



## MASTER THESIS

# Design and Implementation of a High-Precision Temperature Controller

submitted in partial fulfillment of the requirements for the degree of  
Master of Science

supervised by  
Univ.Prof. Dipl.-Ing. Dr.techn. Michael Waltl

performed at the  
Institute for Microelectronics

submitted to the TU Wien,  
Faculty of Electrical Engineering and Information Technology

by

FLORIAN WIMMER BSC

01526746 / E 066 515

Louise-Martini-Weg 9/6, 1030 Wien  
E-Mail: [florian.wimmer@tuwien.ac.at](mailto:florian.wimmer@tuwien.ac.at)

Vienna, October 2025



# Abstract

This thesis presents the design, implementation and experimental validation of a high-precision, high-power temperature controller intended to support low-noise semiconductor device characterisation measurements at the Institute of Microelectronics, TU Wien. Existing commercial systems provide seriously limited voltage range, insufficient electric noise performance, or lack an open interface that would allow the use of efficient control algorithms other than their built-in PID regulators.

The objective of this work was to develop a versatile temperature controller capable of supporting multiple systems, including a resistive oven and a thermoelectric module, while achieving low electrical noise and high temperature accuracy. The system combines a custom high-voltage linear amplifier, low-noise measurement circuitry and a real-time microcontroller firmware with modular experiment modes, USB communication, touch-screen operation and comprehensive safety supervision features such as over current and temperature detection, short-circuit safe operation and more.

The hardware was validated electrically using oscilloscope measurements and thermally by implementing PID-based controllers for both the oven and the thermoelectric setups. The final system achieves an output voltage range from  $-18\text{ V}$  to  $70\text{ V}$  with a voltage noise below  $3.5\text{ mV}_{\text{rms}}$  and supports currents up to  $7\text{ A}$ . The temperature control loop operates at  $5\text{ Hz}$  and achieves a steady state accuracy of approximately  $\pm 35\text{ mK}$  for the oven setup and  $\pm 25\text{ mK}$  for the thermoelectric setup. The temperature sensor provides an accuracy of up to  $\pm 10\text{ mK}$ . All implemented safety mechanisms, including over current detection, fuse protection and firmware level supervision, were verified experimentally.

The results demonstrate that the prototype meets the requirements of most experimental setups at the Institute and provides a flexible platform for future developments. Due to its modular hardware architecture and extensible firmware design, the system also enables the implementation and evaluation of advanced control strategies, offering a promising foundation for further research.



# Contents

<b>List of Figures</b>	<b>v</b>
<b>List of Tables</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Outline . . . . .	3
<b>2 Theory</b>	<b>5</b>
2.1 Fundamentals of Thermal Systems . . . . .	5
2.1.1 Heat Transfer . . . . .	5
2.1.2 Thermal Capacity . . . . .	7
2.1.3 Thermal Actuators and Sensors . . . . .	8
2.2 Electronics and PCB Design . . . . .	11
2.2.1 Electrical Safety Requirements . . . . .	11
2.2.2 Measurement Circuits . . . . .	12
2.2.3 Power Supply and Rectification . . . . .	14
2.2.4 Amplifier Design . . . . .	14
2.3 Software Design . . . . .	21
2.4 Control Design . . . . .	24
2.4.1 System Identification . . . . .	25
2.4.2 PID Controller . . . . .	25
2.4.3 Singular Perturbation Theory . . . . .	26
<b>3 Method</b>	<b>27</b>
3.1 Hardware . . . . .	27
3.1.1 Power Supply . . . . .	28
3.1.2 Backpanel . . . . .	32
3.1.3 Amplifier . . . . .	33
3.1.4 Digital Hardware . . . . .	37
3.2 Software . . . . .	39
3.3 Controller . . . . .	42
3.3.1 Electrical Controller . . . . .	42
3.3.2 Temperature Controller . . . . .	43

<b>4 Result</b>	<b>49</b>
4.1 Electrical Measurements . . . . .	49
4.2 Thermal Measurements . . . . .	50
4.2.1 Oven Measurements . . . . .	51
4.2.2 Thermoelectric Measurements . . . . .	51
4.3 Safety Validation . . . . .	53
<b>5 Conclusions and Outlook</b>	<b>55</b>
5.1 Conclusions . . . . .	55
5.2 Outlook . . . . .	56
<b>References</b>	<b>59</b>

# List of Figures

2.1	Heat Transfer Mechanisms . . . . .	6
2.2	Thermoelectric Module . . . . .	9
2.3	Kelvin Connection . . . . .	12
2.4	Shunt Resistor . . . . .	13
2.5	Common Collector . . . . .	15
2.6	Common Emitter . . . . .	15
2.7	Common Base . . . . .	16
2.8	Class A Amplifier . . . . .	17
2.9	Class B Amplifier . . . . .	17
2.10	Class AB Amplifier . . . . .	18
2.11	Diode Bridge Circuit Diagram . . . . .	19
2.12	Diode Bridge Simulation . . . . .	19
2.13	Soft Start Circuit . . . . .	20
2.14	Soft Start Simulation . . . . .	20
2.15	Capacitance Multiplier . . . . .	21
2.16	Capacitance Multiplier Simulation . . . . .	22
2.17	PID Controller . . . . .	26
3.1	Hardware Overview . . . . .	27
3.2	Hardware Overview . . . . .	28
3.3	Linear Regulator PCB . . . . .	29
3.4	Supply Connector PCB . . . . .	29
3.5	Rectifier PCB . . . . .	30
3.6	Fan Controller PCB . . . . .	31
3.7	Backpanel PCB . . . . .	32
3.8	Amplifier PCB . . . . .	33
3.9	Amplifier Circuit . . . . .	35
3.10	Controller PCB . . . . .	37
3.11	State Machine . . . . .	41
3.12	Display . . . . .	42
3.13	Oven . . . . .	44
3.14	Oven Resistance . . . . .	45
3.15	Thermoelectric Setup . . . . .	46
3.16	Thermoelectric Module Resistance . . . . .	47

4.1	Output Low Noise . . . . .	49
4.2	Output Noise . . . . .	50
4.3	Electric PID . . . . .	50
4.4	Oven PID . . . . .	51
4.5	Oven PID Zoomed . . . . .	51
4.6	Thermoelectric PID . . . . .	52

# List of Tables

1.1	Comparison of considered temperature controllers. . . . .	2
3.1	Overview of used interrupts with their interfaces and priorities. . . . .	40
4.1	Table Name . . . . .	52



# List of Acronyms

<b>AC</b>	Alternating Current
<b>ADC</b>	Analog-to-Digital Converter
<b>ARX</b>	AutoregRessive with eXogenous input
<b>BJT</b>	Bipolar Junction Transistor
<b>DAC</b>	Digital-to-Analog Converter
<b>DC</b>	Direct Current
<b>DMA</b>	Direct Memory Access
<b>DRDY</b>	Data Ready
<b>DUT</b>	Device Under Test
<b>GUI</b>	Graphical User Interface
<b>I2C</b>	Inter-Integrated Circuit
<b>IC</b>	Integrated Circuit
<b>ISR</b>	Interrupt Service Routine
<b>IRQ</b>	Interrupt Request
<b>LDO</b>	Low-DropOut
<b>LED</b>	Light-Emitting Diode
<b>LSB</b>	Least Significant Bit
<b>MOSFET</b>	Metal-Oxide-Semiconductor Field-Effect Transistor
<b>NTC</b>	Negative Temperature Coefficient
<b>OC</b>	Over Current
<b>PCB</b>	Printed Circuit Board
<b>PID</b>	Proportional–Integral–Derivative
<b>PMOS</b>	P-type Metal-Oxide Semiconductor

<b>PWM</b>	Pulse Width Modulation
<b>RTD</b>	Resistive Temperature Detector
<b>RTOS</b>	Real-Time Operating System
<b>SMPs</b>	Switch Mode Power Supply
<b>SNR</b>	Signal-to-Noise Ratio
<b>SPI</b>	Serial Peripheral Interface
<b>TIM</b>	Timer
<b>USB</b>	Universal Serial Bus

# 1 | Introduction

Most modern power electronics focuses on maximizing efficiency to reduce operational costs and overall energy consumption. While this is crucial for everyday electronic devices, such design choices often neglect another important aspect, such as the noise performance of the system [1].

In contrast, high-precision laboratory instruments prioritize low-noise over efficiency to achieve superior measurement accuracy. This has led to the development of a niche market where noise performance is paramount. However, such specialized devices are typically expensive, reflecting high development costs and low production volumes. Consequently, the range of commercially available products is limited, and suitable solutions are often either unavailable or prohibitively expensive [2], [3].

A similar challenge exists at the Institute of Microelectronics, where low-noise temperature controllers are frequently required. Commercial wafer probe stations used for semiconductor device characterization include dedicated temperature controllers. However, experimental or in-house systems often lack such equipment. Furthermore, the requirements for temperature control vary between setups, making it impractical and costly to maintain several different commercial systems. Consequently, developing a single versatile temperature controller capable of supporting multiple applications became a key objective.

Previous research at the Institute already addressed aspects of this problem, including the theses *Temperaturkontrolle für die Defekte Spektroskopie in Halbleiter Transistoren* by Matthias Kratzmann [4] and *Erweiterung des TDDS Temperaturreglers* by Christian Schleich [5]. These works provided valuable insights and a technical foundation for the controller developed in this thesis, but their scope was too narrow to be applied directly to newly arising experimental requirements.

Based on the requirements of current and planned experiments at the Institute, several key design parameters were defined for the implementation of the controller. It should support multiple configurations, including heating and cooling using a thermoelectric element powered at up to  $\pm 18$  V and  $\pm 6$  A, an oven operating at 70 V and 5.8 A and a wafer prober supplied at 65 V and 5 A. Each setup should accommodate at least two temperature sensors and achieve a control cycle frequency of 5 Hz, with a final temperature accuracy of  $\pm 10$  mK, while keeping the voltage noise of the power output below 35 mV<sub>rms</sub>.

Table 1.1: Comparison of considered temperature controllers.

Model	Voltage	Current	Frequency	Accuracy	Noise	Control	Sensors	Price
QuickCool QC-PC-PID-01 [6]	22 V	8 A	3 Hz	$\pm 1 \text{ K}$	–	PID	1	550 €
Meerstetter TEC-1167 [7]	$\pm 56 \text{ V}$	$\pm 25 \text{ A}$	90 Hz	$\pm 0.01 \text{ K}$	$0.12 \text{ V}_{\text{pp}}$	PID	1	850 €
Omron E5AN-H [8]	24 V	3 A	16.6 Hz	$\pm 1 \text{ K}$	–	PID	1	600 €
MB Technologies TC 1010	32 V	4 A	–	$\pm 0.5 \text{ K}$	–	–	1	–
Lake Shore Cryotronics Model 336 [9]	50 V	2 A	10 Hz	$\pm 3.6 \text{ mK}$	$0.12 \mu\text{A}_{\text{rms}}$	PID	4	4500 €

As shown in Table 1.1, only two of the evaluated devices meet the noise and accuracy requirements, but neither supports the required voltage range. Furthermore, all commercially available controllers rely exclusively on PID control, limiting their capability to implement more advanced control strategies.

To overcome these limitations, the decision was made to develop a dedicated in-house solution. This approach allows the design to meet the specific requirements of the Institute regarding voltage range, noise performance, and firmware flexibility while providing a transparent and extensible platform for implementing and testing advanced control algorithms aimed at improving control speed, electrical noise performance and accuracy.

The system should be operable through two interfaces: a Universal Serial Bus (USB) connection, allowing configuration and data logging via a computer, and a touchscreen enabling stand alone operation when only a constant temperature is required.

The hardware of the temperature controller is based on the design of Kratzmann's earlier system, but significantly extended. The output voltage and current ranges are increased, necessitating a complete redesign of the thermal management of the power amplifier. Additional safety features are incorporated to improve operational robustness. The firmware is developed to manage the hardware, support flexible controller development and enable system identification.

The electrical characteristics of the system are verified using oscilloscope measurements, while the functional performance of the temperature controller is evaluated using the finished hardware prototype, simultaneously controlling temperature and logging the physical system state. Two experimental setups currently require temperature control and for both, a suitable control algorithm is implemented on the new hardware platform. The first setup is an oven in which the packaged Device Under Test (DUT) is placed on a ceramic base and heated to the desired temperature to characterise its temperature dependent electrical behaviour. The second setup consists of two thermoelectric modules stacked on top of each other, with the DUT positioned between them. The modules are water cooled, and the entire assembly is located in a vacuum chamber to prevent moisture condensation or freezing at sub-zero temperatures.

## 1.1 Outline

The remainder of this thesis is organised as follows:

Chapter 2 – Theory introduces the theoretical background relevant to thermal modelling, sensor and actuator technologies and the fundamentals of amplifier circuit design. It also discusses software patterns for microcontroller firmware and provides an overview of commonly used control strategies, with a particular focus on PID control and the justification of model simplifications.

Chapter 3 – Method refines and reformulates the previously introduced requirements, so they can be applied during the development of the hardware, firmware and control algorithms. The chapter then explains the hardware and firmware architecture of the temperature controller, describing the power supply, sensor interfaces, safety circuits, USB communication, and touchscreen user interface as well as the design considerations related to low-noise operation. The firmware section covers the control loop implementation, safety supervision, configuration handling and modular experiment modes. Finally, the chapter discusses the rationale behind the implemented control algorithms and outlines the steps taken during their development.

Chapter 4 – Result describes the measurements conducted to validate the temperature controller using both the thermoelectric module and the resistive oven setup. The achieved temperature accuracy, settling behaviour, noise performance and the functioning of all safety mechanisms are analysed.

Chapter 5 – Conclusions and Outlook summarises the main findings of the thesis and provides an outlook on possible improvements and future developments of the temperature controller.



## 2 | Theory

This chapter introduces the theoretical foundations on which the design, implementation, and evaluation of this master's thesis are based. It provides the necessary background in thermal system modelling, electronic hardware design, embedded software structuring, and control engineering.

First, the fundamentals of thermal systems are introduced, which form the physical basis for the controller design. This is followed by key aspects of electronics and PCB design that are relevant for the hardware implementation. The software design section discusses recurring patterns for embedded firmware. Finally, the control design section outlines the mathematical framework used to develop the controller.

### 2.1 Fundamentals of Thermal Systems

The design of a thermal controller requires a sound understanding of the underlying physical processes. Without it, the governing equations of the plant cannot be derived consistently. This section therefore summarises the relevant principles of heat transfer, thermal energy storage, and electrically driven thermal actuators and sensors.

#### 2.1.1 Heat Transfer

Heat transfer requires a temperature difference between two bodies or within a medium. Thermal energy is then exchanged via three fundamental mechanisms: conduction, convection and radiation [10, Ch. 1–3].

##### Conduction

Conduction is the dominant mode of heat transfer in solids and in stationary fluids. On the microscopic level, thermal energy is transported by atomic or molecular interactions and lattice vibrations from regions of higher temperature to regions of lower temperature. If a solid is subjected to a temperature difference, heat will continue to flow until a thermal equilibrium is reached, as illustrated in Figure 2.1.

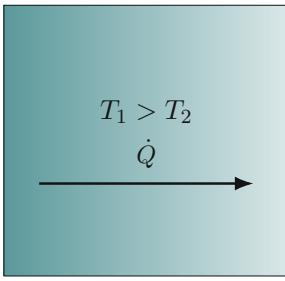
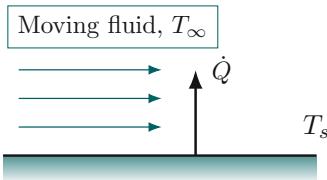
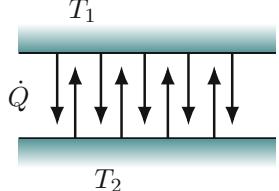
Conduction through a solid	Convection from a surface to a moving fluid or gas	Radiation heat exchange between two surfaces
$T_1$ 	$T_s > T_\infty$ 	

Figure 2.1: Fundamental heat transfer mechanisms: conduction, convection and radiation.  
Adapted from [10, p. 2].

For one-dimensional, steady-state heat conduction through a homogeneous slab of thickness  $L$  and cross-sectional area  $A$ , the heat flow can be expressed in the form of a thermal resistance:

$$R_{\text{cond}} = \frac{T_1 - T_2}{\dot{Q}} = \frac{L}{kA} \quad (2.1)$$

where  $R_{\text{cond}}$  is the thermal resistance due to conduction,  $T_1$  and  $T_2$  are the temperatures at the two surfaces,  $\dot{Q}$  is the heat transfer rate or power and  $k$  is the thermal conductivity of the material.

## Convection

Convection describes heat transfer between a solid surface and a moving or stationary fluid. It is typically composed of two effects: the molecular diffusion of energy near the surface and the bulk transport of energy by the fluid motion also called advection. In natural convection, the fluid motion is driven by buoyancy forces, for example when warm air rises to the top. In forced convection, pumps or fans impose the flow of the fluid, enhancing the heat transfer.

For engineering analysis, convection at a surface can be represented by a thermal resistance:

$$R_{\text{conv}} = \frac{T_s - T_a}{\dot{Q}} = \frac{1}{h_c A} \quad (2.2)$$

where  $R_{\text{conv}}$  is the thermal resistance due to convection,  $T_s$  is the surface temperature,  $T_a$  is the ambient (fluid) temperature,  $\dot{Q}$  is the heat transfer rate,  $A$  is the surface area, and  $h_c$  is the convective heat transfer coefficient, which depends on the fluid properties and the flow conditions.

## Radiation

Radiative heat transfer differs fundamentally from conduction and convection, as it does not require a material medium. It results from the emission of photons, which occurs for all matter at temperatures above absolute zero.

The emitted radiant energy flux of a surface is described by the Stefan–Boltzmann law:

$$P = \varepsilon \sigma T^4 \quad (2.3)$$

where  $P$  is the radiant power emitted per unit area,  $\varepsilon \in [0, 1]$  is the surface emissivity,  $\sigma = 5.67 \cdot 10^{-8} \frac{\text{W}}{\text{m}^2 \text{K}^4}$  is the Stefan–Boltzmann constant, and  $T$  is the absolute temperature of the surface.

When radiative exchange occurs between two surfaces, the process can be linearised and expressed in the form of an equivalent thermal resistance:

$$R_{\text{rad}} = \frac{T_s - T_a}{\dot{Q}} = \frac{1}{h_r A} \quad (2.4)$$

where  $R_{\text{rad}}$  is the thermal resistance due to radiation,  $T_s$  and  $T_a$  are the temperatures of the two surfaces,  $\dot{Q}$  is the net radiative heat transfer rate,  $A$  is the effective radiating area, and  $h_r$  is the linearised radiation heat transfer coefficient. The coefficient  $h_r$  is proportional to  $T^3$ , which follows from the original  $T^4$  dependence of the radiative exchange. Consequently, radiative effects become increasingly significant at higher temperatures.

