Power and Monitor Solution for the Proton Computed Tomography Project

Birger Olsen

Master's Thesis in Physics

University of Bergen
Department of Physics and Technology

June, 2022

Acknowledgements

I want to thank both my supervisors *Professor Johan Alme* and *Professor Kjetil Ullaland* for invaluable guidance and feedback throughout my entire thesis. I would also like to thank *Senior Engineer Shiming Yang* for putting up with my constant barrage of hardware-related questions.

Being a part of the ProtonCT project this last year has been an incredibly fun experience, seeing so many institutions and people come together to work towards a single goal. Having worked alongside *Martin Eggen*, *Jakob Hauser*, and *Håvard Birkenes* has been a joy. Special thank you to *Viljar Nilsen Eikeland* for taking time out of your day to help me test the ALPIDEs, and to *Tea Bodova*, both for your help in understanding pCT, and for not holding back on comments after reading my thesis.

Thank you to my dear friends and peer students in room 312, *Amalie*, *Bendik*, *Håvard*, and *Thomas*, for filling these two years with fun discussions, both academic and casual. Another thanks to both *Bendik* and *Tarje* for reading through and giving great feedback on my thesis.

Lastly thanks to *Stine*, who is always there for me. Thank you for being you.

Abstract

The ProtonCT project is an academic endeavor carried out by the University of Bergen in collaboration with several universities and entities across the world. The end goal of the project is to improve dosage plans by directly measuring the relative stopping power of protons using a digital tracking calorimeter. Directly measuring relative stopping power as opposed to approximating it using CT numbers can provide a more accurate dosage plan. The digital tracking calorimeter will be able to do computed tomography scans of head-sized objects.

The digital tracking calorimeter will utilize pixel detector sensors developed by CERN for the ALICE project. 43 pixel arrays, segmented into layers, measure the angle and energy of proton particles traversing through the layers. With 108 chips per layer, 4644 ALPIDE chips build up all the layers. At full load, the expected power draw is close to 2.5 kW. This thesis explores the design of a user-controllable power delivery and monitoring system. Each layer consists of 12 ALPIDE strings, with 9 ALPIDE chips making up one string.

A power delivery system capable of supplying one layer is realized by using a small form factor switch mode power supply unit. An FPGA design created by peer students connects the 43 power delivery systems to a graphical interface. A filter, monitor, and control solution is designed with a newly released AVR microcontroller unit. A custom PCB, named the *Monitor Board*, is designed to host the filter and the MCU with all its support circuitry.

Using differential signaling, the 43 monitorboards communicate with a Xilinx Kintex UltraScale FPGA responsible for storing and relaying information over IPbus to the user. Each monitor board can switch the strings of its designated layer on or off. Diagnostics and soft startups/shutdowns can be executed through software. The back-biasing of an entire layer is customizable by using the microcontroller DAC and an onboard negative voltage supply. A temperature monitoring solution is designed with the use of a PT1000 element mounted close to the ALPIDE chips.

Contents

List of Figures

1 INTRODUCTION

This chapter outlines the motivation to use proton therapy to treat cancerous growths in high-risk areas of patients. Next, one of the goals of ProtonCT, the development of a high granularity digital calorimeter, is outlined, and the work I have been assigned to contribute is described in full. A chapter summary is present at the end of this chapter.

1.1 Motivation to Use Proton Therapy

Norway is a relatively small country with a population of 5.4 million at the time of writing [1]. From 2016 to 2020, 827 new cases of cancer were reported in the age group of 0 - 14 years old, with a third of those cases being cancer in the central nervous system (CNS) [2]. Over 80% of all CNS cases are brain tumors [3]. Radiation therapy plays a significant role in the treatment of brain CNS cancer [3]. Improving the dose delivery of cancer is crucial in diminishing the risk of developing cancer in these high-risk areas post-treatment. The ProtonCT (pCT) project aims to improve the measurement of the relative stopping power (RSP) of protons to aid in enhancing dosimetry plans for pediatric patients, decreasing the chance of developing more cancer later in their lives.

There are different types of radiation therapy, but this thesis will only focus on external beam therapy, which coincidentally is one of the more conventional treatment forms. External beam therapy comes with the side effect of being capable of producing Radiation-Induced Secondary Malignancies (RISM), and a study done by SEER (Surveillance Epidemiology and End Results) expect that recipients of radiation treatment to the brain will have a 1.9 - 2.4% chance of developing RISM[1]. Most tumors appear within the margins of the treatment area [4].

[1]Keep in mind that secondary malignancies are harmful growths per definition, and not all growths developed after radiation therapy falls into this category. 'Neoplasm' is used to describe both malignant and benign growths, a tumor can be both benign and malignant, but cancer is by definition something that spreads and is therefore malignant. SEER reported a relatively low amount of RISM but did not report the total amount of neoplasms, meaning the number could be higher, and this is a best-case scenario.

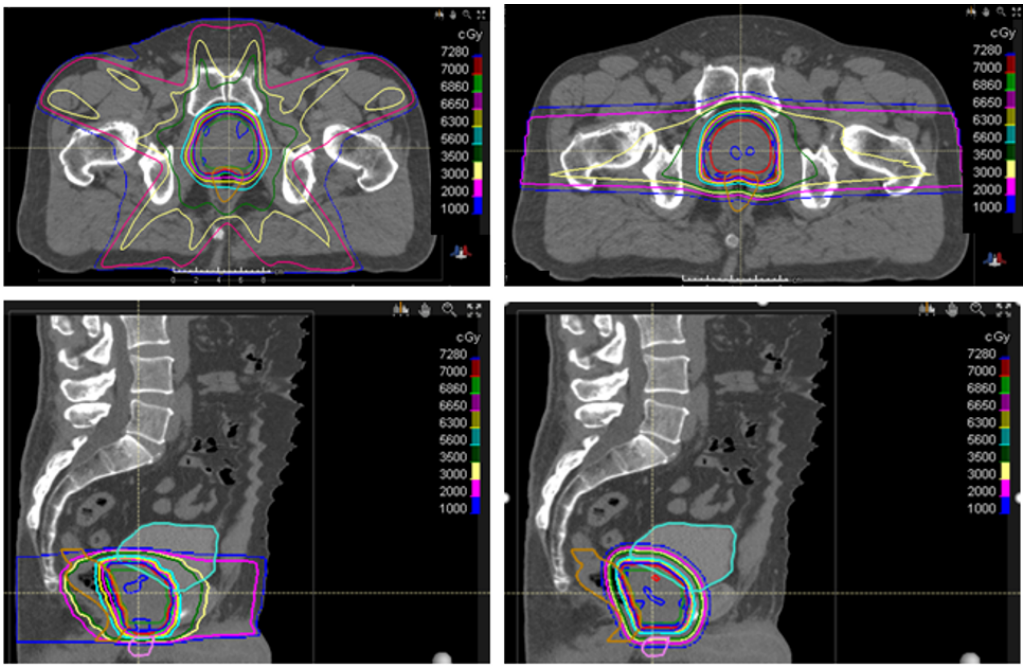

Figure 1.1: IMRT(left) versus proton(right) radiation dose.

Figure 1.1 shows how a multi-beam intensity-modulated radiation therapy ends up irradiating healthy tissue, even when irradiating from several angles to keep the unwanted radiation to an absolute minimum.

Photon radiation therapy delivers energy gradually as it travels through tissue, lacking a distinct energy peak. Intensity-modulated radiation therapy (IMRT) uses several lower energy beams at different angles to limit the radiation delivered to healthy tissue. Proton dose delivery is very different, depositing small amounts of energy at the start of its travel, and delivering most of its remaining energy at the end of its travel. This sharp increase in energy deposition is the *Bragg Peak* of the proton. As a result, a subject can be beamed from a single entry point without it affecting as much healthy tissue before and after the cancerous growth. Figure 1.1 shows how a patient may be treated using particle radiation from a single angle. Beaming a patient using proton beams from several angles would yield the best results.

IMRT is not always an applicable treatment method, malignancies located in and around critical areas might not have enough "entry points" for the photon beams. IMRT was first made commercially available in 1994 [5], and has since increased the number of beams. In figure 1.1 we see five beams used to treat the patient, but in some modern cases, up to 50 beams could be used [6]. A more common approach is to stick with 10 or fewer beam angles [6].

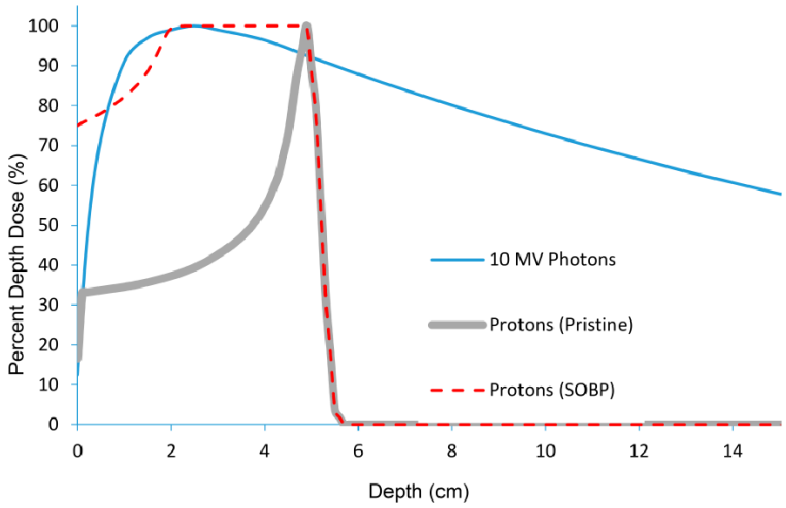

Figure 1.2: Photon versus proton energy deposit in tissue [7].

The proton beam is modulated to deliver the Bragg peak at different depths, irradiating the entire growth while maintaining the characteristic Bragg peak drop-off. Figure 1.2 shows this effect, both the pristine proton and the spread-out Bragg Peak (SOBP) proton beams have the same drop-off characteristic. The photon beam has a sharp increase in energy deposit during the first few centimeters, then it declines slowly traveling through tissue.

1.2 The ProtonCT Project

Utilizing the Bragg Peak effect to treat cancer is not a new concept, and it is entirely possible to obtain proton radiation treatment today. The pCT project aims to further analyze and understand the relative stopping power of protons in human tissue, which will help create more accurate dosage plans.

Proton radiation therapy has been in development since 1946. Robert Wilson was the first to present the idea of utilizing the Bragg peak of protons to treat cancer patients [4]. Proton therapy has the potential to eradicate deeply burrowed cancerous growths while not irradiating nearby healthy tissue, making it safer than traditional photon radiation [8]. Conventional proton radiation therapy first captures a model of the subject by doing a computed tomography (CT) scan, and the information is then analyzed and translated to create a dosimetry plan for proton radiation. However, approximating how protons will behave in nonhomogeneous tissue is complex and creates uncertainty, counteracting the benefits of using proton therapy. The pCT project at UiB explores the use of proton particles for imaging to create a more accurate dosimetry plan, making the irradiating of the patient more accurate and the treatment safer overall.

Conventional CT imaging uses X-rays to create a dosimetry plan. The dosimetry plan tells the system and the radiologist how the patient should be irradiated, e.g where the beam should be pointed and how much intensity it should have. A photon

CT scan can produce very accurate proton dosimetry plans, but anything barring perfect is lackluster when treating cancerous growths in high-risk areas like the brain or lungs. Minimizing the amount of healthy tissue irradiated will decrease the total amount of radiation delivered during treatment, reducing the risk of developing cancer at later stages in the patient's life.

High energy proton beams deposit very little energy when passing through tissue, making it safe to use when imaging, but capturing the data is challenging. The patient is beamed with high-energy proton particles, and the particle's exit angle and energy are recorded using a high granularity digital tracking calorimeter (DTC). The data gathered is the relative stopping power of the proton, which describes how protons interact with the tissue it was sent through. The RSP values are used to create a dosimetry plan.

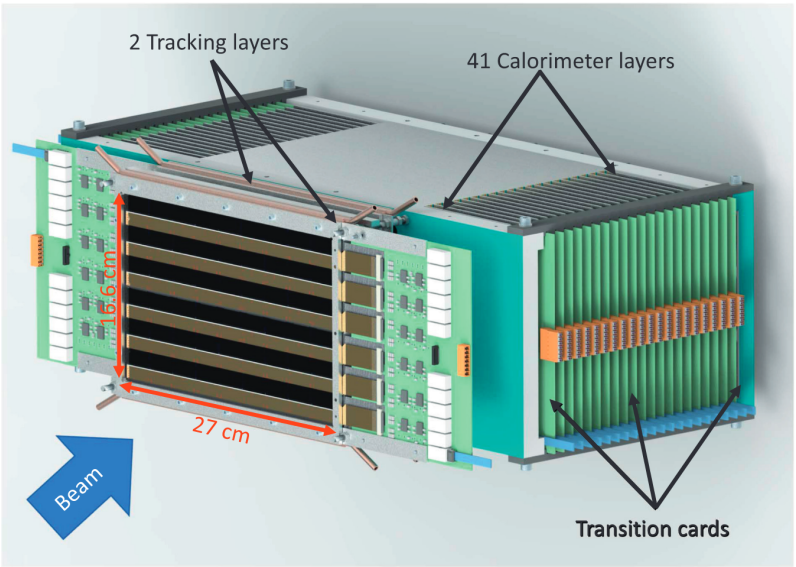

Figure 1.3: 3D model of the DTC [9].

The DTC is in active development by the pCT team; several master's and Ph.D. theses have been dedicated to pCT project. Detailed simulations predicting how protons will behave between the calorimeter and tracking layers are under constant development, and as much as a glue compound change has to be reported to the simulation team.

The chips measuring the presence of protons require a constant voltage of approximately 1.9 V to operate, and ideally, this voltage should be as noise-free and as stable as possible. There are only theoretical designs made for the pCT project power delivery system, and this thesis will explore the creation of a complete power delivery and monitor system.

1.3 Primary Objective of This Thesis

The DTC layers are built up of 12 rows of detector chips, a single row can consume up to 2 W [10], but accounting for the loss in regulators and cables, the total power draw adds up to 4.7 W per row, or 56.4 W per layer. The power delivery system must support a throughput of 2.5 kW (assuming the power delivery system delivers 4.5 V). The total power draw is split into a digital and analog power rail, and a noise requirement of 1 $mV_{p\text{-}p}$ has to be met on the analog rail.

Only radiation-tolerant components should be close to the DTC. The power delivery system connects to the transition card (TC), which handles power delivery to the ALPIDEs and acts as a carrier from the flexible printed circuit (FPC) to FireFly for the data lines. To shield the power supply unit (PSU) from radiation, it should not be closer than five meters to the DTC. The TC is outlined in section 3.2.

The TC requires a minimum of 3.3 V, and to limit waste heat it should not have a supply voltage higher than 5 V. An outline of a monitor, control, and power delivery system will be discussed in detail later in Section 3.

The core functions of the power delivery system are as follows (in no particular order):

1. Supply stable power of at least 3.3 V to the TC, low noise (<1 mV). The power supply solution should be no closer than 5 m to the DTC.
2. Be able to deliver at least 13.2 A of current. 2.4 A to AVDD and 10.8 A to DVDD.
3. Monitoring of the current passing through the AVDD and DVDD supply lines (PWELL optional).
4. Take regular measurements of a temperature sensor.
5. Logic to allow the user to control the voltage regulators on the TC.
6. A communication system, relaying information to a control database, preferably through IPbus.
7. User controllable variable negative voltage.

In addition, the design should ideally be made of commercial components. Future assembly of the design should not rely on custom or expensive one-off designs, but rather it should be easily purchasable. Other notable qualities are scaleability, form factor, customizability, and efficiency. An easily scalable system allows us to expand the DTC with more layers in the future if needed. Efficiency is essential due to the large power draw of the entire machine; even small inefficiencies add up when dealing with large systems. Customizability denotes how easily the system or design can be altered later on, locking the design to be completely dependent on certain components can create issues in the future.

One or several temperature sensors will be placed on the TC, ideally as close to the chips as possible to monitor the temperature. The exact position is not determined as of now. In the event that the temperature rises above a set value, the enable signals shall be pulled low to shut off the voltage regulators on the TC. The monitoring solution must be able to operate autonomously, the user should not be in charge of taking action if the temperature rises. The same preset shutdown has to happen in the event of large current surges.

1.4 Chapter Overview

Chapter 2 - Proton and Photon Radiation *This chapter starts by examining the long-term effects of radiation therapy, then compares the long-term effect of both photon and proton radiation therapy. Further, it provides some basic knowledge on X-ray CT scanning and how it relates to proton treatment. This is beneficial to aid in understanding why we are developing a high granularity calorimeter.*

Chapter 3 - The ProtonCT Project *This chapter will discuss what a digital calorimeter is exactly, and how it relates to the ProtonCT project. How the calorimeter tracks and measures proton energy will be explained. The current theoretical design of the Power Control Unit will be discussed, and the pros and cons will be outlined versus my design choices.*

Chapter 4 - Hardware Design Considerations *This chapter will focus on considerations around building a real-world system. Conventional PCB layout methodology will be outlined, and highlighting how it is relevant for the pCT project. The content will be relevant for the later creation of a custom PCB, and to the creation of the PSU filter.*

Chapter 5 - The Power Supply *This chapter first determines how much power is needed to be delivered to each TC, then outline the maximum voltage that can be delivered and why it matters. Later the two main categories of power supplies will be discussed, highlighting the pros and cons. A power supply is chosen for testing and prototyping based on this discussion. The ALPIDEs have been tested using this new PSU. The tests are explained, and the results are shown.*

Finally, several filters are proposed and tested to best determine the classification of noise and thereby determine the final solution.

Chapter 6 - The MonitorBoard *The monitorboard PCB made specifically for the MAX105 is outlined. The functionality is described, and the working of the current software is explained in broad terms.*

Chapter 7 - Results and Future Work *This chapter evaluates the work described in this thesis. Additionally, it discusses the remaining tasks of the project.*

2 PROTON AND PHOTON RADIATION

This chapter starts by examining the long-term effects of radiation therapy, then compares the long-term effect of both photon and proton radiation therapy. Further, it provides some basic knowledge on X-ray CT scanning and how it relates to proton treatment. This is beneficial to aid in understanding why we are developing a high granularity calorimeter.

2.1 Secondary Neoplasms in Patients

The Journal of the National Cancer Institute examined the well-known *Childhood Cancer Survivor Study* cohort and compiled the incidents of subsequent neoplasms at 5- and 30-year intervals. There were 14 359 survivors enrolled in the study. Of the survivors, 1514 (10.5%) received only radiation therapy, and 7022 (48.9%) received both chemotherapy and radiation therapy [11].

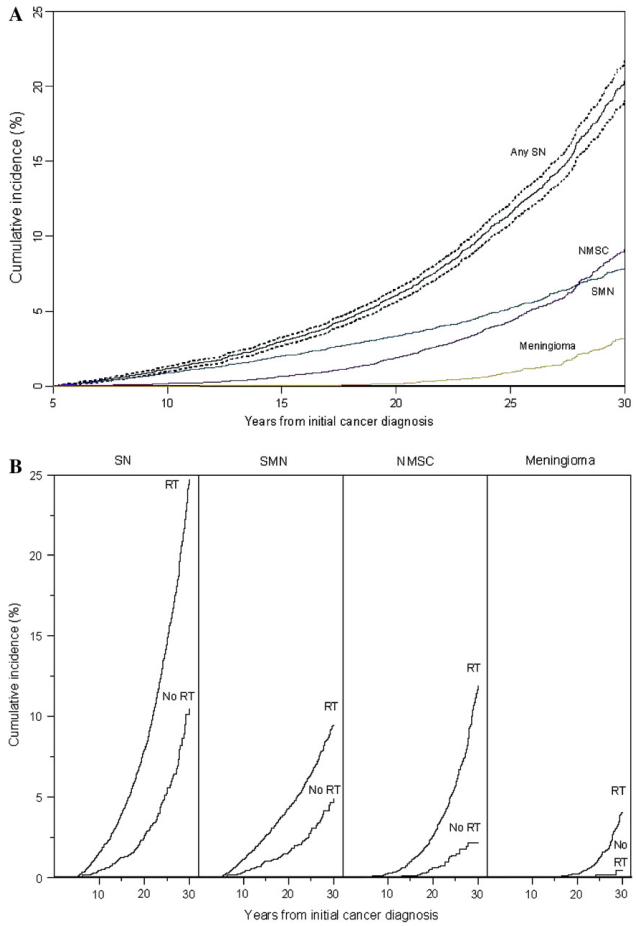

Figure 2.1: **A)** Cumulative incidence of second neoplasms at 30 years after initial cancer diagnosis. **B)** Cumulative incidence of any secondary neoplasm, second malignant neoplasm, nonmelanoma skin cancer, and meningioma [11].

Looking at Figure 2.1, almost a quarter of the patients receiving radiation treatment developed some form of subsequent neoplasm [11]. By effectively reducing the patient's treatment area, improving dosimetry plans to reduce the radiation to healthy tissue can lessen the chance of secondary neoplasms. As mentioned earlier, most malignant growths appear within the radiated area [4], so reducing it will reduce the chance of radiation-induced secondary malignancies.

Long Term Effect of Particle Treatment

The International Journal of Radiation Oncology [12] conducted a study (n=1034) where cancer patients treated with proton radiation therapy were compared to cancer patients who had received similar photon radiation treatment. Of the 1034 candidates receiving proton treatment, 558 cases could be matched to similar cases receiving photon radiation treatment. Unfortunately, several factors are ignored due to the low sample size, such as age and sex. The patients participating in the study received treatment from 1973 to 2001. In the context of proton treatment, the long-term effect of proton radiation versus photon is interesting, so similar cases are being matched against each other to analyze the long-term effect.

Of the 558 proton and 558 photon patients, secondary malignancies occurred in 5.2% and 7.5% of the cases, respectively. The total incident rate of secondary malignancies per 1000 patients was 6.9 for matched proton, 10.3 for matched photon, and 5.6 for unmatched proton. The study concludes with there currently being no signs of increased risk with proton treatment versus photon.

It is important to note that a large portion of the unmatched proton patients received treatment for cancer within or near the skull/central nervous system, a placement not ideal for photon radiation treatment. Proton treatment is in other words already in use in areas where they stray away from using conventional radiation treatment. The use of proton radiation treatment limits the amount of radiation to nearby healthy tissue, which is discussed in detail in Section 2.5.

2.2 X-ray Computed Tomography Scan

The CT scan is a medical technique to create detailed images of the cross-sectional area of a patient's body. The CT scan is entirely non-intrusive and is a widespread tool used in diagnostic medicine. CT scanners usually utilize a rotating X-ray tube together with detectors to measure the X-ray attenuation of the patient. Figure 2.2 shows a patient laying on a motorized table while the X-ray source and detectors circle around the subject. This creates slices of the patients' cross-sectional area.

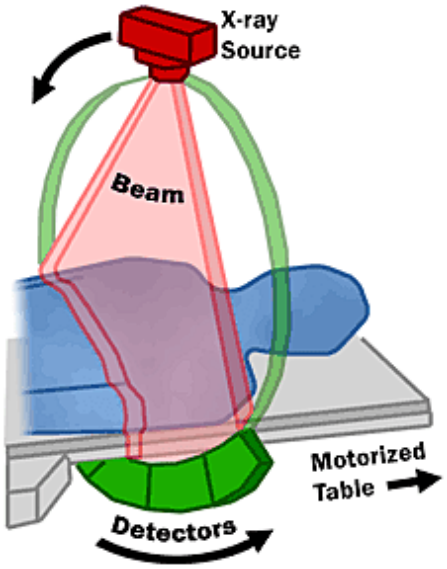

Figure 2.2: Example of a CT system [13].

Conceptually a proton CT machine works the same way as an X-ray CT scanner, using a particle accelerator to charge the protons and utilizing a calorimeter to measure energy.

A CT number is acquired to characterize tissue densities recorded during a CT scan. The CT number is based on the Hounsfield scale, expressed in Hounsfield units (HU), and directly correlates to the radiodensity[2] of the tissue. The radiodensity of water is 0 HU, tissue with a higher radiodensity than water has positive Hounsfield values, while lower density tissue has a negative value [14]. Fat is less dense than water, and would therefore not absorb as much radiation, showing up in a darker color than water would on a CT scan.

A CT image is an ensemble of scans from various angles used to construct a 3D image of the patient. Individual slices can also be analyzed if you want to look at the cross-sectional area of a patient. Each point in the 3D model is a voxel, and each voxel is color based on the Hounsfield value of that particular point [15].

2.3 Proton Computed Tomography

Scanning a patient using proton CT exposes the patient to less ionizing radiation than X-ray imaging does. The proton scan can also be used for cross-calibration of other X-ray CT scanners [15]. However, the most important benefit is directly measuring the relative stopping power (RSP).

RSP is a measure of how much energy protons deposit while traveling through a specific tissue relative to how much energy protons deposit traveling the same distance in water [16]. To perform particle therapy using protons, the RSP has to

[2]Radiodensity is the relative inability for X-ray radiation to pass through tissue. The higher the number, the more radiodense it is.

be either measured or calculated. RSP is to proton particle therapy what Hounsfield units are to massless radiation therapy.

DeJongh et al. 2020 took an X-ray CT scan of a porcine specimen, then measured the RSP of protons traveling through the same specimen. The purpose was to compare the calculated RSP using the X-ray CT scan against the direct measurement. This experiment shared much the same goals as pCT have, improving dosimetry plans by directly measuring RSP.

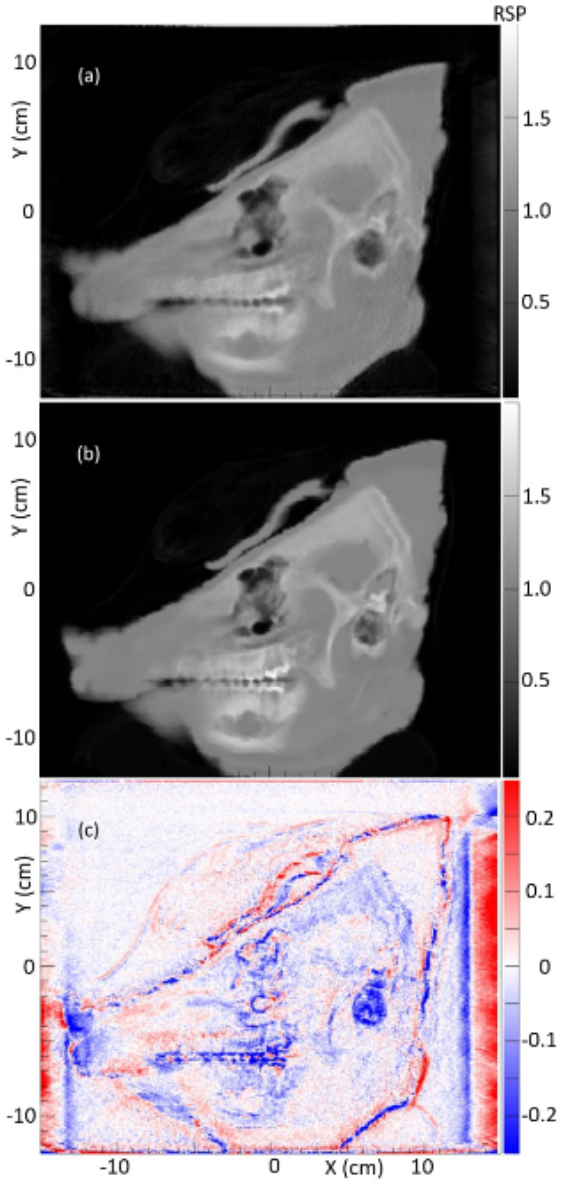

Figure 2.3: Ten-slice averages (a) Pig's head pCT. (b) Pig's head vertical X-ray CT. (c) proton CT and X-ray CT difference [15].

Figure 2.3 (c) shows the difference of the measured and the calculated RSP. The difference was within 1% to 2% for homogenous soft tissues. They found a larger discrepancy in compact bone, where the difference was 7%, and the largest in mixed

areas consisting of air, soft tissue, and bone, with up to a 40% discrepancy.

Using single proton tracking to measure the RSP yields a clear advantage in the accuracy of the RSP and can therefore be used to improve dosage planning by hopefully the same amount.

2.4 Photon Interaction with Matter

Treating and imaging patients using photons require them to be sufficiently charged to enter and penetrate matter correctly. High-energy photons have three distinct interactions with matter based on their energy. Diagnostic radiology usually employs photon energies in the range of 10 keV to 100 keV, with the three interactions being ordered from requiring the least to most energy.

1. Photoelectric effect
2. Incoherent and coherent (Compton) scattering
3. Pair and triplet production

At diagnostic radiology energy levels, the photoelectric interaction is mostly the only one occurring. Here a photon transfers its energy to an orbiting electron, knocking it out of orbit. Ionising an atom this way causes an outer shell electron to fill the spot of an inner shell electron. The electron emits X-rays due to the difference in binding energy [17].

2.5 Proton Interaction with Matter

An important note is that barring Compton scattering, photons do not change course when traversing tissue. Compton scattering can be avoided by keeping the beam energy below 100 keV. On the other hand, protons scatter quite frequently, and this is one of the main hurdles to overcome with proton CT scans.

When a proton traverses matter, its energy decreases as it collides with nearby electrons. Charging a proton to a predetermined energy level as denoted in the dosimetry planning, the proton is charged such that it stops inside a patient's treatment area. When reaching low energies, the proton stops abruptly and violently, depositing all its remaining energy into the tissue in its immediate vicinity. This sharp increase in energy deposition is the Bragg peak of the proton. The energy the proton beam is charged with during treatment is calculated from the RSP data. When utilizing protons for imaging purposes, the proton is charged sufficiently enough to pass through the subject being imaged, so that only a low amount of radiation is inflicted during this stage.

The Bragg peak is caused by the fact that the longer a proton is in the vicinity of an electron, the more charge it deposits. The range of a proton is proportional to kinetic energy squared [4], which means the energy deposited increases squared as the speed diminishes. This interaction is not unique to protons, but it has to do

with the mass of the particle. The more massive a particle is, the sharper the Bragg peak, but protons have been chosen for radiology as their energy also falls close to zero after the peak, not something that is true for all particles. For example, both neon and carbon exhibit a sharper Bragg peak, but both particles also deposit more dose following the peak, making it more difficult to use for radiation therapy [4].

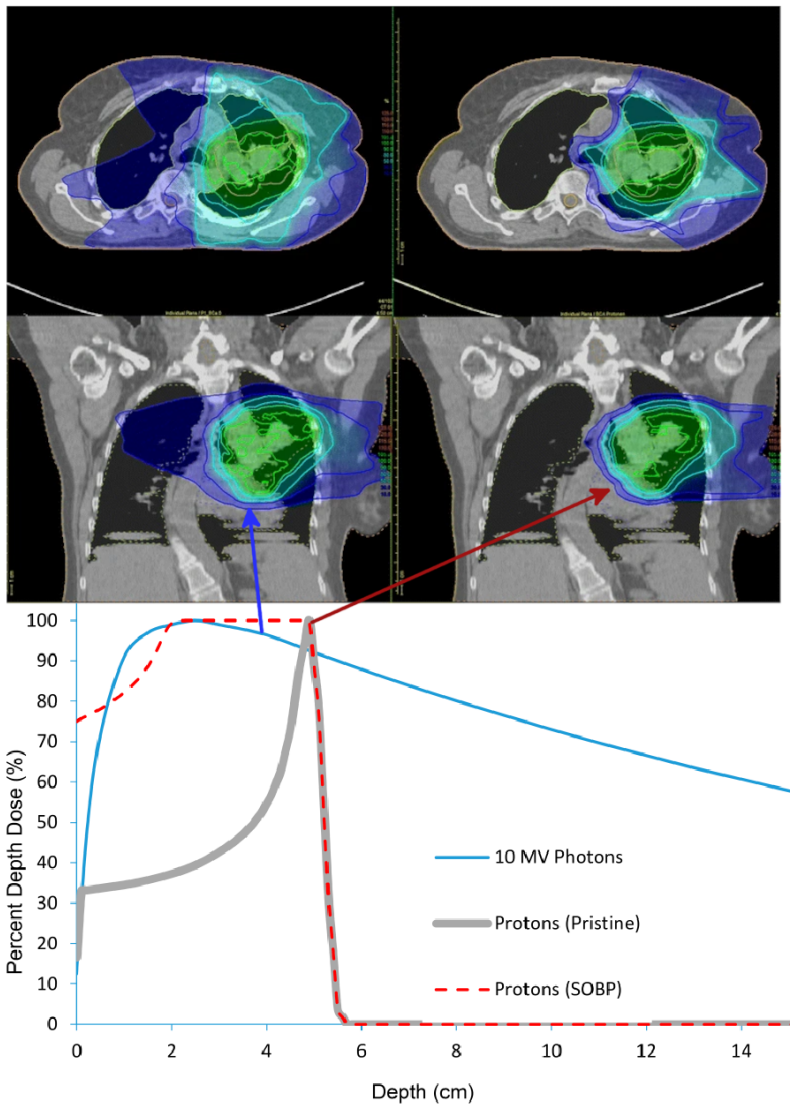

Figure 2.4: Bragg peak of a photon versus proton radiadion. The dose delivery in tissue is shown in a simulation. Top image [18]. Bottom image [19].

On the right of Figure 2.4 we see that proton radiation therapy will deposit some dose as it enters the patient, but after the Bragg peak occurs, the particle stops. Less radiation is delivered after the tumor than what we see with photon radiation on the left. The irradiation of the patient becomes more precise, meaning less overall radiation will enter the body, but the tumor will still be treated.

