Department of Physics and Astronomy

University of Heidelberg

Master thesis

in Physics

submitted by

Marcell W. Gall

born in Krumbach (Schwaben)

2015

Active Magnetic Field Stabilisation for Ultracold

Sodium Lithium Mixtures

This Master thesis has been carried out by Marcell W. Gall

at the

Kirchhoff Institute of Physics

under the supervision of

Prof. Dr. Markus K. Oberthaler

Abstract

This thesis discusses technical improvements done at a sodium lithium mixtures experiment, to enable measurements highly sensitive on the magnetic field.

An active magnetic field stabilisation for low magnetic fields of less then one gauss was tested and added to the experiment. The focus was put on the suppression of field oscillations from the power lines, cancellation of external influences and long term stability. The prototype configuration and its performance is then reviewed and possible future improvements are discussed.

For precise atom population measurements a new control and readout camera software was written and implemented. A systematic evaluation of the settings was performed and subsequently used to optimise the lithium imaging.

Zusammenfassung

Diese Masterarbeit beschreibt technische Verbesserungen an einem Natrium-Lithium-Mischungsexperiment, um magnetisch hochsensitive Messungen zu ermöglichen.

Eine aktive Magnetfeldstabilisierung für kleine Magnetfelder unter einem Gauss wurde erprobt und in das Experiment eingefügt. Der Fokus hierbei lag auf der Unterdrückung von Magnetfeldoszillationen hervorgerufen durch Netzströme, das Ausgleichen von externen Einflüssen auf das Feld und dessen Langzeitstablitiät. Die Konfiguration des Prototypsystems wird besprochen und zukünftig mögliche Verbesserungen diskutiert.

Eine Kontroll- und Auslesesoftware für die Kameras zum präzisen Bestimmen der Atomzahlen wurde neugeschrieben und im Experiment eingesetzt. Nach einer systematischen Untersuchung des Einflusses verschiedener Einstellungen wurde diese dazu benutzt um das Lithium Bildgebungsverfahren zu optimieren.

Contents

1 Introduction

*Headlines, in a way, are what mislead you because bad news is a headline,
and gradual improvement is not.*

— Bill Gates

The NaLi machine, a dual species degenerate quantum gas mixtures experiment, has now existed for over 9 years [26, 40] and during that time over 16 people worked there to gain knowledge on sodium-lithium mixtures. Feshbach physics [31, 37], motional coherence [29] and bath induced energy shifts [25] were researched. This was only possible due to the continuous improvement of a very complex machine. In the sense of the quote above, I would like to present no headline, but technical improvements future members of this research group can base their work on.

The experiment was previously designed to create large and stable magnetic fields for Feshbach resonances. Fields above $2000\,\mathrm{G}$ could be created and used for measurements [37]. Through precise control and regulation of the coil currents the magnetic field could be passively stabilised to below $10\,\mathrm{mG}$ at a field of $1200\,\mathrm{G}$. This stability of a few ten parts per million level for high magnetic fields was necessary to conduct the experiments and reach an understanding of the underlying physics [30].

The current research goal has shifted to the other extreme, where we want to observe hetero-nuclear spin dynamics in very low absolute magnetic field, far smaller than a gauss. This has to be done using new methods and tools, since problems in this regime differ a lot from the high field. One now also has to include the earth magnetic field and other, man-made external fields which are comparable in strength to the field magnitudes we would like to achieve. The main issue here is that their day-to-day and hour-to-hour fluctuations are out of our control and make reaching an absolute stability very hard. Getting to these low fields also leads to additional relevance of background magnetic field noise. All power cables in the lab using $50\,\mathrm{Hz}$ AC current now act as an antenna and produce field oscillations in the order of a few milligauss. The only way to achieve a field stability, comparable to the high field experiments of 100 parts per million, is to actively regulate the field with a feedback control in all directions. The implementation of this active magnetic field stabilisation was one focus of this work.

What complicated the experiments planned even further was the high accuracy with which the field sensitive changes had to be detected. We would like to observe relative spin population changes within only a few percent accuracy. Without the capabilities to detect these changes as precise as possible it would not be possible to characterise the effects of the magnetic field stabilisation on the atoms. Therefore a lot of work was first put into optimising the lithium imaging. For that a new software for the readout of the sodium and lithium camera was written and implemented in our experiment.

This thesis is structured as following:

- In chapter two the basics of light-atom interaction are reviewed and put into context of the near resonant photon absorption of atoms. The magnetic field induced energy level shifts on the atoms is also described. Additionally the basics of feedback regulation is explained, including the specific case of a PID controller.

- Chapter three will focus on the improvement of the imaging sequence and the new camera software. By explaining the absorption imaging method and its issues we can identify important experimental parameters we would like to adjust. The description of the camera program then features these parameters and in a last section we will show the optimisation to the lithium done with them.

- The magnetic field stabilisation will be topic in chapter four. After an initial characterisation of the field with two different sensors, a test system for the active field stabilisation will be explained. Afterwards the prototype added to the experiment will be described. We could achieved an at least four-fold background noise suppression simultaneously in each axis at the sensor position. The relative difference in magnetic field compared to the place of the atoms will be explained and the regulation to less then $4\,\mathrm{mG}$ total field measured with atoms will be discussed.

- The last chapter will give a conclusive summary of the achieved results and describe the future steps to improve this set up to our experimental needs.

2 Light Absorption and Resonance Conditions

Even though the scope of this work is mostly technical, we would like to introduce some theoretical concepts for better understanding of the underlying physics.

2.1 Atom-Light Interaction

One of the most important tools in the field of ultracold atoms is the manipulation of atoms with the use of coherent laser light. The theory behind this is discussed in detail in the literature [21, 32, 15]. The parts relevant to this work will be summarised in this section.

We will look at the simple case of an atom having only two discrete energy levels, an excited state $|e\rangle$ and a ground state $|g\rangle$ which are eigenstates of the atomic Hamiltonian $\mathcal{H}_0$. Their eigenenergys are $\hbar\omega_g$ and $\hbar\omega_e$ respectively. Each pure state $|\Psi\rangle$ of the atom then can be described as:

$$|\Psi\rangle = c_g(t)\,|g\rangle + c_e(t)\,|e\rangle \tag{2.1}$$

$$\langle\Psi|\Psi\rangle = |c_g(t)|^2 + |c_e(t)|^2 = 1 \tag{2.2}$$

Their undisturbed time evolution can be derived with the Schrödinger equation:

$$i\hbar\frac{\partial}{\partial t}\,|\Psi(t)\rangle = \mathcal{H}\,|\Psi(t)\rangle = \hbar\omega_g c_g(t)\,|g\rangle + \hbar\omega_e c_e(t)\,|e\rangle \tag{2.3}$$

$$|\Psi(t)\rangle = c_g(0)e^{-i\omega_g t}\,|g\rangle + c_e(0)e^{-i\omega_e t}\,|e\rangle \tag{2.4}$$

With an additional light matter interaction term $\mathcal{H}_L$ to the non disturbed Hamiltonian $\mathcal{H}_0$ we get the time evolution of these states with $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_L$:

$$i\hbar\frac{\partial}{\partial t}\,|\Psi(t)\rangle = [\mathcal{H}_0 + \mathcal{H}_L]\,|\Psi(t)\rangle \tag{2.5}$$

Using equation (2.4) as an ansatz, by letting c_g and c_e be time dependent again and applying further manipulations we get [21]:

$$i\hbar\frac{\partial}{\partial t}c_g(t) = i\hbar\,\dot{c}_g(t) = c_e(t)\,\langle g|\,\mathcal{H}'_L\,|e\rangle\,e^{-i\omega_{eg}t} \tag{2.6}$$

$$i\hbar\frac{\partial}{\partial t}c_e(t) = i\hbar\,\dot{c}_e(t) = c_g(t)\,\langle e|\,\mathcal{H}'_L\,|g\rangle\,e^{i\omega_{eg}t} \tag{2.7}$$

Where $\omega_{eg} = \omega_e - \omega_g$. Using the semi-classical description and a dipole operator $\vec{d} = -e \cdot \vec{r}$ whereas e denotes the elementary charge and r the position operator, the interaction between the electric field operator $\vec{\mathcal{E}}(t, \vec{r})$ and the atom becomes:

$$\mathcal{H}_L = \vec{d} \cdot \vec{\mathcal{E}}(t, \vec{r}) \tag{2.8}$$

Describing an elementary field as a plane wave travelling along the z axis with an wave vector $\vec{k}$ and a frequency of $\omega_L = |k| \cdot c$, the electric field operator becomes:

$$\vec{\mathcal{E}}(t, \vec{r}) = E_0 \hat{\epsilon} \cos(kz - \omega_L t) \tag{2.9}$$

With E_0 denoting the amplitude and $\hat{\epsilon}$ the polarisation vector . Since the wavelength of visible light $\lambda = 2\pi/k$ is orders of magnitude larger than the extension of an atom we can use the dipole approximation:

$$\vec{\mathcal{E}}(t, \vec{r}) \approx \vec{\mathcal{E}}(t, \vec{r} = 0) = E_0 \hat{\epsilon} \cos(\omega_L t) \tag{2.10}$$

