

Exploration of new Data Acquisition and Background Reduction Techniques for the COBRA Experiment

Dissertation
zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften
der Fakultät Physik
an der Technischen Universität Dortmund



vorgelegt von
DIPL. PHYS. OLIVER SCHULZ

Lehrstuhl für Experimentelle Physik IV
Fakultät Physik



Dortmund, den 11. August 2011

Exploration of new Data Acquisition and Background Reduction Techniques for the COBRA Experiment

Dissertation
zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften
der Fakultät Physik
an der Technischen Universität Dortmund

vorgelegt von
DIPL. PHYS. OLIVER SCHULZ

Teilergebnisse dieser Arbeit waren Gegenstand verschiedener Veröffentlichungen und Tagungsbeiträge. Eine Liste befindet sich am Ende der Arbeit ab Seite 101.

For my mother - you helped me discover my love for science as a child.

In memory of my father - you were always there.

Contents

1	Introduction	9
1.1	Open Questions in Neutrino Physics	9
1.2	Neutrinoless Double-Beta Decay	11
1.3	The COBRA Experiment	13
1.3.1	Isotopes and decays of primary interest in this work	13
1.3.2	Detector Technologies	14
1.3.3	Pixel Detectors	16
2	Background Improvements	17
2.1	Dominating Background Sources	17
2.2	New Detector Contacting System	18
2.2.1	Pin Pressure Contacts	18
2.2.2	Gallium soldering	20
2.2.3	Silver-filled Conductive Lacquer	20
2.3	Nitrogen Flushing System	21
2.3.1	Radon Trap Modifications	22
2.4	Alternative Radiopure Detector Coatings	23
2.5	Background Reduction	28
3	Polaris CdZnTe Pixel Detector System	31
3.1	Design Considerations for Pixel Detectors	31
3.2	The Polaris System	32
3.3	System Properties	33
3.4	Underground Installation	34
3.5	Results	35
4	CPG Detectors with Pulse-Shape Readout	39
4.1	Previous DAQ System	39
4.2	COBRA DAQ Electronics Upgrade	40
4.2.1	New Data-Acquisition Chain	40
4.2.2	Pre-Amplifiers	41
4.2.3	Fast Linear Amplifiers	43
4.2.4	Fast Analog-Digital Converters	45
4.2.5	Scalability considerations	50

Contents

4.3	Pulse-Shape Analysis	50
4.3.1	Anode Pulse Shapes	50
4.3.2	Dithering	52
4.3.3	Signal Calibration	54
4.3.4	Deconvolution	54
4.3.5	Weighted Anode Difference	56
4.3.6	Cathode-Signal Reconstruction	56
4.3.7	Window Functions	56
4.3.8	Pulse Properties	57
4.3.9	Analysis Calibration Parameters	59
4.3.10	Physical Quantities	61
4.4	Software implementation	65
4.5	Data Compression Options	65
4.5.1	Wavelets and Multiresolution	66
4.5.2	Wavelet compression	68
4.5.3	Compression results	69
5	COBRA DAQ System	75
5.1	Requirements	76
5.2	Design Philosophy	77
5.2.1	Network-centric Approach	77
5.2.2	The Actor Model	77
5.2.3	Error handling and supervision	79
5.2.4	JVM-Based	81
5.2.5	Handling Embedded Systems	81
5.2.6	Handling Latency	82
5.3	Data-Acquisition Hardware	82
5.3.1	Physics Data Acquisition	82
5.3.2	SCADA / Slow Control	82
5.4	Software Implementation	83
5.4.1	Choice of Programming Language	83
5.4.2	Communication protocol support	84
5.4.3	Data Management - CouchDB	84
6	Low-Background CPG Results	87
6.1	Pulse-Shape Readout under Low-Background Conditions	87
6.2	Pathological Pulse Shapes	89
6.3	Background Spectrum	90
7	Summary and Outlook	97
	Acknowledgements	99

Contents

Publications	101
Bibliography	103
List of Figures	111
Acronyms	116

Contents

1 Introduction

1.1 Open Questions in Neutrino Physics

The field of neutrino physics has seen rapid changes during the last two decades. Neutrinos had originally been supposed to be massless, though extensions of the Standard Model of Particle Physics existed which allowed for neutrinos with non-zero rest mass. When the existence of neutrino oscillations was confirmed by the Super-Kamiokande Experiment [The98] in 1998, establishing the fact that neutrinos do indeed possess a rest mass, a new chapter in neutrino physics began. The results were confirmed by numerous other experiments during the following years and the field entered an era of precision measurements.

Neutrino oscillations occur because the flavour eigenstates of the neutrino (the electron neutrino ν_e , the muon neutrino ν_μ and the tau neutrino ν_τ) are not identical to its mass eigenstates (ν_1, ν_2, ν_3). Instead, the two eigenstate bases are coupled by the Pontecorvo-Maki-Nakagava-Sakata (PMNS) matrix [MN62]:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}. \quad (1.1)$$

The existence of neutrino oscillations therefore implies the existence of different neutrino masses. However, while oscillation experiments have been able to measure the difference between these masses with high precision, they cannot determine the absolute mass scale. In fact, even the order of the neutrino masses is still unknown (see Fig. 1.1).

A way to directly measure absolute masses is to determine the endpoint of a β -decay energy spectrum. Since neutrino masses are known to be small (a current upper limit for the mass of the ν_e is 2 eV, determined by the Mainz-Experiment [K⁺05]), this is, however, a very challenging task. The Katrin-Experiment [The01], currently under construction, and the Mare-Experiment [A⁺06a], currently in the prototyping stage, will use this approach to probe the mass scale below 2 eV (down to 0.2 eV in the case of Katrin).

The second possible access to the absolute neutrino mass is given by the neutrinoless double-beta decay ($0\nu\beta\beta$ -decay).

1 Introduction




Figure 1.1: The two possible neutrino mass hierarchies, called the normal and the inverted hierarchy [Mü07, p. 11]. The absolute scale of the masses is still unknown.




Figure 1.2: Simple β -decay is energetically forbidden for some isotopes [Mü07]. They can only decay via a two-step process.




Figure 1.3: $2\nu\beta\beta$ -decay and $0\nu\beta\beta$ -decay. In the neutrinoless case, the two decays are coupled by Majorana neutrino undergoing a helicity flip. While the energy spectrum of the $2\nu\beta\beta$ -decay is continuous, the spectrum of the $0\nu\beta\beta$ -decay consists of a single line, since no energy is lost.

1.2 Neutrinoless Double-Beta Decay

Some isotopes cannot undergo simple β -decay

$$(A, Z) \rightarrow (A, Z + 1) + e^- + \bar{\nu}_e \quad (1.2)$$

$$(A, Z) \rightarrow (A, Z - 1) + e^+ + \nu_e \quad (1.3)$$

for energetic reasons, as their decay product would have higher mass than the original nucleus. They can, however, undergo a double decay to a nucleus with $Z \pm 2$ (see Fig. 1.2):

$$(A, Z) \rightarrow (A, Z + 2) + 2e^- + 2\bar{\nu}_e \quad (1.4)$$

$$(A, Z) \rightarrow (A, Z - 2) + 2e^+ + 2\nu_e \quad (1.5)$$

Such decays are called 2-neutrino double-beta decays ($2\nu\beta\beta$ -decays). As they are second-order processes, they have very long half-lives.

Within the Standard Model, the neutrino has to be a Dirac particle. As such, the neutrino and antineutrino (of the same flavour) are different particles with a fixed helicity. Beyond the boundaries of the Standard Model, however, the neutrino may also be a Majorana particle - in this case the neutrinos are their own antiparticles.

Such a Majorana neutrino with a non-zero rest mass can undergo a helicity flip. This gives rise to the $0\nu\beta\beta$ -decay [Rac37, BV87]: In a double-beta decay, a majorana antineutrino emitted by one of the decays can flip its helicity and enter the second decay vertex as a neutrino (see Fig. 1.3), resulting in a process that emits no neutrino at all:

$$(A, Z) \rightarrow (A, Z + 2) + 2e^- \quad (1.6)$$

$$(A, Z) \rightarrow (A, Z - 2) + 2e^+ \quad (1.7)$$

1 Introduction

The half life of this type of decay is directly coupled to the so-called effective Majorana neutrino mass by

$$\left(T_{1/2}^{0\nu}\right)^{-1} = G^{0\nu}(Q, Z) \left|M_{GT}^{0\nu} - M_F^{0\nu}\right|^2 \left(\frac{\langle m_{\nu_e} \rangle}{m_e}\right)^2 \quad (1.8)$$

where $M_{GT}^{0\nu}$ and $M_F^{0\nu}$ denote the Gamow-Teller and the Fermi nuclear matrix elements and $G^{0\nu}(Q, Z)$ is the phase space integral. A main source of uncertainty arises from the numerical calculation of the matrix elements. As $G^{0\nu}(Q, Z)$ scales with Q^5 [Zub04], large Q -values lead to shorter half-lives.

Thus, investigation of $0\nu\beta\beta$ -decays can give access to the absolute neutrino mass scale and at the same time answer the question whether the neutrino is a majorana or a dirac particle.

However, the measurement of such decays is not an easy task as the expected half-lives are on the order of 10^{25} to 10^{26} years. To detect these rare decays, an experiment must have excellent shielding against all types of background radiation: The only effective shielding against the cosmic background is to locate the experiment deeply underground. A sophisticated shielding must protect the setup from environmental radiation, while all parts of the experiment near the detectors must be made from carefully selected radiopure materials. Also due to the long half-lives, the amount of the isotope under study (the source mass) must be high to detect at least a few events during a run time of several years. Finally, the $2\nu\beta\beta$ -decay causes an unavoidable, intrinsic background, since the endpoint of its energy distribution partially overlaps the $0\nu\beta\beta$ -decay peak. The experiment must therefore use detectors with a good energy resolution, to keep this peak as narrow as possible and enable separation between the two decay modes.

The sensitivity of an experiment to the half-life time of a rare decay can be estimated by

$$T_{1/2} \propto \varepsilon \cdot a \cdot \sqrt{\frac{M \cdot t_{live}}{\theta_B \cdot \Delta E}} \quad (1.9)$$

where t_{live} is the livetime of the experiment, ΔE the detector energy resolution, M the source mass, a the natural abundance of the examined isotope, θ_B the background radiation level and ε the detector efficiency.

It follows that a low background level, a high detection efficiency and a good energy resolution are critical for the success of such an experiment.

So far, the existence of $0\nu\beta\beta$ -decays has not been proven experimentally. The Heidelberg-Moscow Experiment has resulted in an upper limit of 1.9×10^{25} on the $0\nu\beta\beta$ -decay half-life of ^{76}Ge [KK⁺01]. A new generation of double-beta decay experiments is currently in the planning or commissioning phase.

1.3 The COBRA Experiment

One of these next-generation experiments is the COBRA experiment [Zub01, Zub10]. Currently in the prototyping stage, COBRA will search for $0\nu\beta\beta$ -decays using cadmium zinc telluride (CdZnTe) semiconductor detectors.

CdZnTe contains nine double-beta decay isotopes (see Table 1.1) - the detector therefore also acts as its own source mass. This approach has significant advantages: As semiconductor detectors are made from very pure materials in order to function, they are intrinsically radiopure. Since the source mass is identical to the detector, the detection efficiency is high, as the range of the β -particles is less than 2mm.

Compared to other semiconductor materials, CdZnTe has the advantage that radiation detectors made from this material can be operated at room temperature. Most other materials, like high-purity germanium, require cryogenic cooling.

Isotope	Decay mode	a [%]	$N_0/1\text{g Cd}_{0.9}\text{Zn}_{0.1}\text{Te}$	Q-value [keV]
^{106}Cd	$\beta^+\beta^+, \beta^+/EC, EC/EC$	1.25	$2.889 \cdot 10^{19}$	2771
^{108}Cd	EC/EC	0.89	$2.057 \cdot 10^{19}$	231
^{114}Cd	$\beta^-\beta^-$	28.73	$6.640 \cdot 10^{20}$	534
^{116}Cd	$\beta^-\beta^-$	7.49	$1.731 \cdot 10^{20}$	2809
^{64}Zn	$\beta^+/EC, EC/EC$	48.6	$1.248 \cdot 10^{20}$	1096
^{70}Zn	$\beta^-\beta^-$	0.6	$1.541 \cdot 10^{18}$	1001
^{120}Te	$\beta^+/EC, EC/EC$	0.096	$2.465 \cdot 10^{19}$	1722
^{128}Te	$\beta^-\beta^-$	31.69	$8.137 \cdot 10^{20}$	868
^{130}Te	$\beta^-\beta^-$	33.80	$8.680 \cdot 10^{20}$	2529

Table 1.1: Decay modes, Q-values, natural abundances and number of nuclei per detector mass of double beta isotopes contained in CdZnTe [Kie05]

1.3.1 Isotopes and decays of primary interest in this work

$0\nu\beta\beta$ -decay of ^{116}Cd

Of the nine $0\nu\beta\beta$ -isotopes contained in CdZnTe, three have high Q-values. ^{116}Cd is of particular interest, since its Q-Value of 2.8 MeV is beyond the limit of naturally occurring gamma radiation (2.6 MeV). The energy region around 2.8 MeV will therefore be the main region of interest within this work.

β -decay ^{113}Cd

There is one radioactive background source intrinsic to the material: CdZnTe contains ^{113}Cd , which has a natural abundance of 12.2 %. ^{113}Cd decays via a four-fold forbidden β^- decay with a Q-value of 320 keV and a half-life time of $8.2 \cdot 10^{15}$

1 Introduction

years [GJK⁺05]. Due to the low endpoint of the resulting β -spectrum, compared to the Q-values of the double- β isotopes, this background does not pose a problem. It has, in fact, shown itself to be very beneficial: Since the ^{113}Cd is evenly distributed throughout the detector volume, it provides a way to continuously measure and monitor the relative efficiency of the detectors (see Fig. 6.3).

1.3.2 Detector Technologies

On the other hand, two drawbacks of CdZnTe as a semiconductor material have to be addressed: Growing large detector-grade CdZnTe crystals is difficult - the largest detectors currently available measure $2 \times 2 \times 1.5$ mm, while most are even smaller. To reach the detector mass required for a $0\nu\beta\beta$ -decay search (on the order of 400 kg of enriched CdZnTe [Wil05]), the experiment is expected to consist of about 60000 detectors. Also, CdZnTe exhibits a very low hole mobility (Table 1.2). As the holes drift much slower through the detector crystal than the electrons, the hole signal is mostly lost during data acquisition [Luk95]. Using a planar detector design, the detector output would depend on the drift length, and therefore the interaction depth of the event, as much as on the actual energy deposition. Two popular detector technologies which work around this problem are coplanar grid (CPG) detectors and pixel detectors. COBRA is evaluating both technologies as candidates for the detector design in the final experiment.

Property	Cd _{0.9} Zn _{0.1} Te	CdTe	Ge	Si
Atomic numbers	48,30,52	48,52	32	14
Density ρ [g/cm ³]	5.78	5.85	5.33	2.33
Band gap E_g [eV]	1.57	1.5	0.67	1.12
Pair creation Energy E_{pair} [eV]	4.64	4.43	2.95	3.63
Resistivity ρ [Ωcm]	3×10^{10}	10^9	50	$< 10^4$
Electron mobility μ_e [cm ² /Vs]	1000	1100	3900	1400
Electron lifetime τ_e [s]	3×10^{-6}	3×10^{-6}	$> 10^{-3}$	$> 10^{-3}$
Hole mobility μ_h [cm ² /Vs]	50 - 80	100	1900	480
Hole lifetime τ_h [s]	10^{-6}	2×10^{-6}	10^{-3}	2×10^{-3}
$(\mu \cdot \tau)_e$ [cm ² /V]	$(3 - 5) \times 10^{-3}$	3.3×10^{-3}	> 1	> 1
$(\mu \cdot \tau)_h$ [cm ² /V]	5×10^{-5}	2×10^{-4}	> 1	≈ 1

Table 1.2: Properties of CdZnTe, CdTe, Ge and Si at room temperature [eM]. The holes have a much lower mobility than the electrons in CdZnTe and CdTe.

Coplanar-Grid Detectors

CPG detectors are an elegant solution to handle the asymmetric drift velocity of the charge carriers. The design, invented by Luke [Luk95], is similar to a planar detector, but uses an anode which is structured into two comb-shaped parts isolated from each other (see Fig. 1.4). One of the anodes, the collecting anode (CA),




Figure 1.4: A coplanar grid (CPG) detector. Note the two comb-shaped anode structures.

is kept at ground potential. The other anode, called the non-collecting anode (NCA) is kept at a fraction of the cathode bias voltage.

The resulting electric potential (Fig. 1.5) in the detector can be utilized to compensate the missing hole signal. As the electrons drift through the bulk of the detector, they induce equal charge on the anodes. Only when the charge cloud finally enters the thin region dominated by the potential difference between the two parts of the anode grid, the electrons drift between the anodes. The difference signal between the CA and the NCA is therefore proportional to the total charge and independent from the interaction depth. To correct for secondary effects such as trapping, a weighting factor has to be introduced to this difference.

Only the two anodes have to be instrumented to read out a CPG detector, so the electronics can be separated from the detectors and can be placed outside of the radiation shielding. Thus, this detector type is well suited for low-background applications. On the other hand, the detector provides only limited information beyond the energy deposition (see 4.3).




Figure 1.5: Charge induction and potentials for a CPG detector. *Left:* Charge induced on the collecting anode (q_A), the non collecting anode (q_b) and the anode charge difference ($q_A - q_B$) as a function of the charge drift distance. *Right:* Weighting potential of a grid electrode [Luk95]

1.3.3 Pixel Detectors

Pixelized detectors have an anode which is structured as an array of mutually isolated pixels. They avoid the problem of the missing hole signal by design - due to the the small-pixel effect [BEB95], the charge deposited on the detector pixels is practically independent of the interaction depth.

Pixel detectors record the projection of the particle track in the detector volume, acting like a semiconductor time-projection chamber. The track information can make it possible to identify the particles involved. This provides a unique and powerful way to discriminate between signal and background events which is not available with other semiconductor detector designs. During the course of this work, the first COBRA physics data was taken using a pixel detector system.

The Cobra Prototype Setup at LNGS

As stated above, a low background experiment like COBRA requires an underground location. The collaboration operates a prototype setup at Laboratori Nazionali del Gran Sasso (LNGS) in Italy, located 1400 m under the Gran-Sasso massif. The setup (described in detail in [Mü07]) is shielded against environmental radiation, mechanical influences and electromagnetic interference.

Designed to house up to 64 CPG detectors of a volume of 1 cm^3 each, the setup has been used to evaluate different detector variants over the past years. Currently, 8 low-background optimised CPG detectors are installed.

A pixel detector setup was added to the CPG setup during the course of this work, to explore the potential of this new technology for COBRA under low-background conditions.

2 Background Improvements

2.1 Dominating Background Sources

Even though a lot of care had been taken in the design of the COBRA setup at LNGS, including a high-quality shielding and using only radiopure materials in the inner parts of the setup, the background level achieved was still orders of magnitude above the level necessary for a successful $0\nu\beta\beta$ -decay search [Mü07]. To get one step closer to the levels required in the final experiment, it was necessary to identify and eliminate the currently dominating background sources.

There was strong evidence [Ree09] that radioactive contaminations in the protective red paint on the detectors and the decay products of the ^{222}Rn in the air around the detectors constituted two main sources of the remaining radioactive background (the intrinsic ^{113}Cd beta spectrum is regarded as a signal, not as background here).

Flushing the setup with filtered nitrogen to remove air-borne contaminations from the immediate surrounding of the detectors, as well as using detectors coated with a new, clear paint of superior radiopurity (see Table 2.1) had been tried in the past. While both improvements resulted in a reduction of the radioactive background [Mü07, pp 59,65], they were not tried in combination and neither did prove decisive on its own. Therefore, it was also uncertain whether the paint and ^{222}Rn were the only major remaining contributions to the current background level.

Another major background-related problem was the contacting of the detectors. Contacting the detectors with industry-standard H20E conductive epoxy resin had proven to be fairly reliable, but was discontinued due to background and maintenance considerations [Mü07, pp 59,65]. One concern was the fact that H20E contains silver, which is a potential background source due to the metastable isotope $^{110\text{m}}\text{Ag}$, which has a half-life of 250 days. $^{110\text{m}}\text{Ag}$ is produced by neutron-activation of ^{109}Ag . The other concern was that H20E forms a bond that can only be removed destructively, necessitating at least a complete rework of the detector crystal, risking permanent damage in case of a faulty connection or when switching to a different contacting method. As a solution, a single-component copper-filled conductive glue had been developed, which was both radiopure and dissolvable. Unfortunately however, the connection provided by this glue was found to be unreliable, with a significant number of connections

2 Background Improvements

failing shortly after assembly, and more failing during long-term operation. Of the 16 detectors installed at LNGS at that time, usually only 12 or less were usable, even though some glue contacts were repaired during operation.

Lastly, a significant part of the background did not stem from radioactive contamination, but was due to signals generated by micro-discharges in the Kapton high voltage (HV) bias cables leading to the detectors. While the problem could be mitigated somewhat by adding a shield to the cables [Mü07], a significant number of these events remained [Ree09].

It was therefore important to find a new detector contacting system, and - using a different cable type for the HV-bias - explore the background reduction potential of detectors with the new, clear coating in a radon-free nitrogen atmosphere.