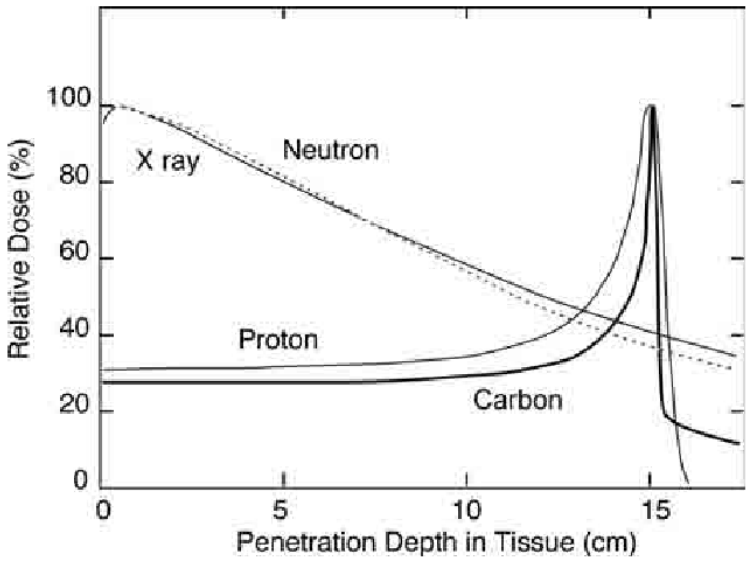

Figure 2.5: Comparative dose distributions for various therapy beams. Carbon exhibit a more prominent Bragg peak and tail compared to proton [20].

Figure 2.5 compares X-ray, proton, and carbon dose distributions. Carbon has a lower initial dose before the Bragg peak, and a sharper peak overall, the extended dose delivery following the peak is greater. The sharper Bragg peak is owed to less multiple scattering [20]. The dose delivery following the Bragg peak is caused by lighter fragments from nuclear reactions of the slowing carbon ions, which are not present in the proton, showing an advantage of using protons as opposed to carbon [20].

As with photons, proton particles has three distinct interactions with matter, *stopping*, *scattering*, and *non-elastic nuclear interactions* [4]. To measure the RSP of a particle, we are most interested in the stopping and scattering effects. non-elastic nuclear interaction is an unwanted and rare interaction during range estimation and won't be discussed [16].

The stopping interaction is accurately described by its name, as protons traverse matter, it interacts with electrons, transferring their kinetic energy and slowing down. This is a very consistent interaction, once again reiterating the necessity for an accurate RSP [4]. Protons are slowed down when traversing tissue, and the denser the tissue, the more electrons the proton interacts with, and thus they slow down faster. Unfortunately, protons don't move in a straight line through non-vacuums.

Proton beams traversing matter are not only slowed down, but through electromagnetic interactions with atomic nuclei, their path is also altered [4]. This phenomenon is known as both "scattering" and "multiple Coulomb scattering". The proton's resulting exit path can be accurately approximated from the input speed and angle [21]. Multiple Coulomb scattering will occur hundreds of thousands of times in the time it takes a proton to travel through a patient, constantly making small adjust-

ments to its path. The scattering observes a Gaussian distribution. Millions of particles are sent through the patient to approximate the most likely path [15]. We can also assume a proton's maximum scattering angle as shown in Figure 2.6, we know the input angle, and we can then assume the proton will scatter in a cone with radius X_0.

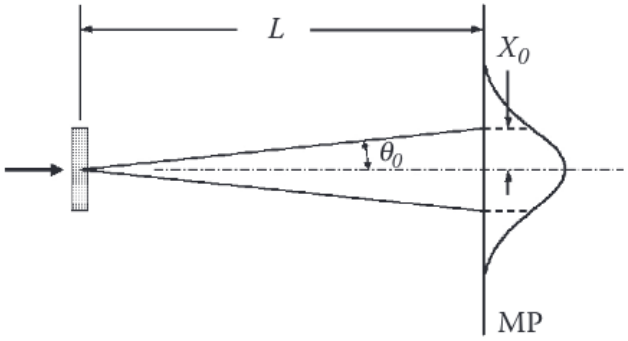

Figure 2.6: Multiple Coulumb scattering in a thin slab. [21]

The current most likely path calculation is done using the assumption that the proton is going to scatter the same way it does when propagating through water [22]. This assumption is reasonable as the human body consists of approximately 70% water, making it precisely predict the stopping point within 1 mm [22]. This is a nice time to study Figure 2.3, where it is seen that the calculation of RSP is best done in soft tissue, which holds a lot of liquid. The largest discrepancies are found in air-filled cavities and compact bone, the two materials in our body which least resemble water.

2.6 Summary

DNA damage from radiotherapy can make up for half of the secondary malignant neoplasms found in cancer survivors, and the majority of these neoplasms occur in the treatment area.

Compton scattering is the main interaction that photons exhibit in diagnostic radiology. In proton therapy, *stopping* and *scattering* are the main interactions protons exhibit. Photons travel in a straight line when going through patients at diagnostic energies, while protons are scattered continuously, requiring us to calculate their most common path. An RSP value is usually extracted from an X-ray CT scan in the preliminary work before proton particle therapy. The sharp Bragg peak of protons allows us to deposit almost all of the charge in a pinpointed area during treatment, delivering substantially less ionizing radiation to healthy tissue. Approximating its path is difficult, but its distribution is Gaussian, and its mean path is calculated when performing CT scans. By using a DTC the direct measurement of RSP can be done to improve dosage planning.

Recalling back to Section 2.3, direct measurements of the RSP of protons have shown to have clear benefits as opposed to calculating them. Improvements of up to 7% could be made by measuring instead of approximating.

3 THE ProtonCT PROJECT

This chapter discusses what a digital calorimeter is, and how it relates to the ProtonCT project. How the calorimeter tracks and measures proton energy will be explained. The current theoretical design of the Power Control Unit will be discussed, and the pros and cons will be outlined versus my design choices.

The earliest thesis submitted to the pCT wiki is dated 2014, since then there have been tens of theses submitted expanding upon the work. The end goal for pCT is to get a better understanding of the RSP of particles, then apply this knowledge to improve dosimetry plans for patients, particularly pediatric patients. A digital calorimeter is in its final stages of development, which will be used to measure the RSP of protons.

More conventional calorimeters in use today report a value of absorbed dose-to-water by measuring the heat change in some medium that absorbs the ionizing radiation. A calorimeter can employ water, graphite, or other mediums with known characteristics [4]. The absorbed dose is measured and used as a reference point for the dosimetry. This is very accurate at measuring absorbed dose, but it gives no information about the individual protons. The pCT DTC wants to track individual protons to measure RSP directly, as opposed to just saying how much radiation was sent through the patient [23]. This is also what differentiates a water calorimeter from a digital calorimeter, as the DTC operates in a binary hit- or no-hit fashion.

The DTC is built up of several layers of monolithic active pixel sensors (MAPS). The chips were originally designed for the "A Large Ion Collider Experiment" (ALICE) project, appropriately called ALICE Pixel Detector (ALPIDE) chips, which are used to detect particles at high resolutions. The DTC has 43 layers, 2 tracking layers up front, and a following 41 calorimeter layers. The tracking layers report the angle of the incoming proton, while the calorimeter measures the remaining energy.

The two tracking layers measure the scattering angle a proton particle inhabits after passing through a subject. The tracking chips are mounted on carbon carriers and do not have any aluminum absorbers, minimizing the material budget to stop the proton from scattering in this stage. A proton's exit angle is measured based on the impact point on both the tracking layers. The ALPIDE chips present on the tracking layers are also manufactured with a slimmer 50 µm thickness, meaning they are twice as slim as the conventional calorimeter layer ALPIDE chips. This is again to minimize the material budget.

Protons will exhibit a known maximum scattering angle given we know its minimum energy and what material it passes through. Figure 3.1 shows an example of particle scattering. By having 41 layers spaced with aluminum absorbers between each layer, the DTC can determine how much energy is left in the particle when it exits the test subject.

Figure 3.1 shows how an example of how a particle can move through the layers of a digital calorimeter. The exit angle out of the subject is compared with the position it was sent from, creating a detailed map of the test subject. This has to be done thousands of times to get a good model, as the particle scattering angle exhibits a Gaussian distribution as discussed in section 2.5.

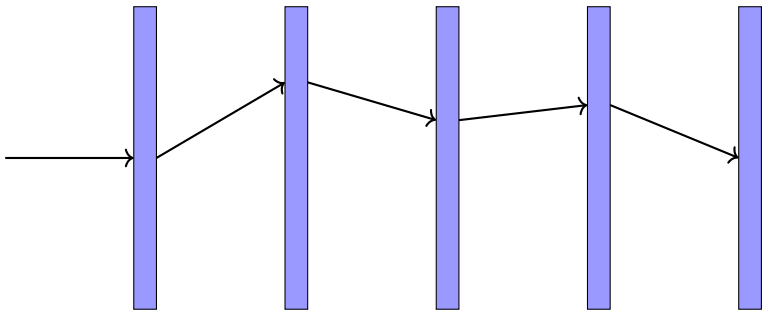

Figure 3.1: Example of particle movement through the calorimeter layers.

The layers have 12 ALPIDE strings, and each string consists of 9 ALPIDE chips connected horizontally, creating a 12 x 9 grid of ALPIDE chips in each layer. Each layer is built up of two half-layers, with each half-layer hosting six ALPIDE strings. One half-layer is mounted front-to-front onto the other, meaning one half-layer is rotated 180 degrees. Since each half-layer is covered with 50% ALPIDE chip and 50% flex cable, the face-to-face orientation ensures that the chips cover the entire layer area. Figure 3.2 shows an ALPIDE mounted on an aluminum absorber.

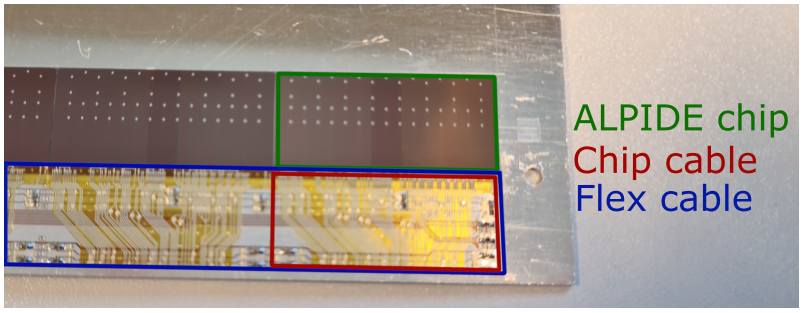

Figure 3.2: View of the ALPIDE chips.

The ALPIDE chips are located on an FPC. Due to their complexity, mounting them on a copper printed circuit board (PCB) is difficult, but it would also increase the material budget. The TC is connected as close to the ALPIDE chips as spacially possible to reduce the flex cable length. Due to radiation, the readout electronics are not placed close to the DTC. The FPC also uses aluminum traces to ensure that the entire DTC is homogeneous in its scattering property. Aluminum has a radiation length almost six times that of copper [24][3]. Aluminum is also beneficial to copper with its high corrosion resistance. However, it is both more brittle and

[3]The radiation length of a material is a characteristic relating to the energy loss high energy particles have when electromagnetically interacting with it.

has a higher electrical resistance, making the need for the TCs to be close to the ALPIDE chips even greater. Figure 3.3 shows that each string is connected to the TC individually through the flex cables.

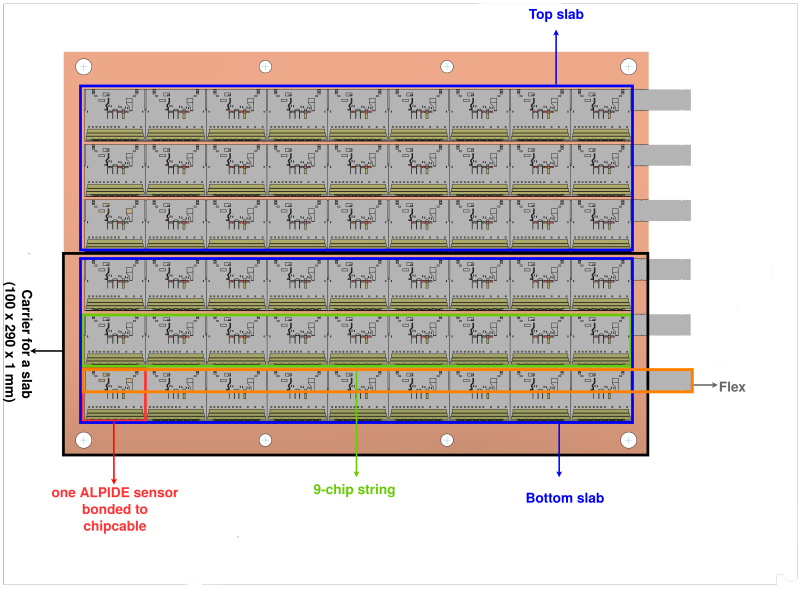

Figure 3.3: View of a pCT half-layer.

3.1 The ALPIDE Monolithic Active Pixel Sensors

CMOS MAPS was developed in the early '90s for the detection of visible light, since then it has become the dominant technology within the imaging market [25]. Since the '90s there have been developed several more specialized CMOS sensors, and the ones used in the pCT project was originally designed for the ALICE Inner Tracking System (ALICE ITS).

ALPIDE chips served as an upgrade to ALICEs ITS, mainly to improve read-out rate, but also to achieve greater resolution in terms of position and momemtum [26]. The ALPIDE chip is 1.5 cm × 3 cm large with a resolution of 512 ×1024. The area of each chip is 28 µm × 28 µm.

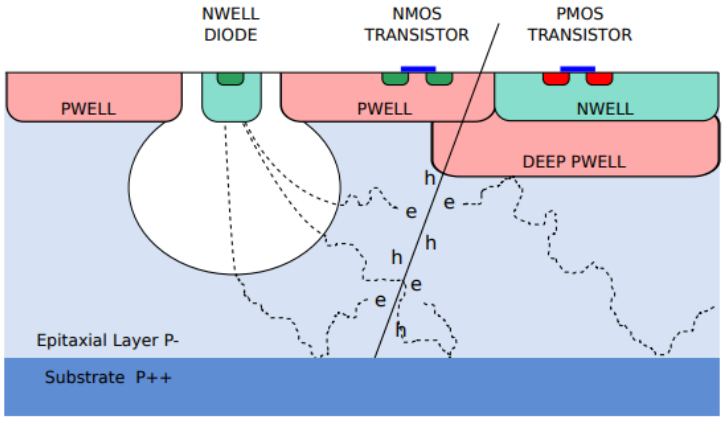

Figure 3.4: Cross section of an ALPIDE pixel as a particle passes through [27].

When a particle traverses the sensor's active volume, it frees charge carriers in the form of electrons and holes. Electrodes pick up the released charge, and the pixels will only report if a charge above the threshold level is registered, as they operate in a binary hit-/no-hit manner. The more energy a particle has will have an effect on the number of pixels triggered on a single layer. This is called cluster size. A large cluster denotes that a proton with a lot of energy traveled the sensor's active volume. The amount of charge needed to register a hit, the threshold voltage, can be altered digitally. The triggered pixels stores the charge until it is read out.

Inside each pixel of the ALPIDE chips there is a sensing diode (the NWELL diode in Figure 3.4), a front-end amplifying/shaping stage, a discriminator, and a digital stage. The sensing stage will forward any current it registers through to the amplifying and shaping stage. If the resulting voltage level rises above a set threshold, the chip stores a hit in the registers and waits for a readout. If the chip is measured to be noisy, the threshold can be adjusted to accommodate this. Changing the threshold voltage and back-bias alters the chip's sensitivity to background noise. The digital stage stores either a one or a zero based on if there was a hit or not.

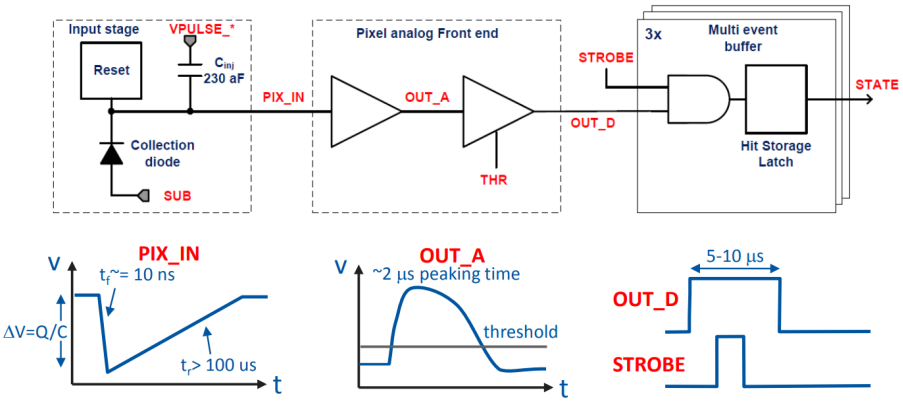

Figure 3.5: Block diagram of the ALPIDE pixel cell

Figure 3.5 shows the three different stages. The OUT_D signal stays high for about 5 µs to 10 µs. The ALPIDE chips also store up to the 3 last hits in the multi-event buffer [28].

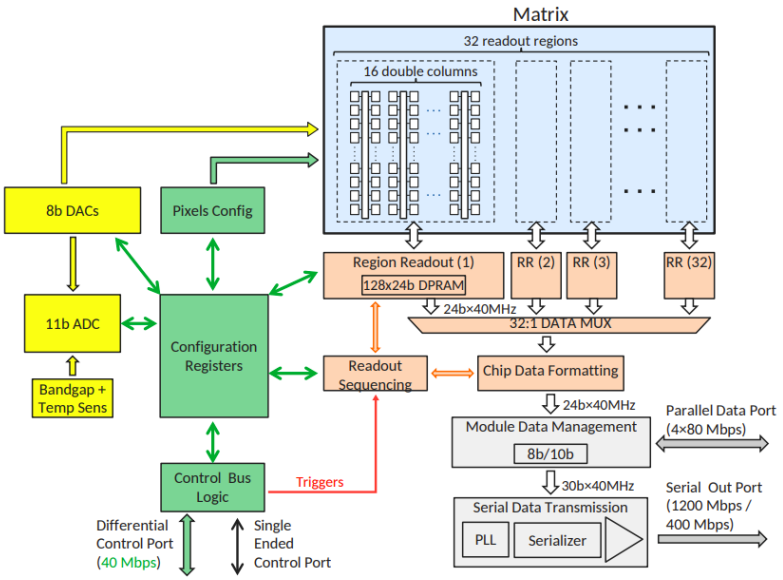

Figure 3.6: ALPIDE chip block diagram [28].

The chips feature an onboard 8b/10b decoding to achieve DC balance[4]. The serial data transmission unit sends data to the readout unit, which talks to a control computer through IPbus.

3.1.1 ALPIDE classification

The electrical function of the ALPIDE chips is tested with an automated test system using a needle probe card before being used by any institute. The test covers basic functions such as checking power draw, and more advanced functions such as testing the threshold values by measuring the S-curve of each pixel.

The tests divide the chips into two categories, OK and NOK (not-OK). After a chip is deemed OK, it is further given a GOLD, SILVER, or BRONZE classification. GOLD is given to the best performing chips and BRONZE is given to the worst-performing chips. For example, a chip with less than 50 broken pixels is given a GOLD classification. A chip with between 2100 and 5243 broken pixels is given a BRONZE classification. If a chip has fewer working pixels than BRONZE it is deemed NOK.

The classification of the chips is used as a baseline later on in Section 5

[4]DC balance can incorrectly charge a coupling capacitor to an unwanted level, which can have a detrimental effect on asynchronous processes as it makes it challenging to execute clock recovery [29].

3.1.2 ALPIDE Powernets

An ALPIDE chip has three power domains that require as stable and noise-free supplies as possible. The domains are explained as follows in the ALPIDE Operations Manual [28]:

AVDD, AVSS: Supply and ground nets of the analog domain. This includes the pixel front-end circuits, the analog biasing circuits, and the ADC block.

DVDD, DVSS: Supply and ground nets of the digital domain. This includes the in-pixel configuration registers, the matrix readout circuits, the peripheral readout circuits, and the chip input and output buffers and transceivers.

PDVDD, PDVSS: Supply and ground nets are exclusively dedicated to the Phase Locked Loop of the Data Transmission Unit.

PWELL: Bias of the p-type wells in the pixel matrix region.

SUB: Bias to the substrate contacts in the seal ring and the periphery region

We have access to three nets from the TC: AVDD, DVDD, and PWELL. Of the three nets, only PWELL should be modulated, both AVDD and DVDD have a set target voltage. PWELL and SUB are internally connected through the conductance of the die substrate, and changing the PWELL voltage changes the back-bias voltage on the charge collecting diodes. Increasing the back-bias voltage on the diodes increases the depletion region, affecting the cluster size of a hit [30]. Unless the application requires enhanced sensor performance, these nets can be shorted to analog ground, but in the scenario where it is needed, PWELL can be regulated down to $-6\,\mathrm{V}$ with reference to AVSS. Increasing the back-bias too much can increase the chip's susceptibility to noise in the substrate substantially [30], which is why a regulatable voltage is ideal during calibration.

3.2 Transition Card

The TC is a custom PCB made by Tea Bodova at UiB. The board has a custom stackup to minimize the material budget as much as possible. The name accurately describes its core function, as it transitions the connection of the ALPIDE strings from flex cable to FireFly.

The board has 12 FireFly connectors and 12 zero injection force (ZIF) connectors. On the board is 24 MIC low dropout voltage regulators, which regulate the power to both the DVDD and AVDD power domains of each ALPIDE string. The TC receives power using a Weidmuller 7 pin connector, rated for $1\,\mathrm{mm}^2$ cross-sectional wires.

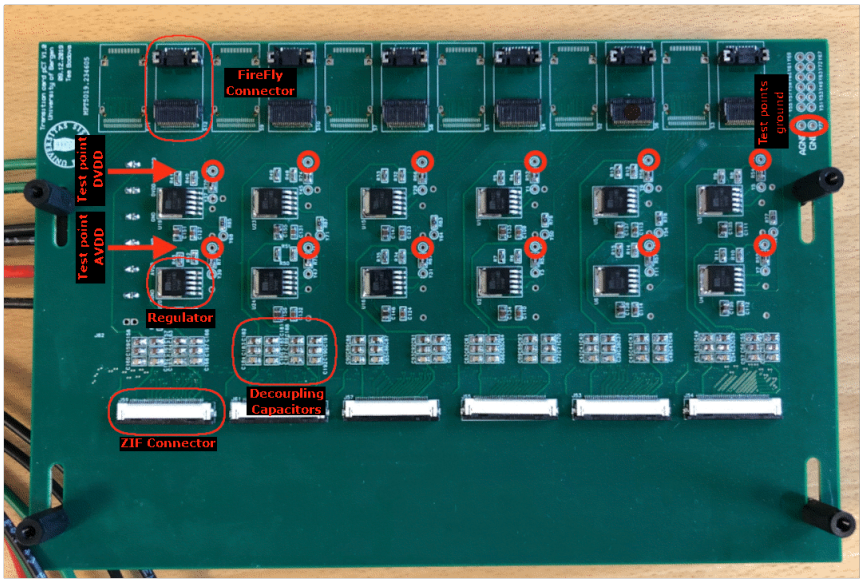

Figure 3.7: Transition card. The ALPIDE chips are connected to the readout electronics using the ZIF and FireFly connectors [10].

Figure 3.7 shows version 1 of the TC. A total of 12 ALPIDE strings connect to the board, 6 on each side. The power integrity will later be measured across the decoupling capacitors of the TC.

3.3 IPbus at pCT

IPbus was developed by J. Mans et al. in 2009, and is a tool to monitor and control field-programmable gate array (FPGA) hardware from software on a computer [31]. IPbus is, as the name suggests, an IP-based protocol, meaning that all hardware that is to be controlled is given an IP address and connected to some subnet. When requesting to read or write data to a device, we communicate with it using the preassigned address. Hundreds of devices can be controlled by a single computer host. IPbus can also be used point-to-point style by connecting the master directly to a single slave.

IPbus is already the tool used for communication between a client software and the readout electronics, the chips can be controlled and monitored in detail from a computer. The power monitor board will also communicate with a computer, and it is preferable that this communication also happens over IPbus to have uniformity in the pCT designs.

3.4 Current Power Control Unit

A conceptual design of the power delivery and monitor system was made by Tea Bodova. The unit was dubbed the power control unit (PCU), and an outline of its specifications was determined. The PCU was going to monitor the currents of the

different domains on the ALPIDE chips, record the temperature of a temperature sensor, and control the enable signals of the regulators.

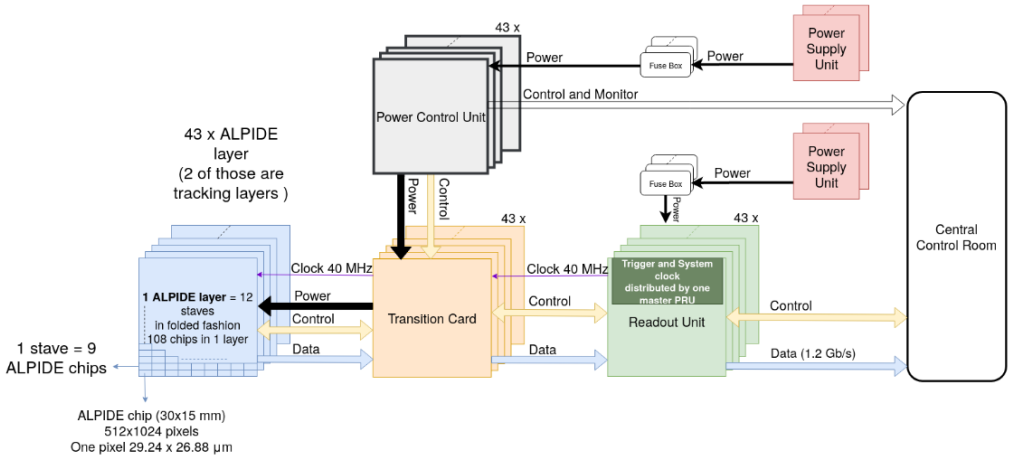

Figure 3.8: Current pCT system as proposed by Tea Bodova.

The current pCT PCU proposed by Tea Bodova is shown above, it suggests that several large PSUs should feed power to a fuse box, then the TCs are controlled and powered through the PCU. A more advanced version was later proposed by Frode Løvbrekke Aase at UiB, shown in figure 3.9. The new design moved away from having one PCU per TC, and instead used an FPGA with large IO banks to control and monitor several TCs at once.

The design would use a Xilinx FPGA mounted on a Trenz breakout module. The FPGA layout procedure would be simplified, as the FPGA power delivery and decoupling would be taken care of. A total of nine layers would be monitored and controlled from each PCU. Each analog and digital domain would be monitored using an INA current monitor, totaling 27 ICs that the FPGA would need to communicate with.

Five large power supplies would supply 43 layers. The power draw of the DTC is 2.5 kW, meaning at 5 V, each supply would need to deliver at least 100 A of current. The MSP-600-3.3 from Mean Well was proposed as a potential PSU, but with a noise of 120 mV ripple on the supply line, it would also need a substantial filter. Neither a filter specification nor the voltage drop across the cables had been outlined in this solution.

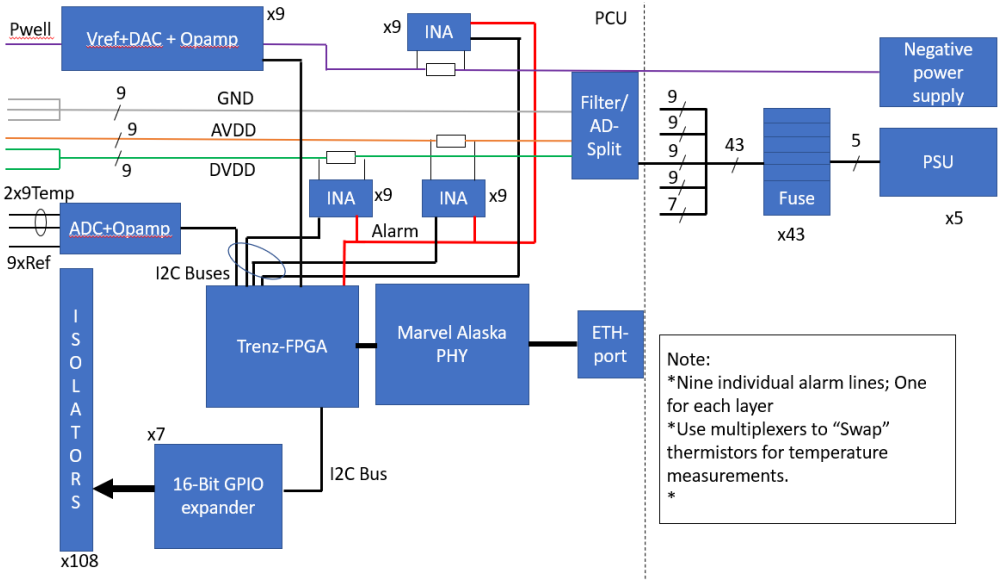

Figure 3.9: Theoretical design of the power control unit (PCU). Made by Frode Loevbrekke Aase

Several different solutions for the PWELL voltage were explored, some using resistance to digital (digital potentiometer) to control a negative voltage regulator, and others explored having a dedicated power supply specifically to generate a negative voltage. The latter plan was exploiting that the PWELL powerline sinks current in the milliampere range, meaning a single low noise supply could be split to supply the back-biasing of all the layers with ease.

IPbus needs to be connected through an ethernet interface and therefore requires some physical layer (PHY) to function. A PHY handles the communication protocol over ethernet or fiber, and the MCU/FPGA/ASIC communicates with the PHY over some media-independent interface (MII) standard. The MII is standardized and having a PHY simplifies the design process down to just choosing an MCU/FPGA capable of handling MII, which should not be an issue.

Between the power control unit and the TC there would be relatively long cables of up to five meters, meaning there would be a significant difference in ground potential. Optocouplers would act as an isolated switch to control the enable signals, of which there would be 108 in total. The optocouplers would be controlled using 16-channel GPIO expanders, reducing the number of pins needed.

3.5 New Power Control Design

Having a couple of large PSUs to supply several layers is a good design in principle, but difficult to realize. No PSU was found to satisfy the low noise requirement while also supplying enough current. FPGA layout and design are also time-consuming, meaning the scope of the assignment was quite large from the get-go.

The new design proposed will use a single smaller form factor PSU to power each

TC individually. The smaller PSUs will all have their own custom PCB, a monitor board (MB), accompanying them, filtering the supply lines, monitoring current, and controlling the enable signals on the TC. Instead of using an FPGA, the MBs use a microcontroller unit (MCU) instead. Since there is still a desire to use IPbus for control and communication, creating an IPbus translation layer has been outsourced as a bachelor's thesis to the Western Norway University of Applied Sciences.

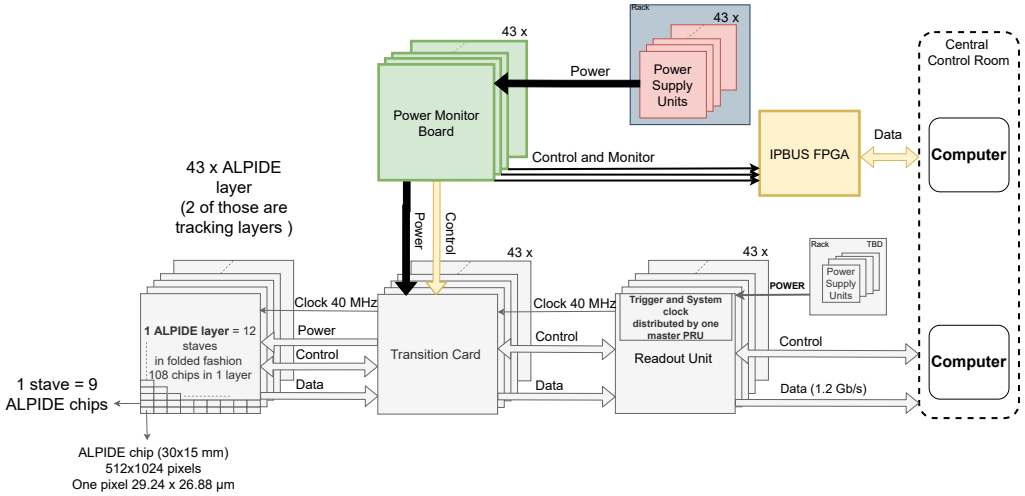

Figure 3.10: Proposed changes to the pCT ecosystem highlighted with colors.

Figure 3.10 has the proposed changes highlighted with color. The 43 MBs talk to an FPGA using a universal synchronous/asynchronous receiver-transmitter (USART) protocol. The FPGA system is outlined in Section 3.5.1, and the communication protocol is outlined later in Section 6.2.2. The "Control Room" square has also been divided into two individual computers, as there is no requirement for the readout electronics and the MB to talk to the same system. The computer hosting and displaying the data collected from the MBs are outlined in Section 3.5.2. The new power control design chain is visualized in Figure 3.12 at the bottom of this chapter.

Breaking the PSU solution into one PSU per layer opens up the design to utilize cheaper and smaller supplies. Having only one MB per layer also requires substantially fewer IO pins on the controller, so a change to an MCU from an FPGA was natural. Since many MCUs have built-in peripherals such as DACs and ADCs, this will reduce the need for external ICs to handle temperature measurements. The switch MCU will be cheaper, which is very beneficial for both prototyping and production at a later stage.

3.5.1 IPbus Translation Layer

A Xilinx Kintex Ultrascale KCU115 evaluation kit is used for the IPbus translation layer. The IPbus translation layer communicates with the MBs over differential USART and relays the information to the database over IPbus. The FPGA design is made by Martin Eggen and Jakob Hauser at Western Norway University of Applied Sciences.

A breakout board containing 43 RJ45 connectors connects all the MBs to the KCU115. Each RJ45 connector uses one twisted pair for the data and one for the clock, using differential signaling for communication to eliminate the need for common ground. The KCU115 functions as the master for all the communication, and the MCU will never initiate any data transaction by itself.