Using this in equation (2.6) and (2.7) leads us to the following:

$$i\hbar \dot{c}_g(t) = -c_e(t) E_0 \, e \, \hat{\epsilon} \cdot \vec{d}_{ge} \frac{1}{2} \left[e^{+i(\omega_L - \omega_{eg})t} + e^{-i(\omega_L + \omega_{eg})t} \right] \tag{2.11}$$

$$i\hbar \dot{c}_e(t) = -c_g(t) E_0 \, e \, \hat{\epsilon} \cdot \vec{d}_{eg} \frac{1}{2} \left[e^{-i(\omega_L - \omega_{eg})t} + e^{+i(\omega_L + \omega_{eg})t} \right] \tag{2.12}$$

with: $\vec{d}_{eg} = \langle e | \vec{d} | g \rangle = \vec{d}_{ge}^*$

The * is denoting the complex conjugate of $\vec{d}_{eg}$. Making use of the rotating-wave approximation (RWA) one can argue that any terms of order $1/\omega_{eg}$ can be neglected compared to terms with order $1/\delta$, where $\delta = \omega_L - \omega_{eg}$. Here δ is the detuning of the laser to resonance.
We hence can come to the result:

$$\dot{c}_g(t) = ic_e(t)(\vec{d}_{ge} \cdot \hat{\epsilon}) \frac{E_0 \, e}{2\hbar} e^{+i\delta t} = ic_e(t) \frac{\Omega_0}{2} e^{+i\delta t} \tag{2.13}$$

$$\dot{c}_e(t) = ic_g(t)(\vec{d}_{ge}^* \cdot \hat{\epsilon}) \frac{E_0 \, e}{2\hbar} e^{-i\delta t} = ic_g(t) \frac{\Omega_0^*}{2} e^{-i\delta t} \tag{2.14}$$

Here $\Omega = (\vec{d}_{ge} \cdot \hat{\epsilon}) \frac{E_0 \, e}{\hbar}$ is the Rabi frequency. Using the time derivatives of equation (2.13) and (2.14) and then substituting with $c_e(t)$ leads to an analytically solvable differential equation system:

$$\frac{\partial^2}{\partial t^2} c_g(t) - i\delta \frac{\partial}{\partial t} c_g(t) + \frac{\Omega^2}{4} c_g(t) = 0 \tag{2.15}$$

$$\frac{\partial^2}{\partial t^2} c_e(t) + i\delta \frac{\partial}{\partial t} c_e(t) + \frac{\Omega^2}{4} c_e(t) = 0 \tag{2.16}$$

With initial conditions of $c_g(0) = 1$ and $c_e(0) = 0$ we get:

$$c_g(t) = \left(\cos \frac{\Omega_{eff} t}{2} - i \frac{\delta}{\Omega_{eff}} \sin \frac{\Omega_{eff} t}{2} \right) e^{+i\delta t/2} \tag{2.17}$$

$$c_e(t) = -i \frac{\Omega}{\Omega_{eff}} \sin \frac{\Omega_{eff} t}{2} e^{-i\delta t/2} \tag{2.18}$$

$$\text{with:} \quad \Omega_{eff} = \sqrt{\Omega^2 + \delta^2} \tag{2.19}$$

This is the solution to the so called Rabi oscillations.

Optical Bloch Equations

We will now use these results in the context of the density matrix description, where the density matrix ρ of a state $|\Psi\rangle$ is defined as:

$$\rho = |\Psi\rangle \langle \Psi| \tag{2.20}$$

And more specifically for the two level system:

$$\rho = \left\{ c_g(t) |g\rangle + c_e(t) |e\rangle \right\} \cdot \left\{ c_g(t) \langle g| + c_e(t) \langle e| \right\} \tag{2.21}$$

$$= \begin{pmatrix} \rho_{ee} & \rho_{eg} \\ \rho_{ge} & \rho_{gg} \end{pmatrix} = \begin{pmatrix} c_e c_e^* & c_e c_g^* \\ c_g c_e^* & c_g c_g^* \end{pmatrix} \tag{2.22}$$

Here one can see that $\rho_{ee} = c_e c_e^* = p_e$ is directly giving the probability to be in the excited state. The same is true for $\rho_{gg} = p_{gg}$ and the ground state.
It's time evolution is given by the Von-Neumann equation:

$$i\hbar \frac{\mathrm{d}}{\mathrm{d}t} \rho = [\mathcal{H}, \rho] \tag{2.23}$$

Which leads to the same result as the solution of equations (2.15) and (2.16). The big advantage of the density matrix description is the easy addition of the spontaneous decay effect with a constant rate of $\gamma = 1/\tau$ with τ being the lifetime of the excited state [21]. One can calculate:

$$\gamma = \frac{\omega_{eg}^3 d_{eg}^2}{3\pi \epsilon_0 \hbar c^3} \tag{2.24}$$

Using also that $\tilde{\rho}_{ge} = \rho_{ge} e^{-i\delta t}$ we get for the time evolution [21]:

$$\frac{\mathrm{d}}{\mathrm{d}t} \rho_{gg} = +\gamma \rho_{ee} + \frac{i}{2} \left(\Omega^* \tilde{\rho}_{eg} - \Omega \tilde{\rho}_{ge} \right) \tag{2.25}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} \rho_{ee} = -\gamma \rho_{ee} + \frac{i}{2} \left(\Omega \tilde{\rho}_{ge} - \Omega^* \tilde{\rho}_{eg} \right) \tag{2.26}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} \tilde{\rho}_{ge} = - \left(\frac{\gamma}{2} + i\delta \right) \tilde{\rho}_{ge} + \frac{i}{2} \Omega^* \left(\rho_{ee} - \rho_{gg} \right) \tag{2.27}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} \tilde{\rho}_{eg} = - \left(\frac{\gamma}{2} - i\delta \right) \tilde{\rho}_{eg} + \frac{i}{2} \Omega^* \left(\rho_{gg} - \rho_{ee} \right) \tag{2.28}$$

These important equations are the so called optical Bloch equations. They directly show that the system dynamics are driven by the coherences ρ_{eg} and ρ_{ge}.

Scattering Rate and Power Broadening

Assuming a coherent light field we can set $\rho_{ge} = \rho_{eg}^*$. Which in combination with the optical Bloch equations will give us for the time evolution of the population difference $w = \rho_{gg} - \rho_{ee}$:

$$\frac{\mathrm{d}}{\mathrm{d}t}\rho_{eg} = -\left(\frac{\gamma}{2} - i\delta\right)\rho_{eg} + \frac{i\omega\Omega}{2} \tag{2.29}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}w = -\gamma w - i(\Omega\rho_{eg}^* - \Omega^*\rho_{eg}) + \gamma \tag{2.30}$$

We define the saturation parameter s and the on-resonance saturation parameter s_0 as follows:

$$s \equiv \frac{|\Omega|^2}{2\delta^2 + \gamma^2/2} = \frac{I/I_S}{1 + (2\delta/\gamma)^2} = \frac{s_0}{1 + (2\delta/\gamma)^2} \tag{2.31}$$

with the incoming intensity I and the saturation intensity $I_S \equiv \pi h c/3\lambda^3\tau$. Knowing that the excitation probability ρ_{ee} and therefore also w will reach an equilibrium after a certain time we can get:

$$\text{steady state:} \quad w = \frac{1}{1+s} = \rho_{gg} - \rho_{ee} = 1 - 2\rho_{ee} \tag{2.32}$$

This then can be used to define an excitation or scattering rate γ_P that uses the fact that in the steady state the decay rate of the excited state is equal to the scattering rate:

$$\gamma_P = \gamma\rho_{ee} = \frac{s_0\gamma/2}{1 + s_0 + (2\delta/\gamma)^2} \tag{2.33}$$

To adapt this in our specific case we can use the values for the fermionic lithium ^{6}Li D2 line, with a decay rate $\gamma = 1/\tau = 36.898 \times 10^6\,\mathrm{s}^{-1} = 2\pi \times 5.8724\,\mathrm{MHz}$ and a saturation intensity of $I_S = 2.54\,\mathrm{mW/cm}^2$ [13]. Figure 2.1 shows the scattering rate γ_P and the power broadening of the absorption probability in several settings for the saturation parameter $s_0 = I/I_S$. We usually work with an imaging laser intensity of around $I \approx 0.25\,\mathrm{mW/cm}^2$, resulting in $s_0 \approx 0.1$. With increasing intensity one reaches a threshold of $\gamma_P = \gamma/2$ on resonance, but the shape of the scattering rate also broadens.

