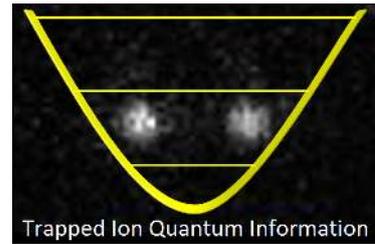


ETH zürich



**Low-Noise DC Amplifier Design for
Trapped Ion Setups
&
Bayesian Schemes for Fast Single Qubit
Detection**

Michael Marti

Master Thesis

Supervisor: Dr. Matteo Marinelli

Principal Investigator: Prof. Jonathan Home

Trapped Ion Quantum Information, Institut für Quantenelektronik, ETH Zürich

April 11, 2022

Abstract

To allow for more flexibility in shaping DC potentials in ion traps, we discuss the development of an amplifier board compatible with the commercially available 32-channel arbitrary waveform generator card Fastino for DC electrode control. We outline criteria for operational amplifiers as main building block for low-noise active filters. Using the outlined criteria we proceed to design a versatile active filter topology for the chosen operational amplifier. A first design will allow users to generate output voltage amplitudes of 70V. This amounts to an increase by a factor of seven compared to using the Fastino alone. We discuss how to adapt the design to fit different experimental specifications and propose improvements to the existing design. Furthermore, we analyze Bayesian schemes for single qubit state detection. The commonly used adaptive maximum likelihood method exhibits a significantly higher average dark state readout time compared to the average bright state readout time. By conditionally applying π -pulses and therefore changing qubit state, we aim to equalize both average detection times. We analyze two methods to determine when to apply a π -pulse. We use hidden Markov models to simulate state readout and benchmark the two methods against the adaptive maximum likelihood method. Under ideal conditions, we achieve a significant speed-up of average dark state detection times, rendering average detection times nearly independent of the prepared state.

Contents

1	Introduction	1
I	Low-Noise DC Amplifier Design for Trapped Ion Setups	2
2	Concepts	3
2.1	Active Filters	3
2.1.1	Operational Amplifiers	3
2.1.2	Second Order Transfer Function	4
2.1.3	Sallen-Key Low-pass Filter	5
3	Amplifier Board Design	8
3.1	Operational Amplifier Selection & Criteria	8
3.1.1	Criteria	8
3.1.2	Comparison	9
3.2	Circuit Topology & Performance	10
3.2.1	Alternative Topologies	12
3.3	Passive Component Selection	13
3.4	Prototype Board	14
3.5	Design Verification	14
3.6	Full Amplifier Board	17
3.7	Power Supplies & Rail Regulation	18
4	Full System Testing	20
4.1	Overview Signal Path	20
4.1.1	Fastino DAC board	20
4.1.2	Board Interconnect	21
4.2	Overview Power Distribution	22
4.3	Preliminary Results	22
4.4	Possible Improvements	24
5	Conclusion & Outlook	26
II	Bayesian Schemes for Fast Single Qubit Detection	27
6	Concepts	28
6.1	Hidden Markov Models	28
6.2	Factor Graphs and Belief-Propagation	29
6.2.1	Factor Graphs	29
6.2.2	Belief Propagation	29
6.3	State Readout	30
6.3.1	Non-Poissonian Statistics	30
6.3.2	Real-Time State Discrimination	31

7	State Discrimination Algorithms	32
7.1	Adaptive Maximum Likelihood	32
7.1.1	Bright State Detection	32
7.1.2	Dark State Detection	33
7.2	π -pulse Methods	34
7.3	Information Gain Estimator	35
7.3.1	Example	36
7.4	Minimum Bins Estimator	38
7.5	Hidden Markov Model MAP Decoding	39
8	Simulation Environment	41
8.1	General Structure	41
8.2	Ion Class	41
8.3	Estimator Class	41
8.4	Environment Class	41
8.4.1	Data Structure	42
8.5	User Example Code	43
9	Estimator Benchmarking	44
9.1	General Setup	44
9.2	Effect Of Non-zero Background	44
9.2.1	Entropy Gain Trajectory Analysis	47
9.3	Influence of π -pulse Infidelity	47
9.4	Effect of Leakage Mechanisms	48
10	Conclusion & Outlook	52
11	Appendix	55
11.1	Gain-Bandwidth Product	55
11.2	Repositories	55

Chapter 1

Introduction

Quantum information processing and quantum computation are active fields of research. After having proven to offer potential computational advantages over classical computer hardware for problems such as solving the discrete logarithm problem and factoring large integers using Shor's algorithm [1] or performing unstructured search as in Grover's algorithm [2], efforts to build a universal quantum computer have been undertaken. Instead of relying on classical bits for performing calculations, quantum computers make use of qubits [3]. There are many different physical realizations of qubits. Trapped ions are one of the most developed and promising candidates for achieving a system which is able to perform universal quantum computation. Trapping ions require a static potential as well as a rotating quadrupole potential. Ion traps feature electrodes to which voltages are applied for generating the aforementioned potentials [4]. While many different types of ion traps exist, we will focus on linear traps where the ions are trapped in a chain forming a 1D crystal. Having trapped ions, lasers are used to manipulate motional and internal electronic states of ions in order to perform computations. Independent of the physical realization, qubits need to satisfy certain qualities to be viable as a quantum computing platform. These criteria have been famously summarized by DiVincenzo [5]. One criterion is the ability to perform quantum gates. In the setting of trapped ions, this means using lasers to address specific ions individually. One approach for individual addressing is to use waveguide arrays. Ion-ion distance in a linear chain is not constant, increasing towards the ends of the chain. This effect is exacerbated if we load an increasing number of ions into the same trap. Waveguide arrays however work best for fixed and equidistant ion spacings. To combat the uneven distribution of ions, one can apply an anharmonic static potential which aims to cancel the mutual Coulomb repulsion between the ions while also being able to set the distance [6]. In order to achieve higher curvatures for anharmonic potentials as the uniform distance potential, it is advantageous to be able to apply high DC voltages to the electrodes generating the static potential. High DC voltages can also be used for micromotion compensation using shim electrodes which can be integrated into the trap geometry. Micromotion is the undesirable motional excitation of ions by the rotating quadrupole potential. Another criterion is the ability to reliably read out qubits after having performed a quantum algorithm. In addition, as physical qubits suffer from noise which make computations error-prone, quantum error correction schemes have been developed. In the stabilizer framework, we will repeatedly read out ancilla qubits which store correlations between qubits used for the actual computation. Depending on the result of the ancilla readout, real-time feedback is applied to the qubits involved in computation. It is therefore crucial that ancilla readout is not only reliable but also fast. These two factors are only examples of many challenges which need to be resolved. The growing field of quantum engineering specializes in finding solutions to such hurdles. In this thesis we will tackle both problems outlined above. We will first upgrade an existing system for DC voltage generation in order to be able to apply higher DC electrode voltages. At the same time we will need to take care of the noise performance of the whole signal path. Noise on the DC electrodes potentially leads to unwanted motional excitations of trapped ions introducing errors. Secondly, we will explore different readout methods in order to speed up state discrimination. We will test the described algorithms in a simulation environment for a first performance assessment.

Part I

Low-Noise DC Amplifier Design for Trapped Ion Setups

Chapter 2

Concepts

In this chapter we introduce operational amplifiers as fundamental building blocks for analog signal processing. We will explore active filters and explore some design methodologies. For this chapter we assume that the reader is familiar with the concept of frequency domain analysis as we only briefly introduce this topic. For an introduction to frequency domain analysis of linear time-invariant systems and related control topics we refer to chapter 1 and 2 of [7].

2.1 Active Filters

2.1.1 Operational Amplifiers

Operational amplifiers (op-amp) are an ubiquitous resource in analog signal processing. The op-amp is a three terminal device in its most idealized description. It features two inputs, an inverting input V_- and a non-inverting input V_+ . The output is the difference between the two inputs amplified by a factor A :

$$V_{\text{out}} = A(V_+ - V_-) \quad (2.1)$$

Ideally, the gain factor A is infinite. Therefore, the output will only be finite, if both inputs are exactly at the same potential. As infinite gain is not physical we need to take into account, that an op-amp is powered by two power rails, namely the positive supply voltage $V_{\text{supp}, +}$ and the negative $V_{\text{supp}, -}$ respectively. This will limit the output voltage as $V_{\text{out}} \in [V_{\text{supp}, -}, V_{\text{supp}, +}]$. The commonly used circuit symbol featuring all terminals is shown in Fig. 2.1.

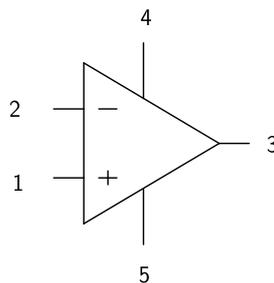


Figure 2.1: Circuit symbol of an operational amplifier. (1) Non-inverting input V_+ . (2) Inverting input V_- . (3) Output of the operational amplifier V_{out} (4) Positive supply voltage $V_{\text{supp}, +}$. (5) Negative supply voltage $V_{\text{supp}, -}$

This behaviour may not seem useful yet. However, if we apply negative feedback ¹ it can be shown that the output will settle to a value such that the voltage difference between the two input terminals is essentially zero. This fact is also known as the virtual ground assumption. To further simplify circuit analysis, one also assumes that no current flows into either of the two input terminals.

In section 3.1 we will explore more realistic models and limitations of non-ideal op-amps.

2.1.2 Second Order Transfer Function

Second transfer functions are an important building block in analog signal processing. The order refers to the number of poles the frequency response exhibits. Compared to first order transfer functions, second order transfer functions can exhibit complex-conjugate pole pairs instead of purely real poles. Simply put, a pole leads to a phase shift of -90° and a roll-off of the amplitude by 20dB per decade for frequencies above the cutoff frequency. The placement of the poles in the complex plane will influence how sharp the cutoff is, if we see resonances in the response, and how quickly the phase response will reach the dictated phase shift. We will only consider low-pass transfer functions for this thesis. The second order low-pass transfer function is given by:

$$H_{LP}(s) = \frac{k\omega_0^2}{s^2 + \frac{\omega_0}{Q}s + \omega_0^2} \quad (2.2)$$

where k is the DC gain of the filter, ω_0 the natural frequency, and Q the quality factor. Sometimes, the quality factor is replaced by the damping factor ζ . We can easily convert between the two parameters using the following relation:

$$Q = \frac{1}{2\zeta} \quad (2.3)$$

Fig. 2.2 shows the frequency response of second order transfer functions varying the quality factor Q . For values of $Q > \frac{1}{\sqrt{2}}$ we observe a resonance peak at the natural frequency. Also, the phase will drop over a shorter frequency span. If $Q = \frac{1}{\sqrt{2}}$ the resulting filter is a second order Butterworth filter. Butterworth filters have a maximally flat amplitude response. For $Q < \frac{1}{\sqrt{2}}$ the response will be under-damped. We lack a sharp cutoff and the phase drop is dragged out over a large frequency range. This is generally undesirable.

¹Less abstractly, this statement just means that the voltage at the inverting input is dependant on the output voltage.

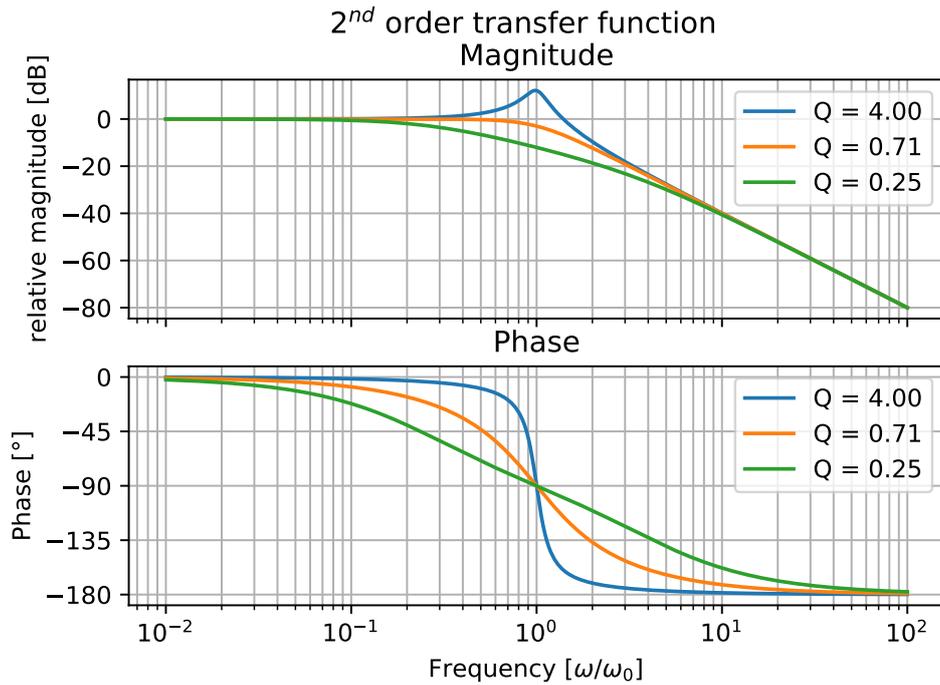


Figure 2.2: Bode plots for different values of the quality factor Q with normalized frequency axis and gain. For high quality factors we have strong resonances at ω_0 . For low quality factors $Q < \frac{1}{\sqrt{2}}$ we observe a slow transition from constant gain to a rolloff of -40dB/dec.

2.1.3 Sallen-Key Low-pass Filter

A popular implementation of arbitrary second order active filters is the Sallen-Key topology [8]. Initially designed to reduce requirements on passive components for low-frequency low-pass filters, it is still widely used in analog signal processing. One of the one of the advantages of the topology is the need for only one op-amp or active device.

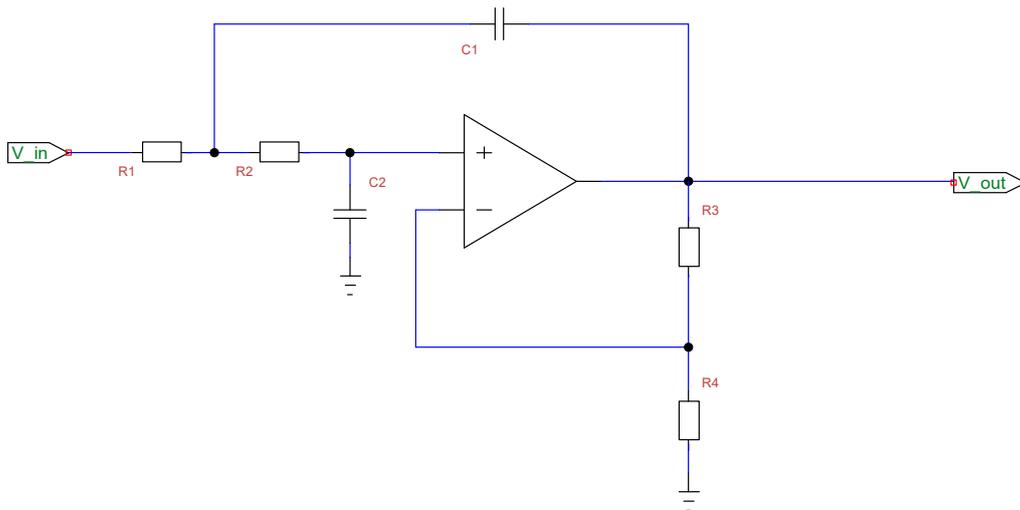


Figure 2.3: The Sallen-Key low-pass filter with gain. With the right choice of passive components, the design parameters A , Q , and ω_0 can be realized.

The parameters introduced in Eq. 2.2 can be expressed in terms of component values as:

$$k = 1 + \frac{R_3}{R_4} \quad (2.4)$$

$$\omega_0 = \sqrt{\frac{1}{R_1 R_2 C_1 C_2}} \quad (2.5)$$

$$Q = \frac{\omega_0}{\frac{1}{R_1 C_1} + \frac{1}{R_2 C_1} + \frac{1-k}{R_2 C_2}} \quad (2.6)$$

If we set the design parameters, we are left with an under-determined system of equations. If not differently specified, we are left with three components to choose freely. In the case of k , only the ratio $\frac{R_3}{R_4}$ is of importance. Therefore, one resistor value can be fixed and the other can be chosen accordingly. For Q & ω_0 we are able to freely choose two components. For a reason much more clearer later, we are interested in fixing both capacitors C_1 & C_2 . Rewriting 2.5 & 2.6 we can define the system of equations:

$$\begin{cases} P := R_1 R_2 = \frac{1}{\omega_0^2 C_1 C_2} \\ S := A R_1 + B R_2 = Q \omega_0 \end{cases} \quad (2.7)$$

with $A = C_2 + (1 - k)C_1$ and $B = C_2$. We observe that P , S , and B are always strictly positive real quantities while $A \in \mathbb{R}$. If $A = 0$, we obtain immediately the following unique solution:

$$\begin{aligned} R_2 &= \frac{S}{B} \\ R_1 &= \frac{PB}{S} \end{aligned} \quad (2.8)$$

For the case that $A \neq 0$ two solution candidates can be found:

$$\begin{aligned} R_1 &= \frac{S \mp D}{2A} \\ R_2 &= \frac{S \pm D}{2B} \end{aligned} \quad (2.9)$$

where $D^2 = S^2 - 4ABP$. As $R_1, R_2 \in \mathbb{R}^+$, we can discard solutions where $D^2 < 0$.

While analytical solutions to the design equations might exist, finding components with the found values is usually impossible. Only specific values of resistance and capacitance are manufactured in large quantities². Resistors with precision tolerance below 1% with respect to the nominal value are readily available. With high precision comes also the fine granularity of values. There are different design strategies to which components one initially fixes and determines the remaining. With modern computing power it is also possible to go through all variations. Inspired by the design note [9], an optimized brute-force algorithm can be run to find the best set of purchasable components.

The runtime of this optimized brute-force algorithm is $\mathcal{O}(|\mathbf{R}_{\text{avail}}|^2 + |\mathbf{C}_{\text{avail}}|^2)$. This result motivates why we wanted to fix the capacitance values in the first place. There are less combinations to go through and the discretization error for resistors is lower. How well a set of components performs is encompassed in the cost function f_{cost} . An intuitive cost function is to sum the squared error of our design parameters Q and ω_0 :

$$f_{\text{cost}}(x = (k_{\text{best}}, R_1, R_2, C_1, C_2)) = \left(\frac{Q_{\text{desired}} - Q(x)}{Q_{\text{desired}}} \right)^2 + \left(\frac{\omega_{0,\text{desired}} - \omega_0(x)}{\omega_{0,\text{desired}}} \right)^2 \quad (2.10)$$

Many more complex functions can be implemented. Another improvement is to not only calculate the analytical deviation from the design parameters but also to simulate the corresponding circuit to assess the performance.

²The specific values are the EX. Each decade is divided into X values logarithmically. The available series are the E12, E24, E48 & E96.

Algorithm 1 Optimal Component Selection algorithm

Output: Optimal components for Sallen-Key low-pass filter

$R_{1,best}, R_{2,best}, R_{3,best}, R_{4,best} \leftarrow 0$

$C_{1,best}, C_{2,best} \leftarrow 0$

$k_{best} \leftarrow \infty$

$Cost_{best} \leftarrow \infty$

for $R_3, R_4 \in \mathbf{R}_{avail}$ **do**

$k_i = 1 + \frac{R_3}{R_4}$

if $\|k_{desired} - k_i\| \leq \|k_{desired} - k_{best}\|$ **then**

$(R_{3,best}, R_{4,best}) \leftarrow (R_3, R_4)$

end if

end for

for $C_1, C_2 \in \mathbf{C}_{avail}$ **do**

calculate $R_{1,analytical}, R_{2,analytical}$

pick $R_1, R_2 \in \mathbf{R}_{avail}$

$Cost_i = f_{cost}(k_{best}, R_1, R_2, C_1, C_2)$

if $Cost_i \leq Cost_{best}$ **then**

$(R_{1,best}, R_{2,best}, C_{1,best}, C_{2,best}) \leftarrow (R_1, R_2, C_1, C_2)$

end if

end for

Chapter 3

Amplifier Board Design

3.1 Operational Amplifier Selection & Criteria

3.1.1 Criteria

- **Maximum supply voltage:** The maximum supply voltage limits the output range of the amplifier circuit. As we aim for higher DC electrode voltages, we want an output range as large as possible. The type of potential wells that we would like to create in the experimental setup require voltages in the order of $\pm 70V$. Therefore, we opt for a rail-to-rail operational amplifier design that matches these requirements.
- **Gain-bandwidth product:** The open-loop transfer function of an operational amplifier is usually designed to exhibit a single dominant pole [10]:

$$A(s) = \frac{V_{out}}{V_{in,+} - V_{in,-}} = \frac{A_{DC}}{1 + \frac{s}{\omega_a}} \quad (3.1)$$

where A_{DC} is the open-loop DC gain and ω_a is the dominant single pole. We can then define the gain-bandwidth product as:

$$GBW = A_{DC}\omega_a \quad (3.2)$$

An ideal op-amp has an infinite open-loop gain and bandwidth. The pole will limit the frequency regime where these assumptions are valid. When using negative feedback to limit the gain of the circuit $G \ll A_{DC}$ it can be shown that (App. 11.1):

$$G\omega_c \leq GBW \quad (3.3)$$

where ω_c is the crossover frequency of the negative feedback circuit. A high gain-bandwidth product is desirable.

- **Input referred voltage noise:** The input referred noise voltage is defined as the applied noise voltage needed at the input of an amplifier circuit to produce the noise present at the output. This allows us to regard the amplifier circuit as noise-free while adding a voltage noise source at the input. This figure of merit makes it easier to compare noise contributions from previous signal sources with the noise added by the circuit itself. In many applications (including trapped-ion experiments) it is vital to reduce noise as much as possible. For our case we aim to keep the input referred noise low compared to the output referred noise of the signal source to be amplified. In our case we can assume the output noise of the signal source to be close to $50 \frac{nV}{\sqrt{Hz}}$ [11].
- **Maximum output current:** The maximum output current could limit the resistive load which can be driven by the amplifier circuit. An ideal op-amp has no output current limit. As we do not intend to drive low resistive loads with the active filter this figure is of lesser importance.
- **Offset voltage:** The offset voltage is another non-ideality of real operational amplifiers. This effect can be modelled as an additional voltage source at the input terminal which will add to the input signal. This figure can vary with time and operating temperature. It and the corresponding temperature variation should be as low as possible.

- **Input bias current:** The input bias current is defined as the amount of current which typically flows into either the positive or negative input terminal of an operational amplifier. In the ideal case this current is zero and the amplifier has infinite input resistance.
- **Maximum slew rate:** The maximum slew rate defines the rate at which the output of an amplifier can change. It is usually given in $\frac{V}{\mu s}$. This non-linearity can impose limits to the bandwidth as the output might not be able to track a signal over its whole output range. Therefore, it should be as high as possible.
- **Area per channel:** The package size will potentially limit the number of channels which will fit onto the printed circuit board (PCB). Optimally, a set of 32 channels can be placed on a standard Eurocard sized PCB (160mm \times 100mm). However, it is clear that the layout is also determined by the pin layout of each component. So the total area might not be well approximated by the package of the amplifier alone. We note that there might be several available packages for the same amplifier. If this is the case, we will use the smallest area as lower estimate.
- **Price per channel:** This pure economic consideration can become relevant if many channels are needed.
- **Retail availability:** The best part has not only good specifications but is also purchasable. As of the time of this venture a general electronics shortage is plaguing designers.

