Johannes Gutenberg-Universität Mainz

Master Thesis

Powering of the Mu3e Detector

Author:
Lucas Sebastian Binn

First Evaluator:
Prof. Dr. Niklaus Berger
Second Evaluator:
Jun.-Prof. Dr. Florian Hug

20. December 2021

Declaration

Hiermit versichere ich, dass ich die vorliegende Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe, dass alle Stellen der Arbeit, die wörtlich oder sinngemäß aus andern Quellen übernommen wurden, als solche kenntlich gemacht sind und dass die Arbeit in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegt wurde.

Mainz, 20.12.2021
Ort, Datum

Lucas Sebastian Binn

Abstract

The Mu3e experiment, located at the Paul Scherrer Institute in Villigen, Switzerland, aims to search for the lepton flavour violating decay $\mu^+ \to e^+e^+e^-$ with a branching ratio sensitivity of 10^{-16}. The detector consists of a lightweight tracker built from high-voltage monolithic active pixel sensors complemented by scintillating tiles and fibres for precise timing measurements. Powering of the detector modules poses unique challenges, requiring voltages in the range of 1.8 V to 3.6 V and currents up to 28 A. The low voltages will be obtained by stepping-down 20 V, provided by commercial power supplies, using custom DC-DC converters located in the magnet, thus reducing cable losses significantly. Converters will be operated in eight crates, with each crate holding 16 boards and equipped with a crate controller and water cooling.

In this thesis, the design of the crate has been developed, and to test the feasibility, a prototype crate, providing space for four boards, has been built. Assessment of the heatsink and water cooling design shows promising results, with an efficiency of (81 ± 6) % and possibly higher numbers reachable using an improved heatsink mounting mechanism. Power consumption of the detector has been estimated to be 11.7 kW, taking power losses in cables fully into account, and measures to reduce this number are proposed. The current iteration of the converter has been thoroughly tested and the output filter optimized, resulting in low-frequency ripple of 1.5 mV peak-to-peak. A newly developed backplane bus design implements geographical addressing, simplifying communication and debugging. With high currents and thus large voltage drops across cables, a voltage regulating mechanism is required. A combination of cable voltage drop compensation and digital voltage regulation has been found practical and will be implemented in the next iteration of the converter.

Zusammenfassung

Das Mu3e Experiment, welches sich am Paul Scherrer Institute in Villigen in der Schweiz befindet, sucht nach dem Leptonfamilienzahl-verletzenden Zerfall $\mu^+ \to e^+e^+e^-$ mit einer Verzweigungsverhältnis-Sensitivität von 10^{-16}. Der Detektor besteht aus einem leichten Spurdetektor, der aus dünnen monolithischen aktiven Pixelsensoren mit Hochspannung aufgebaut ist. Zusätzlich werden szintillierende Kacheln und Fasern für eine verbesserte Zeitauflösung benutzt. Die Spannungsversorgung der Detektormodule stellt eine besondere Herausfordung dar, da Spannungen im Bereich von 1.8 V bis 3.6 V und Ströme bis zu 28 A benötigt werden. Die niedrigen Spannungen werden durch Umwandeln von 20 V, bereitgestellt durch kommerzielle Netzteile, mit Hilfe von DC-DC-Wandlern im Magneten erzielt. Hierdurch werden Verluste in den Kabeln erheblich reduziert. Die DC-DC-Wandler werden in acht Crates mit jeweils 16 Wandlern betrieben, wobei jedes Crate mit einer Steuerplatine und Wasserkühlung ausgestattet ist.

In dieser Abschlussarbeit wurde das Design des Crates entwickelt und eine Machbarkeitsstudie unternommen. Für diese wurde ein Prototyp-Crate gebaut, welches Platz für vier Wandler bietet. Die Analyse der Wasserkühlung zeigt vielversprechende Ergebnisse. Ein Wirkungsgrad von $(81 \pm 6)\,\%$ wurde ermittelt, wobei mit einem verbesserten Befestigungsmechanismus für den Kühlkörper vermutlich noch höhere Werte erreicht werden könnten. Die Leistungsaufnahme des Detektors wurde auf 11.7 kW geschätzt, wobei Leistungsverluste in den Kabeln mit einbezogen wurden. Verschiedene Maßnahmen zur Verlustreduzierung wurden aufgezeigt. Die aktuelle Version des Wandlers wurde ergiebig getestet und der Ausgangsfilter optimiert, wodurch niederfrequente Oszillationen mit einem Spitze-Tal-Wert von 1.5 mV starkt unterdrückt werden. Ein neue entwickelts Backplane-Bus-Design implementiert geografische Adressierung, wodurch Kommunikation und Fehlerbehebung vereinfacht werden. Da bei hohen Strömen große Spannungen über die Kabel abfallen, ist ein Mechanismus zur Spannungsregelung erforderlich. Eine Kombination aus "cable voltage drop compensation" und digitalem Nachregulieren hat sich als praktisch erwiesen und wird in der nächsten Iteration des Wandlers implementiert werden.

Contents

Chapter 1

Physics Beyond the Standard Model

Over the history of humanity, humans have pondered over the nature of matter. The word atom stems from the ancient Greek word *atomos*, meaning "uncuttable". At the time the reasoning for assuming a granularity of matter was of philosophical roots. While the atom proved not to be the smallest element, modern physics has confirmed the suspicion that everything in the universe is made from a few basic building blocks. Knowledge has been accumulated, condensed, and ordered, resulting in one of the most successful theories, the Standard Model.

1.1 The Standard Model

The Standard Model (SM) of particle physics contains 12 elementary particles (and their antiparticles) of spin $1/2$ called fermions (see figure 1.1). The other particles, the bosons, have spin 1 and are referred to as the *force mediating particles*, as they represent the fundamental forces and thus allow interaction between particles. The fermions are divided into quarks and leptons. Whereas quarks have a color charge and interact via the strong force, leptons do not. Leptons can be divided into the charged leptons (electron, muon, tau, and corresponding antiparticles), which interact via the electromagnetic and weak force, and their neutral counterparts, the neutrinos (and antineutrinos), which interact only via the weak force.

The SM does not only describe, but also predicts the outcome of many particle physics experiments. The latest success story being the discovery of the Higgs boson at the Large Hadron Collider (LHC) in 2012, which has been postulated in 1964. While with the discovery of the Higgs particle the SM is thought to be completed, physicists have shifted their gaze at a handful of phenomena that can not be explained by the theory, looking for physics beyond the SM. This includes, for example, dark matter, invisible matter which would explain the measured velocity distribution of stars orbiting in a galaxy. None of the SM particles have the right properties. Other phenomena are, for example, dark energy or the matter-antimatter asymmetry in the universe.

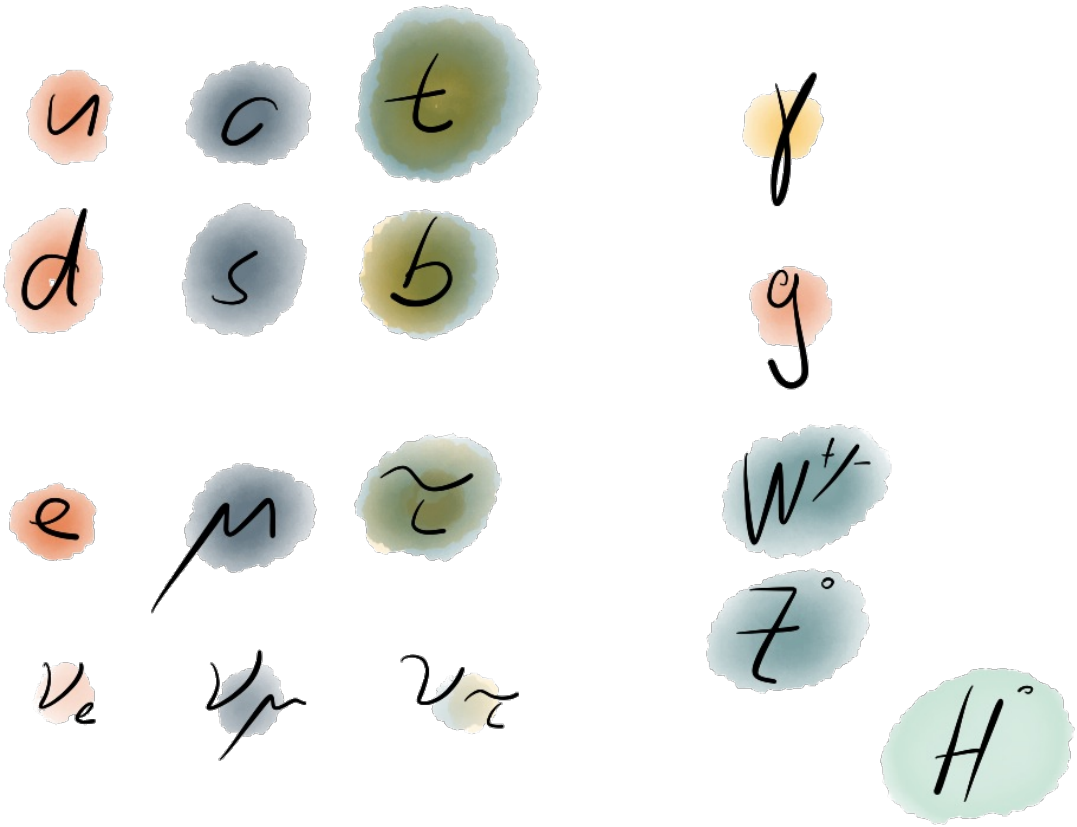

Figure 1.1: The Standard Model of particle physics [1].

1.2 Charged Lepton Flavor Violation

Leptons can be divided, the same way as quarks, into three families. In the minimal SM (neutrinos are massless), a transition between the families is not allowed, i.e. initial and final state of a process must have the same number of leptons per family. The lepton flavour number is separately conserved for any process. By giving mass to the neutrinos, and thus enabling neutrino oscillation, a mechanism for lepton flavour violation (LFV) in the neutrino sector is created. This was first hinted at by Ray Davis's Homestake experiment in the 1960s and has been confirmed by many experiments since then.

LFV for charged leptons, referred to as charged lepton flavour violation (cLFV), has not been observed yet. In the SM with neutrino mixing, cLFV is only possible via higher-order loop diagrams involving neutrino oscillation. These processes are always strongly suppressed (see figure 1.2).

The decay $\mu^+ \to e^+e^+e^-$, for example, has its largest SM contribution coming from the Feynman diagram shown in figure 1.3. With a branching ratio smaller than 10^{-54}, the

decay is essentially predicted to be not observable. With many beyond Standard Model theories significantly increasing the probability of this and similar decays, it is a very clean environment and any observed decay would be a hint towards physics beyond the SM.

The branching ratios of cLFV decays have been investigated by many experiments over the years, so far only setting upper limits (see figure 1.4). The world's leading upper limit for the decay $\mu \to e\gamma$ was set by the MEG experiment in 2016 with a branching ratio of $< 4.2 \cdot 10^{-13}$ and for the decay $\mu \to 3e$ by the SINDRUM experiment in 1988 with a branching ratio of $< 1 \cdot 10^{-12}$.

μ^- DECAY MODES		Fraction (Γ_i/Γ)	Confidence level	p (MeV/c)
$e^- \overline{\nu}_e \nu_\mu$		$\approx 100\%$		53
$e^- \overline{\nu}_e \nu_\mu \gamma$		[d] $(6.0 \pm 0.5) \times 10^{-8}$		53
$e^- \overline{\nu}_e \nu_\mu e^+ e^-$		[e] $(3.4 \pm 0.4) \times 10^{-5}$		53
Lepton Family number (*LF*) violating modes				
$e^- \nu_e \overline{\nu}_\mu$	*LF*	[f] < 1.2 %	90%	53
$e^- \gamma$	*LF*	$< 4.2 \times 10^{-13}$	90%	53
$e^- e^+ e^-$	*LF*	$< 1.0 \times 10^{-12}$	90%	53
$e^- 2\gamma$	*LF*	$< 7.2 \times 10^{-11}$	90%	53

Figure 1.2: Decay modes of the negatively charged muon with measured branching ratios. For LFV decays only upper limits have been measured [2, p. 34].

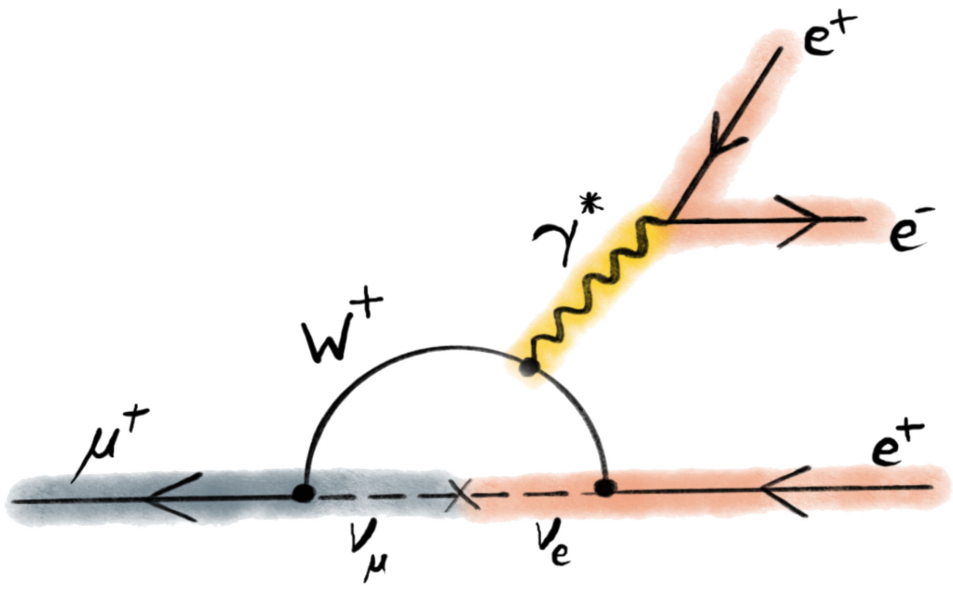

Figure 1.3: The decay $\mu^+ \to e^+e^+e^-$ is allowed in the Standard Model with neutrino mixing, but with a factor of 10^{-54} highly suppressed.

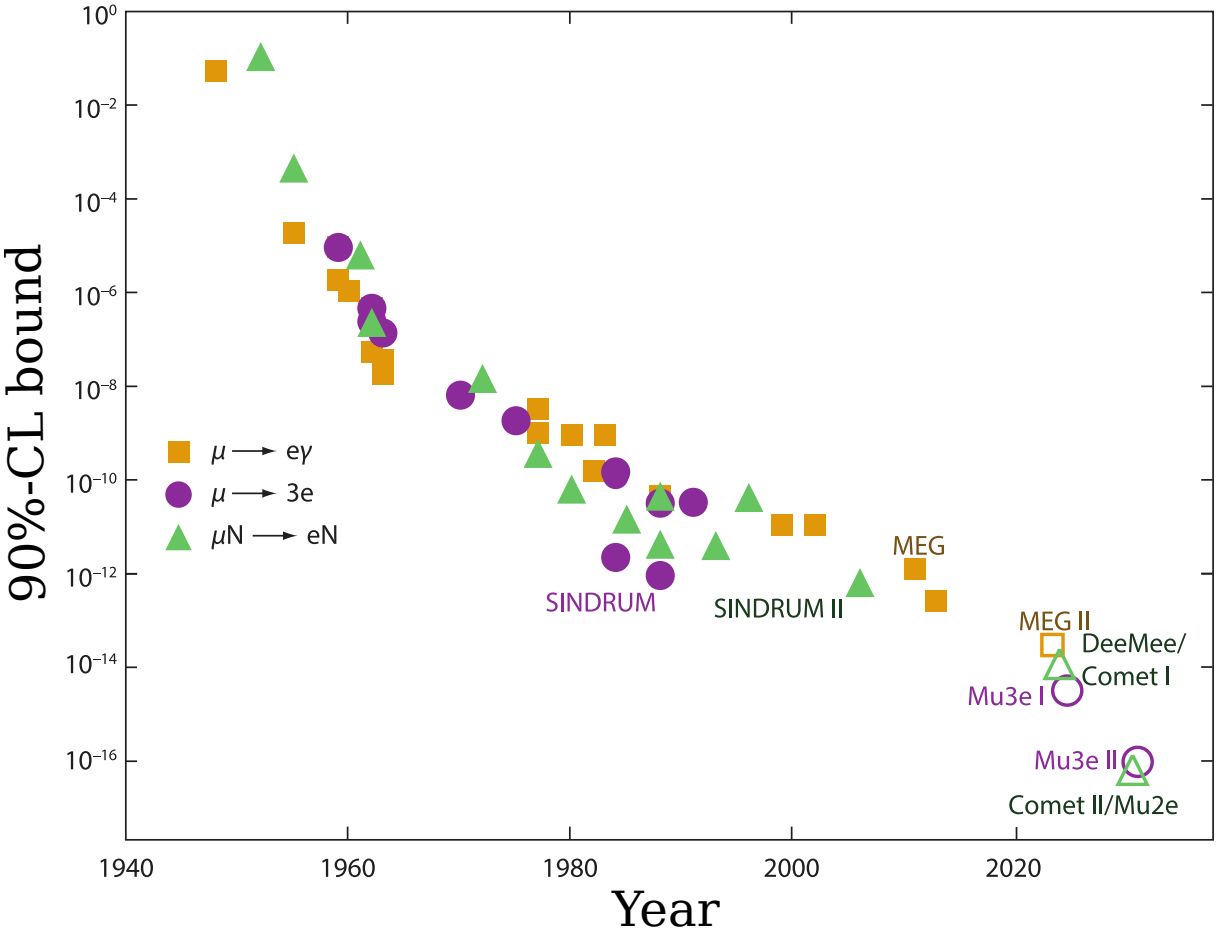

Figure 1.4: Experiments looking for charged lepton flavour violation in muonic decay channels. Solid markers represent past and current experiments, whereas hollow markers represent experiments under development. Adapted from [3].

Chapter 2

The Mu3e Experiment

The Mu3e experiments aims to detect the lepton flavor violating decay

$$\mu^+ \to e^+e^+e^- \tag{2.1}$$

if its branching fraction is larger than 10^{-16} or exlude a branching fraction of $> 10^{-16}$. In the Standard Model this decay is forbidden. Only by adding neutrino mixing it becomes an allowed process, but is still heavily suppressed by a factor of approx. 10^{-54} (see figure 1.3).

To achieve these sensitivities in a reasonable amount of time, a high-intensity muon beam will be used, provided by the Paul Scherer Institute in Villigen, Switzerland. For Phase I of the experiment, the beamline $\pi E5$ will be used, which holds the title for the highest intensity low-energy muon beam in the world, with up to 10^8 muons per second. The goal is to reach a single event sensitivity of $2 \cdot 10^{-15}$, slightly short of the final goal of 10^{-16}. Phase I will be used to establish the novel technologies, set first limits, and prepare for the next phase. Phase II is going to use the *high intensity muon beam* (HiMB), which is currently under study at PSI and is planned to be ready by 2025. With an approx. 20 times higher muon rate, the final sensitivity of 10^{-16} can be reached.

2.1 Signal and Background

The signal is differentiated from the background by looking at the vertices, energies, and momenta of the emitted particles.

The muon is stopped on a target and decays at rest, emitting three particles from the same vertex (see figure 2.1). For that reason, the sum of the three momenta must vanish

$$|\vec{p_{\text{tot}}}| = \left| \sum_i \vec{p_i} \right| = 0\,, \tag{2.2}$$

and the sum of the particles' energies must be equal to the muon mass

$$m_{\text{inv}} = \sum_i E_i = m_\mu\,. \tag{2.3}$$

Furthermore, emitted particles can have energies starting from the electron mass up to half the muon mass (approx. 53 MeV) and must lie in a single plane.

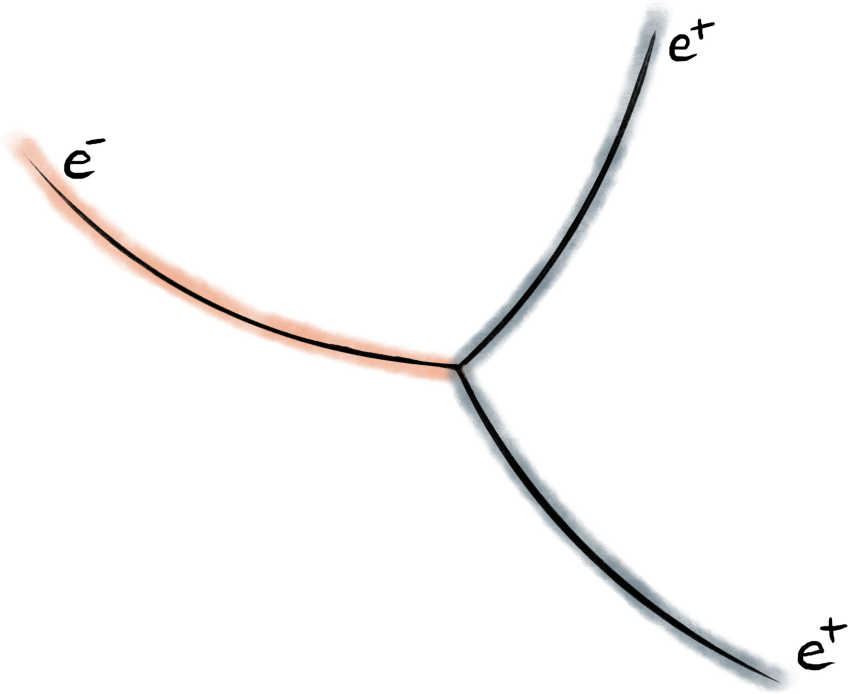

Figure 2.1: Signal Topology

Since the decay is suppressed in the Standard Model (with neutrino mixing) by a factor of around 10^{-54}, it is essentially not observable. This means there is no irreducible physics background, i.e. background which can not be differentiated from the signal process. Reducible background can be categorized in combinatorial background and internal conversion background.

Combinatorial Background

This is background that stems from the combination of processes producing three tracks which by coincidence resemble the signal decay (see section 2.1). It can for example be the Michel decay of a muon ($\mu^+ \to e^+\nu_e\bar{\nu_\mu}$) in combination with a process providing negatively charged particles, e.g. Bhabha scattering, photon conversion or Compton scattering. To suppress the accidental combination of tracks and vertices from different processes, a high vertex, momentum, and time resolution is needed.

Internal Conversion Background

The decay $\mu^+ \to e^+e^+e^-\bar{\nu_\mu}\nu_e$ is allowed in the Standard Model and has a branching fraction of $3.4 \cdot 10^{-5}$ (see figure 2.2). The only way to differentiate it from the signal decay is by noticing the missing energy carried away by the neutrinos (see figure 2.3b). For that reason, it can only be suppressed by momentum resolution.

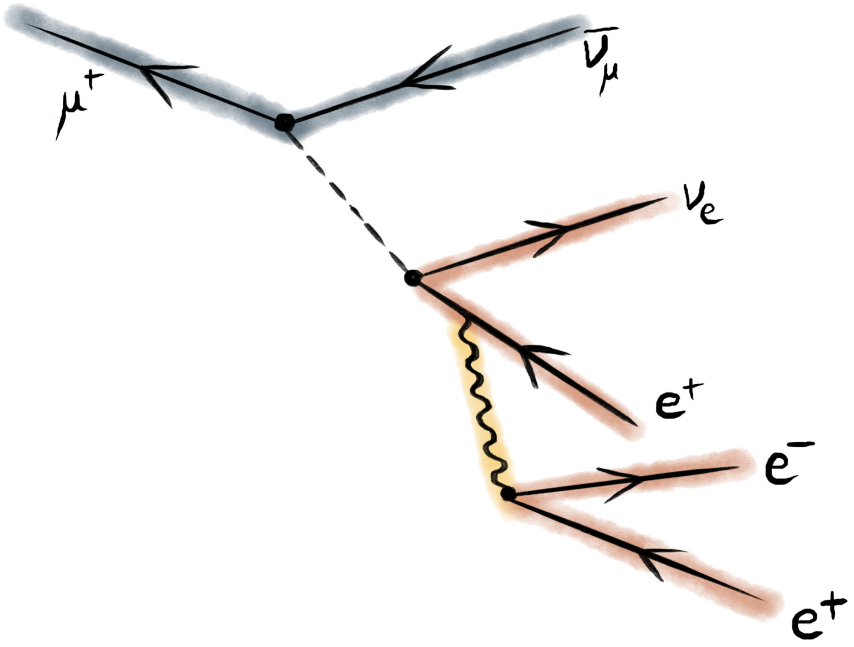

Figure 2.2: Allowed radiative decay with internal conversion.

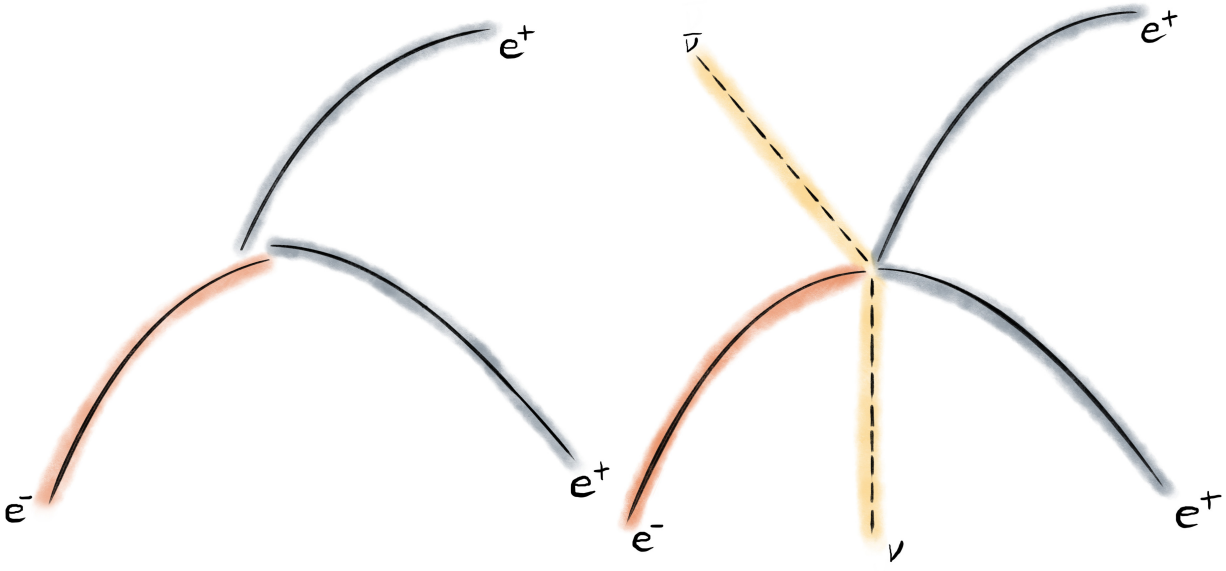

(a) Combinatorial Background

(b) Internal Conversion Background

Figure 2.3: Background mimicking the $\mu \to eee$ decay.

2.2 Challenges

The Mu3e detector focuses on track reconstruction of particles. Using a 1 T solenoidal magnetic field, tracks of charged particles are bent because of the Lorentz force. As a

result charged particles follow helical trajectories. From the direction of curvature, the charge can be determined, and measuring the radius of the track, the momentum of the particle can be calculated (assuming charges of $\pm e$).

As previously described (see section 2.1), vertex, momentum and time resolution are of highest importance. With pixel sizes of $80 \times 80\,\mu m^2$ and low momentum particles, the dominating factor of uncertainty for vertex and momentum resolution is multiple Coulomb scattering in the detector material (see figure 2.4). When traversing the detector material, a particle scatters multiple times, leading to an offset and scattering angle with respect to its incoming trajectory. This leads to uncertainties in radius and thus momentum reconstruction. To minimize this effect, the material budget in the detector must be kept as small as possible. For the tracking detector, special super-thin silicon pixel sensors are used. They have a thickness of 50 µm. With a time resolution of a few ns, they are not able to sufficiently suppress combinatorial background. For that reason, an additional timing detector with resolutions in the 10s to 100s of ps range is used.

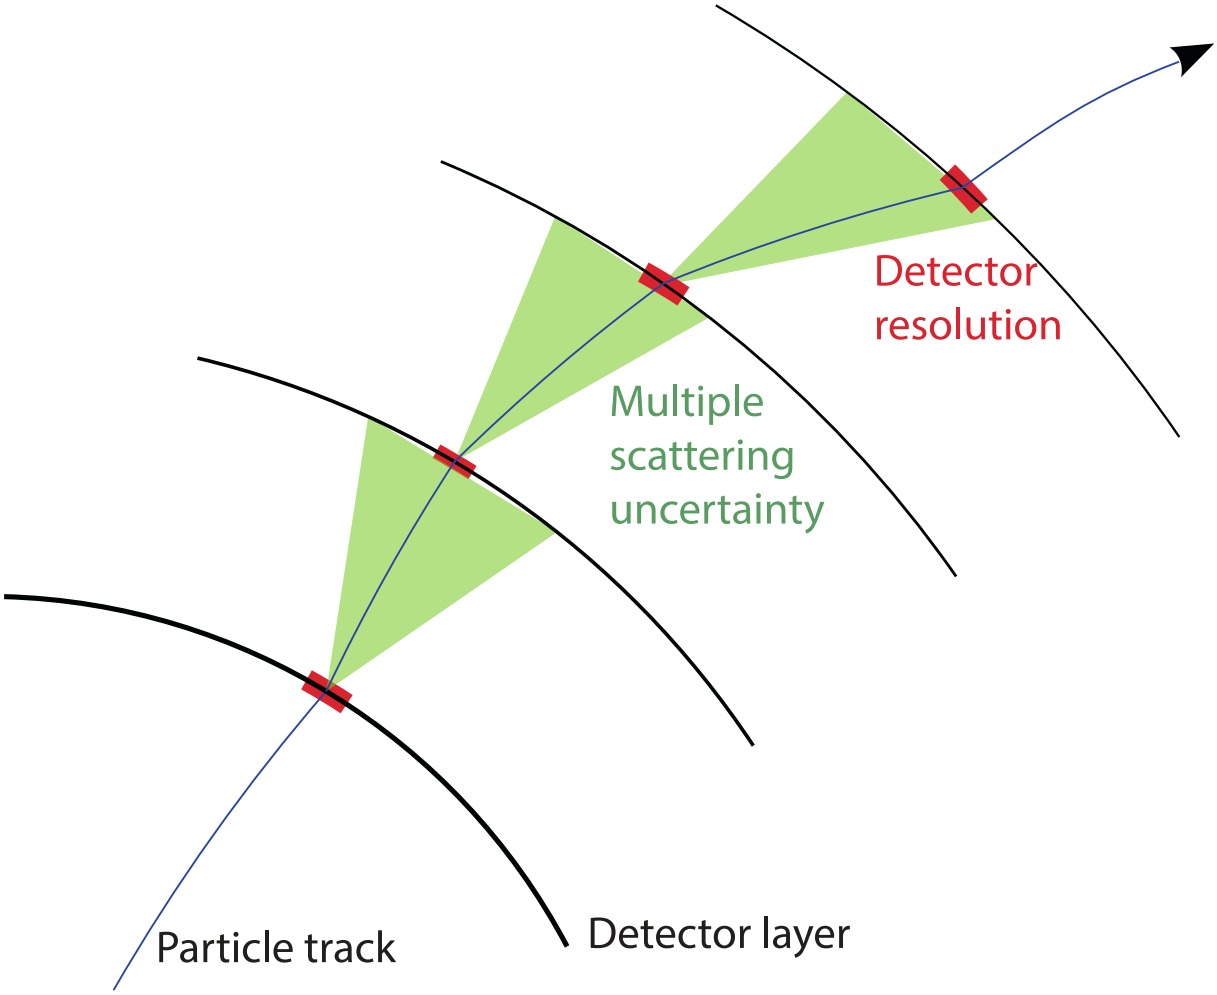

Figure 2.4: Scattering of a low-momenta particle in the detector material. Detector resolution is adequate while uncertainty in scattering angle dominates.

2.3 Detector

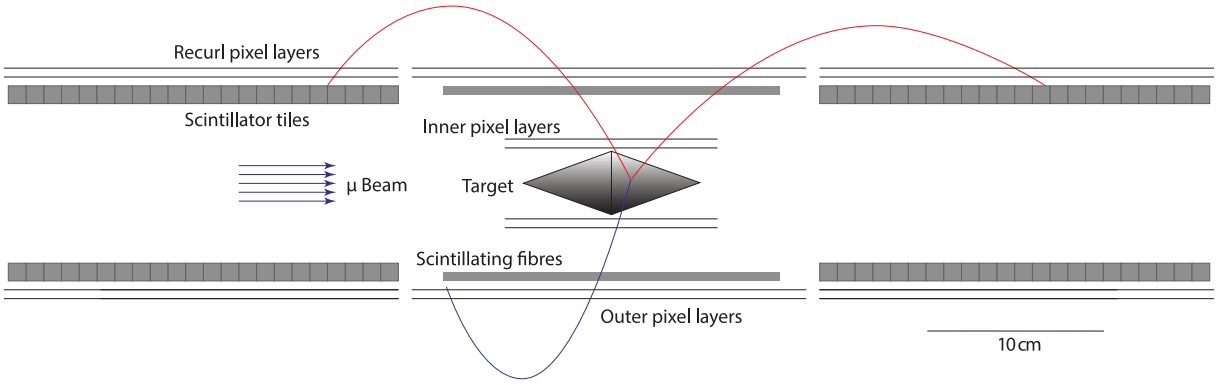

Figure 2.5: Longitudinal cross section of the detector. Positrons tracks in red, electron track in blue [1].

For phase I of the experiment, the detector will consist of three stations, one central and two recurl stations (see figure 2.5).

Muons enter the experiment through the beam pipe upstream and are stopped on a hollow double-cone target located at the center of the central station. Muons then decay at rest and the outgoing charged particles (electrons and positrons) follow helical trajectories in the magnetic field, traversing different layers of detectors on the way. Because of the low kinetic energies and strong magnetic field, a large number of tracks will curl back and be detected again in one of the recurl stations. Reconstruction of the tracks allows the determination of the charge and momentum of particles.

Precise track reconstruction is achieved with a combination of three detector systems, the pixel tracking detector, providing accurate positional information, and two timing detectors, for superior time resolution.

2.3.1 Pixel Tracking Detector

The ultra-thin silicon pixel tracker's smallest building blocks are High Voltage Monolithic Active Pixel Sensors (HV-MAPS), developed by Ivan Perić [4]. Monolithic sensors combine sensor and readout in the same device, thereby reducing the detector material. Sensors are manufactured in a commercial complementary metal-oxide-semiconductor (CMOS) process. Being able to apply voltages of up to 120 V allows for fast charge collection and thin depletion zones. This means the sensor can be thinned down to 50 µm. For the Mu3e experiment, a special HV-MAPS has been developed, the MuPix sensor (see figure 2.6). It has an active sensor matrix of around $20.48 \times 20.00\,\mathrm{mm}^2$, pixels of size $80 \times 80\,\mu\mathrm{m}^2$ and a total chip size of $20.66 \times 23.18\,\mathrm{mm}^2$. The non-sensitive area is used to host peripheral electronics and bonding pads.

The MuPix sensors are glued onto stripes of polyimide foil, on which signal and power

lines are implemented as thin aluminium traces to which the sensor is bonded. Gluing two V-folds along the foil adds stability and creates flow channels for helium to allow cooling of the sensors. This construct is called a ladder and represents the smallest mechanical unit in the pixel tracker. To reduce power losses over the traces, each ladder is powered from both sides. This is referred to as powering half-ladders. Ladders are then combined to mechanically robust modules and modules form detector layers (see table 2.1).

The central station will have four layers, arranged in two double layers, the inner and outer pixel layers. Each pair provides information about track trajectories The recurl stations implement only the outer two pixel layers. Tracks with hits in the recurl stations allow reconstruction with higher purity and momentum resolution.

The MuPix sensors have a time resolution of a few ns, which by itself is not sufficient to suppress combinatorial background at the level needed.

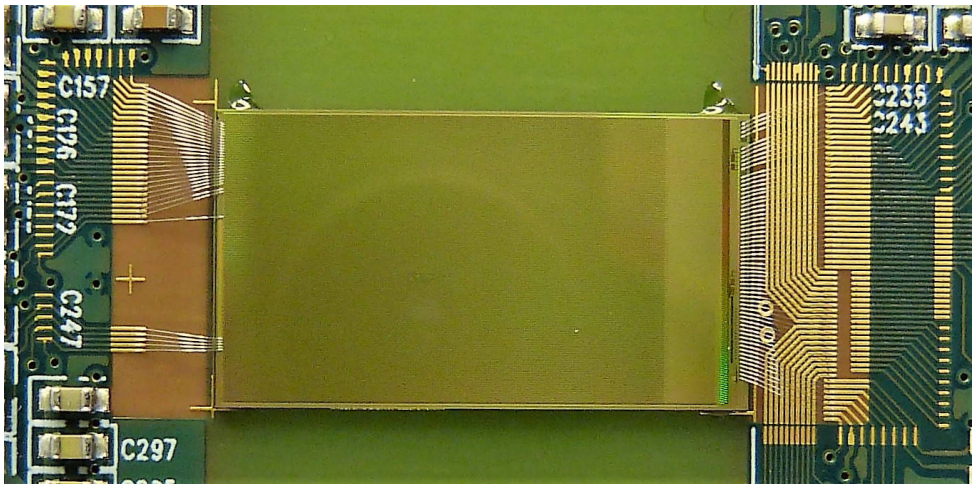

Figure 2.6: MuPix sensor v8 (middle green part) bonded to a PCB. The sensitive area can be seen in dark green, while the peripheral electronic is in the light green zone on the right. Size differes from the final specifications [1].

Table 2.1: The pixel tracker is made out of MuPix sensors, ladders, modules and layers [5, p. 24].

