A Silicon Photomultiplier Based Readout System For A Cosmic Muon Telescope; Design And Implementation

Magne Eik Lauritzen

A thesis presented for the degree of
Master of Science

Department of Physics and Technology
University of Bergen
Norway
June 2017

A Silicon Photomultiplier Based Readout System For A Cosmic Muon Telescope, Design And Implementation

Magne Eik Lauritzen

Abstract

The Cosmic Ray Telescope (CRT) is a $4\,m^2$ scintillator detector segmented into 16 slabs with dimensions of $25 \times 100\,cm^2$. The detector is located at the Department of Physics, University of Bergen, and is intended primarily for educational purposes. As part of this project, the original photomultiplier tubes have been replaced with silicon photomultipliers (SiPMs), as they offered greater benefits than PMTs for this project.

An entirely new SiPM readout system has been developed, enabling low muon crossing timing uncertainty ($0.6\,ns$). The readout system runs custom software on an Arduino microcontroller, and any number of these systems can be linked together and controlled from a computer. The design of this readout system is presented in detail, and its implementation on the CRT is described.

In addition to the development of the readout system, an application was developed for the purposes of controlling the detector and performing on-line analysis of the data it generates in order to locate coincidences between scintillator slabs. Compared to traditional analogue coincidence schemes this offered greater flexibility and cost-reduction. The algorithm and capabilities of the control software are presented.

Acknowledgements

Glossary

ROB Readout Board. 61
RTD Resistive Temperature Detector. 96

SiPM Silicon Photomultiplier. 13
SM Standard Model of particle physics. 15
SNR Signal to Noise Ratio. 58
SPICE Simulation Program with Integrated Circuit Emphasis. 57
STX Start Of Transmission. 105

TCR Temperature Coefficient of Resistance. 91
TINA Toolkit for Interactive Network Analysis. 57
TOF Time Of Flight. 52
TTL Transistor-Transistor Logic. 84

VI Virtual Instrument. 116

WALTA Washington Area Large-scale Time-coincidence Array. 24

Contents

1 | Introduction

This thesis details the work performed in an attempt to resurrect and improve a muon telescope called the Cosmic Ray Telescope (CRT), nicknamed "Lille My" (that's "Little Mu" for those of you that are, understandably, not fluent in Norwegian). The CRT consists of Bicron BC-408 organic scintillators [1], and is able to detect charged particles, fast neutrons, and gamma rays in the energy range 100 keV to 5 MeV. It has a 4 m^2 active area divided on 16 channels, and is located on the premises of the Department of Physics and Technology (IFT) at the University of Bergen in Norway. An illustration of the CRT is seen in Figure 1.1. The CRT has sadly been out of commission ever since the vital Photomultiplier Tubes (PMTs) began failing one after another. The true cause of this was never pinned down, but the PMTs were old, noisy, and had large dark currents. The Particle Physics group at IFT has long wished for its return to operation, and so I took it upon myself to see it done. The PMTs were replaced with modern Silicon Photomultipliers (SiPMs), and a robust control and readout system was designed and implemented.
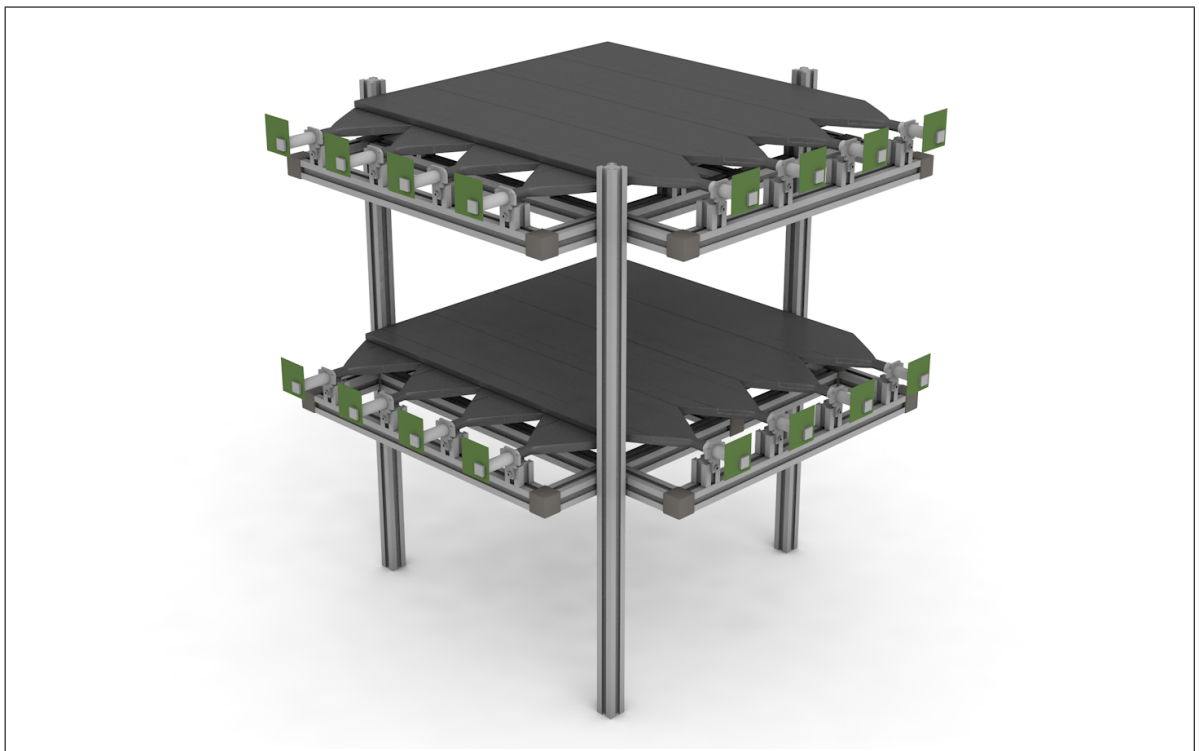
Figure 1.1: A simplified rendition of the CRT. Certain components, such as cables and unnecessary structures, have been omitted.

The CRT was designed, assembled and shown to function (for a while) by my predecessor Øyvind Sætre [2]. I was therefore able to focus my resources and time on redesigning and replacing the malfunctioning parts, namely the photon detectors themselves and their related readout system. The primary focus of this thesis is therefore the design and development of the new readout electronics and the data acquisition system. Additionally, a chapter has been dedicated to the scintillator slabs and light guide configuration, which was analysed with Monte Carlo (MC) simulations to better understand their behaviour and to optimize the light guides. This thesis is also meant as a users guide to whomever wishes to use or modify the CRT. Throughout the design process, decisions were made to ensure that the detector will be user friendly and easy to maintain and upgrade.

The primary function of the CRT is to aid in the education of physics undergraduates at IFT by incorporating it into courses. For many students it would serve as their first hands-on experience with a particle detector, and analysis of the data it produces would introduce them to many of the same concepts in data analysis and statistics that are used in large particle physics experiments. To enable undergraduate students to easily operate the CRT, the control software was made user friendly, but also flexible and powerful enough so that the operator can adjust every aspect of the detector and perform long periods of uninterrupted data taking.

The thesis is divided into 7 chapters. The first chapter will serve as an introduction to cosmic rays and muons. Chapter 2 details simulations performed in order to investigate the characteristics of the photon flux impeding on the SiPM detector. Chapters 3, 4, and 5 discusses the SiPM, the electronics designed for reading its signal, and the data acquisition & control system for the CRT. In Chapter 6 data taken with the CRT is presented, and used among other things to verify the CRT performance and estimate the muon lifetime. Finally, Chapter 7 concludes the thesis and discusses the future of the CRT.

1.1 Background

The intended function of Little Mu is to detect muons (denoted μ, pronounced *mu* - hence its nickname). These elementary particles are constantly generated when highly energetic protons and nuclei from space called *cosmic rays* impacts the upper atmosphere. This generates a far-reaching hadronic and electromagnetic cascade that is called an air shower. This section will go into the details of muons, cosmic rays, and how the generation occurs.

The statements presented in this section are primarily cited from the *Review of Particle Physics*, Particle Data Group 2016 [3], in particular the chapter on cosmic rays and muons, pages 421-428.

1.1.1 The Standard Model, decays, and particle interaction with matter

The Standard Model of particle physics (SM) describes the elementary forces of the universe (except gravity) and classifies the elementary particles. A schematic representation of the SM is seen in Figure 1.2. All matter that is encountered in daily life consists of fermions from the first generation: The electron, and the up and down quark. The other fermions are unstable (with exception of the neutrinos) and will quickly decay, and are normally only produced in energetic particle collisions.

An unstable elementary particle or ensemble of quarks (called a hadron) will decay, and for a certain particle there are often several decay modes. In the production of muons by cosmic rays, only a handful different types of particles are created, and some of these are unstable. Table 1.1 lists the most commonly produced particles in air showers, and their primary decay modes.

Particle	Decay modes	Branching ratio
π^+	$\mu^+ + \nu_\mu$	~ 1
π^-	$\mu^- + \bar{\nu}_\mu$	~ 1
π^0	2γ	0.988
	$\gamma + e^- + e^+$	0.0117
μ^+	$e^+ + \nu_e + \bar{\nu}_\mu$	~ 1
μ^-	$e^- + \bar{\nu}_e + \nu_\mu$	~ 1

Table 1.1: Decay modes of primary air shower constituents

Energetic particles can interact with matter in several ways, losing energy in the process. An air shower consists primarily of charged particles, and for those the methods of interaction are:

- Atom ionization
- Atom excitation
- Bremsstrahlung
- Cherenkov radiation
- Transition radiation

Electrons and positrons constitutes a large fraction of the particles in an air shower. These light charged particles will primarily lose energy through Bremsstrahlung, due to their low mass. Bremsstrahlung is electromagnetic radiation (photons) emitted by the deceleration of charged particles. It is emitted by electrons and positrons as they are deflected from their paths by the Coloumb forces between it and nuclei. The Bremsstrahlung photons may go on to create new electrons and positrons if their energy is great enough. This effect is called pair production and will be explained later.

Another prevalent charged particle in air showers is the muon. It has a much greater mass than the electron, and will not lose much energy via Bremsstrahlung. Its primary energy loss mechanism is atom excitation and ionization. As a muon passes by an atom, an energy transfer occurs between it and the electrons bound to the nucleus, via the Coloumb force. The electron will jump to a higher orbital, and the atom becomes excited. If enough energy is transferred to the electron to overcome its binding energy, it exits its bound state, and the atom becomes ionized.

A large amount of photons are generated in an air shower. They may lose energy through three primary interactions with matter:

- Photoelectric effect $\quad (\gamma + {}^A_Z X \rightarrow {}^A_Z X^+ + e^-)$
- Compton scattering $\quad (\gamma + {}^A_Z X \rightarrow \gamma + {}^A_Z X^+ + e^-)$
- Pair production $\quad (\gamma + {}^A_Z X \rightarrow {}^A_Z X + e^- + e^+)$

These interaction types dominate at different photon energy levels. For low photon energy ($<1\,\mathrm{MeV}$) the photoelectric effect is the primary method of photon interaction. At intermediate energies Compton scattering dominates, and for large energies ($>5\,\mathrm{MeV}$) pair production is dominant. The initial photons generated in an air shower are highly energetic, and thus interact with matter primarily through pair production. In this form of interaction, a photon that passes close by a nucleus may decay into an electron-positron pair, and the nucleus receives some recoil. The electron and positron may go on to produce new photons through the aforementioned Bremsstrahlung effect, and the process repeats.
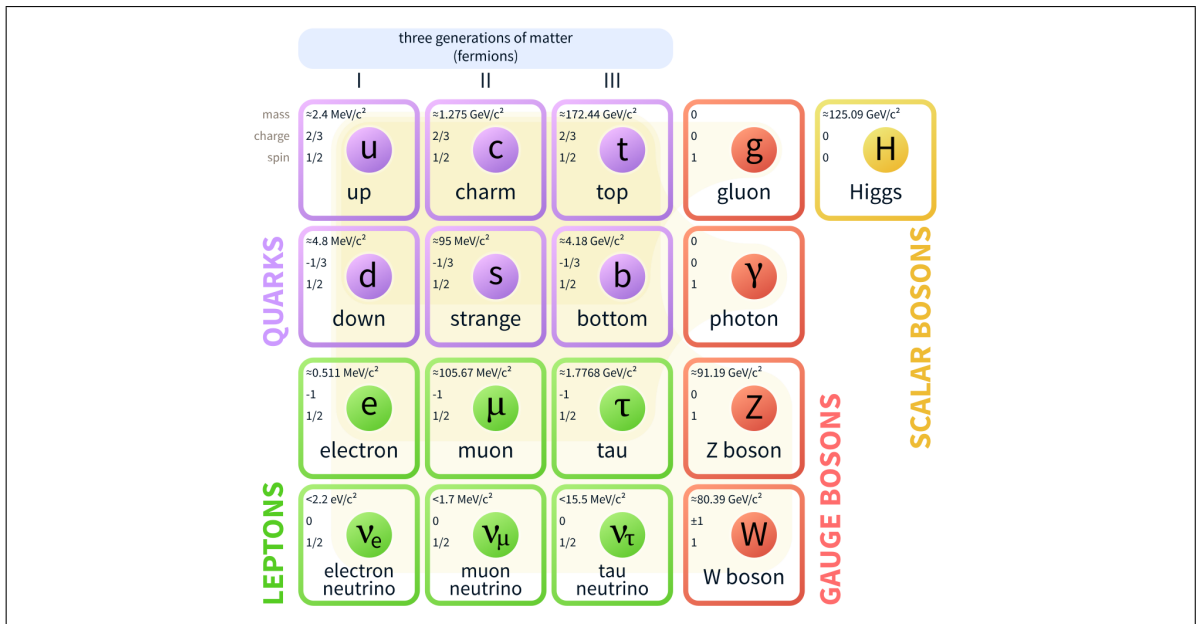
Figure 1.2: A schematic representation of the Standard Model of particle physics. The elementary particles (fermions), force-carriers (gauge bosons), and the Higgs boson are listed.

1.1.2 Muons

Muons (μ^-, antiparticle: μ^+) are elementary particles, very much reminiscent of the ubiquitous electron. Two differences that separate them are its mass, which is some 200 times that of the electron (105 MeV vs. the electrons measly 510 keV), and that it is unstable with a lifetime of 2.2 μs. The muon was discovered in 1937 by C.W. Anderson and S.H. Neddermeyer by exposing a cloud chamber to the by-products of cosmic rays [4]. Muons generated by cosmic rays make up a large part of the natural background radiation, and are the most common free charged particle at sea level.

Due to the large mass of the muon, its long lifetime (by subatomic standards), and the fact that it does not interact with matter through the strong nuclear force, muons of relatively low energies are capable of traversing large amounts of matter. This makes them useful for a range of applications, which are discussed in Section 1.2.

1.1.3 Cosmic rays

Cosmic rays are protons and nuclei constantly impinging on the top of the Earths atmosphere. The name is a misnomer, and was coined by Robert Millikan in the 1920s [5]. He performed measurements on the discharge of electroscopes due to atmospheric radiation at various absorption depths in the atmosphere and under water, and showed that highly penetrating particles were being generated at the top of the atmosphere. He believed they were generated by incoming radiation from space, and believed that these "cosmic rays" consisted of gamma rays emitted by the formation of atoms, hence the name. Later, this hypothesis was shown to be incorrect, and it became clear that cosmic "rays" were in fact particles. The primary

constituents of cosmic rays are free protons (79 %) and helium nuclei (14.7 %). Heavier nuclei (Z>3) constitute 5 %, and electrons 1 %. Cosmic rays originate from outside the solar system, apart from low-energy protons, electrons, and helium nuclei which are constantly emitted by the Sun in the form of solar wind and coronal mass ejections. Pinning down the nature and location of the sources of cosmic rays has been difficult, as the charged cosmic rays are deflected from their path by the galactic magnetic field that exists in the space between stars, and by the terrestrial geomagnetic field that surrounds the Earth. Many potential sources of cosmic rays have been proposed, and include supernovae, pulsars, active galactic nuclei (supermassive black holes), quasars, gamma ray bursts, and other energetic stellar phenomena. Locating the sources of cosmic rays is currently an active field of research.

The energy spectrum of cosmic rays impinging on the atmosphere has been measured to extend beyond 1×10^{20} eV or 1×10^{11} GeV, but cosmic rays of this energy are exceedingly rare. The Pierre Auger Observatory, which surveys a 3000 km^2 area in Argentina, detected only a mere 27 cosmic rays of energies greater than 5.7×10^{19} eV during its initial run between 2004 and 2007 [6]. The peak in the energy spectrum of cosmic rays is located at $\sim$1 GeV/nucleon, and decreases rapidly with increasing nucleon energy, following a power law. Lower-energy cosmic rays are excluded partially by the solar wind and the geomagnetic field. Fluxes of the most common cosmic ray nuclei can be seen in Figure 1.3

More than 50000 cosmic rays with energies greater than 1 GeV impacts each square meter of the Earths atmosphere per second. Life on Earth is protected against this deadly deluge by the atmosphere, which prevents them from reaching the surface of the planet unhindered. As a cosmic ray particle or nucleus encounters the atmosphere, it will quickly begin interacting with the nuclei of atmospheric gases. The energy of the cosmic ray will be distributed among the resulting fragments. These consist mostly of pions that subsequently decay into muons, and fragments of the cosmic ray and nucleus which goes on to collide with new nuclei. This process generates what is called an air shower, and is what leads to the abundance of muons in the atmosphere. In the next section it will be explained in greater detail.
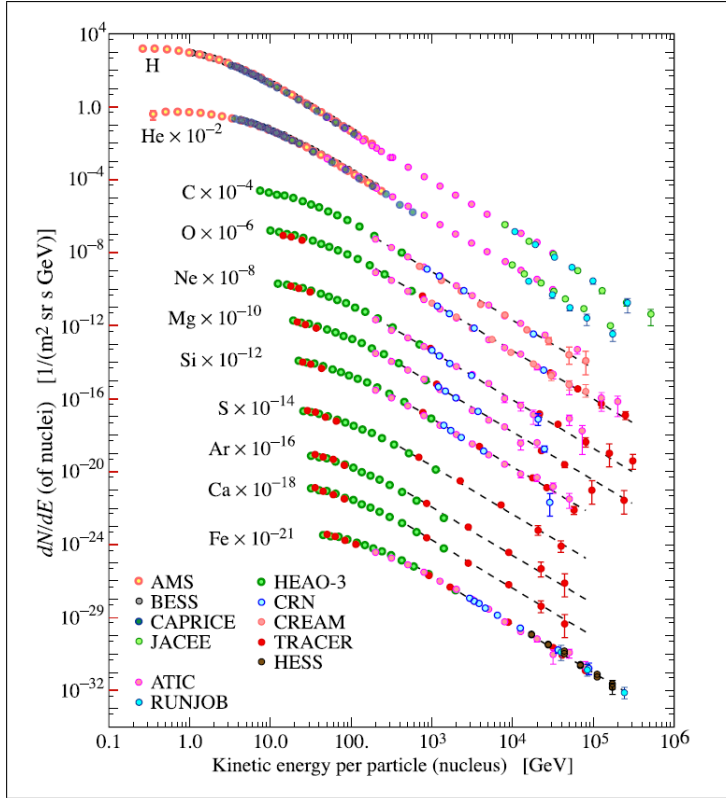
Figure 1.3: The fluxes at the top of the Earths atmosphere of various nuclei which make up cosmic radiation. From [3].

1.1.4 Muon production by cosmic rays

As a cosmic ray enters the atmosphere, the density of the medium it traverses rises exponentially. It quickly collides with a nuclei in the upper atmosphere, and a far-reaching cascade called an *air shower* is initiated. The primary constituents of the particles generated at the initial collision are pions (π^-, π^+, π^0) and a small amount of kaons and various baryons. The kaons are short-lived and will decay into pions or muons. The baryons will either go on to collide with other nuclei and initiate new cascades, or decay.

The neutral pions (π^0), having a lifetime of about 84 as (0.084 femtoseconds), will immediately decay into gamma photons ($\pi^0 \rightarrow \gamma\gamma$). The gamma photons will decay to electron-positron pair through pair production, and the electrons and positrons will generate new gamma photons via the Bremsstrahlung effect. These two processes create a cone of rapidly multiplying electrons, positrons, and gamma photons, called an electromagnetic cascade. This continues until the energy of the photons fall below the pair production threshold, and other energy loss processes begins to dominate.

The charged pions (π^-, π^+) have lifetimes of $\sim$12 ns, and do not decay as fast as their neutral cousin. But 12 ns is still too short of a lifetime to reach the the surface of the Earth, and they instead decay far up in the atmosphere. Their decay products are muons and muon neutrinos ($\pi^- \rightarrow \mu^- + \nu_\mu$ and $\pi^+ \rightarrow \mu^+ + \bar{\nu}_\mu$). The neutrinos have extremely small interaction cross sections, and most will traverse the atmosphere and continue through the Earth without leaving a trace. We therefore ignore them and turn our attention to the muons. The lifetimes of the positive and negatively charged muons are both 2.2 µs. This is nearly 200 times that of the

charged pions, and muons are also some of the longest-lived among the unstable particles. Those muons generated in air showers which have large energies (>2 GeV) are moving at nearly the speed of light, and therefore their Lorentz factor is large (>20). Their long lifetimes, combined with Lorentz contraction of the distance to the surface of the Earth from the reference frame of the muon, means that a great fraction will reach sea level before decaying. There they may be detected by muon detectors such as the CRT.

Most muons in an air shower are generated high up in the atmosphere, typically at a height of 15 km. The muon production energy spectrum is continuous, but the spectrum at sea level is modulated by the energy loss of muons in the atmosphere, and by the decreased Lorentz contraction of low-energy muons. As a result of this, most muons detected at sea level have energies of 2 GeV-4 GeV, which corresponds to the point of minimum energy loss for muons, see Figure 1.5. The atmospheric muon flux is strongly dependent on elevation, as the atmospheric depth the muons must traverse decreases rapidly with increasing altitude.

Energetic muons will lose little energy before reaching sea level, typically around 2 GeV. The amount of energy lost naturally depends on the length the muon has to traverse in the atmosphere, and therefore the muon trajectory will have an impact on the energy loss. Consider two muons generated at a height of 15 km, one with a polar angle of 0°(moving directly towards the surface of the Earth), and the other with a polar angle of 45°. The muon with a vertical path will traverse exactly 15 km of the atmosphere before reaching the surface. The muon with a slant will not only lose more energy due to travelling longer, but will also have a greater probability of decaying before reaching the surface, as the time it spends in flight is greater. We therefore expect that if one set out to count the number of muons arriving at sea level at various polar angles, there would be most vertical muons and least horizontal muons. This is indeed what has been measured. The overall muon flux at sea level has been found to be $\propto cos^2(\theta)$. In Figure 1.6, data on the muon flux versus incident angle as found by several experiments is plotted, and we see that the total muon flux decreases by increasing polar angle.
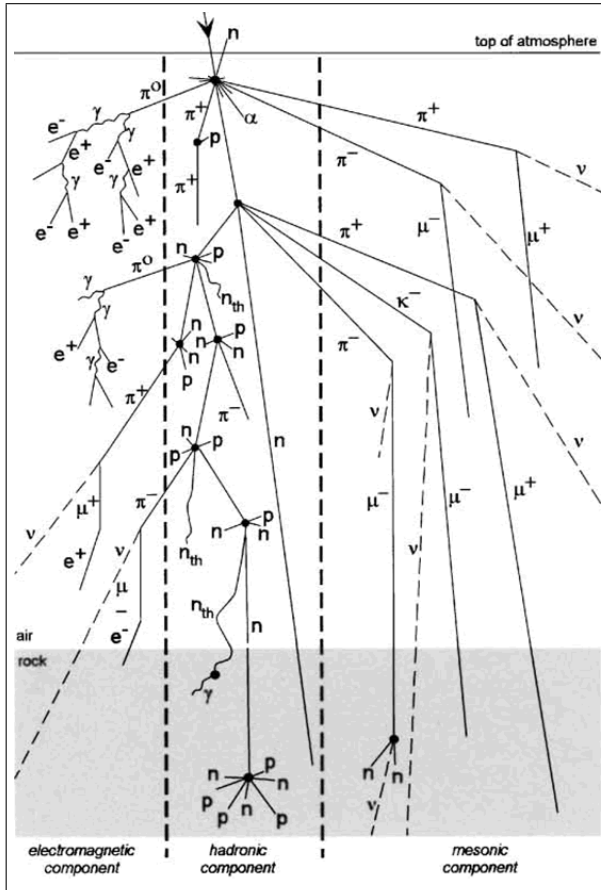
Figure 1.4: Overview of the main processes occurring in an air shower (cosmic ray initiated cascade), and how the generation of particles in the atmosphere and rock occurs. Modified from Allkofer and Grieder, 1984; Clay and Dawson, 1997.
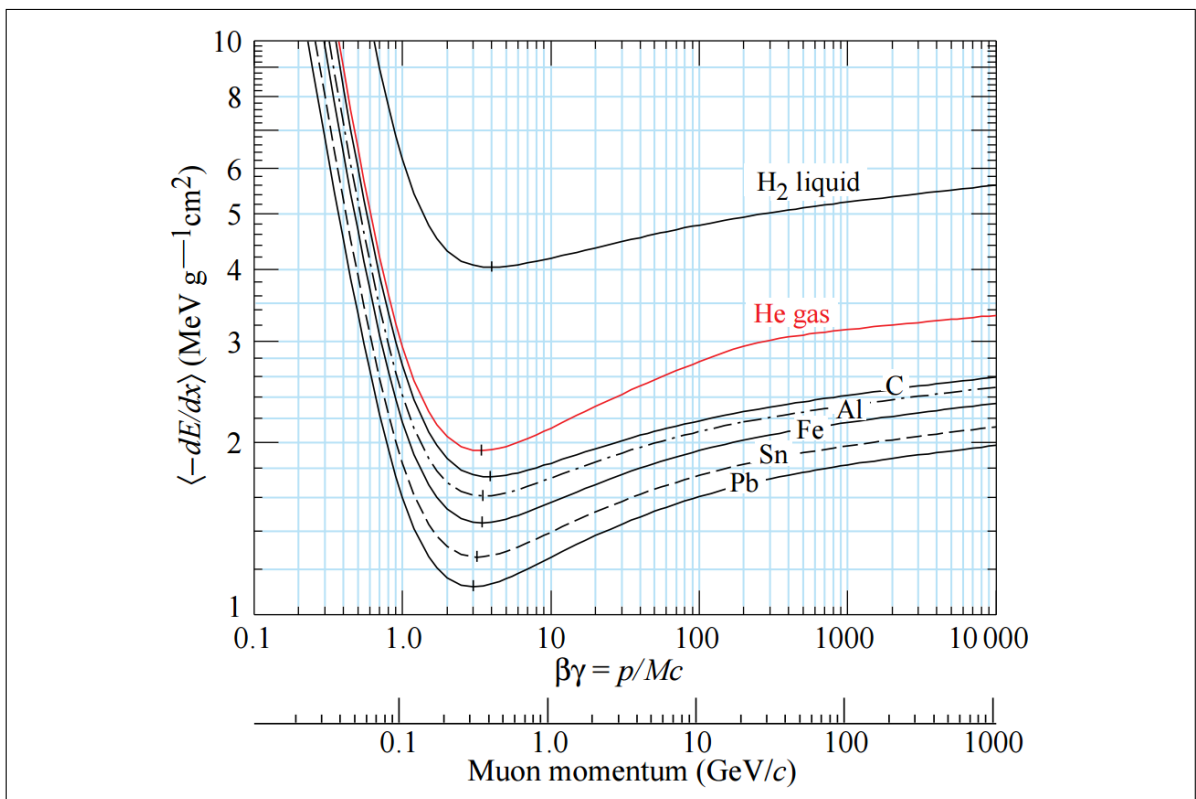
Figure 1.5: Energy loss of a muon due to ionization in various materials as a function of muon momentum. From [3].
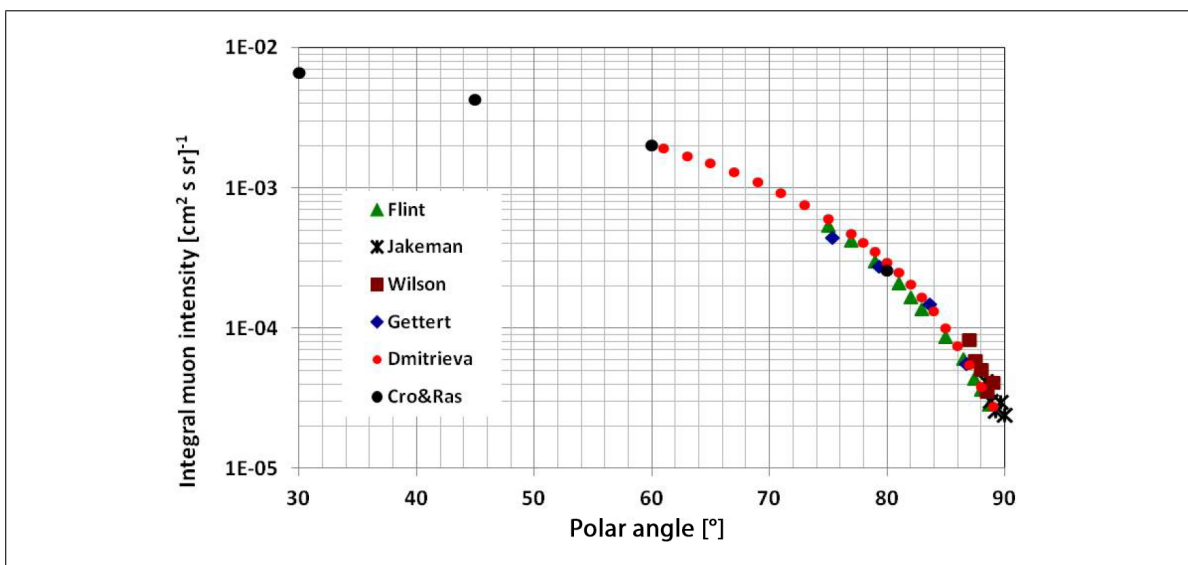
Figure 1.6: Absolute integral intensity of muons at sea level plotted versus polar angle. From [7].

1.2 Muons today

In contemporary science, muons generated by cosmic rays are finding use in several fields. The abundance of muons at sea level, their highly penetrating power, and the ease of which they ionize the medium they traverse makes them a valuable resource. In addition, favourable experimental conditions makes it easy to collect large amounts data over extended periods of time, using relatively simple detectors based on well-tried and established scintillator-photomultiplier methods.

An interesting application is muon tomography, where the passage of muons through solid matter is used to construct a three-dimensional volumetric representation of the object they passed through. Since muons are much more penetrating than X-rays, they can be used to image much thicker and absorptive materials than what is possible with X-ray based computed tomography. The natural abundance of muons means that no artificial source is required to perform muon tomography, and the technique can be used virtually anywhere a detector can be placed, even deep underground. Muon tomography was first used in the 1950s to measure the ice overburden of mining tunnels in Australia, and in the 1960s the technique was used in an attempt to locate hidden chambers within the Pyramid of Chephren in Giza [8]. More recently, muon tomography has been used in an effort to image the interior of Mount Vesuvius [9], a volcano located in Naples, Italy. The same technique will be applicable to volcanoes around the world, with the goal of better predicting when they erupt.

A new technique in muon tomography called muon scattering tomography ("muography") was developed at the Los Alamos National Laboratory in 2003 [10]. This technique reconstructs the path of both incoming and outgoing muons, and can be used to detect dense materials within less dense material by assessing the deflection of muons by Coulomb scattering. Muography has found use outside research, as companies have started using it to detect hidden nuclear cargo and contraband at shipping ports [11, 12]. Muography is also used as a non-invasive method of characterizing nuclear waste storages, and was used in order to assess the status of the Unit 2 reactor core at the Fukushima-Daichii facility in Japan, which was famously damaged by a tsunami in 2011 [13, 14]. Muography is proving a powerful tool for internal imaging of large and/or inaccessible structures.

Apart from the use of muon detectors in tomography, they are also employed in cosmic ray experiments. Notably, an upgrade to the Pierre Auger Observatory in Argentina, which detects ultra-high energy cosmic rays, will see 61 muon counters installed underground, using plastic scintillators and silicon photomultipliers in much the same fashion as the CRT. Cosmic muons are also very useful for commissioning, aligning, and verifying detectors used in large particle physics experiments, such as the A Large Toroidal Lhc ApparatuS (ATLAS) detector built around an interaction point on the Large Hadron Collider at the European Organization for Nuclear Research. During commissioning of the ATLAS inner detector, cosmic muons were used to test the software integration, tune the readout parameters, study its performance, check alignments, and other tasks [15]. After installation in ATLAS, located about 100 m below the surface of the Earth, muons can still reach the detector through a wide vertical access shaft

extending to the surface. This enables operators to continuously monitor the subdetectors for daily verification and adjustment of its alignment parameters.

Lastly, cosmic muons are often used in physics undergraduate courses and in public science outreach programmes. The simplistic nature of scintillator-based muon detectors makes them cheap to manufacture and easy to operate, and the high energy of cosmic muons can be used to provide hands-on experience with relativistic effects. Several types of simple detectors for cosmic muons can be purchased by schools, and usually consists of one to three plastic scintillators with photomultiplier tubes. These systems range from simple analogue do-it-yourself solutions such as the Berkeley Lab Cosmic Ray Detector [16] to more advanced systems such as the distributed Washington Area Large-scale Time-coincidence Array (WALTA) [17], which provides equipment to secondary schools in the Seattle area with the aim of building a large-scale cosmic ray detector.

2 | Detector simulations

In the current chapter the simulations that were performed on the scintillators and light guides are presented. The C++ framework GEometry ANd Tracking 4 (GEANT4) [18] was used to model the geometry under consideration, and a simplified model of the muon flux was employed to investigate the photon response at a virtual detector mounted on the end of the light guide. The results are used to define an optimal light guide end radius, and in order to estimate the probability distribution of the number of photons that reach the detector, based on where the muon impacts the scintillator.

2.1   Introduction

Today, most particle physics experiments employ numerical simulations to gain knowledge of the behaviour of their detectors and the physical processes under consideration. At the core of these simulations is a numerical technique developed in the late 1940s by Stanisław Ulam, called the *Monte Carlo Technique* [19, 20]. This technique rely on the probabilistic nature of many physical events. Instead of attempting to calculate the final state of an event (often a complicated many-body problem) from first principles, the MC technique simulates the event numerically. Each step in the simulation is governed by probabilistic processes, such as scattering, absorption, decays, etc. The end state of the event will thus be but one of a very large set of end states. The simulation is performed several times to obtain meaningful statistics, each simulation normally resulting in a different end state. The collection of MC simulations therefore gives a probability distribution of the end state, where the uncertainty in the results is set by the number of times the event was simulated, and how well the physical models used in the simulation corresponds to reality.

Experimentalists use MC simulations for several reasons. It enables them to understand how data from the experiments should be analysed, investigate optimizations or changes to the detector geometry, optimize the design of event reconstruction software, study the effects of radiation damage, and to find the resolution and other performance parameters of the experiment. Pre-existing software frameworks for MC-based particle simulations such as GEANT4 [18] and FLUktuierende KAskade (FLUKA) [21] are widely used in particle physics experiments, and enables the experimentalist to quickly get started on the task of simulating their experiment with relatively little effort. The geometry of the detector and its physical properties are modelled in the software framework, events are generated, and the results of the simulations can later be analysed offline.

2.2 Motivation

Although the CRT is a relatively simple detector with few elements, having some knowledge of its response to an incident muon is beneficial. The main motivations for performing the simulations are listed in the following paragraphs.

Light guide optimization

The scintillator slabs have already been fitted with acrylic glass light guides, dimensioned to fit a conventional PMT with an aperture diameter of 27 mm. The SiPM chosen for the upgrade of the CRT has an active area of 36 mm^2. This is less than the light guide area by a factor of 16. Therefore, a large fraction of the scintillation photons (just "photons" from here) exiting the light guide will not be registered if it is simply mounted on the existing light guide. A narrowing light guide was required to be fitted on the existing light guide in order to concentrate the photons onto the SiPM. Finding the optimal end radius was not trivial, and therefore it was left to simulations.

Scintillator response

As the scintillator is relatively large, it is easy to imagine that the photon response, i.e. the time distribution of photons impeding on the SiPM, will vary based on where the muon crossed the scintillator. This variation comes from the loss of photons by absorption and scattering in the scintillator, light guide, and wrapping material. It is also expected that the photons will display a grouping behaviour due to reflections within the scintillator. The signal generated by the SiPM will have an amplitude and shape directly dependent on the amount of scintillation photons and their time distribution. Therefore, an understanding of this variation would enable a better analysis of the detector data, and possibly enable more efficient readout electronics to be designed.

2.3   Simulation setup

The detector simulation was created with GEANT4, a toolkit for simulating the passage of particles through matter. "The toolkit provides a diverse, wide-ranging, yet cohesive set of software components which can be employed in a variety of settings. These range from simple one-off studies of basic phenomena and geometries to full-scale detector simulations for experiments at the Large Hadron Collider and other facilities." [18, p.252]. As the CRT in its entirety consist of 16 identical scintillator slabs, only one was simulated and its properties assumed to be valid for the remaining 15. The geometry and materials used for the simulations will now be presented.

2.3.1   Geometry

The simulation geometry consisted of three main components: The scintillator, the light guide, and a detector area. The detector area, located at the end of the light guide, was a cube of silicon resin with a sensitive area encased within. It was not meant to be a realistic representation of the SiPM, but rather a simplified model able to detect all photons exiting the light guide. The sensitive area registers data about the photons impeding on it, and destroys them. The sensitive area and the tapered light guide were separated by a layer of silicone resin, 0.8 mm thick. This layer is present on the real SiPM chosen for the CRT upgrade, and was therefore included in the geometry. Illustrations of the geometry are presented in Figure 2.1, and Figure 2.2 shows an enhanced view of the detector area. The dimensions used for the simulation geometry were taken from measurements of their physical counterparts, and are listed in Table 2.1.
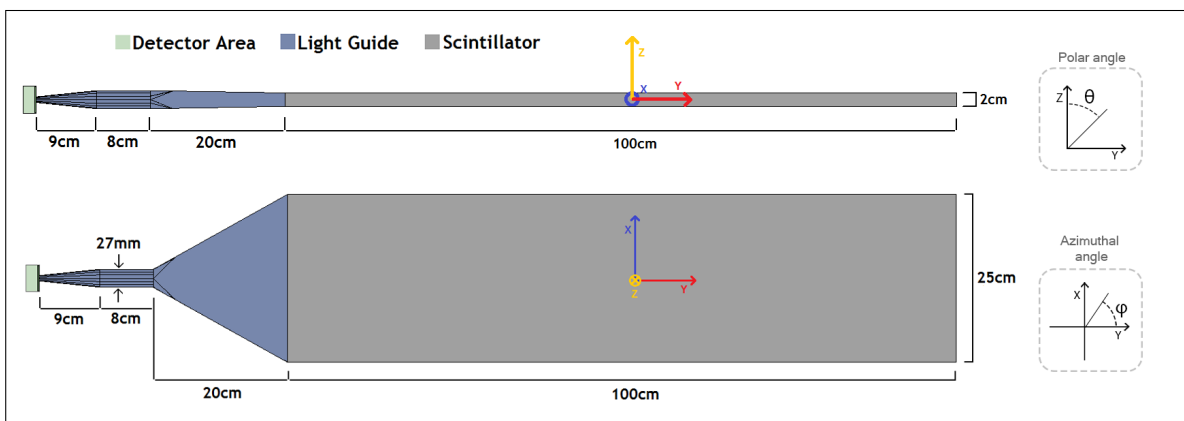
Figure 2.1: Top and side view of the simulation geometry. The origin of the simulation is defined as the middle of the scintillator. The polar angle θ and azimuthal angle φ are defined as shown.
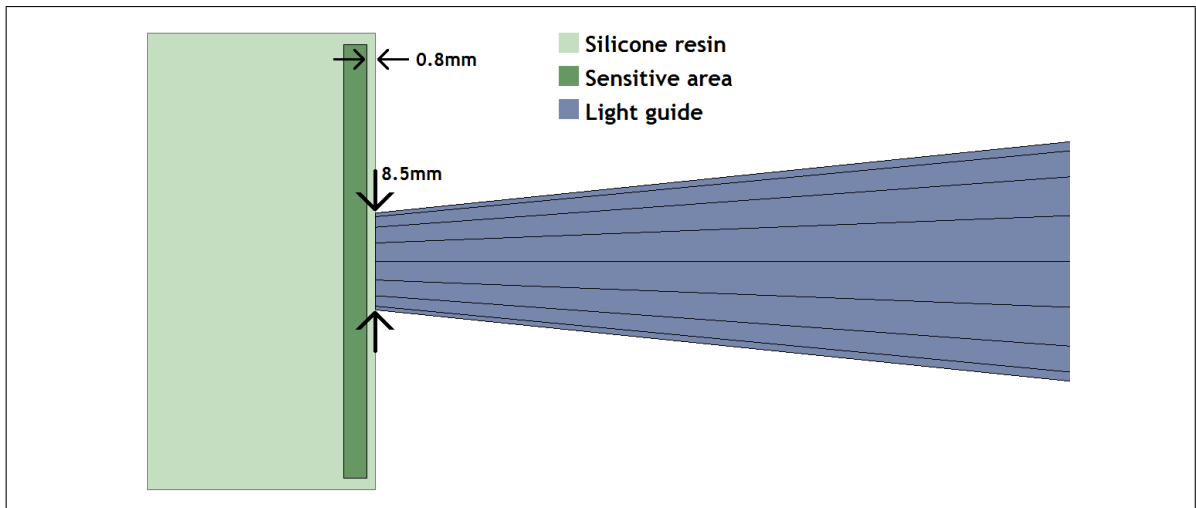
Figure 2.2: Enhanced view of the detector geometry used in the MC simulations. The sensitive area is much larger in the simulation than in reality. This is done so that all the photons exiting the light guide will be registered. Later, a subset of the registered photons can be extracted from the simulation dataset and analysed. This subset can for example be only those photons that would impede on the SiPM. The surface of the sensitive area pointing towards the light guide is located at $x = -87.08\,\text{cm}$.

Scintillator			Light guide					
			Trapezoid			Cylinder		Tapered Cylinder
Length	100cm		Length	20cm		Length	8cm	Length	9cm
Height	2cm		Height, base	2cm		Diameter	2.7cm	Diameter, start	27mm
Width	25cm		Height, peak	2.7cm				Diameter, end	8.5mm
			Width, base	25cm				
			Width, peak	2.7cm				

Table 2.1: Simulation geometry dimensions

2.3.2 Materials and surfaces

The materials of the various geometry objects were defined by the properties that would affect the generation of scintillation photons, their propagation, and absorption. The values were taken from the datasheets of the materials where possible, and estimated if no information existed. The scintillator material, Bicron BC-408 [1], was defined by its scintillation gain, emission spectrum, decay time, attenuation length, and index of refraction. These properties are listed in Table 2.2. The emission spectrum was defined in the code as a set of emission probabilities for various wavelengths, seen in Figure 2.3. GEANT4 interpolates these points and uses them when generating scintillation photons. The acrylic glass of the light guide was defined by its index of refraction and attenuation length for UV radiation. However, the manufacturer of the acrylic glass did not specify this latter property in their datasheets, and were not responsive to inquiries. Therefore, it was estimated by referring to the work done by M. Bodmer et al. [22], whose measurements of acrylic glass attenuation lengths are shown in Figure 2.4. Most types

of acrylic glass appears to have attenuation lengths on the order of one meter near the emission spectrum of BC-408. The attenuation length of the light guide was therefore set to 1 m, with the assumption that the absorption of photons only serves to alter the amount photons registered by the SiPM, and not the time nor probability distribution. This meant that an incorrect estimation for the attenuation length would not greatly impact the simulation results. The silicone resin was defined only by its refractive index.

The photons are contained within the scintillator and lightguide by two methods: The first is total internal reflection, in which a photon will reach the boundary between the bulk it is currently propagating (scintillator or light guide), and the air surrounding this bulk. If the photon strikes this bulk/air boundary with an incident angle greater than some critical angle (with respect to the surface normal), it will be reflected back into the bulk material with the same angle. The value of the critical angle is proportional to the fraction of refraction indexes of the two materials, given by a rearrangement of Snell's law:

$$\theta_C = sin^{-1}\left(\frac{n_2}{n_1}\right) \tag{2.1}$$

Here, n_1 is the index of refraction of the material the photon is propagating, and n_2 is the index of refraction of the material on the other side of the boundary. Total internal reflection is a common optical phenomenon, and is also valid for other types of waves. The second method of containment is a layer of highly reflective Mylar. The Mylar is wrapped around the outside of the scintillator and light guide, serving to reflect any photons not contained in the scintillator by total internal reflection. There is therefore a relatively complicated boundary region, with the bulk material and Mylar wrapping being separated by a thin, but not omittable, layer of air. Luckily, GEANT4 enables simple modelling of such a "bulk + thin layer + wrapping" region. The only parameters needed are the refractive index of the bulk material and intermediate material, how well polished the bulk surface is, the reflectivity of the wrapping, and whether the wrapping is polished or ground. A disadvantage of this model is that one cannot specify how well polished the wrapping is, only whether it is perfectly polished or scattering. Since Mylar is highly reflective and polished, this limitation did not affect us.

The refractive indexes and attenuation lengths were assumed constant for all wavelengths in the emission spectrum of BC-408, which ranges from 350nm to 520nm. The material and surface properties used in the simulation are listed in Table 2.2 and Table 2.3. The scintillator and silicone resin properties were found in the manufacturer datasheets [1, 23].

Material:	Scintillator	Acrylic glass	Silicone resin	Air
Index of refraction	1.58	1.488	1.41	1
Attenuation length	210cm	100cm	Not specified	Not specified
Emission spectrum	See Figure 2.3	N/A	N/A	N/A
Scintillation yield	1.28×10^4 photons/MeV	N/A	N/A	N/A
Decay time	2.1 ns	N/A	N/A	N/A

Table 2.2: Simulation material properties

Surface	Polished	Reflectivity
Scintillator/air	0.99	N/A
Acrylic glass/air	0.99	N/A
Mylar wrapping	1	0.98

Table 2.3: Simulation surface properties
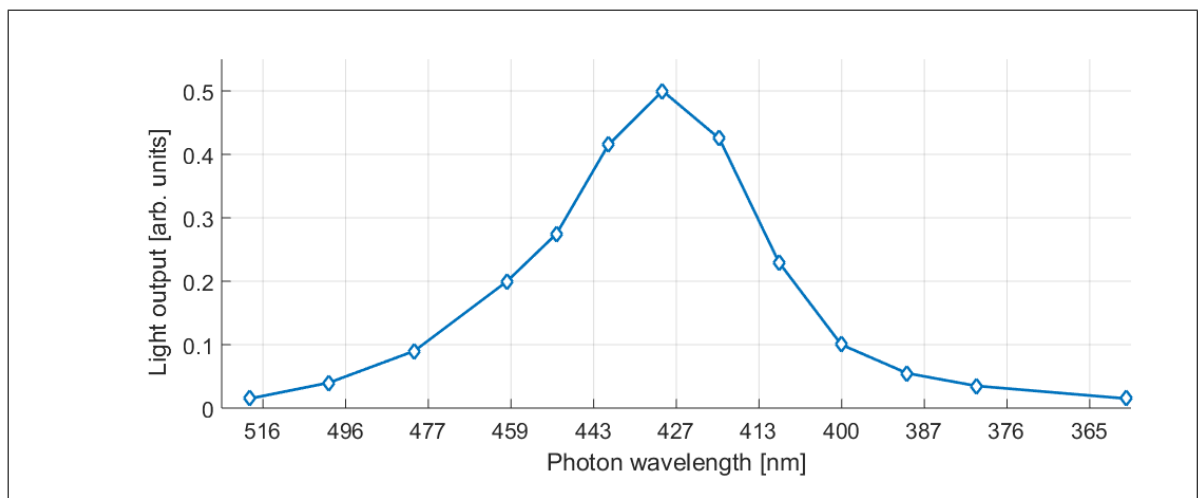
Figure 2.3: Emission spectrum of the scintillator material used in the GEANT4 simulations

Sample	Wavelength (nm)				
	375	405	440	543	632
	Attenuation Length (m)				
Tecacryl #1	0.14 ± 0.01	0.57 ± 0.01	7.30 ± 0.93	11.4 ± 1.7	36.7 ± 28.5
Tecacryl #2	0.12 ± 0.00	0.35 ± 0.01	1.68 ± 0.09	10.1 ± 1.7	10.8 ± 2.3
Tecacryl #3	1.36 ± 0.03	4.03 ± 0.14	5.65 ± 0.38	$-$	30.6 ± 19.6
Tecacryl #4	1.26 ± 0.04	3.87 ± 0.22	4.89 ± 0.30	$-$	24.9 ± 19.7
RPT #1	0.30 ± 0.03	1.52 ± 0.05	2.50 ± 0.07	492 ± 3270	15.6 ± 6.61
RPT #2	0.31 ± 0.02	1.56 ± 0.06	2.51 ± 0.11	56.2 ± 39.4	2.9 ± 3.7
RPT #3	$0.29 \pm .03$	1.55 ± 0.07	2.61 ± 0.11	33.6 ± 8.80	16.5 ± 5.6
RPT #4	0.11 ± 0.00	0.28 ± 0.01	1.60 ± 0.05	20.3 ± 6.1	7.7 ± 1.1
Spartech #1	0.70 ± 0.02	1.25 ± 0.07	1.54 ± 0.07	17.1 ± 2.8	22.0 ± 11.8
Spartech #2	0.56 ± 0.01	1.19 ± 0.06	1.53 ± 0.04	20.1 ± 4.4	16.5 ± 6.9
Spartech #3 (short side)	0.07 ± 0.00	0.23 ± 0.01	2.41 ± 0.92	$-$	24.9 ± 95.0
Spartech #3 (long side)	0.05 ± 0.00	0.26 ± 0.00	3.01 ± 0.23	$-$	11.2 ± 8.0
Spartech #4	0.15 ± 0.00	0.28 ± 0.01	2.74 ± 0.33	230 ± 379	27.8 ± 19.3
McMaster (Spartech #1	0.48 ± 0.03	1.51 ± 0.11	1.80 ± 0.15	$-$	$-$
McMaster (US Cast) #2	0.18 ± 0.00	1.38 ± 0.03	2.46 ± 0.25	$-$	$-$

Figure 2.4: Table of attenuation lengths for various types of acrylic glass. Entries marked as '-' have a calculated transmittance of 100%. Reused with permission from [22].

2.3.3 Muons

The muon flux at sea level is relatively complicated. S. Cecchini and M. Spurio [7] collected data on muons at sea level, and found that they have a wide energy spectrum, from less than 1 GeV to several TeV. Their angular distributions are also not constant, but depends primarily on the energy of the muons. Low energy muons primarily move vertically down, whereas high energy muons are more likely to come from the horizon. In addition, the geomagnetic field plays a role in the flux direction. However, for the simulation only a simplified muon flux model was used. The mean energy of muons at sea level is ~ 4 GeV, and at this energy, the muon angular distribution is $\propto cos^2\theta$ [3], where θ is the polar angle. Therefore, the simulations were performed with these parameters unless otherwise stated.

2.4 Simulation results

In this section the results of the MC simulations are presented. The optimal light guide end radius is found, the illumination homogeneity of the SiPM is verified, and the scintillator response to muons of various angles and crossing positions is studied.

2.4.1 Light guide optimization

The light guide tapering serves to concentrate the scintillation photons onto the SiPM. Finding the best end radius is however not trivial. If the radius is very small, few photons will exit the light guide end, as most will simply be reflected back into the scintillator. Consequently, few photons will hit the SiPM. Increasing the radius of the light guide end will increase the probability that a photon exits it, but if the area of the light guide is large compared to the area of the SiPM, the probability of hitting the SiPM is low. Finding the optimal end radius was one of the main objectives of the MC simulations, and 10 different values were tested. For each value, 5000 muons were homogeneously distributed across the scintillator, each with an energy of 4 GeV and impacting the scintillator with a polar angle $\theta = 0$. Figure 2.5 shows the results of these simulations. From 1 mm to 4 mm, the amount of incident photons increase linearly with the radius. In this region nearly all of the photons exiting the end of the light guide hits the SiPM. However after 4 mm, the increase reduces and produces a maxima near 6 mm before decreasing again. In this region, the two effects mentioned earlier are counteracting each other. Increasing the radius increases the probability of photons exiting the light guide, but reduces the probability of them hitting the SiPM. As the maxima is quite broad, any value between 4 and 10 mm could effectively be used, and for all subsequent MC simulations an end radius of 6 mm was used.
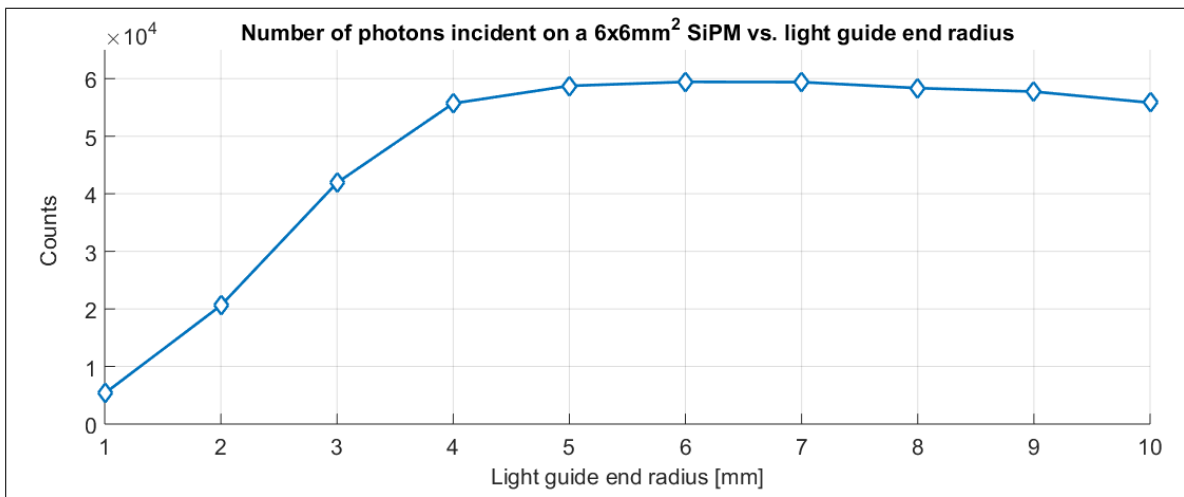
Figure 2.5: Total amount of incident photons on the SiPM as a function varying light guide end radius.
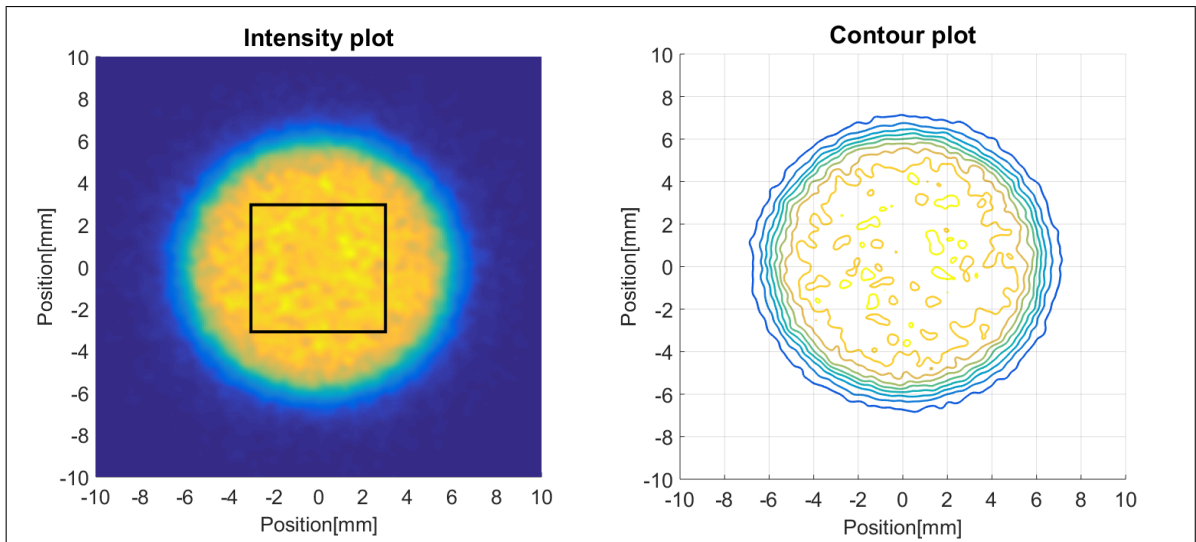
Figure 2.6: Photon position distribution on the detector area. The simulation was performed with a homogeneous muon distribution and the light guide end radius set to 6 mm. The black rectangle denotes the area of a 6×6 mm² SiPM. The contour levels are evenly spaced, with arbitrary units.

2.4.2   Illumination homogeneity

A SiPM consists of a finite amount of G-APDs. When a G-APD is activated by an incoming photon, it generates a signal that is independent on the energy of the photon, and independent on the amount of incident photons. For some time after being activated (on the order of tens of nanoseconds), the G-APD will not be able to become activated again. Therefore a homogeneous illumination of the SiPM is desirable, so as to ensure that the incident photons are spread evenly across its G-APDs. This will reduce the probability of one or more photons hitting the same G-APD, and so be registered as a single photon. Chapter 3 covers SiPM theory, and discusses this in greater detail.

To verify the homogeneity of the flux of photons exiting the light guide, 50 000 muons with $E = 4$ GeV were homogeneously distributed across the scintillator with polar angle $\theta = 0$. The positions of the photons were registered as they exited the lightguide and impacted on the detector area. Figure 2.6 shows the results of the simulation. The entire 6×6 mm² area enclosed by the black rectangle is uniformly illuminated, and this area is the same as the area of the SiPM used for the CRT upgrade. One can also see that the location of the SiPM does not need to be very precise, as it will be homogeneously illuminated even if offset 1 mm from the center in any direction.

2.4.3   Scintillator response

The large dimensions of the scintillators motivated us to investigate how the scintillator response would depend on the location that a muon crosses it. It was expected that the amount of photons that exits the light guide, N_{LG}, and their distribution in time, $A_{LG}(t)$, would depend on the muon position, and this was found to be correct.

Figure 2.7: Positions P_1 ($y = -49.9\,\text{cm}$, $x = 0\,\text{cm}$), P_2 ($y = 49.9\,\text{cm}$, $x = 0\,\text{cm}$), and P_3($y = -48\,\text{cm}$, $x = -10.5\,\text{cm}$), denoting the positions of impacting muons used in the investigation of $A_{LG}(t)$.
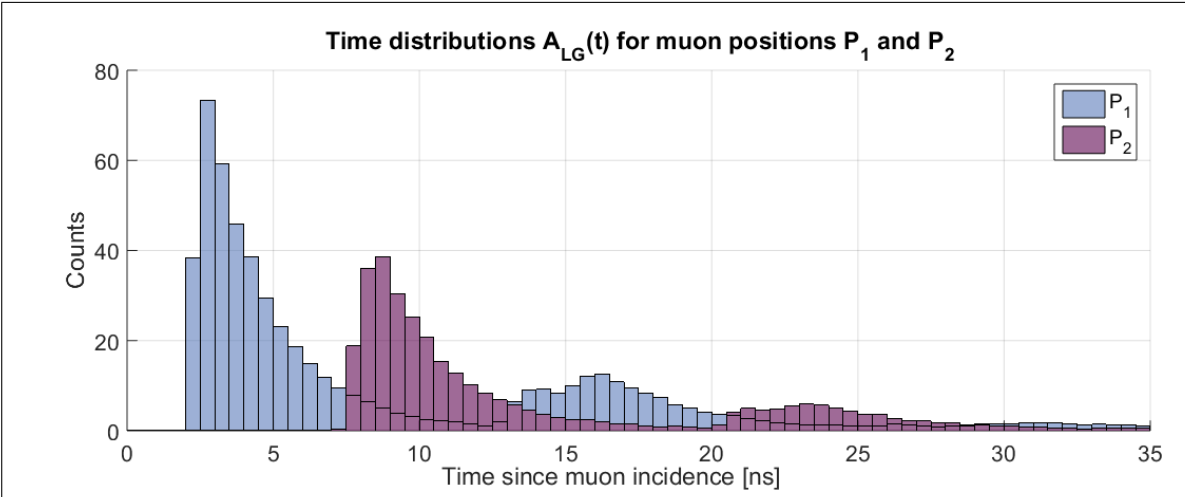
Figure 2.8: Time distributions of scintillation photons exiting the light guide for 200 muons passing through P_1 and P_2. The histogram bin size is 0.5 ns.

Scintillation photons arrival time distributions $A_{LG}(t)$ and $A_{SiPM}(t)$

The photons generated by a muon crossing the scintillator will not reach the SiPM simultaneous since they are emitted equally in all directions from every point on the muon track. Internal reflections and scattering will distribute their arrival times, and this distribution was investigated. First, 200 muons were impacted at opposite ends of scintillator, positions P_1 and P_2 in Figure 2.7, with polar angle $\theta = 0$. The time distributions of all photons exiting the light guide, $A_{LG}(t)$, can be seen in Figure 2.8. As expected, the total amount of photons exiting the light guide (N_{LG}) is lower when the muon impacts the far end of the scintillator, point P_2, than when it impacts the near end, point P_1. The reflections are easily discernible as secondary peaks in the distributions, explained by the photons being reflected back and forth between the two scintillator ends.

To investigate $A_{LG}(t)$ for a natural muon flux, 50000 muons were homogeneously distributed across the scintillator, with energy and angular distribution as described in Section 2.3.3. In addition to measuring the arrival times, the photon locations were recorded as they impacted the detector area. This information was used to sort out those photons that impacted on a $6 \times 6\,\text{mm}^2$ SiPM positioned on the center of the light guide end. These are denoted N_{SiPM}, and their time distribution is $A_{SiPM}(t)$. The two distributions $A_{LG}(t)$ and $A_{SiPM}(t)$, and their ratio $\epsilon(t)$, are plotted in Figure 2.9. The ratio is nearly constant, meaning that the amount of photons registered by the SiPM will be a good indicator of how many photons actually reaches
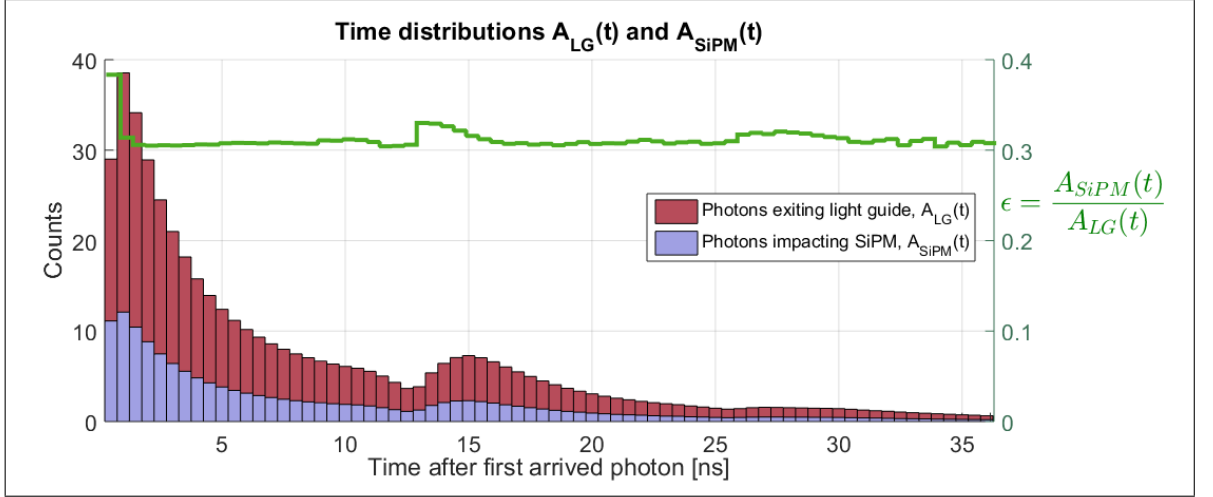
Figure 2.9: Average arrival time distributions of scintillation photons for a natural muon flux, calculated from the simulation results of 50 000 impacting muons. $A_{LG}(t)$ is the time distribution of the set of photons exiting the light guide. $A_{SiPM}(t)$ is the subset of photons that impacts on a $6 \times 6 \, \text{mm}^2$ area mounted on the middle of the light guide end. This area represents an equally sized SiPM. Their ratio, $\epsilon(t)$, denotes the photon detection efficiency of the scintillator-SiPM setup.

the end of the light guide at any given time.

Dependency of N_{LG} on muon crossing position

The amount of photons reaching the end of the light guide, and exiting it (N_{LG}), will depend on a few factors, primarily the muon angle of incidence, and the statistical fluctuations that govern how many scintillation photons are generated. In addition to these factors, the position where the muon crosses the scintillator will also have an impact. This is due to absorption in the scintillator bulk, and we will also see that the geometry of the light guide plays a major role.

The dependence of N_{LG} on the position of an incident muon was investigated by distributing muons homogeneously over the scintillator area with $E = 4 \, \text{GeV}$ and polar angle $\theta = 0$. Each $5 \times 5 \, \text{mm}^2$ square of the scintillator were impacted with 5 muons, resulting in a total of 50 000 muons. N_{LG} was then plotted as a 2D function, $N_{LG}(P_x, P_y)$, of where the muon that generated those photons impacted the scintillator. The result of the simulation is shown in Figure 2.10. One can see that N_{LG} is relatively independent on the muon position for $P_y > -25 \, \text{cm}$. But for $P_y < -25 \, \text{cm}$, the light guide geometry becomes an important factor. In this region, muons impacting near $P_x = 0$ will generate photons that has a clear line of sight to the end of the light guide, and thus a higher probability of reaching it. The closer to the light guide the muon is incident, the more does its x-position P_x matter. Not only does the total amount of photons N_{LG} depend strongly on the muon position in this area, but so too does the arrival time distribution. This can be seen from Figure 2.11, which shows $A_{LG}(t)$ for muons impacting in one of the two blue areas in Figure 2.10, at $P_y < -40 \, \text{cm}$. Most photons generated here must first reflect off the other end of the scintillator in order to reach the end of the light guide. This distribution was generated by impacting muons through point P_3 ($y = -48 \, \text{cm}$, $x = -10.5 \, \text{cm}$), marked in Figure 2.7.
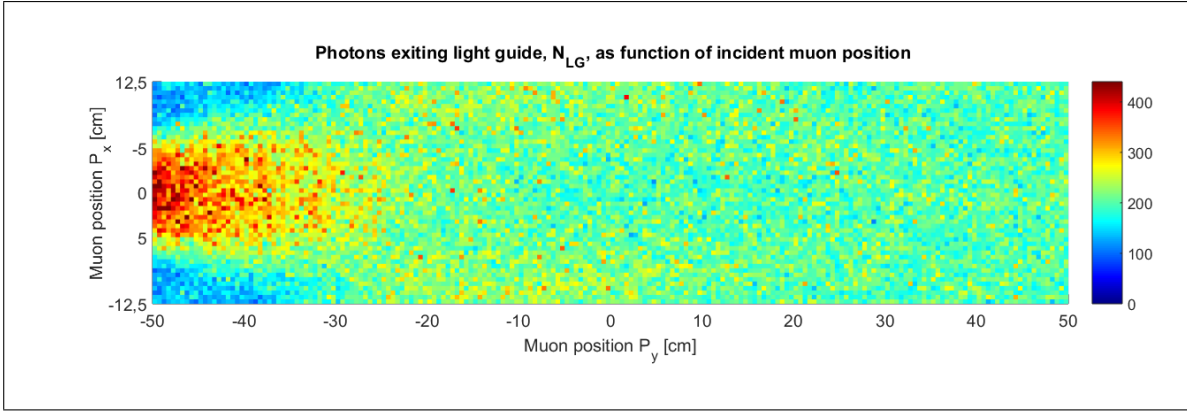
Figure 2.10: Total amount of photons exiting the light guide, N_{LG}, as a function of the muon incident position on the scintillator. The incident muons had $E = 4\,\text{GeV}$ and polar angle $\theta = 0$. The light guide can be thought of as being located outside the plot to the left, beginning at $y = -50\,\text{cm}$.
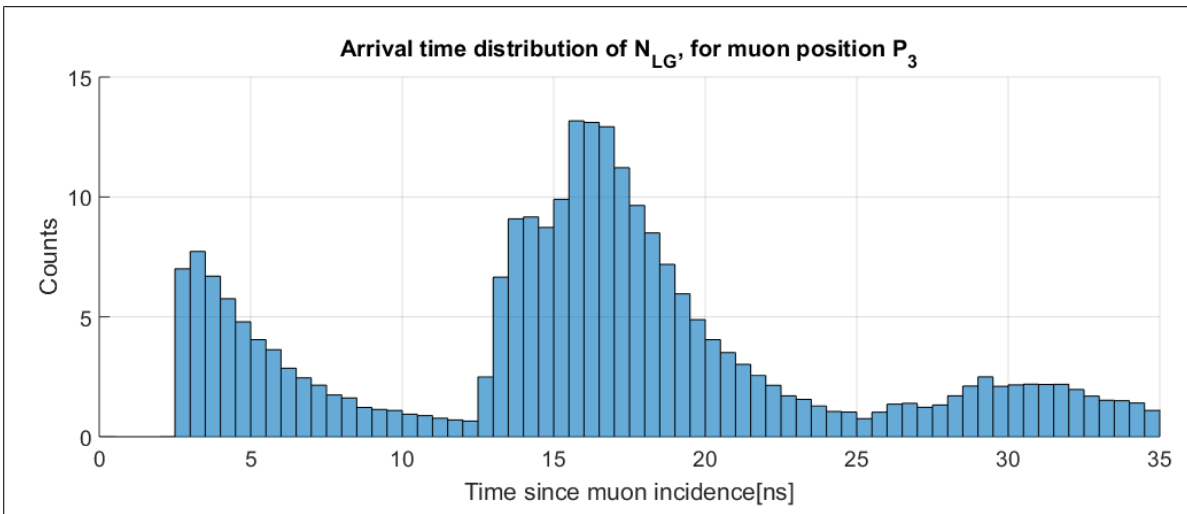
Figure 2.11: Arrival time distribution of scintillation photons exiting the light guide for muons passing through P_3. Comparing this to Figure 2.8, the initial peak has been decreased by a factor of 10, while the second peak remains unchanged.