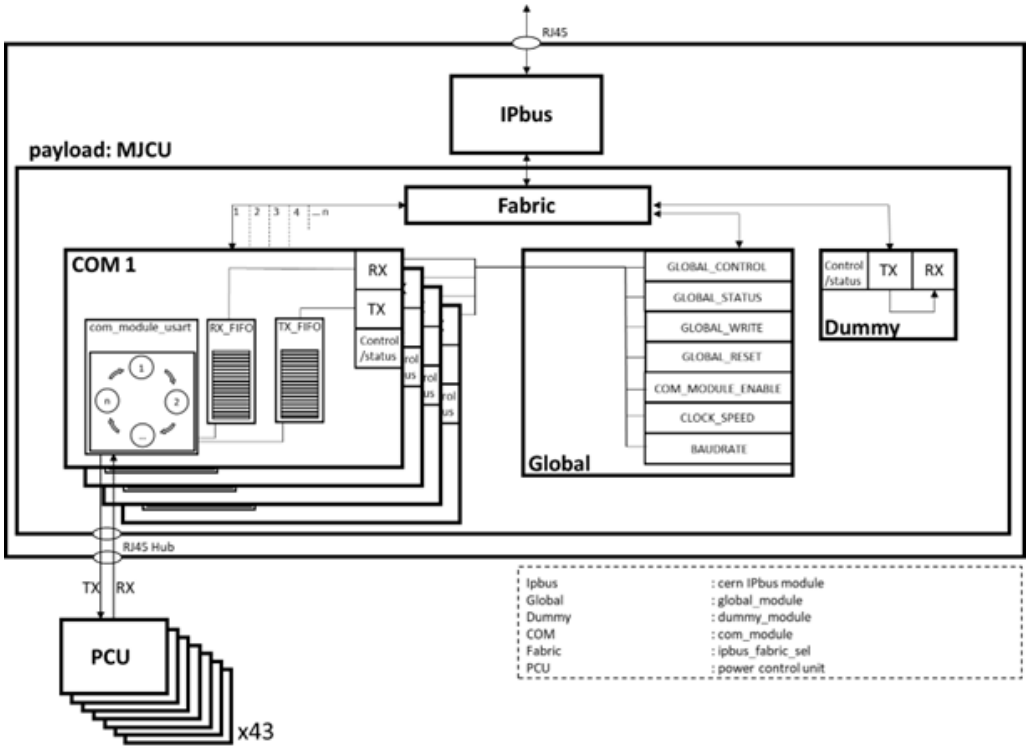

Figure 3.11: Block diagram of the FPGA translation layer. The module named "PCU" is the MB in this context.

The KCU115 has a set of registers that trigger different functions. Figure 3.11 shows that data read is kept in a FIFO register, temporarily storing data before the database reads it. The database can also send multiple bytes of data to the TX_FIFO, which will be written one after another when the MCU on the MB catches up.

3.5.2 Communication and Control Database

A database system talks to the KCU115 over IPbus. The database has a graphical user interface (GUI) to easily write and read from the MBs. The data is visualized over a timeline using Grafana. Figure 3.12 shows a simplified overview of the connection from the DTC to the control room.

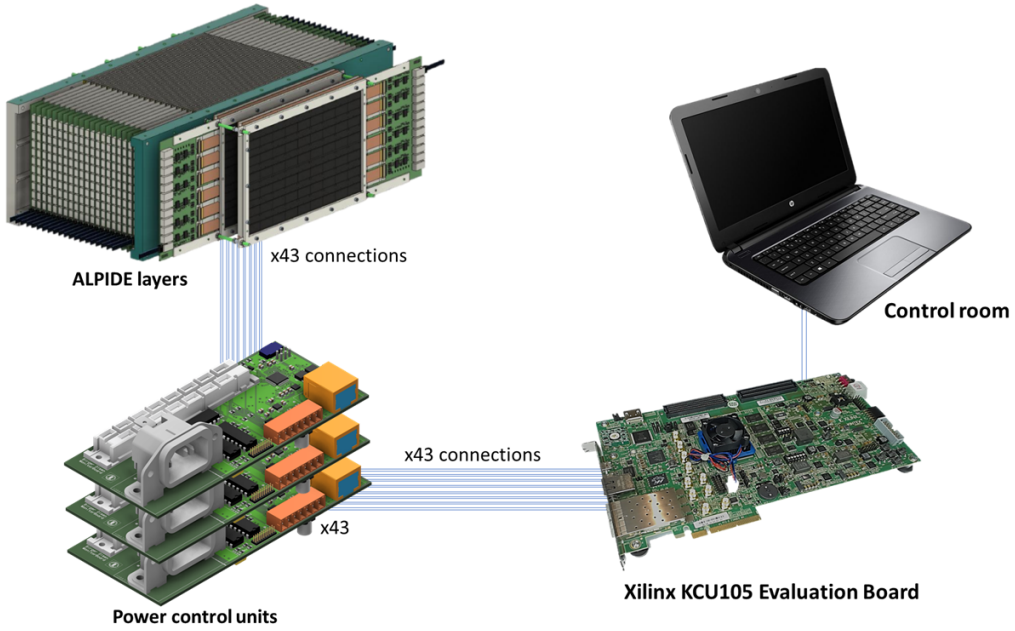

Figure 3.12: Power control system overview. Overview provided by Jacob Hauser and Martin Eggen.

The database GUI will be used to handle the triggering of all soft startup and shutdown procedures, making it vital to design the MB in conjecture with the database. Using a set of predetermined values for all the writeable registers in the MB, the database will globally send these out to all the MB during configurations, then it can initiate a soft startup and diagnostics on all layers. The register values should be customizable by the user if any settings change is wanted. The values control the maximum allowed current draw before shutdown, the PWELL signal voltage, the maximum allowed temperature, and more described in Appendix B.1.

4 HARDWARE DESIGN CONSIDERATIONS

This chapter will focus on considerations around building a real-world system. Conventional PCB layout methodology will be outlined, and highlighting how it is relevant for the pCT project. The content will be relevant for the later creation of a custom PCB, and to the creation of the PSU filter.

> "If you ask ten people a question about why you have EMI in your circuit, you will get ten conflicting opinions and a lot of debate."
>
> *Anonymous internet user*

4.1 Wire Inductance

Any piece of wire has a set amount of resistance which is important to account for when designing your system. A thicker power cable leads to lower resistance, resulting in lower heat dissipation and voltage drop across the cable. But having long wires does not only come with the drawback of added resistance, they also inherit a significant amount of inductance.

The inductance of wiring is an important factor for both high-speed design, and circuits where the rate-of-change of current, called di/dt current, is high [32]. Doubling the cross-sectional area of a wire will halve the resistivity, making it ideal to choose thicker cables when dealing with large currents. Increasing the cross-sectional area has a much smaller effect on the inductance, where length is the primary contributor.

$$\mathrm{L_{wire}} = 2\mathrm{l}\left(2.303\log\left(\frac{4\mathrm{l}}{\mathrm{d}}\right) - 1 + \frac{\mu}{4} + \left(\frac{\mathrm{d}}{2\mathrm{l}}\right)\right) \tag{4.1}$$

Going from a wire size of 0.2 mm to 1.6 mm will reduce the resistance from about 0.5 Ω/m to 8.5 mΩ/m, but using Equation 4.1 shows that inductance only decreases from 1.831 µH/m to 1.415 µH/m [33]. That is a 98% decrease in resistance, but only a 23% decrease in inductance. Changing from one wire to either several wires, or a single cable[5], will greatly decrease the overall inductance.

$$\mathrm{L_{total}} = \frac{1}{\frac{1}{\mathrm{L_1}} + \frac{1}{\mathrm{L_2}}} \tag{4.2}$$

As shown in Equation 4.2, having two identical inductors in parallel yields half the original inductance, and a cable usually has tens of wires, which also increases the flexibility.

[5] A cable is per definition several wires.

When designing a system, it is important to remember how it will be powered, or in the case of designing a power delivery system, how it will power other systems. When powering something with a single meter of standard AWG 12 wire ($2\,\mathrm{mm}^2$), the resistance of the wires is negligible, but the inductance is $1.415\,\mu\mathrm{H}$. If this system were to have a $1\,\mathrm{A}$ load which switched on in the span of a $1\,\mu\mathrm{s}$, the resulting voltage across the wires alone would be almost $1.5\,\mathrm{V}$ as shown in Equation 4.3.

$$\begin{aligned} V_L &= L \cdot \frac{di}{dt} \\ V_L &= 1.415\,\mu\mathrm{H} \cdot \frac{1\,\mathrm{A}}{1\,\mu\mathrm{s}} \\ V_L &= 1.5\,\mathrm{V} \end{aligned} \tag{4.3}$$

For the pCT layers, the digital domain is powered using four cables, two positive and two return cables. The use of several cables negates the inductance problem so much that it ceases to be a problem. A load step response is measured in section 5.3.4, but the timescale is in the millisecond range, and the inductance becomes less relevant.

4.2 Ground Bounce

Ground bounce is a common issue in digital electronics, and shortly put, it stems from a surge in current. An alternative name for ground bounce is induced switching noise since it is often created when large amounts of logic gates switch in sync with the clock [32]. However, ground bounce is not limited to logic gates switching, in fact it can occur when any amount of current surges through the ground plane.

Decoupling is of little use with this issue, as logic gates often have a rise time in the nanosecond range, putting the noise frequency in the gigahertz range. Not only will a decoupling capacitor not help smooth out a ground bounce, but it will also act as an inductor, possibly making the issue worse. This effect, along with decoupling capacitors, will be discussed in detail later.

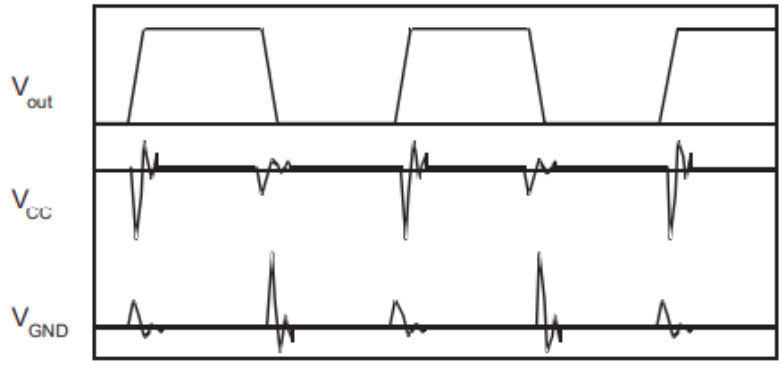

Figure 4.1: Switching noise on power and ground rails [32].

The decoupling capacitor's primary function is to maintain a low dynamic impedance from the supply to ground in a localized area [32], and even though the capacitors are

limited by bandwidth, a high bandwidth interplane capacitance is achievable with a proper stackup. By lowering the impedance of the ground plane, the resulting ground bounce will be diminished, as it is heavily dependent on the resistance of the conducting material of the ground. Having a dedicated ground plane in your PCB is the easiest and best way to lower the impedance.

4.3 Crosstalk

Wires closely packed together will have some capacitive coupling to each other, the tighter the packaging, the larger the coupling. Such a capacitive coupling can create unwanted effects on the signal line. This is called crosstalk [34]. When wires next to each other stay static for some time, the capacitance between them gets charged to some voltage. If one wire is logically high and one is logically low, the voltage difference equals the supply voltage. If the two wires switch opposite of each other, the change in voltage, ΔV, will be twice the supply voltage. The capacitive coupling needs to be charged to this new value for the signals to change, effectively delaying the switching.

Section 4.6 will go into more depth about how the electromagnetic interference (EMI) from PCB traces can affect each other.

4.4 Common Mode Noise

Common mode noise is present on both supply lines simultaneously, not able to be measured by a two-probe setup as there is no change in voltage differential. Lightning strikes are one perpetrator of common-mode noise, as they can create a significant surge of power on both supply lines of a household. A less extreme cause of common-mode noise is having no proper path to ground in a transformer. If a PSU does not have a proper path to ground for stray pickup/capacitances to discharge their unwanted parasitic energy, it will propagate through the system.

In theory, a signal's supply and return line should be perfectly balanced, as all things should be, but this is extremely difficult to achieve. When common mode noise is present on an unbalanced supply/return path, it turns into transverse mode noise [35]. The better way to combat common mode noise is to filter it out, ideally by providing a proper path to ground for it to dissipate.

4.5 Thermal Relief

Through-hole components are often connected directly to power planes, like VDD or GND. Though less usual, SMD components can also be directly connected to a copper pour[6].

[6] A copper pour is an area that is filled with copper. Copper pours are usually used to create power and ground layers.

The solder must be melted to properly flow onto the component lead and into the PCB hole to properly connect a component to the PCB. By having a pin directly connected to a large copper plane, the plane acts as a heatsink, and properly melting the solder becomes tedious and difficult by hand. In addition, having to heat the area around the solder pin will also cause the solder to flow away from the pad onto the pour, creating a bad connection [36].

To circumvent this, thermal ties[7] is used. As shown in Figure 4.2, thermal ties create a guard ring around the pin or pad in question, then connect it to the power plane using one or several trace spokes. The spokes are slim and short, meaning they can carry much current without generating much heat.

A choice has to be made during the design process of any PCB regarding the tradeoff between thermal and electrical connection. Slimmer spokes yield better thermal relief, but worse electrical connection.

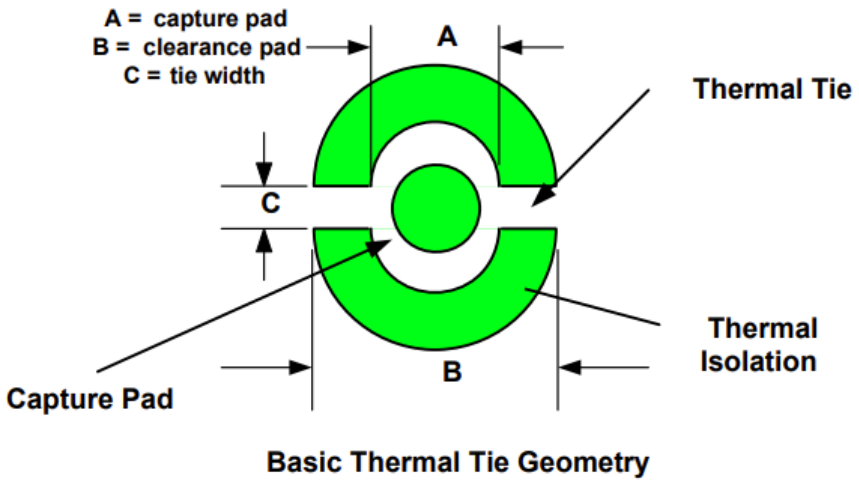

Figure 4.2: A typical thermal tie in a power plane. [36]

[7]Also called thermal reliefs.

4.6 Signal Integrity

Stackup is a large part of signal integrity, and there is a large volume of literature on it. And whilst you could go in-depth on hundreds of different aspects of PCB design, it is essential always to have a clear overview of your PCB to focus on the critical points. Electromagnetic Compatibility Engineering by Henry Ott [37] states six design objectives that should be kept in mind when designing a multilayer PCB. The list is ordered from most to least critical.

1. A signal layer should *always* be kept adjacent to a plane.
2. Signal layers should be tightly coupled (close) to their adjecent planes.
3. Power and ground planes should be closely coupled together.
4. High-speed signals should be routed on buried layers located between planes. The planes can then act as shields and contain the radiation from the high-speed traces.
5. Multiple-ground planes are very advantageous, because they will lower the ground (reference plane) impedance of the board and reduce the common-mode radiation.
6. When critical signals are routed on more than one layer, they should be confined to two layers adjacent to the same plane. As discussed, this objective has usually been ignored.[8]

EMI is a leading cause of high-speed issues in PCB design, it can either be picked up from external sources or it can be produced locally. Henry Otts's list mentions radiation several times, and each point on the list improves EMI performance. As a signal travels through a PCB trace, a magnetic field is created perpendicular to the current flow [38], affecting any nearby signals. If the trace is on the top of the PCB, this magnetic field can also have an effect on external electronics. By having these signals located on an internal layer close to ground planes, most of the EMI the signals are emitting will be absorbed, reducing its effect on nearby traces and electronics. Figure 4.3 shows that fringe fields can branch out and effect other traces.

[8] Discussed in Electromagnetic Compatibility Engineering p. 638 [37]

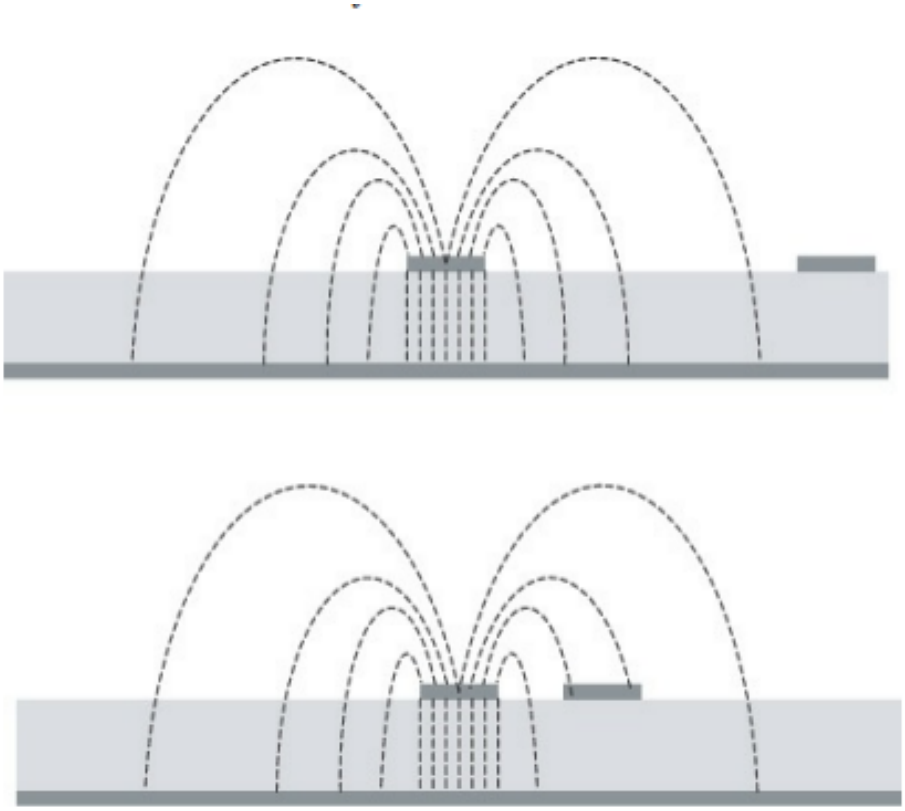

Figure 4.3: Fringe fields from a trace being absorbed by a ground plane. [38]

Interplane Capacitance

Having power and ground layers as close together as possible in the stackup will also provide interplanar capacitance. Remembering back to decoupling capacitors, their main function is to provide a low impedance path from the supply to ground. A large interplane capacitance provides such a path across the entire PCB.

Figure 4.4 shows how steep the capacitance diminishes when the space between layers is increased. Digital designs which utilize any form of digital logic must be wary of bad stackup, as even though you might not have any high-speed signals, rise and fall times have to be met regardless.

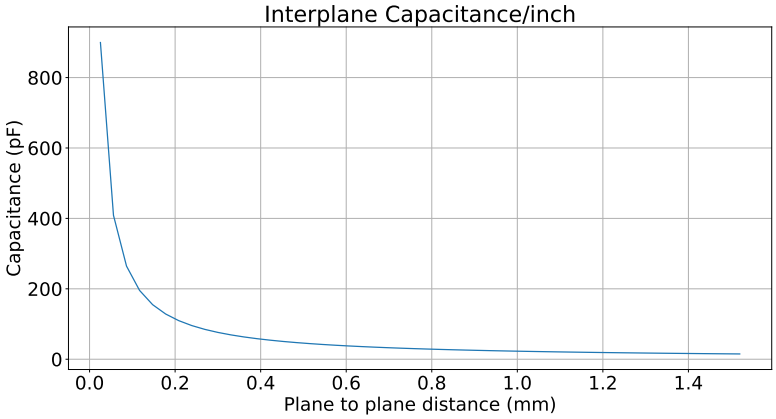

Figure 4.4: Interplane capacitance as a function of distance between the planes [36].

Two power planes are effectively two copper plates with some dielectric in between, which generate a small bit of capacitance. The relative permittivity of FR4, the most common dielectric, is 4.4. The capacitance between two power planes is given by the total overlap area of the planes, the distance between, and the relative permittivity of the material in between.

$$\mathrm{C} = \frac{k\epsilon_0 A}{d} \tag{4.4}$$

The PCBs for high-end products are rarely two-layer boards because routing while also providing power and return paths is too difficult. Ideally, you would want ground around, above, and below a signal trace, but this would require very many layers if you have a complex design. You can customize the stackup spacing between layers quite extensively, but the more you customize, the more it will cost. A standard six-layer PCB has an even spread of 11th (0.28 mm) of FR4 dielectric between each copper layer.

Looking at figure 4.4 and using Equation 4.4, having 0.254 mm between two power planes would mean an interplane capacitance of below 100 pF per square inch. To increase this value, copper pours can be utilized. Filling the unused space in a signal layer with copper makes the PCB function as a multi-plate capacitor. A multi-plate capacitor increases the area by (number of plates - 1) times, meaning you get five times more interplane capacitance out of a six-layer PCB. Important to note is the fact that this is difficult in complex designs which are routing constrained. The top and bottom layers are often crammed with components, making it even more challenging to get a good pour. If there is enough space to stitch the different layers together with vias, then having alternating power and ground is good practice.

"Right the first time" [36] by Lee W. Ritchey has an example of a six-layer board before and after filling copper in the unused routing areas. The board experienced EMC issues before getting a revised copper pour, and had an inter-plane capacitance

of 500 pF. The revised board had 4 nF inter-plane-capacitance, eight times more than the original one, with an extremely quick and easy fix.

5 THE POWER SUPPLY

This chapter first determines how much power is needed to be delivered to each TC, then outline the maximum voltage that can be delivered and why it matters. Later the two main categories of power supplies will be discussed, highlighting the pros and cons. A power supply is chosen for testing and prototyping based on this discussion. The ALPIDEs have been tested using this new PSU. The tests are explained, and the results are shown.

Finally, several filters are proposed and tested to best determine the classification of noise and thereby determine the final solution.

5.1 ALPIDE Power Consumption

An ALPIDE chip with GOLD classification should never draw more than 50 mA digital current and 16 mA analog current, but having all chips be GOLD classified is not always doable. The power consumption of an ALPIDE string with unknown classification was tested during high data transfer, and a maximum digital current draw of 880 mA per string was recorded at UiB [10], a maximum analog current draw of 200 mA is also expected. As a consequence, the TC has as a requirement to deliver 1.9 V and 1.1 A of current to each string in a 12-string layer. The TC uses linear voltage regulators to shift the voltage down from 3.3 V to an appropriate 1.9 V, which means the overall power needed is $3.3\,\mathrm{V} \cdot 1.1\,\mathrm{A} \cdot 12 = 43.5\,\mathrm{W}$. Overall, each TC needs at least 45 W of power. Developing a power supply solution with exactly the correct voltage might become difficult as a voltage drop will occur over the long cables connected to the transition cards from the PSU. The voltage drop will also differ based on the current draw.

There is a slim chance that all the chips in a layer will be drawing their maximum rated current, but designing the system around this max power draw gives us a natural power overhead during system idling.[9]

5.2 Power Delivery Considerations

Choosing a power delivery system has several solutions, and they all have their benefits and tradeoffs. The power supply can either be custom or pre-built, linear or switch-mode, you can have one large form factor or several smaller form factor units. The only requirement is that the PSU should deliver at least 3.3 V and 13.2 A of current in total to each layer. 13.2 A is more than what we expect the ALPIDEs to draw during normal operation, but is still considered the minimum to ensure we have some power and thermal leniency.

[9] During testing done by the ITS team at CERN the chips should draw at most 50 mA per digital chip during idle powering, equating to a power draw of 650 mA per string. Idle analog current draw was not mentioned.

5.2.1 Transition Card Voltage Regulators

The MIC29302A linear voltage regulators placed on the TCs are accurate and provide little fluctuation even when supplying the maximum DVDD load of 0.9 A per string [10]. However, linear voltage regulators are inefficient, dissipating all excess energy as heat as shown in Equation 5.1.

$$\mathrm{P_D} = \mathrm{I} \cdot (\mathrm{V_{in}} - \mathrm{V_{out}}) \tag{5.1}$$

The transition cards have 7 cm of free air on each side. If the TC gets 5 V while being hit by a full load, there will be 33.5 W of energy dissipated as heat across the 12 DVDD regulators. Each regulator has to then dissipate 2.8 W. The MIP29302A specifies a power dissipation a bit higher than the ideal formula as shown in equation 5.1.

$$\begin{aligned} \mathrm{P_{wasted\text{-}DVDD}} &= \mathrm{I} \cdot (1.05 \cdot \mathrm{V_{in}} - \mathrm{V_{out}}) \\ \mathrm{P_{wasted\text{-}DVDD}} &= 0.9\,\mathrm{A} \cdot (1.05 \cdot 5\,\mathrm{V} - 1.9\,\mathrm{V}) = 3.015\,\mathrm{W} \end{aligned} \tag{5.2}$$

With a power dissipation of 3.015 W the junction temperature, $\mathrm{T_J}$, can then be calculated given we approximate the air temperatue. The thermal resistance between the junction and air, θ_{JA}, is given as 35 °C/W.

$$\begin{aligned} \mathrm{T_J} &= (\theta_{\mathrm{JA}} \cdot \mathrm{P_{wasted\text{-}DVDD}}) + \mathrm{T_A} \\ \mathrm{T_J} &= (35\,^\circ\mathrm{C/W} \cdot 3.015\,\mathrm{W}) + 25\,^\circ\mathrm{C} \\ \mathrm{T_J} &= 130.5\,^\circ\mathrm{C} \end{aligned} \tag{5.3}$$

If we approximate an ambient temperature of 25 °C, as done above, the junction temperature would be 130 °C. The datasheet specifies 125 °C as the maximum junction temperature, and to adhere to this, the power dissipation would have to be no higher than 2.57 W [39].

The above calculation of a 3.015 W power dissipation is an absolute worst-case scenario. A more accurate result is calculated by factoring in that the PSU output can be regulated and that there will be a voltage drop across the cables. The voltage drop occurring across the cables will be a couple of hundred millivolts, dropping the power delivered. For example, assuming a 1.5mm^2 cable with 0.0133 Ω/m and five meters of cable. The DVDD lines use two cables, additionally halving the total voltage drop.

$$\mathrm{V_{cable}} = 0.9\,\mathrm{A} \cdot 12 \cdot \mathrm{l} \cdot \mathrm{R_{cable}}/\mathrm{n} \tag{5.4}$$

In Equation 5.4, where l is the length of the cable and n is the number of cables, a total voltage drop of 0.36 V would occur across them. Power supplies often have a customizable output voltage inside a set range, generally +/- 10%, which would

further drop the voltage. Using the voltage drop calculated above and accounting for the customizable output, the regulators receive a voltage of 4.14 V, which requires them to dissipate 2.02 W for a short period. 2.02 W is below the 2.57 W requirement, and we can conclude that a 5 V PSU will not cause the voltage regulators to overheat.

The TCs receive active air cooling in the form of fans that blow air through the DTC chassis which have not been accounted for in the above section. The voltage regulators are also mounted flat onto the PCB, using it as a heatsink, helping even more with cooling.

5.2.2 Linear Power Supplies

Linear power supplies generate no ripple noise and have no switching components. A linear supply uses a transformer to lower the voltage by a desired factor depending on the number of secondary windings on the coil, the sine voltage is then rectified by a set of diodes, before being smoothed out by a large capacitor. Therefore, a linear PSU does not have the switching components employed by switch-mode PSUs and is very suitable for the ALPIDE chips. However, linear style power supplies are both costly and often physically large. There is also a potential 50 Hz noise being introduced as we rectify a pure mains signal, filtering a 50 Hz component of a signal is difficult as it would require very large coils and capacitors.

When dealing with scenarios where we need to regulate down from 220 V to 5 V, and supply tens of amps of power, it becomes difficult to use linear PSUs as it would generate a significant amount of heat, just as we witnessed earlier with the LDOs in Section 5.2. Looking at the linear power supply section of Mouser.com there is only a couple of PSUs for sale which could deliver 5 V and at least 13 A. The PSU proposed in the preliminary power delivery system only had 120 mV ripple, but that is still not low enough to not warrant a proper filter. The need to filter 50 Hz noise is also a drawback.

Having to regulate down from 220 V to 5 V would waste about 98% of the energy as heat, and using a switch mode PSU which regulates down using linear voltage regulators just shifts the issue instead of fixing it. Using a large linear power supply for this thesis is then difficult to achieve, and designing a custom linear PSU solution is a topic you could use years to master. Opting for a small form factor switch mode supply with a relatively low ripple will be the choice from here on out.

5.2.3 Switch Mode Power Supplies

Most power supply units are switch-mode units, which are very efficient, but also create a lot of noise due to fluctuations in the powerline created by the switching [37]. A switch-mode circuit often uses a transistor to regulate the duty cycle of an input signal. A capacitor then smoothens out the resulting square-wave signal in parallel with the load. The design of a switch mode PSU is talked about more in detail later, and shown in Figure 5.8.

Regulating voltage this way is efficient, and a good switch mode supply can have up to 95% efficiency. Unfortunately, these switch mode supplies generate much noise, often in the range of tens, if not hundreds, of millivolts. Any noise on the supply line can be wrongly recorded as a hit on the ALPIDE chips. To circumvent this, the chip's sensitivity can be turned down to be more resilient to noise, but less sensitive overall.

5.3 nVent Schroff MAX105

The MAX105 PSU manufactured by nVent Schroff was recommended by a supplier and should be capable of delivering a steady 5 V signal with relatively low noise under all loads. The voltage can be finetuned ±10%. The PSU is rated for a maximum of 16 A, meaning 2.8 A is left as overhead for the monitoring electronics and thermal overhead. The MAX105 comes in a standard 4U size, making it eligible for easy rack mounting. It features a large DIN 41612 style connector where a custom PCB is mounted in the final design.

According to the datasheet of the MAX105 PSU, it should have a switching noise of $>40\,\mathrm{mV_{p\text{-}p}}$, the datasheet does not mention the frequency of the switching components, so this must be measured.

Connecting the PSU to the ALPIDEs and running several tests is the best way of determining if the noise has any effect on the circuitry. The ALPIDEs have several tests specified in the ALPIDE user manual [28]. The MAX105 is compared to a linear lab bench power supply with a known low ripple output in the tests below.

5.3.1 Digital Tests

In theory, a switching power supply should not introduce any problems for the digital circuitry on the ALPIDE chips, this is easy to validate by running both a DAC test and measuring the signal integrity of the communication eye as done below.

DAC test

The DAC test tells the chip to read out predetermined data while reporting the number of occurring errors. The ALPIDEs line driver current settings are varied to give insight into which setting should be used, the ideal setting might vary on a chip-to-chip basis given the production variations that might occur. The DAC test is named after the DAC register in the ALPIDE chips which alter certain settings.

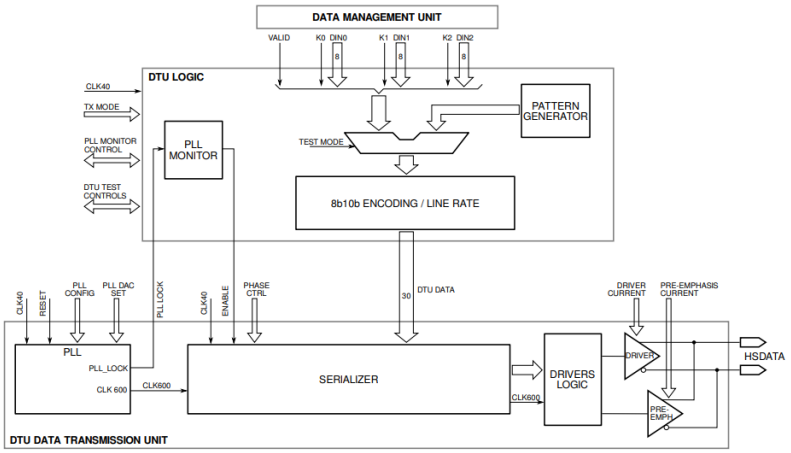

Figure 5.1: Functional block diagram of the Data Transmission Unit (DTU) and the related digital module DTU logic [28].

By writing to the DTU DACs register (0x0015), we can alter the line driver current setting and the pre-emphasis driver current settings, the drivers are shown conceptually in figure 5.1. The data is encoded using an 8b10b module, and when received by the readout module, it is decoded back to 8 bits. If the received 10 bits of data have five or more equal bits in a row, the test fails, as DC balance is not achieved. The amount of fails is reported in the decode test result. By running a DAC test, we are testing the signal integrity between the DTU and the readout unit. This test does not show the sensitivity of the chips, but it is important to check that the communication line is unaffected by ripple.

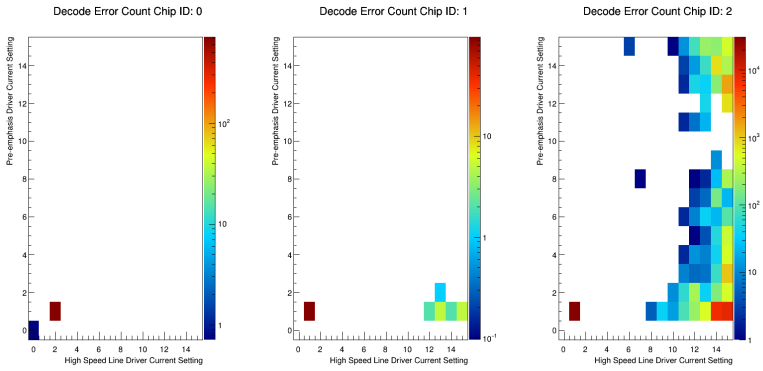

Figure 5.2: 48-hour DAC test decode errors using lab bench linear PSU. The left vertical axis shows the pre-emphasis setting. The bottom axis shows the high-speed line driver setting. The corresponding position on the plot will be colored depending on how many errors were recorded with those settings. The right vertical axis shows which color corresponds to the number of registered errors.