Layer	#Modules	#Ladders per Module	#MuPix per Ladder
1	2	4	6
2	2	5	6
3 Central	6	4	17
3 Recurl	2 x 6	4	18
4	3 x 7	4	18

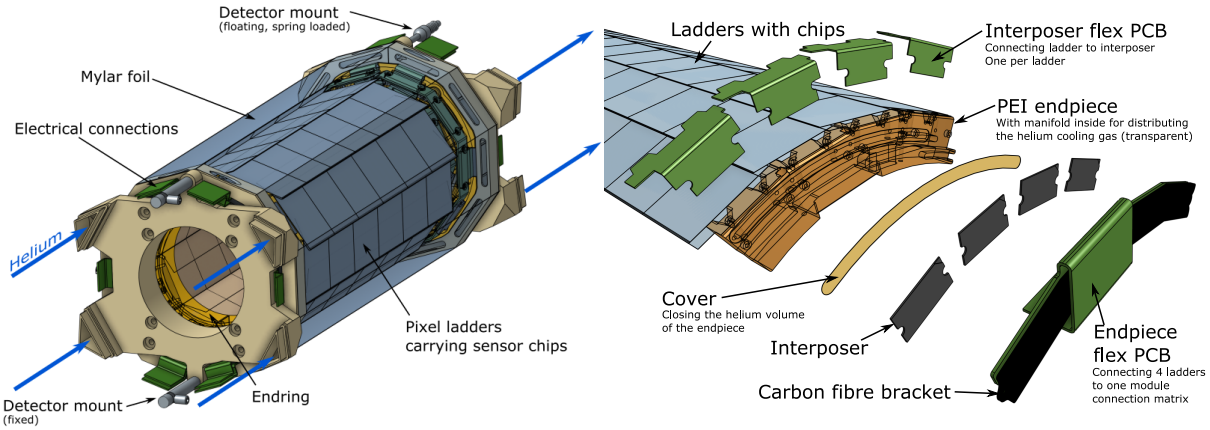

(a) Schematic view of the central layer 1 and 2 pixel detector. Endrings provide mechanical support.

(b) Exploded view of an outer layer module. Interposer connectors and flex PCBs are used to make electrical contact.

Figure 2.7: Structure of the pixel detector [5, pp. 27, 30].

2.3.2 Timing Detector

The time resolution of the detector is increased by an additional timing detector on the inside of the outer pixel layers. This helps to match hits in the pixel detector and allows for superior track reconstruction. Even at high muon stopping rates, combinatorial background can be efficiently suppressed.

While there are different requirements for the central and recurl stations, and thus the scintillating geometry is changed, the readout chain is the same. Light from the scintillators is converted to an electrical signal using silicon photomultipliers (SiPMs) and the signal is read out using a custom chip called MuTRiG. The Muon Timing Resolver including Gigabit-link is a 32 channel, mixed-signal silicon photo-multiplier readout ASIC (Application-Specific Integrated Circuit).

Scintillating Fibres

For the central station, a low material budget is crucial to minimise the effect of multiple Coulomb scattering. For that reason, three layers of scintillating fibres are used (see figure 2.8a). Readout electronic is located at the ends of the fibres. This design allows for a time resolution better than 400 ps (including ASICs).

Scintillating Tiles

At the recurl stations, the requirements are different. Particles passing the recurl stations do not need to be detected afterward. For that reason, scintillating tiles with a volume of 0.5 cm^3 and a large material budget can be used. This results in a larger energy deposition

and higher light yield. Furthermore, each tile can be read out individually (see figure 2.8b). This adds up to an increased timing resolution of better than 80 ps.

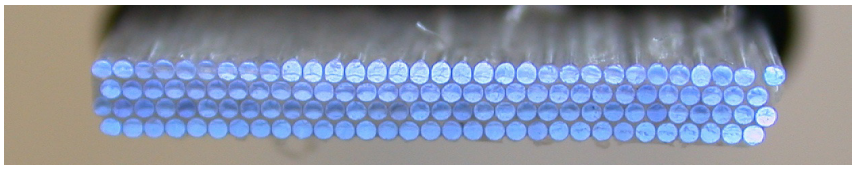

(a) Stacked layers of scintillating fibres. Fibres have a diameter of 250 µm.

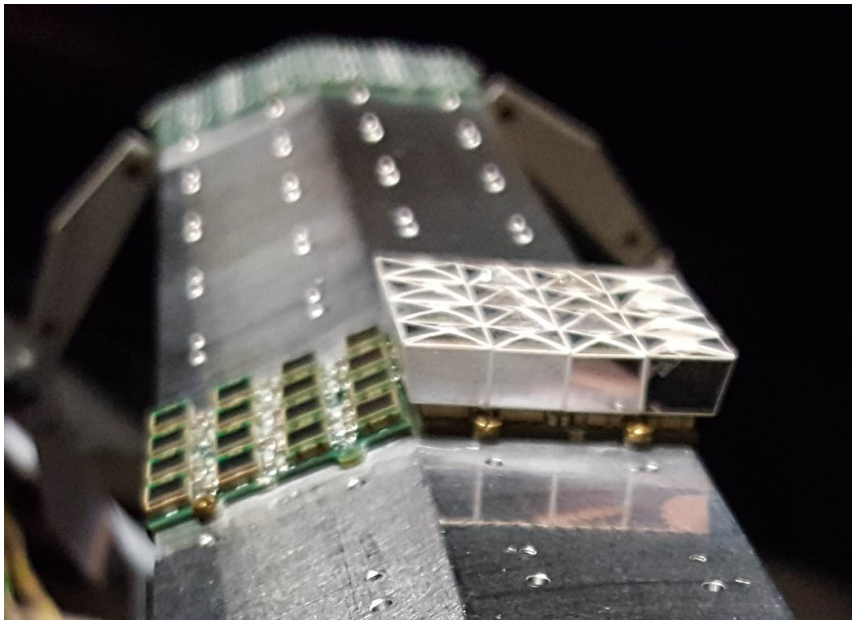

(b) Assembly of a test setup for the scintillating tiles detector [5, p. 66].

Figure 2.8: Subdetector modules of the timing detector.

2.4 Powering Scheme

While the main detector is quite small, with a power consumption of up to 12 kW the power density is considerable and thus a robust powering scheme has been developed (see figure 2.9). Power supplies are located outside the magnet in the power rack and cables are brought into the magnet using a high-density power connector.

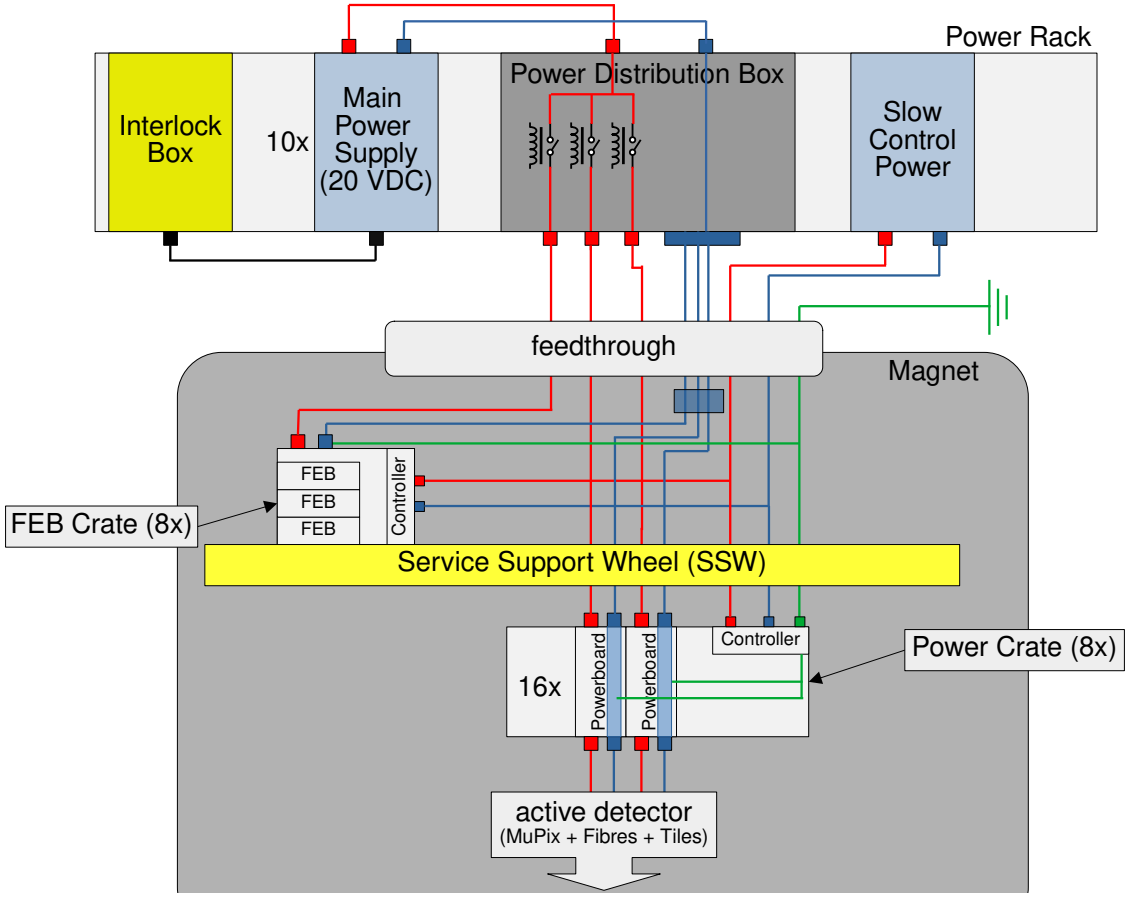

Figure 2.9: Powering scheme. Only one side of the detector is shown, the other side is identical.

2.4.1 Slow Control Power

Slow control power is used to power devices producing "slow data", such as environmental sensors, crate controllers, etc., whereas the high-power detector modules are powered separately. It stays on in case of an interlock trigger or shutdown. This way critical monitoring and control of the experiment is possible at any time, for example, before turning on the high-power detector modules. Since this power is not high-current, cable voltage drops are small and the final voltage can be provided directly by the power supplies.

2.4.2 Main Power

All detector modules need low voltages (2 V to 4 V) and high currents. Providing these directly by the power supplies outside the experiment would result in massive, unflexible cables and losses along the way. For that reason, a higher voltage of 20 V is provided by the power supplies and DC-DC converters located in the magnet close to the detector modules step the voltage down to the final levels.

For the front-end boards (FEB), three converters are embedded on the board itself. The converters' design is based on the TPS548A20 chip with integrated switch from Texas Instruments.

For the active detector (MuPix & Fibre & Tile), a custom DC-DC converter board has been designed, referred to as the powerboard (see section 3.4). Powerboards are located in 8 power-crates, four upstream and four downstream of the target, mounted on the inside of the service support wheels (SSWs) (see figure 2.10). The SSW is a support structure for auxiliary hardware and is located outside of the active detector volume. By mounting the power-crates on the inside, the distance to the detector modules and thus cable losses are minimized. The low-voltage high-current connection from the powerboards to the detector modules is described in section 2.4.4. Each power-crate can hold up to 16 powerboards. Chapter 4 is concerned with the design of the power-crate and a first prototype is made in section 4.3.

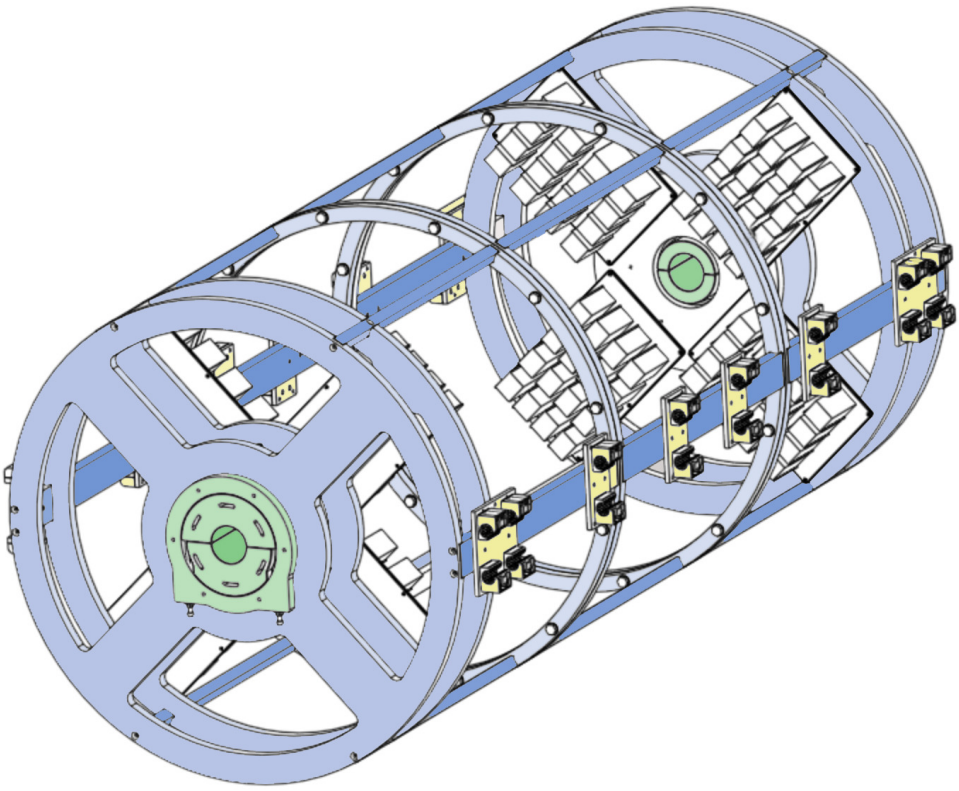

Figure 2.10: Detector cage and service support wheels with 8 power-crates mounted on the inside. Detector, target and beam pipe has been left out [5, p. 72].

2.4.3 Power Partitions

The detector is divided into 120 power partitions (see table 2.2). This is a way of subdividing the experiment into logical segments. With respect to power, this is the smallest

unit. Each partition can be switched on and off individually. The 8 FEB crates, for example, are each a separate partition, rated up to 20 A at 20 V. The other partitions are each powered by a single powerboard, except for the tile detector, where two powerboards share a partition (generating the two required voltages, see table 2.4). These partitions are rated up to 6 A at 20 V. As previously described, MuPix ladders are powered from both sides, referred to as half-ladders (see section 2.3.1). A MuPix partition always powers four half-ladders.

The ability to switch each power partition independently is provided by the power distribution box. Here the outputs of the 10 power supplies are branched to the number of partitions and high side power MOSFETs allow switching. The input current for each partition can be measured.

Table 2.2: Assigned paritions and powerboards for components.

Component	#Partitions	#Powerboards/partition
MuPix		
Layer 1	4	1
Layer 2	4	1
Layer 3 Central	12	1
Layer 3 Recurl	2 x 12	1
Layer 4	3 x 14	1
Fibre (2x SMB)	12	1
Tile (1x TMB)	14	2
FEB	8	0
Total	**120**	**126**

2.4.4 Main Power Distribution

As described in section 2.4.2, the powerboards are located in 8 power-crates mounted on the inside of the service support wheels. Delivery of the power from the output of the powerboards to the detector modules is done using a chain of different conductors. The composition and length of the chain depend on the location of the powered detector module.

For the tile detector, which is located at the two recurl stations and will only be powered from the outer sides, only a short distance needs to be bridged using cables.

For the fibre detector and pixel detector layers, which are powered from both sides, power also needs to be delivered from the power-crates to the ends of the central station along the beam pipe.

In the following the chain is described for powering a MuPix sensor located on a half-ladder at the central layer 4. For this, the whole chain is needed and the voltage drop is substantial. Resistances and voltage drops are given for one-way, meaning they have to be doubled to get the final estimate.

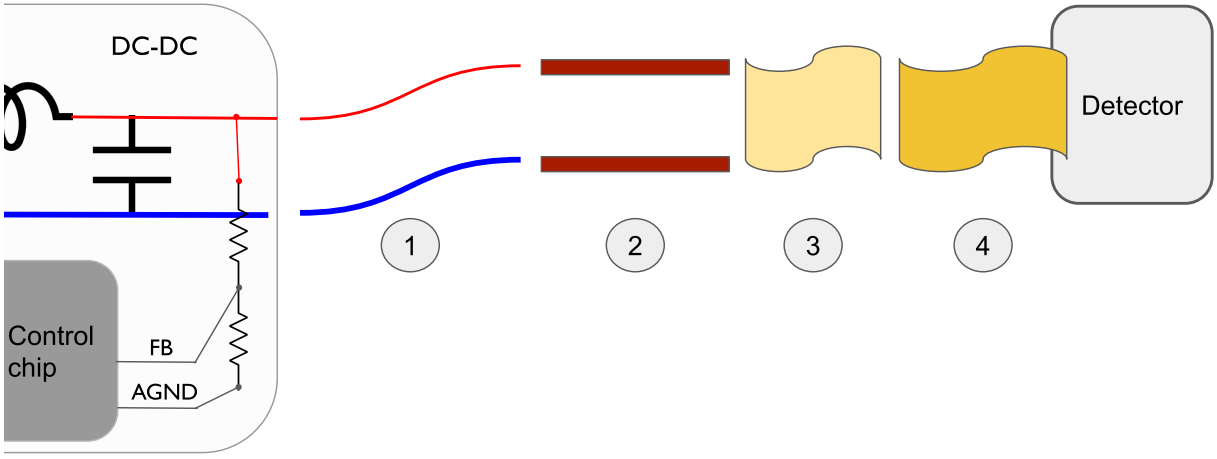

Figure 2.11: Chain of conductors to deliver power from a powerboard to a MuPix sensor located on a half-ladder at the central layer 4. Adapted from [6].

1: Cables

Cables of AWG 14 (2.08 mm^2) are used to connect the powerboard outputs to the copper bus bars, creating a flexible connection. The exact cable model is not specified yet, but resistances are around 9.2 mΩ/m (see for example AlphaWire 6829 BK001). An estimated length of 50 cm will be needed, giving a resistance of **4.6 mΩ**. For 20 A this results in a voltage drop of 92 mV.

2: Copper Bus Bars

Copper rods (5×2.5 mm^2 each), mounted on the beam pipe (see figure 2.13), are used to conduct power to the center endrings of the detector. To avoid any transfer of heat from the powerboards to the detector modules, the rods are cooled using a water-cooled ring. Using the special resistance of copper of $1.72 \cdot 10^{-8}$ Ω m, a cross-sectional area of 12.5 mm^2 and a length of 50 cm, the resistance is around **0.7 mΩ**. For 20 A this means a voltage drop of 14 mV.

3: Interposers and flex PCBs

The connection between endring - module - half-ladder (see figure 2.7) is done using Samtec Z-Ray® Interposers and flex PCBs. The resistance of the connection is not well known and needs further investigation. A rough estimate of 6 mΩ can be given based on the measurements in [7, p. 185]. For 20 A this gives a voltage drop of 120 mV.

4: Half-Ladder

MuPix sensors are located on so-called ladders, where each ladder is powered from both sides. This means only a half-ladder needs to be powered per side. For a layer 4

half-ladder, this means 9 chips. Power is delivered using aluminium traces located on the Kapton flex-prints. They have a thickness of 12 µm and a total width of 20 mm is available. Different ways of routing the traces for the individual chips have been evaluated. The naive approach of powering all chips from the same two power traces (in parallel) has the drawback that the first chip will get the full supply voltage, while the last chip will get the smallest supply voltage because of voltage drops along the traces. To prevent this problem, and deliver the same supply voltage to the inputs of each chip, equal resistance routing can be done (see figure 2.12). This results in an equal input trace resistance for each chip of **100 mΩ**. The resulting voltage drop is thus only dependent on the current assumed for one chip and not the total current of the chips in the same power partition. For ladders with fewer chips, the traces width can be increased, and resistance decreased. For that reason, the number should be taken as an upper limit. The typical current of 0.56 A per chip gives a voltage drop of 56 mV.

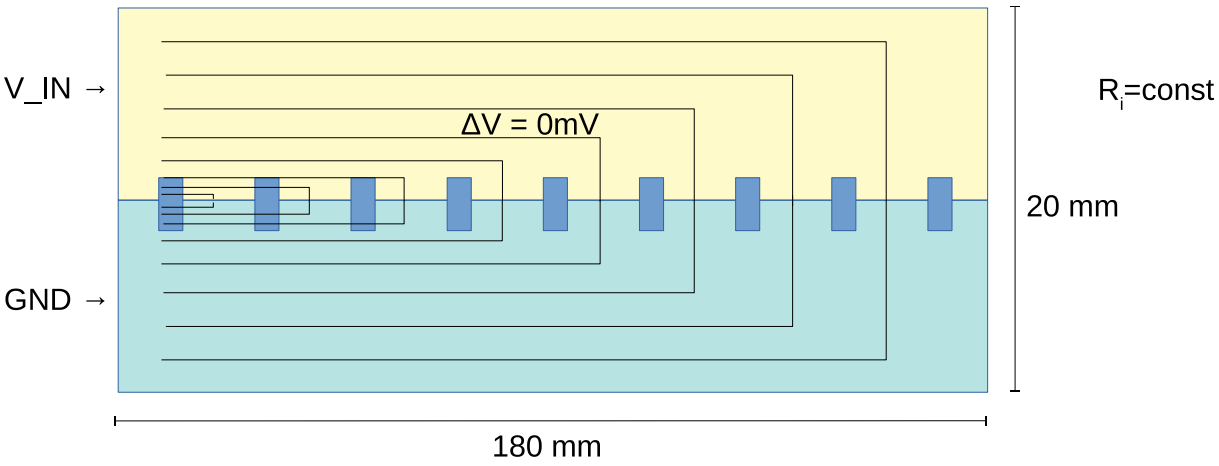

Figure 2.12: Equal Resistance Routing for a half-ladder with 9 MuPix sensors (blue). Traces are split and widths varied resulting in equal supply voltages for all chips [8].

Estimated voltage drop for different power partitions and an in-depth discussion can be found in section 2.5.

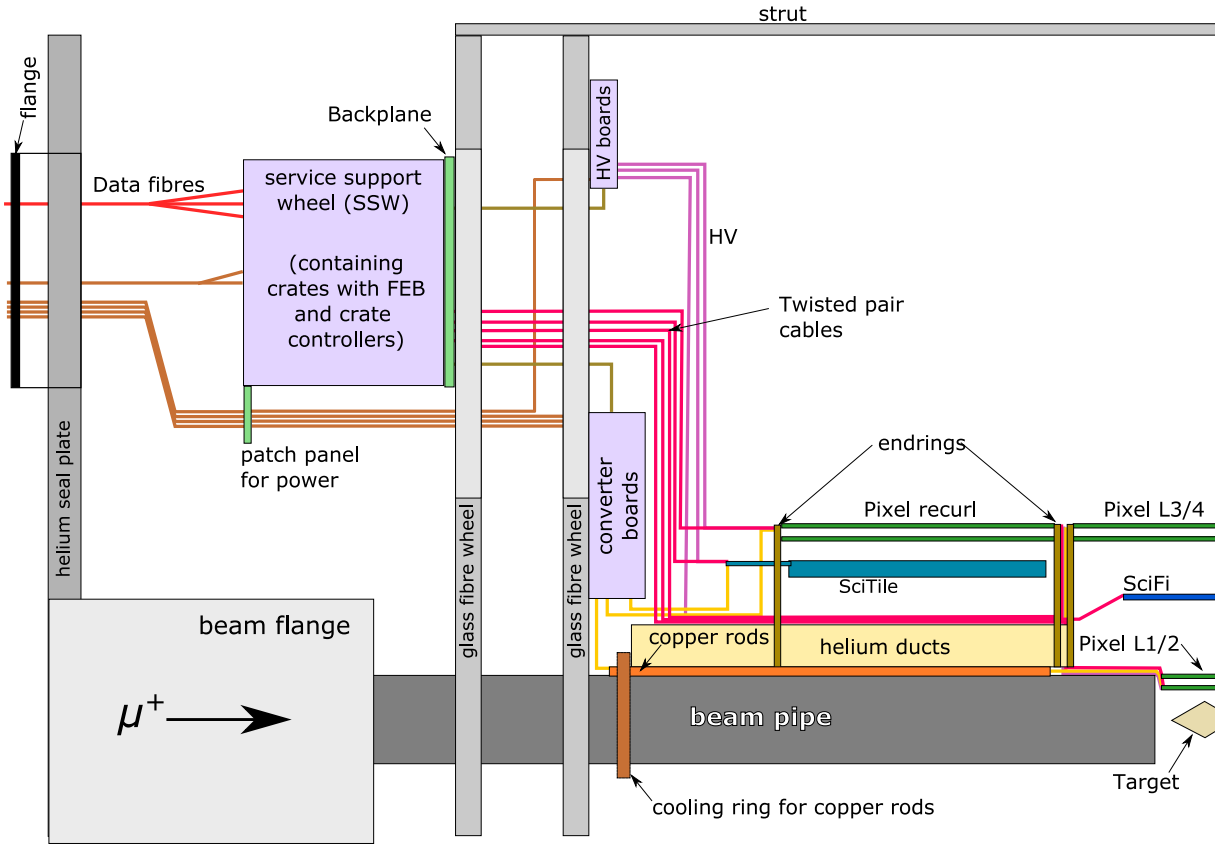

Figure 2.13: Conceptual view of cable routing [5, p. 74].

2.5 Power Requirements

Each FEB requires three different voltages at only a few amps (see table 2.3). These will be provided by buck converters embedded on the board itself.

Table 2.3: Power requirements for the front-end board (FEB). These are provided by embedded buck converters.

Voltage [V]	Typical Current [A]
1.1	2.0
2.5	2.0
3.3	1.0 - 2.3

Table 2.4 shows power requirements for different detector modules. For the MuPix sensors, a voltage of 1.8 V is needed. The typical current per sensor can be calculated assuming a power consumption of 250 mW/cm^2. With $A_{\text{mupix}} = 4\,\text{cm}^2$, the current is $I_{\text{mupix, typ}} = 0.56\,\text{A}$. For the maximum current we assume 350 mW/cm^2, resulting in a

current of $I_{\text{mupix, max}} = 0.78\,\text{A}$ per chip. The central station of MuPix layer 3 has ladders with 17 sensors and thus half-ladders with 8 and 9 sensors. The cable voltage drop for the MuPix partitions is calculated following the discussion in section 2.4.4 using the following equation:

$$\begin{aligned} V_{\text{cable}} &= 2 \cdot ((R_{\text{cable}} + R_{\text{copper bar}} + R_{\text{interposer}}) \cdot I_{\text{partition, max}} + R_{\text{ladder}} \cdot I_{\text{sensor, max}}) \\ &= 2 \cdot ((4.6\,\text{m}\Omega + 0.7\,\text{m}\Omega + 6\,\text{m}\Omega) \cdot I_{\text{partition, max}} + 100\,\text{m}\Omega \cdot I_{\text{sensor, max}}) \,. \end{aligned} \tag{2.4}$$

While the 2.2 V for the fibre is provided by the powerboard's buck converter, the low-current 3.6 V will be generated by an LDO. Space for the LDO has been designated on version 3 of the powerboard, but will only be populated on the 8 required boards. To calculate the cable voltage drop for the fibre partitions equation (2.4) with $R_{\text{ladder}} = 0$ is used.

Since the tile detector is located at the recurl stations and will only be powered from the outer side, cable voltage drops are small. Furthermore, the tile module board (TMB) has LDOs to reduce voltages to the final levels, meaning a slightly higher voltage than needed can be supplied.

Table 2.4: Power requirements per partition. These will be provided by the powerboards. The two-way (total) cable voltage drop is given.

Partition Type (ASIC)	#ASICS/ partition	Demanden Voltage [V]	Max. Cable Voltage Drop [V]	Current [A]	
				Typical	Max.
Pixel (MuPix)					
Layer 1	12	1.8	0.37	6.7	9.4
Layer 2	15	1.8	0.42	8.4	11.7
Layer 3 Central	32/36	1.8	0.79	19.0	28.1
Layer 3 Recurl	36	1.8	0.79	20.2	28.1
Layer 4	36	1.8	0.79	20.2	28.1
Fibre (MuTRiG)	8				
		2.2	0.20	7	9
		3.6	0	0.1	0.5
Tile (MuTRiG)	13				
		2.2	0	9	11
		3.6	0	3.1	4

Table 2.5: Power estimates for partitions and in total. Values calculated using table 2.1 and table 2.4. Power lost on input cables is not included. Converter efficiencies of 75 % are assumed.

Partition Type (ASIC)	Maximum Power per Partition [W]		Total Power [W]
	excluding	including Converter & Cable Drop	
Pixel (MuPix)			
Layer 1	16.8	27.1	108.2
Layer 2	21.1	34.6	138.6
Layer 3 Central	50.5	97.0	1163.9
Layer 3 Recurl	50.5	97.0	2327.8
Layer 4	50.5	97.0	4073.7
Fibre (MuTRiG)	23.5	30.7	367.7
Tile (MuTRiG)	38.6	51.5	720.6
Front-end board	266	350	2800
Total			**11701**

Table 2.5 shows power estimates for partitions and in total. Comparing the total power estimate of 11.7 kW to the estimate given in [5, p. 77] shows an increase by around 2.3 kW. This is mostly related to the cable voltage drops being properly taken into account here. To get a better understanding, the total power dissipated in conductors for partitions that are powered from the central station can be calculated. For these partitions, power is conducted through the water-cooled copper bars. Assuming decent thermal contact between the different conductors, a rough estimate of what cooling power is needed can be made. This includes all layer 1, 2, and central layer 3 partitions, half of the layer 3 recurl partitions and ⅔ of the layer 4 partitions. In total around 1 kW is lost in the conductors for these partitions (max. 1210 kW, typical 859 W), meaning a non-negligible extra 500 W of water-cooling power is need per detector side. Overall the cable voltage drops are quite high, especially for the high-current partitions. The main contributions and points of improvement are the AWG 14 cables and the interposers and flex PCBs connection. Increasing the cable diameter is an easy change, with the main limiting factor being the cable bending radius. The interposers and flex PCBs connection has not been thoroughly analysed yet, but trace width and the number of used pins for connections should leave room for improvement.

The cooling power needed per power-crate, assuming an efficiency of 75 %, can be calculated using the following formula:

$$P_{\text{power-crate}} = (P_{\text{total}} - P_{\text{FEB}}) \cdot 25\,\% \cdot \frac{1}{N_{\text{crates}}} \,. \tag{2.5}$$

With $N_{\text{crates}} = 8$, around 278 W per power-crate or 18 W per powerboard are expected. Assuming a flow similar to that in section 4.4 of 0.63 L/min, this would result in a temper-

ature increase of the cooling-water of 6 °C for a full-crate. This is substantial, but starting with a low base temperature, does not pose a problem.

Chapter 3

DC-DC Switching Regulators

Switched-Mode Power Supplies (SMPS) are used to convert AC/DC input power to regulated stable DC output voltage. They are used in everyday applications, for example, USB wall chargers or computer power supplies. They are valued for their low cost, small size, and high efficiency (up to above 90 %). The main disadvantage is the generation of electromagnetic interference (EMI) and noise from the high-frequency switching, which can be conducted or radiated to the load or nearby devices and might disrupt their stable operation (see section 3.2).

A regulator can be divided into the power and control stages. The power stage does the switching and energy conversion, while the control stage senses the output voltage and adjusts the frequency and/or duty cycle of the switching. The closed control loop allows for the output to be regulated to a stable DC voltage independent of input voltage or output current.

For DC-DC switching regulators, two basic topologies exist, depending on the required output voltage. For lower output than input voltage, a buck converter (step-down converter) is used, while for higher output than input voltage, a boost converter (step-up converter) is used. The Mu3e experiment uses buck converters, for which the operating principle is explained in the following section.

3.1 Buck Converter Operation Principle

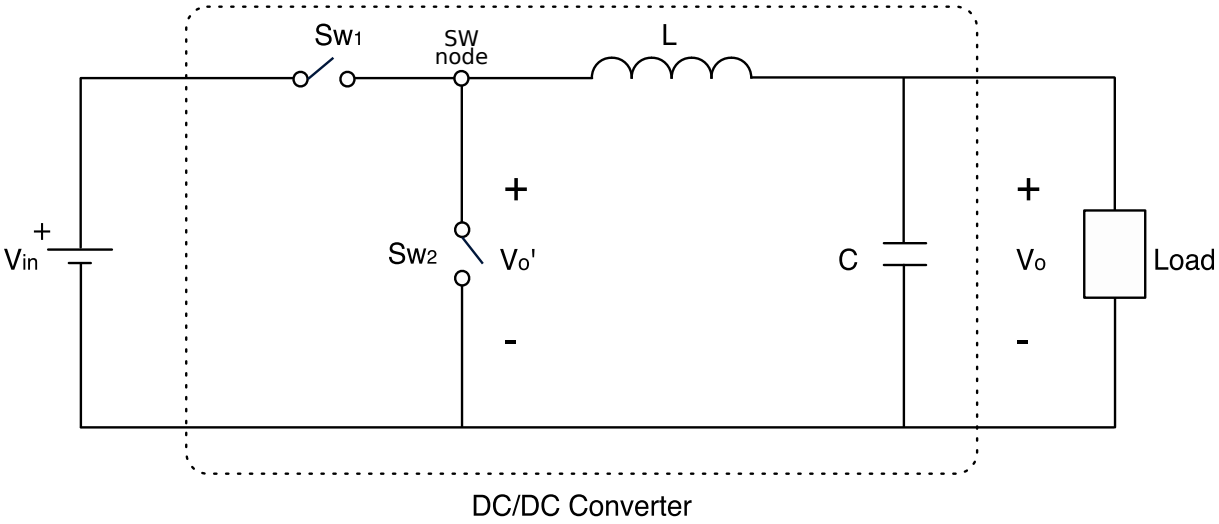

Figure 3.1: Simplified schematic of a synchronous buck converter [9, p. 20].

For a buck converter, the regulated output voltage is lower than the input voltage. The main components of the power stage are the switches and the two energy storing elements, inductor and capacitor (see figure 3.1). When using a high- and low-side switch, it is called a synchronous buck converter, since the switches need to be synchronised to avoid, for example, the case where both switches are conducting and thus shorting the inputs. One can replace the low-side switch with a diode, making the control chip less complex and cheaper, but having drawbacks like introducing more noise and a lower efficiency.

The principle of the buck converter can be easier understood when looking at the two distinct phases during operation, as seen in figure 3.2.

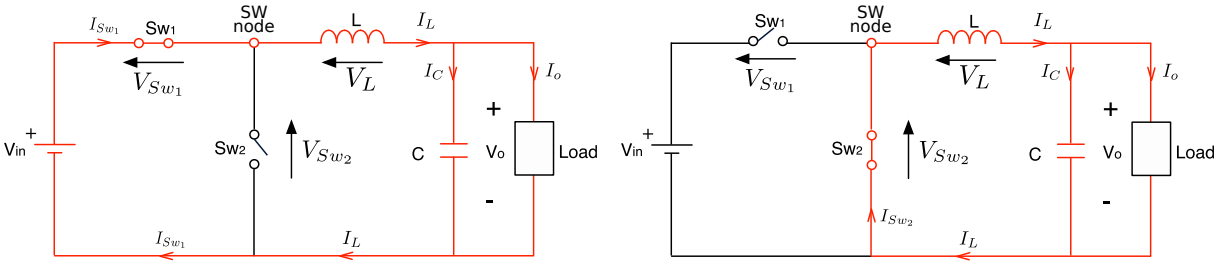

(a) High-side switch conducting.

(b) Low-side switch conducting.

Figure 3.2: The two phases of the buck converter. High-side switch or low-side switch conducting [9].

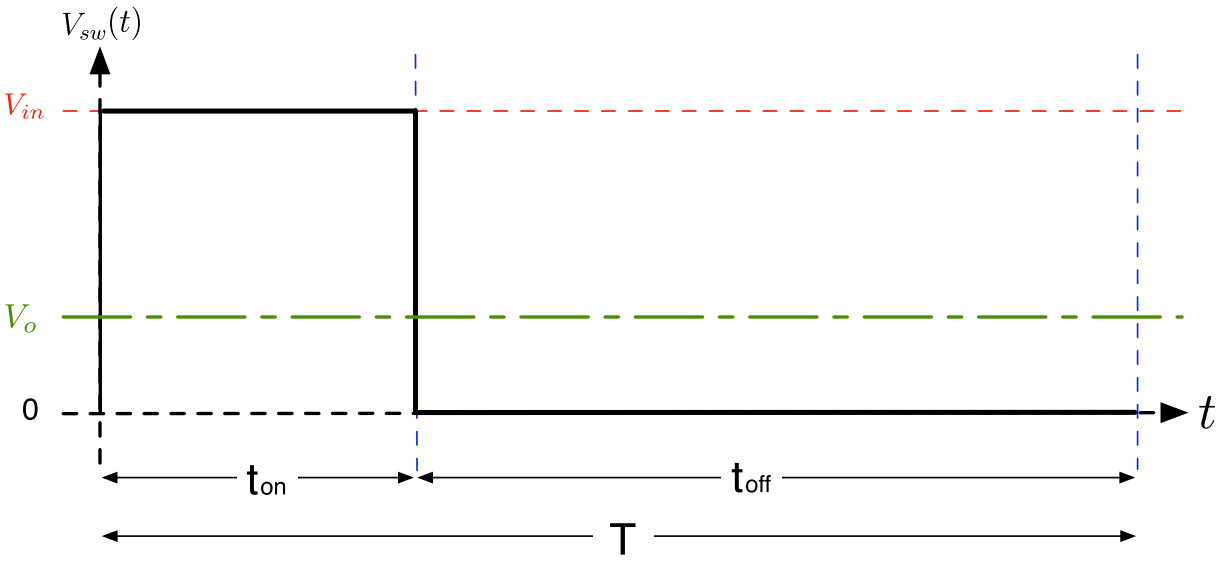

Figure 3.3: Voltage at the switching node over one period [9].

For the first phase of the period, the high-side switch is conducting while the low-side switch is open, giving a voltage of V_{in} for t_{on} at the switching node (see figure 3.3). For the second phase, the high-side switch is open and the low-side switch is conduction, giving a voltage of 0 V for $T - t_{\text{on}}$. The resulting square wave is averaged using the inductor and capacitor, which store energy in their respective fields during the first phase and release energy during the second phase. This LC combination can be seen as an LC filter, hence the naming "secondary LC filter" when talking about an additional output filter in section 5.1. The output voltage is then just the average voltage at the switching node over one period:

$$\begin{aligned} V_{\text{out}} = V_o &= \overline{V_{sw}} \\ &= \frac{1}{T}\int_0^T V_{sw}(t)\,\mathrm{d}t \\ &= \frac{t_{on}}{T}\cdot V_{in} = D\cdot V_{in}\,, \end{aligned} \tag{3.1}$$

with the duty cycle $D = {}^{t_{on}}/_{T}$. To get a regulated output voltage, the control chip keeps a constant switching frequency while adjusting the duty cycle accordingly (i.e. pulse-width modulation).