3.1.2 Comparison

For this project three different amplifiers have been compared: ADHV4702-1 [12] from Analog Devices, LTC6091 [13] from Linear Technology, and PA443 [14] designed by Apex Microtechnology. We will continue to compare the options against the criteria laid out in subsection 3.1.1

Operational amplifier	ADHV4702-1	LTC6091	PA443
Maximum supply voltage range [V]	220	140	350
Gain-bandwidth product [MHz]	10	12	10
Input noise density [$\frac{nV}{\sqrt{Hz}}$]	8 (at 10kHz)	11 (at 10kHz)	figure not available
Maximum output current (continuous) [mA]	20	50	60
Offset voltage (typ/max) [mV]	(0.15/1)	(0.33/1.25)	(50/200)
Input bias current at optimal conditions (typ/max) [pA]	(0.3/2)	(3/50)	figure not available
Maximum slew rate [$\frac{V}{\mu s}$]	24 (74 without protection circuit)	21	32
# channel per package	1	2	1
Area per channel [mm ²]	49	12	>125
Price per channel (approximately) [CHF]	20	7	17
Availability	not available	periodically available	not available

Specification	Value
DC gain	7.8 (17.8 dB)
Bandwidth [Hz]	50 kHz - 100 kHz
Low-pass order	> 2
Desired low-pass type	Butterworth (maximally flat passband)
Maximum overshoot step response [%]	< 5

Table 3.1: Design specification for a general purpose amplifier used in trapped-ion experiments.

Component	Value	Component	Value
R_1 [k Ω]	12	C_1 [pF]	22
R_2 [k Ω]	2	C_2 [pF]	120
R_3 [k Ω]	100	C_F [pF]	22 - 24
R_4 [k Ω]	14.7	C_{out} [pF]	1000
R_{out} [k Ω]	1	-	-

Table 3.2: Resulting values for the passive components.

Going by the most important criteria, the ADHV4702-1 would be the most fitting option. Unfortunately, this part is practically unobtainable in the near future. At the time of writing this thesis, the expected time of availability is in 2023. The PA443 may have a great output range and voltage swing, but it is generally too noisy. The best compromise is then the LTC6091 which is available to purchase and is not ruled out by any of the criteria. This op-amp was also successfully used in [15] for ion shuttling experiments.

3.2 Circuit Topology & Performance

The basis of the final topology is a Sallen-Key low-pass filter with gain. To further reduce noise introduced by the circuit, additional noise-suppressing measures have been integrated. The first is adding a single pole RC filter directly at the output of the op-amp. The crossover frequency has been chosen to be 1 MHz. Secondly, a high-frequency feedback path to the inverting input has been added. The final topology is depicted in Fig. 3.1 These additions lead to a reduction of bandwidth. The maximal achievable bandwidth however is also limited by the slew rate of the op-amp to about 100 kHz over the full output range. We note that the design equations for the pure Sallen-Key topology do longer apply. Parameters of the components need to be tuned by hand for useful results. Using the chosen topology, we can choose the passive components to realize the desired filter. The design parameters are listed in Tab. 3.1 and the found values for passive components are summarized in Tab. 3.2.

We proceed by presenting simulation results which have been generated using the LTSpice VXII simulator by Analog Devices.¹

The transfer function of the resulting filter is shown in Fig. 3.2. The cutoff frequency is roughly at 70 kHz. The magnitude response drops sharply by about 60dB over one decade. This corresponds to a third order low-pass filter. Starting from 8 MHz onward we see a leveling off followed by a less steep roll-off. This is due to introduced zeros which originate from the op-amp itself. We can use passive filtering to increase the roll-off in this frequency region if necessary. The next crucial figure of merit is

¹<https://www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html>

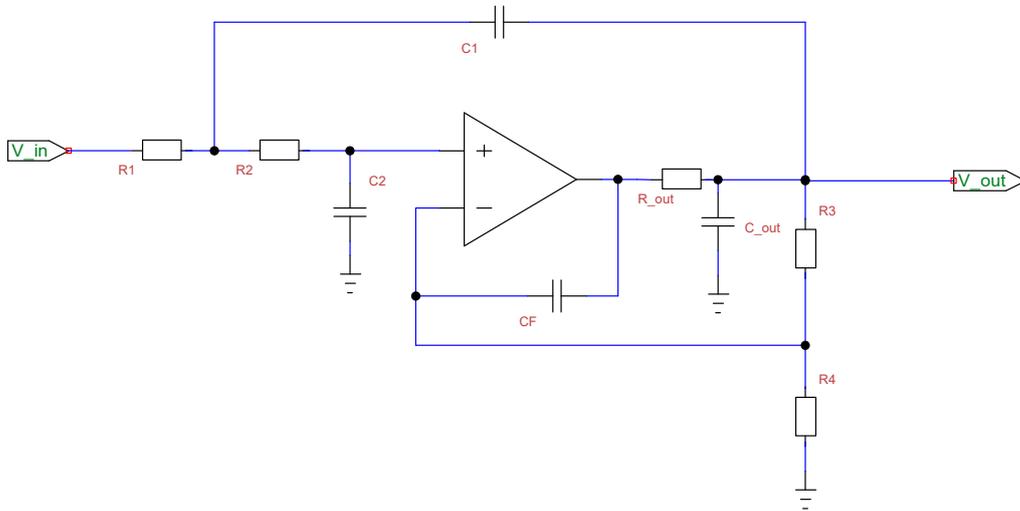


Figure 3.1: The final topology, combining the additional noise-reducing filters to the Sallen-Key with gain topology.

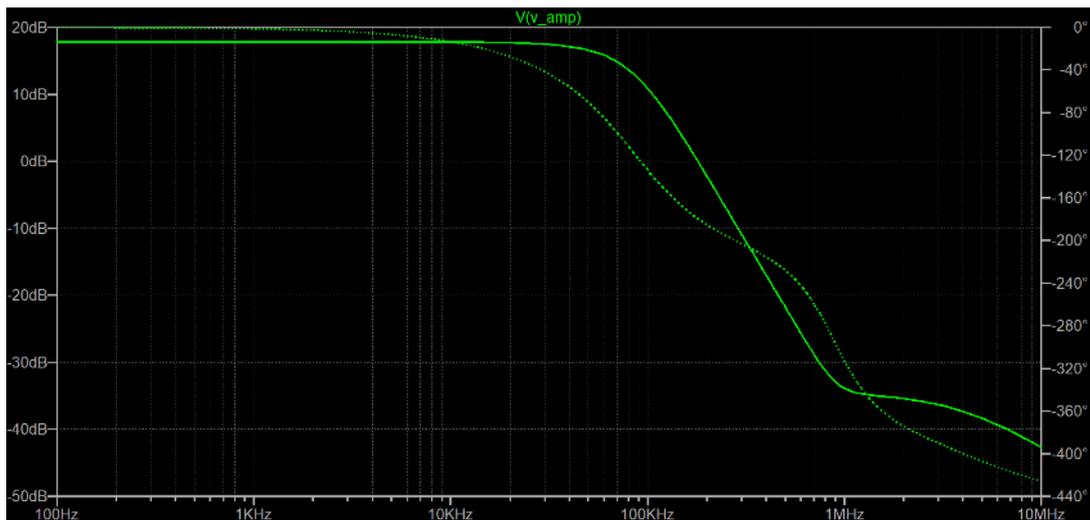


Figure 3.2: Frequency response of the resulting active filter. Continuous line depicts the magnitude of the response. The dotted line mark the phase shift.

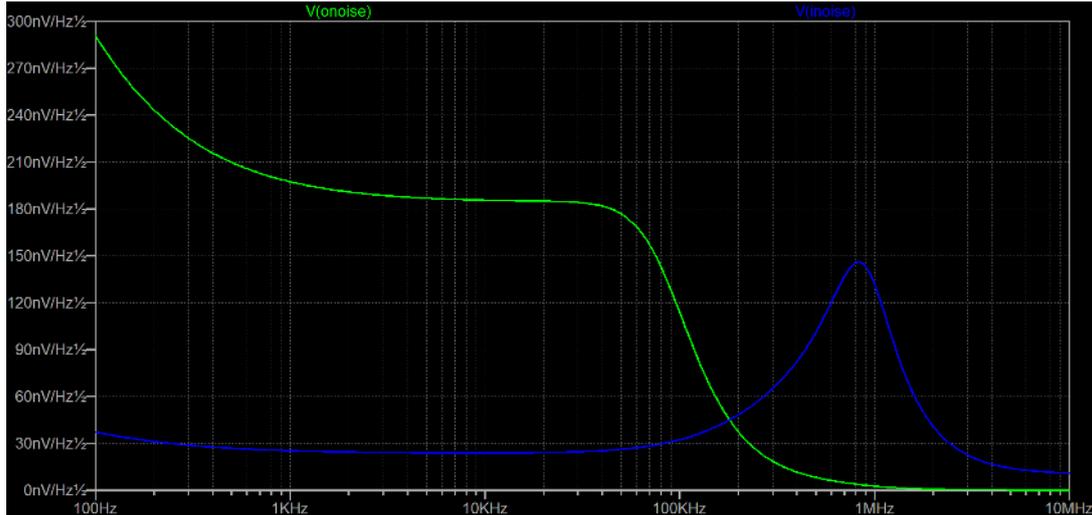


Figure 3.3: Noise spectral density in $\frac{nV}{\sqrt{Hz}}$ (RMS) of the resulting active filter. The green line corresponds to the noise density present at the output while the blue line gives the value of the input referred noise.

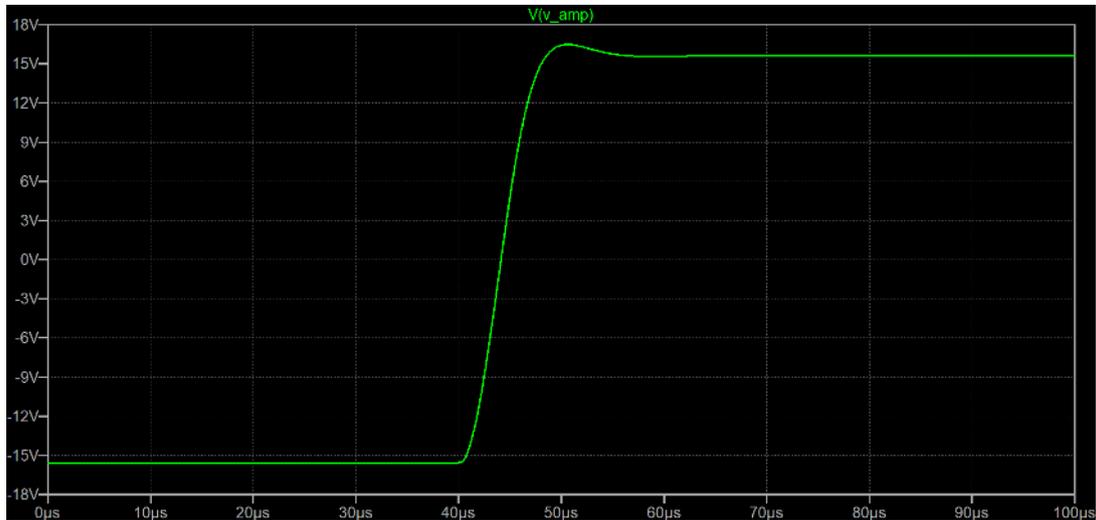


Figure 3.4: Step from -2V to +2V at the input leading to a step at the output from -15.6V to +15.6V. The step at the input occurs at $t = 40\mu s$

the noise spectral density of the amplifier. The goal is to introduce as little additional input referred noise into the pass-band as possible. Noise contributions from frequencies above the cutoff frequency will be filtered by the circuit itself. Additionally, there are further passive filters in the signal path to the trap electrodes. Passive filters will not be considered in this chapter as the implementation is specific to the setup in which they are integrated.

Lastly, we can analyse the step response of the amplifier. Two key characteristics are the overshoot (or undershoot) and settling time. The settling time depends on the allowed error band of the output voltage. Common values are in the range between 1% and 5% error. For this simulation we used a step from -2V to +2V at the input. The response overshoots to 16.54V by about 3%. The response does not display ringing after the initial overshoot. The settling time for 2% error is 12 μs .

3.2.1 Alternative Topologies

By not populating some components or using jumpers to bridge a signal we can also use different topologies with the same PCB layout. Not populating C_{out} and C_F and using a 0Ω jumper for R_{out} results in the standard Sallen-Key topology as in Fig. 2.3. A next possibility is to jump both R_1 and R_2 and not populating C_1 and C_2 . This leaves a non-inverting topology with additional noise suppression

filters.

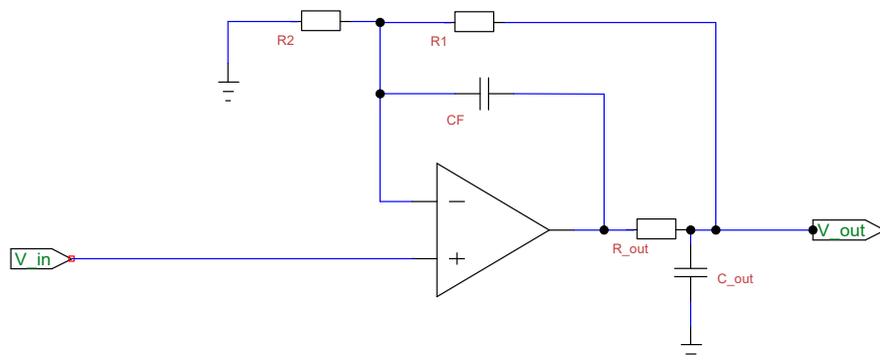


Figure 3.5: Non-inverting amplifier with additional noise suppression filters. This topology is taken directly from the data sheet of the LTC6091 [13]. Note that also the naming convention for the resistors has been adopted accordingly.

3.3 Passive Component Selection

In this section we briefly discuss some key considerations when choosing packages for passive components and proceeding to buying a specific model. The package will determine the footprint on the PCB. Most modern circuit boards use surface mount devices (SMD) instead of through-hole components. Smaller packages are easier to place and route, while also achieving a higher component density. However, there are values that might not exist for a given package. This is especially true for capacitors where there is a limit to the maximum achievable capacitance. It is therefore advisable to go with the smallest package which still offers the flexibility to slightly alter values. When looking for suitable components there are additional factors which one needs to take into account. We start our discussion with resistors. Resistors are usually more *package-agnostic* when it comes to values and tolerances. However, smaller packages also have a lower power rating. The maximum power dissipated by a resistor can be estimated from the schematic or simulated. Low tolerance resistors ($\pm 0.1\%$) are widely available. When choosing capacitors, there are more things to consider. SMD capacitors are usually multi-layer ceramic capacitors (MLCC). The used dielectric determines the temperature coefficient and DC bias derating. Derating describes how the effective capacitance changes with applied DC voltages. Ideally, there would be no dependence on applied voltage. The temperature coefficient describes the change of capacitance for a given temperature change with respect to a control temperature. For many dielectrics the effective capacitance decreases with applied bias. Furthermore, the voltage rating is also affected by the dielectric. Ideally, the maximum applied voltage should be well below the rating. For this design, the main goal was to only reach half of the voltage rating during operation. Fortunately, the C0G (NP0) dielectric is temperature stable, experiences little voltage derating, and is frequently rated for hundreds of volts. Using all considerations, the package size 0805 (imperial, 2012 metric) has been chosen for all passive filter components². The tolerances have been chosen to be 0.1% and 5% for resistors and capacitors respectively. Low tolerances are needed if the design is susceptible to component variations. The added benefit is that this also reduces channel variability.

²Excluded are decoupling capacitors which need to be placed as close as possible to the supply pins of the op-amp. This necessitates the usage of a smaller package.

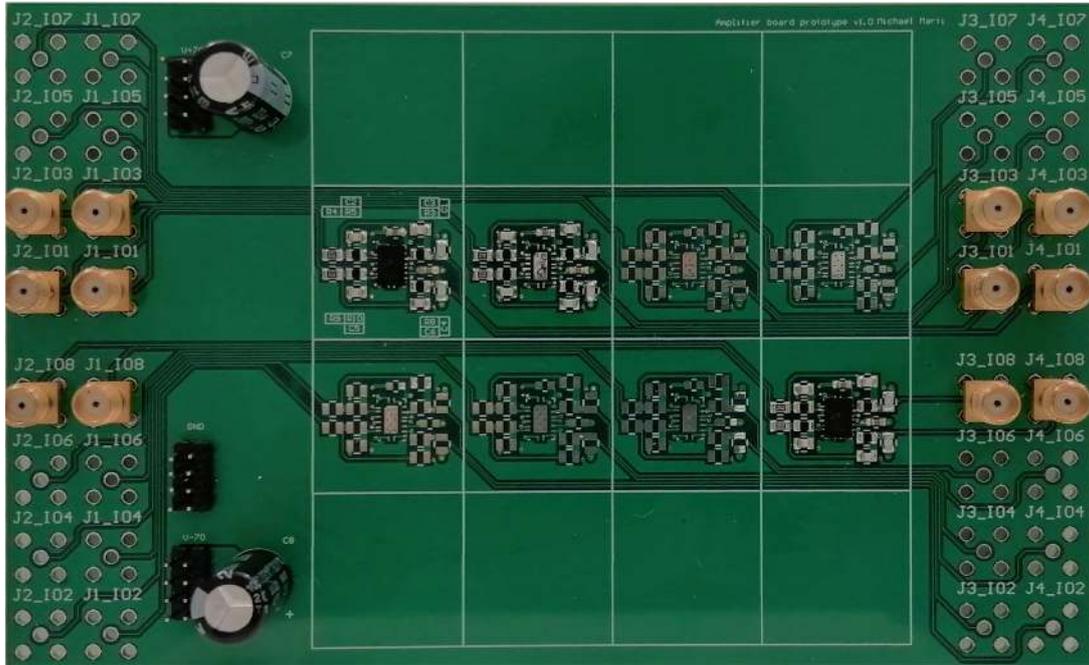


Figure 3.6: Image of the amplifier prototype board. 4 channels have been fully populated. It is noted that the solder side also features components.

3.4 Prototype Board

It is vital to determine how the physical layout of the PCB and the chosen components affect the performance of the design. In order to be able to assess the design, a prototype board has been designed first. Most importantly, the input and output of the amplifier circuit should be easy to interface. Otherwise, the layout should be as final as possible such that it can be carried over to the full design. An image of the prototype board can be seen in Fig. 3.6. Inputs (left-hand side) and outputs (right-hand side) are interfaced using SMA connectors. Power is supplied with pin headers. Four channels have been fully populated mainly for two reasons. Firstly, we can infer channel variability. Secondly, we also need to investigate potential cross-talk between the channels due to coupling between the channel traces on the PCB or coupling between the two channels on one amplifier package.

3.5 Design Verification

In this section we replicate the simulations shown in 3.2 using the prototype board. The frequency and step response has been measured using the Digilent Analog Discovery 2 kit³. The board was powered using a switched bench-top power supply unit (PSU). The positive and negative supply voltage was set to +30V and -30V respectively. The results for all four channels are shown in Fig. 3.7. In addition to the simulated design, the transfer function of a third order Butterworth filter with the same cutoff frequency as the simulated design is plotted for comparison. We can see that the measured magnitude response qualitatively agrees with the simulation. The effects of the layout seem to be beneficial in this case. The overall filter is less over-damped and has a slightly higher roll-off then expected. The channel variability is insignificant. The roll-off however is lower compared to a third order Butterworth filter.

³Digilent Analog Discovery 2 sale page

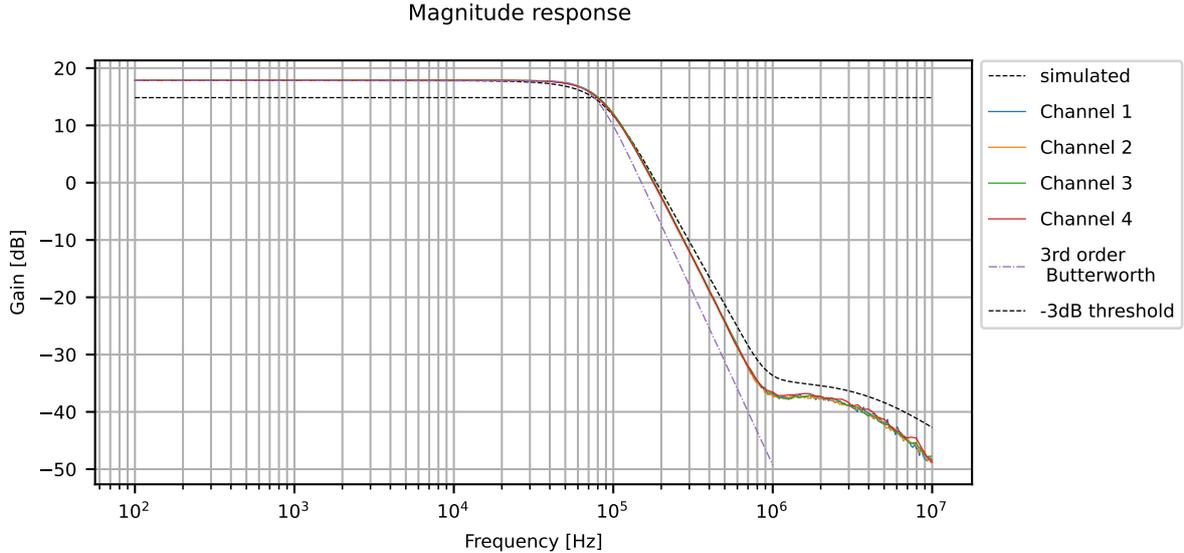


Figure 3.7: Magnitude response of four channels on the prototype board. Added for reference are the simulated result as well as the magnitude response of a 3rd order Butterworth filter.