Variations in N_{LG}

The crossing of a muon will only be registered by the CRT if enough photons reach the SiPM to induce a signal large enough to cross the trigger threshold of the discriminator. Therefore, an analysis of the probability distribution of N_{LG}, denoted D_{LG}, was of interest. This distribution will show how many photons we can expect to exit the lightguide for various events. We have already seen that the amount of photons reaching, and exiting, the light guide (N_{LG}) is a function of the muon position. In addition, statistical fluctuations in the amount of scintillation photons generated and losses in the bulk material will play a role.

For a normal muon flux as described in Section 2.3.3, D_{LG} is relatively wide, roughly following a Poisson distribution, which is expected from the probabilistic nature of scintillation photon generation. That means that for any muon event it will not be possible to say with much certainty how many photons will be registered. However, cuts can be introduced to the data in order to narrow the distribution, and Figure 2.12 shows the results of this. The cuts will now be explained.

The first cut will serve to remove some of the tail of the Poisson distribution. This is done by discarding all those events where the muon impacted with a polar angle $\theta > 64°$. This is the largest polar angle of interest, as any muons with a greater angle than this will not be able to impact both the top plane and the bottom plane of scintillators in the CRT. Since a muon with a large polar angle will cross through more of the scintillator and thus generate more photons, setting this cut removes a small amount of the distribution tail. The second cut we introduce will be to only consider the very first photons arriving at the end of the light guide. As can be seen in Figure 2.9, the distribution of photons follows an exponential decrease, until the photons that has been reflected from the end arrives. By only considering those photons that arrive within ten nanoseconds of the first photon, the statistical fluctuations due to scattering and absorption can be decreased. This cut is also justified by the discriminator, which only triggers on the initial rise of the SiPM signal, and this rise is generated by the first wave of photons that impacts the SiPM. One sees that the distributions becomes narrower and its mode is shifted towards a lower value. The third cut applied will split the dataset into two subsets: One subset for events generated by muons impacting between $-50 \leq P_y < -25$, and another subset for muons impacting between $-25 < P_y \leq 50$. The choices for these values can be seen from Figure 2.10. For low values of P_y the the dependence of N_{LG} on the muon position increases. Therefore it appears necessary to have two distinct D_{LG} distributions: One for photons generated by muons impacting the scintillator near the light guide, and one for muons impacting elsewhere. The right plot in Figure 2.12 shows D_{LG} with these position-based cuts. A tail of low counts manifests, indicating that for some muon positions the amount of registered photons by the SiPM will be very low. This is also seen in Figure 2.13, which plots N_{LG} as a 2D function of the muon position, as in Figure 2.10, but here with cuts applied. The probability of photons reaching the end of the light guide is low if they were generated by muons impacting in the dark blue corners.
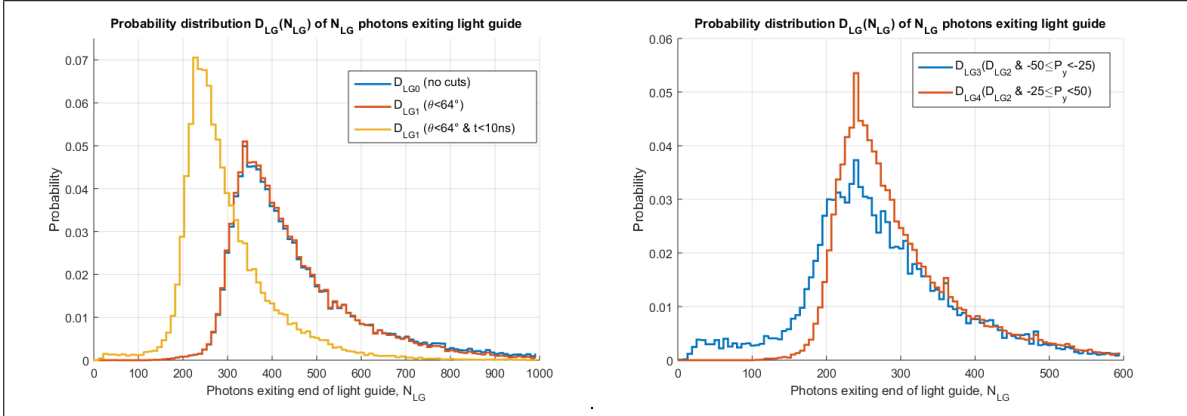
Figure 2.12: Probability distributions of N_{LG}, with various cuts. D_{LG0} is the distribution for all events, with no cuts. It is relatively wide. The tail can be suppressed slightly by introducing the cut $\theta < 64$, which ignores photons from muons impacting the scintillator with a polar angle greater than 64 degrees, as shown by D_{LG1}. By only considering photons reaching the end of the light guide within 10 nanoseconds after the first photon, the distribution can be further narrowed, as shown by D_{LG2}. This introduces a small but continuous distribution nearly reaching 0. This property can be shown to come from muons impacting near the light guide, as illustrated by distributions D_{LG3} and D_{LG4}. These distributions are the same as D_{LG2}, but with additional position-based cuts applied. D_{LG3} is the probability distribution of N_{LG} for photons generated by muons impacting within 25 cm of the light guide, and D_{LG4} is the same distribution except for muons impacting elsewhere. The differences between these two distributions are easily visible. The continuous distribution reaching 0 that appeared in D_{LG2} can be seen to be entirely caused by the muons impacting near the light guide, more precisely in the dark blue areas seen in Figure 2.13
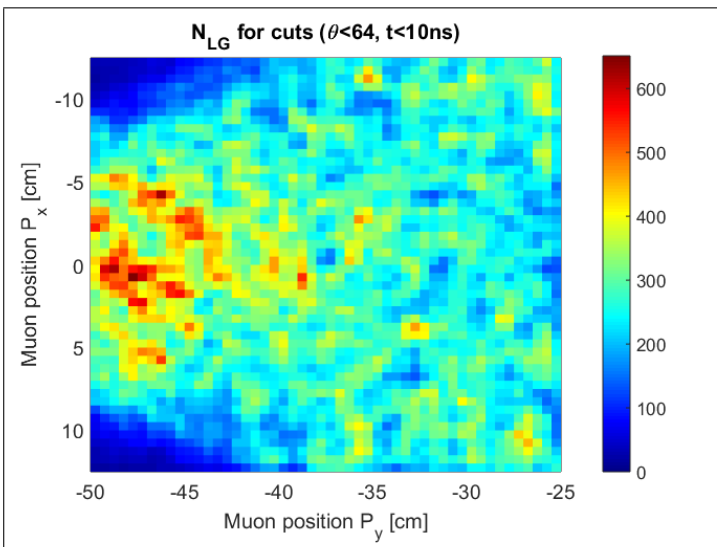
Figure 2.13: Total amount of photons exiting the light guide, N_{LG}, as a function of the muon incident position on the scintillator. The cuts are the same as for D_{LG3} in Figure 2.12. The continuous distribution nearly reaching 0 can be seen to come from muons impacting the scintillator in the dark blue areas of this plot, where the value of N_{LG}(with cuts) is distributed below $\sim$150.

2.5 Conclusions

By creating a simplified model of the scintillator and light guide in GEANT4, an estimate for the scintillator response has been obtained. The photon count probability distribution was found to depend strongly on the muon position within 25 cm of the light guide. In this area the geometry of the light guide favours photons generated by muons impacting close to the light guide and centrally on the X axis. We see a sharp decline in the amount of photons reaching the SiPM for muons impacting in the corners of the scintillator near the light guide. To accurately register these muons, the trigger threshold would have to be set low enough to register the leading edge of the first photon group. This would make the discriminator susceptible to instead trigger on the second photon group, which has a larger amplitude (See Figure 2.11). Additionally, a low threshold would make the detector susceptible to trigger on noise.

The simulations are also not a perfect representation of reality. The models were simplified, and so the results only provide an estimate of what to expect. The trigger thresholds will eventually have to be found experimentally, with the simulation results as a guideline.

3 | The Silicon Photomultiplier

This chapter will introduce the SiPM, and the theory of its operation will be explained. Its properties will be compared to those of traditional photomultiplier tubes, and its benefits and disadvantages discussed. The choice of SiPMs for this project will be justified, and the parameters used for its simulation model will be found. Lastly, sampled signals from a SiPM connected to a CRT scintillator slab will be discussed, and a method for improving the pulse timing presented.

3.1   Introduction

The SiPM is a proportional photon counter, consisting of a large amount (hundreds to tens of thousands) of Avalanche Photodiodes (APDs) operated in Geiger mode. These Geiger-mode Avalanche Photodiodes (G-APDs), or Single Photon Avalanche Photodiodes as they are also called, are connected in parallel via individual quenching resistors. A schematic cross section of a typical SiPM can be seen in Figure 3.1. APDs can be thought of as the semiconductor analogue to photomultiplier tubes. They convert incoming photons to an electrical signal, and amplify it through the phenomenon of avalanche multiplication. APDs require a voltage applied across its terminals to operate. When this applied reverse voltage (or "bias voltage") is larger than what is the called the breakdown voltage, the APD goes from being proportional to digital, called Geiger-mode. A SiPM is operated in this mode, and therefore measures light intensity by counting photons. The amount of photons a SiPM can count per unit time depends on the amount of G-APDs it consists of, and how quickly individual G-APDs recharge after discharging. Since a single G-APD will generate the same output signal regardless of how many photons are simultaneously impeding on it, one must choose a SiPM with enough pixels for the expected photon flux so as not to saturate it.

Development of the SiPM began in Russia during the early 1990s [24]. Researchers succes-fully created small-structured G-APDs and connected them in parallel, forming the first SiPM. Since then, development continued at many separate laboratories and companies, and there are today a wide variety of SiPMs to choose from. Each company often label their SiPMs under a unique name, in order to distinguish themselves from the competition. Examples include Micropixel Avalanche Photo Diode, Multi Pixel Photon Counter, Metal Resistive layer Semicon-ductor APD, and so forth. Even though a large range of SiPMs exists, the technology behind them are still the same.

The G-APDs in a SiPM are microscopic structures, usually spaced with a pitch of 25-100 μm. A large amount of G-APDs can therefore be placed within a small area, and as an effect the SiPM is a physically small device. Commercially available SiPMs are at the time of writing no larger than $6{\times}6\,\text{mm}^2$, with the most common size being $1{\times}1\,\text{mm}^2$. A few typical SiPMs can be seen in Figure 3.2. The maximum area of a single SiPM is limited by various unwanted effects which decreases the signal-to-noise ratio and gain [26]. For this reason, the current largest commercially available SiPM is $6{\times}6\,\text{mm}^2$.
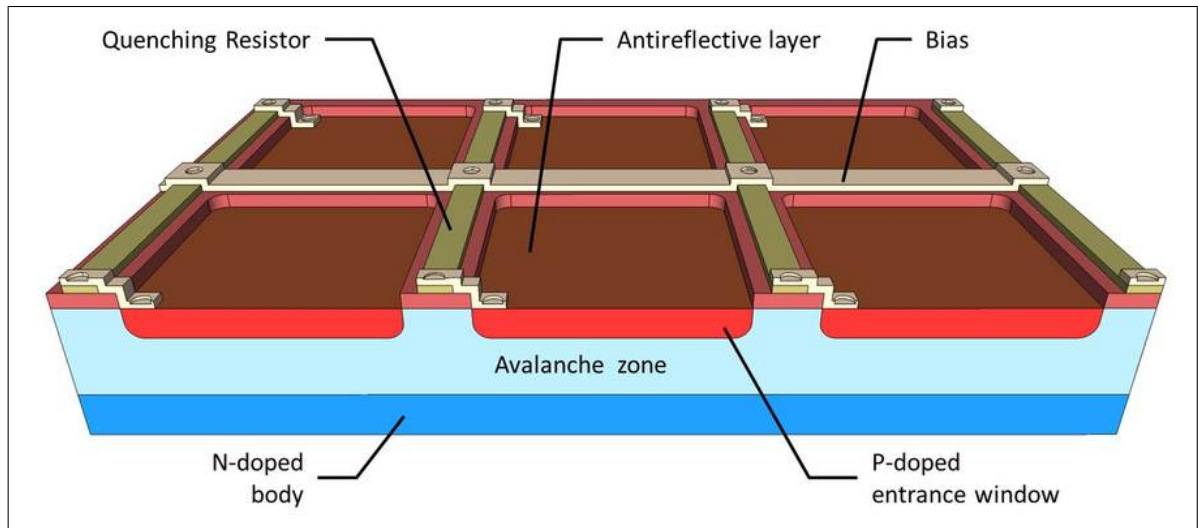
Figure 3.1: Cross section of a SiPM. Ketek GmbH 2017 [25].

Figure 3.2: Photographs of typical SiPMs. Hamamatsu Photonics K.K. 2017 [27, 28, 29].

3.2   SiPM theory

3.2.1   The APD

To explain the SiPM theory of operation, one must start with the APD. A schematic view of an APD can be seen in Figure 3.3. SiPMs consist of silicon APDs, and so they will be discussed here in detail. In general, APDs are semiconductor devices that can be made from germanium, indium gallium arsenide (InGaAs), or indium gallium arsenide phosphide (InGaAsP), in addition to silicon, depending on its required specifications [30]. A typical APD has a protective window of a material transparent to the wavelengths it is sensitive to (usually SiO_2 or silicone), which allows photons to enter the APD. Beyond this window is a *p-n junction*, located on the boundary of a p-doped region and an n-doped region. This is where the avalanche multiplication takes place. The largest region, marked in green in Figure 3.3, is known as the conversion, or π-region.

A photon that enters an APD is most likely to deposit its energy in the large π-region. This is a region within the APD that has been left undoped, and so it has a constant electric field of relatively low strength. When a photon enters this region, an electron-hole pair may be created if the photon energy is greater than the band gap energy of silicon (1.11 eV, λ=1117 nm). The electron will drift towards the positive cathode, and the hole towards the negative anode. As the electron nears the cathode, it enters the aforementioned avalanche region. As this is a reverse-biased p-n junction, there is a large depletion region within it. In this region the electric field increases greatly, enabling a single free electron to undergo rapid multiplication through the avalanche multiplication phenomenon. The electric potential applied across the terminals of a G-APD, called the reverse voltage, needs to be large enough to induce this phenomenon, which begins to take effect at a field strength of 2×10^5 V/cm. If the APD is operated below its breakdown voltage the number of avalanche electrons generated will be proportional to the number of electrons that initiated the avalanche.

3.2.2   APDs operated in proportional mode

When an APD is operated with a bias voltage lower than what is known as the breakdown voltage, it is said to be operating in the proportional mode. The gain of an APD operated in this manner is directly dependent on the applied bias voltage, since the avalanche multiplication factor increases with field strength. This is however only true up a certain point, as the avalanche-generated current must pass through the internal resistance of the APD and its load resistance. This will induce a voltage drop, reducing the bias voltage across the APD. For large bias voltages, this results in a non-linear behaviour where the signal amplitude is not proportional to the amount of incident light. Additionally, the APD gain also depends on temperature. This is because the substrate crystal lattice vibrates more intensely for higher temperatures. This acts as a "resistance" for the accelerating electrons in the multiplication region, which decreases the multiplication factor. If one wishes to maintain stable gain, it is therefore necessary to monitor temperature and adjust the bias voltage accordingly.
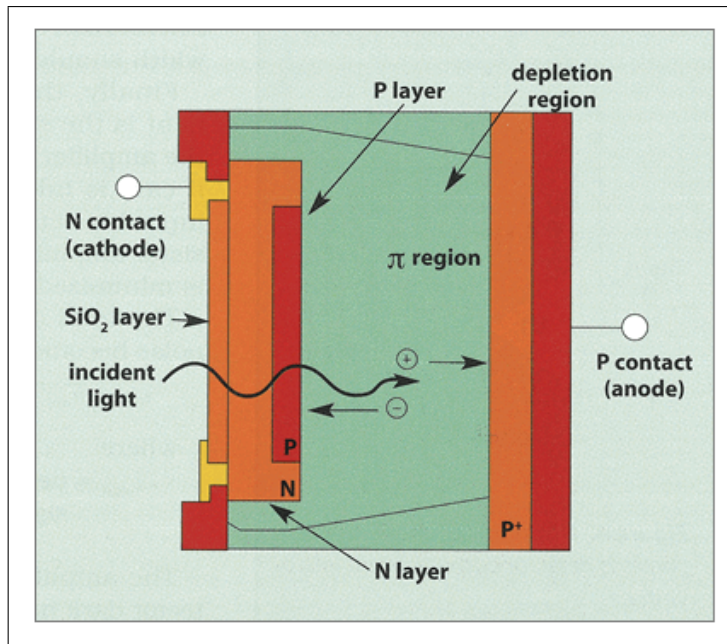
Figure 3.3: Cross section of an APD. Hamamatsu K.K. 2017 [31].

3.2.3   APDs operated in Geiger mode (G-APD)

Normally, an APD will be operated with a bias voltage slightly below its breakdown voltage to maximise gain while maintaining proportionality. In this region, the avalanche is self-quenching, as the reverse voltage is not great enough to allow the avalanche to develop through the entire substrate. The APD gain displays a strong dependence on bias voltage and temperature in this region [32]. Since an APD cannot be operated arbitrarily close to the breakdown voltage while maintaining proportionality, the gain is severely limited (<1000). This requires very low-noise amplifiers to read out and amplify the APD signal further

These limitations can be avoided by operating the APD slightly *above* the breakdown voltage instead. This is called the Geiger mode, and in this mode of operation the avalanche will not be self-limiting. Instead it diverges until all charge carriers in the avalanche region are free, and the G-APD resistance drops significantly. This leads to a constant current flowing through the G-APD. To stop the current, the bias voltage across the G-APD needs to be decreased to below the breakdown voltage. This can be done passively by adding a large ohmic resistance in series with the G-APD. The charge stored on the junction capacitance will flow as a current out of the G-APD when it discharges. The resistor, called a "quenching resistor" prevents the bias voltage supply from immediately recharging the G-APD. The discharge will continue until the voltage across the G-APD junction capacitor drops below the breakdown voltage, where the avalanche can no longer maintain itself and therefore stops.

The benefits of operating an APD in Geiger mode is that it produces a very well-defined output pulse of 10^5-10^7 electrons, depending on the applied bias voltage and junction capacitance. This can be thought of as the gain of the G-APD, and is much higher than the gain of an APD in proportionality mode. All information about the photon rate is lost however, since the G-APD will discharge entirely regardless of the amount of photons which entered it. It can therefore be thought of as an analogue to the Geiger-Müller Tube (GMT), a common radiation detector that

also function based on the same avalanche effect (albeit in gas, not a semiconductor material). In the same manner that a GMT has a deadtime set by the time it takes to recharge, so does the G-APD. The RC time constant, i.e. the time it takes the G-APD to recharge to 63 % of its nominal value, is on the order of a few tens of nanoseconds.

3.2.4   The SiPM

What we have now learned about the G-APD carries over to the SiPM, which is simply a collection of parallel G-APDs. Although the SiPM is a proportional counter, a single G-APD in the SiPM is not. However, by combining the signal of many G-APDs in parallel, and keeping the amount of incident photons low, they will together form a proportional counter. This is the operating principle behind the SiPM. The amplitude of the signal generated by a SiPM will be proportional to the amount of G-APDs that went into breakdown, i.e discharged or "fired". As long as the amount of incident photons is low enough that each photon impacts individual G-APDs, then the signal amplitude will be proportional to the total amount of photons that impacted the SiPM. This proportionality will naturally only hold if the probability that two or more photons impact the same G-APD within its deadtime of a few hundred nanoseconds is negligible. It is therefore important that a SiPM with enough G-APDs is chosen for the expected photon rate in the application.

3.3 SiPM properties

3.3.1 Signal amplitude and photon detection efficiency

When a G-APD fires, a well-defined charge Q_i is released. The amount of charge depends on the G-APD capacitance C and the overvoltage V_{over}, i.e. how many volts above the breakdown voltage it is operated at [33]:

$$Q_i = C(V_{bias} - V_{br}) = C \times V_{over} \tag{3.1}$$

If n G-APDs in a SiPM fires simultaneously, the charge released by the SiPM will then simply be the sum of their charges:

$$Q = \sum^n Q_i \tag{3.2}$$

A SiPM impacted by n photons will not fire n G-APDs, even when the number of photons is much lower than the number of G-APDs. This is because G-APDs does not have a 100 % Photon Detection Efficiency (PDE). The PDE is given by the Quantum Efficiency (QE) of the G-APDs, the fill factor ϵ, and the probability that an absorbed photon triggers a breakdown (P_{tr}):

$$\text{PDE} = \text{QE} \times \epsilon \times P_{tr}(V_{over}) \tag{3.3}$$

The quantum efficiency QE is at most 80%-90%, and will depend on the photon wavelength. It represents the probability that a photon is converted into photoelectrons in the substrate of the G-APD. The fill factor ϵ represents the fraction of SiPM active area to non-active area, and is less than 1 since some space is required between each G-APD to optically and electrically insulate them. Additionally, some of the surface area is also covered by the bias voltage grid and quenching resistors. Generally the fill factor decreases with decreasing G-APD pitch, and can be anything from 30 % to 80 %. P_{tr} depends on where in the G-APD the photon was converted to an electron-hole pair, and its value will increase with increasing bias voltage [24].

A typical PDE curve is shown in Figure 3.4. We see that the efficiency peak is narrow, so when choosing a SiPM, the wavelength of the photons one wish to detect must be considered.

Taking into account the PDE and finite number N_G G-APDs in a SiPM, the number of firing G-APDs of a SiPM impacted by N_γ photons is therefore [24]:

$$N_{fired} = N_G \left(1 - e^{\frac{N_\gamma \times \text{PDE}}{N_G}} \right) \tag{3.4}$$
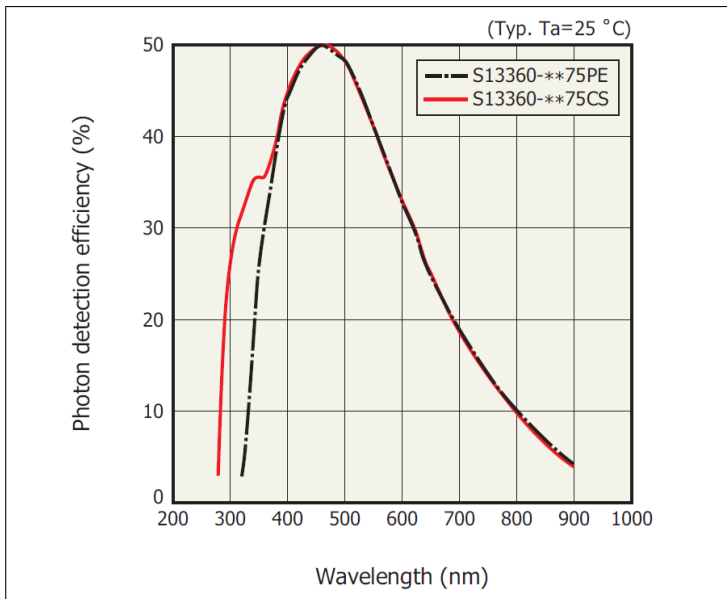
Figure 3.4: Photon detection efficiency of two typical SiPMs at 25°C. From [23].

3.3.2   Dark rate

Any type of electron-hole pair production in the G-APD will initiate a discharge, not only those generated by absorbed photons. Thermal excitation and electric field-assisted tunnelling can generate free carriers in the conversion or avalanche region. This causes the SiPM to produce signals even when in the absence of photons, called dark noise. The rate of which these signals occur is called the dark rate or dark count, and it will depend on temperature, the area of the SiPM, and the bias voltage. One study found that, for a $0.43\,mm^2$ SiPM, the dark rate could be decreased drastically by cooling it with liquid nitrogen and decreasing the bias voltage. This lowered the dark rate from 1MHz at 25°C to 40Hz at -196°C for the same overvoltage [34].

3.3.3   Afterpulsing and crosstalk

During the avalanche breakdown of a G-APD, an average of 2.9 photons with an energy higher than 1.14eV (the band gap of silicon) is generated for every 10^5 charge carriers [35]. These photons may induce new breakdowns in neighbouring G-APDs. This effect is called crosstalk, and introduces noise in the SiPM signal. Manufacturers attempt to inhibit this effect by designing optically insulating grooves between the G-APDs. It can also be decreased by operating the SiPM at lower overvoltage.

Afterpulsing is an effect that can be observed as additional pulses occurring after the primary pulse. It comes from metastable trapping of avalanche carriers, and their subsequent release after the G-APD has had time to recharge. Fast releases that occur before the G-APD has recharged have no effect, but some may be released several microseconds after being trapped, inducing a new breakdown. As with dark noise, this unwanted effect can be reduced by decreasing the overvoltage.

3.3.4   Temperature dependency of breakdown voltage

The SiPM gain, i.e. the amount of charge that is released by a single G-APD as it discharges, is linearly dependent on the overvoltage applied to the SiPM. In addition to this, the gain also has a certain dependency on temperature. This is due to the fact that the breakdown voltage V_{br} of a PN-junction increases with temperature [36, p.3]. For a G-APD with a constant applied bias voltage, an increase in V_{br} reduces the overvoltage. This in turn reduces the gain, as can be seen from Equation 3.3.1. Therefore, if one wishes to maintain a stable gain for a SiPM, the applied bias voltage must be increased proportionally to its temperature in order to maintain a constant overvoltage.

The dependency of V_{br} on temperature is on the order of 10 to $100\,\mathrm{mV}\,^{\circ}\mathrm{C}^{-1}$. Since it is an absolute change, the closer the SiPM is operated to V_{br} the greater of an effect a variation in temperature will have on the gain. A typical SiPM may have a recommended overvoltage of $3\,\mathrm{V}$ and a temperature dependency of $50\,\mathrm{mV}$, which results in a relative gain change of $1.7\,\%\,^{\circ}\mathrm{C}^{-1}$ for small deviations in temperature. But if the designer instead chooses to operate the SiPM at a lower overvoltage of $0.5\,\mathrm{V}$ in order to decrease the dark rate, the relative relative gain change will then be 6 times larger, $10\,\%\,^{\circ}\mathrm{C}^{-1}$. In this case, a mere ten degrees increase in temperature will be enough to bring the bias voltage below V_{br}.

It is therefore clear that the temperature dependency of V_{br} must be taken into account when designing an application that uses SiPMs. The accuracy to which the SiPM temperature can be determined and the bias voltage controlled will set a limit on the gain accuracy and how near V_{br} the SiPM can be operated.

3.4   Comparison to PMTs

One immediately apparent benefit of the SiPM is its small size. This is positive in applications and fields where physical space is a commodity, such as in medical imaging and satellites. A disadvantage to this is that if an area larger than $6 \times 6\,\mathrm{mm}^2$ must be covered, multiple SiPMs have to be mounted in an array and read out either individually or via some multiplexing scheme.

Another important property of the SiPM is its excellent timing resolution. Sub-nanosecond timing is easily obtained [37, 38], with single-photon timing shown to be less than 70ps for some SiPM models [39, 24]. This naturally makes SiPMs attractive for use in Time Of Flight (TOF) applications where the photon arrival time resolution is key. SiPMs have in particular been explored for use in Positron Emission Tomography (PET) where the addition of TOF information allows excellent rejection of undesired effects such as random coincidence events and scattered photons [38]. Much research has been done on the potential use of SiPMs in TOF-PET detectors, with promising results [40, 41].

Other benefits of SiPMs include:

- Low bias voltage (~50V).
- Insensitive to magnetic fields up to 15T.
- Compact and sturdy.
- Cheap compared to PMTs for small light collection areas.
- Low power consumption.
- Not damaged by exposure to ambient light.
- No hysteresis effects (no "memory" of previous configuration).
- Very low nuclear counter effect (low sensitivity to ionizing radiation).
- High quantum efficiency (80%-90%).
- Immune to extreme conditions (temperature, humidity, pressure).
- Very low afterpulsing probability (~1 % for modern SiPMs).

We see that the SiPM excels in many regards when compared to PMTs. PMTs require high bias voltages on the order of 1kV, are sensitive to magnetic fields, expensive to manufacture, have low quantum efficiency, and worse timing resolution.

In many cases, SiPMs are the better choice when it comes to low-light detectors. One might therefore think that PMT manufacturers are going to go out of business, but at least for the time being that is not true. While SiPMs do have many clear benefits compared to PMTs, they also have some disadvantages:

- Small light collection area (currently 36 mm^2 maximum).
- Narrow photon frequency sensitivity.
- Gain has strong dependency on temperature and bias voltage.
- High dark count rate (~1000 times higher than PMT).
- Low gain variability (2-4 vs. ~1000 for PMT).
- Dynamic range limited by number of G-APDs.

PMTs have the benefit of high dynamic range, gain variability, wide photon frequency range, and are cheap compared to SiPMs when large areas must be covered. Of special importance when it comes to single-photon detection is the very low dark count of PMTs, which can only be matched by SiPMs through significant cooling. SiPMs are not the best solution in all cases, and the benefits and disadvantages of both SiPMs and PMTs must be weighed for each application where low-light detectors are needed.

For the upgrade of the CRT, SiPMs were chosen primarily due to their low price and non-lethal bias voltage. Their ruggedness and insensitivity to ambient light exposure also helped reduce the probability of damaging them during testing. Lastly, the small size of SiPMs made measurements in complete darkness easier, as our largest blackbox was barely able to fit a scintillator slab with light guides and SiPM readout electronics mounted. The large sizes of PMTs would have made this impossible without modifications to the blackbox, or by building an entirely new one.

3.5   SiPM model and parameters

F.Corsi et al. developed an electrical equivalent model of the SiPM in 2007 [33]. A single G-APD in this model is represented by the junction capacitance C_d in series with a quenching resistor R_q, and a small parasitic capacitance C_q in parallel with the resistor. C_d is the capacitance of the G-APD, and is on the order of a few hundred femtofarad. It is the charge stored on this capacitance that is released when a G-APD fires. R_q is on the order of hundreds of kiloohm, and the parasitic capacitance between it and the substrate it is in contact with is modelled by the parasitic C_q, normally on the order of tens of femtofarad. A current source in parallel with C_d models the charge delivered when the G-APD fires. The electrical equivalent model is seen in Figure 3.5. It shows the special case when a single G-APD fires, but could easily be extended to the case where n G-APDs fire by calculating the component values for n parallel G-APDs. The unfired G-APDs are represented by their total parallel capacitances and resistances. An additional parasitic capacitance C_g is required to take into account the metal grid that lays on top of the SiPM and connects the quench resistors to the common cathode. R_s represents the resistance of the front-end electronics as seen by the SiPM.

This model has been analysed, verified, and shown to be extremely accurate if a physically correct current source model is used [42, 43]. The original model by F. Corsi et al. used a Dirac delta pulse current source, which was also shown to produce somewhat accurate SiPM pulses. But we found that using Dirac delta pulses did not fulfil our needs, as it introduced spikes in the waveform and could not accurately recreate the observed SiPM waveforms. Therefore a physically correct current source model was used instead, as in the two references cited before.

The SiPM chosen for the CRT was the Hamamatsu S13360-6075CS Multi Pixel Photon Counter (MPPC) [23], which has the following specifications:

- Area: 6x6 mm^2
- Number of G-APDs: 6400
- Total terminal capacitance: 1280 pF
- Fill factor: 82 %

The model parameters (R_q, C_d, C_g, and C_q) for this MPPC were needed in order to have a physically correct electrical equivalent model for use in the design of the MPPC amplifier. Accurately determining the model parameters for a specific SiPM is possible, and the approach has been described by F. Corsi et al. [33] and Stefan Seifert et al. [43]. However, the process for determining any parameter except the quench resistor value is somewhat involved. Therefore all parameters of the MPPC, except the quench resistor value, were estimated.

3.5.1   Determination of quench resistor value

The quench resistor value R_q can be measured by forward biasing the SiPM and measuring the magnitude of the current that flows through it as a function of bias voltage. Since an APD shares many of the characteristics of a diode, it will follow typical diode behaviour. Therefore,
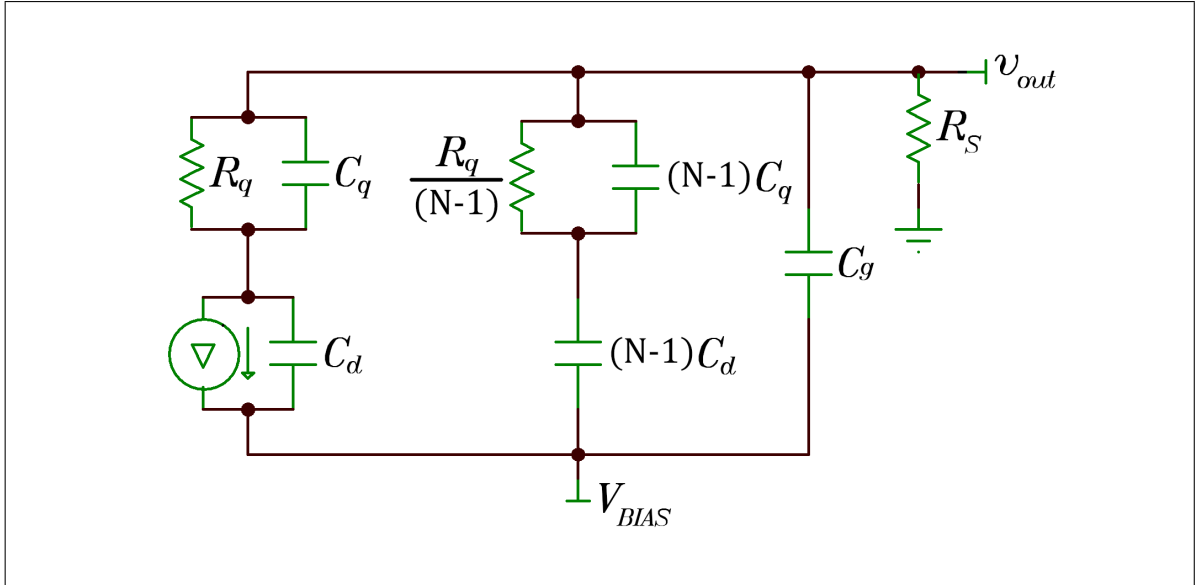
Figure 3.5: An equivalent electrical model of the SiPM. Recreated from [33].

the SiPM will become conductive when the forward bias is greater than the diode voltage drop $\sim650\,$mV. In this mode the resistances of the APDs in the SiPM are negligible compared to the quench resistors, and the change in current passing through the SiPM (ΔI) will depend solely on the combined parallel resistance of all the quench resistors and the change in forward voltage (ΔV_{FW}), from Ohm's law:

$$\Delta V_{FW} = \frac{R_q}{N_{total}}\Delta I$$

$$\Delta I = \Delta V_{FW}\frac{N_{total}}{R_q}$$

(3.5)

Measuring the current and plotting it as a function of voltage produces an IV-curve. To make these measurements for the S13360-6075CS MPPC, a TTI QL355TP power supply was used to generate the forward bias. The current was monitored by a Keithley 2100 multimeter to an accuracy of $\pm4.5\,\mu$A, and the power supply output voltage by an Aim TTI 1604 multimeter to an accuracy of $\pm0.8\,$mV. The results are seen in Figure 3.6. Performing a linear fit on the data points corresponding to the MPPC being in forward mode shows that the parallel quench resistors value is $69\pm0.6\,\Omega$.

However, the measurement setup has a non-negligible parasitic series resistance that must be taken into account. This comes primarily from the Keithley 2100 shunt resistor, which has a value of $5.1\,\Omega$ according to the datasheet. A small amount of resistance will also manifest in the wires and contacts. The total parasitic resistance was measured by first circumventing the MPPC by attaching both bias voltage wires to the same MPPC lead. This ensured that the contact resistance remained the same as when the IV-curve measurements were made. The voltage was then increased until a current of $1\,$mA was flowing, comparable to the current observed during the MPPC measurements. A total series resistance of $6.2\pm0.1\,\Omega$ was measured. The true value of the parallel quench resistors is therefore $62.8\pm0.7\,\Omega$. This means that the
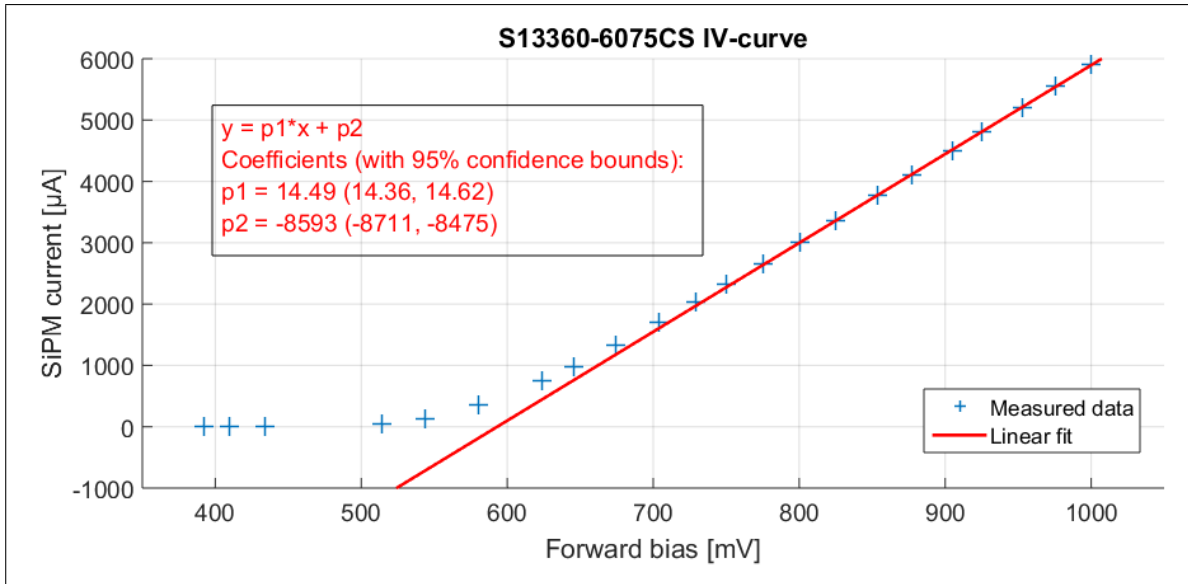
Figure 3.6: Measurement of the IV-characteristics of the Hamamatsu S13360-6075CS MPPC. The data points used for the linear fit was for a forward bias of 775 mV-1000 mV

individual quench resistors have a resistance of:

$$\frac{R_q}{N_{total}} = 62.8 \pm 0.7\,\Omega$$

$$R_q = 62.8 \pm 0.7\,\Omega \times 6400$$

$$R_q = 402 \pm 5\,\text{k}\Omega \tag{3.6}$$

It was later realized that a better method would have been to measure the voltage directly across the MPPC leads instead of at the power supply outputs, as this would have circumvented the need to estimate the voltage drop due to parasitic resistances. But the result found by the method described still resulted in a value for R_q that was deemed to be sufficiently accurate.

3.5.2 Estimation of other parameters

The capacitance parameters C_q (quench resistor parasitic capacitance), C_d (G-APD diode capacitance), and C_g (bulk parasitic capacitance) can be estimated by considering typical parameter values determined for other SiPM models [44, 45, 46, 33, 43]. The fact that we know exactly how many G-APDs our SiPM consists of, and the total terminal capacitance, will also help.

An estimate for C_g was found by assuming that 5 % of the SiPM area is covered by the metal grid that distributes the bias voltage. This assumption is supported by microscope images of SiPMs, which shows that most of the non-active area is covered by quenching resistors or open space to optically insulate the G-APDs. By using a typical value for the metal-substrate capacitance of $30\,\text{pF}\,\text{mm}^{-2}$, we get an estimate of $30\,\text{pF} \times 0.05 \times 6\,\text{mm} \times 6\,\text{mm} = 54\,\text{pF}$. Rounding this off to 50 pF leaves 1230 pF for the G-APD capacitance.

Parameter	Value
C_d	173 fF
C_q	19 fF
C_g	50 pF
R_q	402 kΩ

Table 3.1: Measured and estimated model parameters for the Hamamatsu S13360-6075CS MPPC.

The remaining capacitance will be distributed evenly on the 6400 G-APDs. That results in 1230 pF/6400≈192 fF per G-APD. The G-APD capacitance will be distributed on C_d, the diode capacitance, and C_q, the quench resistor parallel capacitance. The fraction of capacitance distributed on C_q was estimated by considering typical SiPM parameters as measured in the articles cited above. It was found that C_q typically contributes 10%-20% of the total G-APD capacitance.

To summarize, the estimated parameters for the S13360-6075CS MPPC are listed in Table 3.1. These values were used to model the MPPC in Toolkit for Interactive Network Analysis (TINA). TINA is a analogue electronics simulation program based on Simulation Program with Integrated Circuit Emphasis (SPICE). Texas Instruments offers a free version of this program, called TINA-TI, and this version was used extensively in this thesis for the purpose of simulating the MPPC and the electronic circuits that amplifies and shapes its signal.

3.6   Pulse timing improvement

The pulses that the MPPC generates are meant to be amplified and pass through a pulse height discriminator. This discriminator will trigger on a MPPC pulse if it crosses its trigger threshold, outputting a square pulse that signifies a MPPC pulse of sufficient amplitude was registered.

We saw in chapter 2 that the scintillation photons generated by a charged particle crossing the scintillator will not arrive at the MPPC simultaneously. The simulations indicated that they would arrive in "groups" spaced 10 ns to 15 ns apart due to photons being reflected and scattered inside the scintillator. This effect was also observed experimentally, and presented a problem: How do we prevent the discriminator electronics from triggering on the second group of photons? When the second group arrives it can increase the SiPM signal amplitude, "riding" the signal of the first pulse to reach the trigger threshold of the discriminator even if the first pulse did not. An example of such a pulse is shown in Figure 3.7. This will result in greatly reduced timing accuracy, since one can never know if the discriminator triggered on the first or second wave.

A method was needed for reshaping the pulse so that the discriminator will only trigger on the first pulse. A simple yet effective way to do this is by passing the MPPC signal through an amplifier designed as a *bandpass filter*. Such a filter will suppress frequency components the further they are from some center frequency, and if that frequency is chosen to be comparable to the main frequency component in the rising edge of the MPPC pulses, it will result in only the rising edges being amplified. The filtered and amplified pulses will be narrow and have a relatively well defined risetime, which improves timing. And most importantly, the secondary pulse will not raise the MPPC amplitude over the trigger threshold if the first pulse did not.

To find the center frequency best suited for such a filter, a collection of typical MPPC pulses were sampled. The pulses were imported to TINA-TI and applied to narrow second-order bandpass filters with different center frequencies f_C. Two typical pulses and their resulting waveforms after filtering can be seen in Figure 3.8. Setting f_C to a high value results in faster risetimes and thus better timing, but worse Signal to Noise Ratio (SNR) since the filtered pulse amplitude is comparable to the background noise. Setting f_C too low results in large signals and good SNR, but the timing suffers from slow risetimes. From quantitative analysis of many different sampled MPPC pulses it was decided that the filter should have a center frequency between 40 MHz-80 MHz. The exact value was not determined, which provided some freedom when designing the amplifier chain.
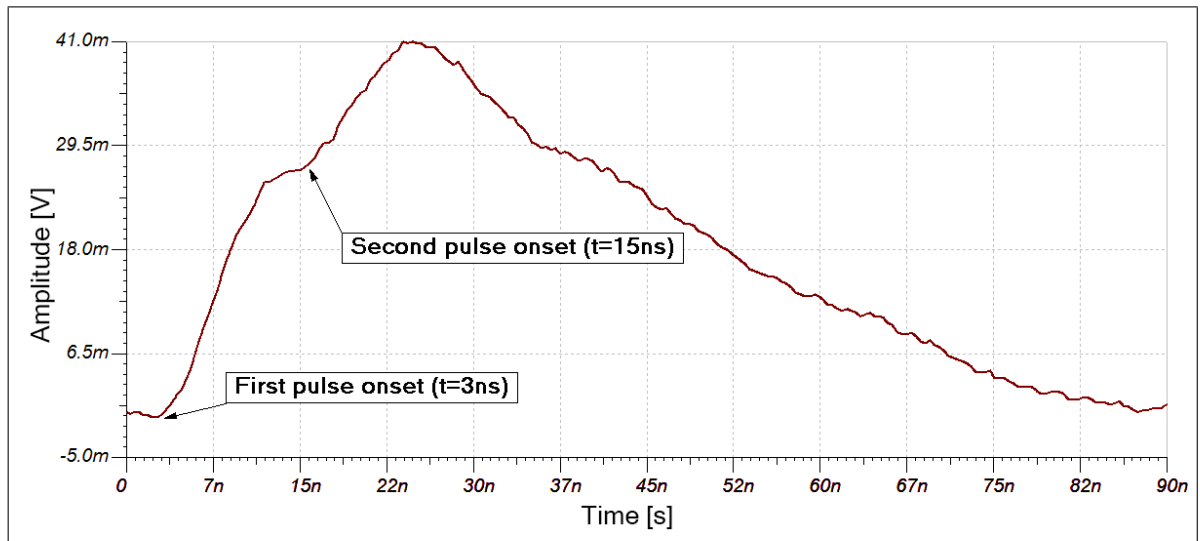
Figure 3.7: Measured MPPC pulse showing how the second group of photons can increase the pulse amplitude.
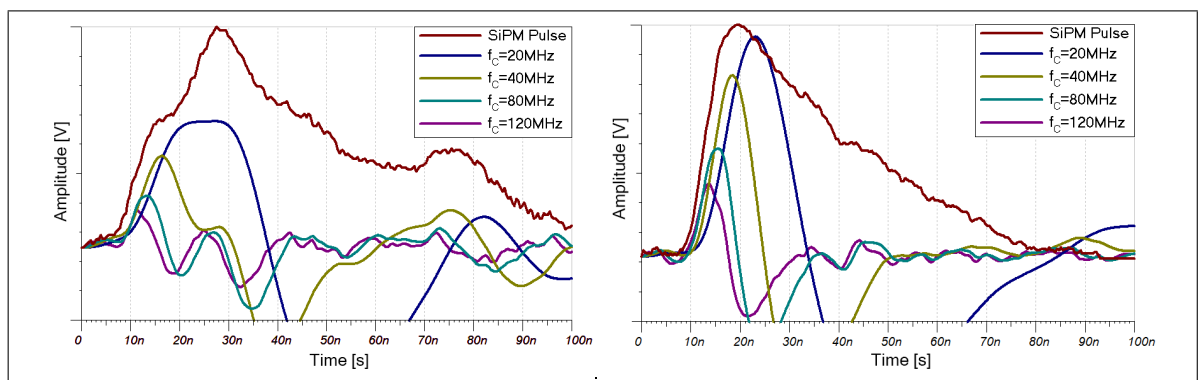
Figure 3.8: Two typical MPPC pulses and the resulting waveform after passing through a second-order narrow bandpass filter with varying center frequencies f_C. The resulting waveforms have been upscaled 200 % to better compare them to the original waveform.

4 | The Readout Board

The current chapter will present the electronics that were designed for the purposes of controlling the MPPC and processing the signals it produces. This Readout Board (ROB) is capable of performing automated temperature-based bias voltage compensation of the MPPC, is controllable from a PC via a custom-designed communication protocol, and will amplify, shape, and trigger on the MPPC pulses. The first part of this chapter describes the design and operation of the ROB, while the last section presents measurements performed to validate its performance.

Throughout the chapter, simplified schematics are presented to illustrate the function of certain circuits. If the reader feel that they are lacking in detail, the complete circuit diagrams of the ROB are found in Appendix D.

4.1   Introduction

As any engineer, scientist, or experimental particle physicist will tell you, simply having a sensor in your possession won't get you very far in your task to measure the quantity that you are interested in. Some amount of supportive elements are required to get the data we desire, enabling the sensor to perform at its maximum potential, and to minimize uncertainties in the measurements. The amount of supportive elements and their complexity normally increases when the signal that is measured decreases. The desired precision of the measurements will also play a large role in choice and configuration of the supportive elements. In the present case we are trying to measure a small charge deposited in a short timespan. The charge released by the SiPM in a typical muon event is on the order of 10^{-12} C, and the charge is deposited on the timescale of nanoseconds. This forces certain requirements on the readout components to have low noise and high bandwidth so that the signal will not disappear in noise or end before the electronics have had time to react.

For the sake of simplicity, the choice was made to implement all the supportive elements for the MPPC on a single Printed Circuit Board (PCB), called a ROB. A single PCB lowers the cost of manufacture, eases assembly, and makes it more practical to set up when doing experiments. The ROB required the following capabilities:

- Collection, amplification, and shaping of the MPPC signal.
- Triggering on the shaped MPPC signal and generating a pulse for off-board processing.
- Monitoring MPPC temperature and adjusting its bias voltage.
- Controlling and monitoring on-board voltage regulators.
- Communication with the CRT control software.

This collection of capabilities presented some design problems, as required a combination of high-speed analogue electronics, high-gain high-bandwidth amplifiers, and digital electronics in a physically small space. Therefore great care was taken to minimize the effect of interference on the analogue circuit elements arising from the digital circuit.

The goal was to make the readout electronics as robust as possible, to minimize the need for future interventions or replacements. Therefore, the following guidelines were followed during the design process:

- Self-containment
- Safety
- Flexibility
- Ease of use
- Performance

The first and second point relates to the readout electronics being a complete, self-contained unit. That means regulating and filtering the raw voltages being supplied to it via the power connectors, enabling and disabling board components and voltages in a safe manner, automatically adjusting parameters, and monitoring parameters to quickly react to problems that might damage components. These tasks are all performed by an on-board Atmel Microcontroller Unit

(MCU), enabling anyone with some knowledge of programming to modify and improve the detector board. The third point, flexibility, is taken care of by a communication scheme that connects all the detector boards to a computer. This enables a user to quickly enable, disable, and set configurations of any set of detector boards, making it very simple to operate the CRT. This capability also ties into the fourth guideline, easy of use, as it makes the ROB simple to control. And since the MCU handles the various automated tasks, such as adjusting the MPPC bias voltage, the user need not handle this themselves. The last guideline, performance, is a wide term that covers the general performance of the ROB. It means that the ROB is capable of accurately measure its on-board voltages and the MPPC temperature, precisely adjust the bias voltage and trigger threshold, prevent excessive noise on the MPPC signal, work within its temperature limits, etc. These requirements are taken care of by paying close attention to component selection, circuit design, and PCB layout to minimize unwanted effects and ensure the ROB will function within the desired specifications.

4.2   Overview

The PCB with all its supportive elements is called a ROB. It is a 10x10cm four-layer PCB that contains all the elements required for autonomous operation of the MPPC. It incorporates a signal amplification chain, test charge inserter, discriminator, voltage regulators, MPPC bias voltage temperature correction, and an MCU. An overview of the ROB can be seen in Figure 4.1. The main blocks are:

- **Preamplifier**: The preamplifier collects the weak signal generated by the MPPC and performs initial amplification and shaping. The preamplifier also shields the high capacitance of the MPPC from the second amplifier so that it can operate as intended. The preamplifier was designed to be low noise and has been shielded from Electromagnetic Interference (EMI).
- **Primary amplifier**: The primary amplifier receives the MPPC signal to amplify it further, and to shape it for the discriminator. It is designed as a narrow bandpass filter centered on the characteristic frequency of the rising edge of the MPPC signal to suppress noise and amplify only MPPC signals.
- **Discriminator**: The discriminator will trigger on an MPPC signal if it reaches a certain user-settable amplitude called the trigger threshold. It is designed as a pulse-extender, and on a trigger it will generate a square pulse with duration 50 ns. This pulse is supplied to a line driver that outputs it over coaxial cable.
- **Voltage regulators**: The ROB regulates and filters the noisy off-board supply voltages to generate the various low-noise voltages required by the on-board components. Each regulator is controlled and monitored by the MCU.
- **Microcontroller**: The MCU controls the ROB, ensuring it is operated in a safe manner that will not damage the components. It also monitors the voltage regulators and MPPC temperature, and allows the user to control the ROB remotely with a computer.

The signal chain is the part of the ROB which is responsible for the amplification, shaping, and discrimination of the MPPC pulses. These are the primary capabilities of the ROB, and all other components and systems are simply supporting its operation. A block diagram of the MPPC signal chain can be seen in Figure 4.3.

Each sub-system of the ROB is described in further detail in the follow sections, and the communications protocol and SiPM bias regulation are discussed in separate sections.
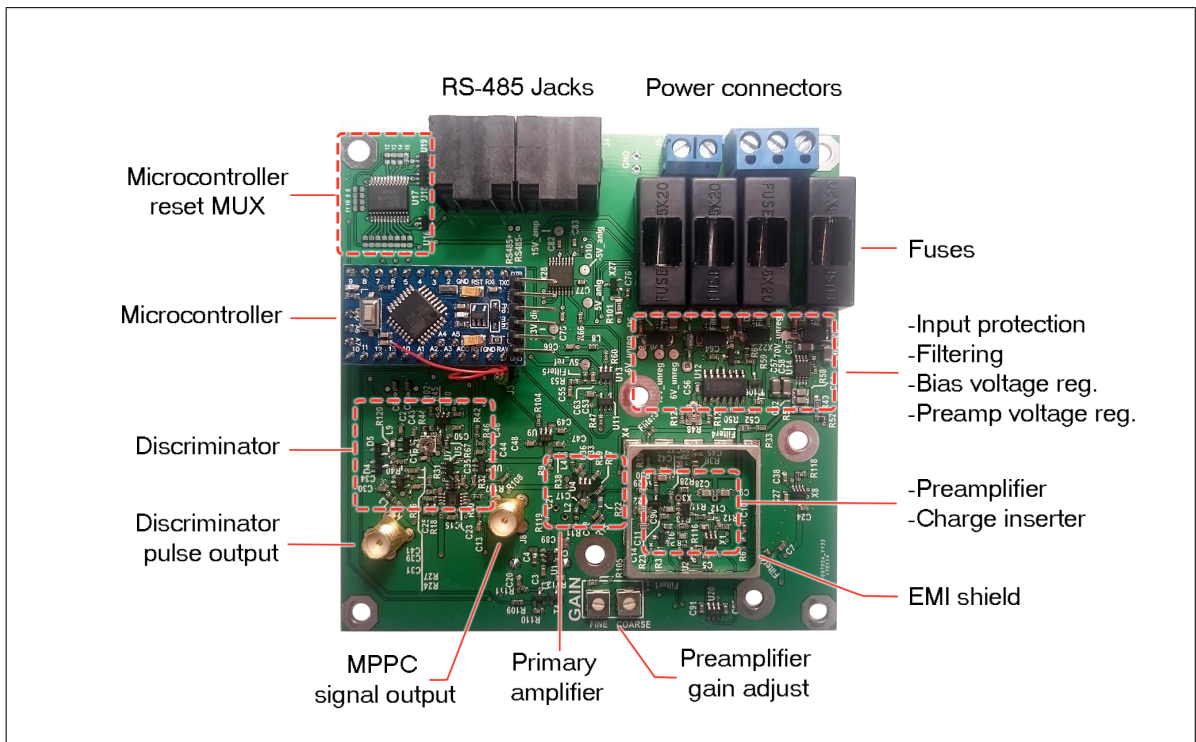
Figure 4.1: A top-level overview of the readout board. Each component block have been labelled to note their function.
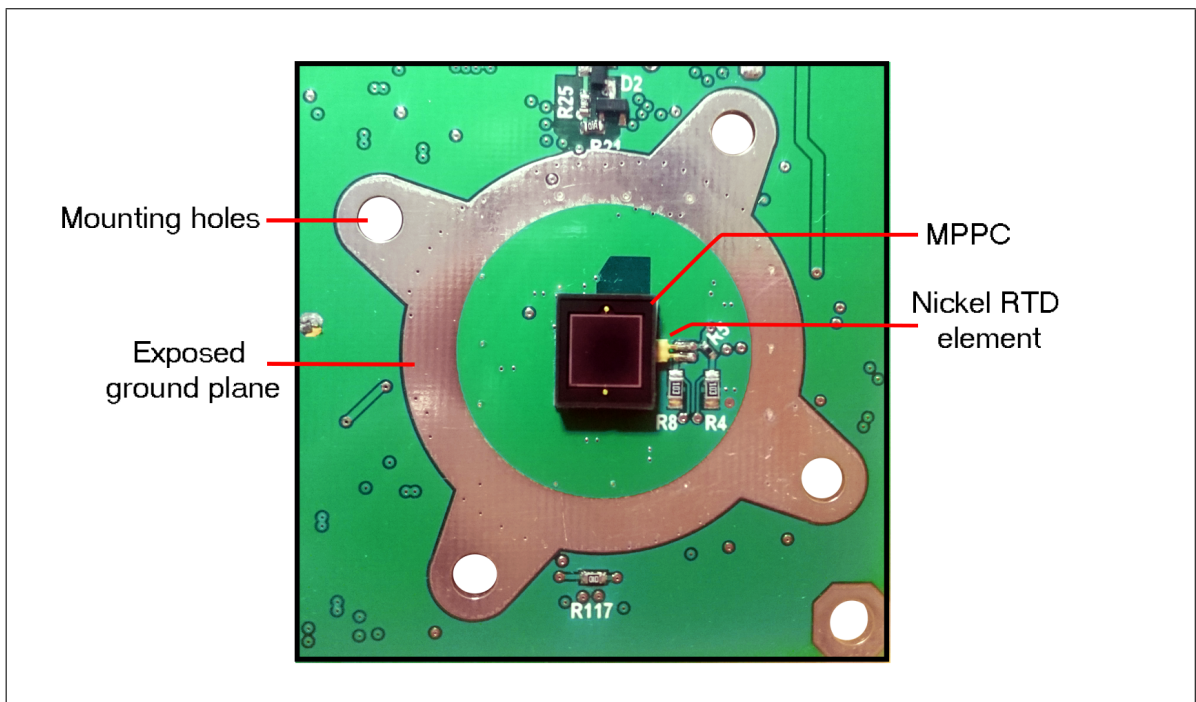
Figure 4.2: Backside of the ROB, showing the MPPC and temperature sensor (RTD). The exposed ground plane is in contact with the circular metal frame that mounts the ROB to the scintillator light guide, improving EMI shielding.
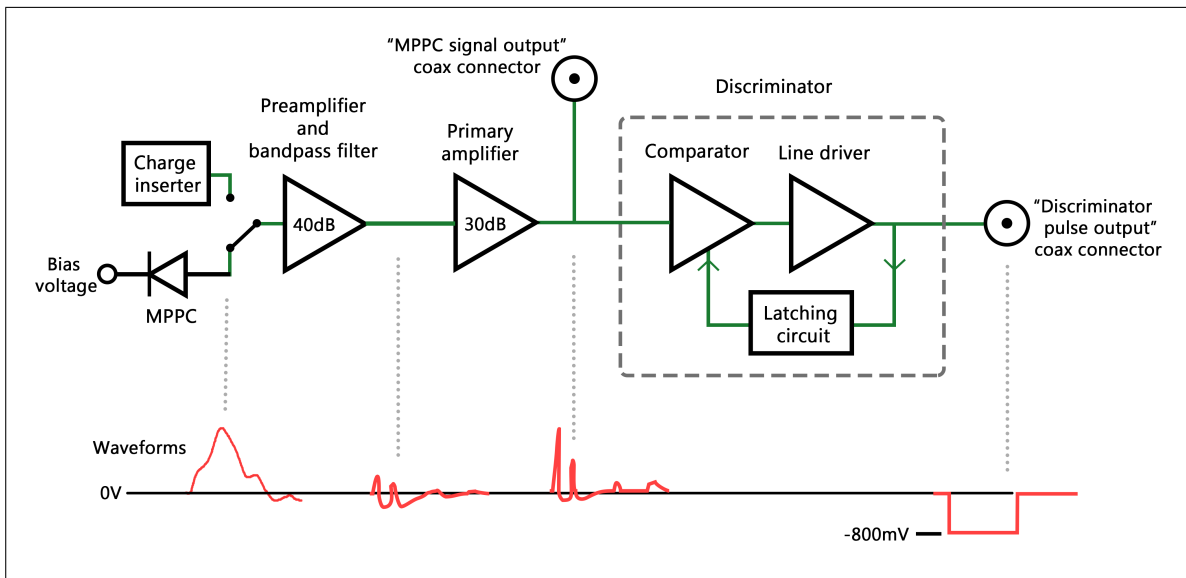
Figure 4.3: Block diagram of the MPPC signal chain located on the bottom half of the ROB in Figure 4.1. Typical waveforms have been added to illustrate the functions of the various elements.

4.3   Preamplifier

The preamplifier is by far the most critical component of the MPPC signal chain. It is responsible for the first amplification of the weak MPPC signal and must therefore be carefully designed to not to degrade the signal. These degradations can manifest themselves as noise, oscillations, the loss of frequency components, non-linearities, or all of them at once.

The need for a dedicated preamplifier becomes clear when one looks at the characteristics of a SiPM. There are two main characteristics that voids the use of traditional operational amplifier (op amp) schemes as used in PMT amplifiers. The first is the fact that SiPMs have a much larger terminal capacitance than PMTs. This can be understood by looking at the physical structure of a SiPM, which consist of a multitude of G-APDs, with each G-APD having a small capacitance. This capacitance is a combination of the diode capacitance of the G-APD, and parasitic capacitances. The total capacitance of a SiPM is directly tied to its area, and a larger detector will therefore have a larger capacitance. For the 6x6mm SiPMs used in this thesis, the anode-cathode capacitance is 1280 pF, and if it was simply connected to a traditional PMT amplifier, its performance would be seriously degraded. When using PMTs, a dedicated preamplifier is usually not required. A firing PMT can be modelled as a near ideal current source [47], and has a low terminal capacitance of 3 pF-10 pF. These characteristics makes amplification simple, and for this task one can use operational amplifiers in a variety of configurations such as current-to-voltage, charge-sensitive, or trans-impedance without giving too much thought to degradation of the signal.

There is also another cause for frustration. It has been observed that the signal from an MPPC is not linearly proportional to the amount of G-APDs that fired if the MPPC output is connected to a resistive load. This effect was explained and recreated in simulations by Stefan Seifert et al. [43]. The short explanation is that when the amount of firing G-APDs increases, the parallel combination of the quench resistors of the firing G-APDs decreases. Since the quench resistors are in series with the unchanging load, that has the effect of increasing the effective quench resistor value. This causes the avalanches to be quenched earlier, and less charge is released by the G-APDs. Therefore, if a SiPM is connected to an amplifier with a fixed non-zero input impedance, a non-linear relationship between pulse height and the amount of G-APDs fired will materialize.

Therefore, normal amplification schemes cannot normally be used directly with the SiPM. Instead, a preamplifier is needed, which serves two purposes: It should shield the main amplifier from the high terminal capacitance of the SiPM so that it can operate efficiently, and present a small input impedance to the SiPM to minimize non-linear effects. Not many solutions that fulfil these requirements exist. One method is to use a Current Feedback (CFB) op amp [48]. Due to their construction, the inverting terminal has an inherently low input impedance of typically 50 Ω to 100 Ω, with some commercially available devices having as low as 2.5 Ω. The effective input impedance is lowered further when the amplifier feedback is taken into account. However, it is not constant with frequency, and will increase to non-negligible values at higher frequencies. When a SiPM is connected to the frequency-dependent non-zero input impedance of a CFB

op amp, they form a reactive impedance that peaks the gain at higher frequencies and thus produce undesired oscillations on the output signal. It is however possible to create a SiPM preamplifier with this arrangement if care is taken to keep the gain peaking under control by using SiPMs with low terminal capacitance and CFBs with stable, small input impedances over the frequency range of the SiPM signal [49].

Another solution is to use a regular NPN transistor in what is called a *common base configuration*. Such a scheme was explored by J. Huizenga et al. [49], who showed that a common-base transistor amplifier could be used as a low-noise high-bandwidth preamplifier for SiPMs. This method was adopted for the design of the ROB preamplifier, since the SiPMs used in this thesis had a large terminal capacitance which made them unsuitable for a CFB preamplifier scheme.

4.3.1   The transistor

To understand how the common base amplifier functions, one must first understand the transistor. Tomes have been written on the theory of transistors, and this thesis can only give a short introduction to the subject, hopefully enough to make the reader sated without feeling overwhelmed.

The transistor is an immensely important electrical component that paved the way for the technological society we enjoy today. It was invented in the late 40's and the first commercial transistors based on silicon came to the market in 1954. The historical, economic, and societal importance of the transistor is not in the scope of this thesis, but suffice to say we would be living in a very different world without it, as it is the basis of nearly every electrical contrivance in our daily lives.

There are two main families of transistors, Bipolar Junction Transistors (BJTs) and Field Effect Transistors (FETs). The difference is in whether it is controlled by a current or voltage. We will ignore the FET, as the transistor used for the preamplifier is a BJT. A schematic image of a BJT is shown in Figure 4.4.

In short, a BJT can be thought of as switch whose state is controlled by a current I_B applied to the terminal of the transistor called the base. Depending on the magnitude of the base current, the current I_C that flows from the collector terminal to the emitter terminal can be controlled. The structure of a BJT resembles two diodes sharing a common terminal. A diagram of this is seen in Figure 4.5, which shows an NPN-type BJT. It is called so due to the fact that it consists of a p-doped region sandwiched between two n-doped regions. The base is connected to the shared diode region, with the collector and emitter connected to each end of the coupled diodes. In the absence of a base current, one of the two diodes will be reverse-biased regardless of the emitter and collector polarities, and no current flows through the device. Biasing the base with a voltage V_{BE} will draw electrons from the emitter region into the base region. Some electrons will recombine with the holes in the base and be extracted, but most of the electrons will be drawn to the collector region, and a current I_C flows.
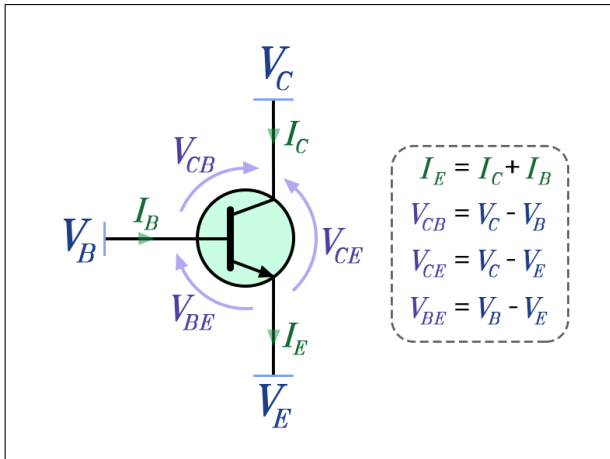
Figure 4.4: Illustration of a transistor, with voltages and currents marked. V_{CB}: Collector voltage with respect to base voltage. V_{CE}: Collector voltage with respect to emitter voltage. V_{BE}: Base voltage with respect to emitter voltage.
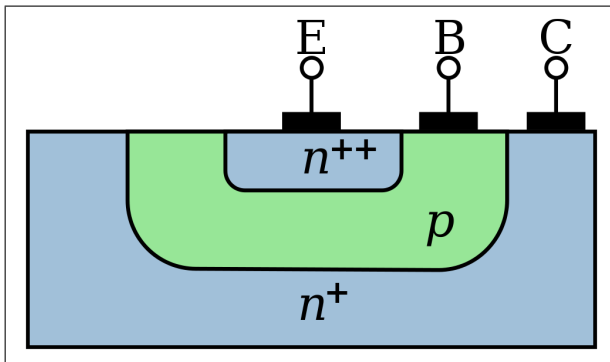
Figure 4.5: Illustration of a planar NPN-type transistor. A lightly p-doped region (the base) separates two n-doped regions. The smaller is strongly doped, and is designated as the emitter. The large region is less doped than the emitter, and acts as the collector. The asymmetry allows large values of β, the current amplification factor.

The type of transistor used in the ROB preamplifier is a heterojunction NPN BJT, a subtype in the BJT transistor family. These types of transistors can handle very high frequencies, up to several hundred GHz, and is therefore often used in Radio Frequency (RF) systems. For this chapter, "transistor" will be used to mean an NPN BJT.