Voltages across different elements and the current through the inductor can be seen in figure 3.4. The inductor voltage for the first phase of the period is $V_{\text{in}} - V_o$. The general relation between voltage across an inductor and the current through it is given by:

$$V_L = -L\cdot\frac{\mathrm{d}I_L(t)}{\mathrm{d}t}\,. \tag{3.2}$$

This can be rearranged to give the current through the inductor:

$$I(t) = I(0) + \frac{1}{L}\int_0^t V_L(t')\,\mathrm{d}t' \tag{3.3}$$

$$= I(0) + \frac{V_{\text{in}} - V_o}{L} \cdot t = I(0) + m_1 \cdot t\ . \tag{3.4}$$

This means the current is increasing with a slope of $m_1 = {}^{(V_{\text{in}}-V_o)}/_L$ for the first phase. For the second phase, the voltage across the inductor is $0\,\text{V} - V_o = -V_0$ and with that the current

$$I(t) = I(0) + \frac{-V_o}{L} \cdot t = I(0) - m_2 \cdot t\ . \tag{3.5}$$

The current is decreasing with a slope of $m_2 = {}^{V_o}/_L$ for the second phase. The inductor helps to keep a stable output current while also limiting the maximum current, giving rise to the inductor ripple current with its typical triangular waveform. For low output current, the inductor current can decrease to 0 A or even get negative. The converter is then said to operate in **discontinous conduction mode** (DCM), which is often unwanted and can result in more noise. For that reason, the inductance needs to be matched to the typical output current and can for example be increased to reduce the amplitude of the inductor ripple current. If the inductor current stays positive at all times, the converter is said to operate in **continous conduction mode** (CCM).

The capacitor is used to stabilise the voltage across the load and ground the unwanted AC component of the inductor ripple current (see section 3.2).

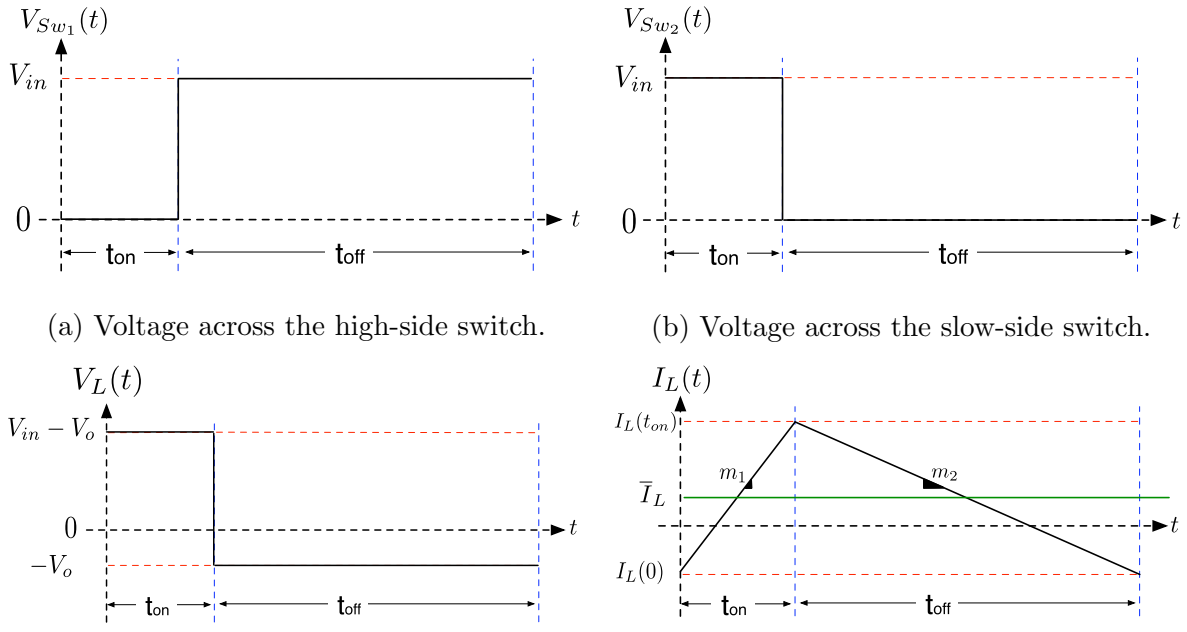

(a) Voltage across the high-side switch.

(b) Voltage across the slow-side switch.

(c) Voltage across the inductor.

(d) Current through the inductor.

Figure 3.4: Important signals of the buck converter [9].

As previously mentioned, the output voltage only depends on the input voltage and the duty cycle. This statement is only true for steady-state operation, i.e. constant input voltage and constant load current, and does not hold for intermediate transient states.

In the following, the transient for a suddenly decreased load resistance is analysed. When the load resistance is decreased, the output voltage will decrease, since the average current through the inductor, i.e. the output current, does not change instantly. This means the voltage across the inductor (during the two phases) changes, resulting in an increase of m_1 and a decrease of m_2. Overall, the average inductor current will increase until the output voltage is $D \cdot V_{in}$ again and a steady-state is reached. The system is thus self-regulating.

This result can also be seen in the simulation (see figure 3.5 and figure 3.6), which uses a Schottky diode as a low-side switch and a 100 kHz square wave with a 50 % duty cycle to switch the high-side MOSFET.

This means a controller is in theory not needed. In reality, even with a stable input voltage, the self-regulation is rather slow and a controller can minimize the sudden voltage drop by increasing the duty cycle temporarily.

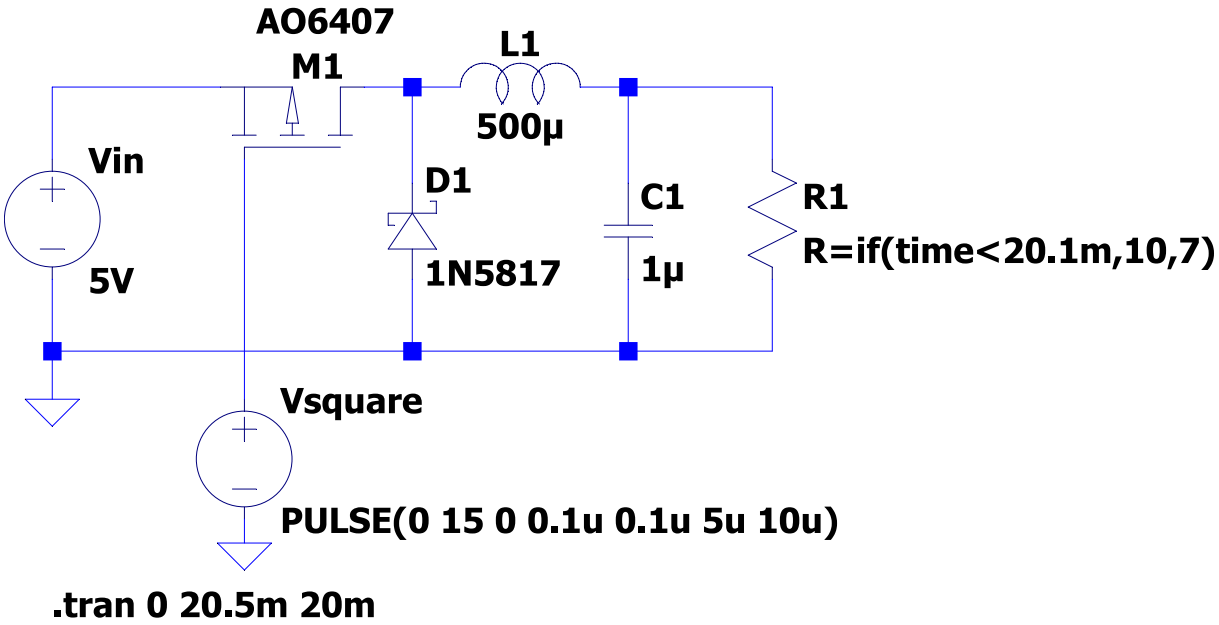

Figure 3.5: Circuit to simulate a buck converter with a sudden decrease in load resistance (at 20.1 ms from 10 Ω to 7 Ω) in LTspice®.

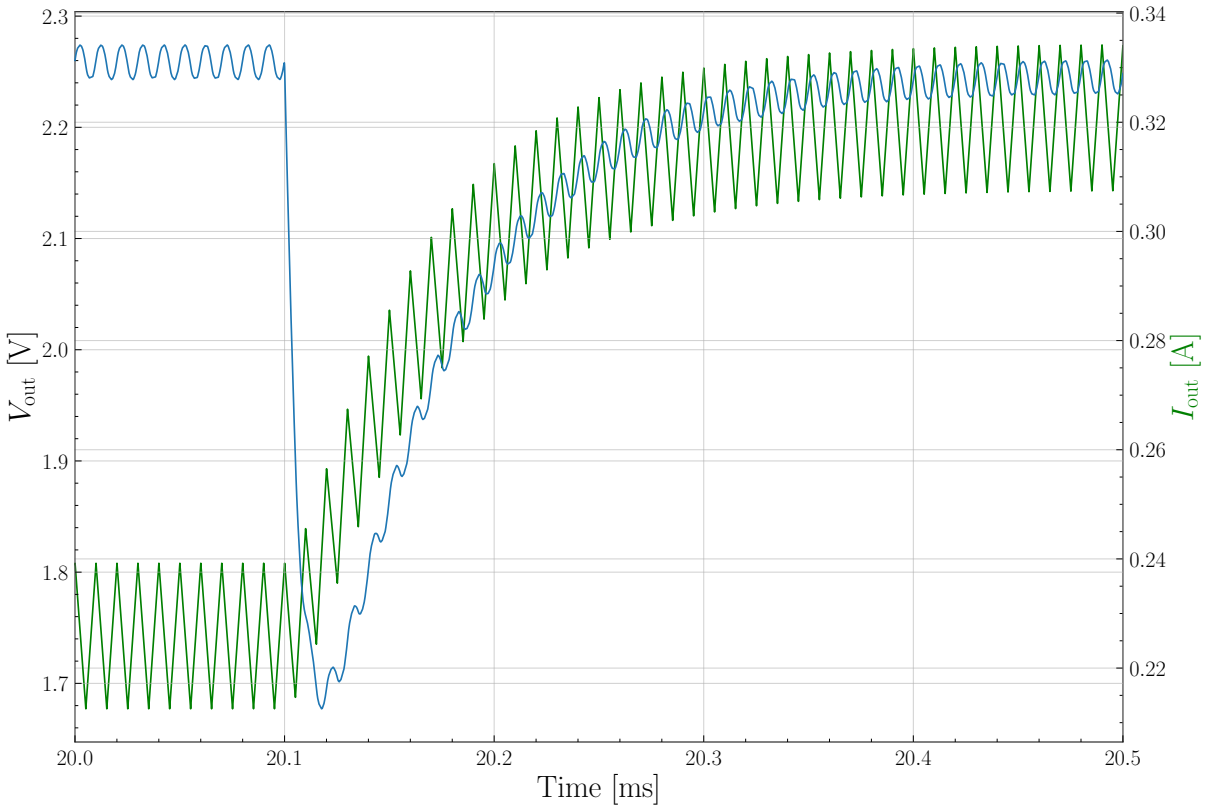

Figure 3.6: Simulation results showing the buck converter self-regulating by increasing the ouput current until $V_{\text{out}} = D \cdot V_{in}$ is fulfilled again.

3.2 Noise

Switched-mode power supplies can reach efficiencies up to 90 % and are for that reason very popular. But they come with the drawback of introducing switching noise to the system (see figure 3.7). With form factors getting smaller and thus smaller inductor and capacitor values, switching frequencies are increasing (up to a few MHz). Without proper consideration and filtering, this can lead to stability problems at the load and emission of EMI.

Noise in SMPS is mostly created by unwanted parasitic elements and can never be fully avoided, just minimized (see figure 3.9). It can predominantly be divided into two categories, low-frequency (LF) ripple, and high-frequency (HF) noise. Each is caused by different processes.

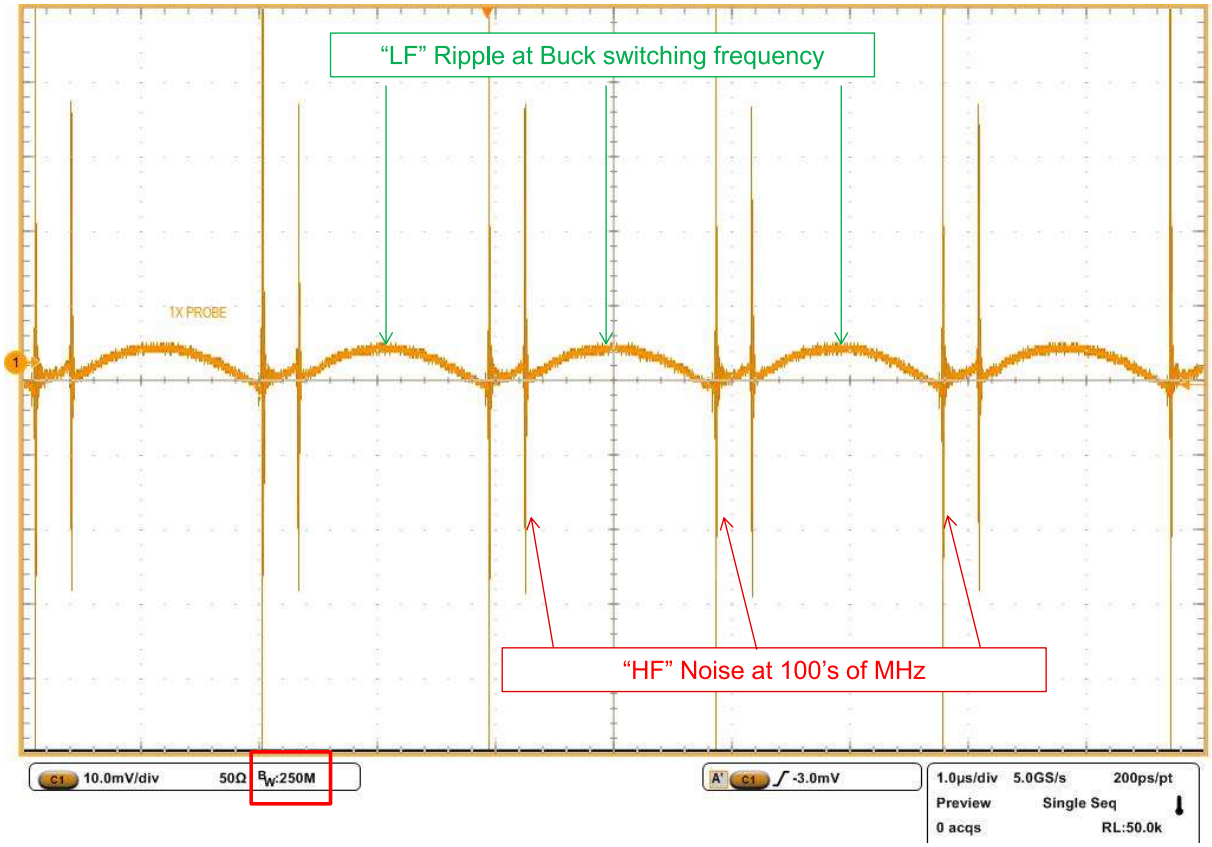

Figure 3.7: Scope trace of the output voltage of a buck converter. LF ripple and HF noise can be seen [10].

3.2.1 Low-Frequency Ripple

LF ripple appears at the switching frequency. It is a result of the inductor ripple current and the output capacitor impedance (see figure 3.8).

The inductor output current has a triangular waveform. The frequency is equal to the switching frequency. While the DC component (I_o) flows to the output, the AC component (I_{ac}) is grounded through the output capacitor. This current causes a time-dependent voltage drop across the capacitor, which modulates the output voltage. The waveform of the modulation is the sum of the voltage drops across the different elements of the output capacitor. For the resistive element, it varies linear with the current. For the capacitive element, it is proportional to the integral of the current and for the inductive element, it is proportional to the derivative of the current:

$$U_R(t) = R \cdot I_{ac}(t), \tag{3.6}$$

$$U_C(t) = \frac{1}{C} \int I_{ac}(t)\,\mathrm{d}x, \tag{3.7}$$

$$U_L(t) = -L \cdot \frac{\mathrm{d}I_{ac}(t)}{\mathrm{d}t}. \tag{3.8}$$

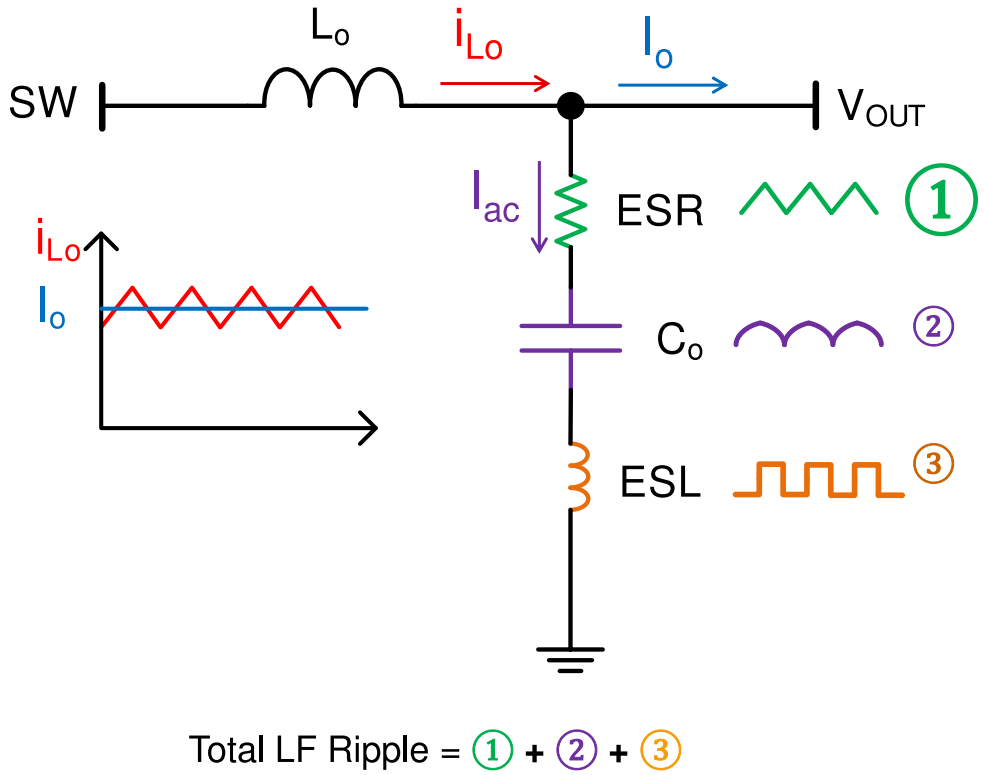

Figure 3.8: LF ripple resulting from inductor ripple current and output capacitor impedance [10].

To reduce LF ripple, the inductance of the main coil or the switching frequency can be increased. This results in a smaller amplitude of the inductor ripple current. Another option is to reduce the parasitic elements of the output capacitor. This can be done by using ceramic capacitors (MLCC) since they have low equivalent series resistance (ESR) and low equivalent series inductance (ESL) compared to other capacitors (for example aluminium electrolytic capacitors have high ESR). To further reduce ESR, the output capacitance can be split between multiple capacitors in parallel. Further ripple reduction can be achieved by actively filtering out noise using a second low-pass LC output filter (see section 3.3 and section 5.1). Due to the power dissipation, an RC output filter is only useful for very low output current applications.

3.2.2 High-Frequency Noise

HF noise appears on the switching node as ringing at each edge with frequencies up in the 100 MHz range (see the top of figure 3.9). It is present twice per period, once for the low-side and once for the high-side switching (as can be seen in figure 3.7).

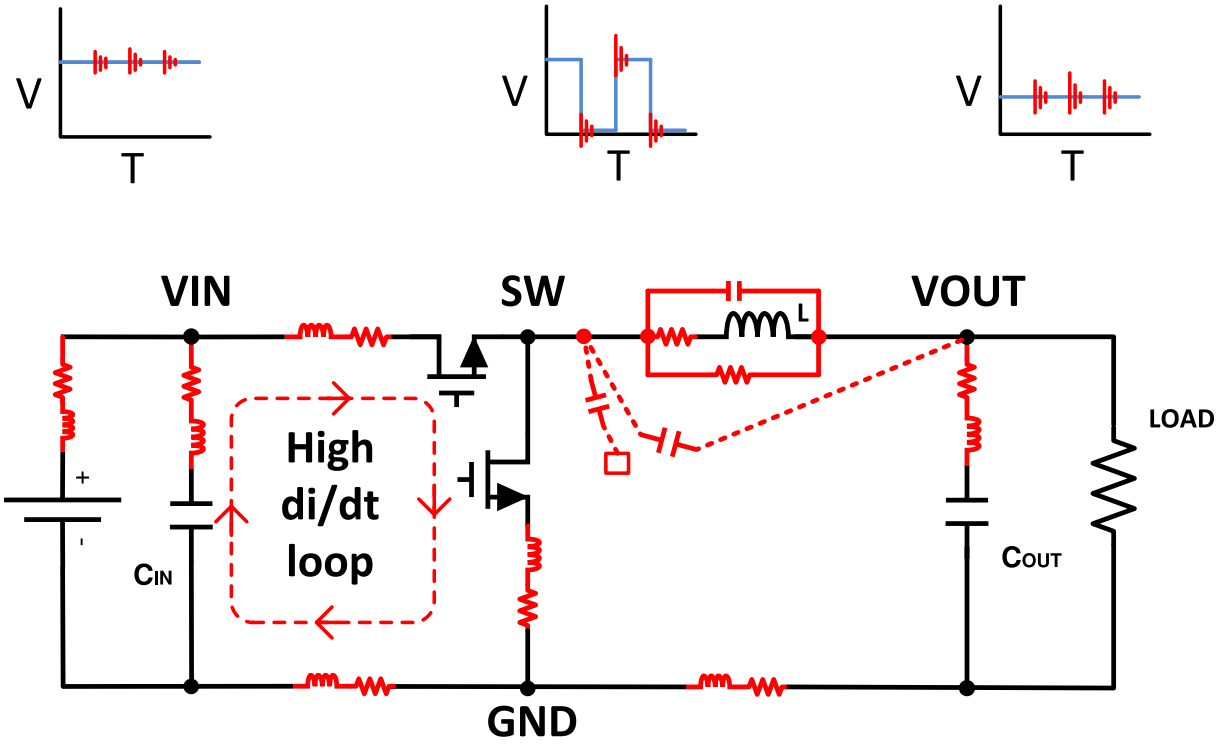

Figure 3.9: Buck converter with parasitic elements in red [10].

Because of the sharp switching, a high dI/dt-loop is formed on the input side of the converter. The loop contains the switches and the input capacitor C_{IN}. Any inductance in this loop will result in voltage spikes and ringing of the switching node. The noise is coupled to the output of the converter by parasitic capacitances between the switching node and output. These arise for example from the inductor or between PCB traces. But also the input of the converter can be polluted by HF noise, a reason why input filters are common for SMPS.

To keep the noise minimal, the inductance of the high dI/dt-loop has to be kept as small as possible. This can be done by placing the input capacitor as close as possible to the switches. The parasitic capacitance between switching node and output can be reduced by choosing an appropriate PCB stack-up and taking care when routing traces.

3.3 Low-pass LC Filter

A low-pass LC filter uses an inductor and a capacitor to filter out high-frequency noise (see figure 3.10). Since it uses two active components, it is a second-order filter. Having only ideal components, the transfer function (depending on the complex frequency $s = \sigma + j\omega$) is:

$$H(s) = \frac{\frac{1}{LC}}{s^2 + \frac{1}{LC}} . \tag{3.9}$$

Comparing this to the generalized transfer function of a 2nd order low-pass filter:

$$H(s) = \frac{w_0^2}{s^2 + \frac{w_0}{Q}s + w_0^2} , \tag{3.10}$$

one can identify the resonance frequency of $f_0 = {}^{w_0}/_{2\pi} = {}^{1}/_{2\pi\sqrt{LC}}$ and the high quality factor of $Q = +\infty$. The denominator in equation (3.9) has two complex poles $(\pm j\sqrt{{}^{1}/_{LC}})$, resulting in a roll-of after the resonance frequency f_0 of $-40\,\mathrm{db/decade}$. For a first order filter this would only be $-20\,\mathrm{dB/decade}$. A disadvantage of the LC filter is the high quality factor/high amplification at the resonance frequency. While of advantage when trying to tune a radio receiver to a specific frequency, for filters it can lead to unwanted oscillations. For that reason different methods of damping can be used. Damping with a resistor in parallel to the inductor results in a smaller quality factor (see figure 3.11).

When passing the resonance in frequency space, a change of phase can be observed. Before the resonance input and output voltage are in phase whereas after the resonance they are shifted by 90°. If the filter is, for example, used in a control loop, this might lead to amplification of certain frequencies.

The LC filter is of interest for high current applications, since the DC signal resistance is ideally zero, resulting in low power dissipation.

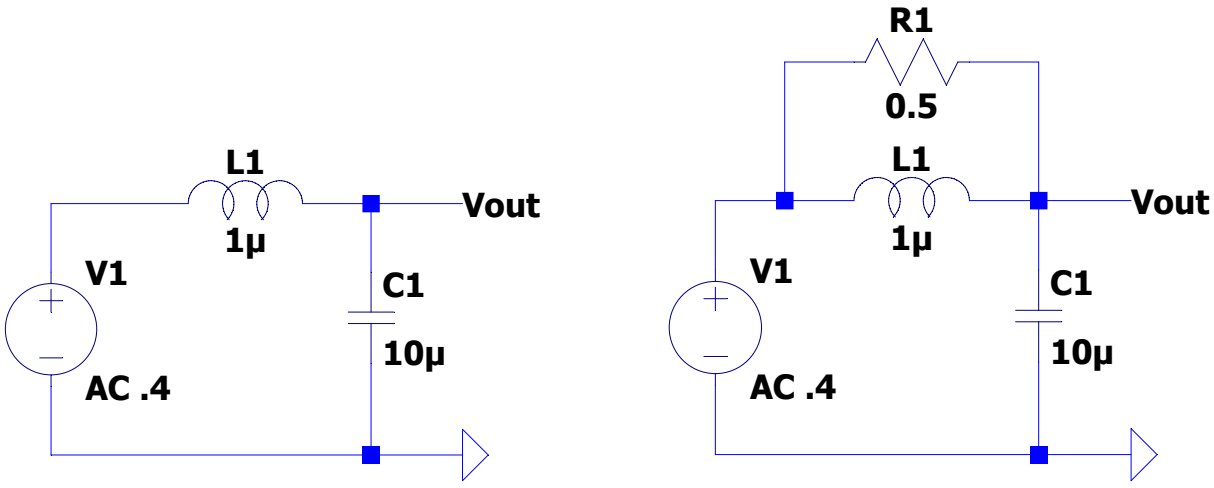

(a) Undamped (b) Damped

Figure 3.10: Circuits to simulate LC filter configurations in LTspice®.

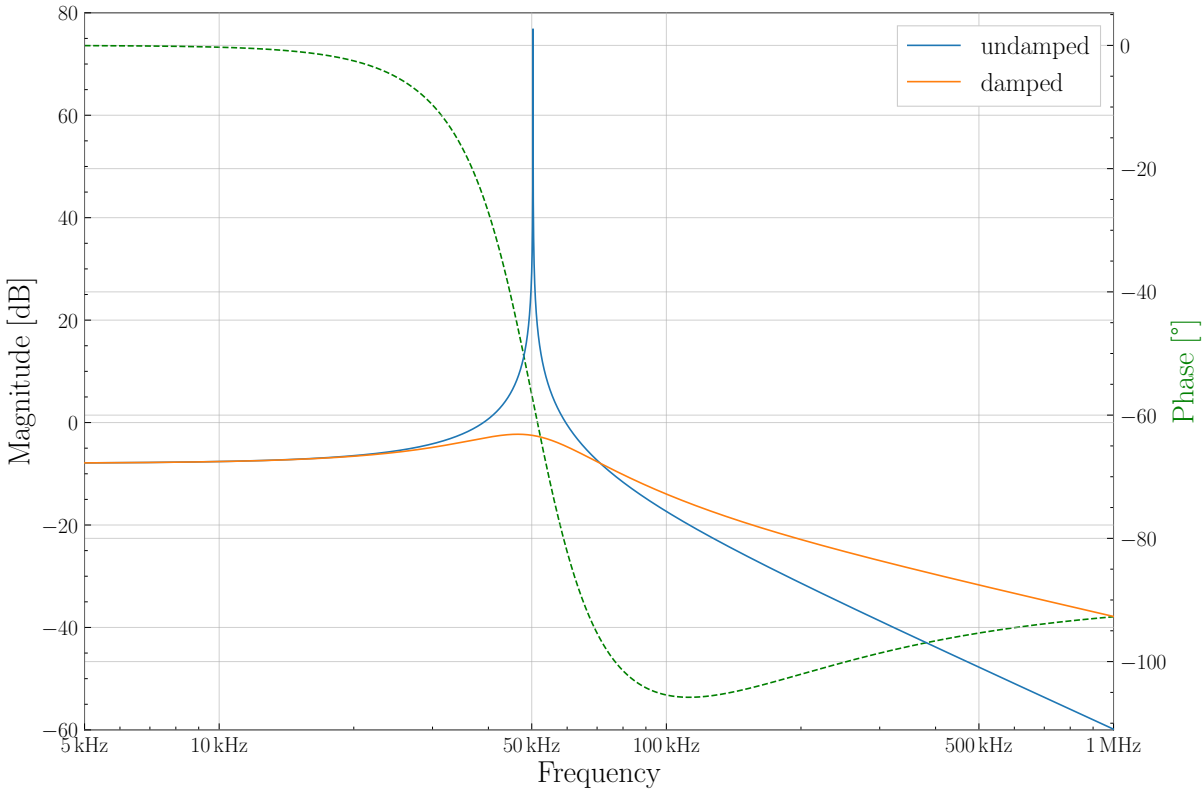

Figure 3.11: Bode magnitude and phase plot for an LC filter in the undamped and damped configuration.

3.4 The Mu3e Buck Converter

The Mu3e buck converter, referred to as powerboard, has been under design since 2019 and two iterations have been produced so far, v1 in 2019 and v2 at the end of 2020 (see figure 3.12). A third iteration is going to be produced at the beginning of 2022.

While the layout and auxiliary functions changed, the main specifications have been fixed. The PCB has a size of $100 \times 50\,\mathrm{mm}^2$. The input voltage is 20 V. For the output voltage, there are two configurations of the board. For one, the output voltage is around 2.2 V and the main inductor has a value of 550 nH. These will be used for powering of the MuPix sensors, fibre detector, and part of the tile detector. For the other configuration, the output voltage is around 3.6 V and the main inductor value has been adjusted to 3.5 µH. These will be used to provide the second required voltage for the tile detector (see table 2.4). Since the majority of boards will be in the first configuration (112 vs 14, see table 2.2), they have been evaluated mainly and if not mentioned otherwise, it can be assumed that a board with the 550 nH inductor is used.

The board has been designed with the maximum currents of table 2.4 in mind (around 30 A), but a long-time study of maximum currents over time has not been done yet, mainly

because version 2 of the board does not have a proper heatsink mount, which makes it difficult to provide adequate cooling. Version 3 solves this problem with a new mounting mechanism (see section 4.4.3).

Some of the main components are introduced in the following.

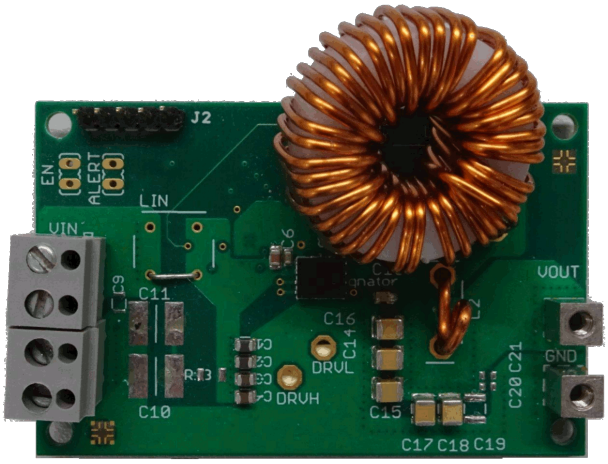

(a) Version 1

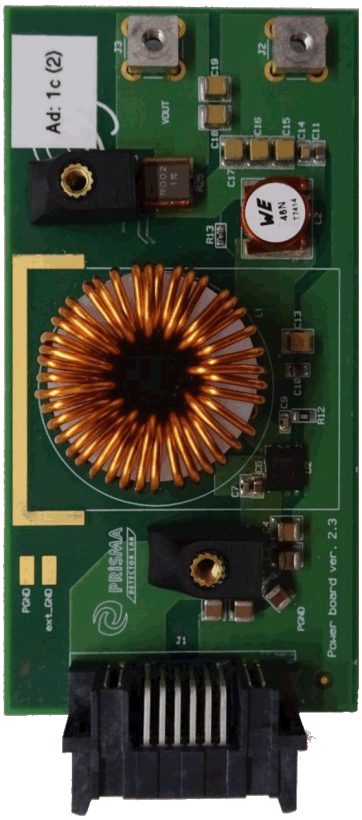

(b) Version 2

Figure 3.12: Iterations of the Mu3e buck converter.

3.4.1 TPS53819A Buck Controller (TPS)

The TPS53819A [11] is a small-sized controller chip designed to be used with a synchronous buck converter. It can be used for converters with input voltages of 3 V to 28 V and output voltages of 0.6 V to 5.5 V. It can drive two external N-channel MOSFETs with a gate drive voltage of 5 V.

A built-in LDO converts the V_{dd} supply voltage of 4.5 V to 28 V down to a regulated 5 V. A maximum current of 152 mA can be drawn from the LDO. For the powerboard, V_{dd} will be around 6 V to avoid significant heat generation in the TPS/built-in LDO.

Communication with the chip is done via the power management bus (PMBus), which contains commands for power supply management and runs over the two-wire I^2C protocol. An address can be set using a voltage divider (16 available, see [11, p. 22]). One

of 8 switching frequencies can be selected between 275 kHz and 1 MHz. For the Mu3e experiment, the highest switching frequency of 1 MHz will be used. The target voltage at the feedback pin is 0.6 V, but can be adjusted via PMBus by ±9 %, meaning the output voltage can be adjusted by this percentage.

The chip contains many more features, like overvoltage and undervoltage protection, allowing a safe and controlled operation of the converter.

3.4.2 CSD86360Q5D Synchronous Buck Power Block (CSD)

The CSD86360Q5D [12] power block contains two n-channel power MOSFETs in a half-bridge configuration and is optimized for 5 V gate drivers. It is specifically designed and optimized for synchronous buck converters. It minimizes many parasitic elements, for example, the impedance, which, as described in section section 3.2.2, can lead to HF noise. It can be operated up to 50 A and can utilize the GND plane as the primary thermal path by means of a thermal pad on the bottom.

3.4.3 Air Coils

The design and manufacturing of the required inductors was one of the first challenges on the way to a buck converter for the Mu3e experiment (see [13]). The powerboards will be operated in a 1 T magnetic field which saturates any inductors with ferrite cores. For that reason, air coils have to be used, i.e. coils that do not use a magnetic core material. A toroidal coil design was chosen over a solenoidal design since measurements showed that toroidal inductors emit significantly less EMI.

The coils required for the phase I experiment have already been manufactured (see figure 3.13).

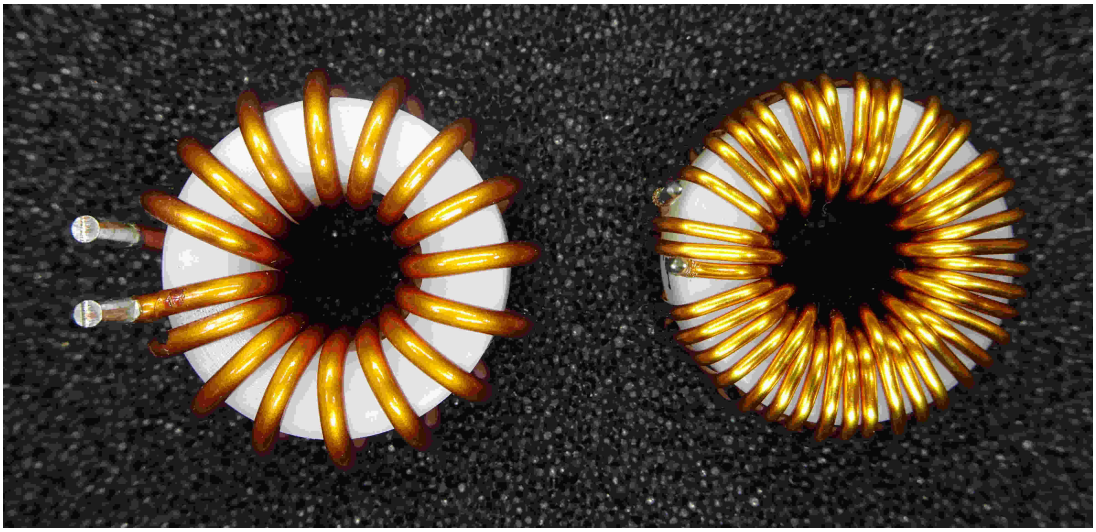

Figure 3.13: Air coils produced for the Mu3e experiment. Left: 550 nH. Right: 3.5 µH.

Chapter 4

Design of the Power-Crate

In the Mu3e phase I, the experiment will be using a total of 126 powerboards (see table 2.2). These will be located in 8 identical crates, four upstream and four downstream of the detector, each able to hold a maximum of 16 boards (see figure 2.10). The crates provide status and control interfaces via the crate controller, water-cooling, input power, and more. In this thesis, the CAD model of the full-crate (the final crate, holding 16 boards) was developed. To test the feasibility of the design on a reduced scale, a prototype crate, the mini-crate, was designed, built, and tested. It provides an easy way of operating and testing up to 4 powerboards v2 and allows a first assessment of the cooling system.

4.1 Full-Crate Design

The full-crate is the crate to be used in the experiment, allowing the operation of 16 powerboards. The developed design uses the powerboard v2 CAD model as a reference, but can be easily adjusted for later iterations. With boards plugged in, it has a size of $345 \times 150 \times 120\,\mathrm{mm}^3$. Technical drawings of the heatsinks and cooling blocks can be found in appendix A.3 and appendix A.4.