Doppler Shift of Accelerated Atoms

In the previous section we discussed the absorption/excitation probability for an atom at rest and described it with the scattering rate γ_P. The light field resonant to the atomic transition $\omega_{eg} = 2\pi \cdot \nu_{eg}$ had to have the frequency of $\nu_L = \nu_{eg}$. If the atom is however moving relative to the observer system with velocity $\vec{v}$ anti-parallel to the light wave vector $\vec{k}$, then the resonance frequency of the atom is shifted:
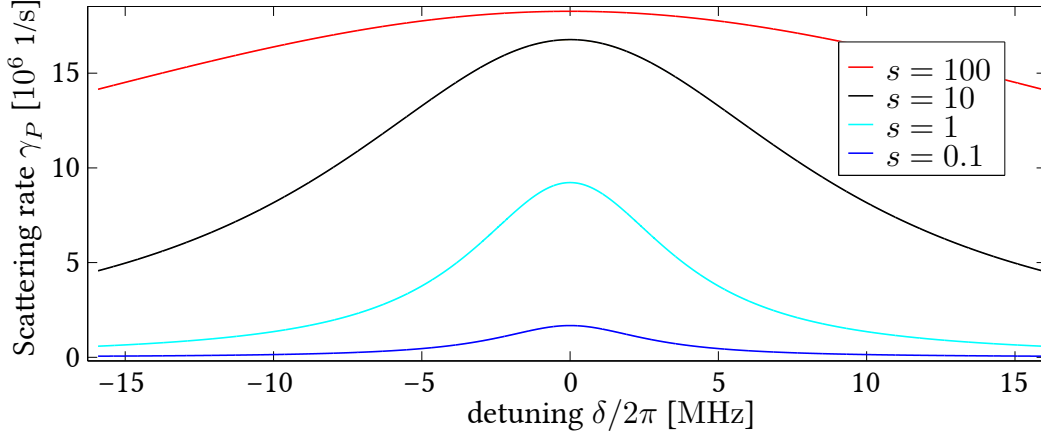
$$\nu(\vec{v}) = \nu_{eg}(1 + \frac{v}{c}) \tag{2.34}$$

Figure 2.1: Scattering rate of ^{6}Li with different saturation parameters $s = I/I_S$. One can see that the absorption line broadens up significantly with intensities I much bigger than the saturation intensity I_S ($s > 1$). For imaging in our lab we are usually at $s \approx 0.1$. The D2 line width is $2\pi \times 5.8724\,\text{MHz}$.

Looking again at a ^{6}Li atom, we can calculate the recoil velocity v_r of an initially resting atom $v_\text{i} = 0$ after absorbing the momentum of a photon resonant on the D2 transition wavelength $\lambda = c/\nu_\text{eg} = 670.977\,\text{nm}$ [13]:

$$\frac{h}{\lambda} = v_\text{r} \cdot m_{^6\text{Li}} \tag{2.35}$$

$$v_\text{r} = \frac{h}{m_{^6\text{Li}} \cdot \lambda} = 9.887\,\text{cm}\,\text{s}^{-1} \tag{2.36}$$

Here we used the atomic mass of ^{6}Li: $m_{^6\text{Li}} = 9.988 \times 10^{-27}\,\text{kg}$ [13].
The absorption of a single photon therefore results in a Doppler shift of:

$$\Delta\nu = \nu(\vec{v}_\text{r}) - \nu(0) = \nu_\text{eg}\frac{v_\text{r}}{c} = 0.1473\,\text{MHz} \tag{2.37}$$

Which is about 2.5 percent of the absorption line width of ^{6}Li. The emission direction of the photon afterwards can be assumed to be mostly random, since the laser intensity is not high enough for stimulated emission. So in average over many excitation cycles it can be considered not to change the velocity. Atoms which were initial at resonance with an incoming laser, are therefore accelerated out of resonance.

We would like to quantify this assuming an ensemble of many atoms using an iterative approach. Integrating the scattering rate from equation (2.33) over a sample time t_s will give us the averaged number of scattering events in that time step depending on the detuning. The change in detuning then can be calculated by equation (2.36) and (2.37). Figure 2.2 shows the results of this simulation for fermionic lithium on the D2 transition. One can identify in figure 2.2(a) the different slopes for resonant and off resonant absorption. Also with initial blue detuning one sees the passing through resonance in figure 2.2(b) which is then translated to the steep slope in the number of scattering events. For
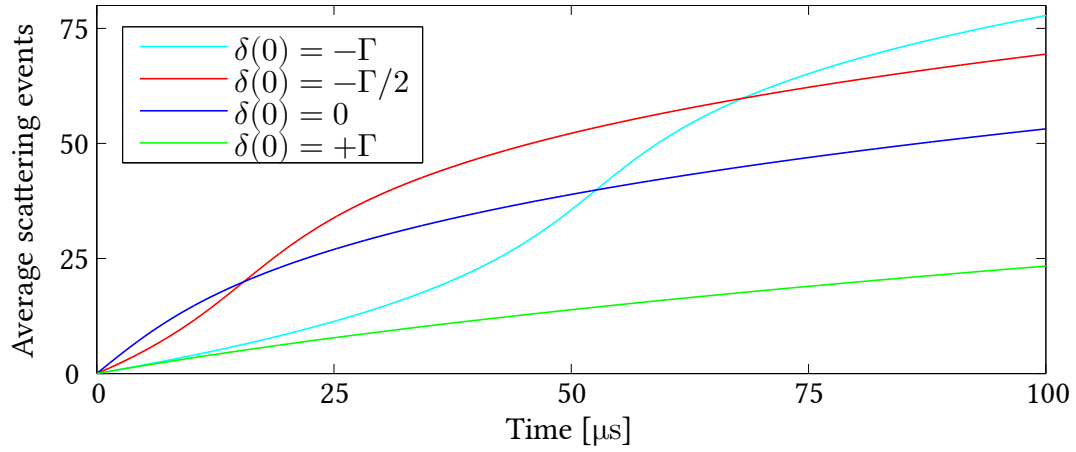
(a) Scattering events dependent on the initial detuning
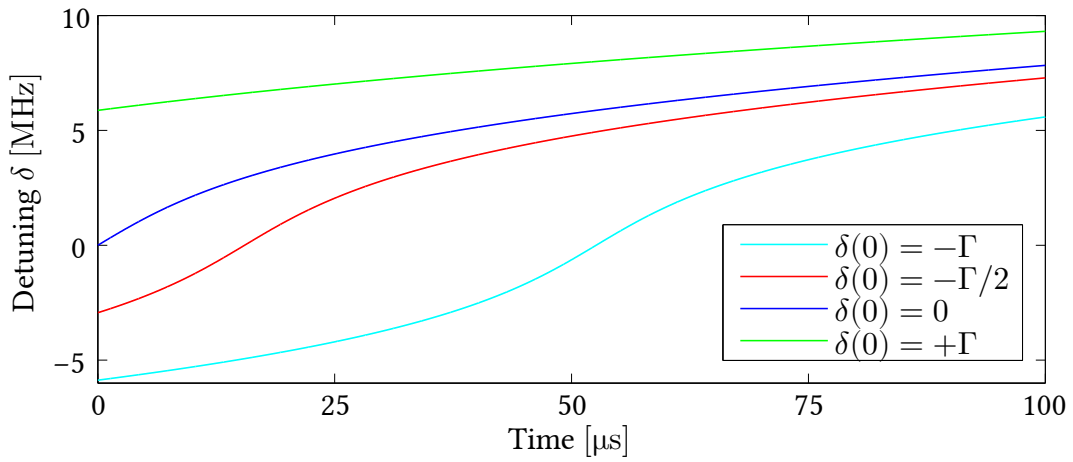
(b) Change in detuning dependent on the initial value

Figure 2.2: Effect of acceleration due to photon scattering processes. The calculation was done with $s = 0.1$ and is shown for different initial detunings. Here $\Gamma = \gamma/2\pi = 5.872\,\mathrm{MHz}$ is for the D2 line of $^6\mathrm{Li}$. One can identify the different slopes for in (a) resonant and far off resonant absorption.

measurements on atom numbers we would like to stay most of the time in the resonant absorption regime with linear scattering slope while still having a lot of scattering events. This can be achieved for example with initial detuning of $\delta(0) = -\Gamma/2$.

2.2 Magnetic Energy Level Shifts

In the previous section the atom was described with only the light wave interacting with the system. The magnetic field however also can influences the atom and for example shift the internal energy levels. We saw already that a shift in energy leads to a different resonance condition for an incoming laser beam and an atom. Magnetic energy shifts are however also extremely important in other processes we want to measure, for example in Feshbach spectroscopy, spin-changing collisions or can be used for trapping the atoms. This section uses [15, 35, 13] and presents their results.

Many atom states posses a non zero magnetic moment $\vec{\mu}$ which interacts with the external magnetic field $\vec{B}$ such that it introduces an additional energy term [15]:

$$E(B) = -\vec{\mu} \cdot \vec{B} \tag{2.38}$$

In the regime of low magnetic fields this is called the Zeeman effect. Here, the magnetic moment comes from the angular momentum $\vec{F}$ of the Hyperfine states with a quantisation axis $\hat{z}$ along the magnetic field direction. The quantum number F can be used to describe the absolute value of the angular momentum: $|\vec{F}| = \hbar\sqrt{F(F+1)}$. This angular momentum is then quantised along the magnetic field direction, and described with the quantum number m_F: $F_z = m_F \cdot \hbar$. The total angular momentum $\vec{F}$ is the result of a magnetic coupling between the electron spin $\vec{S}$, the orbital spin $\vec{L}$, which form together the total electron momentum $\vec{J}$, and the nuclear spin $\vec{I}$, which again can be described the same way with their quantum numbers S, L, J and I.