### 2.1.2 Thermal Capacity

Matter not only transfers heat but also stores thermal energy, resulting in an increase of its internal temperature. When the spatial temperature gradients within a body are sufficiently small, the system can be approximated by the *lumped-capacitance model* [10, Ch. 5], [11, Ch. 3–4], [12, Ch. 10].

This approximation assumes a uniform temperature distribution within the object, which is valid when the Biot number,

$$B_i = \frac{h_c L}{k} \quad (2.5)$$

is much smaller than one. The Biot number compares the internal conductive resistance of the material to the external convective resistance and depends on the characteristic length  $L$ , the convective heat-transfer coefficient  $h_c$ , and the thermal conductivity  $k$ . It therefore quantifies whether internal temperature gradients can be neglected.

For a lumped thermal system, the thermal energy in the body relative to the ambient temperature can be written as

$$Q(t) = C_t (T(t) - T_a) \quad (2.6)$$

where  $Q(t)$  is the stored thermal energy,  $C_t$  is the thermal capacitance of the body,  $T(t)$  is the body temperature, and  $T_a$  is the ambient temperature. The rate of change of the stored energy is therefore  $\dot{Q} = C_t \frac{dT}{dt}$ .

Combining the definitions of thermal convection (Equation 2.2) and thermal capacitance (Equation 2.6), the energy balance of the system  $\dot{Q}_{\text{in}} - \dot{Q}_{\text{out}} = \frac{dE}{dt}$  can be expressed in terms of the required input power as

$$\dot{Q}_{\text{in}} = \frac{T(t) - T_a}{R_{\text{conv}}} + C_t \frac{dT}{dt} \quad (2.7)$$

where  $\dot{Q}_{\text{in}}$  is the thermal input power,  $R_{\text{conv}}$  is the convective thermal resistance, and  $\frac{dT}{dt}$  denotes the rate of temperature change. The first term on the right-hand side represents the steady-state heat loss to the environment, while the second term accounts for the transient energy stored in the material.

### 2.1.3 Thermal Actuators and Sensors

Temperature control requires active modulation of the thermal energy exchanged with the system. To ensure seamless integration between the digital controller and the physical plant, this work focuses exclusively on electrically driven thermal elements, which enable direct conversion from electrical power into thermal energy.

#### Resistive Heater

A resistive heater converts electrical energy into thermal energy through *Joule heating*. In an ideal resistor, the entire electrical input power is dissipated as heat, corresponding to a conversion efficiency of 100 % [13].

The generated thermal power is expressed as

$$P = I^2 R = \frac{U^2}{R} \quad (2.8)$$

where  $P$  denotes the thermal power,  $I$  the electric current,  $U$  the voltage, and  $R$  the electrical resistance of the heating element.

#### Thermoelectric Modules

Thermoelectric modules operate based on the *Peltier effect*, which describes the heat transfer occurring at the junction between two materials with dissimilar Seebeck coefficients. This effect arises because charge carriers transport not only electric charge but also entropy. The amount of heat carried per unit charge is expressed by the Seebeck coefficient of the material. When an electric current flows through the junction, heat is absorbed on one side and released on the other, resulting in simultaneous cooling and

heating, as illustrated in Figure 2.2. Modern thermoelectric devices exploit this principle by employing doped semiconductor materials such as bismuth telluride or silicon germanium, which significantly enhance efficiency [14].

This solid-state heat pump principle enables precise temperature control, fast dynamic response, and high reliability due to the absence of moving parts. However, the relatively low coefficient of performance compared to vapour-compression systems and the need for additional heat dissipation on the hot side restricts the applicability of thermoelectric modules to cases where compactness, reliability, or bidirectional control offer strategic advantages [15], [16].

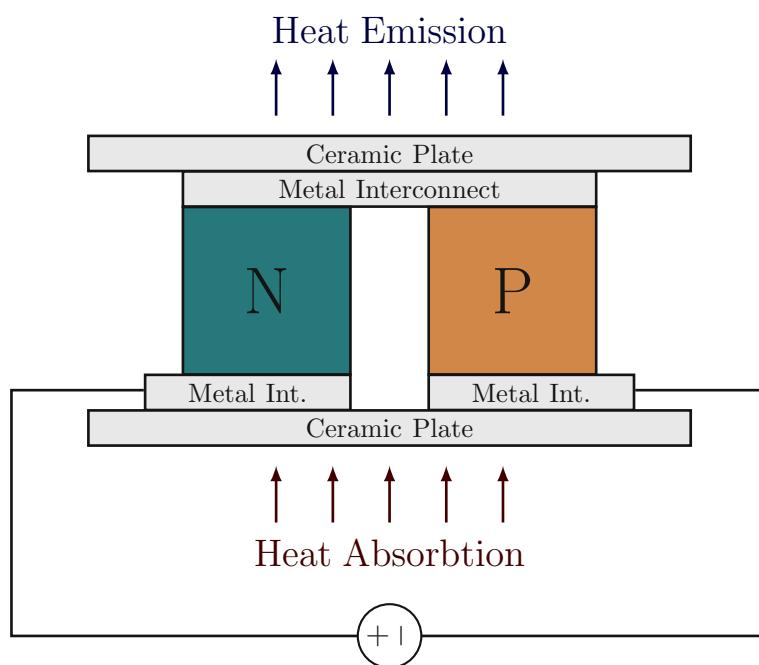


Figure 2.2: Each thermoelectric cell generates a heat flow, whose direction depends on the polarity of the applied voltage.

The heat transfer and electrical power relations of a thermoelectric module can be described by [17]:

$$\dot{Q}_c = S_M I T_c - \frac{1}{2} I^2 R_M - K_M (T_h - T_c) \quad (2.9)$$

$$\dot{Q}_h = S_M I T_c + \frac{1}{2} I^2 R_M - K_M (T_h - T_c) \quad (2.10)$$

$$P_M = \dot{Q}_h - \dot{Q}_c = S_M I (T_h - T_c) + I^2 R_M \quad (2.11)$$

where  $\dot{Q}_c$  and  $\dot{Q}_h$  denote the heat flow rates at the cold and hot sides, respectively,  $T_c$  and  $T_h$  are the corresponding temperatures,  $S_M$  is the Seebeck coefficient,  $I$  is the electrical current,  $R_M$  is the internal electrical resistance,  $K_M$  is the thermal conductance, and  $P_M$  is the electrical power input to the module.

## Temperature Sensors

Integrated-circuit temperature sensors are typically limited in their operating range by the characteristics of their internal semiconductor components and are therefore unsuitable for applications below  $-40^{\circ}\text{C}$  or above  $125^{\circ}\text{C}$ . To extend the measurable temperature range, three analog sensor types are commonly employed: Negative Temperature Coefficient (NTC) thermistors, platinum Resistive Temperature Detectors (RTDs) and thermocouples [18], [19].

**NTC Thermistor** A NTC Thermistor is a resistor composed of metal-oxide materials that exhibits a negative temperature coefficient, meaning its resistance decreases exponentially with increasing temperature. This strong nonlinearity allows for high sensitivity to small temperature variations. However, the resistance value drifts with component ageing, limiting both long-term stability and absolute accuracy, particularly over wide temperature ranges.

**Platinum RTD** A platinum RTD is a resistive temperature device whose sensing element consists of pure platinum, providing a nearly linear and precisely known temperature coefficient for resistance. Platinum RTDs cover a wide temperature range from approximately  $-200^{\circ}\text{C}$  to  $800^{\circ}\text{C}$ . For high-accuracy measurements, their resistance–temperature relationship can be modeled using a polynom of the third order. Due to their wide operating range, long-term stability, and high repeatability, platinum RTDs are considered superior to thermistors, albeit at a higher cost.

**Thermocouple** Thermocouple elements offer the widest measurement range among the discussed sensor types, covering temperatures from approximately  $-185^{\circ}\text{C}$  to  $2300^{\circ}\text{C}$ . They rely on the *Seebeck effect*, in which an electromotive force is generated by a temperature difference between two dissimilar materials. The resulting voltage is typically in the mV range, necessitating precise amplification and compensation techniques for accurate measurement.

## Cooling Techniques for Power Electronics

Power electronic devices such as transistors and diodes often dissipate more heat than can be naturally conducted away through their packaging. Therefore, additional cooling mechanisms are required to maintain safe operating temperatures and ensure long-term reliability. Depending on the power density and system constraints, passive heat sinks, forced-air cooling, or liquid cooling systems can be employed, each relying primarily on thermal conduction and convection to remove excess heat [10, Ch. 7, 11], [11, Ch. 7, 14].

**Passive heat sinks** Passive heat sinks increase the effective thermal capacitance of the device through their added mass and reduce the overall thermal resistance to the ambient environment by providing a larger surface area for heat exchange. To enhance this effect, most heat sinks feature fins or slits to increase the exposed surface area and promote airflow.

**Forced-air cooling** Forced-air cooling, achieved by integrating fans with heat sinks, initially provides similar thermal performance as passive cooling. However, as air velocity increases beyond a critical threshold, the resulting turbulent flow disrupts or thins the thermal boundary layer at the heat sink surface, significantly improving convective heat transfer. The downside of this approach lies in the added mechanical complexity, the need for fan-speed control, and potential reliability concerns associated with moving parts. It also adds electrical noise, which can be mitigated through galvanic separation between the high-precision circuitry and the fan supply.

**Liquid cooling** Liquid cooling systems use water or dielectric fluids as the working medium instead of air. Owing to their higher specific heat capacity and thermal conductivity, liquids enable much higher heat removal efficiency, allowing more compact thermal designs. The coolant is circulated through a closed loop using a pump and may include a reservoir to increase thermal mass. The absorbed heat is transferred to a radiator, where it is dissipated into the environment. Although highly effective, liquid cooling systems are more complex, costly, and require careful isolation to avoid leakage and electrical hazards.

## 2.2 Electronics and PCB Design

The design of a Printed Circuit Board (PCB) requires careful planning and attention to detail. This section focuses on two key aspects of the layout process: critical safety considerations and advanced circuit concepts relevant to the final prototype. A comprehensive discussion of PCB design practices lies beyond the scope of this work, therefore only topics directly related to the presented system are addressed. Additional layout considerations are introduced in the respective chapters where they serve to clarify specific design decisions.

### 2.2.1 Electrical Safety Requirements

According to the Austrian standards for protection against electric shock, the prototype operating at 100 V AC or 140 V DC falls within the category of low-voltage equipment. This classification covers devices with nominal voltages below 1000 V AC or 1500 V DC, but above the safety voltage limits of 50 V AC or 120 V DC. Low-voltage devices are not considered safe to touch and therefore require appropriate protective measures during

design, assembly, and testing. In accordance with the applicable standards, protection against electric shock can be achieved by enclosing conductive parts within a protective grounded housing, and ensuring that all user-accessible outputs remain safe to touch or are automatically disconnected in the event of a fault [20].

Furthermore, the PCB and connector design must provide sufficient clearance between conductors at different potentials to prevent dielectric breakdown or leakage currents. For equipment operating between 30 V DC to 300 V DC, a minimum trace clearance distance of 0.8 mm is required [21].

## 2.2.2 Measurement Circuits

Voltage, current, and temperature are fundamental quantities used to characterize any electrical system. Although their measurement may appear straightforward, achieving high accuracy or bandwidth can be challenging. Several key design considerations can significantly improve the performance of measurement circuits [22, Ch. 3].

### Voltage Measurement

Operational amplifiers inherently measure the differential voltage between their input terminals and are therefore well suited for voltage measurements. To achieve high measurement accuracy, the primary error sources must be carefully addressed during circuit design.

The first source of errors arises from non-ideal amplifier characteristics, most notably the input offset voltage and voltage noise. The offset can typically be compensated digitally by subtracting the mean offset value from the measured signal. Voltage noise, on the other hand, cannot be completely removed and must instead be minimized by selecting amplifiers with a low input noise density and by keeping gain and feedback resistor values as small as practical for the given application.

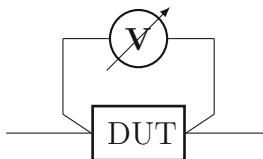


Figure 2.3: Kelvin connection setup used to measure the voltage directly at the DUT terminals.

A second major error source originates from the resistance of the signal path. When a significant current flows through a trace, even small parasitic resistances can introduce a noticeable voltage drop along the conductor. This can be mitigated by separating the current-carrying path from the voltage-sensing path. Since the voltage measurement circuit draws only a negligible current, typically in the microampere range, the associated

voltage drop along the sense lines is minimal. This configuration, known as a *Kelvin connection*, is illustrated in Figure 2.3.

## Current Measurement

Several approaches exist for measuring current in an electrical circuit. Although their implementations differ, all electrical current measurements must be connected in series with the desired current path. Magnetic current sensors can operate in parallel but generally offer inferior Signal-to-Noise Ratio (SNR) and are limited to measuring only the AC component of the signal. For this reason, magnetic methods are not considered in this work.

Electrical measurement techniques can be broadly categorized into two groups: *active* and *passive* methods. Active methods use operational amplifiers to sense or sink the measured current and can achieve higher bandwidth, but their implementation becomes increasingly complex and limited for high-current applications. Passive methods, on the other hand, determine the current indirectly by measuring the voltage drop across a known series resistance, commonly referred to as a shunt resistor. Due to their simplicity, robustness, and suitability for a wide current range, passive shunt-based methods are employed in this design [23], [24].

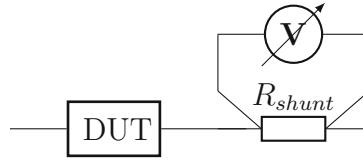


Figure 2.4: Shunt resistor connected in series to measure the current through the DUT.

In this configuration, shown in Figure 2.4, the measured current is proportional to the voltage drop across the shunt resistor according to Ohm's law:

$$R = \frac{U}{I} \quad (2.12)$$

where  $I$  is the measured current,  $U$  is the voltage drop across the resistor, and  $R$  is the resistance value.

## Temperature Measurement

For the RTD temperature sensors discussed in subsubsection 2.1.3, temperature variations cause corresponding changes in the electrical resistance of the sensor. Accurate temperature measurements therefore depend on precise resistance determination, which combines both current reference and voltage measurements according to Ohm's law (Equation 2.12).

To achieve high measurement accuracy, several factors must be considered. First, the excitation current through the sensor should be sufficiently large to maximize the SNR, but small enough to avoid self-heating of the sensing element. Second, the voltage across the sensor should be measured using a Kelvin connection, ensuring that the sensing leads are electrically isolated from the current-carrying leads. Finally, reference current sources typically exhibit drift over temperature and time. To compensate for this, the voltage across a precision reference resistor connected in series with the sensor is measured simultaneously. Any variation in the excitation current will then affect both measurements equally and can be eliminated by ratio-metric evaluation [25].

For thermocouple elements, a high-accuracy voltage measurement is sufficient, as the temperature-dependent electromotive force directly represents the temperature [26].

### 2.2.3 Power Supply and Rectification

Rectifier circuits are fundamental components of virtually all modern power supplies. Several design aspects must be carefully considered in high-power applications, including noise performance, voltage and current ratings, and thermal management. The simplest rectifier implementation consists of a diode bridge followed by a smoothing capacitor. Additional elements such as filters or current limiters may be added to enhance performance or safety. To ensure a low-noise Direct Current (DC) output, many systems ultimately employ a Low-DropOut (LDO) regulator [27, Ch. 9].

### 2.2.4 Amplifier Design

The output stages of typical signal-generator Integrated Circuits (ICs) are not capable of maintaining signal accuracy under load conditions. Therefore, amplifier circuits are employed to adapt the voltage signal to the required amplitude range and to provide sufficient output current for larger loads. A basic amplifier consists of at least one transistor, although the overall complexity may vary significantly depending on the load characteristics, temperature stability, and precision requirements. Regardless of the specific amplifier topology, Bipolar Junction Transistors (BJTs) and Metal-Oxide-Semiconductor Field-Effect Transistors (MOSFETs) are the two most common transistor types. While MOSFETs offer higher energy efficiency and better scalability, BJTs provide superior linearity, making them particularly suitable for analog amplifier stages. Their advantages include higher transconductance, better device matching, a wider linear operating range, and lower  $1/f$  noise. Properties essential for high-fidelity signal amplification [28, Ch. 3], [29, Ch. 2], [30, p. 39].

#### Basic Transistor Circuits

A transistor fundamentally operates as a current- or voltage-controlled power switch, consisting of the three terminals base, collector, and emitter. From the perspective of

small signal behaviour, three basic amplifier configurations are distinguished. In the following, all examples refer to NPN transistors [27, Ch. 2], [31, Ch. 1].

**Common Collector (Emitter Follower)** The common collector circuit, shown in Figure 2.5, uses the base as the input, the collector is tied to the supply, and the output is taken from the emitter.

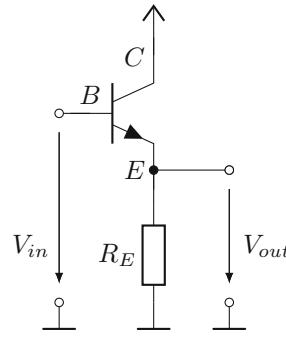


Figure 2.5: Circuit diagram of the common collector stage.

The voltage gain is approximately unity:

$$A_V = \frac{V_{\text{out}}}{V_{\text{in}}} \approx \frac{g_m R_E}{1 + g_m R_E} \approx 1 \quad (2.13)$$

Where  $g_m$  is the transistor transconductance and  $R_E$  the emitter resistor. The current gain equals the transistor current gain  $\beta$ , since  $I_E \approx I_C$ .

This configuration provides a high input impedance, a low output impedance, an excellent current amplification, but only modest voltage and power gain.

**Common Emitter** The common emitter amplifier, shown in Figure 2.6, uses the base as the input, the emitter is grounded, and the output is taken at the collector.

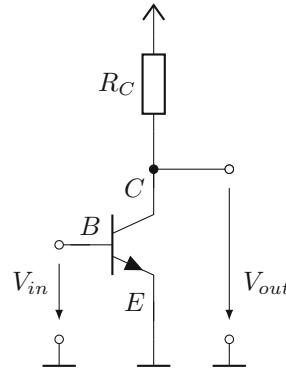


Figure 2.6: Circuit diagram of a common-emitter stage.

The small signal voltage gain for the common emitter circuit is:

$$A_V = \frac{V_{\text{out}}}{V_{\text{in}}} \approx -g_m(R_C \parallel r_{CE}) \approx -g_m R_C \quad (2.14)$$

Where  $R_C$  is the collector resistor and  $r_{CE}$  the intrinsic collector-emitter resistance. Since  $r_{CE} \gg R_C$ , this reduces to  $A_V \approx -g_m R_C$ .

This configuration provides a medium voltage gain, a high power gain, a low input and a high output impedance while also introducing a  $180^\circ$  phase shift.

**Common Base** The common base circuit, shown in Figure 2.7, holds the base terminal at the ground. The input is applied to the emitter and the output is taken from the collector.