2.2 New Detector Contacting System

The new detector contacting system was required to be:

- reliable
- reworkable and removable
- radiopure
- robust under transport or easy to assemble on site
- suitable for large-scale manufacturing

Finding such a system proved to be a very challenging task. In collaboration with T. Köttig and S. Rajek several possibilities were explored, only one of which was found acceptable.

2.2.1 Pin Pressure Contacts

Due to the negative experiences with glue-based contacts, we evaluated a purely mechanical system. In laboratory measurements, the detectors had always been routinely and successfully contacted using probe needles. While the same was not possible in a low-background setup due to the bulk of the probes and the materials involved, the basic concept of a pressure contact between pins or needles and the detector electrodes seemed to be a good solution. Since there is no permanent contact, the detectors are not contaminated, can easily be installed on site and can just as easily be removed or replaced.

A challenge was the positioning of the contact pins for the detector anodes. The anode contact pads are determined by small paint-free areas on the gold-plated anode grids - since the detectors are painted by hand at the manufacturer, the position of these areas varies significantly between each detector. The size of the

pads required a positioning accuracy of $100\ \mu\text{m}$. To further complicate matters, the size of the detectors themselves can vary on the scale of about $500\ \mu\text{m}$.

During the course of this work, a modification of the Delrin original detector holder [Mü07] was developed. It incorporates a layer of springs, cut into the holder material itself, holding small gold-plated brass pins. The detectors are individually positioned under the springs using an alignment layer, also made of Delrin, containing holes to fix each detector in position. A photographic survey of all detectors yielded the information necessary to determine the correct size and position for each hole (see Fig. 2.1). All Delrin machining was done by laser cutting, ensuring a very clean and precise manufacturing process.

While more challenging in design than glue-based contacting systems, this system was easy to manufacture and to use, and was suitable for large-scale automatic manufacturing.

The system performed well during laboratory evaluation. However, a large-scale test with 64 detectors in four layers at LNGS in the fall of 2007 revealed serious flaws. The springs proved to be too stiff and the necessary uniform tension over the whole detector layer could not be achieved under field conditions, resulting in lack of contact on some electrodes and damage due to high pressure on others. Ironically, later analysis of the resulting scratch marks on the anode pads showed that the system had consistently achieved or exceeded the required positioning accuracy of $100\ \mu\text{m}$.

Thus, we finally had to abandon this solution. For details concerning design and performance see [Kö08]. However, the approach to use laser cutting to machine Delrin is now standard for nearly all Delrin components in the experiment.

Recently, the COBRA group at TU Dresden was very successful in contacting detectors using custom-ordered copper beryllium springs. While these springs are currently quite expensive and their radiopurity is unknown, a combination with the alignment method described above might be an interesting option in the




Figure 2.1: Detector position alignment by CAD.

2 Background Improvements

future.

2.2.2 Gallium soldering

CdZnTe detector crystals are very sensitive to heat and should not be exposed to temperatures above 80 °C. This makes it impossible to contact the detectors using standard soldering techniques - soldering at near room temperature, however, seemed a viable option. After careful evaluation of the properties of different solder materials, it was decided to test gallium as a solder material. Over the course of several months, we tried to develop a method for soldering the detectors with gallium and gallium alloys under a nitrogen atmosphere.

While several working contacts were made, we were not able to achieve reliable contacting despite substantial efforts. For details see [Kö08].

2.2.3 Silver-filled Conductive Lacquer

After the failure of both the mechanical and solder contacting scheme, glue was again considered as an option. The COBRA group at Freiburger Materialforschungszentrum (FMF) Freiburg had reported good experience using LS200 conductive lacquer. While LS200 contains silver, it has a much lower viscosity than the H20E glue described before. This makes it possible to use a very small amount for each contact, reducing the mass of silver placed near the detector and alleviating the concerns about $^{110\text{m}}\text{Ag}$.

In the beginning, the short pottime of LS200 made it difficult to apply with high precision. Relentless efforts of S. Rajek finally led to the development of a reliable contacting system based on this lacquer [Kö08, Raj].

The LS200-based contacting system is the current standard in the collaboration. It presents a compromise between the required properties:




Figure 2.2: A detector holder with spring/pin contacting.

The system is reliable - eight detectors with these contacts have now continuously been operated at LNGS for 30 months. All detectors are still working. It is also reworkable, since the contacts can be dissolved with acetone at any time. Contacted detectors can be transported if great care is taken, and assembly on site is possible. The solution is not, in its current form, ready to scale up to a large-scale manufacturing process. This might be overcome in the future though, since it is mostly due to the difficult handling of the LS200. Concerning radiopurity, the small amounts of silver involved do not yet present a problem at current background levels [Kö08]. Since the basis of the lacquer is also available separately, there is also the option of replacing the silver flakes within the glue with gold later on. This option has not yet been evaluated due to the costs involved.

2.3 Nitrogen Flushing System

The radon background contribution can be eliminated by displacing the air in the inner part of the experimental setup by clean nitrogen. Vaporised liquid nitrogen is already considered to be quite clean, and can be cleaned further using a radon trap filter installed within the Dewar containing the liquid nitrogen [Mü07]. Such a flushing system had been designed and built for the collaboration at TU Dortmund and was transported to LNGS in 2007. Necessary changes to make the fixtures on and inside the Dewar compatible with the filling system on site delayed its commissioning, though.

Another problem that precluded the system from regular use was the heater included to control the vaporization rate of the nitrogen. The heater, attached to the radon trap, did not possess any kind of temperature monitoring or over-temperature fail-safe. Since COBRA does not have personnel constantly on site, regular filling of the Dewar is not guaranteed. As a result, the heater exploded




Figure 2.3: LS200 conductive glue detector contacts.

2 Background Improvements

shortly after installation, as the Dewar ran empty - and so did its replacement later on. This made it necessary to clean the inside of the whole Dewar - a very difficult procedure - to prevent future contamination of the nitrogen and therefore the experiment.

To enable regular and stable operation of the nitrogen flushing system, the radon flushing system had to be modified in the course of this work, in collaboration with M. Alex.

2.3.1 Radon Trap Modifications

Improved Heater

The heater was replaced with a $5\ \Omega$ high-power resistor, in direct contact with a PT1000 temperature sensor and a thermal fuse with a trip temperature of $72\ ^\circ\text{C}$. When the Dewar runs empty, the temperature of the heater, no longer immersed in liquid nitrogen, rises. The sensor makes it possible to detect this state and switch off the heater. As a fail-safe, the thermal fuse will permanently disable the heater before its temperature reaches critical levels should the heater not be switched off in time due to sensor or software failure. While this will necessitate the replacement of the heater, it does prevent a contamination of the Dewar from the residue of an exploded heater. Additionally, a redundant second heater was installed for such an eventuality.

Fill Level Monitoring

To collect as much data under radon-free conditions as possible, regular filling of the Dewar is required. Due to the friendly support of the LNGS staff and scientists working on other experiments at LNGS this is often possible even when no COBRA collaborators are on site. It is only possible though, when the fill level of the Dewar can be monitored remotely, so filling can be requested in time but only when necessary. We found no commercial system compatible with the mechanical constraints imposed by the radon trap, and ordering a custom-made system was not financially viable. Since we wanted to bring the flushing system into operation as quickly as possible, there was no time for detailed examination of different fill level monitoring systems under field conditions, so we decided to incorporate two different solutions, providing redundant information.

The first sensing technique is capacitive: The sensor is a vertical cylindrical capacitor which was affixed to the radon trap. The capacitor consists of a steel pipe with an inner diameter R_2 of 6 mm, containing a steel rod with an outer diameter R_1 of 4 mm. As liquid nitrogen N_2 has an ϵ_r of 1.4 [Hay10], while gaseous nitrogen has an ϵ_r of 1.0, the capacitance of the sensor is proportional to the liquid level h :

$$C = const + \frac{2\pi\epsilon_0}{\ln\frac{R_2}{R_1}}\epsilon_{r_{liquid}} h \quad (2.1)$$

The concept is based on a system proposed by Landis et al. [LMG86], but is implemented here as a coaxial, instead of a triaxial design.

The second technique uses the vast difference in thermal conductivity between liquid and gaseous nitrogen: A serial chain of two PT100 resistive temperature sensors is fixed to the radon trap. The first sensor is located near the top of the capacitive sensor, at a level reached by the liquid when the Dewar is full, the second one at the bottom of the capacitive sensor. The chain, having a total resistance of 35Ω at the temperature of liquid nitrogen, is driven with 5 V over a 10Ω pre-resistor, resulting in a power dissipation of about 0.3 W at each resistor. As long as the resistors are immersed in the liquid, their temperature remains virtually unchanged. A PT100 resistor in the gaseous phase, however, heats up sufficiently to detect the temperature difference, resulting in a measurable voltage change between the pre-resistor and the PT100 chain Fig. 2.5. This way, three fill states (full, normal and near-empty) can be detected. This can also be used to calibrate the capacitive sensor.

The capacitive sensor proved itself to be a reliable, linear level monitoring solution during operations (Fig. 2.8). The resistive sensor, while functioning well, provides only three discrete fill states and is therefore only used as a backup system.

The sensor readout electronics, together with a MOSFET circuit to switch the heater on and off and a microcontroller board for remote control, were installed in a small box on top of the Dewar. The only external connections are a power connector for the heater voltage and an Ethernet connector providing data connectivity and power for the electronics via Power-over-Ethernet (POE). If empty, the Dewar is unplugged and refilled at the nearby filling station. The monitoring system then automatically resumes operation after reconnection.

2.4 Alternative Radiopure Detector Coatings

As stated above, the colourless detector coating by eV-Microelectronics has a much higher radiopurity than the red coating used earlier. Four detectors with this coating had been purchased for measurements under low-background conditions. However, the collaboration already possessed over 70 red-coated ones. After the red coating had been shown to be a dominating background source, we therefore looked for a way to refit these detectors with an alternate, radiopure coating [Sch09]. K. Scheiner and S. Rajek developed a technique to remove the red coating using an acetone bath, without degradation of detector performance. Since the colourless coating used by eV-Microelectronics is proprietary and not

2 Background Improvements




Figure 2.4: The reworked radon trap, with improved and redundant heaters and fill level monitoring system.




Figure 2.5: Schematic of the Dewar with radon trap and monitoring sensors [Ale09]

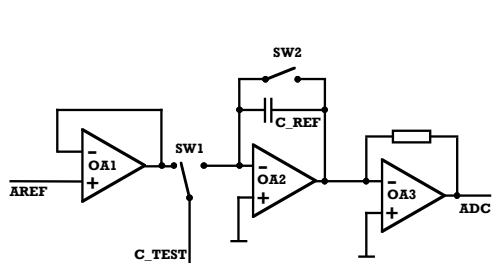


Figure 2.6: The switched-capacitor capacitance sensing circuit [Ale09]




Figure 2.7: The PT sensors are instrumented as a simple voltage divider.

2 Background Improvements




Figure 2.8: Continuous capacitive fill level monitoring of the liquid nitrogen Dewar at LNGS. Several fill cycles are shown - the capacitance provides a good linear measure for the fill state.




Figure 2.9: Waterfall plot of background events at LNGS before and after activation of nitrogen flushing. The radon background is significantly reduced.

Chain	Decay	CdZnTe	Red paint	Clear Paint	Cyclotene
^{232}Th	^{228}Ra	< 12	1100 ± 100	< 180	< 1200
	^{228}Th	< 9	730 ± 70	< 190	< 370
^{238}U	^{226}Ra	< 10	2100 ± 100	< 140	< 1000
	^{234}Th	< 250	1100 ± 300	< 1500	< 5200
	$^{234\text{m}}\text{Pa}$	< 220	1600 ± 100	< 4300	
^{235}U	^{235}U	< 7	170 ± 30	< 15	
	^{40}K	< 91	6900 ± 800	< 15	7700 ± 1500
	^{60}Co	< 5	< 20	< 72	550 ± 120
	^{137}Cs	< 4	< 15	< 56	1004 ± 180

Table 2.1: Activity of detector coatings in mBq/kg [Sch09, p39]

for sale, we looked for an alternate material providing good protection of the detector crystals against moisture and mechanical damage. Following a recommendation by A. Fauler from the COBRA group at FMF Freiburg, Cyclotene 4024-35 was chosen and tested for radiopurity (see Table 2.1, Courtesy of K. Schreiner). Due to the encouraging results, four detectors were recoated with Cyclotene and installed at LNGS in December 2008 [Sch09].

Operating detectors without any coating, entirely removing a potential background source, was deferred to later stages of the experiment since experience showed the detectors to be very susceptible to mechanical damage during handling and transport. The collaboration currently plans to install facilities at LNGS suitable for assembly of future detector layers on site. The author recommends to re-evaluate the option of installing uncoated detectors once on-site assembly has become standard procedure.



Figure 2.10: The liquid nitrogen Dewar with the monitoring and control electronics box.

2.5 Background Reduction

Two test layers, with four CPG detectors each, were run long-term at LNGS under low-background conditions. One layer contains the four detectors with the colourless coating by eV-Microelectronics, the other layer the four detectors coated with Cyclotene. All detectors are contacted using the new contacting system and are still in operation at this time.

The combination of the new detector coatings and the nitrogen flushing system results in a reduction of the radioactive background of over an order of magnitude (see Fig. 2.12) compared to the red-coated detector run in normal air. The detector now achieve a background level of $6 \text{ cts keV}^{-1} \text{ kg}^{-1} \text{ year}^{-1}$ in the region of interest for the $0\nu\beta\beta$ -decay of ^{116}Cd . This is an important achievement, advancing the experiment to a new level of radiopurity.

The new contacting system proved to be reliable. Since no $^{110\text{m}}\text{Ag}$ contamination could be identified in the spectrum [Sch09], it will be used for the next stages of the experiment.

No significant difference could be found in background levels of the two layers [Sch09]. Therefore, recoating the available red-coated detectors with Cyclotene is a viable option to provide the number of radiopure detectors required for the installation of the planned 64-detector array. Using the unmodified red-coated detectors as intended originally, no longer seems expedient considering the results presented here.




Figure 2.11: The two test layers with low-background crystal coatings, eight detectors in total.




Figure 2.12: Background reduction achieved at LNGS by combination of nitrogen flushing of the setup and change of detector passivation coating.

2 Background Improvements

3 Polaris CdZnTe Pixel Detector System

As described in 1.3.3 pixel detectors are a very interesting technology option for the COBRA experiment due to their unique tracking and veto capabilities. These advantages come at the cost of higher complexity compared to CPG detectors, resulting in additional design challenges.

3.1 Design Considerations for Pixel Detectors

The readout electronics of a pixel detector system cannot be separated from the detector crystal due to the high number of signal channels (one per pixel). Instead, the necessary amplifiers and analog to digital converters (ADCs) reside on a so-called frontend chip, an application-specific integrated circuit (ASIC) which is usually permanently bonded to the crystal.

The physical layout of the ASIC determines the size of the pixels, called the pixel pitch. This is crucial, as there is a trade-off between energy resolution and track resolution: Smaller pixels result in tracks with a higher spatial resolution, but have a worse energy resolution than bigger pixels [GAD⁺11]. In addition to that, the thickness of the detector places an upper limit on the track resolution which is due to the increased diffusion of the charge clouds in the detector volume at higher drift lengths. This results in an additional lower limit for the usable pixel pitch. The COBRA group at the University of Erlangen is currently conducting extensive simulations to determine the optimum combination of pixel pitch and detector thickness.

Due to the close proximity to the detector crystal, both the ASIC and the bonding process must be very radiopure and there is little research on the radiopurity of such components at the level an experiment like COBRA requires. The collaboration has started to measure batches of different products at Laboratoire Souterrain de Modane (LSM). The first preliminary results are encouraging, but it is not yet certain whether complete pixel detector systems can be built at the required radiopurity levels.

Numerous other factors, such as power consumption and the resulting cooling requirements or trigger and timing capabilities also have to be considered when selecting a frontend chip. Eventually, a custom ASIC, or a custom modification of an existing ASIC, will have to be designed to fit the requirements of the experiment. This is, in fact, quite usual in pixel detector applications. However, ASIC




Figure 3.1: Pixel structure of the Polaris Detector. Note the steering grid structure between the pixels.

development is costly: A design modification will often result in production costs of over a million EUR and require months or even over a year of development time [Cam11]. The collaboration is therefore currently evaluating several existing frontend designs, each of which partially matches the given requirements, to gain the data and experience necessary for a well founded design decision.

3.2 The Polaris System

During the course of this work, a large-volume CdZnTe pixel detector was installed at LNGS in September 2009 in cooperation with Zhong He and Feng Zhang (University of Michigan). The detector, measuring $22 \times 22 \times 15 \text{ mm}^3$ was successfully run underground from September 2009 to January 2010, collecting 4.32 kg days of data.

Due to the large pixel pitch of 1.8 mm, this detector type cannot resolve track details of double-beta decay events. Instead, the focus with this system was on detector volume and energy resolution.

The detector used was a prototype of the Polaris system in development at Prof. He's group at Michigan. Polaris is intended to be a portable, hand-held 3D position sensitive radiation sensor - the final system will contain an array of $3 \times 3 \times 3$ large-volume pixelated CdZnTe detectors. The prototype we used at LNGS had one detector installed. While the system is not optimized in any way for low-background applications, its technical maturity made it an ideal choice for a first test with a large-volume pixelated CdZnTe detector under low-background conditions.

3.3 System Properties

To achieve a high energy resolution, the anode of the detector crystal is segmented into 11×11 pixels, embedded in a steering grid (see Fig. 3.1). The grid, held at a negative bias, is used to focus the drifting charges onto the pixels, preventing loss of charge in the space between the pixels. The best energy resolution achieved with this type of system so far is 0.5% at 662 keV [WHWZ08], the system we used had an energy resolution of 2% [Zha09].

The Polaris prototype [ZHS07, ZH06] consists of a motherboard (type Moca-9) able to support up to nine detector modules (Fig. 3.2), one of which was installed. The detector box, including mainboard and detector module, is connected to the data acquisition (DAQ) PC via a digital readout board (type CROB-16), which has a parallel TTL interface on the PC side. The detector modules consist of the CdZnTe crystal and a custom VAS_UM2/TAT4 readout ASIC built by Gamma-Medica Inc. The ASIC has 129 channels to read out the anode pixels, the cathode signal, and the anode grid. One channel is dedicated to read out an on-chip temperature sensor, providing temperature monitoring directly at the detector module.

In addition to spatial resolution in the pixel plane, the ASIC is also capable of measuring the interaction depth for each individual pixel, providing fully 3-dimensional hit information [ZHS07]. The depth information is calculated from the ratio between cathode and anode signal and used to virtually slice the crystal into 40 vertical layers, which are calibrated individually. A full calibration of the detector requires the use of several different radioactive sources and takes at least a week under proper laboratory conditions. Calibration of the detector on site at




Figure 3.2: Polaris pixel detector module, consisting of the CdZnTe crystal and the front-end chip. For enhanced modularity the frontend chip is not directly bonded to the crystal. Instead, both are located on small carrier boards with a high-density connector in between.

LNGS was therefore not possible. Instead, the system was calibrated at Michigan and the calibration then verified at LNGS.

3.4 Underground Installation

Because of spatial constraints within the EMI-shielding (Faraday-cage) on site, the Polaris pixel detector was set up on top of the existing CPG setup (see Fig. 3.3). The detector box was shielded with 10 cm of lead on the sides and on top while the massive CPG lead-castle provided even better shielding on the bottom. Since the system was itself not low-background optimized it was decided not to invest in a high quality shielding with a core of copper and clean lead, so standard lead was used.

The calibration and energy resolution of the pixel detector system is temperature sensitive - as the system had been pre-calibrated at 22 °C, it was important to operate it under similar conditions. Due to the heat produced by the ASIC the temperature reached 27 °C during first tests, resulting in an unsatisfactory energy resolution. After thermally coupling the heat sink directly to the lead castle using a specially machined small block of copper (see Fig. 3.4), we were able to achieve a stable operating temperature of 23 °C.

Since the system came to LNGS pre-calibrated and is designed for calibration stability [ZH06], calibration at LNGS could be limited to a verification with a ^{228}Th source. As the source usually used to calibrate the COBRA CPG-setup is not active enough to generate sufficient statistics in all pixels, we had to use an alternate 100 kBq source instead, kindly lent to us by the XENON collaboration. Unfortunately, the source was not of a wire-type design and so could not be inserted using the calibration tubes of the setup. We therefore had to complete the calibration and remove the source before closing the lead shield around the



Figure 3.3: The polaris pixel detector setup at LNGS during construction.




Figure 3.4: The copper block used to thermally couple the detector to the lead castle for cooling.

detector.

3.5 Results

The detector crystal was coated with a red passivation paint very similar to the one found on the old COBRA CPG detectors. This made it a likely candidate for a significant background source, judging from the experiences made with the red-painted CPG crystals [Mü07, Ree09]. To investigate this, it is useful to compare the background spectra of single-pixel events for boundary and non-boundary (inner) detector pixels, since neither the top nor the bottom of the detector are painted. The spectra (see Fig. 3.5 and 3.6) show a difference by several orders of magnitude, especially at higher energies. While the background level of the boundary pixels stays almost constant, there are no background events on the inner pixels above 2.6 MeV. This is consistent with an alpha- and beta-type radioactive surface contamination, matching previous studies of the CPG red paint [Ree09].

A fiducial cut was therefore applied to the data, vetoing all events which affected pixels at the detector boundary.

