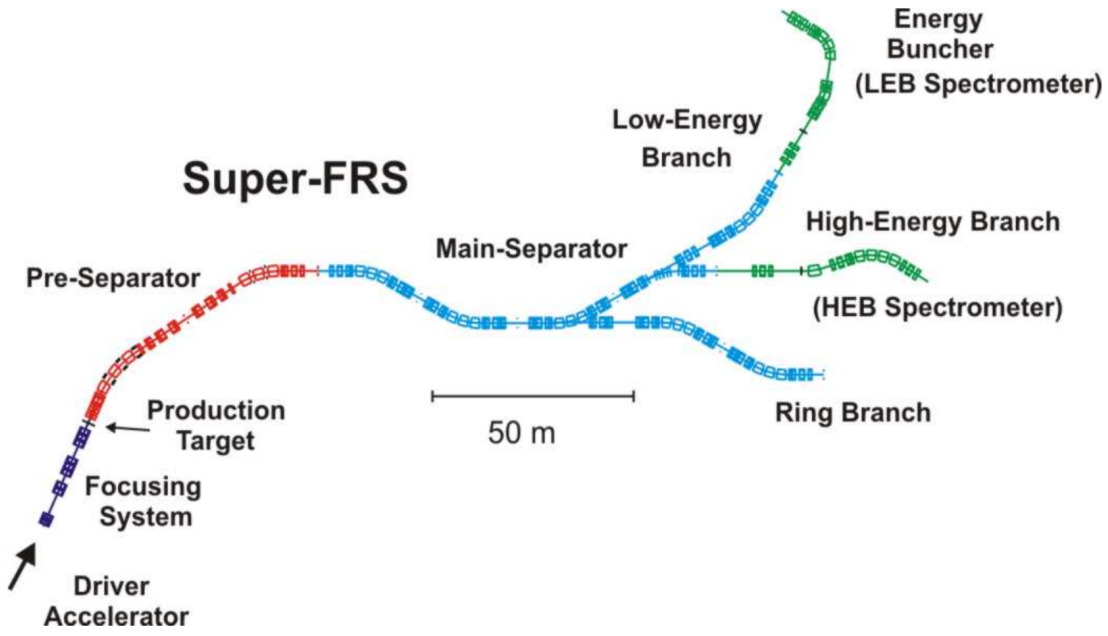

Figure 2: The super fragment separator and beam lines served by it, from [3]. The high energy branch output will lead to the R^3B experiment with the GLAD magnet right in front of the HEB spectrometer.

Fragmentation happens when the primary beam ion grazes a target nucleus. The ion will typically lose multiple nucleons and then evaporate some neutrons. Due to this, the fragment the fragment distribution on average will be more proton rich than the primary beam. Unlike with fission products, the momentum transfer to the fragment will be small.

As shown in figure 3, the production rates from these processes vary over many orders of magnitude. The main advantage of FAIR's SIS100 is the high beam intensity of up to 5×10^{11} uranium ions (at 1 Gev/u) per second.

After the production target, the super fragment seperator (SuperFRS [3]) will select only a few isotopes at a time out of the abundance being produced and forward them to the experiments, such as R^3B. This works by first using a dipole magnet to select nuclei based on their magnetic rigidity:

$$B\rho = \frac{p}{Ze} \sim \frac{\beta\gamma A}{Z} \tag{1}$$

That alone is not sufficient separation because ions with equal $\frac{A}{Z}$ and equal velocity will be deflected in the same direction, as will some ions with different charge to mass ratios and different velocities.

Then the beam is subjected to Bethe-Bloch energy loss in a degrader which is in first order described by

$$\Delta E \sim -\frac{Z^2}{\beta^2} \tag{2}$$

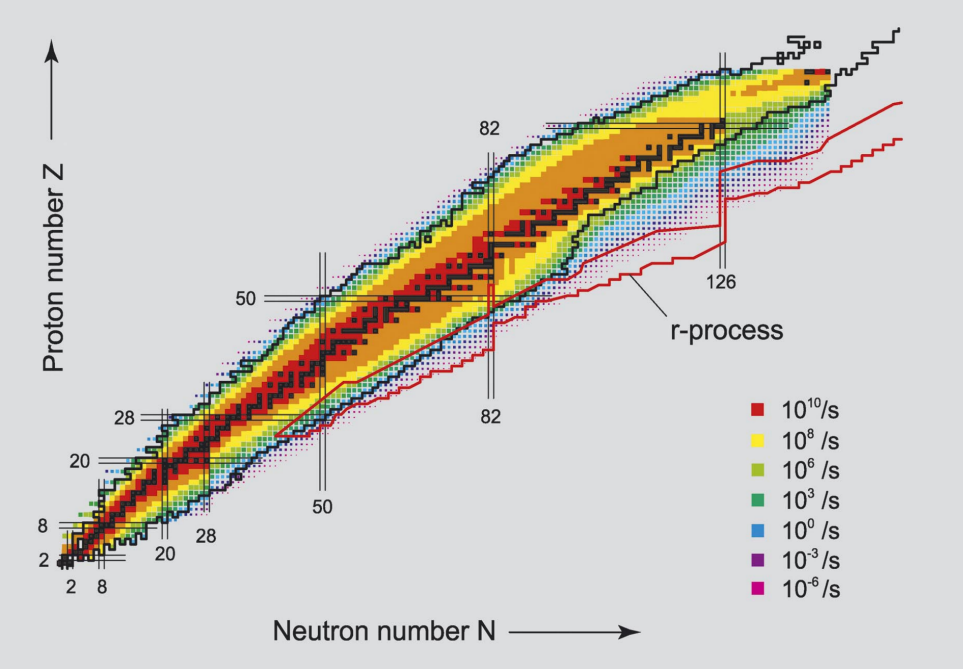

Figure 3: From [4], the estimated production rate of nuclei. The black outline corresponds to the known nuclei. The numbers indicate shell closures. The estimated pathway of the r-process is marked in red. After neutron shell closures, any further neutron captured will be weakly bound and easily removed by one of the ubiquitous photons in a (γ, n) reaction.

. This will lead to a momentum change which is in first order:

$$\Delta p \sim -\frac{E\Delta E}{p} \sim -\frac{EZ^2}{p\beta^2} \tag{3}$$

To compensate chromatic effects in the beam optics, it can be beneficial to use a wedge-shaped degrader instead of one with an uniform thickness.

Finally the filtering by $B\rho$ is repeated – this time selecting on $\frac{p-\Delta p}{Z}$. This will select a cocktail of just a handful of different isotopes – all with similar momentum – out of the thousands of isotopes produced in the production target.

The SuperFRS differs from the GSI FRS [5] in that it can repeat the $B\rho$-ΔE-$B\rho$ step twice – the beam first passes through a preseparator and then through the main separator. This setup is better suited to handle the high intensities provided by SIS100 as the pre-separator allows for the use of in-beam detectors in the subsequent stages to provide event by event PID at such high primary beam intensities.

2.2 Reaction studies with R^{3}B

Having produced exotic nuclei with a reaction target and selected a particular region of the chart using SuperFRS, one can now subject them to different **reactions** to study different properties

of the nuclei.

For collective mode properties, such as dipole resonances, **Coulomb excitation** is used. This involves an exchange of a photon between the ion and a target nucleus which leads to excited states which can either decay by emitting a cascade of γ rays or – for very large exictation energies – emit single nucleons. This type of reaction can be used to study the neutron skin thickness which is important for the equation of state.

To probe single particle properties, **knock-out reactions** can be used. Here, bound nucleons are kicked out of the ion by some target nucleon. Of particular interest is the $(p, 2p)$ reaction in which a free proton in a hydrogen-rich target knocks out a bound proton of the ion, both of which are coincidently detected at large scattering angles due to the large momentum transfer.

Within a nucleus, multiple nuclei can bind together and form clusters. Ikeda argued in 1968 that many nuclei exhibit states in which can be described as one or more α quasi-particles close to the emission threshold interacting with the remaining nucleus. More recently, the Ikeda conjecture [6] has been generalized to clusters of neutrons for very neutron-rich nuclei [7]. For example the ground state of the neutron-rich isotope ^{11}Li is a ^{9}Li core with a two-neutron halo. Unbound states such as ^{26}O also show a resonances near the two neutron emission energy.

The reactions with relativistic, radioactive beams (R^3B) experiment will offer a versatile setup a **kinematically complete measurement** of both the incoming ion as well as any reaction products (including outgoing ions, neutrons, γ rays and both high- and low-momentum protons) for all nuclear reactions. This will allow to reconstruct the momentum transfer to the heavy ion.

Given conservation of momentum and energy and baryon number, a kinematically complete measurement will actually offer redundant information. This is a key feature to select different reaction channels by using different kinematical constraints.

In the Reactions with Relativistic Radioactive Beams (R^3B) experiment, exotic ions will hit the target at a high momentum (e.g. a relativistic velocity $\beta = 0.82$, kinetic energy per nucleon $E/A = 700\,\mathrm{MeV/u}$).

R^3B is a typical magnetic dipole spectrometer for fixed target experiments. Due to the high beam energies all beam like reaction products are focused in forward direction while the target like reaction products are scattered to larger lab angles. The goal of a kinematically complete measurement necessitates a special detector setup discussed in section 2.3.

2.3 R^3B: a versatile detectors setup

Figure 4 gives an overview of some of the possible reaction products which can be produced by a reaction channel of interest as discussed in the previous subsection.

Generally the reaction products can be separated into two classes. **Beam-like fragments** experience only a small momentum change in the reaction and thus roughly travel in beam direction. They are subsequently separated by their magnetic rigidity in the spectroscopic magnet GLAD.

By contrast, **target-like particles** and photons – though still boosted forward in the lab frame due to the beam momentum – will be emitted under much higher angles in the lab frame.

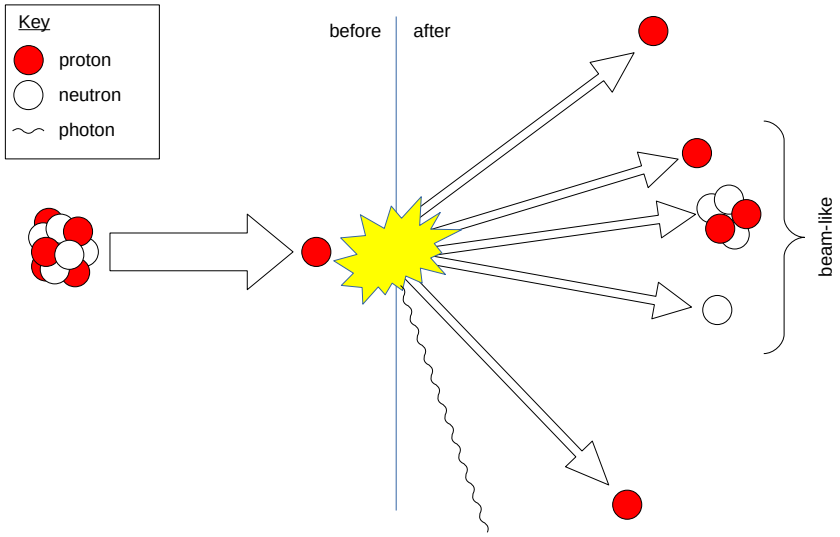

Figure 4: A cartoon of a reaction which might happen in R^{3}B. An incoming ion (produced by the SuperFRS, entering from the left) interacts with a hydrogen nucleus in a target. Two target-like protons are produced. The residual nucleus might deexcite by ejecting neutrons or γ rays.

The goal of R^{3}B is the kinematically complete measurement of nuclear reactions in inverse kinematics. To achieve this, a high acceptance for all the reaction products is required.

As will be shown in figure 5 this requires a dedicated detector setup for the individual tasks. Upstream of the target the particle identification (PID) of the incoming beam is essential to select particles of interest for the reaction. For this task the time of flight (TOF) between the particle in a focal planes of the fragment separator and directly in front of the target is measured, as well as the position at both locations.

For the PID and high resolution momentum measurement of all the reaction products various detector arrays are grouped in different arm configuration around the large dipole magnet GLAD. While the neutrons from the reactions are not deflected in the magnet they require a minimal material budget for all in-beam detectors upstream of GLAD. The Neuland Detector at a distance of 15 m downstream covers the full opening of ± 80 mrad of GLAD to register all high energy neutrons from the reaction with the maximum efficiency.

At the largest bending angles the proton arm registers beam like protons emitted in the reaction.

The most difficult task – especially in case of very heavy ions – is a full mass and charge identification of the heavy residuals. High resolution tracking before and after the magnetic field has to provide absolute positions and angles of the particles passing through the magnetic field. Additionally a large area TOF wall also provides energy loss information for the particles.

The most complex detector configuration is arranged around the target. Here the specific requirements of the different experiments set the stage for very different detector configurations.

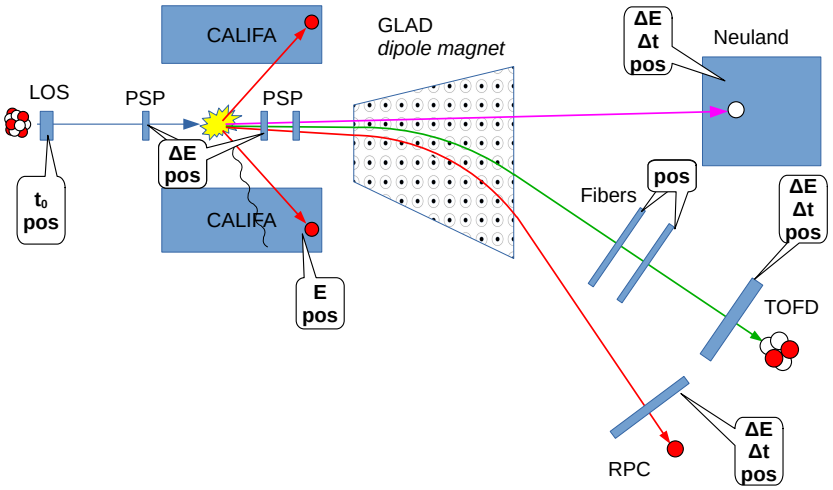

Figure 5: The reaction products from figure 4 are detected with an array of different detectors. Different quantities such as total energy, energy loss, time of flight and position are measured to create a kinematically complete measurement of the reaction.

2.4 Properties of the different detector systems

This goal informs the choice of detectors both before and after the target:

- The incoming exotic ion is produced in the production target of the FRS/SuperFRS, which will forward multiple isotopes with a range of momenta to R^{3}B. Upstream of the reaction target, the precise trajectory of the **incoming** ion, its charge and mass need to be determined. Multiple detectors exist to provide this information.

- At the **reaction target**, the energy and θ of photons emitted by excited nuclei need to be measured. Target-like reaction products (that is to say, reaction products which do not have a momentum similar to the incoming beam) must both be tracked and their energy must be measured. The need to measure photon energies precludes time of flight measurement. Semiconductor detectors like high purity germanium (HPGe) are unsuitable because they do not provide cost-effective stopping power for protons, and would in fact be damaged by them. Thus, CALIFA is a scintillation calorimeter.

- After the target the trajectory, energy, charge and momentum of the residual nucleus after the interaction have to be quantified. Furthermore, beam-like reaction products (those with a similar momentum as the beam) such as protons or evaporated neutrons will also need to be detected. Collectively, all of these reaction products are also called **fragments**. The requirements to determine the charge and mass of these fragments, as well as the need to measure the energy of neutrons makes calorimeters unsuitable here. Instead, a large dipole magnet spits the outgoing fragments according to their $\frac{Z}{A}$ while further charge sensitive tracking detectors provide A. The fragment energy is determined using the time of flight method.

Figure 5 gives an overview of the main detectors of R^3B and what quantities they measure. **Positions** are either measured with separate detectors (e.g. multiple crystals, bars or strips) or by reading out a strip or bar from both ends and comparing signal amplitudes or timing.

Energies are measured using either scintillators or semiconductor detectors. In the beam line, they generally measure the energy loss ΔE which is is proportional to Z^2 at a given β. Thus, ΔE can be used to determine Z.

The time of flight Δt is the time difference between two fast scintillation detectors triggering. As the distance between both of them is known, this allows determining the velocity $v = \frac{\Delta s}{\Delta t}$.

Start time for time of flight detectors

The beam will then enter the R^3B cave with a typical kinetic energy of some 700 MeV per nucleon, which corresponds to $\beta = 0.82$. Before hitting the target, the ion will travel through a thin organic scintillator called LOS[5] – which will generate a precise start point for time of flight measurements – and possibly though some adjustable active veto colimators called ROLU[6].

In-beam tracking and charge determination

Both for the incoming ion and any outgoing fragments, it is essential to determine their charge and mass. This works by a ΔE measurement using Bethe-Bloch and observing $B\rho$ by tracking the particles trajectory through a magnet.

For the ion before the target $B\rho$ is given by the SuperFRS. After the target, the magnetic rigidity of the outgoing ion (or ions) is measured using the GLAD magnet.

Different choices exist for position and charge detection of in-beam particles. [8] The **P**osition **S**ensitive **P**in diode (PSP) is a semiconductor segmented into strips along one direction on the front side and along the other direction on the back side. Each strip is read out from both ends, and the ratio of charges from both ends can determine the exact position of the hit on that strip. The ion charge can be detected by measuring ΔE, e.g. summing the charges detected on both ends of the strip. However, unambigious charge assignment is still tricky to do and typically involves multiple detector layers.

Another approach here are gaseous detectors such as the **MU**lti **S**ampling **I**onization **C**hamber. Compared to semiconductor detectors, these detectors typically sport a larger active area. Charge is again determined via ΔE by measuring the total ionization charge observed. Position along one axis can be inferred by observing the drift times in the gas.

Various fiber detectors are used downstream of GLAD to measure the tracks of the outgoing ions.

Target

The target can either be a conventional solid target (typically as part of a target wheel) such as Pb or CH_2. Also, a cyrogenic liquid hydrogen (LH2) target is available.

[5] Probably from German *los*, meaning start.

[6] From German *rechts, oben, links, unten* meaning right, above, left, below, denoting the four directions of the active vetos with relation to the beam axis.

The advantage of the target wheel is that the target in path of the beam can quickly be switched remotely without opening the scattering chamber. While conventional targets can measure the cross section for proton reactions by using both CH_2 and a carbon target and then subtracting the results of both to get the reactions caused by protons, the liquid hydrogen target allows to measure proton reactions without having to do such a background subtraction.

In general, different target types are preferred for different reaction channels. The virtual photon spectrum used in coulomb excitations strongly favors high Z targets, while proton induced quasi free reaction require a proton target.

Target area calorimeter: CALIFA

When the nucleus undergoes a reaction at the target, not only is the outgoing nucleus (possibly) changed in composition, energy and momentum, but it might also emit both nucleons (and fragments) and photons. An inorganic scintillator calorimeter called CALIFA [9, 10] surrounds the target area. The focus of this work is the detection of γ rays in this detector. This detector is discussed in much more detail in the following section 3.

GLAD magnet

For R^3B, GSI has commissioned a new superconducting large acceptance dipole magnet aptly called the **G**SI **L**arge **A**cceptance **D**ipole. [11]

Dipole magnets can be classified using the integral of the field B along the trajectory of a typical particle. The GLAD magnet has a maximum field integral of 4.8 Tm.

For $(B\rho)_{Magnet} \ll (B\rho)_{Ion}$, the bending angle α (in radians) of an ion in a dipole magnet can be approximated using:

$$\alpha = \frac{(B\rho)_{Magnet}}{(B\rho)_{Ion}} \tag{4}$$

The super fragment separator (SuperFRS) will provide beams with a maximum rigidity of $(B\rho)_{Ion} = 20\,\text{Tm}$. Thus, such ions will be bent by an angle of around 14°.

The large acceptance of GLAD is especially important for neutrons. GLAD will not impede the traversal of neutrons with an angle of less than 80 mrad (4.6°).

Time of flight wall

Downstream of the GLAD magnet and in-beam tracking detectors, the ions will hit the plastic scintillators time-of-flight wall. [8] As the beam widens after the target, this detector has an active area of 1.2m × 0.8m to catch all the outgoing ions.

Of course, the point of the ToF measurement is not to confirm that yes, we actually have a beam-like particle (e.g. $\beta \approx 0.8$), but to determine the *precise* velocity of the particle. The requirement here is $\frac{\sigma_t}{t} < 2.5 \times 10^{-4}$, which means $\sigma_t < 20\,\text{ps}$. The energy deposit ΔE within the scintillator is measured using the time over threshold method, which can be used to determine the charge.

NeuLAND

NeuLAND [12] consists of $2.5\,\mathrm{m} \times 2.5\,\mathrm{m} \times 3\,\mathrm{m}$ of plastic scintillator in 3000 bar-shaped submodules read out on both ends with photomultiplier tubes. A neutron entering NeuLAND has a 95% chance of scattering on at least one hydrogen nucleus – which can then be detected.

Detection alone is not sufficient, however: neutrons evaporating from the ion after a nuclear reaction with the target will have roughly the same momentum per nucleon as the ion. Naturally, this miniscule difference in momentum – which stems from the momentum the neutron carries in the ions rest frame – is what is interesting.

The transverse momentum of the neutron (in the ion frame) can be determined geometrically by extrapolating the trajectory of the ion (before GLAD) to NeuLAND. If the neutron had a transverse momentum, it will not have hit NeuLAND some distance away from that trajectory. The ratio of neutron transverse momentum to beam momentum is equal to the ratio of that distance to the distance of the NeuLAND detector from the target.

The longitudinal (that is, in-beam-direction) momentum of the neutron is measured using the time of flight method.

Many nuclear reactions will also result in the emission of multiple neutrons which have to be detected. As each neutron might scatter on a proton more than once, this is a lot harder than detecting a single neutron. For a four neutron event, the reconstruction efficiency is still 60%. As the aim of R^3B is a kinematically complete measurement, if a neutron is missing, the mass before and after the reaction will not be the same, so these events can be filtered out.

3 The CALIFA calorimeter

For most of the experiments planned for R^3B a full reconstruction of the reaction is essential. This means that the residual heavy fragment has to be fully identified in mass and charge. In Addition the different states of this nucleus populated in the reaction have to be identified by their unique fingerprint from γ rays emitted after the reaction.

As the residual nucleus moves at a relativistic velocity, the γ rays emitted are subject to the Doppler effect. The consequences of the Doppler effect strongly influence the design of CALIFA, so to understand why CALIFA is designed the way it is, understanding these consequences is essential.

After discussing the resulting detector geometry, we will travel along the "data acquisition food chain" to understand how an energy deposit in a crystal is turned into numbers on a hard disk.

3.1 The relativistic Doppler effect

Photons emitted from a nucleus traveling at a relativistic speed[7] β in positive z direction will be boosted forward following the relativistic Doppler effect, a special case of the Lorentz boost. The energy and polar angle are transformed like this:

[7]By convention, β is the velocity as a fraction of the speed of light and $\gamma = \frac{1}{\sqrt{1-\beta^2}}$. For the general case of the Lorentz boost, see section D.

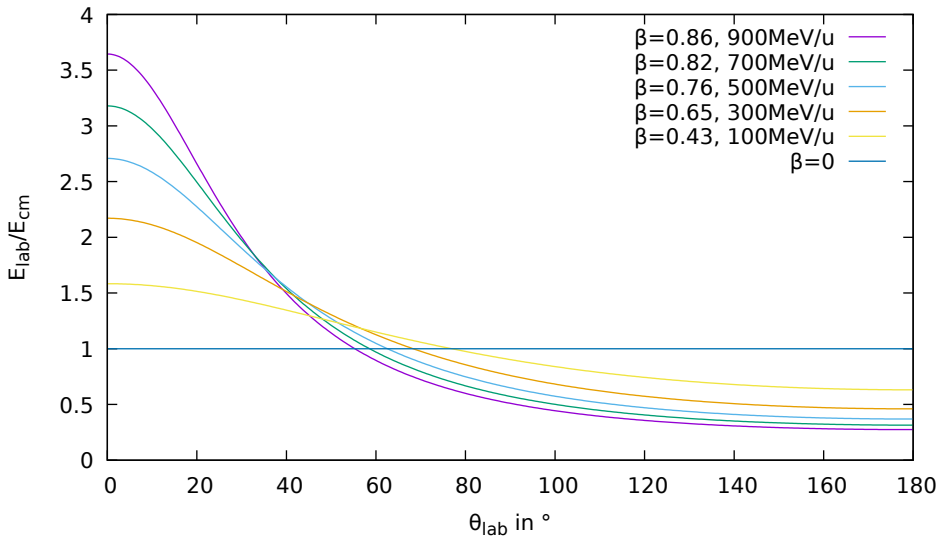

Figure 6: The variation of γ energy due to the Doppler effect as a function of the polar angle of the γ ray detected in the lab frame.

$$E_{lab} = \gamma(1 + \beta \cos\theta_{cm}) E_{cm} \quad (5)$$

$$\tan\theta_{lab} = \frac{\sin\theta_{cm}}{\gamma(cos\theta_{cm} + \beta)} \quad (6)$$

The Doppler effect means that the gamma energy mesured in the lab frame is different from the gamma energy as emitted from the ion. Figure 6 gives the factor as a function of θ_{lab} for different ion energies. The γ energy in the lab frame of reference, as a function of the polar angle in that frame, is

$$E_{cm} = \gamma(1 - \beta \cos\theta_{lab}) E_{lab} \quad (7)$$

This has a huge impact on the design of CALIFA. Different parts of the detector have completely different design requirements. If the γ ray energy is 3 MeV in the ion frame of reference, it can easily reach 10 MeV in the forward CEPA part of CALIFA. As will be shown in figure 31, the dominant interaction of γ rays with such energies is e^+e^- pair production. Electrons and positrons will then emit Bremsstrahlung, possibly creating secondary clusters. By contrast, in the backwards part of the barrel, the dominant processes will be Compton scattering and photoeffect.

Also, to reconstruct the gamma energy in the ion frame of reference, it is not enough to measure the γ ray energy, the polar angle also has to be well defined. Figure 7 shows the dependence of the relative reconstructed γ ray energy resolution $\Delta E_{cm}/E_{cm}$ on the angular resolution $\Delta\theta$. Thus, CALIFA needs a high granularity with respect to the polar angle, especially in the iPhos region around $\theta \approx 30°$. The effects of the angular uncertainty can not be completely

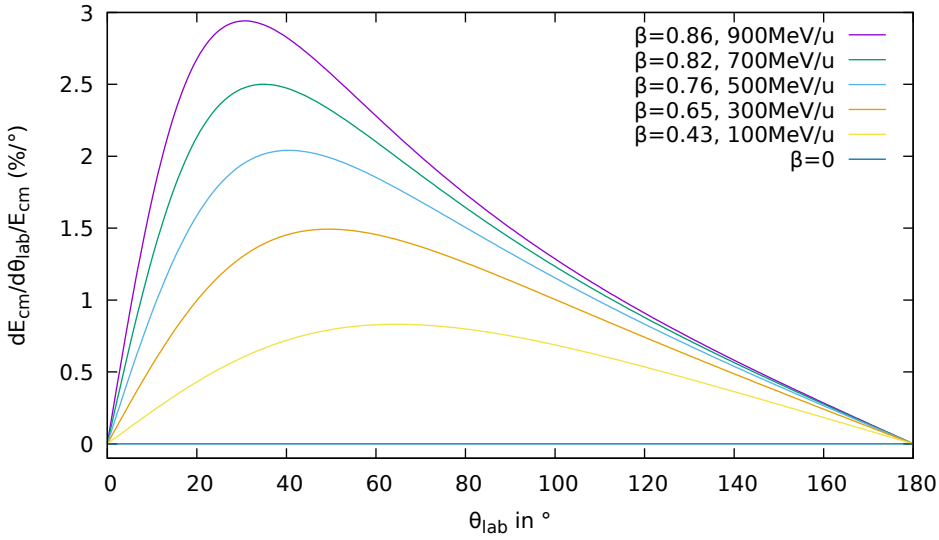

Figure 7: The impact of the angular resolution on the energy resolution of the photons in the center of mass frame.

compensated. The barrel shape puts that critical section further away from the target than a spherical detector arrangement would, thus giving a longer lever arm for measuring angles, but a longer barrel would take more of the space between target and GLAD, which also has to fit in-beam tracking and particle identification detectors, and increasing the distance between target and GLAD will lower the angular acceptance. Within CALIFA, there are also mechanical constraints as well as the fact that light collection would suffer if the iPhos crystals were replaced by multiple thinner crystals. [10]

Figure 8 gives the reduction of θ between both frames of reference as a function of θ_{cm}. Obviously, if the photon travels directly in or against z direction in the ion frame ($\theta = 0$, $\theta = \pi$), the it will also travel in that direction in the lab frame as $\sin\theta = 0$. Around $\theta = 120°$, the shift in θ is highest.

This result has a huge effect on the transformation of solid angles. Assuming that the photons are emitted isotropically in the ion frame, one can calculate the anisotropy in the lab frame. Integrating over the azimuthal angle, the solid angle is generally:

$$d\Omega = 2\pi \sin\theta d\theta \tag{8}$$

Thus, the ratio of the differential solid angles is:

$$\frac{d\Omega_{cm}}{d\Omega_{Lab}} = \frac{\sin\theta_{cm} d\theta_{cm}}{\sin\theta_{Lab} d\theta_{Lab}} \tag{9}$$

$$= \frac{p_{\perp,cm}}{E_{cm}} \frac{E_{Lab}}{p_{\perp,Lab}} \frac{d\theta_{cm}}{d\theta_{Lab}} \quad \Big| \; p_{\perp,cm} = p_{\perp,Lab} \tag{10}$$

$$= \frac{E_{Lab}}{E_{cm}} \frac{d\theta_{cm}}{d\theta_{Lab}} \tag{11}$$

$$= \frac{1}{\gamma^2} \times \frac{1}{(1 - \beta\cos\theta_{lab})^2} \tag{12}$$

Figure 9 visualizes the effect of the Doppler effect on the distribution of photons[8]. For $\beta = 0.82$, the rate of γ rays detected will be very skewed towards small values θ – the chances to see a one in forward detection are a hundred times higher than seeing one in backwards direction. This means that the rate of γ rays – and also light particles – will be vastly higher in the forward-facing part of the detector. Thus, different technologies are appropriate for that critical part and the more relaxed parts of CALIFA covering higher θ angles.

In conclusion, the requirements for γ ray detection depend strongly on the angle θ_{Lab} between the detected photon and the momentum direction of the ion. As the emission is boosted towards low angles, the forward-facing part of the detector is critical and will have to handle a high rate, while the backward-facing part is more of an afterthought as it will see only a small fraction of photons in any case. To undo the Lorentz boost, one has to know the detection angle of the γ ray. This necessitates a high granularity, especially around the $\theta_{lab} \approx 30°$ region. Finally, one has to handle very different photon energies, in forward direction the observed energies may be boosted by a factor of three compared to the ion frame, while in the backward region they can be only a third of the CMS γ ray energy.

[8] Which are assumed to be emitted isotropically in the ion frame of reference.

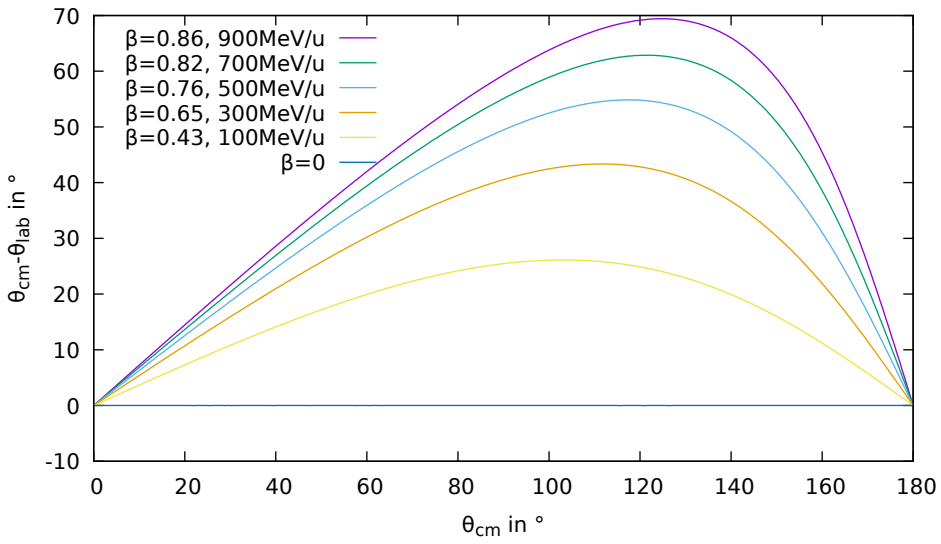

Figure 8: The reduction of the polar emission angle of γ rays due to the kinematic boost.

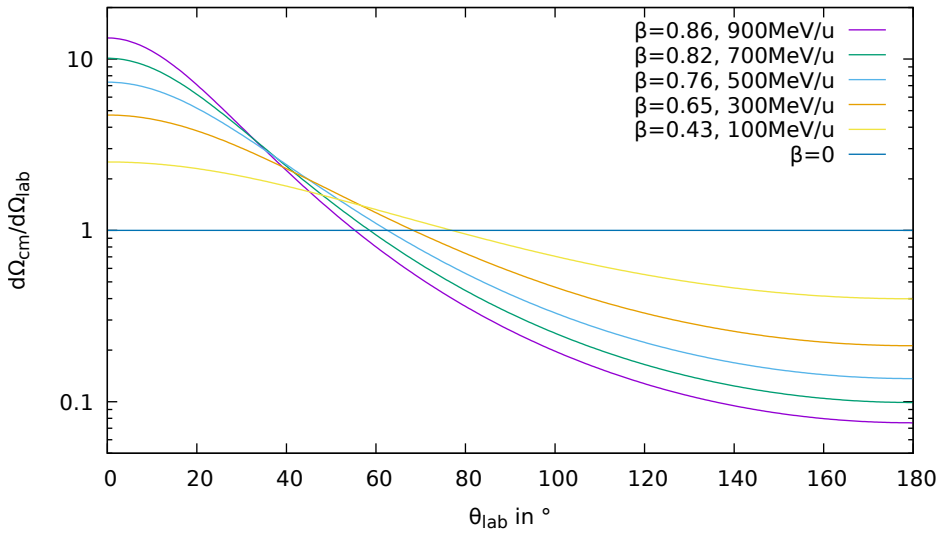

Figure 9: The effects of the Doppler effect on the solid angle. For $\beta = 0.82$, the region near $\theta_{Lab} = 0$ will see ten times the fraction of photons it would see of photons isotropically emitted in the lab system. By contrast, near $\theta = 180°$, the rate will seen by a detector would be ten times lower than what would be expected from lab-isotropically emitted photons.

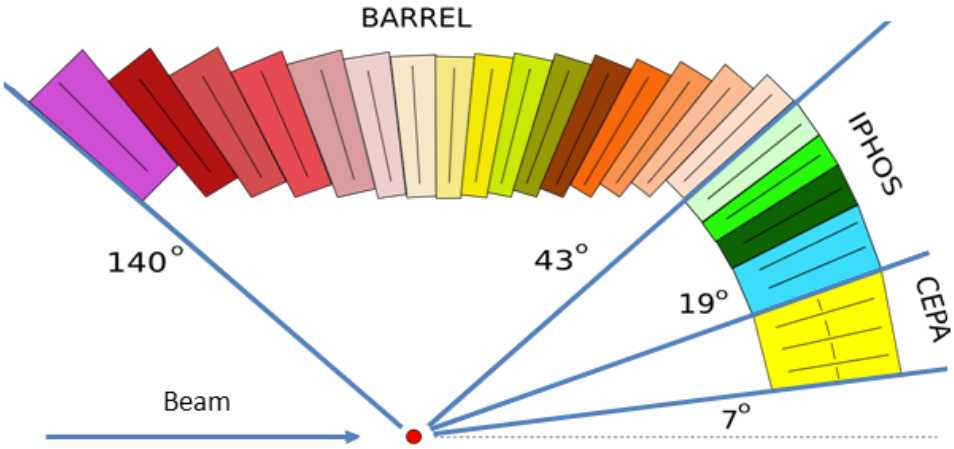

Figure 10: From [10], a visualization of the geometry of CALIFA. Image was modified by author.

3.2 The CALIFA geometry

The geometry of CALIFA [9, 10] is strongly influenced by the Doppler effect discussed in section 3.1. To reiterate:

- The Doppler shift depends on the polar angle of the photon. Hence to correct for the shift, that polar angle has to be measured accurately.
- The Doppler effect changes the angular distribution of the γ rays, strongly pushing photons towards the beam axis in the laboratory frame.
- The forward part will see a much higher rate and much higher energies than the rest.

To determine the polar angle of γ rays accurately, the calorimeter needs a high granularity. As most photons (and particles) are expected to hit CALIFA at small polar angles, different technologies are appropriate in different regions.

For maximum detection efficiency, CALIFA has to cover the target area hermetically. To allow access to the target area (which is most densely populated with detectors), CALIFA is splitable into two halves[9] mounted on a hanging rail system which allows movement of the halves in the XZ plane. The electronics racks are also mounted on the same rail system so that the whole half can be opened as one compact unit and there is no stress on the cables connecting the electronics to the detector during movements.

As illustrated in figure 10, CALIFA is split into three sections:

- The **CALIFA Barrel** covers the region from $\theta = 43°$ up to 140°.
- The **Intrinsic Phoswich** (IPHOS, IPhos) region covers the region from $\theta = 19°$ to $\theta = 43°$.

[9]The halves are commonly referred to as *Messel* and *Wixhausen*.

- The **CALIFA Endcap Phoswich Array** (CEPA) covers the forward region of $\theta = 7°$ to $19°$.

In 2019, the **CALIFA demonstrator** was used. It consists of seven mechanically separate *petals*, each containing a section of 64 barrel crystals, $\Delta\phi = 22.5°$.

At the end of 2019, the forward part ($\theta \in [43°, 90°]$) of the barrel in its fully hermetic 2π coverage in azimuthal angles was equipped with 1024 crystals.

The iPhos region covering the angles $\theta \in [43°, 90°]$ was completed in 2021 with 480 crystals.

The Barrel section

The Barrel section of CALIFA features CsI(Tl) 1952 crystals of six different shapes arranged in a barrel with an inner diameter of sixty centimeters. The aim is to fully absorb most of the γ rays emitted in the reaction while also determining the energies and angles of protons emitted. As of 2022, the Barrel is complete for $< 42°\theta < 90°$, while preparations are ongoing to fill the less occupied pockets in the backward direction.