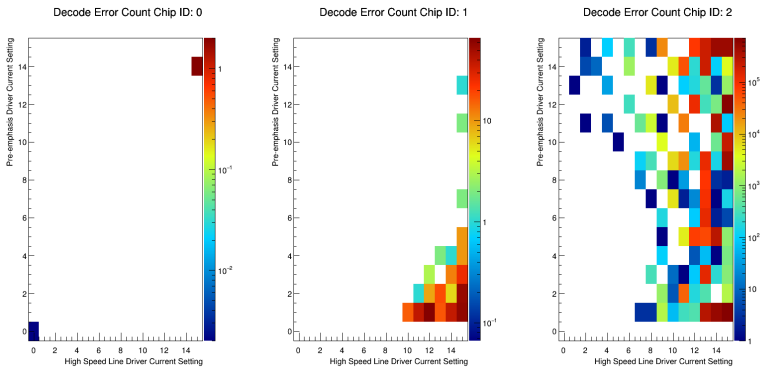

Figure 5.3: 48-hour DAC test decode errors using MAX105 PSU. The axes are the same as in the plot above.

Even though the digital performance of the ALPIDE chips should be indifferent between using a switch mode PSU and a linear one, we can clearly see on chip 2 that more errors were reported. The plots are nearly identical on chips 0 and 1, but the errors are amplified on the more noisy chip 2. Since the tests were performed over a long duration, the number of errors is not significant, and we can expect close to equal performance. The test also reports protocol errors, but it was omitted given its similarity to the decode error report.

Signal Integrity Measurement

Another digital test is measuring the dimensions of the signal eye using an oscilloscope. The greater the eye dimension is, the easier it is for the readout unit to recognize individual bits.

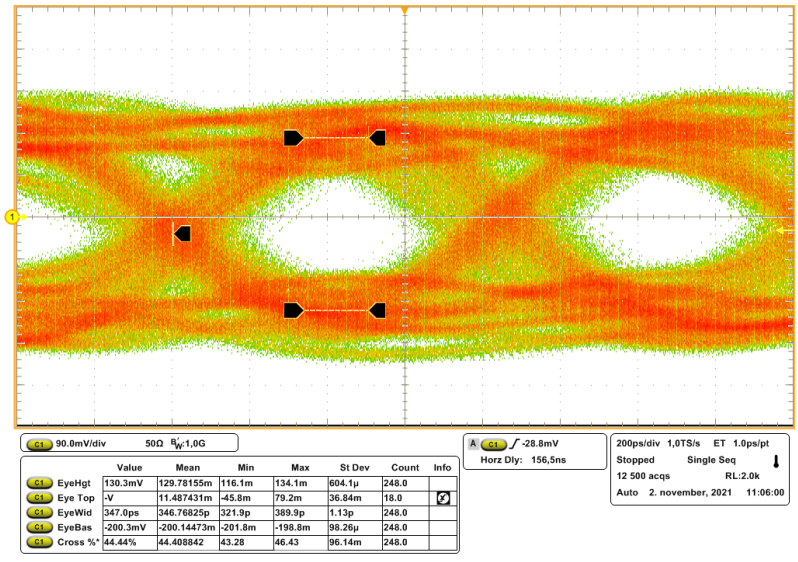

Figure 5.4: Eye measurement using linear PSU

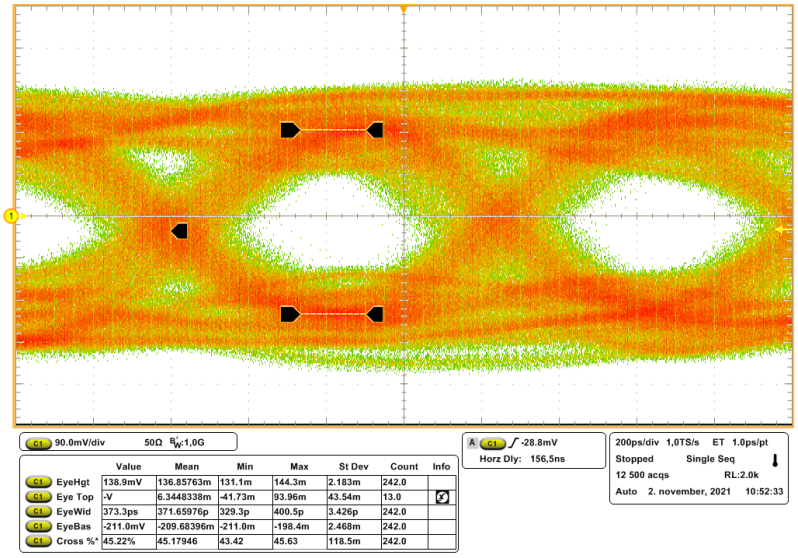

Figure 5.5: Eye measurement using MAX105 PSU

Here the tests turn out exactly the way we predicted they would, with virtually no difference between using a switch-mode PSU and a linear PSU. Any difference in the two eye diagrams is entirely within the margin of error. This should conclude that the signal integrity of the communication should be unaffected when using a switch-mode PSU.

5.3.2 Charge Injection Threshold Test

As mentioned in Section 3.1 the ALPIDE pixels consist of three sections, an input stage, a pixel analog front end, and a multi-event buffer. Figure 3.5 shows the layout and we can see how altering the threshold voltage will affect the sensitivity of the pixels.

The pixels also have a small 230 aF capacitor, which can be charged and discharged by the VPULSE_* signal, effectively creating a fake particle hit. This threshold test will then measure how much charge is needed for a pixel to detect a hit. When noise enters the system, such as ripple noise, the measured signal and threshold signal will also ripple, and a spike in voltage caused by noise might be registered as a hit.

Average Threshold Test			
Run number	Linear lab bench PSU	MAX105 PSU	Unit
#1	9.74	11.05	DAQ
#2	9.6	11.07	DAQ
#3	9.55	10.91	DAQ

Table 5.1: Average amount of charge needed to register a hit in the ALPIDE analog pixel front end. The DAQ unit is an arbitrary unit used by the ALPIDE which correspond to a specific charge. An average threshold level between 9 and 11 is within GOLD classification.

As expected, there is a discrepancy between the linear lab bench PSU and the MAX105 PSU, the difference is just large enough to push the chip out of a GOLD classification. The test shows that we might be experiencing some of the symptoms of having ripple in the supply lines. The effect it has on the sensitivity of the chips should be explored further with a proper filter.

5.3.3 Fake Hit Rate

When running the charge injection threshold test, the VPULSE_* signal is asserted, followed by the STROBE signal, but when doing a fake hit rate test, only the STROBE signal is asserted, and we can test many pixels in quick succession. The STROBE signal only needs to be high for 50 ns and does not inject any charge. Asserting only the STROBE signal is called sending an empty trigger, and in the test, one million triggers were sent to a single ALPIDE chip. Since we are running the test without any proton beam or light exposure, each hit reported will be a fake hit. Dividing the number of hits by the number of pixels times the number of sent triggers gives us the fake hit ratio. If an ALPIDE chip has a fake hit rate of below $10 \cdot 10^{-7}$ hits/pixel/event with the ten most noisy pixels masked, it can be classified with a GOLD standard label.

Fake Hit Rate Test	
Linear lab bench PSU	MAX105 PSU
$4.257 \cdot 10^{-6}$	$4.317 \cdot 10^{-6}$

Table 5.2: Amount of hits registered per pixel per trigger.

The MAX105 test reported a fake hit ratio of $4.317 \cdot 10^{-6}$, and when doing the fake hit test with the linear PSU a fake hit ratio of $4.257 \cdot 10^{-6}$ was recorded, which is only a 1.4% difference, well within the margin of error. Of course, masking the ten most noisy pixels will reduce this number. Since a few noisy pixels often dominate the fake hit rate count, masking away the ten most noisy will probably give the chip a GOLD classification.

After running several of these tests, a trend of some chips performing slightly worse with the MAX105 PSU was apparent. A good filter could prove beneficial for some chips that already struggle with meeting the noise criteria. We can conclude from the threshold and fake hit rate test that the chips will not register more false hits when being supplied by a switch-mode supply, but they will have a higher threshold value.

5.3.4 Load Step Response and Voltage Drop

Most, if not all, switch-mode PSUs have some internal control circuit to monitor and control the pulse width of the output signal. This modulation is essential unless there is an entirely static load. Pulse width modulation (PWM) is a way of digitally reducing a voltage. For example, if a 5 V signal has a duty cycle of 50 %, then it

will effectively be a 2.5 V average voltage signal. A longer duty cycle means more time for the smoothing capacitor to charge up.

The PWM control circuit is in charge of maintaining a stable voltage in the event of a load change. The MAX105 was tested with an instantaneous 8 A load to measure the voltage drop that will occur across the long cables, as well as measuring the "total control time", the time it takes for the signal to stabilize.

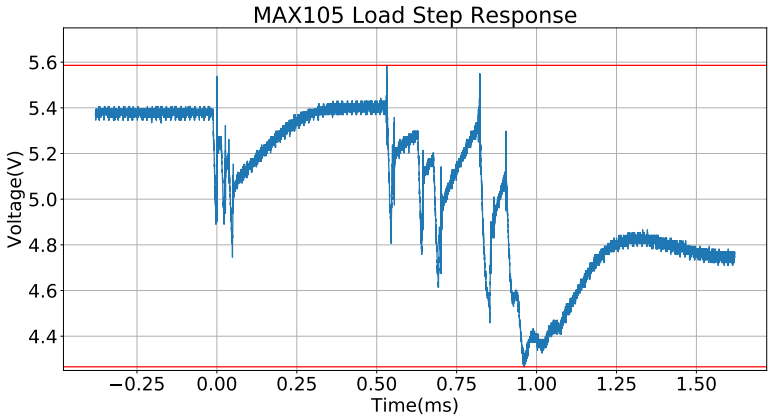

Figure 5.6: Load step response of the MAX105 PSU, measured at PSU terminals.

The MAX105 datasheet reports a total control time of less than 0.8 ms, while ensuring a maximum under or overshoot at 300 mV. The test was done by connecting a resistive load to a dummy TC where the ALPIDEs should have been connected, this way, we get a realistic look at the effect a large surge of power would have on the powerlines.

Figure 5.6 shows that the overshoot stays within the 300 mV boundary, undershoot is negligible as the voltage drop that commences across the cables with increased load is dominating anyways. The voltage drops from 5.4 V to 4.75 V, a 0.7 V drop. The TC is required to have at least 3.3 V in. Accounting for the voltage drop across the cables the PSU then needs to deliver at least 4 V. The rippling is caused by using a mechanical switch in this test.

The test shows that the PSU can handle a large instantaneous load without becoming unstable and without dropping below the minimum voltage of 4 V.

5.3.5 Temperature

Preceding the 48-hour DAC test, the MAX105 had been delivering power to the TCs from Friday to Monday. Upon touching the PSU, it felt uncomfortably hot to hold, but not in a way that caused immediate pain, meaning the temperature should be below 55 °C. Proper temperature tests should be conducted in the future with the MAX105 connected to a high load, preferably when it has been mounted in its chosen installation system.

5.4 Filter Design

5.4.1 Switching Noise Produced by MAX105

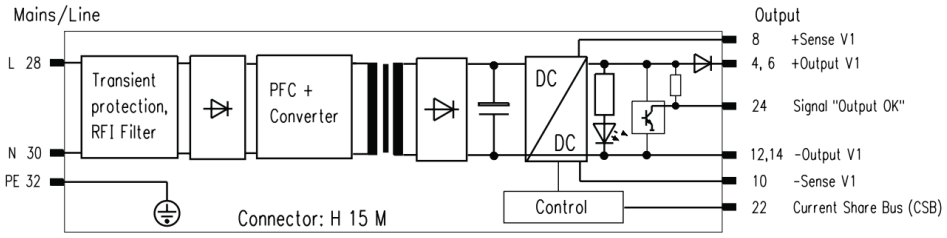

Figure 5.7: Internal layout of the MAX105

The schematic and datasheet provided by Schroff give some general insight into how it operates. The PSU has some transformer to turn the mains 220 V signal into a more suitable voltage, which is then rectified before being stepped down by the DC/DC switching component. The efficiency of the 5 V module is at 63%, while the efficiency of the 24 V module is 77%, a significant difference probably caused by the different switch-mode PSU models using the same core parts, and only changing the values of the DC/DC component from one model to the other.

Since Schroff did not put any identifying information on the schematic for DC/DC components, we have to assume that it uses a basic flyback topology. The efficiency of the 5 V module is relatively low, meaning the switching frequency is probably turned down to achieve the desired voltage, and this creates a more significant ripple since there is a longer charge and discharge period internally in the DC/DC modules capacitors.

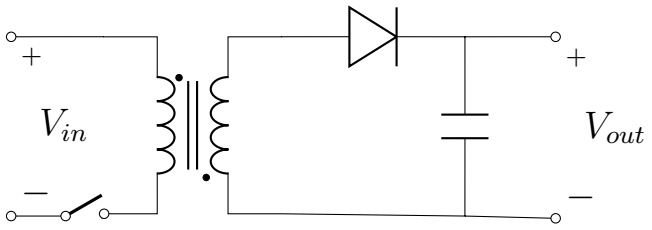

Figure 5.8: Basic flyback topology

During testing, several measurements were taken. The original hypothesis was that the MIC2915X linear regulators on the TC would mitigate most of the ripple, as a voltage regulator's output should be stable, even though the input is not. The PSU is supplying 5 V, but the voltage drop across the cables makes it closer to 4.5 V. The voltage regulators regulate the voltage down to 1.9 V. As the voltage regulators have an overhead of 2.6 V, they should handle differential mode noise without any problem.

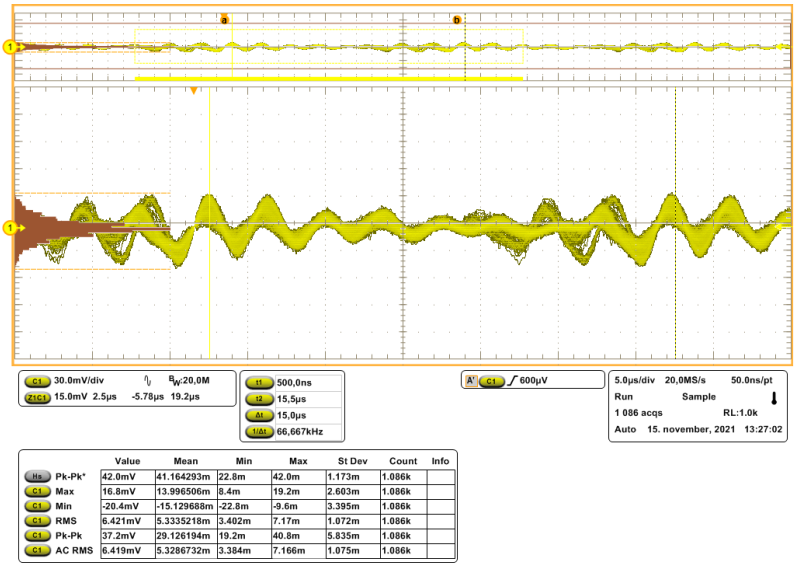

Figure 5.9: 66.6 kHz ripple on the analog rail of the TC, measured on the output to the ALPIDE string.

Looking at the measured signal in Figure 5.9, the linear voltage regulators offer less than ideal attenuation, and the MAX105 datasheet reported a ripple of about $50\,\mathrm{mV_{p\text{-}p}}$, and the measurements of 42 mV reinforce this. The noise switches with a frequency of 66.67 kHz.

5.4.2 Decoupling Capacitors

The choice of capacitors has to be done on a case-to-case basis. Multilayer ceramic capacitors (MLCC) are the most popular choice due to size and cost, especially decoupling capacitors. MLCCs offer low equivalent series resistance (ESR) in extremely small footprints. The low ESR makes it easier to supply short bursts of current. Electrolytic capacitors are generally large, and their use case is reserved for scenarios where large amounts of capacitance are needed. Looking inside a PSU you can often see large electrolytic capacitors as large as a couple of farads. Electrolytic aluminum capacitors are tall and skinny to utilize as little PCB space, but they are not aimed at low ESR applications. Electrolytic tantalum capacitors come in the same footprint as MLCCs, and are a nice sweet spot between size, capacitance, and ESR.

Small ceramic capacitors are also used to filter out noise, commonly the noise that is produced by other digital circuitry. High-grade tantalum capacitors should handle noise better than MLCCs, but the gain is small and often insignificant. High-grade tantalum capacitors will be used to filter out the noise on the transmission lines, as having a low dynamic impedance was prioritized for the filter design.

The impedance of a capacitor is given by

$$\frac{1}{2\pi \cdot \mathrm{f} \cdot \mathrm{C}} \tag{5.5}$$

and ideally, this would mean that the impedance falls linearly with the frequency, but realistically the packaging of any capacitor creates both resistance and inductance across the package. We want the impedance of a decoupling capacitor to be as low as possible regarding the offending frequency. This way, the current surges caused by the rippling are shorted to ground while the DC supply voltage persists.

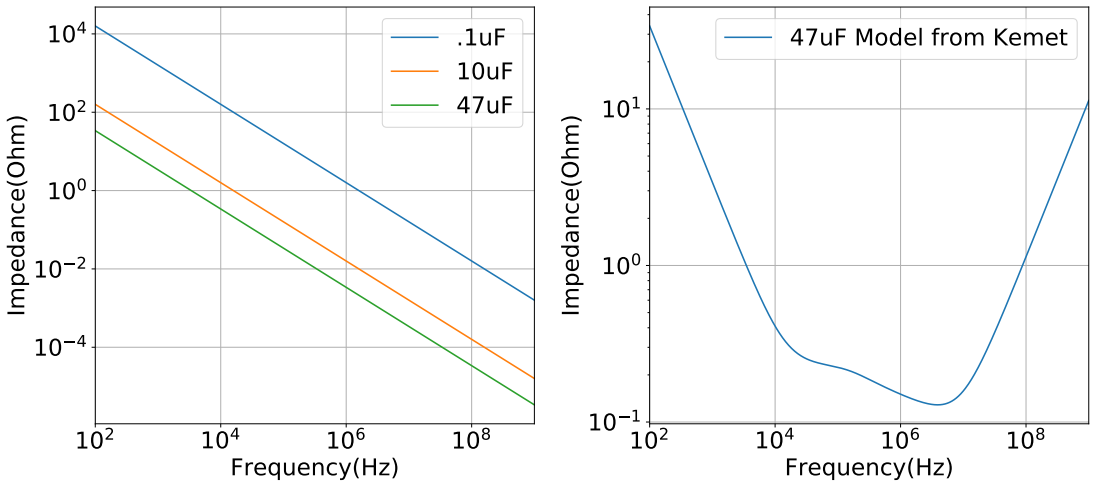

Figure 5.10: a) Ideal impedance of 10 µF, 1 µF, and 0.1 µF capacitors with no ESR or inductance. b) 47 µF tantalum capacitor model supplied by Kemet.

The capacitance denotes the impedance at low frequencies, but the series inductance is the main contributing factor at higher frequencies. The impedance of an inductor is

$$Z = 2\pi fL \tag{5.6}$$

which starts very low and increases with frequency. Combining the ideal capacitor and ideal inductor plots, we end up with a V-shaped figure which bottoms out at a frequency of $\frac{1}{2\cdot\pi\cdot\sqrt{LC}}$ as shown in Figure 5.10. However, choosing any single capacitor to function as the sole decoupler is wrong in all but the slimmest of fringe cases. A combination of 3 or more capacitors is usually chosen, this way signals at several frequencies have a low impedance path to ground.

The standard for over 50 years is to use a 10 µF, a 1 µF, and a 0.1 µF capacitor in parallel. The parasitic inductance of the capacitors will then be reduced by a factor of $\frac{L}{n}$ [37][1].

The TC has decoupling capacitors by the flex cable connector close to the ALPIDEs, each ALPIDE string has 3 MLCCs in case of a surge in power occurs on the chip, and to filter high-frequency noise. These MLCCs are 10 µF, 1 µF, and 0.1 µF large. Using three different capacitors is ideal for filtering out noise at different frequencies.

[1]This only works with identical capacitors, when dealing with different packaging, the series inductance will still be reduced, but not by a factor n.

The collection of 10 µF, 1 µF, and 0.1 µF decoupling capacitors are widely used as they filter out high-frequency pickup quite well while still large enough to handle small current surges.

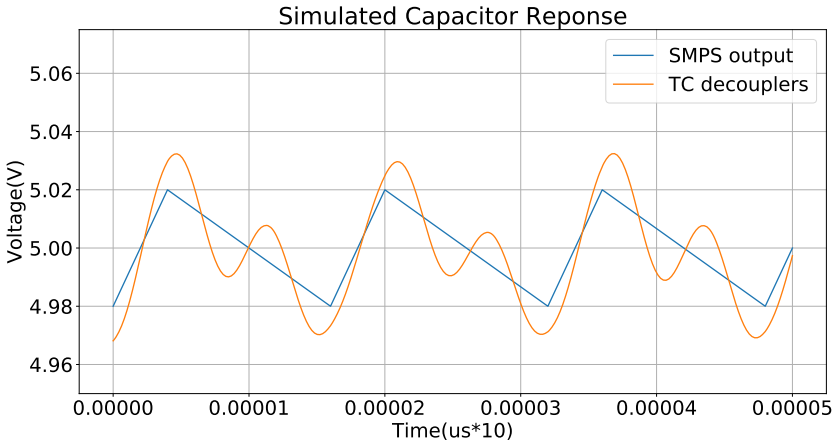

Figure 5.11: Simulated effect 10 µF, 1 µF, and 0.1 µF MLCC capacitors have on a 66 kHz triangle shaped signal.

The simulations in Figure 5.11 show that the decouplers amplify the noise produced at 66 kHz. The input waveform is shaped to simulate a switching supply's charge and discharge phases, it is not smoothed out to exaggerate the effect of the decouplers, but the lack of simulated spikes in current makes up for the lack of smoothing.

5.4.3 LC Filters

Using the characteristics for an ideal inductor and capacitor meantioned earlier, we can understand how the classic LC filter could be designed to have large attenuation against the switching noise we are experiencing. The effective low pass filter shown below will have an attenuation of −20 dB per decade starting from the critical frequency

$$\mathrm{f} = \frac{1}{2\pi \cdot \mathrm{L} \cdot \mathrm{C}} \tag{5.7}$$

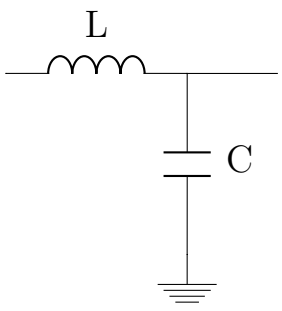

Figure 5.12: Basic LC filter.

5.4.4 Class X/Y Capacitors

Sometimes engineers spend hundreds of hours designing and perfecting their circuit, only to carelessly plug it into the power grid together with a hundred other devices, all of which create noise on the supply line caused by fluctuating loads, switching noise, et cetera. Class X and class Y capacitors filter out the noise on the supply lines. X capacitors are often called "Line-to-Line", and Y capacitors are called "Line-to-ground". Their names are effective in describing their use case. Figure 5.13 shows that X capacitors are present before a common mode coil, while the Y capacitors are placed after. Line-to-line filtering using X capacitors is reminiscent of standard bypass filtering, using capacitors between the supply lines to both filter out EMI/differential mode noise, and make sure the potential between the lines remains consistent [37].

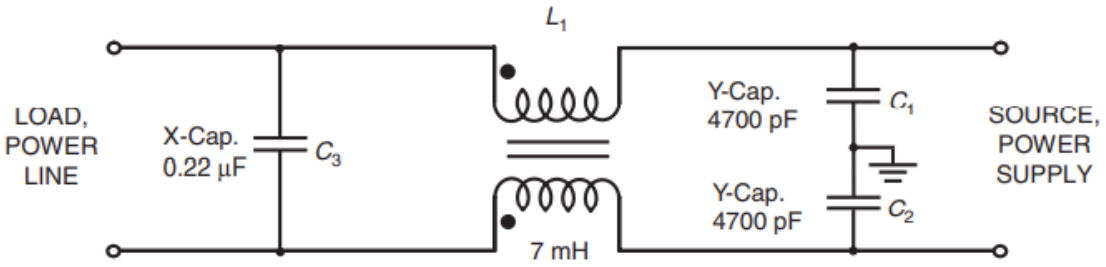

Figure 5.13: Generic power-line filter topology, including typical component values [37].

The Y class capacitors are effective in filtering out some common noise that is present on both supply lines. To do so, they need to be connected to some common voltage neutral point, and we cannot use conventional signal ground since there is no way to decouple a line to itself. This common ground can be the center of a transformer's coil, the chassis of the device, a rod placed in the literal ground, and many other solutions. The two Y capacitors help equalize the supply lines against this common neutral point, effectively filtering out common-mode noise. X capacitors are impervious to common-mode noise as they only react to a difference in potential across the supply lines. Common mode noise is discussed more in Section 4.4.

Naming schemes in engineering literature are often so bad that one should think it was made confusing on purpose. This stands true for different types of ground. Digital/signal ground can be viewed as the signal return path, and as mentioned above, the common neutral point we connect the Y capacitors to can be the literal earth ground if we want. The monitor board neutral ground point is connected to the enclosure of the Schroff PSU, so it will be referred to as chassis-ground.

Connecting anything to chassis-ground must be done with caution, as passing large currents through the chassis **will** cause a shock hazard, and giving your customers involuntary electroconvulsive therapy is bad for business [citation unnecessary]. Sev-

eral institutes across the globe have come up with a maximum amount of leakage current allowable by different instruments, things like dishwashers often get a bit more than conventional computers. NEK IEC 60598-1:2014 specifies that luminaires up to 20 A must not exceed 10 mA [40]. When working with 240 V systems, this requirement might be difficult to adhere to if you are trying to filter large transients of common-mode noise, but when working with 5 V, the leakage current out of a high-quality tantalum capacitor is going to be well below 10 mA

5.4.5 Final Filter Design

All the theory laid out in this chapter has been relevant to the final filter design. Filter design and testing was a new topic for me, and one important thing I learned was to test first and assume later, as even the advice I got directly from Schroff turned out to be based on false assumptions.

Testing the most obvious solutions, such as proper decoupling and a low pass filter, turned out to be futile, which was done before acquiring Figure 5.9, and still, there is close to 40 mV ripple on the supply line. The first proper solution tested was several decoupling capacitors and a proper inductor to create a low pass filter. A 15 µH coil was used in conjunction with one 220 µF and one 47 µF capacitor.

$$\begin{aligned} f_{3db} &= \frac{1}{2 \cdot \pi \sqrt{15\,\mu H \cdot 267\,\mu F}} \\ f_{3db} &= 4116\,Hz \end{aligned} \tag{5.8}$$

Having a critical frequency of 4116 Hz should, in theory, provide more than adequate attenuation to combat the apparent ripple. In practice, the filter had no effect, and a 'before/after' visualization would have been a waste of space as they would have been almost identical. A small common mode choke was later added in conjunction with this solution to see if it affected the noise.

A common mode choke is just two inductors wound together, so noise on one supply line will produce a complementary effect on the opposite line, since both lines should see the same noise, they would then cancel each other out. Common mode chokes by themselves had little to no effect on the low-frequency noise, both in simulation and in practice.

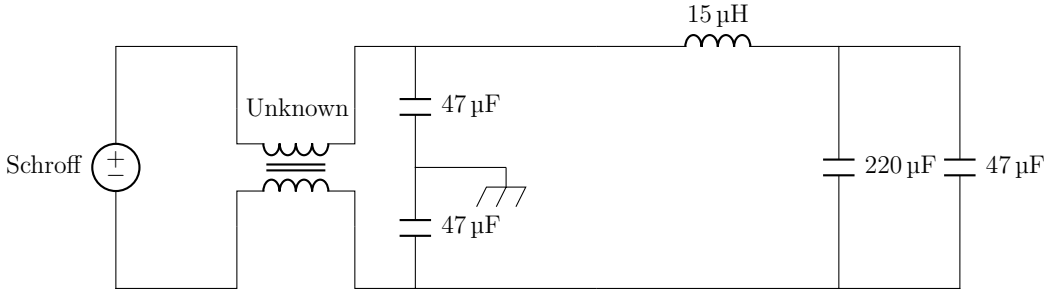

Figure 5.14: Filter prototype design #2.

Using Y capacitors in series with a common mode choke creates a path for the common mode noise to dissipate into, removing it from the supply lines. Prototype filter #2 shown in Figure 5.14 used a custom-wound common mode choke together with two 47 µF ceramic capacitors, equating to a completely unknown filter frequency.

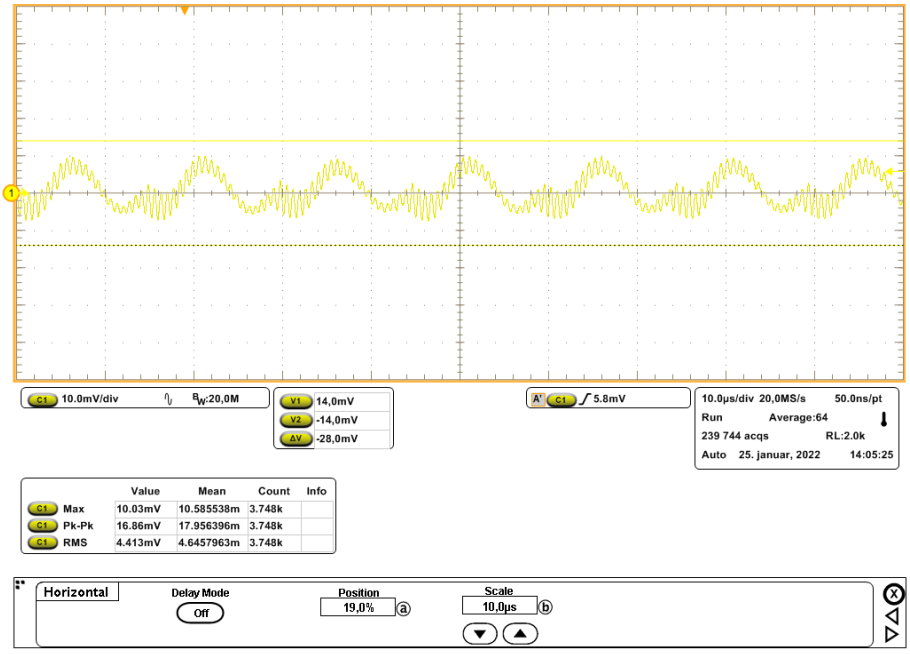

Figure 5.15: Ripple measured across decoupling capacitor on the TC with filter #2.

This filter was the first to have a proper effect on the noise, the ripple was reduced by about 30%, but it also introduced some much higher frequency ringing on the supply line. The ringing had a frequency of 1.4 MHz. Removing the 15 µH inductor yielded an even worse ringing, but identical common-mode noise ripple. Adding two smaller capacitors to handle the high-frequency ringing while using a proper common mode choke removes what is left of the ringing on the supply line. The final filter design is shown in Figure 5.16.

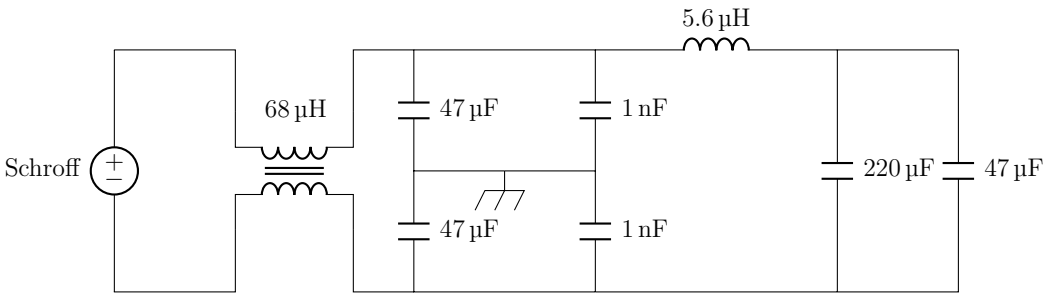

Figure 5.16: Filter prototype design #3.

Adding both a 47 µF capacitor and a 1 nF capacitor after the large common-mode choke creates two individual critical frequencies. By having large attenuation at

two points in the frequency spectrum both the high-frequency ringing and the low-frequency common-mode noise are filtered. The critical frequencies are shown in Equation 5.9.

$$\begin{aligned} f3db &= \frac{1}{2\pi\sqrt{\mathrm{L}\cdot\mathrm{C}}} = 2815\,\mathrm{Hz} \\ f3db_1 &= \frac{1}{2\pi\sqrt{68\,\mu\mathrm{H}\cdot 47\,\mu\mathrm{F}}} = 2815\,\mathrm{Hz} \\ f3db_2 &= \frac{1}{2\pi\sqrt{68\,\mu\mathrm{H}\cdot 1\,\mathrm{nF}}} = 610\,\mathrm{kHz} \end{aligned} \tag{5.9}$$

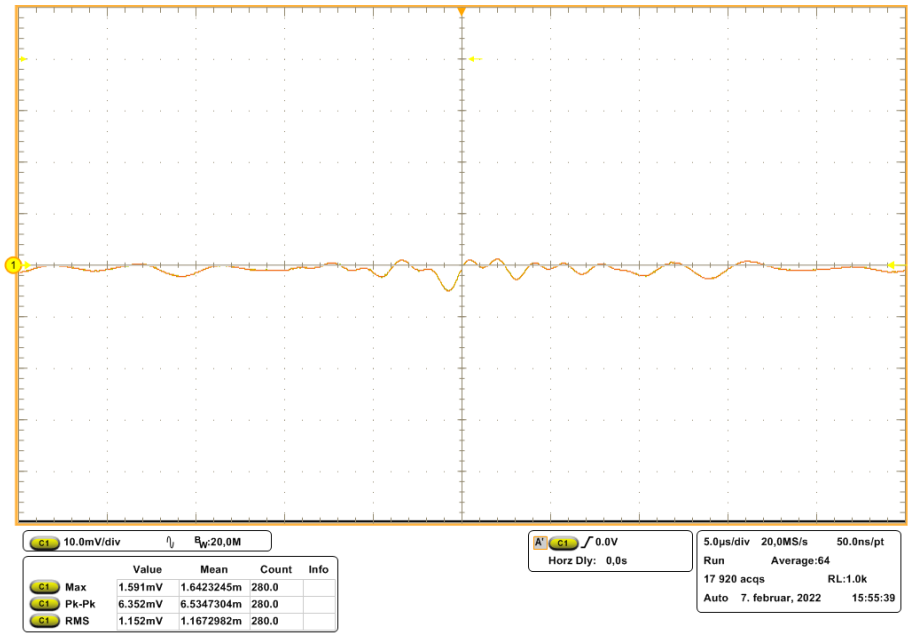

Figure 5.17: Ripple measured across decoupling capacitor on the TC with final filter solution.