4.1.1 Mechanical Design

Five cooling blocks (figure 4.2, dark grey) are firmly mounted through the PCB to the backplate (figure 4.1, dark grey) from the bottom using 7 screws per block (M3, flat head). Each has an 8 mm through-hole with tube fittings attached on both sides. This allows direct water-cooling of the blocks. Cut-outs create four slots between two blocks, resulting in a total of 4×4 slots for powerboards. To further increase stability, the blocks can be connected at the front and back using a metal bar and two screws per block. The cooling blocks are designed to be all the same, resulting in cheaper manufacturing.

The metal backplate provides stability and stiffness to the crate and protects the PCB from damage and excessive forces. Cutouts are needed to avoid interference with through-

hole component leads and the screws are countersunk to allow flush mounting to the service support wheel. The exact mechanism for mounting the crate to the service support wheel is not yet implemented in the design, but simple modifications to the backplate should allow for a multitude of methods.

Heatsinks (figure 4.2, light grey) are firmly mounted onto the powerboards prior to plugging them into the crate (see section 4.2). This connection needs to be stiff, precise, and repeatable over all boards, since it is important for the quality of the thermal coupling of the heatsink to the powerboard. The position of the heatsink on the board and the height of the cooling blocks also constrain the overlap in the powerboard-crate connector. This connector has a few mm of play, allowing a reliable electrical connection and a good thermal connection between heatsink and cooling blocks at the same time. The heatsink with the attached board is then mounted into the appropriate slot using two M3 screws and thermal paste. The output power is provided by the two connectors located at the end of each powerboard (figure 4.2, brown). The crate controller can be plugged into the provided connector on the right side of the crate. It will have eurocard format ($100 \times 160\,\mathrm{mm}^2$), but has not been fully specified yet.

Overall, this design allows for dense operation of many boards and provides the needed cooling. At the same time, each board is easily accessible and can be swapped if needed.

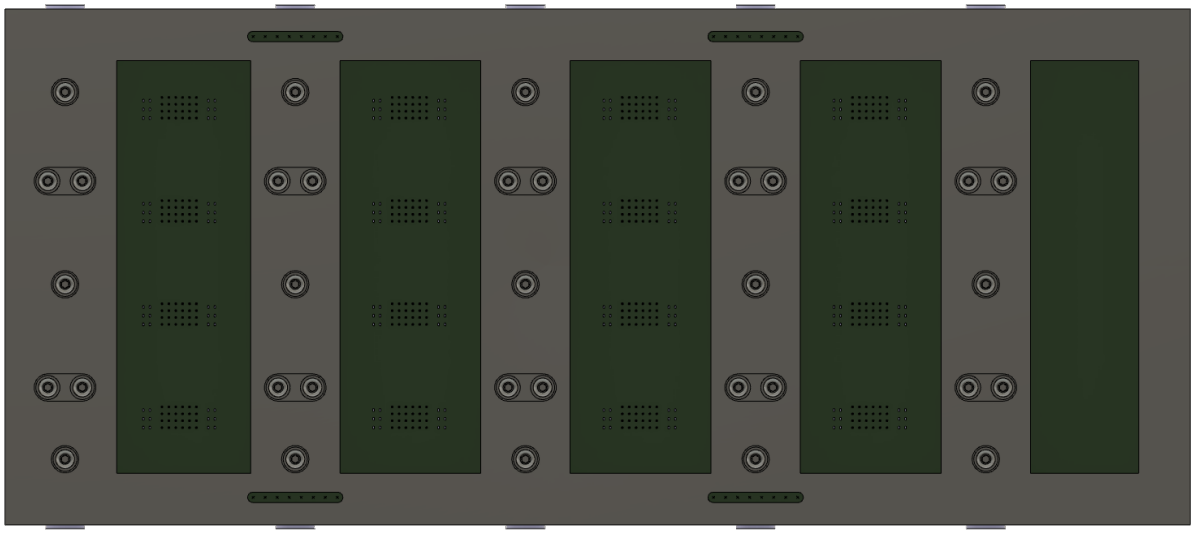

Figure 4.1: Full-Crate Design - Bottom View. Each cooling block is mounted to the backplate using 7 screws. Cutouts are needed to avoid conflicts with through-hole components.

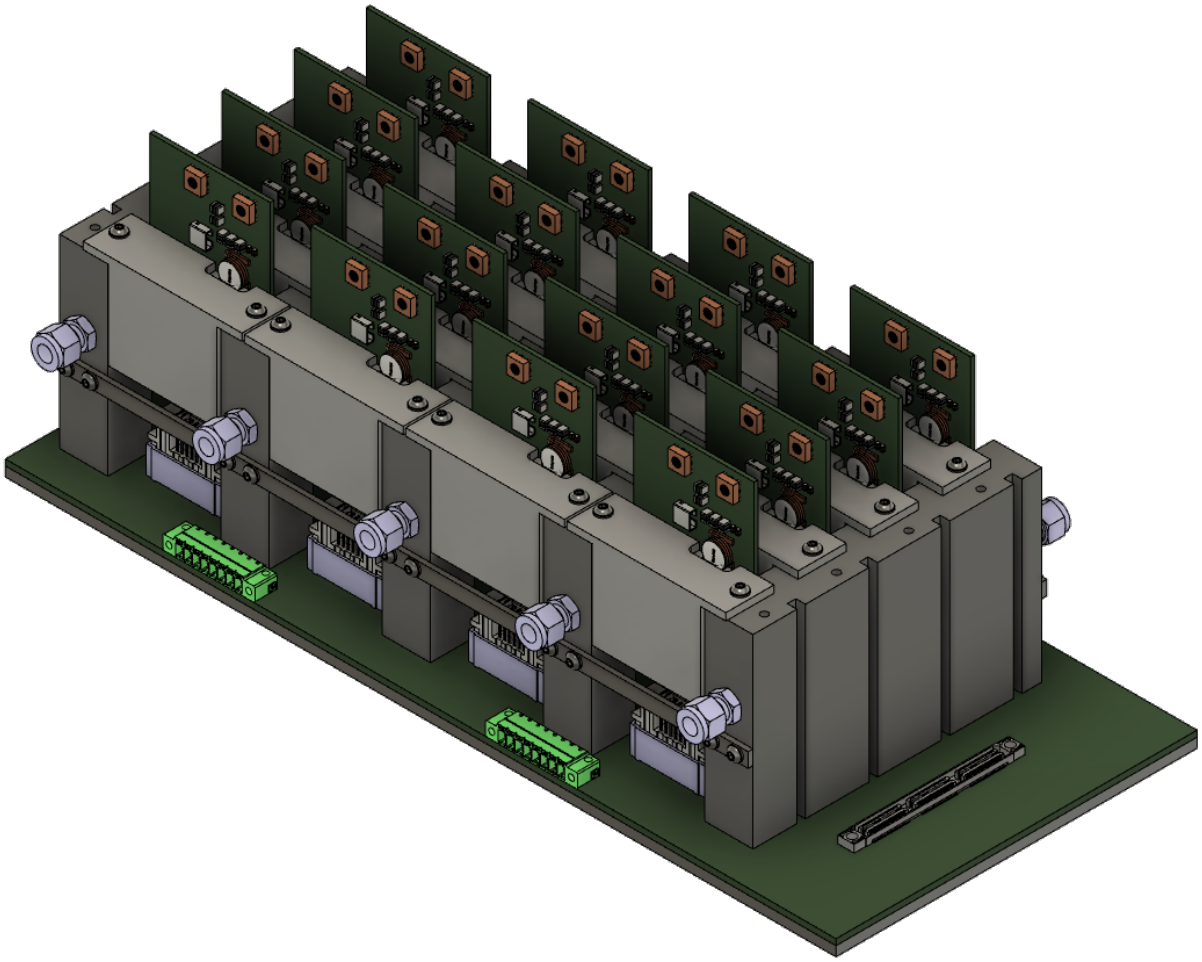

Figure 4.2: Full-Crate Design. The green connectors are used for input power, output power is provided on the two brown connectors on each powerboard. The connector on the right side is used for the MSCB [14] crate controller. Tubing for cooling has been left out for clarity.

4.1.2 Weight Estimation

The total weight of a fully populated crate has been estimated in table 4.1. With an estimated weight of around 5.2 kg, the crate is near the upper limit. Since the second and third heaviest components, heatsinks and powerboards, are not much more optimizable, the main work has to be done on the cooling blocks. Removing material on the bottom middle part of the blocks and fixing it to the backplate with fewer screws would reduce weight significantly, but the mechanical integrity needs to be kept in mind. Aluminium has a significantly lower shear strength than for example stainless steel (130 MPa vs 400 MPa to 690 MPa [15][16]), making it prone to thread stripping. To further reduce this risk, one could increase thread engagement, for example by using larger diameter screws or greater thread depth. Making the blocks more narrow is not practical, since it would result in less

contact area between the heatsink and cooling block and thus increase thermal resistance.

Table 4.1: Estimated weight of the full-crate. Metal parts are assumed to be made from aluminium (except screws).

Component	Weight [g]	Quantity	Percent of total weight [%]
Cooling Block (with tube fittings)	471	5	45
Heatsink (with thermal pads)	62	16	19
Powerboard	55	16	17
PCB (with connectors)	377	1	7
Backplate	293	1	6
Controller	120	1	2
Tubing (1m, filled with water)	84	1	2
Screw	1	67	1
Fixation Bar (with screws)	19	2	1
Total	**5.2 kg**		
Total (including 20 % safety margin)	**6.3 kg**		

4.1.3 Electrical Design

Different connectors have been evaluated and suitable ones chosen (see table 4.2). The backplane uses four 8 pin input connectors, containing seperate input and return lines for each board (figure 4.2, neon green). Conductors are connected using a push-in-spring mechanism and the connector can be locked using two screws. Powerboards are plugged in straight using a Samtec connector with two power pins (up to 28.8 A) and 24 signal pins. The crate controller is also plugged in straight (two power pins, 180 signal pins).

Table 4.2: Connectors used in the full-crate design.

Input	Header	Phoenix Contact MC 1,5/ 8-GF-3,5 - 1843855
	Housing	Phoenix Contact FMC 1,5/ 8-STF-3,5 - 1966156
Powerboards	Female	Samtec MPTC-01-24-01-6.30-03-L-V
	Male	Samtec MPSC-01-24-01-01-01-L-RA-LC
Controller		Samtec QSH-090-01-L-D-A-GP

4.1.4 Water Cooling and Galvanic Corrosion

Cooling of the crate is achieved by directly water-cooling the cooling blocks using the 8 mm through-holes and attached tube-fittings. The heatsinks, and thus powerboards, are in turn in thermal contact with these blocks.

Connecting the blocks to the water-cooling can be done in a serial or parallel configuration (see figure 4.3). The serial configuration is simpler, having automatically an equal flow through all blocks, but results in strongly uneven cooling. Since the powerboards are mostly indifferent to temperature within the specifications, as long as the cooling margin is big enough, this might be tolerable. To get more even cooling, one could also mix the order of cooled blocks, at the expense of more tubing. The parallel configuration results in more even cooling, but the flow has to be evenly split between blocks. If that is possible this configuration has the advantage.

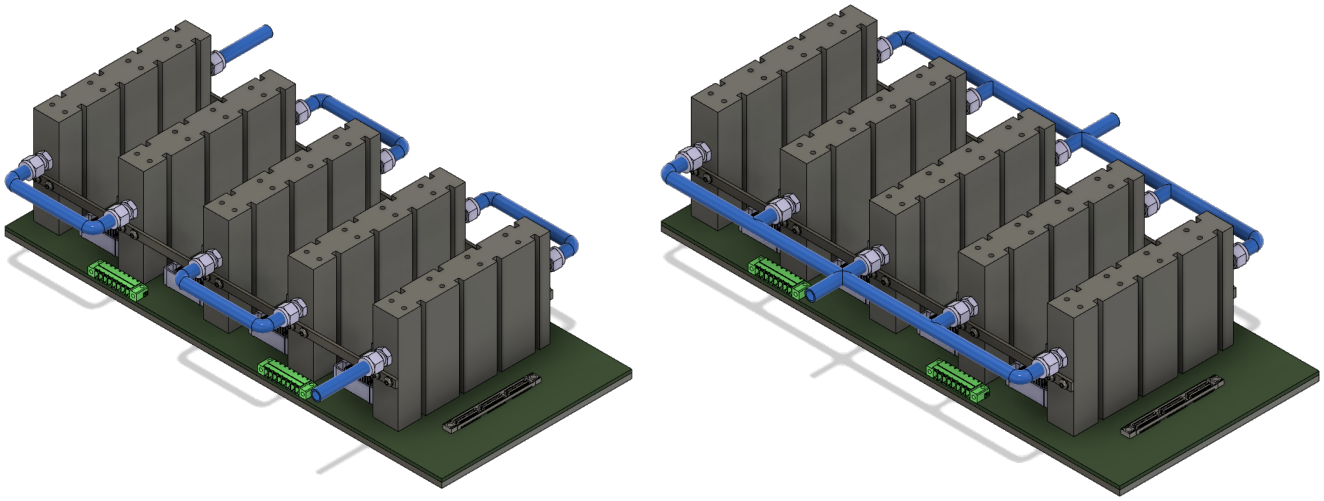

(a) Serial Configuration

(b) Parallel Configuration

Figure 4.3: Two ways of connecting the water-cooling for the full-crate.

Another important topic (as thoroughly discussed in all computer liquid-cooling forums) is galvanic corrosion. Galvanic corrosion is an accelerated form of corrosion. It is an electrochemical process and occurs when two dissimilar metals are in electrical contact in an aqueous electrolyte [17, p. 76]. In that case, the more noble (figure 4.4, left side) metal functions as the anode, while the more active metal (figure 4.4, right side) functions as the cathode (similar to a battery). The anode drives the corrosion of the cathode, where metal ions dissolve into the liquid while the electrons travel to the anode. The anode material stays unharmed. Of importance is the ratio between surface areas of anode and cathode. Having a larger cathode surface spreads the corrosion over a bigger area, resulting in locally less damage.

For the production of the cooling blocks and heatsinks, an aluminium alloy is used, since it has good thermal conductivity, is readily available, machinable, and light. In the experiment the cooling loop will be rather big, containing many different components with different metals, many electrically connected, for example, via a common ground. Since aluminum alloys are rather active, this opens up the door for galvanic corrosion. To mitigate this problem, parts in the cooling loop should if possible also be made out of aluminium. For the tube fittings, this is especially important, since they are in direct

electrical contact with the cooling blocks. Since this is not possible for all parts, it would be advisable to add corrosion inhibitors to the cooling water, which can greatly reduce corrosion. Another option is the use of sacrificial cathodes, for example, made from zinc.

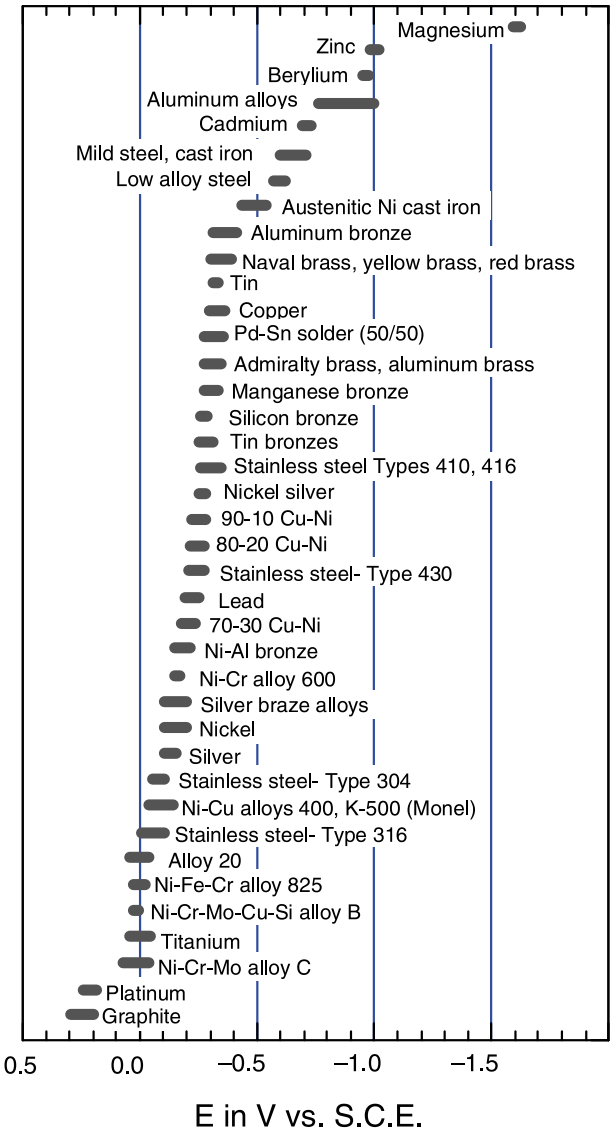

Figure 4.4: The galvanic series in seawater. Metals on the right are more active, metals on the left more noble. Electrode potentials are measured relative to a saturated calomel electrode. Alloys exhibit a range of electrode potentials [17, p. 77].

4.2 Heatsink Design

To design an efficient heatsink, the components with the most heat generation must be known. In [18, p. 34] (see figure 4.5), it was shown that most power is dissipated by the IC containing the power MOSFETs (CSD, see section 3.4.2) and in the main coil by ohmic

losses. In these measurements, the controller chip (TPS, see section 3.4.1) also got quite hot. This mostly relates to the built-in LDO, which steps down the controller power supply voltage V_{dd}, in this case chosen to be the same as $V_{\mathrm{in}} = 20\,\mathrm{V}$, to 5 V. Choosing a separate V_{dd}, for example, 6 V, leads to a negligible temperature increase of the controller chip. This would be the preferred configuration for the experiment. The low voltage could, for example, be generated by LDOs or small buck converters on the backplane, which could, if needed, be connected to one of the cooling blocks.

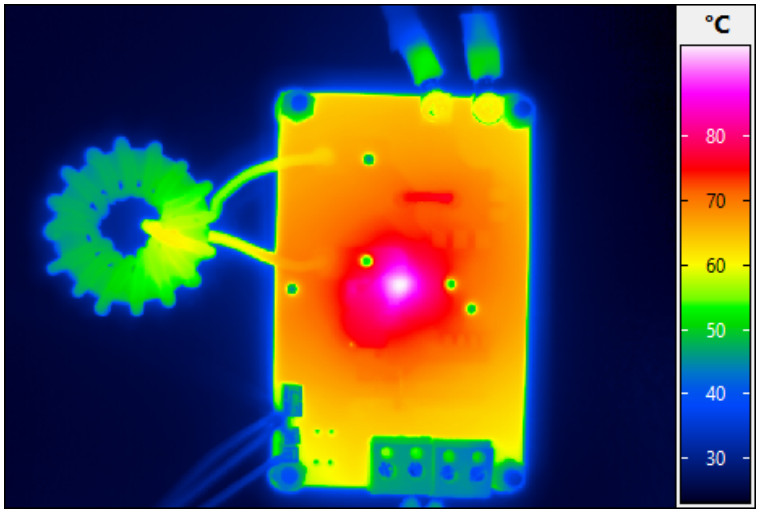

Figure 4.5: Thermal image of the front of powerboard v1 after 10 minutes with $I_{\mathrm{load}} = 13\,\mathrm{A}$ [18, p. 34].

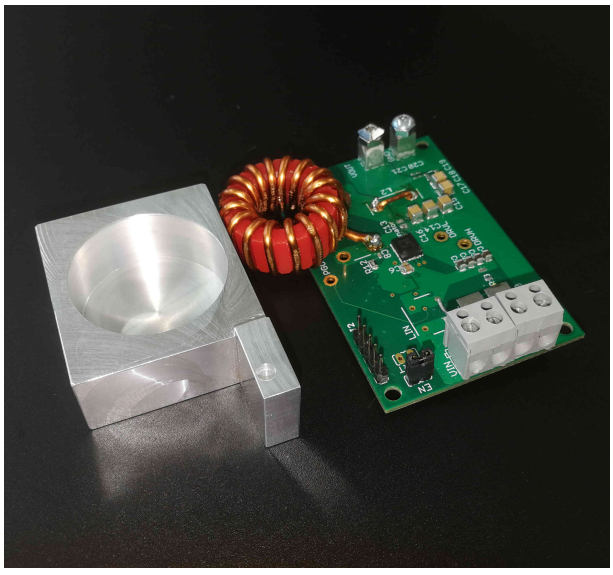

(a) Heatsink with cutout for coil and powerboard v1.

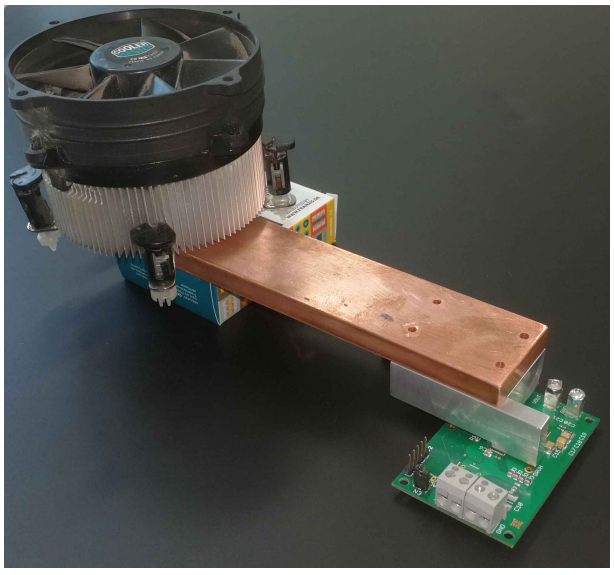

(b) Heatsink connected to heat exchanger via copper bar.

Figure 4.6: The first heatsink was designed for the powerboard v1.

Initial heatsink tests were done with the powerboard v1 (see figure 4.6) and showed promising results.

The next iteration of the heatsink, version 2, was designed for the powerboard v2 and is suitable for being mounted in the crate. The CSD chip has a height of around 1.5 mm and the corresponding cutout in the heatsink was designed to be 2 mm deep. This leaves a gap of 0.5 mm which is filled with a 1 mm thermal pad, giving a compression of 50 %. A thermal pad from Gelid solutions was chosen, which has a recommended compression of 40 % and a thermal conductivity of 12 $^{W}/_{mK}$ [19]. For the coil, the cutout was oversized to allow thermal pads on all sides (see figure 4.7a). Since the coil does not have flat surfaces, one could also use a two-component liquid gap-filling material in later iterations, for example, one of the two-part gap fillers from Bergquist [20]. Version 2 of the powerboard did not provide any mounting points for the heatsink. For that reason, a custom mount has been designed and 3D printed, which can be glued onto the board (see figure 4.7b). This is only a temporary solution since the exerted pressure is not evenly distributed and not high enough to create optimal thermal contact. Temperature measurements demonstrated that remounting of a nonideally mounted heatsink led to operational temperature differences of 5.5 °C, showing the importance of a good and consistent heatsink mounting mechanism.

The efficiency of the heatsink has been tested in section 4.4 and a discussion of the results and improvements for the next iteration can be found in the conclusion (section 4.4.3).

Another function of the heatsink is shielding. So far most work on the heatsink has been done for thermal optimization, but the heatsink is also used for containing the HF switching noise. For this, it needs to be grounded and care needs to be taken not to create ground loops. The shielding functionality of the heatsink needs to be further investigated and a ground scheme established.

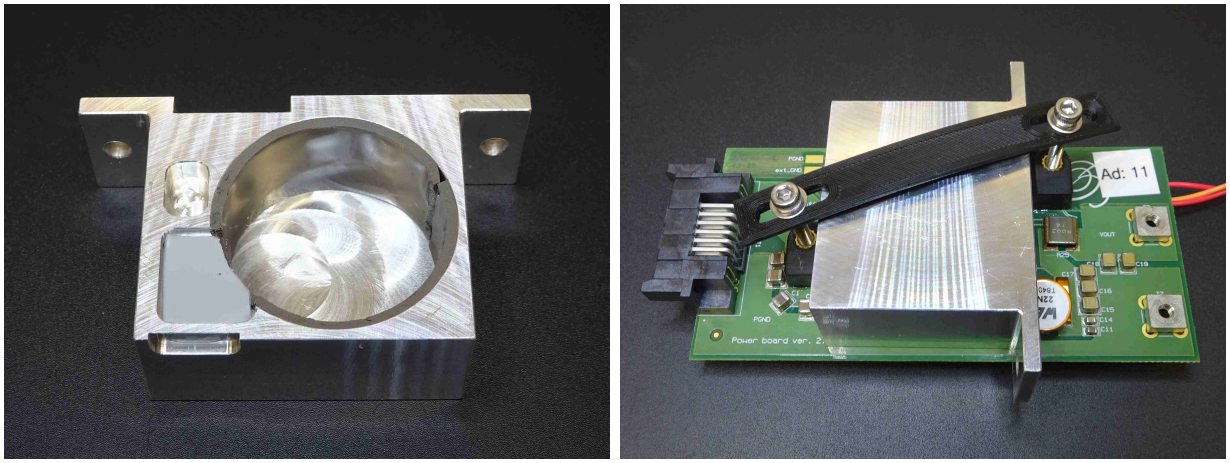

(a) Heatsink with cutouts and thermal pads for coil and CSD chip.

(b) Heatsink mounted to powerboard v2 using glued on threads.

Figure 4.7: Heatsink and mounting mechanism for powerboard v2 (which by design did not provide mounting points).

4.3 Designing and Building the Mini-Crate

The mini-crate is a smaller version of the full-crate, having only one row of slots (4 boards) (see figure 4.8). But since it uses the same components as the full crate, it is a good prototyping environment for the final design. The controller is connected externally, currently a Raspberry Pi is used, and water cooling can be installed if needed. With boards plugged in, it has a size of $170 \times 140 \times 131\,\mathrm{mm}^3$.

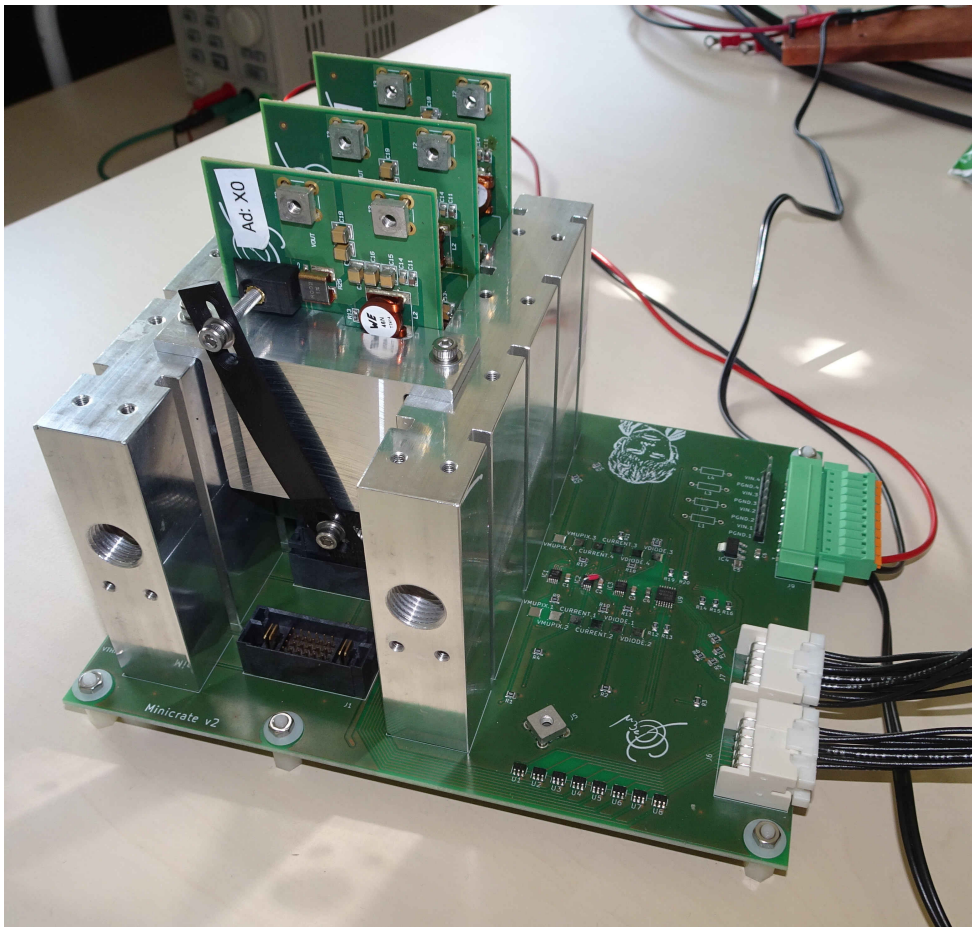

Figure 4.8: Mini-crate v2 with three boards plugged in.

4.3.1 Mechanical Parts

The designs for heatsinks and cooling blocks are taken from the full-crate. Parts were produced from the aluminium alloy EN AW-6060, which has superior thermal conductivity compared to other alloys ($209\,{}^{W}/_{\mathrm{m\,K}}$ [21] versus $172\,{}^{W}/_{\mathrm{m\,K}}$ for EN AW-6082 [22]). Technical drawings can be found in appendix A.3 and appendix A.4. During production, a small mistake in the design of the cooling blocks was found. The G 1/4" threading to attach the tube fittings penetrates the wall in two slots on one end, leading to a leakage. The holes

were sealed using two-component epoxy and, because the slots are greatly oversized, this did not lead to any problems during operation.

4.3.2 Backplane PCB Design

The PCB schematic of the backplane can be found in appendix A.2. It was designed for operating the powerboard v2. A 4-layer stackup was chosen, with signals on the front and first inner layer, ground plane on the second inner layer, and power traces on the back layer (trace thickness outer layer 70 µm, inner layer 35 µm). Cooling blocks are mounted through plated holes (not electrically connected), allowing extra wear resistance. A backplate was not deemed necessary for the weight of only two cooling blocks.

Two I2C busses are used, one is for the TPS53819A chips (buck converter controller, located on each board, see section 3.4.1). The other bus is for communicating with the three ADCs, which sample the analog signals (V_{mupix}, V_{diode}, V_{shunt}) coming from the boards. The separation gives a level of security against a failing device blocking one bus. Test points are provided for various analog signals.

As described in [18, p. 51], each board has an analog enable signal which needs to be pulled to V_{reg} externally to enable the converter. Using an IC with 4 integrated switches (TMUX1511 from Texas Instruments) allows simple enabling and disabling of the converters using the controller's (Raspberry Pi) digital logic pins. Care must be taken to choose an IC which can switch the needed 5 V, even if above supply voltage (in this design $V_{\text{sup}} = 3.3\,\text{V}$).

To generate the $V_{\text{sup}} = 3.3\,\text{V}$ needed for the ADCs and the TMUX chip from the external slow control voltage $V_{\text{slow}} = 6\,\text{V}$, an LDO is used (ZLDO1117 from Diodes Incorporated). Since LC input filters might be useful in future tests, space has been designated for one on each input. A pin header is connected to the input traces which can be used for powering boards from the same input (necessary for the tile detector configuration, where two converters share the same power partition providing the two required voltages, see section 2.4.3). For connecting the control signals, JST XAD headers (S12B-XADSS-N(LF)(SN)) are used on the PCB and TE Connectivity 20 position housings (1-87631-5) on the Raspberry Pi.

4.3.3 Control and Status Interface

The mini-crate can be monitored and controlled using an interface written in Python running on a Raspberry Pi (see figure 4.9). Plugged in boards are automatically recognized on startup and the default board configuration is loaded.

The top right panel allows simple control of the boards. By calling high-level functions, configuration registers can be written and boards can be controlled. The ADCs are continuously read out and values are printed in the left panel. Configuration and status information are regularly updated and shown in the bottom right panel. Here the tem-

peratures of the boards and of various points on the crate are also shown. If needed, data from the ADCs and temperature sensors can be logged to a file at variable time intervals.

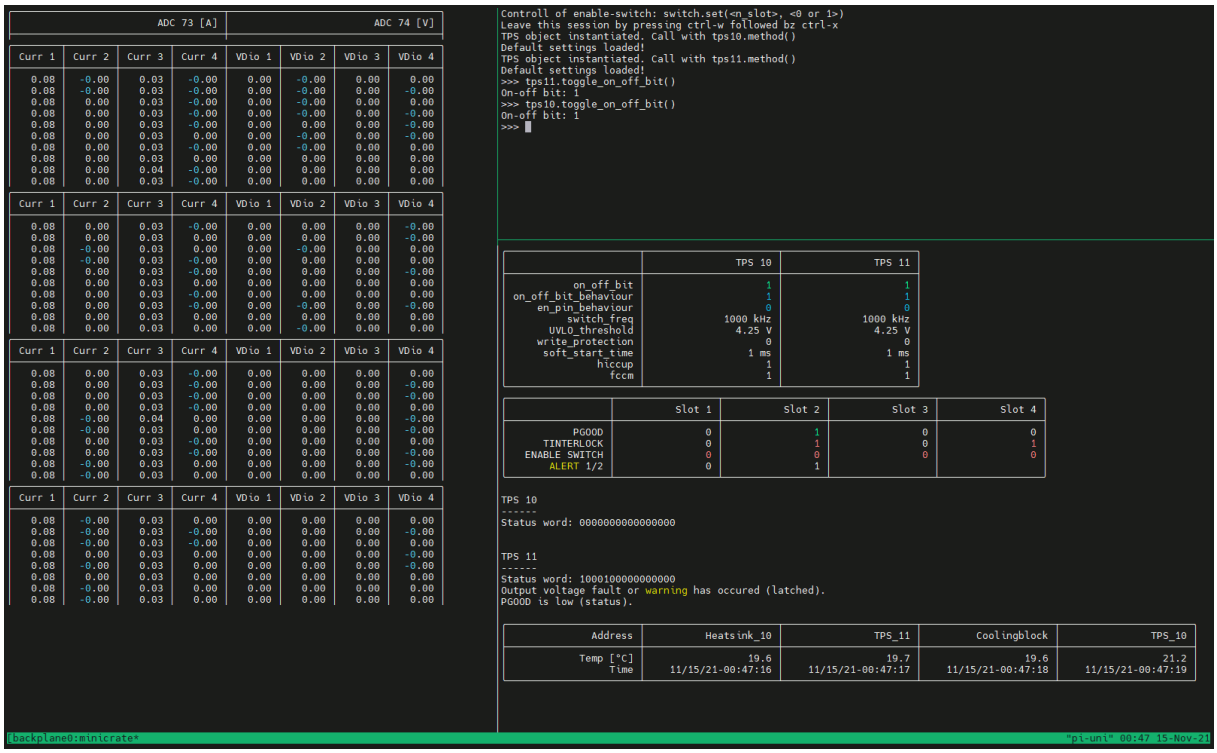

Figure 4.9: Monitor and control interface for the mini-crate.

4.4 Water Cooling Test

The mini-crate was built both for the simple operation of multiple boards and for testing the feasibility of the cooling design on a reduced scale compared to the full-crate. For that reason a cooling loop for the mini-crate was designed and built (for components see table A.2 and table A.1).

4.4.1 Setup

The setup consists of two loops, the primary loop, containing the mini-crate, and a secondary loop, connected to the cooling water of the building. The two loops are in thermal contact using a plate heat exchanger.

The primary loop utilizes a water pump and two temperature sensors, one at the entrance and one at the exit of the heat exchanger. Tubes have an outer diameter of 8 mm and an inner diameter of 6 mm. Tube fittings made out of aluminium are used to reduce corrosion (see section 4.1.4). The water is split between the two cooling blocks using T-pieces (figure 4.10, blue). A differential pressure sensor is used to measure pressure drops

over the crate, an important number to know for the design of the final cooling loop. The flow rate can be measured using an inductive flow meter, but simply disconnecting one hose and measuring water output over time turned out to be more accurate.

The secondary loop runs water from the building. The flow rate can be read and adjusted using valves. Analog thermometers allow the reading of water temperatures. Tubes have an outer diameter of 12 mm and an inner diameter of 10 mm.

Initially, a high flow was chosen for a few minutes until the water temperature stabilised around 13 °C (in the secondary loop), then the flow was reduced to 3.5 $^{L}/_{\text{min}}$.

Multiple temperature sensors (DS18B20 from Maxim Integrated) were glued to different parts of the crate using a custom mount for the TO-92 package (see figure 4.11). Sensors were also glued to the back of the powerboard's PCB opposite of the CSD chip (see figure 4.11). Previous measurements showed this IC to be the hottest location on the board, requiring the most cooling (see section 4.2). Since the IC is connected via thermal vias to the backside's ground plane, mounting the sensor there provides an indirect temperature measurement of the IC without interfering with the heatsink.

Three boards were used (slot 1: board 1b, slot 2: board 10, slot 3: board 11). The boards in slots 1 and 2 were attached to active electronic loads, the board in slot 3 to an ohmic load. They were turned on in the order slot 2 - 3 - 1 and turned off in the order slot 3 - 1 - 2, each adjusted to an output current of around 19.6 A (see figure 4.12).