This results also translates to the step response. The response for an input step from -2V to $+2\text{V}$ and the calculated slew rate are depicted in the right panels of Fig. 3.8. Qualitatively, the response is more aggressive: It exhibits a higher slew rate as well as overshoot. This also negatively affects the settling time. The overshoot increases to approximately 6.8%. If we consider the usage case of this amplifier board this is less of an issue. There will most certainly be additional filter stages which can be used to counter higher overshoots if necessary. We observe that the overshoot is not independent of the step size. The left side panels of Fig. 3.8 show an input step from -0.1V to $+0.1\text{V}$ with all other parameters left unchanged. We see that the measurement series follows the simulation more closely. Next we investigate the potential cross-talks between channels. For this figure we do not have any simulation results nor does the data sheet of the LTC6091 provide any information about cross-talk between the two channels housed inside one single package. Most likely, the cross-talk will then be a result of the signal routing on the PCB. This measurement was performed in the same fashion as the frequency and step responses with the only difference that the measured output channel was mismatched with respect to the input channel. There are two types of neighboring traces. Firstly, there are channels where the input line is routed in parallel to an output line over a longer distance. Secondly, there are channels where the corresponding output traces are routed in parallel over longer distances. The two types are indicated in Fig. 3.9. For the following measurements the supply voltage has been increased to $\pm 40\text{V}$ and the input amplitude has been set to 5V which is the highest possible for the signal generator of the Analog Discovery 2. The maximal output amplitude is therefore approximately 39V . As we do not expect strong coupling between channels we want to apply large signals to the amps to increase the signal-to-noise ratio. The oscilloscope inputs of the kit are only rated for $\pm 25\text{V}$ signals. In order to not damage the device, it is crucial to never correctly match input-output channels for this specific setup. The measured responses for both neighbor cases are shown in Fig. 3.10. The most significant cross-talk can be observed for output-output neighbors which peaks at -56dB in the frequency range around 100kHz . Input-output pairs exhibit the same response qualitatively but with reduced peak. We see a second increase in the magnitude for frequencies above 8MHz . Judging by the noise apparent in the measurements for gains below -70dB these values only suggest that cross-talk is a non-issue. The only effective way to combat cross-talk induced by trace-to-trace coupling is to maximally increase the spacing between traces. Measurements indicate that most importantly, the output signal traces need to be spaced as far apart as possible.

Characterising the noise performance needs to be handled with more care. If switched power supplies are used for providing the op-amp supply voltage, switching noise is able to bleed through and appear at the output. It is therefore crucial to use a non-switching, linear power supply, especially when it is integrated into a experimental setup. Another possibility is to use batteries as power source. For the following measurement series, the prototype board was powered using two $+9\text{V}$ batteries in series per supply leading to a positive (negative) supply voltage of $(-)\text{18V}$. To minimize potential EMI pickup the board was also placed into a metallic box. Measurements were taken using the Stanford Research SR760

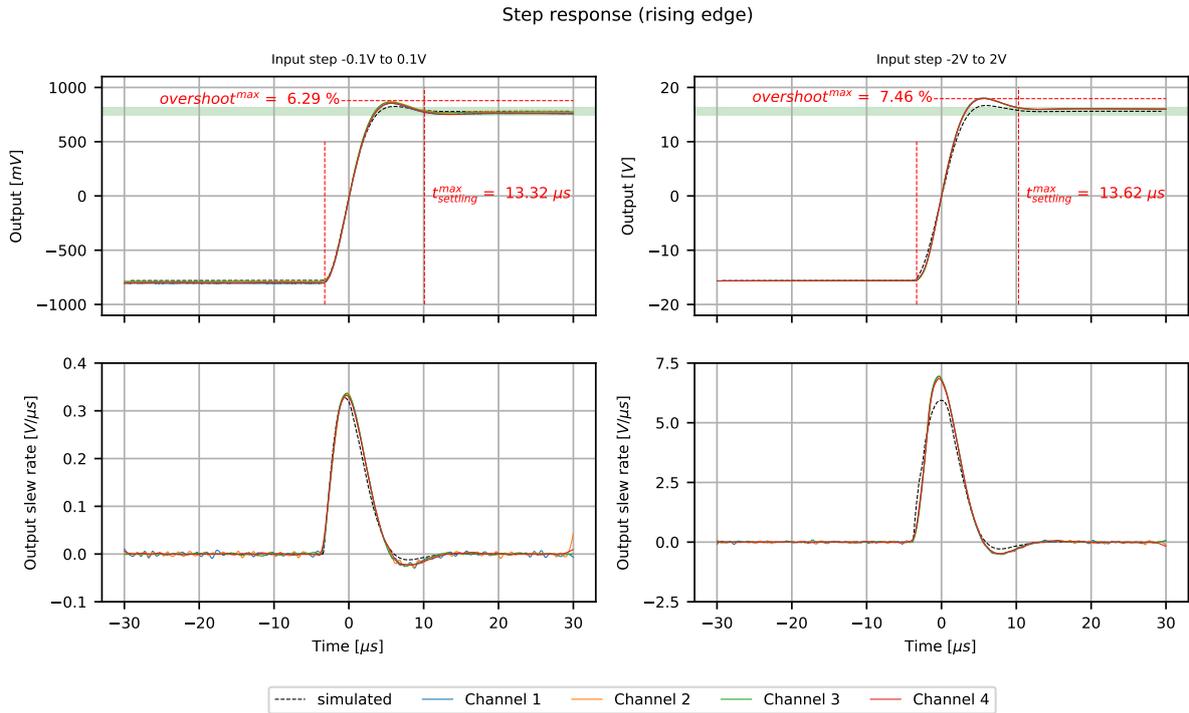


Figure 3.8: **Top:** Step response of the amplifier channel for 0.2V input step (left) and 4V step (right). **Bottom:** Calculated slew rate. We note that the slew rate has been determined numerically using the time series signal and sparsely sampling it. The samples have been interpolated using cubic splines.

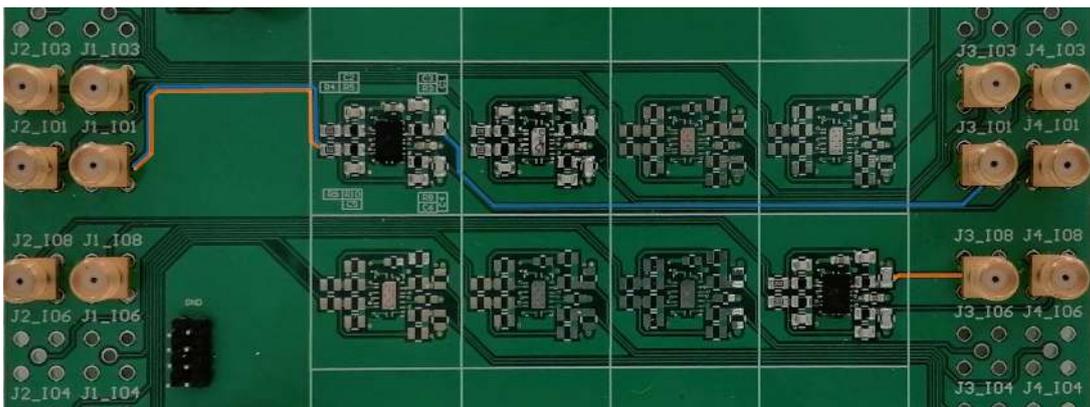


Figure 3.9: The inputs enter an amplifier block in the top left and outputs leave in the bottom right. Amplifier blocks have white borders. Orange overlay: Depiction of the input-output pair with potential cross-talk. The output signal trace of the top left amplifier is routed next to an input signal trace of the bottom right. Blue overlay: depiction of the output-output pair. Both output signal traces are routed in parallel.

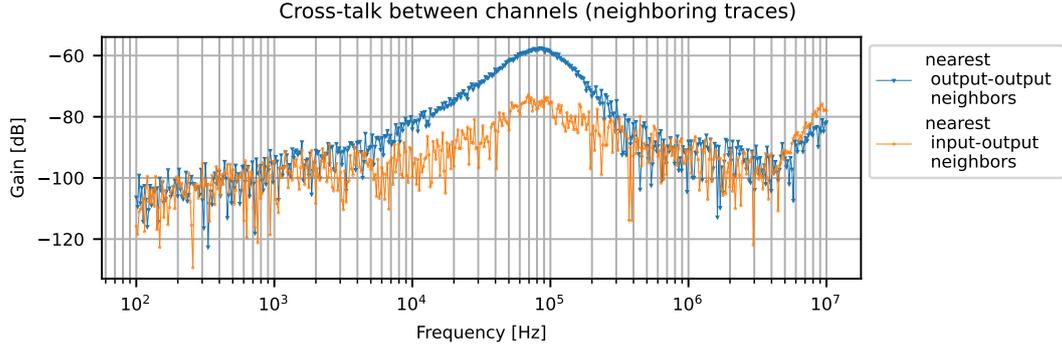


Figure 3.10: Magnitude of cross-talk between neighboring amplifier channels.

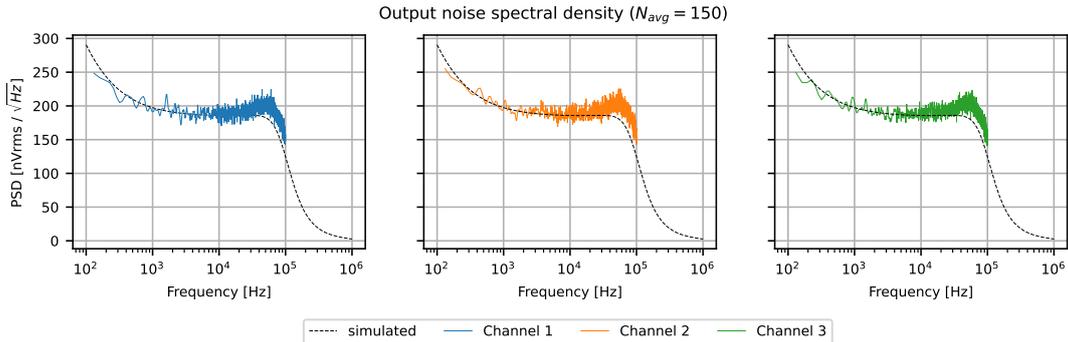


Figure 3.11: Output noise spectral density measurement for three of the four populated amplifier channels. We see similar densities among all channels. While mostly in agreement with the simulation, we observe an increase in output noise in the frequency region close to cutoff.

spectrum analyzer. The noise floor is typically around $7 \frac{\text{nV}}{\sqrt{\text{Hz}}}$ which is two magnitudes below the expected noise level at the output of the amplifier circuit. The noise floor was verified in the frequency range 100Hz to 100kHz by measuring at the noise spectral density of the ground potential at the PCB. This is done to also include the cables used to interface the analyzer. This gives a better picture of the noise floor than simply terminating the analyzer input with 50Ω . Fig. 3.11 shows the noise spectral density for channel 1 to 3. Again, we observe that the simulation and measurement mostly agree. However, we see an increase in noise density in the cutoff frequency region. The easiest way to reduce this peak is to increase C_F . This will in turn decrease the bandwidth. This leads to no suggested changes to the layout.

3.6 Full Amplifier Board

After the specifications have been assessed and verified, we can carry over the layout of the prototype board to the full 32 channel board. We proceed by summarizing the changes and improvements.

- **Power connectors:** The pins have been replaced by screw terminals.
- **Signal inputs:** Instead of using SMA connectors, we use a 2 by 13 pin header. The connectors are placed on the solder side of the PCB.
- **Signal outputs:** All 32 outputs can be interfaced with a SCSI-III 68 position connector. Additionally, mounting holes for a front panel has been added for installation in an 19" rack unit.
- **Trace spacing:** All signal traces have been spaced as far apart as possible. Spacing between output traces have been prioritized.
- **Improved decoupling:** More decoupling capacitors have been placed. Firstly, more bulk capacitance is available close to the screw terminals. Moreover, one LMCC per group of four amplifiers

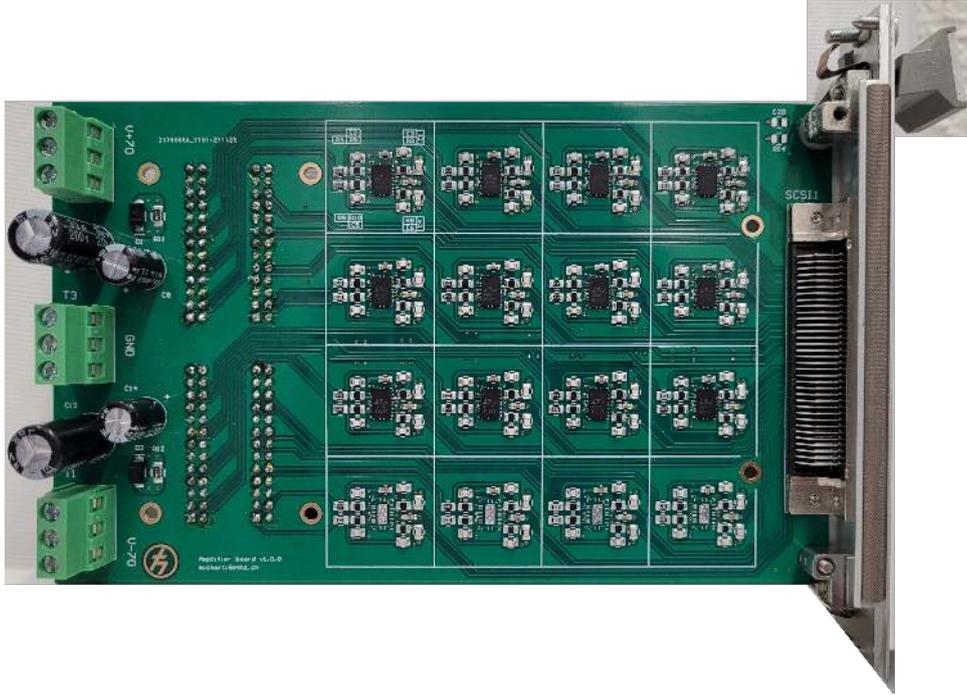


Figure 3.12: Full amplifier board prototype. Only 24 of the 32 possible channels have been populated. As only 22 channels are needed for setups which can make use of this version, we reduce power consumption and the associated heat generation by omitting the eight unused channels.

have been added to the design. The capacitance is chosen to be between the bulk capacitance and the closest decoupling capacitors.

- **Safety & protection:** The prototype board lacks pull-down resistors for the power rails. This can be dangerous as the bulk capacitors only discharge slowly - potentially shocking a user when handling the board even if the power supply is disconnected. To protect the amplifier from damage, diodes have been added to the power rails. This prevents attaching power with the wrong polarity. Even for $\pm 30\text{V}$ supplies, the op-amps are left permanently damaged if polarity is reversed.⁴

An image of an assembled full board can be seen in Fig. 3.12.

3.7 Power Supplies & Rail Regulation

Next, we discuss how to power the amplifiers when installed into an experimental setup. Theoretically, we should be able to support up to 50 mA to all 32 channels simultaneously to make full use of the LTC6091's capability. This would necessitate a power supply which can provide up to 1.6A at $\pm 70\text{V}$. If we consider the use case at hand, this estimate is most likely too high. If we use the amplifier board purely for DC voltage generation, the current demand decreases significantly. Even if ion transport experiments are considered, transient current draw spikes are alleviated with the available decoupling capacitance. As a final argument we also note that it is unlikely that all channels operate at their maximum output current rating. Using all these considerations, we aim for a 1A at 75V linear power supply for each power rail. Options are limited and an Acopian A75MT100M-230 Gold Box was chosen. The same brand also offers higher output current models if needed. Higher output current models come in a larger case and it needs to be made sure that the supply can be integrated into the setup on a case by case basis. We are not using the output of the power supplies as power rail. We will further add line regulation to the design. Using three-terminal regulators, we step down (up) the voltage from (-)75V to (-)70V. We make use of the LM317 for positive rail regulation and the LM337 for negative rail regulation. We note that different manufacturers (such as Texas Instruments and onSemi) produce the regulators under the same base model name. The manufacturer does not matter and one can order

⁴ These values have been experimentally verified, if not on purpose.



Figure 3.13: Image of the regulator board. As with the amplifier board, power connectors are screw terminals. Otherwise, the board is mostly a heat sink for the three-terminal regulators.

the regulators according to availability. The only restriction is to use the package TO220 such that the regulators can be soldered to the PCB. To achieve the best results, these regulators would need to be placed on the same PCB as the op-amps. Due to thermal considerations we decided against trying to fit everything onto one board. The power which needs to be dissipated by the op-amps is already enough to heat the whole PCB considerably. The regulators where in turn moved to a separate PCB. An image can be seen in Fig. 3.13. Heating has two negative effects. Firstly, it increases thermal noise due to resistors. Secondly, it limits the maximum output current of the three-terminal regulators.

Due to the large heat sink, the regulator board could also be used in scenarios with output currents up to 1.5A at the some dropout voltage. Both boards need to be installed close to each other and the wiring between both boards should be kept as short as possible.

Chapter 4

Full System Testing

In this chapter we will discuss how the amplifier board is integrated into the experimental setup. We describe all the connections for the signal path as well as for the power distribution. We will present preliminary results from usage in a different but similar setup.

4.1 Overview Signal Path

4.1.1 Fastino DAC board

Up to this point we have not discussed where the inputs to the amplifier board are generated from. For this we use the digital-to-analog converter (DAC) board Fastino which is part of the Sinara hardware ecosystem [11]. An image of the board is shown in Fig. 4.1. The Sinara ecosystem focuses on developing hardware used by quantum science laboratories. The Fastino features 32 DAC channels with a $3 \frac{\text{Msamples}}{\text{s}}$ update rate (simultaneous on all channels). The featured DAC has a resolution of 16bits while the output range of the output stage is $\pm 10\text{V}$. The outputs of the channels can be interfaced using either 2 by 13 pin headers or by a SCSI-III connector. Each pin header holds a sub-group of eight channels whereas the SCSI-III connector features all 32 outputs in a single connector. The samples are transferred using a serial interface. The Kasli board is a card also part of the Sinara ecosystem able to transfer sample data received from a control computer to the Fastino. The control system developed by the group in which this thesis was written in has been adapted such that it can be used with the Fastino.

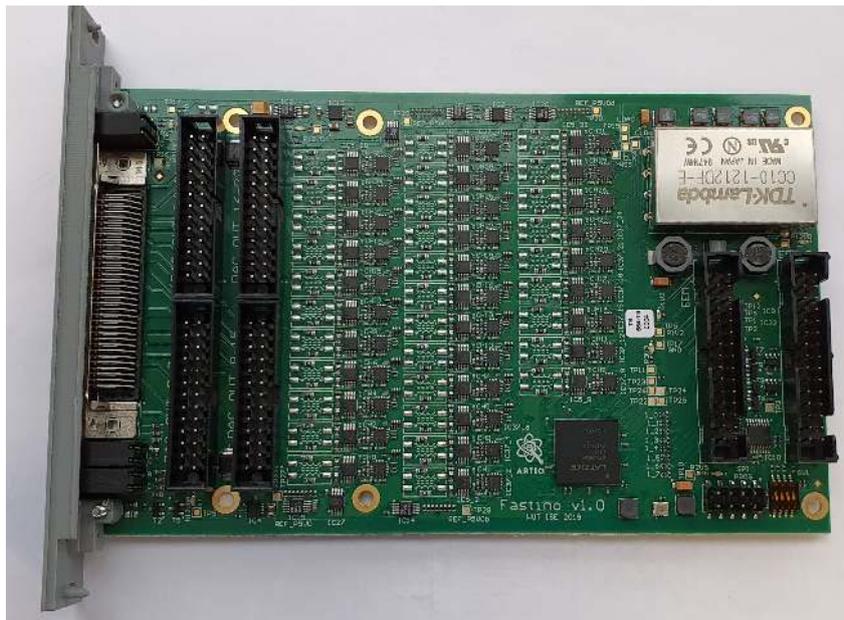


Figure 4.1: Image of the Fastino DAC card taken from [11].

4.1.2 Board Interconnect

All DC voltages will be generated by the Fastino. Ribbon cables connect the output of the Fastino with the input of the amplifier boards. We have made sure that the pin layout is compatible between the two boards. The output of the amplifier boards (SCSI-III connector) is also compatible with the outputs of the Fastino. One important difference is that we have introduced a permutation of channels. This means that the channel-to-pin mapping at the Fastino SCSI-III connector is not the same as the mapping of the connector of the amplifier boards. This can be corrected in software. For the re-mapping we refer to the schematics. Up to this point, there is no use for the SCSI-III output of the Fastino. In principal, one might use the un-amplified signal for monitoring. If not all 32 channel are used for trapping voltages, it is also possible to use them for remote control. In this setup, we only make use of 22 channels for trapping. Four of the remaining channels are used to externally control a piezo driver amplifier. A small custom PCB has been designed to breakout 8 channels from an SCSI-III connector to 8 separate SMA connectors.

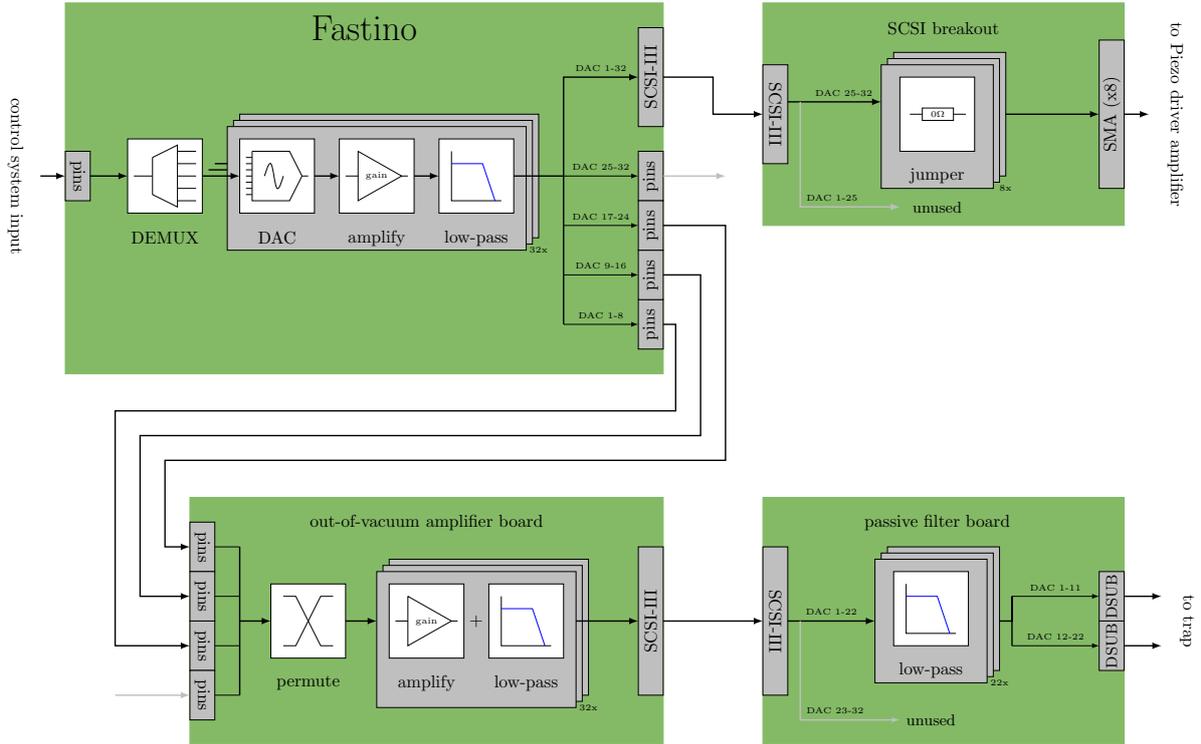


Figure 4.2: Block diagram of the overall system for cavity piezo driver and DC electrode voltage control including the amplifier board.