The relation between the angular momentum along $\hat{z}$ and the magnetic moment in that direction is given by:

$$\mu_z = -g_F \mu_B \frac{F_z}{\hbar} = -g_F \mu_B m_F \tag{2.39}$$

$$\Rightarrow E_{\text{Zeeman}} = g_F \mu_B m_F B_z \tag{2.40}$$

With the Bohr magneton μ_B and the gyromagnetic factor g_F we get [13]:

$$g_F = g_J \frac{F(F+1) - I(I+1) + J(J+1)}{2F(F+1)} + g_I \frac{F(F+1) + I(I+1) - J(J+1)}{2F(F+1)} \tag{2.41}$$

The values of g_J and g_I can are listed in [4] and since for Lithium and Sodium $g_J \gg g_I$ the second term can be neglected most of the time. With B comparable or larger than the fine structure splitting ΔE_{hfs} one reaches a regime where F isn't a good quantum number anymore and the states have to be described again with J and I [35]:

$$E_{PB} = \mu_B \left(g_J m_J + g_I m_I \right) B_z \approx \mu_B g_J m_J B_z \tag{2.42}$$
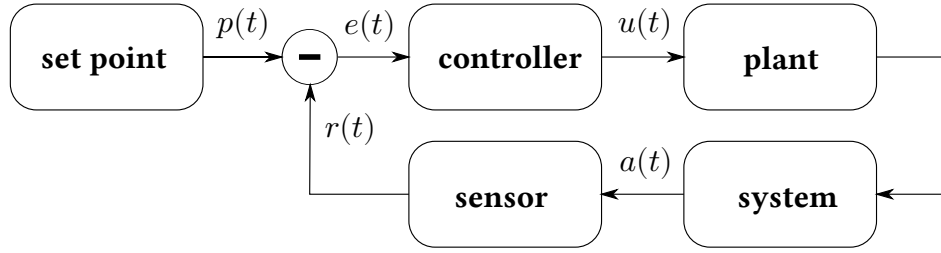
Figure 2.3: Scheme of a basic feedback controller loop. For a real system one has to include disturbances to the system, which do not stem from the plant.

The intermediate regime usually has to be calculated numerically. However for all states with $J = 1/2$ one can use the analytical Breit-Rabi formula [9] to calculate the energy depending on the magnetic field. All alkali ground states fulfil this condition and so we can use for sodium and lithium:

$$E_{BR,J=1/2} = -\frac{\Delta E_{\mathrm{hfs}}}{2(2I+1)} + g_I \mu_{\mathrm{B}} m_F B \pm \frac{\Delta E_{\mathrm{hfs}}}{2}\left(1 + \frac{4m_F x}{2I+1} + x^2\right)^{\frac{1}{2}} \quad (2.43)$$

$$\text{using:} \quad x = \frac{(g_J - g_I)\mu_{\mathrm{B}} B}{\Delta E_{\mathrm{hfs}}}$$

Here ΔE_{hfs} denotes the zero field energy difference of the hyperfine states. With knowing all variables in this equation the energy splittings can be calculated very accurately. On the other hand it is also possible to scan the energy difference with microwaves with very good frequency control and then calculating the magnetic field present at the place of the atom. The atoms themselves are actually the best magnetic field sensors [33, 22].

2.3 Control and Regulation Theory

To change the state of a system e. g. a prameter in our experiment in a controlled way, one will need a regulation. We will summarize the relevant information needed for the later chapters out of [3], [39] and [8].

As an example for a regulation, one can take a human driving a car: The driver will try to control the speed of the car to a desired value he would like to maintain. He can change the speed of the car by pushing the gas pedal and accelerating the vehicle.
We can describe the current state of the car, e. g. the speed, with the actual value $a(t)$. This can't be directly accessed, but by using a sensor, e. g. a tachometer, we get a sensor value of $r(t) = 80\,\mathrm{km/h}$. The set point of our regulation $p(t)$ is the value we would like $r(t)$ to reach, in our case the speed the driver has in mind $r(t) = 100\,\mathrm{km/h}$. The difference between the sensor value and the set point is the error signal $e(t) = p(t) - r(t) = -20\,\mathrm{km/h}$. The error signal is then used by a controller, in our case the human, to generate a controller output $u(t)$, e. g. the driver would increase the force pushing down the gas pedal. This controller output $u(t)$ then leads the system plant to try to

change the system parameter $a(t)$, where in our example the plant would be the motor getting more gas to increase the speed of the car.

This whole process of regulation is called a feedback controller loop, shown in figure 2.3. In the lab, the set point, the calculation of the error signal and controller are usually separated, whereas in our example all three are done and set by the driver. It is however useful to always think of them separately. For us, most of the times $r(t)$, $p(t)$, $e(t)$ and $u(t)$ are voltages, enabling an electric regulation.

We are interested in the time dependence of this regulation, so for each step in the loop we have to look at the time dependent reaction to an input parameter. We are for example interested in how the controller output $u(t)$ depends on the input of $e(t)$.

The PID Controller

A very popular design to construct and to model such controller is the proportional-integral-differential (PID) controller. We will introduce some concepts we can later use for other parts of the feedback loop.

The behaviour of a PID controller can be described with three components:

$$u(t) = \underbrace{K_P e(t)}_{\text{proportional(P)}} + \underbrace{K_I \int_0^t e(t')\mathrm{d}t'}_{\text{integral(I)}} + \underbrace{K_D \frac{\mathrm{d}}{\mathrm{d}t} e(t)}_{\text{differential(D)}} \tag{2.44}$$

Here K_P, K_I and K_D describe the gain of the different controller parts, with which we can tune the controller behaviour.

The P part of the regulation gives an output directly proportional to the input error signal. This is directly following the time behaviour of the error signal. The I part however is integrating over the error signal and gets stronger as the measured regulation value stays below or above the set point. The D part is sensitive to the derivation of the error signal and is used to reduce the steepness of the control output in order to minimize an overreaction of the controller.

The driver in our example would initially react to the observed error signal, e. g. the car is too slow, by pressing down the gas pedal to a certain amount, which would be the proportional part of our description. After observing that the speed is still to slow, the driver would over time gradually increase the gas, corresponding to an integration of the error signal. However if the the change in speed, e. g. the acceleration is too strong and the driver gets pushed in his seat he will reduce the gas again, therefore adding a differential control element.

For a different notation one can use the bilateral Laplace transformation $\mathcal{L}$ on the input $e(t)$ and output signal $u(t)$ of the controller to get the transfer function $G(s)$ which is used in signal theory [8]. This transformation shifts the description from the real time t to the complex frequency s domain, which includes an imaginary part to describe the

frequency dependent change through the transfer in the dynamical system.

$$\mathcal{L}\left\{e(t)\right\} = \int_{-\infty}^{\infty} e(t)e^{-st}\mathrm{d}t = X(s)$$

$$\mathcal{L}\left\{u(t)\right\} = \int_{-\infty}^{\infty} u(t)e^{-st}\mathrm{d}t = Y(s)$$

$$G(s) = \frac{Y(s)}{X(s)} = K_P + K_I\frac{1}{s} + K_D s \tag{2.45}$$

$$= K_P\left(1 + \frac{1}{T_I s} + T_D s\right) \tag{2.46}$$

Where $T_I = K_P/K_I$ is called integration time and $T_D = K_D/K_P$ the differentiation time. It can directly be seen that for high input frequencies s the differential part $K_D \cdot s$ is dominant in the response of the controller, whereas for very small frequencies it is the integral part K_I/s.

Going back to the example with the driver controlling the car speed:
With only a very small proportional part the driver can not reach his set point, but only asymptotically approach an equilibrium speed. To see this one can think of the air friction, decelerating the car. If now at a certain gas input level (controller output) the air friction force is equal to the accelerating force, the speed doesn't change anymore. This output signal is only proportional to the error signal, so with no change in speed, the whole loop stays constant and the set point is never reached. To change this state of equilibrium, the controller needs an integral part, meaning after a while, still not reaching the desired car speed, the driver would increase gas, effectively using an integral part. We therefore always need an integral part to reach the set point desired in our regulation!
If the driver however would react to any deviation from his set point too sensitive, corresponding to a too big amplification K_P, he would give more gas than necessary. The car would soon be accelerated so fast that the car speed is over the desired set point and the driver would break again very strongly. By being too slow the same process repeats. This leads to an oscillatory behaviour in the regulation which is not wanted. If the oscillations don't decrease over time, one has reached the so called critical behaviour, which will be important later in this chapter. With an even bigger amplification $K_P > K_P(\text{critical})$ these oscillations will increase in time, leading to the collapse of the regulation.
This qualitative description includes some parameters of the system-plant behaviour that should be described more accurately, like the delay between reaction of input and reaction of the system.