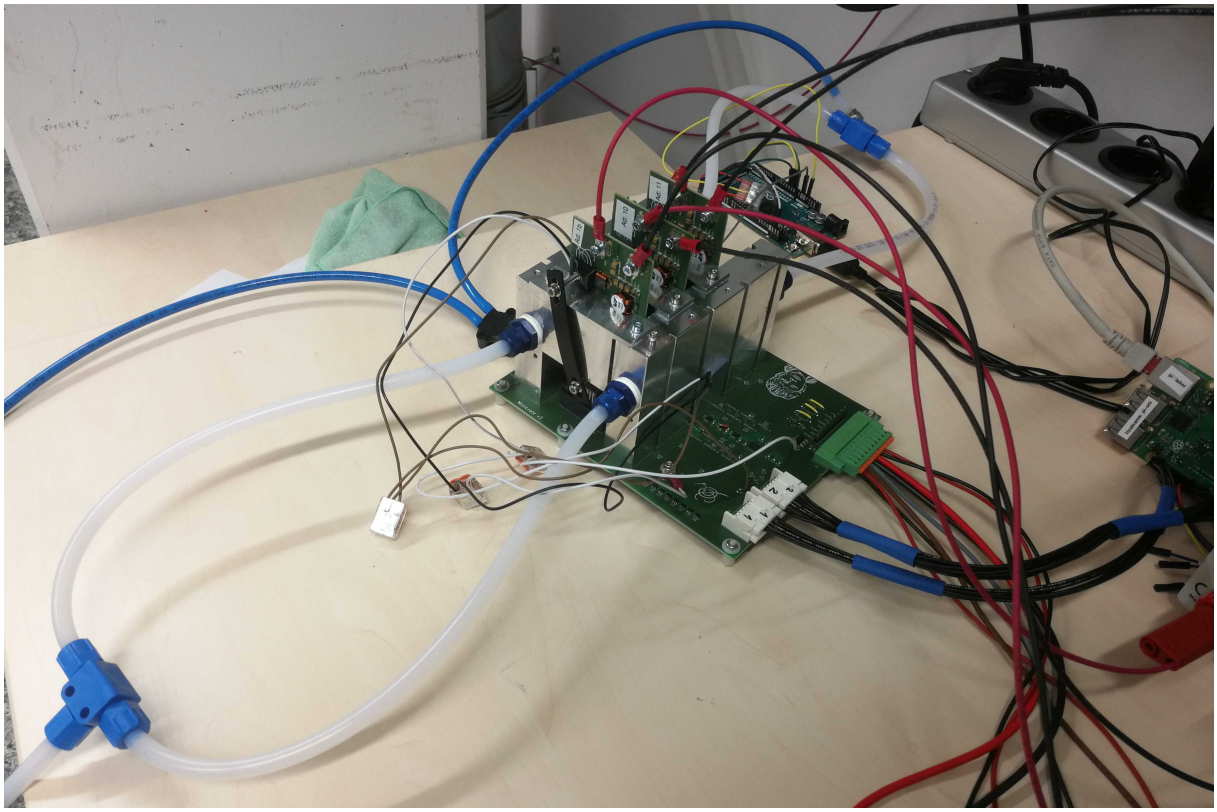

Figure 4.10: Mini-Crate with three boards, temperature sensors, input and output cables, and water cooling connected.

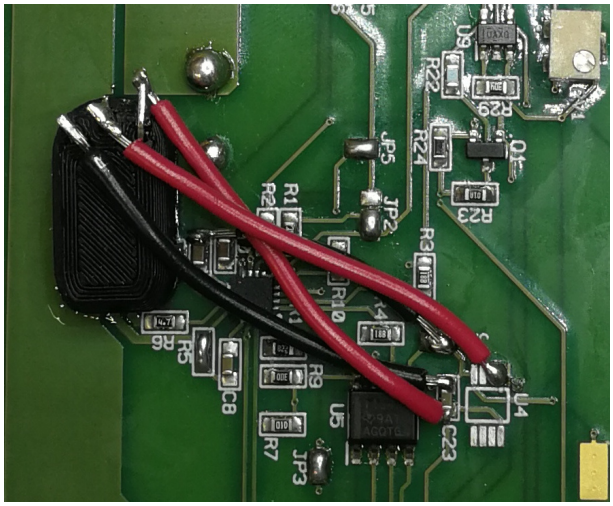

Figure 4.11: Temperature sensor DS18B20 in 3D-printed mount glued to the back of the PCB. In the original design, U4 is a PCB temperature sensor (DS18B20U), which was placed at a thermally uninteresting region of the PCB. Using wires, the position was changed to be opposite of the CSD chip, which generates the most heat.

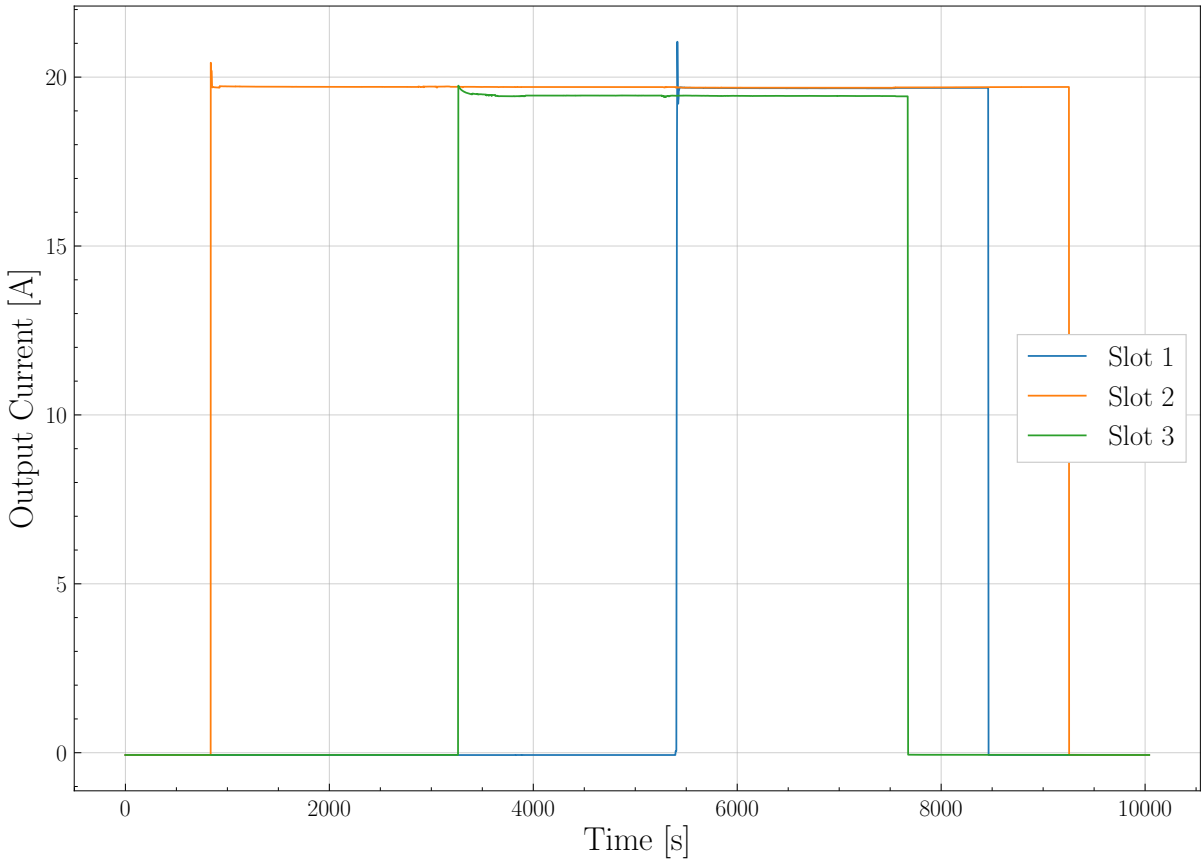

Figure 4.12: Output currents of the three boards measured with the powerboards' onboard shunts.

4.4.2 Measurements and Results

Figure 4.13 shows the water temperatures in the primary loop before and after the heat exchanger. The water temperature after the heat exchange is almost constant, with only a slight increase in the region with more thermal load. This suggests that the cooling power of the heat exchanger is sufficient.

The temperature of the water before the heat exchanger has three plateaus above the base level, corresponding to the thermal load from the three boards. Subtracting the water temperature after the heat exchanger and calculating the mean difference between plateaus, one gets a temperature increase of $\Delta T = (0.16 \pm 0.01)\,°\mathrm{C}$ per plateau. The flow has been measured to be $F = (0.63 \pm 0.02)\,{}^{\mathrm{L}}/_{\mathrm{min}}$. With the specific heat capacity of water ($c = 4184\,{}^{\mathrm{J}}/_{\mathrm{kg \cdot K}}$), one can calculate the thermal power carried away by the water per board:

$$P_{\mathrm{thermal}} = \Delta T \cdot F \cdot c = (7.1 \pm 0.5)\,\mathrm{W}\,. \tag{4.1}$$

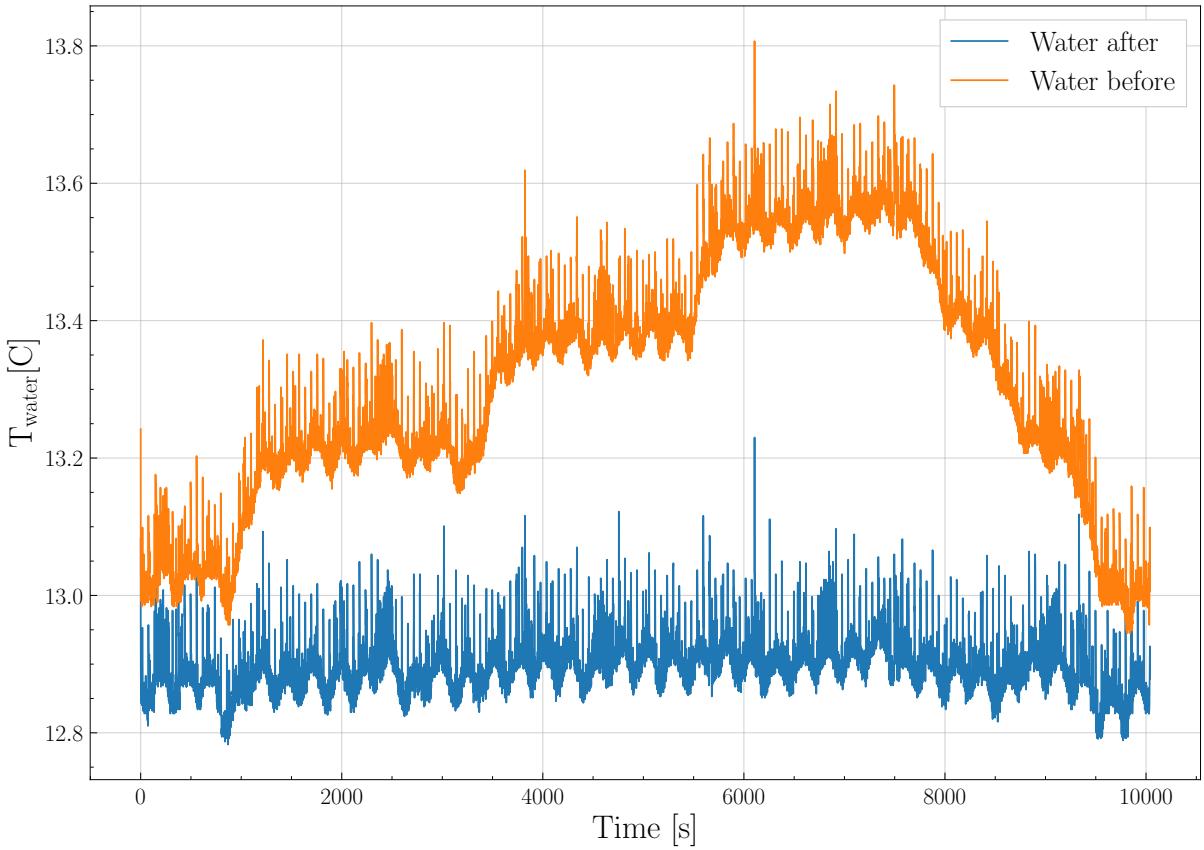

Figure 4.13: Temperature of the water in the primary loop (including mini-crate) before and after entering the heat exchanger.

This can be compared to the dissipated power per board, with losses in cables subtracted (see section 5.4):

$$P_{\text{dissipated}} = P_{\text{in (power supply)}} - P_{\text{input cables}} - P_{\text{out (board)}} = (8.8 \pm 0.3)\,\text{W}\,. \tag{4.2}$$

The measurement implies that around $(81 \pm 6)\,\%$ of the generated heat is successfully extracted by the cooling loop, while $(19 \pm 6)\,\%$ goes to the surrounding. The error on the result is quite big, but overall seems to be in a plausible region.

The temperature of the boards, essentially the temperature of the IC containing the MOSFETs (CSD), can be seen in figure 4.14. Turning the output on increases the temperature of the board by around $(29 \pm 2)\,^{\circ}\text{C}$. But also the turn-on of other boards increases the temperature. For direct neighbours the temperature increase is slightly higher with $(1.4 \pm 0.3)\,^{\circ}\text{C}$ versus $(0.7 \pm 0.2)\,^{\circ}\text{C}$ for non-direct neighbours in the crate. This is plausible, since boards are stacked close together in the crate, and heat exchange is non-negligible.

Using this data, one can extrapolate temperatures to a fully occupied crate, with boards running at 20 A. Since most cooling blocks are used from both sides, a total of 8 boards

per block are assumed. Each board has two direct neighbours and 5 non-direct neighbours. Assuming a maximum increase of 2 °C per board, a turn-on increase of 31 °C and a base temperature of 14 °C, a temperatures of around 59 °C on the backside of the board is reached.

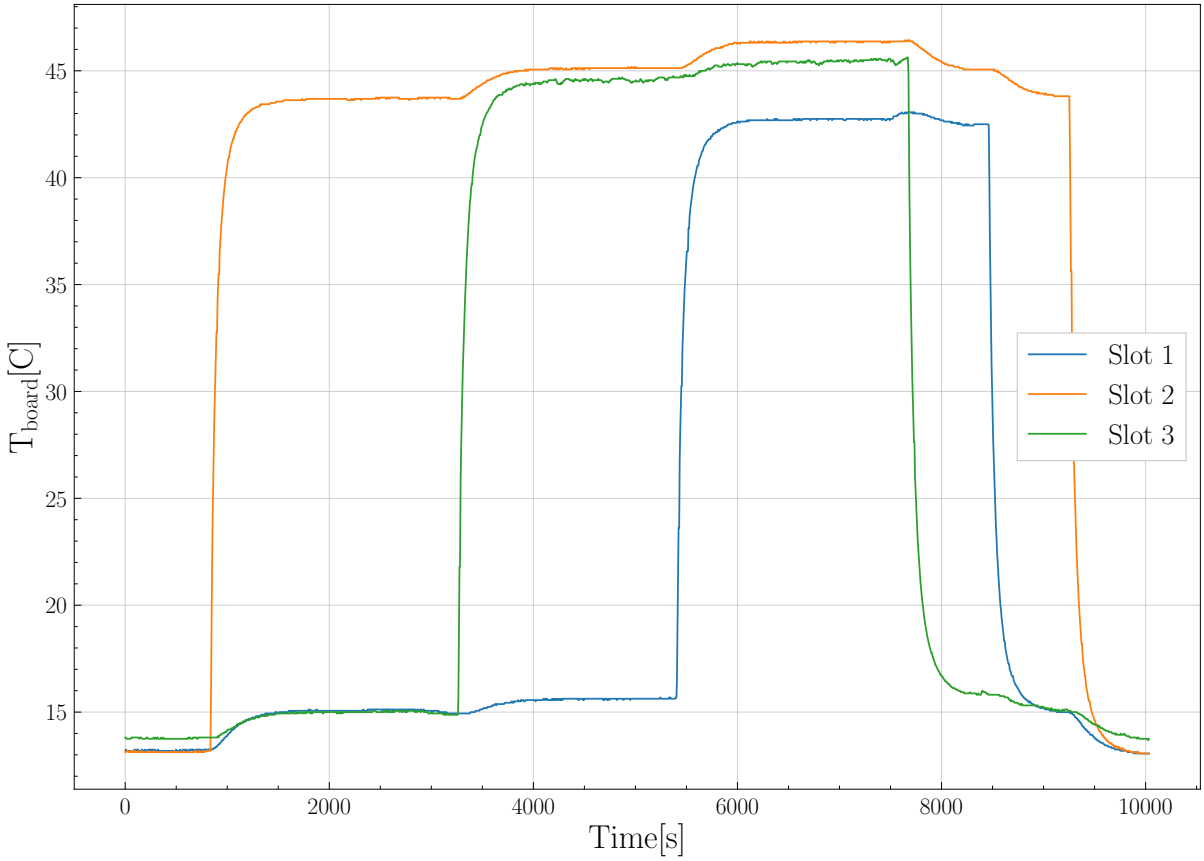

Figure 4.14: Temperatures measured on the backsides of the powerboards.

4.4.3 Conclusion

With version 2 of the heatsink design, around (81 ± 6) % of the generated heat is extracted by the water cooling. Two measures to improve this number are taken in version 3 of the design.

Firstly, a proper mounting mechanism is used. Version 3 of the powerboards is going to include mounting holes, allowing the fastening of the heatsink from the backside of the board with screws (see figure 4.15). This way, an even and high pressure can be exerted, compressing the thermal pad evenly and resulting in better thermal contact.

Secondly, the CSD chip's main thermal path, over the ground plane, is used. The data sheet gives a thermal resistance of 20 °C/w between the junction and the top of the package, and a thermal resistance of only 2 °C/w between the junction and thermal pad on the

bottom. While the thermal pad is soldered to the powerboard's ground plane, version 2 of the heatsink does not actively cool the plane. For version 3 of the powerboard, the plane next to the CSD chip will be exposed, allowing the coupling of the heatsink to it. This should be done using a thermal pad, to compensate for any imperfections, but also to electrically isolate the heatsink from the ground plane. Dissipating heat outside of the cooling loop will always happen. For small amounts of heat, this is not problematic, since the inside of the magnet will receive some cooling as a side-effect of cooling the MuPix sensors using helium. Still, maximizing the efficiency of the water cooling is preferred.

The data sheet of the CSD chip [12] gives an operating junction temperature T_J of $-55\,^\circ\mathrm{C}$ to $150\,^\circ\mathrm{C}$, with a recommended maximum temperature of $125\,^\circ\mathrm{C}$. The junction temperature can in theory be estimated using thermal resistances from the data sheet (a model similar to resistances in electric circuits with temperature drops instead of voltage drops)

$$T_J = T_{\mathrm{PCB\,(bottom)}} + R_{\mathrm{thermal\,pad}} \cdot P_{\mathrm{thermal\,pad}} + R_{\mathrm{PCB}} \cdot P_{\mathrm{PCB}}\,, \tag{4.3}$$

with R_{PCB} and P_{PCB} the thermal restistance and power going through the PCB, respectively. Since many quantities are unknown, only a very rough estimate can be made. Assuming $T_{\mathrm{PCB\,(bottom)}} = (59 \pm 2)\,^\circ\mathrm{C}$, $R_{\mathrm{thermal\,pad}} = 2\,^\circ\mathrm{C/W}$, $P_{\mathrm{thermal\,pad}} = (9 \pm 3)\,\mathrm{W}$, $R_{\mathrm{PCB}} = (8 \pm 5)\,^\circ\mathrm{C/W}$, $P_{\mathrm{PCB}} = (2 \pm 2)\,\mathrm{W}$, the junction temperature calculates to be $(93 \pm 20)\,^\circ\mathrm{C}$, below the maximum specifications. This should be taken with a grain of salt, and with powerboard version 3 available, a long-term study should be run to get a maximum current over time specification.

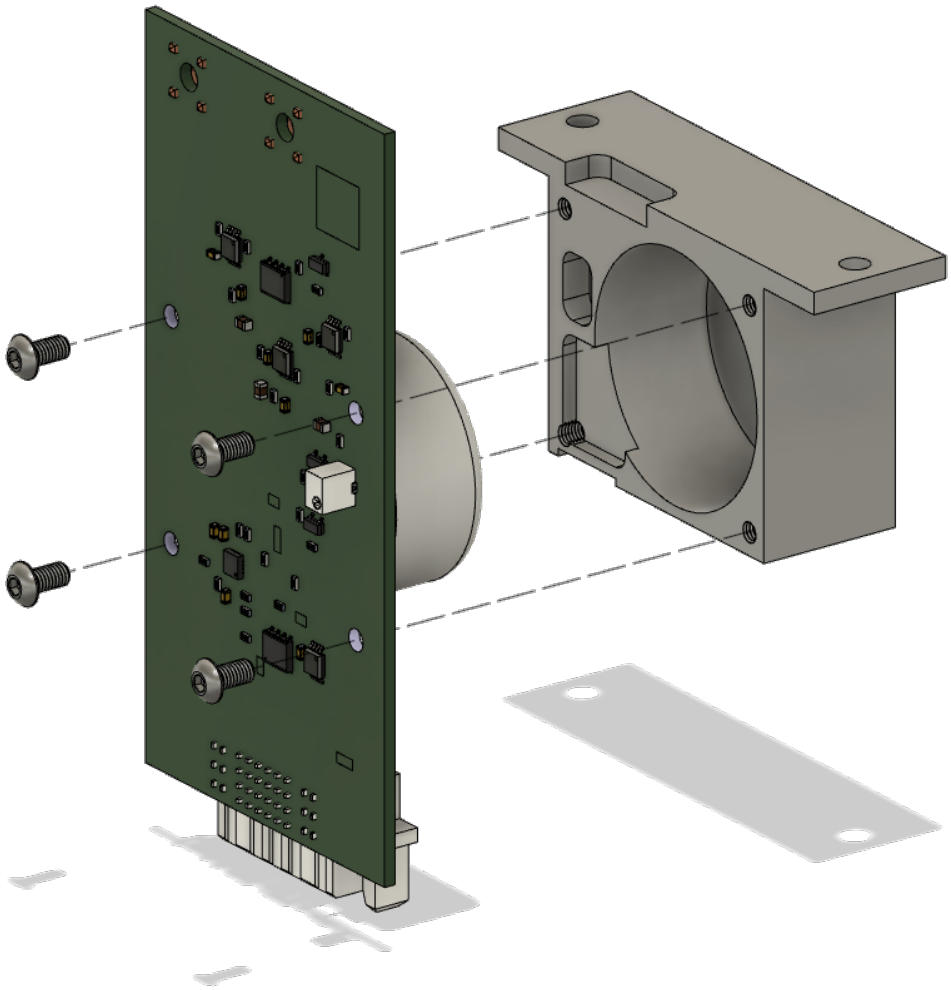

Figure 4.15: Concept for the mechanism to mount the heatsink onto powerboard v3 using four screws. Component placement needs to be change to accommodate space for mounting holes (CAD model of board v2 is used here).

Chapter 5

Powerboard v2

While version 1 of the powerboard was focused on the buck converter itself, version 2 implements many auxiliary features required in the experiment. A batch of 10 boards was produced at the end of 2020. The schematic can be found in appendix B and an introduction to the Mu3e buck converter in section 3.4.

5.1 Output Filter Optimization

The powerboard v2 uses a secondary LC output filter to further attenuate high frequency noise, especially at the switching frequency ($f_{\mathrm{sw}} = 1\,\mathrm{MHz}$) (see figure 5.1). Low-pass LC filters have been analysed in section 3.3. With the initially chosen values ($L_2 = 46.5\,\mathrm{nH}$, $C_{2\mathrm{nd}} = C_{11} \parallel C_{14} \parallel \ldots \parallel C_{19} = 501.1\,\mu\mathrm{F}$, $R_{\mathrm{damping}} = R_{13} = 3.3\,\Omega$) the board enters oscillations at various operational points. Once entered, the board has to be turned off and on to reach a stable state again. The oscillations have peak-to-peak voltages around 65 mV and frequencies in the range from 100 kHz to 200 kHz, depending on the board, switching frequency and output current (see figure 5.2).

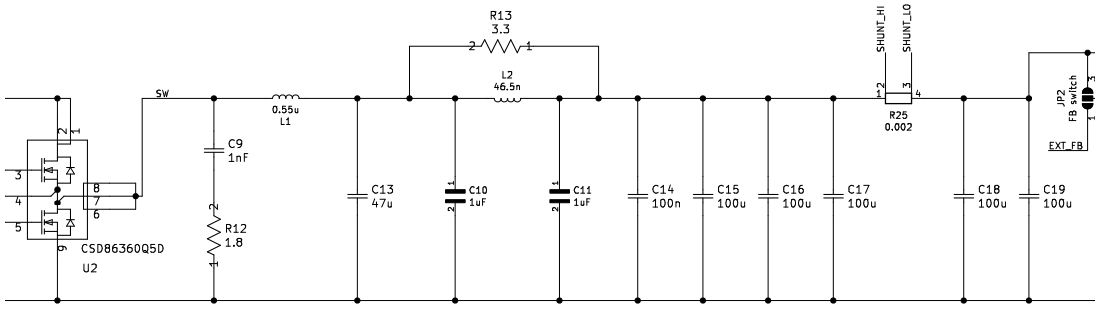

Figure 5.1: Output stage of the powerboard v2. The CSD chip contains the synchronous MOSFETs. C_9 and R_{12} are the snubber circuit. L_1 and C_{13} are the main passive buck converter components (first output filter). L_2 and following capacitances are the second output filter. R_{13} is used to dampen the resonance created by L_2.

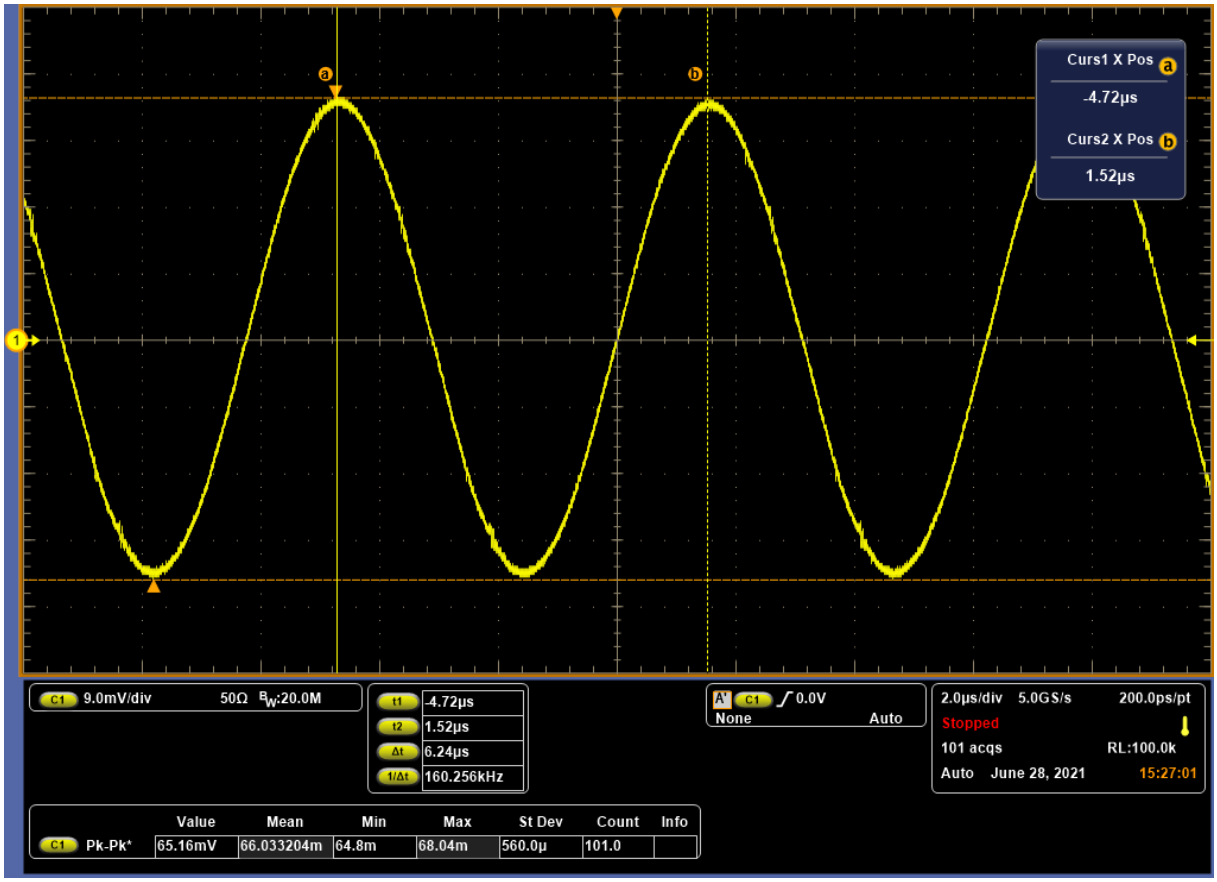

Figure 5.2: Oscillating of the output voltage at 160 kHz with $R_{\text{damping}} = 0.2\,\Omega$, $L_2 = 46.5\,\text{nH}$, $C_{2nd} = 501.1\,\mu\text{F}$, $f_{\text{sw}} = 625\,\text{kHz}$.

To understand the effect of the filter and its components better, the output stage of the board has been simulated using LTspice® (see figure 5.3). Since parasitic elements are not included and not the whole control loop is simulated (TPS and CSD chip are left out), quantitative results might differ from measurements, but qualitative knowledge about the effects of different components can be gained.

The simulation uses two different values for L_2, 46.5 nH and 22 nH, and for each inductance the damping resistor R_{damping} was varied between 0 Ω and 3.3 Ω. The results of the AC sweep can be seen in figure 5.4. As intended, with the second filter stage one can reach an additional attenuation at the switching frequency (around −30 dB/−97 %). This comes at the price of adding a second resonance to the spectrum. For lower inductances, the resonance sits at a higher frequency and is significantly shifted down in magnitude. Slightly unintuitive is the behaviour of the damping resistor. Low values decrease the magnitude of the resonance, while higher values increase peaking. This is an effect of using parallel damping, where replacing the resistor with a short would in theory result in no current going through the coil and no resonance being created. Increasing resistance, in turn, leads to higher currents through the coil and a more distinct resonance peak. The simu-

lation shows that a damping resistor value of 3.3 Ω, as initially chosen, results in a sharp resonance. The calculated resonance frequency for the 46.5 nH inductor is 115 kHz. This is slightly off from the measurements and is probably related to parasitics, for example, inductances from the PCB traces, which are also in the order of nH.

Having a deeper understanding of the filter, measurements with different valued components were done in the lab. Reducing the damping resistor to 0.2 Ω (with $L_2 = 46.5\,\text{nH}$) results in notably fewer oscillations, thus agreeing with the simulation. Increasing the inductance to 79 nH leads to overall more oscillations. In the end, it was decided to use the 22 nH inductor because of the downward shift in magnitude of the resonance (around −6 dB). At 1 MHz the attenuation is only around 6 dB worse compared to the 46.5 nH inductor. The damping resistor was increased to 0.5 Ω since the resonance does not need to be damped that strongly anymore (see figure 5.5). With these values, oscillations are only entered at $f_{\text{sw}} = 325\,\text{kHz}$, far below our nominal switching frequency of 1 MHz.

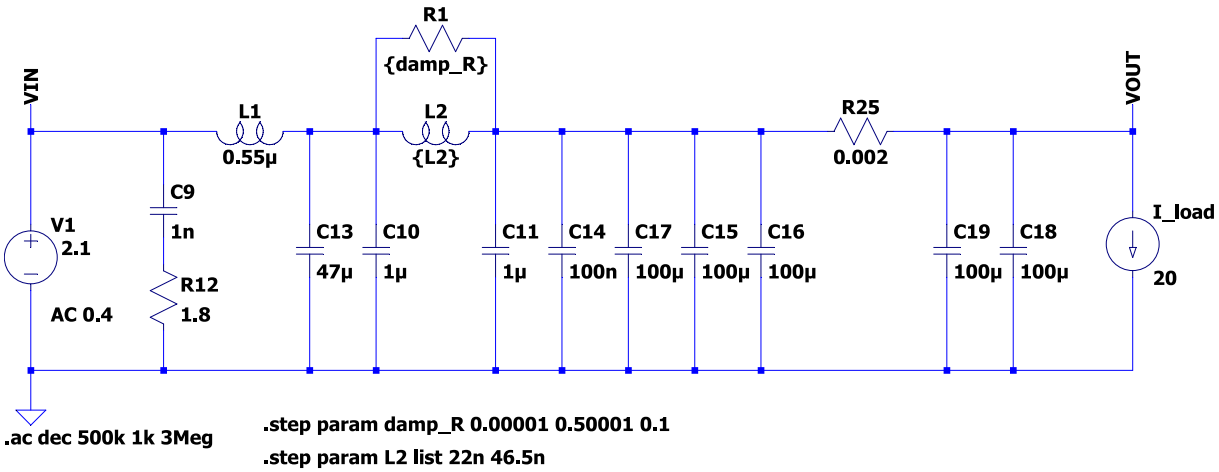

Figure 5.3: Circuit used to simulate the output stage of the powerboard in LTspice®. L2 and damp_R are varied. An AC sweep is done from 1 kHz to 3 MHz.

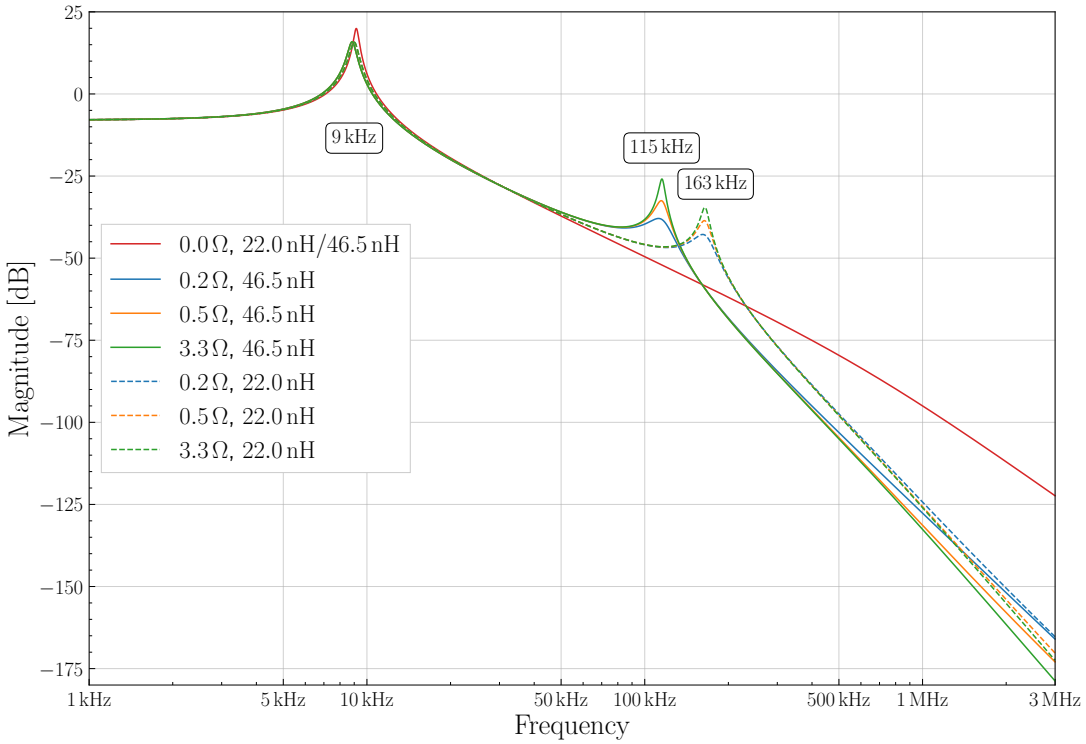

(a) Full range

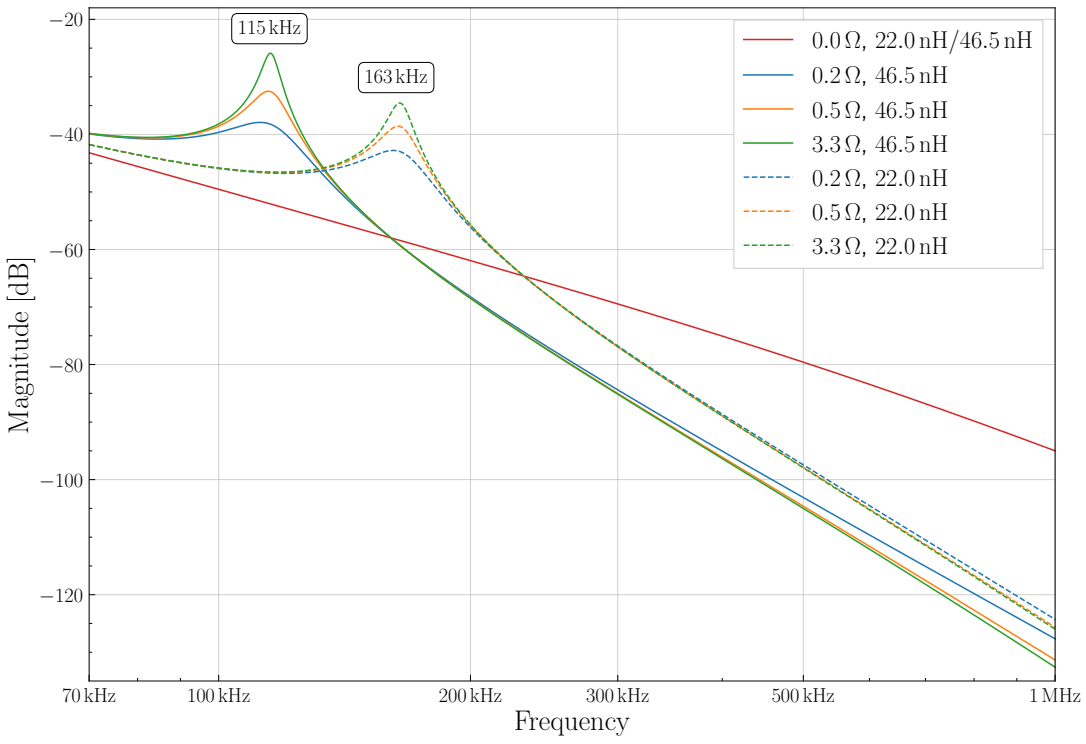

(b) Zoomed

Figure 5.4: Bode magnitude plot of the AC sweep.