This filter proved very effective, removing almost all of the noise on the supply line, with only a small 50 Hz component remaining originating from an unknown source. Note that the signal in Figure 5.17 has been averaged to clearly show the waveform.

Fake Hit Rate with Filter

In Section 5.3.3 a fake hit rate of $4.317 \cdot 10^{-6}$ was reported without masking out any noisy pixels. Rerunning the fake hit rate test with a significantly longer strobe was executed to try to exaggerate the difference between the MAX105 PSU and the linear lab bench supply.

Row/Col #	Col hits	Row hits
1	999968	999968
2	422724	422716
3	361744	361744
4	849013	849063
5	330161	330163
6	904768	904771
7	504771	504769
8	983801	983801
9	501747	501751

Table 5.3: Comparison of the 9 most noisy rows and columns during the MAX105 fake hit rate test. This verifies that a noisy row column combo is dominated by a single pixel.

The original test did not mask any pixels, but even without doing so it managed a fake hit rate low enough for a bronze classification. The new tests shown in Figure 5.18 have had their 10 most noisy columns masked out before plotting the data.

Most often a single faulty pixel dominates the fake hit amount on a column, which can be validated by listing the 9 most noisy columns versus the 9 most noisy rows. The 9 most noisy columns and rows are shown in Table 5.3 was chosen as they all have reported more than 300 thousand hits. Masking an entire noisy column will not affect the data too much when running fake hit rate tests.

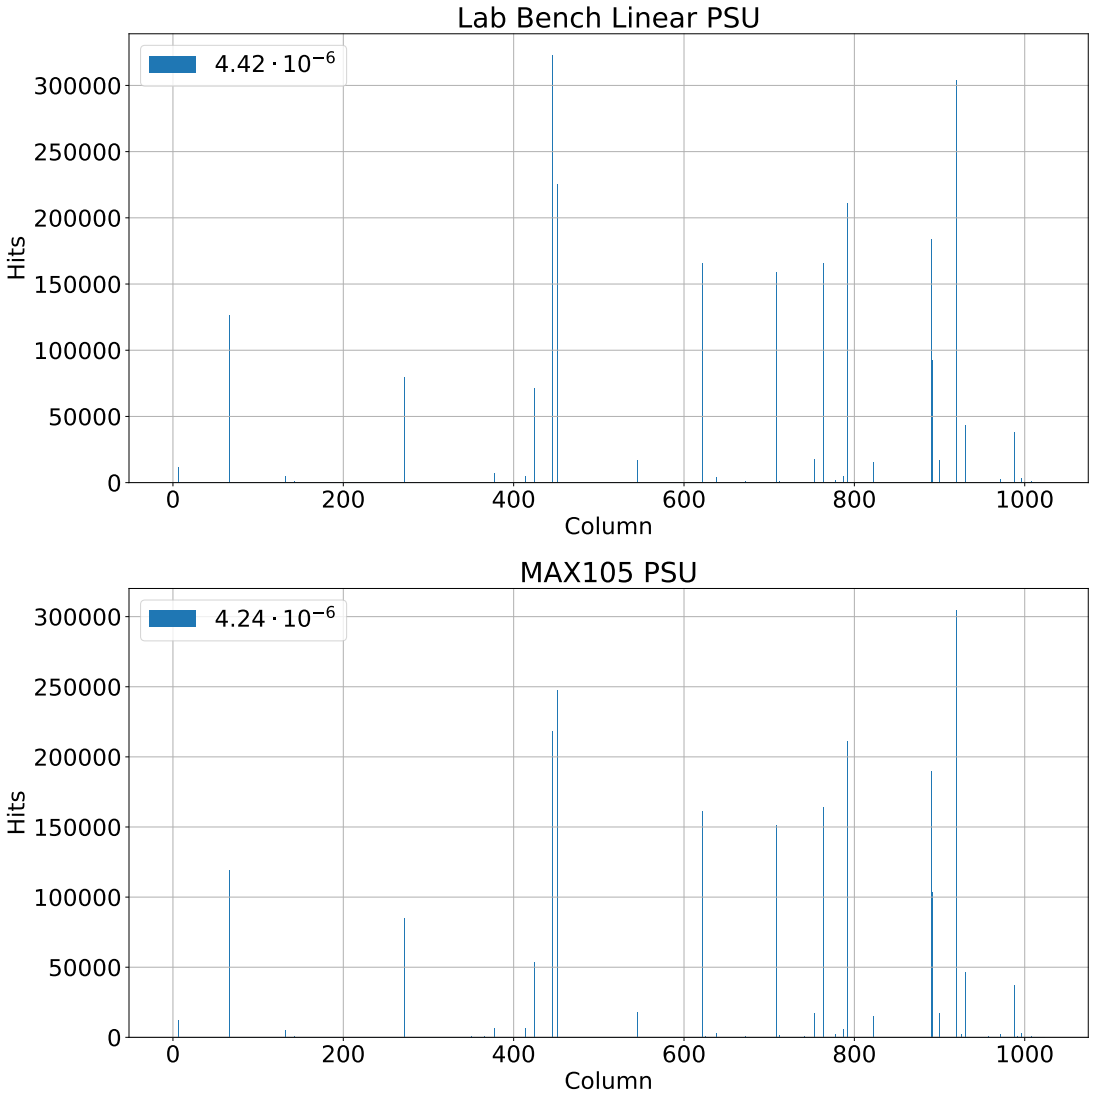

Figure 5.18: Fake hit rate test result of the linear lab bench supply and MAX105 PSU.

Figure 5.18 shows two runs of the fake hit rate test, one with the linear lab bench supply, and one with the MAX105 PSU. Looking at the data shows similar trends of noise in the same pixel columns, with some slight variances in the amount of hits in any given column.

The difference between the lab bench supply and the MAX105 PSU is negligible with the new solution, which was to be expected. The false hit rate was 4.42 false hits per million with the lab bench supply, and 4.24 per million with the MAX105. The difference is well within margin of error and we can assume the two supplies perform the same.

6 THE MONITOR BOARD

The MB PCB made specifically for the MAX105 is outlined. The functionality is described, and the working of the current software is explained in broad terms.

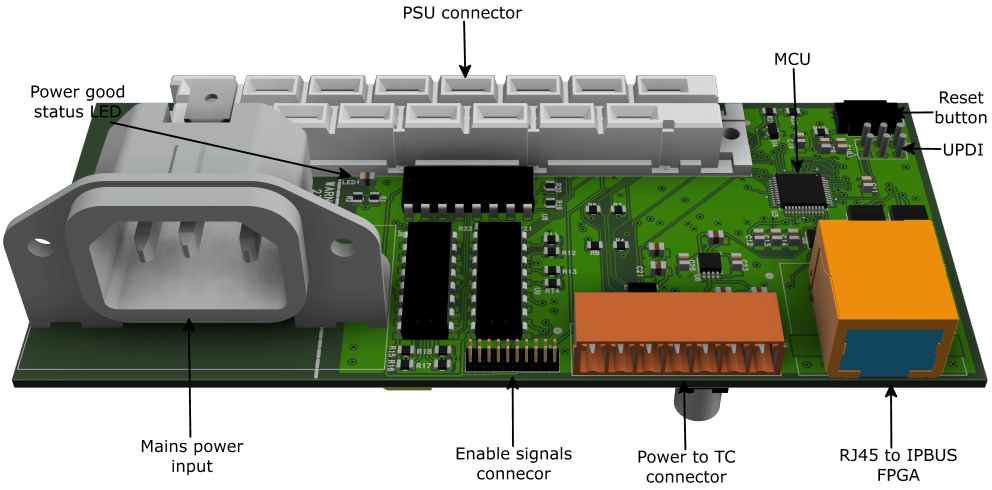

Figure 6.1: 3D model of the MB.

6.1 Requirements

The monitor boards will be directly connected to the transition cards and must satisfy these requirements:

1. The monitor boards must have a Weidmüller 7-pin interface to support powerlines for 2xDVDD, 2xDVSS, AVDD, AVSS, and PWELL.
2. The boards must be able to deliver a minimum of 3.3 V over the DVDD and AVDD cables spanning several meters. The voltage drop must be accounted for. The boards must be able to provide at least 13.2 A to each TC.
3. There should be a soft startup to stop large power surges.
4. Each monitor board must communicate with an FPGA over a serial communication protocol. The FPGA will communicate with some interface that the user can access to control and monitor each TC.
5. The PWELL signal must be regulatable by the user between 0 V and −6 V.
6. All twelve enable signals must be individually controllable by the user. If the chips are drawing too much current, the monitor board should shut off the entire layer and report the error.
7. Temperature should be monitored and reported regularly.

8. On encountering an error, it should be stored until the user issues a read command.

The MBs are connected to the main KCU115 FPGA using RJ45 cables. Each MB has an AVR128DB48 microcontroller for communication, monitoring, and enabling signal control. The MCU communicates with the KCU115 using the RS485 differential signaling USART protocol.

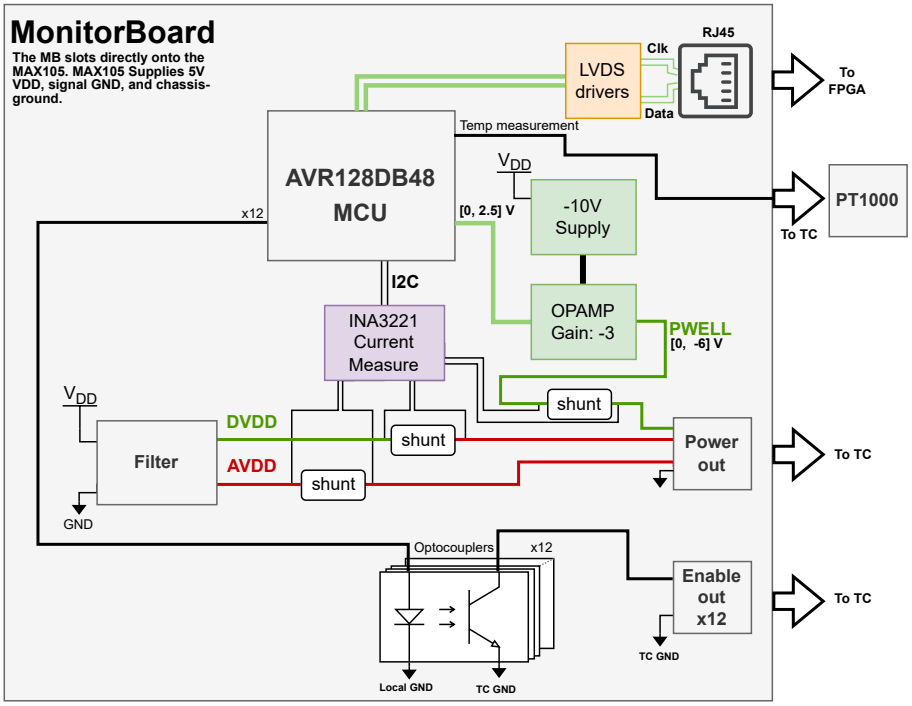

Figure 6.2: Simplified overview of the MB.

Figure 6.2 shows a simplified diagram of the MB. The connector for 230 V and the DIN41612 connector are omitted for simplicity.

6.2 Microcontroller Software

The choice of MCU was heavily dependent on what internal components it offered. Given the broad spectrum of different tasks, a general-purpose MCU is preferred to one that specializes in one area. The MCU needs at least one DAC (Digital to Analog converter) and one ADC (Analog to Digital Converter). In addition, it would be beneficial to have both an I^2C module and a USART module. The AVR128DB ticks all the boxes whilst having a small form factor, and being Norwegian, I am naturally *slightly* biased towards AVR products. The AVR128DB features a clock speed of up to 24 MHz, but since such speed is not crucial, the clock rate is kept at the default 4 MHz. It can be altered later on if needed.

A Curiosity Nano development kit was used to test and verify the communication protocol from the microcontroller, through the FPGA translation layer, and up to the user over IPbus.

6.2.1 Watchdog

An adequately designed watchdog is essential in any autonomous software, it can be implemented in both HW and SW, and the AVR128DB has an internal watchdog module. A watchdog's function is quite simple; if the dog is not "kicked" every so often, it will bark, which restarts the MCU. The watchdog is adjustable to accept kicks within timed windows, requiring the MCU to keep tabs on when it was last kicked. Kicking the dog outside the window will cause the system to restart.

This designs watchdog will function in "normal mode", which denotes that a kick command only has to be issued before the time out timer runs down, as shown in figure 6.3

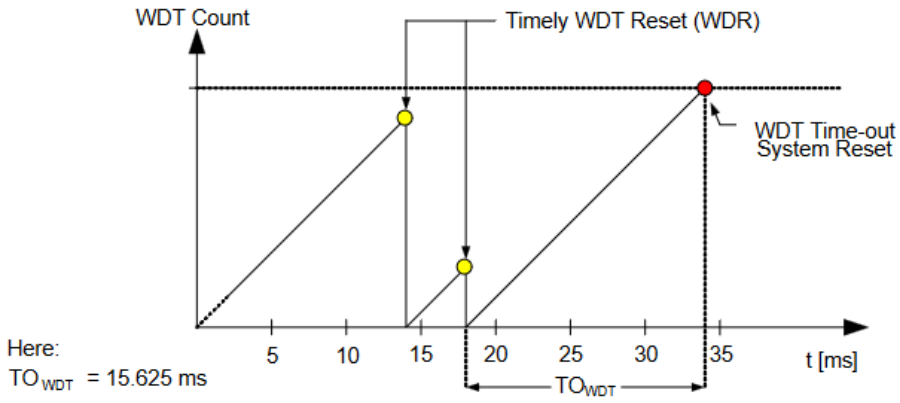

Figure 6.3: Watchdog timeout window [41].

6.2.2 Communication Protocol

The USART interface is either a two (synchronous) or one (asynchronous) wire interface used for debugging and communication. The AVR128DB has several built-in USART modules, which can be set to RS485 mode. RS485 mode uses four pins, one for the clock, two for data, and the last to signal if the chip is receiving or sending data. Even though the MCU has a TX and RX pin, only a single twisted pair will be dedicated to data, meaning the communication will be half-duplex.

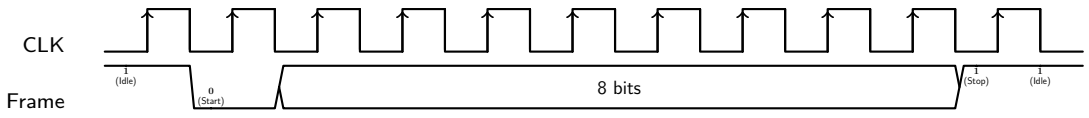

Figure 6.4: Single USART frame

The USART frame of the MCU has one start bit, one stop bit, eight data bits, and no parity or error-correcting checks. The MB uses a twisted pair for the clock. The MCU needs some time to shift the data from the IO register into an internal register, denoting a maximum clock frequency of $\frac{\text{MCU}_{\text{CLK}}}{4}$. The MCU crystal is set to 4 MHz, giving a theoretical maximum baud rate of one megabit per second (Mbps). The

internal crystal has a maximum frequency of 24 MHz, adjustable through software, which could provide a theoretical communication speed of 6 Mbps.

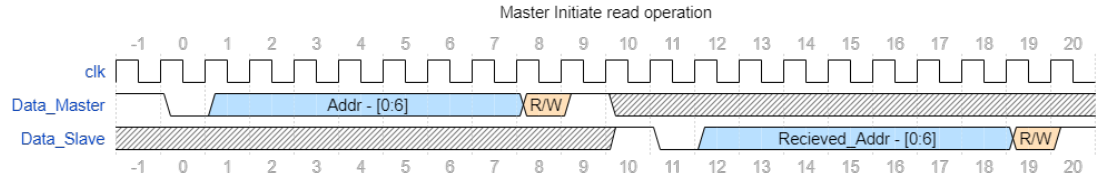

Figure 6.5: Master initiating a read or write operation, MCU responds with the received adress to confirm that communications are functioning correctly.

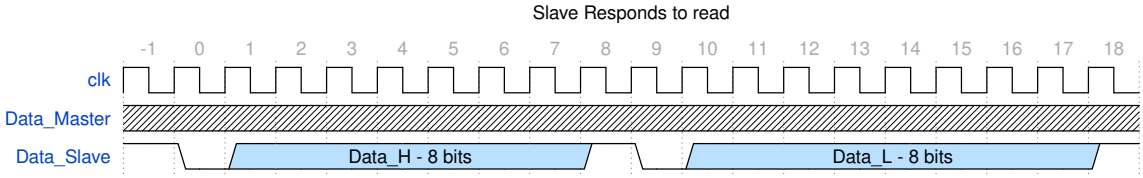

Figure 6.6: Slave responds to a read operation and sends 16 bits spread over 2 frames after confirming the adress.

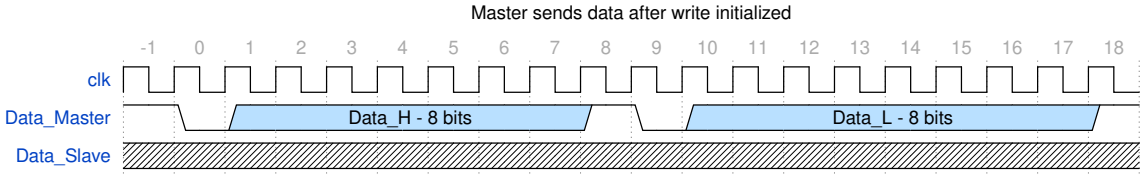

Figure 6.7: Master write operation, if the received adress is correct 16 bits is sent to the MCU spread over 2 frames.

As shown in figure 6.5 to figure 6.7, whenever a transaction starts between the FPGA and the MCU, the FPGA sends the address it wants to interact with, and the MCU responds by sending the address back. This transaction doubles as a handshake; if the FPGA receives the wrong address back, then the communication settings are mismatched, or something is wrong on the communication line. Following a correct handshake, the FPGA either receives or sends 2 bytes. Several of the values saved on the MCU is larger than eight-bit, but since both the MCU and the RS485 protocol are 8-bit, they are split up and sent as two individual bytes. The MCU always sends and receives the eight most significant bits first. The data line is changed on the rising edge and sampled on the falling edge.

In asynchronous operation, a baud rate is specified beforehand so that all devices know when to sample data. As the MCU is running in synchronous mode, it will use the supplied FPGA clock to send and receive data. Having the FPGA supply a clock ensures better communication, and it adds the ability to alter the communication speed from the FPGA, instead of reprogramming all 43 MBs.

The USART protocol in the AVR128DB starts with sending and receiving the LSB first, making it seem flipped when viewing the wave diagram. An example is shown in figure 6.8. Here a read command is issued to register 0x01 (0b00000001), but since the signal is flipped, the signal appears as 0b10000000. The MCU responds with a handshake, and data is then sent in two frames.

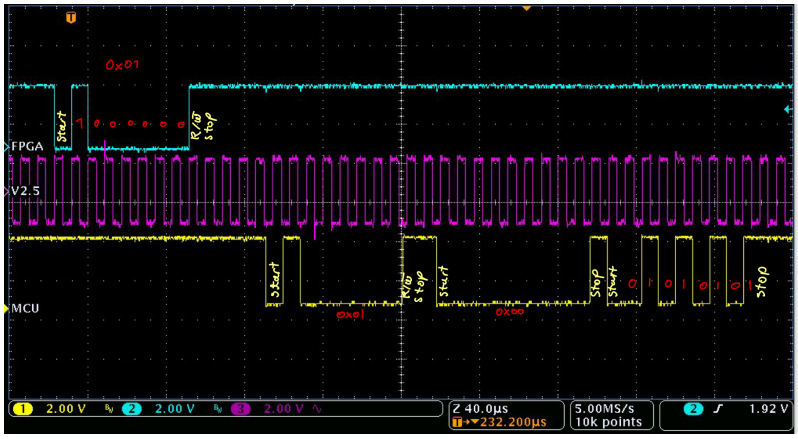

Figure 6.8: Read operation responded to by the MCU. Test was done using IPbus software, KCU115 development board, and a AVR128DA48 on a Curiosity Nano.

6.2.3 MAX3441 Transceivers

The MAX3441 transceivers are single-ended to differential transceivers. The chips support the RS485 protocol and can operate at speeds up to 10 Mbps. There are two MAX3441 ICs present on the MB, one for the data line and one for the clock line. Both ICs are connected to the RJ45 connector, using two twisted pairs to communicate with the KCU115.

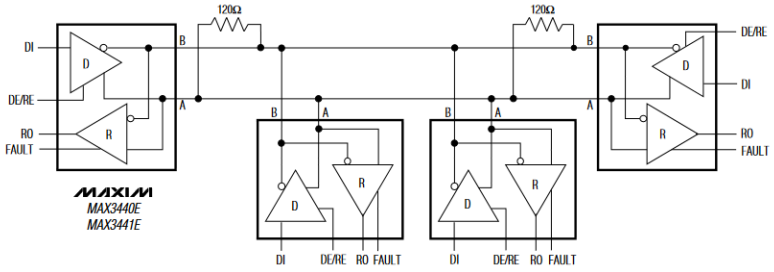

Figure 6.9: Typical RS-485 network [42].

Figure 6.9 shows how multiple units are connected to the same communication line in an RS485 network. To alleviate the need for program-specific addresses for each MB, the MB and the FPGA translation layer will only have one MB per communication line.

6.2.4 Registers

All the MCU monitoring, controlling, and debugging are done through the built-in registers. A complete register map can be found in appendix B.1. There are

25 registers, 20 general-purpose, 3 control registers, and 2 error message handling registers.

Regular Registers

Issuing a read command of any register will always be reciprocated with 16 bits, but few registers use all 16 bits. The registers which do not use all 16 bits will be filled with leading zeroes.

Control Registers

The purpose of the control registers is to have an easy way of executing custom code, which can automate things like scanning the enable registers and recording current draw with one operation. CTRL1, CTRL2, and CTRL3 are all control registers, where each bit in the registers functions as a flag that trigger special functions. After triggering a function by setting the bit to '1', the flag is reset to '0'. The CTRL registers should never be read. Register CTRL3 is entirely empty, and CTRL2 is only used for debugging. More functions can be added to these registers later.

The control registers will use each bit as an individual flag, meaning they support at max eight commands. For example, writing 0x03 to CTRL1 will trigger two commands, first 0x01, then 0x02 (assigned to Set_Enable_Signals_OFF() and Enable_Scan()).

Error Registers

There are two registers dedicated to error handling. One keeps tabs on the number of errors recorded, whilst the other stores the error code. Clearing the ERROR_COUNT will automatically clear the ERROR_MSG register as well. The ERROR_MSG register has 128 bytes reserved for error storing. The error message register is its completely own register, and expanding it to support more messages is trivial.

6.2.5 PWELL register

The PWELL register gets its own entry given that it is a tad more complex than the other registers. The hardware is described in Section 6.4.

The output of the DAC on the MCU is connected to an operational amplifier set to have a gain of $-3\,\mathrm{V/V}$. By outputting $1\,\mathrm{V}$ from the MCU DAC, the PWELL signal will be $-3\,\mathrm{V}$.

Two registers are connected to the DAC, the DAC_VALUE, and the PWELL_VOLTAGE_MCU register. The user can write a value between $0\,\mathrm{V}$ to $6\,\mathrm{V}$ to the PWELL_VOLTAGE_MCU register, in which the MCU will regulate the DAC output such that a complementary signal is created on the PWELL line. The

user can write an integer value to the register, which the MCU interprets as mV. For example, writing 1000 to the register will result in a PWELL value of −1 V.

If the user wants to directly write to the DAC, the DAC_VALUE register can be used, where writing a value between 0x00 and 0x400 will output a signal between 0 V and 2.5 V from the MCU DAC. Here the gain of the LT1991 amplifier will have to be accounted for by the user.

6.2.6 Programming

The 6-pin (2.54 mm) universal program and debug interface (UPDI) is a proprietary interface for AVR microcontrollers. The UPDI programmers require six pins, one pin for data, one for power, and one for ground. Three of the pins are not in use.

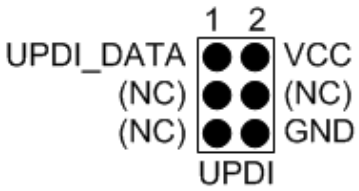

Figure 6.10: Pinout of the UPDI interface [43].

The intended way to program an AVR microcontroller is to use an Atmel-ICE programmer, but the cost seemed a little steep since these cost around 200 USD. On the other hand, UPDI is a slow-speed and straightforward protocol capable of being done from a USB serial adapter. The project SerialUPDI goes more into detail about how it functions, but the important part is that it is an improved version of the pymcuprog tool from Microchip.

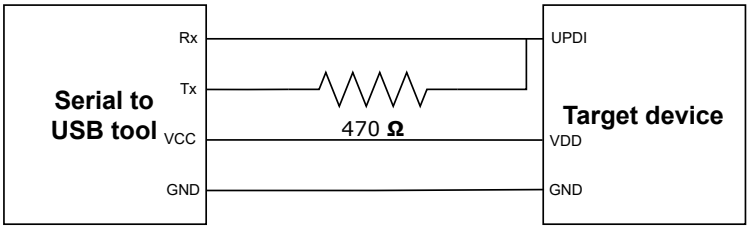

Figure 6.11: Wiring for the SerialUPDI software.

Figure 6.11 shows how to connect the USB serial adapter to the UPDI connector present on the MB.

6.3 INA 3221 Current Measuring

An INA3221 current and power monitor IC has three pairs of inputs. Each input pair will be connected to a shunt resistor used to measure voltage. The total current

draw of the TC's DVDD, AVDD, and PWELL domains can be measured using the voltage across the three shunt resistors.

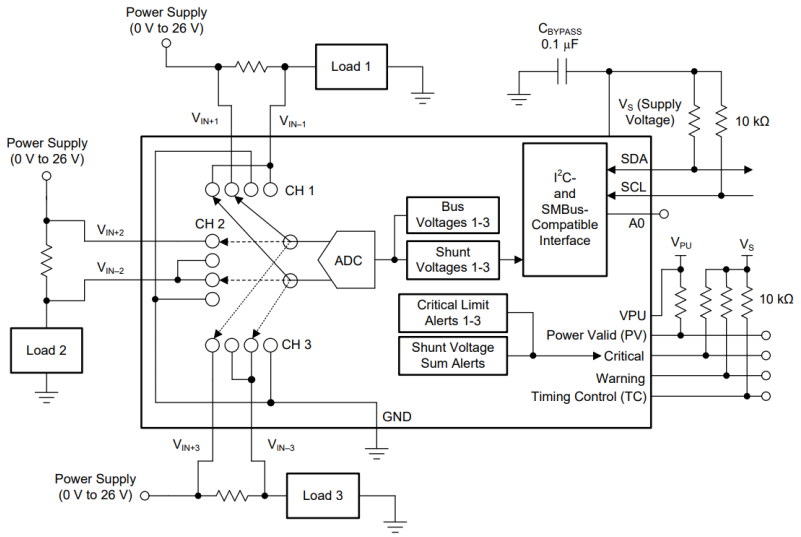

Figure 6.12: Typical application of the INA3221 [44].

Utilizing an ADC compares the voltage across the shunt resistor to the bus voltage, and the current passing through these resistors is stored in the MCU for the user to read out.

The INA IC registers use 13 bits to store the measurements, one bit for the sign, and two bits for configuration. The INA3221 has a resolution of 40 µV and a range of 2^{13}. To utilize the entire bandwidth, the voltage across the shunts should be 328 mV during max load, but using the entire bandwidth is unnecessary as long as we know the circuit is not experiencing unhealthy loads. A smaller value shunt resistor is chosen to minimize the amount of heat generated. For reference, to use the entire bandwidth, 1.73 W of energy would be wasted as heat.

The DVDD lines are split in two to minimize the amount of current we need to carry over a single pin. The two cable setup also halves the voltage drop we experience with long power cables from the PCU to the TC. For the digital lines, two shunt resistors are used with a resistance of 25 mΩ. Given that the absolute maximum power draw we are going to experience over the DVDD lines is 880 mA times 12 strings, the INA3221 will measure a voltage from 0 V to 132 mV. Given the earlier mentioned 1-LSB step size of 40 µV we will have 3300 steps of granularity. Only 0.7 W of energy will be wasted as heat in this configuration.

6.4 PWELL Circuitry

As mentioned earlier in section 3.1.2, to regulate the back biasing of the ALPIDEs, a variable voltage source is needed. The voltage should span 0 V to −6 V, and it has to be user-controllable through software. Since the MAX105 only have a positive 5 V output, a local voltage source had to be included locally on the MB.

MAX865 from MAXIM is a "Compact, Dual-Output Charge Pump". The IC doubles and inverts a given input voltage. MAX865 needs four capacitors to operate, equating to a relatively low total footprint.

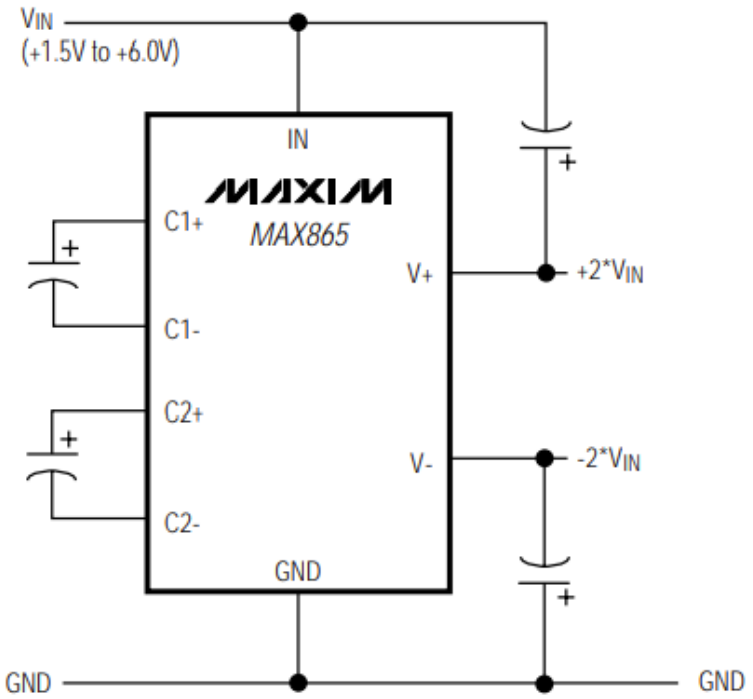

Figure 6.13: MAX865 chip, showing the need for two flying capacitors and two charge storing capacitors [45].

The chip uses the C1 and C2 together with several switches to double and invert the voltage. A capacitor is charged to V_{in}, and then the polarity is switched to produce an inverse potential of what it was charged with. An example of such a flying capacitor voltage pump is shown below. It is called a flying capacitor voltage pump since both the leads of the capacitor are changing simultaneously.

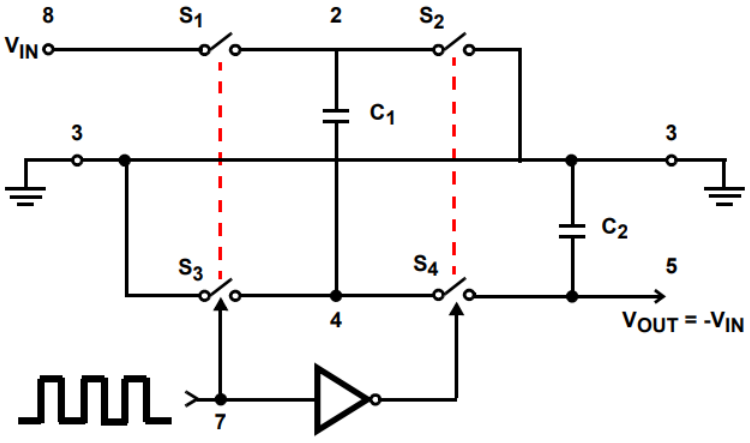

Figure 6.14: An ideal negative voltage pump [46].

In figure 6.14 a capacitor C1 is used to flip the polarity, and a second capacitor C2 stores the charge between the pulses. To create this polarity change, the switches S1 and S3 operate opposite of S2 and S4, meaning that the capacitor lead that is connected to V_{in}, is in the next clock cycle connected to ground. A voltage doubler would charge a capacitor up to V_{in}, then it would flip the switches so that the V_{out} pin had two times V_{in} potential to ground. The MAX865 uses both these concepts, requiring four capacitors to operate. The capacitors are 10 µF, the datasheet shows that larger capacitors yield less ripple on the output. Having larger capacitors than 10 µF would probably not yield any less ripple in our case as the PWELL power lines draw very little current.

The voltage created by the MAX865 is −10 V, and an LT1991 precision amplifier use this as a supply voltage. The LT1991 has built-in precision resistors, making the gain it produces very predictable and accurate from unit to unit. By having this negative voltage supply, the LT1991 can generate a variable negative voltage for PWELL between 0 V and −6 V.