System Plant

We can define in the same way a transfer function for the plant P and the system S and a combined system-plant transfer function H:

$$H(s) = P(s) \cdot S(s) = \frac{\mathcal{L}\left\{a(t)\right\}}{\mathcal{L}\left\{u(t)\right\}} \tag{2.47}$$
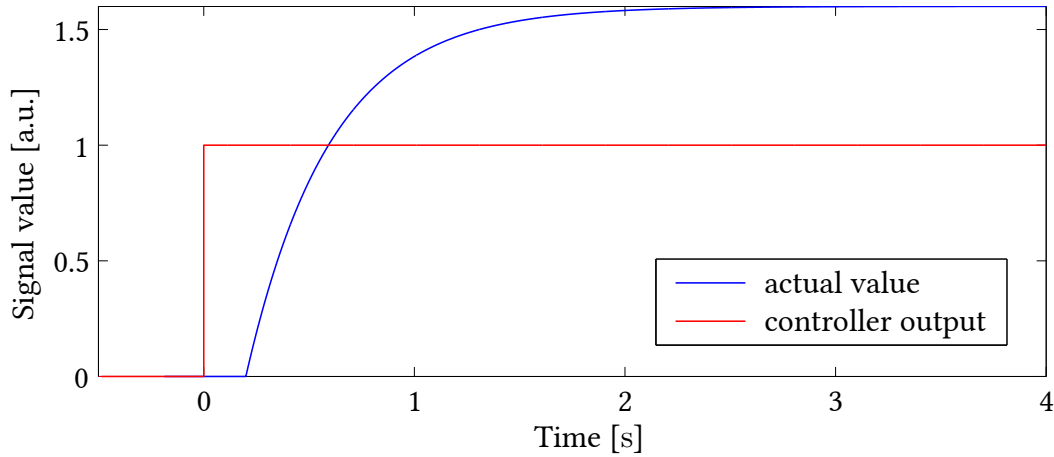
Figure 2.4: System plant response modelled using equation (2.52). One can see the dead time and the of $T_D = 0.2\,\mathrm{s}$ and the then slow system response of $T_L = 0.4\,\mathrm{s}$. These two values limit the acctual speed any regulation process on this system can stably achieve.

This function describes the frequency s dependent reaction of the system $a(t)$ to the controller output $u(t)$.

To start modelling this system-plant transfer function, one is introducing a conversion factor A that describes the amplitude change from system plant input to output signal:

$$a(t) = A \cdot u(t) \tag{2.48}$$

$$\Rightarrow H(s) = \frac{\mathcal{L}\{a(t)\}}{\mathcal{L}\{u(t)\}} = A \tag{2.49}$$

This is determined for example by the gear conversion of engine rotations to the car speed, this is however very simplified.

In a car the motor will only slowly change its rotations per minute, since the combustion processes needs to adapt to the new air-fuel mixture. This slow reaction, translated to the frequency domain, can be described as a low-pass filter, so that fast changes (large s) are attenuated or completely blocked. The transfer function of a low pass L with a time constant T_L is given by:

$$L(s) = \frac{1}{1 + s \cdot T_L} \tag{2.50}$$

combined with (2.49): $\Rightarrow H(s) = \dfrac{A}{1 + s \cdot T_L}$ $\tag{2.51}$

Most systems also have a dead time T_D which describes the time in which, despite changes in the input, the output doesn't change at all. Adding this into the description
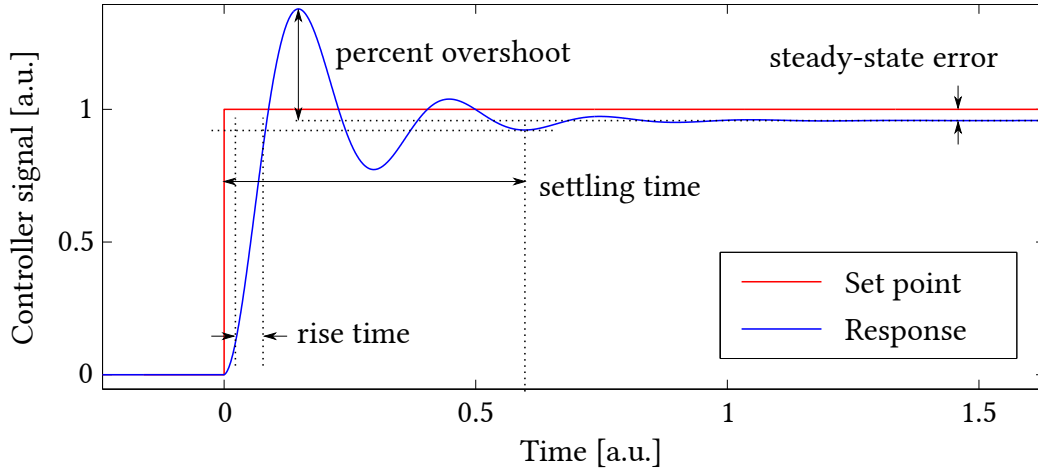
Figure 2.5: Quantised description of a control-system response. Rise time is the duration between 10 and 90 % of the set point signal whereas the settling time is reached when the signal stays between 5 % of the final value.

of H we get:

$$H(s) = e^{-T_D \cdot s} \cdot \frac{A}{1 + s \cdot T_L} \tag{2.52}$$

$$\approx \frac{1}{1 + s \cdot T_D} \cdot \frac{A}{1 + s \cdot T_L} \tag{2.53}$$

Here we used an approximation with the tailor expansion for a very small T_D. Figure 2.4 shows this system behaviour for certain values in real time. This very simple model can describe the behaviour of most systems used in the lab. With this one can already see that any regulation is limited in it's achievable speed by the dead time T_D and the delay time T_L. Since they give the time scale for the regulation process, one wants to always have a fast reaction, e. g. low T_L and only very little, ideally no dead time T_D. For the next sections we will however not include the dead time and only use equation (2.51) for simplicity.

Regulatory Behaviour

Adding the system-plant response and the controller transfer function into a closed loop to model the complete regulatory behaviour is done in appendix A.1. We will use the final result derived there for the closed loop transfer function and model the influences of different controller behaviour.

As already mentioned previously, these gain settings of our controller K_P, K_I and K_D have to be adjusted to the system and its frequency dependent behaviour. The values defined in figure 2.5 are used to characterise the control loop system while changing the set point in a step function. An ideal adjustment would have fast settling and rise time, with nearly no overshoot and a steady state error of zero. Table 2.1 shows the influence
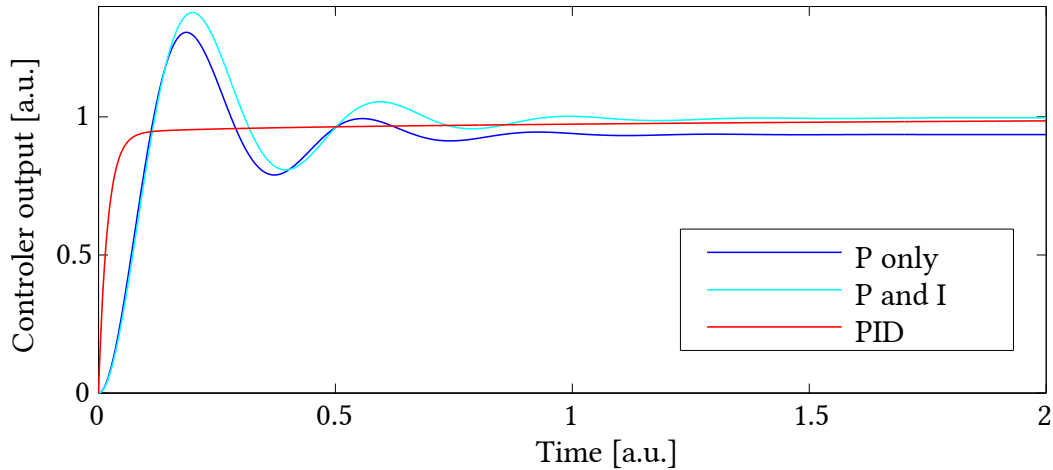
Figure 2.6: Regulatory behaviour of a PID controlled system adjusted with the Ziegler-Nichols method, see table 2.2. One can see that including and optimising all components of the PID regulation leads to a very sharp response with little oscillations.

of increasing the different gain variables in the controller on the system dynamics [41]. Since in most cases K_P, K_I and K_D are dependent on each other this can only be a reference for the determination of the controller parameters.

Table 2.1: Idealised change after optimising the gain of different controller parts individually with the values defined in figure 2.5.

gain value	rise time	overshoot	settling time	steady-state error
K_P	decrease	increase	small influence	decrease
K_I	decrease	increase	increase	eliminate
K_D	small influence	decrease	decrease	no change

As said previously, figure 2.5 doesn't include a dead time, which is the delay of the change in set point to the first change in the system observable.