Following the amplifiers, there is an additional stage of passive filtering. The specific implementation will differ for each setup. The following shall only serve as an example and can be adapted to meet requirements. Generally, these passive filters are placed physically close the trap. Their main objective is to further reduce noise and filter any electromagnetic interference picked up by the connecting cables. For this implementation, the filter board also changes the interface from the SCSI-III connector to two separate DSUB-15 connectors matching the DC feedthrough of the vacuum chamber in which the ion trap will be placed. An enclosure will hold the 100mm by 100mm PCB as is shown in Fig. 4.3.



Figure 4.3: Passive filter board fitted into an enclosure.

4.2 Overview Power Distribution

In this section we discuss how to interconnect all components for power distribution to the amplifier boards. We make use of the two identical power supplies by connecting them in series. The positive output of the negative rail supply will therefore be at the same potential as the negative output of the positive rail. This potential will further be the main ground for all components involved in generating the DC electrode voltages. We note that both the negative and positive output of the power supplies have additional sense inputs. As the name suggests, this input is used to feedback the output voltage for active regulation by the PSU. The Acopian models are shipped with a bracket connecting output and sense line directly. Removing this bracket, one can also use a wire and connect sense and output at the load. With this modification, the voltage drop across the connecting wire can be compensated by the power supply. To avoid oscillations from over-regulation, the power and sense wire need to be twisted to avoid creating a loop which can act as an inductance picking up magnetic fields. To avoid potential damage to the Acopian models, the sense lines *must* be connected using one of the two methods mentioned earlier. The regulator board will then create the amplifier board power rails. It is advised to use as short wires as possible to interconnect amplifier and regulator board. Otherwise, the performance of the three-terminal regulators will deteriorate. This necessitates that the two boards are placed in neighboring slots if installed into a 19" rack unit. A diagrammatic overview of the power distribution is shown in Fig. 4.4.

4.3 Preliminary Results

The whole system consisting of regulator and amplifier board and power supplies has been put into a different setup in the same group where this thesis was written in. The testing capabilities do not change as the two setups are very similar. For example, the used ion trap is identical. This allows for testing of the whole DC signal path as well as the power distribution to the amplifier boards. However, at the time of writing, experiments which rely on low noise DC trap voltages have not been possible yet. Powering the board with the chosen power supplies in tandem with the regulator board is working. We can also connect the Fastino outputs using ribbon cables to apply an input to the amplifier board. Following the signal path, we confirm that the connectivity from the passive filters up to the trap electrodes is as intended. As a first proof of concept attempts of trapping ions using the amplified signal has been successful. In Fig. 4.5 we show an image of a trapped calcium ion using an EMCCD camera.

Up to this point, quantitative statements concerning PCB heating and power dissipation of the amplifier boards have not been disclosed. This is owed to the fact that the ability to dissipate heat to the ambient is highly dependant on the PCB layout. Data sheets may provide estimates, but in the

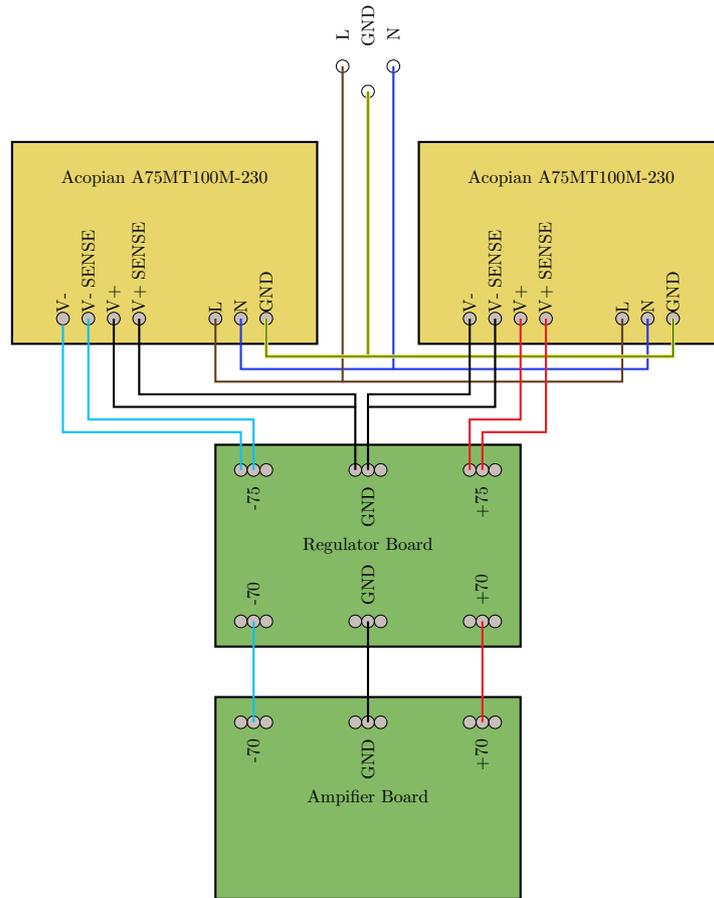


Figure 4.4: Block diagram of the power distribution to the amplifier boards. The outputs and their corresponding sense lines must be twisted to avoid oscillations. This is omitted in the diagram to not impede readability.

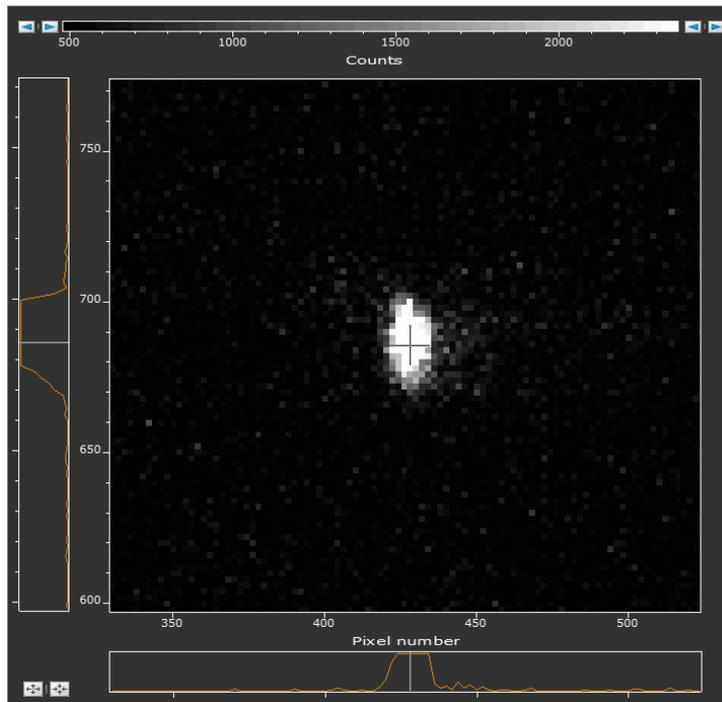
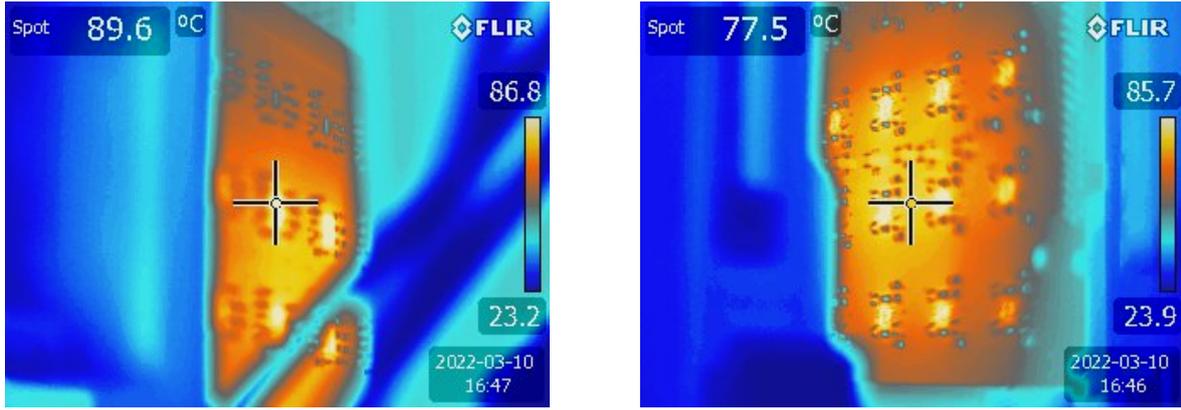


Figure 4.5: Image of an trapped ion using an EMCCD camera. The amplifier board has been used to apply the DC electrode voltages.



(a) Infrared image of the top side of the amplifier PCB. The spot is aimed at the package of an LTC6091.

(b) Back side of the amplifier PCB. The spot is aimed at the solder mask to assess passive component heating. The hottest spots are exposed metal pads.

Figure 4.6: Thermal imaging of the amplifier board when used for trapping ions. This does not include shuttling experiments as they will lead to transient heating responses.

case of the LTC6091 they are of limited use¹. To close this gap, a thermal camera has been used to determine key temperatures of the amplifier board during operation. In Fig. 4.6a we take a temperature reading directly from the amplifier package. The package temperature is about 90°C. The temperature of interest however is the junction temperature of the internal MOSFETs which should not exceed 150°C. To approximate the junction temperature we can use the thermal resistance $\theta_{JC} = 15 \frac{^\circ\text{C}}{\text{W}}$ from junction to case as specified in [13]. If we disregard transients, the average power dissipation is well below 2W. We therefore still have an adequate safety margin before the amplifiers would perform an autonomous thermal shutdown. Next, we also measure the average PCB temperature as seen in Fig. 4.6b. Hot spots of the PCB reach about 80°C at the surface. This indicates that the power dissipation from the amplifier case to the PCB is effective. The PCB area will further increase dissipation to the ambient air. We conclude our discussion on thermals with two remarks. Firstly, we note that the front of the rack unit has not been fully closed to make measurements possible. It is to be expected that the temperature will increase when fully closed as there is less air flow over the amplifier PCB combined with a decreased temperature differential between ambient air inside the enclosure and the amplifier PCB. Perforating the bottom and top lids of rack enclosures is recommended to maximize air flow. Applying this measure should allow for operating the board without any active cooling or heat sink. Secondly, it has not been assessed yet how closing the front will affect the temperatures of the power supplies and regulator board. The main side effect would be a decreased supply current. This needs to be evaluated on case by case basis as the enclosure has a significant impact.

4.4 Possible Improvements

Future iterations of the amplifier board can be improved in many ways. We follow by providing a list of possible changes for the next iteration. We begin the list with the most pressing functional and quality-of-life improvements and end on minor details nevertheless deserving attention.

- **Power Connectors:** Screw terminals are less than ideal to work with. Primarily, it is not guaranteed that wires make good contact after installing the card into the rack unit. Adding to that is the fact that removing the amplifier card is a hassle as it is connected directly to the regulator board. The situation with the regulator board is even worse as it is also connected to the Acopian power supplies. In the future 15 position DIN41612 type H connectors could be used to solve the problems. Both the amplifier and regulator cards will then slot into an receptacle and no hard to reach screws need to be loosened to remove the card.

¹Linear Technology only provides estimates for 2-layer PCBs with only one amplifier package populated. This is quite different from placing up to 16 on a 4-layer PCB.

- **Power Rail Indication LED:** The amplifier board does not have any indication that the power rails are present. Adding LEDs visible at the front panel are an easy addition. Previously custom made panels can be updated by drilling additional holes.
- **Ribbon cable connectors:** The 2 by 13 ribbon cable connectors on the bottom are flipped by 180 degrees. This necessitates an 180° twist/torsion of the ribbon cable which in turn increases the minimal length of the cable. Either the headers are placed on the top or are left at the bottom and rotated by 180°. Leaving them on the bottom side would still allow for the shortest possible connection and the correct orientation will also cancel out the permutation of the channels introduced to make the design easier to route.
- **Thermal shutdown indication:** Each channel of the LTC6091 feature an over-temperature flag output. This would in turn inform users of the thermal limitation of the board. The remaining consideration is how to display this signal to the user. Either, this signal is lead to physical pins for each channel individually and/or all signals are logically OR-ed and the resulting signal is used to drive a status LED integrated into the front panel. One could also group banks of eight channels together (those belonging to one ribbon cable) instead of all 32 channels
- **SCSI-III connector voltage rating:** We note that the used connector for the amplifier outputs has a voltage rating of only 30V ². The board is still safe to use. The absolute maximum applicable voltage difference between two adjacent pin positions is approximately 400V. At this point the dielectric insulator starts to break down.
- **Output stage disable:** It is possible to turn off the output stage of the LTC6091. The main use case for this functionality are either the self-controlled thermal shutdown using the over-temperature flag or multiplexing applications. We note that it is possible to disconnect the amplifiers without immediately losing the trapped ions. This is possible due to the fact that the electrodes lose charge through two possible mechanisms. Firstly, there is a connection to ground through the negative feedback path. If one is able to increase the two resistors setting the DC gain by an order of magnitude without changing the ratio, the capacitor discharge should be minimal during one experimental shot. The second way to lose charge are leakage currents either from the trap electrodes themselves or through the capacitors present in the passive filters which are left connected. It could be evaluated if turning off the output stage would lead to a significant noise reduction. The trap would then need to be periodically recharged which could be controlled using the output stage enable. This possibility has not been explored as in the current layout, the output enable is tied to the over-temperature flag, making the output disable hard to trigger externally.
- **Heat sink:** Upcoming use cases may impose a greater thermal load onto the amplifiers. Mounting holes for a custom heat sink are already available on the PCB. As some of the passive components have a higher height clearance, a custom heat sink must be designed and manufactured. Flat heat sink would only make contact with passive components. One could also try to bridge the gap between a flat heat sink and the LTC6091 package by using thick thermal pads.

² *This type of connector was designed for digital signal transmission, not trapped ion experiments. We want to highlight that it is also possible that vendors and/or manufacturers simply do not bother achieving high voltage ratings if they are not part of the typical application.*

Chapter 5

Conclusion & Outlook

During this part of the thesis, we documented the development of a low noise amplifier design for use in trapped ion quantum experiments. We introduced the basics of analog signal processing and active filters. We proceeded with a discussion of component criteria and a final selection. We verified the performance using a prototype board. As its layout allows for different circuit topologies to be realized, the prototype can be used as a development platform for different active filters. We documented how the amplifier board and its surrounding hardware can be integrated into an experimental setup. We started the process of verifying the designs suitability for trapped-ion quantum information experiments by trapping calcium ions. We concluded this part by giving a comprehensive list of possible improvements to the current design. However, the development and performance assessment is by no means complete. We provide next steps which would advance the design:

- **Motional coherence measurements:** Being able to trap ions is only the first step to assess suitability of the amplifier board. A practical test for the noise performance are motional coherence measurements. Determining the heating rates will enable future developers to compare the design against other solutions.
- **Reliability testing:** At the end of writing this thesis, the amplifier board setup has been running for about a month without incident. To ensure that we are past the early failures ¹, it is key to ensure correct operation of the hardware setup.
- **Implementing improvements:** Functional and quality-of-life improvements should be implemented before deploying the system in other setups.
- **Designing different active filters:** The amplifier layout allows for different circuit topologies to be realized. Using the prototype board one can design filters with different parameters for the to cutoff frequency, output voltage swing, et cetera.

¹Early failures and their rate are also gruesomely referred to as (product) infant mortality.

Part II

Bayesian Schemes for Fast Single Qubit Detection

Chapter 6

Concepts

In this chapter we introduce the concept of Markov chains, hidden Markov models, and qubit state readout in trapped-ion quantum information processing. We assume that the reader has a solid understanding of probability theory. Sections 6.1 & 6.2 are a highly condensed summary of the lecture notes [16]¹ by Hans-Andrea Loeliger while section 6.3 is a summary of chapter 3 in [17]. We suggest to first read the corresponding chapter before returning to the aforementioned sections.

6.1 Hidden Markov Models

Before introducing the concept of a hidden Markov model, we first define under which conditions random variables form a Markov chain. Let $X_{0\dots n}$, $n \geq 2$ be random variables. We say that these random variables form a Markov chain if for any $i \in 1, \dots, n-1$, the joint probability distribution conditioned on X_i has the following property:

$$p(x_0, \dots, x_{i-1}, x_{i+1}, \dots, x_n | X_i = x_i) = p(x_1, \dots, x_{i-1} | X_i = x_i) p(x_{i+1}, \dots, x_n | X_i = x_i) \quad (6.1)$$

Knowing the outcome of the random variable X_i will split the joint probability distribution into two independent factors. The split occurs at the observed random variable. A common notation to indicate that random variables form a Markov chain is as follows:

$$X_1 \leftrightarrow \dots \leftrightarrow X_n \quad (6.2)$$

This allows us to write the joint probability distribution as:

$$p(x_{1\dots n}) = p(x_1) p(x_2 | x_1) \dots p(x_n | x_{n-1}) \quad (6.3)$$

We can think of Markov chains as a system being in a starting internal state $X_0 = x_0$. When simulating the system, we transition non-deterministically from the state x_0 to a state $X_1 = x_1$. The Markov property ensures that the transition probabilities only depend on the current state, independent of the full state history. We can extend this model to the case where it is not possible to identify or measure the internal state directly. In a hidden Markov model, we are able to make observations y_i which either depend on the internal state transition $p(y_{i-1} | x_i, x_{i-1})$ or even simpler only on the state during which the measurement was taken $p(y_{i-1} | x_{i-1})$. This allows us to write the joint probability distribution of internal states $x_{0\dots n}$ and observations $y_{1\dots n}$ as:

$$p(x_{0\dots n}, y_{0\dots n-1}) = p(x_0) \prod_{i=1}^n p(x_i | x_{i-1}) p(y_{i-1} | x_{i-1}) \quad (6.4)$$

There are many questions which one can answer about this structured model. In our case, we are most interested in the initial state x_0 or the current state x_i for reasons which will become clear. Explicitly, we want to estimate the state at any given time index i conditioned on fixed observations $y_{0\dots n-1}$:

$$\hat{x}_i = \arg \max_{x_i} \sum_{x_k \neq x_i} \frac{p(x_{0\dots n}, y_{0\dots n-1})}{p(y_{0\dots n-1})} \propto \sum_{x_k \neq x_i} p(x_{0\dots n}, y_{0\dots n-1}) \quad (6.5)$$

¹Specifically chapter 8

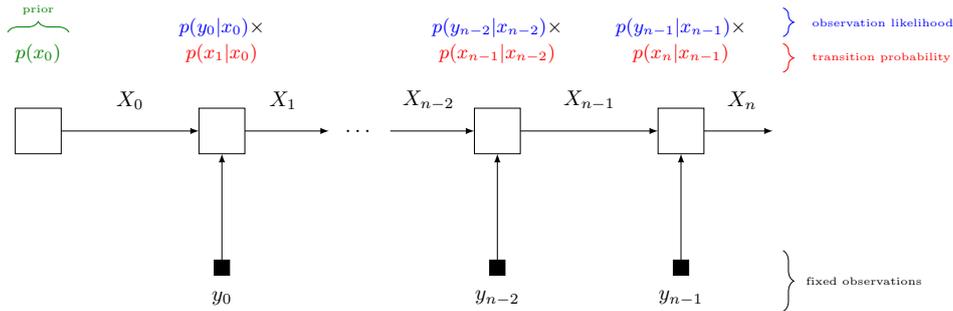


Figure 6.1: Factor graph of Eq. 6.4. The factors which are represented by a box are indicated above. The box corresponding to the prior probability $p(x_0)$ (furthest to the right) is a half box. The edge corresponding to X_n is a half edge. We indicate the fixed observations with smaller filled boxes.

We note that the sums over x_k in Eq. 6.5 are a shorthand notation for marginalizing over all random variables excluding the variable we want to maximize. Without making use of the structure of hidden Markov models, the marginalization is usually computationally infeasible. We proceed to explore a general method to answer this and related questions efficiently.

6.2 Factor Graphs and Belief-Propagation

6.2.1 Factor Graphs

Factor graphs are a graphical representation of the constituent factors of a global function. We will make use of Forney factor graphs for our analysis. Each factor (which is a function in its own right) is represented as a box. Each variable is represented using a unique edge. If the edges are chosen to be directional, the direction of the edge has no influence on the global function. It is however useful for unambiguous notation later on. If a factor depends on a hidden variable, the corresponding edge connects to the box of the factor. This would lead to a strict limitation of the structure of a given global function as a variable could only appear in two factors. There is an easy way to overcome this limitation using the notion of equality constraints. We refer to [16] for a complete and more rigorous introduction to factor graphs. Boxes which only have one edge connected are called leaf boxes. Edges which only connect to one box are referred to as half edges. If we have fixed observations $y_{0\dots n-1}$ we indicate this by adding a small filled box to the edges $Y_{0\dots n-1}$. As an example we can look at the factor graph of the joint probability function of a hidden Markov model as in Eq. 6.4. The corresponding factor graph is shown in Fig. 6.1. We note that this is not the most refined factor graph of Eq. 6.4 as all but one box contain two factors instead of one. This is a different, non-standard method to overcome the limit of variables ability to appear in at most two factors. We proceed with the concept of closing boxes. We can group neighboring boxes together to create one larger box. Edges (and the corresponding variable) which are fully enclosed by the larger box are therefore not variables of this generated factor anymore. Moreover, the global function must no longer depend on enclosed variables. We can perform the elimination of fully enclosed variables in two ways. Firstly, we can close boxes by summing (or integrating) over all possible values of the enclosed variables. If we apply this to joint probability distributions, we effectively take the marginal distribution. Secondly, we can think of taking the maximum marginal over the enclosed variables. We will only make use of the first case and do not further comment on max marginals.²

6.2.2 Belief Propagation

We continue with our discussion on factor graphs by considering the belief propagation algorithm. Belief propagation is also referred to as message passing. A message μ_{X_i} out of a box is a function which only depends on the variable x_i . Edges which connect two boxes therefore have two possible messages, one leaving each box. We can now use the directed edges introduced earlier to indicate the message along the edge as $\bar{\mu}_{X_i}(x_i)$ while the message in the opposite direction is denoted as $\bar{\mu}_{X_i}(x_i)$. We define a message out of a box representing a factor with m variables incoming $g(x_{i_1}, \dots, x_{i_m})$ on the edge of variable X_{i_1} recursively by the following rule:

²The non-standard factor graph in Fig. 6.1 is obtained by closing boxes in the standard factor graph of 6.4.

$$\mu_{\text{out}}(x_{i_1}) = \sum_{x_{i_2} \dots x_{i_m}} \left(g(x_{i_1}, \dots, x_{i_m}) \prod_{x=x_{i_2}}^{x_{i_m}} \mu_{X,\text{incoming}}(x) \right) \quad (6.6)$$

We define messages of half boxes as the factor itself. Messages *entering* a box along half edges are the neutral message. The neutral message is simply $\mu_{\text{neutral}}(x) = 1 \forall x$. Fixed observations are plugged into the connected factors. By the recursive definition, we need to start with leaf nodes and half edges.