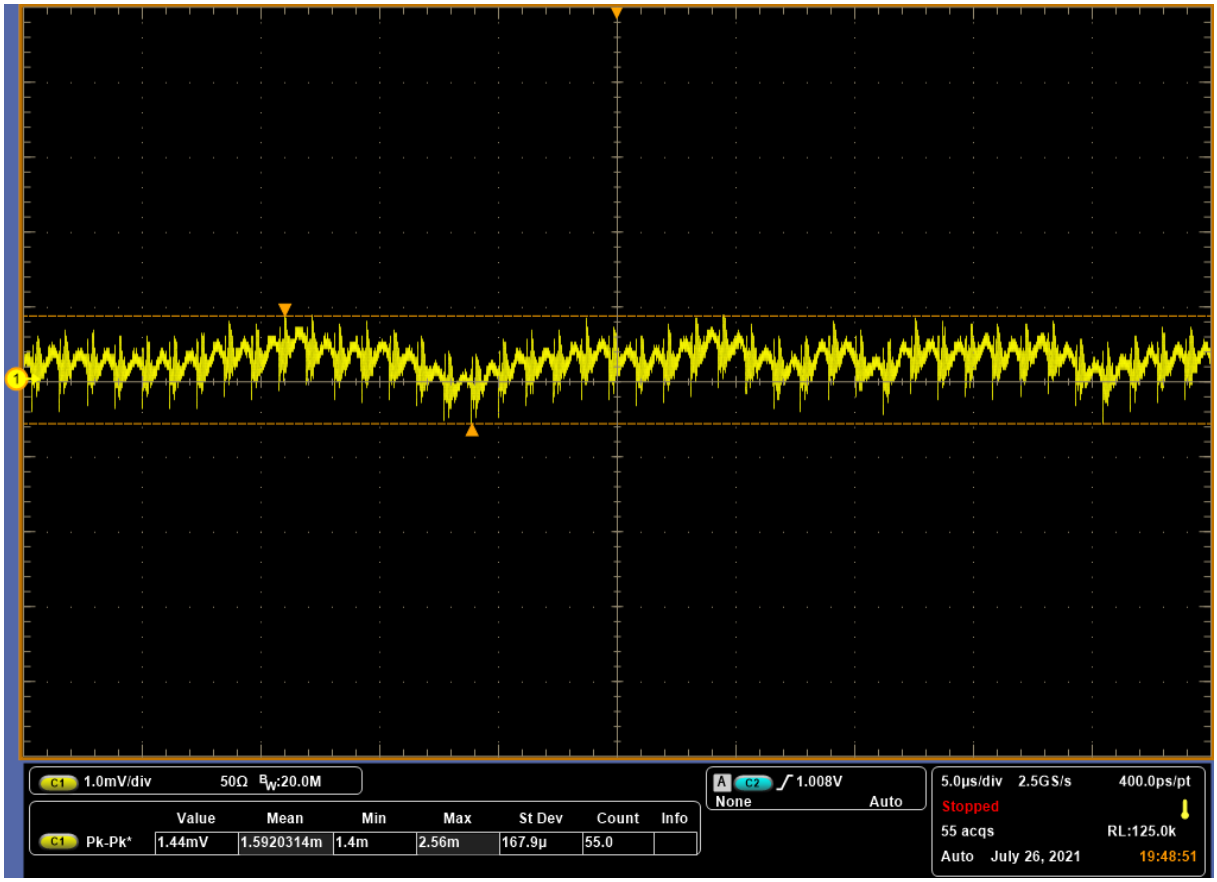

Figure 5.5: Scope trace of the output voltage with the optimized output filter values $R_{\text{damping}} = 0.5\,\Omega$, $L_2 = 22\,\text{nH}$, $C_{2nd} = 501.1\,\mu\text{F}$, $f_{\text{sw}} = 1\,\text{MHz}$. LF ripple is around 1.44 mV peak-to-peak.

5.2 Helium Atmosphere Tests

In the experiment, the powerboards will be operated inside the magnet. This means they have to be functional in a helium environment. To test this, one powerboard (board 11) has been put inside a sealed metal box with attached in- and outgoing tubing (see figure 5.6). A constant flow of around 0.67 liter/min of helium has been run through the box. The board was operated at an output of 6 A at 2.1 V for 26 h. Functionally was tested by communicating with the controller chip (TPS) and checking the output power regularly. The powerboard was responsive and fully functional at all times.

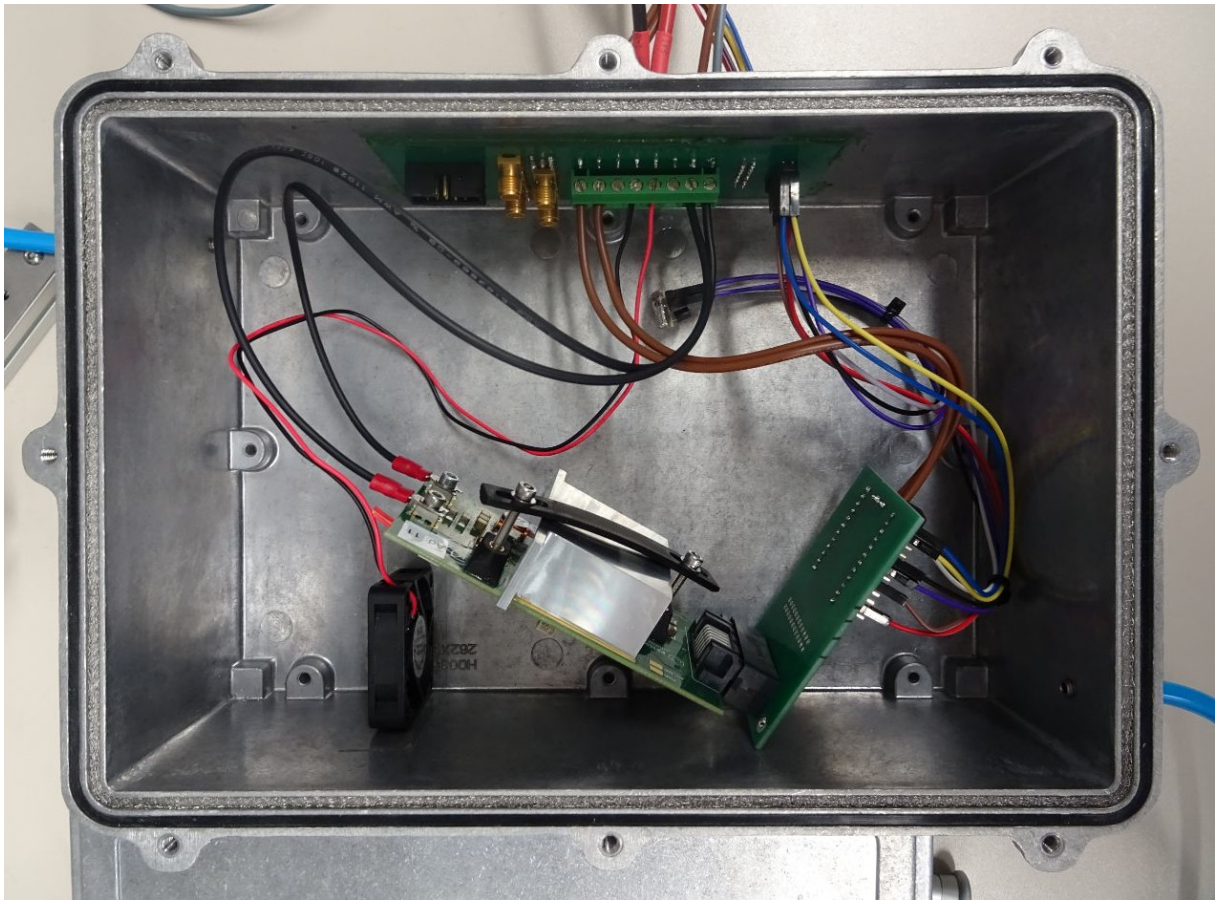

Figure 5.6: Helium atmosphere test setup. The box can be sealed and a fan creates circulation. Helium enters through the tube in the top left corner and exits through the bottom right. Cables are run through a feed-through.

5.3 Noise Figures

The figure of merit for every DC-DC converter is the noise level. To power the MuPix sensors reliable, a peak-to-peak noise of less than 10 mV is aimed for.

Noise from switching converters can be divided into two categories, low-frequency (LF) ripple and high-frequency (HF) noise (see section 3.2). LF ripple is typically shown in data sheets and is measured using a 20 MHz bandwidth. Measuring the HF noise is more difficult. Using a long leaded ground connection on the oscilloscope probe will result in a large ground loop. This can act as an antenna and pick up additional noise, which is not present on the output of the converter. Additionally, the measurement is very sensitive to the exact positions and angles of cables and probes. To minimize these effects, a x1 probe has been soldered using a coaxial cable and a 0.1 µF coupling capactitor (see figure 5.7 and figure 5.8). Normally x1 probes are avoided since they capacitively load the circuit (in the order of 100 pF), but since the output capacitance of the converter is significantly higher,

this does not pose a problem here. The coaxial cable was crimped to be able to easily attach it to the output connectors and a 50 Ω termination is chosen on the oscilloscope. The x1 probe also allows for a sub 1 mV resolution and thus exact measurement of the LF ripple magnitude.

Measurements can be seen in figure 5.9 and figure 5.10. They were taken using the optimized second stage output filter (22 nH and 0.5 Ω, see section 5.1) with heatsink attached and plugged in the mini-crate at 20 A.

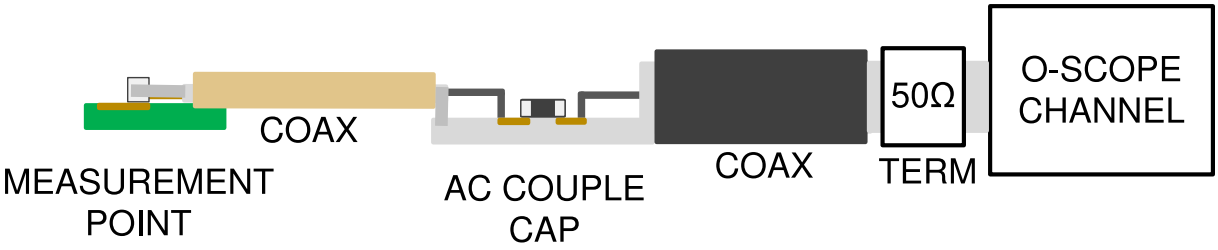

Figure 5.7: Structure of a self-made x1 probe [10].

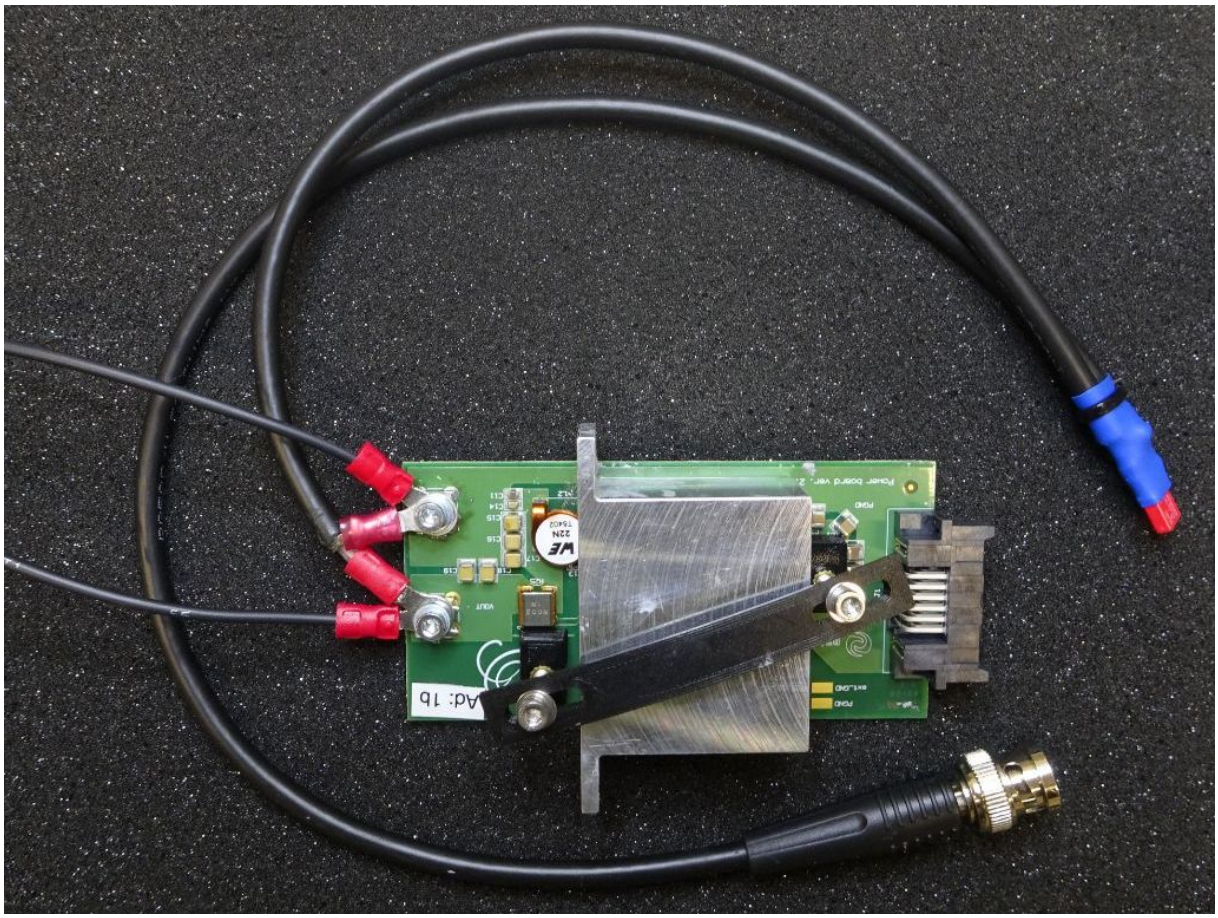

Figure 5.8: Powerboard with output cables and x1 probe connected.

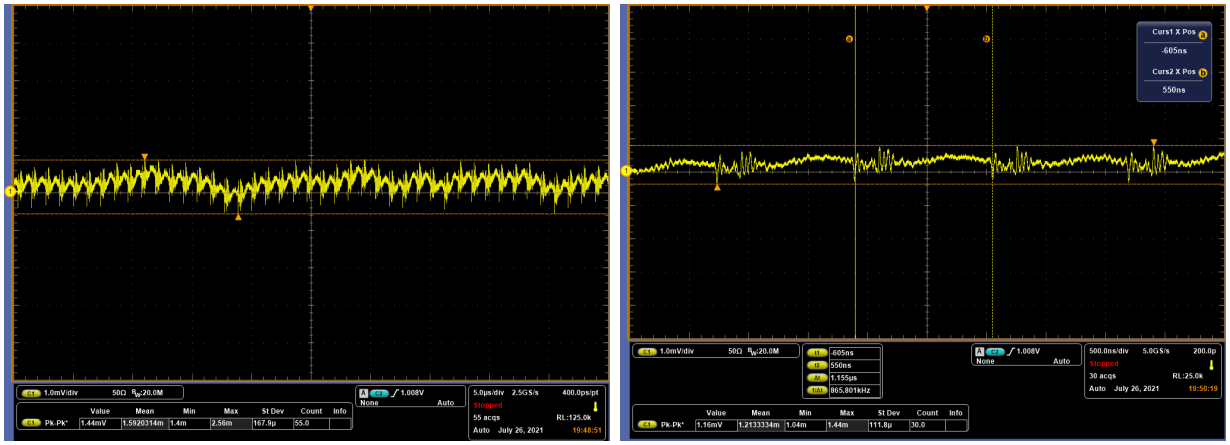

(a) Far (b) Zoomed

Figure 5.9: LF ripple at the switching frequency with peak-to-peak voltage approx. 1.5 mV (20 MHz bandwidth).

(a) Far (b) Zoomed

Figure 5.10: HF noise with peak-to-peak voltage approx. 60 mV (250 MHz bandwidth).

LF ripple has a peak-to-peak voltage of around 1.5 mV with the optimized filter values. This is far below the upper limit of 10 mV and does not need to be reduced further.

The HF noise in figure 5.10 shows two peaks, one for the high-side switching and one for the low-side switching (see section 3.2.2). It is, as expected, with a peak-to-peak voltage of around 60 mV, a lot higher than the LF ripple. This value is not absolute, even with the x1 probe the measurement is very dependent on the positions of the cables. It is also not clear how sensitive the detector modules are to HF noise and how much is transmitted over the cables. Still, the result hints that the PCB layout and components can be optimized to reduce HF noise. In contrast to reducing LF ripple, these changes are very small and

subtle, and a lot of care needs to be taken. For example, the area of the high dI/dt-loop needs to be kept as small as possible (see section 3.2.2). Also, the values of the RC snubber circuit (see figure 5.1, C9 and R12) and the boost resistor (see appendix B, R5) can be fine-tuned to the noise frequency spectrum. When choosing filtering capacitors, packages with small ESL (equivalent series inductance) need to be chosen, since, at high frequencies, ESL is dominating the impedance of a capacitor. ESL depends mostly on the package size and type, not capacitance (see figure 5.11). A 0805 capacitor will have a higher ESL than a 0402 capacitor, independent of capacitance. To further significantly reduce ESL, a low inductance chip capacitor (LICC) can be used, for example, 0306.

The noise figures also show that the switching frequency is with 866 kHz significantly below the configured 1 MHz. This is unexpected and needs further investigation.

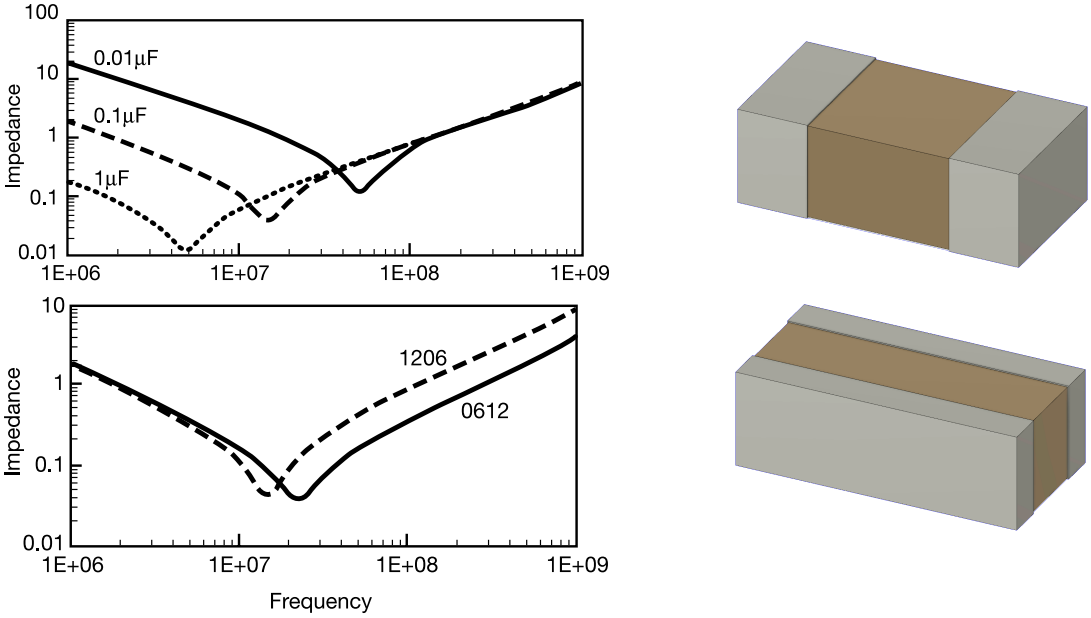

Figure 5.11: Left side: Impedances for different capacitor values and packages [23]. Right side: Model of MLCC (top) and LICC (bottom) capacitor.

5.4 Efficiency Calculation

Knowing the efficiency of the powerboard is not only important for calculating the needed input power, but also for the cooling system, since the lost power is converted to heat, predominantly in the switching chip and output filter, and needs to be transported away. The efficiency is calculated using the following formula:

$$\eta = \frac{P_{out}}{P_{in}} = \frac{U_{out} \cdot I_{out}}{U_{in} \cdot I_{in}} \,. \tag{5.1}$$

To exclude any losses in cables and get the efficiency of only the converter, U_{in} and U_{out} need to be measured directly at the board's input and output respectively. The efficiency

is also dependent on the operation point of the converter. For this reason, scans of input voltage (see figure 5.13) and output current (see figure 5.12) have been done.

The output voltage was measured directly at the board, meaning losses in the output cables are automatically excluded. Since the input voltage was read from the power supply (HAMEG HMP4040), input cables losses are initially included. To subtract these losses, the input cables resistance was calculated by measuring voltage drop across input cables at different currents, resulting in $R_{\text{cables}} = (68.4 \pm 0.4)\,\text{m}\Omega$. The error on the efficiency is mostly dominated by the uncertainty on the output current. The current was partially measured using a multimeter (Uni-T UT161E) and for currents greater 15 A, above the range of the multimeter, using the powerboard's onboard current measurement. Version 2 of the powerboard has problems with the reference ground of the ADCs being shifted, probably since they are located on the backplane instead of on the board itself, giving these measurements an accuracy of only around ±0.2 A.

Looking at the region of interest (7 A to 21 A), an efficiency of 83 % to 87 % is reached for the nominal input voltage of 20 V. Two currents, (9 A and 21 A), where most converters are going to be operated, have been marked and exact values can be taken from table 5.1.

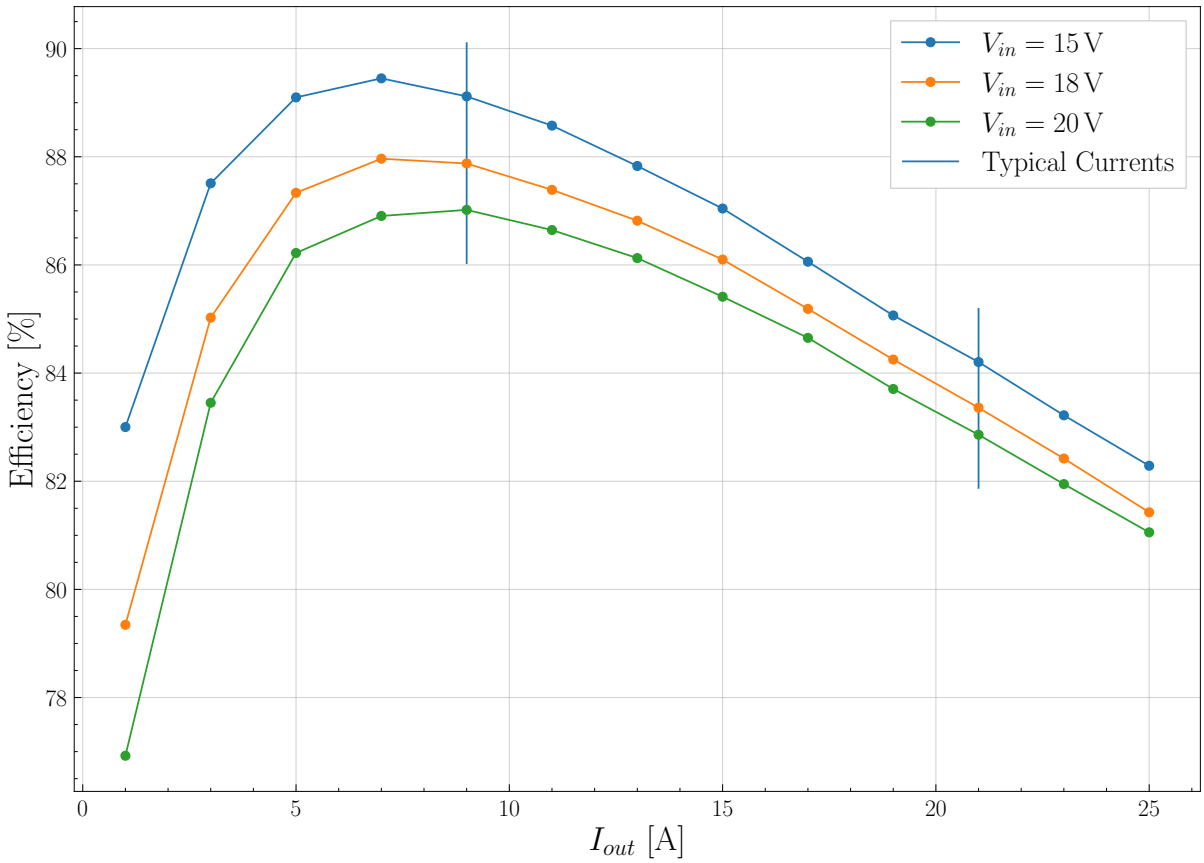

Figure 5.12: Efficiency for different input voltages at varying output currents. Errorbars have been omitted for clarity. Typical currents have been marked (see table 2.4).

Table 5.1: Efficiencies for two typical currents (see table 2.4).

Input Voltage	Output Current	
	9 A	21 A
15 V	(89.1 ± 0.6) %	(84.2 ± 0.5) %
18 V	(87.9 ± 0.5) %	(83.4 ± 0.5) %
20 V	(87.0 ± 0.5) %	(82.9 ± 0.5) %

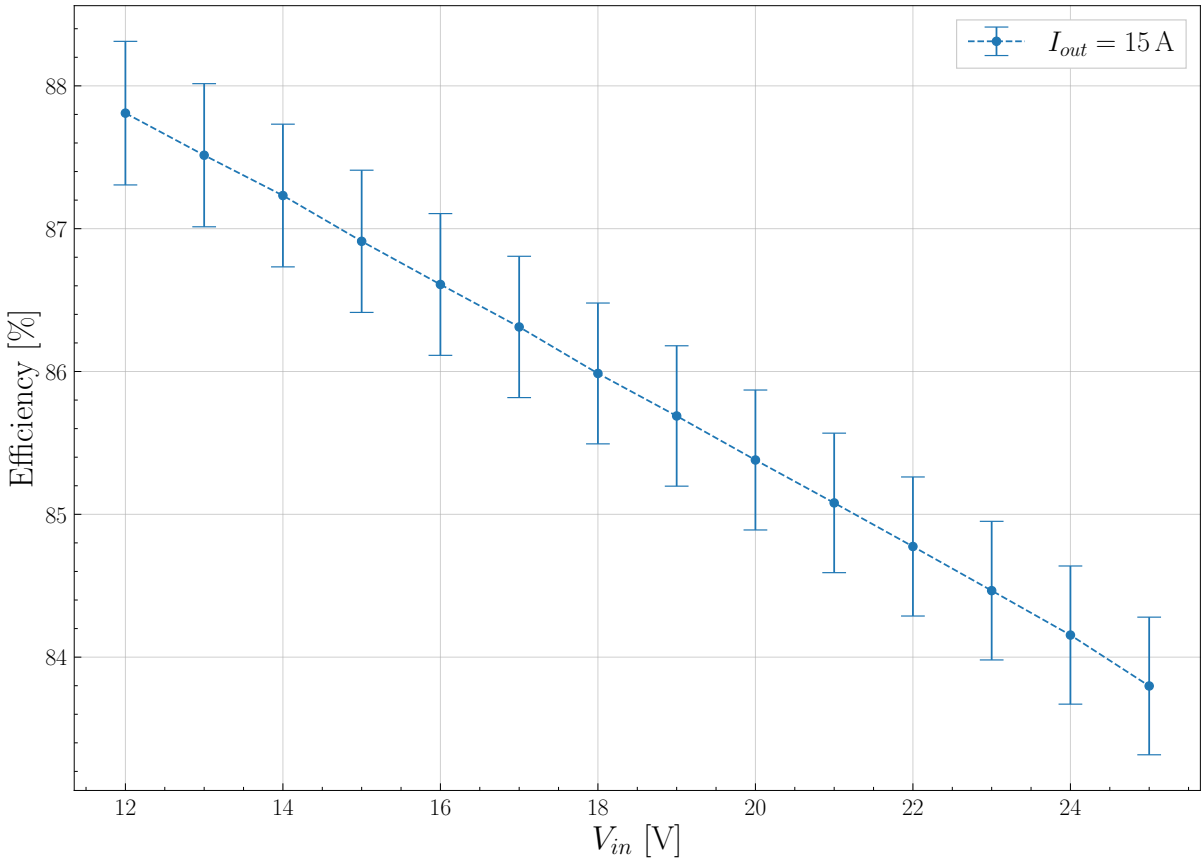

Figure 5.13: Efficiencies for varying input voltage at 15 A output current.

Figure 5.12 and figure 5.13 clearly shows that a lower input voltage will result in higher efficiencies. This is expected since the voltage step is getting smaller. But lower input voltage also means higher input current, resulting in greater losses in the input cables and heating of the cables. Optimization of this problem can be done in later stages of the development when the exact currents and cable resistances are better known.

The measured efficiency can be compared with other buck converters, for example with the FEAST from CERN, which reaches efficiencies of around 80 % [24, p. 7]. Obviously

size and use cases are quite different, but operation conditions are similar. An efficiency of $(82.9 \pm 0.5)\,\%$ at 21 A for the powerboards seems to be a good achievement. Assuming an efficiency of 75 % when calculating cooling power, wire gauges, etc., leaves a moderate safety margin of greater 6 %.

Chapter 6

Powerboard v3

Version 3 of the powerboard is under design and is going to be produced in the first quarter of 2022. The schematic can be found in appendix C.2.

6.1 Voltage Regulation

The powerboards will be connected to the detectors using cables, copper pipes, and flex prints. Even though the resistance is minimal, at high current it will lead to a voltage drop which can not be neglected (see figure 6.1). Since the exact current of a power partition depends on how many of the ASICS are enabled, this drop has to be compensated dynamically. For a MuPix central layer 4 partition running at 20 A it is roughly estimated to be around 564 mV, but the exact value has to be measured during the Mu3e Phase I experiment.

The output of the powerboard can be regulated by the TPS buck converter control chip by up to ±9 %. This needs to be done digitally, by measuring the voltage at the load and then sending the appropriate PMBus commands. The powerboard v2 has a sensing circuit, with which it can measure the voltage at the load using two sense wires (since both GND and VDD will be shifted), a differential amplifier, and an ADC channel (see figure 6.2). Since the regulation via the control chip is not enough (for 2 V one can regulate only ±180 mV), different methods to keep the voltage at the load constant, independent of current, have been tested.

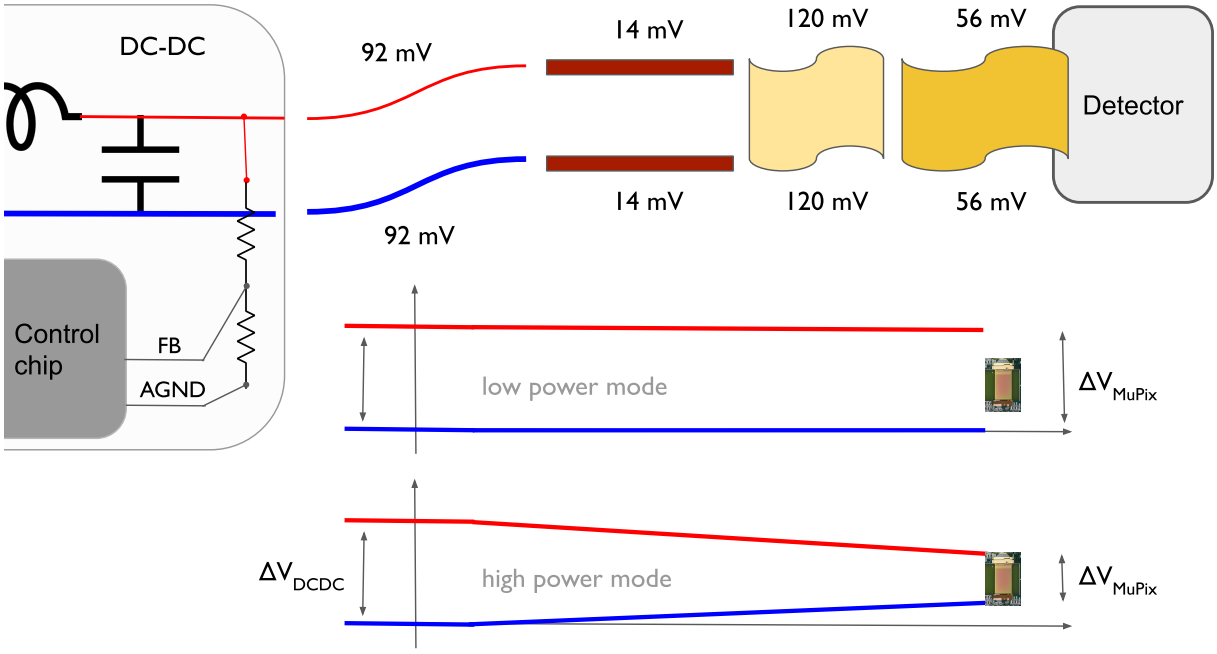

Figure 6.1: Estimated voltage drops of 564 mV for 20 A between converter and MuPix sensor in central layer 4. When disabled, the MuPix sensor is in the low power mode, drawing little current, resulting in a small voltage drop. When enabled, it enters the high power mode, drawing high current, leading to a large voltage drop [6].

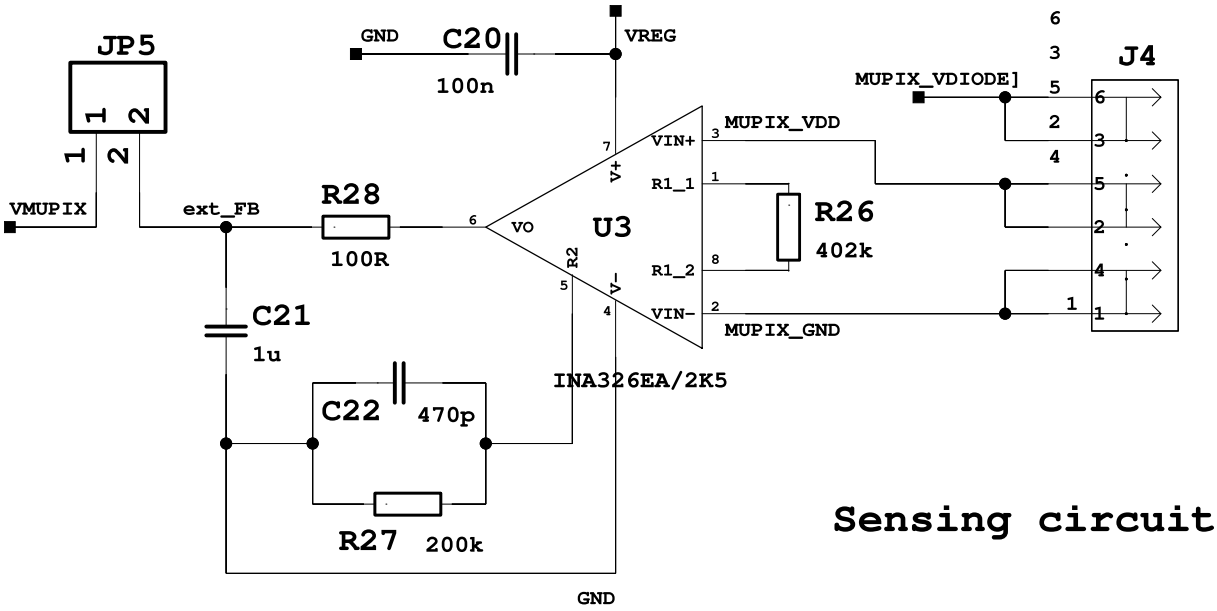

Figure 6.2: Sensing circuit of powerboard v2. MUPIX_GND and MUPIX_VDD are the sense wires. A differential amplifier with gain one and output filter gives the voltage at the load.

6.1.1 Remote Sense Wires using OpAmp

The powerboard v2 default configuration connects the feedback pin to a voltage divider at the end of the output stage. This is a standard approach to voltage feedback. With jumpers, this can be changed and the feedback pin can be connected to the output of the sensing circuit (see figure 6.2), essentially taking the voltage across the load as feedback.

This simple approach did not work in the lab. The sensing circuit introduces a huge time delay ($\mathcal{O}(100\,\text{ms})$) between the output voltage and feedback voltage (see figure 6.3). As a result, the output voltage overshoots, and the overvoltage protection disables the converter. The time delay is caused by the strong RC filter on the output of the differential amplifier. With $100\,\Omega$ and $1\,\mu\text{F}$, it has a cutoff frequency of around $1.6\,\text{kHz}$. Choosing a weaker filter also did not result in a stable feedback loop, since fast oscillations were still not followed by the feedback. Further measurements showed that the feedback loop, in general, is very sensitive to for example delays, noise, oscillations, and should be kept as small and fast as possible. Taking feedback only from the sense wires is for that reason not feasible.

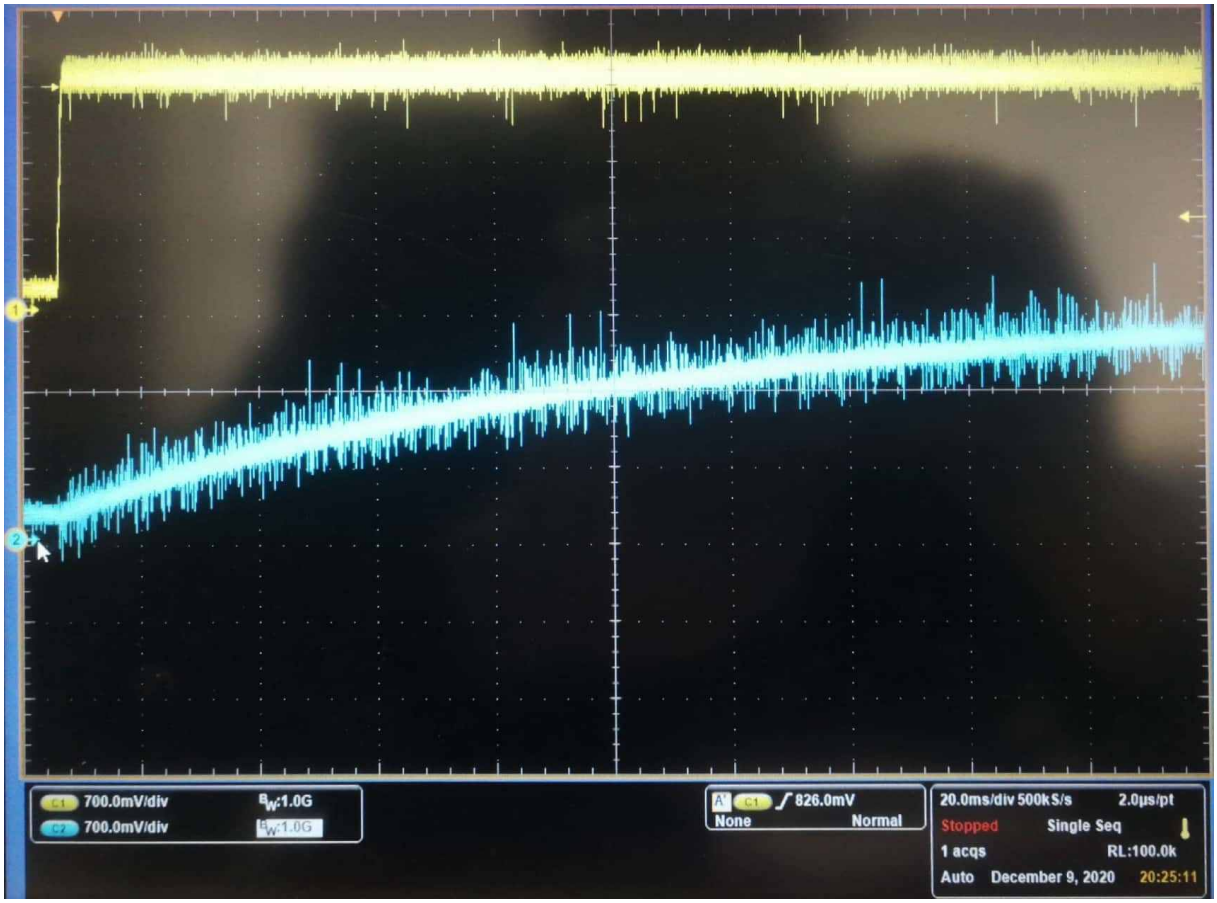

Figure 6.3: Scope trace of output voltage (yellow) and ext_FB node voltage (blue) (see figure 6.2). To observe the delay with a stable output voltage, feedback is taken from the standard voltage divider.