A transistor can operate in four different modes: Reverse Active, Forward Active, Saturation, and Cut-Off. A transistors behaviour is very different based on which state it is in, and the state is determined by the relation of the three input voltages V_B, V_E, and V_C.

Cut-Off

When a transistor is in the Cut-Off mode, it acts as an open circuit. No current flows from the collector to the emitter. The Cut-Off mode is reached by setting the base voltage at a value so that the following relations are true:

$$V_{CB} > 0 \qquad V_{BE} < V_{th} \tag{4.1}$$

Here, V_{th} is the threshold voltage for the transistor. Its value normally lies around 0.5-0.6V, i.e. one diode junction voltage drop. If V_{BE} is below this threshold, the transistor will not conduct current from the collector to emitter, and it acts as an open circuit.

Saturation

When a transistor is in Saturation mode, it closely resembles a short-circuit between the collector and emitter. This mode is reached by setting the base voltage higher than *both* the emitter and collector voltage. But since the transistor is not perfect, there will be a small voltage drop from collector to emitter, so to actually get a transistor into the Saturation mode V_C needs to be a little higher than V_E. In short, the following relations must be true:

$$V_{CB} < 0 \qquad V_{BE} > V_{th} \qquad V_{CE} > V_{CE(sat)} \tag{4.2}$$

Here $V_{CE(sat)}$ is the collector-emitter saturation voltage required for the transistor to enter saturation mode. This value is different for each transistor model, and usually has a value of 0.05 V - 0.2 V.

Forward Active

It is the Forward Active mode (also just called active mode or amplifier mode) that makes transistors such a powerful tool. Whereas a transistor in the Saturation or Cut-Off modes simply acts as a switch, turning on or off the flow of current from collector to emitter, a transistor in the active mode is between these to states and acts as an amplifier. A transistor is in the active mode when the following relations are true:

$$V_{CB} > 0 \qquad V_{BE} > V_{th} \tag{4.3}$$

A small electrical current I_B entering the base is amplified by a constant and sets the amplitude of the collector current I_C. Both of these currents exits the emitter of the transistor, as can be seen in Figure 4.4. The amount of current that enters the collector is almost entirely controlled by amplitude of the base current (with a small dependence on V_C). When a transistor is in active mode there exists a linear relation between the two input currents I_B and I_C, given by:

$$I_C = \beta \times I_B \tag{4.4}$$

This is a simplified form of Equation A.7. A small change of 1 mA in the current entering the base will result in a change of βmA entering the collector. β is a near constant value that is unique to each transistor model and usually lies around 100, although transistors with more unusual β values exists.

Reverse active

The Reverse Active mode is rarely used when designing electrical circuits. It is reached by setting the collector voltage lower than both the base and emitter voltage, and the base voltage lower than the emitter voltage:

$$V_C < V_B < V_E \tag{4.5}$$

A transistor will still operate in this mode as it would in the Forward Active mode, but current flows the opposite direction (from emitter to collector). It will also amplify currents as in
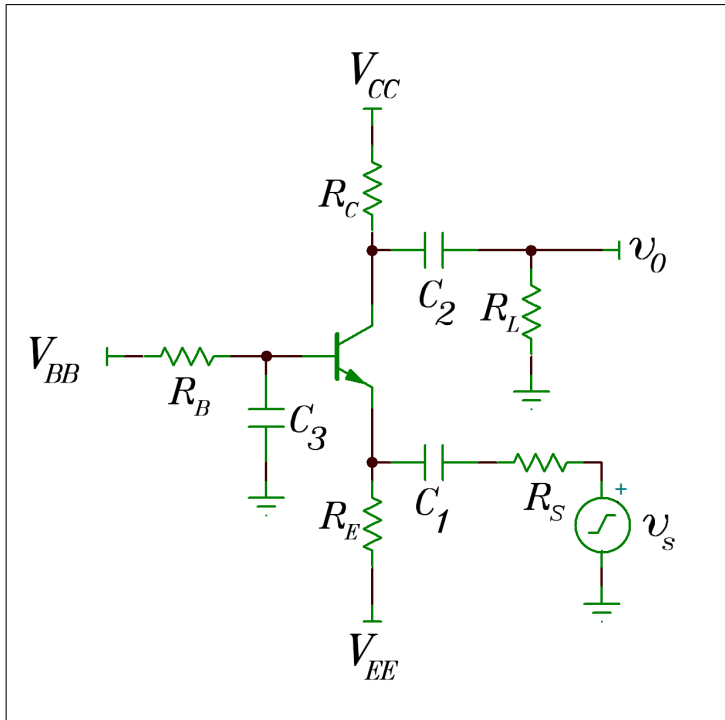
Figure 4.6: Basic circuit diagram of the common base amplifier. The signal we wish to amplify, v_s, is introduced at the emitter, and the amplified signal v_O is read out from the collector.

the Forward Active mode, but the amplification factor β is smaller due to the non-symmetric structure of the transistor as seen in Figure 4.5.

4.3.2   Common base amplifier theory

We can now turn our attention to the Common Base Amplifier (CBA). For a refresher on transistor theory, see Appendix A.1. Figure 4.6 shows a circuit diagram of a general CBA. The NPN BJT transistor is supplied with the voltages V_{BB}, V_{CC}, and V_{EE}. These voltages, as measured at the NPN terminals (V_B, V_C, and V_E), will be different to these since current flowing through the resistors leads to a voltage drop. The signal we wish to amplify, v_s, is introduced to the transistor via the capacitor C_1 connected to the emitter. The amplified signal v_O is read out through a second capacitor, C_2, connected to the collector. The capacitors let Alternating Current (AC) components of the signals pass through, and blocks the Direct Current (DC) component. The amplifier is therefore limited to amplifying AC signals. The resistors R_S and R_L are the resistances of the source and the load, respectively.

The way a CBA functions can be understood by comparing the equivalent DC and AC circuits, and setting $V_{EE} = 0\,\mathrm{V}$ for simplicity. Figure 4.7 shows the DC equivalent circuit of the CBA. Since a DC signal has no frequency components, all the capacitors have been replaced with open circuits and are therefore not drawn. What's left is a circuit that describes the CBA in the absence of an input signal.

Let us now assume that the resistors and voltages in Figure 4.6 that control the CBA DC response has been chosen so that the transistor is in the Forward Active mode. By finding the AC equivalent circuit of the CBA, we can explore its behaviour when subjected to an AC signal
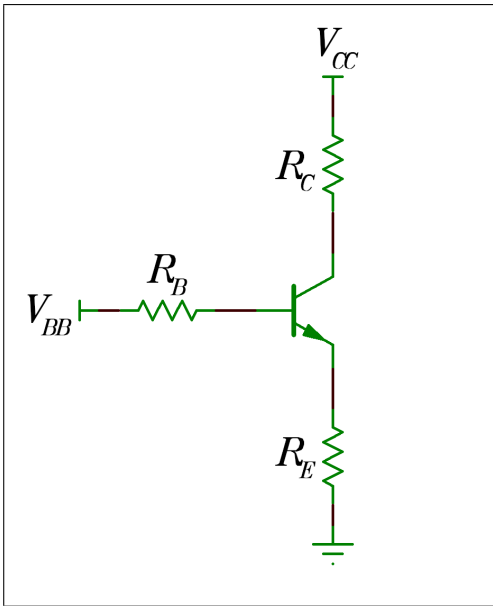
Figure 4.7: DC equivalent circuit for the CBA. Capacitors have been replaced with open circuits. This circuit diagram describes the CBA in the absence of an input signal.
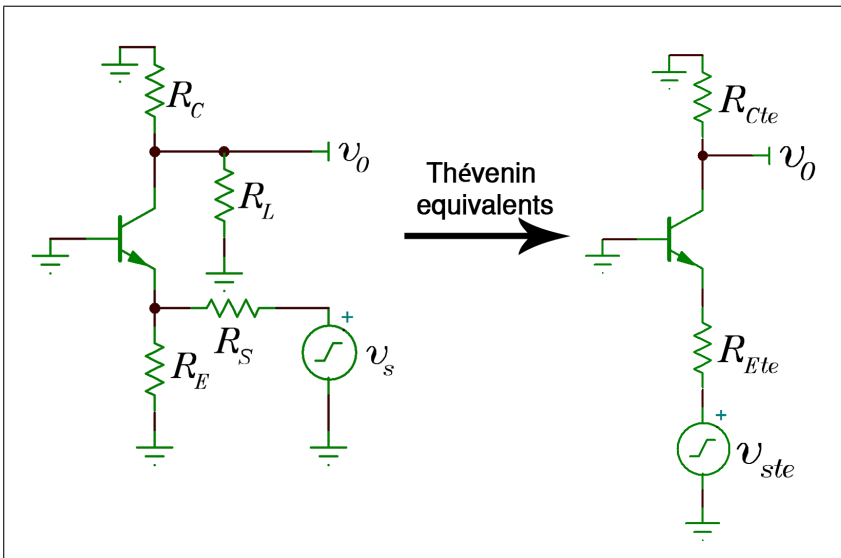
Figure 4.8: AC equivalent circuit for the CBA. Capacitors have been replaced with short circuits, and DC voltages set to 0V. The right-side diagram shows the same circuit, but with the resistors and voltage source replaced with their Thévenin equivalents.

v_s. To find the AC solution, all capacitors are replaced with short circuits, and DC voltages are set to 0V. The resulting circuit can be seen in Figure 4.8. The reason behind the name Common Base Amplifier can be understood by noting that, in the AC solution, the emitter and collector voltages are both referenced to the base voltage. To further simplify the circuit diagram, we use Thévenins Theorem [50, p.6] to replace parallel resistors and the voltage source by their Thévenin equivalents:

$$R_{Cte} = R_C || R_L = \frac{R_C + R_L}{R_C R_L} \tag{4.6}$$

$$R_{Ete} = R_E || R_S = \frac{R_E + R_S}{R_E R_S} \tag{4.7}$$

$$v_{ste} = v_s \frac{R_E}{R_E + R_S} \tag{4.8}$$

Let us now combine the DC and AC solution to understand how the CBA is capable of amplifying signals. The DC equivalent circuit in Figure 4.7 biases the transistor so that is in the Forward Active mode if one chooses suitable values for V_{CC}, V_{BB}, and the three resistors. The transistor in the AC equivalent circuit in Figure 4.8 is therefore also in the Forward Active mode. If the signal source v_{ste} generates a small AC signal, an AC current flows between the base and emitter. Since the transistor is in the forward active mode, this current is amplified in the manner described by Equation 4.4. The AC current drawn by the collector will induce an AC voltage drop across R_{Cte}, and this voltage drop is the amplified signal.

It would make sense to now analyse the CBA further and find the voltage gain, input impedance, output impedance, and other characteristics. But this is not trivial. Certain parameters required for an exact solution are either not available from the manufacturer and difficult to measure, or varies among transistors of the same model. However, in Appendix A.2, a simplified version of the CBA gain has been derived:

$$A_v = \frac{I_C}{I_C + V_T} \times R_C \tag{4.9}$$

We see that the higher the collector resistor R_C, the greater the gain. However, it cannot be made arbitrarily large. The current I_C passing through R_C leads to a voltage drop, decreasing the collector voltage V_C. If it becomes less than the base voltage, the requirement $V_{CB} > 0$ from Equation 4.3 is not met, and the transistor leaves the Forward Active mode. Assuming $V_{EE} = 0\,\text{V}$ and using the equivalent requirement $V_{CE} > V_{th}$, a simplified equation for the maximum value of R_C is deduced:

$$V_{CC} - V_{th} - v_{O-} = I_C \times (R_{Cmax} + R_E)$$

$$R_{Cmax} = \frac{V_{CC} - V_{th} - v_{O-}}{I_C} - R_E \tag{4.10}$$

Here, v_{O-} is the most negative voltage that the output signal must be able to obtain. Equation 4.10 ignores the collector-emitter resistance, but should give a decent guess at the maximum R_C possible. Substituting R_{Cmax} in Equation 4.9 we get:

$$A_v = \frac{I_C}{I_C + V_T} \times \left(\frac{V_{CC} - V_{th} - v_{O-}}{I_C} - R_E \right) \tag{4.11}$$

A plot of this equation for several values of V_{CC} can be seen in Figure 4.9. The CBA gain approaches a limit when the I_C approaches $0\,\text{A}$. This limit will be the maximum theoretical gain of the CBA, and will naturally not be possible to reach in reality since some small amount of current must flow into the collector in order for the CBA to function. In addition will R_C in combination with the BJT parasitic capacitances create a Resistor-Capacitor (RC) lowpass filter. This leads to a lower bandwidth, reducing the frequency range in which the amplifier is useful. This can be seen in Figure 4.10, which shows the frequency response of a CBA simulated in TINA-TI for different I_C. As I_C decreases, larger values for R_C may be used, and the gain approaches the maximum theoretical value but the bandwidth also decreases.
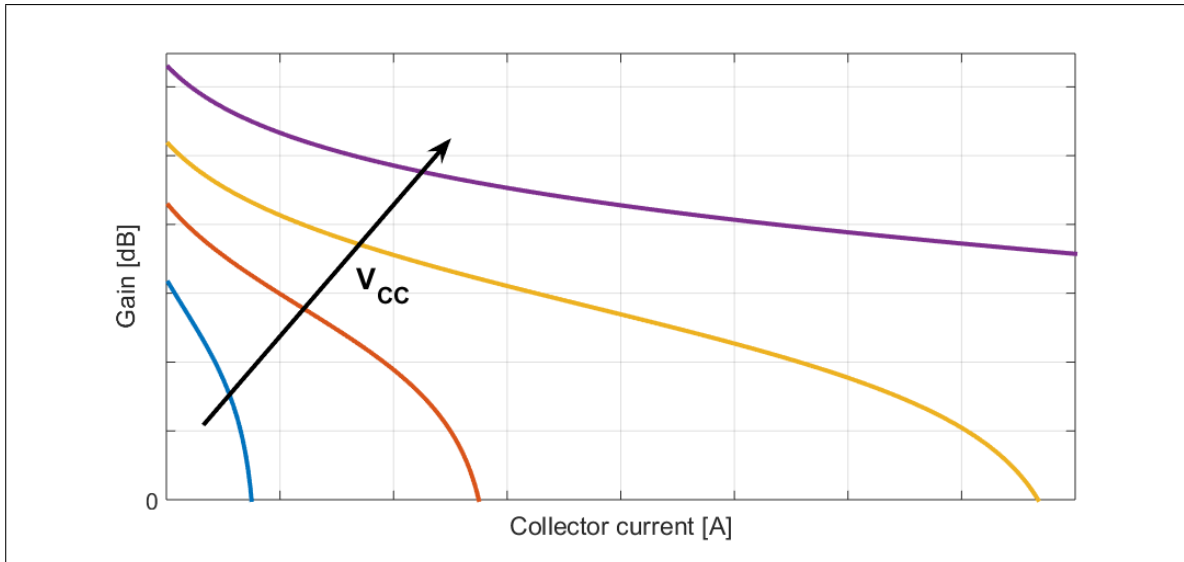
Figure 4.9: Plot of CBA gain as function of I_C, with constant temperature and R_E, and $R_C = R_{C\,max}$ (Equation 4.11). Higher V_{CC} allows for larger maximum values of R_C, raising the maximum possible gain.

It is clear from the simulations that there exists a combination of I_C and R_C that optimises the gain while maintaining a required bandwidth, and that the smaller the BJT parasitic capacitances are, the greater the bandwidth. It is therefore critical that a BJT designed for high-frequency applications is used, since these normally have very low parasitic capacitances. In the end, simulations of the circuit must be used to find a value for R_C that provides the required bandwidth, and then maximize I_C by adjusting V_B.

4.3.3   Noise and shielding

When an electric signal is amplified (be it by a transistor, operational amplifier, or the archaic vacuum tube), some amount of noise is always added to the signal. The sources of noise can be internal to the amplifier (such as the thermal agitation of electrons), or external (such as voltage source variations and electromagnetic interference).

When designing the CBA, certain measures were taken to limit external noise. The entire preamplifier circuit was enclosed in an EMI shield to limit electromagnetic interference from off-board and on-board sources. Off-board sources can for example be radio emitters and badly designed, or faulty, electrical appliances. On-board sources includes any trace that carries a signal with power in the frequency range of the preamplifier. To minimize external noise arising from the voltage sources supplying V_{CC}, V_{BB}, and MPPC bias voltage V_{BIAS}, the traces carrying these voltages were filtered by ferrite beads before entering the EMI shield. A ferrite bead acts as a passive low-pass filter, as its impedance increases with increasing frequency. With the addition of a capacitance in parallel, high frequency components are efficiently filtered out from the voltage supply.
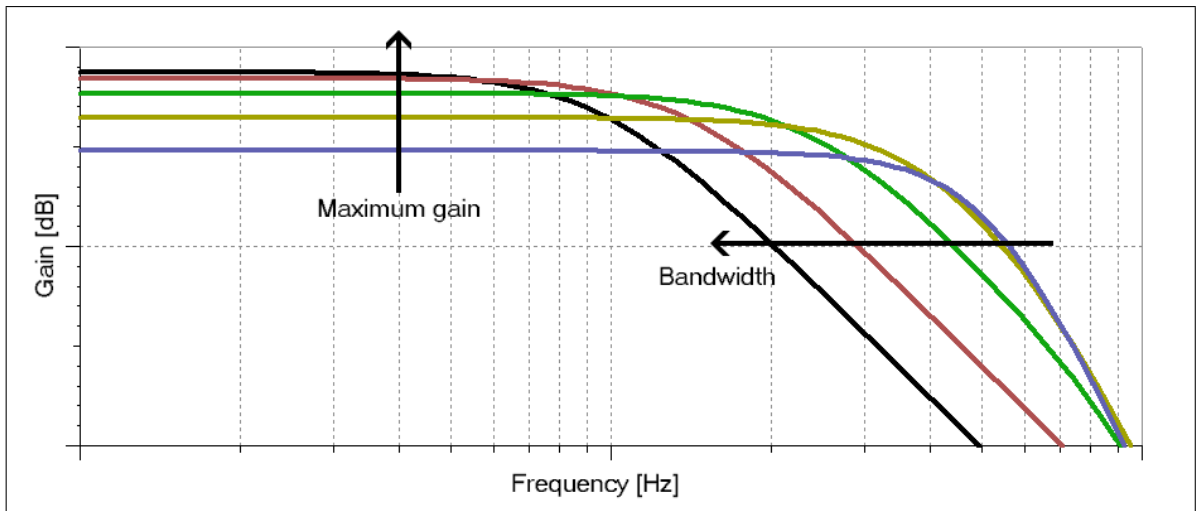
Figure 4.10: Simulated frequency responses of a CBA for varying values of I_C and R_C. As I_C decreases (and thus R_C increases) the CBA approaches its maximum theoretical gain. However, at the same time the bandwidth shrinks due to the effect of transistor parasitic capacitances in parallel with the large-valued R_C.

Great care was also taken to keep the ground plane within the EMI shield free from noise. High frequency ground currents were kept from passing below the preamplifier circuit by attention to power rail routing, and all other components were either placed so that their ground currents did not pass near the preamplifier, or were routed in separate ground current channels if it could not be prevented by placement alone.

4.3.4 Common base amplifier design

An NXP Semiconductors BFU550A was chosen for the preamplifier transistor. Its high bandwidth, low noise, and high maximum collector-emitter voltage made it an excellent choice [51]. In addition, a SPICE model was available from the manufacturer which meant simulations could be used to find the best CBA design for this particular transistor.

R_C was set to $390\,\Omega$. From the BFU550A datasheet we find that the collector stray capacitance C_C is $0.72\,\text{pF}$, which together forms a first-order RC low-pass filter with cutoff frequency f_{-3dB}:

$$f_{-3dB} = \frac{1}{2\pi RC} = \frac{1}{2\pi \times 390\,\Omega \times 0.72\,\text{pF}} \approx 567\,\text{MHz} \tag{4.12}$$

This sets the maximum bandwidth of the amplifier. It will be considerably lower due to parasitic capacitances and resistances internal to the transistor structure, so overshooting the bandwidth is good practice.

Next, a constant I_C of $35\,\text{mA}$ was chosen to investigate the effect of a varying R_E on bandwidth and gain. The source resistance was set to $900\,\text{m}\Omega$ to represent the analogue switch glsic that connects the SiPM to the CBA. R_E was varied and the base voltage adjusted for each value to keep I_C constant. For each value of R_E an AC-analysis was performed to find the CBA gain and the bandwidth, defined as the -3dB cutoff frequency f_{-3dB}. The results can be seen
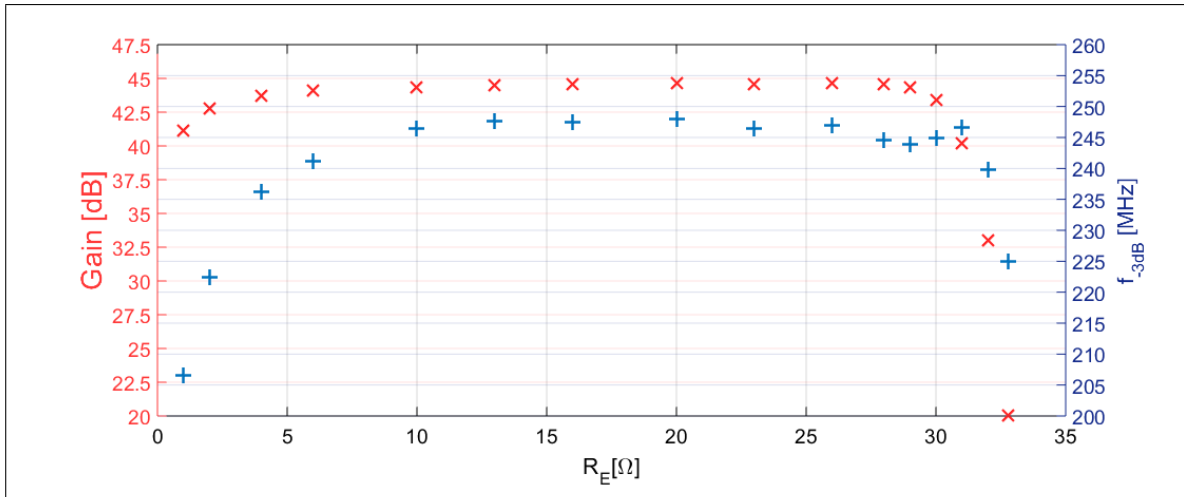
Figure 4.11: Simulated gain and bandwidth of the CBA in Figure 4.6 for varying emitter resistance R_E with $I_C = 35\,\text{mA}$, $R_S = 900\,\text{m}\Omega$, and $R_C = 390\,\Omega$.

in Figure 4.11. We see that for values of R_E between $10\,\Omega$ and $25\,\Omega$ the gain and bandwidth is relatively stable, with the gain increasing slightly with increasing R_E. When R_E nears $33\,\Omega$ the gain and bandwidth both drops rapidly. This corresponds to the collector-emitter voltage V_{CE} approaching the threshold voltage V_{th}. Based on these simulations an R_E value of $22\,\Omega$ was selected.

CBA bandpass filter design

A circuit diagram of the final preamplifier design can be seen in Figure 4.12. The choice of resistor and capacitor values will be explained in the current section.

The CBA was designed to act as a second-order narrow bandpass filter centred around $100\,\text{MHz}$. A narrow bandpass filter is a type of filter which supresses frequency components the further they are from a center frequency f_C. It can be interpreted as consisting of a highpass and lowpass filter with the same $-3\,\text{dB}$ cutoff frequency. An illustration of a narrow bandpass filter is presented in Figure 4.13. $100\,\text{MHz}$ is higher than the $40\,\text{MHz}$-$80\,\text{MHz}$ that was decided for the bandpass center frequency in section 3.6. The reason for this is that simulations showed it would improve the total frequency response of the amplifier chain.

To explain the choice of component values, we must represent the transistor with an equivalent electrical model. Figure 4.14 shows a circuit diagram of the CBA where the transistor has been replaced with a full hybrid-pi model [52] representation, which accurately models the small-signal response of a BJT. A RC high pass filter is generated by the combination of emitter resistance r_π, base spreading resistance r_x, R_E and C_1, suppressing lower frequency components. So too will the combination of C_2 and the load resistance R_L. Together they form the second-order high-pass filter required for the bandpass-filter.

The second-order lowpass filter that makes up the second half of the bandpass filter is generated by the parasitic capacitances present between the base and emitter, and base and collector, C_π and C_μ respectively. The combination of C_μ with R_C and R_L forms a first-order lowpass filter. So does the combination of C_π with R_S. Together they form the second-order
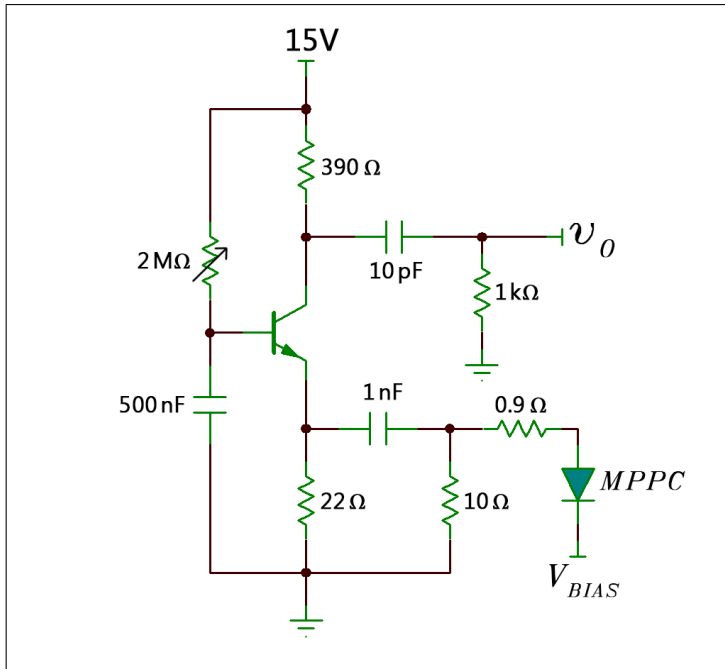
Figure 4.12: Simplified circuit diagram of the final preamplifier design. R_B is adjustable to enable adjustable gain.

lowpass filter that suppresses higher frequencies.

The combination of these two filters results in a narrow bandpass filter as required. The complex behaviour of the transistor meant that the optimal values for C_1 and C_2 would have to be found through simulations to provide an optimal frequency response of the entire amplifier chain. The applied voltages to the transistor, and therefore also the value of I_C, were kept constant in the simulations, since the internal capacitances are voltage dependent.

It was relatively simple to select C_1, and a value of $1\,\mathrm{nF}$ was found to result in an acceptable cutoff frequency. For C_2 it was less straightforward, as two conflicting design goals had to be balanced. Firstly, the CBA AC gain is dependent on the parallel combination of R_C and the load resistance R_L (Equation A.17, Appendix A.2), and therefore a value of R_L that is large relative to R_C is desired so as not to degrade the AC gain. But the larger R_L is, the smaller C_2 must be for the cutoff frequency of the RC highpass filter created by R_L and C_2 to stay constant. This can be seen from Equation 4.12, which is also true for highpass filters. Due to stray capacitances on the PCB traces, C_2 cannot be arbitrarily small, as the parallel stray capacitance will combine with C_2 to form a voltage divider which reduces the overall gain. Estimating a maximum stray capacitance of $1\,\mathrm{pF}$, the optimal combination of C_2 and R_L was found through simulations to be $10\,\mathrm{pF}$ and $1\,\mathrm{k\Omega}$.

An additional $10\,\Omega$ resistor was added to the CBA input to provide a low-impedance path to ground for the charge generated by the SiPM. This reduces the gain by $\sim 8.3\,\%$, but is required, otherwise the SiPM anode will be floating. Using a large-valued resistor would decrease the gain loss, but massively increases the risetime of the MPPC signal, which is undesirable as the goal is amplification of fast MPPC pulses.
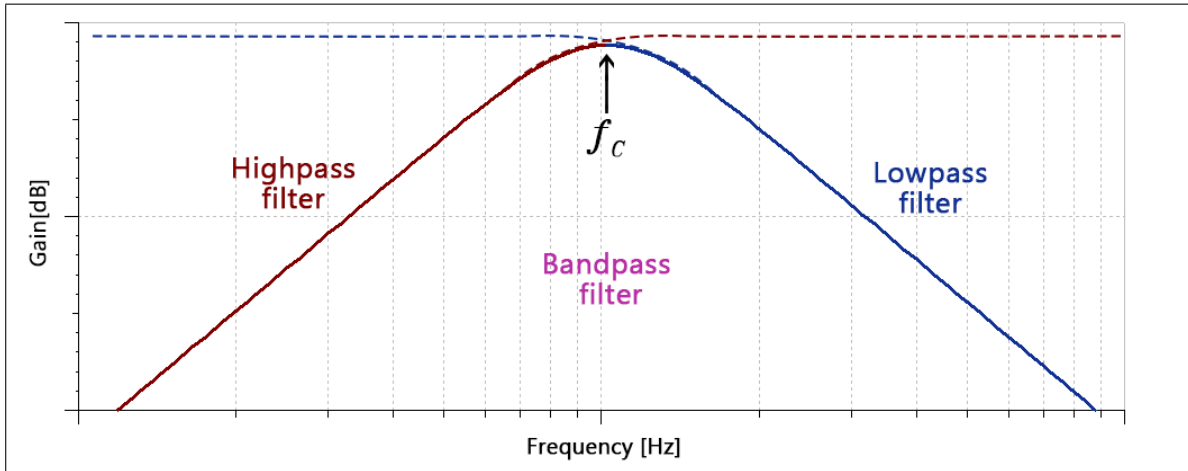
Figure 4.13: Illustration of a narrow bandpass filter. The frequency components of a signal passing through such a filter will be attenuated the further they are from the center frequency f_C.
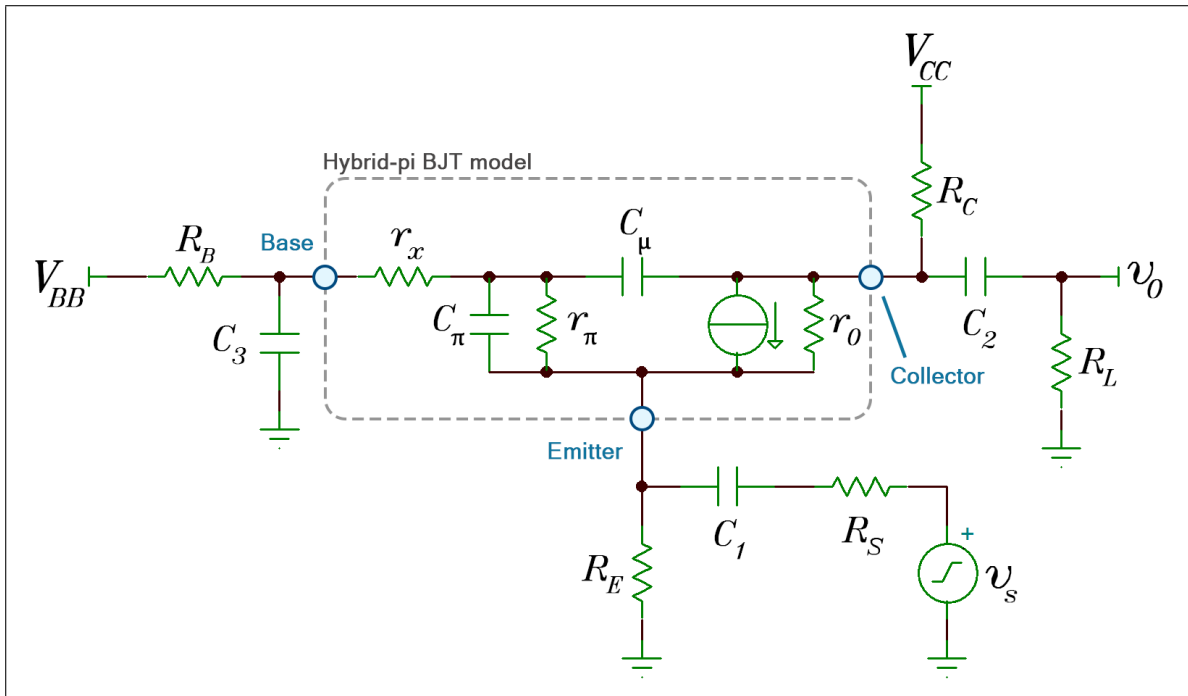
Figure 4.14: Circuit diagram equivalent to the CBA seen in Figure 4.6. The transistor has been replaced with a full hybrid-pi model which accurately models the transistor characteristics, and V_{EE} set to ground as in the subsequent circuit diagrams.

Adjustment of I_C

Being able to adjust the collector current I_C will make it possible to find an optimal trade-off between amplifier gain and MPPC heating. The CBA gain is strongly dependent on I_C (Equation 4.9), but so is the transistor power dissipation given by $P = U \times I$. The larger I_C is, the greater the power dissipation in the preamplifier will be. As the MPPC is physically close to the preamplifier, this will heat it up, increasing the dark rate and reducing the SNR. A second motivation for allowing adjustment of I_C is that due to variations in β among the transistors, there will be a variation in I_C if the base current I_B is constant among the ROBs, as can be seen from Equation 4.4.

These were the motivations for allowing adjustment of I_C, and it was realized by implementing a method of manually adjusting the base current I_B. The base current is supplied by V_{EE} through a drop resistor R_B. By adjusting R_B, the magnitude of the base current can be controlled, and so also I_C.

However, if one decreases I_C too much, the narrow bandpass frequency response of the CBA is degraded. This is apparent from Figure 4.15. As the base resistor R_B increases, I_C decreases, and the peak flattens. The decrease in maximum gain comes simply from the equation for CBA gain, Equation 4.9. The cutoff frequencies of the lowpass and highpass filters are pushed apart due to the current dependency of C_π and r_π [52]:

$$C_\pi = C_{je} + \frac{\tau_F I_C}{V_T} \tag{4.13}$$

$$r_\pi = \frac{V_T}{I_B} \tag{4.14}$$

Here, C_{je} and τ_F are two constants related to the transistor construction (zero-bias capacitance of the base-emitter junction, and forward transit time). What is important is that C_π and r_π depend on the base current and collector current.

A measure of the degradation can be defined as the width of the frequency response, i.e. the distance between the upper and lower -3dB cutoff frequencies, Δf_{-3dB}. Plotting width and power consumption versus R_B shows that above $200\,\text{k}\Omega$ the benefit of increasing R_B is offset by the increasing width Δf_{-3dB}, as can be seen in Figure 4.16. Setting R_B to a value of $100\,\text{k}\Omega$-$300\,\text{k}\Omega$ might prove to be a good middle ground between gain flattening and MPPC heating.
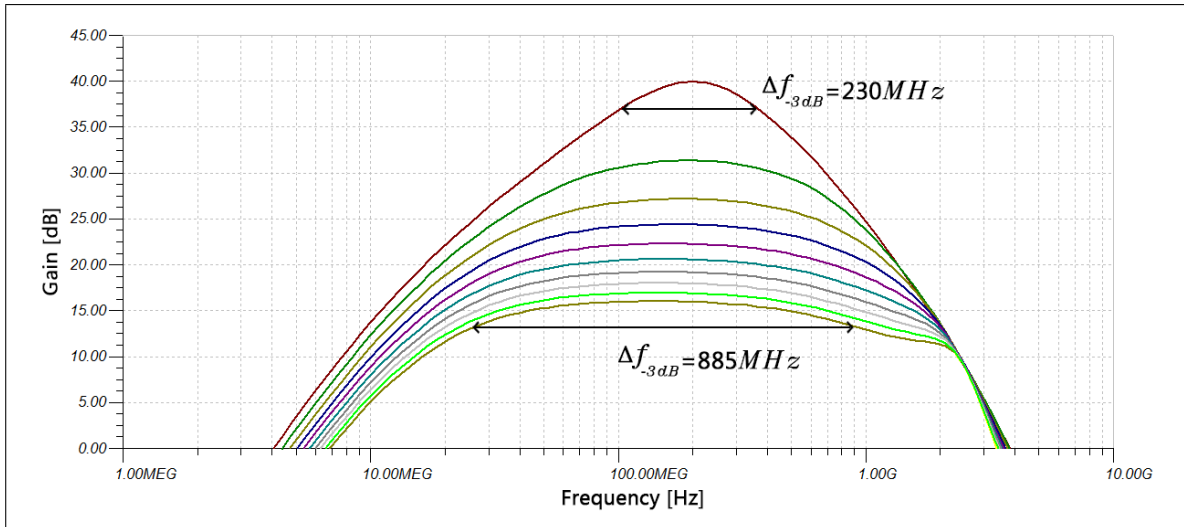
Figure 4.15: Frequency response of the preamplifier in Figure 4.12 for varying base resistance R_B, changing linearly from $R_B = 25\,k\Omega$($I_C = 36\,mA$, top curve) to $R_B = 2\,M\Omega$($I_C = 715\,\mu A$, bottom curve). The variation in I_C widens the peak and decreases the maximum gain.
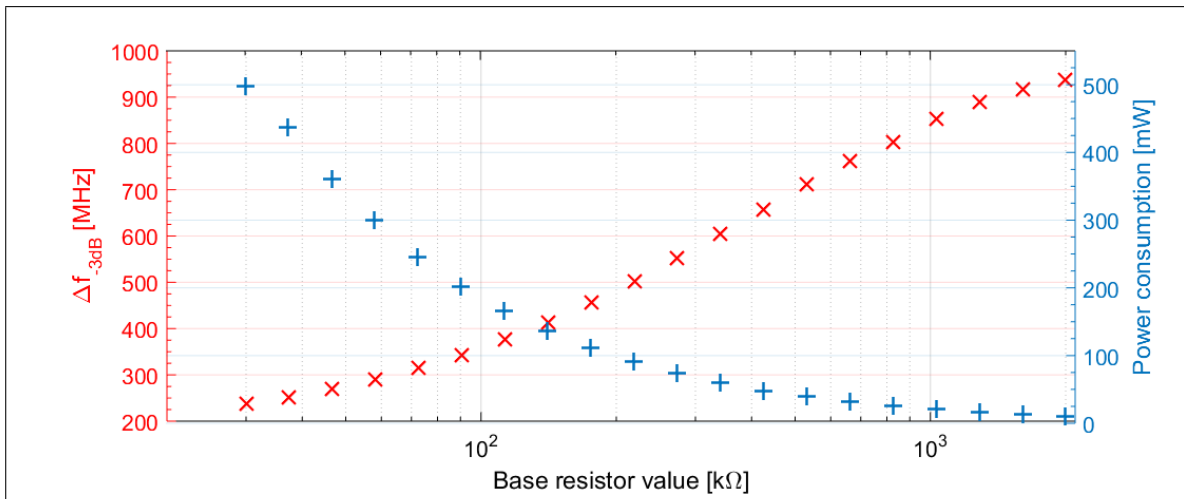
Figure 4.16: Width Δf_{-3dB} of the bandpass filter created by the CBA as seen in Figure 4.15, and the CBA power consumption plotted versus base resistor R_B.

4.3.5   Preamplifier charge inserter

In order to easily evaluate the behaviour of the amplifier chain, a simple charge insertion sub-circuit was added to the preamplifier. This allows the user to insert a small charge into the preamplifier, thereby generating a repeatable well-defined signal. This can be used to test the characteristics of the amplification chain and to adjust its gain.

Before a test charge is inserted, the SiPM is disconnected from the preamplifier to prevent it from interfering with the test charge signal. An analogue switch which connects a precision capacitor to the preamplifier input is then closed. The precision capacitor is normally charged to 3 V, and it discharges through the analogue switch and to ground through R_E. The potential difference V_{BE} is disturbed, and is amplified, generating a large output pulse.

4.4 Primary amplifier

The primary amplifier is the second stage of the amplifier chain. It is a non-inverting precision rectifier, built around the OPA847 voltage feedback op amp [53]. For an introduction to op amps, see Appendix A.3. The rectifier serves two purposes: Further amplifying the MPPC signal, and shielding the discriminator from negative input voltages. Using an op amp for the second stage made it simpler to design than if a single-transistor based amplifier had been used, such as in the preamplifier. The OPA847 was chosen due to its high slew rate, and large 3.9 GHz Gain-Bandwidth product (GBW), which was one of the largest GBWs for op amps in a physical package with leads. op amps with higher GBWs existed, but lacked leads and were therefore only solderable by special equipment.

The primary amplifier was designed to be a precision rectifier, meaning only positive voltages are amplified. A precision rectifier has the benefit over of being able to fully suppress negative voltages, whereas a simple one-diode rectifier has to overcome the diode voltage drop, and will therefore allow negative voltages down to about −650 mV.Using a single-diode rectifier would not be sufficient, since the op amp output is supplied to the input of a comparator which is highly sensitive to negative voltages.

The design of an op amp based precision rectifier involves a few more elements than a basic op amp amplifier. As can be seen in Figure 4.17 there are two feedback paths, instead of one: One for negative voltages and one for positive voltages. The feedback path that the output signal takes is set by diodes in the two paths. When the op amp output voltage is greater than 0V, current flows through diode D_1 and feedback resistor R_{f1}. When it is negative, current flows through diode D_2 and feedback resistor R_{f2}. Generally, $R_{f1} > R_{f2}$, and for certain stable op amps, R_{f2} can be 0 Ω. The disadvantage of precision rectifiers is that the output signal has to cross two diode voltage drops when crossing 0V. This creates undesired effects, warping the output signal at higher frequencies due to finite GBW and slew rate, ultimately limiting the usable frequency range of the rectifier.

The diodes used for the precision rectifier were Avago Technologies HSMS2860, a germanium Schottky diode for use in high frequency applications. Its very low stray capacitance made it an excellent choice for this application, as the op amp is highly sensitive to capacitances on its feedback path and terminals.

To circumvent the problem of the op amp output having to change by two diode voltage drops when crossing 0V, the non-inverting input was positively biased so that the op amp output DC value became positively offset. The offset was set to be greater than the voltage drop of the Schottky diode in the positive-voltage feedback path. This ensured that the leading edge of the preamplifier-filtered SiPM pulse is amplified without delay or distortion, and that the negative part of the output pulse is suppressed if it falls below 0V.

The bias voltage was set with a voltage divider to 5 V on the non-inverting input, whose Thévenin equivalent was 1 kΩ as required by the CBA design. It was tuned to shift the op amp output voltage to ∼250 mV for a constant op amp input bias current of −29 µA. But due to
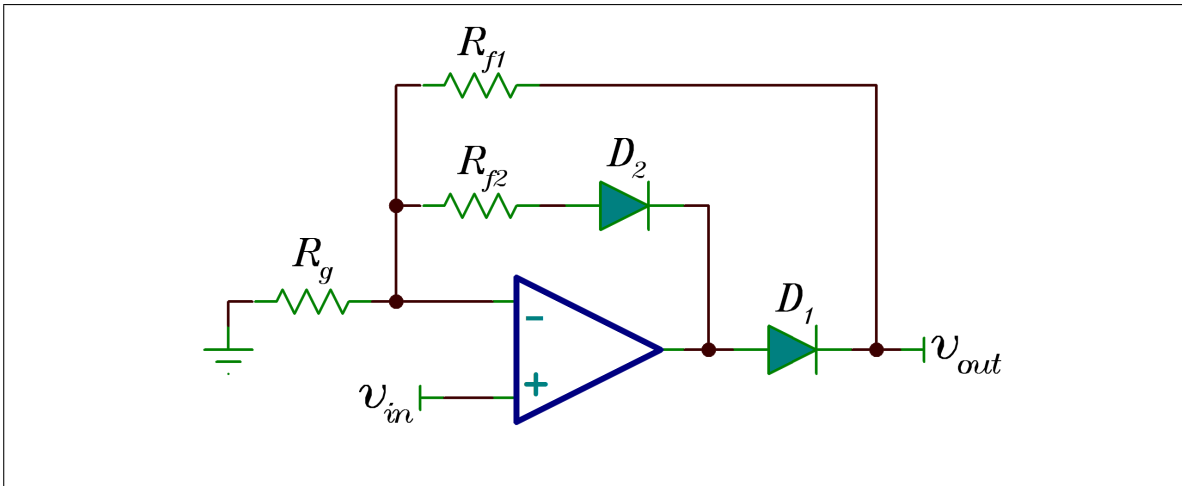
Figure 4.17: The basic configuration of a precision rectifier. v_{out} is shielded from negative voltages by D_1, while D_2 supplies a feedback path for the negative voltages to prevent the op amp output from becoming negatively saturated. Since negative voltages are not of interest, R_{f2} is normally small or omitted.
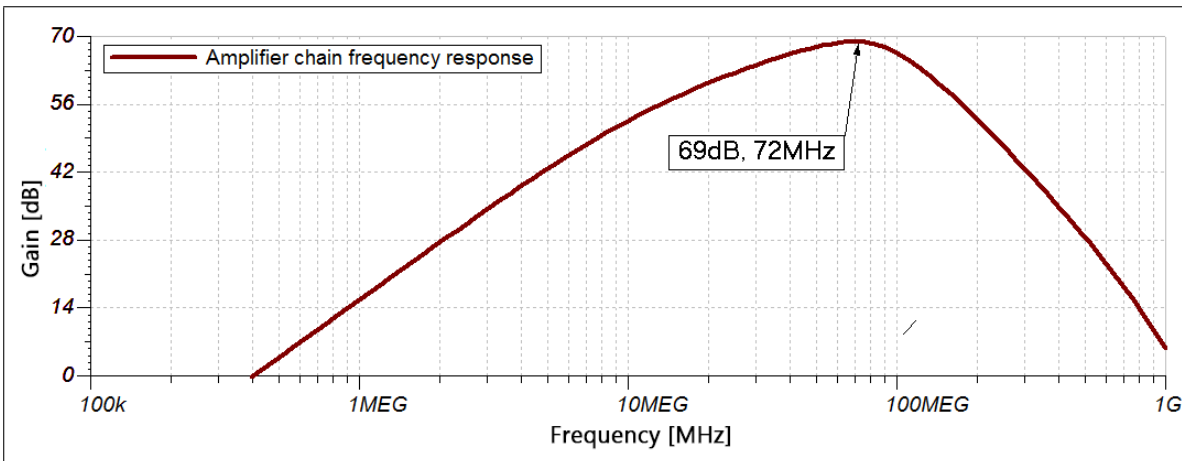
Figure 4.18: Frequency response of the entire MPPC amplifier chain discussed in Sections 4.3 and 4.4. The peak is found at ~72 MHz, with a maximum gain of 68.8 dB.

large variations of the input bias currents of the OPA847 among units ($-19\,\mu$A to $-41\,\mu$A), the voltage divider was later adjusted for each ROB. This resulted in output voltages of each ROB from 200 mV - 350 mV. This variation in the DC component of the output signal was compensated for in the on-board MCU software.

For the selection of the feedback resistors, the OPA847 datasheet recommended a value of $R_g = 39\,\Omega$, and this suggestion was followed. Simulations of the amplifier chain with incrementing R_{f1} showed that setting $R_{f1} = 1.8\,$kΩ resulted in the best compromise of signal risetime and gain. Lowering R_{f1} results in an upwards shifting in frequency of the narrow bandpass filter peak, while increasing R_{f1} shifts it downwards. At 1.8 kΩ the peak is located at 72 MHz, within the 40 MHz-80 MHz that was selected for the bandpass filter design requirements. The final frequency response of the entire amplifier chain can be seen in Figure 4.18.

4.5   Discriminator

The discriminator converts the analogue MPPC signal to a digital yes/no signal. Its input is connected to the amplifier output, as seen in Figure 4.3 and continuously compares the amplitude of the input signal to a fixed voltage, called the trigger threshold. If the input signal becomes greater than this voltage, a square pulse with a duration of 50 ns is generated by the discriminator. This pulse signifies that an MPPC signal of amplitude greater than the trigger threshold has been detected, and it is output over coaxial cable to be counted or processed elsewhere.

4.5.1   Discriminator theory

The discriminator consists in essence of three elements: A comparator, a line driver (or buffer), and a latch circuit. The comparator is an IC that compare two voltages on its two inputs V_+ and V_-, outputting a low or high signal depending on which is larger. The propagation delay of the comparator will set a limit on how quickly it can register a change of the input voltage levels. For example, if a certain comparator has a propagation delay of 10 ns, a short change of which input is larger with a duration of less than roughly 10 ns will not lead to a proper transition of the output. One must therefore choose a comparator with a propagation delay less than the time scale of the pulse one wishes to register.

The line driver is an op amp which buffers the comparator output. Its output goes to the coaxial connector on the ROB, and to the latch circuit. The latch circuit is required for the comparator to output pulses of fixed duration. If no latch circuit is present, an output pulse from the comparator will only last as long as one input was larger than the other. In the case of a comparator comparing pulses to a threshold, as in the MPPC discriminator, this would lead to comparator pulses of varying duration, depending on how long the pulse exceeded the threshold. In our application where the comparator is monitoring the filtered MPPC pulses, this time can be anything from a few picoseconds to tens of nanoseconds depending on threshold level and signal amplitude.

4.5.2   Discriminator design

In this section the design of the comparator and latch circuit will be discussed, and the need for a line driver be explained.

Comparator and line driver

The comparator model chosen was the Analog Devices ADCMP553. It was the cheapest, fastest, human-solderable comparator available, with a propagation delay of 500 ps and minimum input pulse width of 700 ps. It also had a latch input, making it simple to extend the output pulse. The comparing voltage, or "trigger threshold", V_{tt} was set by an analogue output from
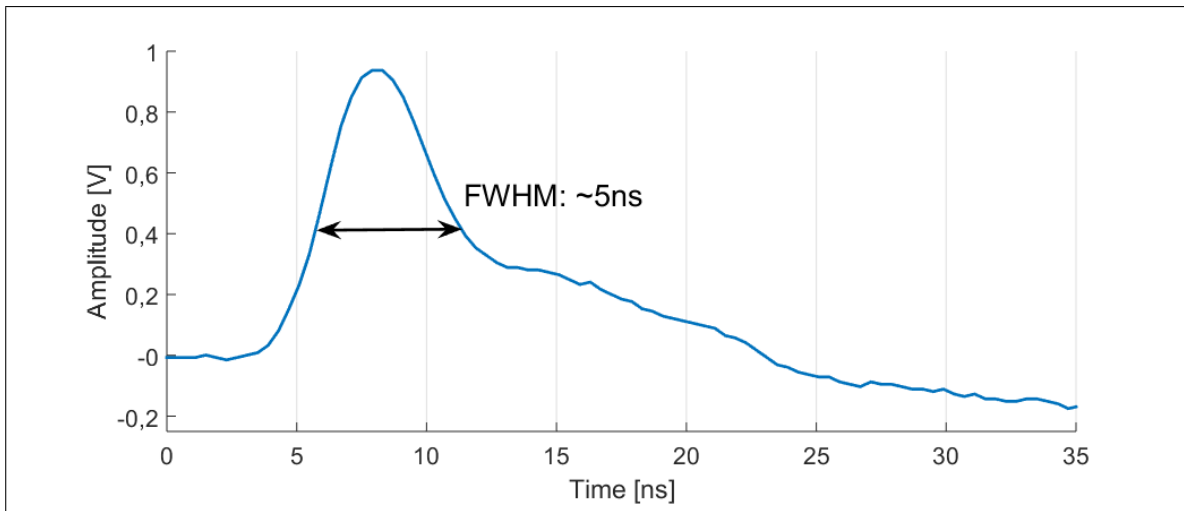
Figure 4.19: Measured typical pulse as seen by the comparator. The small width of the pulse required the use of fast comparator. The pulse was measured at the primary amplifier output, labelled "MPPC signal output" on Figures 4.1 and 4.3.

the MCU, divided and strongly filtered by a voltage divider and capacitor, and buffered by a Microchip MCP6001 general-purpose op amp to prevent the input bias currents offsetting the trigger threshold. The ADCMP553 differential outputs were terminated by connecting them to the common node of a voltage divider created by $130\,\Omega$ to $3.3\,V$ and $82\,\Omega$ to ground. The inverting comparator output was AC-coupled to a Texas Instruments BUF602, a $1\,GHz$ voltage buffer with high output current capability. It served as the line driver for the output coaxial connector.

Due to the short duration of a typical input pulse to the comparator, it was necessary to use a model with very short propagation delay. A recorded MPPC pulse as seen by the comparator is shown in Figure 4.19. $5\,ns$ is already a short propagation delay, and is considered high speed in the world of comparators. But still, using a comparator with $5\,ns$ propagation delay would mean that many pulses would not be registered even if they did cross the trigger threshold, as they would not stay above it for the duration of the propagation delay. It was therefore necessary to select a much faster comparator with a minimum input pulse width of $\sim 2\,ns$ or less. Its physical package would also need extended leads to enable it to be soldered by hand, and be as cheap as possible. The comparator that fitted these requirements best was the ADCMP553.

The negative aspect of using the ADCMP553 is its differential open-emitter outputs, which generates Low Voltage Positive Emitter-Coupled Logic (LVPECL) levels. Unlike Transistor-Transistor Logic (TTL) levels, which normally normally operate with $0\,V$ to $0.8\,V$ as logic 0 and $2\,V$ to $5\,V$ as logic 1, LVPECL levels are a bit odd. Logic 0 is defined as $1.43\,V$ to $1.68\,V$ and logic 1 as $2.23\,V$ to $2.48\,V$. Additionally, for a LVPECL device operating with supply voltage V_S, its outputs requires termination to $V_S - 2\,V$ via $50\,\Omega$. Regardless of the additional design challenges that LVEPCL levels introduced, the ADCMP553 was still the best choice due to its low cost and high speed.

Since the supply voltage of the ADCMP553 is $V_S = 3.3\,$V, its outputs would have to be terminated to 1.3 V via 50 Ω. Instead of adding a 1.3 V voltage rail for this sole purpose, the ADCMP553 outputs were terminated by connecting them to the common node of a voltage divider created by 130 Ω to 3.3 V and 82 Ω to ground. This is the Thévenin equivalent of 50 Ω to 1.3 V [54, p.9].

The line driver was required for two reasons, the first being that the ADCMP553 has a strict termination requirement, and that its output must be AC-coupled to fit the TDC input requirements. The difference between LVPECL low and high logic levels coincided perfectly with the required pulse amplitudes required by the TDC, and no scaling of the signal was therefore necessary.

Latching circuit

The motivation behind implementing the latching circuit was that the comparator pulses were to be registered by the TDC, which requires a minimum input pulse width of 5 ns. This is longer than the minimum pulse width that the comparator is able to output, which is around 700 ps. Additionally, the latching circuit was required to prevent the discriminator from generating several pulses in fast succession, as can easily happen due to the fact that the scintillation photons arrives in "waves". When the comparator is latched, it ignores these later pulses.

In latched operation, the output of a comparator will be in the state that it was in when the latch was activated. In unlatched operation, also called tracking mode, it will behave as a regular comparator. A comparator can therefore be made into a simple pulse extender by connecting the latch input to the comparator output. If the comparator is normally outputting logic 0, the latch is inactive, and the comparator will be able to react to a change of the voltages on its inputs. If it would then output a logic 1, the latch activates, and the comparator output will be logic 1 indefinitely.

The ADCMP553 have complementary latch inputs, requiring LVPECL logic levels. They were connected to the differential outputs through normally activated analogue switches, causing the comparator to be in an autolatching configuration as discussed above. When the comparator becomes latched, a set amount of time passes before it unlatches. The duration of the latch is set by the RC time constant of capacitor C_t discharging through resistor R_t in Figure 4.20. The timing capacitor C_t begins discharging when the comparator triggers on an MPPC pulse and becomes latched, and its voltage is level shifted to fit the trigger range of a Schmitt trigger. When the voltage of C_t descends past the Schmitt trigger low-going threshold, the analogue switches are deactivated in order to disconnect the latch inputs from the comparator outputs, bringing the comparator back to the normal tracking mode. Assuming the comparator input is now below the trigger threshold, the inverted output will then go to logic 1. This quickly recharges C_t by a simple charger circuit made from an inductor and Schottky diode. The C_t voltage crosses the Schmitt trigger high-going threshold voltage, and the analogue switches are reactivated. The comparator can now trigger on a new MPPC pulse.

Figure 4.20: Simplified circuit diagram of the MPPC pulse discriminator, consisting of an ultrafast comparator, trigger threshold setter, latch circuit, and line driver.

There is a possibility that the latch circuit will malfunction and indefinitely latch the comparator. This is due to the fact that the latch circuit is connected to the line driver output, which is AC-coupled from the comparator. The potential therefore exists for the line driver input to drift to a logic 1 level when the comparator inverting output is actually logic 0, if this logic 0 is preserved for more than the ∼5 milliseconds the AC coupling requires to charge. This could easily happen, for example if the trigger threshold is set lower than the DC component of the MPPC signal. The latch circuit will therefore believe the comparator to be unlatched, when it is in fact latched, and it will never reset it. Therefore, a method of monitoring and resetting the latch circuit was implemented. The MCU measures one of the ADCMP553 latch inputs several times per second, and if it is active on many consecutive reads, it has most likely malfunctioned. If so, a forced unlatch is performed by shifting the Schmitt trigger input below its low-going threshold by adjusting the voltage on the input labelled "Latch reset" on Figure 4.20.

4.6   Voltage regulation

Due to the fact that the ROB has a large range of capabilities, voltage regulation was not trivial. Whereas entirely digital systems can make do with a single power supply, or rail, the ROB required 6 unique rails of different voltages:

- **5V(Digital)**: Voltage rail for digital electronics.
- **3.3V(Digital)**: Voltage rail for LVPECL electronics.
- **5V**: Voltage rail for analogue electronics.
- **-5V**: Complementary voltage rail for analogue electronics.
- **15V**: Voltage rail for preamplifier.
- **55V**: Variable voltage rail for MPPC bias voltage

The generation and regulation of these rails were performed in two steps. First, laboratory power supplies generates four separate semi-regulated voltages which are being supplied to all ROBs. These are referred to as the off-board supplies, and generates the power for all ROBs on the CRT. The voltages from these supplies are noisy, and not adjusted accurately. On each ROB, Low Drop-Out linear voltage regulators (LDOs) were added to further filter and perform the final precision regulation of the semi-regulated voltages supplied by the off-board supplies. Several steps of passive filtering were implemented to reduce noise on the final on-board analogue voltage rails.

Fuses on the ROB power inputs protect it from serious damage, possibly igniting, in case of a short circuit or other component failure. Since the CRT will mostly operate without supervision, it was deemed important to have safety measures in place to minimize the probability of this happening.

4.6.1   Voltage regulation considerations

When designing the on-board voltage regulation, the choice stood between using LDOs or Switch Mode Power Supplies (SMPSs). LDOs have the benefit of being very simple to use and have low electrical noise. SMPSs on the other hand require additional external components, and are noisy. The disadvantage to LDOs is the fact that they are inefficient. The load current causes a voltage drop in the LDO analogous to the current passing through a resistor. If a LDO has to regulate $15\,\mathrm{V}$ down to $5\,\mathrm{V}$, and the load draws $1\,\mathrm{A}$, the LDO require $P = (15\,\mathrm{V} - 5\,\mathrm{V}) \times 1\,\mathrm{A} = 10\,\mathrm{W}$ to operate. This is significant wasted energy that will be transformed into heat. In order to minimize wasted energy, LDOs are designed to be able to operate with a very small voltage drop across its terminals, hence the name. SMPSs has a entirely different mode of operation, and can down-regulate high voltages with little wasted energy. However, LDOs had to be chosen due to the fact that SMPSs can generate considerable amount of electrical noise.

The inefficiency of LDOs, and the large variation in required voltage rails ($\sim$55 V to 3.3 V), meant that several semi-regulated voltages had to be supplied by the off-board power supplies. If the ROB voltage regulation had been SMPS-based, a single off-board voltage would have been sufficient. However, the low noise of LDOs made this additional complication justifiable. There are therefore 4 semi-regulated voltages supplied to the ROB: 70 V, 17 V, 7 V, and −7 V. The 70 V line is regulated down and adjusted to fit the bias voltage required by the MPPC. The 17 V line is regulated down to 15 V as required by the preamplifier, and the 7 V and −7 V lines are regulated to supply the remaining voltage rails. The semi-regulated voltages are slightly larger than what is required by the LDOs in order to accommodate for the voltage drop generated by the fuses and the reverse-protection diodes that are present on all four power inputs.

4.6.2   Noise and filtering

The off-board power supplies supplying the semi-regulated ROB voltages are SMPS-based, and therefore noisy. The cables from the power supplies to the ROBs will also have to be relatively long (on the order of a few meters) due to the size of the CRT, and electromagnetic noise is therefore easily picked up. To prevent this noise from making it to the sensitive amplifier chain, a three-step filtering method was employed. In this method, the off-board voltages are first filtered by a combination of a RC low-pass filter and ferrite bead. They are then regulated by the LDOs, which has some noise reduction capabilities. The regulated voltages are then further filtered by a second ferrite bead. Since the MPPC amplifier chain is a narrow bandpass filter centred at 72 MHz suppressing noise around this frequency was the main goal of the filtering. The three stages of filtering will now be explained.

First filtering

The semi-regulated input voltages are first filtered by a low-pass RC filter created by a small-valued resistor, ferrite bead, and capacitor. The ferrite bead is a passive component that exhibits varying impedance by frequency, and is often used to filter voltage rails and power cables. Using it in place of the resistor in an RC filter will greatly increase the filters ability to suppress high-frequency components, as it behaves as a second-order low-pass filter. The ferrite bead chosen for filtering of the semi-regulated voltages was the Murata BLM15AX102SN1D. This ferrite bead has a high impedance of 1 kΩ at 100 MHz and is rated for 350 mA, more than enough for this application.

The capacitors making up the low-pass filter of each voltage input was chosen to be 4.7 μF, 1 μF, and 100 nF depending on the voltage (low to high). Since the cables that carry the voltages from the power supplies have inductance, a small-valued resistor with large power rating was added in series with the inductor to dampen the resulting LCR-circuit.

The strength of the filters at 72 MHz are: −127 dB for the ±7 V lines, −114 dB for the 17 V line, and −94 dB for the 70 V line.

Second filtering

The second filtering was performed by the LDOs themselves. They have a characteristic called Power Supply Rejection Ratio (PSRR). This is a figure of how well the LDO is able to reject frequency components on the supply voltage. The PSRR changes by frequency, and unfortunately it is rarely specified for frequencies above 1 MHz in LDO datasheets. Above these frequencies, the PSRR will be mainly decided by package parasitics and board layout. Therefore, one cannot rely much on the filtering capabilities of LDOs at the frequency range of interest to the amplifier chain. High-valued capacitors on the LDO outputs helps increase PSRR, and is good practice either way to prevent instabilities, and capacitors of values greater than 1 μF were used on all LDO outputs. The LDOs themselves will also generate noise, and low-noise models were therefore used where possible.

Third filtering

The third filter stage was implemented for the components that required particularly noiseless supply voltages. Since ferrite beads are inexpensive, they were added to the voltage rails for all elements in the amplifier chain. This results in an additional noise reduction of 82 dB-106 dB at 72 MHz.

4.7 MPPC bias adjustment

The MPPC bias voltage V_{BIAS} was supplied by a Texas Instruments TPS7A4001, a high-voltage LDO. Active adjustment of this regulator was required mainly due to the dependency of MPPC bias voltage V_{BIAS} to temperature, but also because of the variable breakdown voltage among MPPCs. Being able adjust the MPPC bias voltage would also allow us to perform analysis of the ROB performance for different overvoltages. Therefore, a method of digitally adjusting V_{BIAS} was developed. Additionally, a circuit for monitoring the MPPC temperature was developed to enable the MCU to perform automated gain stabilization.

A simplified schematic of the V_{BIAS} adjustment circuit is seen in Figure 4.21. It illustrates the electronically closed loop regulation scheme that was implemented to digitally adjust V_{BIAS}. In this scheme V_{BIAS} is divided by 14 through resistor divider K_1, and then measured by a high-resolution Analogue to Digital Converter (ADC). The ADC reports the measured value to the MCU, which subsequently increases or decreases a feedback voltage v_{FB} in an attempt to bring V_{BIAS} to some target value. This cycle repeats until the target value for V_{BIAS} is reached.

It was decided that the requirements for V_{BIAS} should be as following:

- Full range: 48 V - 64 V
- Accuracy: ± 50 mV

The range encompasses all possible required MPPC operating voltages, and the accuracy was chosen arbitrarily to be $\sim \pm 0.1 \%$ since any lower than this is rather challenging.

The next sections will describe the adjustment scheme in greater detail, and the accuracy of V_{BIAS} will be found. The temperature measurement accuracy will also be estimated, and together used to find the MPPC gain uncertainty.

4.7.1 Motivation for accuracy calculations

The combination of uncertainty in the temperature measurements together with the uncertainty in determining V_{BIAS} will lead to an uncertainty in the desired overvoltage of the MPPC. Since the overvoltage sets the gain of the MPPC, this will make it difficult to equalize the gain across all ROBs of the CRT. It is therefore of interest to understand what the gain uncertainty is, and a relatively large effort went into estimating it.

To find an expression of the gain uncertainty, we turn to the equation for SiPM gain. From Equation 3.3.1 in Chapter 3 we have that the amount of charge released by a firing G-APD in a SiPM is expressed by

$$Q = C(V_{BIAS} - V_{br}) \tag{4.15}$$

where C is the capacitance of the G-APD. Dividing by the elementary charge unit e yields the gain:

$$G = Q/e = \frac{C}{e}(V_{BIAS} - V_{br}) \tag{4.16}$$
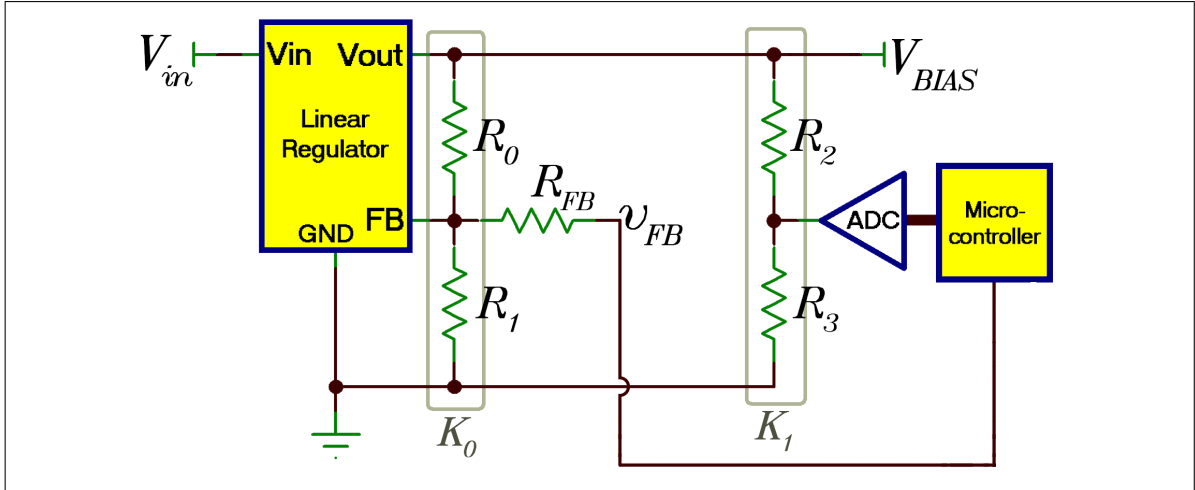
Figure 4.21: Schematic of the MPPC bias voltage regulation circuit, showing the basic components required for measuring and adjusting V_{BIAS}. By adjusting v_{FB}, the MCU adjusts V_{BIAS}.

Since V_{BIAS} is temperature dependent the equation becomes

$$G = \frac{C}{e}(V_{BIAS25} + 54\,\text{mV}\,°\text{C}^{-1}(T - 25\,°\text{C}) - V_{br}) \tag{4.17}$$

Here, V_{BIAS25} is the bias voltage at 25 °C as specified by the manufacturer, and $54\,\text{mV}\,°\text{C}^{-1}$ is the temperature dependency of SiPM used for this project, the Hamamatsu S13360-6075CS MPPC. Performing an error propagation on this equation yields the following dependency on temperature uncertainty σ_T and bias voltage uncertainty σ_{BIAS}:

$$\sigma_G = \frac{C}{e} \times \sqrt{(\sigma_T \times 54\,\text{mV}\,°\text{C}^{-1})^2 + (\sigma_{BIAS})^2}$$

$$\hat{\sigma}_G = \frac{\sigma_G}{G} = \frac{\sqrt{(\sigma_T \times 54\,\text{mV}\,°\text{C}^{-1})^2 + (\sigma_{BIAS})^2}}{V_{BIAS} - V_b r}$$

$$\hat{\sigma}_G = \frac{\sqrt{(\sigma_T \times 54\,\text{mV}\,°\text{C}^{-1})^2 + (\sigma_{BIAS})^2}}{V_{over}} \tag{4.18}$$

We see that the relative uncertainty in gain is the reciprocal of the overvoltage V_{over}, and that the magnitude of the uncertainty scales with the combination of σ_T and σ_{BIAS}. Therefore, if one wishes to be able to precisely set the MPPC gain, these two uncertainties must be low. The remainder of this section is devoted to estimating them.