The IPhos section

The IPhos section consists of 512 CsI(Tl) crystals, e.g. the same material as the barrel. Unlike the barrel crystals, however, these crystals will not be long enough to stop all protons – which will have a higher energy than in the barrel region due to the Lorentz boost.

Thus, the energies of protons will be determined by measuring ΔE and extrapolating to the full energy.

Of course, this poses a challenge: one has to determine for a given hit whether the proton was stopped – and the observed energy is the full energy – or if it punched through the detector, and the full energy has to be extrapolated.

As discussed in [13], CsI(Tl) has two different scintillation states with two different time components. The ratio of those two will be different for stopped and punch-through ions. This allows to decide if the energy has to be calculated as $E(\Delta E)$ or $E = \Delta E$.

The CEPA section

At lower θ angles, the rate of both photons and particles in the detector will increase rapidly due to the Lorentz boost. This means that slow scintillators such as CsI are no longer the optimal choice. In addition, γ energies are also boosted due to the Doppler effect (see section 3.1). As a consequence the ratio of interactions governed by e^+e^- pair creation will increase which favors a detector concept with a segmentation in depth to better select fully contained interaction clusters. As an advantage of the Doppler boost, the derivative of the kinematical energy shift for γ rays is strongly reduced at small polar angles (see figure 7) $LaBr_3(Ce)$ and $LaCl_3(Ce)$ have a similar light yield to CsI, but a much shorter decay time of 16 ns and 28 ns. For the CEPA section, a 7cm thick $LaBr_3(Ce)$ crystal and an 8cm thick $LaCl_3(Ce)$ crystal are glued together with an optical glue and then read out using a single photomultiplier tube (PMT). The signals from both crystal types can then be separated by pulse form analysis in a setup commonly called a phoswich. For a charged particle which punches through the detector, the

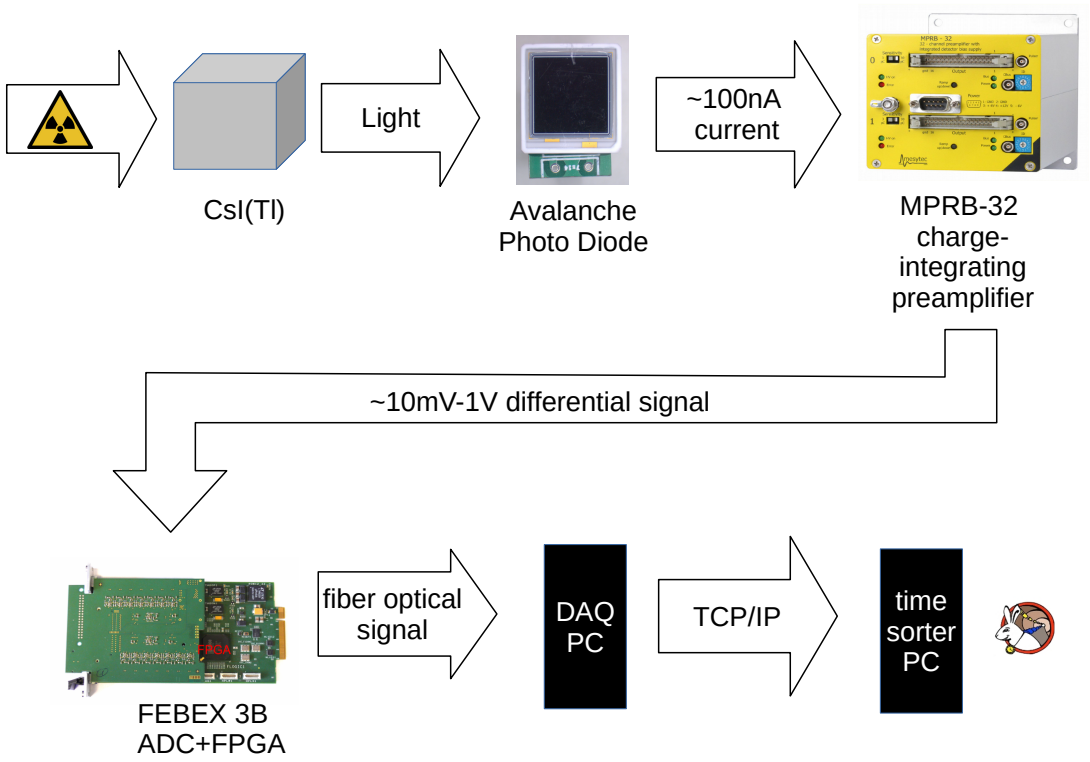

Figure 11: The "food chain" which turns energy deposits into data files.

detector will measure both the energy loss in 7cm of $LaBr_3(Ce)$ ΔE_1 and in 8cm of $LaCl_3(Ce)$ ΔE_2. This allows the reconstruction of $E = E(\Delta E_1, \Delta E_2)$. [10]

The CEPA detector was not yet part of any setup discussed in this thesis.

3.3 Signal processing overview

Figure 11 gives an overview of the signal processing steps involved in CALIFA. Ionising radiation deposited in the CsI(Tl) crystals excites states in that scintillator which decay with the emission of photons. These photons are collected and amplified in an avalanche photo diode (APD). The preamplifier both provides individual bias voltages to the APDs and integrates and amplifies their photocurrent. The resulting signal is then sent (through some more analogue processing) to an sampling ADC. The ADC is read out continuously by an FPGA which will record interesting events and do a pulse shape analysis on them. The resulting data is then collected (either synchronously or asynchronously) by an PCIe card called KINPEX over a proprietary protocol called GOSIP via fiber-optical link. From there, it is (eventually) sent to the time orderer, where the CALIFA hits are merged with events recorded by other detectors based on their white rabbit time stamp.

In the following subsections, these steps will be discussed in detail.

3.4 Crystals

CsI(Tl): an inorganic scintillator

Scintillators are materials that, when exposed to ionizing radiation, emit lower energy photons, e.g. light. They can be separated into two categories: organic scintillators have organic base which typically contain a aromatic rings, and often fluors which will shift the wavelength of the emitted photons. Inorganic scintillators are crystals. Most also contain a dopant which acts as a fluorescent ion – similar to the flours found in organic scintillators.

CsI(Tl) means Cesium Iodide doped with trace amounts of Thallium[10].

CsI(Tl) is an inorganic scintillator with a density ρ of 4.51 $\frac{\text{kg}}{\text{dm}^3}$ – not quite as high as some others (e.g. $PbWO_4$4, $\rho = 8.3\,\frac{\text{kg}}{\text{dm}^3}$), but much higher than organic scintillators ($\rho \leq 1.2\,\frac{\text{kg}}{\text{dm}^3}$). The high density is also the reason for its high stopping power, which is of course important in calorimeter applications where the point is to stop the particle and the amount of space the detector may occupy is limited.

CsI(Tl) is amongst the scintillators with the highest light yields, which means that it provides a large number of photons N_{ph} per unit of energy deposited by ionizing radiation. This is important because the number of photons observed will vary according to some Possion distribution, which means that its variation will scale like $\sqrt{N_{ph}}$. The relative resolution thus scales with $\frac{1}{\sqrt{N_{ph}}}$. [14]

Like many other inorganic scintillators, CsI(Tl) is hygroscopic, but only slightly. This makes it easier to handle than strongly hygroscopic scintillators like $LaBr_3$ or $LaCl_3$.

The light emission of CsI(Tl) peaks at 550 nm. Together with its high light yield, this makes it suitable to be read out by inexpensive avalanche photodiodes (APD).

A speciality of CsI is the existence of two different excited states which will emit scintillation photons, with decay constants of $\tau_f = 600\,\text{ns}$ and $\tau_s = 3.4\,\mu\text{s}$. The total signal can thus be split into a slow and a fast component. For photons, electrons and unstopped ions, these two types of states are both excited proportionally to the energy deposit. A third excited state with a decay constant of 16 μs exists but contributes 15% of the total photons [15] and is not captured by the FEBEX firmware.

Shortly before an ion is stopped, however, it will excite more of the slow component due to the higher ionization density. [16] This effect will scale with both A and Z.

When sampling the scintillator signal with a sufficient time resolution, one can use pulse form analysis methods to estimate the amount of both the fast and the slow component of the signal.

This has two applications:

- Different ion species can be distinguished (particle-ID).
- Stopped and punch-through ions can be discriminated.

[10] Generally, a Tl weight ratio of around 700 ppm seems to give the optimal light yield.

3.5 Avalanche photodiodes

While the classic way to collect scintillation photons involves photomultiplier tubes (PMTs), these tend to be expensive, fragile and bulky. Additionally, they are affected by magnetic fields – such as the stray field from GLAD. As CsI has a very high light yield, one does not actually need the high gain provided by PMTs to read out a CsI crystal especially for detecting sctinillation photons resulting from energy deposits above 50 MeV.

Photodiodes are semiconductor devices suitable for photodetection. A photon creates an electron-hole pair in the space charge region which will result in a current.

Special photodiodes called avalanche photodiodes (APD) allow for a typical gain of $g \approx 50$ already at moderate bias voltages ($U \approx 350\,\mathrm{V}$). Due to a special doping profile in the p-n junction a very large field gradient is created right behind the absorption layer. Charge carriers created by photons will be accelerated in this field and eventually create secondary charge carriers, resulting in an avalanche. This allows to multiply the initial photoelectron produced by the photon. Combined with the large quantum efficiency of the APDs this is sufficient to detect γ ray even below 100 keV in CsI.

3.6 The Preamplifiers

While APDs are powerful light detectors, they also have a few drawbacks. Unlike photomultipliers, their gain is limited to some two orders of magnitude, so the output is not suitable to be transported over long cables.

Due to production issues, each APD reaches a gain of 50 at an individual voltage U_{50}. This neccessitates individual voltages per APD. Also, the gain has a strong temperature dependence due to the mean free path of charge carriers in the space charge region. In addition large area APDs (such as used in CALIFA) have a substantial capacity ($C \approx 500\,\mathrm{nF}$) and charge current of several nA, both resulting in non-negligible contributions to the noise.

The MPRB-32 preamplifier from Mesytec was developed to solve all of these issues. It integrates an APD current signal $I(t)$ to extract the charge $Q(t) = \int I(t)dt$ and amplifies it to a differential current signal suitable for transport over twisted pair cables and measurement with ADCs at a $100\,\Omega$ differential termination. The MPRB-32 also features a temperature sensor and can adjust the bias voltage to keep the gain constant under different temperature conditions.

It can also generates a reverse bias voltage in the order of 400 V and allows to bias channels with individual voltages, thus accommodating the different characteristics of avalanche photo diodes.

For integrating preamplifiers, a general decision has to be made how the preamplifier output will return to zero. While reset preamplifiers will activate a veto and just reset the output voltage when they are at the end of their range, others such as the MPRB-32 have a resistor in parallel to the charge collecting capacitor, causing the signal to decrease exponentially with $\tau = RC$.

The selection of τ is subject to tradeoffs: A long τ will allow precise measurement even of slow rising signals, but limit the rate which the preamp can handle before running into overflow to $O(\tau^{-1})$. For CALIFA, $\tau = 35\,\mathrm{us}$ was selected.

3.7 The CALIFA readout system

The readout system of the CALIFA detector is based on the FEBEX 3B[11] 50 MHz sampling ADC board developed by the GSI EE department.

From the outputs of the MPRB-32s, 34 pin twisted pair cables (sixteen differential pairs and two ground connections) lead to the digitization electronics. Each pair carries the output of a preamplifier channel as an analogue differential signal.

The differential signal first arrives at a so-called AddOn board developed by TUM, where further analogue operations take place: A low-pass filter is applied to eliminate the signal components above the NYQUIST frequency to prevent aliasing. [17] Furthermore, a fixed, user-configurable offset can be added to the signal. This is helpful because the differential ADC has a symmetric range, but the output from the preamplifier only has one polarity. By adding an offset, one can use (almost) the full range of the ADC instead of just the half range, gaining

[11] https://www.gsi.de/work/forschung/experimentelektronik/digitalelektronik/digitalelektronik/module/font_end_module/febex/febex3b

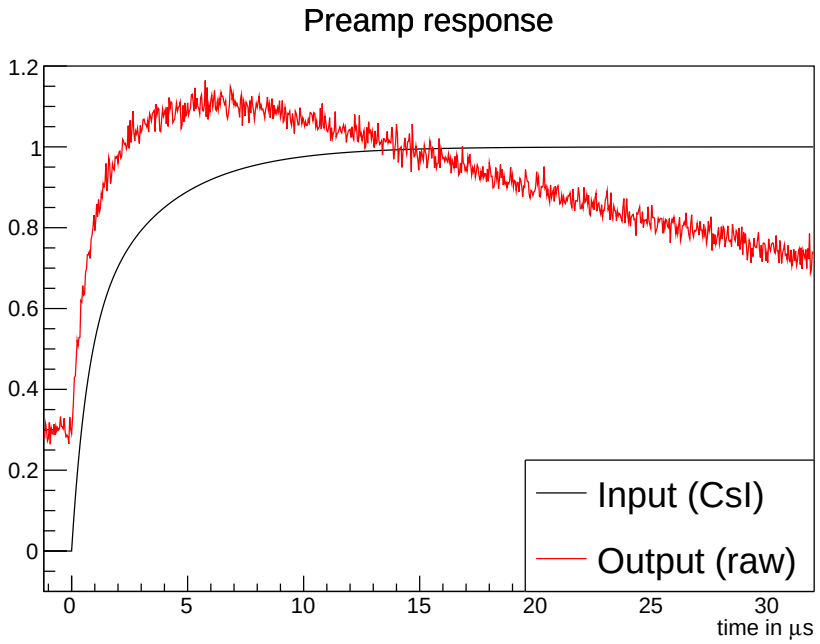

Figure 12: A model of an idealized CsI(Tl) of the integrated current $Q(t)$, and its approximation $\tilde{Q}(t)$ by the preamplifier, both in arbitrary units. In the model of the preamp, noise was added and the RC decay was implemented. An offset was also added to move the baseline from zero to a different value. *In reality, the moving of the baseline is a feature of the Febex Addon Board (which is not explicitly simulated).* Both input and output signals are in arbitrary units. In reality, the incoming current signal might peak around ten μA, giving some 30 pC after integration and resulting in an amplitude of the output signal on the order of 100 mV (over 100Ω).

one bit.

The FEBEX3B board contains two eight channel ADC chips which sample each incoming waveform at $f = 50\,\mathrm{MHz}$. The ADCs are read out by an FPGA (Lattice ECP3).

The common habitat of FEBEX boards are FEBEX crates which fit into standard 19 inch racks commonly used both in nuclear physics and IT.

Besides a 12V supply, FEBEX crates provide a **trigger bus** as well as a daisy-chain topology **2.5Gbit LVDS link**. In the leftmost slot of the FEBEX crate, a special module housing a small form-factor pluggable tranceiver (SFP, a commodity computer network component) converts the LVDS signal into an optical signal connected to the KINPEX1[12] (or PEXOR[13]) PCIe x4 card of a special PC via fiber optical cable. The leftmost card also links the trigger bus to the EXPLODER (see section 5.4) using multiple endpoint low voltage differential signaling (MLVDS).

At CALIFA, each FEBEX crate houses 16 FEBEX cards handling a total of 256 preamplifier channels. As the crates can house 19 cards, a spare card is commonly used to sample the event trigger received from the Main DAQ.

The FEBEX 3B modules of CALIFA are running a custom firmware developed by Max Winkel [16]. Understanding the intricacies of the signal processing is crucial for the latter analysis, so section 4 is decicated to the firmware.

Both sides of CALIFA feature a PC equipped with a GSI KINPEX fiberoptical transceiver card. This card can read out up to four FEBEX crates (256 signal channels). On each side, one KINPEX card is used to read out 256 barrel crystals (with single range preamplifiers) and 128 iPhos channels (with dual range preamplifiers, for a subtotal of 256 channels).

After some further processing, the data is finally sent to the DRASI[14] time orderer, where the streams from both sides are merged with hits from other detectors based on their white rabbit time stamp.

[12] https://www.gsi.de/work/forschung/experimentelektronik/digitalelektronik/digitalelektronik/module/pci_pci_e/kinpex1

[13] https://www.gsi.de/work/forschung/experimentelektronik/digitalelektronik/digitalelektronik/module/pci_pci_e/pexor/pexor3

[14] http://fy.chalmers.se/~f96hajo/drasi/

4 FEBEX FPGA Firmware

The CALIFA Firmware [17, 16] for the FEBEX 3B uses a real time multi layer parallel processing to to collect data, pre-analyze, provide triggers and timing information and interfacing the GOSIP readout.

The firmware was designed to allow channels to self-trigger or act on an external trigger. Providing a timestamped multi-event readout, it avoids the need for a common dead time. It offers high resolution amplitude measurements as well as an online pulse shape analysis (PSA).

Figure 12 shows an example of a typical signal from a preamplifier as it arrives at the FEBEX card. It is continuously sampled at 50 MHz using a 14 bit ADC. In the FPGA, the incoming data is then split into a fast branch (used to generate triggers) and a slow branch (used to determine the pulse shape).

As the configuration is absolutely essential for the performance requirements of experiments, it is worth having a closer look at the processing steps happening in the FPGA.[15]

The early filtering stages process the data unconditionally and regardless of a trigger in every clock cycle. For example, the moving average unit will add its current input (which enters the window) to its internal register and subtract a previous input value (which exits the window), then use that register as its output. This is generally referred to as **pipelining**.

[15] Like many famous scientists of the early twentieth century, Max has opted to write his thesis in German. As luck has it, I am somewhat fluent in that language and thus able to make the inner workings of his firmware available to a wider audience.

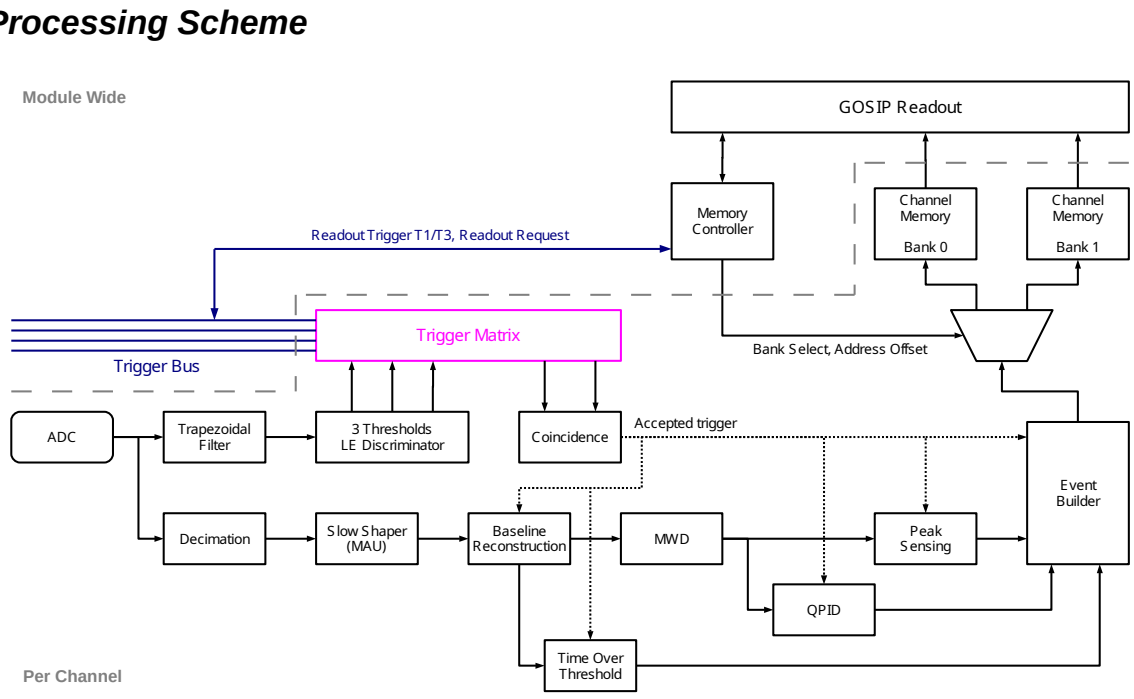

Figure 13: The data processing pipeline in the CALIFA FEBEX3 firmware. Image from [16].

By contrast, stages which do not produce an output in every clock cycle, such as the peak finding or time over threshold measurement are generally implemented as a **state machine** which depends on the event recording trigger.

4.1 Fast branch

In the fast (or trigger) branch, the signal is differentiated using a trapezoidal filter. Afterwards, it is fed to three leading edge discriminators with configurable thresholds. A coincidence of two of these discriminators – or of an external trigger and one of the discriminators – is then used to start data recording for that channel.

The discriminators may be connected to the trigger bus shared by the FEBEX crate.

Typically, one low threshold discriminator provides the trigger timing and a higher threshold discriminator (called "gamma" for historical reasons) provides a validation for event recording.

The third discriminator ("proton") is typically configured to a very high threshold and can be used to provide a proton trigger for ancillary tracking detectors e.g. AMS, FOOT, ALPIDE, L^3T, but also for high level trigger generation in the R^3B central trigger system.

4.2 Slow branch

In the slow (or analysis) branch, the signal is continuously decimated from 50 MHz with 14 bit to 25 MHz with 15 bit and delayed in a FIFO buffer by a configurable amount. The delayed signal is then fed into the MAU filter.

Due to processing element constraints on the chip, only four pipelines for the analysis branch are instantiated on the FPGA. Each of them runs at a clock frequency of 100 MHz and alternately processes four different channels. This is called **interlacing**.

A simplified model of processing will be used to illustrate the effects of the filters. The simulated input is based on the preamp response from figure 12.

The moving average unit (MAU)

A moving average unit (MAU) sums up the last $L = 32$ samples[16] ($1.2\,\mu$s) to reduce higher frequency noise, as seen in figure 14.

Implementation-wise, it should be noted that the arithmetic abilities of FPGAs – speed advantages nonwithstanding – are somewhat similar to second year elementary school students: integers only, addition and subtraction are fine, multiplication is possible and division is right out. However, just like the students could quickly be taught how to divide by powers of ten (discarding the remainder), FPGAs can in fact divide by powers of two as that is equivalent to just discarding some of the rightmost bits.

Thus, the MAU filter does not calculate the average of the last n samples, but the sum. Afterwards, that sum is right-shifted by a configurable amount m before it is passed on to the next processing step. The same is also true for other filter steps involving sums, e.g. the MWD.

[16]The CALIFA FEBEX firmware is very configurable. The parameter controlling this step is called `mau_integration_time` and takes values from 0 (no smoothing) to 1023. Here and in the following sections I will give the default values used by CALIFA in the main text.

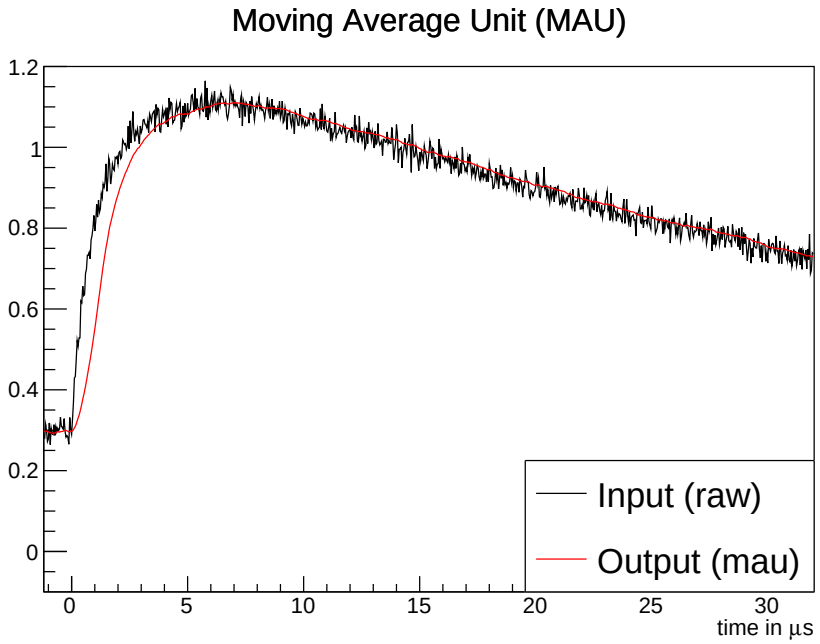

Figure 14: The moving average unit is a simple smoothing filter. The 25 MHz raw trace (black) is averaged over $L = 32$ samples. This supresses high frequency noise (red).

Like virtually every configuration parameter, the selection of the optimal bit-shifts is subject to tradeoffs. If the shift is too large the precision will be lower. By contrast, a shift which is too small might mean that the result overflows. If this happens, an overflow bit is set and the hit can not be used in analysis.

As discussed above, the input of the MAU is a 15 bit number. The data width between subsequent steps of the pulse shape analysis is 18 bits[17]. In the event buffer, only 16 bits are reserved for the result of the PSA (e.g. the signal amplitude), so further shifts may be required to avoid overflows.

Baseline determination and subtraction

The next processing step is meant to get rid of the signal baseline, that is to say the ADC value when there is no ionizing radiation in the scintillator. This is done by keeping a register with the current baseline estimate, and just subtracting this value from every sample. If the discriminator has not fired for 400 μs, the baseline estimate in the register is updated.

The SIS18 accelerator does not deliver a continuous beam but delivers ions in spills (typical length of 1s to 10s) with spill breaks of about 1 s in between them. During these spill breaks, only background activity will hit the detector, so the preamplifier output will return to baseline and the baseline subtraction stage can update the baseline.

[17] This is also the data width of the multiplication units on the ECP3.

This method is preferred with regard to a constant baseline as it will also consider slow baseline changes e.g. due to temperature fluctuations (day/night) in the experimental area.

In the measurements with the very high intensity AmBe source, it was found that the baseline is this method will fail if the trigger threshold is set to low and the time interval without a trigger is never reached due to noise, in which case an undefined value will be subtracted. The first order effects of this can be compensated in the pulse height measurement as discussed below.

A more robust way to reconstruct the correct baseline would be to adjust it based on the output of the moving window deconvolution. In MWD intervals without an energy deposit, the modal signal after the MWD should be zero. This could be used to determine the correct baseline even during an RC decay of the signal.

The moving window deconvolution (MWD)

The output of the charge-integrating preamplifier is exponentially declining with $\tau = 35\,\mu$s. Normally, this would lead to a lower signal after integration as the earlier parts of the signal are already affected by that decline when the peak is reached, a so-called ballistic deficit[18].

For a single event, this would not hurt – the amplitudes measured would simply be slightly reduced by a fixed fraction[19]. However, if a small energy deposit (e.g. from a gamma) shortly after a higher energy deposit (e.g. from a proton), the exponential decline can severely distort the signal. Figure 16 illustrates such a case.

The **moving window deconvolution** (MWD) is meant to compensate this exponential decline. Of course, this can only work if the signal baseline is correctly subtracted, e.g. the input signal can be assumed to decay toward zero rather than an unknown finite value.[20]

The algorithm of the MWD can be derived by assuming that the voltage signal U_i at time step i is given by

$$U_i = kU_{i-1} + mq_i \tag{13}$$

with

$$k := \exp(-\Delta t/\tau_{RC}) \tag{14}$$

The first term is the amount of voltage retained from the previous step, while the second one represents the voltage increase from the charge q_i accumulated in timestep i. m is a conversion factor relating the charge measured by the preamplifier to its output voltage and will be set to one from now on.

Solving this for q_i leads to

$$q_i = U_i - kU_{i-1} \tag{15}$$

[18] "Ballistic deficit" is a strange term for "the RC losses in a charge integrator". It is related to an electromechanic charge integrator called a ballistic galvanometer. Baldinger used the term in 1956 in English, so it clearly predates the unrelated term "missile gap".

[19] Also, the QPID matrix would be more complicated.

[20] [17] cites [18] as a source for the moving window deconvolution algorithm, but uses different variable names. For clarity, we will stick to the varible names from [18].

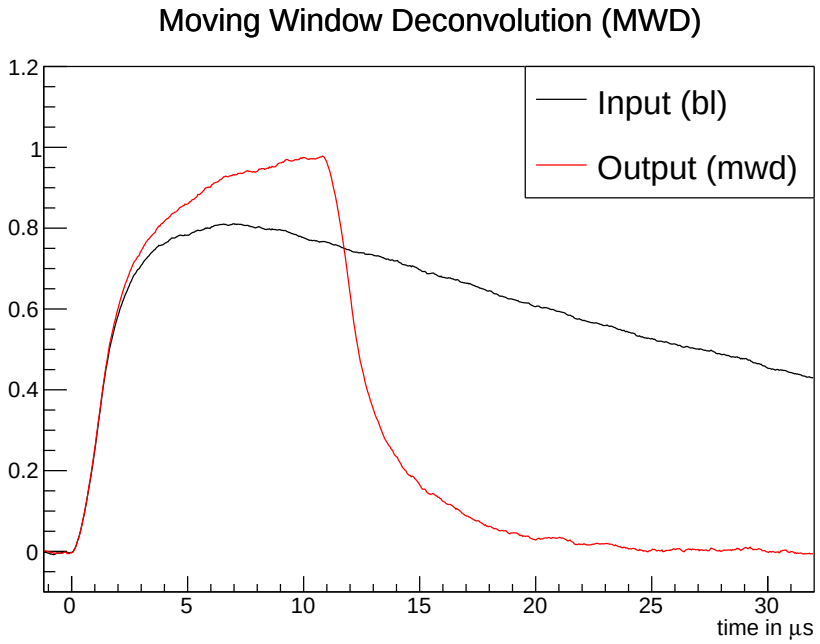

Figure 15: The moving window deconvolution is probably the most intricate processing step in the pulse shape analysis. Here, the window size of the MWD was chosen as 10.8μs. The input is the signal obtained after the baseline subtraction step. Within the window, the MWD counteracts the RC decay of the preamplifier. Any charge collected outside this window is removed by this filter.

For the sum charge over M samples, we arrive at

$$
\begin{aligned}
Q_i &:= \sum_{j=i-M+1}^{i} q_j && (16)\\
&= \sum_{j=i-M+1}^{i} (U_j - kU_{j-1}) && (17)\\
&= \sum_{j=i-M+1}^{i} U_j - k\sum_{j=i-M}^{i-1} U_j && (18)\\
&= U_i - U_{i-M} + \sum_{j=i-M}^{i-1} U_j - k\sum_{j=i-M}^{i-1} U_j && (19)\\
&= U_i - U_{i-M} + (1-k)\sum_{j=i-M}^{i-1} U_j && (20)
\end{aligned}
$$

Assuming that the time difference Δt between samples is small compared to τ_{RC}, one can

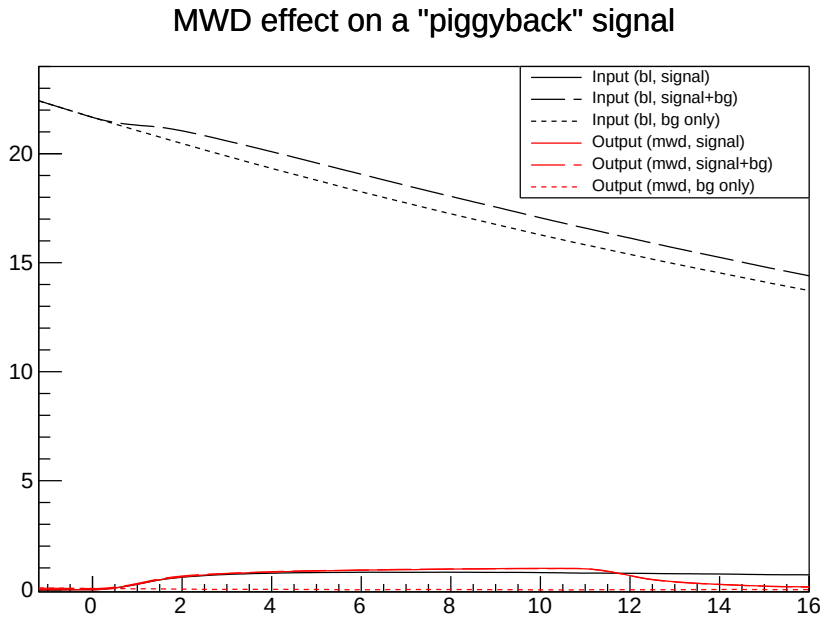

Figure 16: The effect of the moving window deconvolution on a small signal appearing during the decay of a larger "background" signal (dotted). The simulated background signal has 50× the amplitude and precedes the signal of interest by 32 μs. This results in an ADC trace which is monotonically decreasing.[23] However, after the MWD, the pile-up signal is none the worse for wear and it's energy can still be measured accurately.

use a Taylor series expansion of the exponential to simplify this a bit:

$$Q_i = U_i - U_{i-M} + \frac{\Delta t}{\tau_{RC}} \sum_{j=i-M}^{i-1} U_j \tag{21}$$

Which is equation (3.9) in [17] with some terms renamed[21]. As divisions are hard to do on an FPGA, the quantity $\hat{Q}_i := \tau_{RC} Q_i$ is introduced.

The length of the window $M\Delta t$ can be thought of as the memory horizon of the MWD. At any time t, any charge collected by the APD before $t - M\Delta t$ will be irrelevant for the output. Any charge collected within the MWD window will be added, possibly leading to pile-ups[22]. On the other hand, any charge not collected within the window will not contribute to the amplitude. The value widely used in our firmware is $M = 270$ which corresponds to 10.8 μs. Figure 15 demonstrates the effect of the MWD.

The effect of the MWD filter can be seen in figure 15. A more spectacular case is the pile-up shown in figure 16. Here a signal sitting on the tail of a much larger signal was simulated. In reality, the higher signal might be a proton and the lower signal might be a γ ray from a later

[21] $U \to D$, $\Delta t \to 1$, $M \to L$, $j \to k$

[22] Such pile-ups will be detected and flagged by the firmware.

event, so that the RC decay of the background amplitude is larger than the rise of the small signal of interest and the overall trace is decreasing monotonically. However, as the previous big signal follows the exponential decay in the time region of interest, it can be subtracted by the MWD, and the amplitude of the smaller signal can be reconstructed correctly.

For such a signal, a related question is how to trigger on it, as the timing branch uses a simple trapezoidal filter which does not deconvolute the RC decay in the preamp. At the moment, the presence of large signals will lower the trigger efficiency for small signals. One could solve this by using an external trigger without any validation, but then one would have to deal with vast amounts of data from CALIFA. A future firmware version might feature a MWD within the timing branch, but for now an experiment specific evaluation of γ ray trigger efficiency is required.

Effects of an incorrect baseline on the MWD

Suppose that instead of the true baseline, a wrongly estimated baseline is subtracted. Intuitively, this should be fatal: the MWD is meant to compensate the signal decay towards the baseline, which requires the baseline to be known.

Quantitatively, one might model the incorrectly baseline-subtracted input to the MWD as $U_i' = U_i + c$ for some constant offset c. This will introduce an error

$$Q_i' - Q_i = \frac{\Delta t M}{\tau_{RC}} c \tag{22}$$

The most important fact to point out here is that the error is constant, i.e. it does not depend on U_i. In the canonical configuration, the length of the MWD window $\Delta t M$ is $10.8\,\mu$s (three times the life-time of the slow component of CsI) while τ_{RC} is $35\,\mu$s. Thus, the offset after baseline correction should be about a third of c.

In future firmware versions, the output of the MWD could also be used to determine the correct baseline value reliably. As long as there is no signal in most MWD windows[24], the most common output of the MWD should be zero. This would allow to reconstruct the baseline also in scenarios where high rates or noise induced triggers prevent the stronger condition of "$400\,\mu$s without a trigger" from ever becoming true.

The MWD window is so long because CsI(Tl) is a slow scintillator. As discussed in section 3.4, the longer lived excited state has a τ_s of $3.4\,\mu$s. Within the MWD window, one will collect 96% of the photons from that state and virtually all from the short-lived state, for a total of perhaps 98%.

The lower photon statistics impact the relative resolution due to photon statistics by $\frac{1}{\sqrt{0.98}} \approx 1.01$. With relative γ ray energy resolutions on the order of 0.05, this effect is negligible.

However, if one used only half the window, one would collect only 90% of the photons. This would then severely impact the resolution of the γ rays.

Waiting longer to collect every last scintillation photon has the danger of getting pile-ups.[25]

[24] That is a very weak condition. Assuming a Poisson distribution of hits, pile-ups would make useful measurements impossible if this condition was negated.

[25] The expected value of course depends on the beam rate.

ΔtM around 11 μs are thus a reasonable compromise between photon collection and pile-up resistance.[26]

Pulse height measurement ("energy")

After this step, the signal is fed to different units which extract quantities of interest from the 25 MHz trace. The height of the peak signal, commonly identified as the energy deposited in the scintillator is measured by first sending the signal received from the MWD through another MAU unit (L=64, 2.56 μs) and then searching within a 12 μs gate for the maximum, which is taken to be the energy. [27] This is equal to the MWD window – the time during which we collect photons – plus some delay introduced by the MAU plus some safety margin.

As discussed above, in some rare cases, the baseline reconstruction will not work perfectly. In the default configuration, an error in the baseline subtraction will propagate to the energy measurement. However, there is a flag `energy_bias_correct` to activate a 2nd order bias correction in the algorithm. This does not just reconstruct the energy from the absolute value of the peak but introduces a relative measurement of the amplitude between the peak and a reference point taken at an earlier time – the first sample of the peak-sensing MAU after the gate is opened is subtracted from the found maximum, which will avoid any gross amplitude measurement errors from the baseline subtraction. As shown in the previous section, the effect of an incorrect baseline on the MWD results in a constant error term, which will vanish with subtraction.

The major downside of using this flag is that instead of measuring relative to a well determined point of reference, one is now measuring the difference between two quantities with comparable errors, effectively increasing the error by a factor of $\sqrt{2}$ at low rates when pileup within the energy gate is not a dominant effect.

Alternatively, a pulser with a fixed amplitude can be fed to all channels. The amplitude measured by FEBEX will shift if the baseline subtraction is not working correctly.

Quick particle identification (QPID)

The quick particle identification is important not only because it allows identifying different isotopes in CALIFA, but because it allows to differentiate stopped protons (where the energy measured is the energy of the proton) from punch-through protons where the measured energy is ΔE and the actual energy has to be extrapolated (see section 3.2).

The algorithm implemented is called "Quick Particle Identification" (QPID) [17]. It allows to analyse the trace data in the FPGA and estimate how much of the collected charge was due to either the fast or the slow time component of the CsI.