6.1.2 Cable Voltage Drop Compensation

Cable voltage drop compensation increases the output voltage of the converter linearly with the output current, thus if properly calibrated, compensating voltage drops across cables. The slope of increase can be matched to a known cable resistance by choosing the appropriate resistor values. A simplified schematic of the application can be seen in figure 6.4 and a detailed description can be found in Texas Instruments' application report SLVA657 [25].

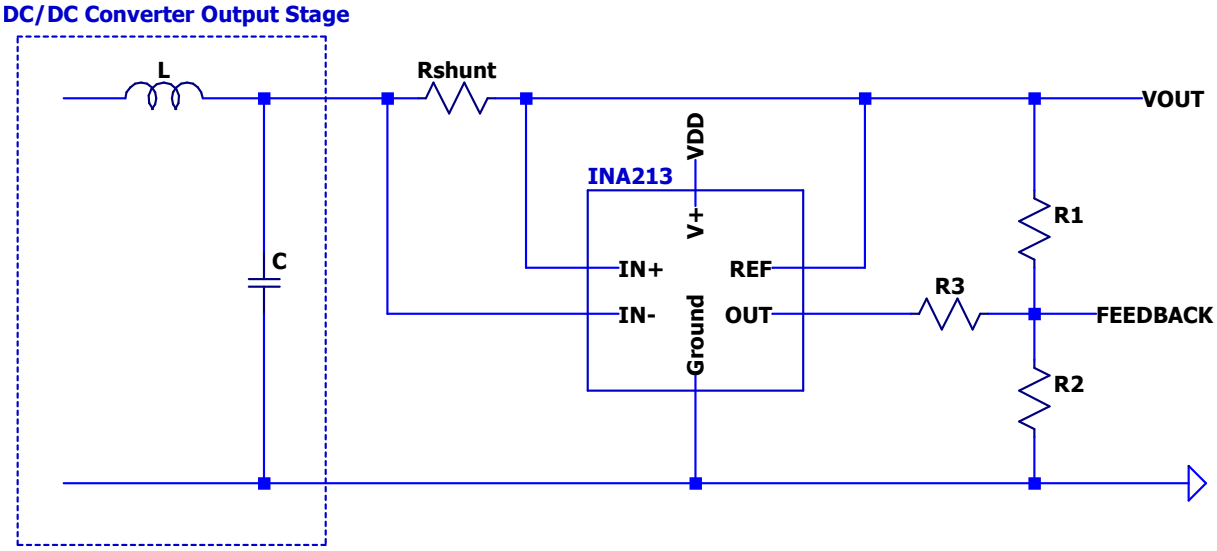

Figure 6.4: Simplified schematic of cable voltage drop compensation.

Analytical Examination

Using a shunt resistor and a differential amplifier (e.g. INA213 with gain 50), the output current is converted into a voltage. By connecting the input of the INA chip reversed (IN- to higher voltage and IN+ to lower voltage) and the reference voltage to V_{out}, the output voltage of the amplifier V_{INA} is:

$$V_{\text{INA}} = V_{\text{out}} - V_{\text{shunt}} \cdot G = V_{\text{out}} - R_{\text{shunt}} \cdot I \cdot G\,. \tag{6.1}$$

The output of the INA chip is connected to the feedback node using the resistor R_3. At 0 A output, both R_3 and R_1 pull up to V_{out}. At higher currents, V_{INA} is smaller than V_{out}. The TPS chip regulates to keep a target voltage of 600 mV at the feedback node, thus has to increase the output voltage, resulting ideally in a constant voltage at the load.

Of the three resistors (R_1, R_2, R_3), one can be chosen freely, for example R_2, defining the magnitude of the other resistors. Then the resistance of the network $R_{13} = R_3 \parallel R_1$ can be calculated from the nominal 0 A output voltage. The value of R_3 defines the slope of the voltage increase and can thus be calculated using cable resistance values. With R_3 and R_{13} known, R_1 can be calculated. The equations given in the application report [25]

are valid if the feedback is connected before the shunt. In the powerboard v2 and v3 schematics, the shunt is included in the feedback loop, since this has the advantage that voltage drops across the shunt are automatically compensated. This means the equations have to be slightly modified:

$$R_{13} = R_2 \cdot \left(\frac{V_{\text{out}}}{V_{\text{fb}}} - 1\right), \tag{6.2}$$

$$R_3 = R_{13} \cdot \frac{G_{\text{cs}} \cdot R_{\text{sh}}}{R_c} - R_{\text{filter}}, \tag{6.3}$$

$$R_1 = R_{13} \cdot \frac{R_3}{R_3 - R_{13}}. \tag{6.4}$$

Here V_{out} is the nominal output voltage (2 V), V_{fb} is the target voltage the controller tries to keep at the feedback node (for TPS53819A 600 mV), G_{cs} is the gain of the differential amplifier (for INA213 50), R_{sh} is the resistance of the shunt (2 mΩ), R_c is the two way cable resistance (60 mΩ), R_{filter} is the value of the resistor of the RC output filter after the differential amplifier (100 Ω). The validity of the equations has been shown by analytically solving the circuit.

Using the numerical values from above (valid for powerboard v3) and setting $R_2 = 10\,\text{k}\Omega$ (recommended in data sheet [11]), the three resistor values calculate to be:

$$R_1 = 58.560\,\text{k}\Omega, \tag{6.5}$$

$$R_2 = 10.000\,\text{k}\Omega, \tag{6.6}$$

$$R_3 = 38.789\,\text{k}\Omega. \tag{6.7}$$

Simulation using TINA-TI™

A circuit implementing drop compensation has been designed (see figure 6.5) and simulated using Texas Instruments' SPICE-based simulation program TINA-TI™ [26]. For simplicity, a buck converter chip with an integrated switch and reference design available is used (TPS53513 from Texas Instruments). It has similar specifications as the control chip used in the powerboards (see section 3.4.1).

To see a transient response of the converter, the enable signal of the chip is initially low and goes high after 0.1 ms. The two power modes of the MuPix sensor are simulated by varying the load resistance over time, with 1.8 Ω for the first 1.5 ms, then 0.3 Ω for 0.5 ms and then 1.8 Ω again. The resistance of the cables was set to 60 mΩ (two-way).

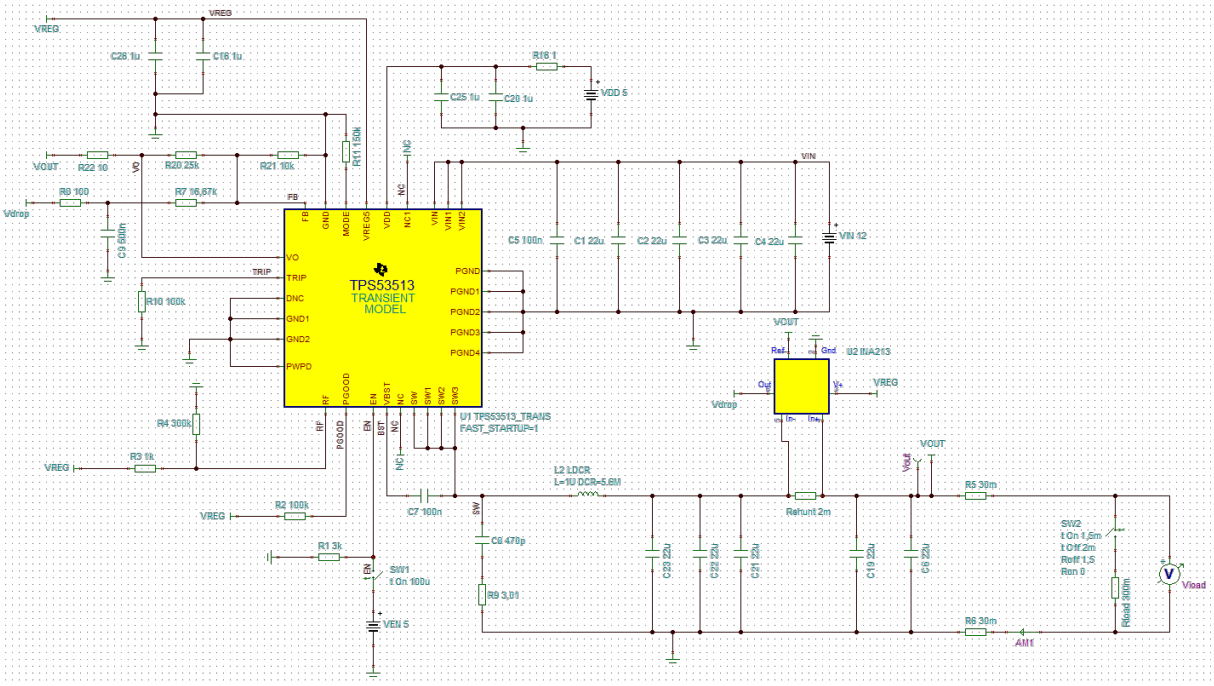

Figure 6.5: Circuit used to simulate drop compensation in TINA-TI™. The TPS53513 chip has an integrated switch and behaves similarly to the chip used in the powerboards (TPS53819A). A low-pass RC output filter is used here for the INA.

Figure 6.6 shows that, with drop compensation disconnected, the voltage stays stable on the converter's output whereas the voltage at the load drops for high currents because of the cable resistance. Connecting the drop compensation circuit without the use of a filter for the INA chip leads to strong oscillations of the output voltage (see figure C.1).

For the filter, two positions were analyzed, on the input of the INA chip and the output. The data sheet proposes input filtering to keep the output impedance low. This has the drawback of introducing a gain error, for 10 Ω resistors around 0.84 % (see [27, p. 17]). Simulation of this configuration, with 10 Ω and 1 µF for filtering, and thus a cutoff frequency of

$$f_{\text{cutoff}} = \frac{1}{2\pi RC} = 15.9\,\text{kHz}\,, \tag{6.8}$$

shows a good transient response with no oscillations (see figure 6.7a).

Since the output of the INA chip is followed by a high ohmic resistor, it is not necessary to design a low impedance output. For that reason, a low-pass RC filter can be placed on the output. To avoid a strong effect of the output impedance on the filter characteristics (see equation (6.8) with $R = R_{\text{output}} + R_{\text{filter}}$), impedance bridging is done, meaning one follows low output impedance with high input impedance. For that reason, R was set to 100 Ω, around a magnitude higher than the usual output impedance of operational amplifiers (1 Ω to 10 Ω). The simulation shows good results with capacitor values around 0.5 µF (see figure 6.7b). For lower values the cutoff frequency is too high, leading to oscillations,

while for higher values the cutoff frequency is too low, leading to slow transient responses and significant overshoots (see figure C.2). Overall input and output filtering have similar capabilities to suppress oscillations, but output filtering does not introduce a gain error and is, therefore, the preferred solution.

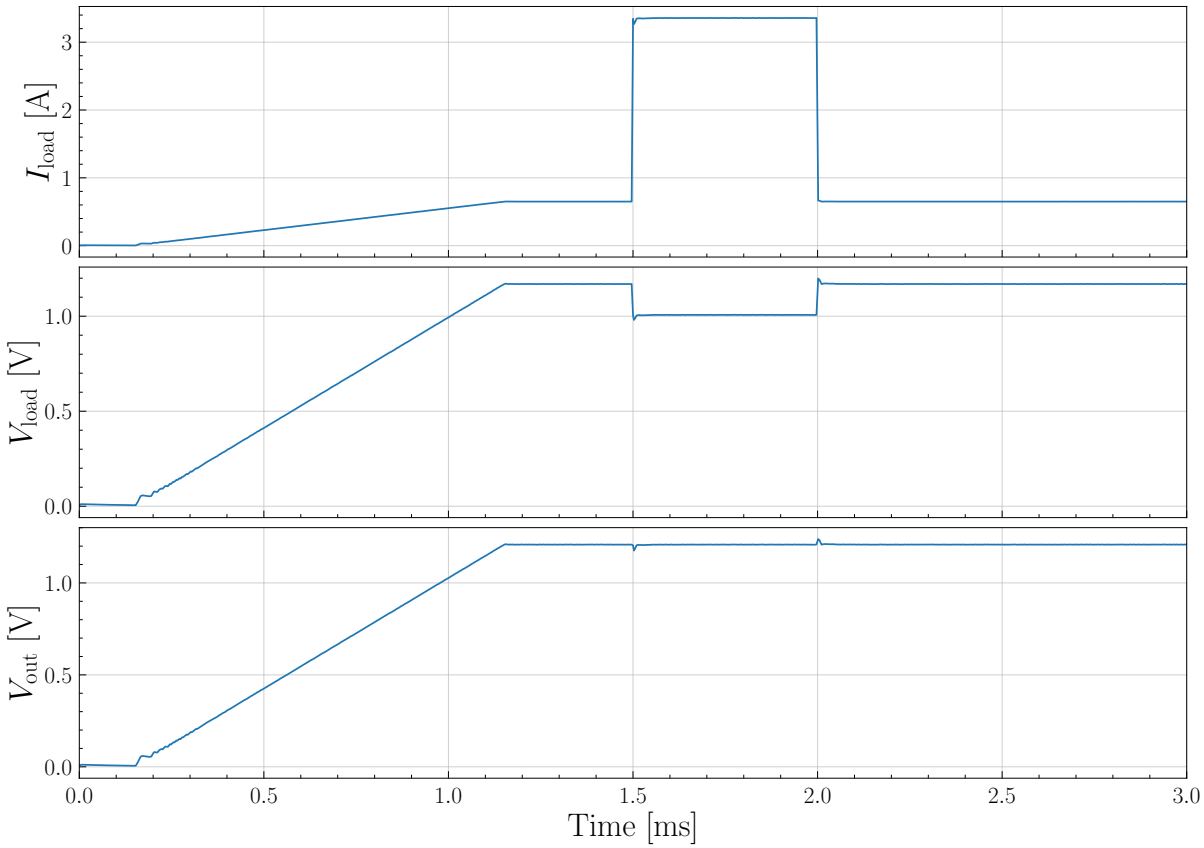

Figure 6.6: Simulation results with drop compensation disconnected.

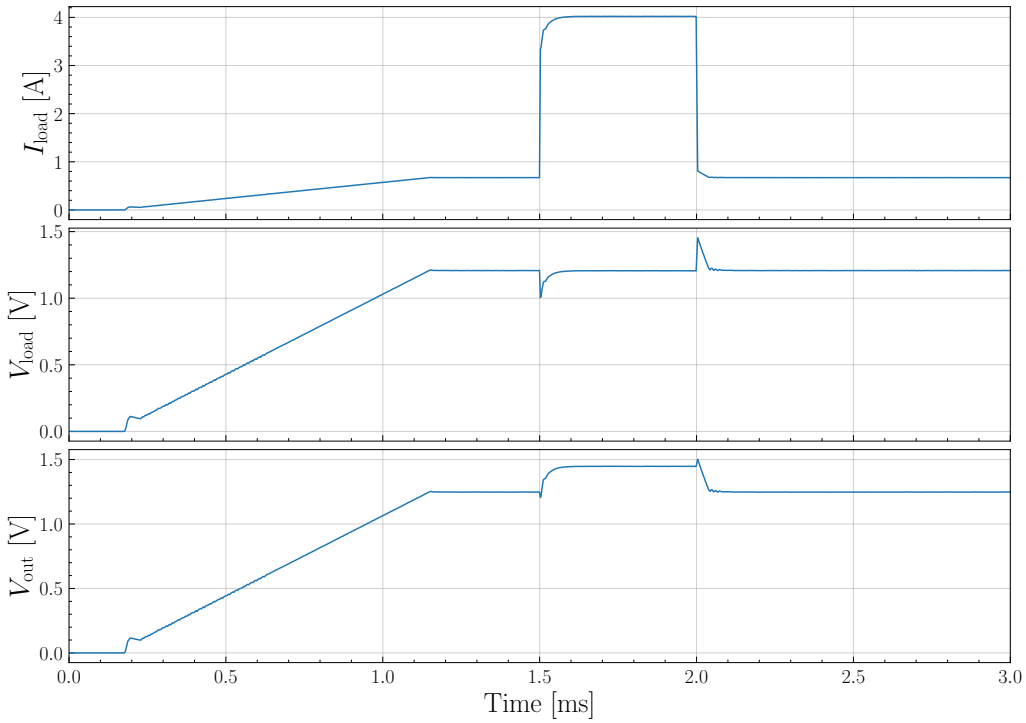

(a) With low-pass RC input filter (10 Ω and 1 µF, $f_c = 15.9$ kHz).

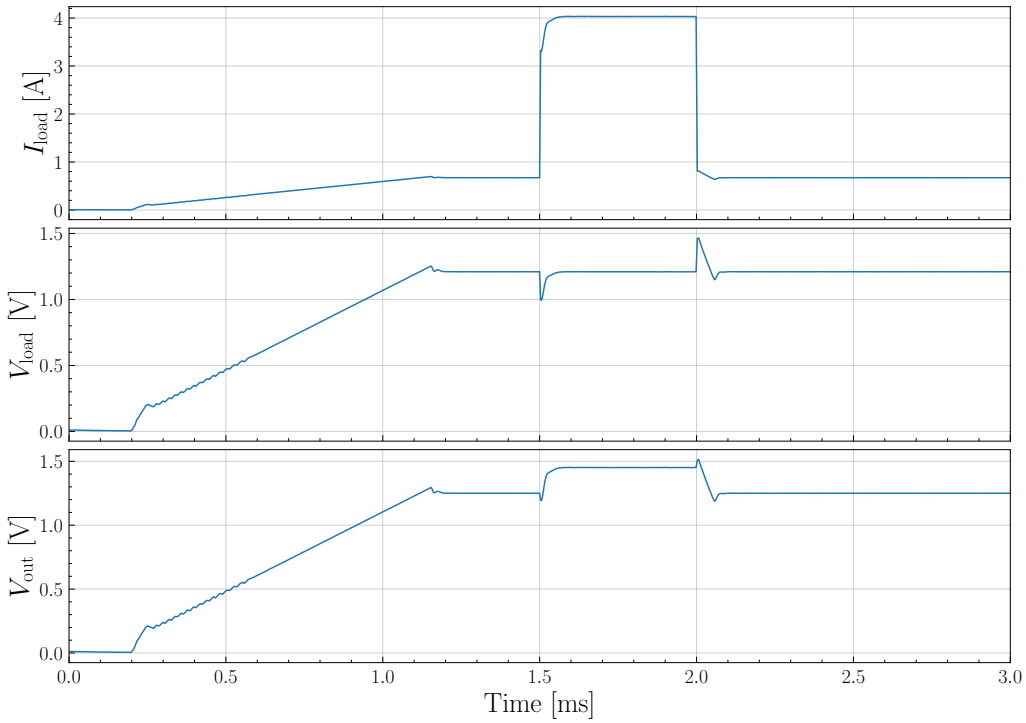

(b) With low-pass RC output filter (100 Ω and 0.5 µF, $f_c = 3.2$ kHz).

Figure 6.7: Simulation result with drop compensation connected and filtering for the INA chip.

Measurements in the Lab

Drop compensation was tested in the lab by soldering the INA213 evaluation module to the powerboard v2 using short wires. The onboard shunt was used as an input for the INA chip and the output was connected to the feedback node. The cable resistance was measured using an LCR meter and the appropriate feedback resistors were calculated and soldered.

Without the use of a filter for the INA chip, the system was unstable and high oscillations on the feedback pin triggered the overvoltage protection, shutting the converter down. Using a low-pass RC output filter with 100 Ω and 0.1 µF (the evaluation module has space designated for filtering), results were the same (f_c = 16 kHz). By increasing the capacitance to 1 µF (f_c = 1.6 kHz), stable operation was reached and the current could be varied up to 25 A.

As can be seen in table 6.1, for 20 A the drop compensation circuit increases the output voltage by 1.07 V, resulting not in a decrease of voltage at the load, but an increase of 40 mV. This can be explained by a slight mismatch between resistor values and cable resistance. The cable resistance was measured to be 52.0 mΩ. Having to use standardized resistor values, components compensating a 53.7 mΩ cable resistance were used. Looking at the voltage drop from table 6.1, the actual cable resistance is 51.5 mΩ. The difference of around 2.2 mΩ explains the 40 mV difference at 20 A.

Table 6.1: Voltage at the output of the board and the load at different currents. The drop compensation circuit slightly overcompensates because the feedback resistor values are slightly off.

I_{out} [A]	V_{out} [V]	V_{load} [V]
0	2.105	2.105
20	3.175	2.145

During operation, an increase of high-frequency noise on the feedback node was measured (see figure 6.8). This might be related to soldering the INA213 evaluation module to the board using wires, which significantly enlarges the area of the feedback loop and can act as an antenna for high-frequency switching noise. Only by designing a new PCB and minimizing the loop area can the final noise level be determined.

Figure 6.8: High-frequency noise measured on the feedback pin with drop compensation connected. Output voltage in yellow and the voltage at the feedback pin in blue.

6.1.3 Power Supply Monitor and Margining Chip

Another option investigated is the use of a second control chip besides the TPS53819A. The LTC2970 from Analog Devices [28] is an I^{2}C-controllable, dual power supply monitor and margining chip. It can servo the converter's output voltage using a single connection to the feedback node, similar to drop compensation. Deactivated, this connection is in a high impedance state, leaving the control loop in its original state. When activated, it connects to the feedback node without changing the output voltage ("soft-connect") and afterward can servo it. Sense wires can be connected to the internal 14-bit differential ADC (which can be multiplexed to 7 channels) and can be used to regulate a stable voltage across the load.

Using the DC2467A Arduino Uno shield, the functionality of the chip has been evaluated. Tests have shown that it stays fully functional when operated in a helium environment for long times. Investigating the sense wire feature of the chip revealed that it is not good at quickly compensating big voltage changes at the load. It operates in two modes, servo voltage and hold voltage mode. The servo voltage mode is used to quickly reach the

target voltage. In this mode, the ADC operates initially in an accelerated 12-bit mode until the target voltage is reached, and then does two more 14-bit conversions for fine-tuning. During this time, ADC conversions for all non-servo channels are temporarily inhibited. When the target voltage is reached, the chip operates in the hold voltage mode. ADC channels are read out sequential and voltage regulation is not prioritized. This results in slow voltage compensation. For a 1 V change the compensation takes around 1 s.

There is no option to stay in the servo voltage mode after the target voltage is reached. That means continuously set voltage commands need to be sent. At the same time, readout of ADC channels should be possible at all times, independent of voltage compensation. For these reasons, the chip was deemed not fitting for the application.

6.1.4 Conclusion

The results show that cable voltage drop compensation combined with voltage regulation via PMBus commands is a suitable solution. Drop compensation responds very quickly and reliable to load changes, while keeping the control loop stable. Cable resistances need to be measured once in the beginning. This should be done by using a high current and measuring the voltage drop. Using this method, heating and thus changing resistance of the cables is also taken into account. Selecting and soldering different resistors to get a fine-tuned drop compensation for every board might be tedious and not recommended. It is also not necessary, since many detector modules have identical power connections up to the final connection. Depending on how much the cable resistance varies, boards can have the same resistor values and fine-tuning of the output voltage can be done using the PMBus voltage regulation command. The voltage at the load will be known since the powerboard v3 includes sense wires for monitoring reasons. How long the feedback loop (ADC readout of load voltage - sending PMBus command - regulate voltage) takes still needs to be investigated, but a first estimate is in the order of 200 ms.

6.2 Backplane Bus and Adressing

To control the powerboards version 2, a shared I^2C bus is used. The address of each TPS chip is set by a voltage divider located on the corresponding board. When working with the mini-crate, this design showed many flaws. Firstly, to get unique addresses, the resistors of the voltage divider need to be changed manually on each board. Secondly, when boards are plugged into the mini-crate, it is not possible to automatically derive which board, and thus address, is located in which slot of the crate. Geographical addressing is missing. This is highly problematic, since temperature sensors, for example, are identified by crate slot and not by powerboard. It also makes debugging eight crates with 16 boards each in the Mu3e phase I experiment very difficult.

To solve this problem, the backplane bus design has been changed drastically. Instead of using only one I^2C bus, each slot gets its own bus, resulting in 16 buses for the full-crate (see figure 6.9). There will be two devices on each bus, the TPS chip and an ADC

(ADS1115) (both located on the corresponding board). Addresses of these devices will be the same across all boards. Since the crate controller only supports one I^2C bus, this bus will be multiplexed using five LTC4314 chips. The LTC4314 chip can multiplex one to four for a 2-wire bus with enable pins for each bus [29]. Using a two-layer stack, with one chip in the first layer and four chips in the second layer, each of the 16 buses can be enabled separately using eight pins of the controller. For the second layer only four pins are needed, since enable pins can be connected between chips. This design solves both of the previously mentioned problems. Unknown is the exact speed at which the two-layer multiplexing works. A high speed is important, especially with the changes made to the voltage regulation (see section 6.1.4), which relies on digital voltage regulation over the PMBus interface (I^2C).

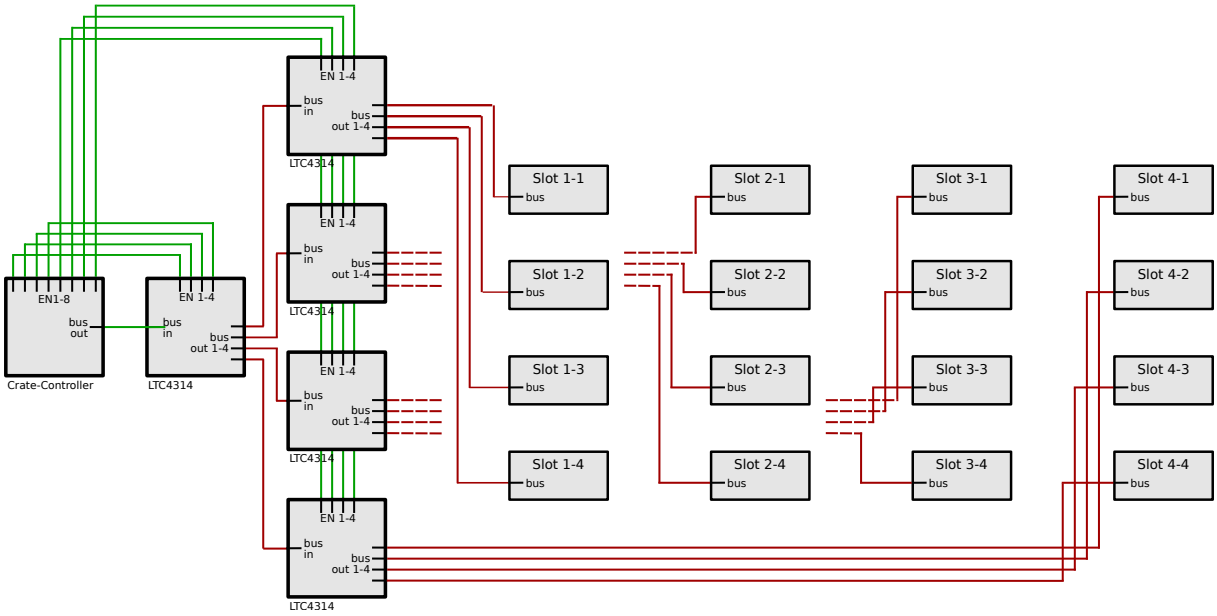

Figure 6.9: Schematic of the backplane bus design.

Chapter 7

Conclusion and Outlook

While the Mu3e detector is small enough to fit on a table, powering is not trivial. With a high power consumption and constrained space, as in most particle physics experiments, operation conditions are not ideal.

Section 2.5 newly estimates the maximum power required by the detector using the most recent measurements. This includes taking voltage drops across cables, which are not negligible, especially for the high current power partitions, fully into account. With 11.7 kW, the new estimate lies around 2.3 kW above previous estimates. This result can be improved by reducing the resistance of the chain of conductors which is used to distribute power (see section 2.4.4). One measure to achieve this is by increasing the diameter of the cables used to connect the powerboards to the copper bus bars, which are currently AWG 14 cables. Another measure is the optimisation of the connection from the copper bus bars to the detector modules, which is done using Samtec Z-Ray® Interposers and flex PCBs. This connection has not been thoroughly analysed yet and with an estimated resistance of 6 mΩ results in the largest voltage drop.

The design of the power-crates, in which the powerboards will be operated, has been developed in chapter 4. The crate features a robust mechanical structure and water cooling. The estimated weight is slightly too high, but can be reduced by removing unnecessary material from the cooling blocks (see section 4.1.2). Studies need to be conducted to fully integrate it into the mechanical and cooling design of the detector.

With the production of the powerboard v2, a first prototype converter conforming to many of the final specifications was available. As a result, the mini-crate has been developed in section 4.3. It is based on the mechanical design of the power-crate, but also includes the electronics required to run up to four boards. Combined with the control interface, simple operation and monitoring of the powerboards have been achieved. Using the mini-crate, the design of the cooling system has been evaluated, giving an efficiency of $(81 \pm 6)\,\%$ (see section 4.4). Operation of three boards at 20 A showed no significant problems and hints that running a full-crate with 16 boards is feasible. The current bottleneck, the mounting mechanism for the heatsink, has been improved in version 3 of the powerboard, possibly

increasing the efficiency of the cooling system. Since version 3 of the powerboard has a different layout and pinout, a new iteration of the mini-crate has to be built. With this available, a long-term study should be done to establish a maximum current over time rating.

Version 2 of the powerboard has been tested extensively and the gained knowledge implemented in the design of version 3. The efficiency of the converter has been calculated for different operation points in section 5.4. For 20 V input voltage and 21 A output current an efficiency of $(82.9 \pm 0.5)\,\%$ has been determined. The output filter has been optimized, resulting in low-frequency ripple of 1.5 mV peak-to-peak (see section 5.1 and section 5.3). When designing the PCB layout for version 3, measures described in section 5.3 should be taken to reduce high-frequency noise. With cable voltage drop compensation and digital voltage regulation, a solution for one of the biggest problems, voltage drops across cables, has been found (see section 6.1). This mechanism still needs to be tested extensively, especially the response time of the digital voltage regulation. The new backplane bus design implements geographical addressing, making debugging and programming of the control interface a lot easier. While the two-layer multiplexing is an elegant solution, the speed at which it can be operated needs to be investigated. It makes sense to produce a prototype PCB for this before finalizing the powerboard version 3 design.

Overall, many things have been learned from this iteration of the powerboard. Big steps have been made concerning the integration of the converters into the experiment. With knowledge obtainable from version 2 almost being depleted, the time for the next iteration has come. Version 3, while posing new challenges, should solve many of the current problems and be significantly closer to the design of the final converter.

Appendix A

Power-Crate

A.1 Mini-Crate Cooling Loop

Table A.1: Secondary Loop (landefeld.de (ld), edelstahl-waermetauscher.de (ew))

Name	**Input**	**Output**	**PN**	**Shop**
Wall Big Lab Out	-	G 3/4" AG		
Filter	G 3/4" IG	G 3/4" IG	SF 34	ld
Reduction Socket	G 3/4" AG	G 1/2" IG	RN 3414 MS	ld
Tube Fitting	G 1/2" AG	12x10 Tube	CK 1210 MSV	ld
Tube	12x10 Tube	12x10 Tube	PA 12X10 NATUR	ld
Tube Fitting	12x10 Tube	G 1/2" AG	CK 1210 MSV	ld
Reduction Socket	G 1/2" IG	G 3/4" IG	MUR 3412 MS	ld
Gasket	G 3/4" IG	G 3/4" IG	DR 34 K	ld
Plate Heat Exchanger	G 3/4" AG	G 3/4" AG	EWT-BE4-13 (20)	ew
Gasket	G 3/4" AG	G 3/4" AG	DR 34 K	ld
Reduction Socket	G 3/4" IG	G 1/2" IG	MUR 3412 MS	ld
Tube Fitting	G 1/2" AG	12x10 Tube	CK 1210 MSV	ld
Tube	12x10 Tube	12x10 Tube	PA 12X10 NATUR	ld
Tube Fitting	12x10 Tube	G 1/2" AG	CK 1210 MSV	ld
Wall Big Lab In	G 1/2" IG	-		

Table A.2: Primary Loop (landefeld.de (ld), edelstahl-waermetauscher.de (ew), conrad.de (cd))

Name	**Input**	**Output**	**PN**	**Shop**
Mini-Crate hot (2x)	-	G 1/4" IG		
Gasket (2x)	G 1/4" AG	G 1/4" AG	DR 14 K	ld
Tube Fitting (2x)	G 1/4" AG	8x6 Tube	CK 146 A	ld
Tube	8x6 Tube	8x6 Tube	PA 8X6 NATUR	ld
T-piece	8x6 Tube	8x6 Tube	FCK 6 K	ld
Tube	8x6 Tube	8x6 Tube	PA 8X6 NATUR	ld
Tube Fitting	8x6 Tube	G 1/4" AG	CK 146 A	ld
Gasket	G 1/4" AG	G 1/4" AG	DR 14 K	ld
Temp Sensor	G 1/4" IG	G 1/4" AG	Phobya 10kOhm	cd
Reduction Nipple	G 1/4" IG	G 3/4" AG	RN 3414 MS	ld
Gasket	G 3/4" AG	G 3/4" AG	DR 34 K	ld
Threaded Socket	G 3/4" IG	G 3/4" IG	MUR 34 MS	ld
Plate Heat Exchanger	G 3/4" AG	G 3/4" AG	EWT-BE4-13 (20)	ew
Threaded Socket	G 3/4" IG	G 3/4" IG	MUR 34 MS	ld
Gasket	G 3/4" IG	G 3/4" IG	DR 34 K	ld
Reduction Nipple	G 3/4" AG	G 1/4" IG	RN 3414 MS	ld
Temp Sensor	G 1/4" IG	G 1/4" AG	Phobya 10kOhm	cd
Gasket	G 1/4" IG	G 1/4" IG	DR 14 K	ld
Tube Fitting	G 1/4" AG	8x6 Tube	CK 146 A	ld
Tube	8x6 Tube	8x6 Tube	PA 8X6 NATUR	ld
Tube Fitting	8x6 Tube	G 1/4" AG	CK 146 A	ld
Gasket	G 1/4" AG	G 1/4" AG	DR 14 K	ld
Water Resevoir	G 1/4" IG	Inlet Pump	500604	cd
Pump	Inlet Pump	G 1/8" IG, special	500733-02	cd
Pump Adapter	G 1/8" AG, special	G 1/4" IG	Eheim Adapter	cd
Tube Fitting	G 1/4" AG	8x6 Tube	CK 146 A	ld
Tube	8x6 Tube	8x6 Tube	PA 8X6 NATUR	ld
T-piece	8x6 Tube	8x6 Tube	FCK 6 K	ld
Tube	8x6 Tube	8x6 Tube	PA 8X6 NATUR	ld
Tube Fitting (2x)	8x6 Tube	G 1/4" AG	CK 146 A	ld
Mini-Crate cold (2x)	8x6 Tube	-		