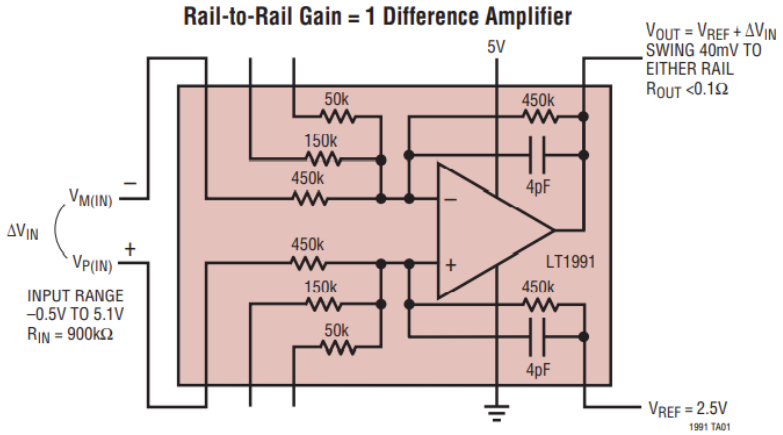

Figure 6.15: LT1991 typical application [47].

The inverting 150 kΩ resistor pin is connected to the AVR128DB, and the three non-inverting pins are connected to ground. With a gain of -3, the MCU can create a PWELL signal between 0 V to −9.960 V, as the LT1991 amplifier can swing within 40 mV of any rail. Supplying a voltage lower than −6 V for the PWELL signal can cause damage to the ALPIDEs. The PWELL signal will have to be limited by software to avoid this. There is also no need to think about the voltage loss over the long cables as the currents will be very low, possibly only a few milliamps.

Setting the PWELL voltage is done by writing to two registers, this is described in detail in Section 6.2.5.

6.5 Temperature Monitoring

A PT1000 element will be connected to the same connector that hosts the enable signals on the TC. The position of this PT1000 element is not yet determined.

The PT1000 element is connected in series with a known resistor. The voltage across the PT1000 will change with the resistance as it changes temperature. There must never be more than 0.5 mA going through the PT1000 to stop it from self-heating due to the dissipation of power.

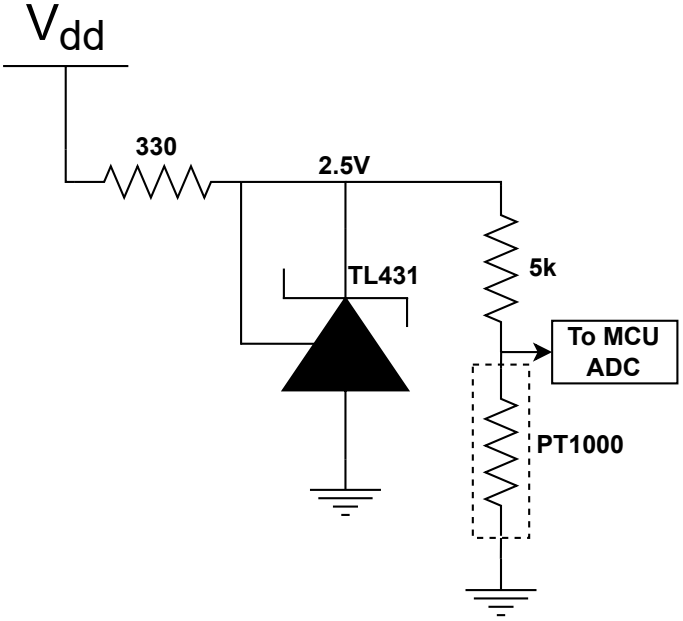

Figure 6.16: PT1000 Wiring.

TL431 from NXP is an adjustable precision shunt regulator. By shorting the reference to the cathode of the unit, it creates a 2.5 V reference. Moreover, having a 5 kΩ resistor in series with the PT1000 ensures that the PT1000 will never self-heat due to excessive current, at 0 °C the total resistance will be 6 kΩ.

$$\begin{aligned} \mathrm{I} &= \frac{2.5\,\mathrm{V}}{5\,\mathrm{k\Omega} + \mathrm{R_{PT1000}}} = \text{current consumption} \\ \mathrm{I_{MAX}} &= \frac{2.5\,\mathrm{V}}{5\,\mathrm{k\Omega} + 1000} = 0.417\,\mathrm{mA} \\ \mathrm{I_{MIN}} &= \frac{2.5\,\mathrm{V}}{5\,\mathrm{k\Omega} + 1385} = 0.392\,\mathrm{mA} \end{aligned} \tag{6.1}$$

PT1000 thermocouples have a relatively linear resistance across the temperature interval in our use case. We can assume that the temperature of the TC will never drop below 0 °C, as then the cooling liquid would risk freezing. If temperatures of 100 °C or more are reached the MBs should shut down the TCs, and we are not interested in knowing the temperatures above this since the MCU is told: "if above 100 → shut down".

The AVR128DB48 MCU features a 12 bit ADC with an adjustable reference voltage. Having the PT1000 element connected locally to ground the PT1000 element can operate in single ended mode, taking advantage of the full resolution (in dual ended

one bit is reserved for sign).

$$\begin{aligned} V_{PT1000\_max} &= \frac{2.5\,V}{6000\,\Omega} \cdot 1000\,\Omega = 0.417\,V \\ V_{PT1000\_min} &= \frac{2.5\,V}{6385\,\Omega} \cdot 1385\,\Omega = 0.542\,V \end{aligned} \tag{6.2}$$

There is a difference of 126 mV over the relevant temperature span, meaning that if the temperature changes 1 °C, the voltage will change 1.26 mV. As mentioned above the MCU has a 12 bit ADC, and the reference voltage is chosen to be 2.5 V, since the measurement area is between 0 V to 2.5 V, the sensitivity will be $\frac{2.5\,V}{2^{12}} = 0.61\,mV$. The ADC will be able to detect a change every 0.48 °C, but this is in an ideal scenario, as long as we detect a change every degree the system is deemed functional. Long cables will add resistance and will lower the sensitivity.

6.6 Optocouplers

An *optocoupler*, or an *optoisolator*, is an IC that shields one part of a system electrically from another by using light-emitting diodes and phototransistors. They usually operate with digital on-off signals but can be used for analog signaling. The AACPL-847-300 features four optocouplers per IC, and three of them are needed to control the 12 enable signals present on the TC. Given the considerable distance from the MB to the TC, a significant ground potential difference will occur, so board-to-board logic is thus unwanted. An alternative to this is using these optocouplers as a set of switches. When the LED side of the optocoupler is powered, the phototransistor will turn on.

Two pins on the enable signal connector are dedicated to ground, and the other enable signals are connected with pull-up resistors on the TC. Setting the optocoupler pin high on the MCU side will close the internal switch and tie the enable signals low on the TC. The enable signals control whether the voltage regulators on the TC are on or off, making it possible for the MCU to control which strings are turned on or off.

6.7 Stackup

Having a total of six planes for routing and power was necessary to get a good layout. It is substantially more difficult to get a good EMC and power design when working with four or fewer layers. Two layers in the design have no routing and are completely filled with VDD and GND copper. The routing tracks are mainly in the top layer, with some collections of signals routed on layers Signal_2 and Signal_3.

The bottom layer is completely reserved for routing and filtering the power sent to the TC. Having this plane at the bottom makes it easier to place components, careful planning has been done to put as little vias as possible in the areas where the TC power is being routed.

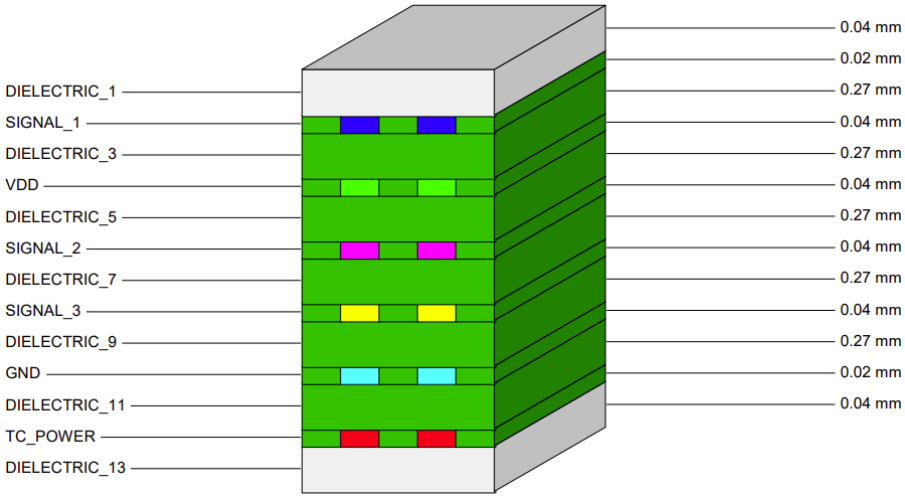

Figure 6.17: Stackup of the MB PCB. Total thickness is 1.6 mm

The signal layers (bottom layer excluded) are filled with ground and VDD pour alternating, creating as much interplane capacitance as possible with the stackup. Another benefit is also having a copper balance in the card, having copper pours on only some layers can cause the board to flex unexpectedly when heated up. The interplane capacitance was measured to be 2.84 nF across the entire pour area, equating to 60 pF per square centimeter. The pour area is 8.6 cm times 5.5 cm.

6.8 Revised Form Factor

An updated revision of the form factor was designed to fit inside a rack-mountable case. The PCBs height is lowered, and its width extended to accommodate for new component placement. To make this possible, the connector for the enable and temperature inputs are moved away from the edge and towards the center of the MB. The changes are visible in Appendix C, revision 0 and revision 1 of the MB.

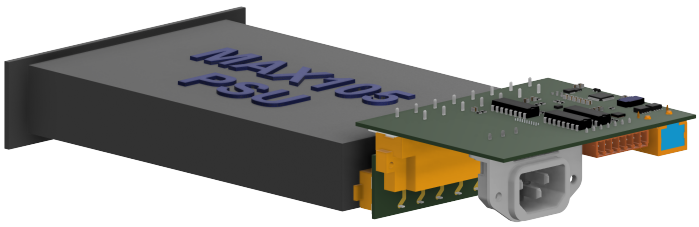

Figure 6.18: MB and PSU connected using an adapterboard.

A 90° adapter board was also made to facilitate mounting the MBs directly onto the MAX105 PSU without needing any cables. The DIN41612 connector holds the

MB in place without needing any additional mounting hardware. Figure 6.18 shows the MB connected to the MAX105 PSU using the adapter board.

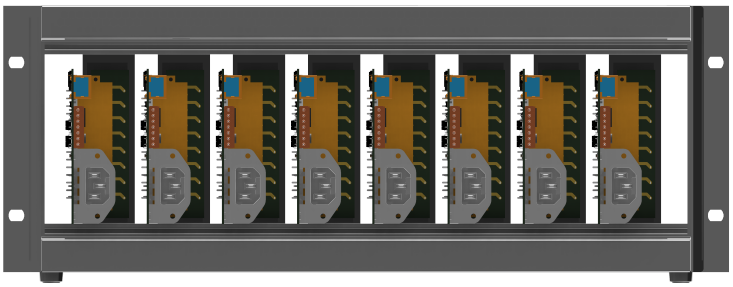

Figure 6.19: Several MBs and MAX105s mounted in a rack case.

Figure 6.19 shows the clearance of the new design. The PSUs are spaced 5 cm from each other. The rack used for demonstration is a RatiopacPRO AIR 19" series case with a width of 426 mm. With eight MBs per case, six cases are needed to house all the PSUs for all the layers. The six cases can all be stored within a single rack tower.

7 SUMMARY, RESULTS, AND DISCUSSION

This chapter evaluates the work described in this thesis. Additionally, it discusses the remaining tasks of the project.

7.1 Summary

The primary objective of this thesis was to choose a suitable power supply for one or several layers of a DTC, while designing a solution to monitor current, monitor temperature, and control several signals going to the TC. The preliminary design this thesis builds upon planned to use IPbus to communicate with a computer control system, but a new design was proposed splitting this task away from the work described in this thesis.

A MAX105 switch mode PSU was chosen as the PSU for its size and cost. One MAX105 is used to power an entire layer. The PSU can deliver 16 A at 5 V, and showed itself to be suitable after a series of tests were run while supplying the ALPIDE chips. The impact the PSU had on both the digital and analog circuitry of the chips was tested and compared to a set benchmark.

The design of a custom PCB to accompany the PSU was made, named the Monitor Board, which slots directly onto the PSU itself. The MB has an AVR128DB48 microcontroller. The MCU monitors current and temperature, generates the PWELL signal, controls the enable signals, and communicates with the user through an IPbus translation layer. A single KCU 115 FPGA talks to all MBs, and the user can access this information over IPbus. The MB communicates with the KCU115 over differential USART using standard ethernet cables.

The MB also features several optocouplers to control the enable signals isolated, a negative voltage supply for the ALPIDEs, a voltage reference for temperature measurements, and a filter to filter out ripple noise in the powerlines going to the TC. The filter was made over several iterations and required careful layout to achieve on the MB.

7.2 Results

The MB design proved functional, but some layout changes were needed to correct mistakes made during the design. The optocouplers had the wrong footprint assigned and were therefore not able to be directly soldered to the PCB, but their functionality was tested using wires. The LT1991 amplifier was also not tested, as there is a chip shortage. The chip was out of stock until this thesis was due. The PWELL negative supply was measured to output -10 V as designed.

The MAX105 PSU has a male DIN41612 style connector, and the original MB design had an angled female DIN41612 connector, making the connection between the MB

and PSU seamless. Unfortunately, the angled female connector is discontinued, and a straight female connector is placed on the MB. A custom adapter board, using one straight female and one angled male DIN41612 connector is used to connect the MB onto the PSU, as having the MB be connected with a 90° bend would make it difficult to have several units side by side. Figure 6.18 shows the MB and PSU connected using the adapter board.

7.2.1 Power Output

The power integrity was measured across the TCs decoupling capacitors with a resistive load instead of an ALPIDE string connected. Measurements of the power integrity after the filter creation were done in Section 5.4.5. Having the filter embedded onto the MB showed identical performance.

7.2.2 Microcontroller

Communication with the AVR128DB48 MCU was successful using the pymcuprog tool. Pinging the MCU gave no error messages, but writing to the flash and then reading it back was somewhat unstable. It seems that the MCU does not always cleanly overwrite old data present in the flash, an advice is to then always wipe the flash before programming.

The MCU was successfully programmed using a HEX file. The AVR toolchain was not set up correctly when compiling the HEX file, meaning the SW could not be verified on the MB, this has to be done in the future.

7.2.3 USART Communication

The USART communication was tested at several baud rates with automated tests written by Håvard Birkenes. The test generates a random 32-bit integer, truncated and written to a 16-bit register on the MCU. After writing a random value to a set register, the test reads the same register 1000 times. The test is done 100 times at the different baud rates. In total, the test performed 100 000 read operations per baud rate. Each read operation sends 24 bits from the MCU to the client SW.

Communication Test

Baud Rate	Bit Error Rate %
57600	0.18%
115200	0.17%
230400	0.18%
460800	0.18%
921600	0
1000000	0

Table 7.1: Communication test with 100 000 read operations per baud rate.

The test shows the system performs surprisingly well at high baud rates and surprisingly bad at low baud rates. Not shown in the table above is the pretty clear divide between good and bad baud rates, which was around 800 000. It seems like the microcontroller skews the waveform for some lower baud rate signals instead of triggering on the rising edge. Further testing was done at a later date, and the clock skew issue was gone but the errors at lower baud rates persisted. Further tests were not conducted due to time constraints.

A 17-hour test was done with a baud rate of 921600, yielding 60 million total write operations. The test concluded with zero errors.

7.3 Future Work

7.3.1 Layout

The layout for the revised form factor design is not completed. The design can change the PCB height and thickness should stay static, but its length can be extended to accommodate for wiring if found difficult. The updated size is outlined in Appendix C.

7.3.2 Microcontroller Software

The software communication was tested and verified, and support structures were written so that implementing functions should be relatively easy. The functions themselves are not final and require further development. Functions such as reading and writing I^2C, setting a PWELL voltage, and reading out the temperature of the PT1000 element, are all yet to be adequately implemented and tested.

When the SW comes closer to a final form, a fine-tuning of the watchdog will have to be done. Currently, it has a large window primarily for debugging. But later, this window should be as tight as possible to minimize the duration for which the MCU can be stuck in a faulty state.

The SW was described such that major redesigns would hopefully not be necessary.

The low baud rate communication should also be thoroughly inspected to shed some light on why low baud rates fail to trigger on the rising edge sent from the FPGA.

7.3.3 Testing and Documentation

A proper long-term test will have to be carried out to assess the functionality of the design properly. Stress tests where a large amount of heat is produced must be done inside a chassis alongside other devices under test.

7.4 Conclusion

This thesis designed and tested a power and filter solution for a single layer of the DTC which was deemed satisfactory. A complete power delivery system design

for all the layers was proposed, designed, and tested in a prototype fashion. The matter of creating a rack-mountable power delivery system for the pCT project was completed.

After creating a prototype MB, both size and layout changes are necessary, which have been properly documented in both this and the former chapter. Appendix C shows both the updated size and an example of the design mounted in a case.

References

[1] SSB.no. *Befolkning*. 2022. URL: https://www.ssb.no/befolkning/folketall/statistikk/befolkning (visited on 05/16/2022).

[2] Kreftregisteret.no. *Kreft i Norge*. 2021. URL: https://www.kreftregisteret.no/Temasider/om-kreft/ (visited on 05/16/2022).

[3] Cancer.gov. *Adult Central Nervous System Tumors Treatment (PDQ) - Health Professional Version*. 2021. URL: https://www.cancer.gov/types/brain/hp/adult-brain-treatment-pdq (visited on 05/16/2022).

[4] Harald Paganetti. *Proton therapy physics*. Boca Raton, FL: CRC Press/Taylor & Francis, 2012. ISBN: 9781439836453.

[5] T S Hong et al. "Intensity-modulated radiation therapy: emerging cancer treatment technology". In: *British Journal of Cancer* 92.10 (Apr. 2005), pp. 1819–1824. DOI: 10.1038/sj.bjc.6602577. URL: https://doi.org/10.1038/sj.bjc.6602577.

[6] Thomas Bortfeld. "The number of beams in IMRT—theoretical investigations and implications for single-arc IMRT". In: *Physics in Medicine and Biology* 55.1 (Nov. 2009), pp. 83–97. DOI: 10.1088/0031-9155/55/1/006. URL: https://doi.org/10.1088/0031-9155/55/1/006.

[7] Abigail T. Berman, Sara St. James, and Ramesh Rengan. "Proton Beam Therapy for Non-Small Cell Lung Cancer: Current Clinical Evidence and Future Directions". In: *Cancers* 7.3 (2015), pp. 1178–1190. ISSN: 2072-6694. DOI: 10.3390/cancers7030831. URL: https://www.mdpi.com/2072-6694/7/3/831.

[8] M. Bryant et al. Curtis. "Consensus Statement on Proton Therapy for Prostate Cancer". In: *International Journal of Particle Therapy* 8 (2021), pp. -. DOI: https://doi.org/10.14338/IJPT-20-00031.1.

[9] Alme. Johan et al. In: *Frontiers in Physics* 8 (2020), p. 460. ISSN: 2296-424X. DOI: 10.3389/fphy.2020.568243. URL: https://www.frontiersin.org/article/10.3389/fphy.2020.568243.

[10] Tea Bodova. "High-Speed Signal and Power Distribution of a Digital Tracking Calorimeter for Proton Computed Tomography". MA thesis. Bergen: University of Bergen, 2020.

[11] Debra L. Friedman et al. "Subsequent Neoplasms in 5-Year Survivors of Childhood Cancer: The Childhood Cancer Survivor Study". In: *JNCI: Journal of the National Cancer Institute* 102.14 (July 2010), pp. 1083–1095. DOI: 10.1093/jnci/djq238. URL: https://doi.org/10.1093/jnci/djq238.

[12] Christine S. Chung et al. "Incidence of Second Malignancies Among Patients Treated With Proton Versus Photon Radiation". In: *International Journal of Radiation Oncology*Biology*Physics* 87.1 (2013), pp. 46–52. ISSN: 0360-3016. DOI: https://doi.org/10.1016/j.ijrobp.2013.04.030. URL: https://www.sciencedirect.com/science/article/pii/S0360301613004598.

[13] US Food and Drug Administration. *Computed Tomography (CT)*. 2019. URL: https://www.fda.gov/radiation-emitting-products/medical-x-ray-imaging/computed-tomography-ct (visited on 01/11/2022).

[14] Anil Kalra. "Developing FE Human Models From Medical Images". In: *Basic Finite Element Method as Applied to Injury Biomechanics*. Elsevier, 2018, pp. 389–415. DOI: 10.1016/b978-0-12-809831-8.00009-x. URL: https://doi.org/10.1016/b978-0-12-809831-8.00009-x.

[15] Don F. DeJongh et al. "A Comparison of Proton Stopping Power Measured with Proton CT and x-ray CT in Fresh Post-Mortem Porcine Structures". In: (2020). DOI: 10.48550/ARXIV.2012.06629. URL: https://arxiv.org/abs/2012.06629.

[16] J.K. van Abbema et al. "High accuracy proton relative stopping power measurement". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 436 (Dec. 2018), pp. 99–106. DOI: 10.1016/j.nimb.2018.09.015. URL: https://doi.org/10.1016/j.nimb.2018.09.015.

[17] D. R. Dance. *Diagnostic radiology physics : a handbook for teachers and students*. Vienna: International Atomic Energy Agency, 2014. ISBN: 9789201310101.

[18] Sebastian Zschaeck et al. "PRONTOX – proton therapy to reduce acute normal tissue toxicity in locally advanced non-small-cell lung carcinomas (NSCLC): study protocol for a randomised controlled trial". In: *Trials* 17.1 (Nov. 2016). DOI: 10.1186/s13063-016-1679-4. URL: https://doi.org/10.1186/s13063-016-1679-4.

[19] Abigail Berman, Sara James, and Ramesh Rengan. "Proton Beam Therapy for Non-Small Cell Lung Cancer: Current Clinical Evidence and Future Directions". In: *Cancers* 7.3 (July 2015), pp. 1178–1190. DOI: 10.3390/cancers7030831. URL: https://doi.org/10.3390/cancers7030831.

[20] Jose R. Alonso and Timothy A. Antaya. "Superconductivity in Medicine". In: *Reviews of Accelerator Science and Technology*. WORLD SCIENTIFIC, Mar. 2013, pp. 227–263. DOI: 10.1142/9789814449953_0009. URL: https://doi.org/10.1142/9789814449953_0009.

[21] M. Prall et al. "High-energy proton imaging for biomedical applications". In: *Scientific Reports* 6.1 (June 2016). DOI: 10.1038/srep27651. URL: https://doi.org/10.1038/srep27651.

[22] Mark D. Brooke and Scott N. Penfold. "An inhomogeneous most likely path formalism for proton computed tomography". In: *Physica Medica* 70 (Feb. 2020), pp. 184–195. DOI: 10.1016/j.ejmp.2020.01.025. URL: https://doi.org/10.1016/j.ejmp.2020.01.025.

[23] Ola Grøttvik. "Design and Implementation of a High-Speed Readout and Control System for a Digital Tracking Calorimeter for proton CT". MA thesis. Bergen: University of Bergen, 2021.

[24] V. M. Borshov et al. "Innovative microelectronic technologies for high-energy physics experiments". In: (2016).

[25] R. Turchetta et al. “CMOS Monolithic Active Pixel Sensors (MAPS): development and future outlook”. In: (), pp. -.

[26] M. Mager. “ALPIDE, the Monolithic Active Pixel Sensor for the ALICE ITS upgrade”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 824 (July 2016), pp. 434–438. DOI: 10.1016/j.nima.2015.09.057. URL: https://doi.org/10.1016/j.nima.2015.09.057.

[27] B Abelev et al and. “Technical Design Report for the Upgrade of the ALICE Inner Tracking System”. In: *Journal of Physics G: Nuclear and Particle Physics* 41.8 (July 2014), p. 087002. DOI: 10.1088/0954-3899/41/8/087002. URL: https://doi.org/10.1088/0954-3899/41/8/087002.

[28] *ALPIDE Operations Manual*. ALICE ITS ALPIDE development team. 2016. URL: http://sunba2.ba.infn.it/MOSAIC/ALICE-ITS/Documents/ALPIDE-operations-manual-version-0_3.pdf.

[29] Wikipedia. *DC Bias*. 2021. URL: https://en.wikipedia.org/wiki/DC_bias (visited on 04/06/2022).

[30] Thomas Økland. “Setup and integration of the Readout Electronics for the Pixel Layers of the New ALICE FoCal Detector”. MA thesis. Bergen: University of Bergen, 2022.

[31] C. Ghabrous Larrea et al. “IPbus: a flexible Ethernet-based control system for xTCA hardware”. In: *Journal of Instrumentation* 10.02 (Feb. 2015), pp. C02019–C02019. DOI: 10.1088/1748-0221/10/02/c02019. URL: https://doi.org/10.1088/1748-0221/10/02/c02019.

[32] Peter Wilson. *The circuit designer’s companion*. Oxford: Newnes, 2011. ISBN: 978-0080971384.

[33] John Rumble, ed. *CRC handbook of chemistry and physics*. 101st ed. CRC Handbook of Chemistry and Physics. London, England: CRC Press, July 2020.

[34] Neil Weste and David Harris. *INTEGRATED CIRCUIT DESIGN*. en. 4th ed. Philadelphia, PA: Pearson Education, Apr. 2010.

[35] C Sankaran. *Power quality*. Boca Raton: CRC Press, 2002. ISBN: 0-8493-1040-7.

[36] Lee Ritchey. *Right the first time*. Speeding Edge, 2006. ISBN: 0-9741936-1-5.

[37] Henry Ott. *Electromagnetic compatibility engineering*. Hoboken, N.J: John Wiley & Sons, 2009. ISBN: 9780470189306.

[38] Eric Bogatin. *Signal and power integrity, simplified*. Boston, MA: Prentice Hall, 2018. ISBN: 9780134513416.

[39] *3A Fast Response LDO Regulator*. Microchip. 2018. URL: https://no.mouser.com/datasheet/2/268/20005897B-1890902.pdf.

[40] Norsk Elektronisk Komite. “NEK IEC 60598-1”. In: 7.0 (2008).

[41] *AVR DB Family*. AVR. 2020. URL: https://ww1.microchip.com/downloads/aemDocuments/documents/MCU08/ProductDocuments/DataSheets/AVR128DB28-32-48-64-DataSheet-DS40002247A.pdf.

[42] *RS-485/J1708 Transceivers*. MAXXIM. 2014. URL: https://datasheets.maximintegrated.com/en/ds/MAX3440E-MAX3444E.pdf.

[43] Microchip. *UPDI Physical Interface*. URL: https://onlinedocs.microchip.com/pr/GUID-DDB0017E-84E3-4E77-AAE9-7AC4290E5E8B-en-US-4/index.html?GUID-9B349315-2842-4189-B88C-49F4E1055D7F (visited on 05/25/2022).

[44] *INA3221 Triple-Channel, High-Side Measurement, Shunt and Bus Voltage Monitor with I2C- and SMBUS-Compatible Interface*. Texas Instruments. 2015. URL: https://www.ti.com/lit/ds/symlink/ina3221.pdf.

[45] *Compact, Dual-Output Charge Pump*. Maxim. 1997. URL: https://datasheets.maximintegrated.com/en/ds/MAX865.pdf.

[46] Maxim. *GUIDE TO INTEGRATED CHARGE PUMP DC-DC CONVERSION*. 2009. URL: https://www.maximintegrated.com/en/design/technical-documents/tutorials/7/725.html (visited on 06/06/2022).

[47] *Gain Selectable Amplifier*. Linear Technology. 2012. URL: https://www.analog.com/media/en/technical-documentation/data-sheets/1991fh.pdf.

Appendices

A ACRONYMS

ALICE A Large Ion Collider Experiment

ALPIDE ALICE Pixel Detector

CMOS Complementary metal-oxide-semiconductor

DAQ Data Acquisition

DTC Digital Tracking Calorimeter

DTU Digital Transmission Unit

EMI ElectroMagnetic Interference

FPC Flexible Printed Circuit

FPGA Field Programmable Gate Array

GUI Graphical User Interface

HU Hounsfield Unit

ITS Inner Tracking System

LVDS Low-Voltage Differential Signaling

MAPS Monolithic Active Pixel Sensor

Mbps MegaBit Per Second

MII Media-Independent Interface

MLCC MultiLayer Ceramic Capacitor

PCB Printed Circuit Board

PHY Physical Interface

PSU Power Supply Unit

PWM Pulse Width Modulation

SOBP Spread Out Bragg Peak

TC Transition Card

UPDI Unified Program and Debug Interface

B SOFTWARE

This appendix has some essential information regarding the software, things like the final address map used, a preliminary error list suggestion, and one suggestion of the control 1 register.

B.1 Adressmap

Register Name	Adress	Number of bits	Unit	Access	Default value	Description
FW_VERSION	0x00	8	-	R	-	Current firmware version.
ADC_VALUE	0x02	12	V/V	R/$\overline{\text{W}}$	-	12-bit ADC value.
PT100_READING	0x04	8	°C	R	-	ADC value converted to degrees celsius.
TEMPERATURE_LIMIT	0x06	8	°C	R/$\overline{\text{W}}$	100°C	On measuring a temperature above TEMPERATURE_LIMIT the enable signals will be set low. Error 0x01.
DAC_VALUE	0x08	10	V/V	R/$\overline{\text{W}}$	0x00	DAC voltage output from the MCU, input from 0x00 to 0x400 creates an output of 0 V to 2.5 V. NOTE: Read about the PWELL generation before changing. DOUBLE NOTE: The DAC utilises the 10 most significant bits!
PWELL_VOLTAGE_MCU	0x0A	13	mV	$\overline{\text{W}}$	0x00	The desired PWELL voltage in millivolt. Writing a value to this registers triggers the MCU to create the complementary voltage on the PWELL line.
DVDD_CURRENT_THRESHOLD1	0x0C	14	mA	R/$\overline{\text{W}}$	0x00	INA3221 DVDD critical threshold. If the DVDD line exceed this current draw the enable signals are set low. Error 0x02
DVDD_CURRENT_THRESHOLD2	0x0E	14	mA	R	0x00	DVDD warning threshold. Error 0x03
DVDD_VOLTAGE	0x10	13	mV	R	-	DVDD shunt resistor measured voltage.
DVDD_CURRENT	0x12	13	mA	R	-	DVDD shunt resistor measured current.
AVDD_CURRENT_THRESHOLD1	0x14	14	mA	R/$\overline{\text{W}}$	0x00	INA3221 AVDD critical threshold. If the AVDD line exceed this current draw the enable signals are set low. Error 0x04
AVDD_CURRENT_THRESHOLD2	0x16	14	mA	R/$\overline{\text{W}}$	0x00	INA3221 AVDD warning threshold. Error 0x05
AVDD_VOLTAGE	0x18	13	mV	R	-	AVDD Voltage
AVDD_CURRENT	0x1A	13	mA	R	-	AVDD CURRENT
PWELL_CURRENT_THRESHOLD1	0x1C	14	mA	R/$\overline{\text{W}}$	0x00	INA3221 PWELL critical threshold. If the PWELL line exceed this current draw the enable signals are set low. Error 0x06

Register Name	Adress	Number of bits	Unit	Access	Default value	Description
PWELL_CURRENT_THRESHOLD2	0x1E	14	mA	R/$\overline{W}$	0x00	INA3221 AVDD warning threshold. Error 0x07
PWELL_VOLTAGE_INA3221	0x20	14	mV	R	-	PWELL voltage measured by the INA3221
PWELL_CURRENT	0x22	14	mV	R	-	PWELL Current measured by the INA3221
ENABLE_SIGNALS	0x24	12	-	R/$\overline{W}$	0x00	Each bit represents an enable line controlling the power to a string. String 0 is tied to LSB.
STRING_DVDD_CURRENT_VALUE[n]	0x26+[2n]	13·12	mA	R	0x00	The DVDD current values for each string after the scan flag has been asserted. In total 12 registers with 13 bytes each. Each string register takes two bytes, and the address for string n is offset by 2n bytes.
STRING_AVDD_CURRENT_VALUE[n]	0x3E+[2n]	13·12	mA	R	0x00	The AVDD current values for each string after the scan flag has been asserted. In total 12 registers with 13 bytes each. Each string register takes two bytes, and the address for string n is offset by 2n bytes.
STRING_PWELL_CURRENT_VALUE[n]	0x56+[2n]	13·12	mA	R	0x00	The PWELL current values for each string after the scan flag has been asserted. In total 12 registers with 13 bytes each. Each string register takes two bytes, and the address for string n is offset by 2n bytes.

Register Name	Adress	Number of bits	Unit	Access	Default value	Description
CTRL1	0x6E	8	-	R/$\overline{\text{W}}$	-	Control register A
CTRL2	0x70	8	-	R/$\overline{\text{W}}$	-	Control register B
CTRL3	0x72	8	-	R/$\overline{\text{W}}$	-	Control register C
ERROR_COUNT	0x74	8	-	R/$\overline{\text{W}}$	0x00	Amount of errors since last clear. Is cleared by writing 0x00 to this register.
ERROR_MSG	0x76	128	-	R	0x00	Error messages stored in sequence, each byte is an error message. The most reccent error is placed in LSB. Automatically cleared by clearing ERROR_COUNT

B.2 Error Codes

Error name	Error code	Description
Temperature limit reached	0x01	The ADC value is reported to be above the temperature set by the TEMPERATURE_LIMIT register.
DVDD critical current	0x02	Critical current reached on DVDD line.
DVDD warning current	0x03	Warning current reached on DVDD line.
AVDD critical current	0x04	Critical current reached on AVDD line.
AVDD warning current	0x05	Warning current reached on AVDD line.
PWELL critical current	0x06	Critical current reached on PWELL line.
PWELL warning current	0x07	Warning current reached on PWELL line.
Write/Read denied	0x08	Tried to write or read a register that should not have been read or written.
String current error	0x09	Large current draws from the strings recorded.
Enable scan error	0x0A	Large current draws from a single string recorded during the enable scan.

B.3 Control Registers

ResRegs				ErrorRes	SoftStart	EnScan	EnOff

Figure B.1: CTRL1 register.

Bit 7 - ResRegs: Reset Registers
Resets all registers back to its default values(check the registermap for default values).

Bit 6 - :

Bit 5 - :

Bit 4 - :

Bit 3 - ErrorRes: Reset Errormessages
Clears the error message registers completely and resets the ERROR_COUNT pointer to 0.