Before the signal is passed to the QPID state machine, it is be shaped by another MAU. Unlike the amplitude measurement, the QPID is very sensitive to the shape of the signal. Thus,

[26] For single event readout, this also has an effect on the dead-time.

[27] As opposed to just subtracting two samples relative to the trigger, taking the maximum over an interval makes the method robust against small uncertainties in the trigger time: that difference will likely be largest just before the rising edge of the signal. It would be interesting to also record the time at which that maximum was found in the data stream so that one can improve on the timing of the signal.

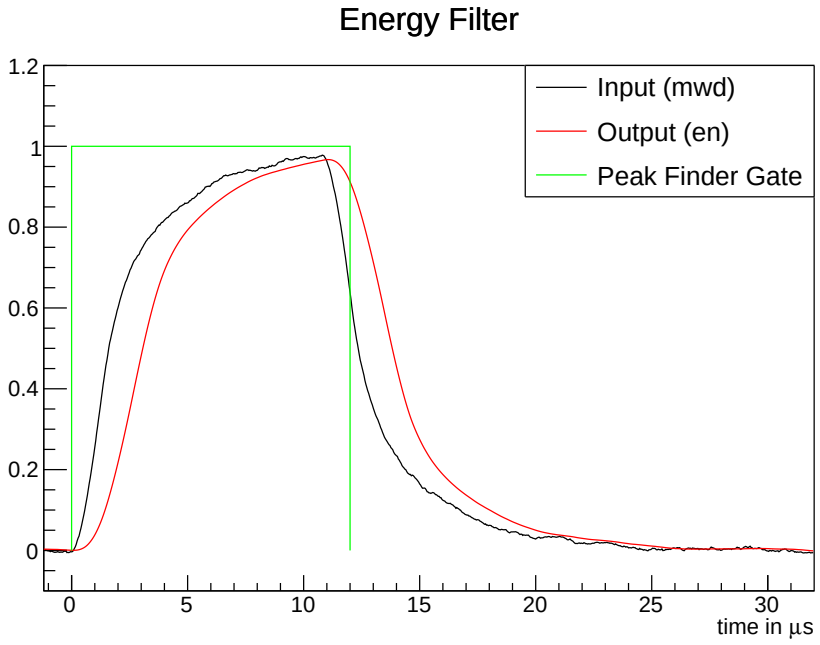

Figure 17: For the energy subbranch, the signal from the MWD is subjected to another longer moving average unit. A peak finder will then search for the maximum within the gate.

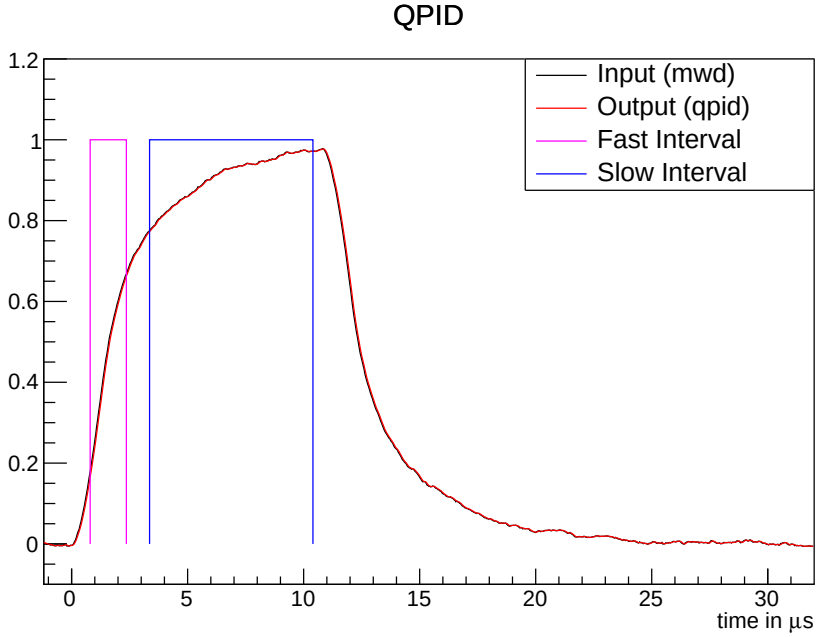

Figure 18: For the QPID, the signal from the MWD is only moderately shaped by another MAU filter. Then, the increase in two different intervals is measured.

only a bit of smoothing (L=4, 160 ns) is applied.

In the QPID, the integrated input signal is sampled at four different times (relative to the trigger time). The increase between the first and the second sample is called Q_F (from Q for charge, F for fast), the increase between the third and fourth sample is called Q_S (S for slow).

For a given amount of light emitted from the fast (N_F) and slow component (N_S) of CsI, the resulting Q_F and Q_S are going to be a linear function of N_F and N_S:

$$\begin{pmatrix} Q_F \\ Q_S \end{pmatrix} = M \begin{pmatrix} N_F \\ N_S \end{pmatrix} \tag{23}$$

M is here a 2x2 matrix taking into account the sampling times, the shaping time and the decay time of the slow and fast component of CsI(Tl).

Of course, in practise, the Q_* are measured and the N_* are desired. By committing the numerical sin of matrix inversion, one obtains

$$\begin{pmatrix} N_F \\ N_S \end{pmatrix} = M^{-1} \begin{pmatrix} Q_F \\ Q_S \end{pmatrix} \tag{24}$$

This step is performed on the FPGA itself using fixed point arithmetic.

QPID and trigger timing

While the QPID should be immune to baseline problems, it relies very much on the timing of the trigger signal, so extra care must be used to set the threshold of the trigger used for timing to a suitable value. Figures 19 and 20 visualize the effects of badly timed triggers in a trace simulation. Here, a function

$$f(t) = \begin{cases} \frac{\tau_F}{2}\left(1 - \exp(\frac{-t+t_0}{\tau_F})\right) + \frac{\tau_S}{2}\left(1 - \exp(\frac{-t+t_0}{\tau_S})\right) & t \geq t_0 \\ 0 & \text{otherwise} \end{cases} \tag{25}$$

was used to simulate the charge signal of an event at time $t_0 = 20\,\mu$s where both time components contribute in equal measure. Figure 19 uses the difference Δt between th start of the event t_0 trigger time t_T as the x axis, while the following plots use t_T directly for technical reasons.

The nominal behavior is observed for the correct trigger time, e.g. $\delta t = 0$ or $t_T = t_0 = 20\,\mu$s. This leads to the maximum value for Q_F.

The number of times the trigger threshold was reached during event recording is also stored in the event as a number called `pileup`. Under some conditions, this will allow the identification of pileup events.

Time over threshold (TOT) measurement

The firmware also includes a time over threshold (TOT) measurement which can measure the pulse height of signals larger than the ADC range, which might happen if a preamplifier channel in high gain ("gamma") mode detects a proton. The main downside is that while the signal

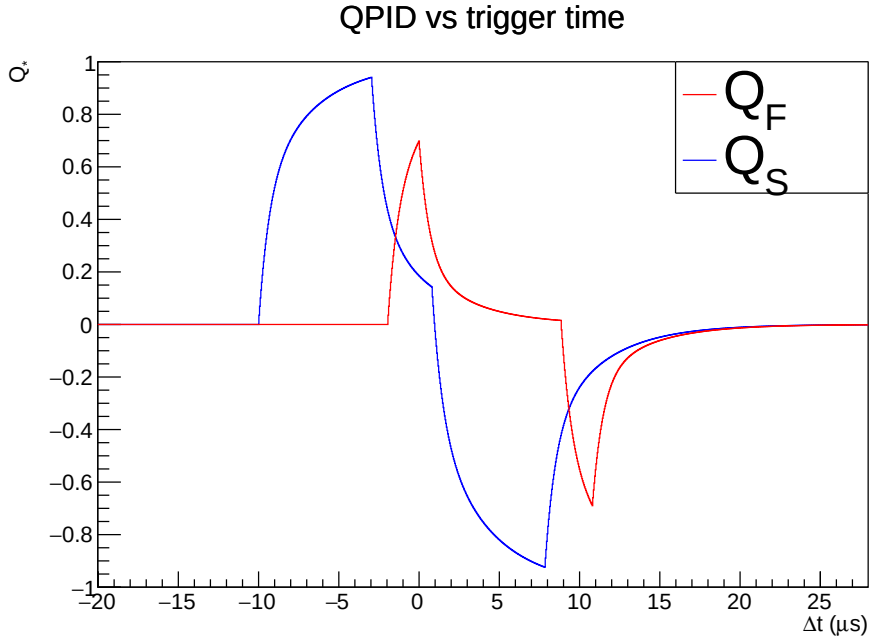

Figure 19: The QPID results for a simulated CsI signal ($N_S = N_F = 0.5$) as a function of the trigger time offset Δt. If the trigger arrives very early ($\Delta t < -10\,\mu$s), both QPID windows will see only baseline. For $\Delta t = -t_4 = -10\,\mu$s to $\Delta t = -t_3 = -3\,\mu$s, the second window catches the rise of the signal, boosting Q_S. At $\Delta t = -t_2 = -2\,\mu$s, the earlier window catches the rise of the signal, while at $\Delta t = -t_1 = 0$ it perfectly captures the rising flank of the signal. As Δt becomes positive, both Q_F and Q_S become smaller. The negative Q_* values for late triggers ($\Delta t > 0.8\,\mu$s) are due to t_4 (and latter t_2) exceeding the MWD length (see figure 15) and thus capturing the "echo" of the rising flank.

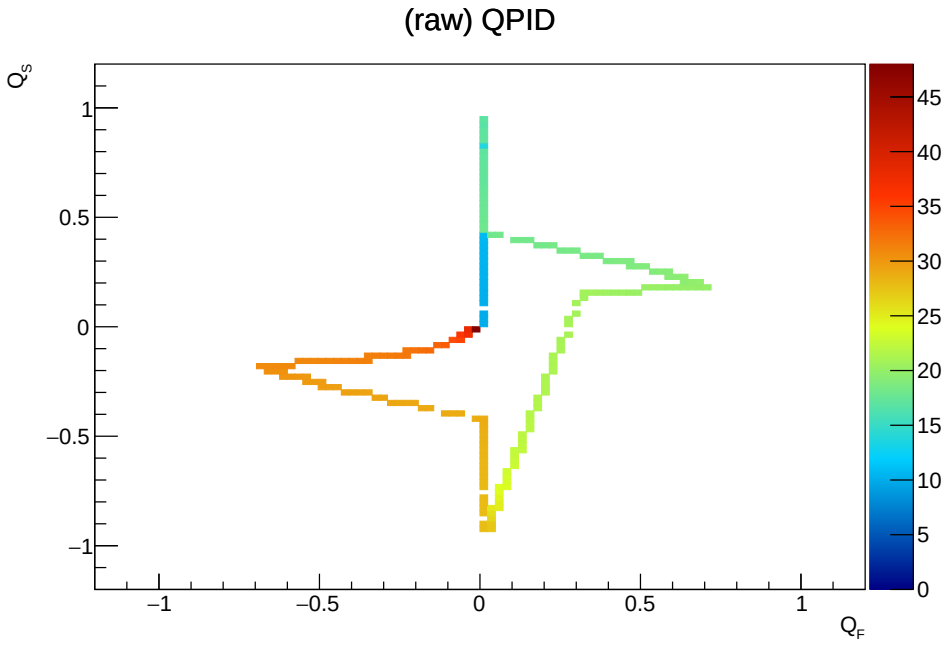

Figure 20: The data from figure 19, plotted as a function of the trigger time. The z axis (color) indicates the artificially shifted time t_T of the trigger in μs. The rising flank of the CsI signal always happens at $t = 20\,\mu$s (previously $\Delta t = 0$), which corresponds to the green part of the curve where Q_F is maximal and the QPID works correctly.

is out of ADC range, it is impossible to detect pile-up events. Any energy deposited into the crystal while the TOT takes place will distort the measurement.

$\tau_{RC} = 35\,\mu$s was chosen to be much larger than the slow time component of CsI. A proton exceeding the ADC range by a factor of $E_0 = 10$ will block that channel for $\tau_{RC} \ln E_0 \approx 80\,\mu$s. If at the end of that $80\,\mu$s, a gamma deposits another $E_1 = 0.1$ of the ADC range, the measured TOT will be increased by $\tau_{RC} \ln(1+E_1) \approx 3.3\,\mu$s. So while one might naively suppose the TOT to just sum up all the energies deposited during the measurement, arriving at $E_0 + E_1 = 10.1$, in fact one arrives at

$$\begin{aligned} E_{TOT} &= \exp(\Delta t/\tau_{RC}) && (26)\\ &= \exp(\ln(E_0) + \ln(1 + E_1)) && (27)\\ &= E_0(1 + E_1) = 11 && (28) \end{aligned}$$

In effect, E_0 acts as a lever arm for E_1, allowing even low energy γ rays to seriously distort the measurement of proton energies.

Also note that while the TOT can substitute for a pulse height measurement in some situations, it will not reproduce the particle identification information provided by the QPID – which will also not work if the signal goes beyond the ADC range.

4.3 Configuration format

The configuration registers of the FEBEX FPGAs are set during the *start acquisition phase* via GOSIP. Changing their values requires a restart of the DAQ.

Traditionally, the FEBEX configuration is stored in a binary file called febex.db from which the desired configuration settings are read. That file is manipulated using a program called `febex_set_param`, which allows the interactive manipulation of settings.

For example, `set 0.5.num_events_readout 42` will set the configuration register describing the number of events in the buffer which should trigger a readout to 42 for the 5th febex module on the 0th fiber optical link number[28] – either index counting from zero. Some wildcards are allowed. For example, `set *.4-5.0-7.discr_threshold_timing 100` will set the "timing" discriminator threshold on all four fiber ports ("*") for modules with indices 4 or 5 ("4-5") of the first eight channels ("0-7") to 100.

Similar to the example `set` commands, `get` commands can query existing settings, while the list command will dump all of the register values of all modules.

During this work, additional commands were added. `slist` will aggregate over all channels and show the range of register values, ordered by the processing stage they affect. `tlist` shows a text based visualization of the trigger matrices[29].

The binary format is not well suited for large scale operations. While different versions of the file can be created by copying, comparing the changes between versions or merging changes

[28] The index of the fiber optical link if often called the SFP number after the SFP+ transceivers used for transmission.

[29] While still requiring the user to specify the trigger matrix as bit fields, this allows to quickly verify that the settings are in fact correct.

from two versions is difficult. To overcome these deficits, the binary files are now generated from text files composed of a sequence of `set` commands.

To simplify complex operation, a macro language is also provided. As most DAQ users are already proficient with the C programming language, the C preprocessor was used to provide this macro language. This allows the following features:

- `#include` statements which split the configuration over multiple files
- `#define` constants allow refering to a range of channels by a name: `#define SPECIAL 3.16.0`
- `#define` macros with arguments can be used to shorten frequently used assignments[30]:
```
#define HT(CH) \
set CH.discr_threshold_timing 200 ; \
set CH.discr_threshold_gamma 300
```
- `#ifdef` can be used to implement branches depending on macros to quickly allow switching configurations.

The top level file is called `febex.h`. The layout of the files is typically the following:

- The correct number of fiber links ("SFPs") are added
- The correct number of modules are added to each SFP
- The default settings are included. They consist of wildcard assignments of good default values to all registers
- For the barrel, an ifdef block defines configurations for both the case where its preamps are running in high gain or low gain mode.
- Sensible default discriminator thresholds are set
- The trigger configuration is loaded (typically from separate files)
- The thresholds for individual higher noise channels is loaded (`thresholds.h`)
- The configuration for any special channels (which might record pulser signals, or the main trigger from the R^{3}B main trigger logic[31])
- Finally, a file `shift.h` is included. This allows the DAQ operators on shift to overwrite any settings without changing the other text files.

When a DAQ restart is requested, the binary febex.db is automatically recreated from the text based configuration. The git[32] version control system is used to automatically store any changes. This means that it is later easy to figure out what settings were used in a particular run.

[30] As the expansion of a macro is can not contain the new line character, the functionality to separate commands by semicolon instead was also added.

[31] The length of that trigger pulse is variable, and the pulse length measurements of multiple systems can be correlated to ensure that the time-stitching is working correctly.

[32] `https://git-scm.com/`

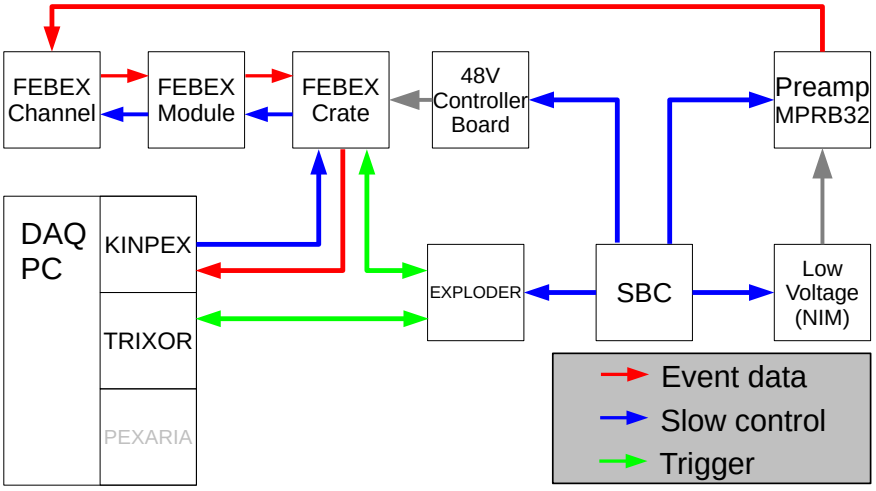

Figure 21: A general overview of the electronics used in CALIFA. The KINPEX card is used both to read event data and set slow control parameters in the FEBEX FPGAs, while the rest of the components are slow controlled by a single board computer (SBC). Control flows are bidirectional, the arrow goes from controller to the controlled device. For more details about the trigger setup, see figure 26.

5 Trigger and Readout System

The trigger and readout system of R^3B and especially the CALIFA detector has to fulfill a series of challenging requirements. Especially the flexibility to adopt for very different experiments with either focus on low rate charged particles, high multiplicity events with a mixture of charged particles and low energy γ rays or high resolution γ spectroscopy initiated the use of a fully digital implementation based on field programmable gate arrays (FPGA) which can be reprogrammed and optimized for different modes of operation. Goal for the R^3B-CALIFA system was to allow for both: a fully triggered system where real time information from the different detectors allows to select a time window of interest and write out all data within this window, or to run a fully self triggered system where individual detectors elements decide on the event selection just using their individual signal patterns.

The whole system is founded on a series of hardware modules developed by the GSI experiment electronics department for the MBS[33] framework which allow for interfacing a large variety of custom made front-ends optimized for different detector applications. Figure 21 gives a conceptual overview on the interplay of the different components, the data stream and the slow control system to set and optimize parameters for a reliable operation. The following section outlines the different components and their implementation and features for the CALIFA detector.

[33]Multi-Branch System, https://www.gsi.de/en/work/research/experiment_electronics/data_processing/data_acquisition/mbs

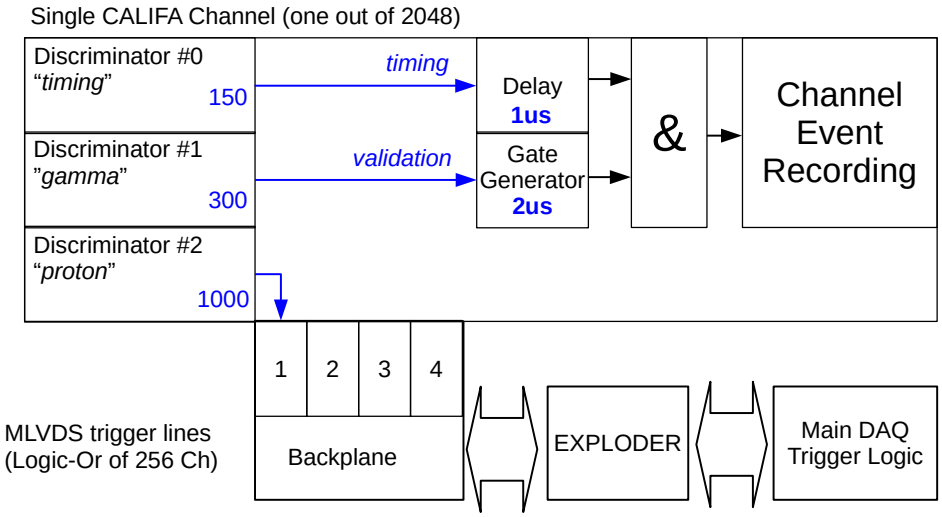

Figure 22: The trigger configuration of a single CALIFA channel. All of the blue arrows and numbers are implemented as slow control parameters which can be easily changed.

5.1 Triggering of CALIFA channels

Figure 22 shows the triggering unit within a single CALIFA channel. As described in section 4.1, three independent discriminators continiously monitor the output of the trapezoidal filter and assert a trigger if their threshold is reached.

The trigger unit which controls event recording has two inputs, a *timing* input that starts event recording and an optional *validation* input which can be used to enforce additional constraints on recording an event. The validation input gets fed to a gate generator, and the timing inputs can only start an event within the gate (or when validation is disabled).

In the case shown, the "timing" discriminator runs with a low threshold and feeds into the timing input of the coincidence unit, while the "gamma" discriminator runs at a higher threshold and serves as validation.

Furthermore, both the outputs and inputs can be connected to a module-wide trigger matrix which is in turn connected to the trigger lines (a shared MLVDS bus) running on the backplane of the FEBEX crate and from there to the EXPLODER, the central trigger module.

As shown in figure 22, the third discriminator can be set to a higher threshold and e.g. provide a proton trigger for other systems of the R^3B experiment. But it is also be possible to use an external trigger in the coincidence unit, for example as validation.

A FEBEX parameter called `logic_gate_size` controls the maximum processing time for one event. It defaults to 16 μs. During that time, the coincidence unit does not accept any further triggers, that individual channel is in *dead time*. [34]

[34] While two events separated by more than this can be recorded, it should be understood that such an event rate (e.g. 50kHz) in one channel is not sustainable with the CALIFA electronics: The $\tau = \frac{1}{RC}$ used in the integrating preamplifiers is 35 μs.

5.2 CALIFA Readout Modes

The CALIFA Data Acquisition System (DAQ System) has to fulfill rather demanding requirements. Operating up to 2432 large scale crystals with a total mass of 1700 kg means a substantial total interaction rate (e.g. 20 kHz) just by the background from radioactives dacays in the experimental hall. When used to detect γ rays, CALIFA has to be sensitive down to 100 keV energy while also having the least possible dead time for individual channels. Additionally, different experiments will have different crystal hit multiplicities expected – ranging from one up to 100. Another essential requirement is the implementation of a triggerless system needed for calibration and efficiency measurement. Here individual γ rays from radioactive sources interact with few detector crystals only and as there is no external signal available the DAQ system has be able to trigger itself.

5.2.1 The triggered mode

To understand how the readout mode of CALIFA works, and why it works the way it does, it is essential to first understand the intended readout method for which the GOSIP protocol was originally intended. This process is visualized in figure 23.

GOSIP is designed for a synchronous readout: the dead time master (the EXPLODER) accepts a physics trigger and broadcasts it to both the FEBEX modules in the dead time domain and the TRIXOR card of the DAQ PC. Every card will start processing the input data of each channel into a hit record stored on a memory bank for readout. There is the possibility to suppress hits where the channel saw nothing interesting.

Meanwhile, the signal arriving at the TRIXOR card wakes up the readout code for the KINPEX card. After a configurable amount of waiting for the FEBEX cards to finish event recording, the KINPEX card starts the readout process. As the FEBEX cards in a crate are linked in a daisy chain, the KINPEX card sends a token to each FEBEX card. The first febex card in a crate will respond by sending the contents of the memory bank back, then pass the token on to the next card and start forwarding any replies received downstream to the upstream (e.g. KINPEX-wards).

Eventually, all of the FEBEX data is collected. The TRIXOR card signals the dead time master that the readout is finished and the dead time (eventually) ends.

This kind of fully synchronous readout scheme is extremely stable but will create a substantial dead time.

5.3 Limitations of the synchronous readout

With sixteen channels per FEBEX module and sixteen modules per FEBEX crate, a single DAQ PC can read out up to 1024 channels over its four fiberoptical data links. Each channel can process events in the kilo Hertz range.

A CALIFA GOSIP readout of 64 FEBEX modules has an overhead of about 150 μs. This comes down to a maximum readout rate of 6 kHz in this mode.[35] If one assumes that most

[35]It should be noted that our `f_user.c` reads out the four links sequentially. By reading them out in parallel, one might gain some speedup, but certainly not more than a factor of four.

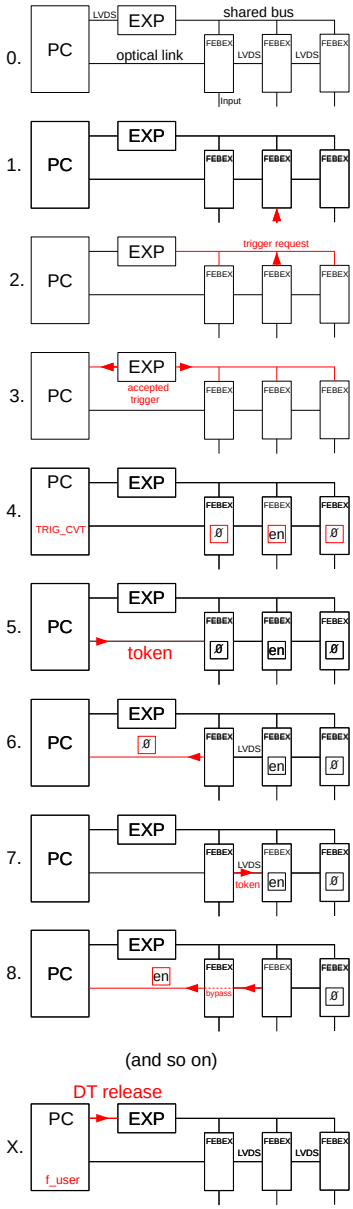

Figure 23: The GOSIP readout process for single event mode, self-triggered. (0) is an overview with the principal key components and links. (1) FEBEX card one receives a signal above the trigger threshold. (2) That FEBEX card asserts the trigger request line in the MLVDS bus. (3) The EXPLODER (labeled EXP), not being in dead time, accepts the trigger and sends that information to both the TRIXOR card and all FEBEX modules. (4) MBS waits for TRIG_CVT while the FEBEX modules start event recording (subject to zero suppression). (5) The KINPEX card sends a token to the FEBEX card zero, requesting a readout. (6). FEBEX card zero did not see anything above the threshold, so it returns an empty frame. (7) The token is passed on to FEBEX card one. (8) FEBEX card zero goes into bypass mode and passes on the frame of FEBEX card one to the KINPEX. This process continues for the other FEBEX cards. (X) After the readout is finished, the dead time is released. In some cases, the dead time may be releases before the readout is finished.[46]

events have a multiplicity of one, we could handle each of the 1024 crystals triggering at a rate of about 6 Hz – perhaps a percent of the rate the crystals could handle, and certainly less than the rate expected from background radioactivity such as ^{40}K.

5.3.1 Multi-event readout and free-running system

To overcome these limitations, we store more than one hit per channel in the memory bank of the FPGA, and only read them out when the buffer becomes too full.

The CALIFA multi-event readout mode runs on top of GOSIP, so neither the EXPLODER nor the DAQ PC needs to be changed. The FEBEX modules store the recorded hits in local memory (up to 254 events per channel) and only sends out a **readout request** when the buffer reaches its configured maximum event number N ($1 \leq N < 256$). From the perspective of the EXPLODER and PC, the readout still happens synchronously in response to a readout request (GOSIP trigger).

Like in every buffered system, finding out which hits belong together becomes non-trivial. Luckily, a timestamping system is available. The EXPLODER uses one of the MLVDS trigger lines to transmit a 20 MHz clock which is counted in every FEBEX module (and automatically reset on certain conditions[36]). The count of this clock is commonly called the **FEBEX timestamp**.

As this mode still works with GOSIP, the real time design is still in place. Only now, instead of a physics trigger, a readout request is sent by the FEBEX card running full, and a readout trigger is sent from the EXPLODER in response. This additional complexity – compared to a simple fire and forget asynchronous transmission, e.g. via IEEE 802.3 (ethernet) – comes from the requirement of also supporting the triggered single event readout mode.

During the first experiments and especially during the R^{3}B commissioning (S444), CALIFA had been operating in fully self triggered mode[37]. However, it turned out that the typical data rate of 500 MiBytes/s from the detector dominate the overall data rate of R^{3}B and especially saturate the server tasked with time sorting and event building. Especially in the experiment S494 using an ^{16}O beam at an intensity of 10^8 particles per second, there was a stringent request to reduce the rate by using an external reaction trigger for validation.

5.3.2 Free-running system with external validation

In the framework of this thesis a substantially new configuration of the CALIFA DAQ as a fully triggered system was implemented. As discussed in section 3.7, the validation of crystal signals is performed in a two stage way. First the pulse shape analysis implemented for the timing discriminator selects a point in time where there could have been a useful signal that should be recorded. In a second stage this assumption is validated by the overlap with a so called validation window. In the self triggering mode this validation window is opened by a different branch of the PSA called the gamma discriminator. Here the same internal trace is evaluated again using a safe threshold, defined by the rate monitoring of each channel in the beam breaks. In the future facility it is planned to use here an adaptive evaluation performed

[36] For the curious, it is the first readout after the acquisition is (re)started.

[37] The Arnaud Amalric approach to DAQ: "Record everything and let the Analysis People figure it out"

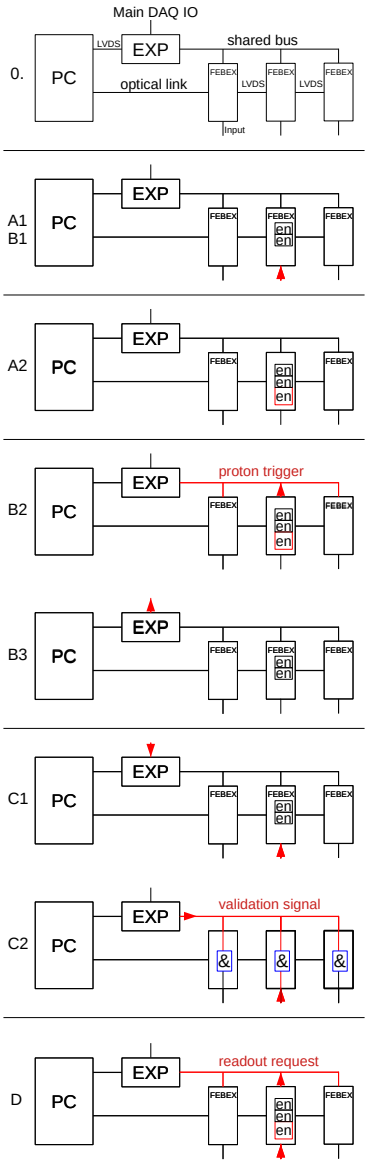

Figure 24: The triggering processes for multi event readout. (0) illustrates the involved modules and their links. Here, the bidirectional connections between the EXPLODER and the Main DAQ trigger logic are also included. If a FEBEX card in free-running mode receives a signal (A1), it is stored in an internal buffer (A2). If the signal is above the proton threshold (B1), a logic signal is sent to the EXPLODER (B2) and from there to the Main DAQ (B3). (C) visualizes the use of a validation trigger: the external validation is received (C1) and sent to the FEBEX modules (C2), where it is used to open a coincidence gate. If a signal surpasses the threshold while the gate is active, the event is recorded. (D) is finally the start of the readout process. A module in which the number of recorded events reaches a predefined threshold sends out a readout request (which is technically equivalent to the single event trigger request). From there, the readout proceeds similar to the process in figure 23.

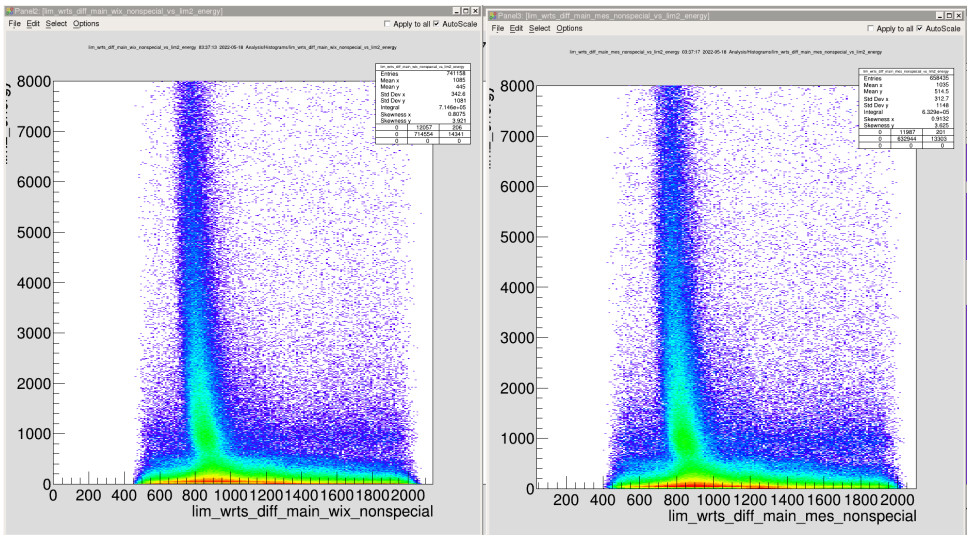

Figure 25: Walk correction plots implemented in the online QA tool for the S522 experiment, the difference of the white rabbit time stamps between the CALIFA hits (interpolated WRTS) and the main DAQ versus the uncalibrated energies. The external validation gate opened by the main DAQ trigger corresponds to the region from 500 ns to 2000 ns region. The hits with higher energies (e.g. above 500 ADC units) clearly fall within the 750 ns to 950 ns range.

in the spill breaks. The length and the delay of the validation windows are adjusted by slow control parameters for each FEBEX module.

Figure 24 gives an overview of the operation of the trigger system in multi event mode.

In the full triggered system the validation trigger (traditionally provided by FEBEX discriminator set to a higher threshold) is replaced by an external signal provided by the R^3B central trigger system. Typically a coincidence of the R^3B start detector with signals from the time of flight wall (TOFD) or the Neuland detector provide here a minimum bias information for any reaction in the target. As the time jitter of this signal is defined by the time resolution of the start detector there is a very precise time delay to the secondary reactions in the CALIFA detector. Only the constant latency of this signal has to be defined experimentally.

Within this work package a procedure to calibrate this latency and adjust a series of slow control parameter for the triggered mode had to be identified tested, implemented in the online monitoring and documented. Several sources of time jitter on the signal path were identified and eliminated. As shown in figure 25 the slow CsI signals above an amplitude of 1000 ADC units fall within a window of 200 ns.

It should be mentioned that especially for a small energy deposit of the particles and gammas in an individual crystal the slow signal rise time of CsI will limit the detection efficiency in this case. This is why the coincidence window is typically adjusted to 4 μs but can be reduced if e.g. only large energy deposits are of interest.

This change in the overall configuration of the CALIFA DAQ could only be performed as already in the stage of development of the system large effort was put to the FPGA code of the

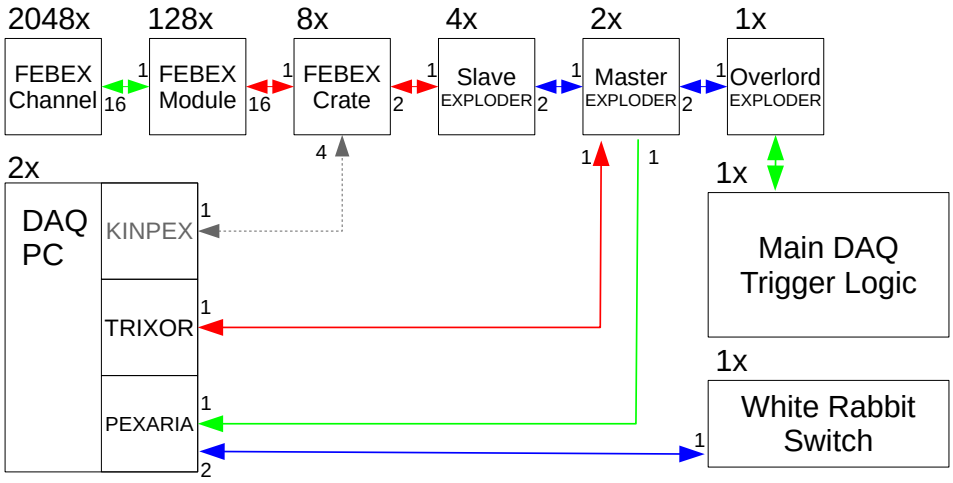

Figure 26: An overview of the DAQ with a focus on trigger and timestamp transmission. Green arrows indicate single ended transmission, red ones (M)LVDS and blue ones transmission over optical fibers. The optical link between the KINPEX and the FEBEX crates is used for readout and does not transmit trigger information. EXPLODERs are shown as logical, not physical units, see figure 27.

individual modules to create a very flexible and versatile trigger system for CALIFA.

5.4 The Trigger System

Trigger handling and timing is managed by a set of two different hardware modules linked together. There is the EXPLODER3[38] module which forwards triggers between digitization modules (such as FEBEX) and their readout system (a server PC with special interface cards), as well external trigger sources for consumers. A shared bus (MLVDS[39]) is used to connect the EXPLODER3 to the FEBEX modules while bidirectional LVDS connection to the server ends at the TRIXOR[40] card, a PCI card which controls data collection and provides signals on the status of operation and initiates the readout process each time a trigger from the EXPLODER is provided.

Besides the MLVDS connectors (which typically go to the FEBEX crates) and the connectors to the TRIXOR card, the EXPLODER also features eight LEMO inputs and outputs as well as four SFP cages for fiberoptical links.

Figure 26 shows the signal distribution of the trigger system and the different connection to the other systems in R^{3}B.

[38] https://www.gsi.de/en/work/forschung/experimentelektronik/digitalelektronik/digitalelektronik/module/font_end_module/exploder/exploder3

[39] Multiple end point low voltage differential signaling

[40] https://www.gsi.de/work/forschung/experimentelektronik/digitalelektronik/digitalelektronik/module/pci_pci_e/trixor

The central features of the trigger system of the CALIFA detector are implemented in the CALIFA FPGA firmware for the EXPLODER.