A very common optimisation scheme of the whole regulated system including controller and plant is the Ziegler-Nichols tuning [41]. For that one switches off any I and D part, using only the P regulation and increasing K_P to the point that it's regulation response is a continuous oscillation, which is called it's critical behaviour. The oscillation period P_C and the gain value $K_C = K_P(\text{critical})$ is then used in table 2.2 to determine good values for a stable regulation, see equation (2.46):

Table 2.2: Gain values for P, PI and PID regulations according to the Ziegler-Nichols method. Using the critical gain K_P and the critical oscillation period P_C.

Control	K_P	T_I	T_D
P	0.5 K_C	-	-
PI	0.45 K_C	P_C / 1.2	-
PID	0.6 K_C	0.5 P_C	P_C / 8

A typical behaviour of a such adjusted regulation with the three possible options of control is shown in figure 2.6.

All this discussion however does not include that there are disturbances introduced to the system outside the control loop. Since the Ziegler-Nichols method we presented here optimises the closed loop behaviour, its values might not be ideal for a disturbed system. A very relevant example for these disturbances are the background oscillations in the magnetic field inside a field stabilised system. Nevertheless they provide very good initial values, which then can be optimised to the behaviour desired. This method was used to find initial control values for the PID controllers in chapter 4.

3 Improvement on the Imaging System

During the work of this thesis it soon became clear that, in order to be able to detect small changes in the population of the Lithium magnetic sub-levels, the imaging had to be optimised. The most suited, non-invasive approach was to re-evaluate the acquisition software and its settings. In this chapter we show the changes made to the existing set-up and introduce the self written camera software. Finally we show how we adjusted and optimised the settings for an improved imaging of Lithium, which allowed us to better observe small population changes in magnetic field sensitive processes.

3.1 Absorption Imaging

The method used for imaging atoms in our experiment is called Absorption Imaging and has been used in the field of ultracold atoms for a long time [11, 19].
It uses the fact that the intensity I of light passing through a cloud of particles with density n is reduced due to the scattering of photons through excitation [21]:

$$\text{therefore:} \quad \frac{\mathrm{d}I(z)}{\mathrm{d}z} = -\hbar\omega_{\mathrm{L}}\gamma_{\mathrm{P}}n \tag{3.1}$$

Here we used the excitation rate γ_P from equation (2.33), scaling with the affected atom number N passed through with $n = N/V$. For near resonant light and low intensities the scattering rate becomes $\gamma_P \approx s_0\gamma/2$ and we can describe this intensity loss as:

$$\frac{\mathrm{d}I}{\mathrm{d}z} = -\sigma_{eg}n(z)I(z) \tag{3.2}$$

Here we identified the cross section σ_{eg} to be [21]:

$$\sigma_{eg} = \frac{\hbar\omega\gamma}{2I_S} = \frac{3\lambda^2}{2\pi} \quad \text{for:} \quad \delta \approx 0; I \ll I_S \tag{3.3}$$

Outside these conditions, the cross-section is dependent on intensity I and the detuning δ to the resonance frequency. We use the approximation that the light beam intensity and atom density is uniform in an area A orthogonal to the beam direction, which might then be projected on a camera pixel. Therefore you can calculate the total number of atoms N in the volume $V = A \cdot z$ by $N = \int_A \int_0^z n(z', x', y')\mathrm{d}z'\mathrm{d}y'\mathrm{d}x' = A \cdot \int_0^z n(z')\mathrm{d}z'$.
Given an initial light intensity $I_0 = I(0)$ we can now calculate the intensity at a certain position z:

$$I(z) = I_0\exp\left(-\sigma_{eg}\int_0^z n(z')\mathrm{d}z'\right) = I_0\exp\left(-\frac{\sigma_{eg}}{A}N(z)\right) \tag{3.4}$$

With known final and initial intensity I_f, I_0 and σ_{eg} we can in turn calculate the number of atoms the light has passed through until detection:

$$N = \frac{A}{\sigma_{eg}} \log \frac{I_0}{I_f} \tag{3.5}$$

This calculation fails of course if no light is passing through the atom cloud $I_f = 0$. This is then called *optically dense* and can be either compensated with an increase in initial light intensity or a reduction in density. We will come back to this problem later in the chapter.

But it is also already harder to calculate σ_{eg} if the intensity I changes too much over the passage or the conditions for equation (3.3) like $\delta \approx 0$ are not met any more. This is the reason why we stated in section 2.1 that for imaging it would be ideal to stay resonant over the interaction time. One also has to keep in mind that the initial approximation used to obtain equation (3.5) was the low intensity limit $I \ll I_S$ and is not valid for high imaging intensities any more.

In our experiment a collimated, monochromatic beam with a separate frequency control and a waist much bigger than the atom cloud size is shined on the atoms. The intensity profile after absorption is then magnified through a lens stack and onto a CCD camera. CCD stands for charge coupling device [7] and it's working principles will be explained later in this chapter. This gives the final Intensity $I_{f,p}$ for each pixel p. We will call this the *atom picture*. According to figure 2.2 and equation (2.36) the atoms then get accelerated so much that they fly out of the focal plane and can't be detected anymore. After a short period of time a second picture is taken with another laser pulse. This picture will be used to determine the initial intensity $I_{0,p}$ and we will call it the *reference picture*. To extract pixel errors and dark currents a third picture without any laser light is taken at the end of each sequence.

As a naming convention we will describe the digital pictures by assigning each pixel p a value V_p.

Since both pictures are separated in time we can't ignore overall intensity fluctuations. This issue is addressed by selecting a *region of interest* (ROI) on the picture which includes the atom cloud and a considerable surrounding volume. Since we can assume that no absorption happens outside of the ROI we can use this area to normalize the ROI part of the picture. For that every pixel value of the atom picture $V_{a,p}$ outside the ROI is summed up and divided by the sum of the pixel values of the reference picture $V_{r,p}$ to get the normalisation factor f_n:

$$f_n = \sum_{p \in \text{ROI}} V_{a,p} \left(\sum_{p \in \text{ROI}} V_{r,p} \right)^{-1} \tag{3.6}$$

$$V'_{r,p} = f_n \cdot V_{r,p}, \text{ for all } p \notin \text{ROI} \tag{3.7}$$

Now equation (3.5) can now be used for each pixel separately if σ_{eg} and A is known. It has to be made sure however that the assumptions for this calculation, e. g. the saturation parameter s, the detuning δ, etc. are met and it's values are know sufficiently. For more information on this specific imaging algorithm one can read [20, 37].

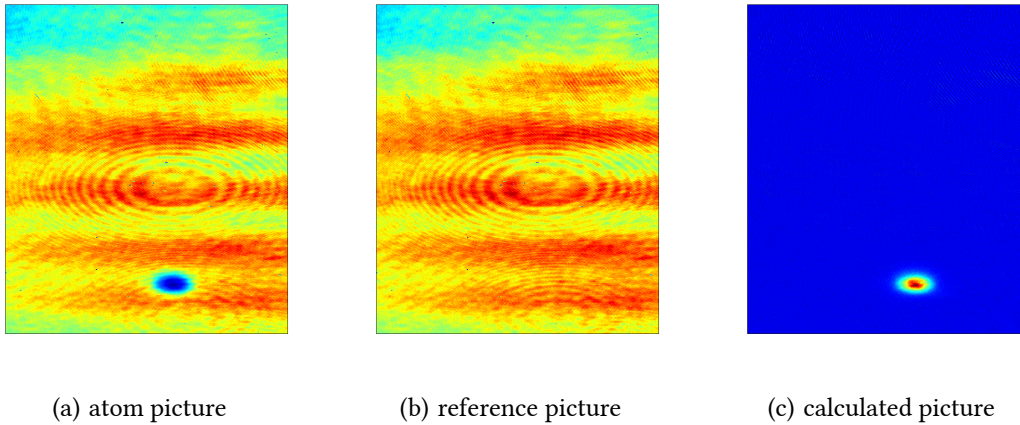| (a) atom picture | (b) reference picture | (c) calculated picture |

Figure 3.1: Calculation of the atomic distribution out of the atom absorption picture (a) and the reference picture (b) without atoms. By calculating the ratio of the initial to the final, absorbed intensity and then taking the logarithm one can calculate the picture (c) that shows the distributions of atoms in the image plane. This particular picture shows a BEC of sodium. Including several scaling factors we can use it as a measurement tool for exact atom numbers.

Limitations and Noise Sources

When used correctly, this method is very robust and successfully used in many experiments. There are however some limitations and issues with this technique.

- Absorption imaging is a destructive measurement, meaning after taking a picture the atoms will be accelerated out of the trapping and imaging region. Any state prepared is therefore destroyed.

- It is crucial for this method that the laser's intensity profile stays the same for the atom and the reference picture. Any deviation will lead to a false calculated atom number.
  The main reasons for shifts or changes in intensity are the mechanical motion of the optical table or the imaging optics, frequency drifts of the imaging laser and air motion (e. g. sound waves) in the laser path.