Bit 2 - SoftStart: Soft Startup
Soft startup initialization. On asserting the bit the MCU turns off all enable signals and writes the current and voltage values to the INA modules. The MCU loops through the strings one by one and saves the current draw. On detection of large current draws from a string error 0x09 is written to the error register.

Bit 1 - EnScan: Enable Scan
Performs a scan of all the strings while logging the current values in the STRING_CURRENT_VALUEn register. The values are not reported back automatically. On completion the enable signals are tied low(off).

Bit 0 - EnOff: Enable Off
Turns off all enable lines. A bit faster than writing to the ENABLE_SIGNALS register as it is stored in a single byte, also ensures that all values are 0.

C TECHNICAL DRAWINGS AND FIGURES

This appendix shows technical drawings of the monitorboard revision 0 and 1, as well as the drawings for the adapterboard. Additional 3D models are present showing the MB and MAX105 pair mounted in a rack.

TOP

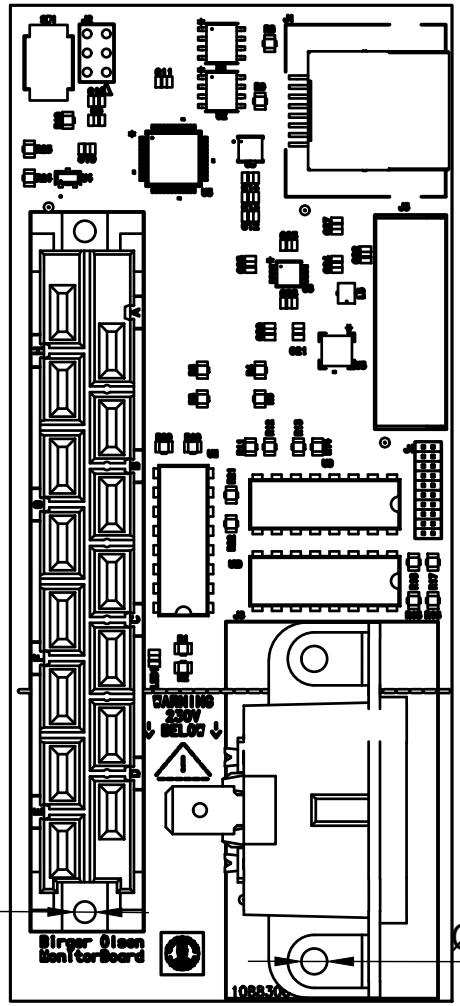

BOTTOM

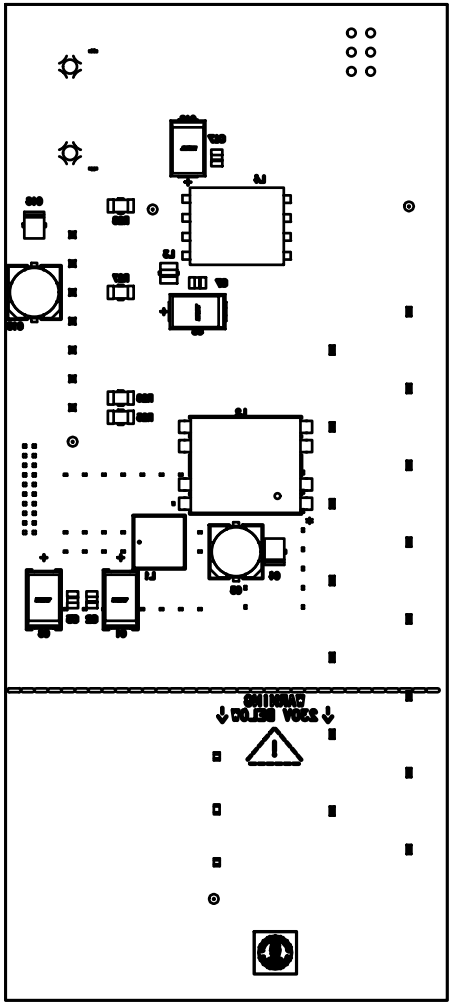

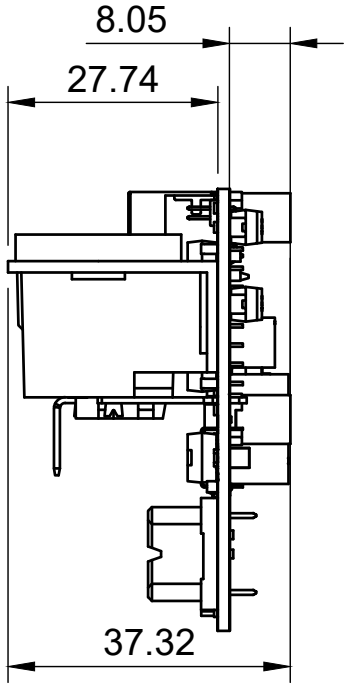

Dept. UiB	Technical reference	Created by Birger Olsen 13/05/2022	Approved by
Size 1:1 All Measurements in mm		Document type	Document status
		Title Board1	DWG No.
			Rev. 0 / Date of issue / Sheet 1/1

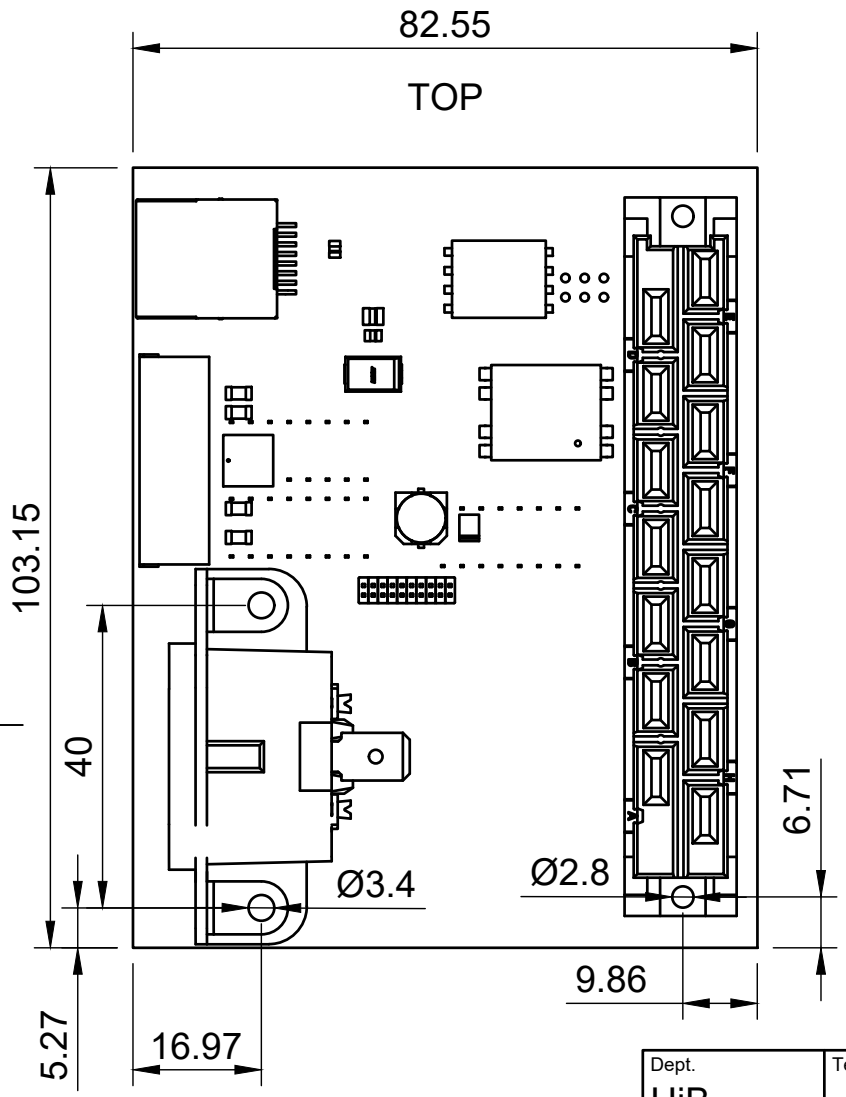

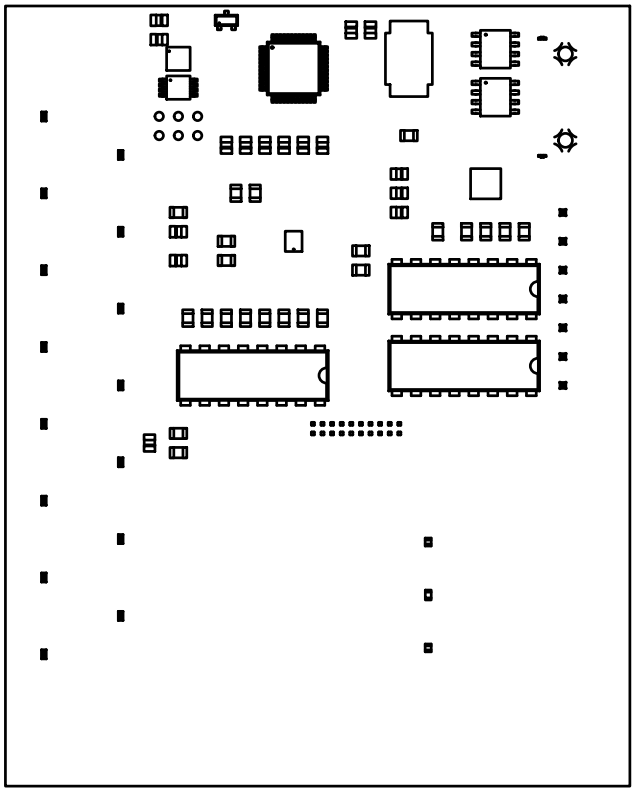

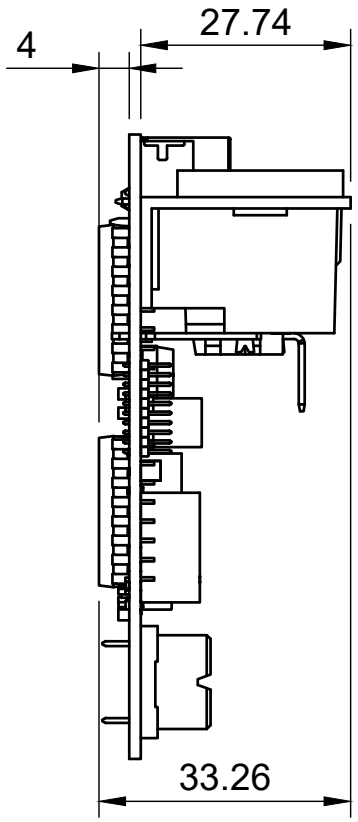

****COMPONENT PLACEMENT NOT FINAL**

Dept. UiB	Technical reference	Created by Birger Olsen 13/05/2022	Approved by
Scale 1:1		Document type Technical Drawing	Document status
		Title MonitorBoard	DWG No.
			Rev. 1 / Date of issue / Sheet 1/1

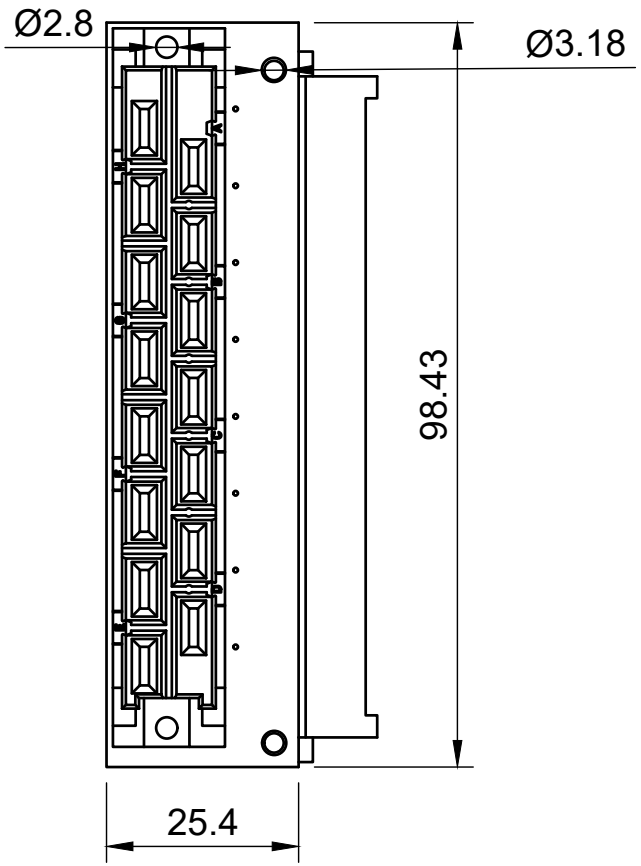

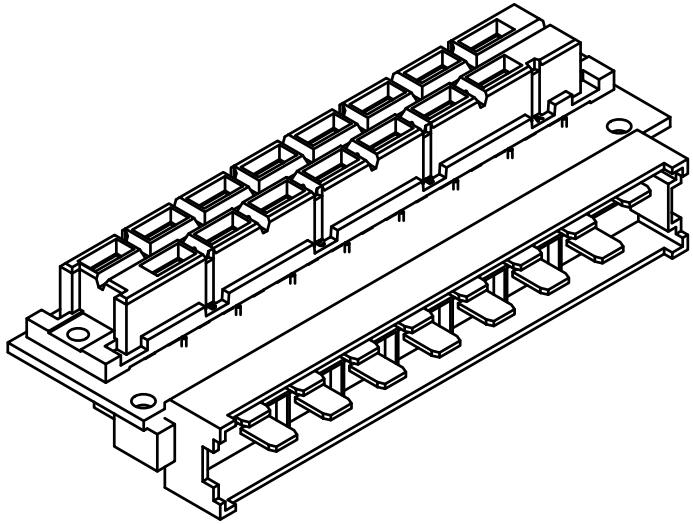

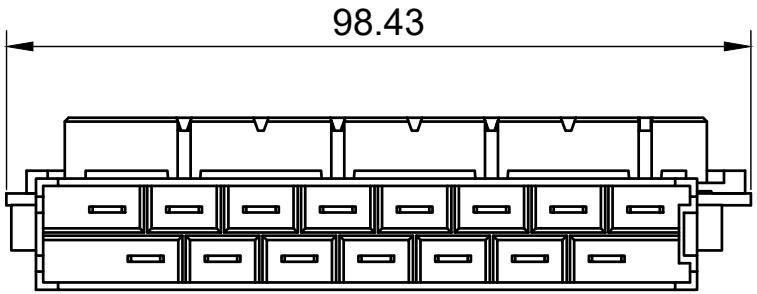

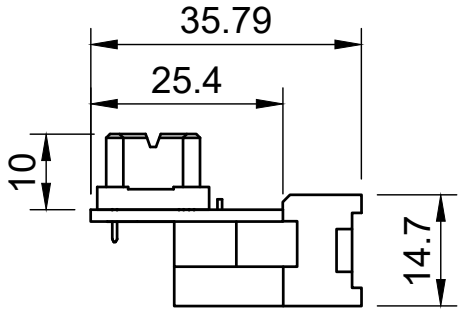

Dept. UiB	Technical reference	Created by Birger Olsen 13/05/2022	Approved by
Size 1:1 2 Layer PCB Measurements in mm		Document type Technical drawing	Document status
		Title **AdapterBoard**	DWG No.
			Rev. 0 / Date of issue / Sheet 1/1

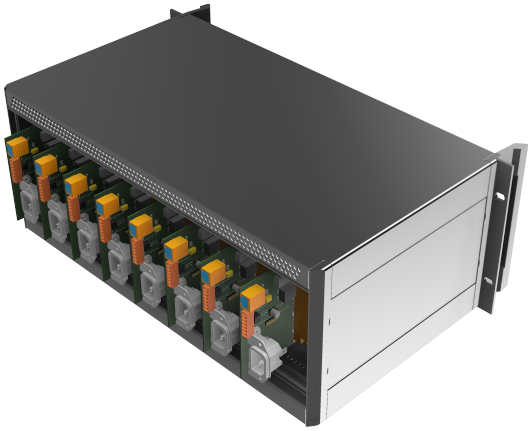

Figure C.1: MB and MAX105 mounted in a rack case.

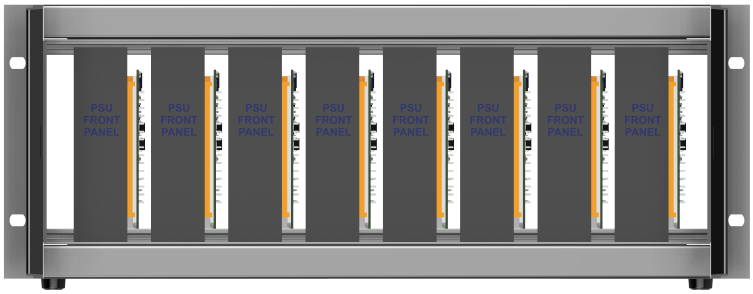

Figure C.2: MB and MAX105 mounted in a rack case. Front view.

D SCHEMATIC

Entire block schematic of the monitor board. Small tweaks have been done since the prototype PCB was manufactured and tested.

InputRippleFilter

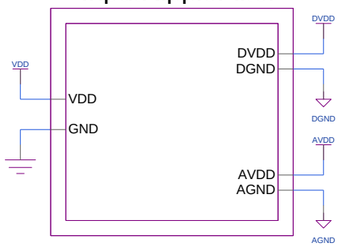

PT100_Sensing

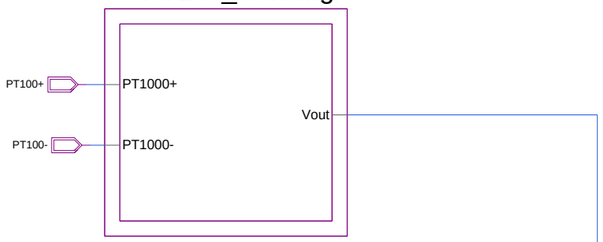

CurrentMonitor

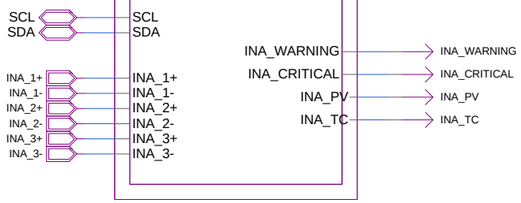

PSU Connector

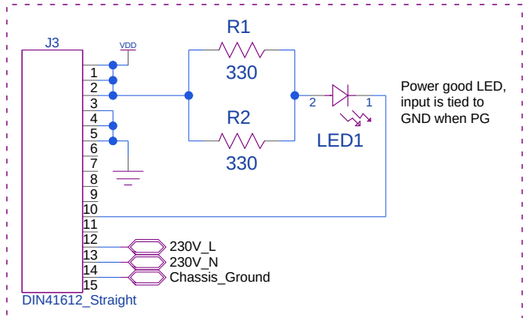

LVDS_Drivers

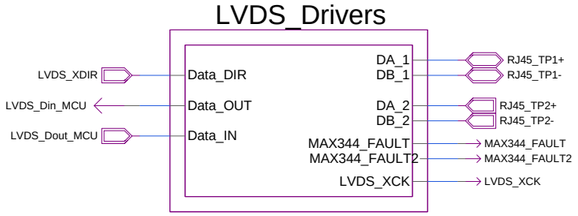

Microcontroller

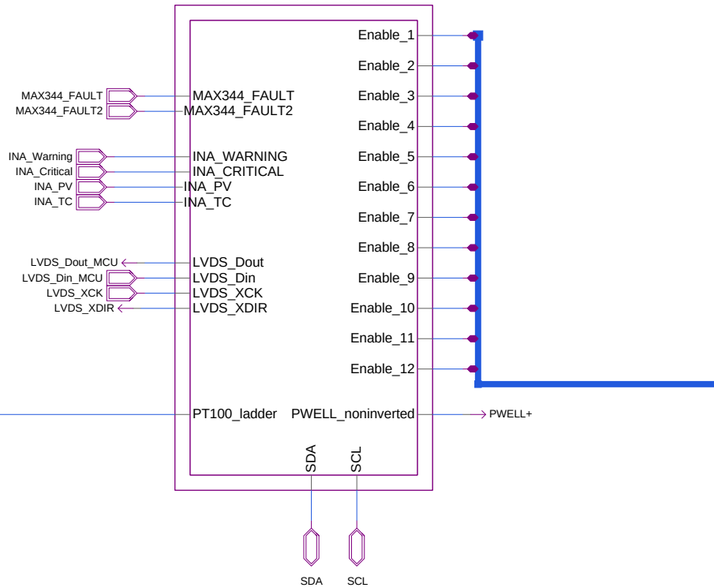

InvertingAmplifier

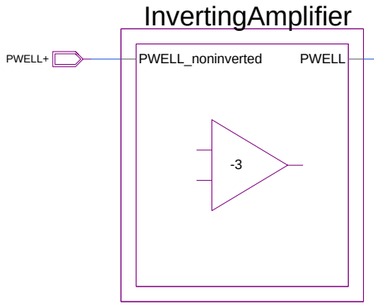

ShuntResistors

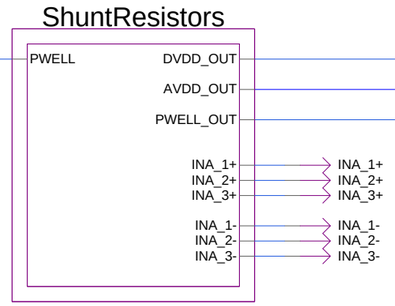

OptoCouplers

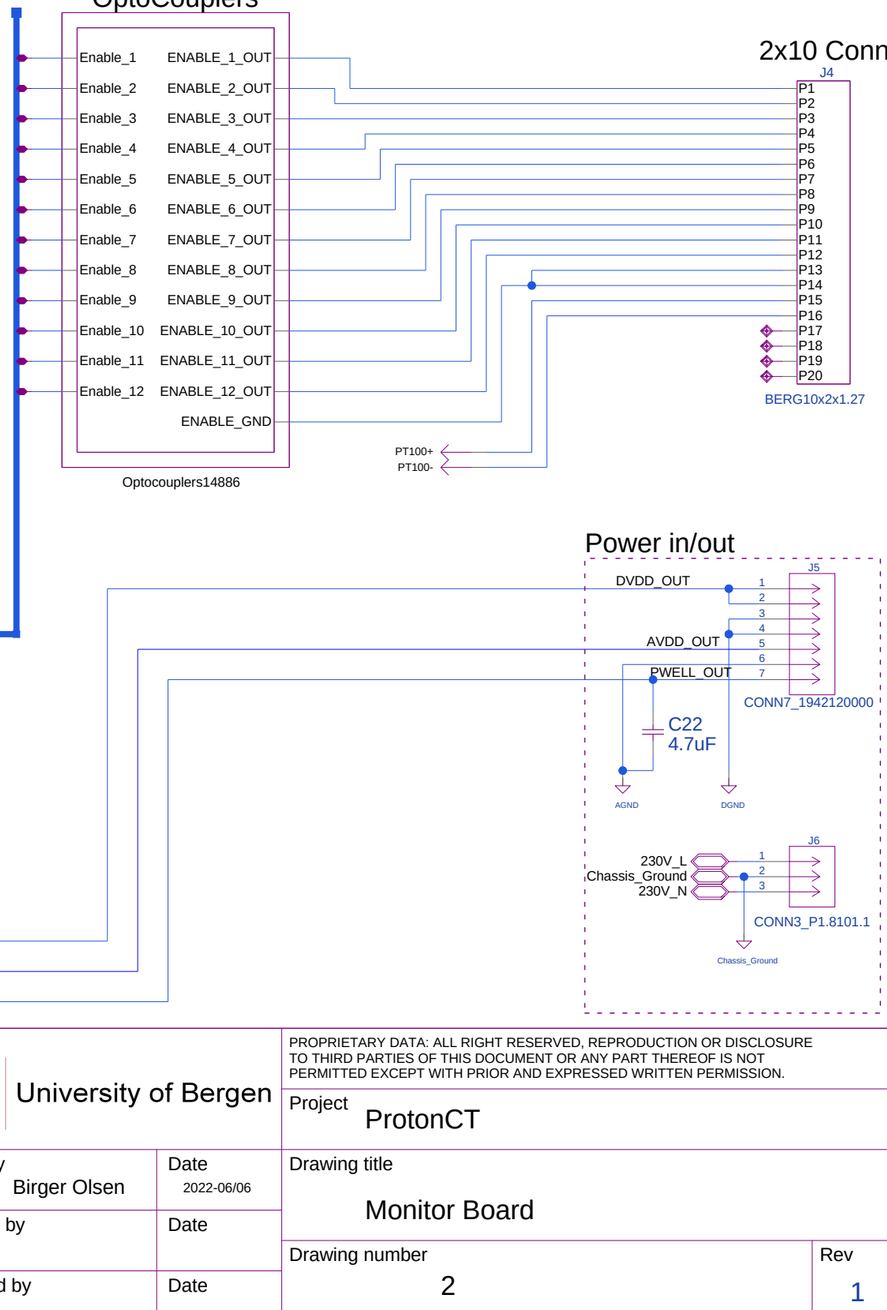

2x10 Conn

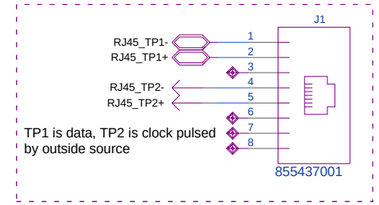

Power in/out

University of Bergen

Project	ProtonCT
Drawn by: Birger Olsen	Date: 2022-06/06
Drawing title	Monitor Board
Checked by	Date
Drawing number	2
Rev	1
Approved by	Date
PA approved by	Date
Sheet title	Schematic1 / 1
Sheet	1/9

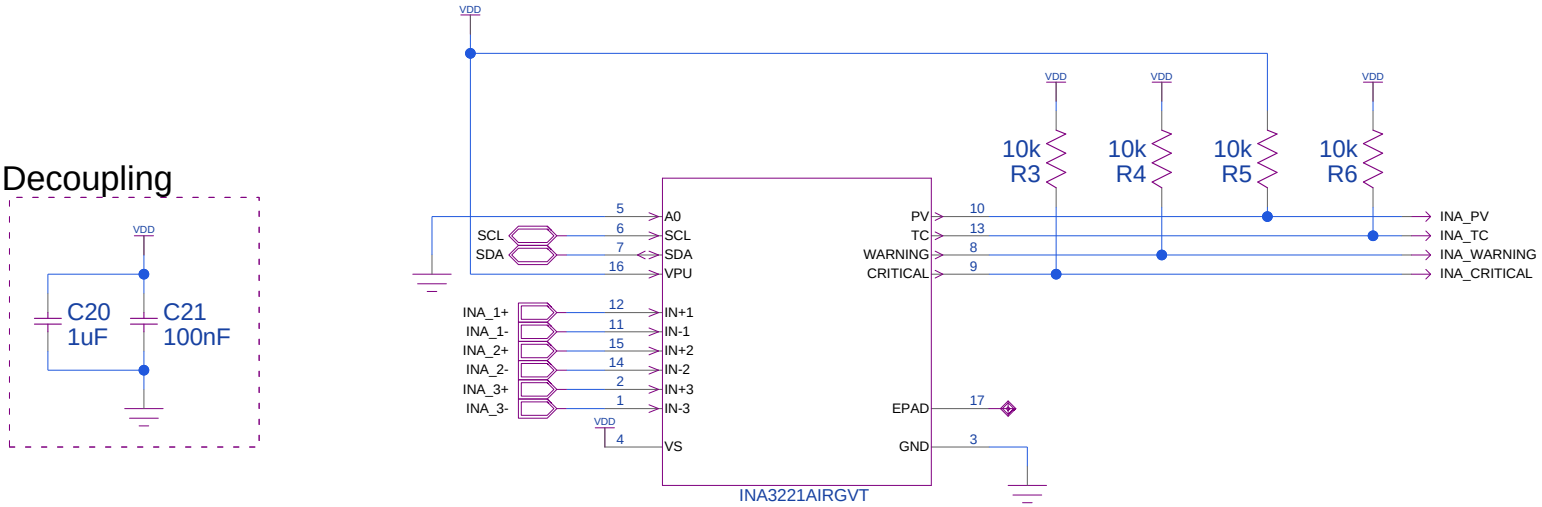

University of Bergen

Project: ProtonCT

Drawn by: Birger Olsen	Date: 2022-06/06	Drawing title: Current Monitor	
Checked by	Date		
Approved by	Date	Drawing number: 1	Rev: 0
PA approved by	Date	Sheet title: CurrentMonitor / 1	Sheet: 2/9

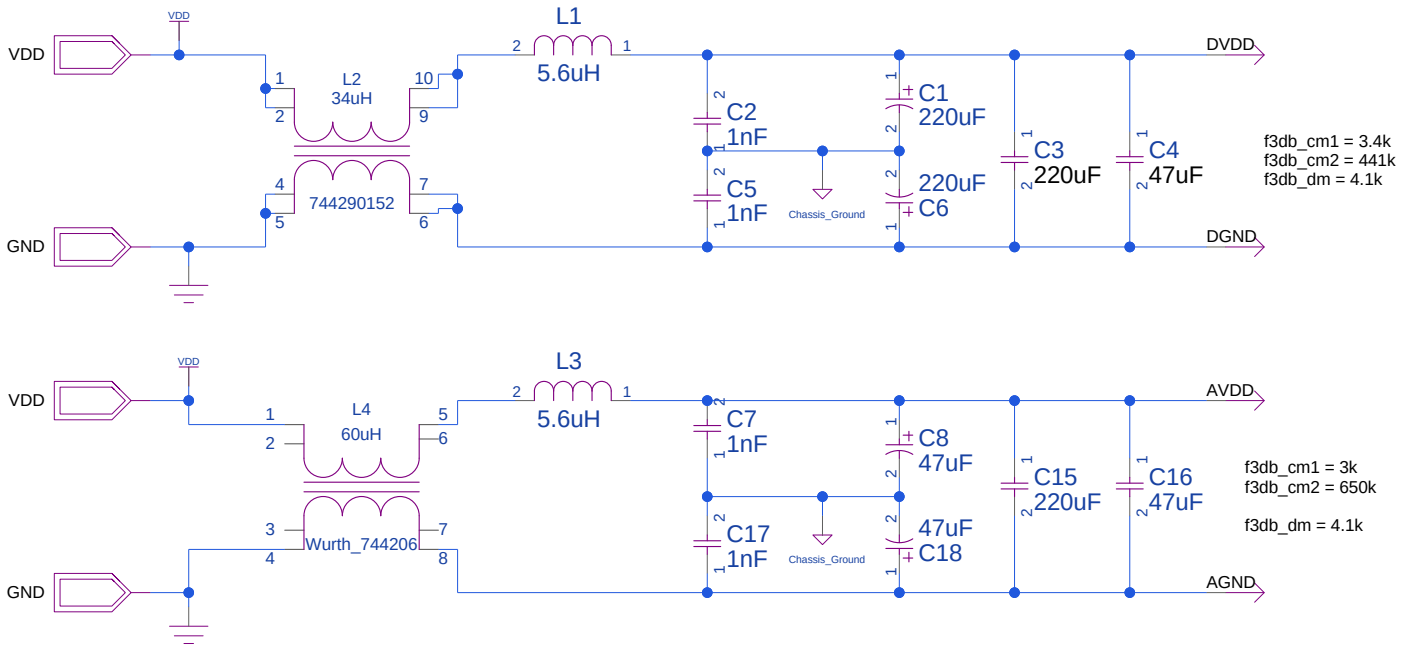

University of Bergen

Project: ProtonCT

Drawn by: Birger Olsen	Date: 2022-06/06	Drawing title:	
Checked by	Date		
Approved by	Date	Drawing number: 1	Rev: 0
PA approved by	Date	Sheet title: Input_Ripple_Filter / 1	Sheet: 3/9

Switched capacitor voltage doubler

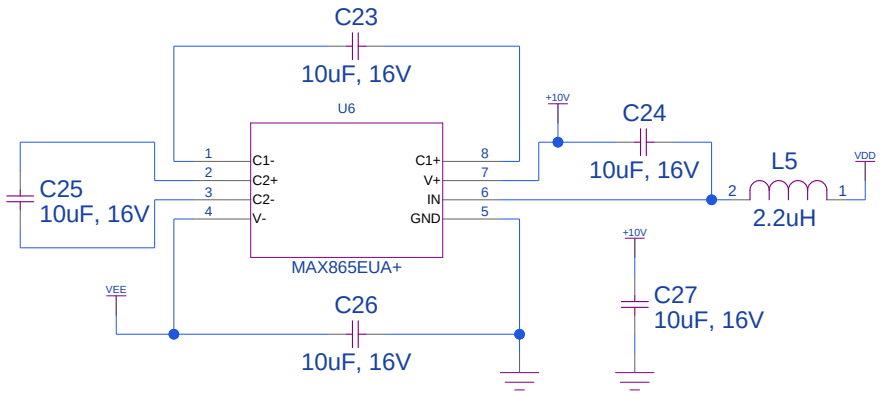

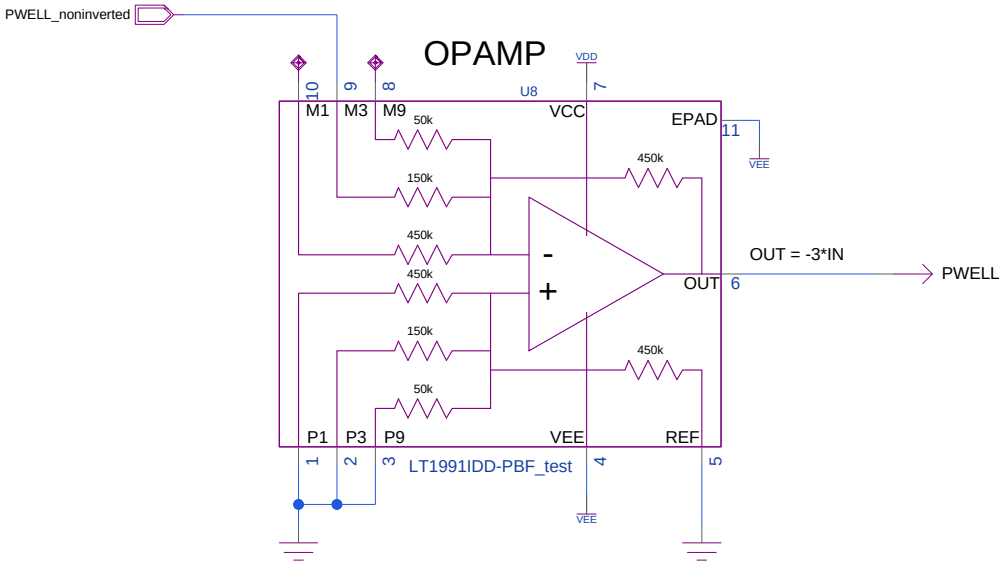

University of Bergen		PROPRIETARY DATA: ALL RIGHT RESERVED, REPRODUCTION OR DISCLOSURE TO THIRD PARTIES OF THIS DOCUMENT OR ANY PART THEREOF IS NOT PERMITTED EXCEPT WITH PRIOR AND EXPRESSED WRITTEN PERMISSION.
		Project: ProtonCT
Drawn by: Birger Olsen	Date: 2022-06-06	Drawing title: Inverting Amplifier
Checked by	Date	
Approved by	Date	Drawing number: 1 — Rev: 0
PA approved by	Date	Sheet title: NAME / — Sheet: 4/8