6.3 State Readout

In trapped-ion quantum information processing, a prominent technique for qubit readout is state-dependant fluorescence. This necessitates that a detection laser resonantly couples one of the qubit states to an excited level through a closed dipole-allowed transition. This allows us to scatter many photons in a short amount of time. Because of the repeated exciting and decay, these transitions are also referred to as cycling transitions. The state which exhibits this cycling transition is called the bright state whereas the other is simply referred to as the dark state. We can use different hardware to detect the scattered photons. Readout devices include photomultiplier tubes (PMT), avalanche photo diodes (APD), and electronmultiplier or CMOS charge coupled devices (EMCCD and CMOS cameras respectively). Cameras allow for imaging multiple ions simultaneously as they provide spatial resolution. Depending on the internal state of the qubit and the used measurement device we expect different photon count distributions. In our discussion we will focus on detection using PMTs. For a summary of photon count distributions of EMCCD and CMOS cameras we refer to [18]. The detection of an emitted photon is usually a rare event, so under the assumption of constant detection rate and time-independent detection events, the photon count distribution follows a Poissonian distribution [17]:

$$p(n | \lambda) = e^{-\lambda} \frac{\lambda^n}{n!} \quad (6.7)$$

The mean λ of the Poissonian distribution is different for the two qubit states. First, we have the mean $\lambda_D = R_D t_{\text{det}}$ for ions in the dark state. R_D defines the rate we detect background photons, whereas t_{det} defines the time interval during which we count photons. Ideally, $R_D = 0$ as no stray photons would be detected. Nonetheless, there is the possibility to detect background counts. Sources of background counts can be scattered photons from trap surfaces, electrical noise during PMT readout and cosmic rays. Secondly, we have the bright state mean $(R_B + R_D) t_{\text{det}}$ which includes the photons scattered by the ion.

6.3.1 Non-Poissonian Statistics

In reality, the photon count distribution is not a pure Poissonian. The main mechanisms driving the deviation are dark-to-bright and bright-to-dark leakage at random times during the detection interval. We briefly state the most important phenomena which lead to leakage [17].

1. **Lifetime:** The excited qubit state may decay spontaneously due to its finite lifetime.
2. **Off-resonant pumping:** The detection laser can drive unwanted qubit transitions, even if the laser is detuned.
3. **Unclosed detection transition:** It is possible that the detection transition is not closed and we drive different transitions or that unwanted decay paths can occur which interrupt the cycling.

In order to take these effects into account we can adjust the mean of the Poissonian distribution. Instead of being constant during the entire detection time we allow the photon count rate to be time-dependent [17].

$$\lambda_{\text{eff}} = \int_{t_0}^{t_0 + t_{\text{det}}} R(t) dt \quad (6.8)$$

where $R(t)$ is the state-dependant photon count rate. As the ion state is time-dependant, the photon count rate implicitly depends on time. If we restrict ourselves to only one type of leakage event, it is possible to derive parametric analytical solutions [17]. Parameters are then fit using experimental data. From these we can usually extract rate equations which model the flow between different states. This

enables us to simulate the ion dynamics using Markov Chain Monte Carlo simulations. Hidden Markov models have been explored in trapped ion detection to account for multiple possible state transitions [19]. The same framework has also been applied to superconducting qubit state readout[20]. Transitions matrices and priors have been learned from unlabeled data. This eliminates the need to determine analytical models and parameter fitting.

6.3.2 Real-Time State Discrimination

After we have retrieved the photon counts, we need to estimate the state. A common technique which can be applied in real-time is thresholding. For this method we define a photon count threshold within a fixed time window. Sequences meeting the threshold or having a higher count are labelled bright. Otherwise, the qubit is labelled as dark. This decision rule is computationally inexpensive and easily applicable in real-time. If we are able to perform floating point operations in real-time, we can implement time of arrival information of the incoming photons, hoping to speed up the detection sequence. To accomplish this, we sub-divide the whole detection window into an integer number of sub-windows. After each sub-window we retrieve the photon counts. We also refer to a detection window and PMT readout as a sub-bin. After each sub-bin we are able to calculate the likelihood of the observation given the ion was in the bright p_B or dark state p_D . How the likelihoods are calculated is discussed in depth in section 7.1. Using Bayes' rule we can calculate error probabilities from the likelihood functions. We define the probability of erroneously labelling a dark (bright) state as bright (dark) as [17]:

$$e_D = \frac{p_B}{p_B + p_D} \left(e_B = \frac{p_D}{p_B + p_D} \right) \quad (6.9)$$

We are able to abort the detection sequence as soon as we satisfy the condition $\min(e_D, e_B) < e_c$ where e_c is parameter which needs to be defined in accordance with the desired fidelity of the readout. Having arrived at that point, we infer that the ion was in the bright (dark) state if $p_B > p_D$ ($p_D > p_B$). A description using pseudo code is given in Alg. 2.

Algorithm 2 Adaptive Maximum Likelihood estimator

```

Output: State of the ion
 $p_B^{(0)}, p_D^{(0)} \leftarrow 0.5$ 
for  $m = 1 \dots n_{\text{bin, max}}$  do
  Detect sub-bin
  Calculate likelihood  $p_B^{(m)}, p_D^{(m)}$ 
  Evaluate errors  $e_B, e_D$ 
  if  $\min(e_D, e_B) < e_c$  then
    if  $p_B^{(m)} > p_D^{(m)}$  then
      return bright
    else
      return dark
    end if
  Abort detection
end if
end for

if  $p_B^{(n_{\text{bin, max}})} > p_D^{(n_{\text{bin, max}})}$  then
  return bright
else
  return dark
end if

```

Using this chapter as a starting point for in-depth analysis of the maximum likelihood method and will explore how we attempt reduce detection time even further in the following chapter.

Chapter 7

State Discrimination Algorithms

In this chapter, we explore various state discrimination algorithms. We start off by characterizing the popular maximum likelihood method. We will further discuss its shortcomings with respect to readout times and propose multiple ways to address them. This will lead to algorithms which take actions on the qubit in order to speed up the state identification sequence.

7.1 Adaptive Maximum Likelihood

The general idea has been outlined in section 6.3. We first define the update rule for the likelihood functions:

$$\begin{aligned} p_D^{(m)} &= p_D^{(m-1)} p(y_{m-1} | \theta_{\text{dark}}) \\ p_B^{(m)} &= p_B^{(m-1)} p(y_{m-1} | \theta_{\text{bright}}) \end{aligned} \quad (7.1)$$

We are now able to analyze the behavior of the maximum likelihood method analytically under ideal conditions. We start off by assuming the following:

1. Photon counts follow a Poisson distribution with either mean $\lambda_D = R_D t_{\text{bin}}$ in the dark state or $\lambda_D = (R_D + R_B) t_{\text{bin}}$. This corresponds to the situation where no transitions which either result in dark-to-bright nor bright-to-dark pumping can occur. t_{bin} is the time interval of the sub-window. It is a whole integer fraction of a total available detection time $t_{\text{det, max}}$.
2. We set $R_D = 0$ assuming no stray photons.
3. We terminate when the threshold $e_c > \frac{\min(p_D, p_B)}{p_D + p_B}$ is met where p_D (p_B) is the likelihood of the dark (bright) state.

This will immediately lead to a first key observation:

$$p_D^{(m)} \propto \prod_{k=0}^m p(n_k | \lambda_D) = \prod_{k=0}^m \delta[n_k] \quad (7.2)$$

where $\delta[\cdot]$ is the Kronecker delta function. Furthermore, if an ion is in the dark state, we will always detect zero photons. We will use this to our advantage.

7.1.1 Bright State Detection

As soon as we detect even a single photon, we will immediately abort the detection sequence as $p_D = 0$. We define:

$$p_{\text{success}} = p(n > 0 | \lambda_B) = 1 - e^{-t_{\text{bin}} R_B} \quad (7.3)$$

We can then easily calculate the probability that we detect the first photon during bin k :

$$P[\text{"first photon at bin } k\text{"}] = (1 - p_{\text{success}})^{k-1} p_{\text{success}} \quad (7.4)$$

This follows a geometric distribution and we can easily calculate the average number of bins \bar{k}_B and with it the average detection time $t_{\text{det, B}}$ it takes to detect the first photon:

$$\begin{aligned}\bar{k}_B &= \frac{1}{p_{\text{success}}} \\ \overline{t_{\text{det, B}}} &= t_{\text{bin}} \bar{k} = \frac{t_{\text{bin}}}{p_{\text{success}}}\end{aligned}\tag{7.5}$$

The average detection time will certainly be at least as long as a bin time. We may investigate when the bin time tends to zero.

$$\begin{aligned}\lim_{t_{\text{bin}} \rightarrow 0} \frac{t_{\text{bin}}}{p_{\text{success}}} &= \lim_{t_{\text{bin}} \rightarrow 0} \frac{t_{\text{bin}}}{1 - e^{-t_{\text{bin}} R_B}} \\ &\stackrel{\text{de l'H}}{=} \lim_{t_{\text{bin}} \rightarrow 0} \frac{1}{R_B e^{-t_{\text{bin}} R_B}} \\ &= \frac{1}{R_B}\end{aligned}\tag{7.6}$$

We see that we are limited by the photon rate $\overline{t_{\text{det}}} \geq \frac{1}{R_B}$.

7.1.2 Dark State Detection

For $R_D = 0$, we will always receive zero photons independent of the sub-bin time. This leads to the following observations:

1. The dark state likelihood at sub-bin index i : $p_D^{(i)} = 1 \forall i$
2. For the bright state likelihood we have in contrast $p_B^{(i)} = p(n = 0 | \lambda_B)^i$
3. Combining both leads to $p_D^{(i)} > p_B^{(i)} \forall \lambda_B > 0$. The estimator will terminate as soon as the bright state likelihood is small enough.

Starting from the termination condition, we can calculate at which sub-bin k the estimator will reach said condition.

$$\begin{aligned}e_D &= \frac{p_B^{(k)}}{1 + p_B^{(k)}} < e_c \\ \Leftrightarrow p_B^{(k)} &= e^{-kt_{\text{bin}} R_B} < \frac{e_c}{1 - e_c} \\ \Leftrightarrow \frac{1}{t_{\text{bin}} R_B} \log\left(\frac{1 - e_c}{e_c}\right) &< k \\ \Rightarrow k_D &= \left\lceil \frac{1}{t_{\text{bin}} R_B} \log\left(\frac{1 - e_c}{e_c}\right) \right\rceil\end{aligned}\tag{7.7}$$

The average detection time can be calculated in the same fashion as in 7.5. Taking again the limit for sub-bin times approaching zero, we can effectively ignore the ceiling operation and we get a similar expression as in 7.6 with an additional penalty factor depending on the threshold:

$$\lim_{t_{\text{bin}} \rightarrow 0} t_{\text{bin}} \left\lceil \frac{1}{t_{\text{bin}} R_B} \log\left(\frac{1 - e_c}{e_c}\right) \right\rceil = \frac{1}{R_B} \log\left(\frac{1 - e_c}{e_c}\right)\tag{7.8}$$

We will finish our analysis by plotting the predicted average detection time. 7.1. For this setup we have set $R_D = 25 \frac{1}{200 \mu\text{s}}$ while assuming no background counts and $e_c = 10^{-4}$. We divide a detection window of 200 μs by an increasing number of sub-bins. Fig. 7.1 shows the prediction using the specified parameters. The main drawback of the adaptive maximum likelihood method is the disparity between the average readout detection time of dark and bright. We will attempt to equalize the average detection time in the following.

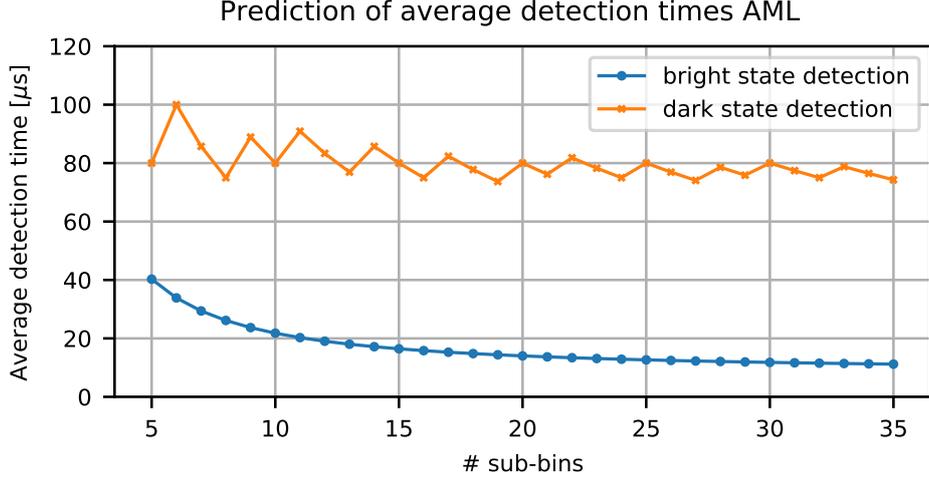


Figure 7.1: Prediction of the average detection time given a number of sub-bins. The oscillations in dark state detection is an artifact of the ceiling operation in Eq. 7.7.

7.2 π -pulse Methods

In order to shorten the average dark state detection time, we can apply a π -pulse to the ion being read out. If we successfully reached the bright state we would expect to see an increase in photon counts and may be able to abort the estimation sequence sooner. Previous studies investigated the idea of using π -pulses during qubit readout [21]. The main difference is that π -pulses were applied between two consecutive measurements of the same state. The information of the two readouts was used to reject contradictory results, improving fidelity. In the case proposed here, the π -pulse is conditionally applied during state readout to shorten the average detection time. To achieve this, we need to determine when we want to apply a pulse. More formally, if we are given a model of the ion with parameters θ and past observations $y_{0\dots n-1}$, we need a objective \mathcal{J} telling us which action a from a set of available actions \mathcal{A} we should perform before the next sub-bin. In the case of trapped ion state discrimination we can identify these formal definitions:

- Parameters θ : This includes photon count statistics given the state and the physical device used to measure and any additional dynamics like pumping rates from the dark to the bright state τ_{db} . We will use this for a convenient abuse of notation: while one usually denotes probability mass functions as $P(x | \theta)$ and probability density functions as $p(x | \theta)$ we will now assume that the distribution and its parameters are all included in θ allowing us to stay oblivious for theoretical considerations.
- Past observations $y_{0\dots n-1} \in \mathcal{Y}$: These are simply the photon counts y_i during sub-bin index i . Therefore, we can identify $\mathcal{Y} = \mathbb{N}$
- Actions \mathcal{A} : For the following state discrimination algorithms we have $\mathcal{A} = \{\pi, \text{no-}\pi\}$.

As state discrimination is a stochastic process, we can assume without loss of generality that we want to optimize for the expectation value of the given objective rather than a specific outcome. Finally, we are able to formulate the problem in a compact fashion:

$$\hat{a} = \arg \max_{a \in \mathcal{A}} \mathbb{E}[\mathcal{J} | \theta, y_{i\dots n}, a] \quad (7.9)$$

An estimator is therefore fully determined by the objective \mathcal{J} . In the special case of only having two actions, we can look at the difference of the expectations given the action and apply a π -pulse given the sign:

$$\hat{a} = \begin{cases} \text{no-}\pi, & \text{if } \mathbb{E}[\mathcal{J} | \theta, y_{i\dots n}, a = \text{no-}\pi] - \mathbb{E}[\mathcal{J} | \theta, y_{i\dots n}, a = \pi] \geq 0 \\ \pi, & \text{otherwise} \end{cases} \quad (7.10)$$

Eq 7.10 leads to decision boundaries where the difference is equal to zero. If the decision boundaries can be pre-calculated when initializing the estimator, we can lower the computational burden of calculating the expectation values in real-time.

As soon as we start to interfere during detection sequences, we need to modify how we update our likelihoods. The dark (bright) state likelihoods are referenced to the start of the detection sequence and are not to be confused with the likelihood of being in the dark (bright) state during a sub-bin. This can be done by keeping track of the number of π -pulses which have been applied during detection. If we have applied an odd number of pulses, we are in the flipped state with respect to the starting point. For an even number, we are back at the initial state. This leads to the following likelihood update rules for sub-bin index m :

$$p_D^{(m)} = \begin{cases} p_D^{(m-1)} p(y_m | \theta_{\text{dark}}), & \text{state not flipped} \\ p_D^{(m-1)} p(y_m | \theta_{\text{bright}}), & \text{state flipped} \end{cases} \quad (7.11)$$

$$p_B^{(m)} = \begin{cases} p_B^{(m-1)} p(y_m | \theta_{\text{bright}}), & \text{state not flipped} \\ p_B^{(m-1)} p(y_m | \theta_{\text{dark}}), & \text{state flipped} \end{cases}$$

We can adapt the maximum adaptive likelihood detection algorithm to incorporate the theory we discussed.

Algorithm 3 π -pulse Methods estimators

Output: State of the ion
 $p_B^{(0)}, p_D^{(0)} \leftarrow 0.5$
state \leftarrow not flipped
for $m = 1 \dots n_{\text{bin, max}}$ **do**
 Detect sub-bin
 Calculate likelihood $p_B^{(m)}, p_D^{(m)}$
 Evaluate errors e_B, e_D
 if $\min(e_D, e_B) < e_c$ **then**
 if $p_B^{(m)} > p_D^{(m)}$ **then**
 return bright
 else
 return dark
 end if
 Abort detection
 end if
 Evaluate 7.10
 Perform action $\hat{a} \in \mathcal{A}$
 Update current state according to action
end for
if $p_B^{(n_{\text{bin, max}})} > p_D^{(n_{\text{bin, max}})}$ **then**
 return bright
else
 return dark
end if

It is crucial to note that up to this point we assumed no leakage mechanisms and a unit π -pulse fidelity. The influence of these unwanted effects are explored in chapter 9.

7.3 Information Gain Estimator

Proposed by Marinelli [17] we can try to maximize the expected entropy of the future normalized likelihoods of the bright or dark state respectively. Intuitively, we want that our actions lead us to a state where sampling (in this case receiving counts for the next sub-bin) offers the most amount of information. The objective function is therefore described by:

$$\mathcal{J} = H_{\text{bin}}\left(\tilde{p}_D^{(m+1)}\right) \quad (7.12)$$

with the normalized dark state likelihood $\tilde{p}_D^{(m+1)}$:

$$\tilde{p}_D^{(m+1)} = \frac{p_D^{(m+1)}}{p_D^{(m+1)} + p_B^{(m+1)}} \quad (7.13)$$

and the binary entropy function:

$$H_{\text{bin}}(p) = -(p \log(p) + (1-p) \log(1-p)) \quad (7.14)$$

As the future observation y_m is not known, we take the expectation value over all future possible outcomes. For the probability of an observation y_i we can use the law of total probability as follows:

$$p(y_i) = p(y_i | \theta_{\text{dark}}) p(\theta_{\text{dark}}) + p(y_i | \theta_{\text{bright}}) p(\theta_{\text{bright}}) \quad (7.15)$$

If we assume that we have a uniform prior probability we can simplify Eq. 7.15 and just note that $p(y_i) \propto p(y_i | \theta_{\text{dark}}) + p(y_i | \theta_{\text{bright}})$. We need to calculate this expectation value for both actions using Eq. 7.11 as update rule. Note that the number of π -pulses increases by one for the case when taking the corresponding action. For this objective function, it is not possible to evaluate the expectation numerically. We need to truncate the infinite series and are left with an approximation. As we are only interested in the sign of the difference, we are able to assume that the likelihoods $p_D^{(m)}$ (and therefore $p_B^{(m)}$) to be normalized. We can proceed to analyze the difference term 7.10 to find zeros. Firstly, if $p_D^{(m)}(p_B^{(m)}) = 0$ directly leads to $p_D^{(m+1)}(p_B^{(m+1)}) = 0 \forall y_m$. Together with the fact that $H_{\text{bin}}(0) = H_{\text{bin}}(1) = 0$ we can conclude that the difference term has zeros for $p_D^{(m)} \in 0, 1$. Using linearity of expectation we can write 7.10 for the case of not being in the flipped state as:

$$\begin{aligned} \mathbb{E}[\mathcal{J} | \theta, y_{i\dots n}, a = \text{no-}\pi] - \mathbb{E}[\mathcal{J} | \theta, y_{i\dots n}, a = \pi] &= \sum_{y=1}^{n_{\text{max}}} -p(y) \times \dots \\ \dots &\left(\frac{p_D^{(m)} p(y | \theta_{\text{dark}})}{p_D^{(m)} p(y | \theta_{\text{dark}}) + p_B^{(m)} p(y | \theta_{\text{bright}})} \log \left(\frac{p_D^{(m)} p(y | \theta_{\text{dark}})}{p_D^{(m)} p(y | \theta_{\text{dark}}) + p_B^{(m)} p(y | \theta_{\text{bright}})} \right) \right. \\ &+ \frac{p_B^{(m)} p(y | \theta_{\text{bright}})}{p_D^{(m)} p(y | \theta_{\text{dark}}) + p_B^{(m)} p(y | \theta_{\text{bright}})} \log \left(\frac{p_B^{(m)} p(y | \theta_{\text{bright}})}{p_D^{(m)} p(y | \theta_{\text{dark}}) + p_B^{(m)} p(y | \theta_{\text{bright}})} \right) \\ &- \frac{p_D^{(m)} p(y | \theta_{\text{bright}})}{p_D^{(m)} p(y | \theta_{\text{bright}}) + p_B^{(m)} p(y | \theta_{\text{dark}})} \log \left(\frac{p_D^{(m)} p(y | \theta_{\text{bright}})}{p_D^{(m)} p(y | \theta_{\text{bright}}) + p_B^{(m)} p(y | \theta_{\text{dark}})} \right) \\ &\left. - \frac{p_B^{(m)} p(y | \theta_{\text{dark}})}{p_D^{(m)} p(y | \theta_{\text{bright}}) + p_B^{(m)} p(y | \theta_{\text{dark}})} \log \left(\frac{p_B^{(m)} p(y | \theta_{\text{dark}})}{p_D^{(m)} p(y | \theta_{\text{bright}}) + p_B^{(m)} p(y | \theta_{\text{dark}})} \right) \right) \end{aligned} \quad (7.16)$$

We can see that for the case $p_D^{(m)} = 0.5 \Rightarrow p_B^{(m)} = 0.5$ we have cancellation of the terms for each individual y which leads to another zero. For the case of being in the flipped state we need to make the substitution $\theta_{\text{dark}} \leftrightarrow \theta_{\text{bright}}$. This will simply lead to an overall sign change.