4.7.2 V_{BIAS} accuracy

The accuracy of which V_{BIAS} can be determined, σ_{BIAS}, is limited by the precision of the voltage divider K_1 and by the resolution of the ADC. It has been estimated to be $\sigma_{BIAS} \approx 27\,\text{mV}$ in Appendix B.2, quite a bit better than the desired 50 mV. The two error sources, K_1 and the ADC, will now be discussed.

Voltage divider

The precision and accuracy of K_1 will depend on the manufacturing uncertainties of R_2 and R_3 and their temperature coefficients respectively. The precision uncertainty was corrected for by measuring the actual division factor when the resistors had been soldered onto the ROB. The Temperature Coefficient of Resistance (TCR) of R_2 and R_3 are therefore the only error sources that will impact the uncertainty of the voltage divider. In the aforementioned appendix, the absolute uncertainty of K_1, σ_{K1}, over the expected ROB temperature range has been estimated to be at most

$$\sigma_{K1} \approx 26\,\text{mV} \tag{4.19}$$

Resistors with lower temperature coefficients could have been chosen, but their high cost made it prohibitive, and an absolute uncertainty of ~26 mV was deemed sufficient. To further increase the accuracy of K_1 the resistors were mounted physically close, and their common nodes soldered onto a large copper pad to act as a heat sink in order to further equalize their temperatures. To minimize self heating, the resistor values chosen were large: $R_2 = 1\,\text{M}\Omega$ and $R_3 = 76.8\,\text{k}\Omega$. By Ohm's law ($P = U^2/R$), at $V_{BIAS} = 64\,\text{V}$ this results in an power deposit of 3.64 mW and 0.28 mW respectively. From the datasheets of the resistors, it is found that this results in a negligible self heating of 2.04 °C and 0.31 °C. The uncertainty in K_1 will be slightly larger if one takes into account the resistors self heating, but as it is simply a rough estimation, we ignore its small addition.

ADC

The ADCs on the MCU are 8-bit, and this resolution is not sufficient. Therefore a dedicated ADC was needed to digitize the analogue voltage on the common node of K_1. A 12-bit Texas Instruments ADC121S021 was chosen for this, as it was cheap and simple. Its input was buffered with a general purpose MCP6001U op amp, as the ADC has a large input leakage current of $\pm 1\,\mu\text{A}$. This current in combination with K_1 would create an uncertainty in V_{BIAS} of $\pm 994\,\text{mV}$.

The MCP6001U has a negligible input bias current of ± 1 pA. However, it has a relatively large input offset voltage of $\pm 4.5\,\text{mV}$. The ADC too has an offset error. To compensate for this, their combined offset is measured and adjusted for in the on-board MCU software to prevent it from negatively affecting the V_{BIAS} measurement accuracy.

In Appendix B.2 the absolute uncertainty of the ADC, σ_{ADC}, has been estimated to be

$$\sigma_{K1} \approx 4.9\,\text{mV} \tag{4.20}$$

We see that the sampling uncertainty is small compared to the uncertainty in K_1, and does not contribute much to the total bias voltage uncertainty σ_{BIAS}.

4.7.3 V_{BIAS} adjustment

Adjustment of V_{BIAS} was realized by a resistor network that connected the LDO output, its feedback pin, and an analogue output from the MCU. The scheme is seen in Figure 4.21, and a more detailed view is presented in Figure 4.22. The values selected for the resistors R_0, R_1, and R_{FB} enables V_{BIAS} to be adjusted from $\sim$46 V to $\sim$64 V, with a resolution of 17.4 mV. The choice of resistor values and range will be discussed in more detail.

Determination of resistor values

The voltage divider k_0 is created by resistors R_0 and R_1. In the absence of R_{FB}, K_0 sets the LDO output voltage V_{BIAS}. The LDO will adjust its output so that the voltage on its feedback pin, V_{FBP}, is always some constant value. For the TPS7A4001, this is nominally 1.173 V. The addition of R_{FB} and a digitally adjustable voltage v_{FB} enables V_{BIAS} to be adjusted.

Determining the values for R_0, R_1 and R_{FB} is relatively simple. First we need an equation for V_{BIAS} as a function of the three resistors and v_{FB}. The LDO input pin has negligible bias current. By Kirchoff's laws, we therefore know that the current flowing through R_0 into the common node equals the sum of the currents flowing out of the common node through R_1 and R_{FB} .

$$I_{R0} = I_{R1} + I_{RFB} \tag{4.21}$$

By Ohm's law Equation 4.21 becomes

$$\frac{V_{BIAS} - V_{FBP}}{R_0} = \frac{V_{FBP}}{R_1} + \frac{V_{FBP} - v_{FB}}{R_{FB}}$$

$$V_{BIAS} - V_{FBP} = \frac{R_0 V_{FBP}}{R_1} + \frac{R_0 (V_{FBP} - v_{FB})}{R_{FB}}$$

$$V_{BIAS} = V_{FBP} \left(1 + \frac{R_0}{R_1} - \frac{R_0}{R_{FB}} \right) - v_{FB} \frac{R_0}{R_{FB}} \tag{4.22}$$

By evaluating this expression at the maximum and minimum desired values for V_{BIAS}, and knowing the range of v_{FB}, the required values of R_1 and R_{FB} can be found for any R_0.

To find the resistor values for this application, V_{BIAS} was first decided to be adjustable from 48 V to 64 V. This range was expected to be sufficient, and in the next section it is shown to be correct. The range of v_{FB} is limited by the MCU from 0 V to 5 V. Fixing R_0 to an arbitrary value of 1 MΩ, this results in $R_1 = 19.856$ kΩ and $R_{FB} = 312.661$ kΩ. High-precision resistors close to these values exists, but the cost can be decreased by using resistors with rounder values which are more commonly used. Setting $R_1 = 20$ kΩ results in $R_{FB} \approx 280$ kΩ, both of which are commonly used resistor values. This change of values has the effect of slightly increasing the range of V_{BIAS}, which ends up being adjustable from 46.16 V to 64.01 V.
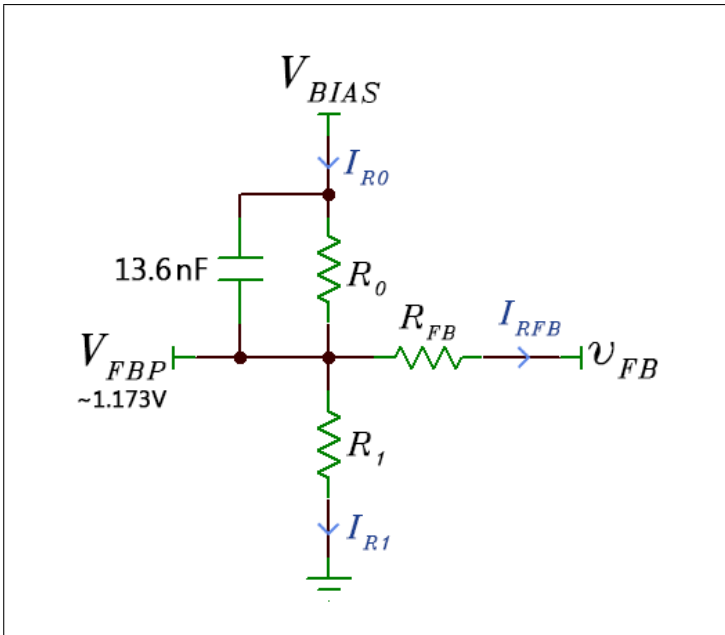
Figure 4.22: Detailed circuit diagram of the V_{BIAS} adjustment resistor network. The LDO adjusts V_{BIAS} to keep the feedback pin voltage V_{FBP} constant. Therefore, v_{FB} sets V_{BIAS}.

Required range

The adjustment range of V_{BIAS} must be able to accommodate for the following:

- Adjustment of MPPC overvoltage (0 V to 5 V)
- Variability in MPPC bias voltage at 25 °C (51.4 V to 52.4 V)
- Variability in expected MPPC temperature (25 °C $\pm$ 14 °C)
- Variability in V_{FBP} (1.173 V $\pm$ 12 mV)

The values in point 2 are the largest and smallest for the MPPC purchased for the CRT upgrade. A simple worst-case analysis can be used to show that the adjustment range is sufficient. The maximum and minimum required values of V_{BIAS} for any MPPC, temperature, and overvoltage are:

$$V_{BIASmax} = 52.4\,\text{V} + 5\,\text{V} + 56\,\frac{\text{mV}}{\text{°C}} \times 14\,\text{°C} \approx 58.2\,\text{V} \tag{4.23}$$

and

$$V_{BIASmin} = 51.4\,\text{V} - 56\,\frac{\text{mV}}{\text{°C}} \times 14\,\text{°C} \approx 50.6\,\text{V} \tag{4.24}$$

Variations in V_{FBP} shifts the range of V_{BIAS}. Using Equation 4.22 the worst case high and low bounds are found to be 63.36 V and 46.81 V. We see that the adjustment range will always be sufficient, regardless of component variations and temperature. The excess range of V_{BIAS} allows for variations in the values of R_0, R_1 and R_{FB}, and their temperature dependencies.

V_{BIAS} resolution

The resolution of V_{BIAS} is directly set by the resolution of v_{FB}. The MCU generates this voltage with an internal 10-bit pulse-width-modulated (PWM) output, resulting in a minimum step size in V_{BIAS} of $(64.01\,\text{V} - 46.16\,\text{V})/2^{10} \approx 17.4\,\text{mV}$. This is sufficient, as it is lower than the uncertainty in V_{BIAS}.

The PWM signal from the MCU is filtered by an RC low-pass filter created by R_{FB} and a

capacitance of 13.6 nF between the LDO feedback pin and LDO output pin. This filter has a −3 dB cutoff frequency of 41 Hz, averaging the high frequency digital PWM signal. If this filter was omitted then high-frequency components would be present on V_{BIAS}. The filter also serves to reduce the rate of change of V_{BIAS} when it is adjusted, which otherwise could have been amplified by the amplifier chain.

4.8   Temperature monitoring

Since the gain of a SiPM is relatively sensitive to temperature, it must be monitored and the bias voltage adjusted accordingly if one wishes to maintain stable gain. One might assume that since temperature is such an important parameter, most SiPMs would already come with temperature sensing capabilities incorporated into the silicon substrate of the detector. However, this is not the case. Therefore, if one wishes to measure the temperature of a SiPM, it must be implemented by the designer of the application.

A temperature sensing circuit was designed that allows the MPPC to be determined with good accuracy. The absolute standard deviation of the error was estimated to be at most $\sigma_T = 0.3\,°\mathrm{C}$ across the expected ROB temperature range of $25°\pm14°$. This section will describe the sensing circuit in greater detail.

4.8.1   Circuit design

The MPPC temperature sensing circuit consists of a Resistive Temperature Detector (RTD) [55] incorporated in a Wheatstone bridge [56]. For more information on the Wheatstone bridge and RTD, see Appendix A.4. The bridge voltage is supplied by a precision voltage source, and the difference voltage that arises as the RTD resistance varies by temperature is amplified by an instrumental amplifier and measured by the MCU. An overview of the circuit is shown in Figure 4.23.

The MPPC temperature will be measured by the MCU with a period of a few minutes, as the MPPC temperatures are not expected to change rapidly. It will then calculate the bias voltage that should be applied to the MPPC for the measured temperature, and adjusts the bias voltage to bring it to the required value.
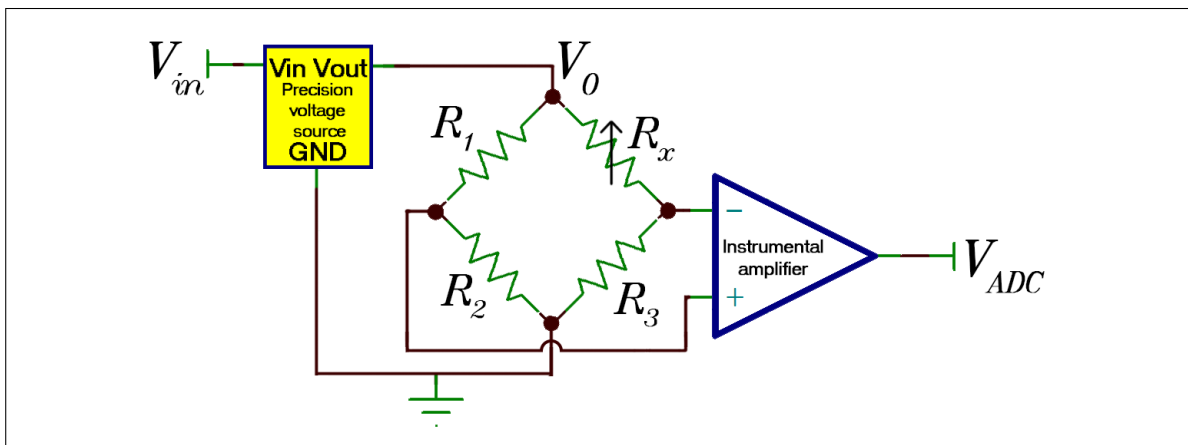
Figure 4.23: Simplified schematic of the MPPC temperature sensing circuit. The difference in voltage between the two Wheatstone legs is amplified by the instrumental amplifier, and its output voltage is digitized by the MCU.
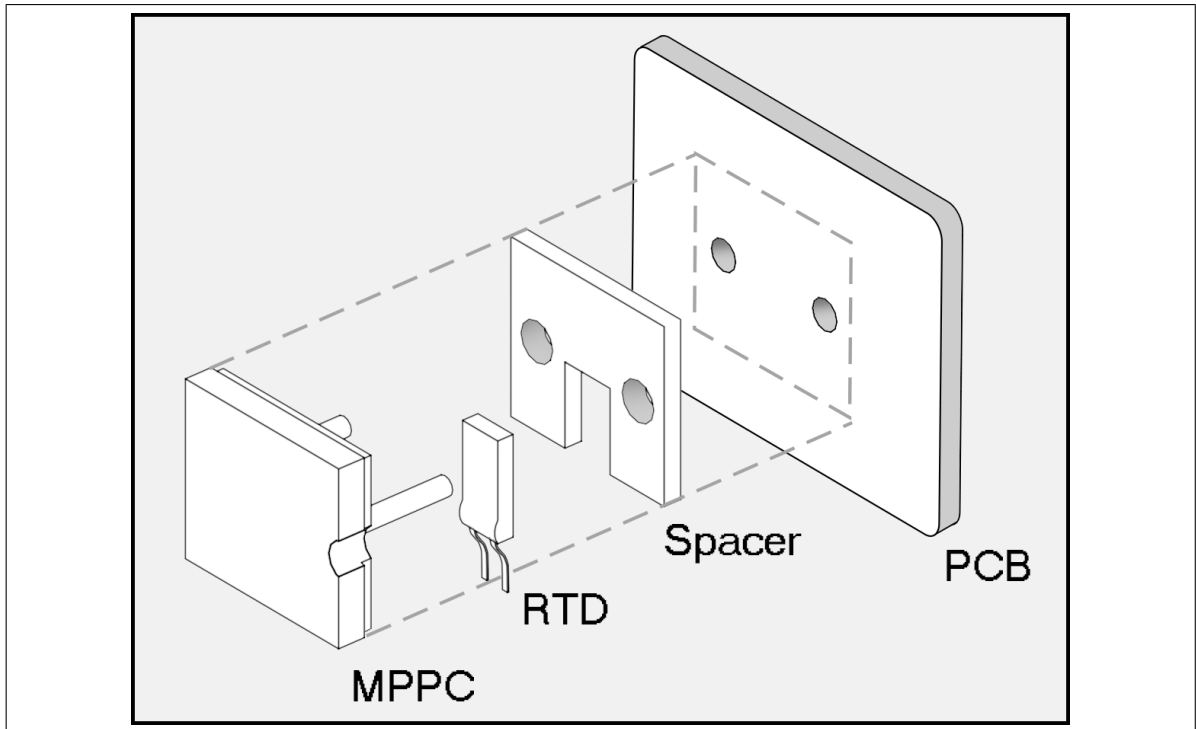
Figure 4.24: Diagram illustrating the mounting method for the RTD and MPPC. The plastic spacer is for support.

An inexpensive nickel RTD, the IST ND1K0.520.2FW.B.007 [57], was chosen for sensing the MPPC temperature. Its small size of $5\,\text{mm} \times 2\,\text{mm}$ and height of 0.65 mm meant it could easily be placed on the underside of the MPPC package. A plastic spacer with the dimensions of the MPPC and with the same height as the RTD was used to raise the MPPC above the PCB and provide structural support. The RTD was placed in a cutout in this plate. It was thermally coupled to the MPPC package with thermal paste, and thermally insulated from the PCB by its own ceramic base. The mounting scheme is illustrated in Figure 4.24.

For the Wheatstone bridge 0.1 % precision resistors with low TCRs of $\pm25\,\text{PPM}\,/^\circ\text{C}$ were selected. R_1 was chosen to be $1050\,\Omega$, balancing the bridge at around $8\,^\circ\text{C}$. R_2 and R_3 were both chosen to be $10\,\text{k}\Omega$, and of the same model. This meant they would both have very similar temperature responses, increasing the accuracy of the Wheatstone bridge.

The Wheatstone bridge voltage was supplied by a precision voltage reference, the Maxim Integrated MAX6071BAUT25 [58]. This voltage reference outputs 2.5 V with an initial accuracy of $\pm0.08\,\%$ and a maximum temperature drift of $\pm8\,\text{PPM}\,/^\circ\text{C}$. It can supply up to 10 mA, more than enough for the bridge which only require less than half a milliampere.

The difference voltage across the bridge, V_D, was amplified by an Analog Devices AD8293G80A [59] instrumental amplifier (in amp) with a fixed gain of 80. This in amp is excellent for precision measurements due to its low input offset voltage of maximum $\pm50\,\mu\text{V}$, low gain error of maximum $\pm1\,\%$ and gain drift of maximum $\pm25\,\text{PPM}\,/^\circ\text{C}$. Typical values were considerably lower.

4.8.2   RTD calibration

Before the RTDs could be used they had to be calibrated, due to variations in the TCR value α and zero point r_0 among the units. The calibration consisted of measuring the resistance of each RTD at a few well-known temperatures, and fitting a linear curve to the measurement points. To execute these measurements, a platinum RTD with an accuracy of $\pm 0.05\,°C$ (PT100, DIN standard 1/10) was used as a reference. It was placed with the RTD to be calibrated in a FLUKE 9102S dry-well in order to control and equalize their temperatures, and their resistances were measured by a FLUKE 1586A SUPER-DAQ precision temperature scanner. By correlating the RTD resistance to the reference probe for several temperatures, a precise calibration of the RTD was performed. The 4-wire measurement technique was used to compensate for lead resistance in the current-carrying wires going to the two RTDs.

The results of the calibrations are shown in Appendix B.5. The actual temperature coefficient of the RTDs were found to be $\sim$5650 PPM/°C, not 6180 PPM/°C as claimed by the manufacturer. To ensure the discrepancy did not come from experimental errors, the RTD resistances were measured with all channels on the SUPER-DAQ, and the temperature of the dry-well was confirmed with a second reference platinum RTD. All results points to the fact that the RTDs do in fact have a temperature coefficient markedly lower than what is expected. This could be due to impurities in the nickel used for the resistive element or other manufacturing errors.

Calibration of the 16 RTDs had to be attempted a few times due to unforeseen error sources which rendered the data unusable. The first error source came from the fact that the RTDs were not sufficiently thermally insulated. It could be seen that the measurement points deviated increasingly from the expected linear response the further from room temperature the measurements had been made. Both the reference RTD and the calibration RTD had measurement leads exiting the dry-well, which acted as thermal sinks and sources, pulling the temperatures of the RTDs towards the temperature outside the dry-well. In subsequent calibrations this error source was successfully minimized by insulating the dry-well entrance and RTD leads with Styrofoam.

The second error source came from self heating. The voltage applied to the RTDs in order to measure their resistances generates a current, which deposits power in the RTDs by the equation $P = U^2/R$. This is called self-heating, and must be taken into account when one designs temperature measurement circuits and when doing calibrations. The SUPER-DAQ applies $\sim$1.2 V when measuring resistances close to, but less than, $1.2\,k\Omega$. This leads to an energy deposit of $P = (1.2\,V)^2/1.2\,k\Omega = 1.2\,mW$. This lead to an increase in the RTD temperature of $\sim$0.4 °C. In the third calibration attempt both of the errors sources that were found in the previous attempts were reduced. The SUPER-DAQ measurement range for the RTDs to be calibrated was manually set to $10\,k\Omega$ to minimize self-heating, and the self-heating of the reference RTD was measured and corrected for. The thermal insulation was increased on both RTDs, and the measurement range decreased in order to further reduce the effect of RTD temperatures deviating towards room temperature.

Error source	Contribution
Wheatstone resistors	46.2%
Voltage source	0.3%
In amp	0.3%
ADC	7.6%
RTD r_0	6.9%
RTD α	38.6%

Table 4.1: Relative contribution to the final uncertainty of each error source in the SiPM temperature measurement electronics. Values have been found from Equation B.3 in Appendix B.3.

4.8.3   Accuracy of temperature measurements

The full mathematical derivation of the uncertainty in the MPPC temperature is performed in Appendix B.3. In these calculations the following error sources are taken into account:

- Constant and temperature-dependent uncertainties in the Wheatstone bridge resistors, instrumental amplifier, and reference voltage source
- Digitization errors
- Uncertainties related to the calibration of the RTD

The final result of the calculations is that we get a value for the absolute standard deviation of the temperature measurements that varies with temperature. The worst case estimate over the entire expected temperature range from $11\,°C$ to $39\,°C$ is found to be:

$$\sigma_T = 0.3\,°C \tag{4.25}$$

The uncertainties related to each component in the measurement circuit contribute to σ_T with different amounts, as can be seen in Table 4.1. We see that it is dominated by the manufacturing uncertainties in the three precision resistors that make up the Wheatstone bridge and by the uncertainty in the RTD temperature coefficient α.

Impact on measurements

We have now found the two variables σ_T and V_{BIAS} that we set out to estimate in the start of this section. From Equation 4.18, the value of relative MPPC gain uncertainty $\hat{\sigma}_G$ is found

to be

$$\hat{\sigma}_G = \frac{\sqrt{(\sigma_T \times 54\,\text{mV}\,°\text{C}^{-1})^2 + (\sigma_{BIAS})^2}}{V_{over}}$$

$$\hat{\sigma}_G = \frac{\sqrt{(0.3\,°\text{C} \times 54\,\text{mV}\,°\text{C}^{-1})^2 + (26\,\text{mV})^2}}{V_{over}}$$

$$\hat{\sigma}_G = \frac{32\,\text{mV}}{V_{over}} \tag{4.26}$$

For a typical overvoltage of 3 V, the relative uncertainty is 1 %. For a lower overvoltage of 0.5 V, the uncertainty is ±6.4 %. This uncertainty in gain translates to an uncertainty in what the correct trigger threshold value should be, and it becomes difficult to accurate choose a trigger threshold for low values of V_{over}.

4.9   The microcontroller

To make the implementation of a MCU on the ROB as simple and least time-consuming as possible, it was decided to use the Arduino platform. Arduino is a collection of open-source hardware and software for MCU implementation, simple enough to serve as a first introduction to programming and electronics for beginners, but also flexible enough to be used for advanced projects. It lowers the difficulty by providing the user with pre-assembled PCBs containing a MCU, voltage regulator, USB communication, and In-Out (IO) pins. Simple software allows the user to quickly start programming the MCU using a supplied C/C++ library. Several types of Arduino boards exists, with varying degree of complexity and capabilities, such as Bluetooth radio, WiFi, microSD slots, LED displays, etc. For the ROB, the MCU only required analogue and digital IO pins, which meant a simple, small, and inexpensive model could be used.

The decision was made to use the Pro Mini model. It has 14 digital IO pins and 6 analogue inputs, a small size, and is inexpensive. This makes it an excellent choice as the ROB control MCU. An image of the Pro Mini can be seen in Figure 4.25. The MCU on this board is an Atmel ATmega328, a high-performance 20 MHz 8-bit MCU with 32KB of flash memory.

The tasks of the MCU consists mainly of monitoring and controlling the ROB. It constantly monitors all voltage rails, the latch circuit status and MPPC temperature. It is able to enable or disable individual LDOs, adjust the MPPC bias voltage, disconnect the MPPC from the amplifier chain, insert test charges, enable or disable the amplifier, set the trigger threshold, and reset the latch circuit. These capabilities will be further detailed in the following sections.

4.9.1   ROB Monitoring

Monitoring of the latch status, MPPC temperature, and most voltage rails was made possible via a multiplexing scheme. The limited amount of inputs on the ATmega328 meant that it was not possible to directly connect all lines of interest to unique MCU analogue inputs. Therefore, a general purpose analogue multiplexer was used to reduce the number of pins required to perform the measurements. It allows switching between 8 analogue inputs with three control pins, reducing the number of pins required for monitoring of the board from 6 analogue pins to 1 analogue and 3 digital pins. Most measurements made by the MCU were done in this manner, except for monitoring of the MPPC bias voltage which required higher resolution than what the ATmega328 is capable of. It was therefore performed by an off-board ADC, as detailed in Section 4.7.2. The digitized voltage measured by the ADC is communicated to the MCU via a 3-wire Serial Interface Peripheral bus.

To reduce the inaccuracy of the analogue measurements that the MCU performs, a dedicated 5 V reference was implemented on the ROB. Since all analogue measurements are referenced to this value, it is of importance that it is stable and precisely known. Therefore, a general-purpose voltage reference ZRB500F01TA was used, and its voltage was determined to good accuracy when mounted on the PCB in order to compensate for offset errors. This way, greater accuracy
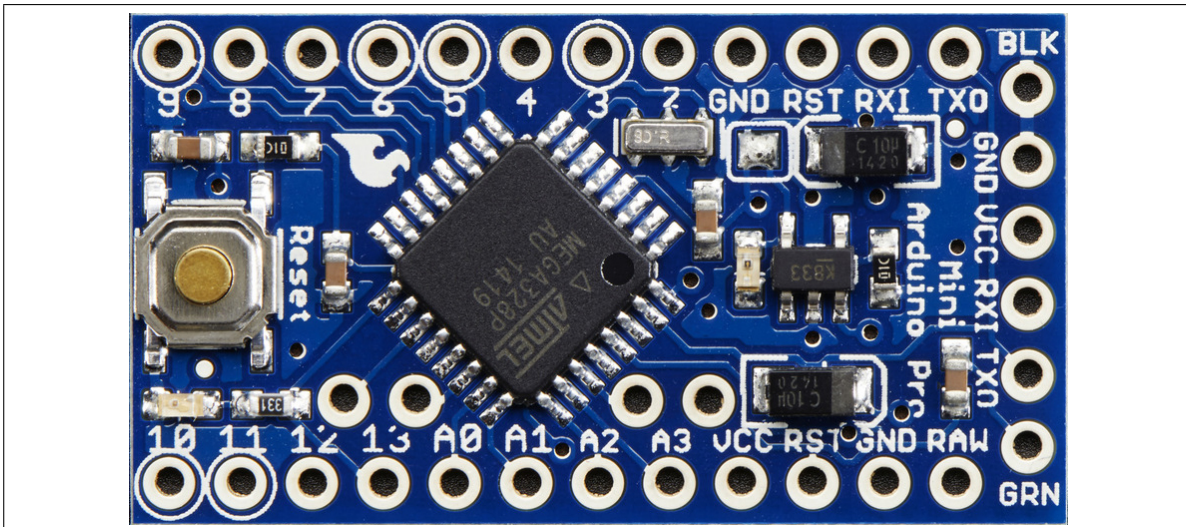
Figure 4.25: The Arduino Pro Mini, part of the open-source Arduino platform for MCU implementation.

is obtained than if the internal voltage reference of the MCU would be used, as there is no information regarding its accuracy nor temperature dependency.

When the MCU measures a voltage, it will compare it to the expected value or expected value range. It then flags the read as either nominal, warning, or error, depending on how far off the expected value or range it is. If the measured value is for a voltage rail, and the read results in an error flag, all voltages on the ROB are immediately turned off.

4.9.2   ROB Control

Control of each individual LDO is necessary due to the fact that different components on the ROB must be turned on or off in a specific sequence in order to not damage them. As an example, the 3.3 V rail must be enabled before the ± 5 V rails so that the sensitive ADCMP553 discriminator is not subjected to input signals from the amplifier before being powered on. It also enables the MCU to disable and restart individual LDOs in case of abnormal behaviour. The only exception is the LDO that supplies the digital 5 V rail, which is located on the Arduino and cannot be disabled as it powers the MCU.

The MCU is able to disconnect the SiPM from the preamplifier input. This is mainly to allow for the insertion of test charges, and to easily halt pulses from individual ROBs without having to fully disable them. Since the digital lines for controlling the insertion of test charges and the MPPC connection enters the RF shield enclosure surrounding the preamplifier, they were heavily filtered in order to prevent high-frequency digital noise from entering the near noiseless preamplifier environment. The resulting risetimes of the digital signals were too slow for the analogue switches which control the MPPC connection and test charges, and Schmitt triggers were therefore added to both lines.

The discriminator trigger threshold can be set by the user, and is controlled by the MCU through a 10-bit PWM output. It is downscaled by a voltage divider with a division factor of 5

to bring the full range down to $0\,V - 1\,V$, resulting in a resolution of $0.98\,mV$. A second 10-bit PWM output adjusts V_{BIAS}, as detailed in Section 4.7.3. Based on the MPPC temperature and base bias voltage at $25\,°C$, the MCU calculates the required value of V_{BIAS} several times per minute. It compares this calculated value with the measured value of V_{BIAS}, and adjusts it if required.

The MCU is able to control the OPA847 amplifier, which is used to disable it during turn-on or turn-off of the $\pm 5\,V$ rails which powers it. Otherwise spurious output signals could been generated by the OPA847. The final controllable element on the ROB is the latch circuit, which can sometimes require a "hard unlatch" by the MCU. This is done by simply bringing an analogue output connected to the latch circuit to $0\,V$ for less than $1\,ms$ (see Section 4.5.2).

4.10   ROB communication

A method of communicating with the ROBs was required in order to allow the user to monitor and control the CRT. The communication scheme needed to be robust, simple, and enable two-ways communication between 16 units and the a PC. Connecting each ROB to a PC with individual USB cables is fine when only a few ROBs are present, but is not practical for the CRT. It was therefore decided that a different method of communication was needed.

A serial communication standard was a natural choice since the ATMega328 has a serial port.   Many serial communication standards exists, with various strength, weaknesses, and difficulty of implementation, but the ones most commonly used for communication between electronic devices are the members of the RS family. These include RS-232, RS-422, and RS-485. Whereas RS-232 only support single-way point-to-point communication, RS-422 support single-way communication from one driver (or "master") to 10 receivers (or "slaves"). This is however not sufficient, as two-way communication between one master and 16 slaves is required for the CRT. The RS-485 standard was therefore the only feasible standard for the CRT, as it support two-way communication between one master and 32 slaves.

4.10.1   The RS-485 standard

RS-485 is a serial communication standard that can be used efficiently across long distances and noisy environments, and support speeds of up to $35\,\mathrm{Mbit\,s^{-1}}$ for distances up to $10\,\mathrm{m}$. There is no limit on the maximum length of an RS-485 bus, and it is therefore widely used in industrial applications to control devices on electromagnetically noisy factory floors. It requires only three wires: Two for the signal, which is sent differentially on a twisted pair, and one for ground. The RS-485 standard does not specify the implementation of the communication protocol, only the electrical characteristics of the transmitter and receiver. There is therefore great freedom in how the standard is implemented, and it can be tailored to fit specific applications.

4.10.2   Physical implementation

The ROBs were connected to a common RS-485 bus controlled by a network master device. The cables used for the bus were common Cat5 Ethernet cables, since they have the required twisted pairs and are ubiquitous and low-cost. The optimal choice of network master would be the controller PC, but this was not possible since PCs rarely have output ports that support RS-485. Instead, an Arduino module was used for the role as network master. The choice fell on the Arduino Micro, as it is the cheapest model with two serial ports which is require for it to be able to simultaneous communicate with the PC over USB and with the ROBs over RS-485. Its role is to mediate commands and data between the PC and the RS-485 network, and the communication protocol is implemented in its software. An illustration of the network can be seen in Figure 4.26.
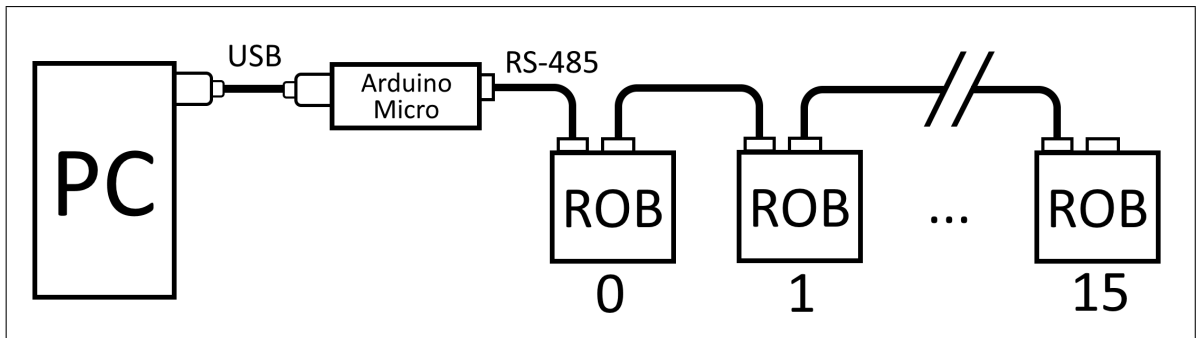
Figure 4.26: The communications network of the CRT. The controller PC communicates with an Arduino over USB, which acts as the master of the RS-485 network. The arduino will translate commands from the PC into data packets for the RS-485 network, and vice versa.

4.10.3 Communication protocol

The protocol for communication between the ROBs and the master required the following:

- Data packets on the order of tens of bytes.
- Addressing, so packets can be sent to specific units.
- Noise immunity, to minimize the probability of corrupted data.
- Error checking to prevent corrupt data from being interpreted.
- Support for master-to-ROB and ROB-to-master communication.
- Handshaking, to confirm data packages were received correctly.

Several of these requirements were implemented in an RS-485 software protocol written by Nick Gammon [60]. It allows for data packages of up to 255 bytes, with rigorous error checking and noise immunity. With this protocol a data packet begins and ends with a Start Of Transmission (STX) and End Of Transmission (ETX) byte, 0x02 and 0x03 respectively. Each byte of the packet is sent in a doubled-inverted form, so for example the byte 0x1B is sent as two bytes, 0x1E and 0xB4. This reduces the number of valid bytes from 256 to 16, and makes it possible to detect corrupt data, since noise only has a 16/256, or 6 %, chance of producing a valid byte. An additional Cyclic Redundancy Check (CRC) is performed on the data packets. The transmitter calculates a unique checksum for the data and sends it with the data packet. When it is retrieved by the target unit, it is recalculated and compared to what the transmitter calculated it to be. If they do not match, the data is known to be corrupted.

The rest of the requirements were implemented by the author. For addressing, all ROBs has a unique ID stored in the code of its MCU. All data packets sent by the master contains an addressing byte at a specific location. When a ROB receives a data packet, it will compare the addressing byte to its own ID, and if it does not match then the packet is ignored. For a detailed description of the CRT communication protocol, see Appendix C.

To ensure that packets are received correctly, a ROB on the RS-485 network will reply with an acknowledgement when it receives a packet meant for it. This makes it possible for the transmitter to confirm that packets are received correctly. If the error-checking algorithm of the protocol discovers that the data is corrupt, it will not reply with an acknowledgement. The transmitter will register this, and attempts to re-transmit the packet up to two times.

To prevent multiple units from transmitting at the same time, the communication protocol was written so that a ROB may only transmit data after the master requests it to do so. After the master has transmitted a packet, it will enable its receiver and wait for the acknowledgement from the ROB. When that has been received, it will then wait a certain amount of time for the ROB to interpret the packet and transmit the results of what the packet requested of it. When the results have been received, it goes back into transmit mode. The ROB is in other words only allowed to transmit under well-defined circumstances. If the master does not receive the results of the data interpretation after some set time-out, an error likely occurred with the ROB. In this case, an error message is returned to the PC.

4.11 MPPC and amplifier chain performance

The preceding sections have mainly focused on the theoretical aspects of the ROB and MPPC. In particular the sections regarding the preamplifier and amplifier, Sections 4.3 and 4.4, relied heavily on simulations of both the circuits and the MPPC. In this section, the actual performance of the readout chain will be compared to its theoretical performance. The amplifier chain frequency response will be presented, and MPPC pulses will be compared to those found through simulations in order to verify that the design choices made were based on a correct understanding of the MPPC. Additionally, dark noise pulse height spectrums are presented and compared to the pulse height spectrums for the ROB mounted on a scintillator.

4.11.1 Amplifier chain response

The amplifier chain consists of the preamplifier, amplifier, and associated filters. It is responsive for amplifying and shaping the pulses generated by the MPPC, and its design and theoretical behaviour were discussed in Sections 4.3 and 4.4. Measurements of its frequency response will now be presented and compared to its theoretical frequency response.

Experimental setup

The measurements were performed by connecting a signal source to the preamplifier input in place of the MPPC, which had been removed for the sake of performing these measurements. Simulations showed that the preamplifier input has an impedance in the range of $3\,\Omega$ to $7\,\Omega$ in the frequency range of interest, and so an additional $47\,\Omega$ SMD resistor was soldered onto the preamplifier input in order to bring the input impedance near $50\,\Omega$. If the input had instead not been impedance-matched to the coaxial cable, reflections in the cable would have occurred as the electrical signal encountered the sudden change in impedance. Interference between the signal and reflection would have caused a frequency-dependent periodic variation in the input amplitude, making it harder to evaluate the amplifier chain frequency response.

The drawback of impedance matching the input to the coaxial cable is that it changes the frequency response of the preamplifier. The combination of an additional $47\,\Omega$ resistor in series with the inherent input impedance of the preamplifier significantly reduces the overall gain and flattens the peak somewhat, as can be seen in Figure 4.27. The measurements therefore have to be compared to this frequency response, and not the original frequency response of the unmodified amplifier chain. However, as it is only the input impedance that is adjusted and not the active components, this comparison will still provide a fair indication of the behaviour of the unmodified amplifier chain.

The output of the amplifier chain was read from the "MPPC signal output" coaxial connector in Figures 4.1 and 4.3. The act of connecting a coaxial cable to this output slightly decreases the gain for higher frequencies and reduces the center frequency f_C, as seen in Figure 4.27. Normally the primary amplifier drives a $50\,\Omega$ load to ground, but by connecting a coaxial cable
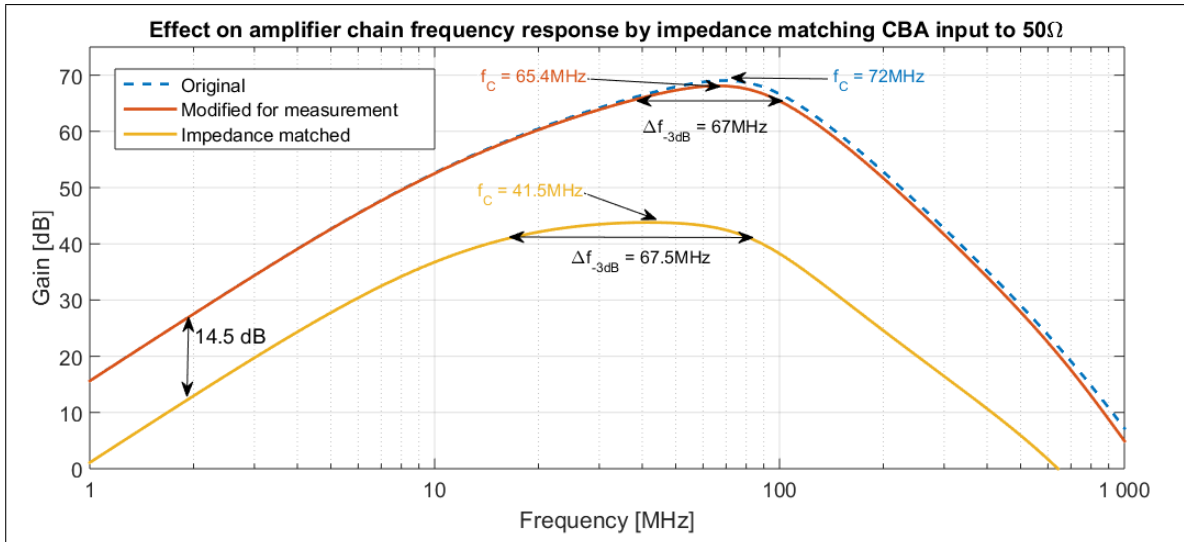
Figure 4.27: Simulated amplifier chain frequency responses for $I_C = 33\,\text{mA}$. The dashed blue plot represents the frequency response of the unperturbed amplifier chain as seen in Figure 4.18. The act of measuring the op amp output halves its load to $25\,\Omega$, decreasing the center frequency f_C. Impedance matching the CBA input significantly reduces the overall gain and further decreases f_C.

to the output, it must instead drive two parallel $50\,\Omega$ loads. The reduction in load resistance decreases its cutoff frequency, and the center frequency decreases too.

Measurements were taken at varying input frequencies, from 1 MHz to 2 GHz. The input signal amplitude was kept small (normally 1 mV) to prevent the amplifier chain from saturating, and most importantly, to prevent the output signal from being clipped or otherwise distorted by the rectifying behaviour of the op amp.

Due to limitations of the equipment available, the measurements had to be performed with two different sets of equipment. Measurements in the frequency range from 1 MHz to 80 MHz were performed manually with an Agilent 33250A arbitrary waveform generator and a Tektronix MSO 4034 mixed signal oscilloscope. The Agilent 33250A is capable of generating sinusoidal waveforms up to a frequency of 80 MHz, which set the upper limit to the frequency range that could be measured with this equipment. The output signal amplitude was measured with the Tektronix MSO 4034, which has a bandwidth of 350 MHz. For measurements performed in the range from 50 MHz to 2 GHz an Hewlett Packard 8722C network analyser was used.

Results

Measurements of the amplifier chain frequency response was performed at two different preamplifier gain configurations. As was shown in Section 4.3.4, the frequency response will vary based on the magnitude of the collector current I_C. Therefore, to confirm this behaviour, the frequency response was measured at two different values of I_C, $33 \pm 1\,\text{mA}$ and $1.7 \pm 0.1\,\text{mA}$.

The results can be seen in Figure 4.28. The overall measured frequency response fits well with the expected frequency response found by simulations. For $I_C = 1.7\,\text{mA}$ the measured frequency response lacks some gain in the frequency range 30 MHz to 100 MHz, but overall it agrees very well with the simulated response.
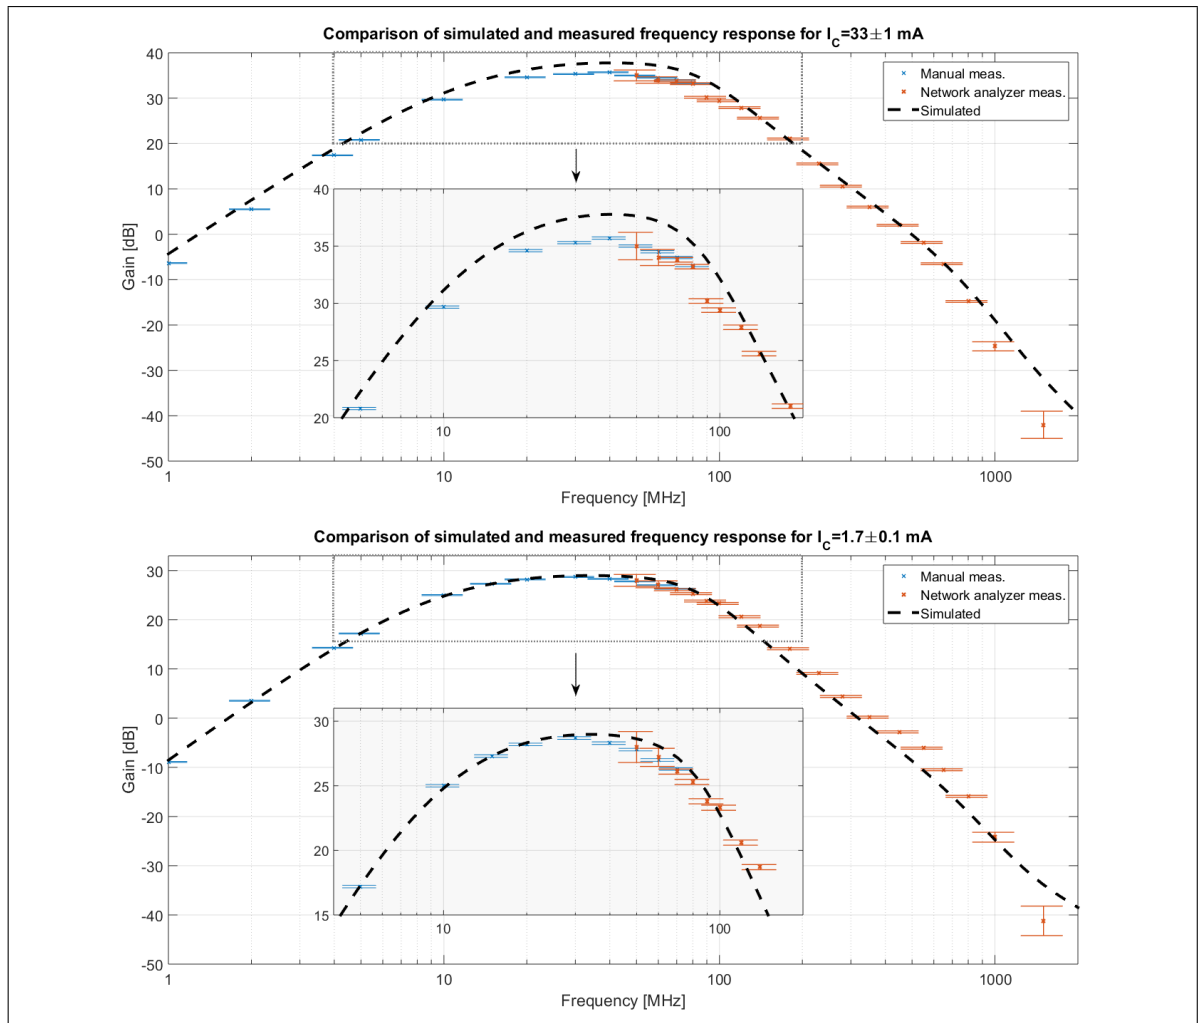
Figure 4.28: Simulated and measured frequency responses for the amplifier chain for two different values of I_C.

For $I_C = 33\,\text{mA}$, the results are not as good. There appears to be an overall lack of gain of $1\,\text{dB}$ to $2\,\text{dB}$ over the entire measurement range, and this lack of gain is greatest at the center frequencies from $30\,\text{MHz}$ to $100\,\text{MHz}$. This was also the case for the measurements made with $I_C = 1.7\,\text{mA}$, but then the lack of gain was not as great.

The discrepancies may be due to stray and parasitic effects which have not been modelled in the simulations. These effects arise from the fact that the passive components that are incorporated in the amplifier chain will all have some small undesired electrical properties. As an example, consider the resistor. There will be some amount of capacitance between its terminals, and between the terminals and PCB. If it is a wirewound resistor, it will also have a non-negligible inductance. The same applies to capacitors and inductors. Since they are real components, it is impossible to eliminate all the undesired properties. They can be difficult to estimate, and are often neglected in simulations, where passive components are normally modelled as being perfect by having only the desired property, whether it be resistance, capacitance, or inductance.

4.11.2   MPPC pulses

MPPC waveforms after passing through the amplifier chain were collected in order to compare them to the waveforms found by simulations of the MPPC and amplifier chain. Of particular interest was the waveform generated from a single G-APD cell firing, also called a 1p.e. (photoelectron) pulse. These pulses are the main constituents of SiPM dark noise, as explained in Section 3.3.2. 1p.e. pulses occur rapidly and their waveforms do not vary much, making them a good candidate for comparison in order to evaluate the MPPC model and the amplifier chain.

Large MPPC pulses generated by impeding scintillation photons were also recorded in order to compare their shape to what is expected from the scintillator simulations in Section 2.4.3 and the effect of filtering MPPC pulses as shown in Section 3.6.

Experimental setup

The MPPC pulses were sampled with a Tektronix MSO 4034 mixed signal oscilloscope connected to the "MPPC signal output" coaxial connector. All measurements were performed on a single ROB mounted on a scintillator, and the MPPC bias overvoltage was set to 3 V. For measurements of 1p.e. pulses, I_C was set to 1.7 ± 0.1 mA and 33 ± 1 mA. The oscilloscope trigger settings were set to only accept pulses of less than 6.5 ns, which ensured that it only 1p.e. pulses were registered.

For measurements of large MPPC waveforms generated by muons crossing the scintillator, the trigger threshold was set to a value of 250 mV in order to not trigger on dark noise, and I_C adjusted to a medium gain of 5.5 ± 0.1 mA.

Results

For both high and low gain, the sampled 1p.e. pulses agrees well with the simulated waveform, as can be seen in Figure 4.29. The overall shape and amplitudes are similar, but in the high gain configuration the risetime is greater than what is expected. This is likely due to a lack of gain from 20 MHz to 100 MHz, as was seen in the measured frequency response for the high gain configuration. This causes low frequencies to be amplified more strongly relative to high frequencies, and the signal risetime increases.

It is unclear why the measured waveforms do not descend as far as the waveforms found by simulations after the initial pulse has occurred. There does in fact appear to be a second, smaller, pulse superimposed on the first pulse 9 ns after it has peaked. Initially, it was thought that this came from reflections in the coaxial cable that carried the pulse from the ROB to the oscilloscope, but this theory was discarded when the coaxial cable length was greatly increased, and no change in the waveform was observed. Another theory was that since the recorded waveforms are the average of many 1p.e. pulses, the additional peak was due to noise or incorrectly triggered pulses. However, every individual 1p.e. pulse also displayed the same additional peak, indicating that it is real and not an artefact introduced by the measuring equipment. It is likely due to some characteristic in the amplifier chain that was not included in the simulation model. But since it does not affect the initial pulse, there was no incentive to
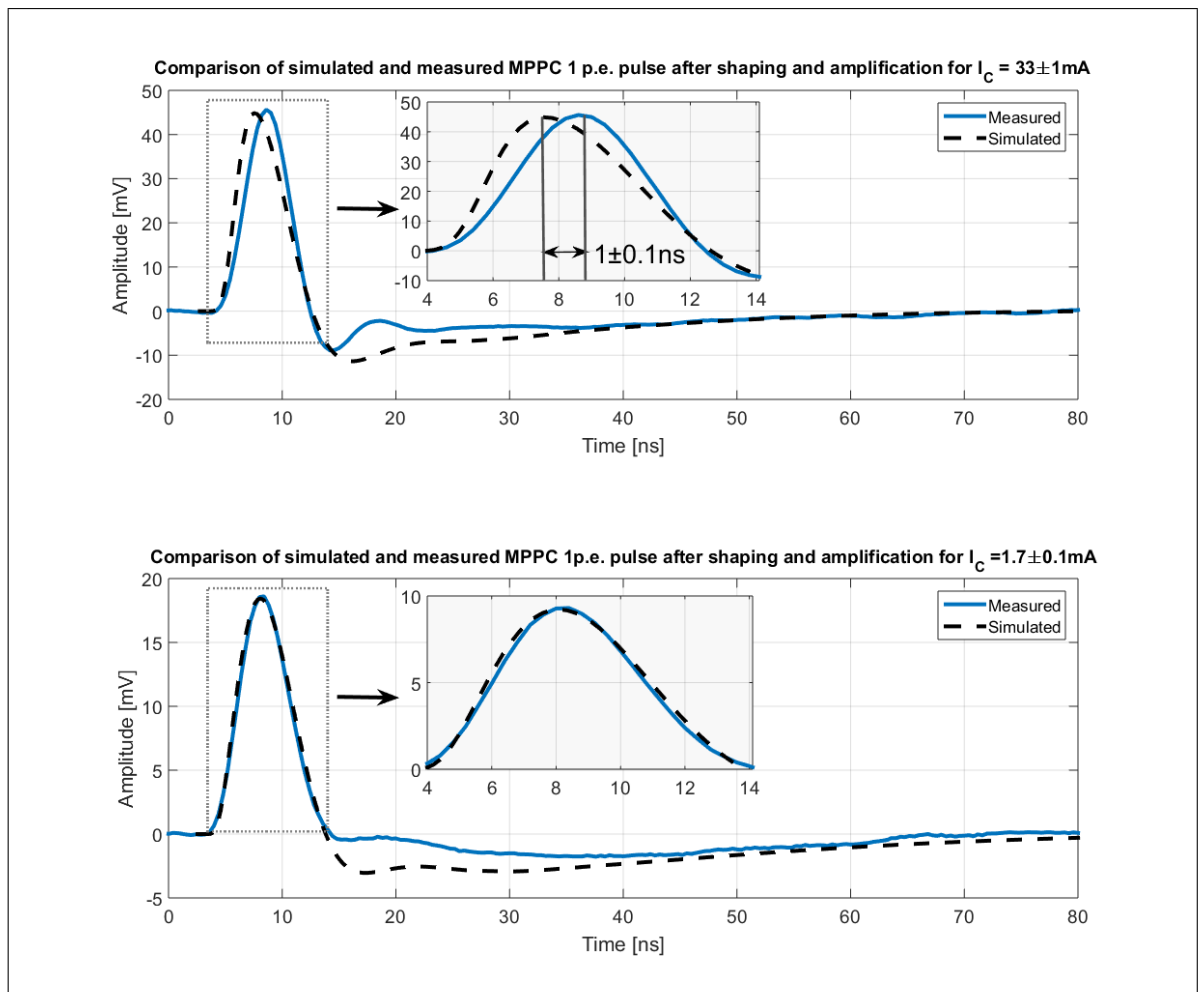
Figure 4.29: Comparison of measured MPPC 1p.e. pulses and their simulated equivalents, after amplification and filtering. The measured pulses are the average of 512 individual pulses. The high gain configuration displays a discrepancy in the signal risetime due to a lack of gain at higher frequencies.

investigate it further.

A typical MPPC pulse generated by impeding scintillation photons, which in turn were generated by a charged particle crossing the scintillator, can be seen in Figure 4.30. The first thing that is apparent are the repeating peaks, which are a result of the fact that the scintillation photons arrives in groups that have been reflected off the scintillator ends. This effect was expected from the simulations discussed in Chapter 2. The periodicity of the peaks agree well with the simulations, which indicated that the reflected groups would have a periodicity of ~12.5 ns. We also see that the MPPC pulse is reminiscent of the waveforms found by simulations in Section 3.6, where the effect of filtering raw MPPC pulses were investigated as a way to improve the trigger timing.
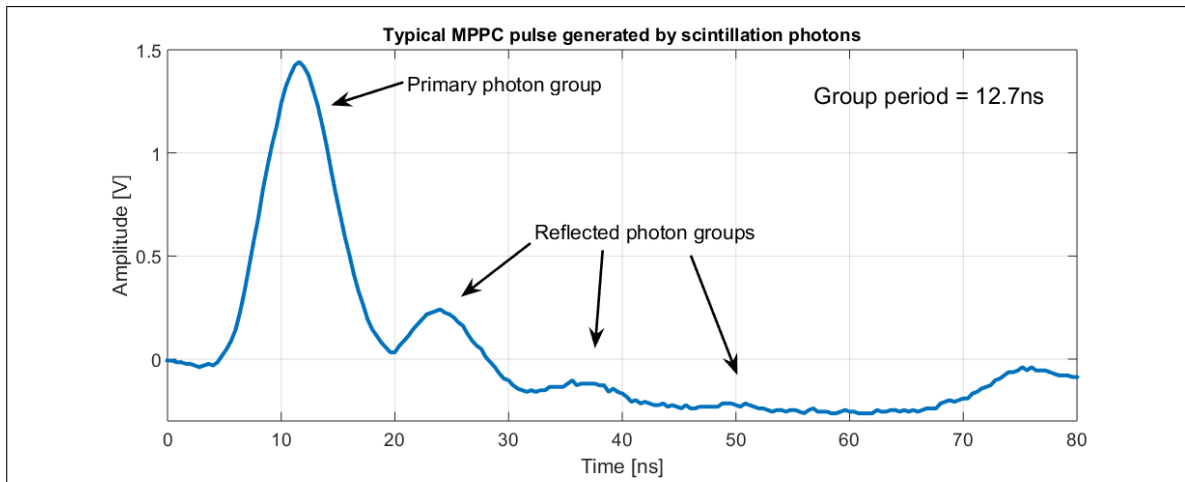
Figure 4.30: A typical MPPC pulse generated by a charged particle, likely a muon, passing through the scintillator of which the MPPC is mounted on. The additional peaks after the first shows the effect of scintillation photons arriving in groups reflected off the scintillator ends.

4.11.3   MPPC pulse height spectrum

The pulse height spectrum of the MPPC pulses after passing through the amplifier chain was measured for various bias overvoltages. This provided us with an understanding of how the MPPC dark noise vary with overvoltage, and how the height spectrum of the pulses of interest (i.e. those generated by scintillation photons) are distributed with regards to the dark noise. Optimally, a clear divide in the pulse heights of dark noise and scintillation pulses would be found. The hope was that the knowledge gained would aid in setting the overvoltage and trigger threshold in order to optimize the compromise between triggered dark noise, gain uncertainty, and detection efficiency.

Experimental setup

A ROB with a mounted MPPC was placed within a darkbox to protect it from ambient illumination. Electrical wires for power, communication, and signals passed through lightproof feedthroughs. The preamplifier gain was kept constant at $I_C = 35 \pm 1\,\mathrm{mA}$ for all measurements, and the MPPC temperature was given time to stabilize before measurements were made.

The MPPC pulses were extracted from the amplifier output via coaxial cable and passed to a CAEN V1729A 14bit 2GS/s digitizer. A Labview application developed by Andreas Tefre Samnøy (PhD Candidate, IFT) for the purpose of characterizing electrical pulses was used to acquire the pulse height spectrum. However, due to limitations in the hardware, the trigger rate was at most 600 Pulses/s, while the dark noise rate of the MPPC was on the order of 2 MPulses/s. This meant that the vast majority of pulses were not registered, and it was therefore not possible to make any assumptions about dark noise rate for any of the measurements. All pulse height spectrums are therefore normalized, and only the shape of the spectrums are of interest, not their absolute amplitudes.

The MPPC pulse height spectrums were measured for various applied overvoltages, and in all cases the CAEN V1729A trigger threshold was set to half the height of a 1p.e. pulse. Two
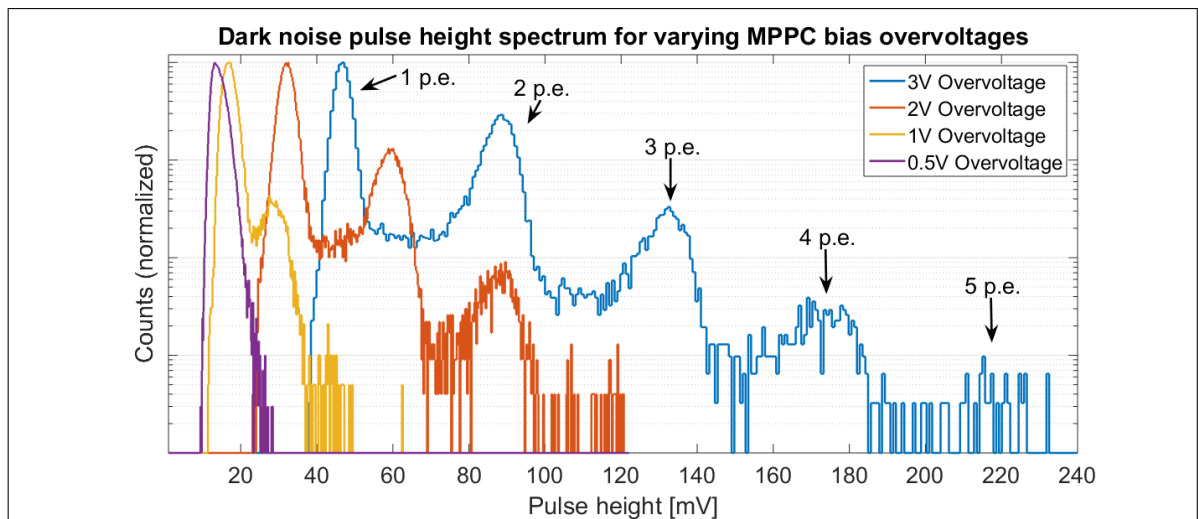
Figure 4.31: Pulse height spectrum of the Hamamatsu S13360-6075CS SiPM for various overvoltages. With decreasing overvoltage, the rate of 2p.e., 3p.e., 4p.e., etc, dark noise decreases relative to the 1p.e. dark noise rate. The spectrums have been normalized to the same 1p.e. peak height.

sets of measurements were made: One with the ROB by itself in the darkbox to measure the dark noise, and one where the ROB was mounted on a scintillator.

It was necessary to perform several individual measurements with increasing digitizer trigger thresholds in order to compare the dark noise pulse height spectrum to the pulse height spectrum obtained with the MPPC mounted on a scintillator. This was due to the exponential dependency of dark noise rate to pulse height, and the great difference in dark noise rate (~ 2 MPulses/s) and pulses generated by scintillation photons (~ 40 Pulses/s). The limited trigger rate of the digitizer of 600 Pulses/s meant that a pulse generated by scintillation photons would only be registered on the order of once every minute, and obtaining enough statistics would take weeks. Therefore, the spectrums were instead obtained in steps of increasing digitizer trigger thresholds. On each increase, the threshold was set to a value between the previous largest and second-largest photo-electron peak. If no distinguishable photo-electron peaks existed, it was set at a value where sufficient overlap would exist between the new and previous spectrum, so that they could later be combined. With this technique, pulse height spectrums spanning over 7 orders of magnitude in amplitude were easily obtained.

Results

The pulse height spectrum of the MPPC dark noise can be seen in Figure 4.31. The spectrum consists of easily discernible individual peaks. These peaks correspond to the number of G-APD cells that fired simultaneously. With decreasing overvoltage, two things occur: The spectrum shifts downwards, and the dark noise rate decreases. The shift is due to the MPPC gain being directly related to the overvoltage, and the decrease in dark rate is in accordance with the theory in Chapter 3 and various literature [34, p.8] [61, p.17-26].

From the pulse height spectrums obtained with the ROB mounted on a scintillator it became obvious that a large fraction of the pulses generated by scintillation photons were indistinguishable from dark noise pulses. It was observed that the noise spectrum and signal spectrum
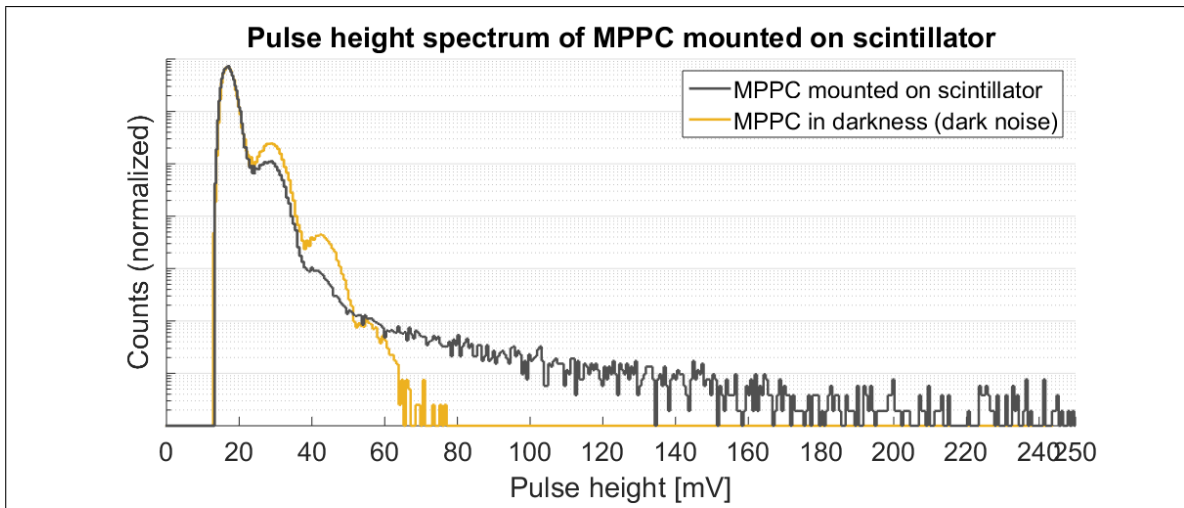
Figure 4.32: Two MPPC pulse height spectrums, one for the MPPC in complete darkness, and the other with the MPPC mounted on a scintillator. The two spectrums have been normalized to the same 1p.e. peak amplitude.

overlapped for all tested overvoltage, and therefore no trigger threshold can be chosen to completely divide dark noise from signals.

Pulse height spectrums for an overvoltage of 2 V can be seen in Figure 4.32. For pulse heights below 3p.e., dark noise is the major contributor the total rate. Pulse height spectrums for all other overvoltages from 0.5 V to 3 V showed similar characteristics. It was not possible to infer an optimal overvoltage and trigger threshold based on these results alone. This prompted the development of a more direct method of assessing the effect of varying overvoltage on the performance of the CRT. In this method the coincidence rate of two overlapping scintillators are characterized. This is done by measuring the rate of pulses that coincidence or not for varying trigger thresholds and overvoltages. This method is discussed in Section 6.2, and it is used to find the overvoltage which optimizes the coincidence rate.

5 | The data acquisition system

The current chapter will detail the system that was developed for acquiring, analysing, and storing the data which the CRT produces. This system is fully digital, and consists of a PC running a Labview application called the CRT Controller (CRTC). The CRTC is an application that incorporates every capability required for the operation of the CRT. The operator of the CRT can manage the detector exclusively via this application.

The pulses generated by the ROBs are transmitted over coaxial cable to the PC, and are acquired with a 16-channel time-to-digital converter (TDC) installed in the PC. This acquisition hardware precisely timestamps the time of arrival of each pulse, and the CRTC analyses these timestamps in order to digitally detect coincidences between scintillators. The coincidences, as well as other data, are permanently stored on harddisks. A top-level overview of this system can be seen in Figure 5.1.
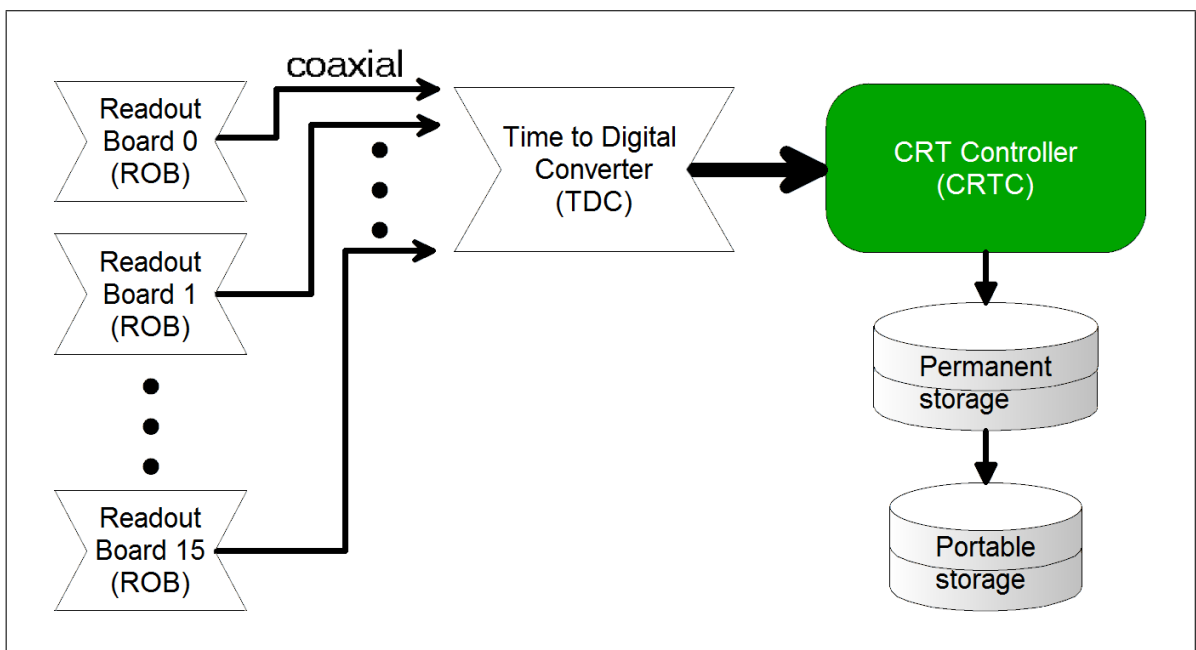
Figure 5.1: Top-level overview of the CRT data acquisition, analysis, and storage. A 16-channel time-to-digital converter timestamps the arrival of pulses from the ROBs. The timestamps are analysed by the CRTC, and coincidences stored on two harddisks.