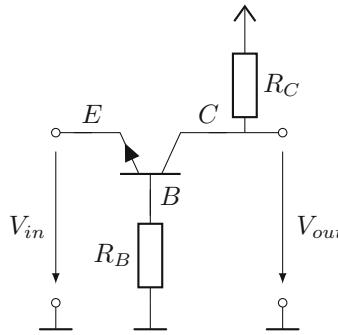


Figure 2.7: Circuit diagram of a common base stage.

The voltage gain is:

$$A_V = \frac{V_{\text{out}}}{V_{\text{in}}} \approx \frac{\beta R_C}{r_{BE} + R_B} \approx g_m R_C, \quad (2.15)$$

where  $r_{BE}$  is the intrinsic base-emitter resistance and  $R_B$  the bias resistor. Assuming  $r_{BE} \gg R_B$ , the final approximation follows from  $g_m = \beta/r_{BE}$ .

This configuration features a very low input impedance, a very high output impedance, a good voltage gain and excellent high-frequency performance due to the absence of Miller-effect feedback.

## Amplifier Classes

For linear amplifier design, several fundamental amplifier configurations can be identified. Each topology offers distinct advantages and limitations, typically involving trade-offs between power consumption, bandwidth, and signal distortion [27, Ch. 2], [31, CH. 2].

The **Class A Amplifier**, shown in Figure 2.8, consists of a single transistor and a load element connected at its output. The transistor is biased such that it remains in the conducting region throughout the entire input signal cycle. This operating region provides the highest linearity, resulting in minimal signal distortion.

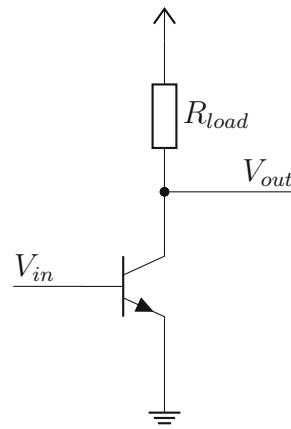


Figure 2.8: Circuit diagram of a class A amplifier, with a transistor and a load resistor.

The main disadvantage of this configuration is its poor efficiency, as the transistor continuously conducts current and therefore dissipates significant power even without an input signal. Depending on the bias point and load conditions, the theoretical efficiency ranges between  $\eta = 25\%$  and  $\eta = 50\%$ . Consequently, Class A amplifiers are typically limited to low-power applications, or cases where linearity and low distortion are prioritized over power efficiency. The continuous power dissipation also necessitates careful thermal management and adequate heat sinking.

The **Class B amplifier**, is shown in Figure 2.9, this configuration employs two transistors to reproduce the input signal at the output, the upper NPN transistor conducts during the positive signal, while the lower PNP transistor conducts during the negative signal. This push-pull arrangement significantly improves efficiency, typically reaching between  $\eta = 50\%$  and  $\eta = 78.5\%$ , depending on the operating point.

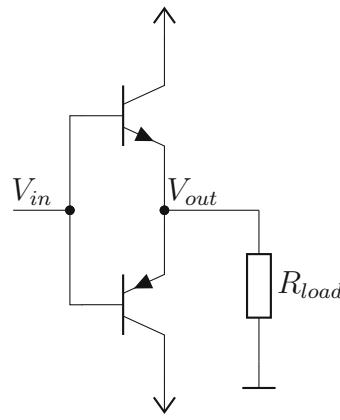


Figure 2.9: Circuit diagram of a Class B amplifier in complementary emitter follower configuration.

The main drawback of the Class B amplifier is the signal distortion in the region around 0 V, where both transistors are non-conductive. Within this region, the amplifier does not respond to small input signals until the base-emitter threshold voltage of either transistor

is exceeded. For sinusoidal inputs, this dead zone results in a segment of the waveform being clipped around zero, producing the characteristic crossover distortion. Even when conduction begins, the absence of biasing voltage causes a discontinuity in the output. Therefore, despite its higher efficiency, the pure Class B configuration is unsuitable for applications requiring high signal fidelity.

The **Class AB amplifier**, shown in Figure 2.10, closely resembles the Class B configuration in its conducting region. Its objective is to combine the high efficiency of Class B operation with the low distortion characteristic of a Class A amplifier. This is achieved by biasing both transistors slightly into conduction, typically through the use of diodes and resistors that establish a small quiescent current. As a result, both transistors remain partially on, even when the output voltage is near zero, allowing the input signal to transition smoothly through the 0 V region and thereby eliminating the crossover distortion present in Class B amplifiers.

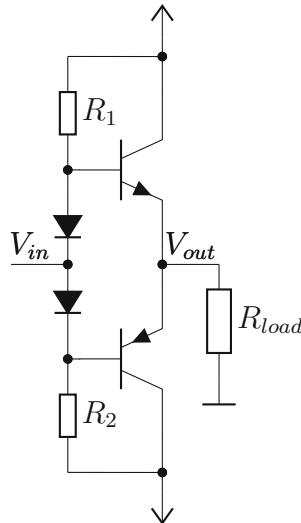


Figure 2.10: Circuit diagram of a Class AB amplifier in complementary emitter follower configuration.

The quiescent current determines the trade-off between efficiency and linearity. With minimal bias current, the amplifier behaves similarly to a class-B stage—highly efficient but with noticeable distortion. Whereas a large bias current approaches Class A behaviour, yielding excellent signal fidelity at the cost of reduced efficiency. Thus, the Class AB design enables a tunable compromise between power efficiency and signal quality, making it suitable for a wide range of analog applications.

## Diode Bridge

The Alternating Current (AC) input voltage of the power supply is first stepped down by a transformer and then rectified by a diode bridge, as shown in Figure 2.11.

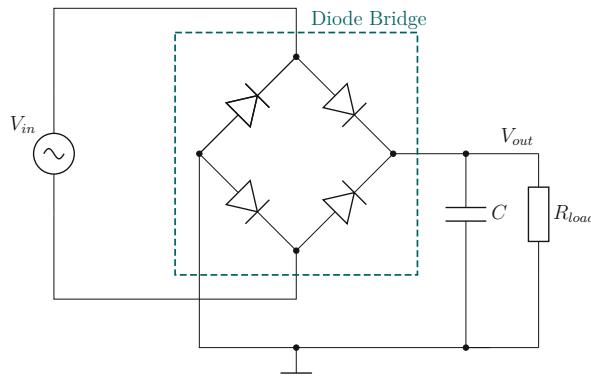


Figure 2.11: Circuit diagram of a full-wave diode bridge with input source, smoothing capacitor, and load resistor.

The diode bridge converts the alternating input voltage into its absolute value, effectively inverting the negative half-cycles of the sine wave. Each diode conducts only when the instantaneous AC voltage exceeds the forward voltage drop of the diodes, resulting in a slightly reduced peak amplitude of the rectified signal.

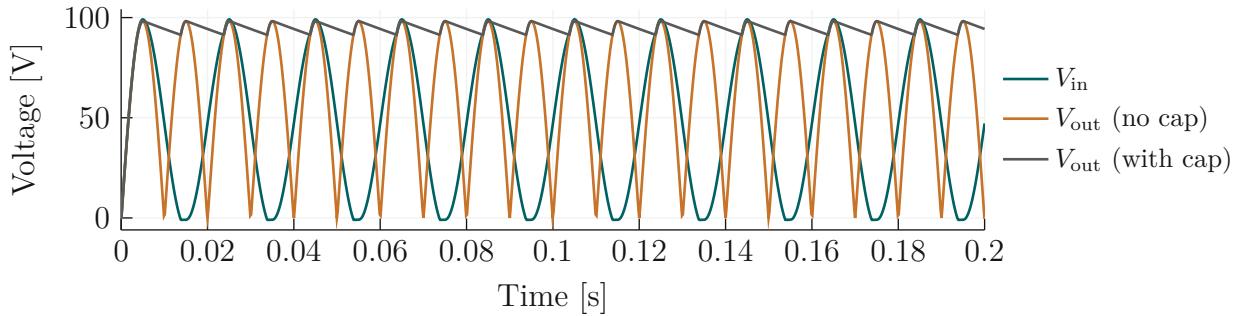


Figure 2.12: Simulation of rectifier input and output voltage obtained using LTspice.

A smoothing capacitor  $C$  stores charge during the voltage peaks and releases it between peaks, producing an approximately DC output voltage. When the load current is high, the capacitor discharges faster, leading to a voltage ripple during intervals where the instantaneous AC voltage falls below the capacitor voltage. This effect can be observed in Figure 2.12.

## Soft-Start Circuit

To limit the inrush current flowing into a large buffer capacitor during power-up, a soft-start circuit can be implemented, as illustrated in Figure 2.13.

At startup ( $t = 0$ ), the capacitor  $C_2$  initially behaves as a short circuit, causing the source and gate voltages of the P-type Metal-Oxide Semiconductor (PMOS) transistor to be equal and thus keeping the transistor in the off state. During the settle in time, the

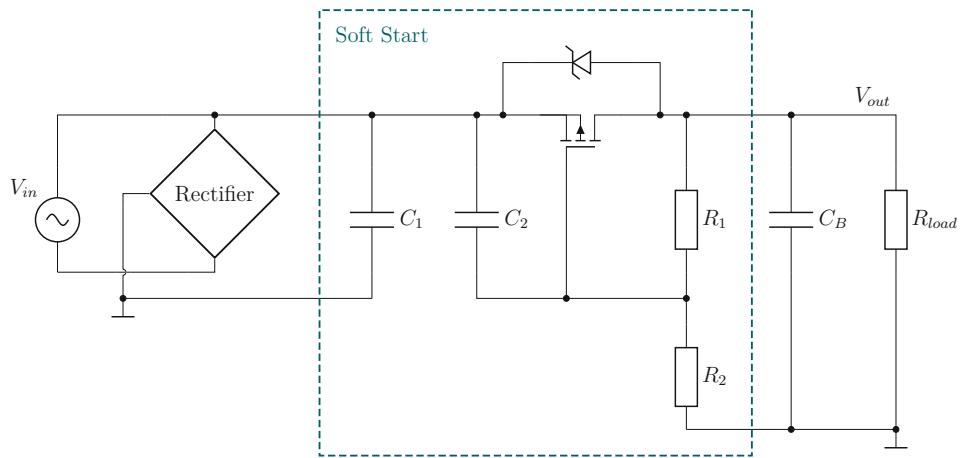


Figure 2.13: Soft-start circuit with rectifier input, buffer capacitor and load resistor at the output.

voltage across  $C_2$  and  $R_2$  gradually changes, slowly turning the transistor on and allowing the buffer capacitor  $C_B$  to charge in a controlled manner, thereby limiting the inrush current.

In steady-state operation,  $C_2$  behaves as an open circuit. The gate voltage of the transistor is then determined by the resistive divider of  $R_1$  and  $R_2$ , keeping the transistor fully on to maintain the charge of the buffer capacitor. The additional capacitor  $C_1$  stabilizes the gate voltage until  $C_B$  is fully charged.

This transient behaviour is illustrated in Figure 2.14, which shows the simulated source, drain, and gate voltages of the transistor during startup, as well as the charging current of the buffer capacitor.

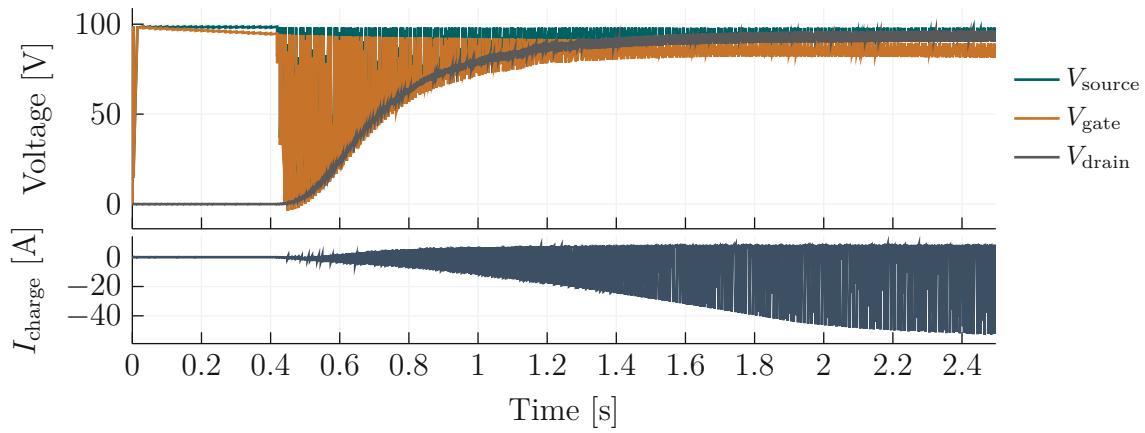


Figure 2.14: Simulated transient response of the rectifier output with soft-start circuit, showing limited inrush current during capacitor charging.

## Capacitance Multiplier Circuit with Feedback

To further reduce the voltage noise of a rectifier output, a capacitance multiplier is an effective method to emulate a large capacitance using smaller physical components and an active transistor circuit. This approach significantly reduces component size and cost, at the expense of increased power consumption in the control circuitry.

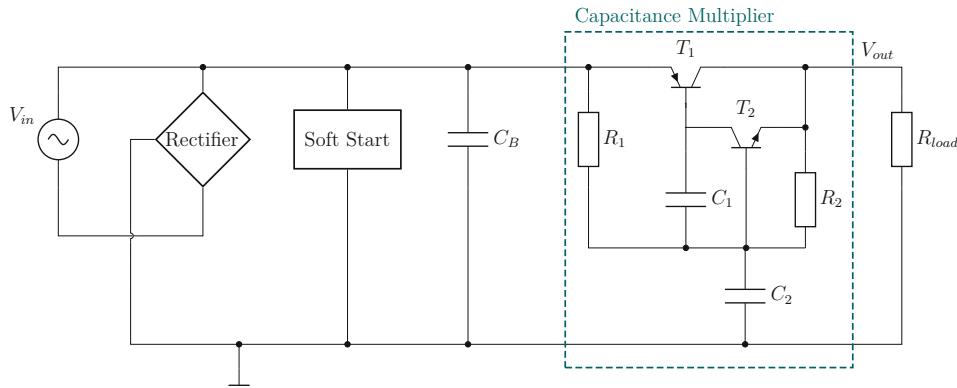


Figure 2.15: Capacitance multiplier integrated within the power circuit.

In the configuration shown in Figure 2.15, the BJTs  $T_1$  and  $T_2$  form a compound (*Sziklai*) pair, which operates similarly to a *Darlington pair* but employs transistors of opposite polarity, either NPN–PNP or PNP–NPN. Compared to the Darlington configuration, the Sziklai pair offers a lower combined threshold voltage, improved thermal stability and enhanced linearity, at the cost of a higher saturation voltage and reduced overall current gain. In this design, the NPN transistor acts as the driver and primarily defines the behaviour of the Sziklai pair [27, Ch. 2, Ch. 9], [32, Ch. 4].

Capacitor  $C_1$  stabilizes the Miller capacitance of the transistor pair, while  $C_2$  is the physical input capacitor whose effective value is multiplied by the transistor gain. The bias resistor  $R_1$  maintains the transistors in their linear operating region and allows adjustment of the bandwidth of the circuit. The feedback resistor  $R_2$  provides a small current from the output, further stabilizing the output voltage. Assuming ideal components, the effective capacitance can be approximated by

$$C_{\text{eff}} = \beta C \quad (2.16)$$

where  $\beta$  denotes the transistor current gain and  $C$  the base capacitor.

The simulated startup response in Figure 2.16 illustrates the voltage at the input and output of the capacitance multiplier, highlighting its noise-suppression capability.

## 2.3 Software Design

To ensure system safety and reduce hardware complexity, certain safety measures can be implemented in software. For the software to contribute reliably to overall system safety,

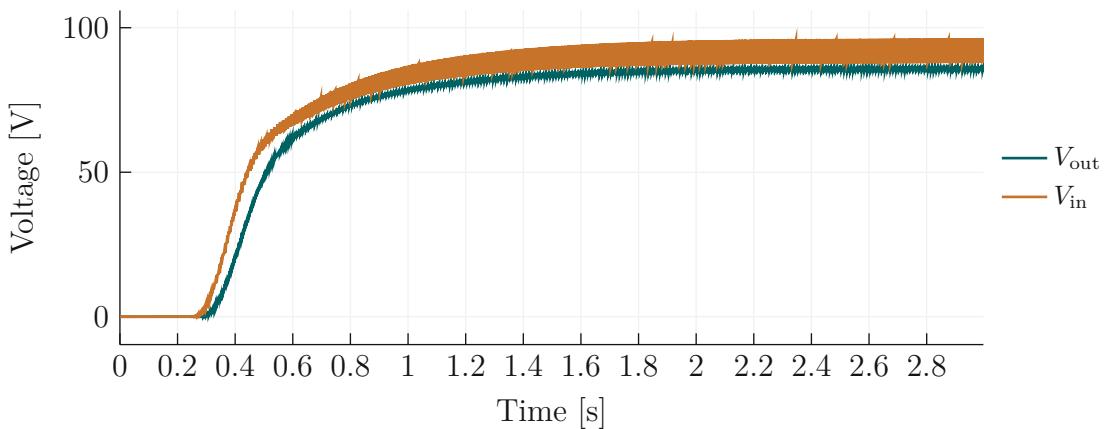


Figure 2.16: Simulated startup of the power supply with capacitance multiplier showing improved noise filtering.

it must meet specific timing and predictability requirements. The most critical of these is the ability to behave as a *real-time system*, a system that executes tasks within a fixed and guaranteed time frame. Whether this time frame is in the range of microseconds or seconds is secondary, the key requirement is that all tasks are completed deterministically and repeatedly within the defined interval. Typical general-purpose operating systems such as Linux, macOS, or Windows do not fulfil this requirement due to their asynchronous task scheduling, which favours performance over determinism. Although Real-Time Operating Systems (RTOSs) such as FreeRTOS exist, this thesis focuses on implementing real-time behaviour within a bare-metal firmware context [33].

To abstract recurring software design problems and provide systematic solutions, several *software design patterns* are introduced. A software pattern is a generalized template for solving a common class of problems rather than a specific implementation. Each pattern is defined by four key attributes: its *name*, the *purpose* it serves, the *solution* it provides, and the *consequences* both advantages and drawback of its use [34].

## Hardware Adapter Pattern

The hardware adapter pattern addresses the tight coupling between firmware and hardware in embedded systems. This coupling makes it difficult to replace or test alternative sensors and actuators. By isolating all hardware-specific code within dedicated hardware modules and exposing only a set of standardized public functions, such as `initialize()`, `configure()`, `get_data()`, and `access()` the pattern improves modularity and testability. While this abstraction introduces a minor execution-time overhead, it significantly reduces the likelihood of errors when modifying or extending the hardware interface.

## Channel Pattern

The channel pattern extends the hardware adapter concept by unifying both data access and data integrity management. It not only standardizes hardware interaction but also ensures that sensor readings are converted to physical units, verified through sanity checks, and error-handled if updates fail. This abstraction enhances fault detection and simplifies system-level diagnostics.