To allow for the most flexible trigger configuration, the CALIFA EXPLODER firmware allows mapping inputs to any of 32 internal lines via the *input matrix*. From there, the 32 internal lines can be connected to any outputs via the *output matrix*.

Two types of internal lines exist:

- **deadtime lines 0-15** participate in a common dead time locking.
- **bypass lines 0-15** unconditionally propagate the signals.

The **input matrix M** is a matrix stored in registers writable via slow control over the Boolean semiring[41] mapping the set of input states P (e.g. LEMO, LVDS, MLVDS, fiber link, locally generated pulsers) to the set of internal states[42] $L = \{0,1\}^{32}$. In particular, $\vec{p} \in P$ is mapped to $\vec{l} \in L$ using:

$$\begin{aligned} \vec{l} &= \mathbf{M} \cdot \vec{p} \\ l_i &= \bigvee_j M_{ij} \wedge p_j \end{aligned}$$

These internal lines then get mapped to the outputs

The **output matrix N** is also a matrix over the Boolean semiring and map any configuration of internal lines $\vec{l'} \in L$ (where l' is the internal state after the dead time locking) to a state of outputs $\vec{u} \in U$ (e.g. LEMO, LVDS, MLVDS, fiber link, scalers):

$$\begin{aligned} \vec{u} &= \mathbf{N} \cdot \vec{l'} \\ u_k &= \bigvee_k N_{ki} \wedge l'_i \end{aligned}$$

For a fixed internal bypass line i and fixed M_{i*} and N_{*i}, the EXPLODER thus resembles a section of a NIM Fan-In/Fan-Out module such as the LeCroy 429A: a trigger on any of the connected inputs (as specified by M_{i*}) is forwarded to all connected outputs (stated in N_{*i}).[43]

For a deadtime locked signal, the process is similar except that a signal has to be accepted by the EXPLODER before it is forwarded to the outputs. Afterwards, the EXPLODER will then reject any other deadtime signal until the TRIXOR card releases the dead-time.

In the frame of this thesis a user friendly interface to configure this three dimensional routing table was developed, tested and already applied in several experiments (S509, S522). YAML[44] was used as a basis. See listing 1 for an example. Instead of writing out the sparse matrices

[41] that is to say, the set $0, 1$ with (inclusive) OR as addition and AND as multiplication

[42] Or, in case of the deadtime locked lines, the inputs to the locking state machine.

[43] Keen students of boolean algebra will notice that the trigger matrices only can combine inputs using the OR operation. The TRLO II firmware will offer more logic options including NOT and AND, allowing implementing coincidences and vetos. Still more advanced designs might implement register-configurable *soft* lookup tables (LUT) with multiple inputs which would be able to trigger on such as *exactly six of the eight inputs are asserted*, not that this would be a likely trigger condition in any experiment.

[44] YAML is a human readable data serialization language. Compared to other such languages such as JSON or XML, it generally has a much more concise syntax.

```
- name: dt0 # readout request => readout (T1)
  in:       # inputs to dt0
  - sfp1.0  # sfp1.0 is multiplexed signal #0 on fiber port 1
  - sfp2.0  # same for fiber port 2
  out:      # outputs from dt0
  - mbs1    # send a T1 trigger to TRIXOR
  - sfp1.1  # send a trigger to both slave exploders ...
  - sfp2.1  #  ... as output signals 1 on sfp 1 and 2
  - lemo7   # lemo output 7 goes to the PEXARIA latch
  - scaler4 # measure the rate of T1 triggers in scaler 4
  label: Califa Master T1
```

Listing 1: The connections to the line named `dt0` (e.g. dead time 0) as used in CALIFA. This means that both fiber cable ports one and two can cause dt0 (subject to deadtime locking conditions). If dt0 is asserted, the lines mbs1 (on an LVDS connector going to the TRIXOR card), a lemo output (connected to the latch input of the PEXARIA card) are asserted in turn. Furthermore, a signal is transmitted towards the slave exploders so that the FEBEX cards can finalize the event. Finally, scaler4 is used to count the number of accepted T1 triggers. The label in the last line will be used when showing the count of that scaler.

directly, the connections are listed in relation to an internal line, which makes the path of the signals easy to follow. Also, the YAML file can be stored automatically in version control systems such as git. This can be used to get a reliable, time-stamped documentation of the changes in the trigger configuration.

The EXPLODER firmware also features eight scalers for the online monitoring of internal and external trigger rates as well as four independent logical pulser signal with adjustable frequency can be used as inputs. Each of the four fiber-optical links can interlace eight different trigger signals in either direction. While the timing jitter of these multiplexed signals is slightly worse, this allows to transmit many signals between different EXPLODERs over a minimum amount of lines. In the current configuration all four SFP fiber interfaces of the module are operated to transfer signals from different slave EXPLODERs to a central "overlord" EXPLODER used as a central trigger distribution system for CALIFA.

As shown in figure 26 this allows to connecting any number of different crates filled with FEBEX digitizers with the same latency to a single point in a star like configuration.

5.5 A scalable configuration system

Like the text based configuration for the FEBEX slow control parameters described in section 4.3, the YAML configuration allows to effectively manage complex configurations consisting of multiple EXPLODERs. Unlike the trigger matrices in the web interface used previously, the connected inputs and outputs can be labeled with text comments for documentation purposes. Version control can be used to make retrieval of historical configurations easy.

This opportunity is exploited for intricate configuration settings described in the following

section.

5.6 Multi-EXPLODER configuration in CALIFA

In CALIFA, one DAQ PC is used to read out four FEBEX crates.

Previously, two EXPLODERS connected using a fiber optical link were used for that: a *master* EXPLODER which makes the readout decision and is connected directly to the PC and a *slave* EXPLODER which relays readout requests from its two FEBEX crates to the master and sends readout triggers from the master back to its FEBEX crates.

This creates a complication, as the RO trigger will arrive earlier for the FEBEX crates connected to the master EXPLODER than the ones connected to the slave EXPLODER. This causes an offset between the time calibration of the FEBEX modules from crates connected to the master and slave.

One way to solve this would be to use three EXPLODERs per PC: one master (connected only to the TRIXOR card of the PC and the other EXPLODERs) and two slave EXPLODERs connected to the FEBEX crates and the master EXPLODER.

The left handed side of figure 27 shows such a setup. If disjoint circuitry is used between the master and slave EXPLODER, then the master and one slave can be combined into a single physical module. As master and slave part are still logically separate, the loop between SFP3 and SFP4 connects these parts with a timing delay similar to the one between two physically separate modules.

5.7 Two Stage Time Stamping

All data recorded by R^3B readout electronics – including CALIFA – is timestamped with White Rabbit Time Stamps (WRTS). This is a key design decision for the R^3B DAQ. It allows different detector systems to be read out out independently which immensely improves scalability The data streams from multiple detectors can then be *merged* into a single stream ordered by their time stamp in software. To reconstruct the full physics events, detector subevents from multiple sources whose time stamps fall within a given coincidence window are then combined in a process called *time-stitching*.

An asynchronous multi-event readout – like the one in CALIFA – implies that the readout blocks can not be directly assigned to physics events. An obvious approach is to also assign each hit in CALIFA a time stamp and then merge the readout blocks from different channels by their time stamp.

As the FEBEX3 platform was not originally designed with multi-event readout in mind, there are some additional technical issues. For KINPEX based readout systems, the GSI standard solution is to have a WRTS receiver card called PEXARIA. This card features a trigger input which will latch the current WRTS when received. The f_user then reads out the PEXARIA as well as the KINPEX.

While the PEXARIA solution works effortlessly in single event readout mode, for multi-event readout mode, things are not quite as straightforward. As the FEBEX modules do not have direct access to the white rabbit time, which is used to synchronize different detectors in R^3B, but only to a hardware clock count called “FEBEX time stamp”.

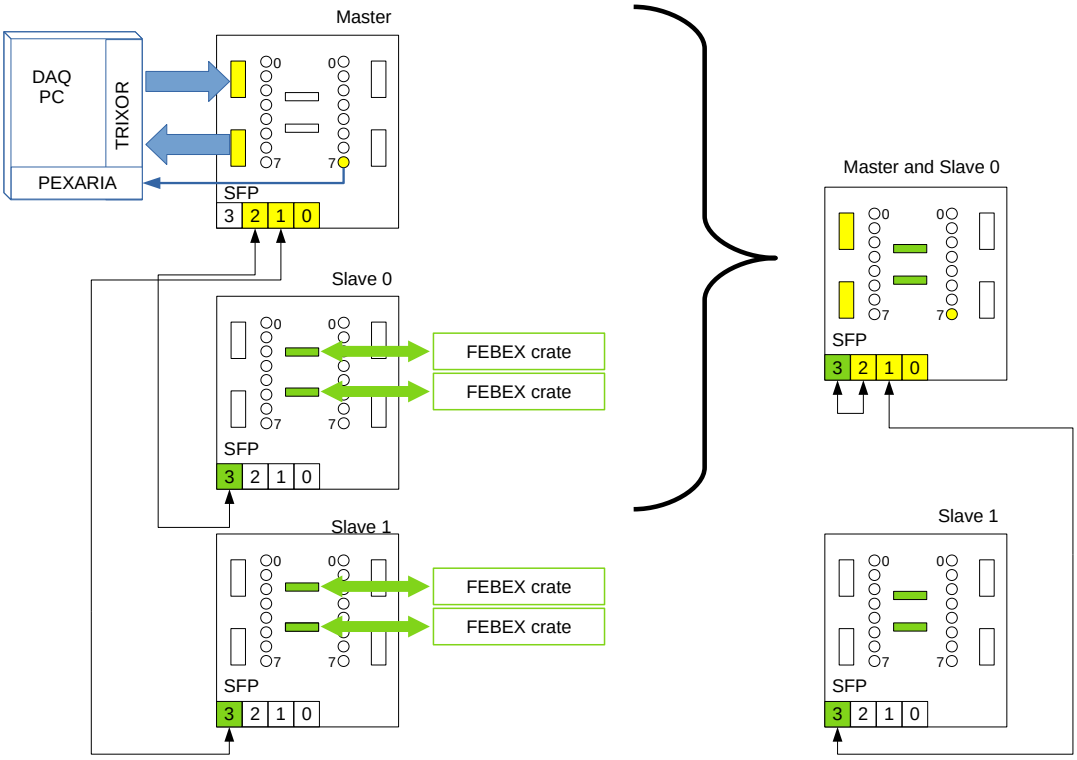

Figure 27: *Left*: a master EXPLODER (connected to the PC) and two slave EXPLODERs (connected to two FEBEX crates each). *Right*: By combining the functionality of the master EXPLODER with a slave EXPLODER, an equivalent two EXPLODER setup can be constructed.

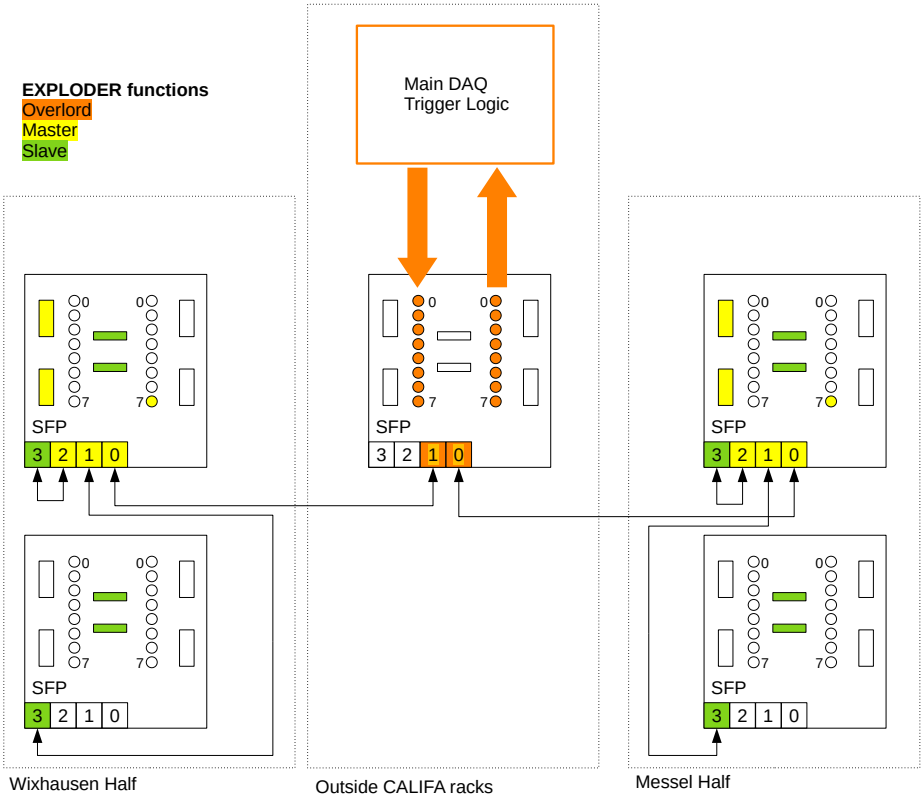

Figure 28: The CALIFA trigger setup as of 2022. The two sides of CALIFA each run a master and slave EXPLODER as described in figure 27. Both master EXPLODERs are connected via fiber optical link to an *overlord* EXPLODER located near the main trigger logic unit of the R^3B Main DAQ. Trigger signals can thus easily be received by and send from CALIFA.

When the readout trigger arrives, the current FebexTS is also stored in the readout buffer. As this FebexTS is coincident with the WRTS recorded by the PEXARIA, this allows interpolating between two WRTS based on the FebexTS and thus assigning a WRTS compatible time stamp to each hit.

This is accomplished by the CALIFA eventbuilder[45] which is based on a modified version of ucesb. As the MLVDS lines typically have some finite bit error rate, the time stamp clocks between modules can slowly drift apart while the DAQ is running. Thus, the interpolation is done on a per module basis. While the interpolation itself is not very computationally expensive, merging 1024 channel buffers – each containing up to 200 hits – is rather expensive. With a Gigabit Ethernet link and a hit size of 56 bytes, the maximum incoming rate is around two million hits per second – more than a single process thread can handle. Ways to increase the throughput by distribution chunks of data to multiple processes and merging the streams afterwards are currently under investigation.

Another related issue is the fact that while other detectors typically emit between zero and one hits per minimum bias trigger, the output of the CALIFA eventbuilder is not timestitched and may thus include many hits for one gamma cluster, and even hits not related to any beam ion from natural background radiation.

As discussed in section 5.3.2, external trigger validation can be used to lower the amount of unrelated hits recorded by CALIFA by only recording hits coincident with a central R^3B event trigger from in-beam detectors.

5.8 White Rabbit Time Stamps (WRTS) in R^3B

The traditional approach to managing readouts of multiple detector systems such as CALIFA, NeuLAND, ToFD, Fibers would be to place all of them in a single dead time domain and provide the same event trigger to all of them. Using such a setup, all detector elements will get the same number of triggers, so the readout system can just collect one event from each subsystem and join them together in one event. No additional sorting is needed and the event header already contains information from all subsystems. The price to pay consists of the creation of a global dead time domain in which all detectors have to wait for the slowest one. Some extra deadtime will be needed to compensate for long cables and to avoid glitches. Such a setup also lacks the flexibility to easily add ancillary detectors with a different readout scheme.

To overcome these limitiations, one might want to use multiple independent dead time domains instead of a single one. However, now data from multiple independent sources must be synchronized. A reasonable way to accomplish that is to establish a global timestamp.

The White Rabbit Project by CERN and GSI provides a way to establish a Nanosecond precision timestamp using specialised fiber optical ethernet links with custom switches and endpoints, and was selected as the timestamp solution for R^3B. White Rabbit timestamps count the number of nanoseconds which has passed since midnight 1970-01-01. Unlike UNIX epoch, which makes special exceptions for leap seconds so that the time format is exactly convertible to the UTC format commonly used on Earth, WRTS more reasonably counts leap seconds just as any other seconds, just like the *temps atomique international* (TAI). This means that basic

[45] Sometimes called "häcksler".

arithmetic can be used to calculate time differences. As WRTS used an unsigned 64 bit integer, there will be eventually a year 2554 problem.

For PC based systems, the GSI *Experimentelektronik* has developed a white rabbit endpoint PCIe board named PEXARIA[46]. It is connected to the White Rabbit network by a single bi-directional SFP+ and has a trigger input. When a trigger is sent to that input, the current WRTS is latched and can be accessed during readout.

As discussed in the following section, subevents from all different dead time domains are tagged with their respective WRTS during readout and can thus be easily merged into a single time ordered data stream. The resulting data can then be time-stitched. This means that subevents within a certain gate (usually a few microseconds) are assembled to a single event suitable for data analysis.

5.9 Timestitching and eventbuilding

In CALIFA, the crystal hits are generally recorded independently. A **timestamp** is attached to each recorded crystal hit This means that after the collection of events by the readout electronics an additional software has to take care that detector hits are sorted and combned in a separate stage. This is performed on a separate computer.

During a process called **merging**, hits from multiple detectors are sorted chronologically by their timestamp and inserted into a single event stream. Merging is accomplished either by using an ucesb [19] empty unpacker with the `--merge` option or by the drasi's [47] lwrocmerge. The former is more suitable for simple merging operations between a few data sources, while the latter is a full-fledged merging system featuring automatic disabling and re-enabling of event sources as well as automatic validation of the synchronicity of time stamps from various sources.

To combine these single detector hits into a physics event, another process called **timestitching** is needed. This works by considering the timestamp of the earliest unprocessed hit, and opening a window of a fixed size Δt_W, which is typically about $2\,\mu$s. Any hits within this window are removed from the input and packaged into a single event. Then the process is repeated for the remaining (e.g. later) hits.

This works well for typical experiments, where the event rate is much lower than the inverse of Δt_W. However, for the AmBe source, the rate of (independent) events was estimated to be one per $2.2\,\mu$s.

Naturally, this means that by using a normal Δt_W, a lot of random coincidences would end up in the same event.

Physical limits for the minimum length of Δt_W are given by the signal rise time of the CsI(Tl) scintillators, its signal to noise ratio and corresponding filter algorithms.

Hits with very low energies are generally recorded at later times than hits with higher energies. This effect is called **trigger walk**. In our FPGA firmware (see section 3.7), the signal for the trigger branch is shaped using a trapezoidal filter which is set to average over 24 samples (480 ns) and differentiate over a `discr_gaptime` 50 samples ($1\,\mu$s). For a very high amplitude signal, the output of the trapezoidal filter might reach the threshold as soon as the first sample

[46] `https://www.gsi.de/work/forschung/experimentelektronik/digitalelektronik/digitalelektronik/module/lwl/pexaria`

[47] `https://git.chalmers.se/expsubphys/drasi`

of the rising edge of the signal enters the filter. For such a signal, the trigger will be early and there will be very little jitter.

On the other hand, the maximum response of the trapezoidal filter is reached `discr_gaptime` -1 samples after the rising edge enters the filter, just before it starts being subtracted from the output. If the energy deposit is so small that the threshold is barely reached, it will be reached almost a microsecond late. Additionally, noise can become important in deciding if and when the threshold is reached. Thus, the trigger efficiency as a function of the energy is not a step function and low energy events have a higher jitter than high energy events.

The fact that both the mean delay and the jitter are energy dependent can be taken under consideration for time stitching. First, the systematical trigger walk is corrected, then, each hit is given an energy-dependent window to account for its jitter. Hits with partially overlapping windows are aggregated into an event.

Using these techniques, the detection of low-energy γ rays could be improved using the EDTW methods described in section 5.10.1.

5.10 Discussion of the event building methods

Ideal timestitching assembles individual hit to actual physics events using a function `SimilarTime(X,Y)` which obeys the axioms of an **equivalence relation** (e.g. reflexive, symmetric, transitive). In practice, for any hits[48] X, Y, Z we would like:

- sensitivity: high probability that hits from one physics event are grouped together
- specificity: low probability that hits from different physics events are grouped together
- transitivity: `SimilarTime(X,Y)` and `SimilarTime(Y,Z)` implies `SimilarTime(X,Z)`
- independence from other hits: `SimilarTime(X,Z)` should not depend on on the existence or properties of any third hit `Y`.
- conservation of hits: every input hit should be grouped into exactly one output hit.[49]
- independence of hit position: for timestitching, only the time stamp and the hit energy should be considered, so the spatial distribution of hits will not become biased

When the mean time between physics events is much larger than the timestamp jitter, these conditions are very easy to fulfill, so the actual algorithm matters little. In the experiment discussed here the time jitter is only one order of magnitude smaller than the mean time between events.

Sensitivity is all about preserving the integrity of the signal. If a cluster of several hits is split into different time stitched events, it will result in two wrong cluster energies instead of one correct one.

[48] "hit" here is meant to include all the information associated with the hit in question: first and foremost, its timestamp, but possibly also its energy

[49] Algorithms which violate this property will not be considered.

Specificity means keeping the background low.

A lack of transitivity introduces the dilemma of a hit Y matching both an earlier hit X and a later hit Z, which are clearly separate. This leaves us various bad options to choose from: we could always group the hit with either (typically the earlier), we could make a choice based on other criteria like energy, or we could even duplicate the hit.

The independence from other hits property comes from the need to estimate random coincidence rates. In principle, we can try to estimate the rate of random coincidences between clusters with energies E_0 and E_1 just from the individual rates of clusters, a time window and the run duration. In practice, many methods will mean that additional hits nearby will influence the efficiency of the cluster reconstruction: non-transitive methods might result in the cluster being split, while transitive methods might add unrelated hits which distort the cluster energy.

The conservation of hits property forbids hit duplication. This is helpful because otherwise, the spectrum would be distorted: lower energy hits – which have a higher timestamp uncertainty – would often be part of separate groups of hits.

The independence of position property rejects the approach of combining time stitching and clustering into one step. After all, when we are in the position to assign a low energy hit Y to either hit X or Z, it would be very tempting to look on whether we could form a nice cluster with either one of them. The downside would be a subtle distortion of the grouping in clusters compared to the grouping between clusters.

For technical reasons, these methods are not implemented directly in UCESB. For beam experiments where in-beam detectors define the timing, CALIFA is configured to record data within a generous coincidence window defined by these detectors and is first stitched using a large window. In the analysis software, individual hits can then be filtered based stringent constraints with regard to time difference before the stitching stage. For pure standalone experiments such as the AmBe source measurement presented in this thesis, the preferred setup feeds non-stitched data into the analysis. A special class then aggregates the events over multiple iterations of the event loop and emits them back to the following stages when it is certain that no more hits should be stitched to the current physics event.

5.10.1 Time stitching in triggered mode

In a typical experiment run of R^3B the central trigger system provides a low level event tag with its high precision white rabbit time stamp. In this configuration the time stitching is straightforward. All detector hits in a time window of $\Delta t_W \approx 4\,\mu$s around this time can be grouped together with regard to the event trigger.

Depending on the beam intensity and background rate, the CALIFA hits recorded under these conditions can still be contaminated with random coincidences. Here, energy dependent cuts can improve the signal to background especially for high energy hits which have low timing jitter.

Fixed time window from earliest hit (FTW)

Here the earliest unprocessed hit at time t_0 is used to create an interval $[t_0, t_0 + \Delta t]$. All hits in this interval are grouped into one event.

This method is understood to be commonly used, but fails the independence property: the earliest hit has a pivotal role. If we only consider transitivity given a fixed earliest hit, it technically holds with the interval being one equivalence class. In the more intuitive setting where we do not take the earliest hit as given, transitivity fails as well.

Hit energy dependent time window around pivot element (EDTW)

As discussed above, the jitter of hits with a high energy is lower than the jitter of low-energy hits. An obvious idea is to use this information to improve timestitching. In the first step a pivot X_0 element is selected. A typical choice would be the hit with the highest energy. Then, all hits Y fulfilling $\texttt{SimilarTime}(\texttt{X}_0, \texttt{Y})$ (which will depend on the energy of either hit) are grouped together.

This approach has a few advantages over the previous one: by exploiting the energy dependence of the time jitter, it allows the improvement of both the sensitivity and the specificity. Using high energy – e.g. low jitter – hits as pivot elements also helps in that regard.

However, both independence and transitivity do not hold, as the pivot element has a *pivotal* role in linking otherwise unrelated events with earlier and later timestamps than itself.

Using the highest energy hit as a pivot is a double edged sword: Apart from making sure that high energy clusters will be reconstructed with optimal efficiency, it also distorts the reconstruction efficiency of lower energy clusters around random hits with high energy, as these will sometimes grab parts of the cluster. This makes it very difficult to estimate background rates.

Hit energy dependent time window, transitive hull (EDTW-TH)

Starting from the earliest hit, add all hits within its time window, then, recursively, add all hits within their time windows.

By construction this method will be transitive. It will also have a good sensitivity. Its specificity will be poorer than the previous one. Lower energy links can link unrelated clusters. With bad luck, a long chain of hits within each being in the previous ones time window might be packed into a jumbo event orders of magnitude longer than the window length.

Random coincidence estimation will have to take into consideration the effective window length of either of the two energy under consideration, which will probably not be derivable from first principles.

Fixed time slices (FTS)

A fixed time window ΔT is chosen. All events are grouped by their timestamp divided by that time window, rounded down: $\left\lfloor \frac{t}{\Delta T} \right\rfloor$.

A tyrant's approach to timestitching: arbitrary, top-down, prescribing fixed rules with no regard of the data involved. Neat clusters ripped apart by unyielding borders, to be grouped with unrelated hits almost ΔT apart instead. Low sensitivity (especially with short ΔT), low specificity (especially with long ΔT). Thus ΔT should be chosen dependent on the requirements of the specific experiment.

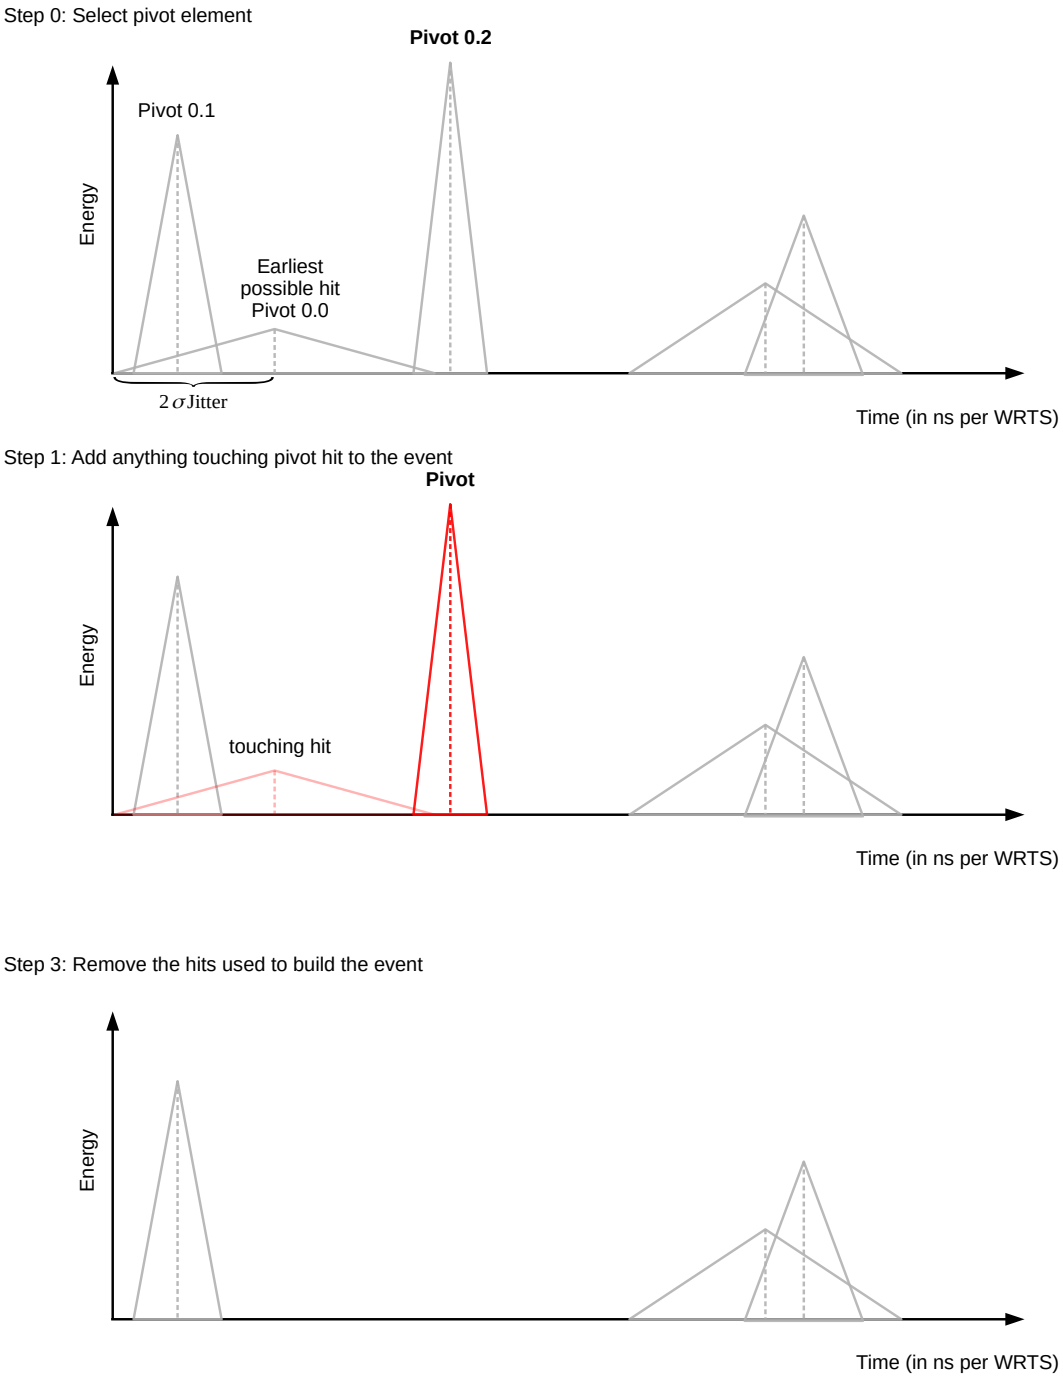

Figure 29: A cartoon of the EDTW timestitching method. The energies of the hits determine the jitter, the triangular shapes are for illustration only, inclusion in step one is determined by the overlapping time intervals.

By forcing the data into the given simplistic form with no regard for its inner structure, the result also make it very easy to reason about rates of random coincidences, as will be shown in the analysis part of this work, *Testing of uncorrelated background estimates* in section 8.3.1.

Transitivity is given, as is independence: every hit is sorted independently of previous or following hits.

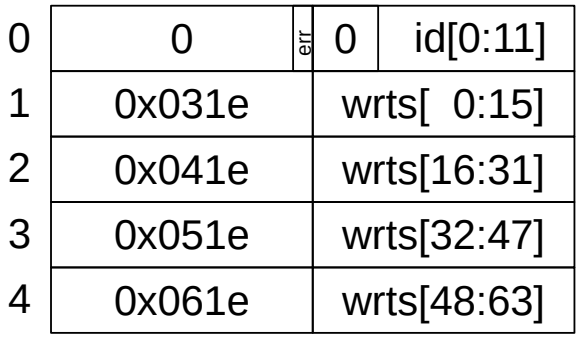

Figure 30: The format used to store a white rabbit time stamp within the LMD file format. The WRTS block consists of five words of 32 bit width. The first word contains an error indicator bit and a 12 bit identifier unique to the system which generated the timestamp. The following four words contain a fixed marker in the upper 16 bits and two bytes of the 64 bit timestamp in the lower 16 bits.

6 Software processing of CALIFA data and clustering

In the final step of event recording, the events are written to an LMD file, which is the format used by MBS and drasi. Each CALIFA hit contains a WRTS followed by a single FPGA hit record.[50] Hits from all detectors are *merged* (see section 5.9) into a single data stream. Any record starts with a WRTS followed by the payload data for that deadtime domain. The WRTS in the LMD file are ordered and use the format described in figure 30.

6.1 Data processing pipeline

The next processing step, also already discussed in section 5.9 is **timestitching**. The goal here is to assemble multiple detector records occurring approximately at the same time into a single physics event. This is done by looking at the earliest unprocessed WRTS, then stitching it with all records arriving in a given time window (typically four microseconds) after it. Care must be taken to use an appropriate window: if it is to short, some jitter might mean that coincidential hits are treated as separate physics events, while too long a window will group multiple physics events together. Because of these trade-offs, timestitching is not performed online, but only during analysis. Time-stitching is commonly accomplished by using the ucesb empty unpacker with the `--time-stitch` option.

The timestitched data (still in LMD format or an equivalent stream) can then be processed with an UCESB unpacker[20, 19]. Such an unpacker converts the memory blocks of FPGA output into a PAW ntuple or ROOT TTree with appropriately named fields. Also, the ucesb unpacker takes care of **mapping** the Febex channels to so called crystal IDs[51]

[50] For technical reasons, sometimes the interpolation is done offline and the readout blocks are written to file instead, in which case the LMD file first has to be converted into another LMD file by applying the FebexTS to WRTS interpolation.

[51] The number of ways to enumerate channels for a detector is approximately equal to the number of people

The output of the unpacker is then read with a software package called R3BRoot. For CALIFA, the pathway of the data starts with R3BCalifaFebexReader[52], which converts the data read into a ROOT TClonesArray (TCA) [53] of R3BCalifaMappedData[54].

The slow and fast components recorded by the firmware are treated analogously to the energy. They are calibrated with the same calibration parameters and also summed up during clustering.

6.2 Application of the calibration in R3BRoot

The TCA of R3BCalifaMappedData is then read by a class called R3BCalifaMapped2CrystalCal[55] to produce a TCA of a class R3BCalifaCrystalCal by scaling the energy, N_f and N_s using a linear function $cal_{ch}(x) = m_{ch}x + c_{ch}$. Also, the overflow flags are checked. Any data which is invalidated by overflows is assigned to the IEEE 754 constant NaN (not a number)[56].

6.3 Gamma clustering in CALIFA

Due to the need to reconstruct the azimuthal angle of photons for doppler reconstruction, CALIFA features a high granularity. This means that most high energy photons will not be absorbed in a single crystal. COMPTON scattering and e^+e^- pair production are both effects which make it likely the full energy of a high energy γ ray will not be deposited in a single crystal but be spread over a cluster of adjecent detector elements. The relevant effects and their crossections are illustrated in figure 31.

However, in total absorption spectroscopy, the energy of the photon can only be reconstructed correctly if all the energy of the photon is deposited within the detector. Thus the energy of a gamma is measured using **addback**: energies detected within a certain time window in a neighborhood are summed up to arrive at the total energy. Roughly speaking, the clustering algorithm starts by selecting the crystal hit with the highest energy. Anything in the

working with this detector. With the advent of dual-range readout, the crystal IDs were extended: crystal IDs between 1 (sic!) and 2432 refer to the high gain readout of that channel, while a crystal ID $n > 2432$ refers to the low gain readout of crystal $(n - 2432)$. This means that the ucesb unpacker has to know the range that the preamplifier was running in.

For the future development of R3BRoot, I would like to point out that the Chinese Remainder Theorem allows for encoding k different properties n_i ($0 \leq n_i < p_k$, p_k pairwise coprime) in an integer N with $0 \leq N < \prod_i p_i$. Using $p_0 = 2437$ for the crystal ID, $p_1 = 2$ for the range, $p_2 = 3$ for spill-condition and so on, this would allow to store many pieces of information in a single 32 bit number. I would assume that this would buy us some years before anyone understands the schema, at which point other results of number theory might inspire additional obfuscation efforts.

[52]That class is not named r3b::califa::FebexReader because we strive to emulate the shining example of software craftsmenship that is ROOT – which uses prefixes instead of C++ namespaces.

[53]TClonesArrays are used throughout R3BRoot. They feature a decided lack of type-safety, arcane syntax and an access time of $O(n)$ to find the properties of any one hit by its crystal ID.

[54]Per C++ Core Guidelines, this class should really be a struct as it lacks any invariants. But R3BRoot prefers to follow the example of R. Brun and F. Rademakers over that of B. Stroustrup or S. Meyers.

[55]The proper name should be R3BCalifaR3BCalifaMapped2R3BCalifaCrystalCal.

[56]A case could be made for using IEEE 754 infinity instead: Any of $<$, $>$, $\leq$, $\geq$ involving NaN results in *false*. This means that `if (en>100) stuff();` is not equivalent to `if (!en<=100) stuff();`, which might be contraintuitive.

neighborhood is added to this hit. If more hits unassigned to a cluster remain, the process is repeated.

Different possible definitions of the neighborhood of a crystal have been implemented in R3BRoot. They generally differ in the shape of the neighborhood and take one or more parameters to set the size of the neighborhood.

Generally, the best side of the region depends on the event rate, background rate and the number of γ rays emitted in a single physics event. In a low background, no-noise environment with a rate of just a few events per second, any hits in the detector within a few microseconds are likely caused by the same nuclear reaction. If that reaction only produced a single photon, adding all of the crystals will yield the highest photopeak efficiency.

In reality, there is a substantial background from the natural radioactivity in the cave, other ions in the beam and electronic noise. So while a photon can get scattered back and end up in a far away part of the detector, two separate photons are a better explanation for such a pattern.

Selecting all crystals in a cone around the given crystal is probably the most obvious way. Other variants include selecting based on the physical distance of the front face centers of the crystals[57], or scaling the accepted solid angle with the energy detected in the crystal. Most of these take at least one parameter.

Future improvements of the clustering code would include energy dependent cluster sizes and taking into account the energy dependent time resolution to exclude or flag hits for which the time difference indicates that they are not actually in coincidence with the event seen by the time of flight detectors (e.g. LOS) or the time of the seed hit of the cluster (for source runs).

6.4 Technical implementation

The clustering/addback algorithm is implemented in the R3BRoot class R3BCalifaCrystalCal2Hit[58].

The first task is to filter the CrystalCal entries, rejecting any value below a certain threshold, normally 200 keV. Also, for any crystals with a dual range readout, one of the ranges has to be picked for clustering. This is done by the following steps:

- A `std::list<R3BCalifaCrystalCalData*>` of candidate hits is created empty.
- A `std::map<uint32_t, R3BCalifaCrystalCalData*>` is created to look up Data based on the crystal ID. This structure is more efficient.[59]
- An iteration over that structure is performed:
  - If there exists a corresponding entry for (crystalId+2432), then the energy of the low gain readout is used to pick the range to use. Typically, the low gain range is preferred above 15 MeV. The selected hit is then added to the list (subject to the threshold).
  - Otherwise, if the crystal is only represented in the data with a single gain setting, it is added to the list (again subject to the threshold).