5.1   Labview

The CRTC was developed with Labview, which is an integrated development environment by National Instruments that enables quick and easy development of data acquisition and control software. It allows the user to create programs with the graphical programming language G, which uses the paradigm of *dataflow*. Unlike traditional programming languages which model a program as a sequence of commands happening in a specific order, Labview and other dataflow languages model a program as a series of connections in which data flows between blocks with inputs and outputs. In Labview terminology, such a block is called a Virtual Instrument (VI). When all the inputs of a block contain valid data the block executes, and when it has finished executing it outputs the results of the actions it has performed. Dataflow languages are inherently parallel, as any number of blocks may execute in unison. This method of programming is especially well suited for data acquisition and analysis, where the focus of the program is the manipulation of large amounts of constantly arriving, or "flowing", numerical data.

A Labview application usually consists of a block diagram where the user develops the program, and a graphical interface, called a "virtual control panel" in Labview terminology. This control panel is designed by the developer and may contain buttons, graphs, tables, indicators, images, and many other types of elements. With this control panel the user can easily, and in realtime, display data and control the program.

The extremely fast development time of Labview applications, and the ease of which data can be displayed, makes it the preferred method for developing laboratory data acquisition and control software. Labview was chosen to develop the CRTC for these reasons.

5.2   Time to Digital Converter

A TDC is a specialized piece of hardware that is capable of accurately time-tagging electrical pulses. Unlike an ADC it does not sample the amplitudes, but instead only accurately registers when an input signal crosses some level. The TDC acts as a link between the ROBs and the controller PC, digitizing the time of arrival for each ROB pulse so that it can be further analysed and compared with the time-of-arrival of other pulses to look for coincidences. We chose to purchase the TDC as a pre-made working solution instead of designing and building it ourselves for the reasons that a TDC is a relatively complicated piece of hardware, has a great impact on the CRT capabilities, and requires high-speed data transfer to a PC. This was deemed above our capabilities.

We were lucky to receive a donation from Cronologic GmbH in the form of two greatly discounted units of their HPTDC8-PCI model TDCs, 8-channel TDCs designed to be placed in a standard PCI expansion slot of a PC. An image of the TDC can be seen in Figure 5.2. This made the TDC implementation quite easy, as it was simply a matter of inserting the two TDCs in free PCI slots on the controller PC and developing Labview wrappers for the drivers which came with them.

Each of the HPTDC8-PCIs are capable of registering up to 4M pulses ("hits") per second with an accuracy up to 25 ps [62]. These capabilities are beyond what is required for the CRT, which is expected to have a rate no greater than a few thousands hits per second per channel, and a timing resolution set in part by the finite risetime of the filtered SiPM pulse, which is roughly a factor 100 greater than the resolution of the HPTDC8-PCI.

5.2.1   Drivers

The driver that was supplied with the HPTDC8-PCI came in the form of a Dynamic Link Library (DLL). Labview wrappers for this DLL file had to be made manually, and took some effort. A library of Labview VIs was developed which enabled access to most of the TDC functions, such as initiating, starting, stopping, reading, configuring, and more. Access to these functionalities were protected by a common semaphore built in to each VI, preventing the possibility of multiple TDC functions being called simultaneously.

An unforeseen bug in the way that Labview receives data from DLL files prevented us from fully utilizing the readout rates of the HPTDC8-PCI. A specific function in the DLL, `ReadTDCHit()`, takes a pointer to an array of `TDCHit` .NET references, and fills this array with timestamps. However, when this function is incorporated in Labview, the array is returned as a copy of the original array, which then has to be dereferenced. This dereferencing takes a finite amount of time which is great compared to the time it takes to receive the TDC data, and it therefore sets a limit on how quickly hits can be read out. For this reason, the readout rate was limited to ~150khps. This proved to not be a problem, as it was the timestamp analysis that in the end proved to be the bottleneck, limiting the maximum readout rate to $\sim 9 \times 10^4 \, \text{s}^{-1}$.

Figure 5.2: The Cronologic GmbH HPTDC8PCI Time to Digital Converter. Two units were used to digitize the time of arrival of the ROB output pulses to an accuracy of 25 ps.

5.3   CRT Controller software

In this section the CRTC will be discussed in more detail. The data path from low-level timestamp acquisition to the storage of detected coincidences will be detailed, and parts of the user interface presented.

5.3.1   ROB control

Through the CRTC, the operator has full control of the ROBs. He or she can adjust all ROB parameters, set ROB states, control individual voltages, insert test charges, troubleshoot problems, and easily obtain information about a ROBs current status through the status display seen in Figure 5.3.

The ROB status display shows information about each ROB, such as the temperature of the MPPC and whether it is within limits. Also the values of the on-board voltage rails are displayed, together with whether those voltages are within their limits and whether they are disabled or enabled. Additionally,the trigger threshold, how many times the latch has been reset, and other types of information is displayed too.

This capability is enabled through the communication network that links all ROBs to the PC, as was discussed in Section 4.10. The CRTC decodes the data packets it receives from the ROB network master and updates the ROB information display as needed. If errors would occur, a message is displayed to the user. This can include such errors as a ROB not accepting a command, being unable to execute a command, a data packet not being delivered correctly, or any other of the 22 errors that can be reported, and which have been listed in Appendix C.2.2. In the case that a ROB automatically shuts off due to out-of-limits voltages or otherwise becomes unresponsive, an email will be sent to a list of recipients so that an operator may intervene.

5.3.2   Scintillator identification scheme

The scintillators, or channels, of the CRT are identified in the CRTC from 0 to 15, counting clockwise and from the bottom plane to the top plane. Scintillators 0-3 and 8-11 are labelled as "primary" channels, while scintillators 4-7 and 12-15 are labelled as "secondary" channels. There is no physical difference between primary and secondary channels, and the designation is made purely for aiding analysis.

The CRT has two planes, and each plane consists of two scintillator layers. The area covered by an overlap of two scintillators is called a pixel. Pixels are numbered from 0-15 on the bottom layer and from 16-31 on the top layer. Coincidences are saved to permanent storage with the pixel ID in which the coincidence was detected. The identification scheme of the CRT can be seen in Figure 5.4.
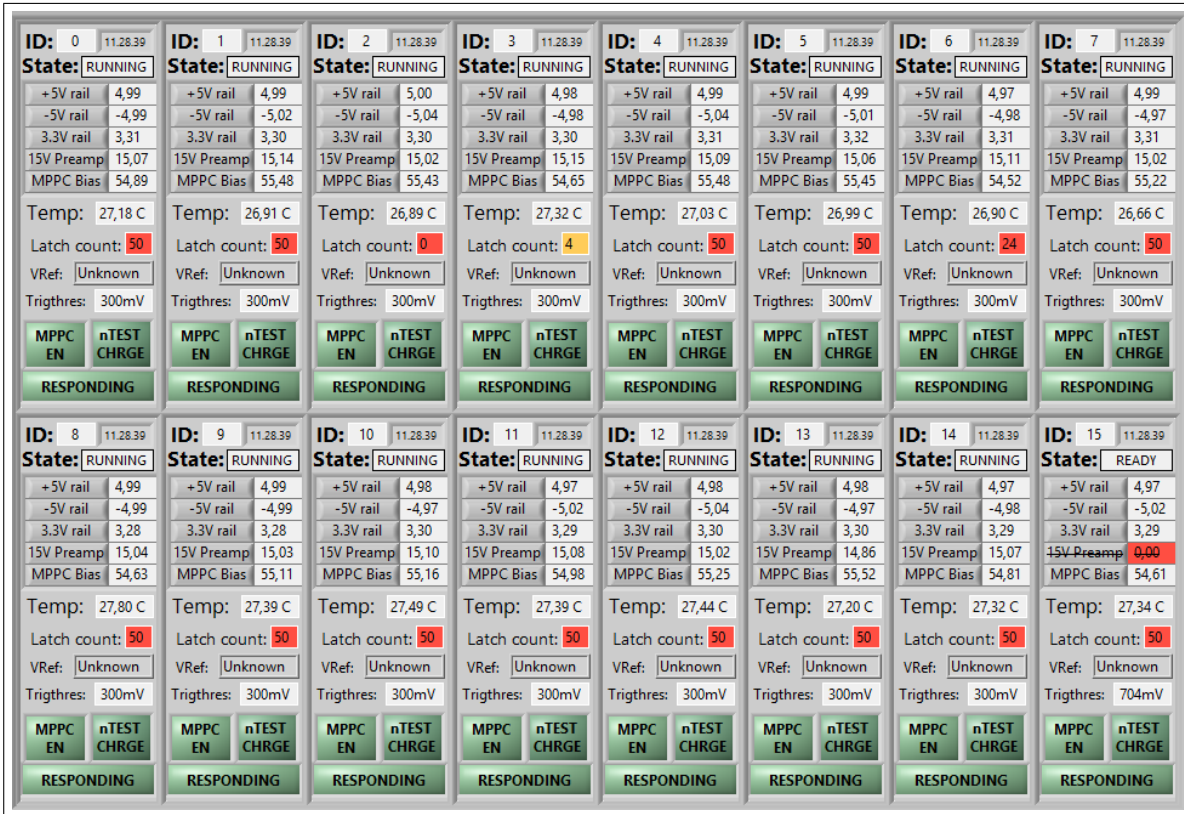
Figure 5.3: The ROB status display found on the CRTC front panel. Information about each ROB is displayed in its own compact container. ROB 15 (bottom right corner) has been set in mode READY to exemplify how it changes to signify that a voltage is out of limits (red background) and disabled(strikethrough).
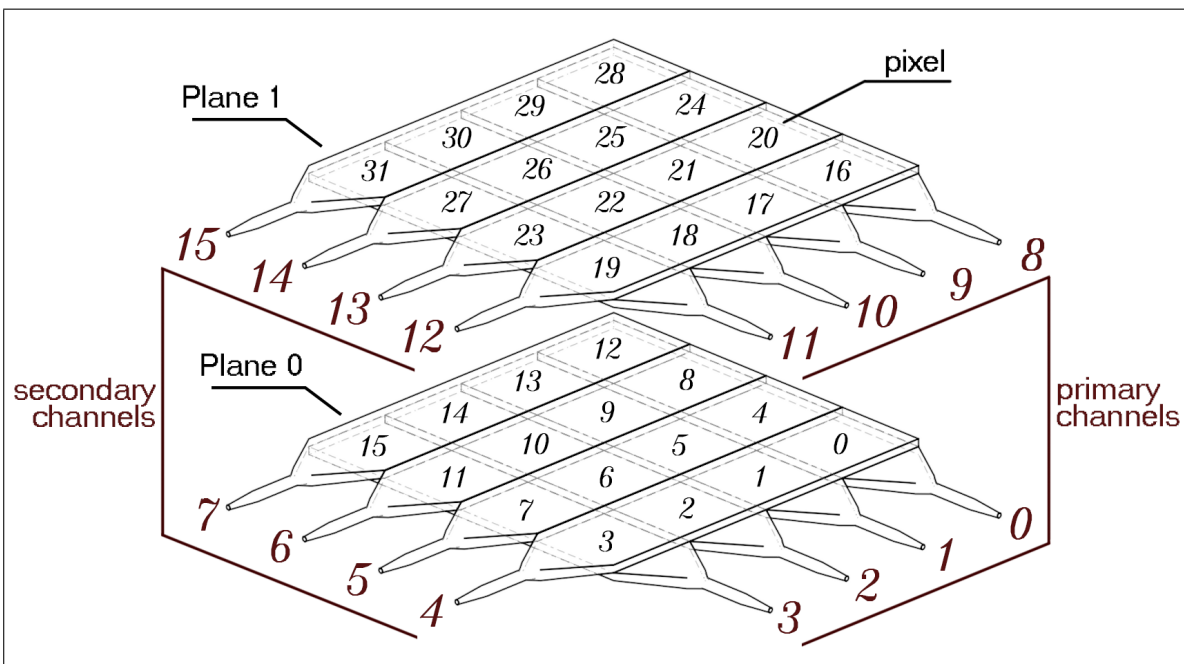

Figure 5.4: A diagram of the CRT showing how the scintillators, their overlapping areas ("pixels"), and the planes are identified.

5.3.3 Readout chain

The part of the CRTC which is responsible for acquiring timestamps, manipulating them in order to detect coincidences, and finally storing and displaying the results, is called the readout chain. The readout chain was designed to be flexible in order to cope with variable muon rates so that no hits would be lost in the case that the readout chain is at its maximum computing capacity. Therefore, it was designed as a series of independent parallel processes linked by flexible buffers called data queues, and in part by global functional storage buffers. Each independent process has its own task, such as reading out the TDC timestamps, sorting them, locating coincidences, and so forth. A process generally executes when it receives data from the previous process through the queue linking them. In the event that a process is unable to finish processing the the previous block of data before new data arrives, the queue or buffer stores the new data until the process is ready. This allows for increased data rates to be supported for some amount of time until the queue overflows and data is lost. In that case a warning is displayed to the user.

The readout chain consists of the following processes:

- TDC timestamp readout, `Read.vi`
- Timestamp channel sorter, `Sort.vi`
- Coincidence finder, `PlanarCoincFinder.vi`
- Coincidence information displayer, `DisplayCoincData.vi`
- Coincidence storage, `SaveCoincData.vi`

They are chained by data buffers as just explained, and in the order presented in the list. Each of these processes will now be further detailed.

Read.vi

`Read.vi` runs continuously, polling the TDC for new data. On each poll, the digitized times of arrival, or timestamps, on all 16 channels are returned. They must then be modified, as the timestamps are signed 64-bit integers with 1 ps resolution, and this value would overflow after 106.75 days of continuous readout. To prevent this, the timestamps are bit-shifted five places to the right, and the five upper bits are used to count the amount of times the original timestamp overflowed. This modified timestamp is an unsigned 64-bit integer with 32 ps resolution, and will overflow after 18.72 years. To make the timestamp absolute, a number equal to the amount of 32 ps intervals that has occurred since 00:00 1. January 2017 is added. This enables the CRTC to uniquely timestamp hits for over 18 years.

Some constant amount of time, on the order of tens of nanoseconds depending on the channel, is then subtracted from each timestamp in order to compensate for the propagation delay in the coaxial cables that carries the electrical pulses from the ROBs to the TDCs. The propagation delays were measured to an accuracy of ± 5 ps by the manufacturer, and those values are used for the compensation. The TDC channel ID of each hit is modified to fit the numbering scheme of the CRT scintillators, and the final timestamp with correct channel ID is then appended to a queue connecting `Read.vi` to `Sort.vi`.

Sort.vi

Sort.vi will sort each timestamp it receives from Read.vi into a functional global variable named TDCHits.vi, implemented as a rolling buffer. This buffer can store up to 10000 of the latest hits on each of the 16 channels. When TDCHits.vi overflows, the oldest hits are overwritten. When TDCHits.vi has been filled to the point where there are more than some specified amount of hits in any channel in one of the scintillator planes, a notification is sent to PlanarCoincFinder.vi, which initiates a coincidence "hunt" in either plane 0 or plane 1, see Figure 5.4. The amount of hits required for a coincidence hunt is flexible, and will vary with the TDC readout rate so that the time between two hunts always lies in the range of 5 ms to 50 ms. This automatic variation ensures that coincidence hunts are issued at a reasonable rate, balancing the need for rapid processing of hits and the overhead related to initiating a coincidence hunt.

PlanarCoincFinder.vi

The process of locating coincidences between two scintillators/channels in a plane is managed by PlanarCoincFinder.vi. Two of these processes are running in parallel, one for each of the CRT planes. When PlanarCoincFinder.vi receives a notification from Sort.vi to perform a hunt it extracts from TDCHits.vi all timestamps from all channels in the plane it handles. This collection of timestamps is called a *data frame*, and can consist of several thousand timestamps per channel. The timestamp of each hit is given a boolean flag, initially set to False, to denote whether it has been counted as a coincidence, and a private array which will be used to contain all those coincidences the specific timestamp becomes a part of. The data frame is then split into two groups, one for the primary channels and one for the secondary channels.

A coincidence hunt is then performed for each of the four primary channels in parallel. To do so, the timestamps of each primary channel is first combined with the timestamps of all secondary channels, called a *collection*. For each of the four collections a "leftover" array, which are those timestamps from the previous processed collection which may form coincidences with timestamps in the current collection, is added to the beginning. This modified collection is then supplied to a process named FindCoincidences.vi. For each primary timestamp, it compares the secondary timestamps to it in order to find two timestamps on a primary and secondary channel which occur within a specified amount of time of each other, called the coincidence window.

A clever algorithm was implemented to make this process faster. A naïve algorithm would, for each primary timestamp, compare each secondary timestamp to it. This is inefficient, as there can be thousands of timestamps, resulting in millions of comparisons. Instead, one can use the fact that all timestamps in the timestamp arrays are sorted in time, i.e. that the timestamps are always increasing in value. Therefore, for each primary timestamp, FindCoincidences.vi will increment through the timestamp array of a secondary channel, comparing newer and newer secondary timestamps to the current primary timestamp. It will do this until it finds a secondary timestamp that is greater than the coincidence window of the current primary timestamp. Since every element is sorted in time, there can be no coincidences between the current primary

timestamp and this secondary channel past this element. The algorithm keeps track of where in each of the secondary timestamp arrays the previous search ended, and for the next primary timestamp it goes backward in that array until it locates an element that is less than the coincidence window of the current primary timestamp before continuing the search. This method not only greatly speeds up the search, since for each primary timestamp only a small subset of each secondary array is searched, but also manages to do so without missing any coincidences.

Before two timestamps are compared, they are modified to compensate for the unequal propagation delay of the scintillation photons. For example, consider a muon passing through scintillators 0 and 4 in Figure 5.4. The photons in scintillator 4 have a longer path to travel than those in scintillator 1 before reaching the MPPC at the end of the light guides. The propagation times are calculated, and subtracted from the timestamps in order to synchronize them. The timestamps are therefore from this point onwards representing the muon crossing time, not the time the ROBs generated them.

When a secondary timestamp t_s is found to be within the coincidence window of a primary timestamp t_p, several things occur. Firstly, the algorithm will verify that the combination of these exact two timestamps have not previously been detected as a coincidence by checking their private coincidence-arrays. This ensures that this combination of timestamps will not accidentally be detected as a coincidence twice, which can happen if they were a part of the "leftover" array which was added to the current dataframe. If the algorithm cannot find this combination of timestamps in their private arrays, it proceeds to build a coincidence object, which contains the channel IDs, coincidence timestamp $(t_p + t_s)/2$, and time difference $t_p - t_s = \Delta t$ of the two timestamps. The coincidence timestamp is the average of the two compensated pulse timestamps, and gives the estimated muon crossing time. The time difference Δt is used to gather statistics on the absolute coincidence timing uncertainty σ_{CT}. The generated coincidence object is added to an array which will later be sent on a queue, and it is also added to the timestamps private coincidence-arrays.

Then, if their flags were not already True, the timestamps are added to a channel-specific array which contains all those timestamps on that channel that were found to be part of one or more coincidences. Their flags are then set to True to denote that they have been counted.

When an instance of `FindCoincidences.vi` has finished running, the leftover array that was added to the beginning is removed, and when the four coincidence hunts are finished their results are combined. The leftovers for the current dataframe, which will be added to the next dataframe, are extracted. The coincidence-object array, together with the arrays of timestamps counted as coincidence and those not counted as coincidence, are sent to `DisplayCoincData.vi` via a queue. To `SaveCoincData.vi` only the coincidence-object array is sent for storage.

DisplayCoincData.vi

Displaying generated CRT data, such as coincidences and total trigger rate, is handled by `DisplayCoincData.vi`. This VI will perform rudimentary analysis of the coincidence data generated by `PlanarCoincFinder.vi`.

For each coincidence that was found, it will sort into a histogram the time difference Δt between the two timestamps that it consists of. There are 32 unique histograms on display in the CRTC, one for each scintillator overlap. They can be used to spot inefficiencies or timing errors in certain ROBs and obtain information about the uncertainty of the detected muon crossing time for each scintillator overlap. A Gaussian distribution is fitted to the histogram, and its standard deviation is displayed in the histogram window. The histograms for the bottom CRT plane can be seen in Figure 5.5

`DisplayCoincData.vi` also plots information about pulse rates. On the front panel there are charts which display the current and historical rate for each channel. The total rate, coincidence rate and non-coincidence rate (noise) are plotted. See Figure 5.6 for such a chart for one channel. The time units are in multiples of the chart update period, which can be set by the user. The rate is given in events per second.

This process also performs a basic grouping of coincidences to analyse how often certain combinations of coincidences occur. By first combining the coincidences found in the top and bottom layers and time-sorting them, the algorithm then iterates through the array, collecting those coincidences that occurred within some specified time of the previous coincidence. For each group, it counts the number of coincidences that occurred in the top layer and bottom layer, and adds the result to a 2D histogram. This allows the operator to spot issues in the coincidence-finding algorithm, and visualizes the problem of *multi-coincidence ambiguity* where two or more simultaneously impacting muons can result in four or more registered coincidences, and which is discussed in more detail in Section 6.4.

An event display window can be opened from the CRTC and used to visualize the tracks of charged particles that passes through the CRT. The event display consists of a 3D model of the CRT, and lines are generated that denote the likely path that a charged particle took. The operator can control the colors of the model, window size, event rate, and the viewport orientation. The event display can be seen in Figure 5.7. `DisplayCoincData.vi` is responsible for extracting coincidence groups with a single unambiguous particle track, and supplying them to the event display.

`DisplayCoincData.vi` will also generate a plot of the polar distribution of detected muons. The CRT is capable of detecting muons up to an polar angle of 0.994 radians, or 57 degrees. It will locate and collect those coincidences which represent single muons passing through both planes, and calculate their polar angles. The number of muons detected for each angle are counted up, and displayed on the CRTC front panel. The counts are scaled to normalize the counting rates to the effective detector area seen for each angle. The scaling coefficients are listed in Appendix D.1.2.
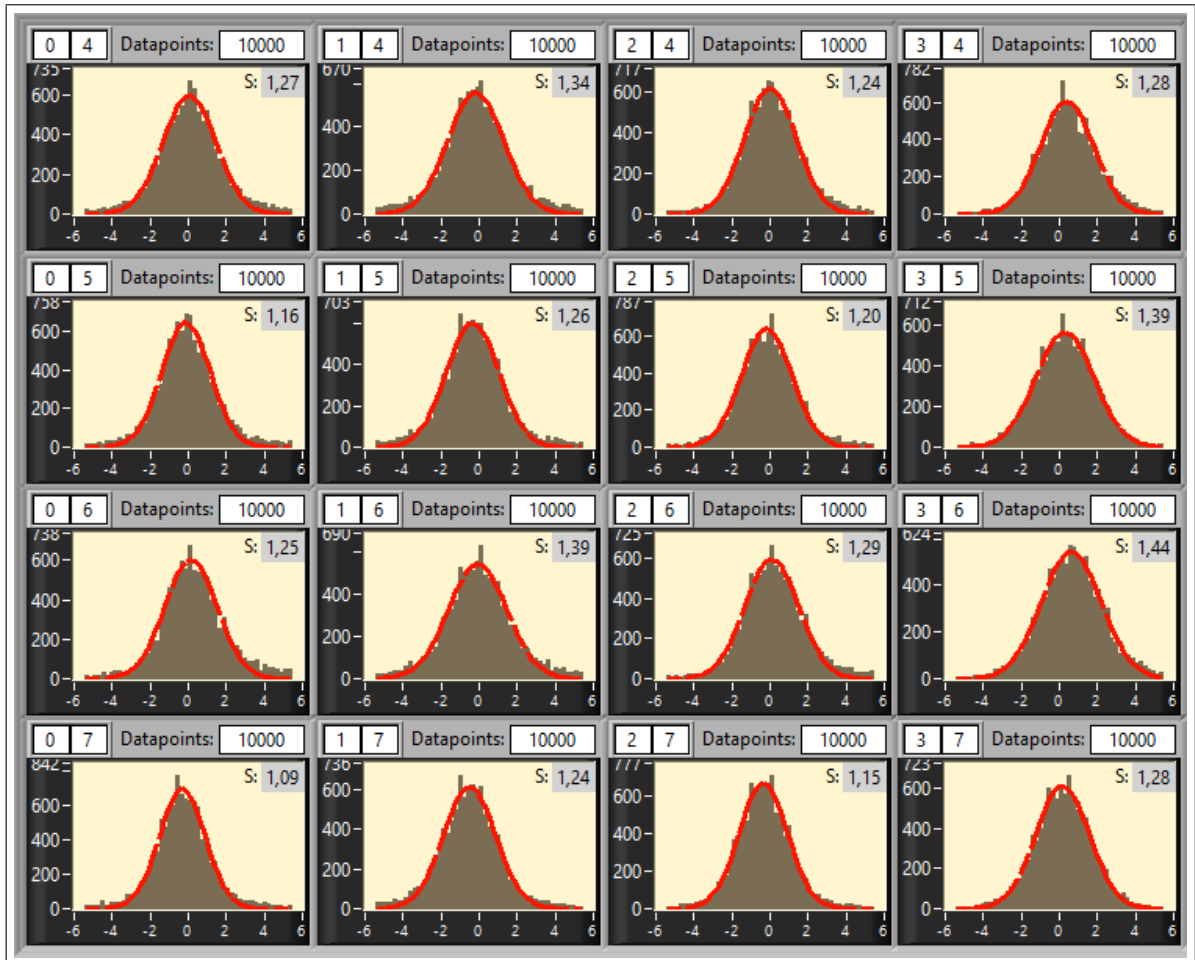
Figure 5.5: Screenshot from the CRTC, showing histograms of the time difference $t_p - t_s = \Delta t$ in nanoseconds. t_p and t_s are the arrival times of the two pulses from a primary channel and secondary channel that makes up a coincidence. The figure shows histograms for all overlapping scintillators in the bottom CRT plane, and the numbers in the top left corner of each histogram denotes which channels are overlapping. The standard deviation of the fitted Gaussian curve is displayed in the top right corner of each histogram, and denotes the 68 % confidence interval in the absolute coincidence time difference uncertainty σ_{CT} of that specific channel overlap.
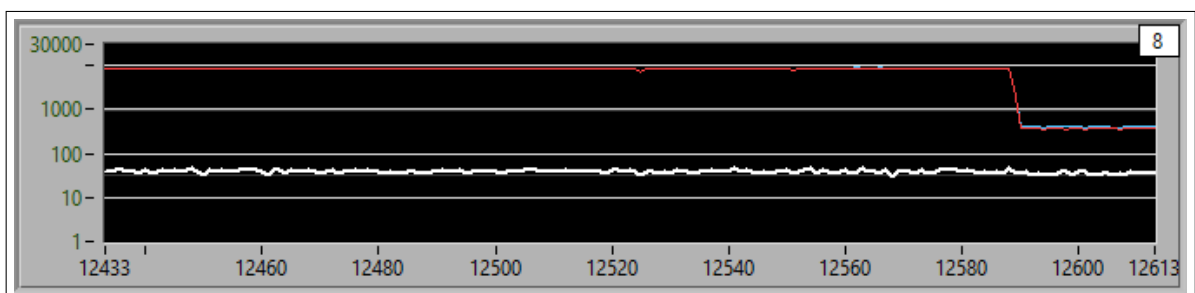
Figure 5.6: One of 16 charts on the CRTC front panel displaying information about the pulse rate for each channel. The time units are in multiples of the chart update period, which can be set by the user, and the rate is given in events per second. This chart is for channel 8, as indicated in the top right corner. The white plot denotes the rate of coincidences generated by this channel. The noise rate, i.e. the rate of pulses from this channel that was not found to be part of a coincidence, is displayed by the red plot. The total rate, in blue, can just barely be seen behind the noise rate, and is the sum of the noise rate and coincidence rate. The drop that can be seen in total rate and noise rate is due to an adjustment of the trigger threshold.
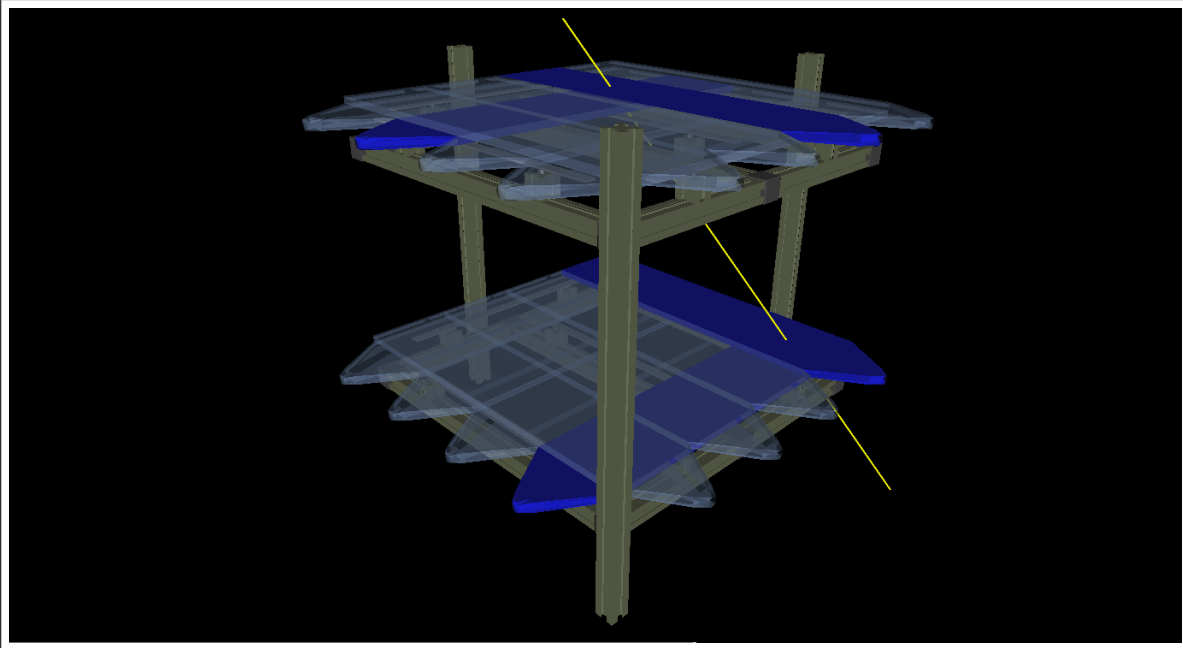
Figure 5.7: The CRTC event display, showing what was most likely a muon passing through the CRT. The solid blue scintillators are those scintillators in which coincidences were found.

SaveCoincData.vi

Saving of coincidences is handled by `SaveCoincData.vi`. This VI continuously obtains coincidence objects from the two `PlanarCoincFinder.vi` processes through a queue linking them to `SaveCoincData.vi`. When it has received more than 1000 coincidence objects, it sorts them by their timestamps, and saves the first 700 elements. The 300 latest coincidences are kept and added to the next collection of coincidences. This is done in an attempt to reduce the probability of one or more coincidences being saved out of order, which can happen since the bottom layer coincidence hunt and top layer coincidence hunt is not synchronized.

`SaveCoincData.vi` also saves information about the ROBs at regular intervals. More information about this is provided in the next subsection.

5.3.4   Data storage

The coincidences are stored together with ROB information in a HDF5 file. HDF5 is a file format designed for the storage and manipulation of large and complex datasets, supporting complex datatypes, grouping, metadata, and hierarchies. It was therefore an excellent format choice for the storage of CRT data, as the CRT generates a relatively large amount of data in the form of hundreds of coincidences per second, and also has to store different datatypes.

The data storage algorithm of the CRTC was designed so that a single HDF5 file is created per week. When a week has ended, the CRTC will close the current file and open a new one. This keeps the number and sizes of the data storage files manageable, and allows users to copy, edit, and otherwise manipulate all stored data up to the current week.

The data storage files are saved on two 2TB harddisks in a RAID 1 configuration. This configuration mirrors the data on each harddisk, functioning as a backup should one harddisk fail. A third external harddisk of smaller storage space stores the last few years of data, and backup software will automatically keep its contents up to date. This external harddisk is meant for easy sharing of CRT data. A scientist can unplug it, bring to their office, and copy its contents over to their own work station for analysis.

Data format

Within the CRT data storage file, data is stored in objects named "RUN [#]", where [#] is the value of the run counter. A single file can contain any number of RUN objects. A run is a time-continuous block of CRT data taking, in which many parameters and settings related to the CRTC are unchangeable by the operator, and the TDC cannot be altered. This ensures that the coincidences stored in a run have been found with unchanging parameters, and without breaks in the collection of hits. If the operator wishes to make any changes, they must first stop the current run. The only exception is that certain ROB parameters can be altered, as those are logged in the data file.

Not only coincidences are saved in the RUN object. With regular intervals, set by the operator, information about each ROB is saved. This includes MPPC temperatures, voltage rail values, statuses, board modes, and so forth. This historical data of the ROBs can be correlated with patterns in the coincidences to explain unexpected behaviour. For example, a sudden decreases in the coincidence rate can be explained by a ROB being disabled, either on purpose or by a fault.

The hierarchy of all CRT datastorage files is as follows, where items in bold are datasets and other items are groups:

- RUN [# -1]
- RUN [#]
    - **coincidences**
    - ROB 0
    - ROB 1
    - .
    - .
    - .
    - ROB 15
        * **latch resets**
        * **rails**
        * **responding**
        * **sipm bias**
        * **sipm temperature**
        * **timestamps**
        * **trigger threshold**

The coincidences are stored in the dataset called `coincidences`, where each element consists of the coincidence timestamp (in multiples of 32 ps intervals since 00:00 01.01.2017), and the ID of the pixel in which the coincidence occurred.

The information regarding the ROBs are stored in individual groups within the HDF5 file. Each ROB has a unique group, named from ROB0 to ROB15. Within these groups are datasets for each type of information that is logged. Every dataset has the same amount of entries, and the Nth element of each dataset was stored at the same time. The time of storage can be found in the dataset named `timestamps`, which contains the timestamp of all the measurements in the standard Windows format (number of seconds since 00:00 01.01.1905).

5.3.5   Channel characterizations

The CRTC is capable of performing automated tests in order to characterize a single channel or two overlapping channels. Characterization of a single channel provides data on the trigger rate as a function of trigger threshold, while the characterization of two channels will provide data on the coincidence rate between those two channels as a function of trigger threshold. This is useful in order to correctly set trigger thresholds.

For both characterizations, the operator can choose which channel(s) he or she wishes to characterize, the minimum and maximum trigger threshold, the trigger threshold step size, and how many seconds of data to be taken for each step. The CRTC will use the settings provided by the operator, decreasing the trigger threshold from the maximum to the minimum value and log the average trigger rate and/or coincidence rate for each trigger threshold step. The results are displayed on a graph, and stored to a file path specified by the operator. The trigger rate and coincidence rate characterizer windows can be seen in Figure 5.8.
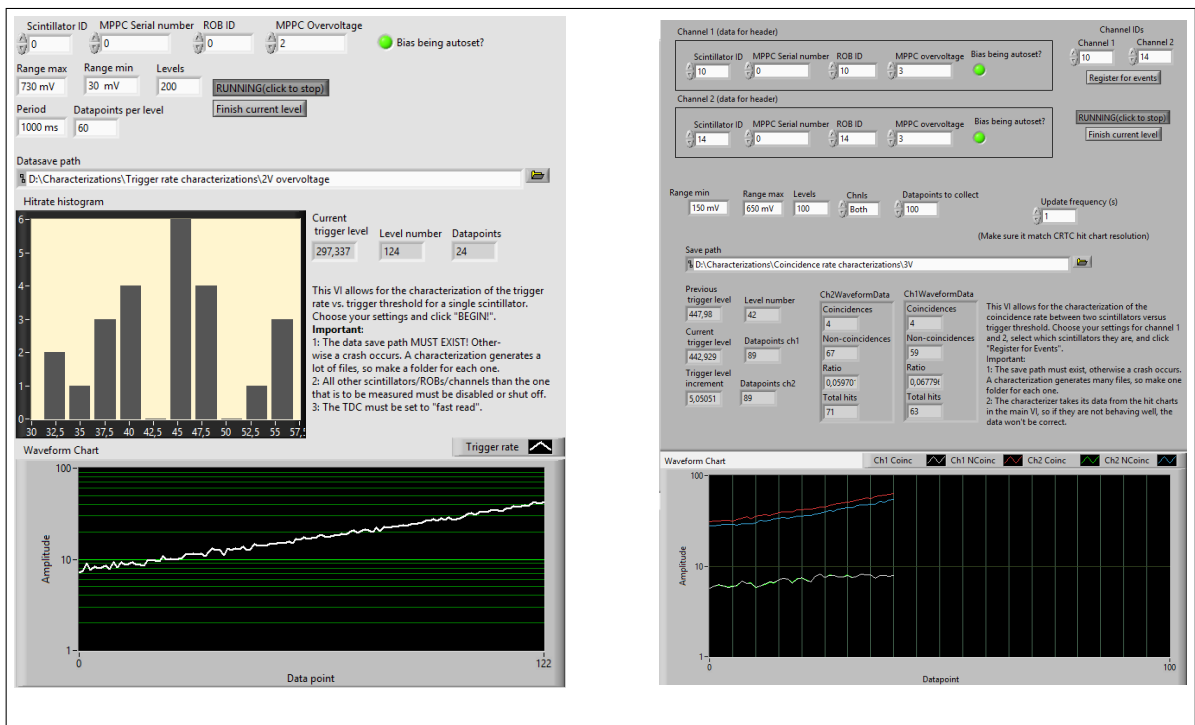
Figure 5.8: The trigger rate and coincidence rate characterization windows of the CRTC. Both are displayed in the process of a characterization.

5.3.6   Damage protection

The CRTC and ROBs have been designed with safety in mind, and therefore it is not possible for an operator to damage the CRT by choosing incorrect settings. It is possible for the operator to control individual ROB on-board voltage regulators, but the ROB MCU software will prevent attempts at enabling them in a manner that may damage components. Since the act of turning on or off the ROB is handled by a single button click in the CRTC, the user will rarely need to do this, and the capability has only been added for debugging.

It is also not possible to damage the MPPC through setting a very large bias voltage. Two layers of protection prevents this. Firstly, the MCU software itself limits the maximum overvoltage to 5 V. Secondly, the CRTC control where the user inputs the desired overvoltage is set to not accept values above 5 V.

There are three off-board laboratory power supplies that generates the semi-regulated voltages supplied to all ROBs on the CRT. If a user were to accidentally adjust these voltages too high, the ROBs will not take damage. This is thanks to reverse-biased Zener diodes on each power input, which will safely divert the current to ground if the input voltage becomes too large. The maximum input voltages of each power input is set by the LDOs that adjust them, which will have maximum input ratings. If a large input voltage is maintained over an extended period of time, the fuse on that input will eventually fail to prevent damage to the Zener diode.

6 | Results, performance, and first data analysis

In the current chapter the performance of the CRT will be presented, and simple analysis of the data it produces will be performed in order to verify its functionality, particularly its ability to detect muons. The coincidence rate of single pixels versus the MPPC overvoltage is presented and used to select a nominal value for the overvoltage. Also, a dataset of coincidences is used to estimate the muon lifetime. The problem of simultaneous two-fold coincidences and a method of estimating the muon multiplicity is discussed.

6.1   Coincidences and timing uncertainty

A coincidence is two or more signals, in our case electric pulses, that are generated within some specific time interval from each other by two or more channels set up to register an event. In order to distinguish between signals generated by noise and signals generated by a real event, we require that an event is only considered "real" if both channels registered a signal within a specified time interval. Coincidence-checking can be seen as a form of mutual vetoing between channels, and is required because there will often be some amount of noise, or "false" signals. If the noise in two channels is uncorrelated, it is unlikely that two or more noise signals will randomly occur at the same time (assuming low noise rate and narrow coincidence window). The lower the noise rate of the channels are, and the smaller the coincidence window is, the less likely it is that two noise events will be registered as a real event. In the case of the CRT, coincidence between electric pulses generated by two overlapping scintillators indicate that a muon crossed through them both, and coincidence between two pixels, one in each plane, indicate that a muon crossed both planes.

The CRT Controller application is capable of detecting two-fold coincidences (between scintillators in a plane), and by analysing these coincidences offline, four-fold coincidences (between pixels in two planes) can be found. In this section the two-fold and four-fold coincidences will be analysed, and used to show that the coincidences that are registered are not due to noise, but due to real events, i.e. charged particles such as muons.

6.1.1   Timing uncertainty

The time difference between two correlated coinciding pulses generated by some event will following a near Gaussian distribution. In the case of the CRT, the time difference originates mainly from three sources:

- The size of the overlapping area of two scintillators.
- The finite risetime of the filtered MPPC signal.
- Comparator propagation delay variation

These sources will now be discussed further.

The overlapping area of two scintillators is a square with dimensions $25\,\text{cm} \times 25\,\text{cm}$. Based on where a muon crosses this area, the photons will reach the MPPCs at the end of the scintillators at slightly different times unless the muon happens to cross the overlap in the center. The expected distribution of pulse arrival time differences is shown in Figure 6.1. This distribution is what would be expected if the only timing uncertainty came from the muon position, assuming that the photons moves in a straight line with no scattering.

This triangular distribution is convoluted with the timing uncertainty from all other sources, primarily the constant risetime of the MPPC signal. Since the trigger threshold is constant, MPPC signals of large amplitudes will cross the trigger threshold earlier than signals of small
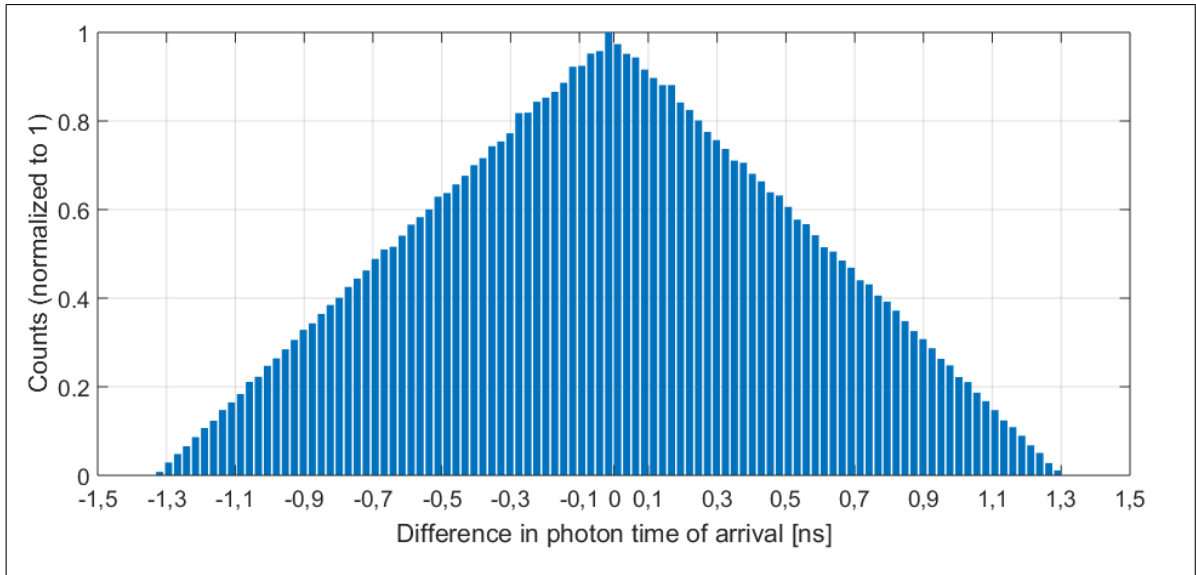
Figure 6.1: Difference in photon time of arrival at the MPPCs for photons generated by muons crossing the overlap of two scintillators. It was found by randomly distribution 100000 points in a $25\,\text{cm} \times 25\,\text{cm}$ square and calculating the time it would take for photons moving at $c/1.58 = 1.89 \times 10^8\,\text{m}\,\text{s}^{-1}$ to reach the ends of the scintillators. 1.58 is the scintillator index of refraction.

amplitudes. This means that there is an uncertainty in the timing of MPPC pulses that depends on its amplitude. However, the impact this uncertainty has on the difference in pulse arrival times can be decreased if a few requirements are met. If the two pulses generated by each MPPC is of similar amplitude, and the trigger thresholds are equal, then both pulses will take equally long to reach the trigger threshold. This is usually the case, since the ROBs of the CRT are operated with nearly equal gain, overvoltage, and trigger thresholds. Additionally, a muon crossing two scintillators will, in general, generate roughly the same amount of scintillation photons in both. That means that the amount of photons reaching both MPPCs is also of the same order, and the MPPC pulses are of similar amplitude. There will naturally be some variation, as the photon production and propagation is probabilistic.

Other sources of timing uncertainty includes comparator propagation delay jitter, photon propagation variations, and variations in gain and trigger threshold among ROBs. In the end, the final distribution of pulse arrival time differences must be found experimentally. In Figure 6.2, this distribution can be seen. The histogram was generated by collecting 10000 coincidences per scintillator overlap, or pixel. For each pixel, a histogram of the time difference between pulses in the coincidences it generated was made. The pulse times have already been modified by the CRTC to compensate for the propagation delays of scintillation photons, and the propagation delays of the coaxial cables that carry the signal. This was mentioned in Section 5.3.3, and means that the pulse times actually represent the muon crossing times. The distribution of pulse time differences for every pixel are therefore centered around 0 ns. However, the compensation is not perfect, and the location of the distribution center in reality has a variation of about $\pm 0.3\,\text{ns}$. The mean of the 32 individual histograms were therefore manually centered on 0, and they were combined to produce the histogram shown in Figure 6.2. It is nearly Gaussian, and by making this simplification, one can use the width to find the timing uncertainty in the muon crossing time.
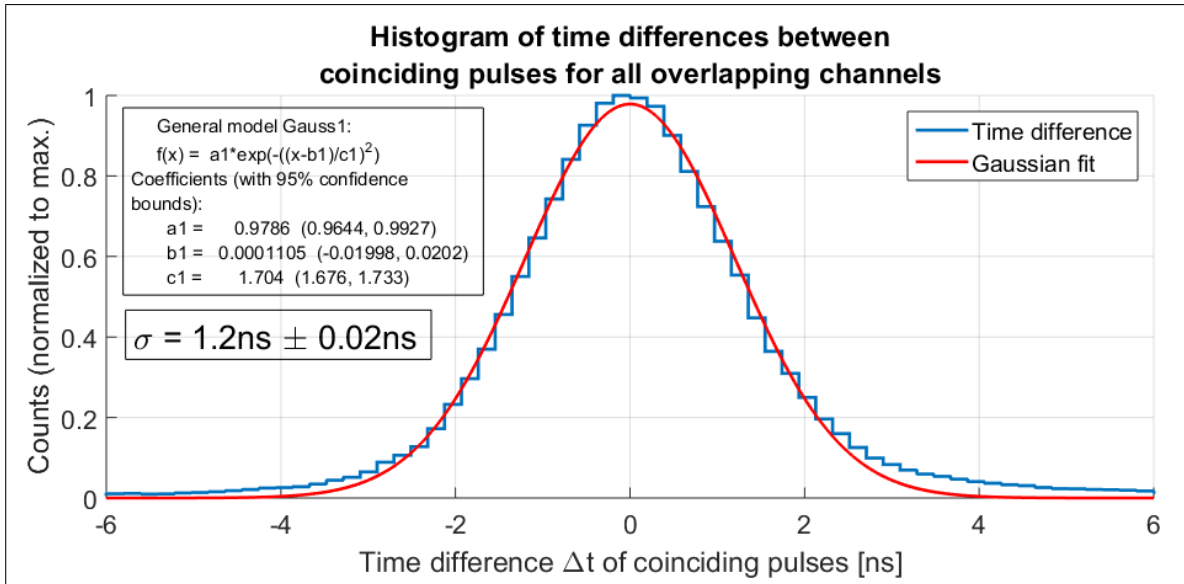
Figure 6.2: Measured time of arrival difference Δt for coinciding pulses generated by all channels of the CRT.

A successful attempt at modelling the measured time difference distribution was performed. The timing uncertainty that arise due to the muon crossing position was disentangled from all other uncertainties by convolving the distribution in Figure 6.1 with the sum of two Gaussian functions. The functions were manually adjusted until the result of the convolution fit well with the measured total timing uncertainty. The results can be seen in Figure 6.3. It appears that the timing uncertainties from all other sources can be modelled quite well as the sum of a wide and narrow Gaussian function, g_w and g_n, with the amplitude of g_w being 0.1 that of g_n, and their standard deviations being $\sigma_w = 3\,\mathrm{ns}$ and $\sigma_n = 0.9\,\mathrm{ns}$. It is not clear what exactly the sources of g_w and g_n are, but they must arise primarily from either the ROB or scintillators, as the TDC has an uncertainty of only $\pm 25\,\mathrm{ps}$. More work will be needed if the source of these uncertainties are to be determined.
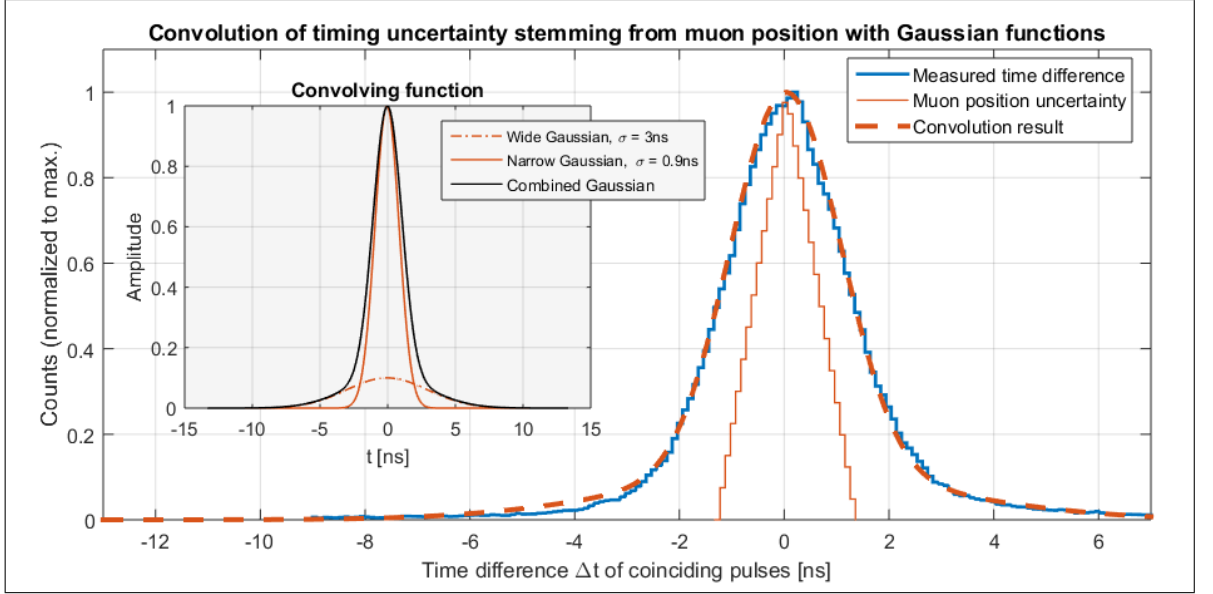
Figure 6.3: Convolution of the pulse time difference distribution due to muon position in a pixel (Figure 6.1) with two Gaussian functions. The inlaid plot shows the narrow and wide Gaussian functions, and their combination which was used in the convolution. The resulting distribution shows excellent agreement with measurements.

Muon crossing time uncertainty

To find the muon crossing time uncertainty $\sigma_{t\mu}$, we first consider how the muon crossing time is defined:

$$t_\mu = (t_1 + t_2)/2 \tag{6.1}$$

t_1 and t_2 represents the muon crossing times estimated from the two pulses that generated a coincidence. They will both have an uncertainty in their values, given by σ_{t1} and σ_{t2}. By making the assumption that these uncertainties are equal, which we can do if the two ROBs that make up the channels are operated with equal settings, we get the general pulse timing uncertainty $\sigma_t = \sigma_{t1} = \sigma_{t2}$. By error propagation, we therefore get the following expression for the uncertainty in t_μ:

$$\sigma_{t\mu} = \sqrt{\left(\frac{\partial t_\mu}{\partial t_1}\sigma_t\right)^2 + \left(\frac{\partial t_\mu}{\partial t_2}\sigma_t\right)^2}$$

$$\sigma_{t\mu} = \sqrt{\left(\frac{1}{2}\sigma_t\right)^2 + \left(\frac{1}{2}\sigma_t\right)^2}$$

$$\sigma_{t\mu} = \frac{1}{\sqrt{2}}\sigma_t \tag{6.2}$$

We must now find an expression for σ_t, the single pulse timing uncertainty. This can be done with the pulse arrival time difference Δt, which can be thought of as having a value of 0 and an uncertainty $\sigma_{\Delta t}$, given by the standard deviation of the Gaussian fit in Figure 6.2.

Since $\Delta t = t_1 - t_2$, the single pulse timing uncertainty σ_t is therefore found by error propagation:

$$\sigma_{\Delta t} = \sqrt{\sigma_t^2 + \sigma_t^2}$$

$$\sigma_{\Delta t} = \sqrt{2}\sigma_t$$

$$\sigma_t = \frac{1}{\sqrt{2}}\sigma_{\Delta t} \tag{6.3}$$

Substituting this expression for σ_t in Equation 6.2, we get the following expression for the muon crossing time uncertainty

$$\sigma_{t\mu} = \frac{1}{\sqrt{2}}\frac{1}{\sqrt{2}}\sigma_{\Delta t}$$

$$\sigma_{t\mu} = \frac{\sigma_{\Delta t}}{2};$$

The value of $\sigma_{\Delta t}$ is found from the Gaussian fit in Figure 6.2 to be 1.2 ns. We therefore have that the general muon crossing timing uncertainty for a CRT plane is

$$\sigma_{t\mu} = 0.6 \, \text{ns} \tag{6.4}$$

6.1.2   Expected false coincidence rate

It is useful to have a method of estimating how many of the coincidences that the CRT detects are from muons, and how many are from uncorrelated noise. Intuitively, this should depend on the trigger rate and coincidence window. As the rate of pulses generated by uncorrelated noise increases on two channels, so too does the probability that two noise pulses will occur within the coincidence window. This probability will naturally also increase with increasing coincidence window.

We will now find a mathematical expression for this. Normally, for two channels with rates μ_1, μ_2, and coincidence window t_{cw}, the equation for coincidence rate is given by

$$E_{coinc}(\mu_1, \mu_2, t_{cw}) = \mu_1 \left(1 - e^{-\mu_2 t_{cw}}\right) \tag{6.5}$$

However, it can not be used for the CRT. This is because it assumes that no matter how many pulses are registered on the two channels within the coincidence window, only one coincidence will be generated. This is correct for analogue coincidence schemes, but the CRTC finds coincidences digitally. For a pulse in a primary channel, a coincidence window is opened up around it. A coincidence will be created between the primary pulse and all other secondary pulses that occur within the window. These pulses can all be from the same secondary channel. An analogue scheme would register only one coincidence, namely that between the primary pulse and the first occurring secondary pulse. The digital coincidence scheme is fundamentally different, and we therefore have to find an equation for coincidence rate that fits the CRT.

For a source generating μ random uncorrelated pulses per second, we have that the probability $P_r(t, \mu)$ of finding r pulses within a time window t is given by the Poisson distribution

$$P_r(t, \mu) = \frac{(\mu t)^r}{r!} e^{-\mu t} \tag{6.6}$$

Let us now imagine that we have one channel generating μ_1 pulses per second, and a second channel generating μ_2 pulses per second. The channels are set up in such away that if two pulses occur within $t_{cw}/2$ seconds of each other, they are registered as a coincidence. By Equation 6.1.2, for every pulse generated by channel 1 there is a $P_r(t_{cw}, \mu_2)$ chance that r pulses also occurs in the second channel within $\pm t_{cw}/2$ seconds of it. In other words, the probability that a pulse in channel 1 generates a single coincidence is $P_1(t_{cw}, \mu_2)$. Likewise, the probability that a single pulse in channel 1 generates two coincidences is $P_2(t_{cw}, \mu_2)$. For "reasonable" values of μ_2 and t_{cw}, the probability of multi-coincidences drops rapidly with increasing r.

So, for each pulse in channel 1, there is a P_1 probability that a single coincidence is generated, P_2 probability that two coincidences are generated, and so on. The expectation value for the number of coincidences per channel 1 pulse is thus:

$$E_{pulse}(\mu_2, t_{cw}) = P_1(t_{cw}, \mu_2) + 2P_1(t_{cw}, \mu_2) + 3P_3(t_{cw}, \mu_2) + \ldots$$

$$E_{pulse}(\mu_2, t_{cw}) = \sum_{r=1}^{\infty} r \frac{(\mu_2 t)^r}{r!} e^{-\mu_2 t_{cw}}$$

$$E_{pulse}(\mu_2, t_{cw}) = \mu_2 t_{cw}$$

We see that the expected number of coincidences per channel 1 pulse is simply $\mu_2 t_{cw}$. Therefore, if channel 1 generates μ_1 pulses per second, the expected coincidence rate is

$$E_{coinc}(\mu_1, \mu_2, t_{cw}) = \mu_1 \mu_2 t_{cw} \tag{6.7}$$

This equation can be used to verify the expected false coincidence rate of the CRT.

6.1.3 Two-fold coincidences

Two-fold coincidences are coincidences between scintillators in a single plane, and are found online by the CRT Controller (CRTC) application's coincidence-finding algorithms (See Section 5.3.3). The CRTC generates histograms of the coincidences found between all channels. These histograms display the time difference between the pulses from two channels that were found to be within the coincidence window, and are normal distributions centred around 0 ns. A screenshot of these histograms can be seen in Figure 5.5 in Section 5.3.3.

To prove the validity of the two-fold coincidences found by the CRTC, the coincidence window was increased from the nominal $\pm 5\,\text{ns}(t_{cw} = 10\,\text{ns})$ to $\pm 250\,\text{ns}(t_{cw} = 500\,\text{ns})$. The MPPC overvoltages were set to 3 V, and 10000 coincidences were collected for each overlapping
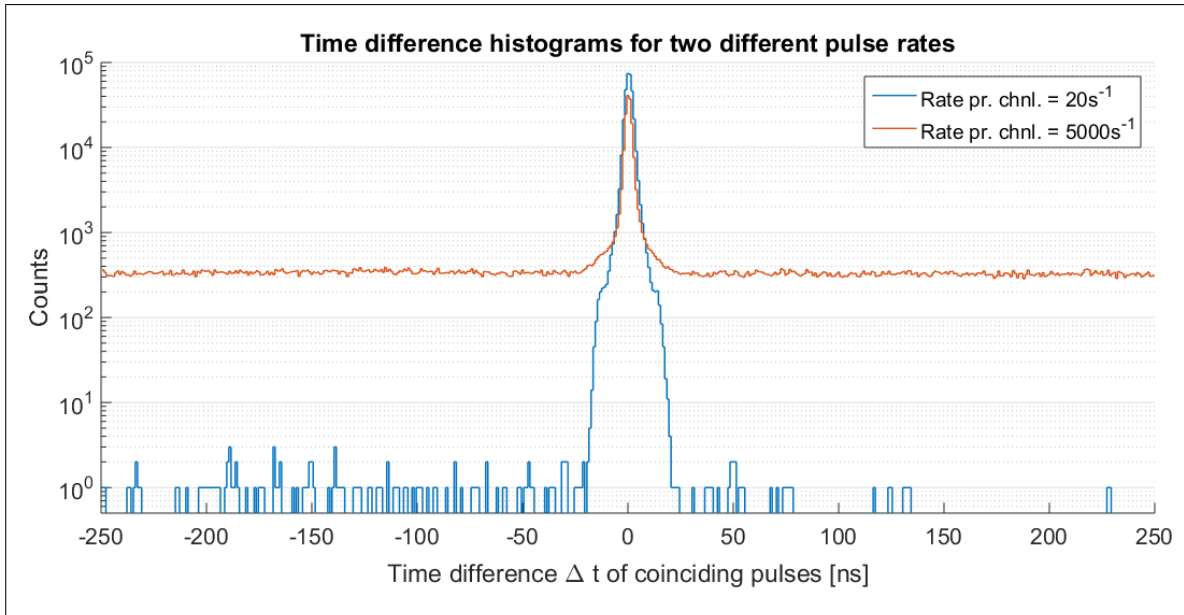
Figure 6.4: Coincidence time difference histograms for two different pulse rates, with the coincidence window $t_{cw} = 500\,\text{ns}$. Each histogram contains 320000 coincidences. When the pulse rate is high, a large fraction of coincidences are from uncorrelated noise, seen as a flat floor.

scintillators for a total of 320000 coincidences. The time difference histograms of each overlapping channel were centered on zero and superimposed. Two such CRT-wide time difference histograms were generated, one for a high trigger threshold, and one for a low trigger threshold, in order to investigate the effect of noise. For the high trigger threshold, it was set to a value well above the dark noise. The average pulse rate per channel was in this case $20\,\text{s}^{-1}$, half the expected muon rate. For the low trigger threshold case, the trigger threshold was set to a value low enough to trigger on a significant amount of dark noise, and the average channel pulse rate was $5000\,\text{s}^{-1}$, out of which an expected $40\,\text{s}^{-1}$ are generated by muons.

In the large trigger threshold case, it is extremely unlikely for a dark noise pulse to obtain an amplitude large enough to exceed the threshold. Therefore, virtually all pulses generated by the channels are from passing muons. We would therefore expect the time difference histograms to consist of a Gaussian distribution centred around 0 ns with little to no noise floor, since all the signals are correlated (the muons). In the second case, where the trigger threshold is low and thus the pulse rate is large, there is a significant amount of uncorrelated pulses generated by dark noise. We therefore expect to see the same normal distribution as for the high trigger threshold case, but with a noise floor of constant amplitude generated by the uncorrelated noise.

The time difference histograms can be seen in Figure 6.4. Exactly what was described in the previous paragraph can be seen. When the trigger threshold is high the only pulses detected are those generated by muons, which are correlated and therefore occur within a few nanoseconds of each other. When the trigger threshold is decreased in order to trigger on dark noise, a background of false coincidences is introduced. These have no correlation, and so their time difference will be random, i.e. any time difference between two pulses will be equally likely. This results in a flat floor being introduced to the histogram.

Verification of rate of false coincidences

We can use the histogram for the high pulse rate to check that the rate of false coincidences is equal to what is expected from Equation 6.1.2. Figure 6.5 shows an enhanced view of the histogram in Figure 6.4 for the high pulse rate of $5 \times 10^3 \, \text{s}^{-1}$. It has been fitted with two Gaussian distributions which represents correlated coincidences, and a constant which represents uncorrelated coincidences, i.e. false coincidences generated by dark noise. Integrating the two Gaussian distributions over the coincidence window gives the number of true coincidences $N_t = 1.51 \times 10^5 (\text{narrow}) + 1.31 \times 10^4 (\text{wide})$. Doing the same for the constant function gives the number of false coincidences $N_f = 1.66 \times 10^5$. We get that the fraction of noise coincidences to the total amount of coincidences is:

$$P_f = \frac{N_f}{N_f + N_t} = \frac{1.66 \times 10^5}{1.66 \times 10^5 + 1.51 \times 10^5 + 1.31 \times 10^4} = 0.5028 \approx 0.5 \qquad (6.8)$$

To compare this ratio with what is expected from Equation 6.1.2 we need to know the rate of true coincidences, and so the coincidence window was decreased to $t_{cw} = 40 \, \text{ns}$. The majority of true coincidences lie within this coincidence window, as was seen in Figure 6.5. A total coincidence rate of $11.1 \, \text{s}^{-1}$ for a pulse rate of $5 \times 10^3 \, \text{s}^{-1}$ was measured. Assuming Equation 6.1.2 is correct, we can find how many of these are expected to be from uncorrelated noise:

$$E_f = \mu_1 \mu_2 t_{cw} = 5 \times 10^3 \, \text{s}^{-1} \times 5 \times 10^3 \, \text{s}^{-1} \times 40 \, \text{ns} \qquad (6.9)$$

$$E_f = 1 \, \text{s}^{-1} \qquad (6.10)$$

Here the approximation $5 \times 10^3 \, \text{s}^{-1} \approx 5 \times 10^3 \, \text{s}^{-1} - E_t$ has been used, resulting in an overestimation of the false coincidence rate by at most $\sim 0.2 \, \%$. Subtracting E_f from the total coincidence rate we get an estimation of the true coincidence rate to be $E_t = 11.1 \, \text{s}^{-1} - 1 \, \text{s}^{-1} = 10.1 \, \text{s}^{-1}$.

The rate of false coincidences expected for a coincidence window of $t_{cw} = 500 \, \text{ns}$ is $12.5 \, \text{s}^{-1}$. We therefore get that the expected fraction of noise coincidences to the total amount of coincidences is:

$$P_f = \frac{E_f}{E_f + E_t} = \frac{12.5 \, \text{s}^{-1}}{12.5 \, \text{s}^{-1} + 10.1 \, \text{s}^{-1}} = 0.5530 \approx 0.55 \qquad (6.11)$$

This is no more than 5 percent off the measured value in Equation 6.1.3, indicating that it is likely correct.
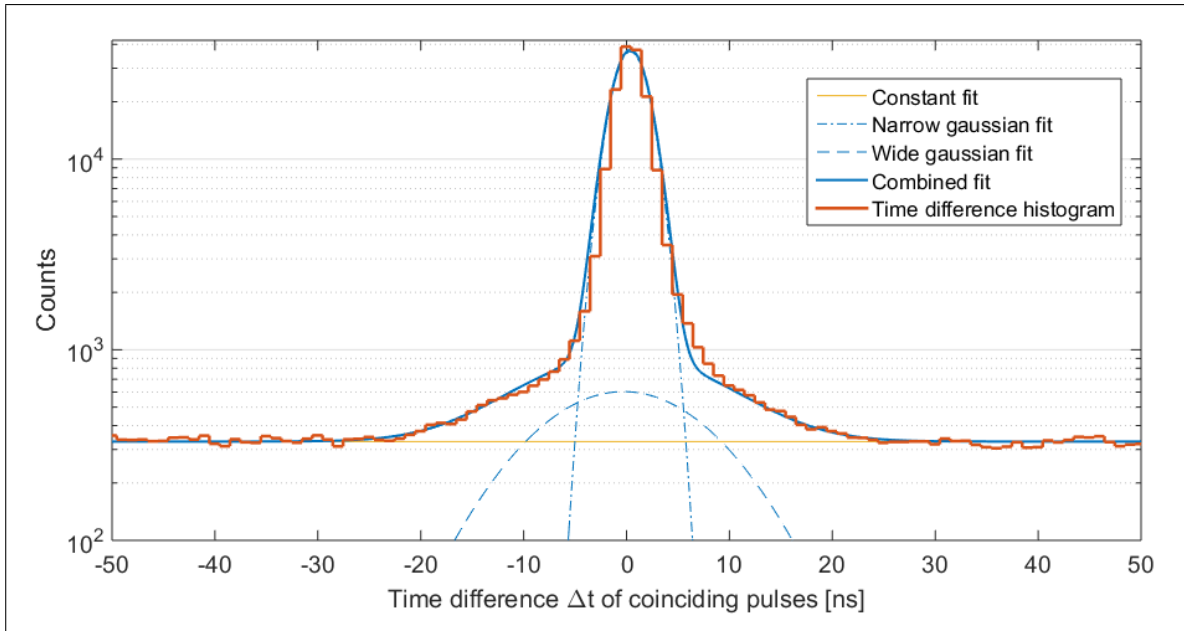
Figure 6.5: Enhanced view of the high pulse rate histogram in Figure 6.4. The histogram has been been fitted with a constant and two gaussian distributions.

6.1.4   Four-fold coincidences

Four-fold coincidences are coincidences between four scintillators, two in each plane. It can therefore also be considered as two-fold coincidence between two pixels, one in each plane. The "signals" are in this case the two-fold coincidences that have been found by online analysis as described in the previous subsection.