A vertical fiducial cut was also evaluated, to veto events near the top or bottom surface of the detector. However, there was no significant further improvement. With a system designed for radiopurity, such a cut - sensitive to α and β radiation sources - should have proven effective. The background of the Polaris prototype




Figure 3.5: LNGS Background Spectrum, single-pixel hits. While there is a high background level on the boundary pixels over the full energy range (shown in red), the background on the inner pixels (shown in blue) is lower by several orders of magnitude and does not reach beyond 2.7 MeV.

3 Polaris CdZnTe Pixel Detector System




Figure 3.6: Distribution of hits in the pixel plane for the LNGS background spectrum. The hit count is much higher for the boundary pixels.

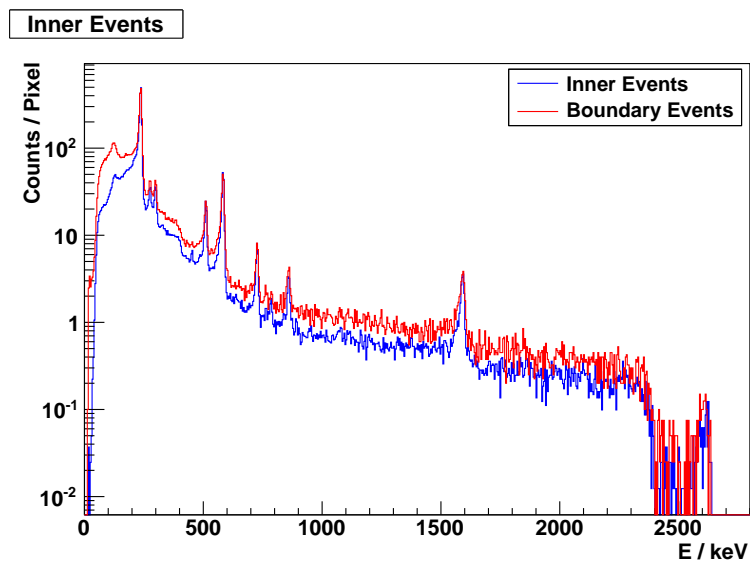


Figure 3.7: LNGS ^{228}Th calibration spectrum. The boundary pixels show a slightly higher efficiency in the full energy peaks and a constantly higher efficiency in the Compton regions. This is caused by the fact that the detector crystal is slightly bigger than the pixel array, adding extra active volume to the boundary pixels.

was therefore dominated by other sources (also compare Fig. 2.12 and 6.10).

To determine the efficiency of the detector after the cuts, ^{116}Cd $0\nu\beta\beta$ -decays were simulated in an 11×11 pixel detector of the same dimensions as the Polaris crystal by T. Köttig, using the COBRA monte-carlo package (Venom). The results (see Fig. 3.9 and Fig. 3.8) show that 72% of all double beta events and 67% of all double beta events with full energy deposition affect less than tree pixels. This selection corresponds to the short track length of $0\nu\beta\beta$ -events in CdZnTe. Thus, as a secondary cut criterion, event selection was restricted to events affecting only one or two directly connected pixels. The absolute efficiency of the detector for full energy deposition of the $0\nu\beta\beta$ -decay of ^{116}Cd after those cuts is 29%.

While the cuts reduce the detector efficiency, the signal to background ratio is significantly increased (Fig. 3.10). The number of events in the region of interest is reduced to zero counts in the interval of 2 FWHMs around 2.8 MeV. The method of Feldman and Cousins [Fel98] yields in an upper limit of 3.2 signal counts at a confidence level of 95 % under the conservative assumption of no background - equivalent to $4 \text{ cts keV}^{-1} \text{ kg}^{-1} \text{ year}^{-1}$.

This is an excellent result compared to the background level achieved with the CPG setup (2.5). While the Polaris system has not been optimized for low-background applications at all, it is able to perform as well or slightly better than the low-background optimized non-pixel detector system, even if only in the region of interest for the ^{116}Cd $0\nu\beta\beta$ -decay.

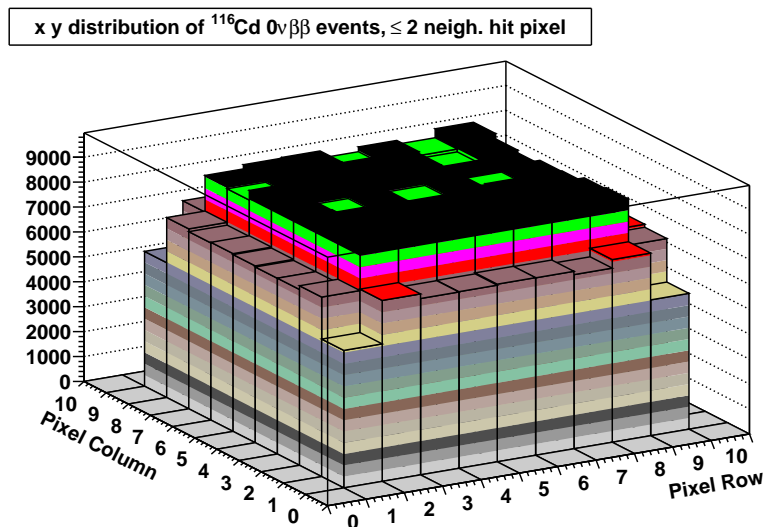


Figure 3.8: Simulation of ^{116}Cd $0\nu\beta\beta$ hit distribution under the cuts applied during the analysis.

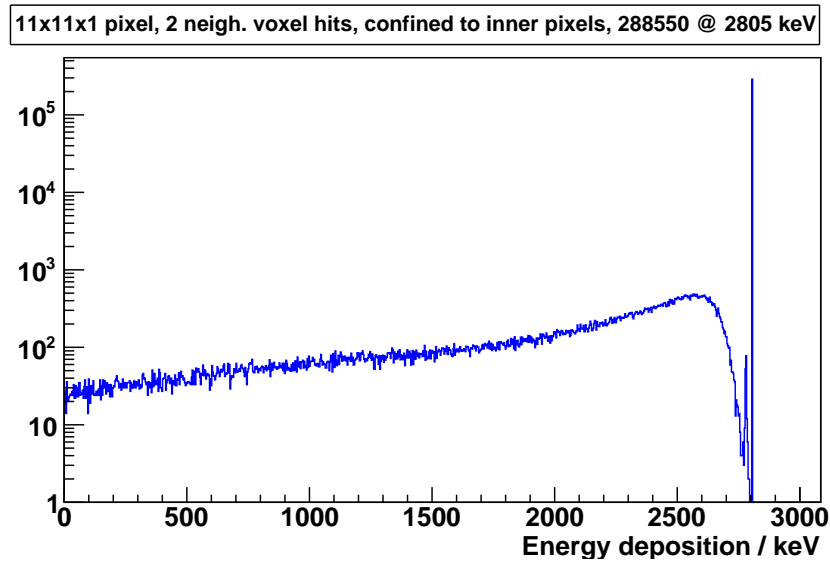


Figure 3.9: Simulation of the ^{116}Cd $0\nu\beta\beta$ spectrum under analysis cuts.

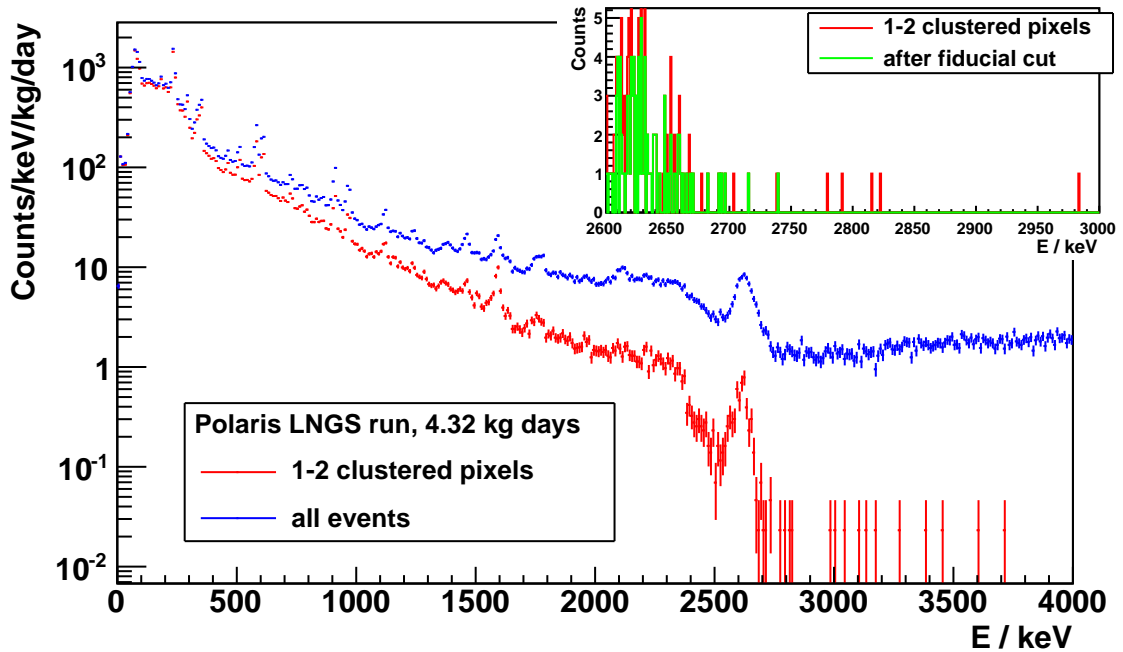


Figure 3.10: Background spectrum of ^{116}Cd before and after analysis cuts. Selecting only events which affect one pixel or two neighbouring pixel reduces the background dramatically. A fiducial cut, excluding all boundary pixels, eliminates all remaining events in the region of interest.

4 CPG Detectors with Pulse-Shape Readout

4.1 Previous DAQ System

The previously used COBRA data-acquisition chain (see Fig. 4.1) was a classic analog spectroscopy chain consisting of charge-sensitive amplifiers (CSAs) for charge integration, a shaping amplifier for noise suppression and for transforming step-like features in the CSA signal into well-formed pulses and - finally - a multichannel analyser for digitisation of these pulses. The weighted difference of the two anode charge signals, required for CPG detectors, was generated by an additional analog circuit between the CSAs and the shaping stage [Mü07, Kie05].

The weighting factor, as well as a final amplification applied within the multi channel analyzers (MCAs) could be configured digitally. The shaping time and amplification, however, had to be set by hand on site within the Nuclear Instrumentation Module (NIM) shaper modules, a very time-consuming process.

Operation of the setup required tuning of 3 main parameters: The cathode bias voltage, the grid bias voltage and the weighting factor. These parameters are interdependent, requiring a scan of a 3-dimensional parameter space to optimise the system's energy resolution, with a calibration spectrum to be taken for each point. The lengthy nature of this process severely limited the parameter combinations that could be explored, resulting in an operation under non-optimal conditions. Especially the weighting factor presented a problem, since its optimum range is very small.

While scalable in principle, we found the system to be difficult to commission and to operate with a larger number of detectors when we tried to install a test setup containing 64 detectors in 2007.

Also, the output of the system, consisting of a single energy deposition value per event and detector, had no degree of redundancy. This left no way to separate valid physic events from electronic disturbances, especially high-voltage micro-discharges, which had been indirectly observed in the old setup but could only be filtered out partially [Ree09].

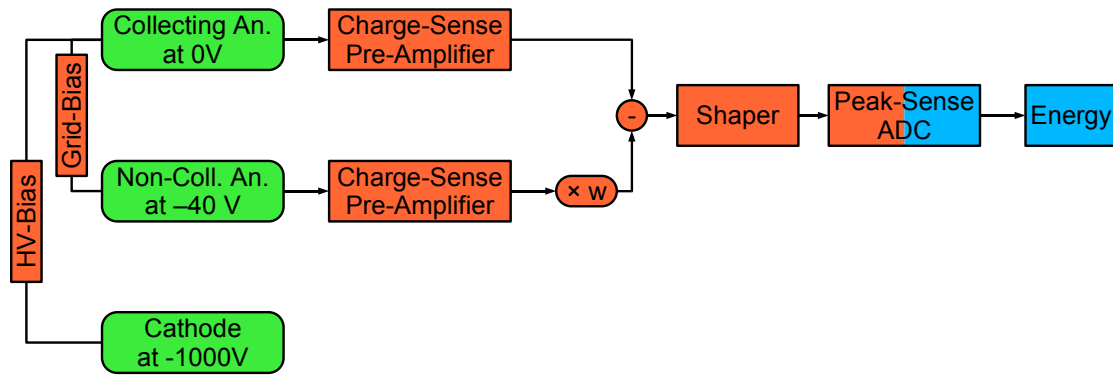


Figure 4.1: Previous COBRA CPG DAQ Chain. Most signal processing is done in the analogue domain, the only output value is a measure for the energy deposition.

4.2 COBRA DAQ Electronics Upgrade

4.2.1 New Data-Acquisition Chain

These weaknesses necessitated major changes to the DAQ Chain. The new system was required to be robust and easy to commission and operate to ensure both scalability and high-quality results. Originally, a system based on separate shaping and digitisation of the anode channels was planned during the course of this work and had shown promise in laboratory tests.

Recording the detailed shape of the preamplifier signal as a high-resolution time series, preserving the full information of the detector output, had always been the preferred option. The large number of fast analog-digital converters required to instrument the experiment had not allowed us to pursue this option in the past, due to the high costs involved. However, the recent decommissioning of the Amanda Experiment offered a sudden chance to acquire a sufficient number of such devices. Prof. K. Zuber, Prof. C. Gößling and the author successfully applied for the reassignment of a significant part of the Amanda ADCs [Wag04] to the COBRA experiment.

The availability of a recording solution for the anode signal transients led to the design of a completely new instrumentation of the experiment (see Fig. 4.2). In contrast to the old system, which did most of its operations in the analog domain, the new system is designed to preserve the detector signals as accurately as possible before digitisation. The preamplifier output is only amplified by a wideband fast linear amplifier to adapt it to the fast analog to digital converters (Fast-ADCs) input range. The signals of both anodes are recorded separately, all operations like generation of the weighted difference and pulse height determination are done digitally on the recorded data. One of the most important consequences is that the weighting factor can now be optimised offline, eliminating the most critical and troublesome parameter from the operation of the system.

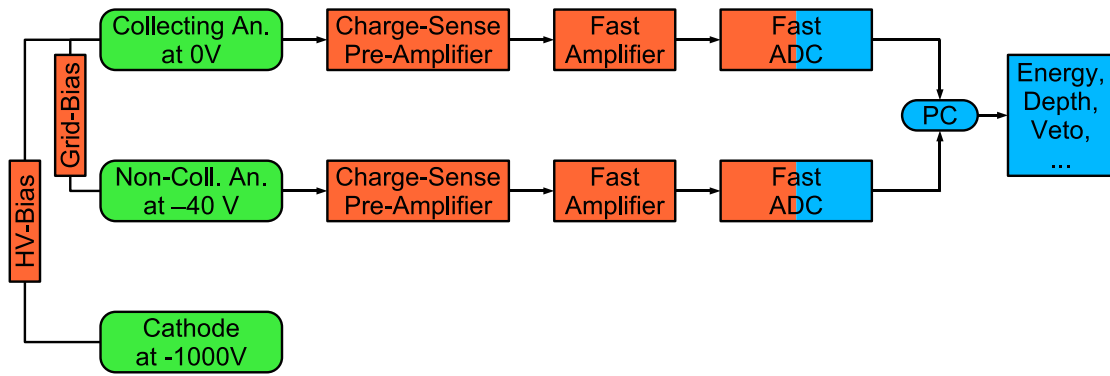


Figure 4.2: New COBRA DAQ Chain. The signal is digitised as soon as possible to conserve the full information contained in the event.

Analysis of the anode pulse shapes allows access to important physical quantities beside the energy deposition. This work will show the possibilities opened by estimating the vertical position of the event within the detectors. Other important characteristics, like measures separating single-site from multi-site events, can be determined from the pulse shapes [ACK⁺08], but require more detailed knowledge of the typical signal shapes for events with known properties. Such pulse shape libraries can be generated by charge-drift simulations, 180° Compton-scattering measurements and scans of the detectors with collimated radioactive sources. Efforts are currently under way within the collaboration to produce the required data, and initial simulations have shown promising results in single-site/multi-site differentiation [J⁺10, McG09].

Within the course of this work, prototypes were developed for all stages of the new DAQ chain, some of which are now already in use at LNGS.

4.2.2 Pre-Amplifiers

The actual charge-sensitive preamplifier circuits had in the past been procured in the form of commercially available modules and deployed in a custom preamplifier box with a module carrier board [Mü07]. Tests of several preamplifier modules had resulted in the selection of the Cremat CR-110 as the CSA module type of choice [Fas]. Since the carrier board also contained the analog subtraction circuit, it did not provide proper buffering of the HV-bias voltage and was not ideally suited for the CR-110. Thus, a new board had to be designed. A prototype was built that has four channels (see Fig. 4.3 and 4.4) and can be used to instrument not only the grid anodes, but also the cathode and guard ring contact of a single detector for research purposes.

Care was taken to find an optimal routing for the sensitive signals and to take the recommendations of the preamplifier manufacturer into account as far as possible under the design constraints. However, one important recommendation, to-

4 CPG Detectors with Pulse-Shape Readout

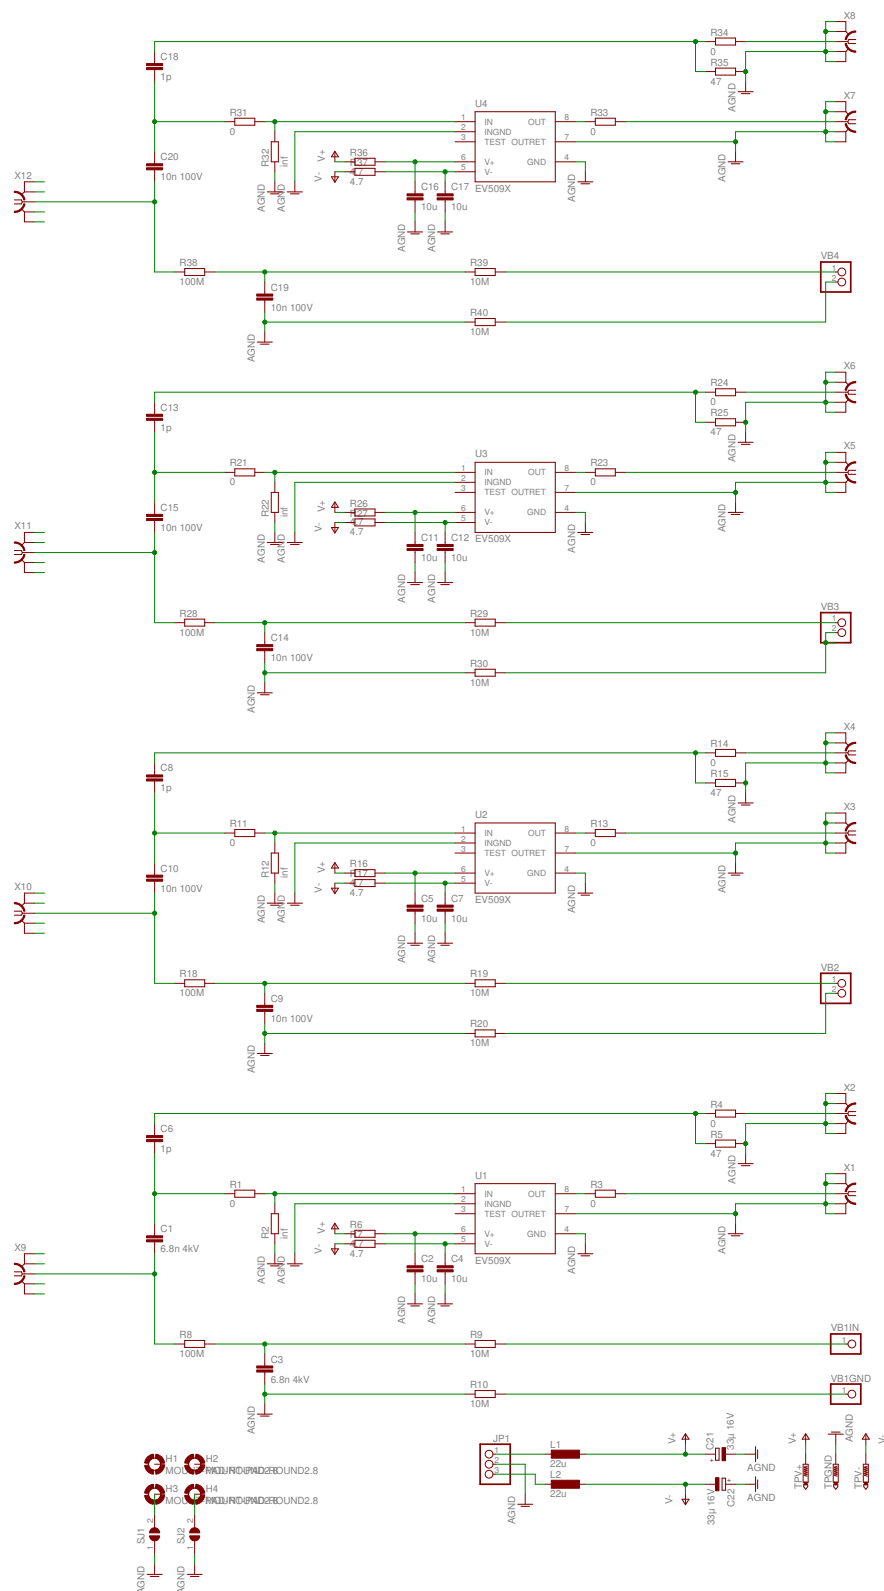


Figure 4.3: Schematic of the prototype preamplifier box.

tal physical separation of the individual amplifier modules, could not be implemented due to the high channel count the experiment requires. Crosstalk measurements currently under way [Teb11] show that this is a minor, but not critical, concern. The next version of the new preamplifier box will incorporate metallic shielding between the modules to eliminate this low-level crosstalk between the channels.