[57] This is different from the cone because CALIFA is not spherical.
[58] For clarity, I advocate replacing Hit with Cluster.
[59] It allows lookups in $O(\log n)$ instead of $O(n)$.

The result is a list of unique hits above the threshold. This list is then sorted, in descending order, by the energy using the STL[60] sort.

Until the list is empty, the front element of the list is used to define a new cluster. Then, every element of the list which is a neighbor of the first element is removed from the list and added to the cluster. The method R3BCalifaCrystalCal2Hit::Match is used to determine the neighborhood relationship between two R3BCalifaCrystalCalData*.

From an algorithmic efficiency perspective, while both the deduplication and sorting of the list take $O(n \log n)$ steps for n entries, the last step is slightly more expensive: it scales as the product of the number of clusters and the number of entries in the list. In the worst case – n single hit clusters – that means $O(n^2)$. This could be improved upon by precalculating the neighborhood of every crystal once and then performing just one map-retrieval (at $O(\log n)$) possible neighbor. More complex algorithms might first sort the calibrated data points into a coarse grid, then perform the neighborhood search only amongst the hits in the nearby grid segments. As the number of crystals hit per event is small in any case an algorithmic improvement is unlikely to be worth the additional code complexity, however.

The cluster level provides the data class R3BCalifaHit TCA which contains the relevant information of clusters and ready for further processing:

- classifying clusters (photons, stopped nuclei, punch-through protons) using the QPID
- for punch-throughs, estimating E_{Kin} from ΔE
- Determining the energy for both γ rays (Doppler) and particles (Lorentz) in the frame of reference in which they were emitted

As these steps are very much depending on the individual experiment and other detectors in the setup, they are done in non-standardized ROOT macros by individual users.

6.5 The need to test CALIFA at high gamma energies

One design goal of CALIFA is the detection of high energy γ rays ($E \approx 10\,\mathrm{MeV}$) emitted from ions moving at relativistic velocities ($\beta \approx 0.8$) [9]. Depending on the angle between photon and the particle, such γ rays can easily reach 20 MeV in the lab frame of reference.

Figure 31 gives a mass attenuation coefficients for photons in CsI.. The important point here is that the process of pair production becomes dominant around 6 MeV.

While parts of CALIFA have been regularly tested with gammas in the 1 MeV range – such as from ^{60}Co sources, the effects of pair production have mostly been studied in Monte Carlo simulations so far.

Pamela Teubig et al [22] used the reaction $^{27}\mathrm{Al}(p,\gamma)^{28}\mathrm{Si}*$ with protons from the CTN tandem accelerator to generate gammas with energies of around 8 MeV. However, the two petals (128 crystals) used in that experiment were earlier prototypes with some light collection issues. Also, the outermost layer of crystals were very likely to have many hits where a 511 keV

[60] STL is the standard template library, it includes type-templated container classes used by most C++ programmers outside the ROOT community.

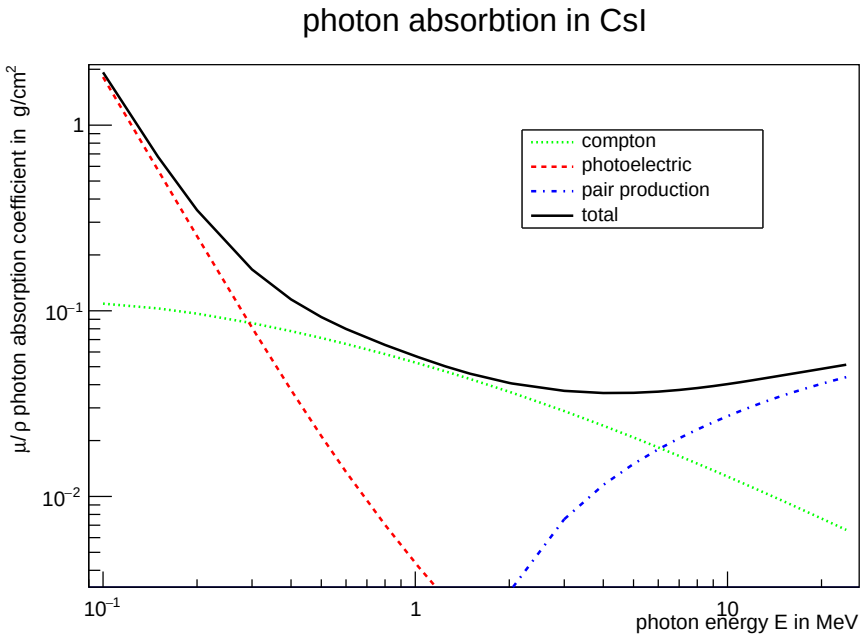

Figure 31: The dominant mass attenuation coefficients for photons in the 100 keV to 20 MeV range. Data taken from the XCOM [21] database.

photon escaped the detector. Given the petal geometry of 4x16 crystals in that experiment, more than half of the crystals are in the outer layer.

After the beam times S444, S473 and S454 at GSI in early 2019 there was a chance to test CALIFA with high energy photons again. As this test is one central part of this thesis, it is discussed separately in chapter 8.

To first understand the influence of the different energy dependent interaction patterns in the highly segmented active material a full GEANT4 simulation was performed. Results are discussed in the following chapter.

7 Monte Carlo simulations

As a key part of this thesis, extended Monte Carlo simulations have been performed to understand the response of a complex, highly segmented detector system like the CALIFA calorimeter. The Monte Carlo method leverages (pseudo-)randomness to solve problems for which an exact solution is typically infeasible. Consider a γ ray traveling through matter which may randomly interact by e.g. Compton scattering, resulting in a secondary photon going in a different direction. The decision tree for this problem grows much to fast to analytically calculate the average energy loss distribution in the matter. Instead of considering all the possibilities at once, each possibility is selected according to its probability. For example, the simulation considers the interaction length for different processes the photon may undergo and selects an interaction type and point according to these probability distributions. Then it goes on to simulate any secondaries to this process, and so forth, until a lower energy threshold for all particles is reached. This is repeated for many events and with a sufficient number of iterations approximates the distributions of interest.

GEANT4 [23] is the *de facto* standard for Monte Carlo simulations regarding the interaction of ionizing radiation with matter. It is widely used in nuclear and high energy physics as well as astrophysics, medical physics or radiation protection. Traditionally, the user defines the geometry by instantiating primitives such as cuboids or spheres in a C++ method during program startup. The generation of primary particles to be simulated is also handled in a method of another class (the *event generator*) which is called for every event. Physical processes which are considered are contained in a *physics lists*. The user can either use a custom list or (more reasonably) use one of several well-tested models provided with GEANT4.

7.1 A simplified model

Seen from the target, CALIFA consists mostly of CsI, but not only. There are reflective foils covering the crystals, and there is a 400μm thick carbon fiber alvioli structure. Behind the crystals, there is a space filled with dry nitrogen which contains the APDs and their cabling. Finally, there is the aluminum casing.

Still, by both weight and volume, massive CsI is a good first order approximation for the inner part of CALIFA. So for getting an estimate of the behaviour of photons and particles in CALIFA, it is reasonable to look at the interaction of them with a massive block of CsI. Of course, one may miss the finer subtleties of the actual CALIFA structure: protons traveling just through the thin carbon fiber without touching the crystals, or a gamma cluster stretching both halves of CALIFA being impeded by the aluminum casing. But only a small fraction of events should be affected by such unlikely occurences.

For a conceptual study in the first place the most simplified model of the detector has been used. A cube of CsI with an edge length of one meter, spanning $x, y \in [-500\,\mathrm{mm}, 500\,\mathrm{mm}]$ and $z \in [0\,\mathrm{mm}; 1000\,\mathrm{mm}]$, surrounded by vacuum. Primary photons are created at the origin and travel in positive z direction right into the block.

The objective here is to study showers in CsI, not to replicate the calorimetric properties of a detector, thus one can simply use the precise information generated by GEANT4 even though there would be no way to record signals at this level of detail in any realistical detector.

In GEANT4, a run consists of a simulation of a series of events. Each event is simulated in steps. A step is a line segment which is part of a particles trajectory. It has a starting point, an end point and the amount of energy the particle has lost during the step[61]. The properties of the particle (e.g. its type and unique identifier) are also accessible. After each step, these pieces of information are written out in a simple JSON format[62].

For analysis, the energy contributions of the steps are then histogrammed according to the position along different axis. z is the axis of momentum of the primary particle, while $r = \sqrt{x^2 + y^2}$ is the radius from this "beam line".

While the distribution of the energy within the crystal is informative, to answer the question "what is a good radius for clustering", it is helpful to show the energy E_{mis} outside that radius r_0:

$$E_{mis}(r_0) = E_{prim} - \sum_{r=0}^{r_0} \Delta E(r) \tag{29}$$

This has the advantage that in the region of interest where most of the energy has been summed up corresponds to small numbers which can be shown in detail using a logarithmic plot, Finally, to compare the distribution for different energies, we normalize on the primary energy:

$$Q(r_0) = \frac{E_{mis}(r_0)}{E_{prim}} \tag{30}$$

For the z axis, $E(z_0)$ is used to calculate $E_{mis}(z_0)$ and $Q(z_0)$.

7.2 Simulation of photons

For photons, the average total energy deposited in our cubic meter of CsI is typically $1-2\%$ lower than the energy of the photon: there is always the possibility of the photon being Compton scattered backwards and escaping into the detector into the half space $z < 0$. Except for summing up all CALIFA crystals, no amount of clustering will catch these escaped photons. To avoid being distracted by that effect, we substitute E_{prim} with E_{tot} – the total amount of energy deposited, on average, in the cubic meter of CsI – in the formulas for E_{mis} and $Q(r)$ in this section. Thus, Q will go towards zero as z and r get large. This ignores all backscattering of primary or secondary particles and simply reflects the exponential decay/absorption visible as straight lines in the logarithmic representation.

Figure 32 gives $Q(z)$, which amounts to what fraction of the energy we can expect to contain given a certain crystal length. This is of course a bit academic: the crystals for CALIFA are already existing, their length is 220 mm for $\theta < 55.5°$ and $170 - 180$ mm for the rest of the forward barrel, and nobody is proposing to replace them with longer ones. Still, it might give a reasonable idea what fraction of non-contained energy could be reached.

[61] It is a bit unclear where exactly the particle has lost the energy. For a process which does not generate secondaries (e.g. ionization), the energy loss is assigned to a random point on the line segment.

[62] *JavaScript Object Notation* (JSON) is a human readable format which features key-value pairs and arrays, both of which can be nested. It can be written from C++ ad-hoc using iostreams with minimal overhead and can be parsed easily with many programming languages.

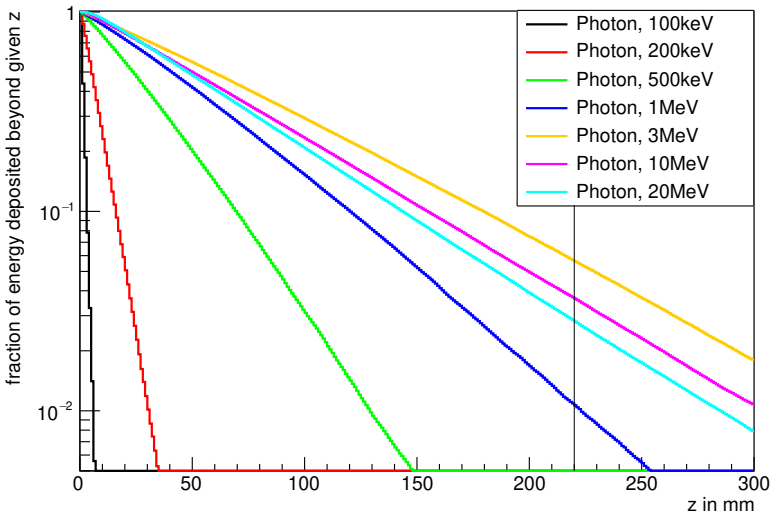

Figure 32: The fraction of energy deposited beyond a given depth z ($Q(z)$) for photons of different energies. The horizontal lines are the 5% and 1% threshold, the vertical line is the typical length of CALIFA crystals. In agreement with figure 31 photons in the 3 MeV region deposit a higher fraction of their energy beyond $z = 300\,\mathrm{mm}$.

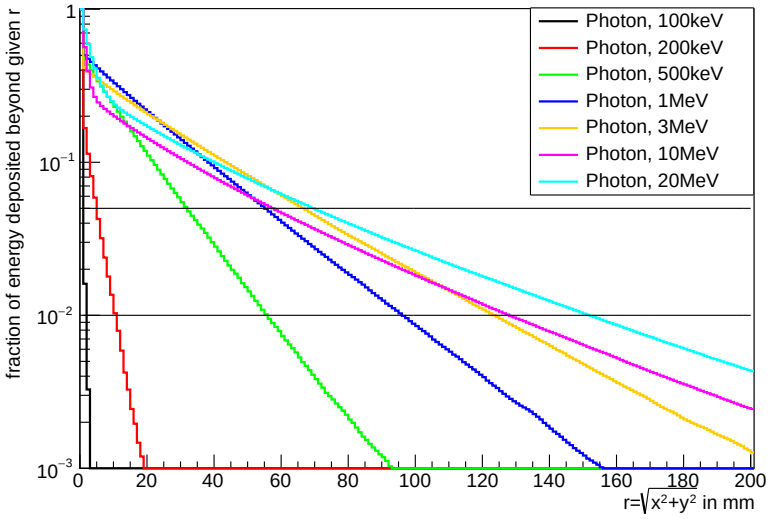

Figure 33: The fraction of energy contained within a cylinder with radius r around the initial photon trajectory ($Q(r)$) for photons of different energies. The horizontal lines are the 5% and 1% threshold. Curves have significantly different shapes for different energies of the incident photon.

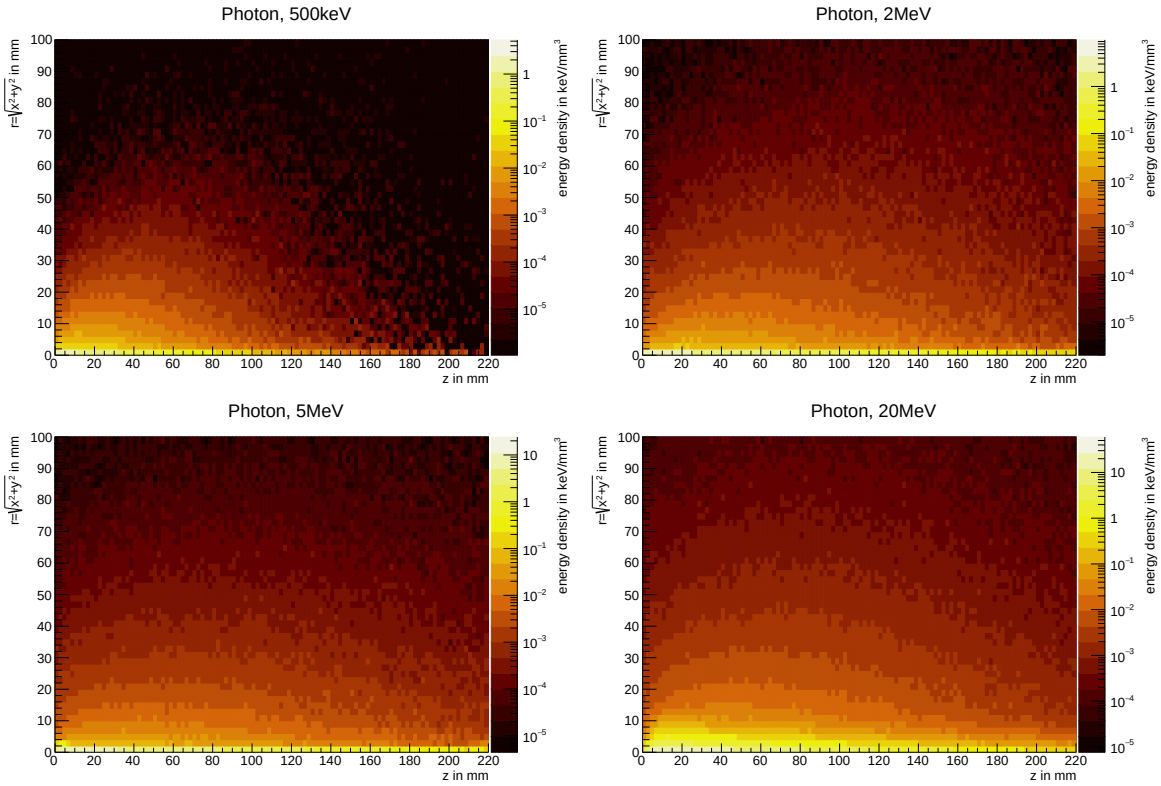

Figure 34: The mean energy density deposited by photons with 500 keV, 2 MeV, 5 MeV and 10 MeV in CsI. In these two-dimensional plots, the fact that $\frac{dA}{dr}$ is not constant is taken into account.

Of more interest is $Q(r)$ given in 33. Just like for $Q(z)$, we see a strong energy dependence. $r = 70\,\text{mm}$ is sufficient to capture 95% of the energy for any of the photon energies, while capturing 99% of the energy would require a radius of $r = 100\,\text{mm}$ for 1 MeV and $r = 155\,\text{mm}$ for 20 MeV photons.

Figure 34 shows the density of the deposited energy along both z and r for different photon energies.

In each case the maximum spread in r direction is reached after some centimeters in z direction. Thus, clustering for CALIFA should be based roughly on the inner edge length of the crystals.

Figure 35 shows the amount of energy absorbed in a CsI. In all cases considered, total absorbtion is the most likely outcome, with the escape of a Compton-scattered or annihilation photon appearing in a minority of events.

7.3 Containment probabilities

While the average fraction of energy deposited in a cluster is one way to grade clustering, it is not the best performance indicator. There is a big difference between the case where each

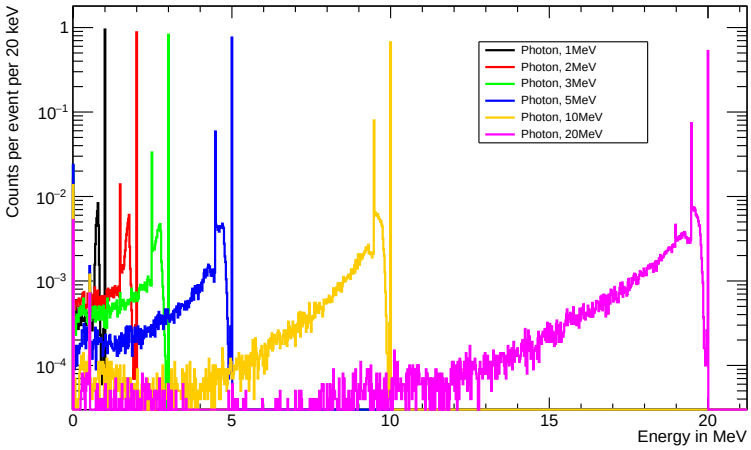

Figure 35: For photons of various energies, the amount of energy which gets absorbed in 22 cm of CsI. Note that this simulation does not include digitization effects which limit the resolution.

photon deposits 90% of the energy and the case where 90% of the photons deposit their full energy and 10% of them do not lose any energy!

As seen from figure 35, there are two clearly distinguishable options: either all of the energy is contained (and clustering will yield the correct energy), or there is a substantial amount of energy missing (and clustering will return a wrong energy value). Clearly the structures are determined by two effects, especially for the higher energies. Firstly the Compton edge of the first scattering separates the full absorption peak, and secondly the single and double escape peaks indicate the contributions from pair creation[63]. There is no need to split hairs over the treatment of the case where a few keV escaped because that case almost never appears.

Figure 36 visualizes the fraction of photons whose energy deposit will be outside the photopeak for various reasons. Compared with figure 32, one notices that while the fraction of energy lost to $z > 220\,\mathrm{mm}$ is maximal around $E \approx 3\,\mathrm{MeV}$, the fraction of events which have some energy deposit at $z > 220\,\mathrm{mm}$ is roughly constant in the 3–10 MeV region. This is not surprising because a higher energy photon can create secondaries in the 3 MeV range. Nuclear reactions includes those in which either ^{133}Cs or ^{127}I is transformed into a different nucleus, for example through (γ, n) or (γ, p) reactions. For the detection of photons discussed here, nuclear reactions within the detector are no concern for CALIFA, but incident nuclei will sometimes react with the target nuclei, thereby depriving the detector of the difference in binding energies.

For different clustering radii, one can now visualize how much the additional radial constraint exacerbates the containment losses. As shown in figure 37, the effects of the cluster radius become more important at higher energies. For energies below 1 MeV, 96% of the longitudinally contained events up to one MeV will also be contained within a clustering radius of 8 cm,

[63] The double escape peak plays a very minor role due to the detector geometry of 1000 mm × 1000 mm × 220 mm, which makes the escape of two back-to-back 511 keV photons unlikely. The double escape peak visible for the 20 MeV photon is likely caused by two pair creation events, each of which has a photon escaping..

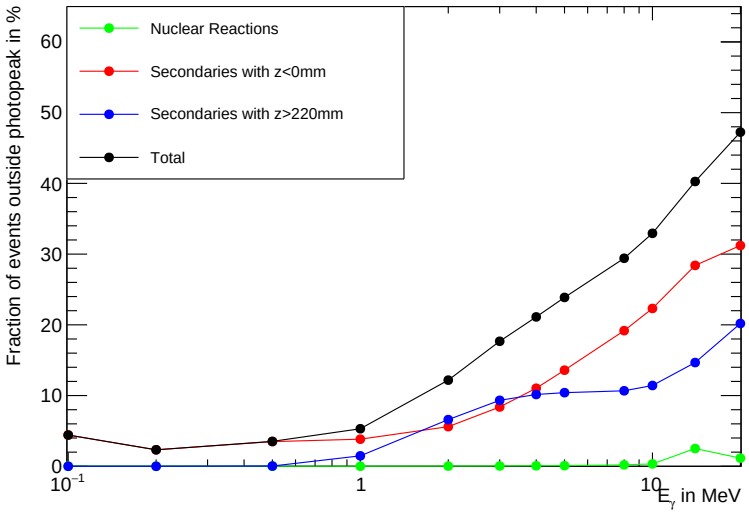

Figure 36: For various photon energies, the fraction of photons which is not contained in the photo peak for a 22 cm thick plane of CsI. Reactions are any events in which the energy of the incoming photon is different from the energy deposited plus the energy of any escaping particles. This happens when GEANT4 rolled a nuclear reaction such as (γ, p) or (γ, n) happening. The red graph represents secondaries lost because they travel into the half-space the photon came out of, while the blue graph represents photons whose secondaries penetrate the 22 cm. As multiple of there effects can appear in the same event, the total can be lower than the sum of the other graphs.

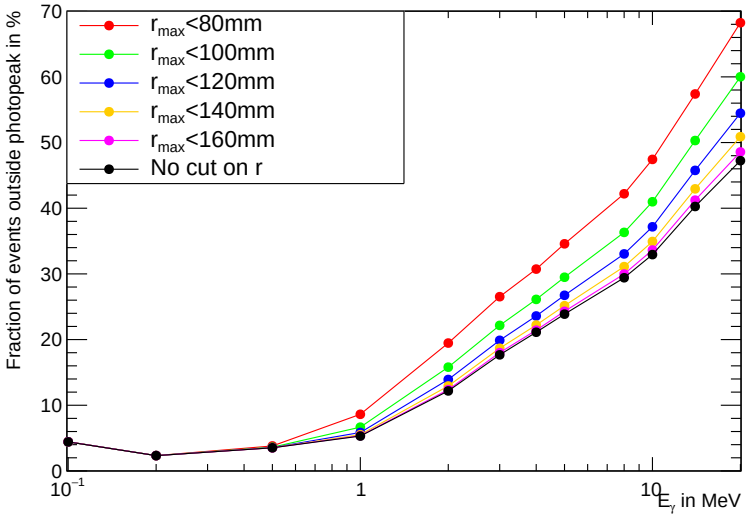

Figure 37: The effect of cluster radii on the containment probability. The black line is identical to the one in figure 36, while the colored line include the effect of having finite cluster radii. For example, at 1 MeV, the chance to contain a γ ray within 220 mm of CsI is 95%. The chance to contain it in a cylinder of length 220 mm and radius 80 mm is 91%.

while 14 cm will yield acceptable results even beyond the gamma energies CALIFA is likely to encounter. The dominant effect which limits the photopeak efficiency is the losses along the longitudinal axis which is very prominent at higher energies.

By histogramming the maximum r and z of any energy deposit within an event, and integrating over the resulting histograms, one can also visualize the effect of a larger variety of cut-offs for both z and r.

The effect of the detector thickness is shown in figure 38. To the degree that the first interaction dominates the mean free path, the penetration of photons should follow an exponential curve which appears linear in the logarithmic plot. As shown in figure 31, the photon absorbtion has a broad minimum in the 3–10 MeV range, so the containment chances in this region are similar. For 20 MeV, the penetration increases. Two facts contribute to this: secondary photons created by such γ rays (e.g. via pair production followed by *Bremssstrahlung*) can themselves be in the broad minimum region discussed previously, and nuclear reactions start to play a role. After a (γ, n) reaction produces a stray neutron, the chances for containment will be slim[64].

Given that CALIFA is already built and the length of the crystals could not be easily changed, the effect of the clustering radius as shown in the corresponding figure 39 is of more interest. Qualitatively, it seems similar to the previous plot. Quantitiatively, there is no plateau, but the penetration depth increases monotonically with the photon energy. This is unsurprising

[64]This is not a problem because after neutron knockout, the energy contributing to the scintillation will not correspond to the original energy of the photon in any case, so we might just write these off and get on with our lives.

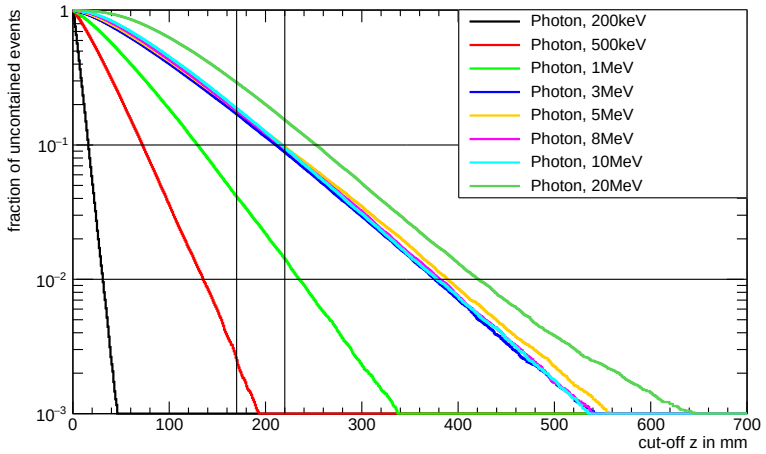

Figure 38: The probability of failing to contain a photon within a CsI detector of thickness z. In this plot, losses from $z < 0$ are not taken into account. The two black vertical lines indicate the 17 cm and 22 cm thickness used in CALIFA. The non-containment chances for photons with energies from 3 MeV to 10 MeV are quite similar and amount to 10% for the longer crystals, consistent with the blue graph in figure 36.

as the primary γ ray travels along the z axis and thus does not contribute to the radius. Instead, the radius is defined by the energies and emission angles of any secondaries.

If one takes into account the fact that the thickness of the detector is 22 cm, increasing the cluster radius hits sharply diminuishing returns for a cluster radius of 14 cm (figure 40) as any photon not contained in that radius appears to be also unlikely to be contained in 20 cm CsI in z direction. While figure 38 implied a similar penetration depth for all energies between 3–10 MeV, the asymptotic losses increase with the photon energy. This comes from unavoidable secondary photons created in e^+e^- annihilation escaping in negative z direction. As discussed below, at 20 MeV the impact of (γ, n) nuclear reactions becomes important.

As the clustering radius is obviously subject to experiment-dependent trade-offs (i.e. a higher photopeak efficiency has to be balanced against a higher probability to add uncorrelated hits or background), increasing the cluster radius towards the plateau would only make sense in low rate experiments.

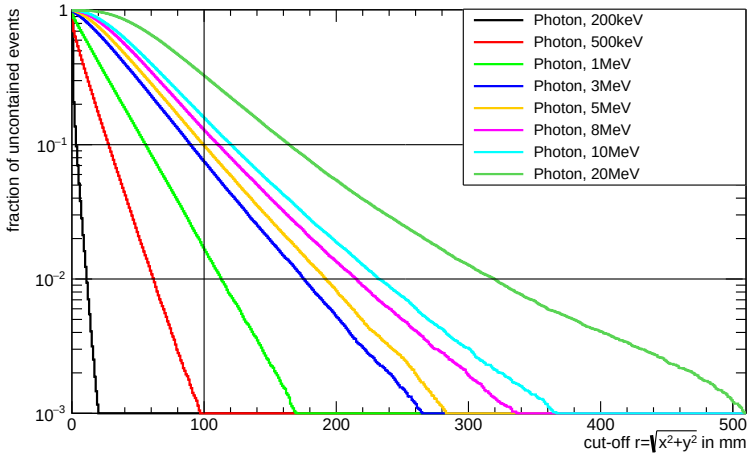

Figure 39: The fraction of events escaping a given cluster radius. Secondaries escaping due to the z axis boundaries of the detector are not taken into account.

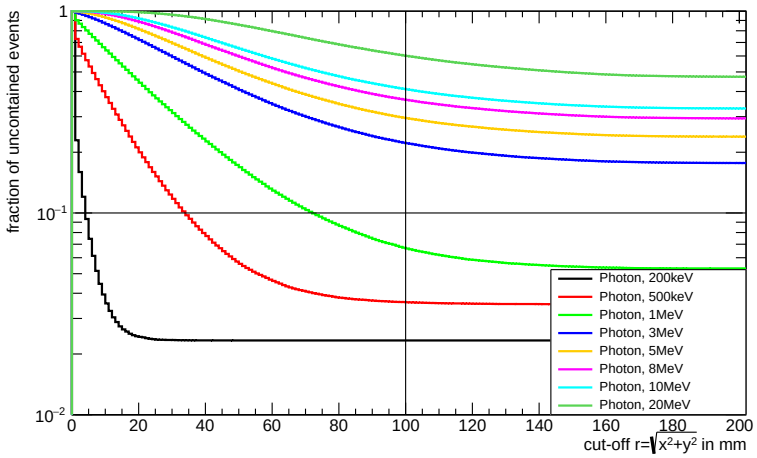

Figure 40: As the previous plot, but events not contained in $z \in [0\,\text{mm}; 220\,\text{mm}]$ are counted as non-contained. It becomes apparent that any gains beyond $r_{max} > 140\,\text{mm}$ are minimal as the losses through the z surfaces of the detector dominate.

7.4 A more realistic detector model

7.4.1 Regarding thresholds for energy detection

Due to the requirements regarding the Doppler reconstruction discussed in section 3.1, CALIFA is made from discrete crystals, and each crystal can only be read out if the energy deposit in it was above a trigger threshold determined by the noise in the readout chain, typically about 100-200 keV. If all of the energy deposits around $r = 80$ mm was from secondary photons, electrons with just a few keV, there might not be any triggerable energy deposits at all. However, neither photons nor electrons below 200 keV are very penetrating. Figure 32 included photons with such energies, they deposit 80% of their energy within a range of 10 mm . But also even a high energy secondary photon might not deposit much energy in a given crystal.

One method to cluster would be to look at the eight nearest neighbors of the crystal containing the first interaction of the incident photon[65], and pick then add the neighbors of any crystals with worthwhile energy deposits, recursively. The limit of this method is given by the fact that a secondary photon going out with a large scattering angle may well go through the thin dimension of one crystal without interacting and then interact in the crystal after that. The most likely photons to behave this way are the 511 keV caused by electron-positron annihilation. Emitted effectively isotropically in the lab frame, they are very likely to go into such inconventient directions. Such a photon has a 66% probability of travelling through 15 mm (a typical length of the crystals along $\vec{e}_\theta$) of CsI while depositing less than 200 keV in it. In fact, it is likely (p=0.55) not to interact at all.

7.4.2 The block model

To study these effects in more detail, the front center of the detector block was split into separate 21 × 11 crystals of 15 mm × 30 mm × 220 mm each. This means that the front faces of the simulated crystals are quite similar to standard crystals used in CALIFA, while also preserving the advantages of a simplified cubic model[66]. The surrounding volume is left as a bulk and used to catch the energy not deposited in the crystals. Figure 41 shows a visualization of the resulting geometry.

Two different position settings for the origin of the primary photon were used. In the **centered** setting, the particle travels directly along the z axis and hits the central crystal head on. In the **smeared** setting, the particle moves parallel to the z axis and hits central crystal at a uniform random point of its face.

7.4.3 Clusters

For clustering, any crystal with an energy deposit below 100 keV (or 200 keV) was rejected. The set of remaining hits was used to analyzed with regard to energy sum, multiplicity and cluster shape.

[65] In practice, one assumes that the crystal with the highest energy deposit is that hit, because it is impossible to know which crystal was hit first.

[66] While CALIFA is symetrical along the ϕ, the barrel shape means that there are offsets between the front faces of the crystals along the θ axis, leading to a jagged line.

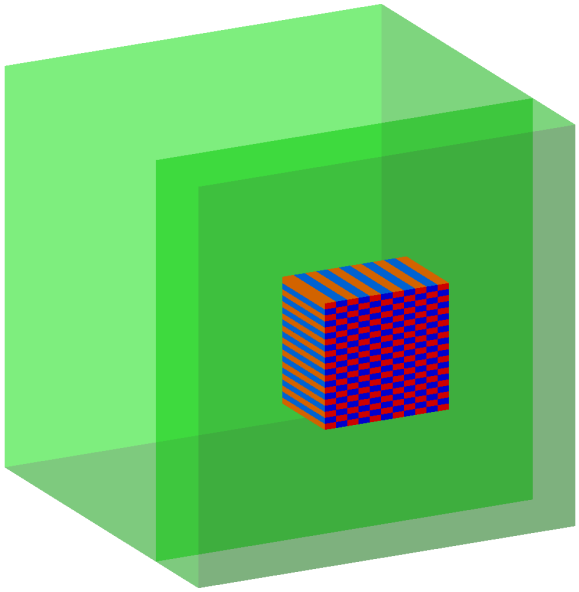

Figure 41: A simplified model for a granular detector. The red and blue boxes are crystals with front faces similar to CALIFA crystals.

The easiest way to categorize clusters is by their multiplicity N. Assuming that each hit i has a crystal position index $(x_i, y_i) \in \mathbb{Z}^2$, one can also calculate the mean position and the variance:

$$\begin{aligned} \bar{x} &= \frac{1}{N}\sum_{i=0}^{N-1} x_i \\ \mathrm{Var}(x) &= \frac{1}{N}\sum_{i=0}^{N-1} (x_i - \bar{x})^2 \\ &= \left(\frac{1}{N}\sum_{i=0}^{N-1} x_i^2\right) - \bar{x}^2 \end{aligned}$$

$\bar{y}$ and $\mathrm{Var}(y)$ are defined analogously. While one might weight the average by the energy by adding E_i as a factor under the sums and replacing the $\frac{1}{N}$ by $\frac{1}{E_\Sigma}$, the equal weight approach has the benefit that $N^2\mathrm{Var}(x)$ is an integer amount. By histogramming the two quantities $N^2\mathrm{Var}(x)$ and $N^2\mathrm{Var}(y)$, we can count which types of clusters appear.

Of course, these two variances do not uniquely describe the cluster shape[67]: mirroring a cluster along either axis will not change them. Likewise, exchanging the x_i (or y_i) of two hits

[67] Beyond the rather trivial $N = 1$

will not change variances. For $N > 2$, there exist some values of $N^2\mathrm{Var}(x)$ for which more than one multiset of differences $(x_i - x_0)$ would create that variance[68].

In appendix C, figures 61 to 64 show the probability of shapes with given mutiplicities obtained this way.

For these 10 MeV clusters, the probabilities of the various crystal multiplicities are:

Multiplicity	0	1	2	3	4	5	6	7	higher
Probability	0.013	0.034	0.131	0.225	0.233	0.174	0.102	0.0513	0.035

7.5 Missing energy

Figure 42 shows the *mean fraction* of the energy which is not captured in the crystals due to different effects. This is different from the containment probability which was used in previous plots. The reason for that is that once one starts to to apply energy thresholds cuts on crystals, the chances to lose some energy to these cuts are rather substantial. At 10 MeV with a 200 KeV crystal energy threshold, the chance to lose *some* energy to that cut are rather substantial[69]. But losing a few hundreds of keV out of 10 MeV does not mean that the clustering failed, it just indicates that the energy will be slightly (but systematically) wrong.[70] [71]

Having a clustering threshold of 100 keV will only result in a small loss of clustering efficiency if the cluster energies are above 1 MeV, while a threshold of 200 keV will suffice for cluster energies of at least 2 MeV. Around cluster energies of 15 MeV, a small fraction of events will miss some of the energy because the giant dipole resonance (GDR) of the detector nuclei is excited and they emit their energy through the evaporation of an undetected neutron.

7.6 Identifying the position of the incoming photon and protons

Identifying the position of the first interaction of a γ ray is important for two different reasons. For clustering, we need a provisional primary interaction crystal which defines the neighborhood considered part of the cluster. Identifying the polar angle of emission θ of the photon is even more crucial because the Doppler correction which needs to be applied is dominating the overall energy resolution of the calorimeter in the typical inverse kinematic experiments over a wide range of energy. This is e.g. why the front faces of CALIFA crystals have a typical length of 30mm along $\vec{e}_\phi$ but only 15mm along $\vec{e}_\theta$ (θ, ϕ denoting the polar coordinate system).

For the definition of the cluster neighborhood, the cluster position is usually taken to be the position of the crystal which saw the highest energy deposit. This is not always accurate: figure

[68] For N=3, both $(0, 0, 7)$ and $(0, 3, 8)$ have a variance of $\frac{98}{9}$. For N=4, $(0, 0, 0, 2)$ and $(0, 0, 2, 3)$ have the same variance: $\frac{27}{16}$.

[69] Per 31, the mean free path of 200 keV photons in CsI is about 1.8 cm. Given that the crystal dimensions are just 1.5 centimeters along the short axis, it seems likely that some low energy photons escape into a neighboring crystal which stays below the threshold.

[70] Applying a model for the energy dependent energy resolution of the CALIFA detector was beyond the scope of this simulation project.