To prove that four-fold coincidences can be found, and that they correspond to charged particles passing through the CRT, a dataset of 28 million coincidences was analysed offline. For each two-fold coincidence (from here on called a signal) in the top layer, a signal was looked for in the bottom layer within a coincidence window of ± 250 ns. If two signals were found in the top layer and bottom layer within the coincidence window, their time differences were plotted in a histogram. We expect to see a distribution that peaks near 2.3 ns, and is nearly zero elsewhere. 2.3 ns is the time it takes a muon moving vertically at near the speed of light to cross the space between the top and bottom layer. Since vertically moving muons are the most common, that is also where we would expect the time difference histogram to peak. The results of the analysis can be seen in Figure 6.6. We see that the distribution is greatest close to 2.3 ns, and nearly zero elsewhere, as expected. This is a strong indicator that the CRT is capable of accurately detecting charged particles (muons) as they cross the two layers. It is unclear what the additional peak at 110 ns is due to, but one explanation is that some pixels on the top layer has a tendency to generate an additional coincidence $\sim$110 ns after the first. This will have to be investigated further.
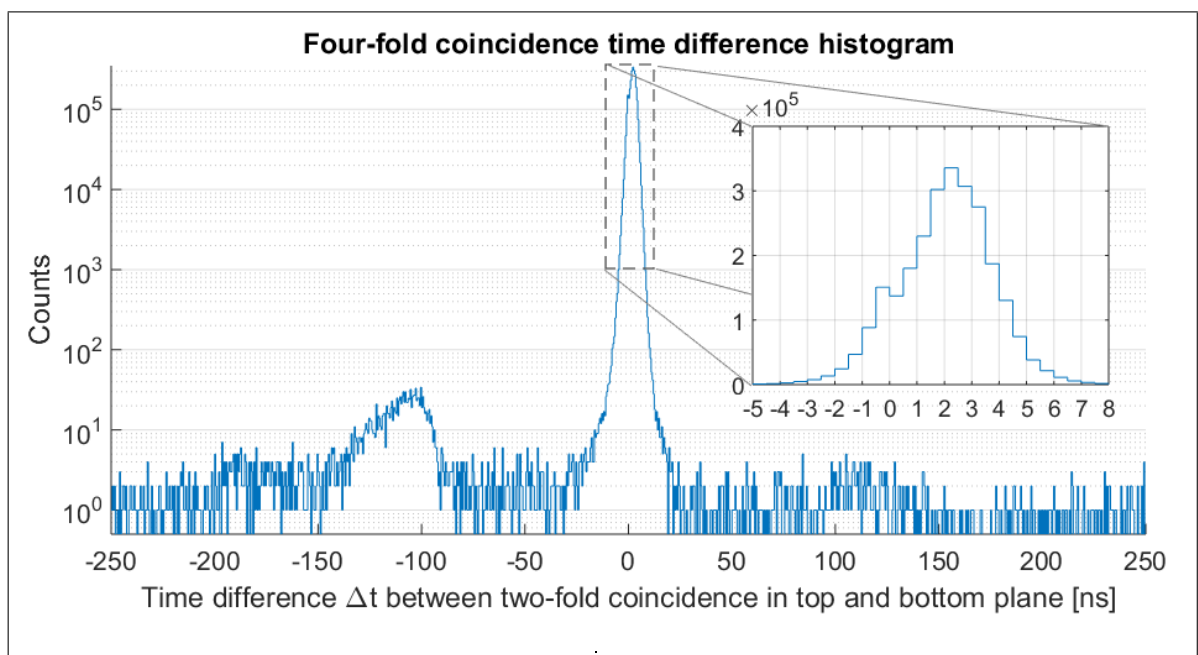
Figure 6.6: Coincidence time difference histogram for four-fold coincidences. A two-fold coincidence in the top layer and a two-fold coincidence in the bottom layer are registered as a four-fold coincidence if they occur within 250 ns of each other. The difference in their time of occurrence $\Delta t = t_{top} - t_{bot}$ is shown

6.2   Determination of optimal MPPC overvoltage

The MPPC overvoltage V_{over} is an important parameter that can be adjusted from the CRT Controller software. V_{over} is the amount of volts above the breakdown voltage of a SiPMs that it is operated at, and directly influences the gain of said device.   This parameter was discussed in Section 3.3, where the problem of increasing dark noise with increasing overvoltage was mentioned. In an attempt to find the optimal overvoltage for the CRT MPPCs, the pulse height spectrum of an MPPC at various overvoltages were recorded (Section 4.11.3) in the hopes that for low enough overvoltages the dark noise and photon-induced signals would be distinguishable in amplitude. However, it was found that determining the optimal overvoltage from these spectra alone was not possible, and a different method would have to be used. In this section the method and its results are presented.

The method is relatively simple, and is fully automated via the Coincidence Characterization sub-program built into the CRTC (see Section 5.3.5). Most ROBs in a plane are disabled, except for two connected to overlapping scintillators. An illustration of this is shown in Figure 6.7. The trigger threshold of the two ROBs are kept equal and decreased in fixed steps. The rate of coincidences and non-coincidences on these two channels are measured over extended periods of time for each trigger threshold value. Coincidences will come primarily from muons that cross their overlapping area, with a small amount coming from uncorrelated noise at high pulse rates. Non-coincidences are pulses which were not part of a coincidence, and at high trigger threshold they are mostly from muons that only cross a single scintillator. For low trigger thresholds, the non-coincidence rate is dominated by dark noise. By performing the characterizations for different overvoltages, it is possible to evaluate which overvoltage is the most optimal, i.e. which overvoltage that maximizes the CRT sensitivity while preventing dark noise from becoming a limiting factor.

6.2.1   Results

There is some variation in the trigger rate among scintillators in the CRT due to uncertainties related to the trigger thresholds and gain settings of each ROB. Therefore, not just any channels
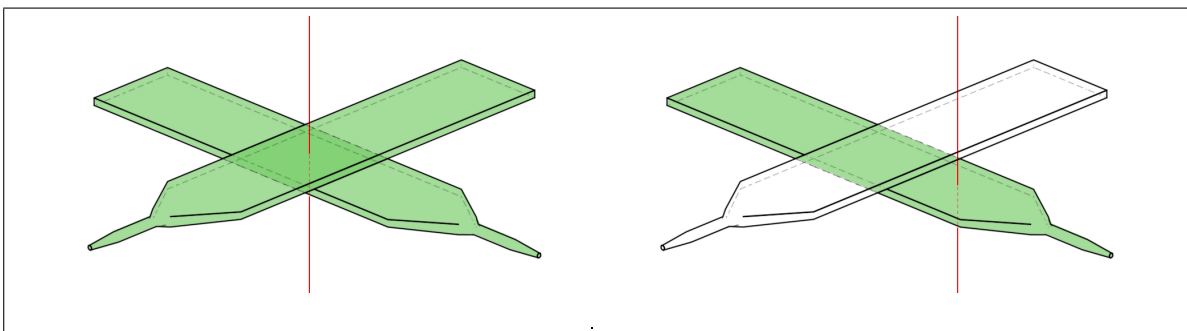
Figure 6.7: Illustration of the coincidence characterization setup of two scintillators. Disabled scintillators in the plane have not been drawn. A coincidence is registered if a muon crosses both scintillators. A muon crossing a single scintillator will be counted as a non-coincidence pulse.
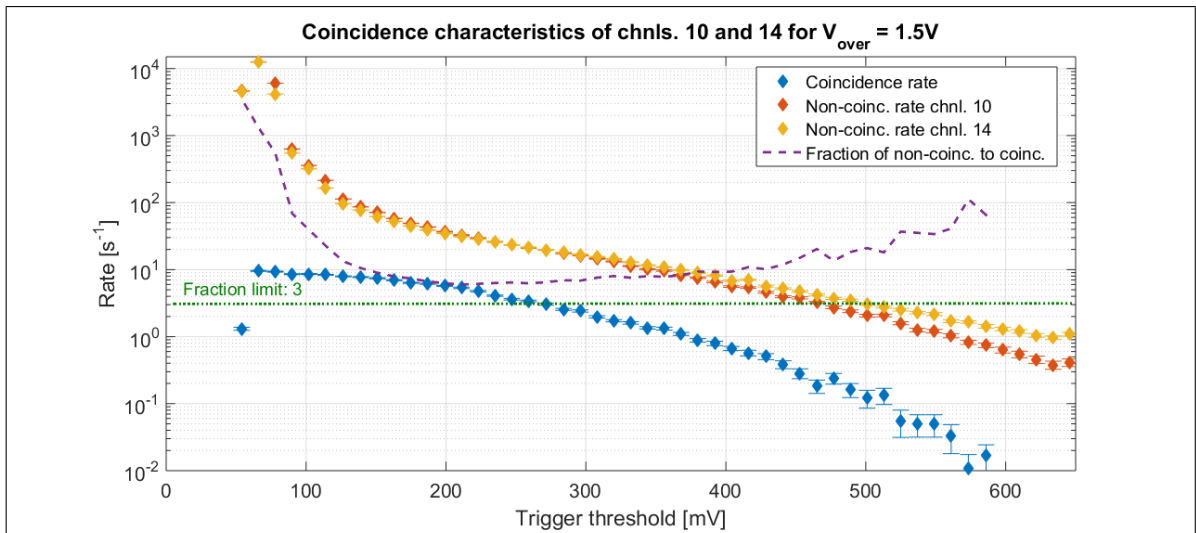
Figure 6.8: Coincidence characteristics of CRT channels 10 and 14 for an MPPC overvoltage of 1.5 V. Dark noise can be seen to start dominating for trigger threshold values less than ~150 mV.

could be chosen for the measurements in this section, and instead two that matched each other well had to be selected. The scintillators that correspond to channel 10 and 14 were chosen because they had shown low light leakage and very similar trigger rates versus trigger threshold.

The coincidence rate of these two overlapping channels were characterized for 6 different MPPC overvoltages, from 0.5 V to 3.5 V in equidistant steps of 0.5 V. For each overvoltage, the trigger threshold was kept equal for the two channels, and adjusted from a maximum value (up to 650 mV) to a lower value (varying with overvoltage) in equidistant steps. For each trigger threshold, the rate of coincidences and non-coincidences on both channels were counted.

In Figure 6.8 the result of one such characterization can be seen, with the overvoltage in this case being 1.5 V. The first thing to notice is that the fraction of coincidences to non-coincidences is not constant. One would expect from a simplistic model of the detector that the fraction stayed constant at a value of 3, since the for a single scintillator the area of non-overlap is three times larger than the area of overlap. And yet, this is not what is observed. This effect will now be explained.

The fraction increase at low trigger threshold can easily be explained by the increase in dark noise. The fraction appears to be approaching the theoretical value 3, but doesn't reach it before dark noise becomes dominant. But also for increasing trigger threshold does the fraction increase, and this is not as easily explained.

One hypothesis for this effect is as follows: When the trigger threshold is increased, the amount of photons a muon must generate in order to be detected increases. Since the majority of muons are minimum-ionizing, this increase must come from muons with large polar angles, as they will cross more scintillator material and therefore deposit more energy. Increasing the trigger threshold is therefore in effect the same as cutting off muons with small polar angles. This can be seen in Figure 6.9, where the polar distribution of muons registered by the CRT can be seen to shift upwards with increasing trigger threshold. Since the muon rate $\propto cos^2(\theta)$ [3, p.423], this can explain the decrease in overall pulse rate, but not the reason for the upwards
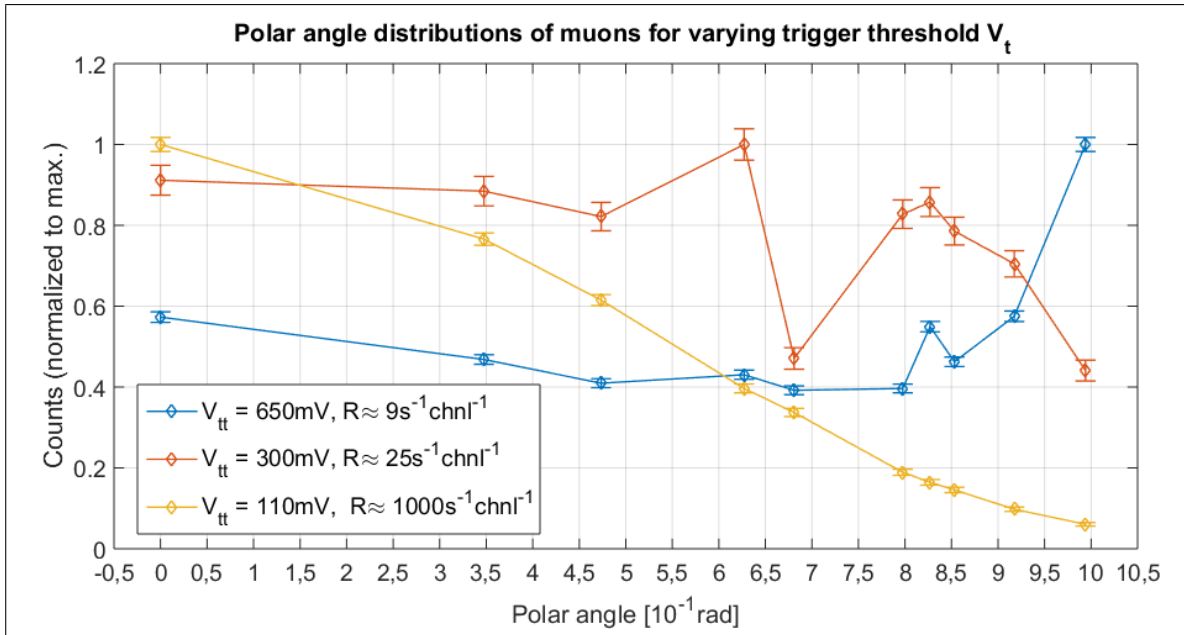
Figure 6.9: Polar angular distribution of muons detected by the CRT for different CRT-wide trigger thresholds. The counts have been compensated for the effective detector area and scaled to their maximum values. An explanation for the outliers at 0.62 and 0.68 radians for $V_{tt} = 300\,\mathrm{mV}$ has not been found. It is possible it arises due to the building that the CRT is housed in, as it will modulate the muon flux.
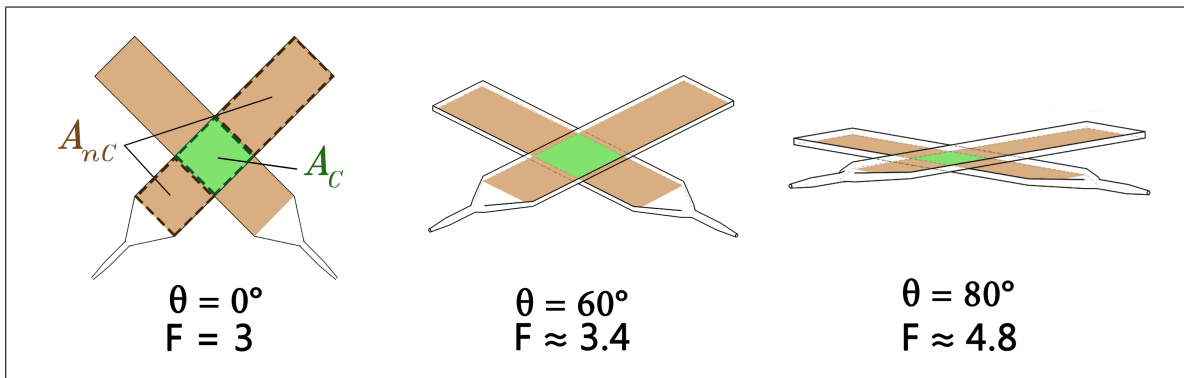
Figure 6.10: Fraction F of non-coincidence area A_{nc} to coincidence area A_c with increasing polar angle, given the requirement that the muon crosses through the entire scintillator thickness.

deviation of the ratio. To do that, we need to consider the effective scintillator area seen by a muon. This will now be discussed.

Due to the non-zero thickness of the scintillator slabs, a muon with an polar angle greater than 0 will cross through less scintillator material if it impacts the edges of the slab. Similarly, the effective scintillator area, i.e. that area seen from the muons point of view where it will cross through the entire scintillator thickness, decreases with increasing polar angle. By increasing the trigger threshold, we are effectively setting the requirement that muons cross this active area to be detected. This in turn causes the fraction of effective non-coincidence area to effective coincidence area to increase. An illustration of this is shown in Figure 6.10. Therefore, by increasing the trigger threshold, the fraction of non-coincidences to coincidences increases. This might be able to explain the increase of the fraction in Figure 6.8.

Figures of all the coincidence characterizations that were performed can be found in Appendix D.1.3. We see that by increasing the overvoltage, the noise rate becomes dominant at larger and larger trigger thresholds. This is expected not only because of the increase in gain, but also the increase in dark noise, as was found in Section 4.11.3 and is in agreement with SiPM theory. An interesting thing to note is that, for increasing overvoltage, the coincidence rate reaches larger values before dark noise takes over. This is because the PDE is a function of overvoltage. It is therefore possible to improve the sensitivity of the CRT by increasing the overvoltage, as the increase in PDE outweighs the increase in dark noise. The overvoltage can naturally only be increased up to a certain limit, which is set by the bias voltage electronics, and it is also possible that too high bias voltages can damage the MPPCs. One can also see from the coincidence characterizations that if the overvoltage is increased too much, the pulse heights that are of interest will be too large for the trigger threshold, which can be set to at most $\sim$650 mV (with some variation among channels).

This problem can however be circumvented by decreasing the gain on the ROBs themselves, which would serve to "shrink" the pulse height spectrums and allow for larger overvoltages. In Figure 6.11 the coincidence characterization for $V_{over} = 5$ V and minimum preamplifier gain can be seen. It is very similar to the coincidence characterization performed with normal preamplifier gain and $V_{over} = 2.5$ V (Appendix D.1.3), but with the added benefit of a smoother dark noise rate increase, and a greater PDE which means the coincidence rate is somewhat larger. It is also not possible to distinguish the steps that correspond to dark noise single p.e. peaks, and this can be beneficial as it will make it easier to equalize the trigger rate among all CRT channels. Another benefit of a large overvoltage is that its uncertainty is minimized ($\sim$0.64 % from Equation 4.18). Based on these results, the CRT is recommended to be operated with minimum preamplifier gains, $V_{over} = 5$ V, and $V_{tt} \approx 150$ mV for normal operation. This value for the trigger threshold will set the total trigger rate for each channel to $\sim$2000 s^{-1}. This allows for headroom in the maximum supported trigger rate for the CRTC, and maximizes the coincidence rate while keeping the probability for noise-induced coincidences low (2.4 min^{-1} pixel^{-1} for a coincidence window of 10 ns).

Decreasing the preamplifier gain does not affect the CRT muon timing uncertainty. This can be seen from Figure 6.12, which shows the histograms of coinciding pulse time arrival differences for a high and low preamplifier gain. For both measurements, the trigger threshold was adjusted to equalize the total trigger rates.
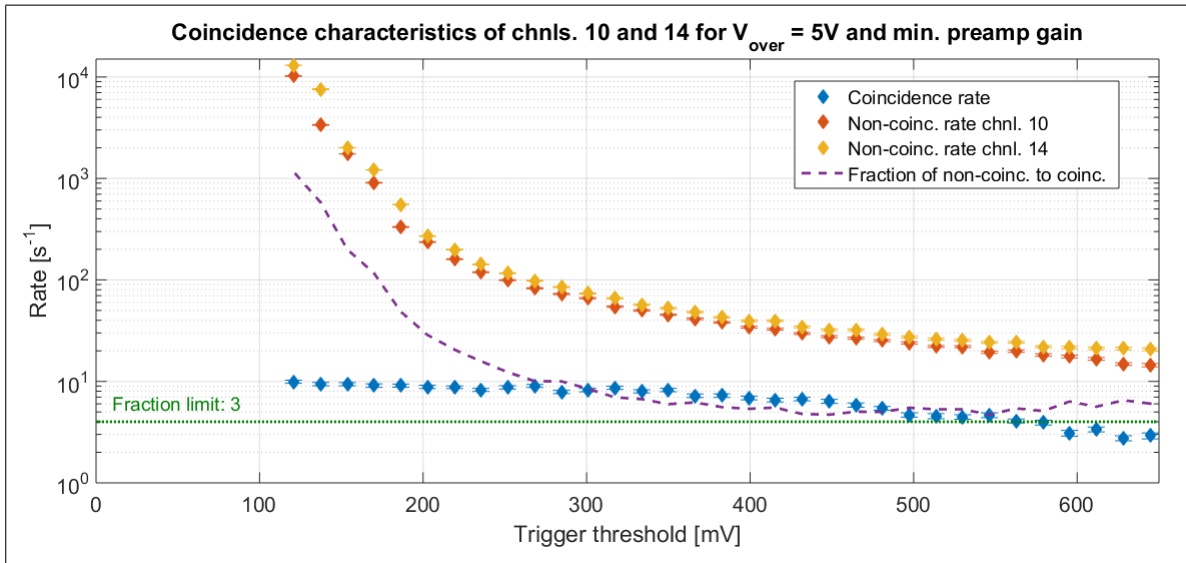
Figure 6.11: Coincidence characteristics of CRT channels 10 and 14 for an MPPC overvoltage of 5 V, and the ROB preamplifier gains set to their minimum values.
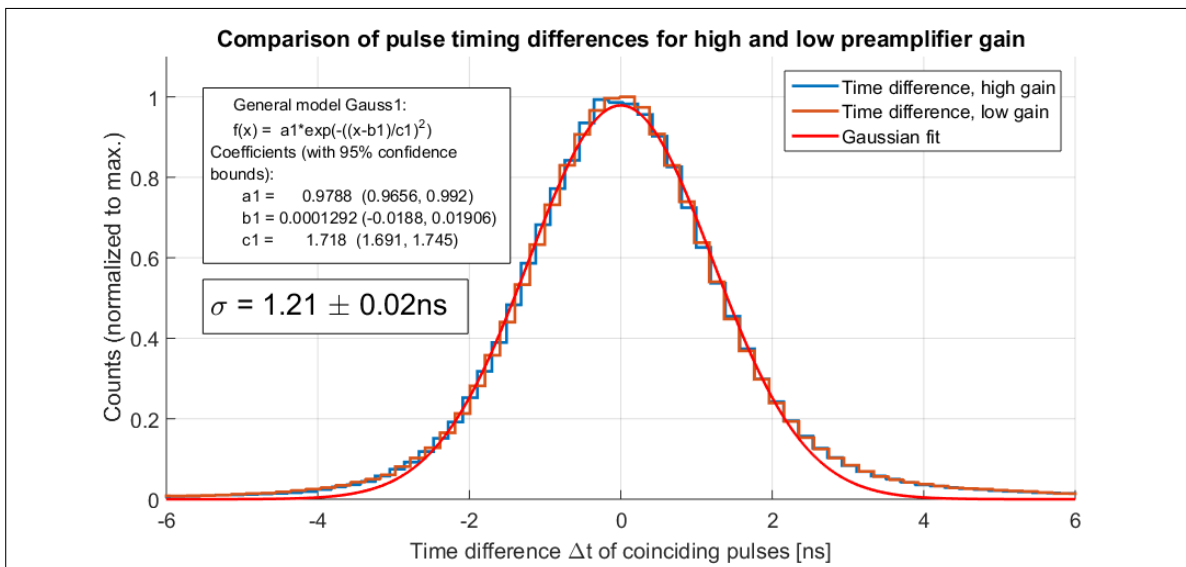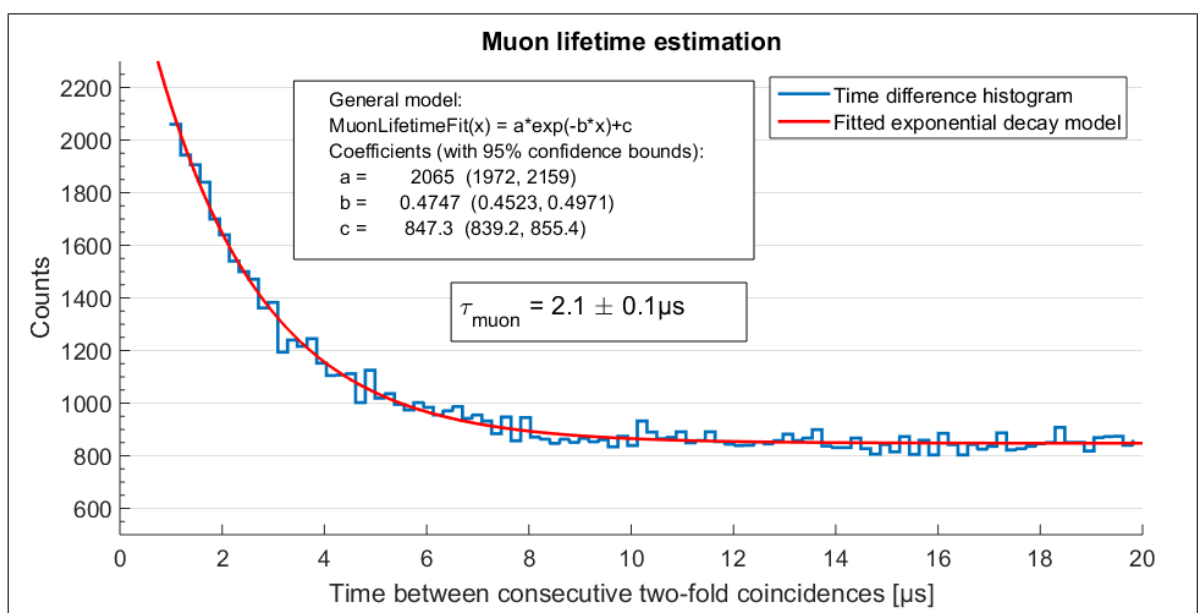

Figure 6.12: Coincidence time difference histograms for two different ROB preamplifier gain configurations.

6.3 Estimation of the muon lifetime

If a muon of low enough energy enters a scintillation slab, it will deposit all its remaining energy in the material come to a rest within it. There it will stay until it decays into three new particles: An electron, an anti electron-neutrino, and a muon-neutrino. Similarly, if an anti-muon comes to rest in the scintillator material, it will decay into the antimatter versions of the aforementioned particles. The electron/positron that is created in the decay is an energetic charged particle, and will therefore generate scintillation photons if it passes through a scintillator. In much the same way as the muon, it will therefore generate a two-fold coincidence. To find the lifetime of the muon, one need therefore simply collect statistics about how long it takes between consecutive two-fold coincidences in the CRT. It is also necessary to ignore those that occur within a few nanoseconds of each other, as those will be from single muons crossing through the top and bottom plane. Since the decay of the muon follows an exponential decay law, one expects to see a histogram that follows the same pattern.

To perform a first estimation of the muon lifetime with data from the CRT, a dataset of 28 million two-fold coincidences was analysed. The time between each two-fold coincidence was sorted in a histogram, and the lowest values were omitted in order to ignore those coming from muons passing through both planes. The result of the analysis is seen in Figure 6.13. The muon lifetime is found to be $\tau_{muon} = 2.1 \pm 0.1\,\mu s$. This value lies within the expected from literature, which lists the muon lifetime at $\tau_{muon} = 2.196\,\mu s$.

Figure 6.13: Estimation of muon lifetime by analysing the time difference between consecutive two-fold coincidences.

6.4 Multi-coincidence ambiguity

If more than one muon impacts a plane within a certain time period, we say they are simultaneous. This time period will depend on several variables. If the coincidence window is much greater than the timing uncertainties involved, the simultaneous muons simply have to impact within the coincidence window:

$$t_{max} - t_{min} \leq t_{cw} \tag{6.12}$$

Here, t_{max} is the time of impact for the last muon, t_{min} is the time of impact for the first muon, and t_{cw} is the coincidence window. However, if the coincidence window becomes on the order of the compensations and uncertainties involved, this relation is no longer true. A compensation is performed in order to convert the ROB pulse timestamps to muon impact timestamps, and is described in Section 5.3.3. This compensation may shift two timestamps together so that they lie within the coincidence window of each other, even if they did not do so before the compensation. The closest two timestamps may be shifted relatively to each other via this compensation is $\sim 5\,\text{ns}$. Additionally there is an uncertainty in the ROB muon timing on the order of $1\,\text{ns}$. The "effective" coincidence window for multiple muons therefore becomes:

$$t_{max} - t_{min} \leq t_C + 6\,\text{ns} \tag{6.13}$$

Simultaneous muons present a problem in accurately determining their numbers, as n muons are not likely to generate $N = n$ two-fold coincidences. This can be understood by imagining the simplest case where two muons are crossing a CRT plane at the same time. Three distinct cases are possible:

1. Both muons impact the same pixel
2. The muons impact two separate pixels on same axis
3. The muons impact two separate pixels on two separate axes

These cases are illustrated in Figure 6.14. Depending on which case occurs, the amount of coincidences registered by the CRTC will differ. If both muons impact the same pixel, a single coincidence will be registered. If they impact separate pixels on the same axis, two coincidences will be registered. If they impact separate pixels on separate axes, four coincidences will be registered. An illustration of this can be seen in Figure 6.14. The fact that multiple simultaneous muons can generate a range of coincidences we call *impact ambiguity*, as the number of impacting muons cannot be determined exactly from the amount of coincidences that are registered.

In this section an expression is presented that finds the probability of N coincidences being registered for n simultaneous muons, a probability which we denote $P_N(n)$. This may be used in future analysis of the CRT data in order to estimate the most probable number of muons that crossed a CRT plane when multiple coincidences are found within a short time span.
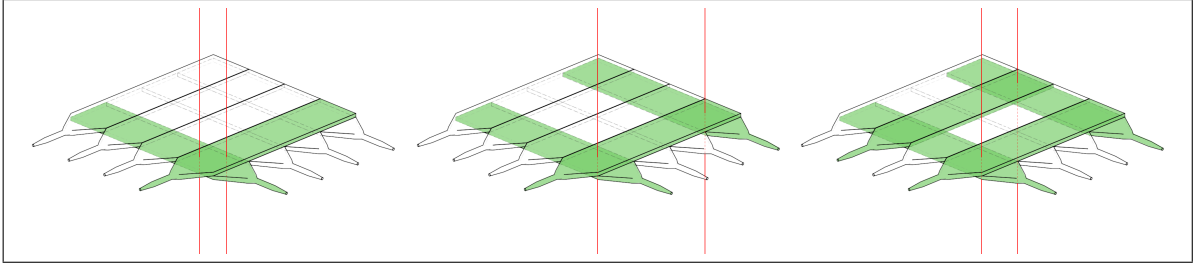
Figure 6.14: Three distinct cases for two muons impacting a CRT plane. Green scintillators have muons passing through them, and will generate a pulse. Where they overlap, a coincidence will be found. One cannot precisely infer the amount of muons that crossed the plane from the amount of coincidences that were registered.

In order to find this expression, we must first find the probability of k scintillators on the same axis being impacted by a muon if n muons are impacting the plane. In other words, if n muons are distributed homogeneously across 4 non-overlapping scintillators, what is the probability that exactly k of those scintillators will contain at least one muon? We denote this probability $P_k(n)$. To answer this question, we must find the total number of ways that n muons can be distributed across 4 scintillators so that exactly k of those contain at least one muon, called N_k. Dividing this number by the total number of ways that n muons can be distributed gives us the probability $P_k(n)$:

$$P_k(n) = \frac{N_k}{N_{tot}} = \frac{N_k(n)}{4^n} \tag{6.14}$$

$N_k(n)$ needs to describe all possible ways that n distinct muons can be distributed in k distinct non-empty scintillators, for any ordering of those k scintillators in a set of four. We have that the *Stirling numbers of second kind* gives the number of ways to partition a set of n labelled objects into k non-empty unlabelled subsets. Multiplying this number with $k!$ gives the number of ways to do it with labelled subsets. We must finally multiply the resulting number with the amount of ways that k out of four scintillators can be ordered, given by the binomial coefficient. We get:

$$N_k(n) = \binom{k}{4} \begin{Bmatrix} n \\ k \end{Bmatrix} k! \tag{6.15}$$

Substituting this in Equation 6.14 gives the final expression for $P_k(n)$

$$P_k(n) = \binom{k}{4} \begin{Bmatrix} n \\ k \end{Bmatrix} \frac{k!}{4^n} \tag{6.16}$$

Getting from this expression to one that describes the probability of N coincidences for n simultaneous muons is straight-forward. We have that if j scintillators in one axis and i scintillators in the other axis are impacted by a muon, the number of coincidences generated is their sum, $N = j \times i$. To find the probability for N coincidences being generated, we therefore must multiply the probabilities for j and i for all combinations of j and i so that $j \times i = N$. This is expressed in the following equation:

$$P_N(n) = \sum_{i=1}^{4} \sum_{j=1}^{4} \left[P_i(n) \times P_j(n) \right] \quad \text{for} \quad i \times j = N \tag{6.17}$$
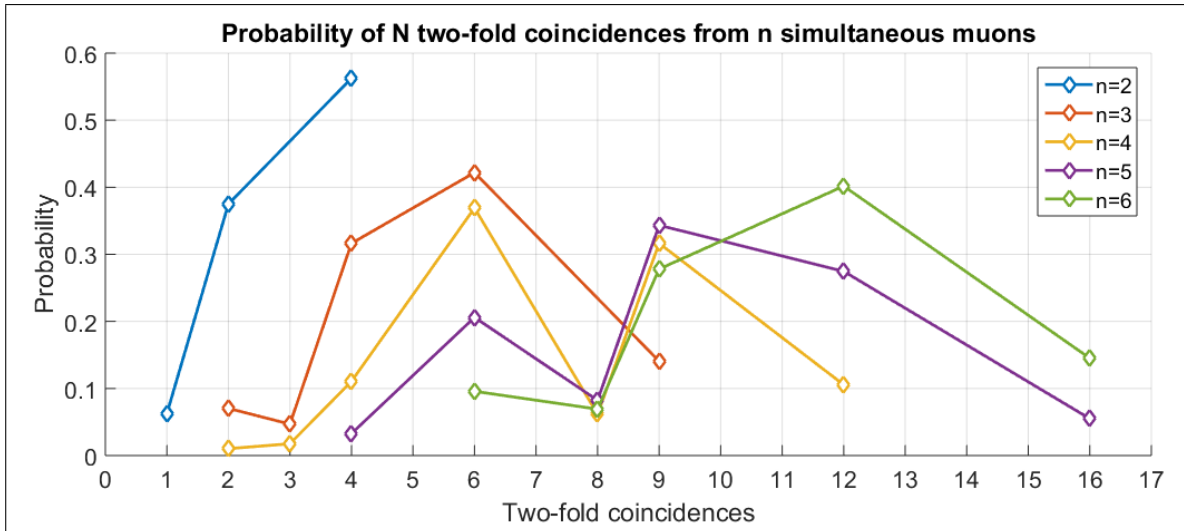
Figure 6.15: Probability of generating N two-fold coincidences for n simultaneous muons, for different values of n. Values less than $0.1\,\%$ have been omitted. Values of N that cannot be divided by two numbers less than four are not possible (5, 7, 10, 11, 13, 14, 15, >16).

And the expectation value for the amount of coincidences is

$$E_N(n) = \sum_{c=1}^{16} \left[c \times \sum_{i=1}^{4} \sum_{j=1}^{4} \left[P_i(n) \times P_j(n) \right] \right] \quad \text{for} \quad i \times j = N \tag{6.18}$$

These two equations were compared to simplistic MC simulations and found to be correct. In the simulations, N points representing muons were distributed on a 2D representation of a CRT plane. After the points had been distributed, the number of "scintillators" in each layer of the plane that contained at least one point were multiplied to give the number of coincidences. The process was repeated 10 000 times for each value of N from 2 to 21 and compared to the equation. The results from the MC simulation and the equation agreed within the uncertainties.

The probabilities of N coincidences for $1 < n < 21$ simultaneous muons were found by solving the equations, and the results are listed in Table D.9 in Appendix D.1.1. For $n > 20$ the probability of 16 coincidences approaches unity, and for $n = 1$, the only possible amount of coincidences is 1. One may use this table to estimate the most probable amount of coincidences that would be registered if n muons crossed it within the time span in Equation 6.4.

However, this cannot be used in reverse, i.e. to find the most probable amount of muons that crossed the detector if N simultaneous coincidences are detected. This is because probability distribution (PDF) for multi-muon events is not unity. As an example, consider an event where 9 simultaneous two-fold coincidences are detected in a plane. We see from Figure 6.15 that this could have come from a multi-muon event where $n > 2$. If the PDF was unity, we could easily say that the most likely value was $n = 5$, as this multiplicity has the greatest probability for creating 9 coincidences, followed by $n = 4$, $n = 6$, and finally $n = 2$. However, the PDF is not flat. In fact, the probability of multi-muon events decreases rapidly with increasing n, and therefore the most likely value for n is 2. The rate of multi-muon events was found by

the ALEPH and ALICE experiments at CERN by analysing data from cosmic rays. The results found by these experiments, in conjunction with Monte Carlo simulations of the CRT and the equation developed in this section, may in the future be used to develop a method of estimating the most likely muon multiplicity n from N simultaneous two-fold coincidences. It is, however, possible to use Equation 6.4 to say that *at least* some amount of simultaneous muons impacted the plane from the number of coincidences registered, N. This is seen from Figure 6.15. If $N = 3$, we know that the muon multiplicity must have been at least 3, as this is the lowest multiplicity that can generate three coincidences. Similarly, if $N = 8$, we know that the muon multiplicity must have been at least 4.

Equation 6.4 and the MC simulations that was used to verify it is for a "perfect" setup, assuming no propagation delays, uncertainties, variable muon distributions, ROB deadtimes, etc. It is unclear how well it holds up to reality. The real probability distributions must be found with more rigorous MC simulations of the CRT in GEANT4, but the equation may be used to give an idea of what is expected for an imaginary, perfect detector. For now, it is recommended that only those two-fold coincidences that occur "alone" (i.e. no other coincidences occurring within $t_{cw} + 6\,\text{ns}$) are used when analysing CRT data. This is because for $N = 1$, one can say with great certainty that $n = 1$. Not only is $n = 1$ the most probable muon multiplicity, but it is also that multiplicity that has the highest probability of generating a single two-fold coincidence.

6.5   ROB issues

The ROBs worked excellently in most aspects, and particularly the characteristics of the amplifier chain agreed well with the expectations found from simulations in TINA-TI. The automatic bias adjustment functioned exactly as intended, as did the fast discriminator.   The communication scheme was also successful, and thanks to it the CRT can reliably be controlled.   However, two minor issues with the ROBs became apparent after they had been in operation for a while.

The first relates to the discriminator and its latch circuit.   It became clear that it has a tendency to enter an indefinite latch if the trigger rate is high.   An indefinite latch prevents the discriminator from triggering on new MPPC pulses until it is reset.   The possibility of indefinite latches was foreseen when designing the discriminator, and a method of forcing an unlatch was implemented, as described in Section 4.5.2.   However, the rate of occurrence was much greater than what was foreseen.   An analysis of the latch rate versus trigger rate was not performed, but a quantitative guess sets it at around one latch per minute for a trigger rate of $1000\,\mathrm{s}^{-1}$. This is not a large issue, as such a latch may exist for at most 100 ms before being reset by the MCU. This might need to be investigated in the future.

The second issue, also not of great importance, relates to the on-board generation of the 15 V rail that is used for the preamplifier. It was observed that it is not stable, but varies over time. It is unclear whether it is due to temperature instabilities in the adjustment potentiometer used to adjust the rail voltage, or if the effect comes from other sources. The magnitude of the instability has been observed to be $\pm 300\,\mathrm{mV}$. Simulations show that this has a minimal effect on the preamplifier gain, on the order of $\pm 0.25\,\mathrm{dB}$. It is therefore not necessary to take actions to compensate for it.

7 | Conclusions and outlook

The contents of this thesis covers two years worth of work in an attempt to restore a cosmic muon telescope to its intended function. The project was a success, and a complete readout and control system for the CRT was designed and implemented, covering all aspects of its operation. The readout electronics that were designed to gather, filter, and trigger on the MPPC pulses worked as expected, with but a few issues. With the aid of MC simulations of the scintillators, the bunching effect of photons was foreseen and verified with measurements. These measurements, in conjunction with the simulations, were used to design the MPPC pulse shaper to trigger on the initial photon bunch, improving the CRT timing accuracy.

A communications protocol was designed to allow two-way communication between the ROBs and software running on a PC, enabling control and monitoring of the CRT. An all-in-one software solution was developed in Labview, called the CRTC, to interact with the CRT. The CRTC was a success, and allows the operator to among other things easily perform data taking runs, adjust & monitor the ROBs, monitor the trigger rate. Coincidences were found digitally by the CRTC and stored in weekly data storage files, sorted in an easily navigable hierarchical structure for later analysis.

The CRT, with its new readout system, has been shown to be flexible and easy to operate. Autonomous operation over several days were performed without problems, and only stopped due to operator intervention. Data taken over a few days were analysed and used to illustrate that the CRT is capable of accurately trigger on muons, with a timing uncertainty of $\sigma = 0.61\,\text{ns}$.

7.1 Outlook

The intended function of the CRT is educational. Therefore, several future projects with the CRT has been considered, of varying degree of difficulty. These projects may be simple enough to become integrated into undergraduate experimental courses, or require enough work for a master thesis. There are also some aspects of the CRT which can be improved. A few of the proposed project ideas will now be discussed.

7.1.1 Muon lifetime measurements

The lifetime of the muon was already found to acceptable precision with simple data analysis is Section 6.3. In order to increase the precision of the muon lifetime found by the CRT, more data is needed. One can either increase the data taking time, or attempt to increase the rate of stopping muons. One way to this is by adding an absorber between the two CRT planes. This absorber could be of steel or another high-Z material. Optimizing its location and thickness with Monte Carlo methods may be a fitting Bachelor project. When the optimal absorber specifications have been found, taking data and analysing it to find the muon lifetime can be a fitting experimental course project, especially since the data can be taken once and re-used each year.

7.1.2 Measurement of the daily muon rate

The muon rate has been shown to vary by several parameters. Atmospheric conditions [63] and elevation [64] play a major role, and the location of the moon and sun has been shown to reduce the rate and direction of muons by blocking high energy cosmic rays [65].

An interesting project would be to perform a long data run, on the order of a few months, and correlate the results with meteorological data from the same period. Corrections can be performed on the muon rate based on this data to remove the effect of atmospheric conditions known to affect the muon rate. If this proves successful, an attempt could then be made to find correlations between the moon and sun visibility and muon rate. When visible, they can block cosmic rays from reaching the atmosphere near the CRT, and the rate should decrease.

7.1.3 Large scale muon detection

There exists several large-scale distributed scintillator-based muon detectors for the detection of cosmic rays. Examples include the Stockholm Educational Air Shower Array (SEASA) [66] and WALTA [17]. These projects distribute simple scintillator-based muon detectors to high schools and middle schools, and in so doing builds a large area air shower detector. The detectors are synchronized via the Global Positioning System (GPS), and data is collected over the Internet. It is possible to build such an array of detectors in Bergen, and with a few

modifications to the CRT its data can be synchronized with the array via GPS.

The large scale arrays of SEANSA and WALTA rely on the timing differences between the individual distributed detectors to infer the direction of the cosmic rays that initiate detected air showers. If the CRT was part of such an array in Bergen, its capability of determining muon directions could be used to more accurately determine the direction of the cosmic ray that generated them.

7.1.4   Muon tomography of the Institute of Physics and Technology

In the introduction chapter, the use of muons for tomography was described. It is not far-fetched to imagine that the CRT could be used for this purpose. It is located on the third floor of the five-floor building of IFT, and attempting to perform muon tomography on the overhead structure could be an interesting project. The expected response of the isolated CRT (with no overhead structure) would first have to be found with MC methods, and the discrepancy in measurements used to estimate the amount of material in various directions. The results could be then be compared to the expected response by modelling IFT in GEANT4 and performing more MC simulations. This method could also be used to find the angular resolution of the CRT by placing heavy absorbers of various areas on the floors above it.

7.1.5   CRT fixes, improvements

Although the CRT is functioning very well, there are a few things that needs to be finished.

The CRTC is still in source code form, and should be compiled into an executable. This has not been possible due to unsolved problems that arise during the compilation, rendering the application non-executable. The CRTC functions in its current state, but it is vulnerable, as anyone can edit it. It also has a tendency to become unresponsive if top bar menu elements are accidentally clicked during operation.

A bug in the ROB MCU software prevents it from removing logged latches after some specified time, as it is intended to do. This renders the latch counters on the ROB info display in the CRTC useless, as they will eventually reach maximum and stay there. This bug should be fixed.

The room which the CRT is located in has no air flow, and therefore gets warm. The longevity of the electrical components of the CRT and operator PC might be increased by improving the room temperature by exchanging the air. Taking it one step further, temperature stabilization would decrease the possible error source coming from varying ambient temperature when analysing muon rates.

Bibliography

[1] Bicron. *BC-400/BC-404/BC-408/BC-412/BC-416 Premium Plastic Scintillators*. Datasheet. Mar. 4, 1998. URL: `http://www.phys.ufl.edu/courses/phy4803L/group_I/muon/bicron_bc400-416.pdf` (visited on 05/31/2017).

[2] Ø. Sætre. "Construction, Testing and First Data Analysis with the Cosmic Ray Telescope". MA thesis. University of Bergen, Nov. 15, 2007.

[3] C. Patrignani et al. (Particle Data Group). "Review of Particle Physics". In: *Chinese Physics C* 40.10 (2016).

[4] R. A. Millikan I.S. Bowen and H. Victor Neher. "New Evidence as to the Nature of the Incoming Cosmic Rays, Their Absorbability in the Atmosphere, and the Secondary Character of the Penetrating Rays Found in Such Abundance at Sea Level and Below". In: *Physical Review*. Second 53.3 (Feb. 1, 1938).

[5] G.F.S. "New results on cosmic rays". In: *Journal of the Franklin Institute* 205 (5 May 1928), pp. 742–745.

[6] L. J. Watson et al. "A Bayesian analysis of the 27 highest energy cosmic rays detected by the Pierre Auger Observatory". In: *Monthly Notices of the Royal Astronomical Society* 418 (2011), pp. 206–213.

[7] M. Spurio S. Cecchini. "Atmospheric muons: experimental aspects". In: *Geoscientific Instrumentation Methods and Data Systems* 1 (2 Nov. 21, 2012), pp. 185–196.

[8] L. W. Alvarez et al. "Search for Hidden Chambers in the Pyramids. The structure of the Second Pyramid of Giza is determined by cosmic-ray absorption." In: *Science* 167 (3919 Feb. 6, 1970), pp. 832–839.

[9] F. Ambrosino et al. "The MU-RAY project: detector technology and first data from Mt. Vesuvius". In: *13th TOPICAL SEMINAR ON INNOVATIVE PARTICLE AND RADIATION DETECTORS*. 2013.

[10] W. C. Priedhorsky et al. "Detection of high-Z objects using multiple scattering of cosmic ray muons". In: *Review of Scientific Instruments* 74.10 (Oct. 2003), pp. 4294–4297.

[11] Decision Sciences Corporation. *How MMPDS Works*. URL: `https://www.decisionsciences.com/ourproduct/technology/`.

[12] K. Gnanvo et al. "Imaging of high-Z material for nuclear contraband detection with a minimal prototype of a Muon Tomography station based on GEM detectors". In: *Nuclear Instruments and Methods in Physics Research A* 652 (1 Oct. 1, 2011), pp. 16–20.

[13] Tokyo Electric Power Company Holdings, Inc. *Locating Fuel Debris inside the Unit 2 Reactor Using a Muon Measurement Technology at Fukushima Daiichi Nuclear Power Station.* Presentation. July 28, 2016. URL: `http://www.tepco.co.jp/en/nu/fukushima-np/handouts/2016/images/handouts_160728_01-e.pdf` (visited on 05/04/2017).

[14] H. Miyadera et al. "Imaging Fukushima Daiichi reactors with muons". In: *AIP Advances* 3.5 (2013), p. 052133.

[15] H. Sandaker. "SCT and TRT Performance from Cosmic Ray Runs". In: *Topical Workshop on Electronics for Particle Physics.* Sept. 2007, pp. 168–172.

[16] L. Wolfley M. Collier. *Assembly Manual for the Berkeley Lab Cosmic Ray Detector.* 2.1. Mar. 2015. URL: `http://cosmic.lbl.gov/documentation/CosmicDetector2-1.pdf`.

[17] H. G. Berns et al. "The Washington Large Area Time Coincidence Array". In: *28th International Cosmic Ray Conference.* 2003, pp. 1065–1068.

[18] S. Agostinelli et al. "Geant4—a simulation toolkit". In: *Nuclear Instruments and Methods in Physics Research A* 506 (2003), pp. 250–303.

[19] R. Eckhardt. "Stan Ulam, John von Neumann, and the Monte Carlo Method". In: *Los Alamos Science Special Issue* 15 (1987), pp. 131–143.

[20] N. Metropolis S. Ulam. "The Monte Carlo Method". In: *Journal of the American Statistical Association* 44.247 (Sept. 1949), pp. 335–341.

[21] A. Ferrari et al. *FLUKA: A Multi-Particle Transport Code.* Tech. rep. Stanford Linear Accelerator Center, Oct. 2005.

[22] M. Bodmer et al. "Measurement of Optical Attenuation in Acrylic Light Guides for a Dark Matter Detector". In: *Journal of Instrumentation* 9 (Feb. 2014).

[23] Hamamatsu Photonics K.K. *S13360 series. MPPCs for precision measurement.* Datasheet. Aug. 2016. URL: `https://www.hamamatsu.com/resources/pdf/ssd/s13360_series_kapd1052e.pdf`.

[24] D. Renker. "Geiger-mode avalanche photodiodes, history, properties and problems". In: *Nuclear Instruments and Methods in Physics Research A* 567 (1 Nov. 1, 2006), pp. 48–56.

[25] KETEK GmbH. *Cross section of micro cells processed in Standard Technology.* Photograph. URL: `http://www.ketek.net/ketek/ketek-33.jpg` (visited on 12/01/2017).

[26] P. Buzhan et al. "Large area silicon photomultipliers: Performance and applications". In: *Nuclear Instruments and Methods in Physics Research A* 567 (1 Nov. 1, 2006), pp. 78–82.

[27] Hamamatsu Photonics K.K. *S13360-1375PE.* Photograph. URL: `https://www.hamamatsu.com/eu/en/product/category/3100/4004/4113/S13360-1375PE/index.html` (visited on 12/01/2017).

[28]   Hamamatsu Photonics K.K. *S13360-6050CS*. Photograph. URL: `https://www.hamamatsu.com/eu/en/product/category/3100/4004/4113/S13360-6050CS/index.html` (visited on 12/01/2017).

[29]   Hamamatsu Photonics K.K. *S13360-6050VE*. Photograph. URL: `https://www.hamamatsu.com/eu/en/product/category/3100/4004/4113/S13360-6050VE/index.html` (visited on 12/01/2017).

[30]   Perkin Elmer Incorporated. *Avalanche photodiode. A User Guide.* 2010. URL: `http://www.perkinelmer.com/CMSResources/Images/44-6538APP_AvalanchePhotodiodesUsersGuide.pdf` (visited on 05/31/2017).

[31]   Hamamatsu Photonics K.K. *Cross section of an avalanche photodiode*. Illustration. URL: `http://www.hamamatsu.com/sp/hc/osh/osh_0101_000_figure02.gif` (visited on 01/12/2017).

[32]   A.N. Otte. "The Silicon Photomultiplier - A new device for High Energy Physics, Astroparticle Physics, Industrial and Medical Applications". In: *International Symposium on Detector Development for Particle, Astroparticle and Synchrotron Radiation Experiments.* 2006.

[33]   F. Corsi et al. "Modelling a silicon photomultiplier (SiPM) as a signal source for optimum front-end design". In: *Nuclear Instruments and Methods in Physics Research A* 572 (1 Mar. 1, 2007), pp. 416–418.

[34]   P. K. Lightfoot et al. "Characterisation of a silicon photomultiplier device for applications in liquid argon based neutrino physics and dark matter searches". In: *Journal of Instrumentation* 3.10 (Oct. 6, 2008), P10001.

[35]   A. L. Lacaita et al. "On the Bremsstrahlung Origin of Hot-Carrier-Induced Photons in Silicon Devices". In: *IEEE Transactions on Electron Devices* 40.3 (Mar. 1993), pp. 577–582.

[36]   A. Goetzberger et al. "Avalanche Effects in Silicon p—n Junctions. II. Structurally Perfect Junctions". In: *Journal of Applied Physics* 34.6 (1963), pp. 1591–1600.

[37]   P. Buzhan et al. "Timing by silicon photomultiplier: A possible application for TOF measurements". In: *Nuclear Instruments and Methods in Physics Research A* 567 (1 Nov. 1, 2006), pp. 353–355.

[38]   R. Vinke et al. "Optimizing the timing resolution of SiPM sensors for use in TOF-PET detectors". In: *Nuclear Instruments and Methods in Physics Research A* 610 (2009), pp. 188–191.

[39]   G. Collazuol et al. "Single photon timing resolution and detection efficiency of the IRST silicon photo-multipliers". In: *Nuclear Instruments and Methods in Physics Research A* 581 (1-2 Oct. 21, 2007), pp. 461–464.

[40]   A. N. Otte et al. "A test of silicon photomultipliers as readout for PET". In: *Nuclear Instruments and Methods in Physics Research A* 545 (3 June 21, 2005), pp. 705–715.

[41] A. Ronzhin et al. "A SiPM-based TOF-PET detector with high speed digital DRS4 read-out". In: *Nuclear Instruments and Methods in Physics Research A* 703 (2013), pp. 109–113.

[42] D. Marano et al. "A New Accurate Analytical Expression for the SiPM Transient Response to Single Photons". In: *21st IEEE International Conference on Electronics, Circuits and Systems (ICECS)*. 2014.

[43] S. Seifert et al. "Simulation of Silicon Photomultiplier Signals". In: *IEEE Transactions On Nuclear Science* 56.6 (Dec. 2009), pp. 3726–3733.

[44] D. Marano et al. "Silicon Photomultipliers Electrical Model Extensive Analytical Analysis". In: *IEEE Transactions on Nuclear Science* 61.1 (Feb. 1, 2017).

[45] D. Marano et al. "A New Simple and Effective Procedure for SiPM Electrical Parameter Extraction". In: *IEEE Sensors Journal* 16.10 (Feb. 18, 2016).

[46] D. Marano et al. "Improved SPICE electrical model of silicon photomultipliers". In: *Nuclear Instruments and Methods in Physics Research A* 726 (Oct. 21, 2013), pp. 1–7.

[47] Hamamatsu Photonics K.K. *Photomultiplier Tubes. Basics and Applications*. 3rd ed. Hamamatsu. 2006. Chap. 5, pp. 83–123. URL: https://www.hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE.pdf.

[48] Texas Instruments Incorporated. *Current Feedback. Amplifier Analysis and Compensation*. Application Report. SLOA021A. Mar. 2001.

[49] J. Huizenga et al. "A fast preamplifier concept for SiPM-based time-of-flight PET detectors". In: *Nuclear Instruments and Methods in Physics Research A* 695 (Dec. 11, 2012), pp. 379–384.

[50] D. H. Johnson. *Origins of the Equivalent Circuit Concept*. Rice University, Aug. 11, 2001. URL: http://tcts.fpms.ac.be/cours/1005-01/equiv.pdf (visited on 05/15/2017).

[51] NXP Semiconductors N.V. *BFU550A NPN wideband silicon RF transistor*. Datasheet. Jan. 13, 2014. URL: http://www.nxp.com/documents/data_sheet/BFU550A.pdf.

[52] W. Marshall Leach, Jr. *The BJT*. Georgia Institute of Technology, School of Electrical and Computer Engineering, 2010. URL: http://leachlegacy.ece.gatech.edu/ece3050/notes/bjt/BJTBasicsSu10.pdf (visited on 05/31/2017).

[53] Texas Instruments Incorporated. *OPA847 Wideband, Ultra-Low Noise, Voltage-Feedback Operational Amplifier with Shutdown*. Datasheet. Dec. 2008. URL: http://www.ti.com/lit/ds/symlink/opa847.pdf (visited on 05/27/2017).

[54] ON Semiconductor. *AND8020/D, Rev. 6. Termination of ECL Logic Devices*. Application Note. ON Semiconductor, Apr. 2017. URL: https://www.onsemi.com/pub/Collateral/AND8020-D.PDF.

[55] Acromag Incorporated. *CRITERIA FOR TEMPERATURE SENSOR SELECTION OF T/C AND RTD SENSOR TYPES. The Basics of Temperature Measurement Using RTDs Part 2 of 3*. https://www.acromag.com/sites/default/files/RTD_Temperature_Measurement_917A.pdf, Accessed on 20.04.2017. May 2011. URL: https://www.acromag.com/sites/default/files/RTD_Temperature_Measurement_917A.pdf.

[56] S. Ekelof. "The Genesis of the Wheatstone Bridge". In: *Engineering Science and Education Journal*. 1st ser. 10 (Feb. 2001), pp. 37–40.

[57] Innovative Sensor Technology AG. *ND1K0.520.2FW.B.007*. Datasheet. URL: `http://www.farnell.com/datasheets/1643986.pdf?_ga=1.110993157.1678754166.1488889661` (visited on 04/20/2017).

[58] Maxim Intregrated. *MAX6070/MAX6071*. Datasheet. URL: `https://datasheets.maximintegrated.com/en/ds/MAX6070-MAX6071.pdf` (visited on 04/20/2017).

[59] Analog Devices Incorporated. *AD8293G80/AD8293G160*. Datasheet. URL: `http://www.mouser.com/ds/2/609/AD8293G80_160-877846.pdf` (visited on 04/20/2017).

[60] N. Gammon. *RS485 communications*. Nov. 14, 2011. URL: `https://www.gammon.com.au/forum/?id=11428`.

[61] J. Schumacher. "Characterization Studies of Silicon Photomultipliers: Noise and Relative Photon Detection Efficiency". Bachelor's thesis. Rheinisch-Westfalischen Technischen Hochschule Aachen, Feb. 2011.

[62] Cronologic GmbH & Co. KG. *HPTDC8-PCI Time to Digital Converter with 25ps Resolution*. Rev 3.3.0. Oct. 2011. URL: `http://www.cronologic.de/products/time_measurement/tdc/hptdc/HPTDC_description.pdf` (visited on 05/31/2017).

[63] K.M.Cheng H.M.Mok. "The Day-night Variation of Cosmic Rays Intensity at Sea Level Under the Influence of Meteorological Fronts and Throughs". In: *IRPA-10 Proceedings of the 10th international congress of the International Radiation Protection Association on harmonization of radiation, human life and the ecosystem*. Radiation Healt Unit, 3/F., Saiwanho Health Centre, HKSAR Government, 28 Tai Hong st., Saiwanho, Hong Kong, China. 2000.

[64] N. Ramesh et al. "Flux Variation of Cosmic Muons". In: *Journal of the Arkansas Academy of Science* 65 (2011).

[65] E.W. Grashorn. "MINOS Observations of Shadowing in the Muon Flux Underground". In: *30TH INTERNATIONAL COSMIC RAY CONFERENCE*. 2007.

[66] M. Pearce. *SEASA: The Stockholm Educational Air Shower Array*. Design Report. Kungl tekniska högskolan, Oct. 14, 2002. URL: `http://gluon.particle.kth.se/SEASA/project.pdf` (visited on 05/31/2017).

[67] D. Chattopadhyay. In: *Electronics (fundamentals And Applications)*. 7th ed. New Age International, 2006. Chap. 7, p. 129.

[68] P. R. Gray et al. *Analysis and Design of Analog Integrated Circuits*. 5th ed. John Wiley & Sons Inc, 2009. Chap. 1.

[69] Jr. Dr. W. Marshall Leach. *The Common-Base Amplifier*. Lecture notes. URL: `http://leachlegacy.ece.gatech.edu/ece3050/notes/bjt/cbamp.pdf`.

[70] K. F. Knott R. T. Unwin. "Comparison of methods used for determining base spreading resistance". In: *IEE Proceedings I - Solid-State and Electron Devices* 127.2 (Apr. 1980), pp. 53–61.

[71]   G. F. Strouse. *Standard Platinum Resistance Thermometer Calibrations from the Ar TP to the Ag FP*. NIST Special Publication. National Institute of Standards and Technology, Jan. 2008. URL: `https://www.nist.gov/sites/default/files/documents/calibrations/sp250-81.pdf`.

[72]   Analog Devices. *Low Cost, Zero-Drift In-Amp with Filter and Fixed Gain AD8293G80/AD8293G160*. Datasheet. 2008. URL: `http://www.analog.com/media/en/technical-documentation/data-sheets/AD8293G80_160.pdf` (visited on 06/01/2017).

[73]   Micrel. *MIC5205*. Datasheet. Feb. 2006. URL: `http://www.farnell.com/datasheets/29605.pdf`.

[74]   Microchip Technology Incorporated. *MCP6001/2/4*. Datasheet. 2004. URL: `http://www.mouser.com/ds/2/268/21733e-41017.pdf`.

Appendices

A | Electronics

The following sections are meant as additional reading material for those who would like to know more about certain electronic components that was used for the ROB, and how the simplified gain equation for the CBA (Equation 4.9) was derived.

A.1   Transistor theory

The following is additional reading material for Section 4.3.

A full mathematical treatment on the theory of transistors would likely push this appendix well outside the boundaries of what is deemed acceptable lengths. Therefore, only the basics will be presented. Most of the material in this section is based on [52] and [67].

For an NPN BJT transistor operating in the forward active mode, its collector current I_C is given by:

$$I_C = I_S \left(e^{\left(\frac{V_{BE}}{V_T} \right)} - 1 \right) \tag{A.1}$$

V_{BE} is the voltage of the base compared to the emitter (see Figure 4.4). I_S is the *saturation current* of the transistor and V_T is the thermal voltage, both given by:

$$I_S = I_{S0} \left(1 + \frac{V_{CE}}{V_A} \right) \tag{A.2}$$

$$V_T = \frac{kT}{q} \tag{A.3}$$

Here, V_A is the Early voltage, a parameter set by the transistor structure. I_{S0} is the zero bias value of the transistors saturation current. Both these values can often be given by the transistors datasheet, or obtained from measurements. V_T is given by the temperature T in kelvin, Boltzmann's constant k, and the elementary charge constant q (also known as e) . For room temperature (290 K), the value of V_T is 25 mV.

By substituting I_S in Equation A.1 with the expression for I_S in A.2, we get the following equation relating I_C to the base-emitter and collector-emitter voltages:

$$I_C = I_{S0} \left(1 + \frac{V_{CE}}{V_A} \right) \left(e^{\left(\frac{V_{BE}}{V_T} \right)} - 1 \right) \tag{A.4}$$

We see that the collector current is strongly dependent on the base-emitter voltage and the temperature, and less dependent on the collector-emitter voltage. A plot of this equation can be seen in Figure A.1, with three graphs showing how I_C varies with temperature. The collector current is effectively zero until V_{BE} reaches a threshold value. This value is the same threshold value V_{th} mentioned earlier in this section. A variation in V_{CE} will cause I_C to increase faster, but the effect will be small, and is not shown in the plot.

We have so far only looked at how I_C varies with the base-emitter voltage V_{BE}. But a BJT transistor is technically current-controlled, so it would make sense to find an equation for I_C as a function of the base current I_B instead. We have that the collector current is set by the base current through the multiplication factor β:

$$I_C = \beta \times I_B \tag{A.5}$$
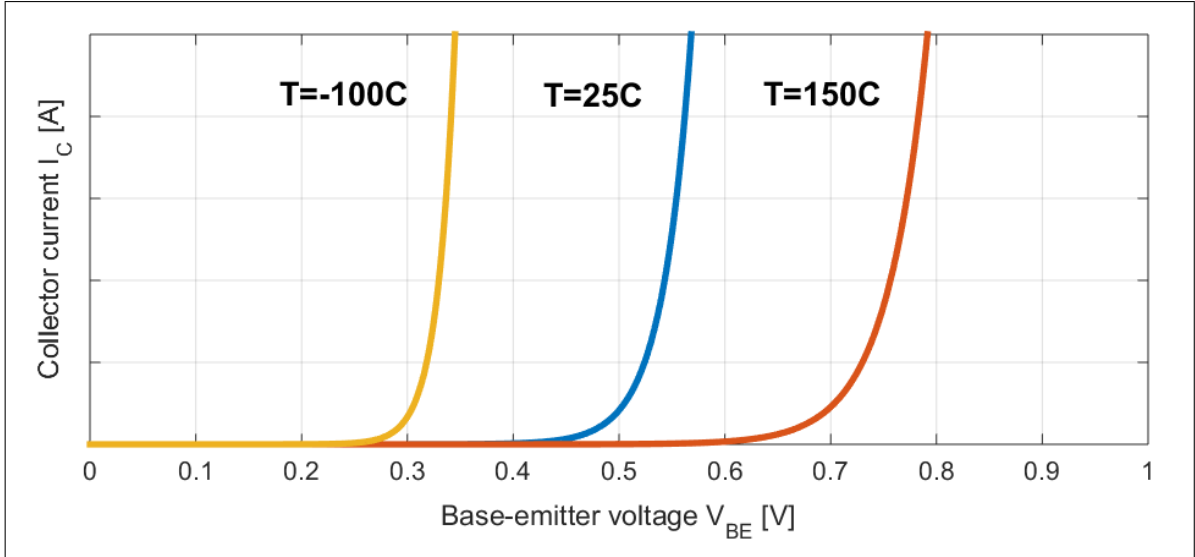
Figure A.1: Plot of a transistors collector current I_C as a function of the base-emitter voltage V_{BE} (Equation A.4). Curves for three different temperatures are shown, to illustrate the effect of temperature on the threshold voltage.

Beta is dependent on the collector-emitter voltage as:

$$\beta = \beta_0 \left(1 + \frac{V_{CE}}{V_A} \right) \tag{A.6}$$

Combining this with Equation A.5, we get:

$$I_C = I_B \beta_0 \left(1 + \frac{V_{CE}}{V_A} \right) \tag{A.7}$$

The value of V_A is typically much larger than V_{CE}, and so the collector current is is mainly controlled by the base current. Equation A.7 only holds when the transistor is in the active mode. For small $V_{CE} < V_{BE}$ the transistor is said to be in saturation mode, and this simplified equation no longer holds. Instead, in saturation mode, I_C quickly drops to 0 as V_{CE} approaches 0. If one were to extrapolate Equation A.7 to its zero-crossing point, one would find the Early voltage V_A. The *Ebers–Moll* model [68, p.19] is used to describe the transistor in the forward active and saturation mode, and this model yields very accurate results when compared to data. The equation for I_C is given by this model as

$$I_C = I_S \left(e^{\left(\frac{V_{BE}}{V_T} \right)} - 1 \right) - I_S \left(1 + \frac{1}{\beta_R} \right) \left(e^{\left(\frac{-V_{BC}}{V_T} \right)} - 1 \right) \tag{A.8}$$

where β_R is the reverse current gain. Plots of this equation is shown in Figure A.2.

To summarize, a transistor can operate as a current-amplifying device, as long as it is biased to be in the Forward Active mode. In this mode, the voltage potential between the base and emitter sets the base current, which in turn controls the collector current. The collector current is, in a simplified model, the same as the base current multiplied by the constant β, which normally takes on a value around the vicinity of 100. A more accurate model takes into account the fact that β is somewhat influenced by the voltage potential between collector and emitter.
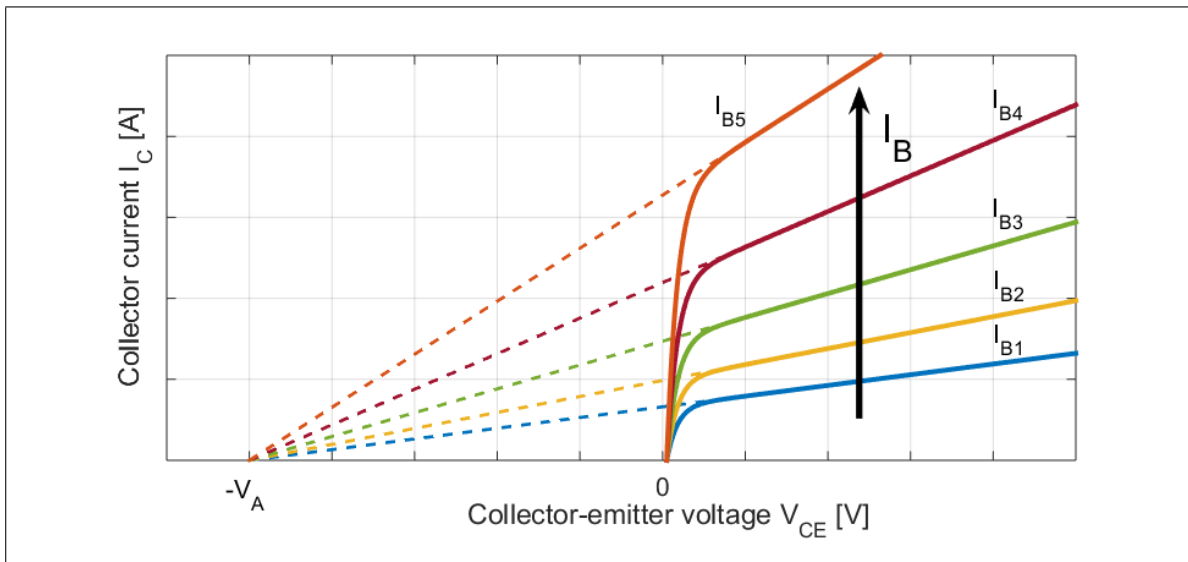
Figure A.2: Illustration of the transistor collector current I_C versus collector-emitter voltage V_{CE}, as described by Equation A.1. Plots for multiple base currents I_B are shown. When V_{CE} reaches a threshold, I_C decreases quickly. This is called the saturation mode. Extrapolation the curves from the Forward Active mode, one finds the transistors Early voltage.

A.2 Derivation of the CBA gain equation

In Section 4.3.2 the theory behind the operation of the CBA (Figure A.3) was introduced. The derivation of a simplified gain equation for the CBA is technical, and is covered in this appendix section.