4.2.3 Fast Linear Amplifiers

Due to the high per-channel cost of commercial fast linear amplifiers providing the required bandwidth of 50 MHz and a maximum amplification of 40x to 100x, it was decided to custom build a COBRA fast linear amplifier, matched to our requirements. This offers the opportunity for optimal adaption to the ADC inputs. The amplifiers are built as NIM form factor modules, since 19" NIM crates are readily available in the collaboration.

As a first prototype, an 8-channel version with single-ended input and output was built to gain experience with pulse shape recording and processing both in the laboratory and under low-background conditions. The amplification factors for the channels are set via potentiometers on the front panel of the device. The amplifier performs well and is currently in use both at the LNGS underground setup, and at TU Dresden for detector characterisation measurements.

The preamplifier box has to be located close to the detectors, just outside the lead shielding, to keep the anode wire capacitance small [Mü07, Fas]. The fast amplifiers, on the other hand, are located near the ADCs in the DAQ rack, both for ease of access and to ensure an optimal ADC input signal. This makes the con-




Figure 4.4: The prototype preamplifier box.

4 CPG Detectors with Pulse-Shape Readout




Figure 4.5: Schematic of the single-ended fast linear amplifier (1 of 8 channels). It is built in two stages, both using a TI THS4211 high speed, low-distortion amplifier IC. The first stage provides a $4\times$ fixed amplification, the amplification of the second stage is freely adjustable. The amplifier is usually operated at a total amplification factor of 20 to 40. The two-stage design results in a bandwidth of 100 MHz even at $40\times$ amplification.




Figure 4.6: Custom built 8-channel fast amplifier for pre-amplifier signal conditioning.

nection between the amplifiers the longest in the system (over 3 m) and so makes it very susceptible to irradiated noise and ground loops. These had already been a problem with the previous DAQ chain, alleviated somewhat by the filtering effect of the shapers. The new requirement of producing clean, low-noise signals suitable for pulse shape analysis lead to the decision to switch to differential signalling on this part of the chain. Also, the potentiometer-defined amplification of the amplifier prototype offered an increased risk of calibration drift and could not be changed remotely.

Therefore, a second prototype was designed, consisting of a single-ended to differential converter (Fig. 4.8), to be included in the final version of the preamplifier, and an amplifier with differential inputs and single ended outputs for the ADC and digitally programmable amplification factors (Fig. 4.9). The voltage regulation and support circuitry is located on a separate module to allow for individual design changes (Fig. 4.7).

Standard Cat6 networking cable was chosen as the standard for the differential transmission line between the amplifiers as it is commercially available in almost every length at high quality and very low costs. The previous single-ended prototype had inherited the cabling of the old DAQ chain. This consisted of special and expensive connectors fitted to many thin coaxial cables which had to be assembled by hand. Also, in addition to the robustness of a differential signal, Cat6 cabling offers a substantially better shielding effectiveness of 60 dB to 70 dB [Tyc07], compared to coaxial cabling which provides about 40 dB [Flo08]. Another advantage is that due to the clear separation of signals and shield, this type of cabling is well suited for feedthrough into the electromagnetic shielding (Faraday cage) around the experiment without damaging the integrity of the shielding.

4.2.4 Fast Analog-Digital Converters

The Fast-ADC modules reassigned from the Amanda collaboration are Struck SIS3300 VME-bus modules [SIS]. Each module has eight input channels with a resolution of 12 bit at a maximum sampling frequency of 100 MHz.

For storage of the sample data, the SIS3300 has two separate memory banks, each with a capacity of 131072 samples per channel. The buffers can be subdivided - here, we acquire 1024 samples per event (spanning 10.24 μ s), resulting in a capacity of 128 events per memory bank. Once a memory bank is full, the module switches to the other bank while the DAQ system reads the acquired data. This multi-event, double-buffered operation allows for dead time free operation as long as the average data rate does not exceed the readout speed of the DAQ system.

4 CPG Detectors with Pulse-Shape Readout

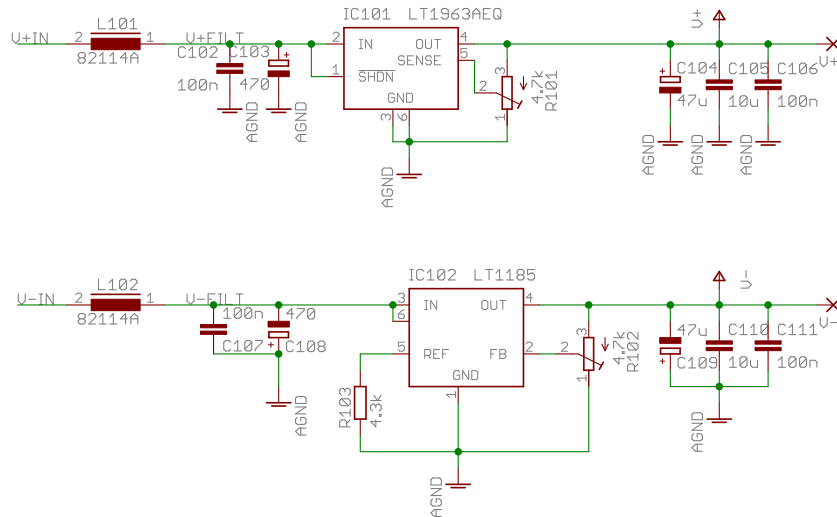


Figure 4.7: Schematic of the voltage-regulator module. Low-noise, low-dropout regulator ICs are used to realise a clean, stable power supply up to 1.5 A at an output voltage of 5 V from an input voltage of only 6 V. The output voltage is fully adjustable.

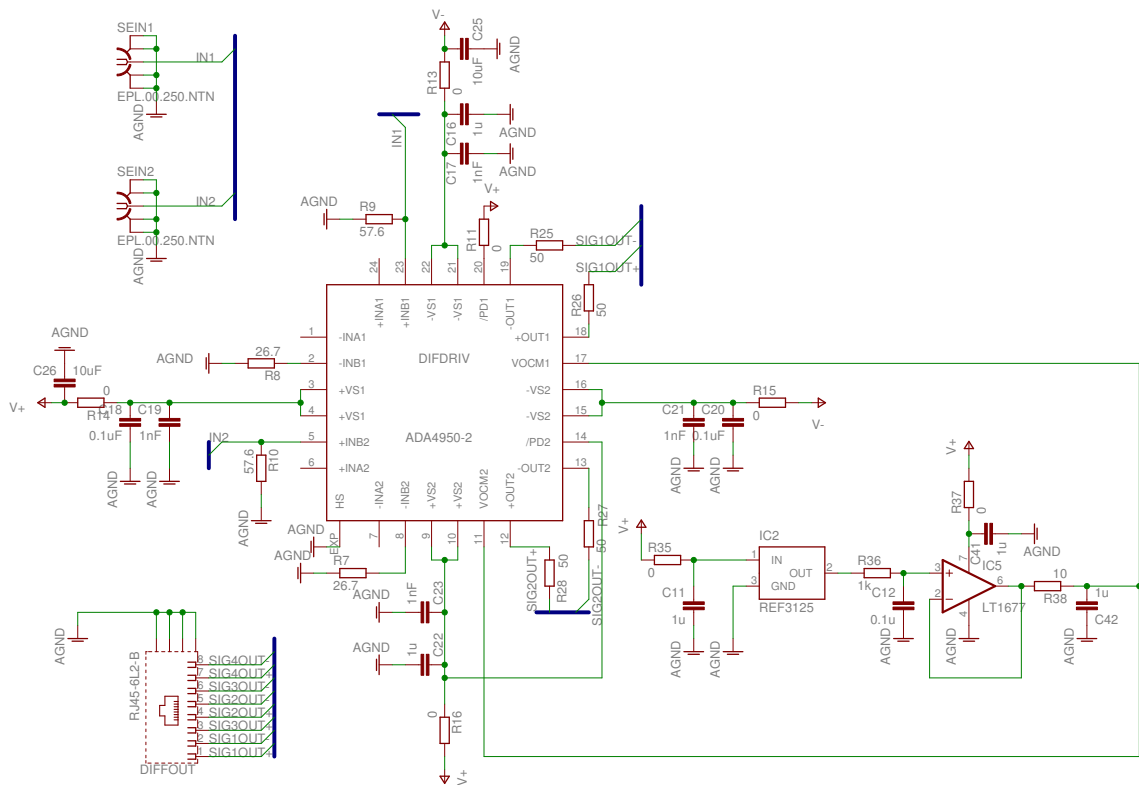


Figure 4.8: Schematic of the single-ended to differential circuit - a two-channel version is shown here. The circuit is based on the two-channel ADA4950-2 differential driver, which features a wide bandwidth of 750 MHz and an extremely low distortion of 108 dB at 10 MHz. A REF312 precision voltage reference generates the common-mode reference voltage of 2.5 V, which is buffered by an LT1677 amplifier for lower impedance.

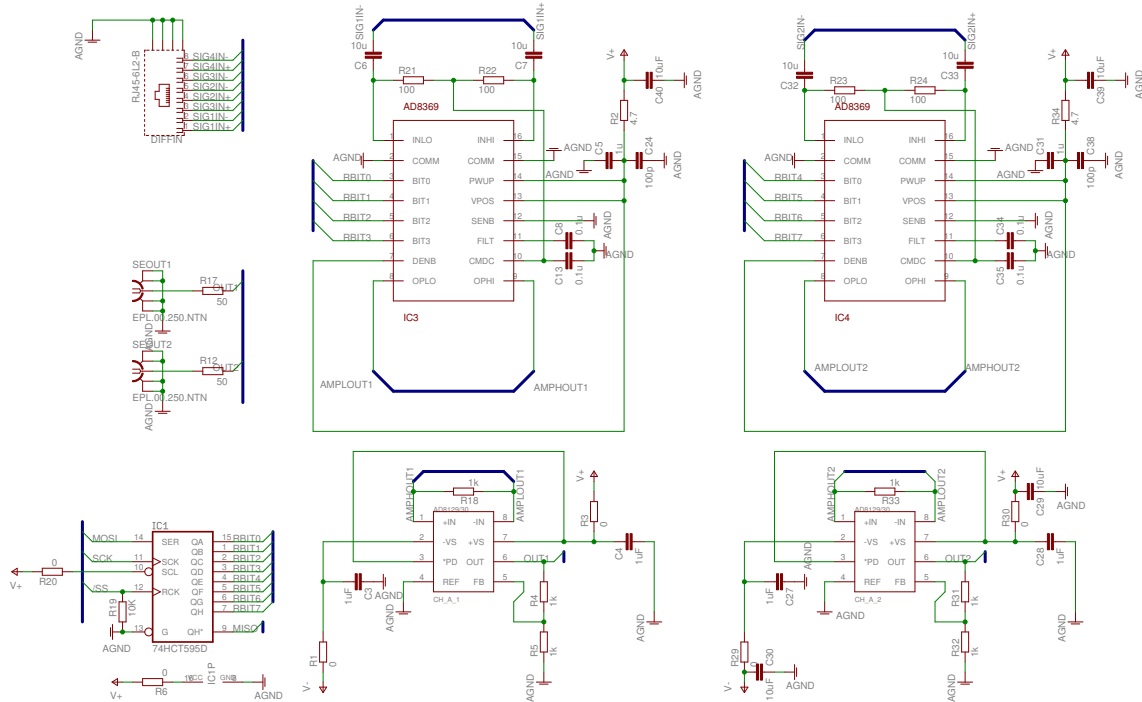


Figure 4.9: Schematic of the differential to single-ended conversion circuit with digitally configurable $0.5x$ to $100x$ amplification - a two-channel version is shown here. The two-stage design consists of a fully differential amplification stage, based on the AD8369 variable gain amplifier, followed by an AD8130 differential receiver which drives the single-ended output. The amplification factor of the two AD8369 ICs is set digitally via a serial shift register. The total bandwidth of the complete circuit is over 250 MHz and is independent of the amplification factor.




Figure 4.10: The differential amplifier prototype with two single-ended to differential and two differential to single-ended channels. For the prototype, both circuits are located on the same board to allow for maximum flexibility during testing. Usually, one board is deployed near the preamplifier prototype box as sender, while a second one in the DAQ rack functions as receiver and amplifier.

4 CPG Detectors with Pulse-Shape Readout




Figure 4.11: A Struck SIS3300 Fast-ADC VME Module

Anti-Aliasing

An important factor to consider when using time series samples is aliasing. When taking discrete samples with a sampling frequency f , signals of frequency ν and $\nu + nf$ become indistinguishable due to the Nyquist–Shannon sampling theorem [Wat01]. This is because the Fourier spectrum of a signal consisting of delta peaks with intervals of $1/f$ in time is periodic with a period f in frequency space. The result is that it is not possible to separate the signal in the frequency range of interest from other signals and noise contributions at higher frequencies.

In reality, the data samples are not acquired at points in time, but are averages over short time intervals. These intervals, however, are significantly smaller than the interval between the successive samples, so aliasing does occur. It is therefore important to limit the input signal to frequencies smaller than $1/f$. The input stage of the SIS3300 contains an analog 50 MHz filter, so in case of sampling at the maximum rate of 100 MHz, as done here, no additional anti-aliasing measures have to be taken.

Trigger generation

Another important prerequisite for a successful ADC data acquisition is the selection of a suitable trigger mechanism. A simple level/edge trigger, as used in most oscilloscopes, is not suitable for the output of a CSA, as the event signals will not always start out from the same baseline. If two events occur within the relaxation time of the CSA output (for example within 50 μs), the second signal will start off from an elevated baseline and would not activate a level-trigger. While COBRA is a low rate experiment, some important background sources (e.g. ^{214}Bi from the ^{238}U decay chain) generate two events with a time coincidence signature of just this type.

Fortunately, a special firmware for spectroscopy applications is available for

the SIS3000 Fast-ADC. The trigger activates on steps within the signal pulses - the minimum pulse height and polarity can be configured individually for each channel. The trigger input p_i is the difference of two moving averages over N samples each, with a separation of M samples between the averages:

$$p_i = \frac{1}{N} \sum_{j=i}^{i+N-1} y_j - y_{j+N+M} \quad (4.1)$$

Here, p_i is a pulse height measure, proportional to the charge collected by the anode connected to the channel. It is the optimal trigger criterion, since the charge deposition is the central quantity of interest in the experiment. In effect, a digital shaper is applied to the input signal and the trigger is activated once the shaper output exceeds the trigger level. This is the digital equivalent of the shaper and trigger in a classic analog spectroscopy DAQ chain.

The ADC samples and digitises the input signal continuously. After the trigger criterion is met, the ADC continues sampling for a preset number of samples before switching to a new data buffer for the next event. The digitised signal therefore contains both pre- and post-trigger information.

The data presented here was acquired using a value of 128 for N and M , with 384 post-trigger samples to ensure that sufficient number baseline samples are recorded at the beginning and end of each event.

Custom firmware

The SIS3300 generated a timestamp for each event. The timestamp is based on the sampling clock, with a configurable pre-scaler, so the trigger time for each event can be recorded with a precision of 10 ns. This timing information is essential to align the data produced by separate ADC modules during analysis, a necessary prerequisite when looking for coincidences and crosstalk between detectors.

Unfortunately, the SIS3300 firmware allows only for a timestamp resolution of 24 bits, leading to an overflow of the timestamp counter every 0.17 s. Also, the timestamp is reset on each memory bank switch (every 128 events in this case). To remedy these problems, a custom modification of the firmware was ordered from the manufacturer during the course of this work. The new firmware features 48-bit timestamps which are not reset on bank switch, allowing for continuous recording over 32 day without timestamp overflow. This more than meets the requirement of the experiment. The new firmware was delivered in June 2011 and is already in daily use at TU Dortmund. It will be deployed at LNGS and the other COBRA institutes in the fall of 2011.

Custom hardware modifications

The Amanda SIS3300 modules are configured for a peak-to-peak input range of 5 V which is then attenuated internally to the native input range of the actual

4 CPG Detectors with Pulse-Shape Readout

ADC IC, which is 2 V. Amplifying the COBRA preamplifier output levels (about 300 mV for at 10 MeV) to this input range - only to have them attenuated again - is not desirable regarding both signal quality and power consumption. Also, the adjustable range of the input range was not sufficient for our requirements. Furthermore, the firmware of the modules can not be written with current hardware (XILINX USB Cable II), requiring an improved Joint Test Action Group (JTAG) termination on the modules.

Changing certain resistors on the SIS3300 modules allowed for the following changes to be made:

- Proper JTAG termination for firmware exchange
- Change of the peak-to-peak input range to 2V
- Extending the input offset range to over -1 V to +1 V

The changes were introduced and tested step by step, and so have not yet been applied to all modules in use.

4.2.5 Scalability considerations

The whole DAQ chain is, in principle, scalable to an arbitrary number of detectors. This will, at some point, require substitution of the modular preamplifiers with integrated, high density versions. Also, for thousands of detectors, the fast-amplifier / ADC combination will have to be replaced with a high density solution. However, operation on a larger scale will also provide the production capabilities necessary to do this. Small integrated Fast-ADCs are now commercially available and the COBRA group at FMF Freiburg has developed a small embeddable preamplifier circuit [D⁺10].

4.3 Pulse-Shape Analysis

4.3.1 Anode Pulse Shapes

The shape of the charge signal registered at the anodes is sensitive to the location of the event within the detector volume, especially on the depth of the interaction.

Typically (Fig. 4.14), the CA signal rises monotonously while the electrons drift from the interaction point in the detector volume to the collecting anode and the charge-sensitive pre-amplifier integrates the current induced on the anode. The charge contribution of the drifting holes is mostly lost due to their low velocity.

The NCA signal, at first, rises along with the CA signal: Before the charges enter the region near the anode grid dominated by the grid potential difference, the detector behaves similar to a planar design. The NCA signal then falls sharply

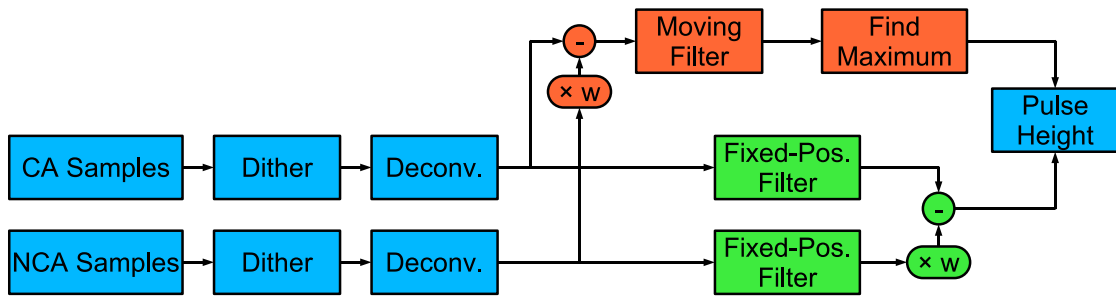


Figure 4.12: Digital Signal Processing chain. After dithering and deconvolution, the are two possible processing paths.




Figure 4.13: Anode pulse shapes before calibration and deconvolution: The pulses decay exponentially. As the decay time is much longer than the record time, a linear approximation can be used.

4 CPG Detectors with Pulse-Shape Readout

as the grid potential dominates over the cathode potential near the anodes, and the charge drift from the NCA to the CA. The CA signal matches this with an increase in slope.

After the charge collection process is complete, the anode signals decay exponentially due to the resistive reset in the pre-amplifier (see Fig. 4.13).

4.3.2 Dithering

During digitisation, the ADC converts the continuous values of the analog input signal to a discrete series of digital values (0 to 4095 for the 12-bit ADC used). The resulting difference between the analog value and its digital representation is called quantisation error and is (ideally) uniformly distributed in the interval $[-1/2, 1/2]$ [Wat01]. If one is only interested in the raw shape of the signal, or if only integer calculations are performed on the sample values, this error can usually be ignored. However, if the values are to be binned (i.e. requantised), for example during histogramming, the new quantisation has to be an even multiple of the original one, so that each new bin encompasses the same number of original quantisation bins. Otherwise, quantisation artefacts will appear in the results since some bins will collect greater parts of the input range than others. In a histogram, these artefacts are visible as peaks with regular intervals, a highly undesirable feature.

If linear operations are performed on the signal (i.e. scaling), or linear metrics of the signal are calculated, the quantisation error range is scaled. Once

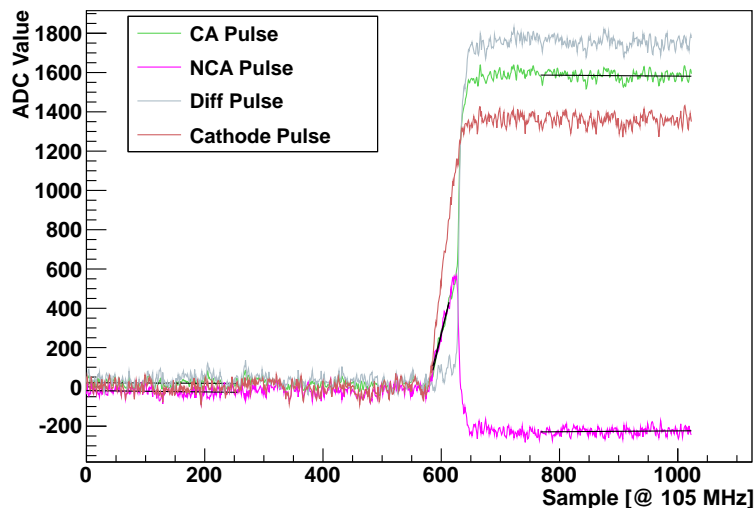


Figure 4.14: Anatomy of the CA and NCA pulse shapes of an event in the main bulk of the detector crystal after deconvolution. The cathode signal has been reconstructed from the anodes signals (see 4.3.6).