## Debouncing Pattern

The debouncing pattern mitigates a common problem with many physical input devices, such as mechanical buttons or touch sensors. When a switch changes state, its mechanical contacts can produce unstable oscillations, known as bouncing that persist for several microseconds to milliseconds. Instead of responding immediately to such transitions, the software waits for a predefined stabilization period before sampling the input, ensuring reliable state detection even with inexpensive or noisy sensors.

## Polling Pattern

The polling pattern manages the timeliness of data acquisition by reading sensor values whenever they are required by the process. This ensures that the most recent measurement is always used, but at the cost of increased latency and reduced process execution speed. Moreover, for time-critical signals, polling may be unsuitable, as missing a single measurement could lead to unsafe behaviour.

## Interrupt Pattern

The interrupt pattern enables the system to react to external events with minimal delay. It requires both hardware and firmware support for interrupt-driven execution. In this pattern, the microcontroller suspends the current process to execute an Interrupt Service Routine (ISR) in response to a high-priority event. Interrupts may be handled sequentially or prioritized, allowing high-priority interrupts to preempt lower-priority ones. While this pattern provides rapid responsiveness, high frequency interrupts or poor prioritization can lock up system resources, starving low-priority tasks.

## State Machine Pattern

The state machine pattern organizes firmware behaviour into well-defined operational states and transitions. Each state implements a consistent set of interface functions governing its execution and exit conditions. Two main types can be distinguished: *OR*-type state machines, where only one state is active at a time and *AND*-type state machines,

where multiple concurrent states are possible. This structured approach simplifies control logic and enhances maintainability, particularly in systems with complex operational sequences.

### **Monitor-Actuator Pattern**

The monitor-actuator pattern ensures safe operation in critical control applications. In this configuration, each actuator is accompanied by an independent monitoring sensor that verifies its output. If abnormal behaviour is detected, corrective or protective actions can be initiated immediately. Although this approach increases system complexity, the added redundancy substantially improves fault tolerance and operational reliability.

### **Watchdog Pattern**

The watchdog pattern is a lightweight supervisory mechanism used to verify that a specific task completes within a defined time frame. If the task fails to report completion before the timeout expires, the watchdog triggers a system reset or recovery routine. This ensures compliance with hard real-time constraints and enhances system robustness, particularly when combined with critical sensor or actuator processes.

## **2.4 Control Design**

With the hardware and firmware completed, the actuators can now be precisely controlled, and the sensors continuously monitor the physical processes. However, for a usable and autonomous system, an appropriate control design is still required. Control design is independent of the specific physical implementation of sensors or actuators and instead focuses on the mathematical description of the physical system and the strategy to steer its parameters toward a desired state. It governs how the system responds to actuation and environmental disturbances, ensuring convergence to the target state as quickly and efficiently as possible while respecting actuator limitations. The control design thus has a decisive impact on the performance of the system, potentially improving efficiency, accuracy, stability, and dynamic response.

This section discusses the key methods applied in this work for system identification and control design. More advanced approaches, such as *Kalman filtering* and *Lyapunov-based control*, were considered conceptually. However, their practical implementation would require analytical efforts beyond the scope of this thesis and is therefore only discussed briefly in the outlook.

## 2.4.1 System Identification

System identification encompasses a broad range of methods used to derive mathematical models of dynamic systems. For the thermal systems analysed in this work, an autoregressive model with exogenous input (Autoregressive with eXogenous input (ARX)) is sufficient, allowing parameter estimation via the least-squares method. In this approach, the system output is expressed as a function of its input and unknown physical parameters. Each measurement contributes one equation to the overall system of equations, where the unknowns are the model parameters. For a model with three unknown parameters, at least three independent measurements would be required for the equation system to be solvable. In practice, many more measurements are taken to account for parameter variations and measurement noise, leading to an overdetermined system.

The least-squares method then computes an optimal parameter set that minimizes the mean squared error between the measured output and the model prediction. While the least-squares criterion may not always provide the most robust weighting for all error types, it guarantees that the obtained solution minimizes the mean squared error globally. Mathematically, this is achieved by differentiating the squared error function with respect to the parameters and setting the result to zero, which yields the global minimum for a quadratic cost function.

Despite its advantages, several limitations must be considered. First, the error function must be linear in the unknown parameters, otherwise the minimum found may be only local. Second, the model must accurately represent the relevant physical dependencies, if essential influences (such as ambient temperature) are omitted, the identified parameters will not generalize well, leading to systematic errors. Finally, the accuracy is affected by measurement noise, but in autoregressive systems, only input signals with a mean component cause lasting output effects. As a result, random noise such as white noise can often be averaged out through sufficient data collection [35].

## 2.4.2 PID Controller

The Proportional–Integral–Derivative (PID) controller is a widely used and easily tunable method for controlling linear systems. As shown in Figure 2.17, it consists of three terms that operate on the error signal, defined as the difference between the desired and the actual system output.

The proportional term multiplies the instantaneous error by the proportional gain  $K_P$ . While this term reduces the error quickly, it cannot eliminate it completely, since the controller output becomes zero when the error reaches zero. To compensate for steady-state errors, an integral term is introduced. The integral term accumulates the error over time and, multiplied by the integral gain  $K_I$ , produces a non-zero output even when the instantaneous error is zero, correcting for low-frequency or constant disturbances. Finally, the derivative term reacts to rapid changes in the error, proportional to its time derivative,

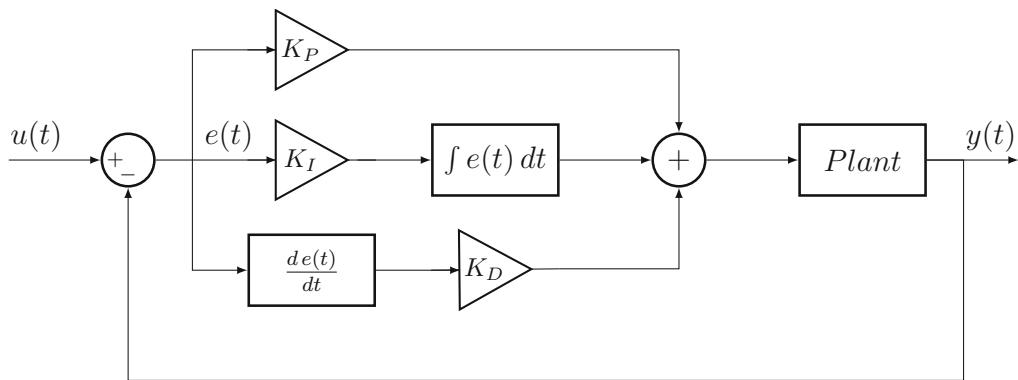


Figure 2.17: Block diagram of a PID controller showing the desired value  $u(t)$ , the error  $e(t)$ , and the output  $y(t)$ .

and is scaled by the derivative gain  $K_D$ . This term improves system stability and damping by counteracting fast error dynamics.

Each of these three components affects the closed-loop gain and phase of the system. Without proper tuning, the overall loop gain may remain greater than one, when the phase lag reaches  $180^\circ$ , causing instability. Careful parameter selection is therefore critical to ensure stability and optimal transient behaviour [36, Ch. 2, 3], [37, Ch. 6].

### 2.4.3 Singular Perturbation Theory

For non-linear dynamic systems, containing both fast and slow subsystems, singular perturbation theory provides a systematic way to reduce model complexity. The method assumes that the fast dynamics reach a quasi-steady state much quicker than the slow dynamics. Under this assumption, the fast subsystem can be replaced by its steady-state approximation, allowing the system to be described solely in terms of the slow variables. This reduced-order model remains valid as long as the transient effects of the fast subsystem are negligible [38, Ch. 4].

In practice, the validity of this approximation depends on the time-scale separation between the two subsystems. As a rule of thumb, the slow subsystem should be at least an order of magnitude, or at least 10 times slower than the fast subsystem to ensure accurate approximation and controllability.

# 3 | Method

This chapter describes the hardware and software implementation of the temperature controller, as well as the design of the electrical and thermal control algorithms that achieve the requirements defined in Chapter 1 – Introduction.

The temperature controller should provide an output voltage of  $-18\text{ V}$  to  $70\text{ V}$  with a voltage noise below  $35\text{ mV}_{\text{rms}}$ , a maximum current of  $\pm 6\text{ A}$ , a control cycle frequency of  $5\text{ Hz}$  and a final temperature accuracy of  $\pm 10\text{ mK}$ . For complex projects such as this, it is best to split the work into larger milestones that can be completed in isolation. Each milestone can have multiple goals and associated tests that verify its completion, while milestones themselves can either be developed in parallel or depend on each other.

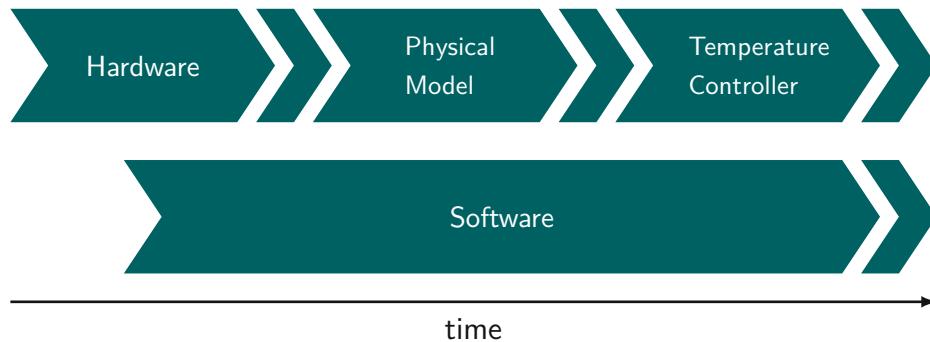


Figure 3.1: Hardware diagram to give an overview about the functions implemented.

The milestones are designed as shown in Figure 3.1. Since the desired electrical characteristics are already defined, the first milestone is the hardware system. The software to control the hardware is developed in parallel during the commissioning of the hardware, and once both are completed, the electrical characteristics can be verified. With the hardware available, the physical properties of the oven and the thermoelectric module can then be characterised. Using this data, the design of the control algorithms can be accelerated and later verified on the real systems.

## 3.1 Hardware

The hardware of the temperature controller is divided into several functional submodules, each with its own design goals and requirements, as shown in Figure 3.2. The system is

centred around two main modules: the amplifier, which forms the analog core and delivers the output power, and the microcontroller, which monitors all sensor inputs, supervises safety conditions and executes the control algorithms.

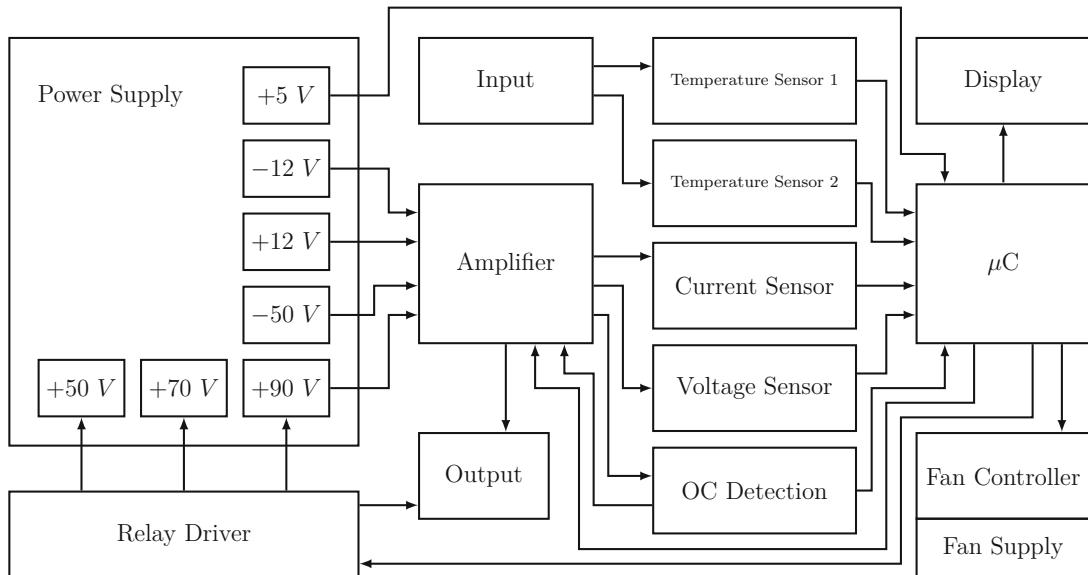


Figure 3.2: Functional overview of the hardware architecture, illustrating the power supplies, amplifier, sensors, safety circuits, relay driver, fan module, and microcontroller, as well as their interconnections.

Based on these two core modules, the remaining hardware can be grouped into three categories: analog, analog-to-digital and digital. The analog group includes the amplifier, all power supplies except the 5 V rail, and the electrical input and output interfaces. The analog-to-digital group contains the temperature sensors, the current sensor, the voltage sensor, and the Over Current (OC) detection circuitry. The digital group contains the microcontroller, the display, the fan controller, fan supply and the relay driver.

Although these are logical functional groups, the individual components may be distributed across several PCBs if this improves the overall performance of the system, for example in terms of thermal behaviour, noise performance or modularity. In the following sections, each PCB is discussed in detail, and all components are referenced according to their logical submodule.

### 3.1.1 Power Supply

The power supply consists of three subsystems: the low-voltage power supply, the high-voltage power supply and the fan power supply. Each subsystem is designed to meet the specific noise and efficiency requirements of the overall system.

## Low Voltage Power Supply

The low-voltage rails are supplied by a PS2127 linear power supply from OEP which provides output voltages of  $\pm 15$  V. These rails are fed into the linear regulator PCB shown in Figure 3.3 where they are converted into 5 V, 12 V and  $-12$  V. The conversion is performed by LDOs followed by an LC low-pass filter. The capacitors consists of an electrolyte capacitor for bulk energy storage and a foil capacitor to suppress high-frequency noise ensuring a low-noise supply for the analog circuits. For easier debugging, each output rail includes an indicator Light-Emitting Diode (LED) that shows whether the voltage is present and stable.

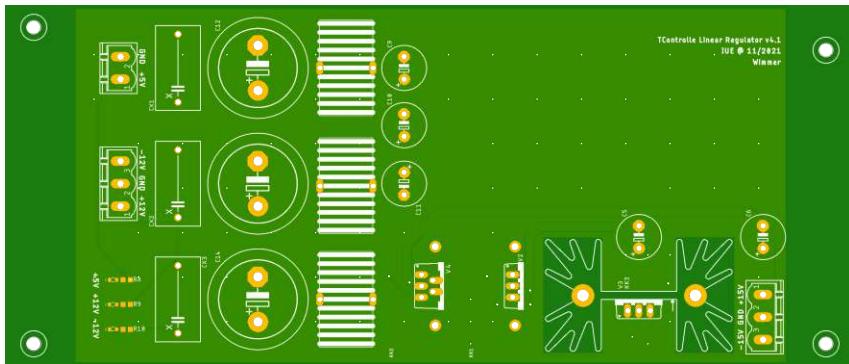


Figure 3.3: Unpopulated linear regulator PCB providing the low-voltage rails for the system.

## High Voltage Power Supply

The high-voltage power supply is generated by a custom-wound transformer designed specifically for this system. It provides four output pairs: two pairs of 35 V AC or 49.50 V<sub>pp</sub> and two pairs of 15 V AC or 21.21 V<sub>pp</sub>. The transformer is connected to the supply-connector PCB shown in Figure 3.4 where 10 A fuses protect it from short circuits.

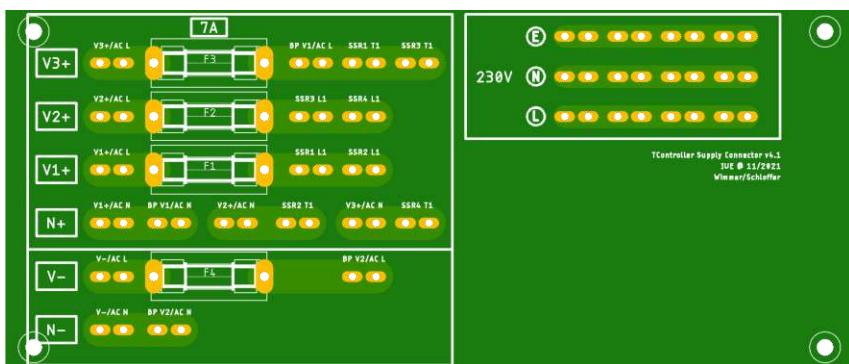


Figure 3.4: Unpopulated supply connector PCB, to connect the transformer input with the backpanel and the solid state relays.

The supply connector not only routes the transformer outputs to the backpanel, but also connects them to the solid-state relays. These relays switch between different transformer taps. In the basic configuration the rectifier receives  $49.50 \text{ V}_{\text{pp}}$ . If the load can tolerate higher voltages the positive rail can be increased to  $70.71 \text{ V}_{\text{pp}}$  or  $91.92 \text{ V}_{\text{pp}}$ .

The rectifier board shown in Figure 3.5 implements the diode bridge, the soft-start circuit and the capacitance multiplier described in Section 2.2 – Electronics and PCB Design. All semiconductors and passive components are selected to withstand the peak voltages present in the circuit and to carry currents exceeding the fuse rating. As specified in Subsection 2.2.1 – Electrical Safety Requirements, the trace clearance is at least 0.8 mm and the traces are designed to safely conduct 10 A. Early prototypes included additional fuses on the rectifier board, however these have become redundant due to the fuses on the supply-connector board and can be omitted in future revisions. The rectifier board connects to the backpanel through a 48-pin connector which also carries the high-voltage AC input.

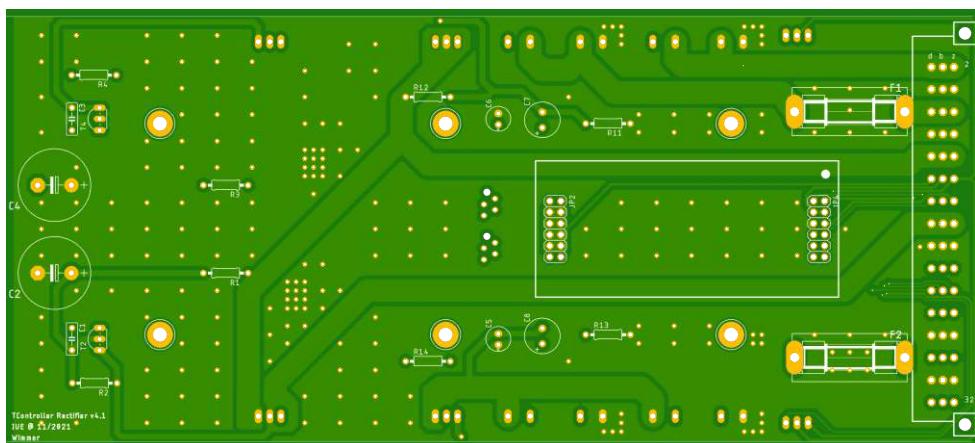


Figure 3.5: Unpopulated rectifier PCB implementing the diode bridge, soft-start circuit and capacitance multiplier.