It can be shown [52, 69] that the voltage gain of the CBA in Figure A.3 is given as:

$$A_v = \frac{v_o}{v_s} = \frac{r_\pi}{R_S + r_\pi} \times G_{me} \times r_{ic} \| R_C \| R_L \tag{A.9}$$

where G_{me} is the transconductance given by

$$G_{me} = \frac{1}{R_{Ete} + r_\pi' \| r_0} \frac{\alpha r_0 + r_\pi'}{r_0 + r_\pi'} \qquad r_\pi' = \frac{r_x}{1 + \beta} + r_\pi \tag{A.10}$$

and r_{ic} is the collector resistance when looking into the collector node, given by

$$r_{ic} = \frac{r_0 + r_\pi' \| R_{Ete}}{1 - \alpha R_{Ete}/(r_\pi' + R_{Ete})} \tag{A.11}$$

Here, r_0 is the small-signal collector-to-emitter resistance, r_x is the *base spreading resistance*, and r_π is the base-emitter resistance. r_x is a constant given by the transistor structure, while r_π and r_0 are given by:

$$r_\pi = \frac{V_T}{I_E} = \frac{V_T}{\frac{1}{\alpha} I_C} = \frac{\alpha V_T}{I_C} \tag{A.12}$$

$$r_0 = \frac{V_A + V_{CE}}{I_C} \tag{A.13}$$

Solving this mess is no easy task. Luckily, we can make a few assumptions that greatly simplify the voltage gain expression (A.9). First of all, r_0 is usually much larger than r_π', R_{Ete}, R_C, and R_L. With this in mind, Equations A.10 and A.11 simplifies to

$$r_{ic} \approx r_0 + \frac{1}{1 - \alpha R_{Ete}/(r'e + R_{Ete})} \tag{A.14}$$

$$G_{me} \approx \frac{\alpha}{R_{Ete} + r_\pi'} \tag{A.15}$$

Also, the last term in Equation A.9) will reduce to $R_C \| R_L$. This is due to the denominator in Equation A.14 always being less than 1, and so $r_{ic} > r_0$. Since we already stated that r_0 is much greater than R_C and R_L, the parallel contribution of r_{ic} to $R_C \| R_L$ can be ignored. Next in our simplification we ignore the base spreading resistance r_x. This resistance is rarely specified by manufacturers, cannot be calculated from other transistor specifications, and is not trivial to measure. Although several techniques exists [70], its value is usually assumed to be on the order of 50 Ω-150 Ω, and ignored in small-signal analysis [52]. This leads to $r_\pi' \approx r_\pi$. We are then left with the following expression for the CBA voltage gain:
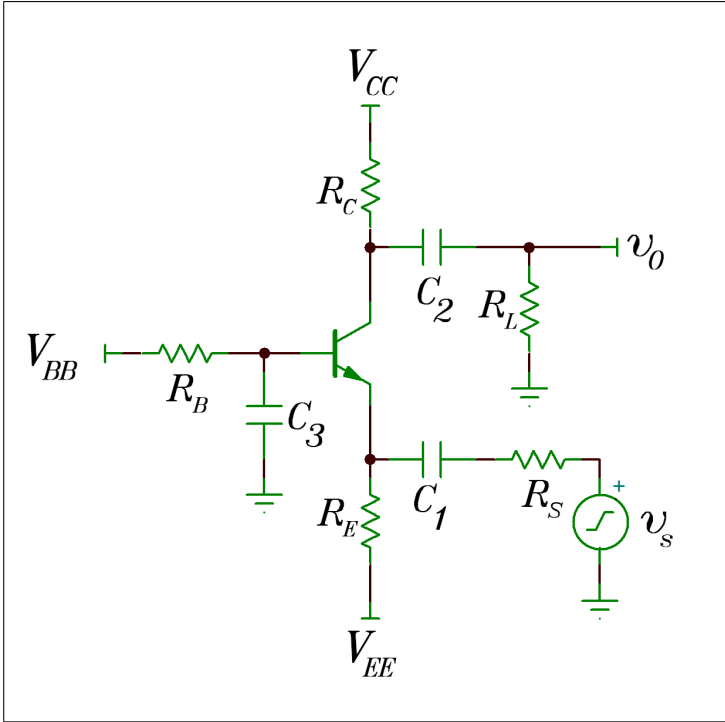
Figure A.3: General circuit diagram of the common base amplifier.

$$A_v = \frac{r_\pi}{R_S + r_\pi} \times \frac{\alpha}{R_{Ete} \times \frac{\alpha V_T}{I_C}} \times R_C \| R_L \tag{A.16}$$

This expression is already a good deal simpler to grasp than what we started out with, but we can still simplify it further. To do that, we first assume $\alpha \approx 1$. α is the reciprocal of β, and is already quite close to 1, no more than a few percent off. Setting it equal to 1 will result in a slight overshoot in the gain expression, but the effect is negligible. The second assumption we make is that $r_\pi \gg R_S$. This leads to the first term in Equation A.16 equalling ~1. The second assumption we make is that the source resistance is much less than the emitter resistor. By requiring that $R_E \gg R_S$, we get that $R_{Ete} \approx R_S$. We finally end up with the following simple expression for the CBA voltage gain:

$$A_v = \frac{1}{R_S + \frac{V_T}{I_C}} \times R_C \| R_L = \frac{I_C}{I_C R_S + V_T} \times R_{Cte} \tag{A.17}$$

We can see that the voltage gain is dependent on three variables: The collector current I_C, the source resistance R_S, and the total external resistance of the collector, R_{Cte}. In order to maximize the gain of a CBA, one should therefore attempt to increase I_C and R_{Cte} while minimizing the signal source resistance R_S.

The validity of Equation A.17 was tested against simulations in TINA. Temperature was held constant, while I_S, R_S, and R_{Cte} were varied while maintaining $r_\pi \gg R_S$ and $R_E \gg R_S$. Each parameter was varied over three orders of magnitude, and the results are shown in Figure A.4. The results show that the simplified model fits within 5 percent and/or 3 dB of the simulated results in 24 of the 27 cases. The model loses its validity especially for high I_S and low R_{Cte}.

$R_S = 0.1\Omega$								
I_S	$R_{Cte} = 10\Omega$		Error (%)	$R_{Cte} = 100\Omega$		Error (%)	$R_{Cte} = 1000\Omega$		Error (%)
	Eq.	Sim.		Eq.	Sim.		Eq.	Sim.	
1mA	-8 dB	-8 dB	-1,8	11,70 dB	11,11 dB	5,3	31,70 dB	31,24 dB	1,5
10mA	11 dB	10 dB	10,6	31,40 dB	30,08 dB	4,4	51,40 dB	50,57 dB	1,6
100mA	29 dB	22 dB	29,7	48,90 dB	41,02 dB	19,2	68,89 dB	64,11 dB	7,5

$R_S = 1\Omega$								
I_S	$R_{Cte} = 10\Omega$		Error (%)	$R_{Cte} = 100\Omega$		Error (%)	$R_{Cte} = 1000\Omega$		Error (%)
	Eq.	Sim.		Eq.	Sim.		Eq.	Sim.	
1mA	-8,60 dB	-8,75 dB	-1,7	11,40 dB	11,24 dB	1,4	31,40 dB	31,22 dB	0,6
10mA	8,90 dB	8,39 dB	6,1	28,90 dB	28,38 dB	1,8	48,89 dB	48,27 dB	1,3
100mA	18,00 dB	16,74 dB	7,5	38,00 dB	36,71 dB	3,5	58,00 dB	56,32 dB	3,0

$R_S = 10\Omega$								
I_S	$R_{Cte} = 10\Omega$		Error (%)	$R_{Cte} = 100\Omega$		Error (%)	$R_{Cte} = 1000\Omega$		Error (%)
	Eq.	Sim.		Eq.	Sim.		Eq.	Sim.	
1mA	-11,10 dB	-11,20 dB	-0,9	8,89 dB	8,80 dB	1,0	28,89 dB	28,78 dB	0,4
10mA	-2,00 dB	-2,16 dB	-7,4	17,99 dB	17,84 dB	0,8	37,99 dB	37,80 dB	0,5
100mA	-0,22 dB	-0,40 dB	-45,0	19,77 dB	19,59 dB	0,9	39,77 dB	39,50 dB	0,7

Figure A.4: Results of simulations performed in order to verify the validity of the simplified equation for CBA gain (Equation A.17). The three variables R_S, R_{Cte}, and I_C were varied over three orders of magnitude, and the maximum gain found by simulations in TINA-TI. The gain for the same configuration of parameters was found by Equation A.17, and their relative differences computed. The simplified equation shows to be valid for most configurations, displaying a relative error of less than 5 % or an absolute error of less than 3 dB in most cases.

This simplified model for the BCA gain was used as the starting point for deciding the preamplifier specifications. A few values had already been set beforehand: V_{EE} was supplied by a 15V voltage regulator, and R_S is defined by an analog switch with an on-resistance of 900 m$\Omega \approx 1\,\Omega$. If we assume that R_L is at least a few times greater than R_C, its contribution to R_{Cte} can be ignored. This leads to an even simpler form of A_v:

$$A_v = \frac{I_C}{I_C + V_T} \times R_C \tag{A.18}$$

A.3   Operational amplifiers

The following is additional reading for Section 4.4.

An operational amplifier (op amp) is in effect an amplifier with extremely high gain, typically between 10 V/mV to 200 V/mV. It will amplify the difference between its two input terminals V_+ and V_- by this amount, called its *open loop gain* A_{OL}:

$$V_{out} = A_{OL}(V_+ - V_-) \qquad (A.19)$$

Since the open loop gain is so high, an op amp cannot be successfully operated without some form of negative feedback to limit the gain, as it would be almost impossible to keep its output from saturating.

The basic negative feedback system is shown in Figure A.5. The gain A of such a system is given as:

$$A = \frac{V_{out}}{V_{in}} = \frac{\alpha}{1 + \beta\alpha} \approx \frac{1}{\beta}, \quad \text{for } \beta\alpha \gg 1 \qquad (A.20)$$

For an op amp, α is the open loop gain and β a feedback value set by external circuitry such that $\beta\alpha \gg 1$, and the approximation holds. The negative feedback constant β is given by a voltage divider set by two resistors R_f and R_g in the feedback path. The resulting gain of the op amp is called the *closed loop gain* A_{CL}, and is given as:

$$A_{CL} = 1 + \frac{R_f}{R_g} \qquad \text{Non-inverting operation}$$

$$A_{CL} = \frac{R_f}{R_g} \qquad \text{Inverting operation}$$

An illustration of the inverting and non-inverting op amp configurations can be seen in Figure A.6. Whether one should use an inverting or non-inverting configuration depends on the application. In most cases their characteristics are the same, except for the fact that the output signal is inverted if operated in inverting mode.
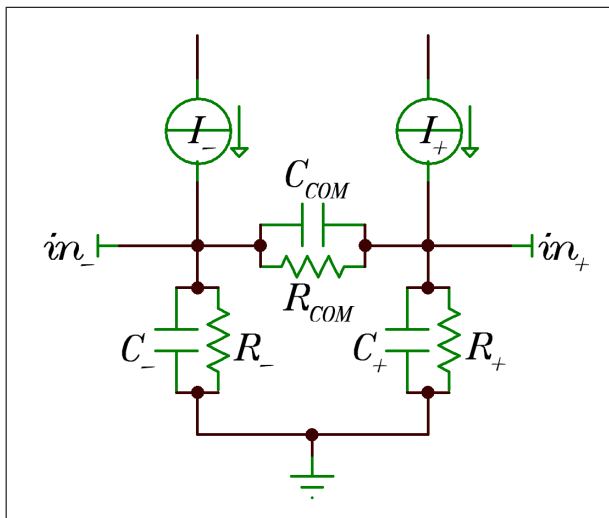
The maximum bandwidth of an op amp is given by its *gain bandwidth product*. It is called so because the gain and -3dB cutoff frequency f_{-3dB} of an op amp is related by the following
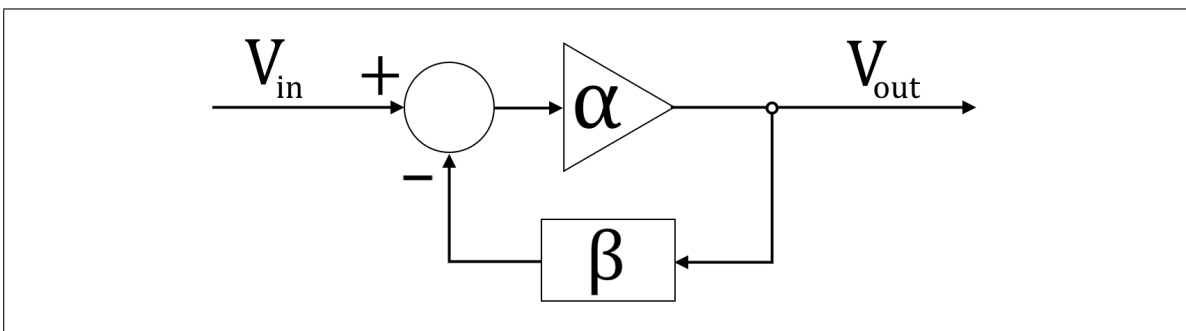
Figure A.5: The basic negative feedback system. If $\beta\alpha \gg 1$ then $V_{out} = V_{in}/\beta$
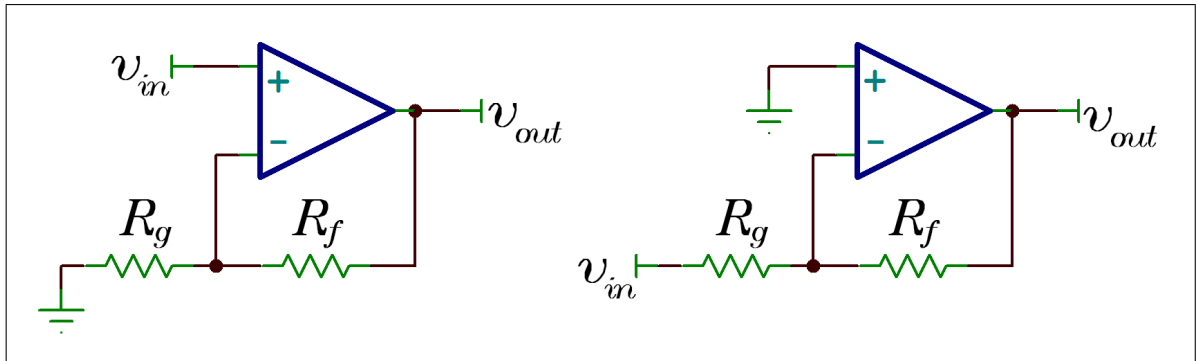
Figure A.6: The two basic operational amplifier configurations. Inverting (right) and non-inverting (left).

Figure A.7: Electrically equivalent model of the inputs of a voltage feedback operational amplifier. The bias currents I_- and I_+ arise from the operation amplifier design, and are undesirable, as they will induce a voltage on the inputs.

equation:

$$A_{CL} \times f_{-3dB} = \text{GBW} \tag{A.21}$$

This means that if gain is increased, a penalty on bandwidth will be introduced to keep the product of gain and bandwidth constant. For some op amp models, the GBW may depend based on whether it is operated in inverting or non-inverting mode.

The inputs of a voltage feedback op amp can be modelled as current sinks (or sources, depending on the model), each with a large resistance and small capacitance to earth. A lower-valued resistor and a capacitor connects the two inputs (Figure A.7). The current sources represents what is called the *input bias currents*. The input bias current of an op amp can vary greatly, from a few fA for high-precision voltage meters to tens of μA for high-GBW devices. If these current sources are not taken into account when designing the op amp-based circuit, undesired effects may arise. For example, if the designer forgets about the current sources and does not implement a low-impedance path to ground on the noninverting input, the current must flow through the high-valued internal resistor to ground. This causes the noninverting input voltage to shift, and this shift is amplified and output by the op amp. The higher the input bias currents are, the more of an issue they can be.

A.4   Wheatstone bridge and RTD

The following is additional reading material for Section 4.8.

A.4.1   The Wheatstone bridge

The Wheatstone bridge is an electrical circuit consisting of two balanced voltage dividers, made with three resistors of well-known values and one resistor of unknown value. The value of the unknown resistor can be determined to great accuracy by measuring the voltage difference between the common nodes of the two dividers, and is often used in strain gauge and resistor-based temperature measurements where small changes in resistances must be measured. A circuit diagram of a Wheatstone bridge can be seen in Figure A.8.

Using Kirchoff's rules, the difference voltage V_D between the two legs of the Wheatstone bridge in Figure A.8 can be found to be

$$V_D = V_0 \left(\frac{R_2}{R_1 + R_2} - \frac{R_3}{R_x + R_3} \right) \tag{A.22}$$

The accuracy to which R_x can be determined is limited by the accuracy of the voltage supply and the other resistors in the Wheatstone bridge. High-accuracy resistors with low temperature coefficients are therefore normally used. Since the uncertainties in the resistor values will depend on temperature, the accuracy must be calculated at the expected maximum or minimum temperature.

A.4.2   The RTD element

The RTD is a temperature-dependent resistive element that differ from the more common thermistor by the fact that it is made of a strip or coil of metal instead of a polymer or ceramic.
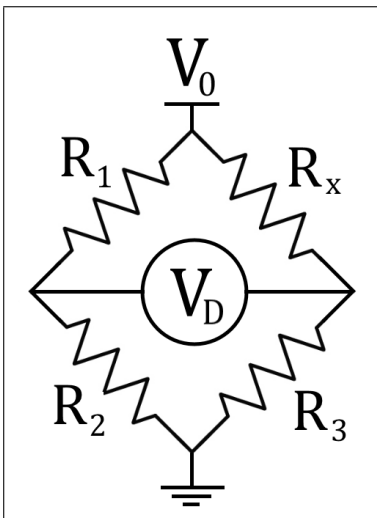
Figure A.8: Circuit diagram of a Wheatstone bridge. Resistors R_1, R_2 and R_3 are of accurately known values and R_x is unknown. Measuring the voltage V_D allows one to determine R_x to great accuracy.

Its temperature dependency is therefore nearly linear, and can be used over a much larger range of temperatures than the thermistor. RTDs can be made from any type of metal, but high purity platinum or nickel is commonly used due to their high resistance to corrosion and linear temperature dependency over a large range of temperatures. Platinum is used for the highest accuracy RTD elements due to its chemical inertness and highly linear response. RTD elements made from high-purity platinum can be calibrated to extremely high accuracies, for example by the National Institute for Standards and Technology (NIST) in the USA, and are used as reference temperature sensors in laboratories around the world [71].

B | Uncertainty calculations

In the following sections, the calculations related to the derivation of the uncertainties in V_{BIAS}, MPPC temperature and trigger threshold are shown in detail. The reader will first be reminded of the mathematical rules that govern error propagation.

B.1   Introduction

The following rules for propagation of uncertainty assumes that the uncertainties represents the standard deviation σ of the random variable which selects the error, and that they are uncorrelated and random.

For $Q(a, b, ...)$, where a, b, ..., are variables with absolute uncertainties σ_a, σ_b, ..., we have the following rules:

Summation and subtraction

If $Q = (a \pm \sigma_a) + (b \pm \sigma_b) + ...$, the error σ_Q is found by

$$\sigma_Q = \sqrt{(\sigma_b)^2 + (\sigma_b)^2 + ...} \tag{B.1}$$

The equation is the same for subtraction.

Multiplication and division

If $Q = (a \pm \sigma_b) \times (b \pm \sigma_v) \times ...$, the *relative* error $\hat{\sigma}_Q$ is found by

$$\hat{\sigma}_Q = \frac{\sigma_Q}{|Q|} = \sqrt{\left(\frac{\sigma_a}{a}\right)^2 + \left(\frac{\sigma_b}{b}\right)^2 + ...} \tag{B.2}$$

Here, σ_a/a, σ_b/b, etc, are the relative errors $\hat{\sigma}_a$, $\hat{\sigma}_b$, etc. The equation is the same for division.

General equations

For more complicated equations one can use the general rule which the above rules are special cases of. If $Q(a, b, ...)$ is any equation of the variables a, b, ..., the absolute error σ_Q is found by

$$\sigma_Q = \sqrt{\left(\frac{\partial Q}{\partial a}\sigma_a\right)^2 + \left(\frac{\partial Q}{\partial b}\sigma_b\right)^2 + ...} \tag{B.3}$$

Standard deviation of the uniform distribution

A uniform distribution is a symmetric probability distribution. Its value is $f(x) = 1/(B - A)$ at all points on the interval $A < x < B$, and is zero elsewhere. The standard deviation of the uniform distribution is given as:

$$\sigma = \sqrt{\frac{(B - A)^2}{12}} \tag{B.4}$$

This is often used when performing error analysis with a variable that has an error distribution set by an unknown random variable. By generalizing the error distribution to be a uniform random variable, the standard deviation can be found and used in error propagations.

B.2 Derivation of the uncertainty in V_{BIAS}

In this appendix section, calculations of the uncertainties related to the determination of V_{BIAS} is covered. The uncertainties in the voltage divider K_1 and ADC are found and combined to find the total accuracy of V_{BIAS}.

Uncertainty of voltage divider

Voltage divider K_1 (Figure 4.21) consists of resistors R_2 and R_3. Their constant errors are measured and compensated for in the ROB microcontroller software, and so we need only consider the uncertainty in K_1 arising from their temperature dependencies. Since $R_2 = 14 \times R_3$, we are far from the center point, and we can simply add their TCRs (α) to get the TCR of K_1. The TCRs of R_2 and R_3 are both $\pm 25 \times 10^{-6}\,°C^{-1}$.

$$\alpha_{K1} = \alpha_{R2} + \alpha_{R3} = \pm 25 \times 10^{-6}\,°C^{-1} + \pm 25 \times 10^{-6}\,°C^{-1} = \pm 50 \times 10^{-6}\,°C^{-1} \qquad \text{(B.5)}$$

The TCRs of resistors are given by manufacturers as maximum and minimum limits, with no information about the distribution. We make the worst-case realistic assumption that they follow a uniform random variable, and the standard deviation of the temperature dependency error is therefore given as:

$$\hat{\sigma}_{\alpha K1} = \sqrt{\frac{(50 \times 10^{-6}\,°C^{-1} + 50 \times 10^{-6}\,°C^{-1})^2}{12}} \approx 29 \times 10^{-6}\,°C^{-1} \qquad \text{(B.6)}$$

Evaluating this at the maximum expected MPPC temperature range of $25\,°C \pm 14\,°C$, we get the maximum relative standard deviation of the error of K_1:

$$\hat{\sigma}_{K1} = \hat{\sigma}_{\alpha K1} \times 14°\,°C$$

$$\hat{\sigma}_{K1} \approx \pm 0.04\,\%$$

The largest value of the bias voltage is $V_{BIAS} = 64\,V$. We therefore get the maximum absolute uncertainty in K_1:

$$\sigma_{K1} = \hat{\sigma}_{K1} \times 64\,V$$

$$\sigma_{K1} \approx 26\,mV$$

Uncertainty of ADC and combined uncertainty

The ADC used for the V_{BIAS} sensing circuit has 12 bits of resolution. It can digitize input voltages up to $5\,V$, which equates to a maximum bias voltage of $V_{BIAS} = 5\,V \times 14$, where 14 is the division factor of K_1. This results in a resolution of $70\,V/2^{12} \approx 17\,mV$, and therefore a digitization uncertainty of $\pm 8.5\,mV$.

The PDF of the digitization error is a uniform random variable, and the standard deviation

is found:

$$\sigma_{ADC} = \sqrt{\frac{(8.5\,\text{mV} + 8.5\,\text{mV})^2}{12}}$$

$$\sigma_{ADC} \approx 4.9\,\text{mV}$$

The combined uncertainties gives the total absolute standard deviation of the error in V_{BIAS}.

$$\sigma_{BIAS} = \sqrt{\sigma_{ADC}^2 + \sigma_{K1}^2} = \sqrt{(4.9\,\text{mV})^2 + (26\,\text{mV})^2}$$

$$\sigma_{BIAS} \approx 27\,\text{mV}$$

B.3 Derivation of the uncertainty in MPPC temperature

This section describes how the uncertainty in the measurements of the MPPC temperature was estimated. We first find an equation for the RTD resistance R_x as a function of the in amp output voltage, and perform an error analysis. We already have an equation for the difference voltage V_D, Equation A.22. We begin by reordering it to get an equation for R_x.

$$V_D = V_0 \left(\frac{R_2}{R_1 + R_2} - \frac{R_3}{R_x + R_3} \right)$$

$$\frac{V_D}{V_0} - \frac{R_2}{R_1 + R_2} = -\frac{R_3}{R_x + R_3}$$

$$R_x + R_3 = \frac{R_3}{\frac{R_2}{R_1+R_2} - \frac{V_D}{V_0}}$$

$$R_x = \frac{R_3}{\frac{R_2}{R_1+R_2} - \frac{V_D}{V_0}} - R_3 \tag{B.7}$$

Next we must find an equation that describes V_D, the difference voltage between the bridge legs. This voltage is amplified by the in amp by the gain factor G, and digitized by the microcontroller, resulting in the measured voltage V_{ADC}. We therefore have

$$V_{ADC} = V_D \times G$$

$$V_D = \frac{V_{ADC}}{G} \tag{B.8}$$

Substituting this into Equation B.7 gives the full equation for R_x:

$$R_x = \frac{R_3}{\frac{R_2}{R_1+R_2} - \frac{V_{ADC}}{V_0 G}} - R_3 \tag{B.9}$$

All the variables in this equation have uncertainties. The resistances, V_0, and G have relative errors and temperature dependencies, and the digitized signal has finite resolution. Understanding how these uncertainties propagate through the equation for R_x to give the final uncertainty on its value can be done by using the method of error propagation.

We have that the standard deviation in the error of a function $Q(a, b, ...)$, also called uncertainty, is given by

$$\sigma_Q = \sqrt{\left(\frac{\partial Q}{\partial a} \sigma_a \right)^2 + \left(\frac{\partial Q}{\partial b} \sigma_b \right)^2 + ...} \tag{B.10}$$

We therefore need to find the partial derivatives of R_x for the six variables it depends on. But before doing so, we must change Equation B.9 slightly. To understand why, let us consider the resistors R_2 and R_3. They are of the same model, and therefore it is very likely that they have equal temperature dependencies. To prevent the error estimation of R_x from counting the temperature-dependent uncertainties once for both R_2 and R_3, we implement them into a second variable, R_T. To do so, we first look at the general equation that describes a resistor

with an uncertainty ϵ and temperature dependency α:

$$R = r(1 \pm \epsilon)(1 \pm \alpha T) \tag{B.11}$$

Here r is the ideal value of R. Multiplying together the parentheses and following the rules for multiplication of uncertainties gives us

$$R = r \left(1 \pm \epsilon \pm \alpha T \pm \sqrt{(\epsilon)^2 + (\alpha T)^2} \right) \tag{B.12}$$

For most resistors, and definitely for the precision resistors used in the Wheatstone bridge, $\epsilon \ll 1$ and $\alpha T \ll 1$. The square root error term will therefore be negligible compared to the two other error terms, and R can be approximated as

$$R \approx r(1 \pm \epsilon \pm \alpha T) \tag{B.13}$$

We can now begin the process of extracting R_T from R_2 and R_3, which are given as

$$R_2 = r_2(1 \pm \epsilon_{r2} \pm \alpha_{r2}T) \qquad R_3 = r_3(1 \pm \epsilon_{r3} \pm \alpha_{r3}T) \tag{B.14}$$

r_2 and r_3 represents the ideal values of R_2 and R_3. ϵ_{r2} and ϵ_{r3} are uncorrelated, and represents the uncertainty in resistance at $25\,°C$ specified by the manufacturer. However, the temperature dependencies α_{r2} and α_{r3} will likely be equal and therefore not uncorrelated, for the reasons stated earlier. We reorder the expressions for R_2 and R_3 to get the temperature dependencies alone:

$$R_2 = r_2(1 \pm \epsilon_{r2}) + r_2(0 \pm \alpha_{r2}T) \qquad R_3 = r_3(1 \pm \epsilon_{r3}) + r_3(0 \pm \alpha_{r3}T) \tag{B.15}$$

The reason for the zeros in the second terms will be explained soon. We then introduce two new variables r_T and α_{rT}. We can set $r_2 = r_3 = r_T$ and $\alpha_{r2} = \alpha_{r3} = \alpha_{rT}$ for the second terms of Equation B.15 because R_2 and R_3 are of the same model. By doing this we create the new independent variable R_T which represents the shared and equal temperature-dependent uncertainties of the two resistors.

$$R_2 = r_2(1 \pm \epsilon_{r2}) + r_T(0 \pm \alpha_{rT}T) \quad R_3 = r_3(1 \pm \epsilon_{r3}T) + r_T(0 \pm \alpha_{rT}T)$$
$$R_2 = r_2(1 \pm \epsilon_{r2}) + R_T \qquad\qquad R_3 = r_3(1 \pm \epsilon_{r3}) + R_T \tag{B.16}$$

We see that R_T is an imaginary resistor that consists only of the equal temperature-dependent uncertainties of R_2 and R_3. To simplify the notation, we make the following declarations:

$$R_2' \equiv r_2(1 \pm \epsilon_{r2}) \implies R_2 = R_2' + R_T \tag{B.17}$$

$$R_3' \equiv r_3(1 \pm \epsilon_{r3}) \implies R_3 = R_3' + R_T \tag{B.18}$$

Substituting these into the equation for R_x (Equation B.9) yields:

$$R_x = \frac{R_3' + R_T}{\frac{R_2' + R_T}{R_1 + R_2' + R_T} - \frac{V_{ADC}}{V_0 G}} - R_3' + R_T \tag{B.19}$$

By Equation B.10 we get the following expression for the uncertainty in R_x:

$$\sigma_{Rx} = \left[\left(\frac{\partial R_x}{\partial R_1} \sigma_{R1} \right)^2 + \left(\frac{\partial R_x}{\partial R_2'} \sigma_{R'2} \right)^2 + \left(\frac{\partial R_x}{\partial R_3'} \sigma_{R'3} \right)^2 + \left(\frac{\partial R_x}{\partial R_T} \sigma_{RT} \right)^2 \right.$$
$$\left. + \left(\frac{\partial R_x}{\partial V_0} \sigma_{V0} \right)^2 + \left(\frac{\partial R_x}{\partial G} \sigma_G \right)^2 + \left(\frac{\partial R_x}{\partial V_{ADC}} \sigma_{ADC} \right)^2 \right]^{1/2} \tag{B.20}$$

To perform an error propagation of R_x, we must therefore find the partial derivatives of R_x with respect to all the variables it depends on, as well as the uncertainties of the variables.

We begin with the partial derivatives. The calculations are tedious and uninteresting, and have been listed in a separate area of the appendix, Appendix B.4.1. A term that often occurs in the solutions is

$$\frac{R_2'}{R_2' + R_1} - \frac{V_{ADC}}{V_0 G} \tag{B.21}$$

Due to the way that the temperature sensing circuit is designed, the largest value for the second term is $5\,\text{V}/(2.5\,\text{V} \times 80) = 0.025$. The first term is set by the resistor values, and is ≈ 0.905. We can therefore neglect the second term, which will not contribute much. This approximation is also justified by the fact that it makes the subsequent mathematical treatment of σ_{RX} much easier, as many of the partial derivatives are greatly simplified. Performing this approximation, we get the following partial derivatives:

$$\frac{\partial R_x}{\partial R_1} = \frac{R_3'}{R_2'} \qquad \frac{\partial R_x}{\partial R_T} = \frac{R_1}{R_2'} - \frac{R_1 R_3'}{R_2'^2} \qquad \frac{\partial R_x}{\partial V_{ADC}} = \frac{1}{V_0 G} \frac{R_3'(R_2' + R_1)^2}{R_2'^2}$$

$$\frac{\partial R_x}{\partial R_2'} = -\frac{R_3' R_1}{R_2'^2} \qquad \frac{\partial R_x}{\partial V_0} = -\frac{V_{ADC}}{V_0^2 G} \frac{R_3'(R_2' + R_1)^2}{R_2'^2} \qquad \frac{\partial R_x}{\partial G} = -\frac{V_{ADC}}{V_0 G^2} \frac{R_3'(R_2' + R_1)^2}{R_2'^2} \tag{B.22}$$

$$\frac{\partial R_x}{\partial R_3'} = \frac{R_1}{R_2'}$$

Substituting each variable for their ideal values gives:

$$\frac{\partial R_x}{\partial R_1} = \frac{10^4\,\Omega}{10^4\,\Omega} = 1$$

$$\frac{\partial R_x}{\partial R'_2} = -\frac{10^4\,\Omega \times 1.05 \times 10^3\,\Omega}{(10^4\,\Omega)^2} = -0.105$$

$$\frac{\partial R_x}{\partial R'_3} = \frac{1.05 \times 10^3\,\Omega}{10^4\,\Omega} = 0.105$$

$$\frac{\partial R_x}{\partial R_T} = \frac{1.05 \times 10^3\,\Omega}{10^4\,\Omega} - \frac{1.05 \times 10^3\,\Omega \times 10^4\,\Omega}{(10^4\,\Omega)^2} = 0 \tag{B.23}$$

$$\frac{\partial R_x}{\partial V_0} = -\frac{V_{ADC}}{(2.5\,\text{V})^2 \times 80}\frac{10^4\,\Omega(10^4\,\Omega + 1.05 \times 10^3\,\Omega)^2}{(10^4\,\Omega)^2} \approx -24.4 V_{ADC}\,\Omega\,\text{V}^{-1}$$

$$\frac{\partial R_x}{\partial G} = -\frac{V_{ADC}}{2.5\,\text{V} \times 80^2}\frac{10^4\,\Omega(10^4\,\Omega + 1.05 \times 10^3\,\Omega)^2}{(10^4\,\Omega)^2} \approx -0.8 V_{ADC}\,\Omega$$

$$\frac{\partial R_x}{\partial V_{ADC}} = \frac{1}{2.5\,\text{V} \times 80}\frac{10^4\,\Omega(10^4\,\Omega + 1.05 \times 10^3\,\Omega)^2}{(10^4\,\Omega)^2} \approx 61\,\Omega\,\text{V}^{-1}$$

The benefit of using equal resistor models for R_2 and R_3 becomes apparent, as the partial derivative of R_x with respect to R_T equals 0. In other words, the temperature dependent uncertainties of R_2 and R_3 cancel each other out. We can therefore ignore R_T in the subsequent error analysis.

Now that we have the partial derivatives required for the error analysis, we must estimate the uncertainties in the parameters. For Equation B.20 to hold, the uncertainties σ_{R1}, $\sigma_{R'2}$, etc., must be equal to the standard deviation of the random variable that sets the total absolute errors related to the variables R_1, R'_2, etc.. As we will see in the following paragraphs, each variable often have more than one error component, and they are rarely given by the manufacturers as standard deviations, and are usually relative errors. Finding expressions for the combined absolute standard deviation of all error components for each variable is the goal of the following sections. We begin by finding the standard deviations for each error component in all variables.

Estimation of uncertainties in resistors

The resistors R_1, R'_2 and R'_3 are given as

$$R_1 = r_1(1 \pm \epsilon_{r1} \pm \alpha_{r1}T) \qquad R'_2 = r_2(1 \pm \epsilon_{r2}) \qquad R'_3 = r_3(1 \pm \epsilon_{r3}) \tag{B.24}$$

The fixed relative uncertainties (ϵ) represents manufacturing uncertainties, and are given as percentage maximum and minimum limits on the error of the resistor value. It provides no information on the probability density function (PDF) of the error, only that it is confined to the interval given by the manufacturer. The fixed relative uncertainties are therefore clearly not representing standard deviations and we cannot use them in Equation B.10.

In order to transform the fixed relative uncertainties into standard deviations we need to

make an assumption about the PDF of the random variable that chooses ϵ. The worst possible realistic case is that any error within the limits is equally likely, i.e. that it follows a uniform distribution. R_1, R_2' and R_3' all have the same manufacturing errors of $\epsilon \pm 0.1\,\%$, resulting in an uncertainty of

$$\hat{\sigma}_{\epsilon R} = \hat{\sigma}_{\epsilon R1} = \hat{\sigma}_{\epsilon R'2} = \hat{\sigma}_{\epsilon R'3} = \sqrt{\frac{(-1 \times 10^{-3} - 1 \times 10^{-3})^2}{12}} = 0.057735...\% \approx 0.06\,\% \quad \text{(B.25)}$$

The hats are added to remind us that these standard deviations are for *relative* errors.

The temperature dependencies (α) of resistors are also given as maximum and minimum limits by the manufacturers, with no information about the distribution. Again we make the worst-case realistic assumption that they follow a uniform random variable, and find the standard deviation of the temperature dependency of R_1 which has α given as $\pm 25 \times 10^{-6}\,°\text{C}^{-1}$ from the manufacturer:

$$\hat{\sigma}_{\alpha R_1} = \sqrt{\frac{(-25 \times 10^{-6}\,°\text{C}^{-1} - 25 \times 10^{-6}\,°\text{C}^{-1})^2}{12}} = 14.43375... \times 10^{-6}°\text{C}^{-1}$$

$$\hat{\sigma}_{\alpha R_1} \approx 14 \times 10^{-6}\,°\text{C}^{-1} \quad \text{(B.26)}$$

Estimation of other uncertainties

The variables V_0, V_{ADC} and G have uncertainties related to their ideal values. Starting with the gain G, we have from the in amp datasheet [72] that it has a constant error ϵ_G and a temperature dependency α_G, exactly as the resistors. Both of these errors are $\ll 1$ and the simplification from Equation B.13 holds. We therefore have the following expression for G:

$$G = g(1 \pm \epsilon_G \pm \alpha_G T) \quad \text{(B.27)}$$

We are given that a typical value for ϵ_G is $\pm 0.3\,\%$, and the maximum value is $\pm 1\,\%$. For α_G it lists $\pm 2.5 \times 10^{-6}\,°\text{C}^{-1}$ as a typical value and $\pm 25 \times 10^{-6}\,°\text{C}^{-1}$ as the maximum. But as before, we cannot make any assumption of the distributions. One cannot know if the typical values are averages, modes, standard deviations, or something else entirely. We therefore do as in the previous paragraph, and estimate the probability distributions to be the worst possible realistic case, and get the standard deviations for these relative errors as:

$$\hat{\sigma}_{\alpha G} = \sqrt{\frac{(-25 \times 10^{-6}\,°\text{C}^{-1} - 25 \times 10^{-6}\,°\text{C}^{-1})^2}{12}} = 14.43375... \times 10^{-6}°\text{C}^{-1}$$

$$\hat{\sigma}_{\alpha G} \approx 14 \times 10^{-6}\,°\text{C}^{-1} \quad \text{(B.28)}$$

and

$$\hat{\sigma}_{\epsilon G} = \sqrt{\frac{(-1\,\% - 1\,\%)^2}{12}} = 0.57735...\% \approx 0.6\,\% \quad \text{(B.29)}$$

To estimate the uncertainties for V_0, we get from the datasheet for the precision voltage

source [58] that it has an initial accuracy ϵ_V of $\pm 0.08\,\%$ and a temperature dependency α_V of at most $\pm 8 \times 10^{-6}\,°C^{-1}$. Using the same reasoning as before, we have the following worst-case estimations of their standard deviations:

$$\hat{\sigma}_{\alpha V0} = \sqrt{\frac{(-8 \times 10^{-6}\,°C^{-1} - 8 \times 10^{-6}\,°C^{-1})^2}{12}} = 4.618802... \times 10^{-6}°C^{-1} \qquad (B.30)$$

$$\hat{\sigma}_{\alpha V0} \approx 5 \times 10^{-6}\,°C^{-1}$$

and

$$\hat{\sigma}_{\epsilon V0} = \sqrt{\frac{(-0.08\,\% - 0.08\,\%)^2}{12}} = 0.046188...\% \approx 0.05\,\% \qquad (B.31)$$

Finally, estimating the uncertainties in V_{ADC} is a little more involved. It has a constant digitization error δ_{ADC} which comes from the fact that the digitization operates with a finite resolution. If a certain true value of V_{true} is digitized by an ADC with a resolution of n bits and a reference voltage of v_{ref}, then any true value of $V_{true} \pm v_{ref}/2^{(n+1)}$ will result in the same digitized value. In our case, this results in an accuracy of:

$$\delta_{ADC} = \pm\frac{v_{ref}}{(2^{(n+1)}} = \pm\frac{5\,V}{2^{(10+1)}} = \pm 2.4414... \times 10^{-}3V \approx \pm 2.4\,mV \qquad (B.32)$$

We know that the PDF of this error is indeed a uniform random variable, since any value of V_{true} is equally probable within the resolution limits when $\delta_{ADC} \ll \Delta V_{true}$, where ΔV_{true} is the expected variation in V_{true}. V_{true} is on our case the in amp output, which is expected to vary several volts over the SiPM temperature range, and the requirement holds. We can therefore find the standard deviation of δ_{ADC}:

$$\sigma_{\delta ADC} = \sqrt{\frac{(-2.4\,mV - 2.4\,mV)^2}{12}} = 1.38564...mV \approx 1.4\,mV \qquad (B.33)$$

An additional error related to V_{ADC} comes from the fact that the in amp has an offset error which manifests itself as an uncertainty ϵ_{off} on the input voltage V_D. This error is then multiplied with G. To see why, we find the expression for V_D as in Equations B.3 and B.8, but this time also considering uncertainties:

$$V_{ADC} = (V_D \pm \epsilon_{off}) \times g(1 \pm \epsilon_G \pm \alpha_G T) \pm \delta_{ADC}$$

$$\frac{V_{ADC} \pm \delta_{ADC}}{g(1 \pm \epsilon_G \pm \alpha_G T)} = (V_D \pm \delta_{off})$$

$$V_D = \frac{V_{ADC} \pm \delta_{ADC} \pm \delta_{off} \times g(1 \pm \epsilon_G \pm \alpha_G T)}{g(1 \pm \epsilon_G \pm \alpha_G T)}$$

$$V_D \approx \frac{V_{ADC} \pm \delta_{ADC} \pm \delta_{off}g}{g(1 \pm \epsilon_G \pm \alpha_G T)} \qquad (B.34)$$

The approximation performed in the last step assumes that the multiplication of δ_{off} with the uncertainties in G has a negligible effect, since $\epsilon_G \ll 1$ and $\alpha_G \ll 1$.

The offset voltage is given in the datasheet for the in amp as having a typical value of $9\,\mu V$ and a maximum value of $50\,\mu V$. It is very likely that the PDF for this error source follows a normal distribution, but to be safe we assume it is a uniform random variable and get a worst-case estimation of the standard deviation in δ_{off}:

$$\sigma_{\delta off} = \sqrt{\frac{(-50\,\mu V - 50\,\mu V)^2}{12}} = 28.86751...\mu V \approx 29\,\mu V \tag{B.35}$$

The combined absolute uncertainty in V_{ADC} is then

$$\sigma_{ADC} = \pm\sigma_{\delta ADC} \pm \sigma_{\delta off}g$$

$$\sigma_{ADC} = \pm\sqrt{\sigma_{\delta ADC}^2 + (\sigma_{\delta off}g)^2}$$

$$\sigma_{ADC} = \pm\sqrt{(1.4\,mV)^2 + (29\,\mu V \times 80)^2} = \pm\sqrt{1.96\,\mu V^2 + 5.3824\,\mu V^2} = 2.70968...mV$$

$$\sigma_{ADC} \approx 2.7\,mV \tag{B.36}$$

Absolute standard deviations

The uncertainties in Equation B.20, represented by standard deviations of the random variable that selects the error, are absolute. But many of the standard deviations that we have estimated in the two previous paragraphs are relative, meaning that they are represented as percentage. The only exception is σ_{ADC}, which is absolute. Before we can use them in Equation B.20 to find the uncertainty of R_x, we must first make them absolute. Since many of the relative uncertainties scale with temperature, we set $T = 14\,°C$ to get the worst case. A maximum and minimum temperature deviation from $25\,°C$ of $14\,°C$ is what the ROB has been designed for.

Following the rules for addition of uncertainties, we get the following values for the absolute

uncertainties of each variable aasdasdasda:

$$\sigma_{R1} = r_1 \sqrt{\hat{\sigma}_{\epsilon R}^2 + (\hat{\sigma}_{\alpha R} \times T)^2}$$

$$= 1.05 \times 10^3 \, \Omega \sqrt{(6 \times 10^{-4})^2 + (14 \times 10^{-6} \, °C^{-1} \times 14 \, °C)^2} \approx 0.66 \, \Omega$$

$$\sigma_{R'2} = r_2 \times \hat{\sigma}_{\epsilon R} = 10^4 \, \Omega \times 6 \times 10^{-4} = 6 \, \Omega$$

$$\sigma_{R'3} = r_3 \times \hat{\sigma}_{\epsilon R} = 10^4 \, \Omega \times 6 \times 10^{-4} = 6 \, \Omega \tag{B.37}$$

$$\sigma_{V0} = v_0 \sqrt{\hat{\sigma}_{\epsilon V0}^2 + (\hat{\sigma}_{\alpha V0} \times T)^2}$$

$$= 2.5 \, V \sqrt{(5 \times 10^{-4})^2 + (5 \times 10^{-6} \, °C^{-1} \times 14 \, °C)^2} \approx 1.3 \times 10^{-3} \, V$$

$$\sigma_G = g \sqrt{\hat{\sigma}_{\epsilon G}^2 + (\hat{\sigma}_{\alpha G} \times T)^2}$$

$$= 80 \sqrt{(6 \times 10^{-4})^2 + ((14 \times 10^{-6} \, °C^{-1} \times 14 \, °C)^2} \approx 5 \times 10^{-2}$$

Now we have everything needed to find σ_{Rx}. We substitute the partial derivatives and absolute errors in Equation B.20 with the values that have been calculated, and which are listed in B.23 and B.37. We get:

$$\sigma_{Rx} \approx \big[(1 \times 0.66 \, \Omega)^2 + (-0.105 \times 6 \, \Omega)^2 + (0.105 \times 6 \, \Omega)^2 + (-24.4 V_{ADC} \Omega \, V^{-1} \times 1.3 \times 10^{-3} \, V)^2$$
$$+ (-0.8 V_{ADC} \Omega \times 5 \times 10^{-2})^2 + (61 \, \Omega \, V^{-1} \times 2.7 \times 10^{-3} \, V)^2 \big]^{1/2}$$

$$\sigma_{Rx} \approx \sqrt{ \underset{R_1}{0.436 \, \Omega^2} + \underset{R'_2}{0.397 \, \Omega^2} + \underset{R'_3}{0.397 \, \Omega^2} + V_{ADC}{}^2 (\underset{V_0}{1} + \underset{G}{1.6} + \underset{V_{ADC}}{27}) \times 10^{-3} \, \Omega^2} \tag{B.38}$$

We are able to see which variables account for the largest error sources in R_x. The uncertainties in the three Wheatstone bridge resistors account for roughly the same amount, and dominates for small V_{ADC}. As it increases, the gain error starts to become non-negligible. We can estimate the greatest expected V_{ADC} by evaluating Equation A.22 with R_x having the value of the RTD element at $T = 25 \, °C + 14 \, °C$, and multiplying the result with the in amp gain. We then get $V_{ADC} \approx 2.74 \, V$, and the worst-case estimate for the uncertainty of R_x becomes:

$$\sigma_{Rx} \approx \sqrt{0.436 \, \Omega^2 + 0.397 \, \Omega^2 + 0.397 \, \Omega^2 + 2.74^2 (1 + 1.6 + 27) \times 10^{-3} \, \Omega^2}$$

$$\sigma_{Rx} \approx \sqrt{1.23 \, \Omega^2 + 7.5 \times 29.6 \times 10^{-3} \, \Omega^2}$$

$$\sigma_{Rx} \approx 1.2 \, \Omega \tag{B.39}$$

B.3.1   Uncertainty in temperature

Now that we have an estimation of the uncertainty in the measurement of the RTD, we turn our attention to the RTD itself in order to find the uncertainty in its temperature. The

equation that describes its resistance is given as:

$$R_x = r_0(1 + \alpha T) \tag{B.40}$$

And thus the following equation for its temperature:

$$T = \frac{R_x - r_0}{r_0 \alpha} \tag{B.41}$$

Here r_o is the resistance of the RTD at $0\,°C$ and α its temperature coefficient. We can now finally find an estimate for the uncertainty in the measured temperature by performing an error propagation on this equation. Following Equation B.10 as before, we get the σ_T, absolute uncertainty in T:

$$\sigma_T = \sqrt{\left(\frac{\partial T}{\partial R_x}\sigma_{Rx}\right)^2 + \left(\frac{\partial T}{\partial r_0}\sigma_{r0}\right)^2 + \left(\frac{\partial T}{\partial \alpha}\sigma_\alpha\right)^2} \tag{B.42}$$

$$\sigma_T = \sqrt{\left(\frac{1}{r_0\alpha}\sigma_{Rx}\right)^2 + \left(\frac{-R_x}{r_0^2\alpha}\sigma_{r0}\right)^2 + \left(\frac{r_0 - R_x}{r_0\alpha^2}\sigma_\alpha\right)^2} \tag{B.43}$$

From calibrating the RTDs we know the average values of r_0 and α to be $1000\,\Omega$ and $5650 \times 10^{-6}\,°C^{-1}$ (see Section 4.8.2). The absolute uncertainty in these variables, σ_{r0} and σ_α, are $68\,\%$ confidence bounds obtained from the fitting algorithm and are different for each RTD. By choosing the largest fit uncertainties among all the RTDs to represent σ_{r0} and σ_α in the error analysis we will overestimate σ_T for most of the RTDs, but that is preferable to underestimating it.

The largest values found for σ_{r0} and σ_α are $0.35\,\Omega$ and $2.6 \times 10^{-5}\,°C^{-1}$ respectively. Substituting these values as well as the average values for r_0 and α in Equation B.43 gives

$$\sigma_T = \sqrt{\left(\frac{1}{5.65\,\Omega\,°C^{-1}}1.2\,\Omega\right)^2 + \left(\frac{-R_x}{5650\,\Omega^2/°C}0.35\,\Omega\right)^2 + \left(\frac{1 \times 10^3\,\Omega - R_x}{31.9 \times 10^{-3}\,\Omega/°C^2}2.6 \times 10^{-5}\,°C^{-1}\right)^2}$$

$$\sigma_T = \sqrt{45.1 \times 10^{-3}\,°C^2 + R_x{}^2 \times 3.8 \times 10^{-9}\,°C^2/\Omega^2 + (1 \times 10^3\,\Omega - R_x)^2 6.6 \times 10^{-7}\,°C^2/\Omega^2}$$

$$\sigma_T = \sqrt{R_x{}^2 \times 6.64 \times 10^{-7}\,°C^2/\Omega^2 - R_x \times 1.32 \times 10^{-3}\,°C^2/\Omega + 0.705\,°C^2} \tag{B.44}$$

The uncertainty in T follows the square root of a parabola which has its minimum in $R_x \approx 993\,\Omega$, which corresponds to a temperature $T \approx -1°C$. Evaluating Equation B.44 for the largest expected temperature of $T = 39\,°C$, which corresponds to $R_x \approx 1220\,\Omega$, we get the final over-estimation of the uncertainty in the temperature measurements of the SiPM:

$$\sigma_T = 0.3\,°C \tag{B.45}$$

B.4   Calculations of partial derivatives

B.4.1   Temperature measurement

The calculations of the partial derivatives in Appendix B.2 are performed here. When a partial derivative has been found, R_T is set to 0 to simplify further. This is allowed because the parameters represents the ideal values with no uncertainty, and R_T consists only of uncertainty, as explained in subsection 4.8.3. We have the following equation for R_x:

$$R_x = \frac{R_3' + R_T}{\frac{R_2' + R_T}{R_1 + R_2' + R_T} - \frac{V_{ADC}}{V_0 G}} - R_3' + R_T \tag{B.46}$$

Partial derivative of R_x with respect to R_1

$$\frac{\partial R_x}{\partial R_1} = \frac{\partial}{\partial R_1}\left[\frac{R_3' + R_T}{\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G}} \right] - 0 = \frac{0 - \frac{\partial}{\partial R_1}\left[\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right](R_3' + R_T)}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2}$$

$$= \frac{-\left(\frac{-(R_2' + R_T)}{(R_2' + R_T + R_1)^2} - 0 \right)(R_3' + R_T)}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2} = \frac{(R_3' + R_T)(R_2' + R_T)}{\left(R_2' + R_T - (R_2' + R_T + R_1)\frac{V_{ADC}}{V_0 G} \right)^2}$$

Setting $R_T = 0$:

$$\frac{\partial R_x}{\partial R_1} = \frac{R_3' R_2'}{\left(R_2' - (R_2' + R_1)\frac{V_{ADC}}{V_0 G} \right)^2} \tag{B.47}$$

Partial derivative of R_x with respect to R_2'

$$\frac{\partial R_x}{\partial R_2'} = \frac{\partial}{\partial R_2'}\left[\frac{R_3' + R_T}{\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G}} \right] - 0 = \frac{0 - \frac{\partial}{\partial R_2'}\left[\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right](R_3' + R_T)}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2}$$

$$= \frac{-\left(\frac{-R_1}{(R_2' + R_T + R_1)^2} - 0 \right)(R_3' + R_T)}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2} = \frac{R_1(R_3' + R_T)}{\left(R_2' + R_T - (R_2' + R_T + R_1)\frac{V_{ADC}}{V_0 G} \right)^2}$$

Setting $R_T = 0$:

$$\frac{\partial R_x}{\partial R_2'} = \frac{R_1 R_3'}{\left(R_2' - (R_2' + R_1)\frac{V_{ADC}}{V_0 G} \right)^2} \tag{B.48}$$

Partial derivative of R_x with respect to R'_3

$$\frac{\partial R_x}{\partial R'_3} = \frac{\partial}{\partial R'_3}\left[\frac{R'_3 + R_T}{\frac{R'_2+R_T}{R_1+R'_2+R_T} - \frac{V_{ADC}}{V_0 G}} - R'_3 + R_T\right] = \frac{\partial}{\partial R'_3}\left[(R'_3 + R_T)\left(\frac{1}{\frac{R'_2+R_T}{R_1+R'_2+R_T} - \frac{V_{ADC}}{V_0 G}} - 1\right)\right]$$

$$= \frac{\partial}{\partial R'_3}\left[R'_3 + 1\right]\left(\frac{1}{\frac{R'_2+R_T}{R_1+R'_2+R_T} - \frac{V_{ADC}}{V_0 G}} - 1\right) + (R'_3 + 1)\frac{\partial}{\partial R'_3}\left[\frac{1}{\frac{R'_2+R_T}{R_1+R'_2+R_T} - \frac{V_{ADC}}{V_0 G}} - 1\right]$$

$$= (1+0)\left(\frac{1}{\frac{R'_2+R_T}{R_1+R'_2+R_T} - \frac{V_{ADC}}{V_0 G}} - 1\right) + (R'_3 + 1)(0 - 0) = \frac{1}{\frac{R'_2+R_T}{R_1+R'_2+R_T} - \frac{V_{ADC}}{V_0 G}} - 1$$

Setting $R_T = 0$:

$$\frac{\partial R_x}{\partial R'_3} = \frac{1}{\frac{R'_2}{R_1+R'_2} - \frac{V_{ADC}}{V_0 G}} - 1 \tag{B.49}$$

Partial derivative of R_x with respect to R_T

$$\frac{\partial R_x}{\partial R_T} = \frac{\partial}{\partial R_T}\left[\frac{R'_3 + R_T}{\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}}\right] - 1$$

$$= \frac{(0+1)\left(\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right) - (R'_3 + R_T)\frac{\partial}{\partial R_T}\left[\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right]}{\left(\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right)^2} - 1$$

$$= \frac{\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G} - \frac{R_1(R'_3+R_T)}{(R'_2+R_T+R_1)^2}}{\left(\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right)^2} - 1$$

$$= \frac{1}{\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}} - \frac{R_1(R'_3 + R_T)}{R'_2 + R_T - (R'_2 + R_T + R_1)\frac{V_{ADC}}{V0G}} - 1$$

Setting $R_T = 0$:

$$\frac{\partial R_x}{\partial R_T} = \frac{1}{\frac{R'_2}{R'_2+R_1} - \frac{V_{ADC}}{V_0 G}} - \frac{R_1 R'_3}{R'_2 - (R'_2 + R_1)\frac{V_{ADC}}{V0G}} - 1 \tag{B.50}$$

Partial derivative of R_x with respect to V_0

$$\frac{\partial R_x}{\partial V_0} = \frac{\partial}{\partial V_0}\left[\frac{R'_3 + R_T}{\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}}\right] - 0 = \frac{0 - \frac{\partial}{\partial V_0}\left[\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right](R'_3 + R_T)}{\left(\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right)^2}$$

$$= \frac{-(0 - (-\frac{V_{ADC}}{V_0^2 G}))(R'_3 + R_T)}{\left(\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right)^2} = -\frac{V_{ADC}}{V_0^2 G}\frac{R'_3 + R_T}{\left(\frac{R'_2+R_T}{R'_2+R_T+R_1} - \frac{V_{ADC}}{V_0 G}\right)^2}$$

Setting $R_T = 0$:

$$\frac{\partial R_x}{\partial V_0} = -\frac{V_{ADC}}{V_0^2 G} \frac{R_3'}{\left(\frac{R_2'}{R_2' + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2} \tag{B.51}$$

Partial derivative of R_x with respect to G

$$\frac{\partial R_x}{\partial G} = \frac{\partial}{\partial G} \left[\frac{R_3' + R_T}{\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G}} \right] - 0 = \frac{0 - \frac{\partial}{\partial G} \left[\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right] (R_3' + R_T)}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2}$$

$$= \frac{-(R_3' + R_T)(0 - (-\frac{V_{ADC}}{V_0 G^2}))}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2} = -\frac{V_{ADC}}{V_0 G^2} \frac{R_3' + R_T}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2}$$

Setting $R_T = 0$:

$$\frac{\partial R_x}{\partial G} = -\frac{V_{ADC}}{V_0 G^2} \frac{R_3'}{\left(\frac{R_2'}{R_2' + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2} \tag{B.52}$$

Partial derivative of R_x with respect to V_{ADC}

$$\frac{\partial R_x}{\partial V_{ADC}} = \frac{\partial}{\partial V_{ADC}} \left[\frac{R_3' + R_T}{\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G}} \right] - 0 = \frac{0 - \frac{\partial}{\partial V_{ADC}} \left[\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right] (R_3' + R_T)}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2}$$

$$= \frac{-(R_3' + R_T)(0 - \frac{1}{V_0 G})}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2} = \frac{1}{V_0 G} \frac{R_3' + R_T}{\left(\frac{R_2' + R_T}{R_2' + R_T + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2}$$

Setting $R_T = 0$:

$$\frac{\partial R_x}{\partial V_{ADC}} = \frac{1}{V_0 G} \frac{R_3'}{\left(\frac{R_2'}{R_2' + R_1} - \frac{V_{ADC}}{V_0 G} \right)^2} \tag{B.53}$$

B.5   Results from RTD temperature sensor calibrations

Here follows the results from the calibrations of the RTD sensors used for sensing the SiPM temperature. The calibration method is discussed in Section 4.8.3.

Standard deviation all temperature measurements: ±2mK							
Standard deviation resistance measurements at 15C and 35C: ±16mΩ							
Standard deviation resistance measurements at 25C: ±25mΩ							
RTD 0		**RTD 1**		**RTD 2**		**RTD 3**	
T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]
15,078	1084,97	15,075	1086,08	15,077	1085,68	15,081	1085,25
25,030	1141,37	25,028	1142,43	25,032	1142,09	25,026	1141,71
34,993	1197,54	34,990	1198,65	34,987	1198,33	34,998	1198,14
RTD 4		**RTD 5**		**RTD 6**		**RTD 7**	
T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]
15,078	1085,20	15,078	1084,97	15,075	1085,95	15,080	1085,52
25,028	1141,74	25,026	1141,43	25,028	1142,02	25,027	1141,85
34,991	1197,97	34,992	1197,81	34,990	1197,92	34,988	1198,11
RTD 8		**RTD 9**		**RTD 10**		**RTD 11**	
T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]
15,079	1085,47	15,075	1088,17	15,082	1085,45	15,078	1085,07
25,028	1141,89	25,028	1144,48	25,025	1141,84	25,026	1141,51
34,992	1198,05	34,994	1200,59	34,992	1198,13	34,990	1197,85
RTD 12		**RTD 13**		**RTD 14**		**RTD 15**	
T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]	T[°C]	R[Ω]
15,075	1084,98	15,077	1084,96	15,063	1085,12	15,080	1085,49
25,028	1141,40	25,027	1141,31	25,027	1141,58	25,027	1141,92
34,986	1197,30	34,991	1197,66	34,989	1197,92	34,995	1198,20

Figure B.1: Result of the measurements performed on each RTD for the calibration process. All values are averages over 2-3 minutes, with 120-180 measurements performed. The given standard deviations are the highest recorded for all RTDs at each temperature.

RTD	R_0 [Ω]	α [$^{\circ}C^{-1}$]x10^{-6}	RTD	R_0 [Ω]	α [$^{\circ}C^{-1}$]x10^{-6}
0	999,9	5650	8	1000,4	5650
1	1001,0	5640	9	1003,3	5620
2	1000,5	5650	10	1000,3	5650
3	999,9	5660	11	999,8	5660
4	1000,4	5650	12	1000,2	5640
5	999,7	5660	13	999,8	5660
6	1001,4	5610	14	1000	5660
7	1000,4	5650	15	1000,3	5650

Figure B.2: Results of performing a linear fit on the data in Table B.1 using the equation $R(T) = R_0(1 + \alpha T)$.
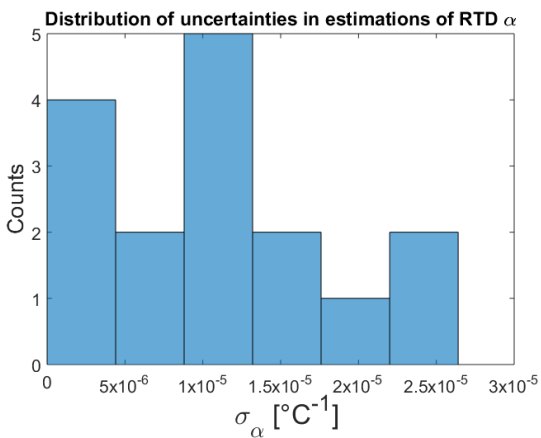

Figure B.3: Distribution of the uncertainties related to the estimation of RTD temperature dependency α for all RTDs. The uncertainties represent 68 % confidence intervals. The largest uncertainty is $\sim 2.6 \times 10^{-5}\,^{\circ}C^{-1}$
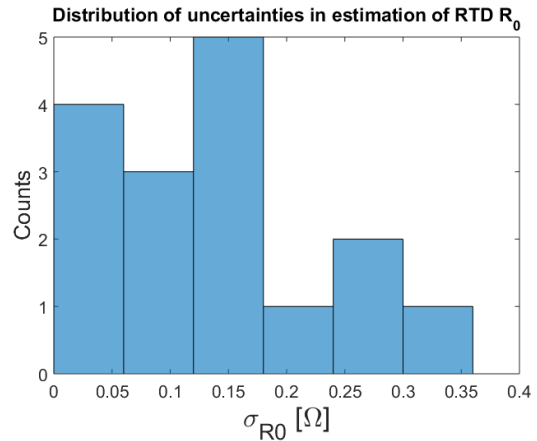

Figure B.4: Distribution of the uncertainties related to the estimation of RTD base resistance R_0 for all RTDs. The uncertainties represent 68 % confidence intervals. The largest uncertainty is $\sim 0.35\,\Omega$

B.6 Estimation of uncertainty in trigger threshold

The discriminator trigger threshold value is set by the circuit seen in Figure B.5. The trigger threshold V_{tt} is given by the following equation:

$$V_{tt} = (V_{DAC} \times K_1 \times B_G + B_{ofs}) \times K_2 - V_{DC} \tag{B.54}$$

Where V_{DAC} represents the voltage generated by the MCU DAC, K_1 represents the division factor of the voltage divider before the buffer, B_G represents the gain of the buffer, B_{ofs} represents the offset of the buffer, K_2 represents the division factor of the voltage divider after the buffer, and V_{DC} represents the DC component of the MPPC signal.

The absolute uncertainty in V_{tt} is given as

$$\sigma_{tt} = \left[\left(\frac{\partial V_{tt}}{\partial V_{DAC}} \sigma_{DAC} \right)^2 + \left(\frac{\partial V_{tt}}{\partial K_1} \sigma_{K1} \right)^2 + \left(\frac{\partial V_{tt}}{\partial K_2} \sigma_{K2} \right)^2 + \left(\frac{\partial V_{tt}}{\partial B_G} \sigma_{Bg} \right)^2 \right.$$
$$\left. + \left(\frac{\partial V_{tt}}{\partial B_{ofs}} \sigma_{Bo} \right)^2 + \left(\frac{\partial V_{tt}}{\partial V_{DC}} \sigma_{DC} \right)^2 \right]^{1/2} \tag{B.55}$$

Solving the partial derivatives gives

$$\sigma_{tt} = \left[(K_1 B_G K_2 \sigma_{DAC})^2 + (V_{DAC} B_G K_2 \sigma_{K1})^2 + ((V_{DAC} K_1 B_G + B_{ofs}) \sigma_{K2})^2 \right.$$
$$\left. + \left(V_{DAC} K_1 K_2 \sigma_{Bg} \right)^2 + (K_2 \sigma_{Bo})^2 + (-\sigma_{DC})^2 \right]^{1/2} \tag{B.56}$$

And then we substitute for known constant values

$$\sigma_{tt} \approx \left[(0.19723 \sigma_{DAC})^2 + (0.99 V_{DAC} \times \sigma_{K1})^2 + (0.19920 V_{DAC} \times \sigma_{K2})^2 \right.$$
$$\left. + \left(0.19723 V_{DAC} \times \sigma_{Bg} \right)^2 + (0.99001 \sigma_{Bo})^2 + (-\sigma_{DC})^2 \right] \tag{B.57}$$

We must now find the values of the uncertainties σ_{DAC}, σ_{K1}, etc.. Let us first consider the voltage dividers K_1 and K_2. They are made with resistors R_0, R_1, R_2, and R_3, which all have the same manufacturing error $\epsilon_R = \pm 1\%$ and TCR $\alpha = \pm 1 \times 10^{-4}\,°C^{-1}$. We can ignore the manufacturing uncertainty, as the actual division factors of K_1 and K_2 were measured with a precision voltmeter after assembly and are compensated for in the MCU software. This leaves the temperature dependencies. We can assume that the dependencies add since the division factor of K_1 is relatively far from the middle point, and very far from the middle point for K_2. The temperature coefficients of K_1 and K_2 are therefore $\alpha_K = \pm 2 \times 10^{-4}\,°C^{-1}$. Assuming an uniform random variable for the TCR probability distribution (as in Appendix B.3) and $T = 25\,°C + 14\,°C$ we get a worst-case 68% confidence interval on the relative uncertainty in the division factors of K_1 and K_2 to be $\hat{\sigma}_K = \pm 0.16\%$, and the absolute errors $\sigma_{K1} =$
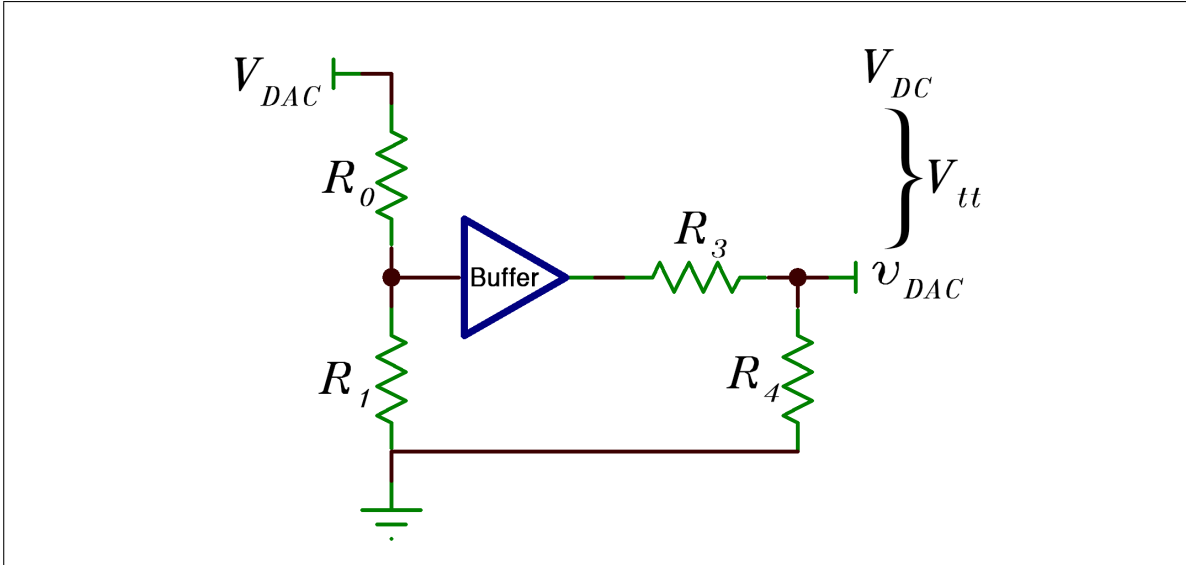
Figure B.5: The circuit that sets the discriminator trigger threshold V_{tt}. V_{DAC} is generated by the microcontroller, and is divided by 5 by R_0 and R_1. The buffer is needed in order to meet the input bias current requirements of the discriminator. R_2 and R_3 slightly scales V_{DAC} by 0.99. The trigger threshold is the scaled V_{DAC} signal, v_{DAC}, minus the MPPC signal DC component V_{DC}.