7.3.1 Example

We will proceed with an specific example to show that there are potentially more zeros than the trivial ones discussed up to this point. The conversion from the purely theoretical notation to the actual parameters are summarized in the table below:

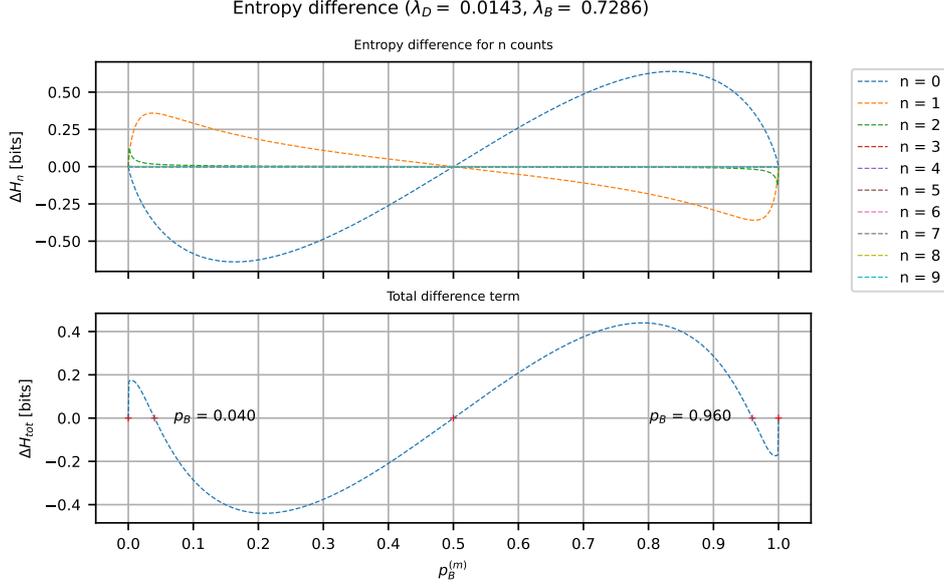


Figure 7.2: Top: Difference term for individual $n = 0, \dots, 9$ weighted by $p(y = n)$ as in 7.16. The contributions for $n > 2$ are negligible for the overall result. Therefore, they could be truncated as well. Bottom: Approximation of the expectation. We can see that there are two additional zeros introduced.

Generic notation	Specific setup
$y_{0\dots n-1}$	Photon counts n using PMT as detection device
$p(y \theta_{\text{bright}})$	Poisson distribution with mean λ_{bright}
$p(y \theta_{\text{dark}})$	Poisson distribution with mean λ_{dark}
Parameters	Value
$t_{\text{detection}}$ [μs]	200
# sub-bins	25
λ_{dark}	$\frac{0.1}{200 \mu\text{s}} \frac{1}{25} t_{\text{detection}} = 0.0143$
λ_{bright}	$\frac{25}{200 \mu\text{s}} \frac{1}{25} t_{\text{detection}} = 0.7286$
Method of approximation of expectation	Truncating 7.16 at maximum photon count $n_{\text{max}} = 9$

Table 7.1: Specific example using reasonable parameters found in experimental setups.

We make use of the fact that the probabilities $p_B^{(m)}$ & $p_D^{(m)}$ are normalized to write $p_D^{(m)} = 1 - p_B^{(m)}$ which leaves us with only one parameter. This enables us to pre-calculate the decision boundaries which takes away the need to recalculate the expectation after each sub-bin. Figure 7.2 shows the contribution of each term in the sum of 7.16.

This shows an interesting quirk of the entropy gain objective. Even though we are confident that we are in the bright state, right before crossing the threshold we will still apply a π -pulse and switch to the dark state. As the sign of the difference changes in the next sub-bin, it is likely that we then stay in the dark state for the rest of the detection sequence. This will potentially prolong the average detection time.

A possible explanation of this behavior is the nature of the objective function itself. Mathematically, it maximizes the expected acquired information which does not immediately imply that the main goal of

reducing the average the detection time is captured. It is also noted that this entropic quantity might be sensitive to the future time horizon. The method proposed in this thesis only looks ahead one sub-bin. It is possible that if the horizon is increased to more than one future sub-bin, the additionally introduced zeros would vanish. This would effectively cancel the need for a π -pulse.

7.4 Minimum Bins Estimator

As our main goal is to minimize the average detection time it only feels natural to use an objective function which explicitly tries to accomplish this task. We will therefore define:

$$\mathcal{J} = - \# \text{ sub-bins left} \quad (7.17)$$

where the minus sign turns in the maximization of the sub-bins left into a minimization. We denote the probability of being in the dark (bright) state after observation y_m as p_D^{curr} (p_B^{curr}). These probabilities depend on the number of π -pulses we have applied and can be found as follows:

$$p_D^{\text{curr}} = \begin{cases} p_D^{(m)}, & \# \pi\text{-pulses even} \\ p_B^{(m)}, & \text{otherwise} \end{cases} \quad (7.18)$$

As a first assumption, we are split the expectation in 7.9 into the two contributing terms:

$$\begin{aligned} \mathbb{E} [\# \text{ sub-bins left} \mid \theta, y_{i\dots n}, a] = \\ p_D^{\text{curr}} \mathbb{E} [\# \text{ sub-bins left} \mid \theta, y_{i\dots m}, \text{currently dark}] \\ + p_B^{\text{curr}} \mathbb{E} [\# \text{ sub-bins left} \mid \theta, y_{i\dots m}, \text{currently bright}] \end{aligned} \quad (7.19)$$

As a further simplification, we will assume that the next π -pulse is the final one. Therefore, we stay in the state after the action has been applied for the rest of detection sequence. Applying a π -pulse will simply interchange the expectations in Eq 7.19. This allows us to write the difference term in Eq. 7.10 as:

$$\mathbb{E} [\mathcal{J} \mid \theta, y_{i\dots n}, a = \text{no-}\pi] - \mathbb{E} [\mathcal{J} \mid \theta, y_{i\dots n}, a = \pi] = -(p_D^{\text{curr}} - p_B^{\text{curr}}) (E_{\text{dark}} - E_{\text{bright}}) \quad (7.20)$$

where we defined:

$$\begin{aligned} E_{\text{dark}} &= \mathbb{E} [\# \text{ sub-bins left} \mid \theta, y_{i\dots m}, \text{currently dark}] \\ E_{\text{bright}} &= \mathbb{E} [\# \text{ sub-bins left} \mid \theta, y_{i\dots m}, \text{currently bright}] \end{aligned} \quad (7.21)$$

Note that the negative sign is again an artefact of the maximization. Using the same reasoning when analyzing the entropy gain estimator, we assume normalized p_D^{curr} and p_B^{curr} and rewrite Eq 7.20 as:

$$-(1 - 2p_B^{\text{curr}}) (E_{\text{dark}} - E_{\text{bright}}) \quad (7.22)$$

The final simplification we can make is to assume no matter what previous observations have been made, we have $E_{\text{dark}} > E_{\text{bright}}$. This is not unreasonable as the premise of investigating these methods was to speed up dark state detection. The assumption that we apply a final π -pulse can be interpreted as performing a standard adaptive maximum likelihood method with non-uniform prior, namely the current likelihoods. In section , we have shown theoretically that on average and under ideal conditions, dark state detection takes longer. Performing the adaptive maximum likelihood method in real experiments confirm the theoretical findings [22]. In this particular scenario with comparable error threshold e_c , the average bright state detection time was found to be 72 μ s while the average dark state detection time was 219 μ s. Provided all assumptions are reasonable, we arrive at a simple final decision rule:

$$\hat{a} = \begin{cases} \text{no-}\pi, & p_B^{\text{curr}} \geq \frac{1}{2} \\ \pi, & \text{otherwise} \end{cases} \quad (7.23)$$

We note that this is only holds true if the dark and bright state distributions do not exhibit a considerable overlap. For pure Poissonian distributions, this translates to $R_D \ll R_B$. This also excludes scenarios when detecting multiple ions using a detection device without spatial resolution.

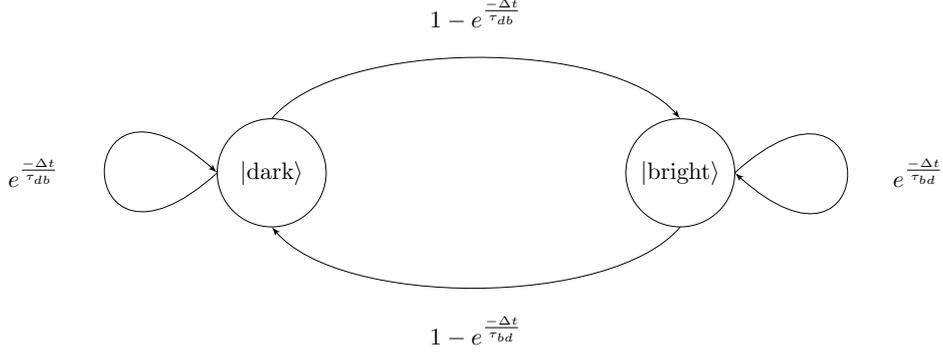


Figure 7.3: Transition diagram showing the conditional probabilities of moving to a next state given the current. Pumping events are modelled using rate equation, giving as the corresponding probabilities.

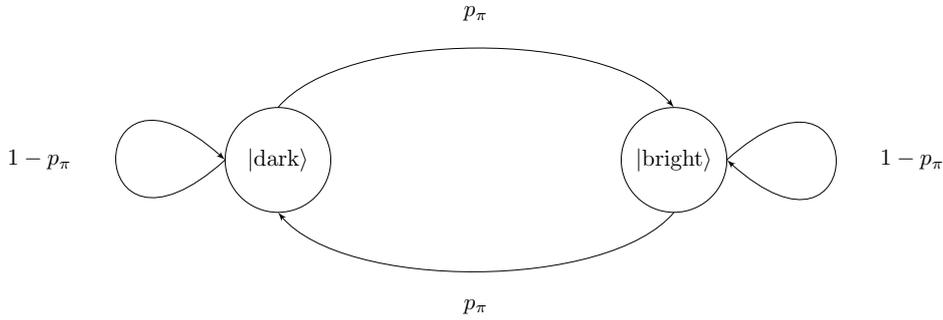


Figure 7.4: Transition diagram showing the conditional probabilities of moving to a next state given the current. We can model π -pulse fidelity with a success rate leading to a bit flip.

7.5 Hidden Markov Model MAP Decoding

Up to this point we did not take any dark or bright state pumping or π -pulse infidelities into account. Hidden Markov models are a valid representation of state discrimination algorithms which can easily incorporate both effects. Hidden Markov models were employed to further develop the adaptive maximum likelihood algorithm and the π -detection method [19]. Moreover, with not too much computational overhead, we can also infer an maximum a posteriori estimate. To simplify the model, we assume that an ion does only decay or gets pumped once a detection sub-window is over. This assumption is justified if leakage rates $\frac{1}{\tau}$ satisfy $\frac{t_{\text{bin}}}{\tau} \ll 1$ for both bright-to-dark leakage rates $\frac{1}{\tau_{bd}}$ and dark-to-bright leakage rates $\frac{1}{\tau_{db}}$. A π -pulse would then be applied immediately after potential leakage. We model both effects as transitions which can be represented graphically as shown in Fig. 7.3.

We will further define the corresponding state transition matrix:

$$T_{\text{no-}\pi} = \begin{pmatrix} e^{-\frac{t_{\text{bin}}}{\tau_{bd}}} & 1 - e^{-\frac{t_{\text{bin}}}{\tau_{db}}} \\ 1 - e^{-\frac{t_{\text{bin}}}{\tau_{bd}}} & e^{-\frac{t_{\text{bin}}}{\tau_{db}}} \end{pmatrix} \quad (7.24)$$

where we represent messages as column vectors. The first entry holds the value associated with the bright state whereas the second entry is dedicated to the dark state. The state transition diagram for applying a π -pulse is shown in Fig 7.4.

Using our assumptions, the transition matrix is defined as:

$$T_\pi = \begin{pmatrix} 1 - p_\pi & p_\pi \\ p_\pi & 1 - p_\pi \end{pmatrix} T_{\text{no-}\pi} \quad (7.25)$$

We therefore only need to keep track of before which sub-bin we have applied a π -pulse. We proceed to define the forward and backward messages for this hidden Markov model. We start of with the forward messages by defining:

$$\vec{\mu}_{X_0} = \begin{pmatrix} p(\theta_{\text{bright}}) \\ p(\theta_{\text{dark}}) \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} \quad (7.26)$$

The recursion for future forward messages is given by:

$$\vec{\mu}_{X_n} = T_i^{(n-1)} \begin{pmatrix} \vec{\mu}_{X_{n-1}}(\text{bright}) \\ \vec{\mu}_{X_{n-1}}(\text{dark}) \end{pmatrix} \odot \begin{pmatrix} p(y_{n-1} | \theta_{\text{bright}}) \\ p(y_{n-1} | \theta_{\text{dark}}) \end{pmatrix}, \quad i \in \{\pi, \text{no-}\pi\} \quad (7.27)$$

For the backward messages starting with the last observation y_{n-1} we have:

$$\bar{\mu}_{X_{n-1}} = \begin{pmatrix} p(y_{n-1} | \theta_{\text{bright}}) \\ p(y_{n-1} | \theta_{\text{dark}}) \end{pmatrix} \quad (7.28)$$

Note that the message $\bar{\mu}_{X_n}$ is the neutral message so we skipped it in the recursion definition. The backward recursion can be computed as follows:

$$\bar{\mu}_{X_{n-1}} = T_i^{T(n-1)} \begin{pmatrix} \bar{\mu}_{X_n}(\text{bright}) \\ \bar{\mu}_{X_n}(\text{dark}) \end{pmatrix} \odot \begin{pmatrix} p(y_{n-1} | \theta_{\text{bright}}) \\ p(y_{n-1} | \theta_{\text{dark}}) \end{pmatrix}, \quad i \in \{\pi, \text{no-}\pi\} \quad (7.29)$$

As the forward messages do not give us the likelihood at the end of each bin but at the start, we are not able to use them for decision making. Furthermore, we cannot calculate the forward messages before knowing our action. We need to modify the recursion such that we get the current state likelihoods at the beginning of each sub-bin:

$$\vec{\mu}_{X_0} = \begin{pmatrix} p(\theta_{\text{bright}}) \\ p(\theta_{\text{dark}}) \end{pmatrix} = \begin{pmatrix} \frac{1}{2} \\ \frac{1}{2} \end{pmatrix} \quad (7.30)$$

$$\vec{\mu}_{X_n} = T_{\text{no-}\pi} \left(\begin{pmatrix} p(y_n | \theta_{\text{bright}}) \\ p(y_n | \theta_{\text{dark}}) \end{pmatrix} \odot \begin{pmatrix} 1 - p_\pi & p_\pi \\ p_\pi & 1 - p_\pi \end{pmatrix} \mathbb{1}_{[\pi\text{-pulse sub-bin } n-1]} \begin{pmatrix} \vec{\mu}_{X_{n-1}}(\text{bright}) \\ \vec{\mu}_{X_{n-1}}(\text{dark}) \end{pmatrix} \right) \quad (7.31)$$

With these messages at hand, we are able to modify the minimum bins estimator to produce a MAP estimate. In contrast to the previously introduced methods, we are not able to make use of previous calculations and need to compute the backward messages from scratch after each sub-bin. Moreover, the number of computations needed increases linearly with the sub-bin index. It may not be feasible to perform the computation in real time. However, this modification allows us to benchmark simpler algorithms against this optimal¹ estimator. When finish writing this thesis, the concept to use π -pulses not only to increase fidelity but also reduce average detection times has been explored in [23]. State detection was described as a partially observable Markov decision process (POMDP). Possible solutions to solve POMDPs employ machine learning techniques. The notation used in section 7.2 is mostly in accordance with notation used in the field of reinforcement learning.

¹Optimal when modeling state identification as a hidden Markov model.

Chapter 8

Simulation Environment

In this chapter we will discuss the simulation environment used to benchmark the estimators described in chapter 7.

8.1 General Structure

For this thesis, a Python package has been developed to streamline estimator bench-marking. It uses an object-oriented approach where the user can instantiate ions and estimators, pair them freely, and pass the them to an environment. The environment will repeat the detection sequence for a specified number of times with user defined parameters. In the following, we will quickly introduce each of the three mentioned base classes.

8.2 Ion Class

The Ion class encapsulates the photon emission statistics of a single ion during readout. To account for non-Poissonian statistics, we use Markov Chain Monte Carlo methods. Leakage mechanisms can be modelled using rates which will naturally lead to transition diagrams as in Fig. 7.3. While we have assumed that the ion will only experience leakage mechanisms at the end of each sub-bin when describing estimators, we will further divide sub-bins for a more fine-grained simulation. The photon statistics are calculated using a discretized version of Eq. 6.8:

$$\lambda_{\text{sub-bin}} = \sum_{n=1}^{\frac{t_{\text{sub-bin}}}{dt}} R(t_0 + n dt) dt \quad (8.1)$$

At the end of each sub-bin we sample from a Poissonian distribution to get the raw photon count. Depending on the detection device, these statistics may need to be altered.

8.3 Estimator Class

Using the photon statistics provided by the paired ions, the estimator will go through either Alg. 2 or Alg. 3, depending in the estimator at hand. To allow for detailed analysis, the estimator will store statistics for post-analysis. Furthermore, it is possible to also log the trajectory of the detection sequence where we store a snapshot of estimator parameters for each sub-bin. Statistics are a summary of the readout process. This most importantly includes total received photon counts, number of sub-bins it took to terminate detection, and for π -pulse methods also the number of applied pulses. In contrast, trajectory data will for example store the number of photons during each sub-bin or if an π -pulse was requested. This allows for even more detailed analysis.

8.4 Environment Class

The Environment class coordinates the simulation. It will distribute the photon counts according to the pairing and will execute the π -pulses after each estimator update. The environment is able to store the

statistics and trajectory data of each repetition into a comma separated values or JSON file respectively for long-term storage of simulation results.

8.4.1 Data Structure

We briefly want to introduce the data structure of the files created and stored by the simulation environment. The .csv file containing the statistics has a simple database structure. Each row describes one readout sequence, so the data frame has $n_{\text{repetition}}$ rows. Each column defines a stat. Each column name has a numerical suffix to identify to which estimator it corresponds to. The indexing starts at zero. The order of the estimators is extracted from the pairing list provided to the environment. If an estimator is passed more than once in this list ¹ it will be skipped and only unique estimator entries are indexed and stored. The JSON file structure is more complex as it does not have fixed data types and array lengths. The structure of the nested dictionaries is shown in Fig. 8.1 for a π -pulse method estimator. More simple estimators such as the adaptive maximum likelihood estimator will lack some field entries. To be more specific, all entries concerning π -pulse statistics are not present.

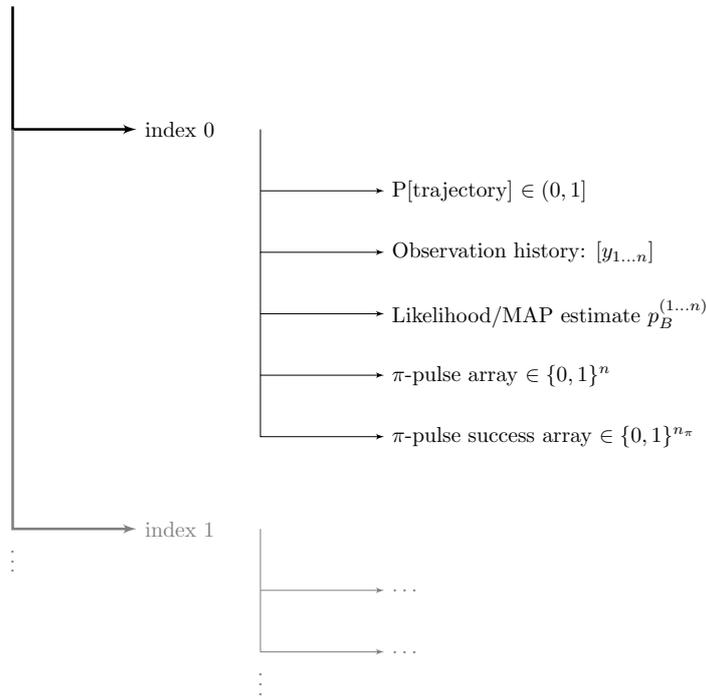


Figure 8.1: JSON file structure showing the hierarchy of the nested dictionaries and their contents.

We proceed to briefly explain the content of each field.

- **index i :** Numerical identifier of the trajectory.
- **P[trajectory]:** Probability of the trajectory occurring during simulation. This is calculated by counting how many times this trajectory was sampled divided by the number of repetitions.
- **Observation history:** Array of the photon counts during sub-bin m .
- **Likelihood:** Array of the normalized bright state likelihood (or MAP estimate) at the end of sub-bin m .
- **π -pulse array:** Boolean array indicating if a π -pulse was requested after sub-bin m .
- **π -pulse success array:** Boolean array indicating the success of requested pulses, ordered chronologically.

¹Meaning it was paired with multiple ions and therefore receives photons from all paired ions.

8.5 User Example Code

We conclude this chapter by providing a short code excerpt which sets up the simulation environment. We initiate the base class `ion` which has no leakage mechanisms. The state discrimination algorithm to be bench-marked is the entropy gain estimator. The attribute `n_counts_max` determines where the infinite series 7.16 is truncated. We both save the summary and the more detailed trajectory data at the end of the simulation.

```
from PyonDetect.ion.ion import Ion
from PyonDetect.estimator.estimator import EntropyGainEstimator
from PyonDetect.environment.environment import SimulationEnvironment

# initialize simulation environment
R_D = 0.1 / 200 # per us
R_B = 25 / 200 #per us
initial_state = 0 # bright
n_subbins = 25
t_detection = 200 # us

ion = Ion(initial_state, R_dark=R_D, R_bright=R_B)
estimator = EntropyGainEstimator(R_D=R_D,
                                 R_B=R_B, e_c=1e-4,
                                 n_subbins_max=n_subbins,
                                 n_counts_max=10,
                                 save_trajectory=True)

env = SimulationEnvironment([ion],
                           [estimator],
                           detection_time=t_detection,
                           n_subbins=n_subbins,
                           n_repetition=10000)

# let simulation run
env.run()

# save data to files
env.save_to_csv()
env.save_to_json()
```

With that we end our brief introduction of the bench-marking framework and continue to discuss the performance of the developed estimators.

Chapter 9

Estimator Benchmarking

In this chapter we will make use of the simulation environment to assess the performance of the derived estimators. In the beginning we assume ideal circumstances, steadily increasing performance-decreasing effects.

9.1 General Setup

For all of the following tests, it was assumed that ions are detected using PMTs. We note that the theory also allows for differing detection devices. The main change that needs to be made is to adapt the observation likelihood function in the simulation of the ion and estimator. If not otherwise stated, a full detection window is $t_{\text{det, max}} = 200 \mu\text{s}$ and the error threshold is set to $e_c = 10^{-4}$. Every simulation will run for an ion initially in the dark and bright state an equal number of times. By default we will analyze three figures of merit:

1. **Average detection time:** Reducing the average detection time was the main goal of our analysis. Apart from the average detection time, we may also be interested in the distribution of the detection time. It is likely that the distribution for a given number of divisions into sub-bins is not symmetric and will exhibit a tail towards higher detection times. We note that we have excluded the π -pulse duration for estimators which make use of them. The rationale behind this decision is the fact that even for a high number of sub-bins, the sub-bin time is still an order of magnitude higher than the pulse duration.
2. **Dark/bright state detection accuracy:** Synonymous with the readout fidelity - the ratio of miss-identified states to total read out states. The speedup in detection should not come at the cost of a degrading fidelity. We will analyze accuracy separately, depending on the initial state. Owing to the fact that we prepare as many ions in the bright and dark state, the average accuracy will simply be the arithmetic mean of both error figures.
3. **Average number of π -pulses applied:** Using a large number of π -pulses comes with two issues: The time it takes to perform said pulse and the introduced possibility for failure. It seems likely that a failed π -pulse will lead to a false identification. Ideally, the estimator will only apply a π -pulse if the ion starts in the dark state.