[71] From figure 42, the difference of the 200 keV threshold curve (black) and the 0 keV threshold curve (green) is about 1% at 10 MeV, which amounts to 100 keV. This could mean that every other event loses one photon of almost 200 keV to a neighboring crystal.

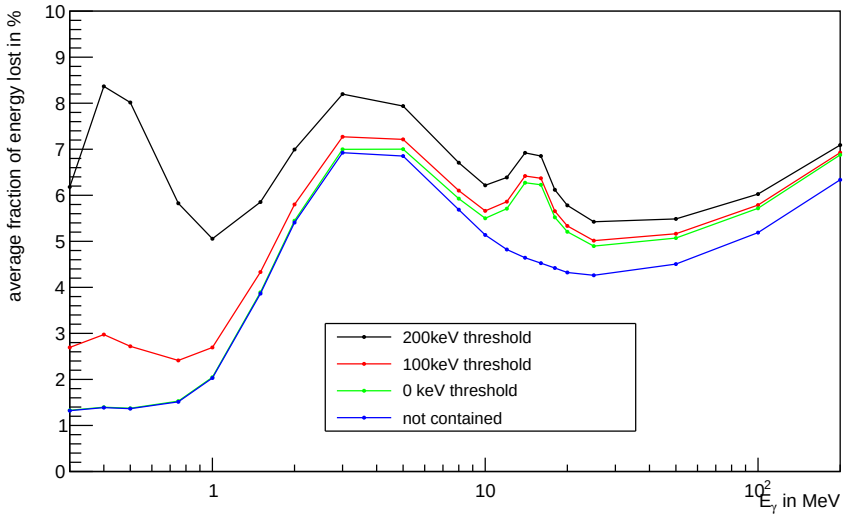

Figure 42: The average fraction of the original photon energy not detected in the crystals. This is caused by (a) secondaries being lost as they travel to $z < 0$ or (b) to $z > 220\,\text{mm}$ (both of these are included in the lowest blue line labeled "not contained"), (c) energy deposits below the trigger threshold or (d) nuclear reactions (at very high energies).

43 shows that even if the photon strikes the crystal in the middle, 25% of the 1 MeV photons will have a higher energy deposit in another crystal.

Different methods could be used to estimate the position after the cluster has been defined. Keeping the position of the crystal with the highest energy is a standard solution. The unweightened mean position of all crystals in the cluster is also simple, while the energy weightened mean position is a more stable estimate with regard of the influence of small fluctuations of the energy distribution within the cluster. For 1.5 MeV photons, figure 44 gives the difference between the actual position and the reconstructed position (i.e. the residual) along direction in which the crystal size provides the higher resolution[72].

In the region where Compton scattering dominates (such as for 1.5 MeV photons), a Compton scattered photon can travel through multiple crystals before eventually depositing all of its energy. While this will affect all three position estimates, for the averaging estimators, the initial energy loss in the compton scattering is at least taken into account while for the position of the crystal with the highest energy can be far from the crystal of the compton scattering indeed. This is why in figure 44, the standard deviation for the maximum energy crystal method is a bit higher than for the weightened mean approach.

However, for *most* events the maximum crystal energy approach is clearly better, even though it is sometimes far off. The question becomes if one wants to optimize for the worst case (Compton scattering at a low angle, most of the energy gets deposited in the neighboring

[72] As mentioned, in CALIFA this is $\vec{e}_\theta$, while in this simulation the x, y direction was arbitrarily mapped to the pitches of 1.5 cm and 3.0 cm.

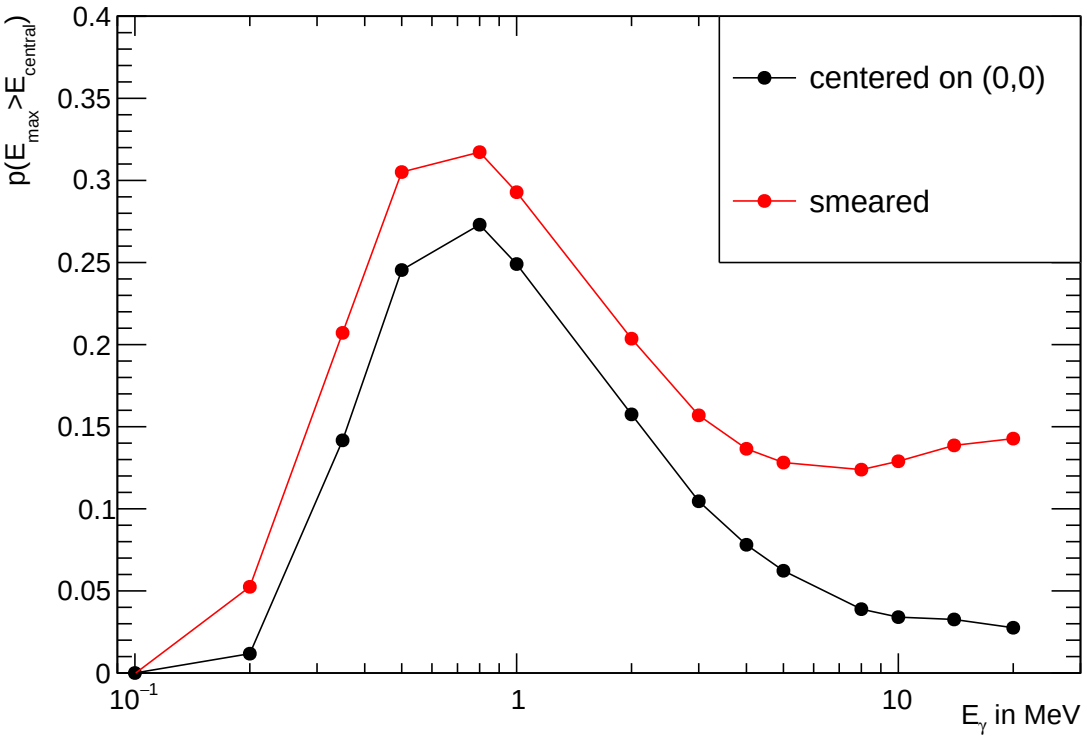

Figure 43: The probability of the crystal directly hit by the photon not ending up with the highest energy deposit. The black curve is the result if the photons are hitting the crystal in the center, the red curve is the result if the photons are allowed to illuminate all of the face of the crystal. For the question *how good is the position of the highest energy crystal as an estimator for the position of first interaction*, the black curve represents a lower bound (counting only the crystals where the that estimator differes from reality by at least one edge length), while the red curve can be seen as an upper bound.

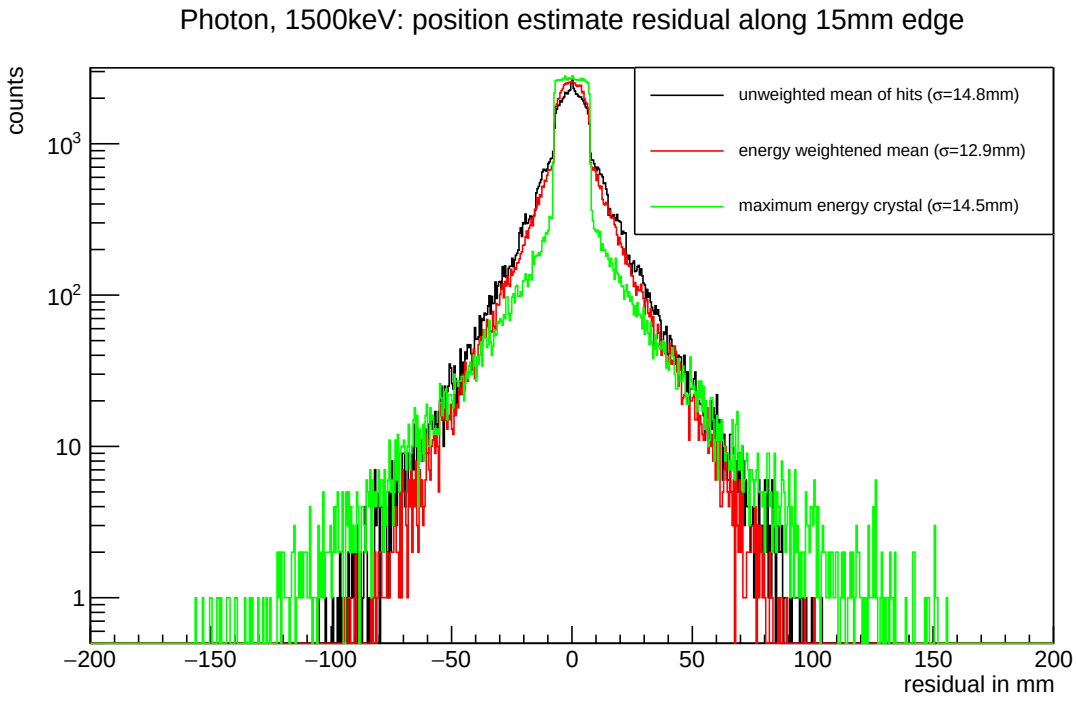

Figure 44: The residual of the position estimate obtained by various methods to obtain the cluster center ((smeared)). Note that the standard deviations here are calculated over the whole range of the histogram. When moving from σ to σ_{98} (as explained in the text), the ranking of the three methods changes. Under this metric, the best resolution is obtained with the maximum energy crystal ($\sigma_{98} = 10.2\,\mathrm{mm}$) followed by the energy weighted mean ($\sigma_{98} = 10.5\,\mathrm{mm}$) while unweightened mean still remains unfavorable ($\sigma_{98} = 12.5\,\mathrm{mm}$).

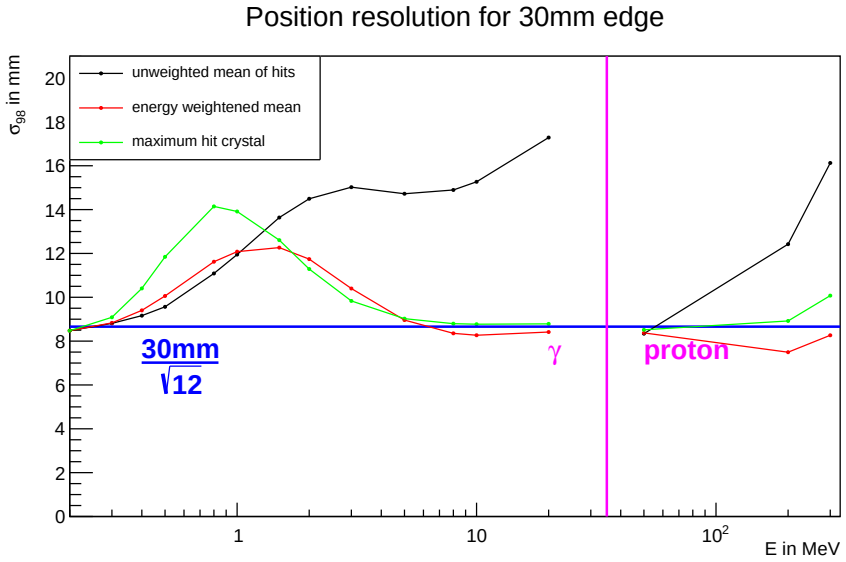

Figure 45: The position resolution in the direction of the 30 mm edge for position randomized photons and protons. The blue line is the theoretical resolution limit when all energy is deposited in the primary crystal – the resolution of a strip detector with a pitch of l is $\frac{l}{\sqrt{12}}$. Here, the energy weightened average position of the cluster offers competitive performance for every energy scale. Around 20 MeV, photons can excite collective modes in the detector nuclei which decay via γ cascades. This results in high crystal multiplicities and worsens the performance of the unweightened mean approach. For protons with 200 MeV, the energy weighted mean resolution is slightly better than the strip detector because the secondary hits provide additional information.

crystal.

To avoid being dominated by rare tail events, the following plots will calculate the residual only for the 98% events which are central, arbitrarily rejecting 2% of the events.[73] The resulting quantity will be denoted as σ_{98}.

Taking the position of the maximum energy hit will often result in picking the correct position within $l = 15\,\text{mm}$, but obviously does not narrow down the position within that crystal. If the correct position is always taken, the resulting distribution is a constant probability distribution, its standard deviation would be $\frac{l}{\sqrt{12}}$. The other methods can in theory have a better resolution, but only if the energy is spread somewhat evenly between the crystals.

Figures 45 and 46 give the position resolution for different methods, as well as the limits one would expect for single hits based on the crystal thickness[74].

[73] From the same 1.5 MeV simulation, it turns out that about 1.4% of the photons do not interact within the 22 cm of the crystal at all. Thus, I feel that rejecting a similar percentage of events as "probably not reconstructable" is still conservative.

[74] The fact that we only take the SD for the central 98% was not taken into account for these reference lines, this is why the outcome for $E_\gamma = 200\,keV$ is slightly below the line.

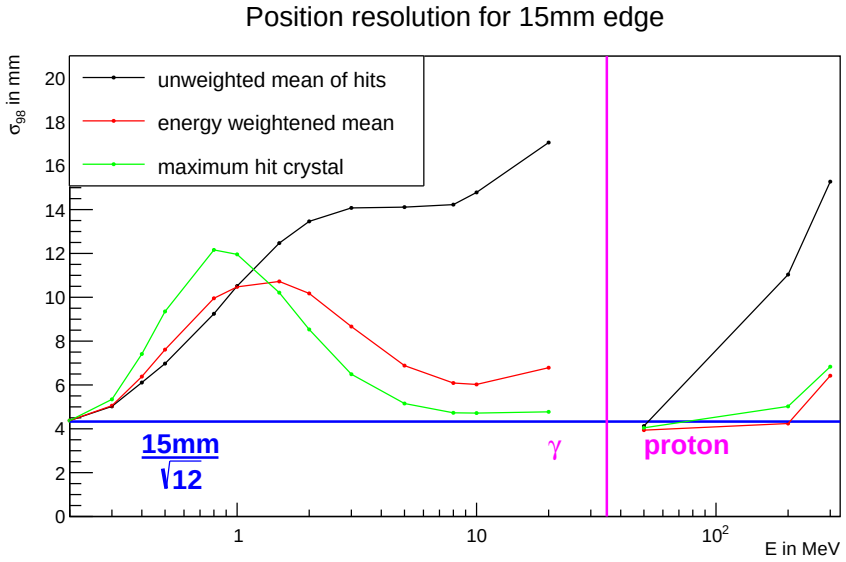

Figure 46: The position resolution in the direction of the 15mm edge for position randomized photons and protons. Compared to figure 45, the

In CALIFA the granularity in $\vec{e}_\theta$ (15 mm) is about twice the granularity in $\vec{e}_\phi$ (30 mm) because the θ angle is essential for the Doppler correction. Comparing figures 45 and 46 can now estimate the advantage in position resolution by going to 15 mm (and thus doubling the amount of channels). For very low energies ($E < 400$ keV, mostly photo effect) and very high energies ($E > 3$ MeV, mostly pair production), the advantage of higher granularity is clear. Even in the intermediate region where Compton scattering dominates, a moderate increase in the resolution is possible: at 1 MeV, one gains 20% (from 12 mm to 10 mm). Of course, this number depends very much on the fraction of photons one deems recoverable.

CALIFA is not only used to detect photons but also light charged particles such as protons. While CALIFA is the only detector to provide a position measurement of γ rays, for protons the situation is more relaxed because semiconductor vertex and strip detectors can provide a much better angular resolution than CALIFA can. However using CALIFA for position resolution in tandem with silicon trackers is an important strategy in case of heavy ion beams where δ electrons create a lot of background in the trackers. Furthermore, for high beam intensities, CALIFA offers higher rate capabilities[75] and a larger radiation hardness than current tracking detectors.

Therefore the position resolution for protons is also shown in figures 45 and 46 for comparison.

See figure 65 in appendix C for the mean energy distribution for a 200 MeV proton.

[75] especially with the multi-event readout mode

8 Experimental studies

For a stable and reliable operation of the CALIFA calorimeter – as well as for its calibration and performance monitoring – the detector response of all the crystals is tested experimentally before and after each experiment using standard γ ray sources. In addition γ rays from the room background are monitored during the run period. One of the major tasks of CALIFA is the spectroscopic analysis of collective excitations with γ ray energies up to 30 MeV in the lab frame (see section 3.1).

While there are many isotopes emitting γ rays of energies around 1 MeV after an α or β decay of their parents, the choice becomes quite limited at higher energies. α and β decays to highly excited states are rare in nature. One special case is the decay radiation of ^{56}Co that features a variety of possible gamma emissions in the 3 MeV region, but its half-life of 77 days is also a bit short as a standard calibration source.

8.1 The Americium Beryllium $(\alpha, \mathrm{n}\gamma)$ source

Besides nuclear decay, nuclear reactions can introduce secondary γ radiation. The $^{9}\mathrm{Be}(\alpha, \mathrm{n})^{12}\mathrm{C}$ reaction[76], for example, has a Q value of 5.7 MeV[77]. In this reaction not only the ground state of ^{12}C is populated by this reaction, but also the first 2^+ state with an energy of 4.44 MeV. In that case, the nucleus de-excites by emission of a gamma with an energy of 4.44 MeV directly to the ground state. For α energies of roughly 5.5 MeV, the cross section for getting the first excited state of ^{12}C is especially large – more than 300 mb [26].

^{241}Am is the second most long-lived isotope of that element, with a reasonable half-life of 433 years and a dominant α energy of 5.49 MeV.

Combining ^{241}Am with ^{9}Be to build a so-called so called AmBe source is a well known technique. The one available at GSI for testing CALIFA with high energy gammas contains 37GBq (1Ci) of ^{241}Am activity and provides an integrated neutron flux of 2.2 M/s. The rate of 4.4 MeV gammas was not given, but is generally 50-75% of the neutron rate, depending on the source geometry. [27, 28]

As the ^{241}Am α energy is clearly higher than four MeV, [26] suggests a 10% branching ratio for the population of the second 0^+ state of ^{12}C at 7.65 MeV. This is the famous Hoyle state involved in triple-α nucleosynthesis. A direct decay to the 0^+ ground state is forbidden, and its main decay mode is the emission of an α, and an isomer transition (via the first 2^+) only has a branching ratio of some 0.04% (see figure 47). Given that the energy of the first gamma would be 3.2 MeV, separating this from the double escape 4.4 MeV is impossible with scintillator detectors.

8.2 Experimental setup

For the calibration measurements, the CALIFA demonstrator was used. It consisted of three single petals of 4 × 16 crystals and two double petals of 8 × 16 crystals – or $\frac{7}{16}$ of the total forward barrel.

[76] Also famous for being involved in the discovery of the neutron [24]

[77] This can be calculated from the mass differences, see also [25]

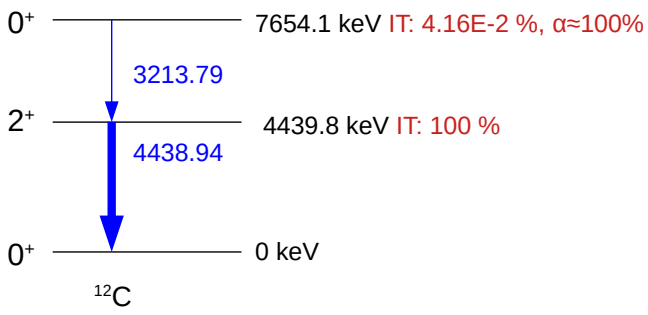

Figure 47: The relevant section of the level scheme of ^{12}C, from NuDat 3.0 by the National Nuclear Data Center by BNL. The branching ratio of the 7.6 MeV state strongly favours the alpha emission.

The AmBe source was placed near the target wheel's nominal target position[78]. A weaker ^{60}Co source was also left in the vicinity for reference.

All crystals were recorded with preamplifiers set to high-gain mode with a total active range of about 30±3 MeV energy deposit per crystal.

A typical spectrum of the signals recorded in a single crystal is shown in figure 48. The main features are the ground state transition at 4.4 MeV and corresponding single and double escape peaks as well as their Compton edges. In addition the 511 keV line from e^+e^- pair production is very prominent.

For a first analysis gaussian functions with a linear background were fitted to the different features. Thus a linear calibration of the individual crystals is accomplished.

Also visible in the spectrum two smaller peaks at about 640 keV and 850 keV whose origin will be discussed in section 8.4.

Figure 49 shows the calibrated spectra obtained for the individual crystals. The 511 keV peak as well as the single and double escape peaks of the 4.4 MeV gamma from the reaction are clearly visible for most crystals. While care was taken to put the source near the nominal target position[79], the size of the source made it difficult to get the position exactly right. Every ring of crystals has different conditions regarding the solid angle it covers when looked at from the source and how much of the crystal is in the shadow of other crystals. The amount of 3 MeV entries seen due to Compton scattering is very different between the rings. Because the *arbitrary crystal numbers* are related to the geometrical position of the crystals in some complex way there are semi-regular structures along the x-axis in this plot, such as bands of eight crystals which see fewer events in the 3 MeV region.

See the appendix for more information on calibration procedure and outcomes. Figure 50 gives the resolution measured for both the 511 keV and 4.4 MeV peak.[80]

[78] As the target wheel could not be removed at the time, and the source used is significantly larger than a target, the center of the source was downstream of the nominal position by about 1.5 cm.

[79] This is not the focal point of CALIFA because different rings are focused towards different points along the beam axis.

[80] In the appendix, figure 60, there is a version of this plot with the *squared* resolutions. This can be helpful because the different error contributions are added quadratically to arrive at the quadratic resolution. So in the

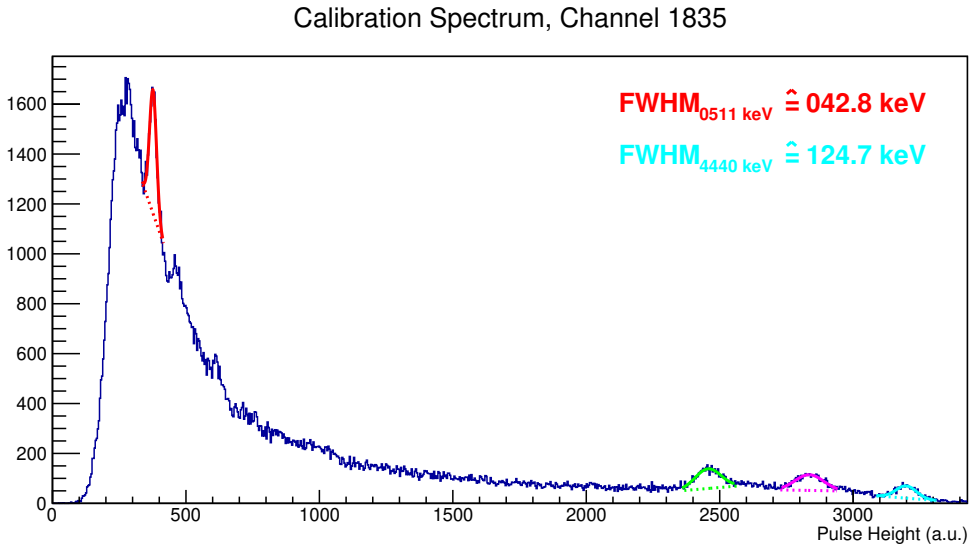

Figure 48: The uncalibrated spectrum for channel 1835. This channel has the best overall resolution. The 4.4 MeV photo peak, single and double escape are clearly visible, as as the 511 keV from positron annihilation.

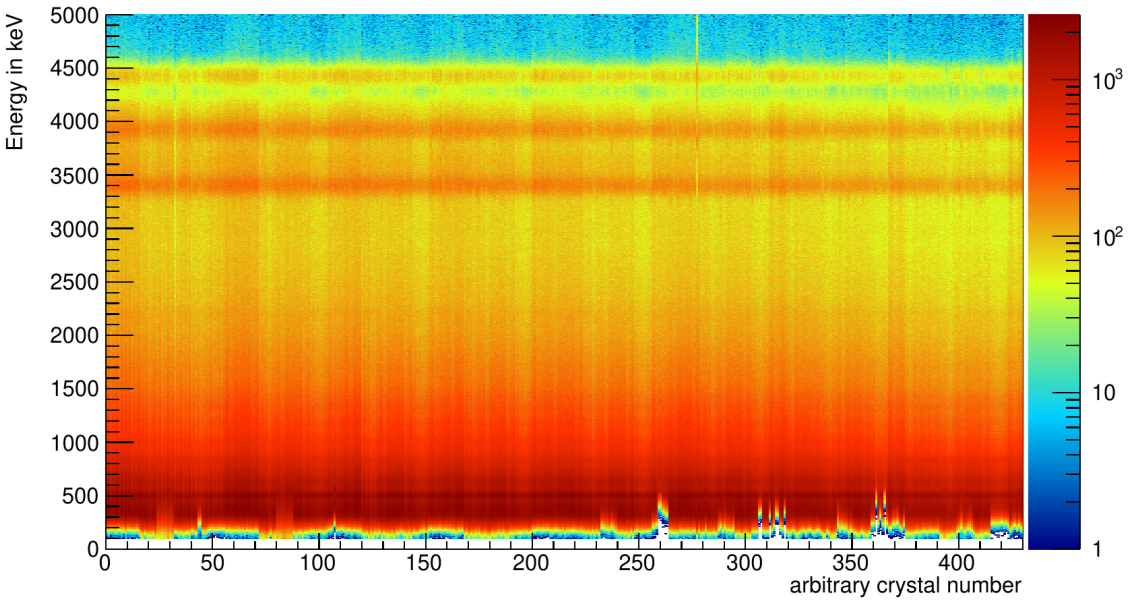

Figure 49: The energy distribution recorded by CALIFA demonstrator crystals exposed to the AmBe source, after applying the calibration.

There are different effects which can limit the resolution of a calorimeter[81]:

- Random noise not related to the event will limit the accuracy of the measurement. As the noise is not depending on the energy, this contribution is constant with regard to the energy of the gammas.

- There are statistical processes involved in turning energy deposits from ionizing radiation into electrical signals. The dominant effect are the intrinsic properties of the CsI(Tl) scintillator. Especially in the energy range CALIFA is operating the non-linear light output response to low energy secondary electrons with $E < 100\,\mathrm{keV}$ affects the statistical sampling and limits the energy resolution achievable [29] [30]. Afterwards, the number $\overline{N}$ of scintillation photons which are generated and then amplified in the APD is proportional to the energy deposit E, but random. As N is approximately Poisson distributed, the standard deviation of N is $\sqrt{\overline{N}}$. Thus $\Delta E_{stat} \sim \sqrt{E}$.

- If the overall gain is not stable, that will affect events in proportion to their energy. The dominant effect of this category is the non-linearity of the light transport. Due to the trapezoidal shape of the crystals the light from the tip of the crystal is typically more efficiently transported to the APD sensor than light emitted closer to the end of the frustrum. This effect is partly compensated by lapping the side faces of the crystals during production, but this is a manual process and typically not perfectly performed for all crystals. As γ rays, emitted from the target area do have a much higher variation in penetration depth at higher energy, this effect can worsen the relative resolution for higher energies.

The blue lines in figure 50 indicate the ratios the energy resolutions at these two different energies if all of the resolution was only due to one of these effects. They are $\Delta E_{4.4\mathrm{MeV}} = \Delta E_{511\mathrm{keV}}$, $\Delta E_{4.4\mathrm{MeV}} = \sqrt{\frac{4.4\mathrm{MeV}}{511\mathrm{keV}}}\Delta E_{511\mathrm{keV}}$ and $\Delta E_{4.4\mathrm{MeV}} = \frac{4.4\mathrm{MeV}}{511\mathrm{keV}}\Delta E_{511\mathrm{keV}}$.

There is little correlation between both resolutions: the correlation factor is just 0.18. For crystals with good resolution (where the FWHM for the 511 keV is below 65 keV and for the 4.44 MeV peak is below 190 keV respectively), the correlation between the resolutions is even smaller, $r = 0.065$. This means that for good crystals, the variance of the statistical contribution does not explain most of the observed variance, the points form a blob on the $\Delta E \sim \sqrt{E}$ line instead of running along that line. This indicates that for most good crystals, differences in statistical effects are minimal. This is expected as the dominant statistical effect is due to the low energy electrons, whose statistics do not depend on the particular crystal as the doping concentration, which effects photon statistics, is well monitored by the manufacturer.

However, the crystal no 1859 features a low gain and a suboptimal resolution at both energies while also having a low gain. This might be due to bad light coupling between the crystal and APD resulting in low photon statistics.

squared plot, one can use vector addition to visualize how these contributions add.

[81] This list is not meant to be exhaustive. Optical and electrical cross-talk, light yield inhomogentities and other such effects may be much more subtle.

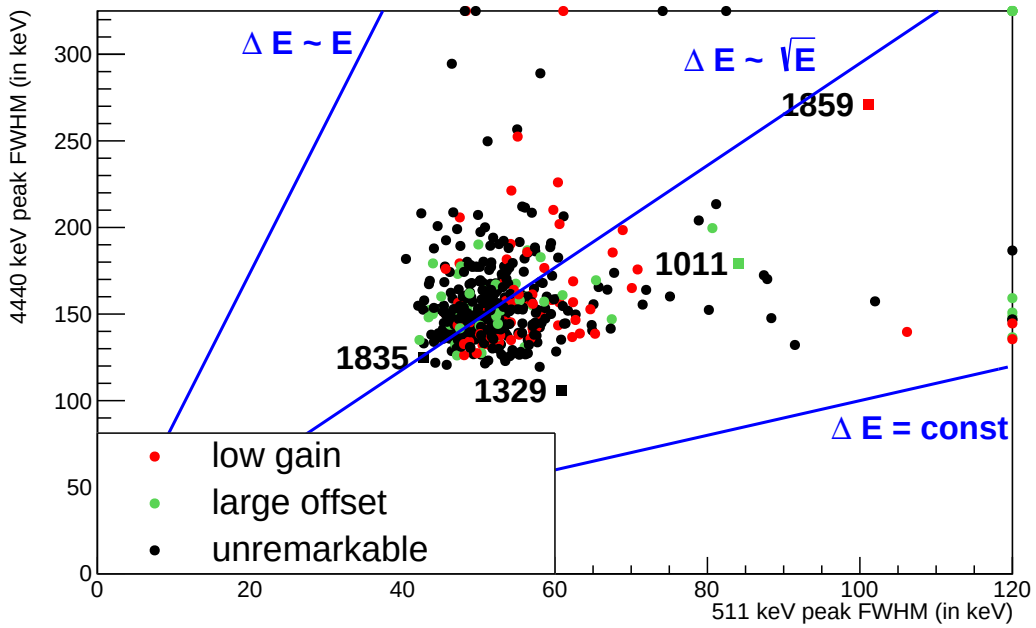

Figure 50: The resolution of the 511 keV peak and the 4.4 MeV for all the 431 crystals for which a calibration could be obtained. If one of the gaussian fits failed, the corresponding value was set to the upper bound of the shown range. In these 33 cases, the single/double escape peak was used as a reference point to obtain a calibration. The colors indicate non-standard calibration parameters. While the large offsets come from a misconfiguration of FEBEX which does not have an impact on the resolution, the channels with low gains exhibit a worse resolution on average.

Near the $\Delta E \sim E$ line, one can find some crystals which have an unremarkable resolution at 511 keV but a very bad resolution at 4.4 MeV. This can appear due to crystal inhomogenities where the light yield is position dependent or non-uniform light transportation.

On the other side, near the $\Delta E = \text{const}$ line, there are crystals which are good at 4.4 MeV but have a bad resolution at low energies. The reason for that is thought to be constant noise caused by front-end electronics.

Every channel is subject to all these influences to some degree. With only two energies measured, the contributions of the three different effects can not be uniquely attributed.

Furthermore, the different contributions have to be added quadratically. To simplify the visualization of this, the squared data points are replotted in figure 60 in the appendix.

In conclusion, plotting the resolution at different energies is a good method to identify specific issues with individual crystals. In particular, this allows one to identify whether the issue is related to the crystaline properties, the APD or the electronics.

8.3 Event correlation and time stitching

8.3.1 Clustering in time, event mixing and background subtraction

As discussed in section 5.10, before clustering can happen, hits on the crystal level – each with their own time stamp – have to be assembled into an event.

For beam experiments with moderate rates of ions (e.g. 10-50 kHz), the details are not critical because the background rate is low, so counting everything in a window within a few microseconds will work fine. For high rate experiments especially using heavy ion beams or heavy targets the atomic background created from the beam interacting with electrons in the target creates a substantial load to the detector not related to the reaction of interest. This is especially important in the case of SIS18 beams not Poisson distributed in time but exhibiting a pronounced microstructure on the microsecond time scale.

The AmBe source emits radiation at a much higher rate in the order of $10^6\,\mathrm{s}^{-1}$. With the limited time resolution especially for small energy deposits, this means that the average time window will contain more than one nuclear reaction.

If one wants to find correlations between energy deposits from the same reaction, one has to carefully subtract the random coincidences.

Due to the high event rate and limited solid angle coverage, random coincidences will always play a major role in any coincidence cuts. This means that we need good models to estimate the background rate for events with given properties.

For coincidences, we commonly split the detector into two halves. Typically, we use one half of the detector to apply a cut, e.g. *a 4.44 MeV gamma cluster was detected*. Then, we can look at the signal of the other half of the detector given that cut.

To get to the rate of random coincidences between two sides of the detector a high energy physics method called **event mixing** is a good approach. The idea is to artificially change the time stamps of one side of the detector by an amount much larger than the length of any reasonable coincidence window[82]. This will mean that any coincidences between the sides will now certainly be random, and a good background estimate should describe them reasonably. This works especially well with the constant time slices method where the presence of hits does not affect the stitching window and produces a robust estimate for the background from random coincidences.

Suppose that the energy of cluster number p for side $s \in \{0, 1\}$ in time slice $n \in \mathbb{N}_0$ is called $E_{s,n,p}$. For the observed coincidences, you would histogram for every time slice n histogram the tuple $(E_{0,n,p}, E_{1,n,q})$ with p and q running over the number of clusters of the respective sides. To estimate the background, you would instead histogram $(E_{0,n,p}, E_{1,n-k,q})$, where k is a fixed constant offset. To obtain the non-random coincidences, you can then just subtract that second histogram binwise from the first.[83]

[82] In practice, 10 μs seems to suffice.

[83] In principle, producing the two dimensional background histogram as the outer product of of the respective one-dimensional histograms would be a similar approach. One advantage would even be less statistical fluctuation because the statistics per bin in a 1-dimensional histogram are of course much higher. However, scaling this background correctly is because sometimes the DAQ became overwhelmed and did not record data for long periods of time, so the effective number of time slices is doubtful. For the event mixing method, a long drought will just reduce the background by $2k$ events as compared to the total. Still, special care was taken to identify

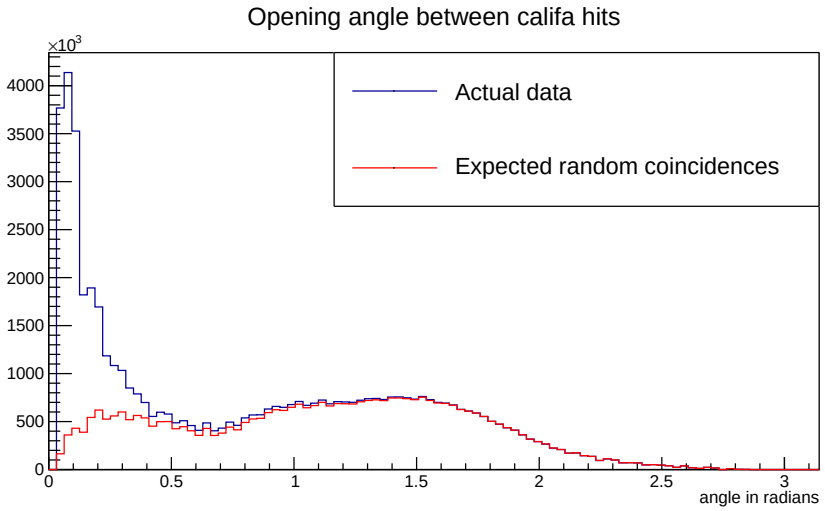

Figure 51: The opening angle between CALIFA hits. This is defined as the anglular difference between crystal centers as seen from the nominal target point using the nominal geometry. For very large opening angles, the background was scaled to match the observed pattern.

8.3.2 Determining the optimal clustering parameters

To understand the optimal clustering parameters in space, statistical considerations play an important role. For background rate zero and γ ray multiplicities of one, a clustering of all crystals would be preferable. In cases of a high background rate, random coincidences within the coincidence window would dominate the addback if the neighborhood is not tightly constrained. For the data taken in this analysis, this was investigated by using the nominal R3BRoot geometry of the CALIFA crystals to calculate the observed distribution of opening angles between coincident crystals. The result is the blue curve shown in figure 51.

Of course, the tail is dominated by random coincidences. These were estimated with the following method: For each time slice, each crystal would independently generate a hit with the same probability[84] as observed in the experiment, independently. For all pairs of crystals which were selected by the random number generator in that time slice, the opening angle between the crystals would then be histogrammed. After rescaling so that the curves match for large opening angles, the result is the result is the red curve. Its shape is dominated by the geometry of the demonstrator (petals of tightly packed crystals with large gaps between the petals), see section 8.2.

Physics correlations at such large opening angles dominated by pair creation leading to a single or double escape from the cluster are explicitly not considered here to allow special considerations of these events in the later analysis.

such regions and skip their borders (where exactly one of n or $n-k$ has zero entries).

[84]The average number of hits per time slice of length 1.5 μs was 1.31, which means that the demonstrator saw 872 thousand hits on the crytal level per second. So the probability of a typical crystal seening a hit in one time slice would be around 0.3%.

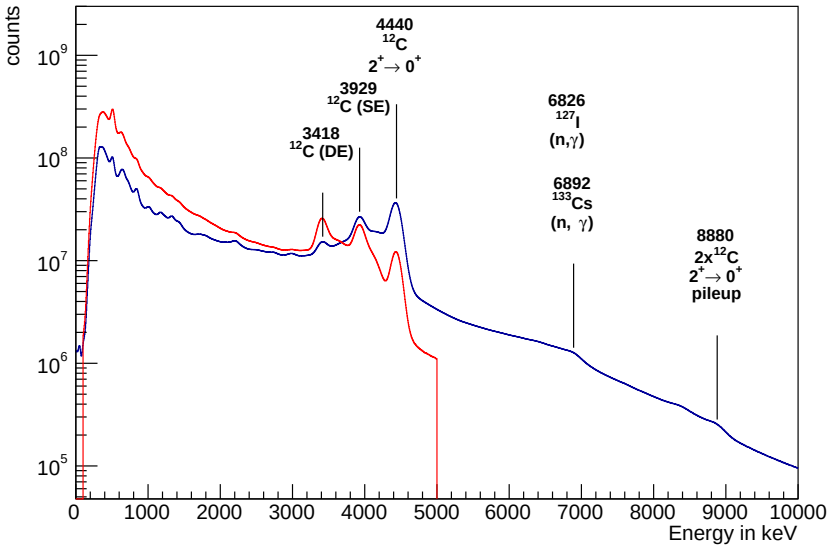

Figure 52: The energy spectrum from the AmBe source obtained after clustering, full range. For comparison, the red histogram indicates the sum of all single crystal spectra (without clustering, up to 5 MeV). For the identification of the low energy peaks, refer to the following figure.