The positive and negative supply rails are generated independently. The AC input is first rectified by the diode bridge and then passed to the soft-start circuit before charging the external 100 mF buffer capacitors. These capacitors are mounted externally due to their large size and are connected with the rectifier board through the backpanel. Because they can retain charge even after power off, additional safety features were implemented. Indicator LEDs on the backpanel show whether the capacitors are still charged and their series resistors ensure slow discharge without any external power. The output of the soft-start circuit also feeds the capacitance multiplier, which further reduces voltage ripple.

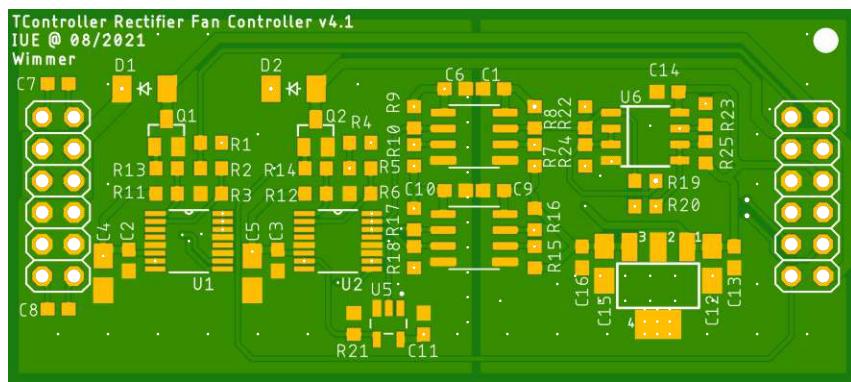
With the information from Section 2.1 – Fundamentals of Thermal Systems, the required heat sink for the rectifier board can be dimensioned to ensure that all components remain within their safe operating temperature. The highest thermal load occurs in the diodes and transistors that conduct the current of the high-voltage power supply. In the worst-case scenario, approximately 90 W of waste heat is generated on one side of the rectifier.

To keep the heat sink temperature below 80 °C, the required thermal resistance can be calculated as

$$R_t = \frac{T_{hs} - T_{amb}}{P_{in}} = \frac{80 - 20 \text{ °C}}{90 \text{ W}} = 0.66 \text{ K/W} \quad (3.1)$$

where  $R_t$  is the thermal resistance,  $T_{hs}$  the maximum allowed heat sink temperature,  $T_{amb}$  the ambient temperature and  $P_{in}$  the dissipated thermal power.

Any heat sink used in this stage therefore needs a thermal resistance of 0.66 K/W or lower. The LAM4DK10012 heat sink from Fischer Elektronik with a thermal resistance of 0.5 K/W under forced convection satisfies this requirement and was selected for the rectifier board.



### 3.1.2 Backpanel

The backpanel shown in Figure 3.7 serves as the central interface between the individual PCBs and allows the system to be developed in a modular way. It distributes the power rails, sensor signals and digital control lines to the amplifier, rectifier and controller boards.

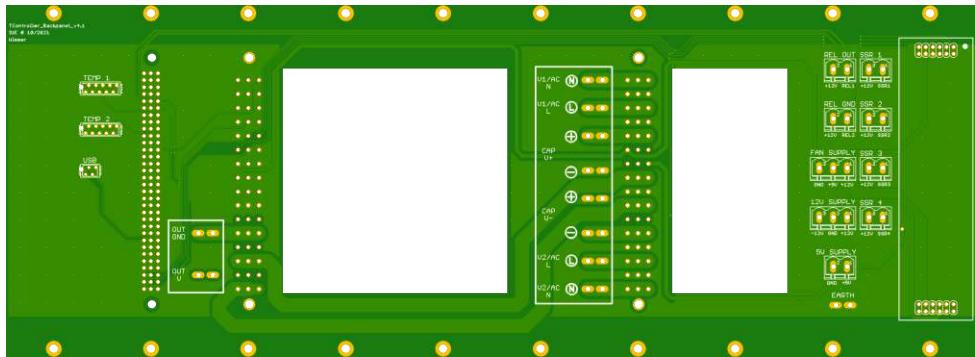


Figure 3.7: Bottom side of the unpopulated backpanel PCB which connects the power supply, rectifier, amplifier and controller boards.

On the front side the backpanel provides connectors for

- controller board
- amplifier board
- rectifier board

The bottom side provides access to

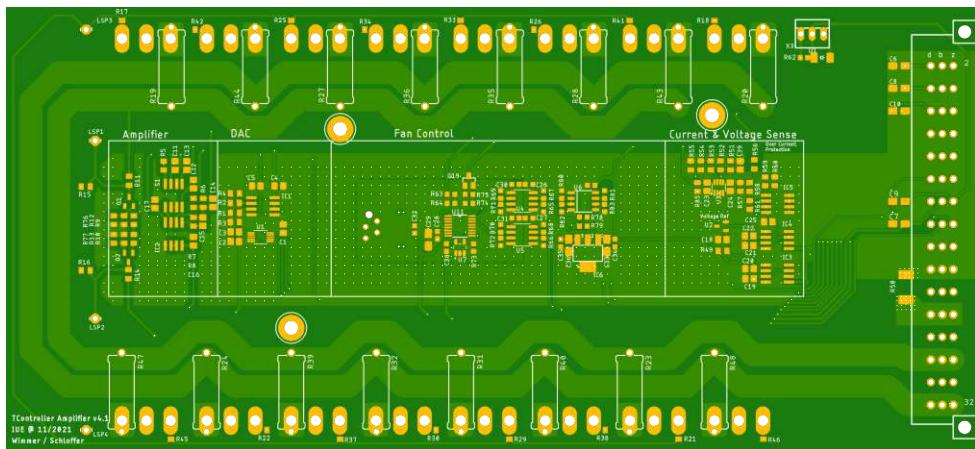
- USB connector
- relay control outputs
- two temperature sensor inputs
- high-voltage input
- buffer capacitor interface
- low-voltage supply input
- fan supply input
- additional fan driver module

Special attention was given to the high-current traces that route the high-voltage input, the connection to the buffer capacitors, the rectified supply from the rectifier board to the amplifier board and the amplifier output to the external connector. The minimum trace clearance is 0.8 mm and the traces are sized to safely conduct 10 A.

To support airflow across the heat sinks, two rectangular cutouts were placed in the backpanel as part of the thermal design.

### 3.1.3 Amplifier

The amplifier circuit shown in Figure 3.8 forms the analog core of the temperature controller and is the most complex subsystem. Its performance depends strongly on the noise of the high-voltage and low-voltage power supplies, which is critical for achieving a stable and precise output. The amplifier can generate a bipolar output voltage in the range of  $-30\text{ V}$  to  $70\text{ V}$  and delivers currents up to  $7\text{ A}$ .



## DAC

The AD5693R DAC was selected because it meets all requirements for this system. With a sampling rate of 25 kHz it is fast enough to support the intended control-loop frequency of at least 5 Hz. Its 16-bit resolution provides 65536 steps, which allows for fine control of the output voltage. The DAC outputs a unipolar 0 V to 5 V signal, which is converted to a bipolar  $\pm 2.5$  V control signal using a differential amplifier and the internal 2.5 V reference of the DAC.

A rough estimation can be used to demonstrate that the DAC resolution is adequate and to check the ideal theoretical accuracy of the system. This estimation provides a useful upper bound although it ignores several practical effects such as amplifier noise, Least Significant Bit (LSB) noise and thermal interactions with the environment, which all contribute significantly to the actual accuracy of the temperature controller. The amplifier is tuned for a maximum bipolar output of 100 V, resulting in a nominal LSB of about 3 mV. With a  $10\Omega$  resistive load this corresponds to a power resolution of about 0.9  $\mu$ W.

To evaluate what this means in practice, consider an aluminium cube with density  $\rho = 2700 \text{ kg/m}^3$ , specific heat capacity  $c = 900 \text{ J/(kg K)}$  and volume  $V = 1 \text{ cm}^3$  [10, Ch. 1]. The energy required to increase its temperature by 1 K is

$$E = \rho c V = 2700 \text{ kg/m}^3 \cdot 900 \text{ J/(kg K)} \cdot 1 \text{ cm}^3 = 2.43 \text{ J/K}$$

Applying only the minimum LSB power of 0.9  $\mu$ W yields

$$t = \frac{E}{P_{\text{LSB}}} \approx \frac{2.43 \text{ J}}{0.9 \mu\text{W}} \approx 31.25 \text{ days.}$$

This back of the envelope calculation shows that the effective resolution of the DAC is far higher than required for practical temperature control.

## Analog Amplifier

The core concept of the amplifier follows the Class AB topology described in Subsection 2.2.4 – Amplifier Design. Several extensions were added to meet the required voltage and current ratings and to ensure stable operation under high thermal load. These modifications transform the basic topology into a high-power stage.

To achieve the required current capability, the output stage uses multiple transistors connected in parallel. In total eight devices are used for the positive rail and eight for the negative rail and they are arranged symmetrically across the heat sink to ensure an even distribution of the thermal load.

Directly connecting the bipolar transistors in parallel, leads to uneven current sharing, because device tolerances introduce small variations in their operating points. A bipolar

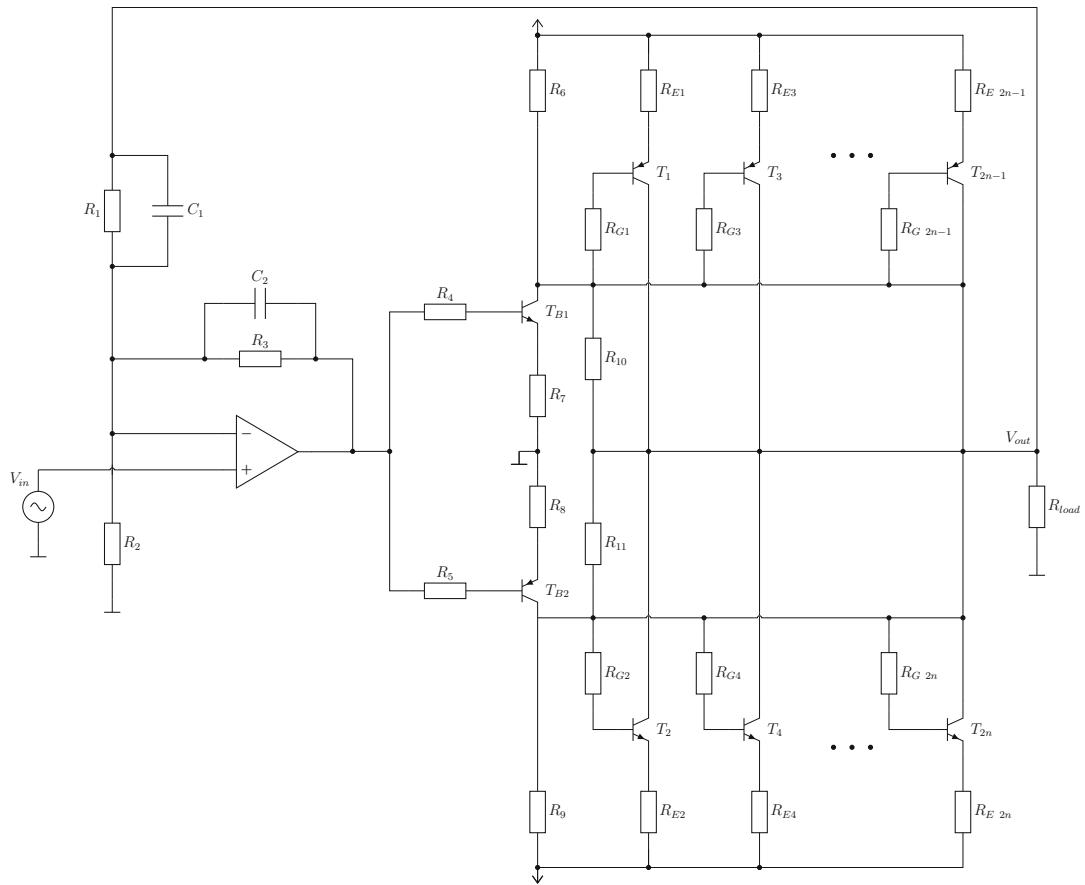


Figure 3.9: Amplifier circuit, with differential amplifier and transistor driver.

transistor that initially conducts slightly more current will heat up faster which in turn increases its conductance and draws even more current. This self reinforcing behaviour can ultimately lead to thermal runaway. To avoid this, each transistor is connected through a  $470\text{ m}\Omega$  emitter resistor that stabilises the operating point and equalises the current distribution across all devices.

Each output device is implemented as a Darlington pair using the MJH11021G and MJH11022G transistors. These devices support up to 250V and 15A and can dissipate up to 150W under ideal cooling conditions. The high voltage rating is essential for the negative side to withstand the full bipolar supply voltage, when the amplifier output is driven to the positive maximum and vice versa.

The output stage itself has a voltage gain close to one because it operates as a common-collector circuit. Therefore the required voltage swing must be generated in the preamplifier stage. This is realised with transistors  $T_{B1}$  and  $T_{B2}$  in a common-emitter configuration which provide the necessary voltage amplification.

The first analog stage is an OPA2277 operational amplifier in an inverting configuration with two feedback paths. A local high-frequency feedback loop stabilises the op-amp behaviour at higher frequencies. The dominant low-frequency feedback comes from the amplifier output node which determines the overall gain. The gain is set by the ratio of

$R_1$  and  $R_2$ . The low pass characteristic formed by  $R_1$  and  $C_1$  ensures loop stability, by making the open loop gain fall below unity before the phase shift reaches  $180^\circ$ .

## Analog Measurement

The measurement circuitry supervises the output of the amplifier and consists of both voltage and current monitoring. The output voltage cannot be measured directly with the Analog-to-Digital Converter (ADC) because it exceeds the allowable input range for a typical operational amplifier. A resistor divider scales the voltage down to a safe level, which is then buffered and level shifted by an operational amplifier before being converted into a digital signal on the controller board.

The current measurement is performed using a shunt resistor. The voltage drop across the resistor is amplified by the AD8418, a current-sense amplifier with a fixed gain of 20. Its reference input allows the output to be shifted so that the resulting signal remains within the 0 V to 2.4 V input range of the ADC on the controller board. The shunt resistor has a value of  $4\text{ m}\Omega$ . This value represents a compromise between measurement accuracy and power dissipation. A higher resistance would improve sensitivity, but it would also increase the power loss and therefore influence the measurement due to the temperature dependence of the shunt resistance.

## Over Current Detection

Relying solely on digital supervision of the output current is insufficient for ensuring system safety. The digital monitoring depends on the microprocessor and therefore only remains effective as long as the firmware is running correctly. If the microprocessor crashes, the amplifier is driven exclusively by the static output voltage of the DAC together with the feedback loop of the analog amplifier. In such a failure scenario, if the user connects a load that causes an over current condition, this fault would no longer be detected, at least not until the fuse breaks.

To address this issue, and especially to protect hardware during firmware development, an additional analog over current detection stage was implemented. The voltage monitor LTC2912 is used to detect a defined over-voltage or under-voltage threshold derived from the current-sense amplifier output, which converts the load current into a proportional voltage. The threshold values are chosen such that an over-current condition at 7.5 A triggers the monitor, pulling either the OV or UV open-drain output low. This newly generated digital fault signal is independent of the microprocessor and is further processed by the external logic on the controller board. The resulting signal is then used as an enable line for components that are critical to the safety response of the system.

On the amplifier board, this enable signal controls the analog switch DG419DY. Under normal operating conditions, the switch connects the amplifier input to the DAC output. If a fault is detected, it instead connects the input to ground, effectively forcing the amplifier output 0 V and ensuring a safe shutdown.

### 3.1.4 Digital Hardware

The controller board, shown in Figure 3.10, consists of the microcontroller, the safety logic, the relay driver, the ADCs, and the display connector. Separating the digital circuitry from the analog amplifier prevents digital noise from degrading the analog performance and allows the analog and digital parts of the system to be developed independently.

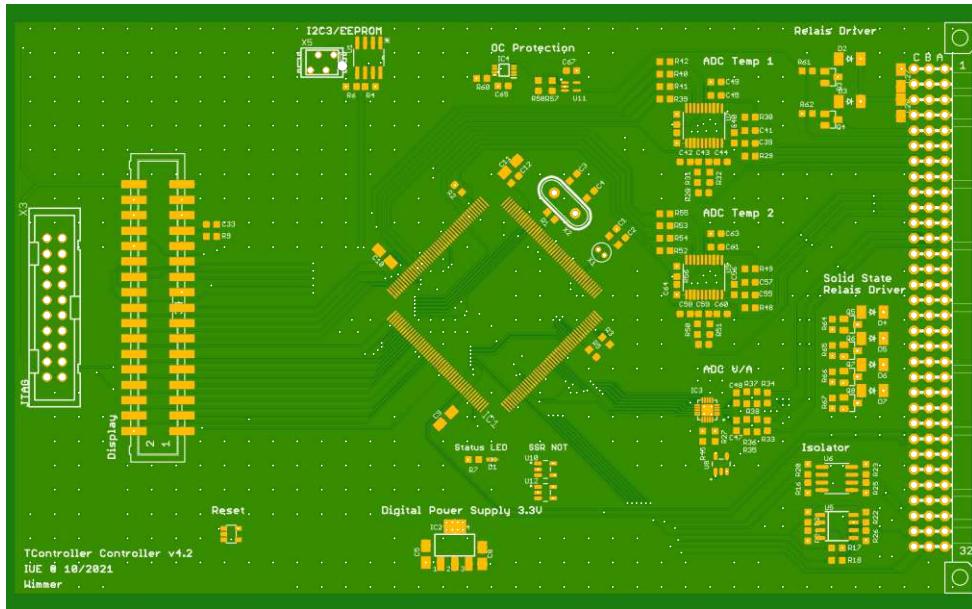


Figure 3.10: Unpopulated controller PCB, featuring the microcontroller, ADCs, relay driver, and display connector.

#### Microcontroller

The STM32F746IGT6 microcontroller operates at a clock frequency of 230 MHz, which is sufficient for the required control loop update rate of 5 Hz. It provides enough digital I/O pins to interface with the ADCs as well as with the display. In addition, this device is widely used in projects within the Institute of Microelectronics, meaning that many interfaces and firmware components are already available, which significantly accelerates development.

A more detailed description of the microcontroller firmware is given in Section 3.2 – Software. Only two digital communication lines are routed from the controller board to the backpanel, both using the Inter-Integrated Circuit (I2C) protocol. One bus programs the output voltage of the DAC that controls the amplifier, while the other bus is connected to all fan controllers in the system, enabling configuration and status monitoring. To store configuration data of the temperature controller, an EEPROM IC is also connected to the microcontroller via I2C.