Lastly, we note that for no leakage mechanisms and unit π -pulse fidelity, hidden Markov models and belief propagation methods do not differ from the π -pulse methods described in sections 7.3 and 7.4.

9.2 Effect Of Non-zero Background

For this section, we assume no leakage mechanisms and unit π -pulse fidelity. Included in this scenario is the ideal case of zero background photon counts. We proceed to analyze the effects of increasing background. The photon rates indicated are always referenced to a detection window of $200 \mu\text{s}$. We assumed that the bright state count rate to be $R_B = 25 \frac{1}{200 \mu\text{s}}$ for the rest of this section. We assess the performance of the Entropy Gain estimator and the Minimum Bins estimator against the standard adaptive maximum likelihood method.

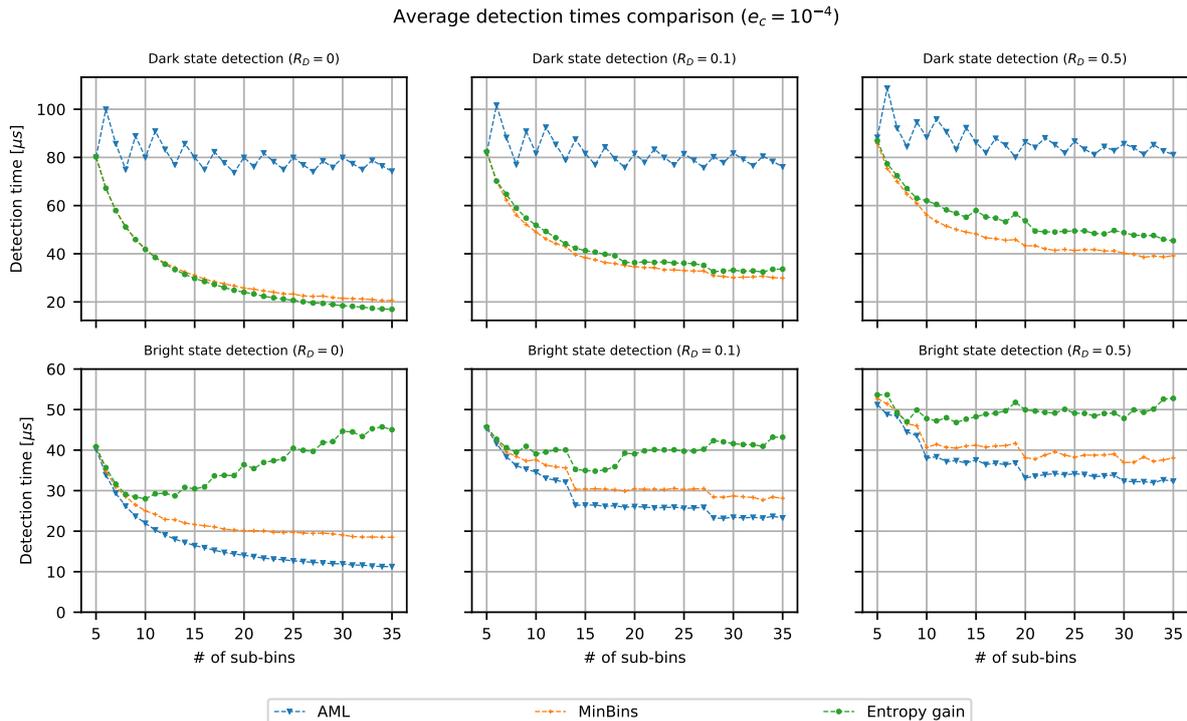


Figure 9.1: Average detection times against number of sub-bins introduced into the $200 \mu\text{s}$ detection window. Number of repetitions for each panel $n = 10000$. **Top row:** Each panel shows the average detection times for dark state readout. Starting at the most left with zero background, we increase dark state photon counts moving right. **Bottom row:** Same as the top row but for bright state readout.

Fig. 9.1 compares the average detection times of the competing estimators for increasing background photon count. The ideal case corresponds to the first column. Note the different time scales for bright and dark state detection respectively. We observe a major speedup in dark state detection for both the entropy gain and minimum bins estimator compared to the standard adaptive maximum likelihood method. The average detection time starts to saturate after approximately 30 subdivisions. Looking at the bright state detection we observe that the standard adaptive maximum likelihood method outperforms both π -pulse methods. The minimum bins estimator for instance lags approximately one sub-bin time behind. The entropy gain estimator starts to exhibit the behaviour discussed in section 7.3.1. We predict that this is due to fact that the entropy gain estimator may request π -pulses even for high bright state likelihoods as discussed in section 7.3, reducing the chance of fast termination of the detection sequence. The analysis shown in section 9.2.1 confirms this hypothesis. For the ideal number of sub-bins, average dark and bright state detection are equal. For the minimum bins estimator, this would correspond about 25 sub-bins resulting into an average detection time of $20 \mu\text{s}$. For increasing background counts we see an overall degradation of performance, affecting all estimators. We proceed by inspecting the accuracy of each estimator.

Fig. 9.2 shows the absolute number of errors depending on the number of sub-bins. We first observe that in general, bright state detection is more error-prone than dark state detection. The error rate for any case is low and is with high certainty below 0.05%. Even with increasing dark count the error rate does not increase dramatically. We also note that when performing MCMC simulations to infer probability distributions we generally overestimate the probability of rare outcomes. With that in mind, this result indicates that the error rate does not strongly depend on the background count. We will conclude the accuracy discussion at this point as leakage mechanisms impose a far greater challenge to the accuracy of the estimators.

Lastly, we analyze the number of π -pulses the estimators applied on average. Fig. 9.3 uses the same structure as the comparison of the average detection time. We first note that ideally, the estimators would not apply a π -pulse when starting in the bright state and only one when starting in the dark state. For the case of no background, the entropy gain estimator closely follows the aforementioned ideal case. We only see a slight deviation for bright state detection. With increasing number of sub-bins, the probability

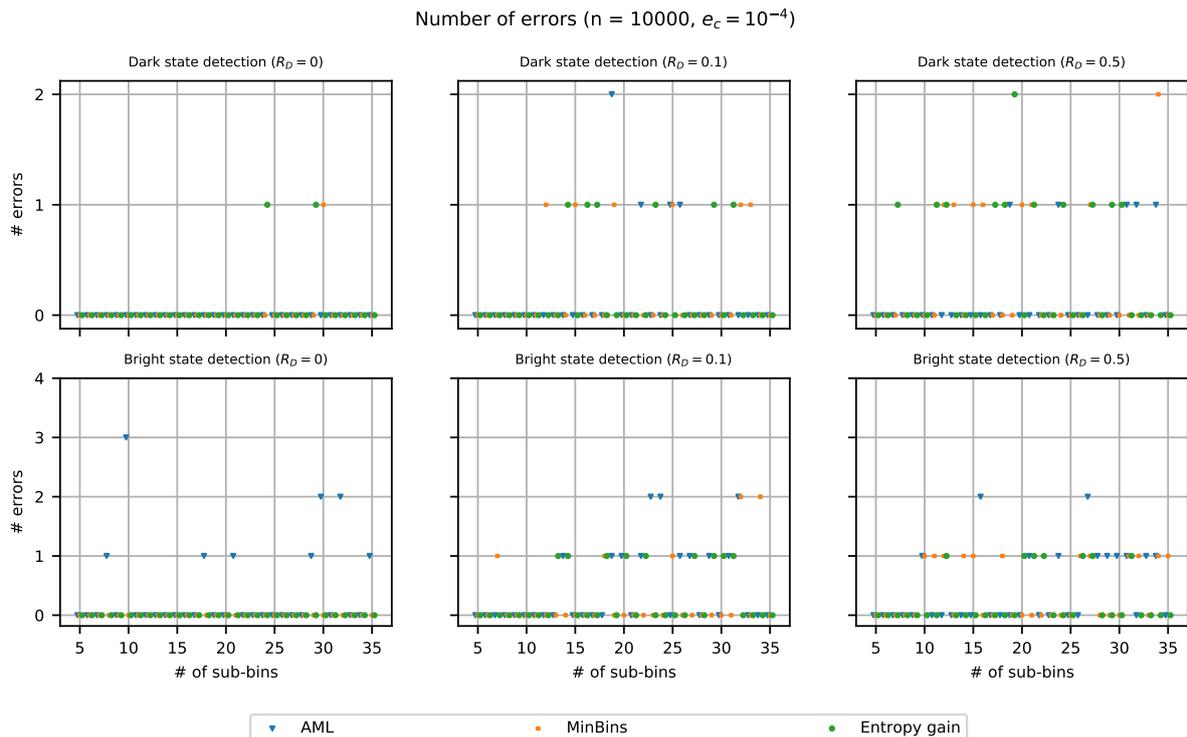


Figure 9.2: Number of errors against number of sub-bins introduced into the $200 \mu\text{s}$ detection window. Number of repetitions for each panel $n = 10000$. **Top row:** Each panel shows the number of miss-classifications for dark state readout. Starting at the left with zero background, we increase dark state photon counts moving right. **Bottom row:** Same as the top row but for bright state readout.

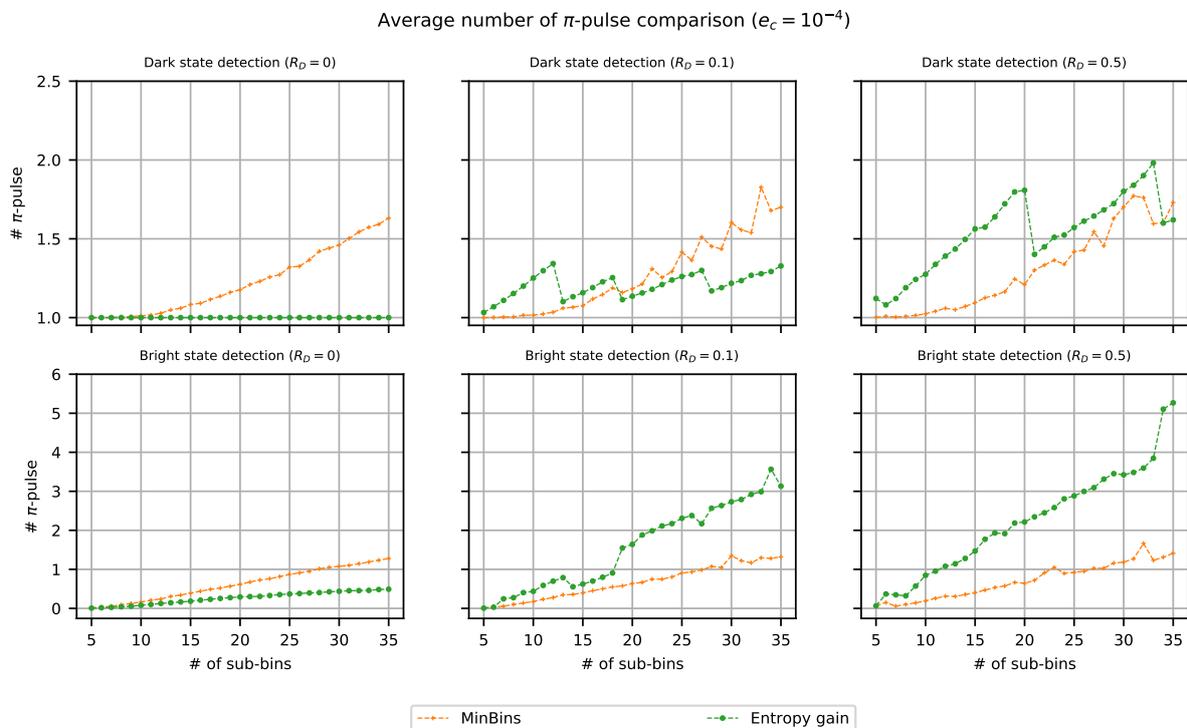


Figure 9.3: Number of applied π -pulses against number of sub-bins introduced into the $200 \mu\text{s}$ detection window. Number of repetitions for each panel $n = 10000$. **Top row:** Each panel shows the number applied pulses for dark state readout. Starting at the left with zero background, we increase dark state photon counts moving right. **Bottom row:** Same as the top row but for bright state readout.

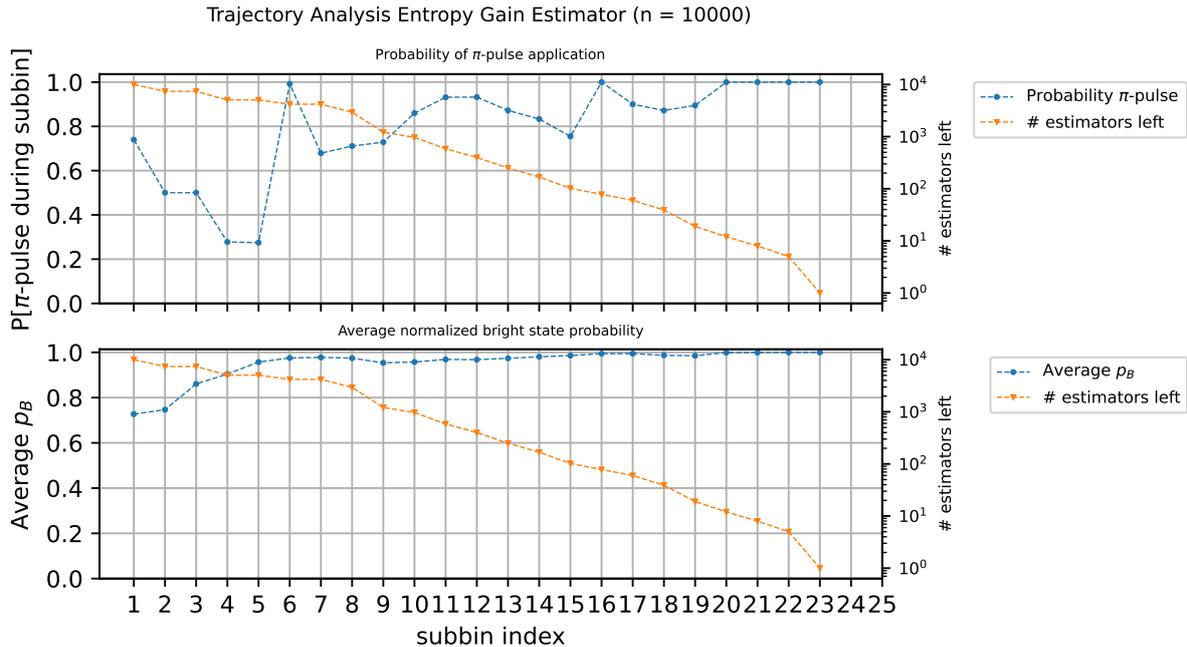


Figure 9.4: The orange trace always indicates the number of remaining estimators during sub-bin index. **Top row:** The fraction of remaining estimators requesting a π -pulse. **Bottom row:** Average normalized bright state likelihood $p_B^{(m)}$ during sub-bin with index m .

to observe zero photons also increases. Therefore it is likely that the entropy gain estimator confuses some rare low-photon-count events when in the bright state with being in the dark state. Clearly, the entropy gain estimator outperforms the minimum bins estimator if no background counts are present. For increased background counts however, the entropy gain estimator starts to apply more π -pulses. The minimum bins estimator in contrast does not increase the frequency of π -pulses.

9.2.1 Entropy Gain Trajectory Analysis

We will demonstrate the trajectory logging capabilities of the simulation environment by analyzing the behavior of the entropy gain estimator. We restrict ourselves for the case of dividing the detection window into 25 sub-bins and only consider bright state detection. Using the trajectory data, we can use more fine-grained statistics to confirm our hypothesis stated earlier namely that the entropy gain estimator will apply π -pulses for bright state likelihoods close to pass the threshold. For this run we start with a population of 10000 entropy gain estimators. Apart from monitoring the estimators still actively detecting, we also keep of track of the average bright state probability among the remaining as well as the fraction of estimators which requested a π -pulse after each sub-bin. The resulting statistics are shown in Fig 9.4.

One would expect that if the normalized bright state likelihood is close to one, no π -pulse application would be necessary. As already theoretically discussed in sub section 7.3.1, the entropy gain estimator will request pulses after passing a certain threshold defined by the total entropy difference term. If this pulse threshold is reached before the detection error threshold e_c is met for the errors e_B or e_D , the pulse will be applied, possibly prolonging the detection sequence.

9.3 Influence of π -pulse Infidelity

Up to this point we have assumed that π -pulses result in a perfect population inversion. To judge the performance of π -pulse method estimators, it is necessary to take this infidelity into account. For this purpose, we again fix the number of sub-bins to 20 while keeping the detection time at 200 μ s. We only consider the case for no background and low background photon counts. We perform a linear scan of increasing infidelity starting at perfect inversion going up to an infidelity of 0.5 %. We analyze the effect on the number of π -pulses applied and the resulting fidelity.

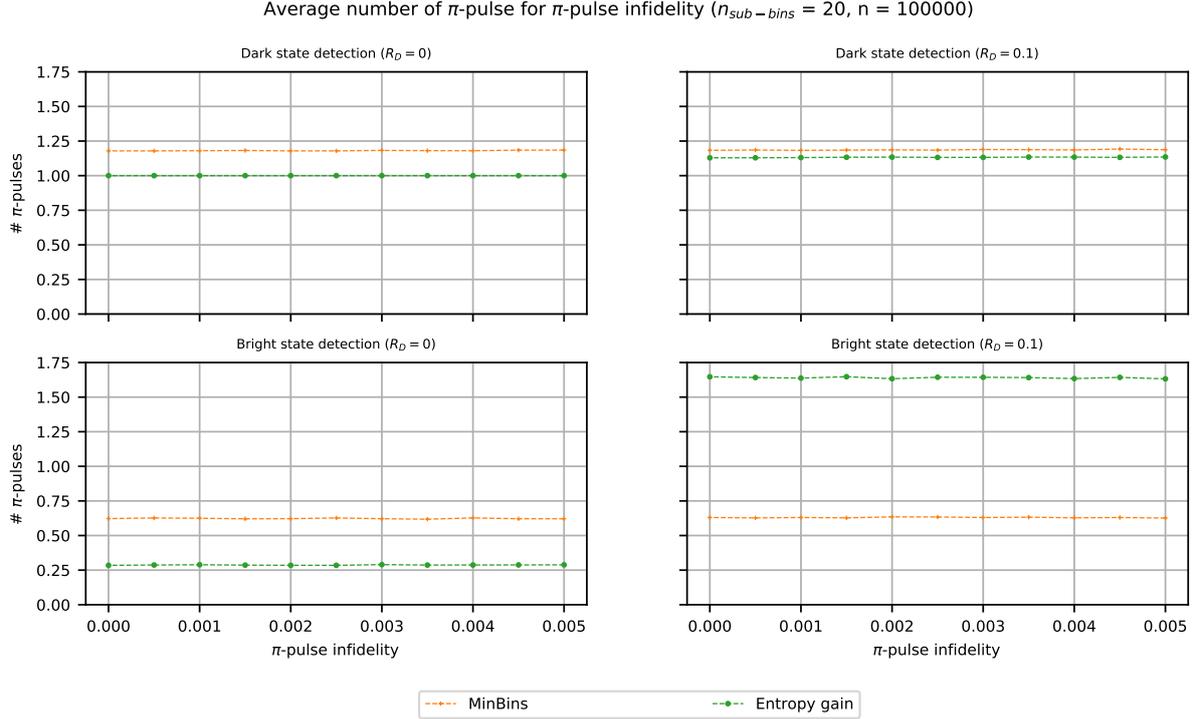


Figure 9.5: Average number of π -pulses applied during a detection sequence for a given infidelity. **Top row:** Dark state detection with zero background counts (left) and low background (right). **Bottom row:** Same as top but for bright state detection

Fig. 9.5 shows the average number of applied π -pulses for increasing infidelity. There is no apparent effect on the number of pulse requests. We focus our attention to the state readout errors introduced by failed population inversion as shown in Fig. 9.6.

In the case background-free detection, we see that the error in population inversion bleeds through to the readout error. It is therefore safe to assume that a single failed pulse will lead to an erroneous detection as the estimator is not able to recover. The effect is amplified if we apply more than one π -pulse.

9.4 Effect of Leakage Mechanisms

We end off your discussion by considering the same scenario as in the discussion of the effect of increasing background counts. We further introduce bright-to-dark as well as dark-to-bright leakage mechanisms. The experimental parameters are summarized in Tab. 9.1.

Parameter	Value
$R_D \left[\frac{1}{200\mu\text{s}} \right]$	0.1
$R_B \left[\frac{1}{200\mu\text{s}} \right]$	26.34
$\tau_{db} [\mu\text{s}]$	12700
$\tau_{bd} [\mu\text{s}]$	2915
e_c	10^{-4}
$n_{\text{repetition}}$	20000

Table 9.1: Parameters used in the scenario for estimator benchmarking in section 9.4

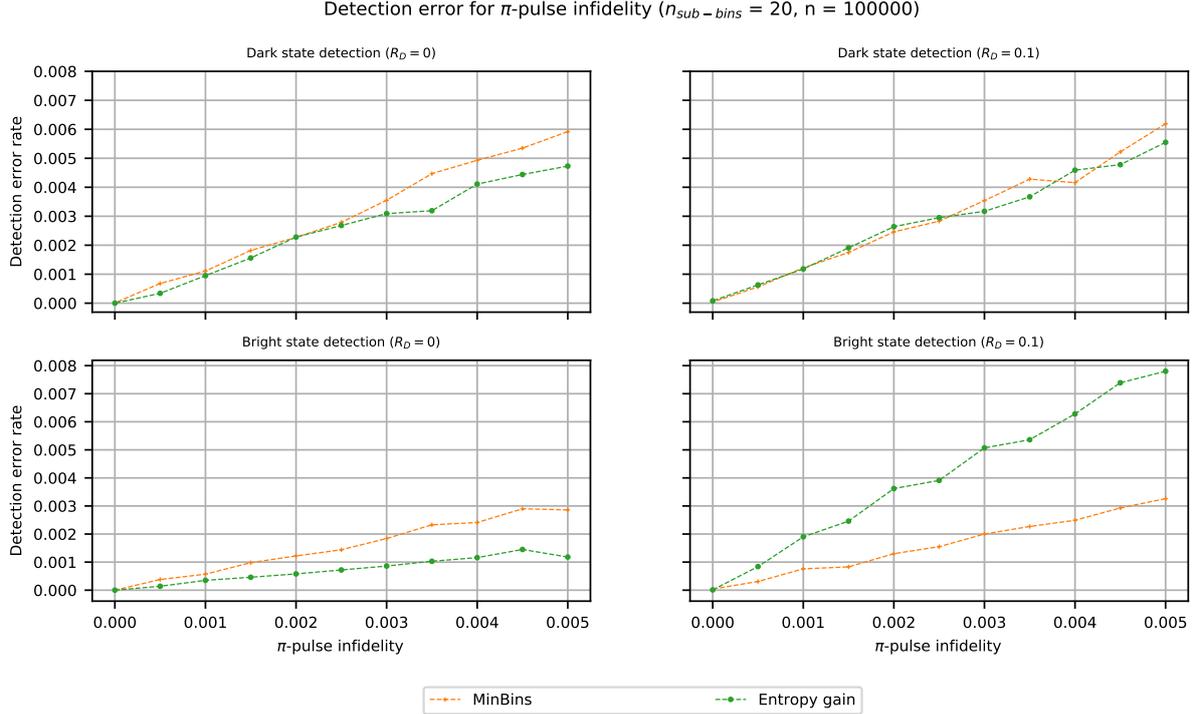


Figure 9.6: readout errors during a detection sequence for a given population inversion error. **Top row:** Dark state detection with zero background counts (left) and low background (right). **Bottom row:** Same as top but for bright state detection.