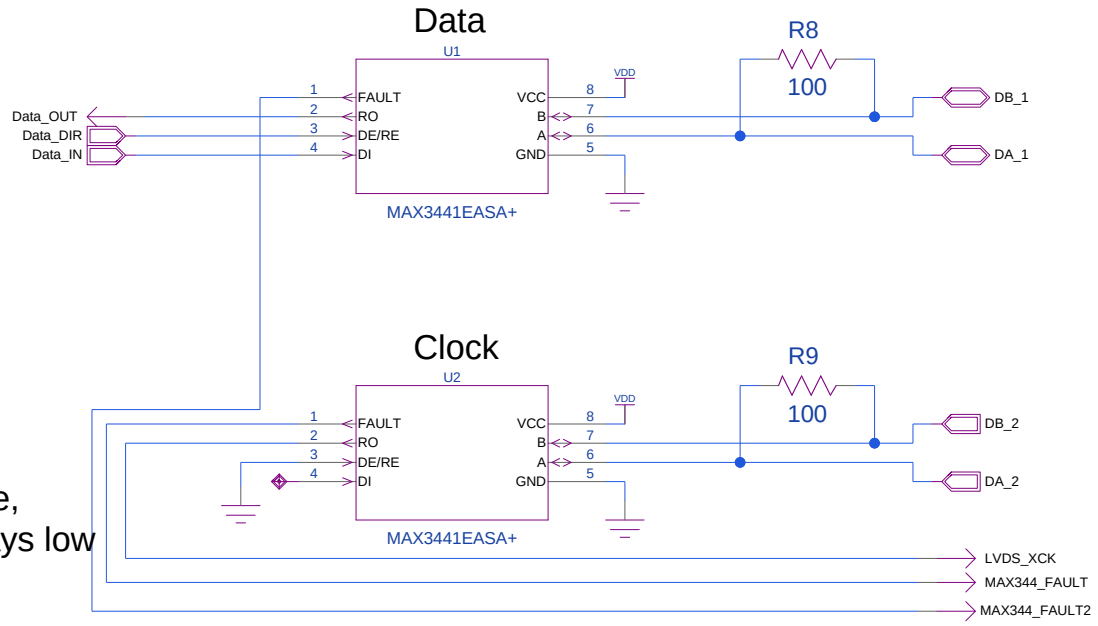

The MCU should never be sending a clock pulse, DE/RE is therefore always low

University of Bergen

Project: ProtonCT

Drawn by: Birger Olsen	Date: 2022-06/06	Drawing title: Differential Signaling Drivers
Checked by	Date	
Approved by	Date	Drawing number: 1 / Rev: 0
PA approved by	Date	Sheet title: LVDS_Drivers / 1 / Sheet: 5/8

Decoupling

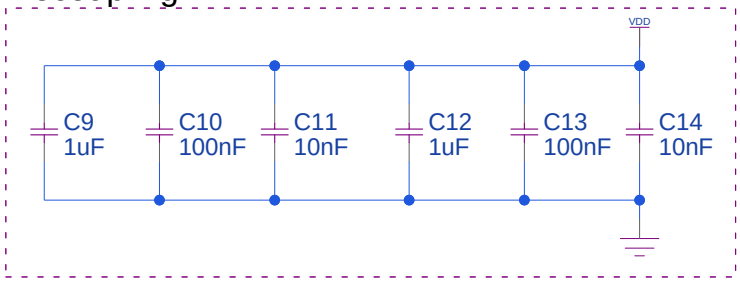

Reset switch

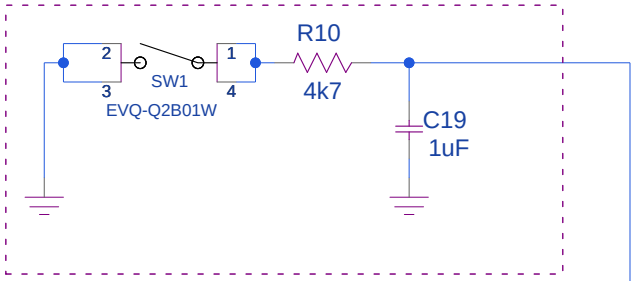

Programming

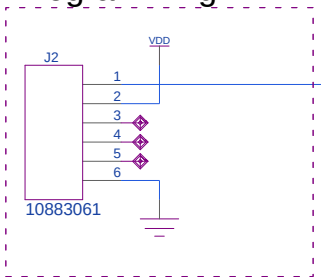

MCU

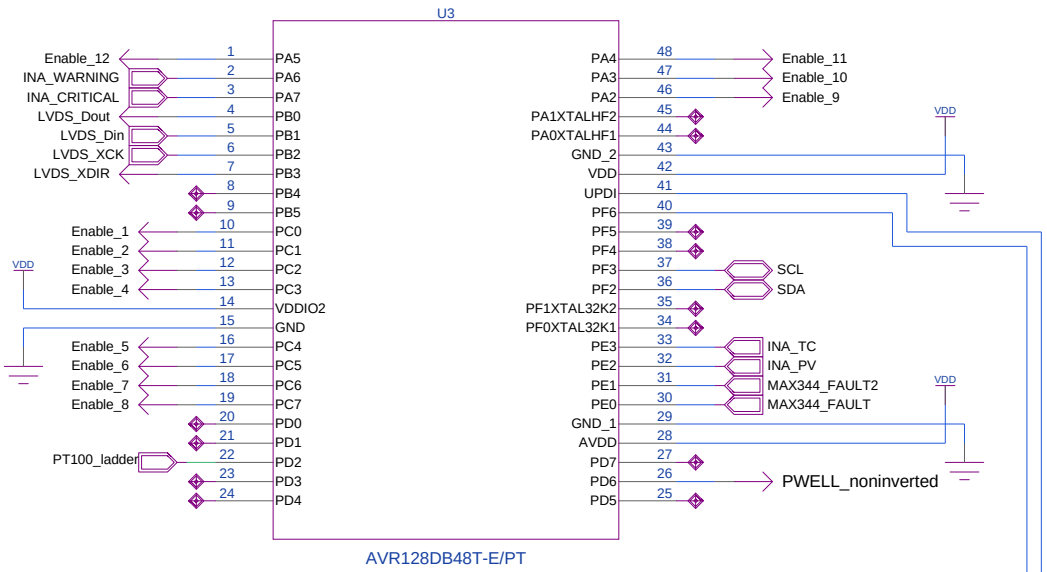

University of Bergen

Project: ProtonCT		
Drawn by: Birger Olsen	Date: 2022-06/06	Drawing title: Microcontroller
Checked by	Date	
Approved by	Date	Drawing number: 1 — Rev: 0
PA approved by	Date	Sheet title: Microcontroller / 1 — Sheet: 6/8

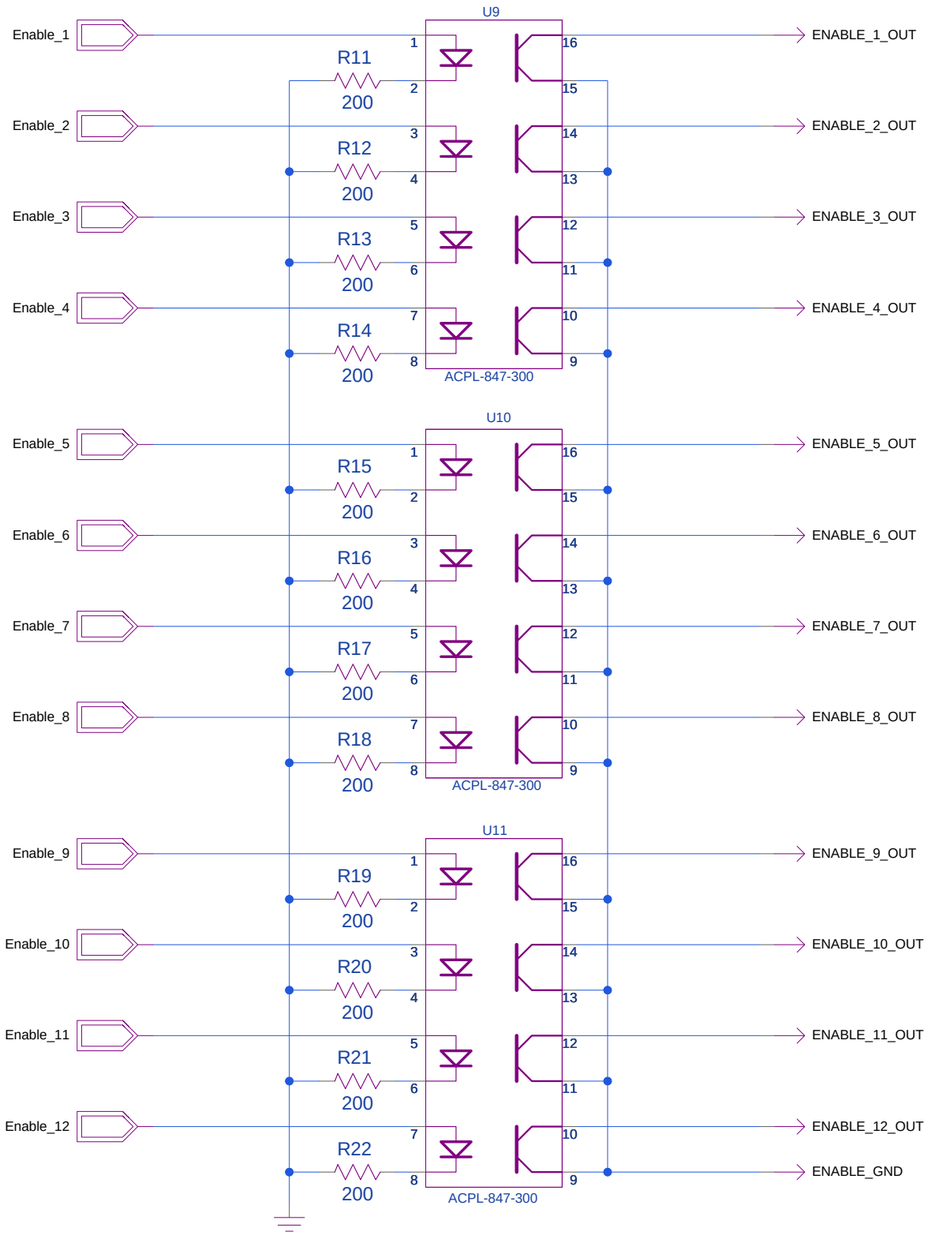

If = 20mA
Vf = 1.2V

V/R = 3.8/200 = 19mA

University of Bergen

Project: ProtonCT

Drawn by: Birger Olsen	Date: 2022-06/06
Checked by	Date
Approved by	Date
PA approved by	Date

Drawing title: Optocouplers

Drawing number: 1

Rev: 0

Sheet title: NAME /1

Sheet: 7/8

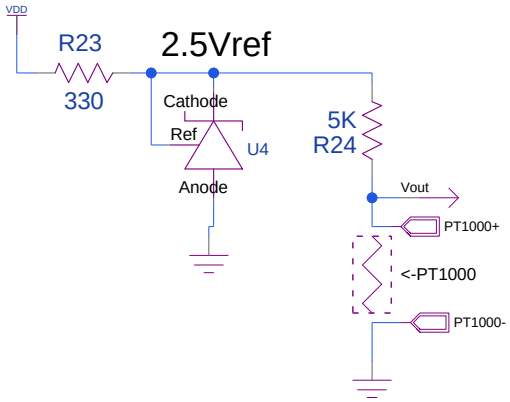

i_min_100deg = 2.5V/(5k+1385) = 0.39 mA
i_max_0deg = 2.5V/6k = 0.4166 mA

Voltage diff(PT1000) from 0 to 100 = 124mV
ADC sensitivity = 0.61mV

A 0.49 degree change will be noticeable by MCU

University of Bergen

Project: ProtonCT

Drawn by: Birger Olsen	Date: 2022-06-06
Checked by	Date
Approved by	Date
PA approved by	Date

Drawing title: Temperature Sensing

Drawing number: 1

Rev: 0

Sheet title: Sensing / 1

Sheet: 8/8

E LAYOUT

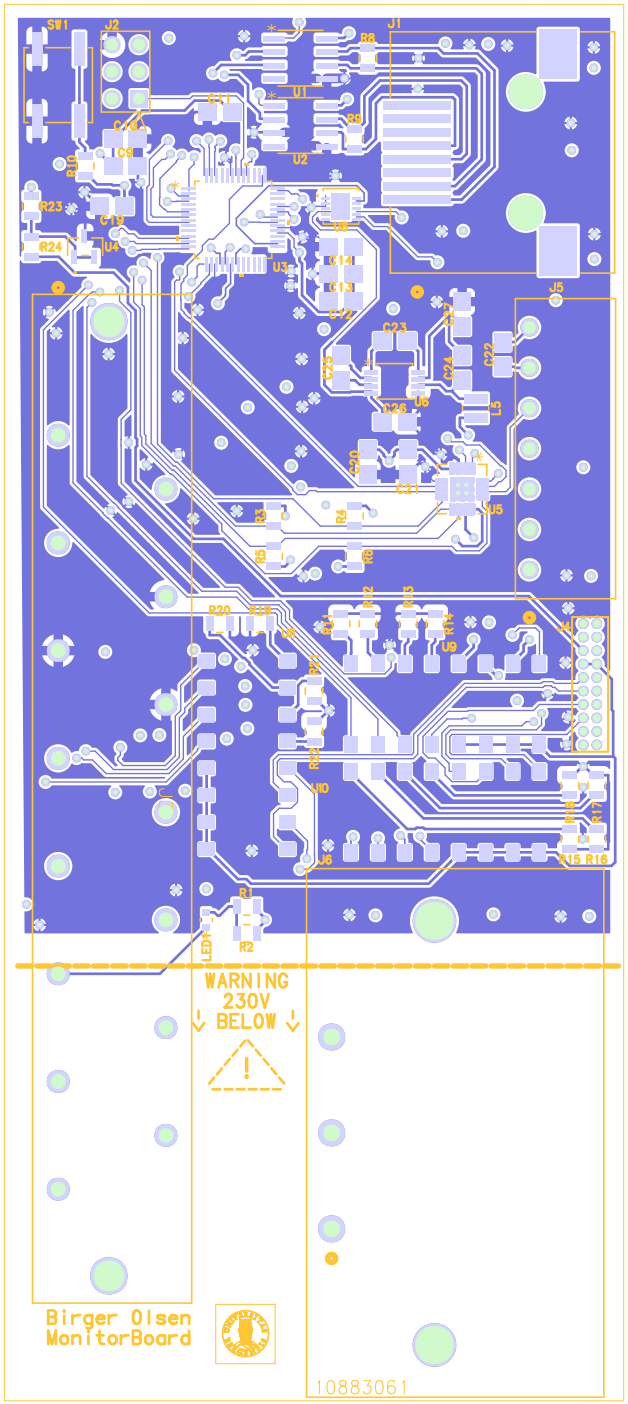

Figure E.3: Top layer, signals and components with ground filled where possible.

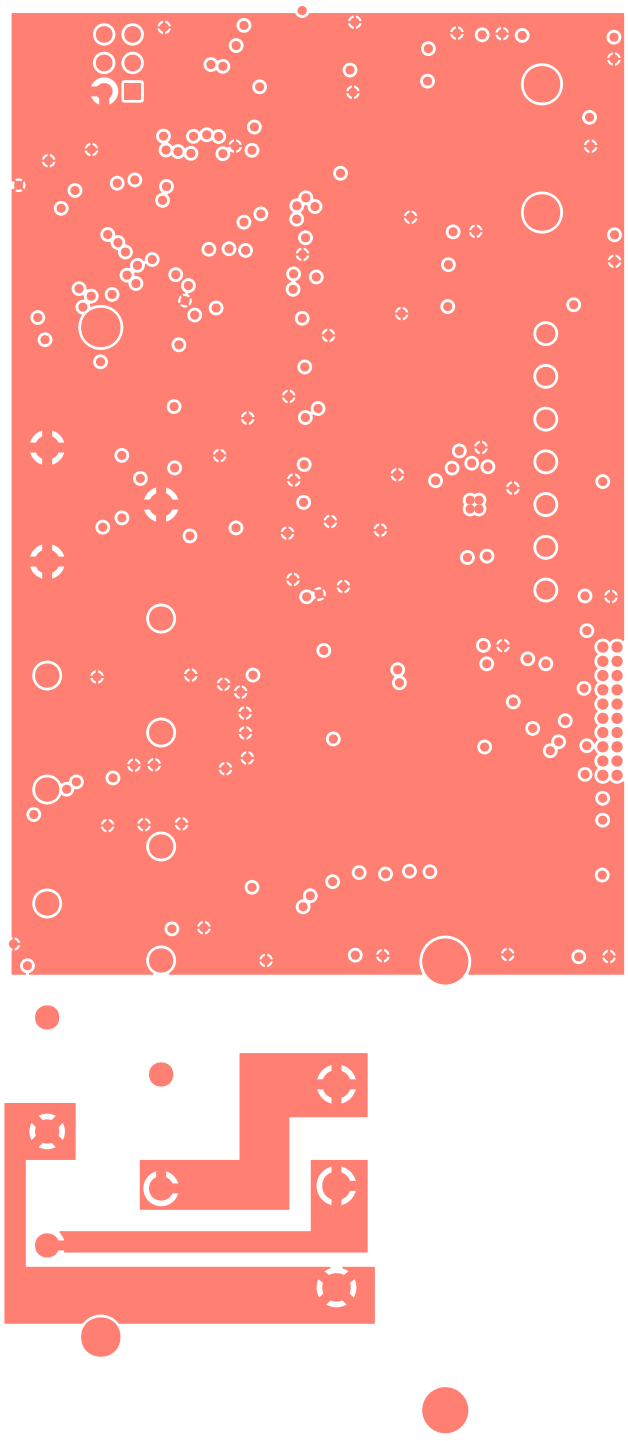

Figure E.4: Layer 2, continous power layer.

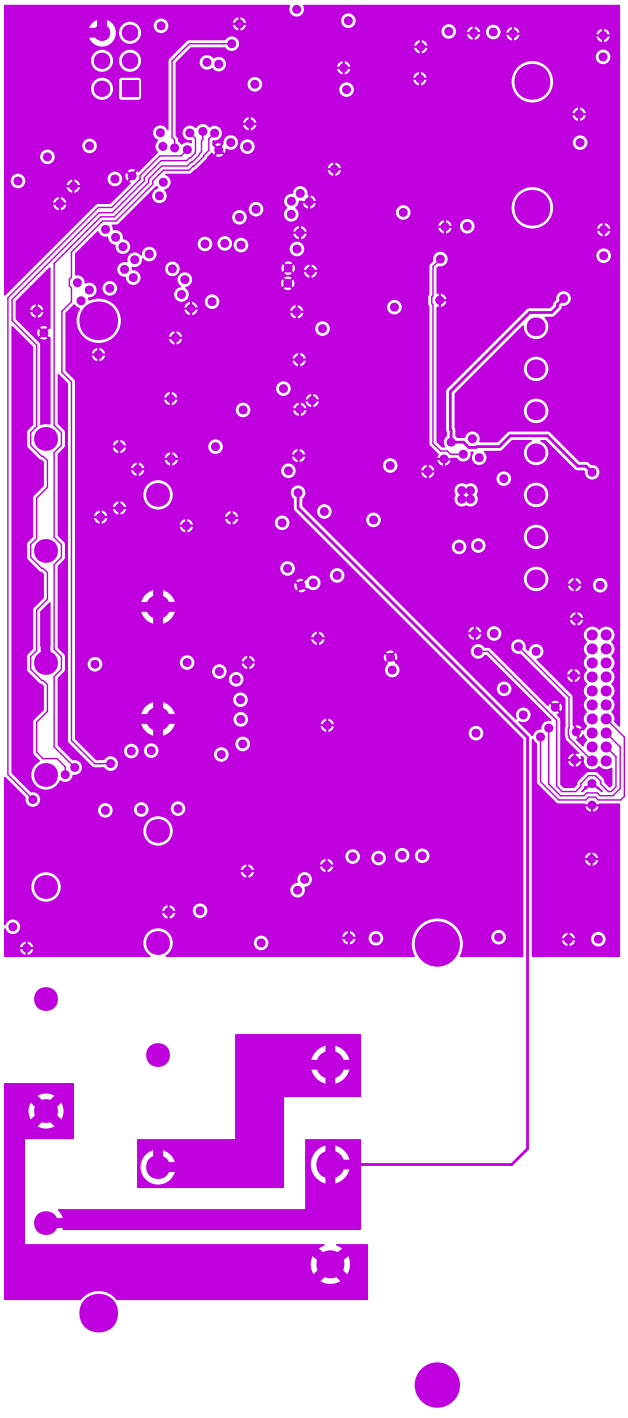

Figure E.5: Layer 3, ground with signals.

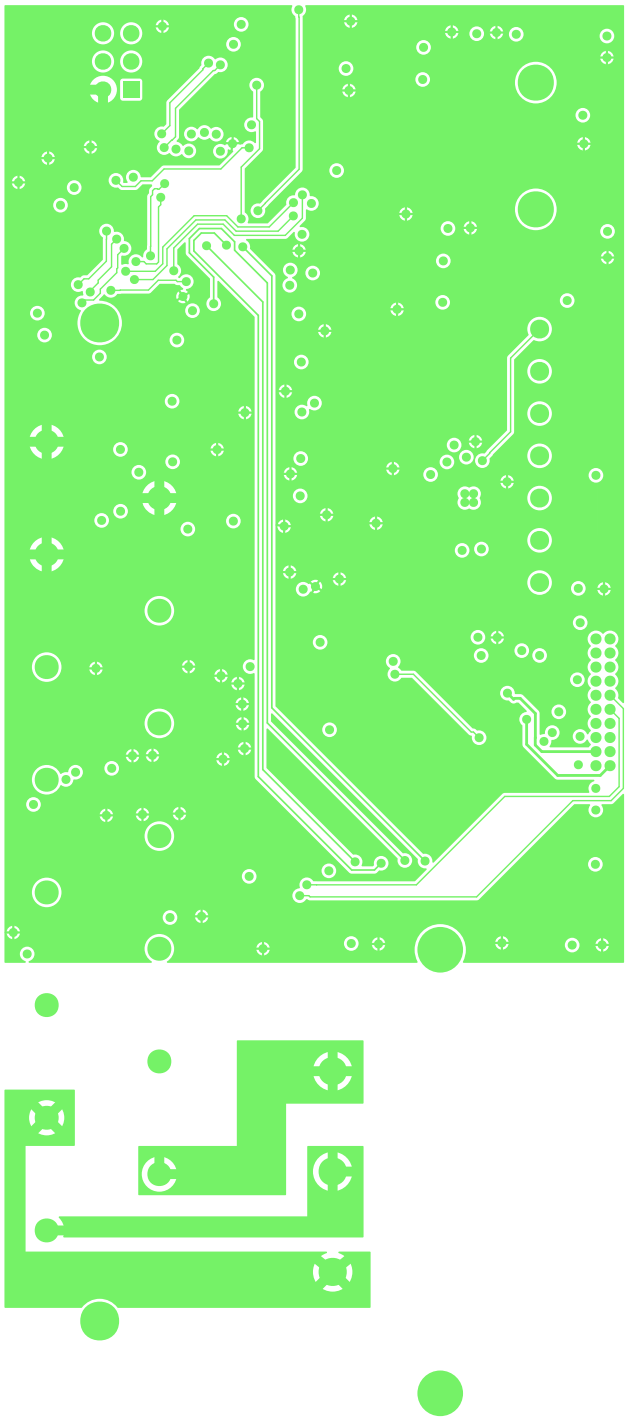

Figure E.6: Layer 4, power with signals.

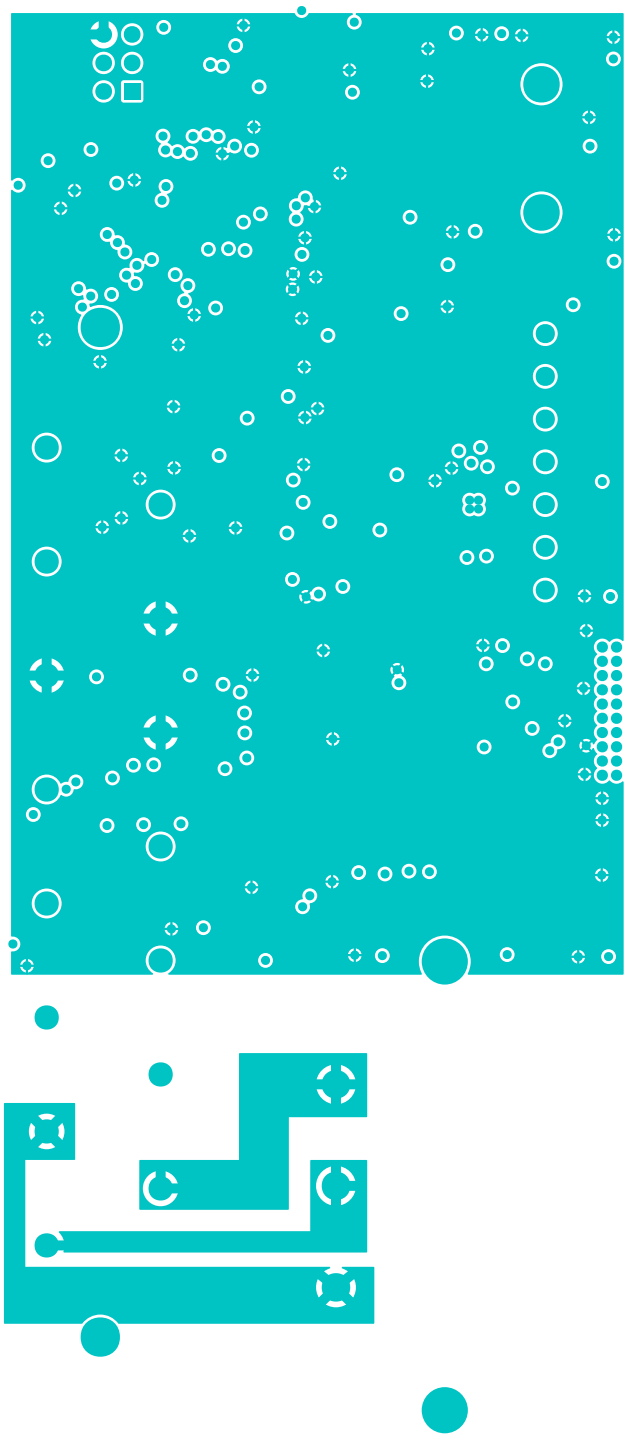

Figure E.7: Layer 5, continous ground layer.

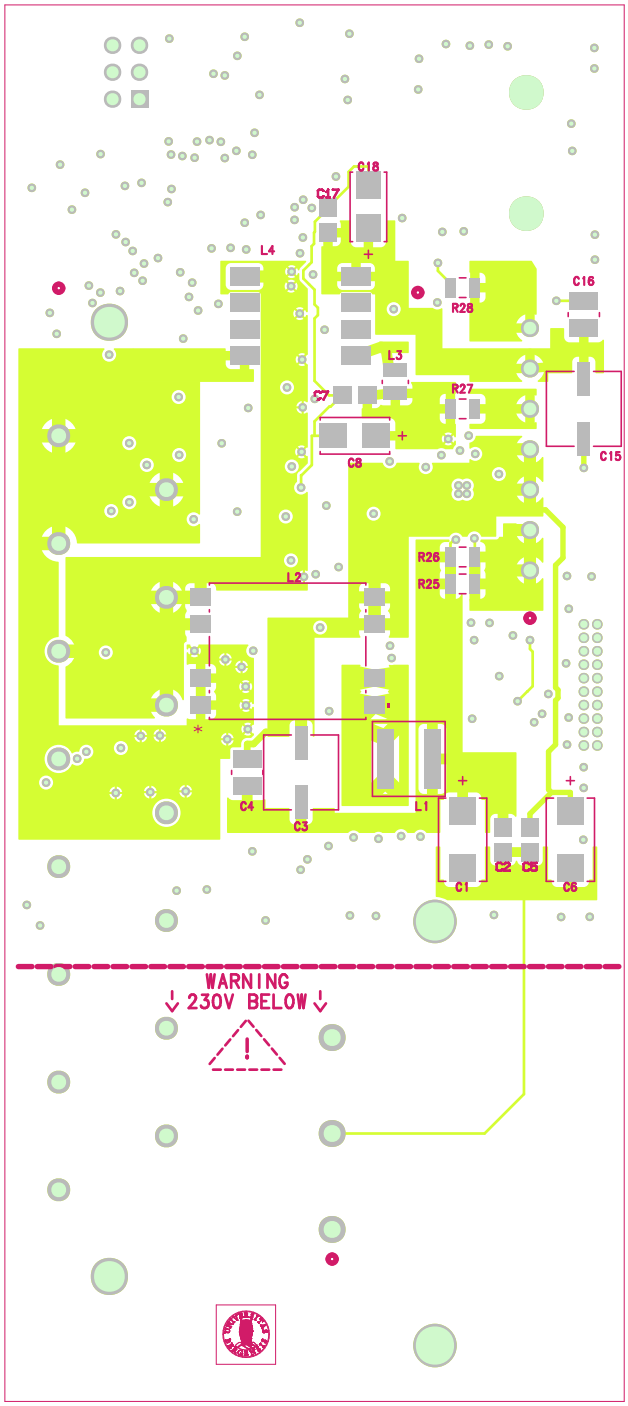

Figure E.8: Bottom layer, reserved for filter.

F BILL OF MATERIALS

Table F.1: Bill of materials monitorboard

Part	Geometry	Quantity	Value	Ref Des	Comment
T495X227K016ATE100	CAPMP7_6X4_3_4_3N_KEM	2	220uF	C1 C6	—
0805Default	C0805	3	100nF	C10 C13 C21	—
0805Default	C0805	2	10nF	C11 C14	—
0805Default	C0805	5	10uF, 16V	C23 C24 C25 C26 C27	—
0805Default	C0805	4	1nF	C2 C5 C7 C17	—
0805Default	C0805	4	1uF	C9 C12 C19 C20	—
0805Default	C0805	1	4.7uF	C22	—
EXV227M016A9HAA	CAP_EXV_9H_KEM	2	220uF	C3 C15	—
GRM32EC81C476KE15L	CAP_GRM32EC81C476KE15L_MUR	2	47uF	C4 C16	—
T491C476K016AT	CAPMP6_3X3_2_2_8N_KEM	2	47uF	C8 C18	—
RES_0805	RES_CRCW_0805	3	330	R1 R2 R23	—
RES_0805	RES_CRCW_0805	2	100	R8 R9	—
RES_0805	RES_CRCW_0805	4	10k	R3 R4 R5 R6	—
RES_0805	RES_CRCW_0805	12	200	R11 R12 R13 R14 R15 R16 R17 R18 R19 R20 R21 R22	—
RES_0805	RES_CRCW_0805	1	4k7	R10	—
RES_0805	RES_CRCW_0805	1	5K	R24	—
LRMAM1206-R002JT5	RES_1206_TTE	4	0.025	R25 R26 R27 R28	—
855437001	CONN_85543-7001_MOL	1	-	J1	—
10883061	CON_10883061_MOL	1	-	J2	—
DIN41612_Straight	CONN_9062152821_HAR	1		J3	—
BERG10x2x1.27	CONN10x2	1		J4	—
1942120000	CONN7_1942120000_WED	1	-	J5	—
GSP1_8101_1	CONN3_P1_8101_1_SAI	1	-	J6	—
150060GS75000	LED_150060GS75000_WRE	1	-	LED1	—
SRP6060FA-4R7M	IND_SRP6060FA_BRN	1	5.6uH	L1	—
744290152	SOIC_90152_WRE	1	34uH	L2	—
74438324100	IND_324100_WRE	1	5.6uH	L3	—
Wurth_744206	Wurth744206_CommonMode	1	60uH	L4	—
74438323022	IND_74438323022_WRE	1	2.2uH	L5	—
EVQ-Q2B01W	EVQQ2_1_PAN	1		SW1	—
MAX3441EASA+	21-0041B_8	2	-	U1 U2	—
AVR128DB48T-E/PT	TQFP48_PT_MCH	1	-	U3	—
TL431_2	SOT_DBZRQ1_TEX	1		U4	—
INA3221AIRGVT	RGV16_2P16X2P16	1	-	U5	—
MAX865EUA+	21-0036K_MXM	1	-	U6	—
LT1991IDD-PBF	DFN-10_DD_LIT	1	-	U8	—
ACPL-847-300	ACPL-847-300	3		U9 U10 U11	—