Figure 4.15: Anatomy of the CA and NCA pulse shapes of an event near the anode grid. The CA and NCA pulse are almost mirror images of each other, since the charges drift only between the anodes. Almost no current is registered at the cathode.

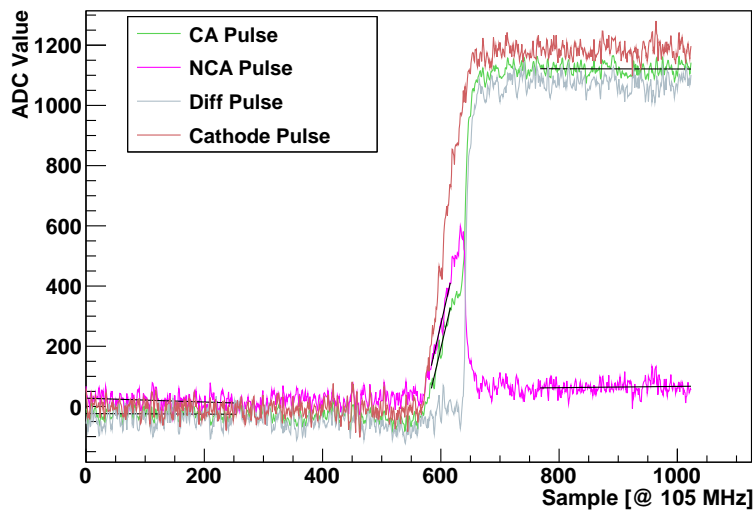


Figure 4.16: Anatomy of the CA and NCA pulse shapes of an event near the cathode. The electrons drift over the full height of the detector crystal. Since the signal has no hole component, no charge is lost and the cathode pulse reaches a similar height as the anode difference pulse. The integral of the drift current is positive on both anodes.

differently scaled signals are combined, or non-linear operations are performed, the quantisation error takes a complex shape which makes it impossible to find matching bin sizes for requantisation or histogramming.

The source of this problem is the rectangular distribution of the quantisation error and the discrete representation of the digitised values. A common and effective solution [Gra93] is to transform the input values to continuously distributed values by adding artificial, low-level noise to the signal. This process is called dithering. The added noise effectively randomises the quantisation error, making the signal suitable for arbitrary processing free of quantisation artefacts.

For dithering time domain signals one commonly uses noise with a triangular distribution function in the interval $[-1, 1]$ [Wat01]. The noise level may have to be increased if the quantisation properties of the ADC are far from ideal.

Random numbers with a triangular distribution can be efficiently computed from equally distributed random numbers: Integration and inversion of the distribution functions yields

$$r_{triang} = \begin{cases} -\sqrt{2-2r} + 1 & \text{if } r \geq 0.5 \\ \sqrt{2r} - 1 & \text{else} \end{cases} \quad (4.2)$$

4.3.3 Signal Calibration

After dithering, the dynamic range of the ADC signal is unchanged. The concrete digital representation of the charge collected at the respective anode depends on the input offset configured at the ADC, the amplification factor configured at the linear amplifier and manufacturing tolerances in all analogue components of the DAQ chain. While an absolute calibration of the signal is not necessary, since the signal processing output will be calibrated afterwards anyway, some of the following processing steps depend on an offset calibration and a calibration of the charge amplification of the CA and NCA channel in relation to each other. A simple linear calibration is sufficient to transform the digital values accordingly, so that a value of 0 corresponds to the preamplifier output baseline and so that equal values represent an equal charge deposition on both channels.

4.3.4 Deconvolution

As detailed above, the input signal decays exponentially after reaching its peak. This results in errors, for example when determining the height of the pulse by comparing the signal before and after its rise, since the height is dependent on the position at which it is taken. This error is commonly termed the ballistic deficit [Leo93].

For scenarios involving only the pulse height of a single, monotonous signal, for example when dealing with a planar detector, the decay of the signal can be

corrected and its height calculated in a single step using the Moving Window Deconvolution (MWD) algorithm [GGL94, SSGG96]. MWD is the common method currently used for pulse height determination for HPGe detectors. It is, however, not suited to the signal of coplanar grid detectors, since multiple signals - generated by preamplifiers with different decay times due to manufacturing tolerances - have to be combined. Therefore, the signal decay has to be corrected by individual deconvolution of the CA and NCA output before further processing.

The preamplifier is, from a signal processing point of view, just an active low-pass filter with a time constant τ . Its effect can, in the digital domain, be described as an IIR-filter with input values x_i and output values y_i :

$$\alpha = \frac{1}{1 + \frac{\tau}{dt}} \quad (4.3)$$

$$y_i = \alpha x_{i-1} + (1 - \alpha)y_{i-1} \quad (4.4)$$

While the integration property of this filter is desired, the signal decay is not. To correct for this, we apply the following deconvolution to the signal:

$$b_i = b_{i-1} + \alpha x_{i-1} \quad (4.5)$$

$$y_i = x_i + b_{i-1} \quad (4.6)$$

Eq. 4.5 compensates the exponential decay (see Fig. 4.17) by accumulating the total loss in b_i . Since the exponential relaxation converges to the preamplifier output baseline, a correct offset calibration of the ADC channel, as described above, is necessary for this to function.




Figure 4.17: Deconvolution of preamplifier signal, compensating the exponential decay [Qua10].

4.3.5 Weighted Anode Difference

As described in section 1.3.2, the energy deposition in the detector is calculated by determining the pulse height of the weighted difference between the CA and NCA signal. The resulting difference signal y^{Diff} is easily computed as

$$y_i^{Diff} = y_i^{CA} - f_W y_i^{NCA} \quad (4.7)$$

Apart from its height, its detailed shape is also of interest because it differs for single- and multi-site events. Since it rises quickly as the charge cloud reaches the anode grid, it is also a good basis to determine the end of the drift time.

4.3.6 Cathode-Signal Reconstruction

Also of interest is the signal of the detector cathode. While the cathode is not instrumented, its signal is necessarily equal to the sum of the anode signals due to Kirchhoff's law. Once the CA and NCA channels have to be calibrated correctly relative to each other (described above), with a relative charge amplification of f_A , the cathode signal can be calculated:

$$y_i^{Diff} = f_A y_i^{CA} + y_i^{NCA} \quad (4.8)$$

4.3.7 Window Functions

A number of the following methods are based on summing or averaging over a contiguous subset of N signal samples, or a function of them - for example calculating the average of number of samples at the beginning or the end of the signal to determine the baseline levels. The aim of this aggregation is always to extract some feature from the signal while suppressing noise-based influences on the result. Aggregating the samples of the subset with equal weight can also be described as aggregating all samples, weighted with a rectangular function (also discrete in time) of length N with an amplitude of $1/N$ for the subset samples and 0 for the rest.

This type of function is called a window function, since it defines a view on a part of the signal. A rectangular window function will not result in an optimal noise suppression, though. Better results are achieved by using low weights for the samples at the beginning and end of the subset and high weights for those in the middle. There are many window functions in use in signal processing in general, with different potential regarding noise-suppression, resolution and sensitivity. Here we choose the Hamming window, defined by the weights:

$$w_i = 0.54 - 0.46 \cos\left(\frac{2\pi i}{N-1}\right) \quad (4.9)$$

The Hamming window is a popular choice for narrow-band applications [Wat01], as is the case here, since it represents a good compromise between resolution and sensitivity.

4.3.8 Pulse Properties

Baseline Levels and Slopes

Determination of the level and slope of the baseline before and after the rise of the signal is a necessary prerequisite for one form of pulse height computation as well as for the estimation of several calibration parameters described below. The baseline level is determined by an average of the baseline samples using the window function values w_i :

$$B_l = \frac{\sum_{i=0}^{N-1} w_i y_i}{\sum_{i=0}^{N-1} w_i} \quad (4.10)$$

The window is typically selected to encompass the first, resp. last, 128 to 256 samples of the signal. The baseline slope B_s is determined using linear fit over the same samples (Fig. 4.14).

Pulse Height

The pulse height P of the signal can be computed as the difference between the signal baseline before and after rise:

$$P_{BD} = B_{post} - B_{pre} \quad (4.11)$$

However, a better resolution can be achieved using a moving-window technique, the same way the ADC generates its trigger signal:

$$P_{MW} = \max_i \frac{1}{N} \sum_{j=i}^{i+N-1} y_j - y_{j+N+M} \quad (4.12)$$

Unfortunately, this method is only applicable for monotonous signals. For non-monotonous signals, like the NCA signal (and seldom the CA signal), with two maxima of converse polarity, the difference between the maxima influences the result. Therefore, we need to use both methods, the choice depending on the type of signal.

As a direct measure for the energy deposition, the pulse height of the weighted anode difference signal y^{Diff} is the single most important property of interest. It can, in principle, be determined in two ways: Directly from the difference signal using the P_{MW} function, or as the weighted difference of the pulse heights of the CA and NCA signal.

4 CPG Detectors with Pulse-Shape Readout

Since the difference signal is monotonous, its height can be calculated using the moving-window technique, yielding the best resolution. This method is used for the final results.

The pulse heights of the CA and NCA signals, being non-monotonous, have to be calculated from their baseline differences, resulting in a lower resolution. They can, on the other hand, be calculated independently without prior knowledge of the weighting factor - this is important during the calibration process as shown below.

We aim to achieve the best energy resolution possible. It should only be limited by the detector performance, not by the resolution of the ADC. When configuring the system for an input range up to 8 keV, the native 12-bit resolution of the ADC results in a digital resolution of 2 keV. While this would be acceptable, the effective resolution of the system is in fact even higher, since the pulse heights are computed here from a mean over 128 or more samples. This technique, called oversampling, results in an increased resolution described by the oversampling theorem, which states that to gain an additional n bits of resolution

$$N = 2^{2^n} \quad (4.13)$$

samples have to be averaged. An oversampling factor of > 64 , which is always achieved in our case, therefore results in a resolution increase of 3 bits with a total dynamic range of 15 bits, corresponding to a theoretical resolution 0.25 keV for an input range of 8 MeV.

The digital pulse height estimation was tested using a Berkeley Nucleonics PB5 research pulser, configured to resemble a typical CA signal. For the results, see Fig. 4.18. The measured resolution naturally does not quite reach the theoretical value, due to the analog noise inherent to the physical components. When directly connecting the pulser to the ADC, a resolution of 0.3 bit is achieved. With the differential signal transmission chain in place, the resolution is 0.9 bit for an amplification factor of 1, and 1.0 bit for a factor of 20 (with the pulser set to $20 \times$ attenuation). These results demonstrate the excellent performance of the new differential signal transmission and amplification.

Rise time

The rise time of a digital signal can be estimated as the difference between the point where the signal intersects a minimum percentage (e.g. 10 %) and the point of intersection with a maximum percentage (e.g. 90 %) of its final height [ACK⁺08]. While simple in principle some care has to be taken to make the result robust against noise-based fluctuations of the signal, since they may result in multiple points of intersect. The minimum intersection should be searched for starting from the end of the signal, so that it can't be confused with fluctuations on the pre-rise baseline. For the same reason, the maximum intersection should

be searched for from the beginning of the signal. Robustness against momentary noise spikes during the rise itself can be achieved by requiring the signal to stay above, resp. below the specified threshold for a certain number of samples after intersecting it. Low-pass filtering of the signal cannot be used as a noise-reduction measure here as it would affect the rise time itself.

4.3.9 Analysis Calibration Parameters

The computation of the pulse shape properties described above depends on the following calibration parameters:

- Total amplifier and ADC offset
- Pre-amplifier signal relaxation time constant
- Relative charge calibration between CA and NCA
- Optimal weighting factor for the anode difference signal

We aim to estimate all these parameters from the events recorded during detector calibration with a γ -source like ^{137}Cs or ^{228}Th . The sources used have a moderate activity, resulting in an event rate of less than 1000 /s.

Offsets

We define the total offset of the DAQ chain as the ADC output for the mean dark current input into the CSA. Since the digitized pre-amplifier signal will exponentially relax to this value after an event, this is the zero level that must be used in the deconvolution step. This level can be estimated as the mean of the pre-pulse baseline distribution, since the CSA signal will usually fully relax between events due to the moderate activity of the calibration source (see Fig. 4.19).




Figure 4.18: Resolution of the digital pulse height estimation for a pulser signal, with and without the new differential amplifier.

4 CPG Detectors with Pulse-Shape Readout




Figure 4.19: Distribution of the pre-pulse baseline level. The offset of the DAQ chain can be estimated from the mean.

CSA Relaxation Time-Constant

The post-pulse baseline fit described above spans only a small part (usually $2.5 \mu\text{s}$) of the relaxation time constant (usually $> 50 \mu\text{s}$). Because of this, it can be described by a first order approximation

$$a e^{-t/\tau} \approx a - t \frac{a}{\tau}, \quad t \ll \tau \quad (4.14)$$

Therefore, the time constant τ can be determined from the relation between the pulse height and the slope of the post-pulse baseline (see Fig. 4.20).

$$\tau \approx -\frac{B_l^{post}}{B_s^{post}} \quad (4.15)$$

Relative Amplification

Estimation of the relative charge amplification between the CA and the NCA channel over the whole DAQ chain is more complicated. The relative pulse




Figure 4.20: Distribution of the post-pulse baseline slope as a function of the pulse-height. The CSA time constant τ is determined from a linear fit.

heights depend on the grid bias voltage as well as the interaction position in the detector, and so cannot be used. However, for events originating near the cathode, both CA and NCA register the same current during the first part of the charge drift. Such events feature small negative, or even positive, NCA pulse heights with a strong and positive initial rise (see Fig. 4.16).

The slope of this rise, and the corresponding section of the CA signal, is proportional to this drift current. The relative current and charge amplification of the two channels can therefore be determined from the ratio of these slopes. The method was verified by measurements with preamplifier modules from the same production run and manually calibrated settings for the linear amplifier, resulting in a relative amplification of almost 1 (See Fig. 4.21), as expected.

4.3.10 Physical Quantities

Energy Deposition

As stated before, the energy deposited within the detector volume is proportional to the pulse height of the correctly weighted anode signal difference. The energy deposition can therefore be calculated as

$$E = a_E + b_E P_{MW}^{diff} \quad (4.16)$$

The energy response of the system has been tested over a wide energy range (see Fig. 4.22 and justifies the use of a linear calibration model. The calibration offset is typically very small, since the pulse height is computed from a difference in the digitized signal and therefore free from analog offsets. For laboratory applications such as detector qualification measurements, a calibration with a single-line γ -source is therefore usually sufficient. For low-background measurements, the complete detector system is calibrated with a Th-228 source, providing




Figure 4.21: Relation between slope of the initial rise of the NCA and CA signal. Due to carefully chosen analog parameters, the relative charge amplification is almost 1 here.

4 CPG Detectors with Pulse-Shape Readout




Figure 4.22: Multi-point calibration of a CPG Detector at LNGS with a Th-228 source. The plot shows the good linearity in the energy response of the detector and DAQ chain.

up to 8 energy lines (depending on the detector resolution). Since this also provides a measurement of the energy resolution itself over a wide range, the Th-228 calibration is also used to determine an energy resolution function. For details pertaining to the calibration process see [Kö08].

The new DAQ chain and the digital pulse processing have resulted in a significantly increased energy resolution. Results with the the previous DAQ chain were often erratic and uncorrelated with the manufacturer's specification [Mü07, Fig. 3.14]. Also, the shape of the energy peaks was often non-gaussian, due to the difficulty of determining the optimal weighting factor in the analog solution. The new method delivers results which are consistently better [Qua10] than both previously measured values and the resolution specified by the manufacturer (See e.g. Fig. 4.23).

This is highly important, since the energy discrimination characteristics of the detectors are a central factor in the success of the COBRA experiment 1.3.

Interaction depth

Another interesting quantity to be determined is the interaction depth. Here, we define the interaction depth as the distance of the event within the detector volume from the anode grid plane.




Figure 4.23: Cs-137 Spectrum of a spectrometer-grade evProducts CPG detector using the new DAQ chain and digital pulse processing. The photo peak has an excellent energy resolution of 1.95% and a symmetrical, Gaussian shape.

4 CPG Detectors with Pulse-Shape Readout

As shown in [HKW⁺96], the charge C^{Cath} collected at the cathode of the detector is proportional to the interaction depth D and the total energy deposited E :

$$C^{Cath} \sim ZE \quad (4.17)$$

The interaction depth is then determined by

$$D = b_D \frac{C^{Cath}}{E} \quad (4.18)$$

or

$$D = b_D \frac{P_{Cath}}{P_{Diff}} \quad (4.19)$$

In case of a multi-site event, the result will be the mean depth of the individual interactions, due to the mostly independent drift process of the separate charge clouds.

The calibration parameter b_D is not 1 as one might expect, since P_{Cath} and P_{Diff} are not calibrated in relation to each other because of the weighting factor. In practice, b_D usually takes values between 0.8 and 0.9.

Measurements have shown that the resulting interaction depth is not equally distributed (see Fig. 4.24). While a substantially reduced efficiency near the anode grid is expected [HKW⁺96], the slope over the main detector bulk is unexpected. It has been suggested that this is due to trapping and due to a partial contribution of the drifting holes to the anode signals [HKWM97]. The depth measure




Figure 4.24: Distribution of the interaction depth in a CPG detector illuminated by a ^{137}Cs source.

presented here therefore results in values of 0 for anode-side and 1 for cathode-side events, but is probably not linear in between. Nonetheless, it is already a valuable tool to identify events near the electrodes (see Fig. 6.7).

Higher order corrections might be necessary to get a linear measure of the interaction depth. It is also possible to derive depth information from the charge drift time [McG09], so a combination of the two methods should be explored in the future. Work is currently underway at the COBRA group at TU-Dresden to make detailed scans of several CPG detectors with a collimated source. These scans will yield absolute depth information for individual events, resulting in a verified and linear depth estimation method.

4.4 Software implementation

The pulse-shape processing and analysis algorithms described above were implemented in form of a software package base on the CERN ROOT System [A⁺09]. The package, named Multiple-Analysis Toolkit for the COBRA Experiment (MAnTiCORE) is now available to the collaboration. The software provides special support to manage and track the numerous processing settings and calibration parameters. While MAnTiCORE currently supports only CPG pulse shape data, care was taken to keep its design flexible enough so that processing for other data types, e.g. pixel detector data, can be added without conflicts.

4.5 Data Compression Options

The switch from single-value pulse height recording to pulse shape recording increases the data volume by 3 orders of magnitude. During low-background data taking, the amount of data collected is still small, currently about 500 MB per detector and year. A single calibration of the planned 64 detector array, will, however produce 250 to 500 GB of raw pulse-shape data. The effort of scanning the detectors with a collimated source will result in 50 to 150 GB per detector. The total expected data volume over the next years will therefore lie between 10 and 50 TB for characterisation, optimisation and regular calibration of a 64 detector setup. Since this is valuable data, the necessary replication and/or backup will multiply this by an additional factor of 2 to 3. A precise prediction is difficult, but depending on the outcome the data volume may easily stretch or even exceed the current resources of the collaboration. A further increase in the number of detector will make this even more problematic.

For these reasons, a significant reduction of the pulse shape data will likely be necessary in the future to maintain the scalability of the experiment. Deleting the raw pulse shapes and keeping only analysis results is a possible but undesirable solution: It would preclude any future improvements and extensions of the

calibration and analysis algorithms from being applied to previously recorded data. The raw data is already losslessly compressed by the ROOT file format, and while slightly better compression ratios could be achieved by using an optimized format and better compression algorithm, tests have shown that the data can at best be compressed a further factor of 2. A lossy compression of the pulse data, preserving pulse shape and energy resolution while discarding unnecessary high entropy information (i.e. noise) should therefore be seriously considered.

Since the anode pulses are transient, non-periodic in nature, a Wavelet-transformation seems to be a better fit than a Fourier-transformation based compression scheme. While the development of an optimized Wavelet-compressor with ideal coefficient coding and compression is beyond the scope of the work, the principal feasibility of a Wavelet coder for the given problem shall be explored here.

4.5.1 Wavelets and Multiresolution

Wavelets have become a popular and powerful tool for signal processing and analysis in recent years. Wavelet theory is closely related to multiresolution theory - the wavelets used here, like most wavelets of practical interest, stem directly from a multiresolution analysis, so we consider both concepts together in the following.