## Safety Logic and Relay Driver

The external logic for over current detection consists of a flip-flop IC, ensuring that an over current event is latched even if it occurs only briefly. The flip-flop can be reset manually via the microcontroller or by performing a complete system restart. A second stage of the safety logic is an AND gate combining the over-current status with the enable signal. This ensures that both the over current detection and the microcontroller must agree before enabling the output and that either can disable it if necessary.

The enable signal is used not only for the analog switch on the amplifier board but also to control the output relays, that connect the amplifier output to the external connector on the enclosure. The High-voltage relays for switching the transformer inputs are controlled from the microcontroller. The digital control signal is amplified by a MOSFET in a common-source configuration, functionally similar to the common-emitter configuration of a BJT.

## Analog-to-Digital Converter

Up to the controller board, all measurement signals are analog and therefore cannot be stored or used directly by the digital controller. These signals must be converted into digital form.

The load current and voltage measurement signals, both scaled to a 0 V to 2.4 V range, are sampled by the ADS131M02, a two-channel, 24-bit sigma-delta ADC with a sampling rate of 64 kHz. The data is transmitted to the microcontroller via an Serial Peripheral Interface (SPI) interface.

The temperature measurement is implemented using both Kelvin connection and ratio-metric evaluation, as described in Subsection 2.2.2 – Measurement Circuits. These techniques are supported by the ADS1247 ADC, which provides 24-bit resolution and a sampling rate of 2 kHz, also using SPI. The analog sensor interface is connected with four wires to an RJ45 input, supporting both Pt100 and Pt1000 temperature sensors.

## Display

To enable stand-alone operation without an USB connection, the HY32D touchscreen display from Waveshare was integrated as a user interface. The display uses a parallel interface for pixel data and provides an interrupt line for signal touch events to the microcontroller. The touch position is then read via I2C from the capacitive touch controller.

## 3.2 Software

The software of the temperature controller can be divided into two principal components, each with distinct requirements and therefore different implementation strategies. The microcontroller firmware, written in C, is responsible for the complete management of the hardware, including all safety relevant functions. This guarantees that no user command can place the system in an unsafe state. Communication with an external computer is implemented via USB, using C on the microcontroller side and Python on the host side. This approach enables the development of flexible Python applications that interact with the microcontroller, while leveraging the computational capabilities and rich library ecosystem of the connected computer.

The firmware architecture follows the design patterns introduced in Section 2.3 – Software Design. Hardware access is abstracted through a hardware adapter and channel pattern. For the microcontroller itself, these abstractions are provided by the STMicroelectronics HAL library. Each additional hardware component is encapsulated in a dedicated C module in which only the `init`, `configure`, `get_data`, and `set_data` functions are exposed as public interfaces.

To ensure deterministic handling of time-critical safety functions while simultaneously processing user inputs, an interrupt-driven architecture is employed. The assigned interrupt priorities are summarized in Table 3.1.

The interfaces used are USB, I2C, Interrupt Request (IRQ), Timer (TIM), Direct Memory Access (DMA), SPI and Data Ready (DRDY). USB, I2C and SPI provide communication with external devices, while TIM generates periodic interrupts based on the internal timer. DMA enables peripherals such as SPI and I2C to directly access memory, reducing processor load. DRDY signals that new ADC data are available, while IRQ is used for general interrupt requests.

The interrupt priorities reflect the relative importance of the tasks. Touch input is assigned the lowest priority. Each DRDY event initiates an SPI transaction, which then stores the result via DMA. To avoid blocking, each step requires a strictly higher priority than the next. Together with USB (Standard Mode), this determines the priorities of all data-acquisition tasks. The timer interrupts and the fan error interrupt form the next priority level, ensuring correct operation of the control loop and rapid detection of cooling system failures. Over current detection, fan controller communication, and output control are assigned higher priority to guarantee the fastest possible shutdown during fault conditions. The highest priority is reserved for USB (Error Mode), allowing developers to recover, reprogram, or reset the microcontroller in any state.

With the IO side of the firmware handled, the control loop of the firmware still need to be implemented, whose behaviour is defined by a state machine architecture described in Figure 3.11. After booting, the initialization routines configure the DAC, ADCs, EEPROM, USB, fan controller, over-current detection circuit and relay drivers. If all initialization steps are successful, the microcontroller transitions to the controller initialization state,

Table 3.1: Overview of used interrupts with their interfaces and priorities.

Interrupt	Priority	Interrupt Source
<b>USB</b>		
USB Error Mode	0	USB
USB Standard Mode	4	USB
<b>DAC / Safety</b>		
DAC Output	1	I2C
Fan Controller	1	I2C
Over Current Detection	1	IRQ
Measurement Timer	2	TIM
Safety Measurement Timer	2	TIM
Fan Controller Error	2	IRQ
<b>ADCs</b>		
ADC Current & Voltage	3	DMA
	4	SPI
	5	DRDY
ADC Temperature 1	3	DMA
	4	SPI
	5	DRDY
ADC Temperature 2	3	DMA
	4	SPI
	5	DRDY
<b>User Interface</b>		
Touchpanel	6	IRQ

where the control loop is configured, the external temperature sensors are detected, and the periodic timers are activated.

Upon successful completion, the firmware enters the operational state in the *Ready* sub-state. Three operational modes: *Resistor*, *Oven* and *Peltier* can be selected. Each mode features its own configuration and may impose restrictions on the temperature control, including full blocking it. The *Error* state is entered whenever the firmware detects an irrecoverable condition. The system can only exit this state through a restart, triggered either by the firmware or the physical power switch. For development and diagnostic purposes, a *Debug* state is available. It is accessible only after controller initialization and exclusively through a firmware update. This prevents users from entering a state in which software safety supervision is disabled.

The measurement timer is executed every 1ms. It retrieves the latest data from the ADCs, processes the measurements using the control algorithms, and updates the volt-

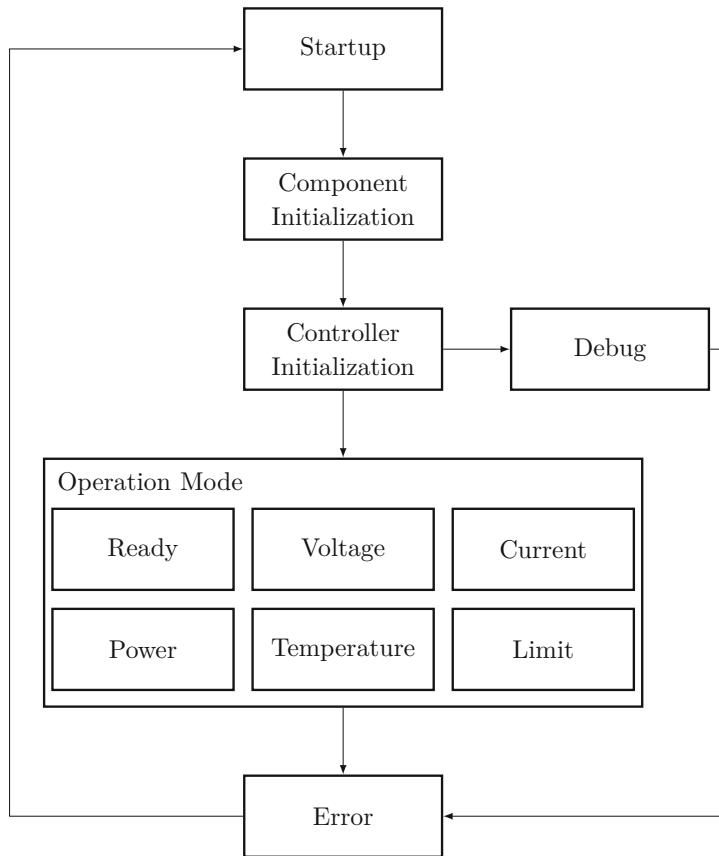


Figure 3.11: State diagram showing all defined firmware states and their transitions.

age command for the amplifier accordingly. A second timer, the safety measurement timer, is executed every 1s. It monitors internal heat sink temperatures and resets certain overshoot-detection variables, effectively filtering short-duration transients using a debouncing mechanism. For each operational mode, individual limits for voltage, current, external temperature, and internal temperature can be configured. This enables the safe operation of systems that cannot tolerate the full output range of the temperature controller.

The Graphical User Interface (GUI) runs in the main loop and therefore has the lowest priority. It is executed only when no other tasks are pending. A representative screen is shown in Figure 3.12. The central display area indicates the current mode, and the set point can be modified using the increment and decrement buttons. Below, the measurement values are shown, the modes can be changed using the arrow keys or the tab button selection.

The touchscreen interrupt handler implements a debouncing mechanism to prevent multiple unintended touch events from being registered during a single interaction.

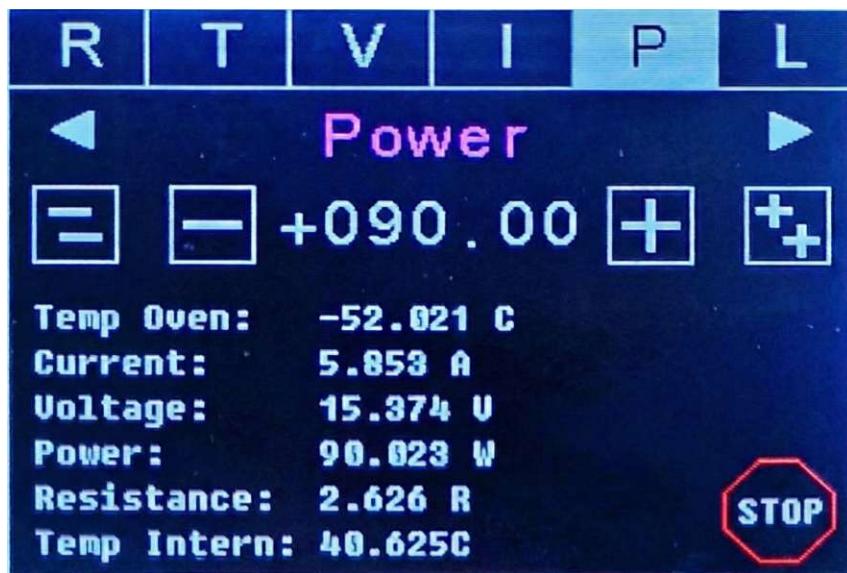


Figure 3.12: Prototype display image. Reflections were reduced and clarity improved using Gemini (Nano Banana Pro).

### 3.3 Controller

With both hardware and software validated, the system can reliably control its output voltage. Additional control objectives can be achieved by incorporating sensor data or model-based parameters, thereby extending the functionality beyond simple voltage regulation. The controllers are organised into two groups, the *electrical controllers* and the *temperature controllers*.

#### 3.3.1 Electrical Controller

Operating the amplifier in pure feed-forward mode, where the output voltage is set directly by the DAC, is sufficient for basic operation. However, to achieve higher accuracy, the output must be corrected for component tolerances such as offset voltages in the operational amplifiers or variations in resistor values. Instead of manually calibrating the entire analog signal chain, this error can be automatically compensated by a PID controller that uses the measured output voltage as feedback. In this configuration, only the sensors require calibration, while the controller compensates for linear deviations during operation. The trade off is an increase in computational complexity and a slower step response compared to a fully calibrated feed-forward controller, but the response remains sufficiently fast for temperature control applications.

Using the system as a current source introduces an additional dependency, the output current is determined by the resistive value of the connected load. Attempting to regulate current directly with a universal PID controller would require the controller to be tuned for the smallest expected load resistance to avoid overshooting, resulting in a very slow

response for larger resistances. A more efficient and load-independent approach is achieved by using a resistive model to compute the required output voltage from the desired output current. The system then regulates voltage rather than current. This model-based control is valid whenever the load behaviour is dominated by its resistive component. To ensure correct operation, the load resistance must be known. The firmware therefore estimates the load resistance during controller initialization, and a default value can be calibrated. For output voltages above 0.1 V, the system continuously updates the estimated resistance using the measured voltage and current values.

Power control exhibits the same dependency on the load resistance, but the nonlinearity is more pronounced because electrical power scales quadratically with the voltage. Using a Joule heating model, the system computes the voltage required to achieve a desired power output. Linearisation of this model allows a PID controller to be used to correct runtime deviations. The Resistance handling is identical to the current controller. In practice, continuously updating the resistance estimation is particularly advantageous due to the temperature coefficient of most resistive loads.

The limit mode combines the voltage and current controllers, while imposing user defined limits on both quantities. For example, applying 10 V to a  $1\Omega$  load without limits would cause the output to rise to 7.5 V, triggering the over current protection and forcing the system into an error state. In limit mode, the user may specify independent voltage and current limits. For example 10 V and 1 A, the controller increases the output voltage until the current reaches its limit, at which point it transitions to current control. Limit enforcement also works in the opposite direction, with limits of 0.5 V and 1 A, the controller will stop at 0.5 V even if the resulting current is only 0.5 A.

### 3.3.2 Temperature Controller

In contrast to the electrical controllers, which exhibit fast dynamics and can therefore be tuned by trial and error, temperature controllers require a model-based approach. The oven setup in particular shows extremely slow thermal dynamics, with settling times exceeding 15 h after a step change in power. Empirical tuning is therefore impractical. Instead, a physical model is required to parametrise and validate the controller before experiments are performed on the actual system.

This thesis focuses on two experimental configurations, the oven setup capable of heating packaged semiconductor devices up to 300 °C and the thermoelectric system integrated into a vacuum chamber, enabling temperature control between –50 °C and 100 °C. Although both systems act as temperature actuators, their distinct thermal characteristics necessitate different modelling approaches and different controller structures.

#### Oven Controller

The oven consists of a ceramic base, as illustrated in Figure 3.13, into which the packaged devices are placed. The heating coil and a temperature sensor are embedded inside an

aluminium frame, which is thermally insulated from the external aluminium case. For electrical precision measurements, the entire assembly is placed inside a steel shielding enclosure. The output of the temperature controller is connected to a power connector on the PCB located at the rear of the oven.

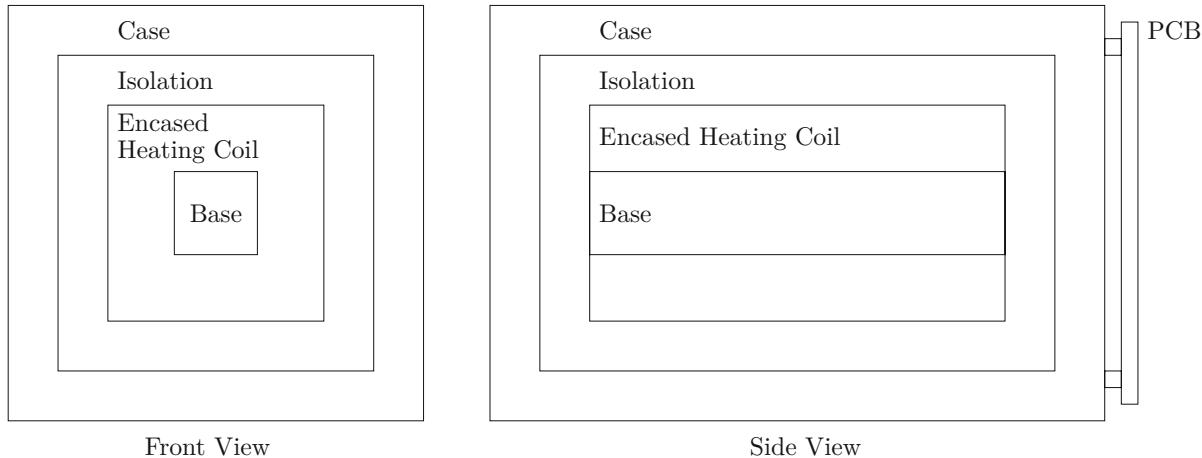


Figure 3.13: Front and side view of the oven setup.

Because the oven contains only a single actuator and a single temperature sensor and the exact geometric relation between heater and sensor is not known, temperature gradients within the structure cannot be compensated. However, the ceramic base and the enclosed heating coil are sufficiently well insulated that a lumped mass model provides an adequate approximation for the control purposes.

The lumped mass model reduces the system to two key parameters: the effective thermal resistance and thermal capacitance. These can be identified by applying a step change in power and analysing the corresponding temperature step response. The measurement data is then fitted to Equation 2.7, using the system identification method described in Subsection 2.4.1 – System Identification.

Initial tests with several sets of PID parameters showed that conventional PID control either resulted in substantial overshoot or in excessively long rise times. Because of the slow system dynamics, the derivative term offers no practical advantage and tends to destabilise the controller. To accelerate the transient response without overshoot, a feed-forward power offset was introduced, computed as

$$\dot{Q}_{\text{in}} = \frac{T(t) - T_a}{R_{\text{conv}}}, \quad (3.2)$$

where  $\dot{Q}_{\text{in}}$  is the input heating power,  $R_{\text{conv}}$  the effective thermal resistance,  $T(t)$  the oven temperature and  $T_a$  the ambient temperature. This offset compensates for heat losses to the environment. The temperature dependence of  $R_{\text{conv}}$  is shown in Figure 3.14. The remaining PI controller then only needs to compensate for the steady-state error arising from the thermal capacitance.

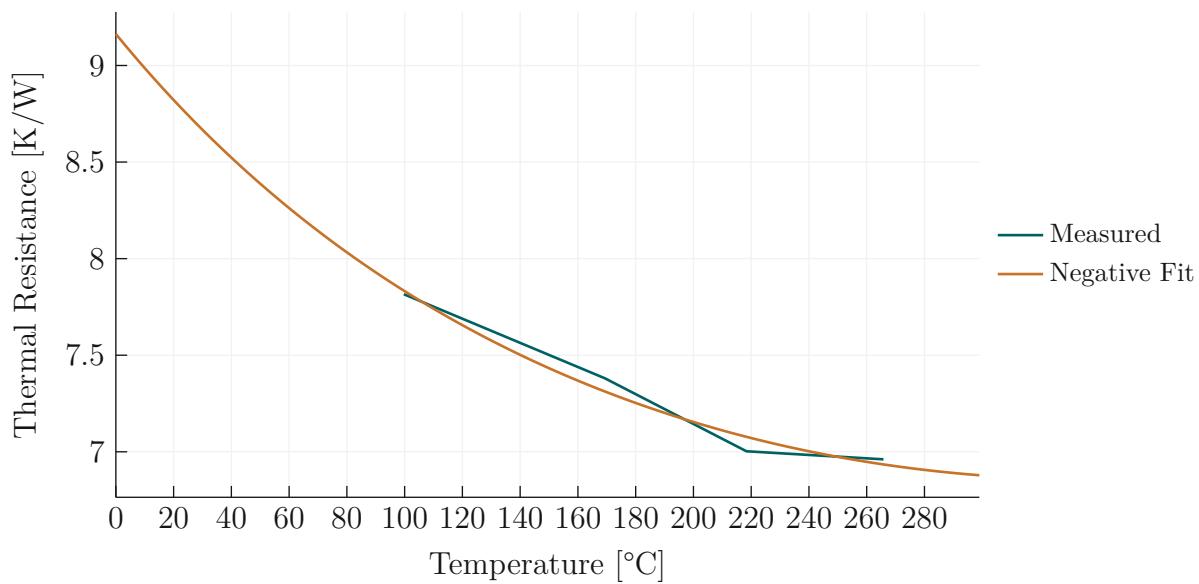


Figure 3.14: Temperature dependence of the effective thermal resistance of the oven, representing the power required to maintain a steady-state temperature.