The numbers in Tab. 9.1 are taken from [17] for beryllium detection. It imposes more of a challenge than calcium as leakage mechanisms are negligible in this case. We will also only consider π -pulse method estimators. In addition to the entropy gain and minimum bins estimator we also introduce the modified minimum bins estimator using belief propagation to infer a MAP estimate of the initial state and the current state before π -pulse application. It will be referred to as the MAP estimator in the following. A slight modification to the MAP estimator is needed to decrease the average readout time. Due to the mixing of the bright and dark states, the a posteriori probabilities will not reach the confidence needed to satisfy the threshold $e_c = 10^{-4}$. We start with a higher threshold of $e_{c, \text{MAP}} = 10^{-3}$ and will multiply it by a factor after each sub-bin. The factor has been chosen such that at the last sub-bin, the threshold satisfies $e_{c, \text{MAP}} = 0.5$ - equivalent to not having to simply guess a state. We go through the three categories analogously as in section 9.2. The average detection times, even in the case of leakage mechanisms, shows a similar behavior to that studied in previous sections. Generally speaking, the detection times have only been slightly increased. The threshold tuning of the MAP estimators results in outperforming both standard estimators. This is not surprising as the unmodified minimum bins and entropy gain estimators use pure Poissonian statistics as the underlying model. Leakage mechanisms are therefore not considered. A leakage event will lead to miss-identification in a similar fashion as the non-ideal population inversion. We expect to see the effect of the missing dynamics in the entropy gain and standard minimum bins estimator most pronounced in the readout fidelity.

In Fig. 9.8 we can see a high increase in readout errors over all test cases. For background-free dark detection we see a significant decrease in readout errors for increasing sub-bin number. This is due to the fact the the average detection time decreases with increasing sub-bin numbers. We are therefore less likely to experience a leakage event if the detection time is kept shorter. The decrease levels off as soon as we reach the performance of background-free bright state detection. We note that bright state detection fidelity deteriorates for shorter sub-bin times. As soon as we introduce background counts, the performance worsens. The MAP estimator has the lowest error rate. This benefit comes with the increased computational burden. Lastly, we analyze the effect of leakage mechanisms on the average number of π -pulses.

A disadvantage of the MAP estimator becomes apparent when regarding the frequency of π -pulse application for dark state detection as seen in Fig. 9.9. The number of applied pulses is at least as high

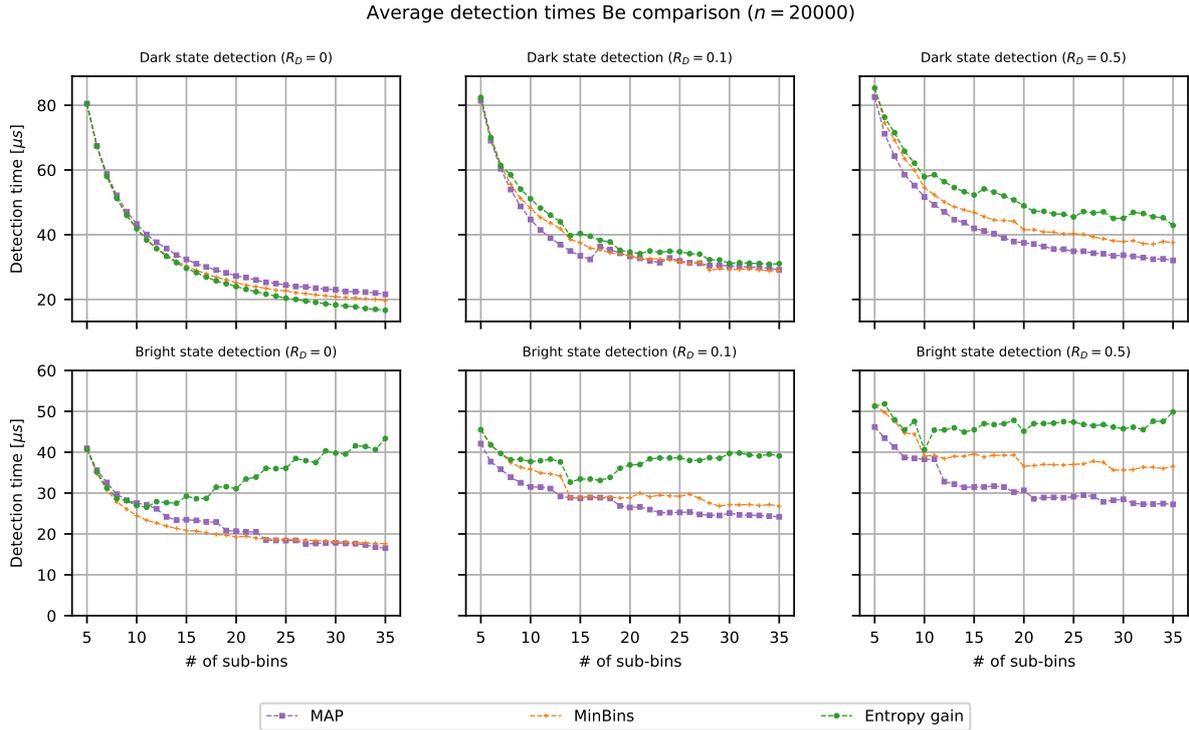


Figure 9.7: Average detection time versus number of sub-bins. Experimental parameters observed in Beryllium state readout have been used to compare the three methods in this scenario. **Top row:** Dark state detection for increasing background photon counts (increasing from left to right). **Bottom row:** Same as top for bright state detection.

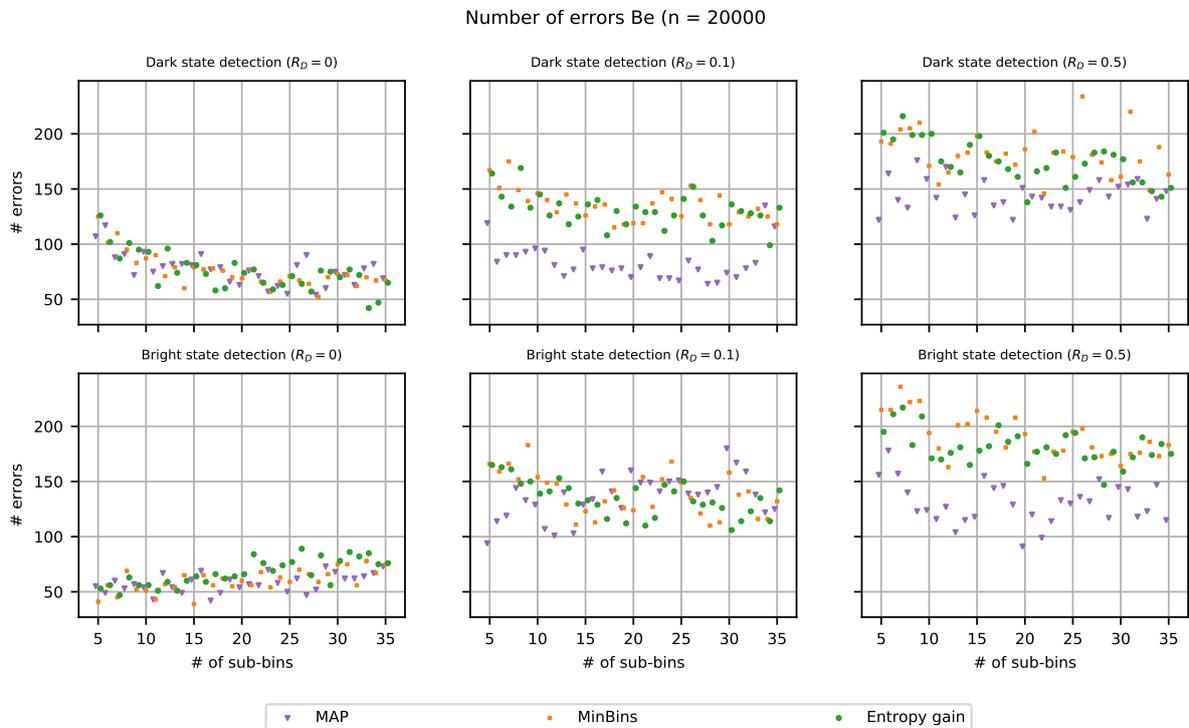


Figure 9.8: Readout errors for a given division into sub-bins. **Top row:** Dark state detection starting with zero background on the left, increasing moving right. **Bottom row:** Same as top but for bright state detection.

Average number of π -pulse Be comparison (n = 20000)

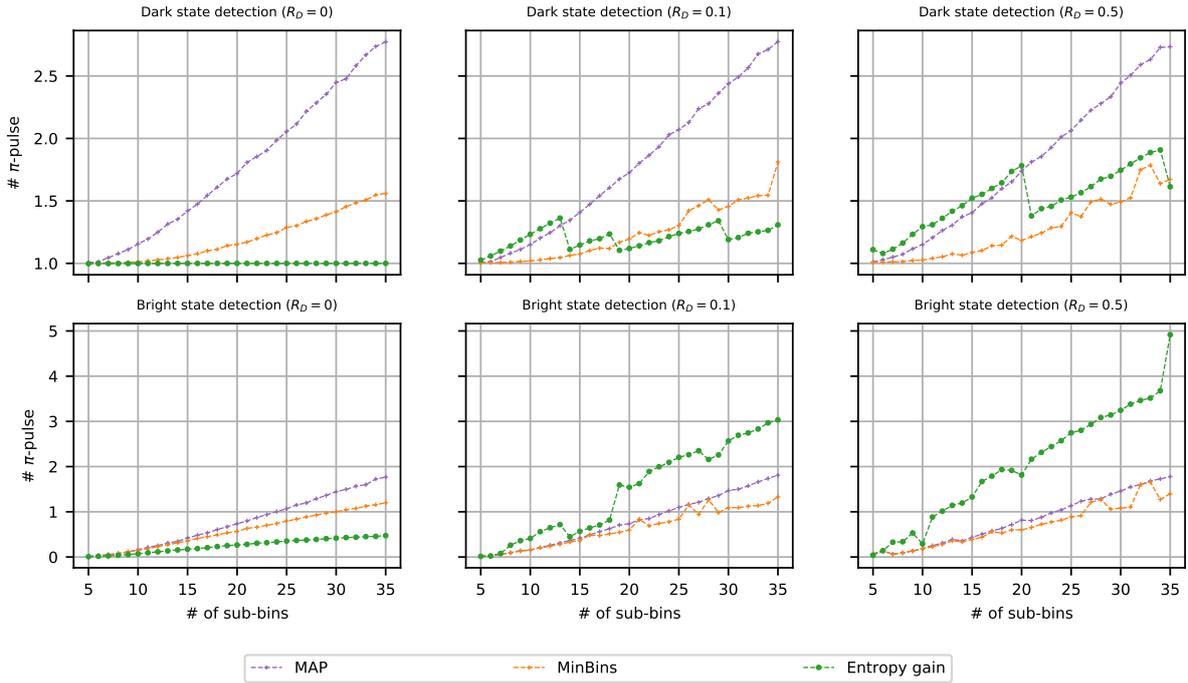


Figure 9.9: Average number of applied π -pulses for a given division into sub-bins. **Top row:** Dark state detection starting with zero background on the left, increasing moving right. **Bottom row:** Same as top but for bright state detection.

as both the entropy gain and the minimum bins estimator, increasing comparatively strongly for higher sub-bin numbers. In the case of dark state detection, the performance of the MAP estimators is on par with the minimum bins estimator. A nice feature of both the standard and modified minimum bins estimator is the independence if the average number of applied pulses with respect to the background present.

Chapter 10

Conclusion & Outlook

For this part of the thesis, we introduced hidden Markov models, the factor graph framework including belief propagation, and trapped ion qubit readout techniques. After analyzing the adaptive maximum likelihood method, we further developed the notion of qubit interacting state discrimination algorithms to speed up detection sequences. A simulation environment was developed to benchmark different estimators. We have shown that under ideal circumstances, π -pulse methods speed up the average dark detection time considerably by a factor of 4. We analyzed the robustness of estimators against non-ideal experiment parameters and concluded that the speed up persists although the relative performance decreases. With the groundwork being laid, we list future steps to further develop π -pulse methods:

- **Adding different detection devices:** Currently, the simulation environment only support detection using PMTs. It should be straightforward to add different detection devices like APD, EMCCD cameras, and CMOS cameras. The main change would be to adapt the count distributions and observation likelihoods.
- **Unsupervised learning of hidden Markov models:** Instead of relying on fitted parameters one could explore learning techniques which also estimate model parameters.
- **Implementing algorithms on control hardware:** While simulations are able to provide estimates, to properly assess performance of the analyzed estimation techniques they need to be implemented and tested in actual trapped-ion quantum experiments.

Bibliography

- [1] Peter W. Shor. “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer”. In: *SIAM Journal on Scientific and Statistical Computing* 26 (1997), p. 1484. URL: <http://www.citebase.org/abstract?id=oai:arXiv.org:quant-ph/9508027>.
- [2] L. K. Grover. “A fast quantum mechanical algorithm for database search”. In: *STOC '96*. 1996.
- [3] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. 10th. New York, NY, USA: Cambridge University Press, 2011. ISBN: 1107002176.
- [4] Grégoire Franck Mario Tomassi. “Concept analysis and fabrication of an ion trap to couple negative and positive ions”. ETH Zurich, 2021.
- [5] David P. DiVincenzo. “The Physical Implementation of Quantum Computation”. In: *Fortschritte der Physik* 48.9-11 (Sept. 2000), pp. 771–783. ISSN: 1521-3978. DOI: 10.1002/1521-3978(200009)48:9/11<771::aid-prop771>3.0.co;2-e. URL: [http://dx.doi.org/10.1002/1521-3978\(200009\)48:9/11%3C771::AID-PROP771%3E3.0.CO;2-E](http://dx.doi.org/10.1002/1521-3978(200009)48:9/11%3C771::AID-PROP771%3E3.0.CO;2-E).
- [6] M Johanning. “Isospaced linear ion strings”. In: *Applied Physics B* 122 (2016), pp. 1–17. DOI: 10.1007/s00340-016-6340-0. eprint: <https://doi.org/10.1007/s00340-016-6340-0>. URL: <https://doi.org/10.1007/s00340-016-6340-0>.
- [7] S. Skogestad and I. Postlethwaite. *Multivariate Feedback Control - Analysis and Design*. 2nd. The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ: John Wiley & Sons Ltd, 2005. ISBN: 978-0-470-01167-6.
- [8] R. P. Sallen and E. L. Key. “A practical method of designing RC active filters”. In: *IRE Transactions on Circuit Theory* 2.1 (1955), pp. 74–85. DOI: 10.1109/TCT.1955.6500159.
- [9] K. A. Kuhn. *Sallen-Key Low-pass Filter*. https://www.kennethkuhn.com/students/ee431/sallen_key_low_pass_filter.pdf. A. Jan. 2016. (Visited on 01/12/2022).
- [10] H. G. Dimopoulos. *Analog Electronic Filters*. Springer, 2012. Chap. 10. ISBN: 978-94-007-2189-0.
- [11] Greg Kasprovicz. *Fastino*. <https://github.com/sinara-hw/Fastino/wiki>. Sept. 2021. (Visited on 04/04/2022).
- [12] Analog Devices. *ADHV4702-1, 24 V to 220 V Precision Operational Amplifier*. <https://www.analog.com/media/en/technical-documentation/data-sheets/ADHV4702-1.pdf>. B. Jan. 2020. (Visited on 01/01/2022).
- [13] Linear Technology. *LTC6091, Dual 140V, Rail-to-Rail Output, Picoamp Input Current Op Amp*. <https://www.analog.com/media/en/technical-documentation/data-sheets/6091fa.pdf>. A. May 2015. (Visited on 01/01/2022).
- [14] Apex Microtechnology. *PA441/PA443, High Voltage Power Operational Amplifier*. <https://www.apexanalog.com/resources/products/pa441u-pa443u.pdf>. F. June 2020. (Visited on 01/20/2022).
- [15] V. Kaushal et al. “Shuttling-based trapped-ion quantum information processing”. In: *AVS Quantum Science* 2.1 (2020), p. 014101. DOI: 10.1116/1.5126186. eprint: <https://doi.org/10.1116/1.5126186>. URL: <https://doi.org/10.1116/1.5126186>.
- [16] H.-A. Loeliger. *Signal Analysis, Models, and Machine Learning*. 2019.
- [17] Matteo Marinelli. “Quantum information processing with mixed-species ion crystals”. PhD thesis. ETH Zurich, 2020.
- [18] Nick Schwegler. “Towards Low-Latency Parallel Readout of Multiple Trapped Ions”. ETH Zurich, 2018.

- [19] Sabine Wölk et al. “State selective detection of hyperfine qubits”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 48.7 (Mar. 2015), p. 075101. DOI: 10.1088/0953-4075/48/7/075101. URL: <https://doi.org/10.1088/0953-4075/48/7/075101>.
- [20] Luis A. Martinez, Yaniv J. Rosen, and Jonathan L. DuBois. “Improving qubit readout with hidden Markov models”. In: *Phys. Rev. A* 102 (6 Dec. 2020), p. 062426. DOI: 10.1103/PhysRevA.102.062426. URL: <https://link.aps.org/doi/10.1103/PhysRevA.102.062426>.
- [21] B Hemmerling et al. “A novel, robust quantum detection scheme”. In: *New Journal of Physics* 14.2 (Feb. 2012), p. 023043. DOI: 10.1088/1367-2630/14/2/023043. URL: <https://doi.org/10.1088/1367-2630/14/2/023043>.
- [22] A. H. Myerson et al. “High-Fidelity Readout of Trapped-Ion Qubits”. In: *Phys. Rev. Lett.* 100 (20 May 2008), p. 200502. DOI: 10.1103/PhysRevLett.100.200502. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.100.200502>.
- [23] Shawn Geller et al. *Improving quantum state detection with adaptive sequential observations*. 2022. DOI: 10.48550/ARXIV.2204.00710. URL: <https://arxiv.org/abs/2204.00710>.

Chapter 11

Appendix

11.1 Gain-Bandwidth Product

While we do not present a proof for the general validity of 3.3 we will show a motivating example. We will focus on the non-inverting amplifier topology which is part of many of the circuits outlined in thesis in order to set the DC gain. The transfer function is frequency-independent and has the following form.

$$H_{\text{ideal}}(s) = 1 + \frac{R_1}{R_2} = G_{\text{non-inv}} \quad (11.1)$$

If we allow for finite, frequency-dependent open-loop transfer function as in 3.1 we have a set of equations:

$$\begin{cases} V_{\text{out}} = (V_+ - V_-)A(s) = (V_{\text{in}} - V_-)A(s) \\ V_- = \frac{R_2}{R_1 + R_2}V_{\text{out}} \end{cases} \quad (11.2)$$

If we solve for the transfer function by substitution of V_- into the first equation we get:

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{R_1 + R_2}{R_2 + (R_1 + R_2)A^{-1}(s)} \quad (11.3)$$

We proceed to divide by R_2 to re-introduce the ideal closed-loop gain we finally find:

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{G_{\text{non-inv}}}{1 + G_{\text{non-inv}}A^{-1}(s)} = \frac{G_{\text{non-inv}}}{1 + \frac{s}{\omega_a} \frac{G_{\text{non-inv}}}{A_{DC}} + \frac{G_{\text{non-inv}}}{A_{DC}}} \quad (11.4)$$

Assuming $A_{DC} \gg G_{\text{non-inv}}$ and defining $\omega_c = \omega_a \frac{A_{DC}}{G_{\text{non-inv}}}$ we have:

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{G_{\text{non-inv}}}{1 + \frac{s}{\omega_c} + \underbrace{\frac{G_{\text{non-inv}}}{A_{DC}}}_{\rightarrow 0}} \approx \frac{G_{\text{non-inv}}}{1 + \frac{s}{\omega_c}} \quad (11.5)$$

which has the same form as the open-loop transfer function of an operational amplifier. We have reduced the gain by the same factor as we have increased the bandwidth. Therefore, also the closed-loop transfer function satisfies the GBW.

11.2 Repositories

In this section we will briefly list the repositories and the containing design files.

1. <https://gitlab.phys.ethz.ch/tiqi-projects/qchub/out-of-vacuum-amplifier-boards>

- Measurements: Contains all the measurements used to create the corresponding plots in this thesis. A jupyter notebook can be used to regenerate the plots if necessary.
- Orders: Contains GERBER files used to order PCB. Also, it contains some orders for passive components for the amplifier board.

- PCB: Contains Altium design files of the amplifier prototype board, the full prototype, and the regulator board.
 - SPICE: Contains simulation files for the amplifier circuit design. The whole chain from Fastino to a lumped trap model can be used to simulate the design.
 - Script: Contains a jupyter notebook which implements automatic component search for a Sallen-Key low-pass filter. It also serves as an example to automate SPICE simulations using python.
2. <https://gitlab.phys.ethz.ch/tiqi-projects/qchub/scsi-to-dsub-passive-filter>
 - PCB: Contains Altium design files for the passive filter board.
 3. <https://github.com/marti-m/ion-detect-mcmc>
 - Contains the python package which has been used to perform MCMC simulations. All batch files which were run for estimator performance assessment can be found.



Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Declaration of originality

The signed declaration of originality is a component of every semester paper, Bachelor's thesis, Master's thesis and any other degree paper undertaken during the course of studies, including the respective electronic versions.

Lecturers may also require a declaration of originality for other written papers compiled for their courses.

I hereby confirm that I am the sole author of the written work here enclosed and that I have compiled it in my own words. Parts excepted are corrections of form and content by the supervisor.

Title of work (in block letters):

Low-Noise DC Amplifier Design for Trapped Ion Setups & Bayesian Schemes for Fast Single Qubit Detection

Authored by (in block letters):

For papers written by groups the names of all authors are required.

Name(s):

Marti

First name(s):

Michael

With my signature I confirm that

- I have committed none of the forms of plagiarism described in the '[Citation etiquette](#)' information sheet.
- I have documented all methods, data and processes truthfully.
- I have not manipulated any data.
- I have mentioned all persons who were significant facilitators of the work.

I am aware that the work may be screened electronically for plagiarism.

Place, date

Bern, 11.04.22

Signature(s)

M. Marti

For papers written by groups the names of all authors are required. Their signatures collectively guarantee the entire content of the written paper.