Formally, a multiresolution analysis is a sequence of nested subspaces on L^2

$$\dots \supset V_{-1} \supset V_0 \supset V_1 \supset V_2 \supset \dots \quad (4.20)$$

satisfying particular requirements. In the context of data processing, these subspaces present a view of the data at different scales, resp. levels of detail. They¹ are defined by a so-called scaling function ϕ , which is stretched by a factor of 2^i and translated by a distance of $j2^i$ to provide a basis for each subspace V_i :

$$\phi_{ij}(x) = \sqrt{2^{-i}} \phi(2^{-i}x - j) \quad (4.21)$$

It is sufficient that the functions ϕ_{ij} are linearly independent [Rei95] (though they may even be orthogonal), so that every function f_i in V_i may be decomposed as

$$f_i = \sum_j s_{ij} \phi_{ij} \quad (4.22)$$

The coefficients s_{ij} are called the scaling coefficients of f_i . Further requirements for a multiresolution analysis are, that the spaces V_i must not be redundant and must approximate L^2 with arbitrary precision:

¹The spaces are sometimes indexed in reverse order in the literature, we follow the indexing introduced by Daubechies [Dau92] here.

$$\bigcap_i V_i = \{0\} \quad \text{and} \quad \bigcup_i V_i = L^2 \quad (4.23)$$

The V_i must provide invariance regarding translations proportional to 2^i

$$f(x) \in V_i \iff f(x - 2^i j) \in V_i \quad (4.24)$$

and finally, they must satisfy the multiresolution condition

$$f(x) \in V_i \iff f\left(\frac{x}{2}\right) \in V_{i+1} \quad (4.25)$$

The latter results in the fact that V_{i+1} only provides half the level of detail in comparison to V_i . In fact, V_i can (in our case) be written as the direct sum of V_{i+1} and a complimentary wavelet subspace W_{i+1} :

$$V_i = V_{i+1} \oplus W_{i+1} \quad (4.26)$$

In other words, W_{i+1} contains the difference in detail between V_i and V_{i+1} . The wavelet subspace W_{i+1} is spanned by the scaled and translated copies ψ_{ij} of the wavelet function ψ

$$\psi_{ij}(x) = \sqrt{2^{-i}} \psi(2^{-i}x - j) \quad (4.27)$$

and the functions in W_i can be decomposed into their wavelet coefficients

$$f_i = \sum_j w_{ij} \psi_{ij} \quad (4.28)$$

The scaling coefficients s_{ij} describe the data at the scale i , while the wavelet coefficients w_{ij} contain the difference in detail to the scale $i - 1$. The scaling function is of an integrating nature, it smooths details and can also be viewed as a low-pass filter. The wavelet function is of a differentiating nature, it exposes detail and is therefore a high-pass filter. In the case of the the most basic wavelet, the Haar wavelet, scaling and wavelet functions are actually true numeric integration and differentiation operators (see Fig. 4.25).

The strength of multiresolution analysis is that the scaling functions and the wavelet functions are local. Thus, the resulting coefficients provide local information at different scales, whereas a Fourier analysis only yields global information.

In practice, data is always acquired with a limited level of detail - in our case the data consists of the pulse shape samples y_j , acquired with a limited time resolution. If we call this level of detail the scale $i = 0$, then the data samples are equal to the scaling coefficients at this scale:

$$s_{0j} = y_j \quad (4.29)$$

4 CPG Detectors with Pulse-Shape Readout

The data can then be decomposed into a projection to a scale L of lower resolution (with fewer scaling coefficients), plus all the detail in between the two scales:

$$V_0 = V_L \oplus W_L \oplus W_{L-1} \dots \oplus W_1 \quad (4.30)$$

The number of coefficients is equal on both sides, so no information is lost. Due to Eq. 4.26, decomposition can be done in a recursive fashion (see Fig. 4.27).

Unfortunately, the Haar wavelet is the only orthogonal symmetric wavelet with limited support. In practical numeric applications it is therefore often necessary to sacrifice orthogonality in favour of other required properties. Biorthogonal wavelets [CDF92] provide limited support, symmetry and other important properties and are very popular for this reason [URB97, UB03]. Since they are not orthogonal, different functions are used for the decomposition of data (called the analyzing wavelet and scaling functions) and the reconstruction of the original data from the coefficients (reconstruction wavelet and scaling functions). In the following, we will use the biorthogonal CDF(2,2) wavelet (see Fig. 4.26) which has the additional advantage that it can be used to decompose integer data, resulting in integer coefficients.

A detailed introduction into wavelet and multiresolution theory is beyond the scope of this work, see [Gra95] or [Mal99] for an excellent introduction into these topics.

4.5.2 Wavelet compression

Here, wavelets will be used to transform the data in a way that separates the low-entropy information of interest - the shape of the pulses in the signal - from high entropy information that can be discarded. This is achieved by discarding all




Figure 4.25: The Haar scaling function (top) and wavelet function (bottom).




Figure 4.26: Top: The analyzing CDF(2,2) scaling and wavelet function. Bottom: The reconstruction CDF(2,2) scaling and wavelet function.

thresholds	size / MB	compr. ratio	Error / keV
original data	216.6	1	0
$j_{thresh} = 5, w_{thresh} = 75$	21.5	10.1	0.45
$j_{thresh} = 7, w_{thresh} = 70$	21.6	10.0	0.96
$j_{thresh} = 6, w_{thresh} = 250$	10.2	21.2	1.21
$j_{thresh} = 7, w_{thresh} = 200$	9.9	21.3	1.96

Table 4.1: Different wavelet compression thresholds for compression ratios of 10 and 20. All data was stored in the ROOT file format. The error given here is the mean square error in pulse height between the original and the compressed data (see Fig. 4.31) and almost independent of the energy deposition (see Fig. 4.32)

wavelet coefficients below a preselected threshold w_{thresh} on all scales $j > j_{thresh}$ with a high level of detail. The remaining data is then compressed using a regular lossless compressor - in this case, the gzip compression intrinsic to the ROOT file format.

For the actual wavelet transformation, the blitzwave [Sch] wavelet library was used. Blitzwave implements the fast wavelet transform using the lifting scheme [Swe95, URB97], delivering high numerical performance.

4.5.3 Compression results

Different Wavelet coefficient thresholds were tested on a dataset containing high-resolution ^{137}Cs calibration data and the parameters resulting in the least impact on the energy resolution and pulse shape for compression factors of 10 ($j_{thresh} = 5, w_{thresh} = 75$) and 20 ($j_{thresh} = 6, w_{thresh} = 250$) were determined. See Table 4.1 for selected results.




Figure 4.27: Recursive Wavelet Decomposition and Recomposition. V_0 is the space of the original data. The subspaces V_i represent the data at progressively lower resolution, the difference is preserved in the wavelet spaces W_i . The number of scaling coefficients is halved in each step.

4 CPG Detectors with Pulse-Shape Readout

Fig. 4.28, 4.29 and 4.30 show the impact of the lossy wavelet compression on the pulse shapes. While the noise information is lost, the shape and height of the pulses is preserved, and even very steep rises of the signal are faithfully reproduced. The statistical and systematic difference between the event energy before and after compression (Fig. 4.31) is very small compared to the energy resolution of the best CPG detectors available, and independent of the energy deposition (Fig. 4.32). The energy spectrum itself is not affected by the compression at all (Fig. 4.33).




Figure 4.28: CA pulse shape before and after compression by a factor of 10. While the noise information is lost for the most part, the shape of the signal is preserved.




Figure 4.29: NCA pulse shape before and after compression by a factor of 10.




Figure 4.30: Preservation of rise-time and timing information under wavelet compression (with a compression factor of 10). A CA signal with a sharp rise was selected here, to show that even very short rise-times are reconstructed perfectly after decompression.

4 CPG Detectors with Pulse-Shape Readout




Figure 4.31: Error introduced into the energy deposition values by the Wavelet-compression, for compression factors of 10 and 20. The statistical error is 0.45 keV with a systematic shift of -0.21 keV for a compression factor of 10, and 1.21 keV (statistical) resp. -0.16 keV (systematic) for a compression factor of 20. The errors are small compared to the energy resolution of the detector of 13 keV (at 662 keV).




Figure 4.32: Distribution of the error in energy ΔE induced by Wavelet compression (with a compression factor of 20) as a function of the event energy. The error is almost constant over the energy range (shown here for a ^{137}Cs -Spectrum).




Figure 4.33: Effect of the wavelet compression on the energy spectrum. The spectra of the original and compressed data are identical apart from minimal binning effects (shown here for a ^{137}Cs -Spectrum and a compression factor of 10).

4 *CPG Detectors with Pulse-Shape Readout*

5 COBRA DAQ System

The COBRA data-acquisition system in use when this work was begun [Kie05, chapter 3, p55] was monolithic (Fig. 5.1). A VME CPU board acted as the main controller for both physics data acquisition, data recording and also as a supervisory control and data acquisition (SCADA) (also called slow-control in physics experiments) system to control voltages, to set amplifier parameters and to measure environmental conditions. Changes to the software required a shutdown of all these systems, including the detector bias voltages, resulting in the invalidation of the detector calibration. Extensions of the system were limited by the I/O ports provided by the CPU board, often causing serious inconvenience. Since the software was written specifically for the used CPU board, switching to different hardware was also not possible.

The planned upgrades of the experimental setup were simply not possible by just extending this system, it was therefore decided to design and implement a new, more flexible DAQ system.

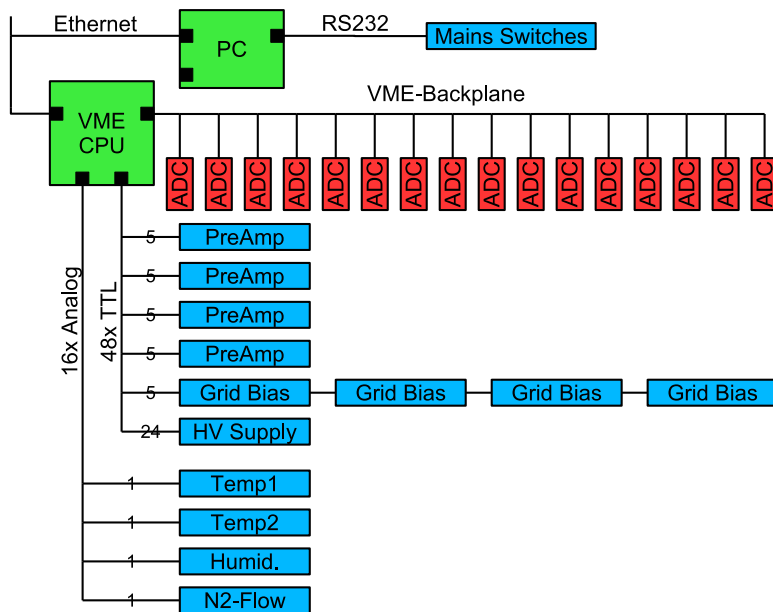


Figure 5.1: Previous monolithic COBRA DAQ structure. Almost all components are directly connected to a single VME CPU board.

5.1 Requirements

The new DAQ system has to be able to handle the current and midterm requirements of the COBRA collaboration. It should also, at least in its basic design, be able to scale to a full sized COBRA experiment, consisting of tens of thousands of detectors.

Both physics data acquisition and slow-control/SCADA functionality should be provided - however, to enable independent shutdown and upgrade, these aspects of the system should be able to run on different computer systems while still being able to communicate.

The COBRA experiment is still in the prototype phase - the hardware components used in the setup, the basic design of the setup itself and the required software features are therefore subject to regular change. Also, the new system should ideally not only support the main physics setup at LNGS, but should also be useable for the COBRA R&D efforts at the collaboration institutes. These R&D-setups, since they differ in purpose, also differ in structure from each other and change even more frequently.

The change to a pulse shape based data acquisition for the CPG detector increases the data volume to be handled by three orders of magnitude. While COBRA is a low-background, and therefore in principle a low data rate experiment, with pulse shape recording this is no longer true during calibration. A data volume of 250 to 500 GB for a single calibration of the planned 64-detector CPG array, will -assuming a calibration time of 5 hours - result in an incoming data rate of 15 to 30 MB/s. While this can still (barely) be handled by a single computing host, further extensions of the setup will soon require a distributed, multi-host system.

As a lesson learned from experiences with the old system, the core of the new DAQ system should avoid any hardware dependencies.

While the CPG-detector data acquisition system has always been Linux-based, two of the three Pixel detector systems evaluated by the collaboration so far only ran on MS Windows. If it is to support future plans of combining CPG and pixel detectors, the new DAQ software must be able to run on multiple platforms so it can act (at least) as a rudimentary bridge to these 3rd-party systems.

Since COBRA is a comparatively small collaboration, the system must require little manpower for operation and maintenance. Full control of the system must be possible using a non-graphical command-line interface while the implementation of a graphical or web-based control interface, deferred for now, should be well supported by the underlying technology.

The DAQ system must, of course, be able to interoperate with the data formats already in use within the collaboration, especially the file format of the CERN ROOT System [A⁺09] [The11a].

In summary, the new DAQ system is required to be

- modular and flexible
- hardware independent
- highly scalable
- multi-platform capable
- robust and fault tolerant
- easy to operate and maintain
- compatible with existing data formats
- able to provide a command-line interface

5.2 Design Philosophy

5.2.1 Network-centric Approach

To ensure both modularity and scalability on the hardware side, it was decided to use a network-centric design approach. At the top level, all system components are connected via Ethernet, using the IP protocol - no hardware is directly connected to the DAQ hosts. Systems without native Ethernet support are connected via bridges (Fig. 5.2).

This Ethernet-based approach has nowadays become highly popular and successful in small and large scale industrial data acquisition and control systems [Pot99], an application example can be found in [Heg06].

Since Ethernet networks are fast, easily and cheaply extensible, and familiar to the intended users, such a design fits the requirements very well. The complete hardware abstraction of the DAQ hosts also enables the DAQ software to run on modern, virtualized computing infrastructures, which eases maintenance and upgrading. Due to the ready availability of Gigabit-Ethernet in most modern systems, and due to the easy extensibility of networks with high-bandwidth switches, this structure can serve the requirements of a full-scale COBRA setup without problems.

5.2.2 The Actor Model

During the last years, the previously exponential rise in per-CPU processing power has slowed down significantly. It has, instead, been replaced by a rise in parallelism in everyday computing [A⁺06b] [A⁺08]. Nowadays, applications designed for rising demands in the future must be able to exploit multi-CPU hardware and often even support distributed, multi-host operation. This requirement is, unfortunately, not well supported by many traditional programming

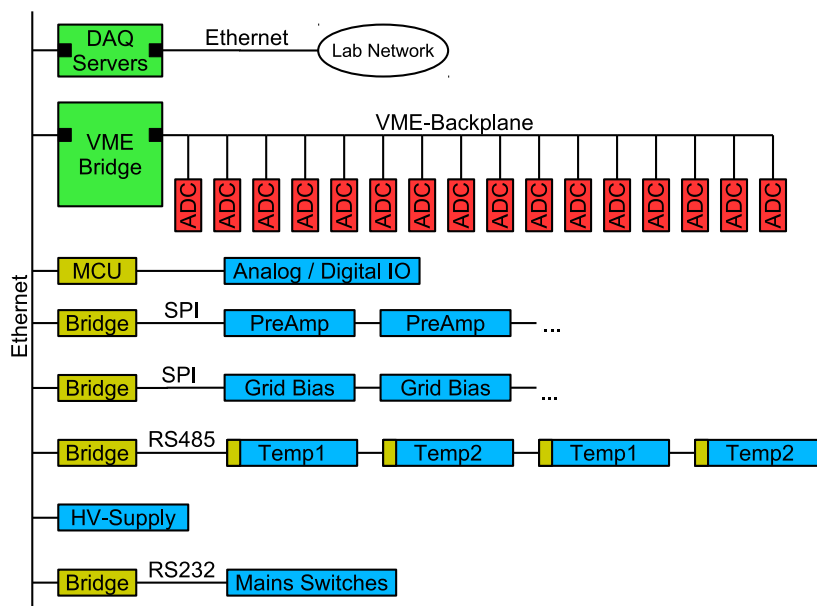


Figure 5.2: New, modular COBRA DAQ structure. All subsystems are connected to an Ethernet backbone.

paradigms like classic object-oriented programming [Pat10], a situation that is discussed extensively in the programming community as the so-called multi-core crisis [Wol08] [Far08].

Communication-centric computing models, on the other hand, are well suited to run on parallel systems. One such computing model in particular, the actor model, has become popular for this reason [KSA09]. It is inherently modular and scalable, fitting the requirement formulated above very well. The actor model is based on independent parallel processes, the actors, which exchange messages in an asynchronous fashion. This model is a natural fit for a DAQ application, since it can easily model the interaction of the individual components in the system.

The actor model has several major advantages over classical lock-based multi-threaded programming [Agh85]: Lock-based programming requires careful analysis to prevent deadlocks and race conditions during data access by concurrent threads. The actor model, based on message passing with no mutable data shared between the actors, is free from deadlocks and race conditions by design. Also, even modern operation systems can only efficiently manage a few hundred co-existing threads. The actors, on the other hand, are not bound to specific system threads, but are executed on a dynamically managed thread pool on demand, which allows tens of thousands of actors to coexist.

5.2.3 Error handling and supervision

Support for the actor model has been an important criterion in the design of several modern programming languages, one of which is Erlang [Arm07]. An interesting property of Erlang is its error handling strategy, which is very suitable for the given application and has already been included in other programming languages [Comb]. Instead of using defensive programming techniques - trying to handle each error in detail and as close to its source as possible - the supervision model embraces failure as an unavoidable, regular occurrence in a large system. This acknowledges the fact that detailed and correct error handling becomes highly nontrivial and is often impossible when dealing with a large concurrent system [AVWW96].

The actors are linked to supervisors, which are actors themselves and monitor one or several other actors. If an error occurs during the execution of one of these actor processes, the affected actor crashes. Its supervisor then restarts the crashed actor, or all its linked actors (see Fig. 5.3), depending on its configuration. The supervisors can have supervisors themselves, forming a so-called supervision hierarchy. If a supervisor experiences repeated crashes of a linked actor, the supervisor itself crashes, thereby delegating the problem to a higher level in the hierarchy and causing the restart of a bigger subsystem (see Fig. 5.4).

Supervision-based error handling has proven to be very robust [AVWW96, chapters 7,8] and also reduces the amount of error handling code to be written.

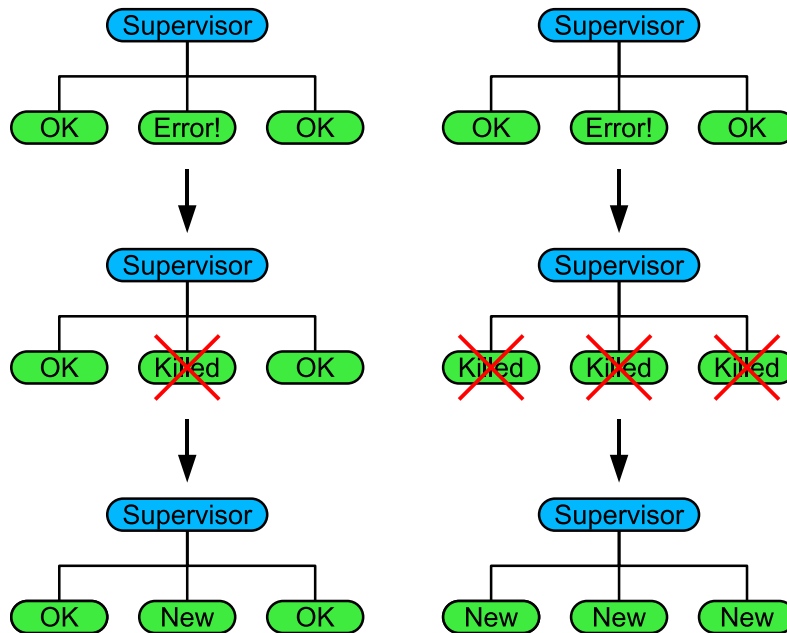


Figure 5.3: Actor supervision: If an actor fails, the supervisor restarts either the failed actor (one-for-one supervision, left), or all linked actors (all-for-one supervision, right).

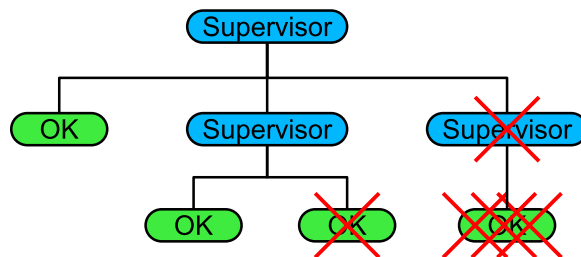


Figure 5.4: Supervision hierarchies: If a supervisor cannot handle the failure of an actor, it is restarted itself.

It is designed for large systems like a full-size COBRA experiment, with tens of thousands of components.

5.2.4 JVM-Based

To ensure easy operation and maintenance, the DAQ software must, first of all, be easy to install, ideally on multiple operating systems. This requirement is very hard to fulfil with natively compiled languages in this specific case - as the software will be required to interface with many different components, it cannot avoid having a large number of dependencies on other software packages and libraries. These will have to be present, or installed manually on all intended target systems. On system upgrades, the software must often be adapted to be compatible with newer version of these packages. In the author's experience, all this either tends to consume a lot of manpower or results in a situation where the software, in the end, only runs on one specific system configuration, updates of which are then avoided as much as possible. Neither outcome is desirable here. Interpreter based platforms (e.g. Python), while flexible and easy to maintain, do not provide enough performance for natively processing large amounts of numeric data (e.g. pulse shapes). Modern virtual machine based platforms with just-in-time compilers, however, provide both easy maintainability (since the software, and even more importantly its dependencies, are binary-compatible across even different operation systems) and high-performance [LN03]. It was therefore decided to base the new DAQ system on the Java Virtual Machine (JVM) platform. The JVM is available for all relevant operating systems and provides access to a software ecosystem on a par with C/C++.