Due to the slow thermal response, even moderate integral action can accumulate significant error over long heating intervals, again leading to overshoot. To prevent integral wind-up, the integral term is activated only when the temperature error is less than 0.5 K. This allows for stronger integral correction without destabilising the system.

Experimental validation revealed that PI parameters tuned for low temperatures resulted in slow dynamics at higher temperatures. This behaviour is attributed to the temperature dependence of the effective thermal capacity, raising the oven temperature from 150 °C to 151 °C requires significantly more power than increasing it from 50 °C to 51 °C. To compensate for this effect, the PI controller parameters are implemented as linear functions of the measured oven temperature.

### Thermoelectric Controller

The thermoelectric system, as illustrated in Figure 3.15, is significantly more complex than the oven. It incorporates two sensors and two actuators. Two thermoelectric modules are arranged vertically with the DUT positioned between them, and each module is driven by a separate temperature controller. The temperature is measured inside the DUT using two PT1000 sensors. Each thermoelectric module is connected to a water cooled heat sink, and the water loop is temperature stabilised by a Cloudray CW5200 industrial chiller. To prevent ice formation at low temperatures, which would damage the DUT and disrupt heat transfer, the entire assembly operates inside a vacuum chamber. Electrical signal, power signals and cooling connections are routed through vacuum tight connections.

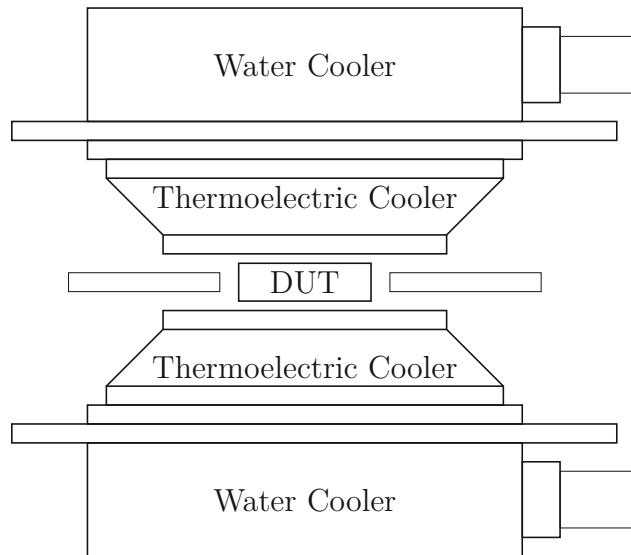


Figure 3.15: Thermoelectric setup with the DUT positioned between two thermoelectric modules, each connected to a water cooler.

Although a lumped mass model could again be used, the significantly smaller thermal mass leads to settling times within minutes rather than hours. Therefore, a separate simulation model is not required for controller design. As in the oven controller, a feed-forward term is introduced to compensate the steady-state heating or cooling power required to hold a given temperature. This term is derived from the characteristics shown in Figure 3.16. In this case, the emphermal resistance does not correspond to a physical convection path but is instead a simplification of the thermoelectric equations derived in subsubsection 2.1.3 – Thermoelectric Modules. Because the cooling water temperature varies only by 1 °C to 2 °C, ambient temperature tracking is unnecessary.

In the control algorithms, both temperature controllers are initiated simultaneously and their temperature readings are averaged once per second. This averaging serves to compensates for sensor tolerances and device to device variations. It also prevents the controllers from counteracting each other if the two temperature readings differ, if for example one controller heats while the other is cools.

Despite steady-state compensation, the controller dynamics still depend on the DUT temperature, the cooling water temperature and the desired set point. These dependencies give rise to four distinct operating states:

1. Heating with  $T_{\text{DUT}} > T_{\text{water}}$ ,
2. Cooling with  $T_{\text{DUT}} > T_{\text{water}}$ ,
3. Heating with  $T_{\text{DUT}} < T_{\text{water}}$ ,
4. Cooling with  $T_{\text{DUT}} < T_{\text{water}}$ .

As shown in Equation 2.9 the individual thermoelectric contributions: Seebeck term, Joule heating, and Fourier conduction, either heat or cool depending on the operating condition. For example, if the DUT temperature is below the cooling water temperature and the desired temperature is higher, all three terms contribute to heating. However, if

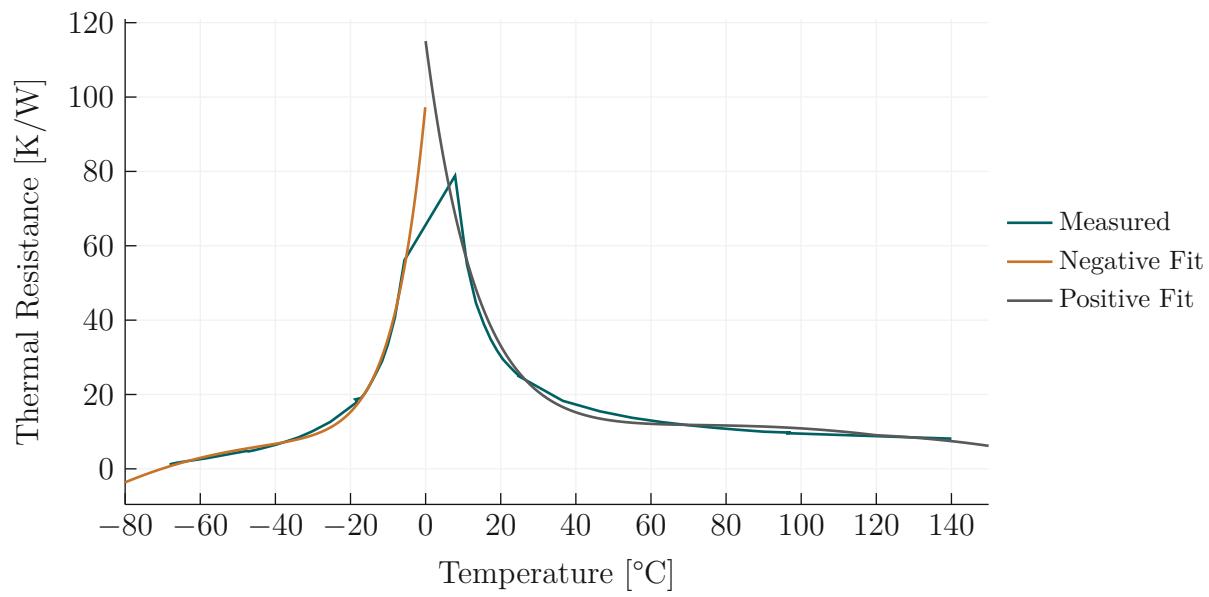


Figure 3.16: Effective thermal resistance of the thermoelectric module used to determine the steady-state heating or cooling power.

the set point is lower than the DUT temperature, only the Seebeck term contributes to cooling, while the other terms reduce the achievable cooling rate.

To compensate for these regime dependent dynamics, each of the four operating states is assigned a separate set of controller gains. This enables the controller to adapt to the instantaneous thermal behaviour of the thermoelectric modules and ensures stable, fast, and accurate temperature regulation across the entire operating range.



# 4 | Result

This chapter presents the measurements conducted to verify the claims made in Chapter 3 – Method, compare the different control approaches, evaluate improvements, and explain the methodology used for data acquisition.

## 4.1 Electrical Measurements

All electrical measurements were performed using a Tektronix TDS2024C oscilloscope, monitoring the voltage at the amplifier output. Figure 4.1 compares the output noise at 60 V with the cooling fan enabled and disabled. The measurement confirms that, after electrically isolating the fan circuit from the analog supply, the fan no longer introduces noise into the amplifier output. The remaining 100 Hz component observed in both traces originates from the high-voltage power supply, resulting in a voltage noise below 3.5 mV<sub>rms</sub>.

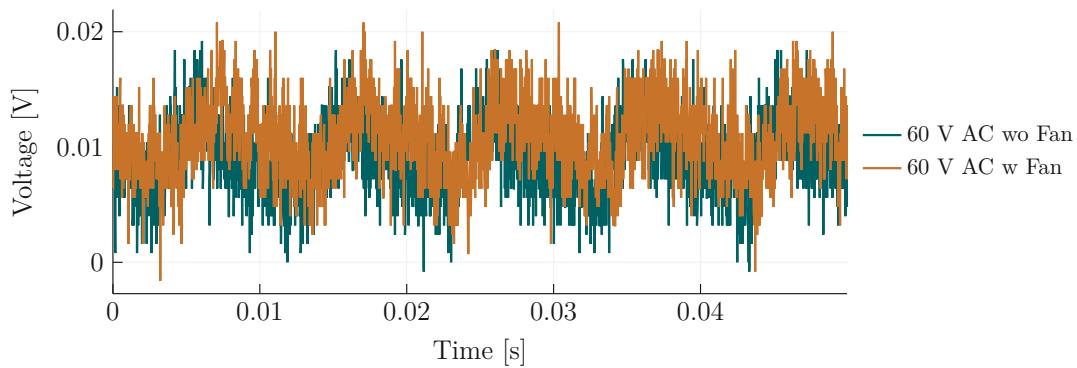


Figure 4.1: Noise comparison of the amplifier output with and without fan operating.

Figure 4.2 compares the noise behaviour at 60 V, where the amplifier output is still stable, with the output at 70 V, where the system is no longer stable. The noticeably increased 100 Hz ripple in the unstable case again originates from the high-voltage supply and demonstrates that, as the amplifier approaches its rail voltage, the buffer capacitor can no longer supply sufficient energy to maintain a stable output.

The final electrical measurement, shown in Figure 4.3, evaluates the step response of the three electrical PID controllers. All three controllers settle within approximately

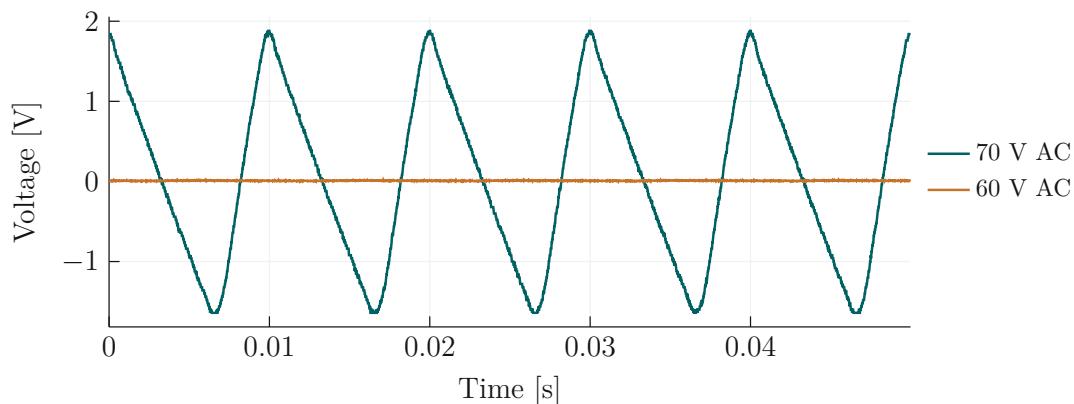


Figure 4.2: Noise comparison between the maximum stable output voltage and an unstable operating point.

8 ms, corresponding to an effective electrical control update rate of at least 125 Hz. This dynamic response is significantly faster than any thermal behaviour of the system and can therefore be neglected when designing the temperature controller. The 1 ms processor timer is clearly visible in the step response, as output updates only occur at multiples of this interval. The settling behaviour of the internal analog feedback loop can also be observed, reaching its final value within 100  $\mu$ s to 200  $\mu$ s.

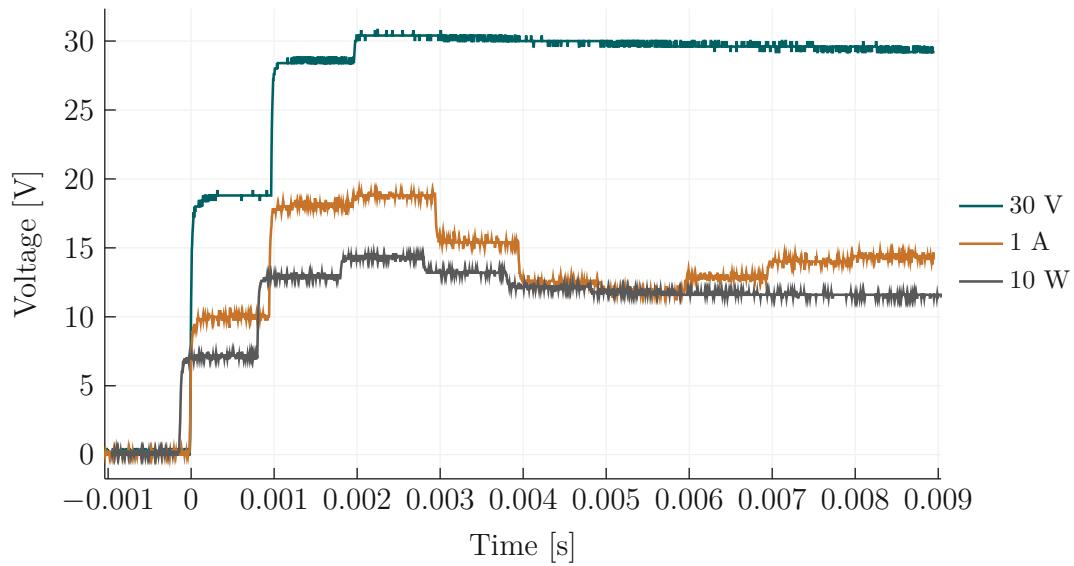


Figure 4.3: Step response of the amplifier output for the three electrical PID controllers.

## 4.2 Thermal Measurements

All thermal results were obtained using the full experimental setups. The temperature controller executed the control algorithms in real time while simultaneously recording the system response.

### 4.2.1 Oven Measurements

As shown in Figure 4.5 and Figure 4.4, even with the temperature controller optimised for speed, the oven requires approximately 30 min to reach the target temperature. After settling, the controller maintains the temperature within  $\pm 35$  mK. The cooling transitions take substantially longer and cannot be accelerated without reducing thermal insulation or introducing an active cooling mechanism.

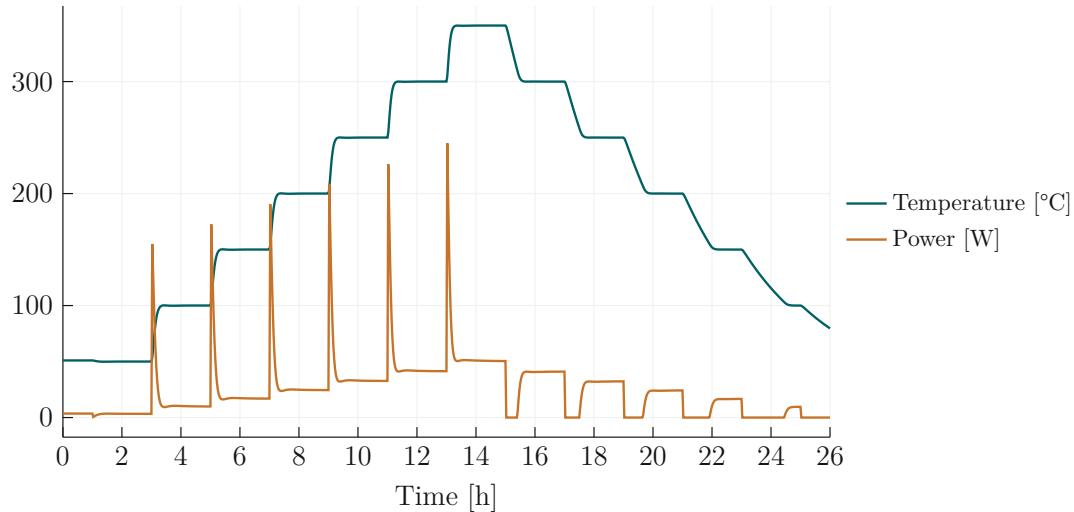


Figure 4.4: Oven temperature controller evaluated over multiple set points.

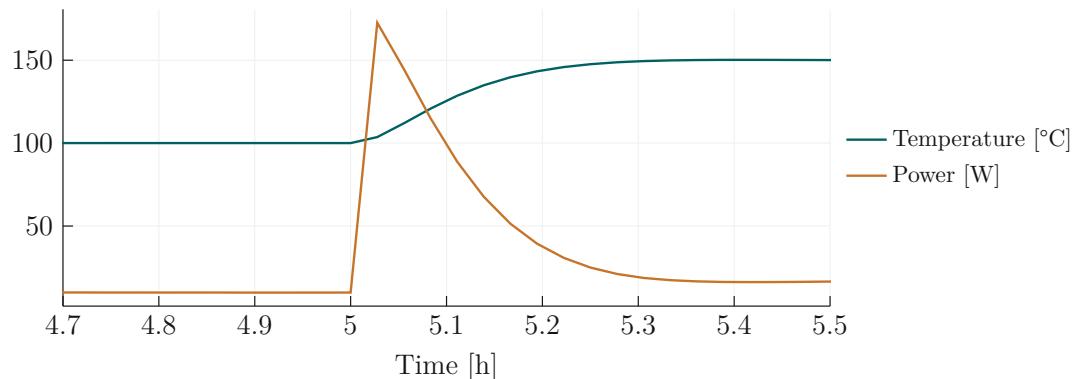


Figure 4.5: Detailed view of the step response from 100 °C to 150 °C.

### 4.2.2 Thermoelectric Measurements

Figure 4.6 and Table 4.1 summarise the results obtained from the thermoelectric setup. Various set points and temperature step sizes were tested to evaluate the controller performance.