A.2 PCB Schematic Backplane Mini-Crate v2

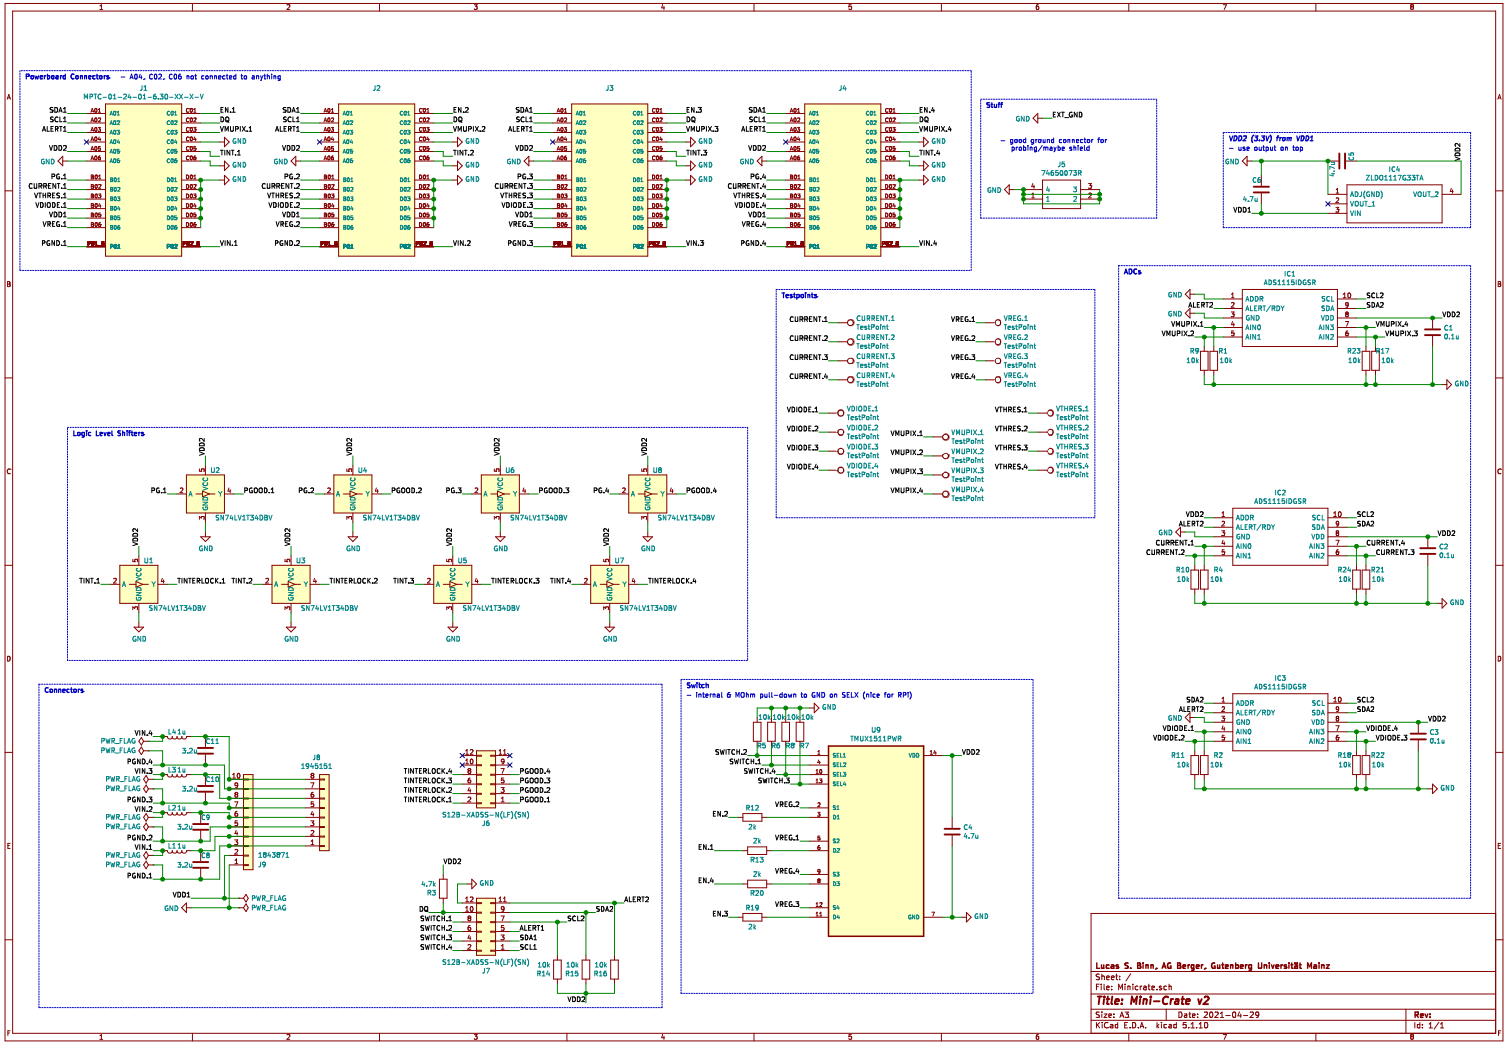

A.3 Drawing Heatsink v2

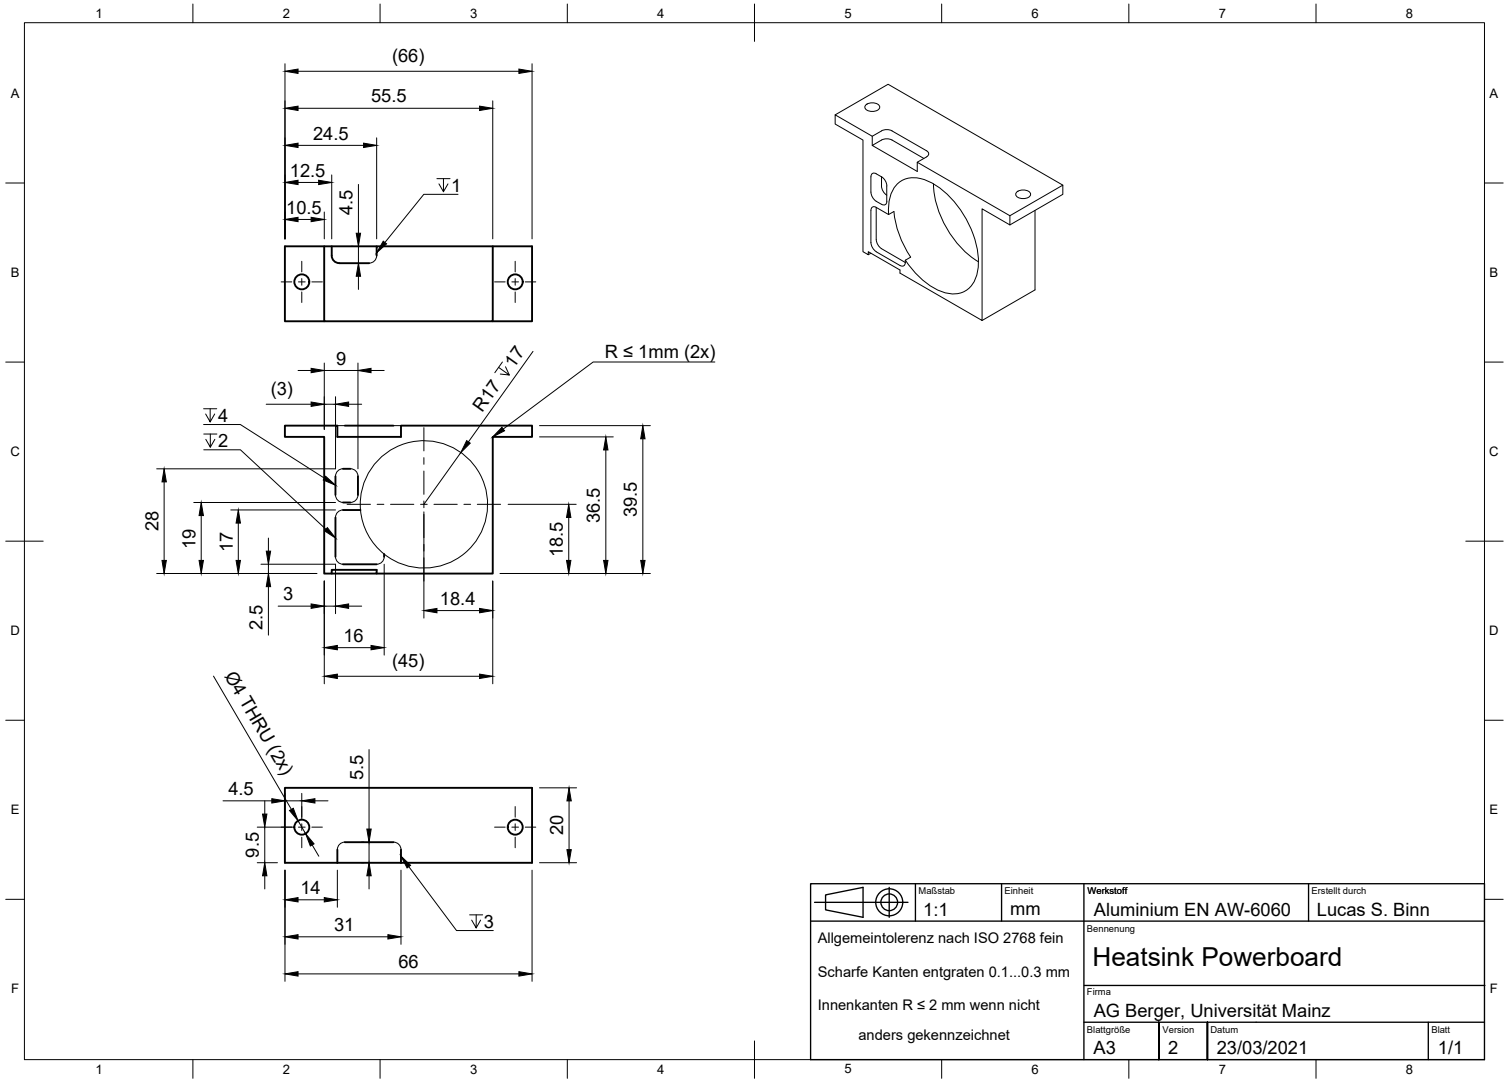

A.4 Drawing Cooling Block v1

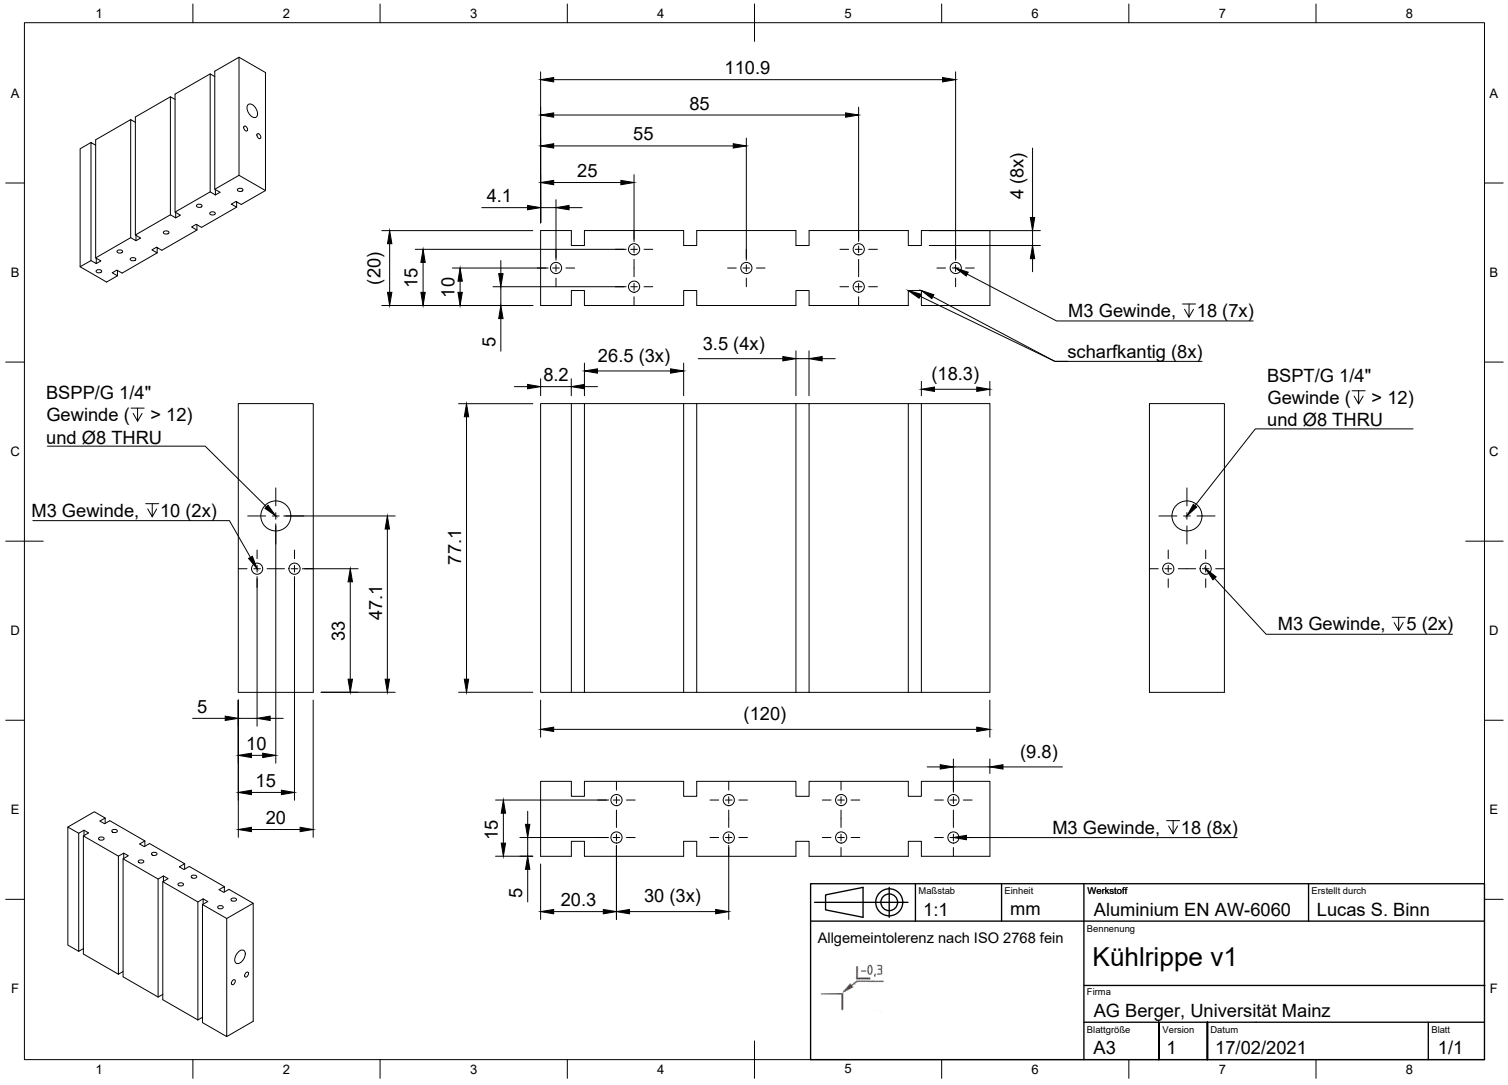

Appendix B

PCB Schematic Powerboard v2

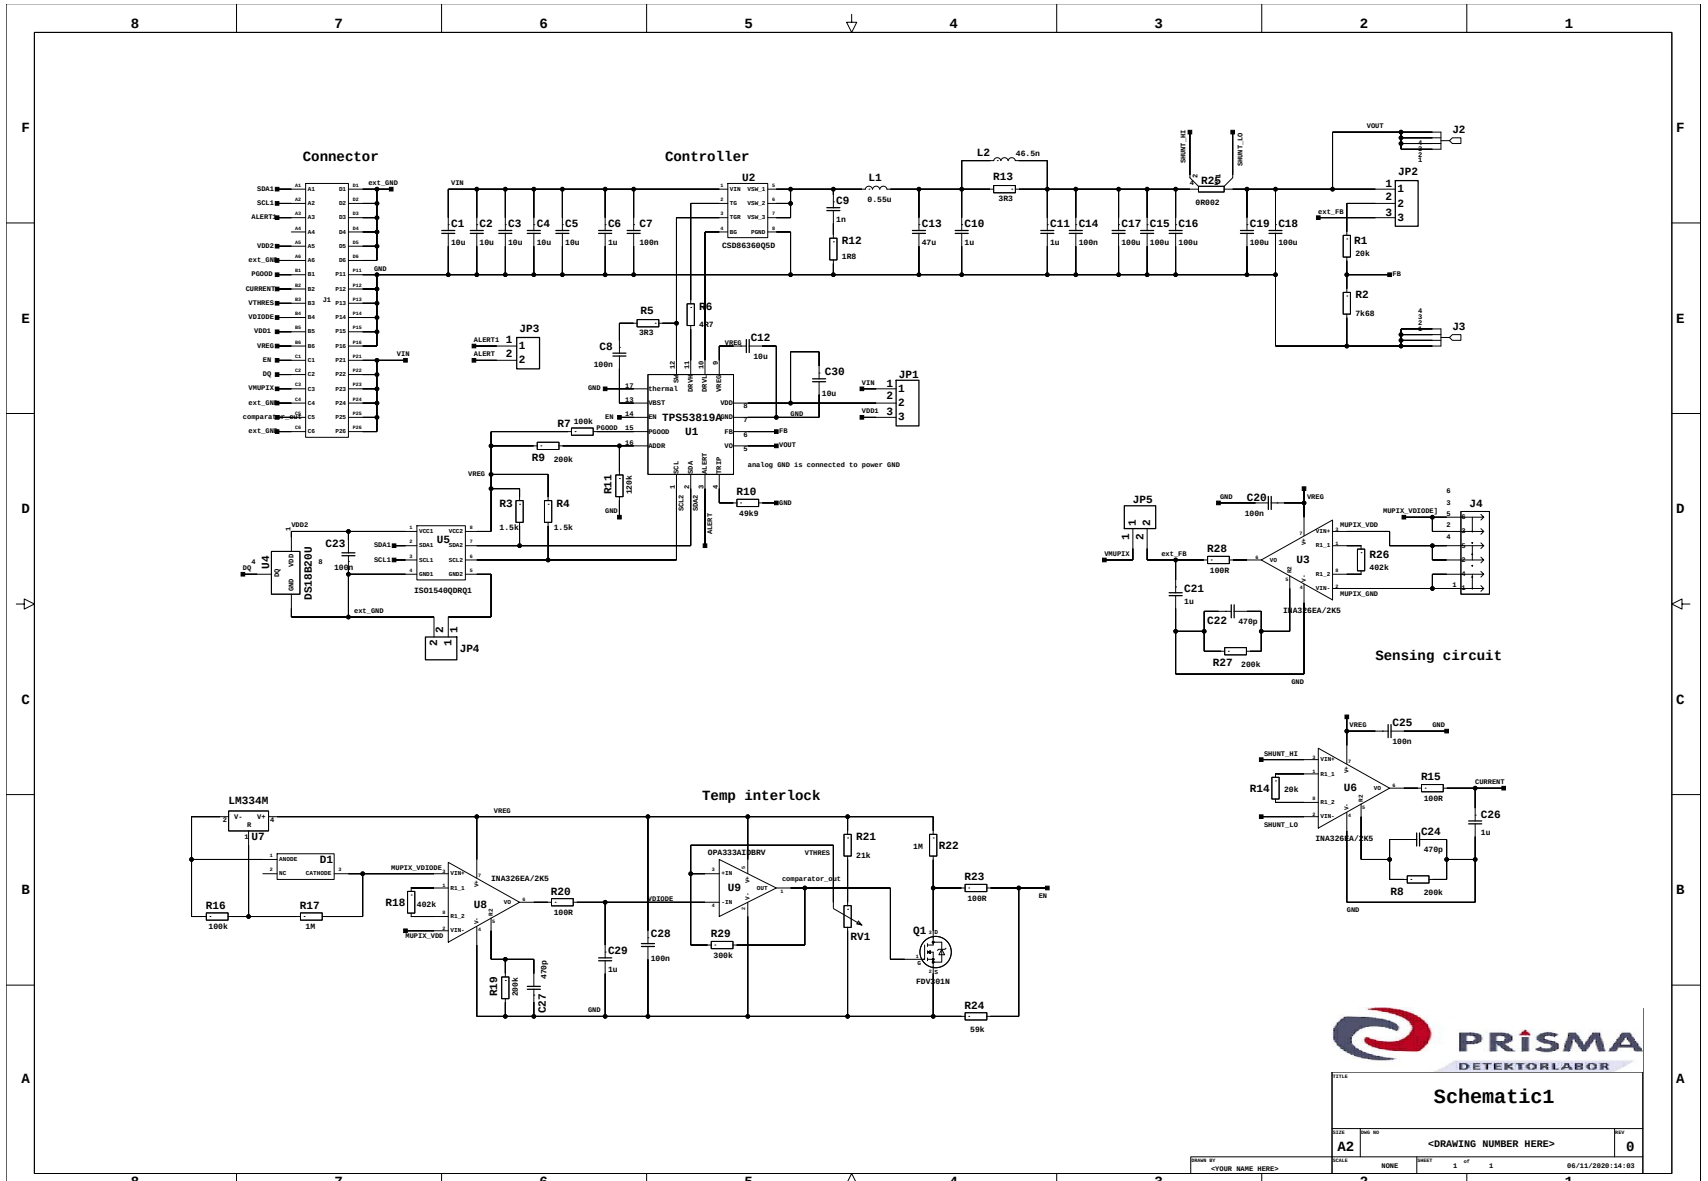
Connector
Controller
Sensing circuit
Temp interlock
Schematic1
PRISMA
DETEKTORLABOR
<DRAWING NUMBER HERE>
<YOUR NAME HERE>
analog GND is connected to power GND
TPS53819A
CSD86360QSD
INA326EA/2K5
ISO1540QDRQ1
DS18B20U
LM334M
FDV301N
OPA333AIDBVR

Appendix C

Powerboard v3

C.1 Cable Voltage Drop Compensation

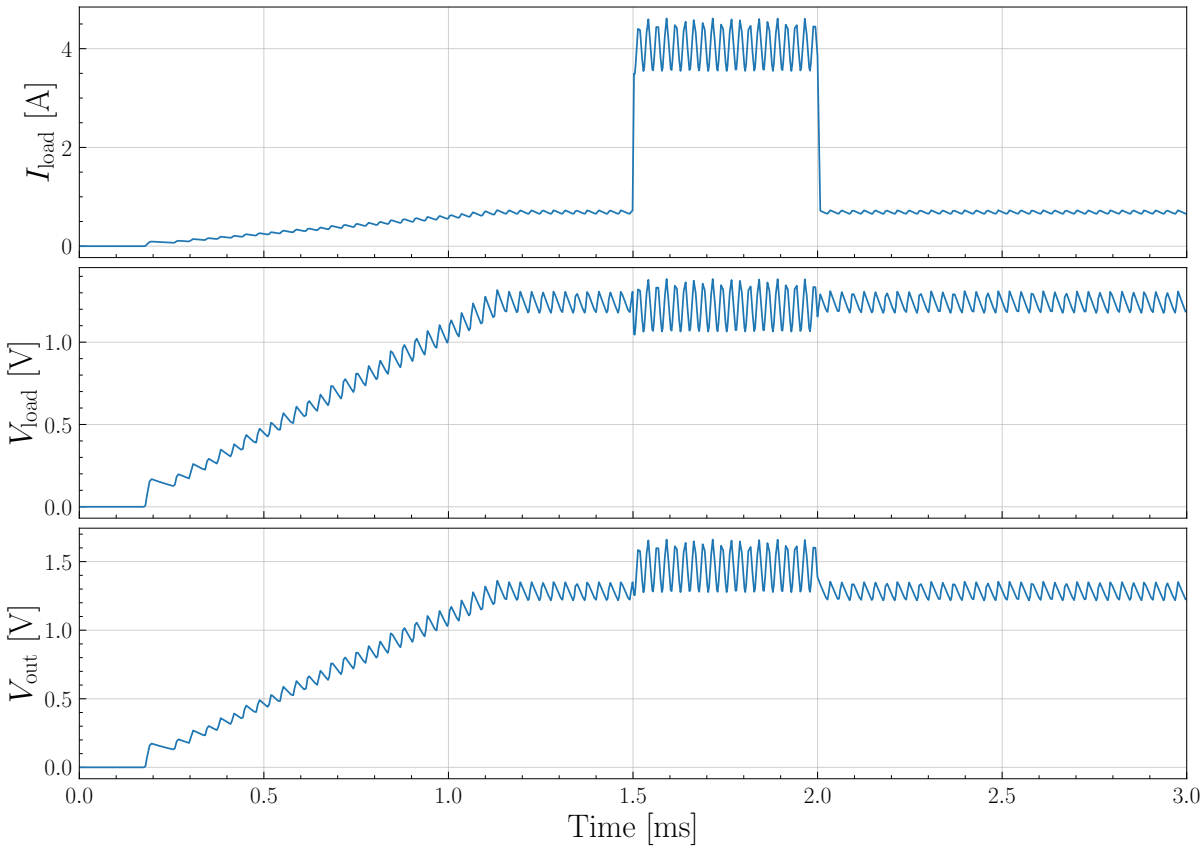

Figure C.1: Results from simulating drop compensation with TINA-TI. Drop compensation connected without the use of a filter.

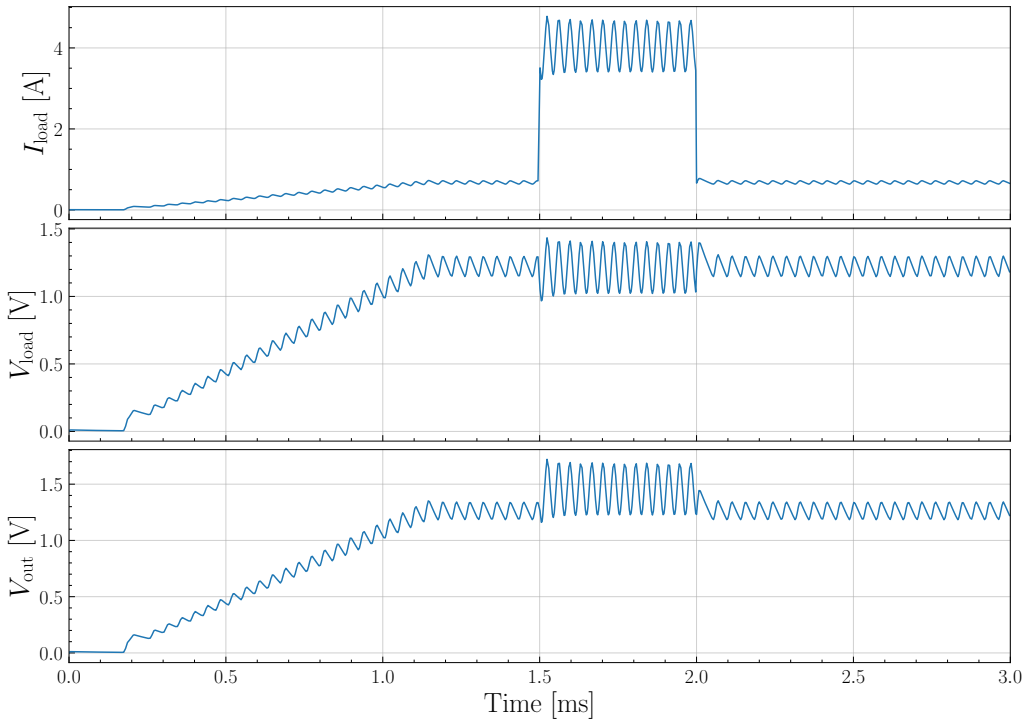

(a) 100 Ω and 0.1 µF ($f_c = 15.9\,\text{kHz}$)

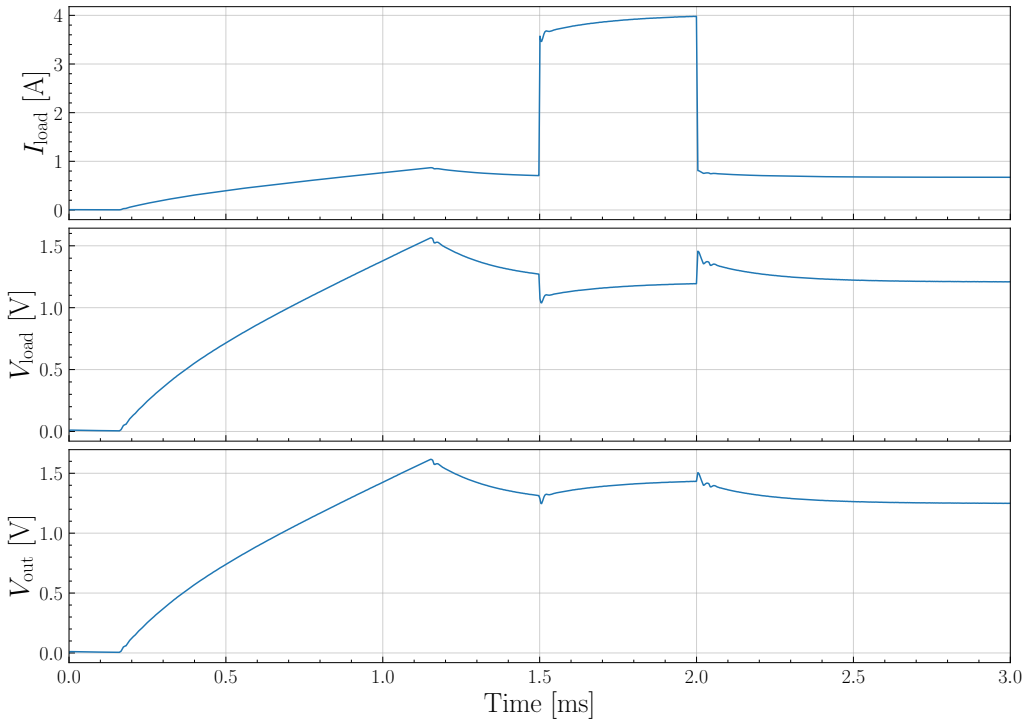

(b) 100 Ω and 5 µF ($f_c = 0.3\,\text{kHz}$)

Figure C.2: Results from simulating drop compensation with TINA-TI using an low-pass RC output filter after the INA.

C.2 PCB Schematic Powerboard v3

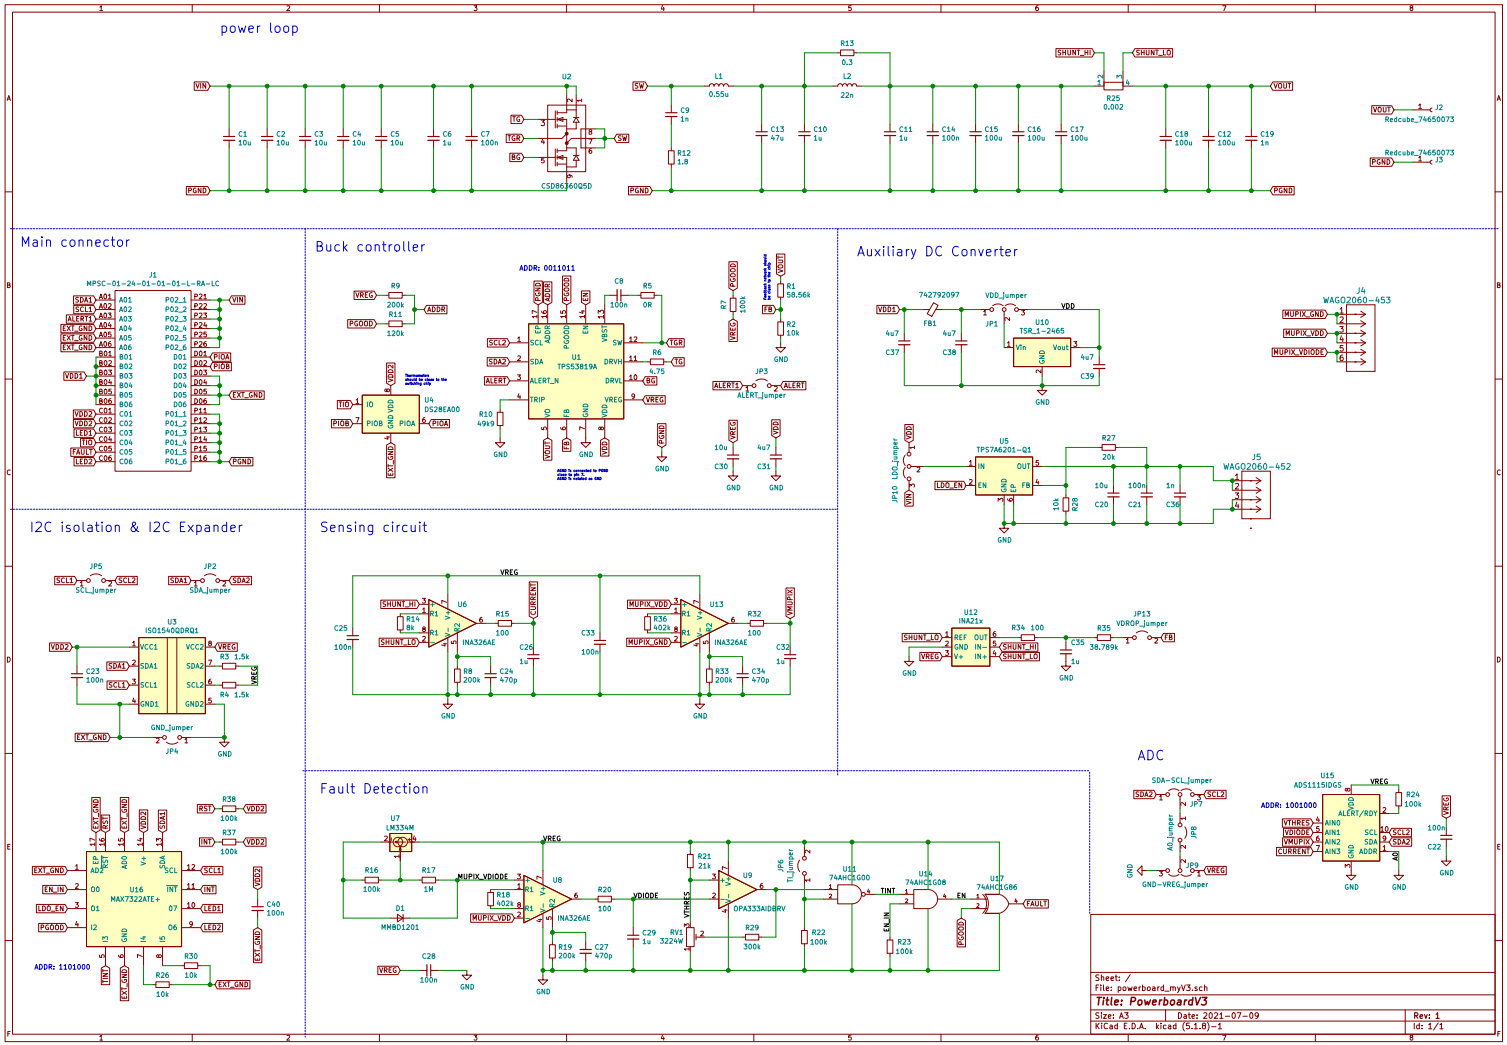

Bibliography

[1] The Mu3e Collaboration. *Mu3e Internal Wiki*. 2021.

[2] P.A. Zyla et al. "Review of Particle Physics". In: *PTEP* 2020.8 (2020), p. 083C01. DOI: 10.1093/ptep/ptaa104.

[3] William J. Marciano, Toshinori Mori, and J. Michael Roney. "Charged Lepton Flavor Violation Experiments". In: *Annual Review of Nuclear and Particle Science* 58.1 (2008), pp. 315–341. DOI: 10.1146/annurev.nucl.58.110707.171126.

[4] Ivan Perić. "A novel monolithic pixelated particle detector implemented in high-voltage CMOS technology". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 582.3 (2007). VERTEX 2006, pp. 876–885. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2007.07.115. URL: https://www.sciencedirect.com/science/article/pii/S0168900207015914.

[5] K. Arndt et al. "Technical design of the phase I Mu3e experiment". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1014 (Oct. 2021), p. 165679. ISSN: 0168-9002. DOI: 10.1016/j.nima.2021.165679. URL: http://dx.doi.org/10.1016/j.nima.2021.165679.

[6] Frederik Wauters. *Power Distribution*. Mu3e Collaboration Meeting, Wengen. 2020.

[7] Marco Zimmermann. *Technical Design Development for the P2 Silicon Pixel Tracking Detector*. 2019.

[8] André Schöning. *Pixel Tracker Voltage Drops and Power Consumption*. Mu3e Collaboration Meeting at UZH. 2019.

[9] Cristian Alejandro Fuentes Rojas. *Optimization of the design of DC-DC converters for improving the electromagnetic compatibility with the Front-End electronic for the super Large Hadron Collider Trackers*. 2011.

[10] TexasInstruments Denislav Petkov. *Understanding, Measuring, and Reducing Output Noise in DC/DC Switching Regulators*. https://training.ti.com/understanding-measuring-and-reducing-output-noise-dcdc-switching-regulators. Accessed: 2021-11-25. 2018.

[11] Texas Instruments. *Data Sheet TPS53819A*. https://www.ti.com/lit/ds/symlink/tps53819a.pdf. Accessed: 2021-11-22. 2019.

[12] Texas Instruments. *Data Sheet CSD86360Q5D*. https://www.ti.com/lit/ds/symlink/csd86360q5d.pdf. Accessed: 2021-08-23. 2018.

[13] Moritz Hesping. *Air Coils for Powering the Mu3e Experiment*. 2019.
[14] Stefan Ritt. *Midas Slow Control Bus (MSCB)*. https://elog.psi.ch/mscb/. Accessed: 2021-11-08. 2001.
[15] makeitfrom. *Properties of Aluminium 6060-T*. https://www.makeitfrom.com/material-properties/6060-T6-Aluminum. Accessed: 2021-10-21.
[16] makeitfrom. *Properties of Stainless Steel 304*. https://www.makeitfrom.com/material-properties/AISI-304-S30400-Stainless-Steel. Accessed: 2021-10-21.
[17] E. McCafferty. *Introduction to Corrosion Science*. Springer New York, 2010. ISBN: 9781441904546.
[18] Sophie Gagneur. *Developement of a DC-DC Converter for the Mu3e Detector*. 2020.
[19] GELID solutions. *GP-Extreme Thermal Pad*. https://gelidsolutions.com/thermal-solutions/accessories-gp-extreme/. Accessed: 2021-10-26.
[20] Bergquist. *Thermal Gap Fillers*. https://www.henkel-adhesives.com/de/en/products/thermal-management-materials/thermal-gap-fillers.html. Accessed: 2021-10-29.
[21] Alumeco A/S. *Data Sheet EN AW 6060*. https://www.alumeco.com/media/3806/6060-profiles-incl-anodizing.pdf. Accessed: 2021-08-24.
[22] Alumeco A/S. *Data Sheet EN AW 6082*. https://www.alumeco.com/media/3809/6082-profiles.pdf. Accessed: 2021-08-24.
[23] AVX Jeffrey Cain. *Parasitic Inductance of Multilayer Ceramic Capacitors*. https://www.avx.com/docs/techinfo/CeramicCapacitors/parasitc.pdf. Accessed: 2021-11-25. 2020.
[24] CERN. *FEAST Data Sheet*. https://espace.cern.ch/project-DCDC-new/Shared%20Documents/FEAST%20datasheet.pdf. Accessed: 2021-10-25. 2014.
[25] Texas Instruments. *Application report SLVA657: Step-Down Converter with Cable Voltage Drop Compensation*. https://www.ti.com/lit/an/slva657/slva657.pdf. Accessed: 2021-09-28. 2014.
[26] Texas Instruments. *TINA-TI*. https://www.ti.com/tool/TINA-TI. Accessed: 2021-11-21.
[27] Texas Instruments. *Data Sheet INA213*. https://www.ti.com/lit/ds/symlink/ina213.pdf. Accessed: 2021-09-28.
[28] Linear Technology. *Data Sheet LTC2970*. https://www.analog.com/media/en/technical-documentation/data-sheets/29701fe.pdf. Accessed: 2021-10-25.
[29] Linear Technology. *Data Sheet LTC4314*. https://www.analog.com/media/en/technical-documentation/data-sheets/4314f.pdf. Accessed: 2021-12-19.

Acknowledgments

I would like to thank Prof. Dr. Niklaus Berger for giving me the opportunity to work in the Mu3e collaboration and for always having time to answer my questions. Many thanks to Dr. Frederik Wauters for the provided guidance and intense discussions. Thanks to Sophie for making long days in the lab feel shorter and to Lars for insights into the world of electronics. Also many thanks to the rest of the AG Berger group for the companionship during this time.

Thanks to my university dream team Anna, Aaron, and Daniel, for the good times spent.

Lastly, I wanna thank my family, Andrea, Karl-Heinz, Larissa, for their constant support and patience with me.