5.2.5 Handling Embedded Systems

The COBRA DAQ setup contains several embedded systems, which require a dedicated control software. Examples are the VME CPU Boards controlling the fast ADCs and the microcontroller boards controlling the amplifier and grid bias voltage settings. Embedded system development is more complex and specialized than standard system programming [SBW01]. Thus testing and debugging are more difficult and development turnaround times are longer. In keeping with the overall design used here, all embedded software was kept as simple as possible, moving most of the functionality into the DAQ software itself. This also makes the embedded components fairly simple to replace. We were, for example, able to exchange the VME CPU board we had used with a completely different product within two days, including all necessary software changes.

5.2.6 Handling Latency

One challenge of a network-based solution, in contrast to a hard-wired one, is that while enough bandwidth is readily available, request/response times have a lower limit of several milliseconds or more, called latency. If the DAQ system was to always wait for a response before issuing the next request to a connected component, its performance would become limited by latency. This can be avoided by pipelining requests, i.e. issuing several non-interdependent requests without waiting for the responses. This is not always easy to implement - fortunately however, the actor model natively supports this communication scheme [Agh85].

5.3 Data-Acquisition Hardware

5.3.1 Physics Data Acquisition

Concerning the physics data acquisition, the main task of the new DAQ system is currently to handle the VME Fast-ADCs. In keeping with the design philosophy, the DAQ communicated with the VME crate over TCP/IP. A small program on a CPU-board in the VME crate acts as a VME to TCP bridge (see Fig. 5.2) and allows remote access to the VME bus. The DAQ can run a different system on a separate system, handling with multiple VME crates at the same time. Since it is capable of distributed operation, it can even run in part on the VME CPU board itself - taking advantage of the faster local communication - and exchange preprocessed and possibly compressed data with the rest of the DAQ running on other systems.

The following components are controlled by the physics DAQ:

- VME CPU boards, acting as a VME-TCP bridge
- Fast ADCs (SIS 3300)

5.3.2 SCADA / Slow Control

On the SCADA side, the DAQ currently handles the following components:

- High voltage supply (Iseg standalone or MPOD system)
- VME crate (Wiener)
- COBRA Fast Amplifiers (custom built)
- COBRA Preamplifier boxes (custom built)

- Setup-specific components, like a collimator scanning table at TU Dresden or the liquid nitrogen Dewar at LNGS

This slow control DAQ can run within the same process as the physics DAQ for quick short-term laboratory applications, but can also run completely separate for modular long-term operation as required for the main setup at LNGS.

5.4 Software Implementation

The new DAQ software framework was named Data-Acquisition and Control Environment (DAQCorE).

5.4.1 Choice of Programming Language

Considering the basic system requirements and design philosophy, the Scala language was chosen for the implementation of DAQCorE. Scala is a very modern programming language on the JVM platform which supports both object-oriented and functional programming, and was designed for building scalable systems [OSV08]. Scala also provides a text console which allows for direct interaction with objects defined in the DAQ code. A scripting mode also allows for setup and measurement scripts, also written in Scala, to easily control and run the DAQ. This means that detailed scriptable configuration and full command-line control of the DAQ are provided for free without any additional implementation. New features added to the DAQ become automatically accessible to configuration and control.

Scala also natively supports the actor model, has extensive support for safe multi-core/multi-threaded programming and is fully interoperable with Java code, so that any code written for Java can be used directly. Since there is a wealth of Java software libraries freely available, this reduces the amount of code to be custom written significantly.

Since Scala is a comparatively young language, it is not as widely known as, for example, C++ and Java. Choosing such a language for a project of this kind should be done with care, since most new personnel working on the software will have to learn the basics of Scala. This is, however, not a serious problem, since Scala is easy to learn, being very similar to Java and C++ in syntax. Also this drawback is more than compensated by the rapid development possible with Scala. The basic design of the complete new DAQ system, ready for first usage at the LNGS setup, was done in about four man-months. After an additional four months of development, the system was already in use at several locations in the collaboration. In the author's experience, this would not have been possible using C++ or Java, certainly not when aiming for a clean and extensible design.

Furthermore, Scala is not an obscure language. Developed by Prof. Martin Odersky at EPF Lausanne (EPFL), author of the Java reference compiler, Scala has seen significant adoption in the industry (the Twitter backend systems for example are written in Scala [The08]). Scala is backed both by EPFL and significant EU-funding [The11b] (and recently, additional venture capital), ensuring long-term development and support.

For the actor model, the Scala Akka framework [Coma] was chosen, since it already implements supervision hierarchies [Comb] and has strong support for distributed operation.

5.4.2 Communication protocol support

DAQCorE supports a variety of communication protocols: Currently Actors for Raw TCP, Simple Network Management Protocol (SNMP), VXI-11, Standard Commands for Programmable Instruments (SCPI) and ROOT Message communication are implemented. Due to the modular nature of the software, other protocols can easily be added as needed.

5.4.3 Data Management - CouchDB

Storage of physics data, currently using the ROOT file format, was straightforward. However, not all data is suitable for sequential storage as provided by this format. Especially meta-data, configuration and conditions data (while it may in principle be stored in a sequential fashion) usually needs to be accessed in a random fashion, especially during data analysis.

For example, while sequentially processing the physics events of the individual data taking runs, the analysis configuration and calibration data for each current run has to be retrieved. These are necessarily stored separately from the run data, as they are usually not even available during data taking. Also, some manually configured data taking parameters will often not be stored with the run data. And lastly, during complex analysis, even SCADA data like temperatures and other environmental conditions may have to be taken into account. Since SCADA data taking is independent from the physics event data taking running on a different time scale, this information is not stored on a run-by-run or event-by-event basis and has to support random look-up by time or other query parameters.

Thus, a database-like storage solution was required for this kind of data. Using a standard relational database, however, requires considerable effort: Either all DAQ and analysis applications have to natively implement the database schema and then be modified on every change of this schema. Or a middleware application has to be developed which handles the database access and provides an application specific interface to the data. Also, the database must be constantly available to all sites in the collaboration. A disruption of the network connection to the LNGS underground lab (not uncommon) for example, would prevent

the DAQ at LNGS to write run meta-data and conditions data. Lastly, since COBRA is still a prototype experiment, the data requirements frequently change. A relational database, having a fixed data schema, is not well adapted to such a situation.

Thus, database system was needed which is schemaless, which supports fully distributed and replicated operation and which is freely available. After thorough investigation of the available products, CouchDB [Fou] was selected as the systems that fits all of these requirements well.

CouchDB stores JavaScript Object Notation (JSON) [Cro] documents with arbitrary structure. Simple Javascript procedures stored in the database allow for arbitrary indexing and lookup of these documents without requiring a fixed document schema. All communication with the database used the HTTP protocol, so no special software is required on the client side - the database can even be queried using standard command-line tools. The data can also be reformatted on the server side (again using procedures stored in the database). Thus the database output can be dynamically adapted any format already supported by the client applications, instead of adapting the client applications to the database.

Distributed operation is supported by CouchDB by design: Multiple replicas of the database can exist at different sites, providing offline read and write access to the complete data. The replicas can be synchronised either continuously or periodically, depending on the requirements of the individual sites.

For these reasons, CouchDB has become popular in applications of this type. The Data Management and Workflow Management (DWMW) system of the CMS experiment, for example, is also built on top of CouchDB [Cou].



Figure 5.5: CouchDB screenshot

5 COBRA DAQ System

The COBRA collaboration also uses this database now to manage and track components like detectors, ADCs and DAQ electronics. This integration of all different types of meta-data improves consistency and allows for automatic combination of data not possible before. For example, the masses of the detector crystals can now be directly queried for efficiency calculations during analysis.

6 Low-Background CPG Results

6.1 Pulse-Shape Readout under Low-Background Conditions

In August 2010, the old DAQ system for the CPG setup at LNGS was decommissioned and a prototype version of the new DAQ system (4.2) installed. The aim was to gain experience with low-background CPG detector pulse shape data to aid in the development of the final new CPG electronics and DAQ software. We installed 16 ADC channels to instrument the eight CPG detectors with colourless coating (2.4) on site. The prototype system consists of:

- 8 CPG CdZnTe detectors, 4 with the colourless eV-coating, 4 with cycloethene coating
- 1 modified preamplifier box (old design)
- 2 prototype 8-channel fast linear amplifiers (single ended input)
- 2 SIS3300 Fast ADCs with the SIS spectroscopy optimized firmware

Since the new differential fast amplifier and preamplifier carrier board were still in the design phase, a preamplifier box of the previous design was modified to extract the unmodified output signal of all CSA modules (see Figure 6.1).

Since the majority of the SIS3300 Fast ADC modules were not available until November 2010, the modules at LNGS could not be used for long-term operation until then, but were also needed for DAQ software development. Low-background physics data taking started in December 2010 and was resumed after the move of the setup in May 2011. It is planned to operate the setup in this configuration until the fall of 2011, when the new differential electronics chain will be ready for installation.

In total 6.7 kgd of data were collected above an energy threshold of 200 keV during the test phase of the new pulse-shape based DAQ from December 2010 to June 2011. At lower energies, the exposure is energy-dependent, as the detectors are operated at different threshold due to their individual noise level. The new data management system (5.4.3) makes it possible to automatically calculate the

6 Low-Background CPG Results




Figure 6.1: The modified version of the old preamplifier box. All active electronics apart from the preamplifier modules have been removed and additional cabling installed to route the preamplifier output signals out of the box. It was unavoidable to route several channels near the CSA modules themselves.




Figure 6.2: Exposure as a function of the energy deposition. The computation of the exposure takes the individual detector properties and operational parameters into account.

exposure at each energy resulting from the individual detector masses, livetimes and threshold settings.

The exposure was not corrected for the individual efficiency of the detectors (per mass), as they had not been determined in the past. A comparison of the number of ^{113}Cd events registered in the detectors shows that they do indeed have different background count rates even though their mass is almost the same (see Fig. 6.3). The next generation of detectors - scheduled for installation at LNGS in the fall of 2011 are currently characterised at TU Dresden, so that in the future the necessary efficiency information will be available.

6.2 Pathological Pulse Shapes

As expected, long-term operation under low-background conditions led to the observation of effects not seen during the short-term surface laboratory measurements during the DAQ hardware and software development.

One such effect are noise-spikes on the signal baseline (Figure 6.4). Signal quality had always been a problem at LNGS [Ree09] due to the long cable length of the single-ended CSA signal cables and the problematic design of the old EM-shielding. Since the frequencies involved are high compared to the shaping (resp. averaging) time scales during pulse shape processing, there is no danger of these spikes to be confused with real physics events at higher energies. Often, they do not even trigger the ADCs due to the shaping filter in the ADC trigger system (see 4.2.4). They do, however, result in some degradation of the energy resolu-




Figure 6.3: ^{113}Cd background activity of the individual detectors (see 1.3.1). The figure displays the number of background counts between 160 keV and 320 keV. Detectors 4, 5 and 6 show reduced efficiency.

6 Low-Background CPG Results

tion, as random coincidence with physics signals has been observed. The new EM-shielding, combined with the new differential signaling chain, is expected to significantly reduce or even eliminate the occurrence of these noise spikes.

As had already been well studied with the previous DAQ chain [Ree09], the noise level of the detectors was not constant. While the reason for this is still not clear, the occasional fluctuation of the noise level over the trigger threshold resulted in runs with an excessive number of events. As dead time free operation of the DAQ is not certain under these conditions, such runs were culled during analysis and are not included in the data shown here.

Another effect not observed in the laboratory is the occurrence of pulses with almost identical shape of the CA and NCA signal (see Figure 6.5). Such pulses indicate that equal charge has been transferred from the cathode to the anodes, with no charge transfer between the anodes. Due to the potential difference between the anodes, charges drifting in the crystal always lead to such a charge transfer. These events are therefore not physics events. A likely explanation are fluctuations or discharge effects of the cathode voltage. Fortunately, the signature of these events also leads to non-physical values for the interaction depth (Figure 6.7), making them easy to identify and exclude during analysis.

6.3 Background Spectrum

A study of the relationship between the energy deposition and the interaction depth yields very interesting results in general. Not only can the non-physical events described above be identified, but it is possible to identify and locate radioactive background sources (see Figure 6.6). Events of obviously non-physical




Figure 6.4: Baseline noise spikes at LNGS. The spikes are a regular feature on the signal baseline, as observed with an oscilloscope during commissioning, but seldom trigger the DAQ by themselves due to the filtered trigger - an exception with high amplitude is shown here.

nature, signified by depth values exceeding 2.0, are excluded from further analysis here.

The background stemming from radon in the air around the detectors is mainly observed as a peak around 5 MeV at the cathode. ^{222}Rn decays to ^{218}Po , which has a half-life of 3 min. The ionised ^{218}Po , when positively charged, is attracted by the cathode and is captured there. Also, the thickness of the protective coating on the detectors is beyond the range of the alpha particles stemming from the radon decay chain on all detector sides, except on the uncoated cathode. A simulation the background generated by a ^{218}Po deposition on the cathode by T. Köttig shows the same peak around 5 MeV (Fig. 6.9).

Another background contribution localized in depth is observed at the anode grid. As the energy values at this depth are subject to major distortion [HKWM97], determination of the origin of this background is beyond the scope of this work. However, events at such depths can easily be excluded during analysis.

This leads to the definition of two further depth-based cut criteria for physics analysis:

- $D < 0.04$: Anode grid events. Cut due to unknown background and distortion of the spectrum at this depth.
- $D > 0.94$: Events near the cathode. Cut due to domination of the radon-induced background.

The distribution of the interaction depth itself (Fig. 6.7) shows the same slope over the depth range observed during laboratory measurements (compare Fig. 4.24). As almost all events are due to the intrinsic ^{113}Cd beta-decay spectrum (Figure 6.10), homogeneously distributed within the detector crystals, this

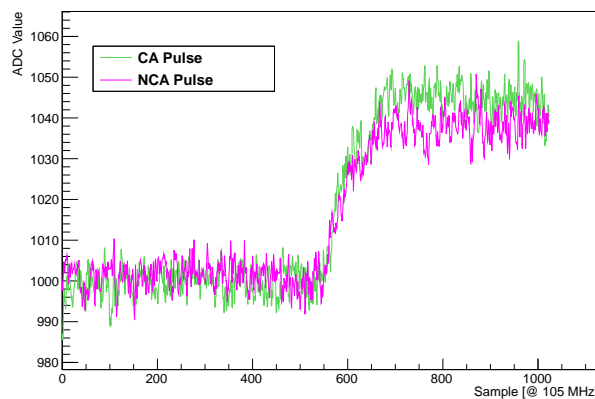


Figure 6.5: Non-physical event. A possible explanation for the identical shape of CA and NCA is a fluctuation of the HV Bias voltage.




Figure 6.6: Interaction-Depth dependence of the energy deposition. The plot shows three interesting features: At depths above 2.0, we observe events with a different spectrum than the main spectrum between -0.5 and 2.0 - such depths are non-physical - a manual examination of the anode pulses of these events showed them to be electrical disturbances of the type displayed in Figure 6.5. At a depth of 1 and energies around 5 MeV, these events, concentrated at the detector cathode, are a good match for alpha decays of isotopes from the ^{222}Rn decay chain (compare Fig. 6.9). Finally, there is an unidentified background contribution concentrated at the anode grids ($D = 0$).

is another indication that the current depth estimator is not linear between 0 and 1 and requires higher order correction.

Since the depth measure is not linear, the precise effect of the analysis cuts on the total efficiency cannot be determined precisely at this point. However, since the measure is not grossly nonlinear, the nonlinearity effect at the depths in question will be small compared to the precision of the results presented here and can be neglected.

The cuts result in an effective background reduction: While number of counts in the intrinsic ^{113}Cd is only reduced by 10% (Fig. 6.10), other background contributions are reduced significantly over the whole energy range, especially at higher energies (Fig. 6.8). The background is reduced by 55 % in the main region of interest (6.11). Since the number of events in this region is small, they were manually inspected for irregularities, but none were found. All nine events remaining after analysis cuts have well-formed anode pulse shapes, though, and must be considered to be valid physical events.




Figure 6.7: Distribution of the interaction depth at LNGS. The depth measure is not linear.

6 Low-Background CPG Results




Figure 6.8: Full LNGS background spectrum with pulse shape based readout before and after analysis cuts.




Figure 6.9: Simulation of the background generated by a ²¹⁸Po deposition on the CPG detector cathode.




Figure 6.10: The intrinsic ^{113}Cd low-energy background spectrum (see 1.3.1).




Figure 6.11: LNGS background spectrum with pulse shape based readout, region of interest for the $0\nu\beta\beta$ -decay of ^{116}Cd (see 1.3.1). Nine counts remain after analysis cuts.

6 *Low-Background CPG Results*

7 Summary and Outlook

This work has contributed improvements to several key aspects of the COBRA double-beta decay experiment.

A new data acquisition chain was developed for pulse-shape based readout of the COBRA coplanar grid (CPG) detectors. Prototype electronics for detector signal transmission and amplification were developed, as well as a software package with algorithms for pulse shape analysis of CPG detector signals. The pulse shape data readout has already resulted in significant improvements of the detector energy resolution. Pulse shape analysis has added interaction depth information to data output of the experiment, which has already shown itself to be an effective tool to suppress part of the background. A wavelet-based data compression technique has been developed to cope with the significant increase of data volume inherent to the recording of the full detector pulse shapes.

A flexible data acquisition software framework was developed along with the new DAQ hardware chain. It is designed for both research and development applications and long-term physics data collection and is now in production use at several sites in the COBRA collaboration. The framework was also designed to scale up to the requirements of a large-scale experiment and its modular nature ensures that future components can be integrated easily.

This work has contributed to the reduction of the radiation background level by over an order of magnitude, achieved by the combination of new CPG detector coatings and contacting methods with nitrogen flushing of the COBRA setup to create a radon-free environment.

For the first time, low-background physics data was taken with a CdZnTe pixel detector. The results show that pixel detectors present an exciting option for the future of the experiment.

Work is currently under way to scale up the new CPG electronics to a high channel count. The COBRA CPG setup at LNGS is scheduled to be upgraded to a higher number of detectors in 2011, which will all be run using the new DAQ system. A detailed study of the detector pulse shapes using simulations and collimated scans, currently in progress, will yield the necessary basis for advanced CPG pulse shape analysis. This will, in the future, enable new form of background suppression like separation between single-site and multi-site events.

7 Summary and Outlook

Acknowledgements

I owe thanks to many people who supported me in this work or made it even possible.

In the first place, I thank my thesis supervisor Prof. Dr. Claus Gößling for taking me on and entrusting me with the coordination of the COBRA activities in his group. He is also a really great boss. I am equally thankful to Prof. Dr. Kai Zuber, father of the COBRA experiment and speaker of the collaboration, a great physicist who always had the answer when I had a question. He kindly agreed to be my second examiner. I would also like to express my thanks to Dr. Bärbel Siegmann for taking part in my examination.

My old friend Dr. Daniel Münstermann roped me into this exciting project, knowing that I would love it. The former COBRAs Dr. Jaime Dawson, Dr. Jeanne Wilson, Dr. Christopher Reeve, Dr. John McGrath and Dr. Benjamin Janutta shared many a day (and late night) with me on shift deep under the Gran Sasso mountain at LNGS. Dr. Matthias Junker at LNGS made many things possible. Prof. Dr. Dr. Wolfgang Rhode, Dr. Ralf Wischnewski and the Amanda collaboration got us the Fast-ADCs without which essential parts of this work could not have been done. The students at the COBRA group in Dortmund have enriched the experiment so much with their enthusiasm: Thank you Markus, Katrin, Holger, Michael, Thomas, Julia, Magdalena and especially Jan.

I would also like to thank all the other great people at the Chair of Prof. Dr. Gössling, especially PD Dr. Rainer Klingenberg, Jennifer Jentzsch, Georg Troska, Silke Altenheiner, André Rummeler, Andrea Teichmann and my good friend Markus Alex.

I can hardly express my gratitude to my parents who always supported me on this path. My girlfriend, Birgit, I cannot thank enough for her love, patience and the countless hours she spent proofreading the manuscript.

I owe many thanks to Silke Rajek - I don't know how we would managed to assemble any detectors at all without her ingenuity and magic hands.

Finally, special thanks go to my colleagues and partners in crime Tobias Köttig and Till Neddermann. To work with people of your talent is a privilege - with you two it has also been a lot of fun. It is not without a twinge (or two) of regret that I now take my leave from the project, but I do so knowing that I leave things in the best of hands.

Acknowledgements

Publications

Articles

M. Schwenke, . . . O. Schulz et al.

Exploration of Pixelated detectors for double beta decay searches within the COBRA experiment

Nuclear Instruments and Methods in Physics Research A, in press (2011)

J. McGrath, . . . O. Schulz et al.

Detecting multi-hit events in a CdZnTe coplanar grid detector using pulse shape analysis: A method for improving background rejection in the COBRA $0\nu\beta\beta$ experiment

Nuclear Instruments and Methods in Physics Research A 615, 57-61 (2010)

J.V. Dawson, . . . O. Schulz et al.

Experimental study of double- β decay modes using a CdZnTe detector array

Physical Review C 80, 025502 (2009)

J.V. Dawson, . . . O. Schulz et al.