The State PID controller outperforms both the PI and model-based controllers in terms of overshoot and settling time, with the exception of the 15 °C set point. This temperature

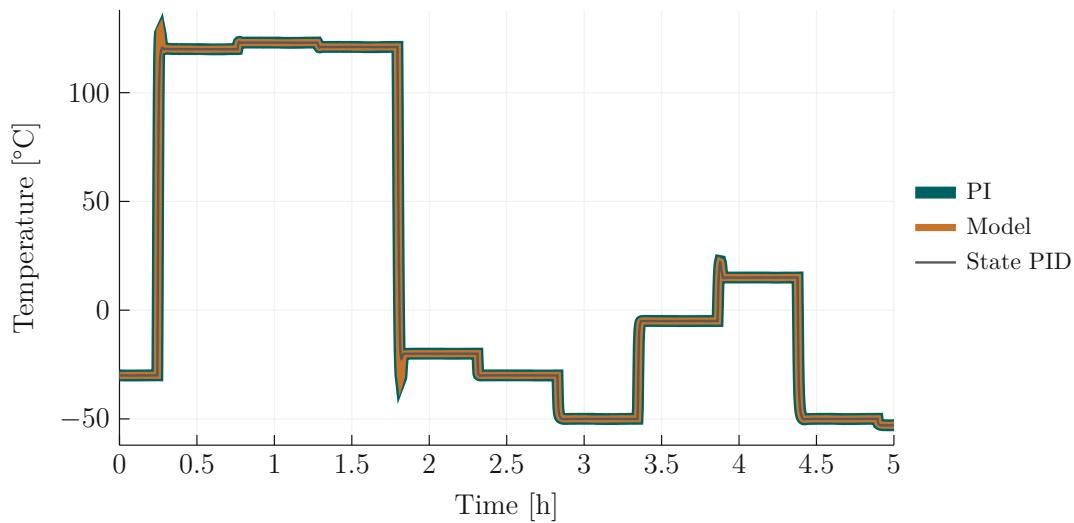


Figure 4.6: Performance comparison of the three control algorithms. The background lines are drawn with increased line width to remain visible in overlapping regions.

lies close to the cooling water temperature, which is not precisely known during runtime. This uncertainty makes it difficult to determine the optimal switching point between heating and cooling, leading to increased overshoot for all algorithms. In addition, the State PID algorithm exhibits superior accuracy, staying within  $\pm 25$  mK after the settle in time. Contrary to expectations, the model based algorithm does not outperform the simple PI controller. Considering the additional parameter identification effort required, its use is therefore not recommended.

Table 4.1: Comparison of overshoot, settling time, and standard deviation after settling time for selected temperatures.

	Temperature [°C]	Overshoot [K]	Settling Time [s]	Standard Deviation [K]
PI	120	7.885	117	0.091
	15	7.669	140	0.025
	-53	-0.083	36	0.078
	-40	-0.65	100	0.050
Model	120	9.082	115	0.092
	15	7.344	130	0.030
	-53	-0.081	38	0.080
	-40	-0.596	96	0.051
State PID	120	0.711	79	0.118
	15	6.905	158	0.026
	-53	-0.047	29	0.050
	-40	-0.615	78	0.031

## 4.3 Safety Validation

All implemented safety functions were validated experimentally. Due to the complexity of the required setups and the transient nature of many failure events, the tests were not recorded. In most cases, measurements would not have provided additional technical insight beyond the instrument readings observed during the experiments.

The over current detection was evaluated by deliberately exceeding the configured current limit and verifying that the device transitioned into the appropriate error state with the output relays opening as intended. A similar method was used to test the fuse protection. However, the over current detection had to be disabled for this test, and the fuse was irreversibly damaged as expected.

The firmware based safety limits were validated by deliberately violating each limit at least once and observing whether the controller entered the correct error state. User inputs were checked for proper range validation, and the state-transition logic was tested to ensure that only permissible system states could be selected.

One of the more severe failures encountered during development occurred due to insufficient thermal robustness of the output amplifier. Local overheating of one transistor caused a short circuit. With the short located within the feedback loop, the amplifier was no longer able to regulate the output voltage, and the system only shut down after the fuse opened.

A potentially dangerous user error can still occur on the experiment side if the external temperature sensor is not thermally coupled to the heating element. In such a case, the controller detects a valid sensor and assumes the system is within the allowed temperature range, while the actual setup may already be overheating. Depending on the configuration, this can lead to hazards ranging from fire risk to component rupture, particularly in the case of smaller compact components.



# 5 | Conclusions and Outlook

This chapter summarises the main findings and achievements of the presented work. It provides an integrated assessment of the complete temperature controller that was developed, implemented and experimentally validated, including the hardware design, firmware architecture, control algorithms and thermal actuation. The key results are evaluated with respect to the initial objectives and their suitability for use in experimental setups at the Institute of Microelectronics. Furthermore, the chapter discusses the limitations observed during operation and testing, offering a realistic perspective on the capabilities and boundaries of the system.

## 5.1 Conclusions

The system is suitable for use as a temperature controller and as a power supply in low-noise experimental environments. These capabilities are particularly relevant for semiconductor device characterisation experiments conducted at the Institute for Microelectronics [39], [40].

When comparing the obtained system parameters with the initial design targets, it becomes evident that not all objectives were fully achieved. The intended output voltage range of  $-25\text{ V}$  to  $70\text{ V}$  can be set. However, above  $60\text{ V}$  the output voltage exhibits fluctuations that render it unsuitable for low-noise operation. Importantly for the use as temperature controller, high precision measurements are carried out only after the setup has reached its target temperature, where the controller primarily compensates for thermal losses. In this operating regime, the output voltage lies within the stable range, exhibiting a noise level below  $3.5\text{ mV}_{\text{rms}}$ . The output current control meets the specified requirements, although the conducted validation did not include a detailed analysis of current noise under varying loads or a possible influence of the controller on the precision measurements.

The temperature control loop achieves the targeted cycle frequency of  $5\text{ Hz}$ . Combined with the electrical control frequency of  $125\text{ Hz}$ , this supports the assumption that the thermal system is effectively regulated through instantaneous power modulation. The control loop frequency could be increased through firmware modifications if required, yet given the relatively slow thermal dynamics of the experimental setups, characterised

by settling times ranging from several seconds to tens of minutes, such an adjustment currently offers no practical benefit.

The desired relative temperature accuracy of  $\pm 10$  mK was achieved only at the sensor level. The oven setup reached an accuracy of approximately  $\pm 35$  mK, while the thermoelectric setup achieved  $\pm 25$  mK once thermal equilibrium was reached. These values are consistent with expectations for systems lacking detailed thermal models accounting for ambient influences or enhanced thermal isolation, provided by cryogenic enclosure for example [15, 41, 42].

Sensor calibration is currently performed manually. For broader deployment or long term use, this procedure should be improved and automated. While the system provides accurate relative temperature measurements, the absolute accuracy cannot be verified without a dedicated temperature calibration setup, which is not yet available.

From a software perspective, the firmware architecture is largely complete. It provides clearly defined interfaces that simplify the integration of additional experiment modes. The microcontroller available computational resources and flexible timing configuration enable the implementation of more advanced control strategies, such as Kalman filtering or Lyapunov based controllers. Importantly, the safety logic remains independent of the specific control algorithm, ensuring robust operation even in experimental configurations or when used by inexperienced personnel.

Overall, the system offers a solid foundation for further research and development. Most modifications can be carried out without compromising operational safety. Only alterations to the high-voltage supply or the power amplifier circuitry require elevated caution and should be performed exclusively by experienced hardware developers due to inherent safety risks.

## 5.2 Outlook

Several avenues exist to further enhance the performance and use of the temperature controller. One current limitation of the current prototype is the low energy efficiency, although acceptable during development, this inefficiency may hinder wider deployment within the Institute. Modern low-noise laboratory systems often rely on Switch Mode Power Supplies (SMPSSs), with the output noise mitigated through appropriate filtering and LDO regulators, ultimately achieving noise characteristics comparable to those of linear power supplies.

Future revisions could benefit from restricting the output voltage range to  $\pm 48$  V. The extended voltage range is rarely required for temperature control and contributes to increased system complexity. Reducing the range would have a negligible impact on heating performance, but would significantly simplify development and improve operational safety, as the circuit could then be powered from a low-voltage supply that is safe for direct contact.

Enhancements to temperature accuracy may be achieved either through improvements to the physical experimental setup or through the adoption of more sophisticated control algorithms. More accurate thermal modelling, including precise actuator and sensor placement, could allow model based control approaches that account explicitly for real world thermal behaviour.

Beyond temperature control, the system offers additional potential use cases within the laboratory. With minor firmware modifications, it could function as a programmable laboratory bench power supply, provided that the electrical PID controller is bypassed and the analogue output stage is calibrated accordingly. Alternatively, the amplifier circuit could serve as a high-power signal generator, though this would require more substantial adjustments depending on the target application.

The existing documentation is sufficient for reproducing the system using the current prototype as a reference. However, if larger quantities are desired, the documentation must be expanded considerably. Since the initial validation phase, the system has already been reconstructed three times, resulting in four operational temperature controllers currently available at the Institute of Microelectronics.



# References

- [1] S. Li, A. Hu, and J. Liu, “Research in Improving the Conversion Efficiency of LED Power Supply,” *Proceedings - 2020 19th Distributed Computing and Applications for Business Engineering and Science, DCABES 2020*, pp. 178–181, 10 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9277831>
- [2] N. Beev, M. Cerqueira Bastos, M. Martino, D. Valuch, L. Palafox, S. Member, and R. Behr, “Design and Metrological Characterization of a Digitizer for the Highest Precision Magnet Powering in the High Luminosity Large Hadron Collider,” *IEEE TRANSACTIONS ON INSTRUMENTATION AND MEASUREMENT*, vol. 73, p. 6501710, 2024.
- [3] Keithley, “Keithley Series 2600B Datasheet,” Keithley, Tech. Rep., 2025. [Online]. Available: [https://download.tek.com/datasheet/1KW-60906-0\\_Series\\_2600BSystem\\_SourceMeter\\_SMU\\_Datasheet\\_011525.pdf](https://download.tek.com/datasheet/1KW-60906-0_Series_2600BSystem_SourceMeter_SMU_Datasheet_011525.pdf)
- [4] M. Kratzmann, “Temperaturkontrolle für die Defekte Spektroskopie in Halbleiter Transistoren,” Ph.D. dissertation, Technical University of Vienna, 2019.
- [5] C. Schleich, “Erweiterung des TDDS Temperaturreglers,” Ph.D. dissertation, Technical University of Vienna, 2017.
- [6] Quick-Ohm Küpper & Co. GmbH, “QC-PC-PID-01 Peltier-Controller Datasheet,” Quick-Ohm Küpper & Co. GmbH, Tech. Rep., 2021. [Online]. Available: [https://www.quick-ohm.de/pdf/peltiercontroller/operating-manual-qc-pc-pid-01\\_en.pdf](https://www.quick-ohm.de/pdf/peltiercontroller/operating-manual-qc-pc-pid-01_en.pdf)
- [7] Meerstetter Engineering, “TEC Controller TEC-1167 Datasheet,” Meerstetter Engineering, Tech. Rep., 2023. [Online]. Available: <https://www.meerstetter.ch/customer-center/downloads/category/29-latest-datasheets?download=645:tec-1167-tec-controller-datasheet>
- [8] Omron, “Advanced Digital Temperature Controller E5AN-H/E5EN-H Datasheet,” Omron, Tech. Rep., 2025. [Online]. Available: [https://files.omron.eu/downloads/latest/datasheet/en/h05e\\_e5an-h\\_e5en-h\\_advanced\\_digital\\_temperature\\_controller\\_datasheet\\_en.pdf?v=2](https://files.omron.eu/downloads/latest/datasheet/en/h05e_e5an-h_e5en-h_advanced_digital_temperature_controller_datasheet_en.pdf?v=2)
- [9] Lake Shore Cryotronics, “Temperature Controller Model 336 Datasheet,” Lake Shore Cryotronics, Tech. Rep., 2024. [Online]. Available: [https://www.lakeshore.com/docs/default-source/default-document-library/336\\_a4.pdf?sfvrsn=762c14c2\\_5](https://www.lakeshore.com/docs/default-source/default-document-library/336_a4.pdf?sfvrsn=762c14c2_5)

- [10] T. Bergman, A. Lavinve, F. Incropera, and D. DeWitt, *Introduction to Heat Transfer*, 6th ed. Wiley, 6 2011. [Online]. Available: <https://www.amazon.com/Introduction-Heat-Transfer-Theodore-Bergman/dp/0470501960>
- [11] Y. Cengel and A. Gahjar, *Heat and mass transfer fundamentals and applications*, 5th ed. New York: McGraw-Hill Education, 2015.
- [12] G. Fasching, *Werkstoffe für die Elektrotechnik*, 4th ed. Vienna: Springer-Verlag, 12 2005.
- [13] A. Prechtl, *Vorlesungen über die Grundlagen der Elektrotechnik Band 1*, 2nd ed. Vienna: Springer Vienna, 11 2005.
- [14] D. M. Rowe, *Thermoelectrics Handbook Macro to Nano*, 1st ed. CRC Press, 2006.
- [15] M. D. Thakor, S. K. Hadia, and A. Kumar, “Precise temperature control through Thermoelectric Cooler with PID controller,” *2015 International Conference on Communication and Signal Processing, ICCSP 2015*, pp. 1118–1122, 11 2015.
- [16] D. Enescu and E. O. Virjoghe, “A review on thermoelectric cooling parameters and performance,” *Renewable and Sustainable Energy Reviews*, vol. 38, pp. 903–916, 2014.
- [17] T. Nunnally, D. Pellicone, N. Van Velson, J. Schmidt, and T. Desai, “Thermoelectric performance model development and validation for a selection and design tool,” *Thermomechanical Phenomena in Electronic Systems -Proceedings of the Intersociety Conference*, pp. 1404–1411, 9 2014.
- [18] K. Alsniae, S. M. Ahmed Ghaly, and M. A. Ali, “Study and Design of a Multi-range Programmable Sensor for Temperature Measurements,” *Engineering, Technology and Applied Science Research*, vol. 12, no. 6, pp. 9601–9606, 12 2022.
- [19] W. Kester, *Practical Design Techniques for Sensor Signal Conditioning*. Analog Devices, 1999.
- [20] OVE Österreichischer Verband für Elektrotechnik, “ÖVE/ÖNORM EN IEC 61140:2016: Schutz gegen elektrischen Schlag - Gemeinsame Anforderungen für Anlagen und Betriebsmittel,” OVE, Vienna, Tech. Rep., 2016.
- [21] —, “ÖVE/ÖNORM EN IEC 61010-2-033: Sicherheitsbestimmungen für elektrische Mess-, Steuer-, Regel- und Laborgeräte,” OVE, Vienna, Tech. Rep., 2023.
- [22] E. Schrüfer, L. Reindl, and Zagar Bernhard, *Elektrische Messtechnik*, 10th ed. Fachbuchverlag Leipzig, 2012.
- [23] D. Ying and D. A. Hall, “Current Sensing Front-Ends: A Review and Design Guidance,” *IEEE Sensors Journal*, vol. 21, no. 20, pp. 22329–22346, 10 2021. [Online]. Available: <https://ieeexplore.ieee.org/document/9475045>
- [24] S. Ziegler, R. C. Woodward, H. H. C. Iu, and L. J. Borle, “Current sensing techniques: A review,” *IEEE Sensors Journal*, vol. 9, no. 4, pp. 354–376, 4 2009.

- [25] J. Wu, "A Basic Guide to RTD Measurements," Texas Instruments, Tech. Rep., 2023. [Online]. Available: [www.ti.com](http://www.ti.com)
- [26] —, "K-Type Thermocouple Measurement Circuit With Precision ADC," Texas Instruments, Tech. Rep., 3 2019.
- [27] P. Horowitz and W. Hill, *The Art of Electronics*, 3rd ed. Cambridge University Press, 2015.
- [28] P. Gray, P. Hurst, S. Lewis, and R. Meyer, *Analysis and Design of Analog Integrated Circuits*, 5th ed. Wiley, 2009.
- [29] B. Razavi, *Design of Analog CMOS Integrated Circuits*, 2nd ed. McGraw - Hill Education, 2017.
- [30] M. von Haartman, "Low-frequency noise characterization, evaluation and modeling of advanced Si- and SiGe-based CMOS transistors," Ph.D. dissertation, KTH Stockholm, Stockholm, 2006.
- [31] Horst Zimmermann, "Schaltungstechnik," Technical University of Vienna, Tech. Rep., 2019.
- [32] D. Self, *Audio Power Amplifier Design Handbook*, 4th ed. Elsevier, 2006.
- [33] P. A. Laplante and S. J. Ovaska, *Real-Time Systems Design and Analysis*, 4th ed. IEEE Press, 2012.
- [34] B. Powel Douglass, *Design Patterns for Embedded Systems in C*, 1st ed. Elsevier, 2011.
- [35] W. Kemmetmüller and A. Kugi, "Regelungssysteme," Institut für Automatisierungs- und Regelungstechnik, TU Wien, Tech. Rep., 2024. [Online]. Available: [https://www.acin.tuwien.ac.at/file/teaching/master/Regelungssysteme-1/Regelungssysteme\\_VO\\_2024.pdf](https://www.acin.tuwien.ac.at/file/teaching/master/Regelungssysteme-1/Regelungssysteme_VO_2024.pdf)
- [36] A. Kugi, "Automatisierung," Institut für Automatisierungs- und Regelungstechnik, TU Wien, Wien, Tech. Rep., 2025. [Online]. Available: <https://www.acin.tuwien.ac.at/file/teaching/bachelor/automatisierung/AutomatisierungVO.pdf>
- [37] R. Munning Schmitd, G. Schitter, A. Rankers, and J. van Eijk, *The Design of High Performance Mechatronics*, 3rd ed. Delft University Press, 2020.
- [38] A. Deutschmann-Olek and A. Kugi, "Nichtlineare dynamische Systeme und Regelung," Institut für Automatisierungs- und Regelungstechnik, TU Wien, Tech. Rep., 2025. [Online]. Available: [https://www.acin.tuwien.ac.at/file/teaching/master/VO\\_Nichtlineare\\_dynamische\\_Systeme\\_und\\_Regelung/SS2025/NDSC\\_notes\\_v6.pdf](https://www.acin.tuwien.ac.at/file/teaching/master/VO_Nichtlineare_dynamische_Systeme_und_Regelung/SS2025/NDSC_notes_v6.pdf)
- [39] P. Stampfer, F. Roger, L. Cvitkovich, T. Grasser, and M. Waltl, "A DLTS Study on Deep Trench Processing-Induced Trap States in Silicon Photodiodes," *IEEE Transactions on Device and Materials Reliability*, vol. 24, no. 2, pp. 161–167, 6 2024.

- [40] M. Waltl, Y. Hernandez, C. Schleich, K. Waschneck, B. Stampfer, H. Reisinger, and T. Grasser, “Performance Analysis of 4H-SiC Pseudo-D CMOS Inverter Circuits Employing Physical Charge Trapping Models,” *Materials Science Forum*, vol. 1062, pp. 688–695, 2022. [Online]. Available: <https://repositum.tuwien.at/handle/20.500.12708/136907>
- [41] Z. Xu, M. Xu, W. Cheng, H. Peng, and Y. Ding, “High-precision, temperature control based on grading-structure and PID-feedback strategies,” *Transactions of the Japan Society for Aeronautical and Space Sciences*, vol. 61, no. 2, pp. 51–59, 2018.
- [42] A. Demann, S. Mueller, and S. B. Field, “1 K cryostat with sub-millikelvin stability based on a pulse-tube cryocooler,” *Cryogenics*, vol. 73, pp. 60–67, 1 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0011227515001642>