$0.16\,\% \times K_1 \approx 3.187 \times 10^{-4}$ and $\sigma_{K2} = 0.16\,\% \times K_2 \approx 1.584 \times 10^{-3}$.

The voltage generated by the MCU DAC has an uncertainty consisting of the uncertainty of the MCU 5 V voltage regulator, and an uncertainty related to the finite resolution of the DAC. We have from the datasheet of the voltage regulator [73] that its manufacturing uncertainty is $\pm 1\,\%$, and that its temperature dependency is $\pm 4 \times 10^{-5}\,°C^{-1}$ which results in a maximum relative uncertainty of $\pm 0.06\,\%$ in the expected temperature range of the ROB. We ignore the temperature dependency, as it is negligible compared to the manufacturing uncertainty. The resolution of the MCU DAC is 10 bits, which results in a relative resolution of $1/2^{10} = 1/1024 \approx 0.1\,\%$, and so we ignore this too. The total relative uncertainty in V_{DAC} is therefore $\hat{\sigma}_{DAC} = \pm 1\,\%$, resulting in an absolute error of $\sigma_{DAC} = \pm 0.01 V_{DAC}$.

The input voltage offset B_{ofs} of the buffer is found from the datasheet [74] to follow a gaussian distribution with a standard deviation of roughly $\sigma_{Bo} = \pm 1\,mV$. As it is operated in a voltage buffer configuration, there is no gain error, and $\sigma_{Bg} = 0$.

The uncertainty in V_{DC} is difficult to estimate correctly. Its value for all ROBs was measured at $21\,°C \pm 3\,°C$, and has been observed to vary with $\pm 1\,mV$ over time. An estimated guess sets its uncertainty at $\sigma_{DC} = \pm 2\,mV$.

Substituting for the uncertainties in Equation B.57 results in:

$$\sigma_{tt} \approx \Big[(0.19723 \times 0.01 V_{DAC})^2 + \big(0.99 V_{DAC} \times 3.187 \times 10^{-4}\big)^2 + \big(0.19920 V_{DAC} \times 1.584 \times 10^{-3}\big)^2$$

$$+ (0.99001 \times 1\,mV)^2 + (-2\,mV)^2 \Big] \quad (B.58)$$

Cleaning up and collecting terms gives

$$\sigma_{tt} \approx \sqrt{4 \times 10^{-6} \times V_{DAC}^2 + 4.98\,\text{mV}^2} \tag{B.59}$$

Evaluating Equation B.59 shows that σ_{tt} has its global minima at $V_{DAC} = 0\,\text{mV}$, where its value is $\sigma_{tt} \approx 2.2\,\text{mV}$. As V_{DAC} increases, it begins to dominate the uncertainty and σ_{tt} increases approximately linearly from $V_{DAC} = 1\,\text{V}$ and upwards.

Since the DC component of the MPPC signal is around $250\,\text{mV}$, V_{DAC} will not be set to lower values than $\sim 250\,\text{mV}/K_1 \approx 1.25\,\text{V}$, and the uncertainty can be assumed to be a linear function:

$$\sigma_{tt} \approx 2 \times 10^{-3} V_{DAC} \quad \text{given} \quad V_{DAC} > 1\,\text{V} \tag{B.60}$$

The equation can be rewritten by assuming that $V_{DC} = 250\,\text{mV}$, and using the fact that $V_{tt} \approx V_{DAC}/5 - V_{DC}$:

$$\sigma_{tt} = 2 \times 10^{-3} \times 5(V_{tt} + V_{DC})\sigma_{tt} \quad = 2 \times 10^{-3} \times 5(V_{tt} + 250\,\text{mV})\sigma_{tt} = V_{tt} \times 10^{-2} + 2.5\,\text{mV} \tag{B.61}$$

C | Communications protocol

In this section of the appendix the communication protocol for the RS-485 network, and the communication protocol between the RS-485 master and the PC, is explained in more detail. Data in both of these protocols are sent as well-defined packets. These packets are sent from the master to the ROBs, from the master to the PC, and vice versa. The data packets can contain commands, variables, error flags, and more. See Section 4.10.3 for more details on the protocol. This appendix chapter will only provide details on which commands can be sent, and how the data packets are structured.

C.1   Data packets on the RS-485 network

All bytes on the RS-485 network are sent in double inverted form. If the byte 0xAB is to be sent, it is first split into two bytes, 0xA0 and 0xB0. The first nibble in each byte is then changed to the inverted of the original nibble, and the two bytes becomes 0xA5 and 0xB4. For the remainder of this section, this conversion is implicit and ignored. When we talk of bytes in a data packet, it is the original byte, not the double inverted forms.

A data packet begins with a STX byte 0x02 and ends with an ETX byte 0x03. After the ETX byte comes the CRC byte that will be compared to the CRC byte generated by the receiver. If they are equal, there was likely no errors in transmission.

C.1.1   Data packets from Master to ROB

A data packet from the Master to a ROB must contain the following bytes, ignoring ETX, STX, and CRC bytes, and this order:

- Master identifier byte
- ROB address byte
- Command byte
- Value bytes (optional, used when setting bias voltage and other variables)

The master identifier byte is 0xCB. This identifier byte will be checked by a ROB on reception of a data packet, and if it is not correct, the packet is ignored. This allows packets from ROBs and packets from the master to be distinguished. The ROB address byte identifies which ROB the data packet is intended for. It must be equal to the ROB ID, and goes from 0 (0x00) to 15 (0x0F). Finally, the command byte signifies what the ROB is requested to do. They are:

- 0xA0: Report MPPC temperature. Measures the MPPC temperature an returns it.
- 0xA1: Report all measured board voltages. Measures all the voltages the MCU monitors, and returns their values.
- 0xA2: Report status. Returns the status (nominal, warning, error) and enabled/disabled state of all measured voltages, as well as the status of the latch counter and MPPC temperature, and the current board mode (standby, forced standby, ready, running).
- 0xA3: Report latch counter. Returns the current size of the latch reset array.
- 0xA4: Report trigger threshold. Returns the value of the trigger threshold.
- 0xB0: Set MPPC bias voltage. Sets the MPPC bias voltage to the value specified by the Value bytes, given as a four-byte float. Automatic bias adjustment must first have been disabled with command 0xF5.
- 0xB1: Set trigger threshold. Sets the trigger threshold to the value specified by the Value bytes, given as a four-byte float.
- 0xB1: Set MPPC bias overvoltage. Sets the bias overvoltage to the value specified by the Value bytes, given as a four-byte float. The automatic bias adjust algorithm will use this value when adjusting the MPPC bias.

- `0xC0`: Set the board in mode Standby.
- `0xC1`: Set the board in mode Ready.
- `0xC2`: Set the board in mode Running.
- `0xD0`: Connect or disconnect the MPPC from the amplifier. Mode set by Value byte, given as a boolean. 0: Disconnect, 1: Connect.
- `0xD1`: Insert a single test charge.
- `0xD2`: Enable or disable continuous test charge insertion. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.

The following commands are for debugging:

- `0xF0`: Echo. Replies back the to the Master with the Value bytes it receives. Used for testing if the ROB is working and is able to receive commands.
- `0xF1`: Latch reset. Forces a reset of the trigger latch.
- `0xF2`: Enable or disable the ±5 V rails. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.
- `0xF3`: Enable or disable the 3 V rail. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.
- `0xF9`: Enable or disable the 15 V rail. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.
- `0xF4`: Set the voltage reference to the default internal or the external precision reference. Mode set by Value byte, given as an integer. 0: Internal, 1: External.
- `0xF5`: Enable or disable automatic MPPC bias voltage adjustment. It must be disabled if one wish to manually set a fixed bias voltage with command `0xB0`. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.
- `0xF6`: Enable or disable the amplifier. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.
- `0xF7`: Set debug verbosity. Verbosity level set by Value byte, given as an integer from 0 to 7.
- `0xF8`: Enable or disable the automatic latch reset routine. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.
- `0xFA`: Enable or disable rail monitor, which will force the board to STANDBY if a rail is out of bounds. Mode set by Value byte, given as a boolean. 0: Disable, 1: Enable.

C.1.2 Data packets from ROB to Master

When a ROB has performed the actions requested of it by the Command byte in a data packet received by the master, it will reply back with the results. A data packet being sent from a ROB to the Master must contain the following, ignoring ETX, STX, and CRC bytes, and this order:

- ROB identifier byte
- ROB ID byte

- "Command recognized" byte
- Value bytes

The ROB identifier byte is `0xDB`, and marks the data packet as originating from a ROB, preventing it from being interpreted by other ROBs. The ROB ID byte is the ID of the ROB sending the data packet. The "Command recognized" byte is `0xAA` if the ROB recognized the command, and `0xFF` if it did not.

The value bytes contains the information that the ROB is expected to return after interpreting a specific command. For example, if it received a command to report the MPPC temperature, the Value bytes will be four consecutive bytes representing a 32-bit floating point value. It is also used to report back the result of executing certain commands that may fail, such as setting board modes or inserting test charges. The Master and ROB must agree on what is expected to be returned for every command. If a command does not require the ROB to send any Value bytes, the ROB will send a single `0x00` Value byte, due to the way the command-sending function is written.

C.2 Data packets between the Master and PC

Communication between the operator, using a PC, and the ROBs on the RS485-network is mediated by the Master. The operator can send data packets to the Master through a USB port, which it will interpret and execute. The Master will send data packets back to the PC in order to inform the user about the results of the command execution, and in order to mediate the data that was returned by the ROB.

C.2.1 Data packets from PC to Master

The data packets sent from the PC to the Master by the user are relatively simple. No STX, ETX, or CRC bytes are needed. The data packets must consist of the following, and this order:

- Command byte
- ROB address byte
- Value bytes

The command bytes are equal to the list of commands in Subsection C.1.1, with the addition of the command `0xAA`. This special command tells the Master to simply ship the rest of the data packet onto the RS485 network, and is used for debugging.

The ROB address byte identifies which ROB the Master should send the command to. It goes from 0 (`0x00`) to 15 (`0x0F`), with the addition of `0xFF`. This address byte means "all ROBs" and causes the Master to send the command to every ROB in succession. Not all commands have this capability, especially the debug commands, but most of the commonly used commands do support it.

The Value bytes will be sent to the ROBs with no modifications, and are those discussed in Appendix C.1.1.

C.2.2 Data packets from Master to PC

The data stream from the Master to the PC consists of a human-readable text interspersed with data packets for the CRTC. The data packets contain information about the execution of the command received by the PC by the Master, and the results of the subsequent execution of those commands by the ROBs.

A data packet from the Master to the PC begins with an STX byte (`0x02`) and ends with an ETX byte (`0x03`), which allows the CRTC to sort them out from the data stream. Within are the following, and in this order:

- Data ID byte
- ROB ID byte

- Data Type byte
- Data bytes

Data ID Bytes

The Data ID byte identifies what kind of data is present in the Data bytes. All the different types of data that can be returned to the PC has its own unique Data ID byte. This enables the CRTC to handle the data and display it on its front panel. The Data ID bytes are as follows:

- 0x00: MPPC temperature.
- 0x01: Positive 5 V rail value.
- 0x02: Negative 5 V rail value.
- 0x21: 3.3 V rail value.
- 0x03: 15 V rail value.
- 0x05: MPPC bias voltage value.
- 0x06: Trigger threshold value.
- 0x07: Number of latch resets.
- 0x08: Board mode.
- 0x0A: Positive 5 V status and enabled/disabled.
- 0x0B: Negative 5 V status and enabled/disabled.
- 0x0C: 3.3 V status and enabled/disabled.
- 0x0D: 15 V status and enabled/disabled.
- 0x0F: MPPC bias status and enabled/disabled.
- 0x10: MPPC temperature status.
- 0x11: Latch reset status.
- 0x12: MPPC connected/disconnected.
- 0x13: Continuous test charge enabled/disabled.
- 0x14: Voltage reference internal or external.
- 0xFE: Error identifier.
- 0xFF: Transmission result identifier.

Data Type Bytes

The Data Type byte signifies what format the Data bytes has, enabling the CRTC to correctly typecast them. Only four types of data can be transmitted from the Master to the PC, identified by the following Data Type bytes:

- 0x00: 32bit signed integer.
- 0x01: 32bit signed float.
- 0x02: Boolean.
- 0x03: Byte array of arbitrary length.

Data Bytes   The Data bytes contain the actual data, and are either of the following:

1. Data returned from the ROBs.

2. Error or Transmission result bytes.

The Master will send data packets to the PC during its execution of the command it has received, using the Error and Transmission result identifier bytes. These data packets lets the user know if something went wrong during the execution and transmission of the command, either by the ROB or the Master. Each possible error and transmission result has a unique byte value, which are sent as single-byte values together with the Data ID 0xFE or 0xFF, and the Data Type 0x03, as listed earlier.

The Error bytes are as follows:

- 0x00: Data packet received from the PC exceeded the maximum length of 12 bytes.
- 0x01: The Value bytes in a data packet from the PC has an incorrect format, and could not be typecast.
- 0x02: Error reported by ROB when setting MPPC bias voltage.
- 0x03: Error reported by ROB when when setting trigger threshold.
- 0x04: Error reported by ROB when when attempting to enter board mode Standby.
- 0x05: Error reported by ROB when when attempting to enter board mode Ready.
- 0x06: Error reported by ROB when when attempting to enter board mode Running.
- 0x07: Error when attempting to connect or disconnect the MPPC to the amplifier.
- 0x08: Error reported by ROB when when attempting to insert test charge.
- 0x09: Error reported by ROB when when attempting to enable or disable continuous test charge.
- 0x0A: Error reported by ROB when when attempting to enable or disable a voltage rail.

The second element in the byte array will often contain additional information, such as the ID of the ROB which reported the error. The negation of these error bytes are created by adding 0x70. So if no error occurred when setting the MPPC bias voltage, the byte 0x70+0x02 = 72 is sent instead.

The transmission result bytes are as follows:

- 0x00: Master recognized command from PC.
- 0x01: Master did not recognize command from PC.
- 0x02: The correct ROB acknowledged data packet.
- 0x03: An incorrect ROB acknowledged the data packet.
- 0x04: Master did not receive an acknowledge-reply from any ROB.
- 0x05: ROB did not acknowledge data packet.
- 0x06: ROB did not return data as expected after command execution.
- 0x07: ROB returned data after command execution, but did not begin with 0xDB.
- 0x08: An incorrect ROB returned data while waiting for command execution result.
- 0x09: The size of the data packet returned by a ROB was of the incorrect expected size.
- 0x0A: ECHO OK.
- 0x0B: ECHO not OK.
- 0x0C: Master is ready for new commands from the PC.

D | Additional data

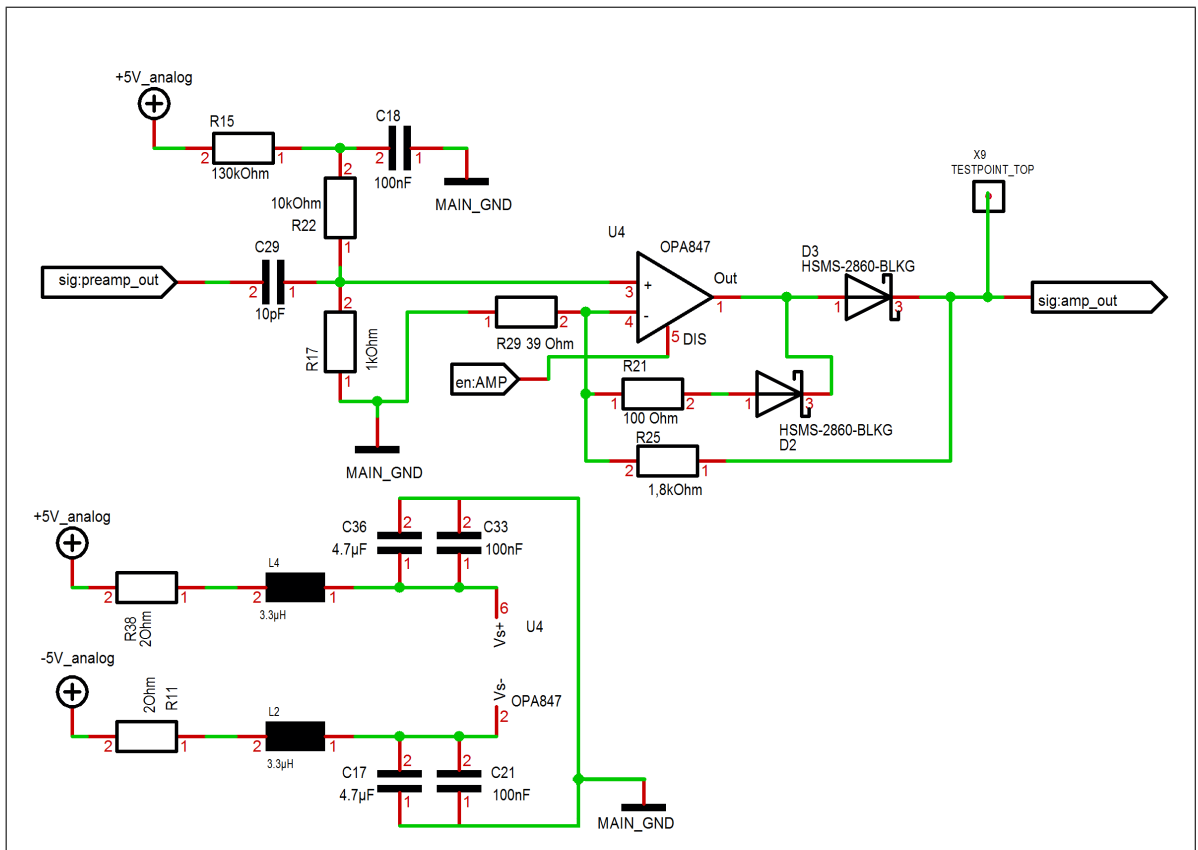This appendix chapter will list figures, tables, and other forms of data that have been referenced in the main text.

D.1 ROB circuit diagrams

Figure D.1: Circuit diagram of the ROB preamplifier

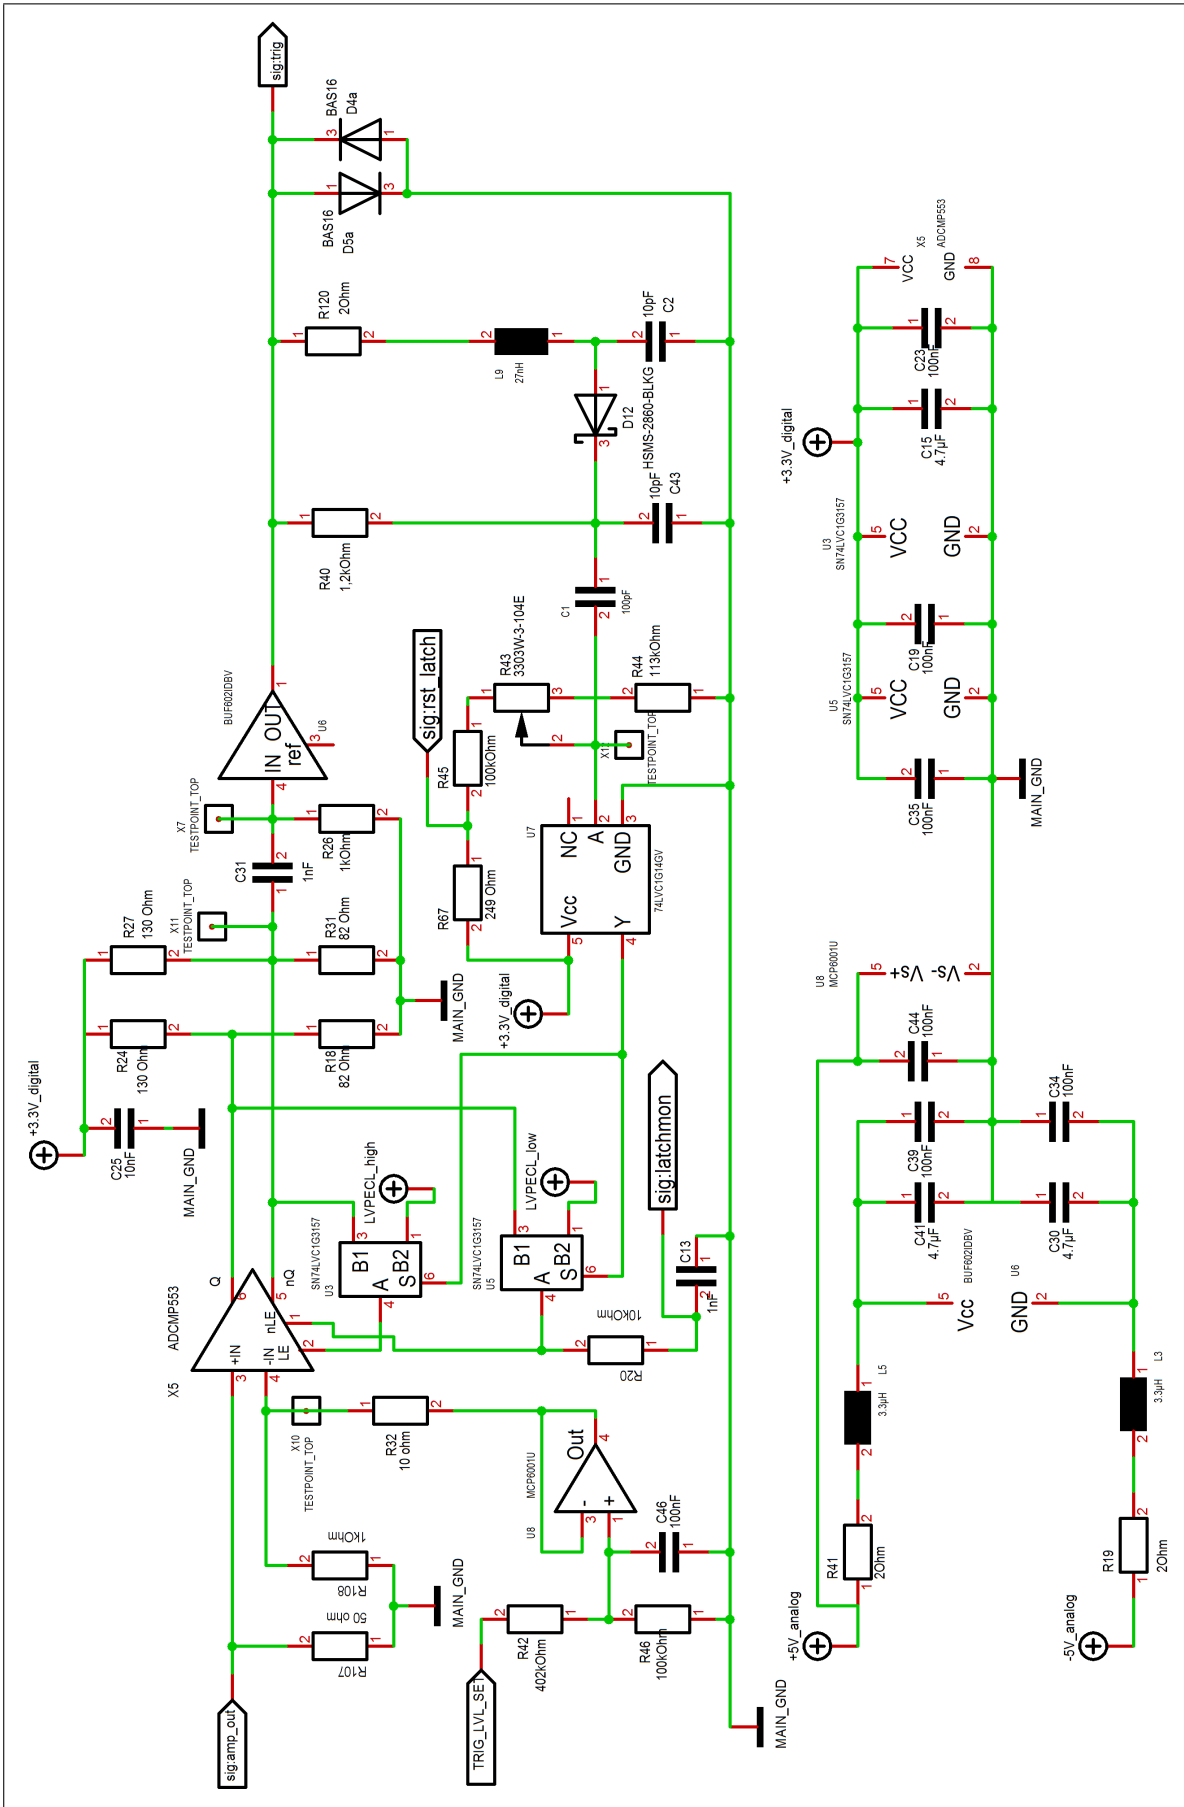**Figure D.2:** Circuit diagram of the ROB primary amplifier

Figure D.3: Circuit diagram of the ROB MPPC pulse discriminator

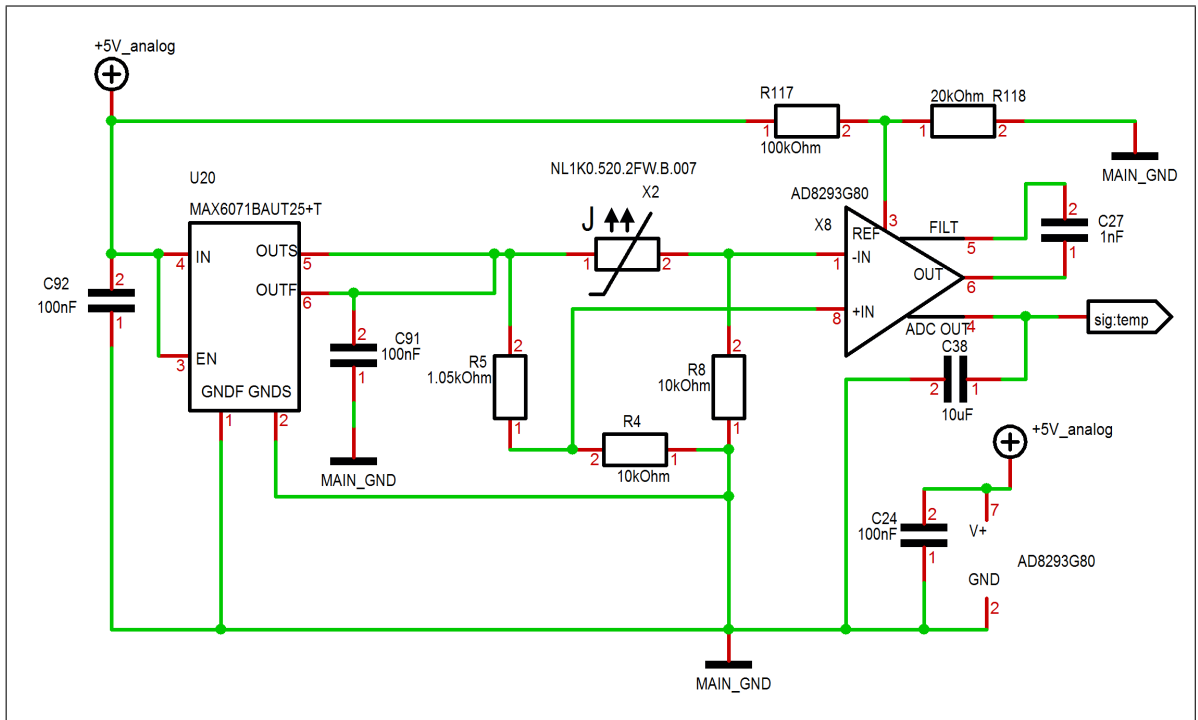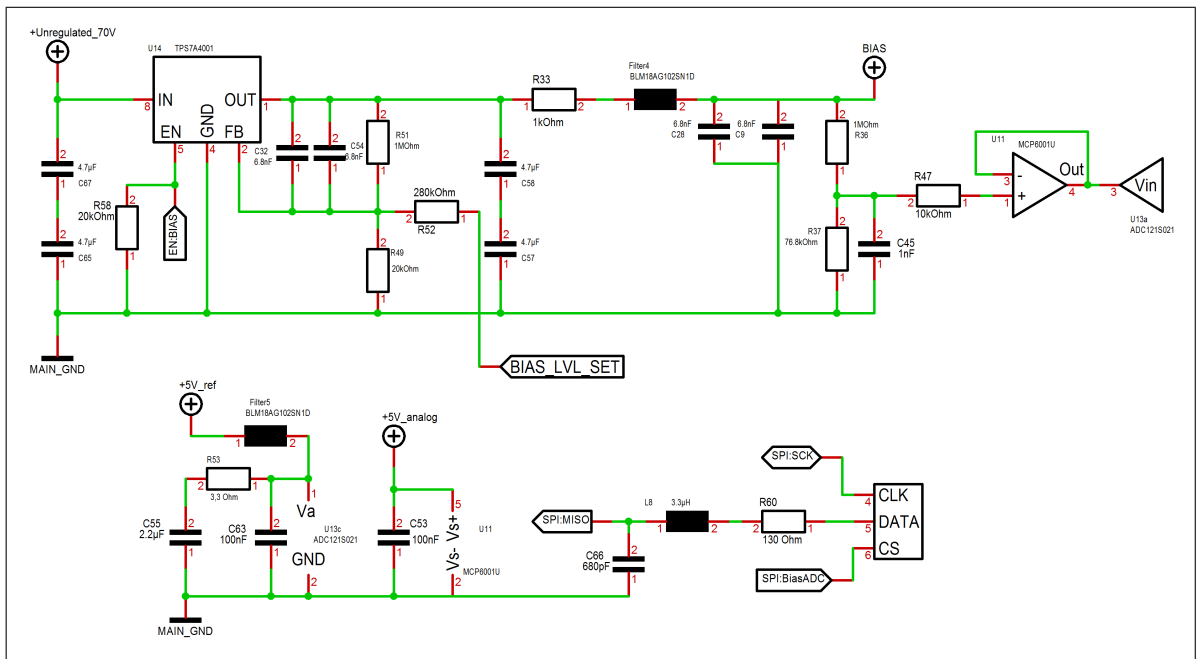**Figure D.4:** Circuit diagram of the ROB MPPC temperature sensing circuit

Figure D.5: Circuit diagram of the ROB MPPC bias voltage regulator circuit
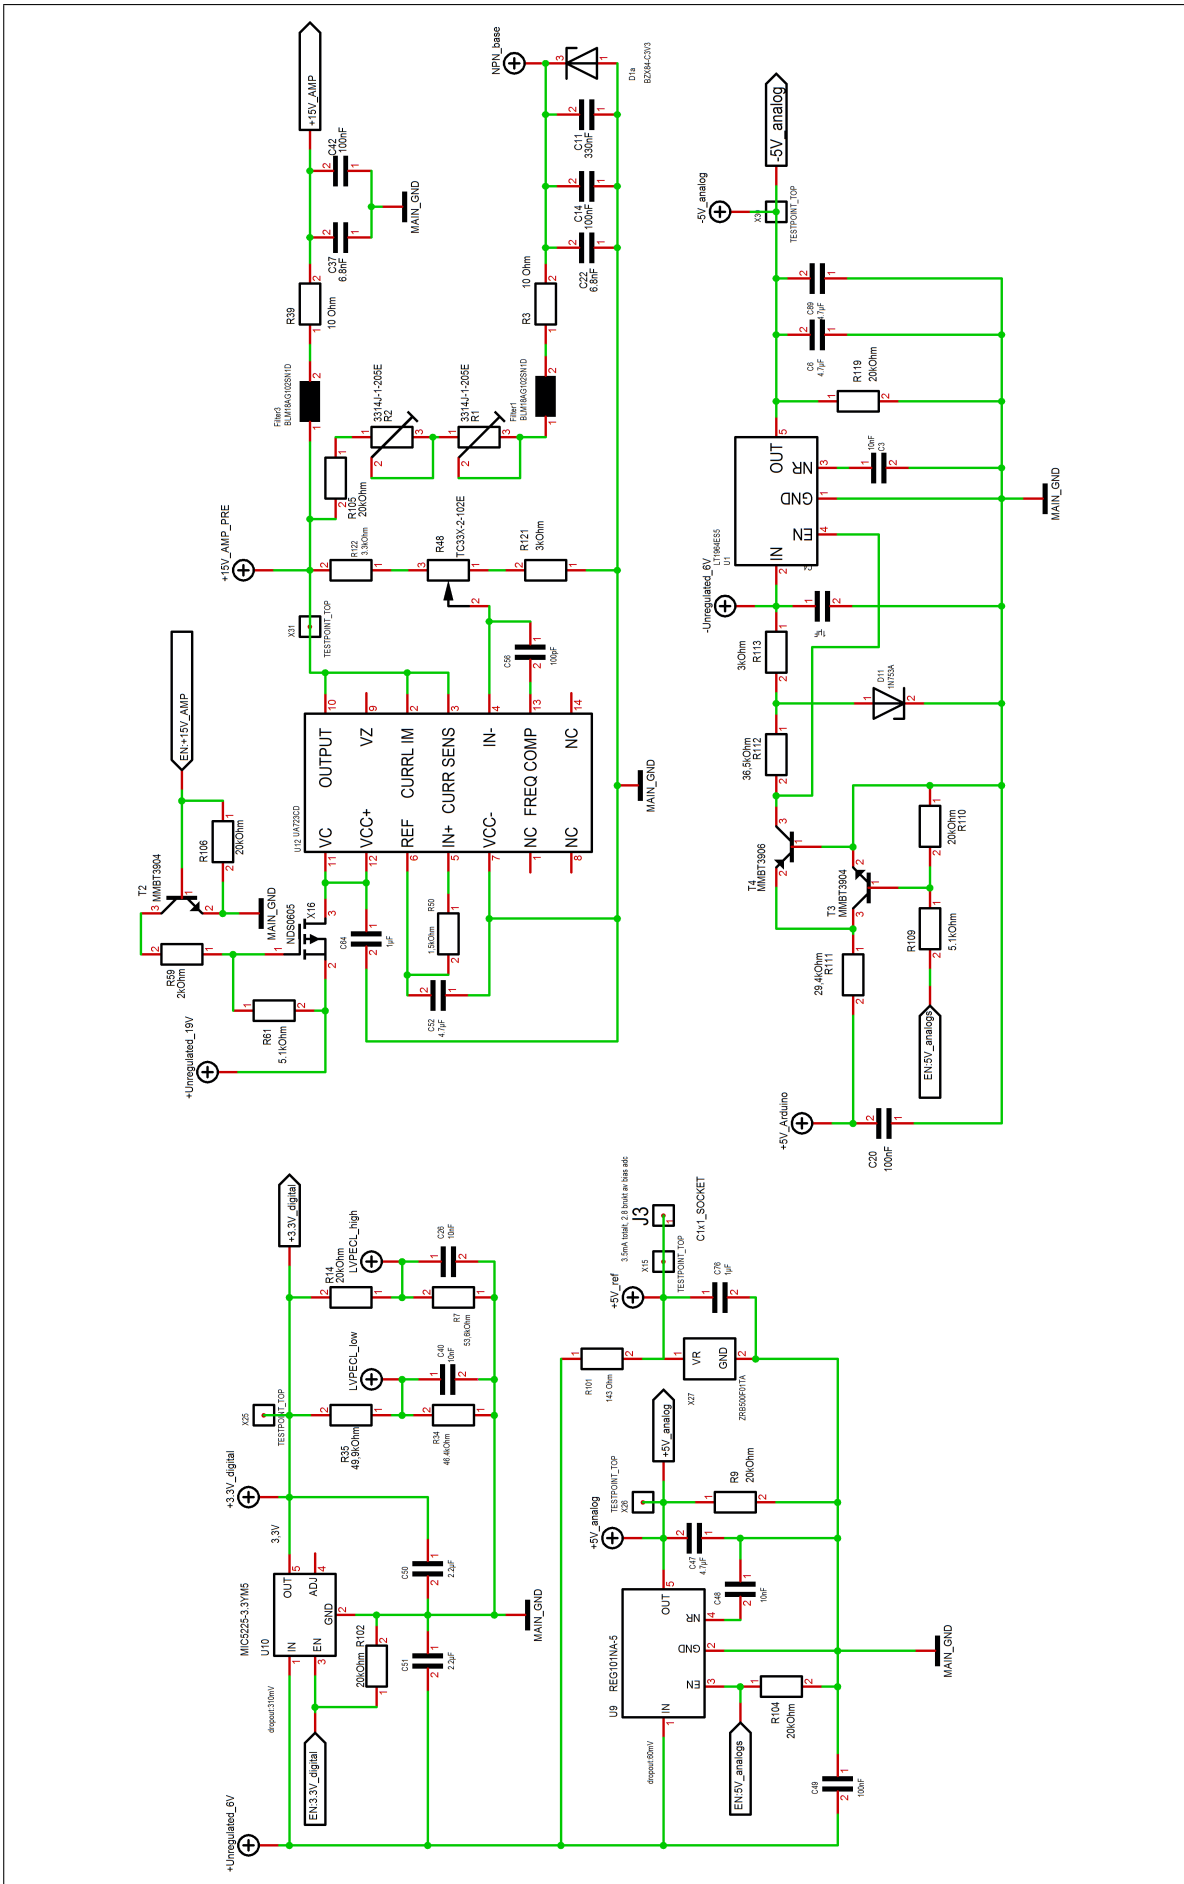
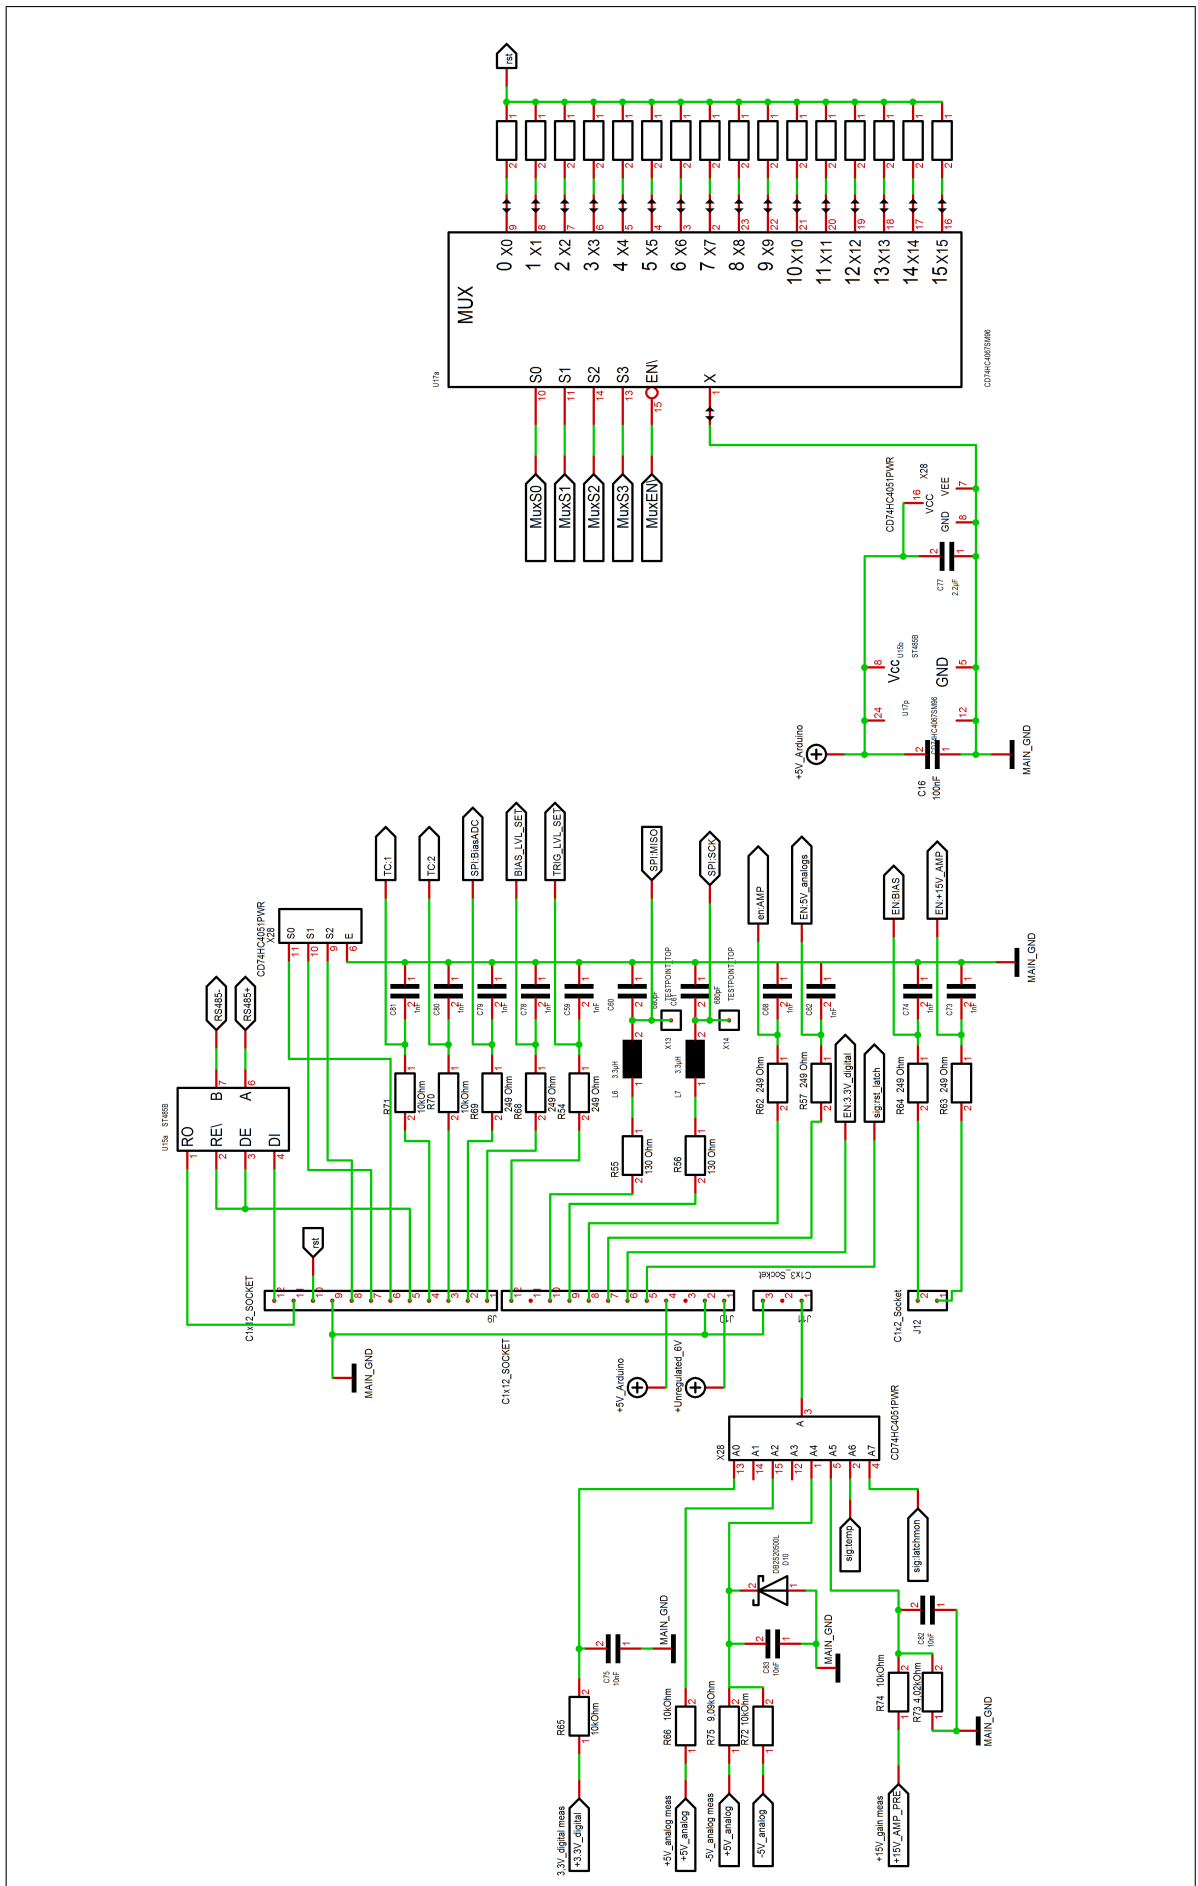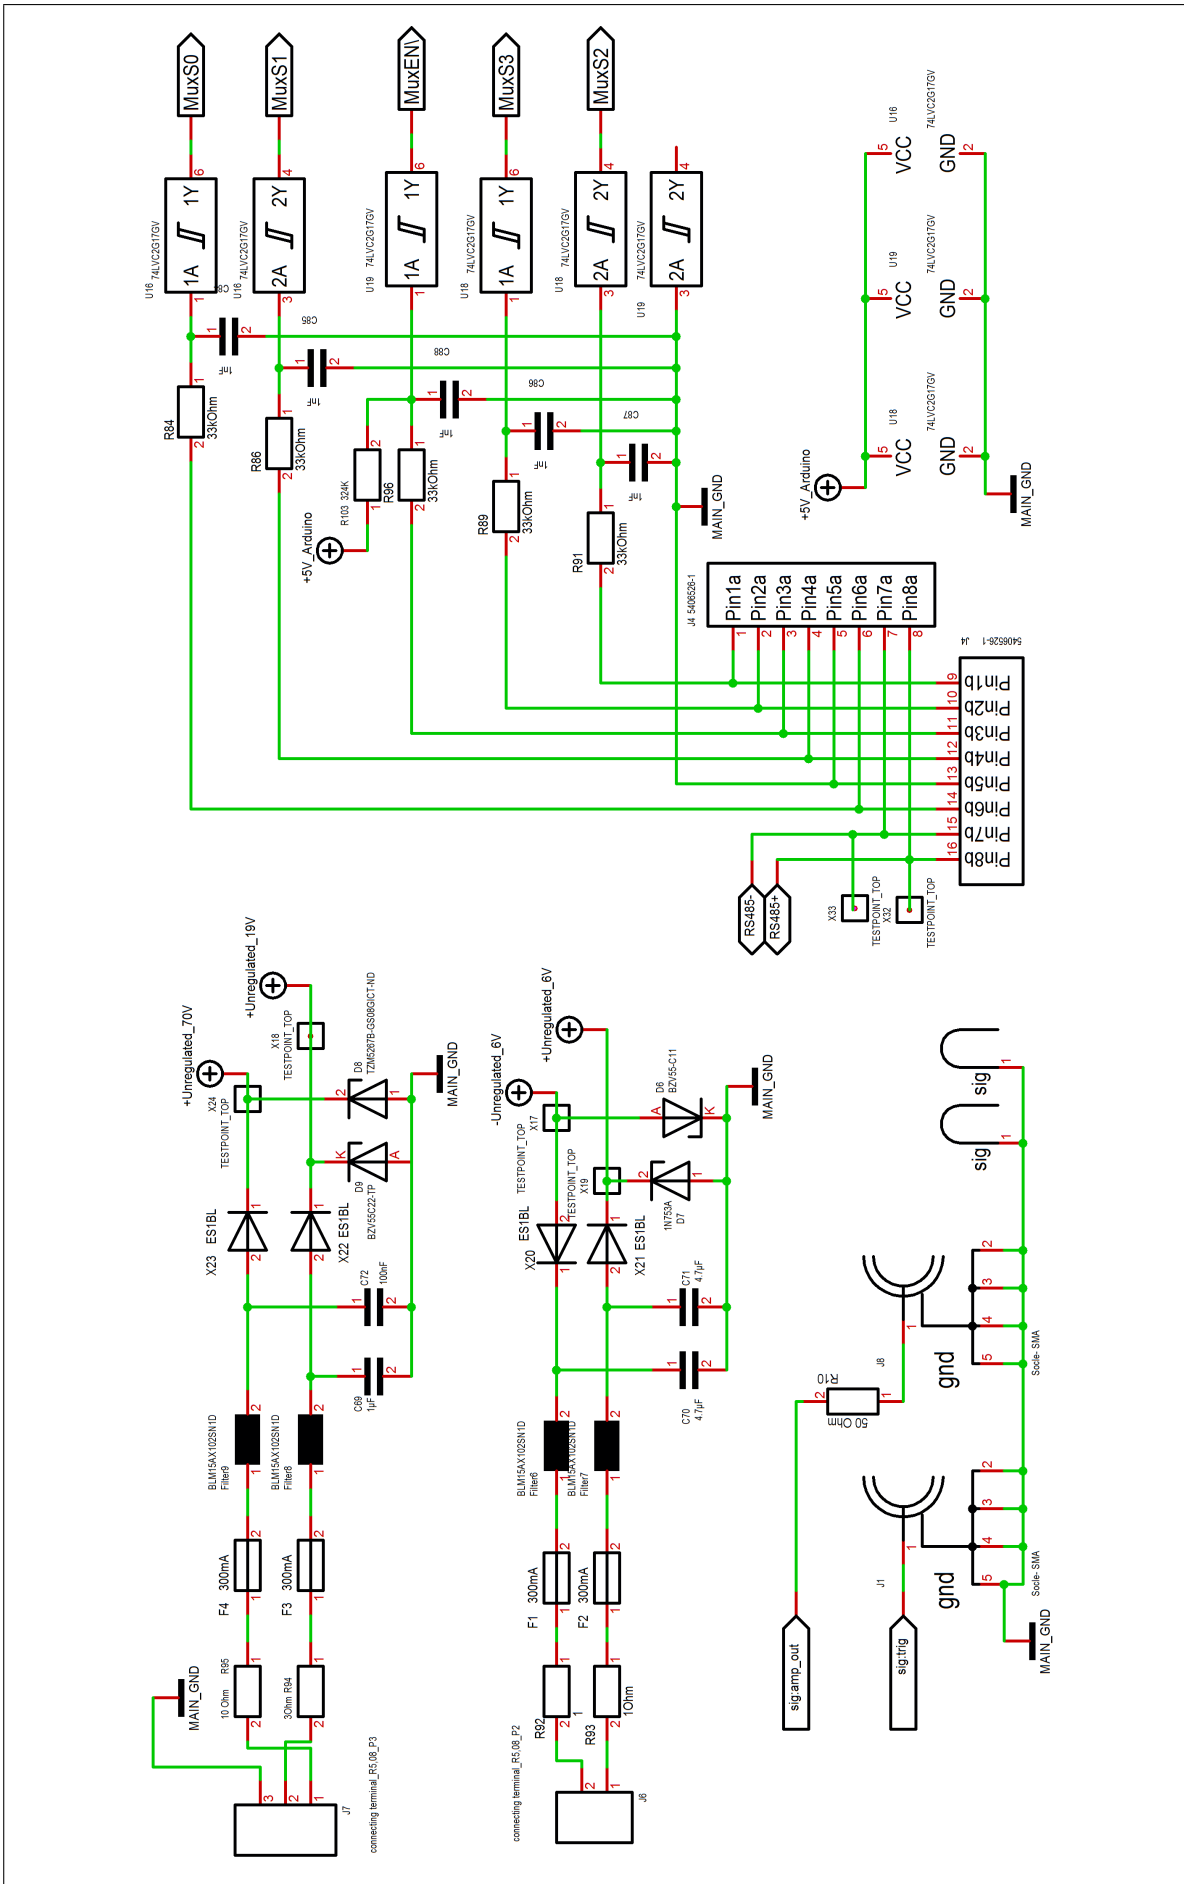**Figure D.6:** Circuit diagram of the ROB voltage regulators

Figure D.7: Circuit diagram of the ROB MCU inputs, outputs, and reset circuit.

Figure D.8: Circuit diagram of the ROB power inputs, input protection, signal outputs, and inputs for the MCU reset mux.

D.1.1 Multi-coincidence ambiguity

																		Probability of generating N two-fold coincidences from n simultaneous muons.

Probability of generating N two-fold coincidences from n simultaneous muons.
The most likely values are marked in dark green. Probabilities below 10^-3 % are denoted ~0. Dark grey coloumns are impossible values for N.

Coincs. / Sim. muons	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Expectation value
2	6,3	37,5	0	56,3	0	0	0	0	0	0	0	0	0	0	0	0	3,06
3	<1%	7,0	4,7	31,6	0	42,2	0	0	14,1	0	0	0	0	0	0	0	5,35
4	<1%	1,0	1,8	11,1	0	36,9	0	6,2	31,6	0	0	10,5	0	0	0	<1%	7,48
5	<1%	<1%	<1%	3,3	0	20,6	0	8,2	34,3	0	0	27,5	0	0	0	5,5	9,31
6	~0	<1%	<1%	<1%	0	9,6	0	6,9	27,8	0	0	40,2	0	0	0	14,5	10,81
7	~0	<1%	<1%	<1%	0	4,1	0	4,7	19,4	0	0	45,2	0	0	0	26,3	12,01
8	~0	~0	<1%	<1%	0	1,6	0	2,9	12,5	0	0	44,1	0	0	0	38,8	12,96
9	~0	~0	~0	<1%	0	<1%	0	1,7	7,7	0	0	39,4	0	0	0	50,6	13,69
10	~0	~0	~0	<1%	0	<1%	0	<1%	4,6	0	0	33,3	0	0	0	60,9	14,25
11	~0	~0	~0	<1%	0	<1%	0	<1%	2,7	0	0	27,2	0	0	0	69,6	14,68
12	~0	~0	~0	~0	0	<1%	0	<1%	1,5	0	0	21,7	0	0	0	76,5	15,00
13	~0	~0	~0	~0	0	<1%	0	<1%	<1%	0	0	16,9	0	0	0	82,0	15,25
14	~0	~0	~0	~0	0	<1%	0	<1%	<1%	0	0	13,1	0	0	0	86,3	15,43
15	~0	~0	~0	~0	0	<1%	0	<1%	<1%	0	0	10,1	0	0	0	89,6	15,58
16	~0	~0	~0	~0	0	~0	0	<1%	<1%	0	0	7,7	0	0	0	92,2	15,68
17	~0	~0	~0	~0	0	~0	0	<1%	<1%	0	0	5,82	0	0	0	94,09	15,76
18	~0	~0	~0	~0	0	~0	0	<1%	<1%	0	0	4,40	0	0	0	95,55	15,82
19	~0	~0	~0	~0	0	~0	0	<1%	<1%	0	0	3,32	0	0	0	96,65	15,87
20	~0	~0	~0	~0	0	~0	0	<1%	<1%	0	0	2,50	0	0	0	97,48	15,90

Figure D.9: Table of probabilities for generating N coincidences for multiple simultaneously impeding muons, and the expectation values for each muon multiplicity. Values calculated with Equations 6.4 and 6.4.

D.1.2 Area and N scalers

Unique polar angles, the number of ways they can be registered (N), N scalers (16/N), effective detector areas (A), and the area scalers (1/A).

Polar angle (rad)	N	N scaler	Effective area [m^2]	Area scaler
0	16	0,25	1,000	1,000
0,3476	48	0,083333333	2,820	0,355
0,4735	36	0,111111111	2,002	0,500
0,6271	32	0,125	1,620	0,617
0,6809	48	0,083333333	2,331	0,429
0,7976	16	0,25	0,698	1,433
0,827	16	0,25	0,677	1,477
0,8532	24	0,166666667	0,986	1,014
0,9175	16	0,25	0,608	1,645
0,994	4	1	0,136	7,353

Figure D.10: Table showing the ten different unique polar angles the CRT is able to detect. For a given polar angle, it can be detected *N* different ways. The rate of muons for each polar angle must be scaled with the Area scaler value in order to compensate for the difference in effective detector area.

D.1.3   Coincidence characterizations

Following are the full results of the coincidence characterization performed in order to find the optimal bias voltage, as discussed in Section 6.2.
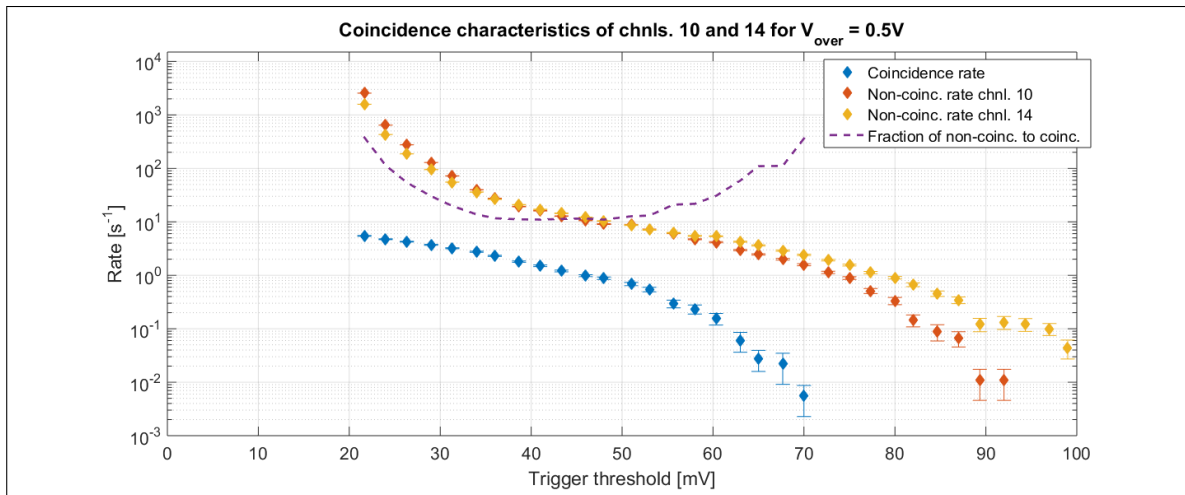

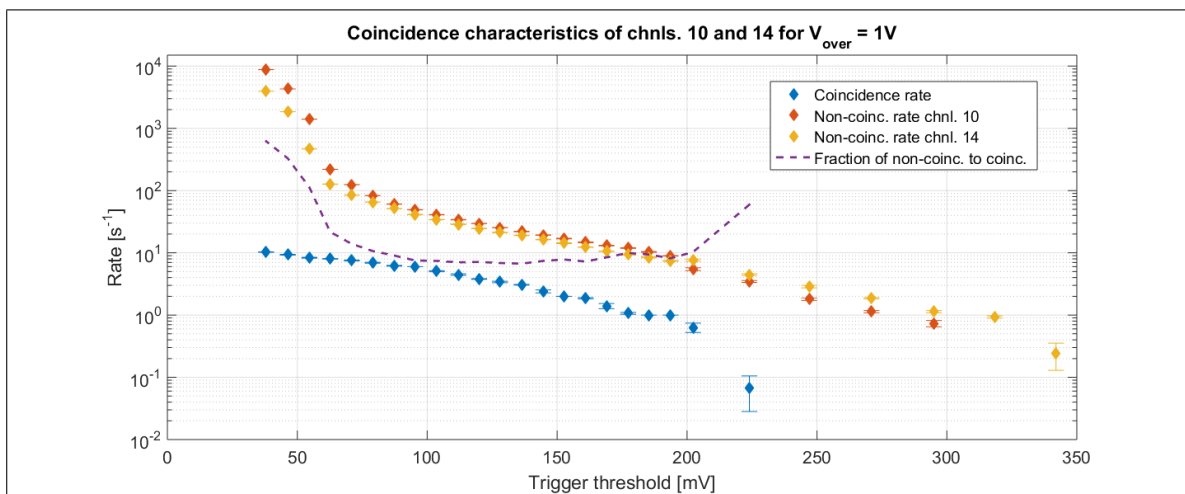**Figure D.11:** Coincidence characterization for 500 mV MPPC overvoltage.

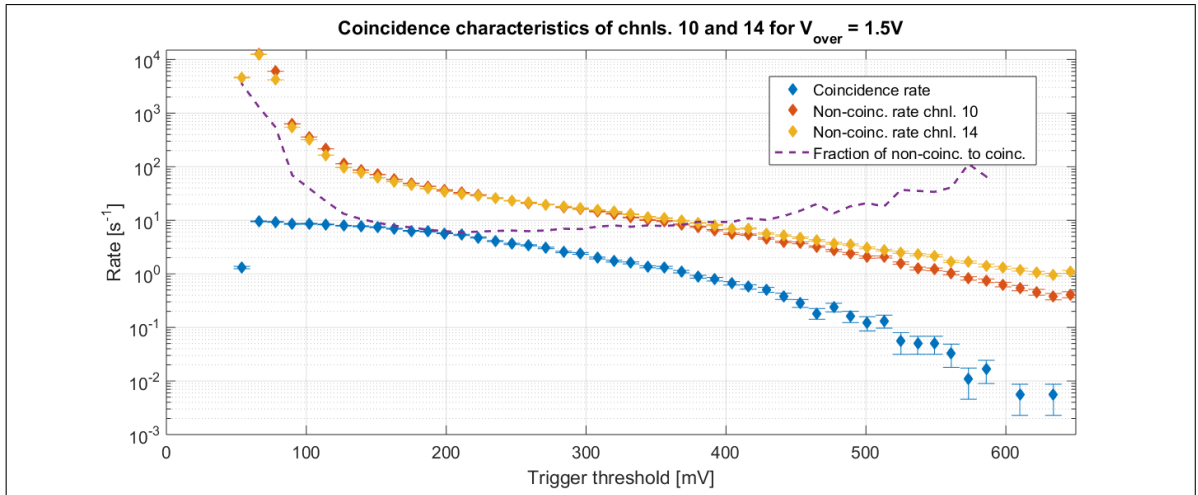**Figure D.12:** Coincidence characterization for 1000 mV MPPC overvoltage.

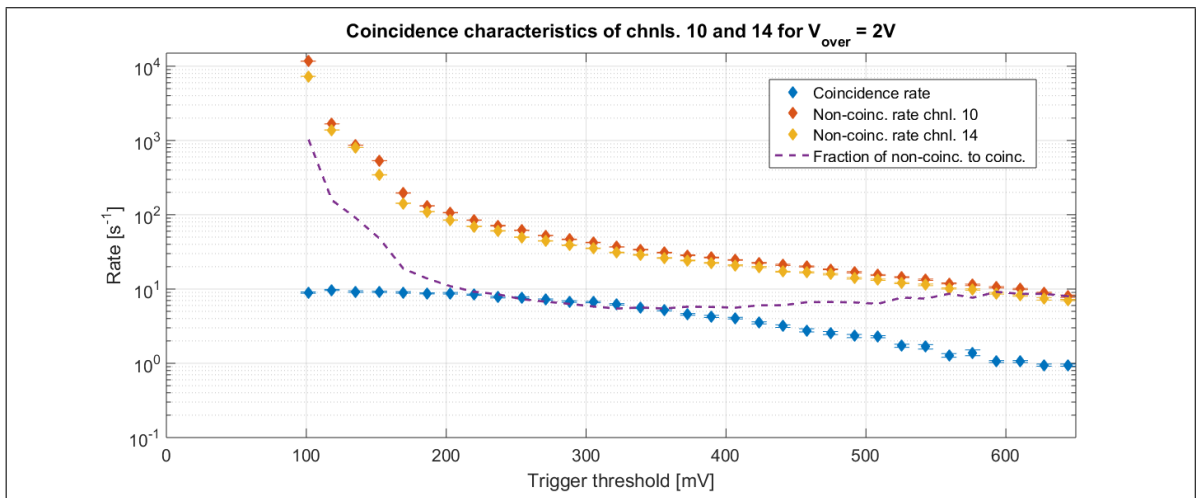**Figure D.13:** Coincidence characterization for 1500 mV MPPC overvoltage.

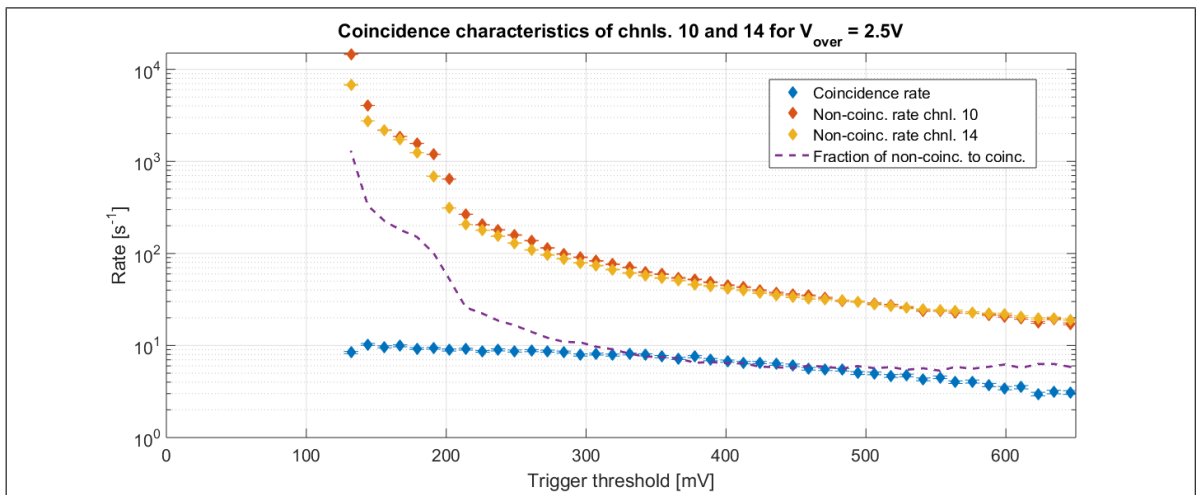**Figure D.14:** Coincidence characterization for 2000 mV MPPC overvoltage.

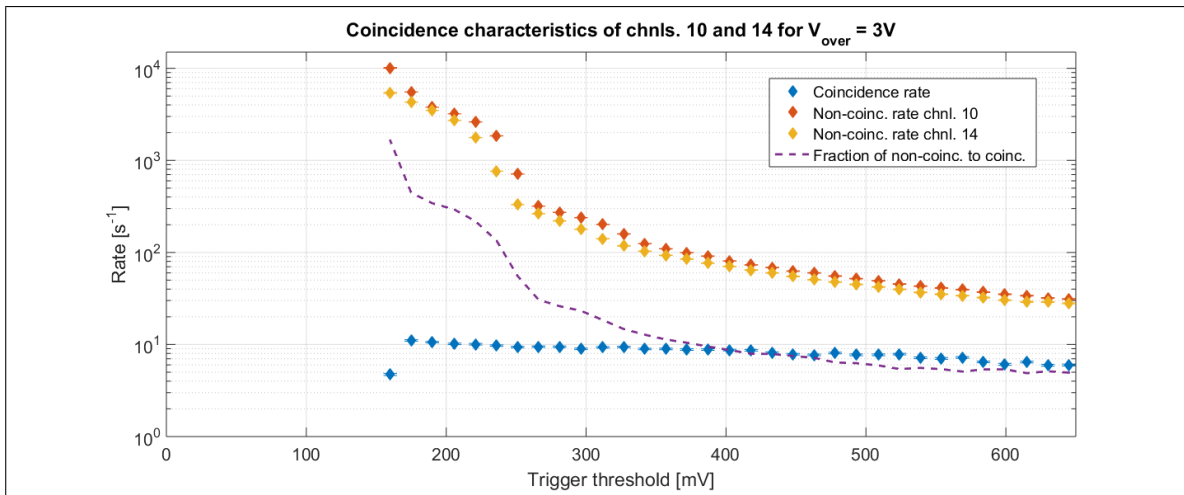**Figure D.15:** Coincidence characterization for 2500 mV MPPC overvoltage.

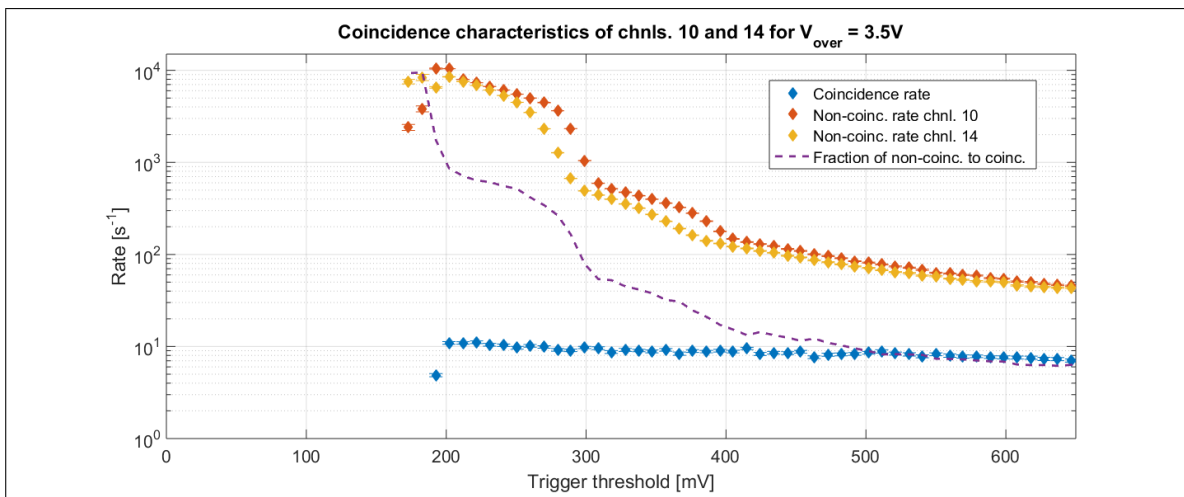**Figure D.16:** Coincidence characterization for 3000 mV MPPC overvoltage.

Figure D.17: Coincidence characterization for 3500 mV MPPC overvoltage.