An investigation into the ^{113}Cd beta decay spectrum using a CdZnTe array

Nuclear Physics A 818, 264-278 (2009)

Conference Talks

O. Schulz on behalf of the COBRA collaboration

Pulsform-Analyse von CdZnTe CPG Detektoren am COBRA-Experiment

Particle Physics DPG Spring Meeting

March 28 – April 1 2011, Karlsruhe (Germany)

O. Schulz on behalf of the COBRA collaboration

Status des COBRA Experiments

Particle Physics DPG Spring Meeting

March 15 – 19 2010, Bonn (Germany)

O. Schulz on behalf of the COBRA collaboration

The COBRA double beta decay experiment

2nd LSM Extension Workshop

Publications

October 16 2009, Modane (France)

O. Schulz on behalf of the COBRA collaboration

Searching for Ultra-Rare Decays with CdZnTe Detectors

**IEEE 16th International Workshop on Room Temperature Semiconductor X-
and Gamma-Ray Detectors (RTSD)**

October 19 – 25 2008, Dresden (Germany)

Bibliography

- [A⁺06a] A. Monfardini et al. The microcalorimeter arrays for a rhenium experiment (mare): A next-generation calorimetric neutrino mass experiment. *Nuclear Instruments and Methods in Physics Research Section A*, 559(2):346–348, 2006.
- [A⁺06b] K. Asanovic et al. The Landscape of Parallel Computing Research: A View from Berkeley. Technical Report UCB/EECS-2006-183, Electrical Engineering and Computer Sciences, University of California at Berkeley, Dec. 2006.
- [A⁺08] S.V. Adve et al. Parallel Computing Research at Illinois: The UPCRC Agenda. White Paper, Nov. 2008.
- [A⁺09] I. Antcheva et al. ROOT - A C++ framework for petabyte data storage, statistical analysis and visualization. *Computer Physics Communications*, 12:2499–2512, 2009.
- [ACK⁺08] I. Abt, A. Caldwell, K. Kröninger, J. Liu, X. Liu, and B. Majorovits. Test of pulse shape analysis using single compton scattering events. *The European Physical Journal C*, 54:425–433, 2008.
- [Agh85] G.A. Agha. *Actors: A Model Of Concurrent Computation In Distributed Systems*. PhD thesis, Massachusetts Institute of Technology (MIT), Artificial Intelligence Laboratory, 1985.
- [Ale09] M. Alex. A microcontroller based solution for remote laboratory data acquisition and control. Diploma thesis, TU Dortmund, 2009.
- [Arm07] J. Armstrong. *Programming Erlang. Software for the Concurrent World*. Pragmatic Bookshelf, 2007.
- [AVWW96] J. Armstrong, R. Virding, C. Wilkstrom, and M. Williams. *Concurrent Programming in Erlang*. Prentice Hall, 1996.
- [BEB95] H. H. Barrett, J. D. Eskin, and H. B. Barber. Charge Transport in Arrays of Semiconductor Gamma-Ray Detectors. *Phys. Rev. Lett.*, 75(1):156–159, Jul 1995. doi:10.1103/PhysRevLett.75.156.

Bibliography

- [BV87] F. Boehm and P. Vogel. *Physics of Massive Neutrinos*. Cambridge University Press, 1987.
- [Cam11] M. Campbell. Personal communication. 2011.
- [CDF92] A. Cohen, I. Daubechies, and J. Feauveau. Bi-orthogonal bases of compactly supported wavelets. *Comm. Pure Appl. Math.*, 45:485–560, 1992.
- [Coma] Akka Developers Community. Akka. <http://akka.io>.
- [Comb] Akka Developers Community. Fault Tolerance Through Supervisor Hierarchies (Scala). <http://akka.io/docs/akka/1.1.1/scala/fault-tolerance.html>.
- [Cou] Couchbase, Inc. CERN – CouchDB and Physics. <http://www.couchbase.com/case-studies/cern>.
- [Cro] D. Crockford. Request for Comments (RFC): The application/json Media Type for JavaScript Object Notation (JSON). <http://tools.ietf.org/html/rfc4627>.
- [D⁺10] C. Disch et al. Stacked coplanar grid detectors. Prague, March 2010. COBRA Collaboration Meeting.
- [Dau92] Ingrid Daubechies. *Ten Lectures on Wavelets*. Vol. 61 of CBMS-NSF Regional Conference Series on Applied Mathematics. SIAM, Philadelphia, 1992.
- [eM] eV Microelectronics. Semiconductor Detector Material Properties. URL: http://www.evmicroelectronics.com/pdf/material_prop.pdf [accessed 2011-08-09].
- [Far08] K. Farnham. Interesting multicore crisis graph and analysis, 2008. URL: <http://software.intel.com/en-us/blogs/2008/01/17/interesting-multicore-crisis-graph-and-analysis/> [accessed 2011-08-09].
- [Fas] Fast ComTech GmbH, Oberhaching, Germany. *CR-110 charge sensitive preamplifier*.
- [Fel98] G.J. Feldman. Unified approach to the classical statistical analysis of small signals. *Physical Review D*, 57(7):3873–3889, 1998.
- [Flo08] Florida RF Labs, Inc. *Cable Catalog Volume 7*, 2008.

- [Fou] Apache Foundation. The Apache CouchDB Project. <http://couchdb.apache.org>.
- [GAD⁺11] T. Gleixner, G. Anton, J. Durst, M. Filipenko, and T. Michel. Simulations for design decisions. Hamburg, June 2011. COBRA Collaboration Meeting.
- [GGL94] A. Georgiev, W. Gast, and R.M. Lieder. An analog-to-digital conversion based on a moving window deconvolution. *IEEE Transactions on Nuclear Sciences*, 41(4):1116–1124, 1994.
- [GJK⁺05] C. Goessling, M. Junker, H. Kiel, D. Muenstermann, S. Oehl, and K. Zuber. Experimental study of ^{113}Cd β decay using CdZnTe detectors. *Physical Review C*, 72:064328, 2005.
- [Gra93] Robert M. Gray. Dithered quantizers. *IEEE TRANSACTIONS ON INFORMATION THEORY*, 39(3):805–812, May 1993.
- [Gra95] Amara Graps. An introduction to wavelets. *IEEE Computational Science & Engineering*, 2(2):50–61, 1995.
- [Hay10] W.M. Haynes, editor. *CRC Handbook of Chemistry and Physics*. CRC Press, 91st edition, 2010.
- [Heg06] G. Hegyesi. Ethernet Based Distributed Data Acquisition System for a Small Animal PET. *IEEE Transactions on Nuclear Science*, 53(4):2112–2117, 2006.
- [HKW⁺96] Z. He, G.F. Knoll, D.K. Wehe, R. Rojeski, C.H. Mastrangelo, M. Hammig, C. Barrett, and A. Uritani. 1-D position sensitive single carrier semiconductor detectors. *Nuclear Instruments and Methods in Physics Research A*, 380:228–231, 1996.
- [HKWM97] Z. He, G.F. Knoll, D.K. Wehe, and J. Miyamoto. Positive-sensitive single carrier CdZnTe detectors. *Nuclear Instruments and Methods in Physics Research A*, 388:180–185, 1997.
- [J⁺10] J. McGrath et al. Detecting multi-hit events in a cdznte coplanar grid detector using pulse shape analysis: A method for improving background rejection in the cobra $0\nu\beta\beta$ experiment. *Nuclear Instruments and Methods in Physics Research A*, 615:57–61, 2010.
- [K⁺05] Ch. Kraus et al. Final results from phase II of the Mainz neutrino mass search in tritium beta decay. *Eur. Phys. J.*, C40:447–468, 2005. arXiv:hep-ex/0412056.

Bibliography

- [Kie05] H. Kiel. *Determination of the Half Lives of Rare Decays of Cd, Te and Zn Isotopes for the COBRA Experiment*. PhD thesis, University of Dortmund, Germany, 2005. URL: <http://hdl.handle.net/2003/21509>.
- [KK⁺01] H. V. Klapdor-Kleingrothaus et al. Latest results from the Heidelberg-Moscow double-beta-decay experiment. *The European Physical Journal A*, 12:147, 2001.
- [KSA09] R.K. Karmani, A. Shali, and G. Agha. Actor frameworks for the JVM platform: a comparative analysis. In *Proceedings of the 7th International Conference on Principles and Practice of Programming in Java*, 2009.
- [Kö08] T. Köttig. Optimisation of Contacting and Energy Resolution of CdZnTe with TCT. Diploma thesis, TU Dortmund, 2008.
- [Leo93] W.R. Leo. *Techniques for Nuclear and Particle Physics Experiments*. Springer Verlag, 2 edition, 1993.
- [LMG86] D.A. Landis, N.W. Madden, and F.S. Goulding. Double Beta, Liquid Nitrogen, Fill Level. *IEEE Transactions on Nuclear Science*, 33(1):399–402, 1986.
- [LN03] J.P. Lewis and U. Neumann. Performance of Java versus C++. <http://scribblethink.org/Computer/javaCbenchmark.html>, Jan 2003.
- [Luk95] P. N. Luke. Unipolar charge sensing with coplanar electrodes - application to semiconductor detectors. *IEEE Transactions on Nuclear Science*, 42:207–213, 1995.
- [Mal99] S. Mallat. *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [McG09] J. McGrath. *Research and development of CdZnTe semiconductor detectors in a low background environment for the COBRA neutrinoless double beta decay experiment*. PhD thesis, University of York, 2009.
- [MN62] Z. Maki and S. Nakagawa, M. Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 28:870, 1962.
- [Mü07] D. Münstermann. *Construction of a Low Background Facility for the Cobra Experiment*. PhD thesis, TU Dortmund, 2007.
- [OSV08] M. Odersky, L. Spoon, and B. Venners. *Programming in Scala*. Artima Press, 2008.
- [Pat10] D. Patterson. The trouble with multi-core. *IEEE Spectrum*, 47(7):28–32, 53, 2010.

- [Pot99] D. Potter. Using Ethernet for industrial I/O and data acquisition. In *Proceedings of the 16th IEEE Instrumentation and Measurement Technology Conference (IMTC)*, pages 1492–1496, May 1999.
- [Qua10] T. Quante. Auswertung der pulsform von coplanar-grid-cdznte dektoren. Bachelor thesis, TU Dortmund, 2010.
- [Rac37] G Racah. Sulla Simmetria Tra Particelle e Antiparticelle. *Nuovo Cimento*, 14:322–328, 1937.
- [Raj] S. Rajek. *Thesis - In preparation*. PhD thesis, TU Dortmund.
- [Ree09] C. Reeve. *Data Analysis and Background Studies for the COBRA Neutrinoless Double-Beta Decay Experiment*. PhD thesis, University of Sussex, Brighton, UK, 2009.
- [Rei95] L. Reissell. Multiresolution and wavelets. In Alain Fournier, editor, *SIGGRAPH '95 Course Notes - Wavelets and their Applications in Computer Graphics*, pages 37–69. ACM SIGGRAPH, 1995.
- [SBW01] H. Sthamer, A. Baresel, and J. Wegener. Evolutionary Testing of Embedded Systems. In *Proceedings of the 14th International Internet & Software Quality Week*, pages 1–34, 2001.
- [Sch] Oliver Schulz. Blitzwave c++ wavelet library. URL: <http://blitzwave.sourceforge.net> [accessed 2011-08-05].
- [Sch09] K. Schreiner. Evaluation of Mounting and Passivation Techniques for Low-Background CdZnTe Detectors. Diploma thesis, TU Dortmund, 2009.
- [SIS] SIS GmbH, Rehlingen, Germany. *SIS3300/SIS3301 65/100 MHz VME FADCs*.
- [SSGG96] J. Stein, F. Scheuer, W. Gast, and A. Georgiev. X-ray detectors with digitized preamplifiers. *Nuclear Instruments and Methods in Physics Research B*, 113:141–145, 1996.
- [Swe95] W. Sweldens. The lifting scheme: A new philosophy in biorthogonal wavelet constructions. In A. F. Laine and M. Unser, editors, *Wavelet Applications in Signal and Image Processing III*, pages 68–79. Proc. SPIE 2569, 1995.
- [Teb11] Jan Tebrügge. In preparation. Master's thesis, TU Dortmund, 2011.

Bibliography

- [The98] The Super-Kamiokande Collaboration: Y. Fukuda et al. Measurements of the Solar Neutrino Flux from Super-Kamiokande's First 300 Days. *Physical Review Letters*, 81:1158–1162, August 1998. arXiv:arXiv:hep-ex/9805021, doi:10.1103/PhysRevLett.81.1158.
- [The01] The KATRIN collaboration. KATRIN: A next generation tritium beta decay experiment with sub-eV sensitivity for the electron neutrino mass, 2001. arXiv:hep-ex/0109033.
- [The08] The Scala Development Team. Twitter message queues move to Scala. <http://www.scala-lang.org/node/1008>, Jul 2008.
- [The11a] The CERN ROOT Team. ROOT-Homepage, Aug. 2011. URL: <http://root.cern.ch> [accessed 2011-08-07].
- [The11b] The Scala Development Team. Scala Team Wins ERC Grant. <http://www.scala-lang.org/node/8579>, Jan 2011.
- [Tyc07] Tyco Electronics. *AMP CO Plus System Overview & Design Application Guide*, 2007.
- [UB03] M. Unser and T. Blu. Mathematical properties of the jpeg2000 wavelet filters. *IEEE Transactions on Image Processing*, 12(9):1080–1090, 2003.
- [URB97] G. Uytterhoeven, D. Roose, and A. Bultheel. Wavelet Transforms Using The Lifting Scheme. Technical Report ITA-Wavelets-WP1.1, Kath. Univ. Leuven, Dept. C.S., 1997.
- [Wag04] W. Wagner. *Design and Realisation of a new AMANDA Data Acquisition System with Transient Waveform Recorders*. PhD thesis, TU Dortmund, 2004.
- [Wat01] John Watkinson. *The art of digital audio*. Focal Press, 2001.
- [WHWZ08] W. Wang, Z. He, C.G. Wahl, and F. Zhang. Detecting shielded sources using 3-D CdZnTe detectors. In *IEEE Nuclear Science Symposium Conference Record*, 2008.
- [Wil05] J. Wilson. Predicted Sensitivities for COBRA using Maximum Likelihood. Technical report, University of Sussex, Brighton, UK, 2005. COBRA memo 2005-06.
- [Wol08] A. Wolfe. Intel blog warns of multicore crisis, 2008. URL: http://www.informationweek.com/blog/main/archives/2008/01/intel_blog_warn.html [accessed 2011-08-07].

- [ZH06] Feng Zhang and Zhong He. New Readout Electronics for 3-D Position Sensitive CdZnTe/HgI₂ Detector Arrays. *IEEE Transactions on Nuclear Science*, 53(5):3021–3027, Oct. 2006. doi:10.1109/TNS.2006.879761.
- [Zha09] F. Zhang. Personal communication, 2009.
- [ZHS07] F. Zhang, Z. He, and C.E. Seifert. A Prototype Three-Dimensional Position Sensitive CdZnTe Detector Array. *IEEE Transactions on Nuclear Science*, 54(4):843–848, Aug. 2007. doi:10.1109/TNS.2007.902354.
- [Zub01] K. Zuber. COBRA—double beta decay searches using CdTe detectors. *Physics Letters B*, 519(1-2):1–7, 2001.
- [Zub04] K. Zuber. *Neutrino Physics*. Institute of Physics Publishing, 2004. ISBN 0-7503-0750-1.
- [Zub10] K. Zuber. The status of the COBRA double-beta-decay experiment. *Progress in Particle and Nuclear Physics*, 64(2):267–269, 2010.

Bibliography

List of Figures

1.1	Normal and inverted neutrino mass hierarchies	10
1.2	Isotope mass parabola	10
1.3	Neutrinoless double-beta decay schema	11
1.4	CPG detector	15
1.5	Weighting potential for CPG detectors.	15
2.1	CAD detector position alignment	19
2.2	Detector holder with spring/pin contacting	20
2.3	LS200 conductive glue detector contacts	21
2.4	Reworked radon trap	24
2.5	Schematic of the Dewar	25
2.6	Switched-capacitor capacitance sensing circuit	25
2.7	Instrumentation of PT sensors	25
2.8	Fill level monitoring of the Dewar	26
2.9	Background events before and after activation of nitrogen flushing	26
2.10	Liquid nitrogen Dewar with monitoring and control	27
2.11	Two test layers with low-background crystal coatings	28
2.12	Background reduction achieved at LNGS	29
3.1	Polaris detector, pixel structure	32
3.2	Polaris detector module	33
3.3	Polaris pixel detector setup	34
3.4	Polaris pixel detector cooling	34
3.5	Polaris detector, LNGS background spectrum	35
3.6	Polaris detector, hit distribution	36
3.7	LNGS 228-Th calibration spectrum	36
3.8	$0\nu\beta\beta$ hit distribution simulation of 116-Cd under analysis cuts . . .	37
3.9	$0\nu\beta\beta$ spectrum simulation of 116-Cd under analysis cuts	38
3.10	Background spectrum before and after analysis cuts	38
4.1	Previous COBRA CPG DAQ Chain	40
4.2	New COBRA CPG Chain	41
4.3	Schematic of the prototype preamplifier box	42
4.4	The prototype preamplifier box	43
4.5	Schematic of the single-ended fast linear amplifier	44

List of Figures

4.6	Custom built 8-channel fast amplifier	44
4.7	Schematic of the voltage-regulator module	46
4.8	Schematic of the single-ended to differential circuit	46
4.9	Schematic of the differential to single-ended conversion circuit	47
4.10	The differential amplifier prototype	47
4.11	Struck SIS3300 Fast-ADC VME Module	48
4.12	Digital Signal Processing chain	51
4.13	Anode pulse shapes before calibration and deconvolution	51
4.14	Anatomy of CA and NCA pulse shapes of an event in the main bulk	52
4.15	Anatomy of CA and NCA pulse shapes of an event near the anode grid	53
4.16	Anatomy of CA and NCA pulse shapes of an event near the cathode	53
4.17	Deconvolution of preamplifier signal	55
4.18	Test of the digital pulse height estimation	59
4.19	Distribution of the pre-pulse baseline level	60
4.20	Distribution of the post-pulse baseline slope	60
4.21	Relation between slope of the initial rise of the NCA and CA signal	61
4.22	Multi-point calibration of a CPG Detector with a Th-228 source	62
4.23	Cs-137 Spectrum of a spectrometer-grade evProducts CPG detector	63
4.24	Distribution of the interaction depth in a CPG detector	64
4.25	Haar wavelet functions	68
4.26	CDF(2,2) Wavelet functions	68
4.27	Recursive Wavelet Composition	69
4.28	CA pulse shape before and after compression by a factor of 10	70
4.29	NCA pulse shape before and after compression by a factor of 10	71
4.30	Preservation of rise-time and timing information under wavelet compression	71
4.31	Error in energy deposition values after wavelet compression	72
4.32	Distribution of the error in energy ΔE induced by Wavelet compression	72
4.33	Effect of the wavelet compression on the energy spectrum	73
5.1	Previous monolithic COBRA DAQ structure	75
5.2	New, modular COBRA DAQ structure	78
5.3	Actor supervision principle	80
5.4	Supervision hierarchies in the actor model	80
5.5	Screenshot of the CouchDB database	85
6.1	Modified version of the old preamplifier box	88
6.2	Energy-dependent exposure at LNGS	88
6.3	Cd-113 Background activity in the individual detectors	89
6.4	Baseline noise spikes at LNGS	90
6.5	91

6.6	Interaction-Depth dependence of the energy deposition	92
6.7	Distribution of the interaction depth at LNGS	93
6.8	LNGS background spectrum with pulse shape based readout . . .	94
6.9	Simulaton of Po-218 on cathode	94
6.10	Intrinsic 113-Cd background spectrum with pulse shape based readout	95
6.11	Cd-116 background spectrum with pulse shape based readout . . .	95

List of Figures

Acronyms

$0\nu\beta\beta$ -decay neutrinoless double-beta decay. 7, 9–12, 26

$2\nu\beta\beta$ -decay 2-neutrino double-beta decay. 9, 10

ADC analog to digital converter. 29, 38, 41, 43, 46–48, 50, 52, 53, 55–57, 79, 80, 83, 85, 87

ASIC application-specific integrated circuit. 29, 31, 32

CA collecting anode. 12, 13, 48, 50, 52–59, 69, 88, 89

CPG coplanar grid. 12–14, 26, 29, 32, 33, 35, 37, 63, 68, 74, 85, 92

CSA charge-sensitive amplifier. 37, 39, 46, 57, 58, 85, 87

DAQ data acquisition. 31, 38, 39, 41, 43, 45, 47, 48, 52, 57, 58, 61, 73–75, 77, 79–83, 85, 87, 88

DAQCorE Data-Acquisition and Control Environment. 81

DMMW Data Management and Workflow Management. 83

EPFL EPF Lausanne. 81

Fast-ADC fast analog to digital converter. 38, 43, 46–48, 80

FMF Freiburger Materialforschungszentrum. 18, 25

HV high voltage. 16

IC integrated circuit. 42, 44, 45, 48

JSON JavaScript Object Notation. 82

JTAG Joint Test Action Group. 48

JVM Java Virtual Machine. 79, 81

Acronyms

LNGS Laboratori Nazionali del Gran Sasso. 14–17, 19, 20, 24–26, 30–32, 34, 39, 41, 47, 74, 80–82, 87

MANtiCORE Multiple-Analysis Toolkit for the COBRA Experiment. 63

MCA multi channel analyzer. 37

MWD Moving Window Deconvolution. 52, 53

NCA non-collecting anode. 13, 48, 50, 52–59, 88, 89

NIM Nuclear Instrumentation Module. 37, 41

POE Power-over-Ethernet. 21

SCADA supervisory control and data acquisition. 73, 74, 80, 82

SCPI Standard Commands for Programmable Instruments. 82

SNMP Simple Network Management Protocol. 82