This allows to experimentally determine the optimum opening angle of the cone used for clustering. Hits with an opening angle larger than 0.25 (e.g. 29°) have a likelihood of more than 50% to be random coincidences. Thus, the a cone with that opening angle was used for clustering in the following sections.

The angle of the cone has to be adjusted to the background conditions for each experiment. For beam emitted γ rays which have much higher density in CEPA-CsI due to the Doppler effect, a *theta*-dependent clustering method might be preferable.

8.3.3 Energy spectra after clustering

Figure 52 shows the full spectrum obtained with the AmBe source both before and after clustering. To quantify the entries in the peaks related to the 4.4 MeV γ ray, gaussian functions with a linear background were fitted to the various peaks. The results are thus:

	Single Crystals			Clusters			$\Delta \int$
	Amplitude	σ/keV	$\int$ Gauß	Amplitude	σ/keV	$\int$ Gauß	
photo peak	7.8 M	60	238 M	24 M	72	895 M	+658 M
single escape	9.7 M	67	328 M	9.3 M	67	315 M	-13 M
double escape	11 M	62	353 M	2.7 M	61	85 M	-268 M

In the single crystal spectra, the high energy range is dominated by the double escape peak. After clustering, the ratio between single escape, double escape and photo peak significantly change: some of the 511 keV γ rays are found in nearby crystals, upgrading both the single and

double escape peak towards the photo peak. While the double escape is nearly depleted, the single escape is slightly increased and the photo peaks increases by a factor of 2.76. Clustering thus offers a substantial improvement for energy reconstruction efficiency. This was with just the CALIFA demonstrator where about half of the crystals are sitting on an outer border instead of being fully surrounded by other crystals[85], the results for the full CALIFA detector will be better still!

It will be later shown that there is a strong coincidence between having a single escape cluster and a stray 511 keV in another part of the detector. For experiments in which only a single high energy γ cluster is detected, this opens up the possibility to increase the photopeak even more by just adding 511 keV to the main cluster in analysis.

The 4.4 MeV photo peak is clearly visible, while the single and double escape peaks are less prominent after clustering. Figure 53 is the same plot zoomed to the low energy region. The most prominent peak here is the 511 keV from positron annihilation.

The ^{60}Ni originate from a weak ^{60}Co source was placed near the AmBe source for monitoring.

^{56}Fe is the most common isotope in the source casing and aluminum is widely used as an engineering material in the target area. Hydrogen is present in the moderator surrounding the source. Inelastic neutron scattering on ^{56}Fe and ^{27}Al can explain the peaks around 850 keV and at 1 MeV. While the formation of deuterium from neutron capture might explain the peak at

[85] This is also why double escape is still a plausible process.

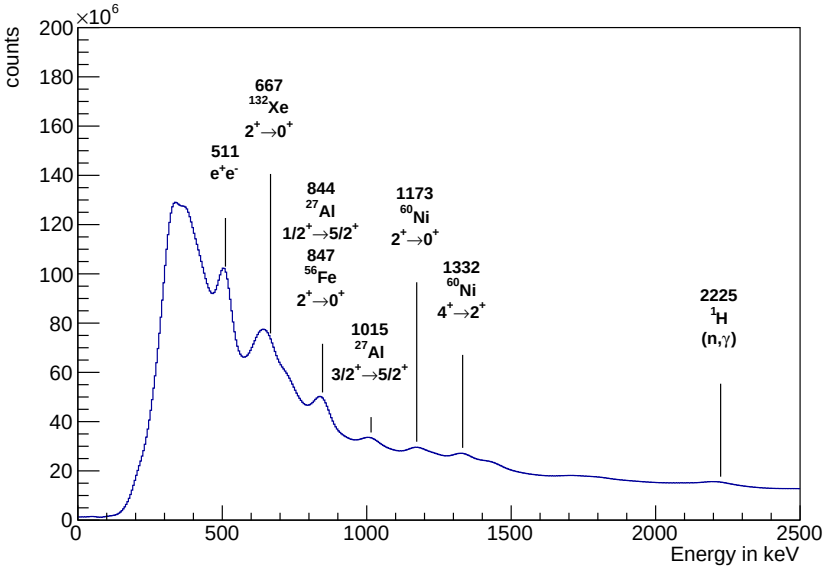

Figure 53: The energy spectrum from the AmBe source obtained after clustering, lower energy region. The excited states of the Fe and Al are produced by inelastic neutron scattering. The ^{60}Ni states come from decays of a ^{60}Co source, see section 8.2. The Xe is produced in the decay of ^{132}Cs which was produced during the previous beam time by knock-out reactions.

2.2 MeV (which is clearer visible in the previous logarithmically scaled plot).

Before the measurement with the AmBe source, CALIFA was exposed to a high intensity beam in the S494 experiment which studies the fission of ^{16}O. It is believed that the broad peak observed at ca. 650 keV might partly be caused by the decay of ^{132}Cs (electron capture decay to ^{132}Xe, $T_{1/2} = 6.5$ d, emission of 667 keV γ ray).

Figure 52 also indicates candidates for the high energy peaks observed in the spectrum. Neutron capture in both ^{133}Cs and ^{127}I would produce a γ ray with energies around 6.8 MeV, and there is a suggestive bump in the spectrum at that energy. By contrast, there is little evidence for neutron capture by either iron or nickel (which are less abundant in the setup).

8.4 Investigation of γ, γ corelations

Given the limited resolution of CsI (especially when compared to Germanium detectors) it is not trivial to identify the various overlapping peaks in the energy spectra. To compensate for limited resolution, we can look at the correlation spectra instead of the one dimensional spectra, using coincidences between multiple γ rays to distinguish separate peaks.

This works because many of the expected γ lines will appear coincident with other γ rays. For example, the peaks caused by the 4.4 MeV will be coincident with a neutron which may cause other nuclear reactions which will result in γ ray emission. So to understand which physics content they have, we can look to γ, γ corelations. One especially expects significant contributions from the 511 keV escape but also from reactions producing a cascade of γ rays.

As the source intensity was very high, this data presents a unique opportunity to study correlations between clusters in adverse conditions in the presence of a high rate of random coincidences apart from helping to understand the lower energy part of the spectrum.

Appendix section B discusses different approaches on histogramming n clusters in one half of CALIFA versus k clusters in the other half of CALIFA. For simplicity and robustness, the E_0 vs E_1 plots shown here are based on the **All pairs** approach (i.e. $n \times k$ entries in the histogram). Figure 54 shows the result of that. The upper plot shows the observed totals. Here, random coincidences clearly dominate. For example observing a 4.4 MeV cluster in side zero is correlated with also observing a 4.4 MeV cluster in side one. The coincidence between the 511 keV and the 4.4 MeV photo peak looks as strong as for the single escape peak.

As discussed in section 8.3.1, in high energy physics the random coincidences are typically investigated by the methods of event mixing, which means you combine randomly events from one interaction with events from another one, to see which part of your spectrum is not physics related. Using this method to subtract the uncorrelated background removes 81% of the events and results in the lower subfigure of figure 54.

The resulting spectrum is well-explainable by coincidences which share a physical process of origin. The coincidences between 4.4 MeV photo peak, single escape or double escape on both sides are gone. (4.4 MeV, 511 keV) is very faint compared to the original spectrum. The remainder is likely from the coincident neutron producing a high energy γ ray in the 3 MeV range

(511 keV, 511 keV) is the dominant peak. This is the result of pair production in passive material (such as the metal table) with both of the annihilation photons escaping and hitting the detector.

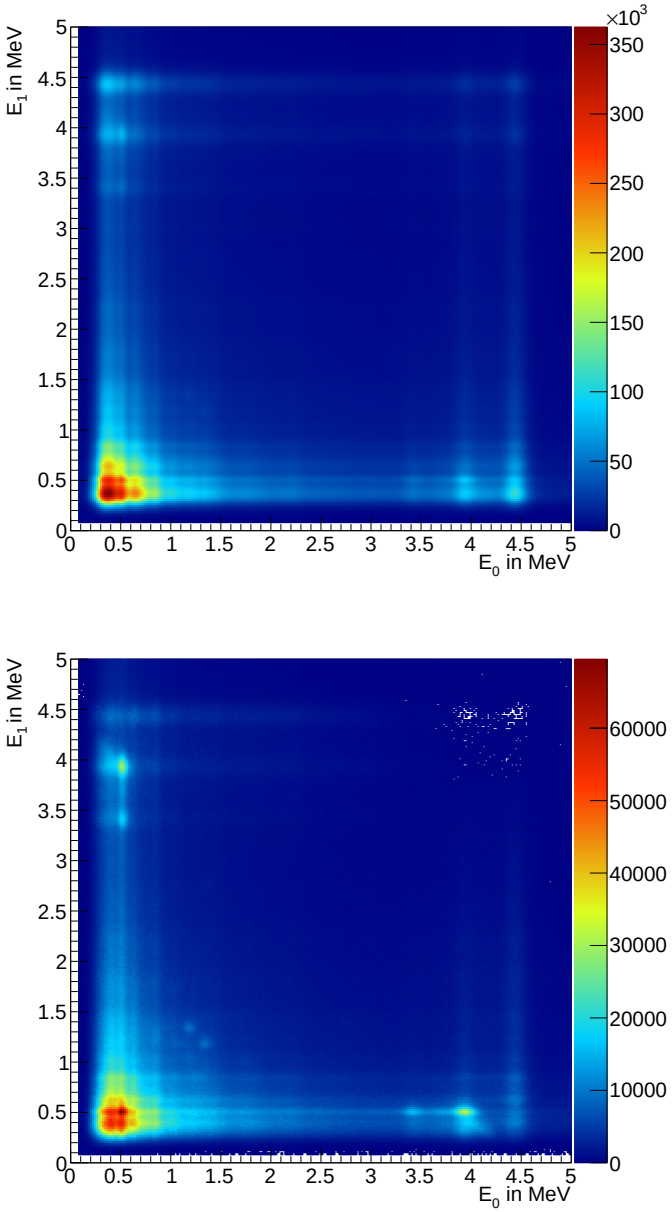

Figure 54: The correlation of the energies of hits on either side of CALIFA. The upper histogram are the observed totals, the visible area contains around 971 M entries. The lower one is what is left after the background obtained from event-mixing (785 M entries) is subtracted and only contains 187 M entries. While random coincidences dominate the upper plot, the lower plot clearly shows the expected physics coincidences. Single and double escape are strongly correlated with 511 keV. 511 keV is correlated with itself, which corresponds to pair creation outside the detector, for example in holding structures. The 1.1 MeV, 1.3 MeV coincidence from the much weaker ^{60}Co source is clearly distinct.

The correlation between the 4.4 MeV escape peaks and the 511 keV peak are clearly visible. This opens up the possibility to fix single escape clusters by opportunistically adding stray 511 keV clusters detected elsewhere in CALIFA to them.

Both coincident peaks of the ^{60}Co source are also clearly distinguished in the coincidence plot. Finally, there are a lot of correlations near 850 keV which will be investigated in subsequent sections.

8.5 Energy spectra with cuts on the 4.44 MeV peak

γ rays with an energy of 4.44MeV are produced by the deexcitation of the ^{12}C$*$ after its production in the ^{9}Be$(\alpha, \mathrm{n})^{12}C*$ reaction. If one cuts on that energy, the events remaining are likely to be caused by neutron interactions.

Cutting on 4.44 MeV will deemphasize anything not caused by the reaction producing ^{12}C$*$, such as other reactions of the alpha particle or environmental background events, as well as the ^{60}Co source.

As lamented above, the event rate of the AmBe source is very high. Thus, random coincidences will always be a given. Fortunately, given the time stitching method in 5.10.1, one can estimate the rate of random coincidences. By subtracting the random coincidences, we can in fact eliminate them at the cost of some statistical accuracy.

Figure 55 shows the observed events in one half of the detector given the existance of a 4.4 MeV hit in the other half, as well as the scaled spectrum after clustering for comparison. For energies above 3.5 MeV, virtually all detected events can be explained as random coincidences, while at lower energies there is an increasing excess of measured events. Many of the peaks identified as neutron-induced γ rays in figure 53 are now strongly enhanced.

8.6 Effect of neutrons

According to [26], the neutrons emitted together with the 4.4 MeV (called n_1 neutrons) have an energy between two and six MeV.

While electrons and γ rays cause electromagnetic showers, neutrons interact with the nuclei of matter. In general, the likelier reactions for these neutrons are:

- elastic scattering: $\mathrm{X}(n, n)\mathrm{X}$
- inelastic scattering: $\mathrm{X}(n, n)\mathrm{X}*$
- absorption: $^{A}\mathrm{X}(n, \gamma)^{A+1}\mathrm{X}$

The most common nuclei near the AmBe source are:

- ^{27}Al: scattering chamber, petal housing
- Fe, Cr, Ni: stainless steel used in source housing and CALIFA mounting
- ^{127}I, ^{133}Cs: detector
- ^{1}H and ^{12}C from plastic and carbon fiber

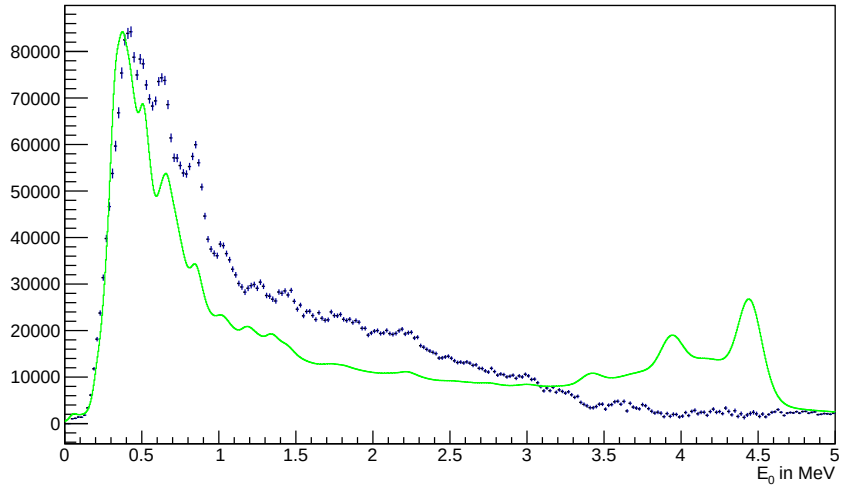

Figure 55: In blue, the γ rays observed coincident (after background removal) with the 4.4±0.1 MeV. The green curve is a scaled version of the spectrum after clustering without any cuts for comparison. The y-axis indicates the counts for the blue curve, the green curve is arbitrarily scaled. With the cut, the 4.4 MeV peaks are gone. The continuous background from the neutrons is enhanced, as are the γ peaks likely caused by neutrons, such as the 2.2 MeV from ^{2}D or the peaks from ^{56}Fe and ^{27}Al discussed in figure 53. With the peaks from the ^{60}Co source gone, there is some evidence for a weaker peak at 1.2 MeV. The limited resolution and much lower statistics make it difficult to clearly identify most of the weaker peaks here. Note that the height of the blue crosses indicates the statistical error of the bin after background subtraction, which is different from the Poission statistics implied by a histogram drawn without error bars.

CALIFA is not a detector meant for detection of slow neutrons. The maximum energy transfer for elastic scattering of nonrelativistic neutrons n on heavier nuclei A is

$$\Delta E_n = E_n \times \frac{4 m_n m_A}{(m_n + m_A)^2} \tag{31}$$

For a six MeV neutron and iodide, this comes to around 188 keV, which is mostly below the trigger threshold used in the experiment. Furthermore, one would not expect a peak: a continuous spectrum of neutrons convoluted with a continuous momentum transfer function will result in a broad distribution of energies, which is unsuitable to identify a specific process.

In the case of inelastic scattering or (n, γ), a much larger fraction of the neutron energy could be transferred to excite the collision partner.

The cut on the 4.4 MeV photopeak used in figure 55 is still not the best one can do to distinguish all of these neutron induced γ rays: sometimes neutrons will excite higher states in nuclei which decay via characteristic cascades of γ rays. A prime example of such a family of cascades will be discussed in the following section.

8.7 Energy spectra with cuts on the dominant Fe-56 line

Figure 56 shows the level scheme of Fe[56] which has about 92% of the natural abundance. As all even/even nuclei, the first excited state is a 2^+ state with an energy of 847 kev in this case. Most of the higher lying states would decay through this state to the ground state.

Even after background subtraction, the two dimensional histogram shows a large continuous background in the region of the 847 keV line due to other interactions of both the neutron. To focus on the coincidences with the small peak, a side band subtraction was applied. From the events coincident with 847±50 keV, half of the events coincident with the background dominated regions $847 - 100 \pm 50$ keV and $847 + 100 \pm 50$ keV were subtracted.

Figure 57 is the resulting spectrum, the green . The red lines are gamma energies of ^{56}Fe emitted during the decay to the first 2^+ state. The first 4^+ and the second 2^+ state indeed correspond to significant peaks in the data, with the correlation with higher states being more debatable.

This should be taken as strong evidence that ^{56}Fe was either present or produced in near the neutron source.[86]

[86] Unfortunately, my paper *Tentative γ-spectroscopic evidence for the presence of ^{56}Fe in AISI 321 stainless steel* was rejected by both *Science* and *Nature*.

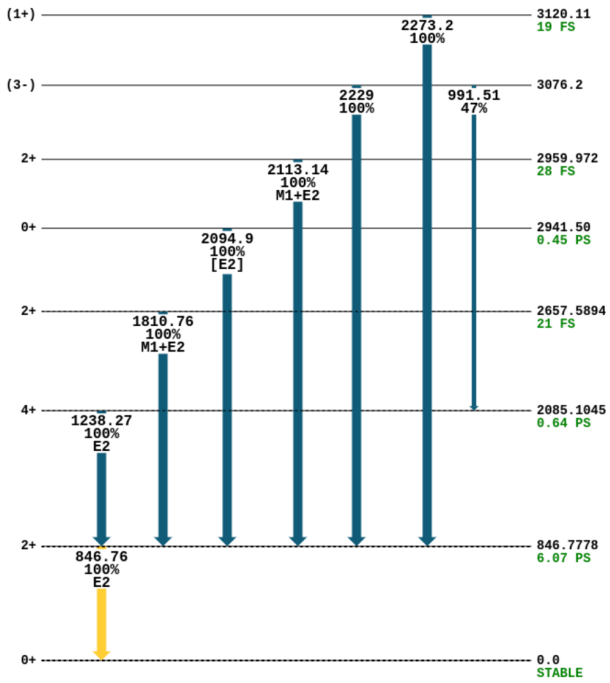

Figure 56: From [31], the level scheme of the lowest state of ^{56}Fe. Note that all of these de-excite via the first 2^+ state.

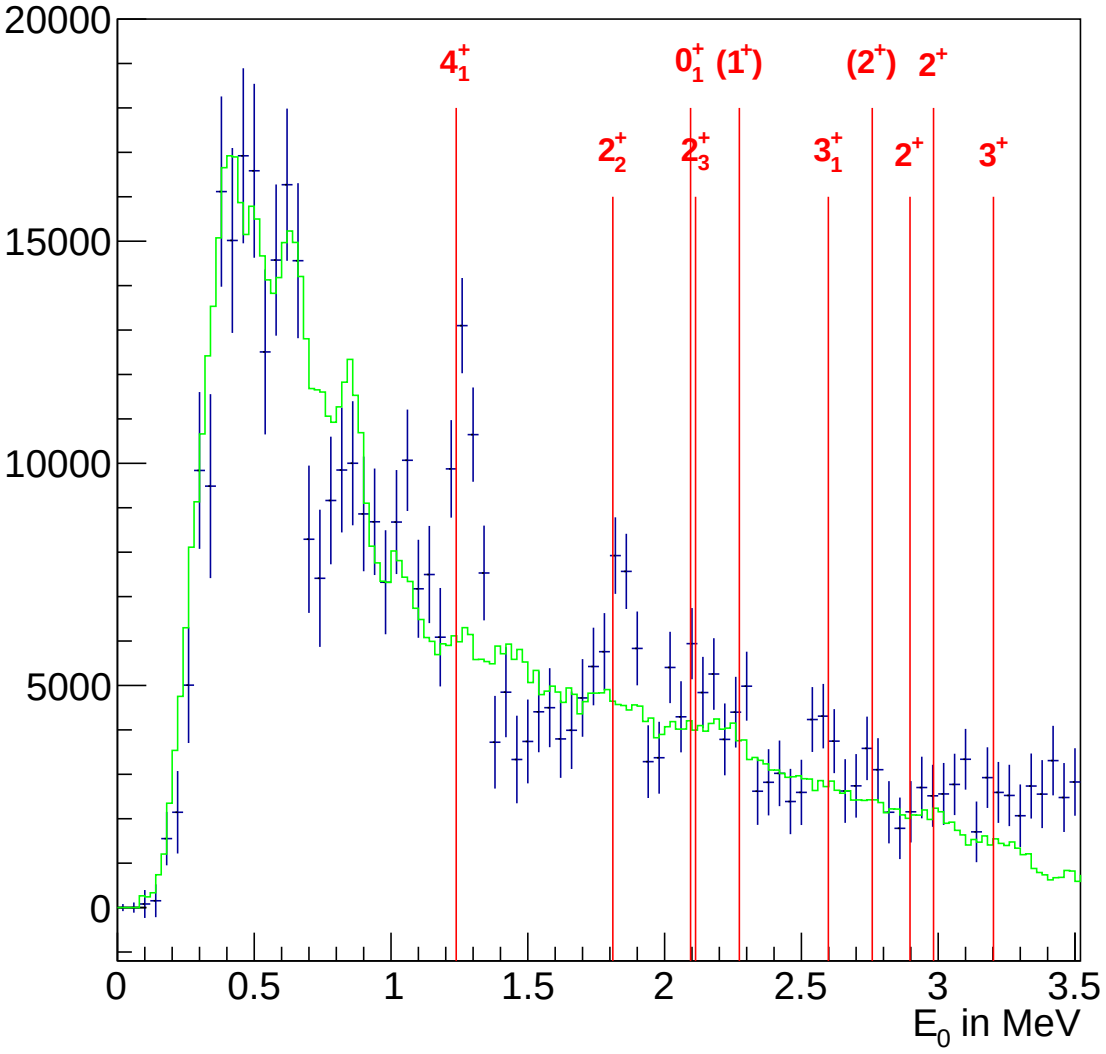

Figure 57: In blue, the excess events in one half of the detector given that the other half saw an energy of 847 keV (which is the energy of the first 2^+ state of ^{56}Fe). The red vertical lines mark the energies of γ rays of decays into that first exicted state (see figure 56). The green histogram are the coincident γ rays for the 4.4 MeV cut, for comparison.

9 Outlook

Since the start of the analysis presented in the previous chapter, the CALIFA demonstrator has given way to the full CALIFA holding structure. The forward part of the barrel and the iPhos section as well as a part of the backwards barrel section have been filled with crystals. Due to the success with iPhos, CsI is now also used for the most forward part of the calorimeter, replacing the $LaBr_3$ $LaCl_3$ phoswitch CEPA section with CEPA-CsI.

For high energy γ rays, this allows experimental studies which focus more on clustering efficiency: while in the demonstrator almost half of the crystals were on the fringe of the petals – and thus very likely to lose some of the γ energy to the side, the final version of CALIFA has no such gaps.

Different sources and geometries

While the 1 Ci AmBe source was well suited for obaining significant statistics per crystal for the photo peaks for the 4.4 MeV γ rays, its high rate makes correlation and cluster analysis difficult due to random coincidences. Increasing the distance between the source and the halves of CALIFA would reduce the background from both unrelated 4.4 MeV γ rays and neutrons. The price to pay would be that for any source position, most crystals would not be hit perpendicular on the front face of the crystal which completely changes the clustering geometry, so only a small part of the barrel could be studied per measurement. Another option is to use a weaker AmBe source, which will cut down on the random coincidences.

In a self triggering system there is quite limited range of activity for a source which would provide a substantial improvement. On one hand the random coincidences should be minimized while on the other side the natural room background from the Uranium and Thorium decay chains produce a substantial rate in such a large volume instrument.

As an alternative a Plutonium Carbon (PuC) may be a valuable extension for further investigations. PuC sources work similar to AmBe sources, only instead of ^{12}C the reaction produces ^{16}O from ^{13}C in an (α, n) reaction. The upside is that the first excited state of ^{16}O has a significantly higher excitation energy of 6.13 MeV, the downside is that the probability to populate this excited state is much lower, which leads to a much lower observed γ to neutron ratio: while a typical γ to neutron ratio for AmBe is 0.56–0.75 [28, 32], that ratio is about 0.05 for PuC [33].

Beyond that energy, neutron capture reactions can be used to reach even higher energies. [34] describes a setup to generate 9 MeV γ rays by moderating neutrons from an AmBe source and subjecting them to a (n, γ) reaction in nickel. This requires two nuclear reactions[87], which will make the effective production rate of the γ rays much smaller. A lot of material is required to moderate the neutrons, which will make the source bulky and not very point-like. While neuton capture crossections of nickel are in the range of few barn other material in the source will also contribute to the gamma emission which will lead to a rather crowded spectrum of such a source.

[87] Not counting the Am decay

Finally, naturally occurring Ni consists of five different stable isotopes, so instead of a single γ line, the spectrum consists of multiple peaks which will be difficult to resolve in CsI detectors.[88]

Further developments in clustering

Due to the continuing upgrade of the CALIFA detector array the clustering studies should be repeated using the tools prepared in the framework of this thesis. With only the data for the demonstrator being available available for this thesis, the insights from the simulation could not be fully exploited in the demonstrator clustering. The full CALIFA detector offers the opportunity to apply these learnings for the clustering of data. An implementation of the clustering algorithm which separates neighborhood definitions from the clustering itself is available[89].

Machine learning with artificial neural networks ("artificial intelligence", AI) might be another avenue to consider. The main challenge here is that the net has to be trained with simulation data, because training a NN with discrete γ energies from calibration sources will result in the network assigning one of the energies it saw during training to the cluster (overfitting). In the author's estimation, the hardest part of such a project would be to quantify the individual crystals well enough with regard to factors such as noise level, resolution at different energies, non-uniformity of light yield, background rate from natural radioactivity and so on that a sufficiently realistic model for Monte Carlo simulations can be obtained. The MC simulation can then train the neural network. While machine learning might help to improve γ clustering within the jagged geometry of CALIFA crystals, neural networks are also black boxes which tend to fail catastrophically when working on data which is dissimilar from their training distribution.

Besides having the neural network do the whole add-back, a more legible method might be to have the network just select which crystals get added to the cluster as a function of the crystal with the highest energy hit and the energy seen in it. The results could then be regularized to form rules defining for each theta angle of the crystal around which clustering happens and for certain energy ranges what the neighborhood should be.

CALIFA is meant to quantify γ rays emitted by fragments moving at relativistic velocities. In particular, the energies of the photons in the frame of reference in which they are emitted are of interest. As discussed *ad nauseam* in section 3.1, this means it is not enough to precisely measure the energy in the lab frame as then angle between the direction of ion movement and the photon is also required. Machine learning approaches might thus also be used to further improve the position assigned to a cluster. Nevertheless this approach is not as straightforward as expected as the beam energies, nuclear properties of the beam, individual detector configurations around the target and many other parameters may strongly influence the average pattern expected in the detector. This might mean that a separate neural network would have to be trained for every experiment and beam energy.

[88] Even if one were to invest in a few grams of isotopically pure ^{58}Ni, manufacturers of strong reaction sources insist on encasing them in thick metal containers for some reason. These metals can also undergo (n, γ) reactions.

[89] `https://github.com/R3BRCDevs/R3Brc`. Note that for political reasons, this is not part of R3BRoot but part of a fork called *R3B root channel*, which is not maintained.

Quantifying CALIFA for light charged particles

Calibrating a γ detector is comparatively easy as there are various γ ray sources available for long term offline measurements. Calibrating a detector for charged particles is a bit harder.

Aligning CALIFA so that a primary beam from the accelerator can pass through the nominal target position and a single crystal at a time would be ideal, but it is not feasible to add the required two rotational degrees of freedom from an engineering point of view.

$(p, 2p)$ reactions with ion beams have the problem the scattered proton carries a Fermi momentum. In principle, the momentum transfer to the ion can be measured using the in beam detectors, but this means that the CALIFA calibration would be dependent on the calibration of the in beam detectors. Another unknown in a $(p, 2p)$ reaction is the binding energy of the proton which is shell-dependent. Finally, $(p, 2p)$ reactions are often the reactions of interest. Using the same reaction channel for calibration and analysis seems non-ideal.

Fortunately, nature provides us with pure hydrogen which we can use as a target material. Besides some niche uses in stellar fusion, chemistry and airships, ^{1}H is notable for being an ideal projectile for calibrating CALIFA with light charged particles. A proton beam scattering on a proton target is a reaction which only has two degrees of freedom: the momentum transfer in the xy plane determines the ϕ angle of the reaction plane and the momentum transfer in z direction determines the asymmetry between the θ angles of both protons[90]. The reaction kinematics can be described fully with just the conservation of momentum and energy – no Fermi momentum or nuclear binding energies to complicate life: the incoming beam momentum and the scattering angle define the energy of the scattered proton.

The remaining difficulty is to measure the scattering angle. Ideally, one would use multiple layers of tracking to reconstruct the reaction vertex and determine the angles precisely. As the target area tracking is still in the process of being finalized, one possibility would be to use a small target so that one point of the tracks is already fixed. Heiss [35] showed that elastic proton-proton scattering can be used to calibrate CALIFA using a small target.

One way would be to use a small[91] plastic (CH_2) target held in place by very thin wires. However, the downside of using CH_2 is that it also contains carbon nuclei, so only a fraction of the two proton events will originate from hydrogen scattering. One hare-brained idea of the author is to use an organic scintillator as a target instead and using thin fibers for both holding the target and transporting scintillation photons. The idea would be measure the recoil of the boron nucleus in a $^{12}\mathrm{C}(p, 2p)^{11}\mathrm{B}$. This would allow discrimination between reactions in carbon and scattering of hydrogen. Alternately the contributions from carbon can be subtracted by having a pure carbon target which contains the same amount of C as the CH_2 target.[92]

[90] Naturally, the conservation of energy enforces a relationship between the momentum transfer in z direction and perpendicular to it.

[91] Perhaps a sphere with $\varnothing \approx 2\,\mathrm{mm}$. Of course, making the target to small will increase the fraction of scattering events from other material in the beamline, like vacuum windows or stray particles hitting the side of the beamline. This imposes a practical limitation on target size.

[92] Assuming low momentum spread in the incoming beam, in-beam tracking would not be required. This means that the beam rate could be very high. The time resolution of CALIFA is good enough that 10^6 reactions per second should be acceptable, which would produce sufficient statistics for each crystal in short order. For the scintillator target proposal, such high rates would probably be a surmountable challenge given that it is enough to know if a CFD threshold was reached or not while for tracking detectors, such a rate would not be practical.

An additional benefit of such a measurement would be the opportunity to not only obtain a high quality characterization of the detector response to stopped protons, but also with regard to punch-throughs. As outlined in the technical design report [10] the energy measurement strategy in this region is to measure ΔE and calculate the total energy from that. Of course, to calculate $E(\Delta E)$, it is important to know the distance the particle traveled through the detector. This would be easy if CALIFA was a hollow sphere of constant thickness surrounding the target center, but a careful observation of figure 10 shows that this is not actually the case. While the final configuration of R^3B will without doubt contain multiple layers of target area tracking which can be used to precisely determine the distance the ion traveled through each crystal[93], an extensive proton run would also be a great chance to validate the $E(\Delta E)$ relationship for the case where the particle passes through a single crystal in radial direction.[94]

Muons produced by cosmic radiation are another possibility for furthering the understanding of the CALIFA detector. Compared to elastic proton scattering, they have a lot of drawbacks: they appear at much lower rates, have an unfavorable angular and spatial distribution and are (in the best case) minimum ionizing particles which will go straight through CALIFA. Their main selling point is that nature provides them around the clock, thereby bypassing the requirement to get a PAC proposal accepted and organizing a beam time. While running CALIFA standalone to record muons is possible, it is not very efficient because most events will be from muons which pass through multiple crystals along an underdetermined trajectory. However, in conjuncture with the target area silicon trackers, muon tracks could be used both to align the detectors and to study the energy deposit in the crystals given that the trackers can constrain the distance the muons travelled in each of the crystals. This method was already tried by the collaboration, but the ambiguities in the energy deposits did not allow for a unique tracking of the muons without the inner tracking system. For the moment an alternative approach should be considered. By placing a rather small cylindrical scintillator at the nominal target position which provides a trigger for CALIFA one could focus on the events where the muon only passed through a single crystal on its way to or from the target, in which case the distance it traversed in the crystal is likely to be the length of the crystal.

Particle identification

Besides determining the energy of particles, CALIFA can also distinguish different stopped particles from each other and from particles which were not stopped. This is an essential part of the iPhos method which requires determining if a particle was stopped. However, the QPID method depends on the timing of the channel trigger. A recent analysis has shown that the time over threshold (ToT) can provide a competitive measure for the particle identification. This is surprising because the ToT is dominated by the $\tau_{RC} = 35\mu$s decay time of the preamplifier,

[93] Limited by the effect of angular straggling. Also, the position of the crystals in radial direction is determined by the tapering of the carbon fiber pockets and may vary on the millimeter scale between crystals.

[94] Another thing to keep in mind is that different rings of CALIFA are facing different points on the beam axis, with the focal point of the end cap being a few centimeters upstream of the focal point of the barrel around 90°. This means that no matter where the target is located, there are always rings where most tracks which go through the inner, target-side face of a crystal will not pass through the far face of that crystal, but instead pass through neighboring crystals.

while the dependence on the shape of the rising signal (which depends on the $N_f : N_s$ ratio) is much smaller, perhaps on the order of a few microseconds. So far only the first two emission processes from CsI(Tl) have been analysed in the CALIFA firmware for the purpose of PID [36, 16], But there is a third one which is much slower than the other two showing a time constant of 16 μs [15]. With the standard MWD window length of 10.8 μs only half of this time component will be captured in the measured energy amplitude.

At the moment, the timing branch does not provide a maximum amplitude for events. Besides this being a useful feature for adjusting noise thresholds[95], the amplitude should be sensitive to the $N_f : N_s$ ratio. The widely used `discr_gaptime` setting for this filter is 1 μs, which is longer than the time constant for the fast component but slower than the time constant for the slow component. Adding an amplitude measurement to the timing branch should be cheap in term of FPGA resources.

The FEBEX FPGA firmware used in CALIFA has three different leading edge discriminators. If the relative timing of these (before any walk correction) was measured[96], this might be sensitive to how fast the signal is rising.

The firmware allows the recording of the ADC samples after *decimation*, i.e. with a sampling rate of 25 MHz. Each sample increases the event size by two bytes, which makes recording of long *traces* (e.g. 10 μs$\hat{=}$500 Bytes) prohibitively expensive in terms of storage and bandwidth. Giving a standard event size of 56 bytes, *short* traces remain relatively cheap, though. Recording three samples baseline and then the first seventeen samples of the signal would give direct access to the amplitude in the interval $[0, \tau_f]$ in which the fast component dominates while also allowing to precisely reconstruct the starting point of the signal independently from the trigger timing. The price would be a doubling of the storage requirements. The main advantage of this method is that it does not require any modifications in the FPGA firmware but merely some minor additions to the analysis code.

The QPID itself could be improved by removing the reliance on the trigger timing. The amplitude determination works by taking the maximum of all the values encountered after the energy MAU filter, and there is no reason why a similar concept should not work on the fast component. This would define the timing of the first QPID window and by extension of all the QPID, with the second window simply having a fixed offset.[97]

Duplicating the MWD - MAU - energy stages would allow to measure both the charge collected within the nominal integration window and separately the charge collected within the first microsecond or so while keeping the neat separation between the fast branch and the slow branch.

Once redundant information about the shape of the signal (e.g. amplitude, QPID, ToT, trigger filter amplitude, trace start) are available, machine learning could be employed to distinguish different particles.

[95] Basically, a noise event (identified by being in the pedestal after the slow energy filter) would provide the threshold which would have been required to reject it.

[96] Naturally in steps of 20 ns.

[97] An earlier version of the CALIFA firmware had a flag which enabled peak-sensing for the fast component.

Getting ready for FAIR early science in phase 1

CALIFA performed very well in the 2024 experiments. With the external validation trigger, the system was generally found to be robust and provides a good foundation for running in the FAIR high energy cave.

On the mechanical side, it would be helpful to align both halves of CALIFA – which are movable in the xz plane – more precisely to the target, ideally with an electronic system to measure offsets.

On the DAQ side more improvements should be done to accommodate the high event rates expected from the SuperFRS. One area of concern here is the timestamp interpolation from FebexTS to WRTS described in section 5.7. At very high rates, this is difficult to accomplish online due to CPU requirements.

The FebexTS is just a clock count recorded in each module. Due to bit errors, the count between modules will slowly drift over time, so the interpolation to WRTS is done on a per-module basis.

Instead of just using a clock count, one could use a more sophisticated scheme which periodically encodes the full time stamp and thus prevents bit errors from accumulating. Directly connecting the FEBEX FPGAs to the White Rabbit network using fiberoptical links is not feasible but also not required.

The rataclock protocol[98], is a monodirectional protocol to encode both a clock and the current WRTS over a single digital line. By design, it does not offer the compensation of cable lengths or the ultra low jitter of the White Rabbit network. But this also means that it can be transported using NIM fan-outs and LEMO cables and decoding it on an FPGA does not take a lot of the FPGA resources. This protocol is already used for multiple detector systems within R^3B.

One of the trigger lines on the backplane of the FEBEX crates is currently used to transmit the 20 MHz clock which drives the FebexTS. Transmitting rataclock instead would require some minimal changes to the EXPLODER firmwares as well as the addition to the receiver core to the CALIFA FEBEX firmware. Merging hits from 2600 inputs by time stamp could then happen using generic, well-optimized tools maintained by a larger user base.

[98] https://fy.chalmers.se/~f96hajo/rataser/

A Supplemental calibration spectra

Figure 58 shows the parameters obtained from the calibration. The fact that most channels have a slope of about 1.4 is by design. During a process called **gain-matching**, the APD voltage is adjusted for every channel so that the 1333 keV peak of a ^{60}Co calibration source appears roughly in the same channel. One benefit of this procedure is that trigger threshold are somewhat universal. Also, the energy range (naturally limited by ADC overflows) is then equal for every crystal. Some channels have a lower gain (thus a higher m_{ch}). This is because increasing the APD voltage is not without trade-offs: while a higher voltage means a higher gain, voltages near the breakthrough voltage will also have a worse resolution due to a very steep gain curve at this point.

The offsets c_{ch} are displayed in figure 58. While most channels have very small offsets, for some the offsets are very large with absolute values of multiple thousands of ADC units. The likely explanation for this is the baseline reconstruction failing. Baselines are not updated after an internal trigger threshold is reached for a time to prevent the signal distorting the baseline estimate. If the trigger rate for a channel is sufficiently high, it may well be that the baseline of a channel is effectively never updated in the FEBEX module, leading to a systematic shift of all detected energies. During data recording, this can be prevented by setting the `energy_bias_correct` firmware parameter as discussed in section 4.2.

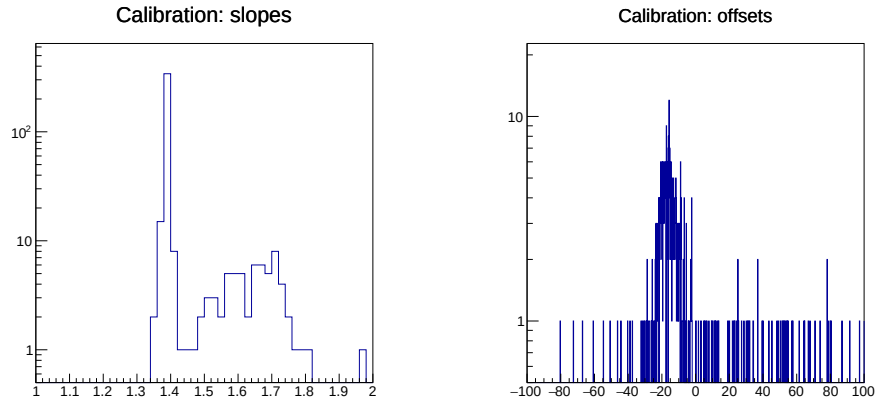

Figure 58: The calibration parameters obtained from the AmBe data. Linear functions $cal_{ch}(x) = m_{ch}x + c_{ch}$ were used to convert the arbitrary hardware units to keV.

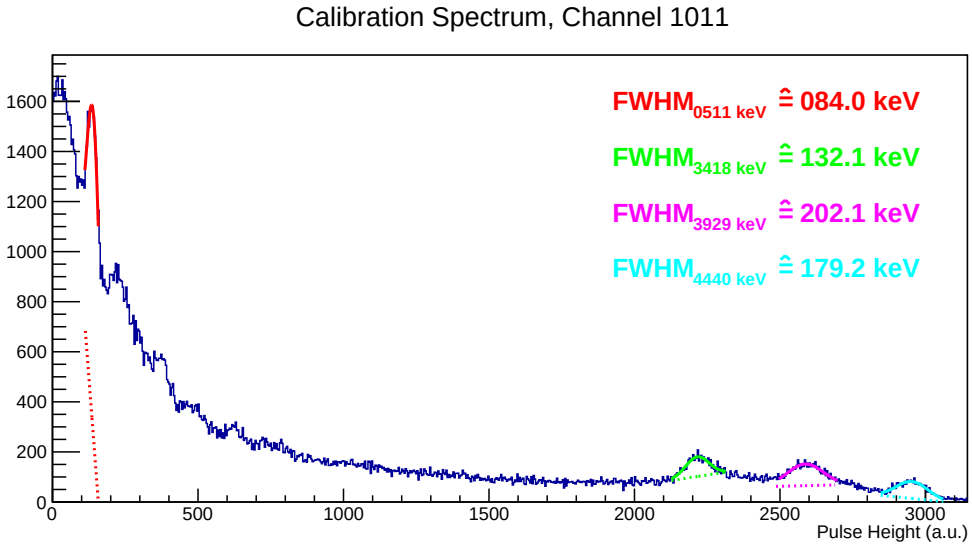

Figure 59: The uncalibrated spectrum for channel 1011, with gaussian fits used for calibration. This channel has a large offset and a substandard resolution at 511 keV, but works fine in the 4 MeV region.

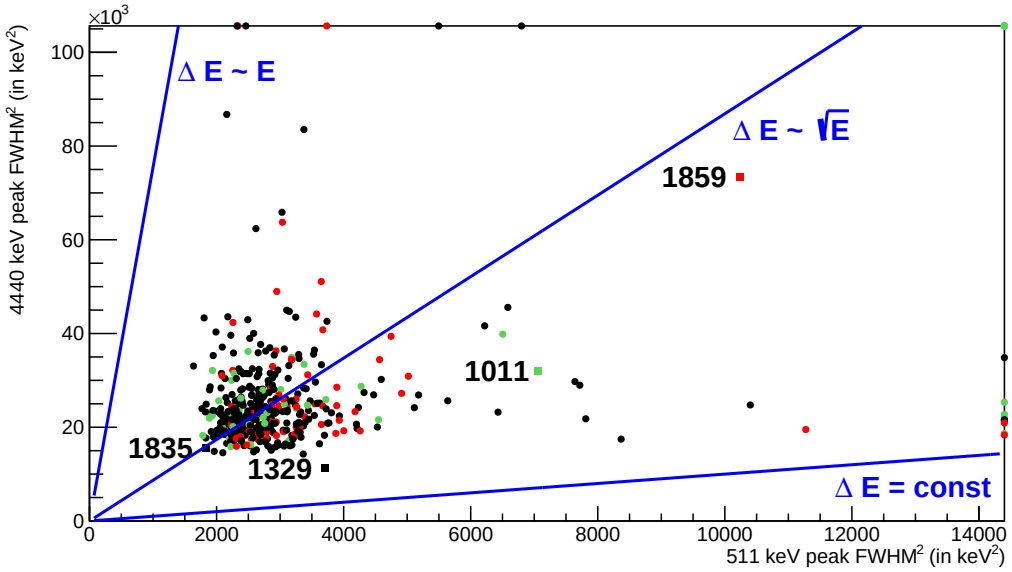

Figure 60: Figure 50, with both axes squared. The benefit of this is that given a decomposition of the resolution due to different effects, i.e. $\Delta E_{tot}(E)^2 = A^2 + (B \times \sqrt{E})^2 + (C \times E)^2$, the different components can be represented as unit vectors along the blue lines and one could use vector addition to reach the total resolution. Alas, which linear combination of three vectors goes to a point in $\mathbb{R}^2$ is underdetermined. If one had the resolution at a third energy on a third axis, one could uniquely decompose any crystal's resolution and determine A, B and C for it by solving a linear system of equations.

B Handling multi-hit events in 2d histograms

Similar considerations apply to 2d histograms visualizing coincidences: A quantity – such as the hit energy – observed in half of the detector will be plotted against a quantity in the other half.

This works very easily as long as there is at most one hit per detector half, but if there is more than one hit, complications ensue.

If we have, within our coincidence window, $N_0 > 1$ hits on the one side and $N_1 \geq 1$ hits on the other side, different approaches could be followed:

- **Summation**: Just take the sum of all the energies for each side, and plot both sums against each other.
- **All pairs**: Fill all $N_0 N_1$ pairs in the histogram.
- **All pairs, scaled**: Scale all pairs by $\frac{1}{N_0 N_1}$
- **Random pair**: Just select one random hit from each side,
- **Highest energy pair**

The best approach depends on the intrinsic hit multiplicity of the underlying process and the detector geometry. With a 4π detector, splitting the volume into two parts and summing over these two would obviously prevent the detection of any processes which result into three separate hits, as two of these will always be part of the same sum.

In general, the marginal distributions (or *projections*, in ROOT slang) will not be identical to the histograms one would obtain when directly filling 1 dimensional histograms. Energies which often coincide with high multiplicities will be overrepresented in the marginal distributions.

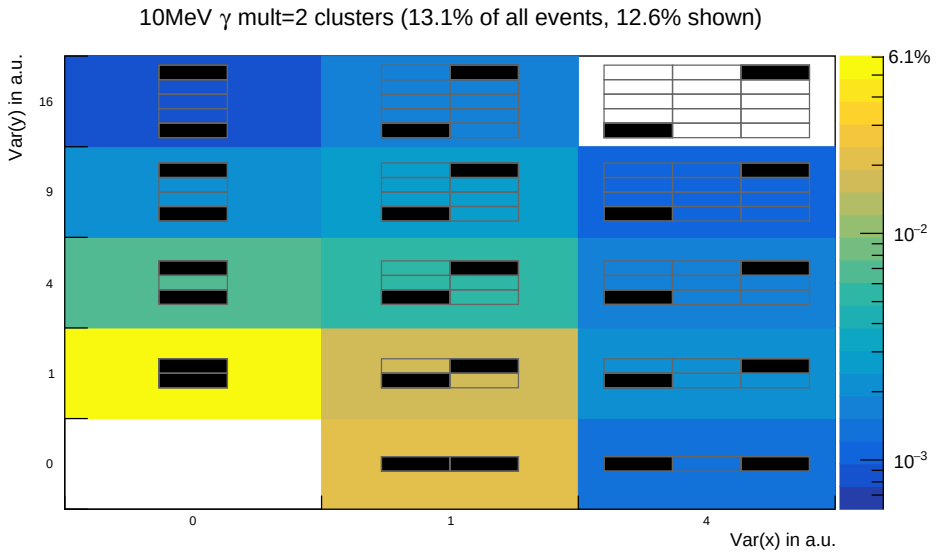

Figure 61: Multiplicity two clusters for 10 MeV photons. Here, the variance uniquely describes the cluster shape.

C Supplemental figures from the MC simulation

Within each histogram bin, a representative of such a cluster was drawn. All of the cluster plots were produced with the (x, y) position of the primary *smeared* as defined in section 7.4.2.

Figure 65 gives the average energy distribution in CsI for a 200 MeV proton. In general, the radii of proton clusters are generally smaller than gamma cluster radii, so clustering protons is an easier task.

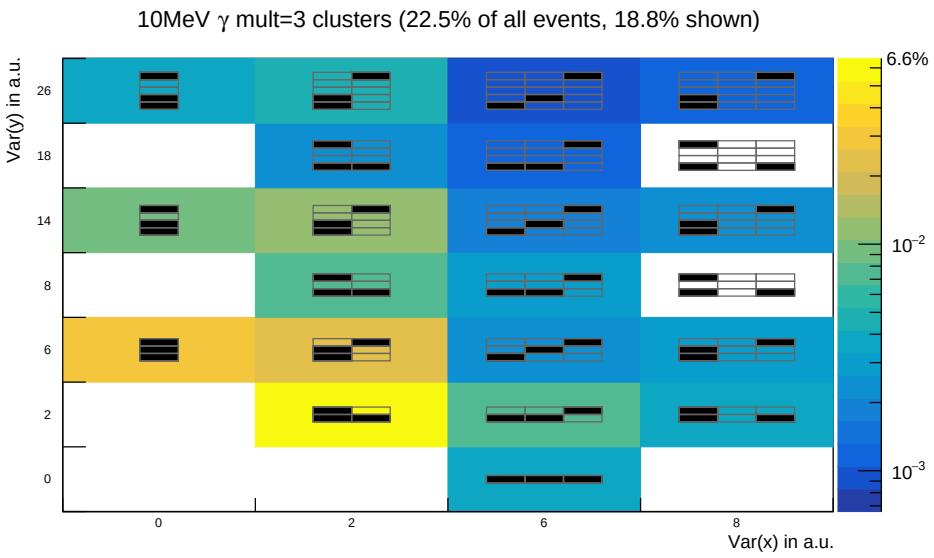

Figure 62: Multiplicity three clusters for 10 MeV photons.

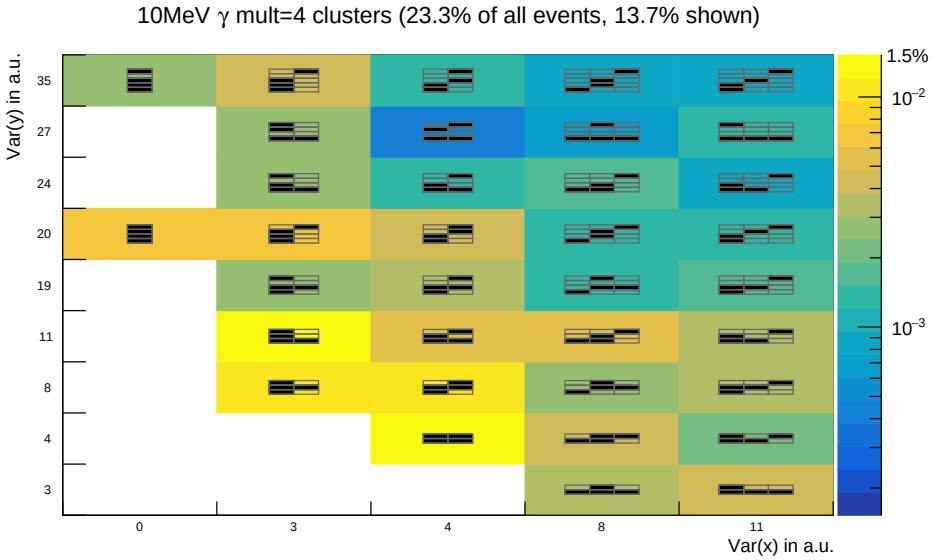

Figure 63: Multiplicity four clusters for 10 MeV photons.

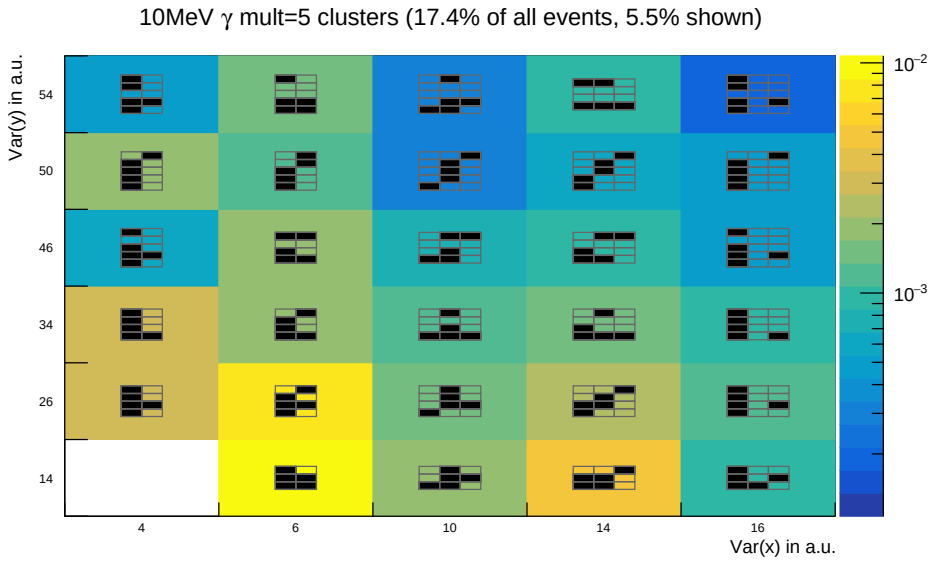

Figure 64: Multiplicity five clusters for 10 MeV photons.

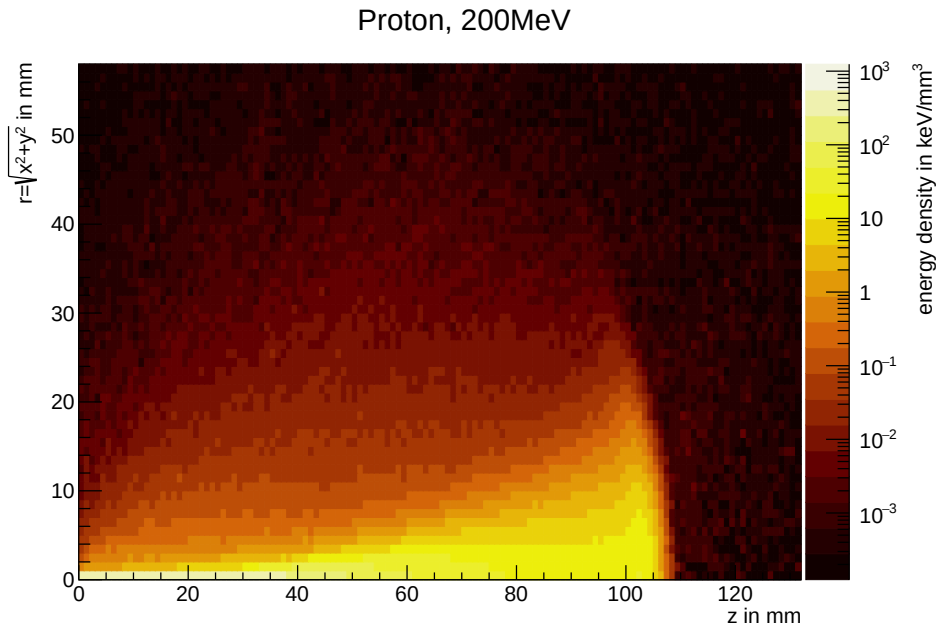

Figure 65: The energy density distribution for a 200 MeV proton. Clearly visible is the Bragg peak and the angular straggling in the material. Faint energy deposits beyond the main features are attributed to nuclear reactions.

D Lorentz boosts

The quantities in the center of mass frame of reference (CMS) are of physical interest while the quantities in the lab frame can be measured by detectors. As the ion is generally travelling with relativistic velocities (as indicated by one of the **R**'s in R^3B), the conversion between quantities between these two frames of reference is nontrivial and involves the Lorentz transformation.

A particle[99] emitted in the beam frame with energy E_{CMS}, three-momentum p_{CMS} and an emission angle θ_{CMS} to the beam axis will be observed in the lab frame of reference with

$$\begin{aligned}
E_{Lab} &= \gamma\left(E_{CMS} - \beta p_{z,CMS}\right) \\
&= \gamma\left(E_{CMS} - \beta |p_{CMS}| \cos\theta_{CMS}\right) \\
p_{\perp,Lab} &= p_{\perp,CMS} \\
p_{z,Lab} &= \gamma(p_{z,CMS} - \beta \frac{E_{CMS}}{c}) \\
\tan\theta_{Lab} &= \frac{p_{\perp,Lab}}{p_{z,Lab}} \\
&= \frac{p_{\perp,CMS}}{\gamma(p_{z,CMS} - \beta \frac{E_{CMS}}{c})} \\
&= \frac{|p_{CMS}| \sin\theta_{CMS}}{\gamma(|p_{CMS}| \cos\theta_{CMS} - \beta \frac{E_{CMS}}{c})} \\
&= \frac{\sin\theta_{CMS}}{\gamma(\cos\theta_{CMS} - \beta\sqrt{1 + \frac{m^2c^2}{|p_{CMS}|^2}})}
\end{aligned}$$

[99] *Particle* here emphatically includes photons, for which the application of a Lorentz boost is also known as a (relativistic) Doppler Shift.

E Two ways to measure energies

Calorimeters

Calorimeters determine energies of particles[100] by stopping them in an active detector material where part of the energy deposited by the particle will be, well, detectable. Different properties of materials can be used for that. Some transparent materials emit light (that is to say, low energy photons) when subjected to ionizing radiation in a process called **scintillation**. Reverse biased diodes will react with a current flow, which is used in **semiconductor detectors** such as high purity germanium (HPGe) spectrometers. More exotically, matter cooled to nearly absolute zero may be warmed up by a detectable amount in **cryogenic detectors**.

Calorimetry is suitable for particles ($m = 0$, $m \neq 0$) which mostly react with either electromagnetically in matter (e.g. photons, charged particles) or by many nuclear interactions (e.g. a relativistical pion causing a hadronic shower).

As these interactions are quantum mechanical processes, the outcome is always governed by statistics. The amount of photons produced by a scintillator, for example, will follow some statistical distribution, with the mean hopefully proportional to the energy deposited in the detector. Thus, the effective resolution is governed by the mean amount of photons (for scintillators) generated per unit of energy absorbed.

The main drawback of calorimeters is that the particle must actually be stopped in the detector, this generally requires huge and heavy detectors when the particle energy is high.

Time of flight

For any massive particle, the relation

$$E = \frac{m}{\sqrt{1 - \left(\frac{v}{c}\right)^2}}$$

holds. If the mass m of a particle is known, its energy E can thus easily be calculated from its velocity v. This is less helpful in the highly relativistic case because tiny differences between v and c will have a huge impact on E, but works great if v is around $0.8c$.

To apply this, the mass of the particle must be known. This should not be considered a drawback because the mass is in any case one of the key paramters in any analysis. Luckily, nature is accomodating here by only providing a discrete spectrum of particle masses, so the mass does not have to be measured exactly, but just at better resolution than 1 u. For ions, measuring the mass number A typically involves measuring both its magnetic rigidity $\frac{B\rho}{\beta\gamma} \sim \frac{A}{Z}$, from the bending angle in a known magnetic field and its charge from the energy loss in a thin tracking detector using Bethe-Bloch[101]: $\Delta E \sim Z^2$.

[100] Again, while the term "particle" has historically been used in less enlightened times to refer to massive particles, I use it to also include massless particles, e.g. photons. Devices which measure the energy of photons this way are also called spectrometers.

[101] For simplicity, the dependencies on β have been ommited here. This is reasonable because $\frac{\Delta\beta}{\beta}$ is very small for beam-like fragments

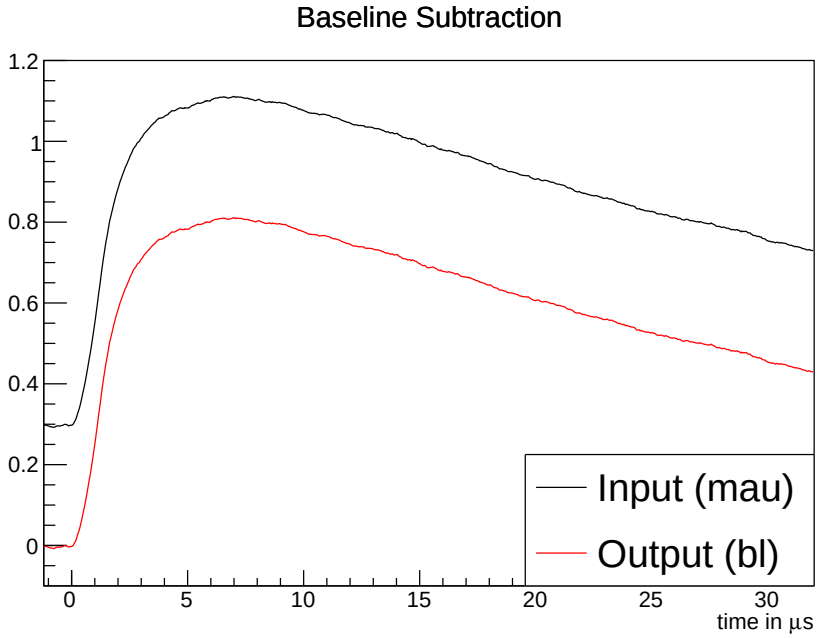

Figure 66: A model of the baseline subtraction operation. The smoothed trace resulting from the MAU is shifted by a fixed amount, which is intended to bring the baseline to zero.

In inverse kinematics, a beam-like fragment has, by definition, roughly the velocity of the beam. While its x, y momentum can be measured by tracking its position relative to the beam, measuring its z momentum in the CMS thus requires a very high resolution measurement of its velocity. As the time measurement is limited by the response time of the detectors[102], this generally requires a large distance over which the time of flight is measured.

[102]Effectively some tens of picoseconds for plastic scintillators

Thanks

Stepp'd in so far that, should I wade no more,
Returning were as tedious as go o'er.

Now that this thesis is finished and I find myself in the position of *having to learn enough nuclear physics to have a chance to survive my defense*, plus *an extra serving of dealing with bureaucracy*, as well as the possibility of gainful employment should I pass, I think I should take a moment to look back on the meandering process of my PhD and the people who supported me.

Roughly in order of appearance.

My parents had fostered my enthusiasm in science from early on and have supported me throughout my lengthly stay in academia. Thank you! Just one quick remark on names: Philipp is a great first name, but if you had had the foresight to give me a second name starting with the letter D – Daniel perhaps – then I could have applied at GSI as *Klenze, Ph. D.* and might have avoided this lengthly endeavor.

My former girlfriend Liz recommeneded I should write to chair E12 about a PhD position a decade ago, which I did. As specifically requested by her, I want to thank her for her finite patience and support.

At TUM, the first and foremost person I should thank is Roman. He took me on as a PhD student and has supported me well throughout it, especially during the writing. His feedback was instrumental in turning my initial drafts into a proper thesis.

Over my PhD, I have worked with a few generations of TUM Califa PhD students. I want to thank Michael, Max (whose FEBEX firmware was mentioned once or twice in this thesis), Benny, Lukas and Tobias. Michael Böhmer from the electronics department was very helpful with the design of circuit boards, as well as commentary about the quality of various readout systems. The late professor Shawn Bishop taught me a colorful Canadian idiom for *avoiding work which leads to a productive outcome* (such as a PhD – as I realized belatedly). Thanks go out to all of them!

When I transitioned to GSI for what I like to call *phase one* of my PhD, I was supported by a lot of colleagues. Leyla was leading the effort to get me to wade a bit further down the path to my thesis, and was supported by Andrea and Haik in these efforts, thanks go out to all of them. Many other colleagues helped in creating an interesting working environment for what might be my future job, I want to thank Anna-Lena, Bastian, Håkan, Hans, Luke, Martin and Michael for that.

Kathrin tried her best to teach me scientific writing. While my thesis was too far along when I started her course to fully benefit from her training, I would generally recommend it to other PhD students.

I want to thank Alexandra E. for library services.

Finally, I want to thank my prospective committee: professor Laura Fabbietti – who has graciously inherited me as a PhD student – as well as professor Lothar Oberauer and professor Norbert Kaiser.

Bibliography

Care was taken to provide a link to a freely accessible download for each reference. In cases where this was not possible, the digital object identifier (DOI) was linked instead. Rumors suggest the existence of some "hub" website where publications can be easily downloaded given the DOI.

[1] J. Chadwick. Possible Existence of a Neutron. *Nature*, 129, 1932. https://doi.org/10.1038/129312a0

[2] International Atomic Energy Agency. Nuclear Data Section. https://www-nds.iaea.org/

[3] H. Geissel et al. Technical Design Report on the Super-FRS. 12 2008. http://repository.gsi.de/record/54552/files/GSI-2013-05264.pdf?version=1

[4] H. H. Gutbrod, K. D. Gross, W. F. Henning, and V. Metag. An international accelerator facility for beams of ions and anti-protons. Conceptual design report. 11 2001.

[5] Hans Geissel, P Armbruster, K Hetal Behr, A Brünle, K Burkard, M Chen, H Folger, B Franczak, H Keller, O Klepper, et al. The GSI projectile fragment separator (FRS): a versatile magnetic system for relativistic heavy ions. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 70(1-4):286–297, 1992.

[6] Kiyomi Ikeda, Noboru Takigawa, and Hisashi Horiuchi. The Systematic Structure-Change into the Molecule-like Structures in the Self-Conjugate 4n Nuclei. *Progress of Theoretical Physics Supplement*, E68:464–475, 07 1968. https://doi.org/10.1143/PTPS.E68.464

[7] Jacek Okołowicz, Marek Płoszajczak, and Witold Nazarewicz. On the Origin of Nuclear Clustering. *Progress of Theoretical Physics Supplement*, 196:230–243, 10 2012. https://academic.oup.com/ptps/article-pdf/doi/10.1143/PTPS.196.230/5342456/196-230.pdf

[8] Stefanos Paschalis et al. Technical Report for the Design, Construction and Commissioning of the Tracking Detectors for R3B. 11 2014. https://edms.cern.ch/ui/file/1865815/2/TDR_R3B_TrackingDetectors_public.pdf

[9] D. Cortina et al. Technical Report for the Design, Construction and Commissioning of The CALIFA Barrel. Technical report, R^3B, 11 2011. https://edms.cern.ch/file/1833500/2/TDR_R3B_CALIFA_BARREL_public.pdf

[10] D. Cortina et al. Technical Report for the Design, Construction and Commissioning of the CALIFA Endcap. Technical report, R^3B, 08 2015. https://edms.cern.ch/file/1833748/2/TDR_R3B_CALIFA_ENDCAP_public.pdf

[11] B. Gastineau, C. Mayri, B. Baudouy, C. Berriaud, G. Disset, A. Donati, J.-E. Ducret, D. Eppelle, P. Fazilleau, P. Graffin, J.-L. Jannin, D. Loiseau, J.-P. Lottin, M. Massinger, C. Pes, Y. Queinec, Z. Sun, P. Charon, P. Contrepois, and H. Neyrial. Progress in Design and Construction of the R^3B -GLAD Large Acceptance Superconducting Dipole Spectrometer for GSI-FAIR. *IEEE Transactions on Applied Superconductivity*, 20(3):328–331, 2010. https://doi.org/10.1109/TASC.2010.2040169

[12] T. Aumann et al. Technical Report for the Design, Construction and Commissioning of NeuLAND: The High-Resolution Neutron Time-of-Flight Spectrometer for R^3B. 11 2011. https://edms.cern.ch/file/1865739/2/TDR_R3B_NeuLAND_public.pdf

[13] JOEL Pouthas, B Borderie, R Dayras, E Plagnol, MF Rivet, F Saint-Laurent, JC Steckmeyer, G Auger, Ch O Bacri, S Barbey, et al. INDRA, a 4π charged product detection array at GANIL. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 357(2-3):418–442, 1995.

[14] P.A. Zyla et al. Review of Particle Physics, Chapter 35: Particle Detectors at Accelerators. *PTEP*, 2020(8):083C01, 2020. https://pdg.lbl.gov/2021/reviews/rpp2021-rev-particle-detectors-accel.pdf

[15] Agnieszka Syntfeld-Kazuch, Marek Moszynski, Lukasz Swiderski, Wlodzimierz Klamra, and Antoni Nassalski. Light Pulse Shape Dependence on *gamma*-Ray Energy in CsI (Tl). *IEEE Transactions on Nuclear Science*, 55(3):1246–1250, 2008.

[16] Max Oliver Winkel. *Komplexe Pulsformalgorithmen und Teilchenidentifikation zur Echtzeit-Implementierung in CALIFA*. Dissertation, Technische Universität München, München, 2016.

[17] Max Oliver Winkel. *Implementierung und Erprobung einer digitalen Pulsformanalyse zur Auslese von Kalorimetern*. Diploma thesis, Technische Universität München, München, 2011.

[18] A. Georgiev and W. Gast. Digital pulse processing in high resolution, high throughput, gamma-ray spectroscopy. *IEEE Transactions on Nuclear Science*, 40(4):770–779, 1993. https://doi.org/10.1109/23.256659

[19] H. Johansson et al. ucesb - unpack & check every single bit. http://fy.chalmers.se/~f96hajo/ucesb/

[20] Håkan T Johansson. *Hunting tools beyond the driplines*. PhD thesis, Chalmers Tekniska Hogskola (Sweden), 2010. http://fy.chalmers.se/~f96hajo/phd/htj_thesis.pdf

[21] Stephen Seltzer. XCOM-Photon Cross Sections Database, NIST Standard Reference Database 8, 2010. http://www.nist.gov/pml/data/xcom/index.cfm

[22] P. Teubig, P. Remmels, P. Klenze, H. Alvarez-Pol, E. Alves, J. M. Boillos, P. Cabanelas, R. C. da Silva, D. Cortina-Gil, J. Cruz, D. Ferreira, M. Fonseca, D. Galaviz, E. Galiana, R. Gernhäuser, D. González, A. Henriques, A. P. Jesus, H. Luís, J. Machado, L. Peralta, J. Rocha, A. M. Sánchez-Benítez, H. Silva, and P. Velho. Challenging the Calorimeter CALIFA for FAIR Using High Energetic Photons. In José-Enrique García-Ramos, María V. Andrés, José A. Lay Valera, Antonio M. Moro, and Francisco Pérez-Bernal, editors, *Basic Concepts in Nuclear Physics: Theory, Experiments and Applications*, pages 245–246, Cham, 2019. Springer International Publishing. https://doi.org/10.1007/978-3-030-22204-8

[23] Sea Agostinelli, John Allison, K al Amako, John Apostolakis, H Araujo, Pedro Arce, Makoto Asai, D Axen, Swagato Banerjee, GJNI Barrand, et al. GEANT4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003. https://www.sciencedirect.com/science/article/abs/pii/S0168900203013688

[24] W. Bothe and H. Becker. Künstliche Erregung von Kern-γ-Strahlen. *Zeitschrift für Physik*, 66(5-6):289–306, May 1930. https://doi.org/10.1007/bf01390908

[25] HR Vega C and SA Martinez O. Neutron spectra and dosimetric features of isotopic neutron sources: a review, 2015. https://inis.iaea.org/collection/NCLCollectionStore/_Public/47/032/47032302.pdf

[26] K.W. Geiger and L. Van Der Zwan. Radioactive neutron source spectra from ^{9}Be(α, n) cross section data. *Nuclear Instruments and Methods*, 131(2):315–321, December 1975. https://doi.org/10.1016/0029-554x(75)90336-5

[27] Zhenzhou Liu, Jinxiang Chen, Pei Zhu, Yongming Li, and Guohui Zhang. The 4.438MeV gamma to neutron ratio for the Am–Be neutron source. *Applied Radiation and Isotopes*, 65(12):1318–1321, December 2007. https://doi.org/10.1016/j.apradiso.2007.04.007

[28] Isao Murata, Iehito Tsuda, Ryotaro Nakamura, Shoko Nakayama, Masao Matsumoto, and Hiroyuki Miyamaru. Neutron and gamma-ray source-term characterization of AmBe sources in Osaka University. *Progress in Nuclear Science and Technology*, 4:345–348, 2014. https://doi.org/10.15669/pnst.4.345

[29] DW Aitken, BL Beron, G Yenicay, and HR Zulliger. The fluorescent response of NaI (Tl), CsI (Tl), CsI (Na) and CaF2 (Eu) to X-rays and low energy gamma rays. *IEEE Transactions on Nuclear Science*, 14(1):468–477, 1967.

[30] M. Pârlog, B. Borderie, M.F. Rivet, G. Tăbăcaru, A. Chbihi, M. Elouardi, N. Le Neindre, O. Lopez, E. Plagnol, L. Tassan-Got, G. Auger, Ch.O. Bacri, N. Bellaize, F. Bocage, R. Bougault, B. Bouriquet, R. Brou, P. Buchet, J.L. Charvet, J. Colin, D. Cussol, R. Dayras, A. Demeyer, D. Doré, D. Durand, J.D. Frankland, E. Galichet, E. Genouin-Duhamel, E. Gerlic, S. Hudan, D. Guinet, P. Lautesse, F. Lavaud, J.L. Laville,

J.F. Lecolley, C. Leduc, R. Legrain, M. Louvel, A.M. Maskay, L. Nalpas, J. Normand, J. Péter, E. Rosato, F. Saint-Laurent, J.C. Steckmeyer, B. Tamain, O. Tirel, E. Vient, C. Volant, and J.P. Wieleczko. Response of CsI(Tl) scintillators over a large range in energy and atomic number of ions. Part I: recombination and δ-electrons. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 482(3):674–692, 2002. https://www.sciencedirect.com/science/article/pii/S0168900201017107

[31] National Nuclear Data Center. NuDat database. https://www.nndc.bnl.gov/nudat/

[32] Zhenzhou Liu, Jinxiang Chen, Pei Zhu, Yongming Li, and Guohui Zhang. The 4.438 MeV gamma to neutron ratio for the Am–Be neutron source. *Applied Radiation and Isotopes*, 65(12):1318–1321, 2007.

[33] Andre Robert, Maurice Blondel, Jean Morel, and Claude Thomas. A 6.13 MeV gamma reference source, measurement of the emission rate. Technical report, CEA Centre d'Etudes Nucleaires de Saclay, 1977.

[34] Joel G Rogers, Mark S Andreaco, Christian Moisan, and Ian M Thorson. A 7–9 MeV isotopic gamma-ray source for detector testing. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 413(2):249–254, 1998. https://www.sciencedirect.com/science/article/pii/S0168900298000977

[35] Benjamin Heiss. *Development of a CALIFA demonstrator benchmark experiment using (p,2p) reactions in direct kinematics.* Dissertation, Technische Universität München, München, 2018.

[36] Michael Bendel. *Entwicklung einer neuartigen Nachweismethode hochenergetischer Teilchen im CALIFA-Kalorimeter.* Dissertation, Technische Universität München, München, 2014.

[37] B. Pietras, M. Winkel, H. Alvarez-Pol, M. Bendel, E. Casarejos, J. Cederkäll, D. Cortina-Gil, G. Fernandez, R. Gernhäuser, P. Golubev, D. González, A. Hartig, P. Izquierdo, P. Klenze, T. Le Bleis, E. Nácher, A. Perea, P. Remmels, G. Ribeiro, P. Teubig, J. Vilan, and P. Yañez. First testing of the CALIFA Barrel Demonstrator. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 814:56 – 65, 2016. https://www.researchgate.net/publication/296716124

[38] Joochun Park, B. Heiss, A.-L. Hartig, H. Alvarez-Pol, G. Bruni, E. Casarejos, Joakim Cederkäll, D. Cortina-Gil, P. Díaz Fernández, D. Galaviz, R. Gernhäuser, Pavel Golubev, A. Heinz, H. T. Johansson, P. Klenze, Alexander Knyazev, T. Kröll, T. Nilsson, A. Perea, H. B. Rhee, O. Tengblad, P. Teubig, and M. Zieblinski. Tests of CALIFA Barrel modules at CCB in Krakow. *GSI-FAIR Scientific Report*, 2018(1), 2018. https://lup.lub.lu.se/search/ws/files/61678599/GSI_Report_2018_Krakow_test.pdf