There are some ways to reduce the intensity fluctuations, like air damping of all optical tables, housing of optics to reduce air flow and very stable laser locks.
Another very effective way is to take atom and reference picture with an extremely small delay between each other. This is limited by the capabilities of the camera and it's control software.
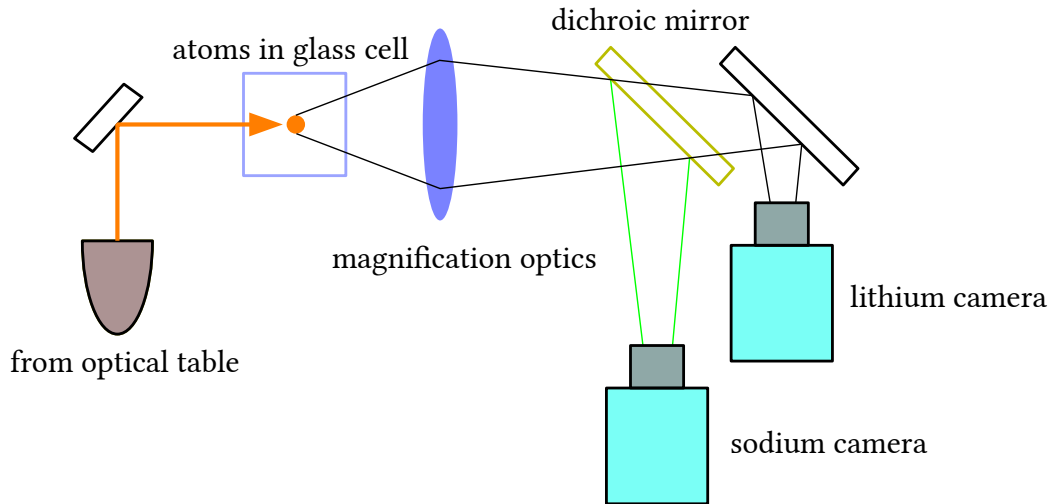
Figure 3.2: Schematic the absorption imaging system. The same optical fibre is used for sodium and lithium light. The light profile shortly after the absorbing atoms is then projected with a lens stack (for simplicity replaced with a lens symbol) onto the CCD chip of the cameras to give a sharp absorption picture. A dichroic mirror is used to separate the different wavelengths for the sodium and lithium camera.

3.2 Hardware and Configuration

The whole imaging as a system can be described in three separate parts:

1. Laser preparation and collimation on the atoms

2. Magnification optics and camera

3. Acquisition software and experimental control

In the following we will give a short description of the parts necessary in understanding the contents of this thesis and the changes done to it. For information on the hardware set up one can read [20, 37, 40, 38].

Imaging Laser and Magnification Optics

We use a dye laser for the creation of 589 nm light for sodium and a diode laser with a tapered amplifier for 670 nm light for lithium. Both lasers are locked to transition lines in the relative spectra and frequency adjusted with acoustio-optical modulators (AOMs). Therefore we can stably tune their frequency in the range of a few MHz close to the resonance of the atoms. After this preparation the two imaging laser beams are joined via a dichroic mirror and then coupled into an optical fibre. The laser system in our experiment is described in more detail in [40] and [38].

The basic scheme of the imaging next to the vacuum cell is shown in figure 3.2. After a

fibre out-coupler the imaging beam is sent onto the square glass cell onto the atoms. This beam then gets magnified by an objective from Zeiss. This lens system is designed for microscopy with a working distance of 87 mm and a focal length of 100 mm. This results in a projection magnification of 2.6x. The two wavelength components then get separated over a dichroic mirror which reflects the sodium light and transmits the lithium light. Both cameras are placed and adjusted to be able to sharply image the atomic cloud. Additionally to the initial placement the lithium camera is mounted on an electronically controlled translation stage and can therefore be moved along the beam path to adjust the focus on the atoms.

The CCD Cameras

We use two Retiga EXi fast (non-cooled) from the company QImaging. A conventional CCD camera consists of an two dimensional array of photosensitive, isolated diodes called pixels. Each incoming photon has the possibility to create an electron hole pair in the semiconductor depending on the wavelength of the photon and the material used in the chip. The probability to create a mobile charge carrier, the electron, upon photon impact is called *quantum efficiency*. The electrons created during the exposure of a pixel directly get pulled by an applied voltage into a deeper lying storage region of the semiconductor creating a potential well. Usually this well has a certain maximal capacity to store electrons, which is called the *linear full well capacity*. If more electrons get created the pixel reaches saturation and is not sensitive to more photons anymore [7].

There are several ways to read out the amount of electrons stored in each pixel, but we will only describe the method used in our camera: in a so called *progressive-scan interline* CCD chip each pixel has its own storage region next to the photosensitive semiconductor area. This is done by adding a small opaque mask on the pixels, where light cannot pass through. They are situated between the lines of pixels, hence the word *interline*. The electrons accumulated on the pixels can then be shifted by applying another voltage to the non sensitive region of the pixel. This shift is done in parallel for all pixels at once and usually works in the time scale of some microseconds. The pixels then get read out pixel by pixel. This happens by shifting the electron packages to a readout area, where they are then detected and converted to a digital value by a *analogue-to-digital converter* (ADC) [18]. Important characteristics for a CCD camera are therefore the quantum efficiency or the sensitivity in the relevant wavelength and the capabilities of the ADC in terms of noise and digital output range. The ADC in our camera translates charges into a 12 bit value, and is therefore distinguishing between 4096 digital values. The conversion factor of electrons to a digital count is adjusted during production to cover the full dynamical range. In table 3.1 we list the most important characteristics of the Retiga EXi fast [24].

One of the camera was previously used in the rubidium BEC experiment of our group and was characterised there [23]. Out of his work we know that with standard gain, the conversion factor of ADC is around 7.29 electrons for one digital count. Combined with the quantum efficiencies for the wavelength we can deduce that one digital count of a pixel corresponds to roughly 12 photons for sodium imaging light and 16 photons for

Table 3.1: Characteristics of the cameras Retiga EXi fast from QImaging, which we use in our experiment.

CCD sensor	Sony ICX285 (monochrome)
Pixels	1392×1040 (1.4 million)
Quantum efficiency	$\sim 60\%$ @ $589\,\mathrm{nm}$ (sodium)
	$\sim 45\%$ @ $670\,\mathrm{nm}$ (lithium)
Pixel size	$6.45\,\mathrm{\mu m} \times 6.45\,\mathrm{\mu m}$
Linear full well	$18,000$ electrons
Digital output	12 bit
Frame rate	10 fps full resolution
Transfer data rate	$40\,\mathrm{MB/s}$ (FireWire)

lithium light. We will however introduce a tool to change this factor later in this chapter. The readout of the camera was done with a proprietary software called QCapture Pro 6, which runs on Windows XP and uses a hardware key dongle to protect against software piracy. Only one software-dongle combination can be used on a computer at the same time. Therefore two separate computers were needed to control the cameras.

The manufacturer also provides a software development kit for the programming languages C\C++ and C#. This enabled us to write a new application, communicating with the camera on a native level.

Experimental Control and Image Acquisition

The camera can be triggered by an external TTL pulse in two modes: by opening the camera shutter while the trigger signal is high and then sending the picture on the falling edge of the signal (*hardware exposure*) or by triggering on the rising edge and opening during a software defined exposure time (*software exposure*). The trigger signal sequence shown in figure 3.3 of the experimental control allows for a switch in these modes without bigger changes:

At the set time of the image capture the TTL trigger goes high and therefore the camera shutter is opened. $125\,\mathrm{\mu s}$ later the imaging laser is turned on for another $125\,\mathrm{\mu s}$. This is done to avoid the atoms being already accelerated out of resonance without the camera being able to capture the absorption properly. $125\,\mathrm{\mu s}$ later the TTL trigger goes low again. The $125\,\mathrm{\mu s}$ step duration is set by our control hardware. If hardware exposure is enabled the chip is exposed over $375\,\mathrm{\mu s}$ with the full $125\,\mathrm{\mu s}$ of the laser pulse. With software exposure one can define a value of the exposure time which then starts from the high edge of the TTL trigger and can also capture only a part of the pulse. The settings we have chosen for Lithium will be further discussed in the next section. In the experimental control, the time span between image capture and turning off of the crossed dipole trap is set. During that time the atoms can expand partially freely, so it is called *time of flight* (TOF). We can choose this timespan individually for Sodium and Lithium. The TOF can be used to reduce the density of our atom cloud by increasing
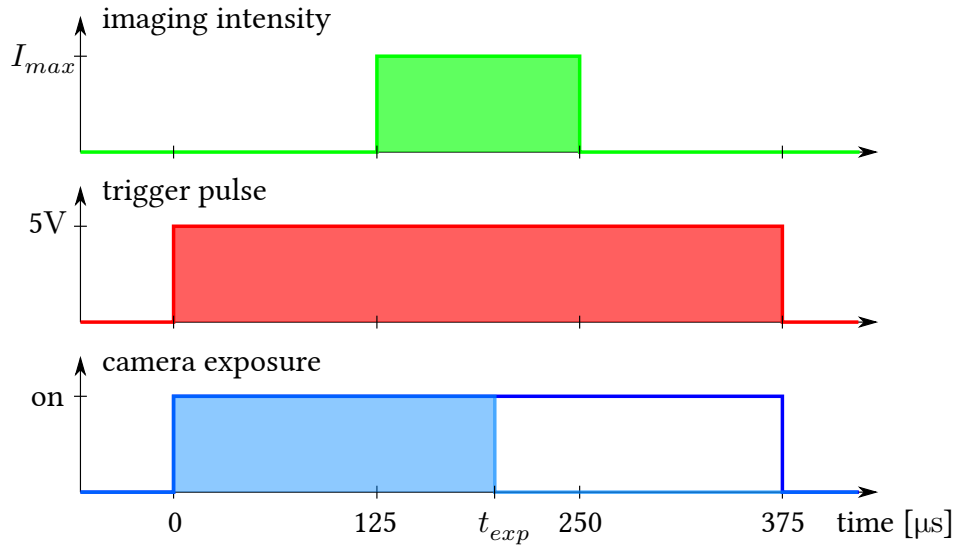
Figure 3.3: Imaging sequence of one species. At time zero the imaging process starts, with the trigger pulse of to the camera going high and remaining high for the next $375\,\mu s$ (red shaded area). The light is only shined on the atoms $125\,\mu s$ later for a period of $125\,\mu s$ (green shaded area). In the previous settings the camera shutter was open for the whole time the trigger pulse (dark blue line). By using software exposure one can freely choose the time of exposure t_{exp} (light blue shaded area).

the size. As explained in section 3.1 this can enable better imaging and is also used in the separation of different magnetic sub-states via a Stern-Gerlach pulse [37]. A typical expansion time for imaging a Sodium BEC is $15\,ms$ and 3 to $5\,ms$ for Lithium.

The reference picture is taken after a certain delay time, which was previously $650\,ms$, which ensured that the camera is ready again and the readout of the old picture has been completed. The dark picture is then taken at the end of the whole sequence, around 2 seconds later, with no light on. The trigger and exposure sequence stays the same for all three images, with of course the light pulse only being present in the first two.

After the end of each experimental run three Sodium and three Lithium pictures are sent to another computer to run the imaging algorithm using equation (3.5), calculating the pictures with calibrated atom counts per pixel.

3.3 New Software and Graphical User Interface

In the course of this thesis a new program was written to replace the commercial software using the manufacturer supplied software development kit (SDK). It is aimed to address several issues and questions:

- The software previously used could only ran on an outdated operating system (OS), so two different OS had to be used in the lab.

- Two computers with the task of reading out only one cameras were necessary. If possible we wanted to reduce the amount of hardware needed for the experiment.

- The commercial software was not easy to use for our special application and many features were hard to access. It was for example optimised to do post processing of the images with filters and smoothing functions, which were not needed in our case.

- The time delay between atom and reference picture was quite long which can lead to intensity fluctuations, see section 3.1. With good knowledge on the low level software side and its adjustment settings one could reach a regime where fringes on the calculated image due to air movement or frequency drifts could be minimised.

Due to previous experience and known support for the future it was decided to write the program in C# using the .NET 4.0 Framework. The graphical user interface was constructed with the Windows Presentation Foundation framework and written in XAML. Microsoft supplies a free integrated development environment (IDE) with it's community version of Visual Studio. Since this IDE is downward compatible and free to use, future changes to the program can be made very easily [10].
In the next section we want to explain operation and the functions of the camera, while not going into every detail of the program code. It is well documented and should be easy to understand.

System architecture

C# is a object oriented programming language, therefore also the internal structure of the program can be understood and described with abstract software objects. These can store information, interact with each other and control hardware connected to the computer, e. g. a camera. Figure 3.4 shows their interaction possibilities and basic functions.
 The *camera object* is the first thing to be initialised by the software on the PC. It includes a collection of all relevant information on the camera and used for all communication with it. In a first step one creates this object and searches for a physical camera connected via the FireWire bus. If one is found it can initialize itself with the information taken from the real, physical camera and is now used for all communication with it. All orders and changes are sent to this software camera object which handles the communication in the background. It can handle for example a *setting object*, a list of all camera settings. It can be adapted to our needs inside the software, and then sent to the camera via the camera object. This settings object also needs to be initialised with information taken from the camera. To let the camera take a picture, one has to initialize a *picture object* and allocate space on the main memory (*picture buffer*) in the exact size of the picture one would like to take. Then a reference to the picture object gets handed over to the camera object, which sends an order to the camera to take a picture with the current settings.
In the standard implementation the picture object is sent to the camera object and waits
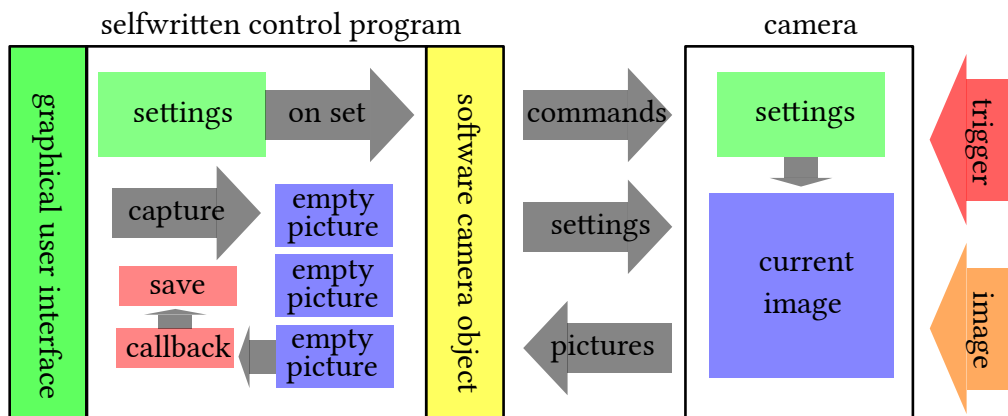
Figure 3.4: Scheme of the software that controls the camera. The user interacts over the graphical user interface (GUI), whereas a camera software object is used as an internal interface for the communication and control of the camera. The settings one chooses over the GUI are only sent to the camera after a separate order *Set Settings* is executed and are then stored in the camera. This includes the settings for the camera trigger. With the order *Start Capture* on the GUI one creates a queue of three picture objects, that have their own memory allocated. These three picture objects then order the camera consecutively to take pictures depending on the settings currently set. These pictures are then filled with the pictures the camera captures in a parallel thread, not blocking the rest of the software. A separate callback function is called after the third picture object is filled with an image, which then saves all three pictures on the hard drive for further processing. The process of capture then can be repeated by again putting three images into the queue.

until the camera sends the picture back, writes it on the pre-allocated memory and confirms completion of the job. This means the whole program is frozen during that time and can't respond nor communicate with the camera further (e. g. asking for a second picture to take). If one wants to take several pictures in a fast succession this could cause a serious delay, or in the worst case the missing of a triggered picture.

To avoid that, one use a parallel computing approach, creating different computation threads inside the same program. By using a *picture queue* connected to a camera object, where several picture objects are put into, we can start a different thread which automatically sends one picture after the other to the camera. With their own already allocated memory they get sent back from the camera on completion, signalling the next picture to be taken. This doesn't block the software, since the main computational thread doesn't have to wait until completion, but just opens a *callback function* when the picture capture is complete. This callback function runs on another parallel thread in the software and can be used to determine further processing of the picture. It is even possible to reuse the picture object for example after saving the picture to the hard drive by sending it back into the queue.

This use of parallel computing can greatly increase the possible frequency of image captures and is therefore deployed in the new software. One has to consider though that sending a picture request to the camera doesn't indicate an immediate capture. The picture is only taken and sent back to the PC once the camera was triggered. We will later come to the different trigger types used in our experiment.

Since we want to take three pictures (atom, reference and dark picture) within each camera a queue with three pictures is formed. The first two will run an empty callback function once completed and don't disturb the software further. The third picture's callback function however initialises the processing of the three, now captured images. A temporary image is created from each picture in the buffer and used in a software display for observation. The pictures then are saved in a predefined folder, where they will be fetched by the evaluation software. Once this callback is completed, the three picture buffers are deleted, reset and the queue is filled up with the three picture objects for the next run.

User Interface and Possible Settings

The graphical user interface (GUI) of the software has to fulfil several purposes: it should control and initialize the camera-PC interaction, set settings on the camera, manage capture and storage of pictures and also display the pictures taken to be able to directly check if the sequence worked and give a picture preview. A picture of the GUI is shown and described in figure 3.5. In the following the main functions and their underlying mechanisms will be described. Checkboxes are used to activate specific settings, if those are unchecked a standard value will be chosen.

The left side is used for adjusting settings and controlling the camera, whereas the right side is used for a preview of the images taken.

a) To use the camera one has to first initialize the camera by clicking the *Initialize*
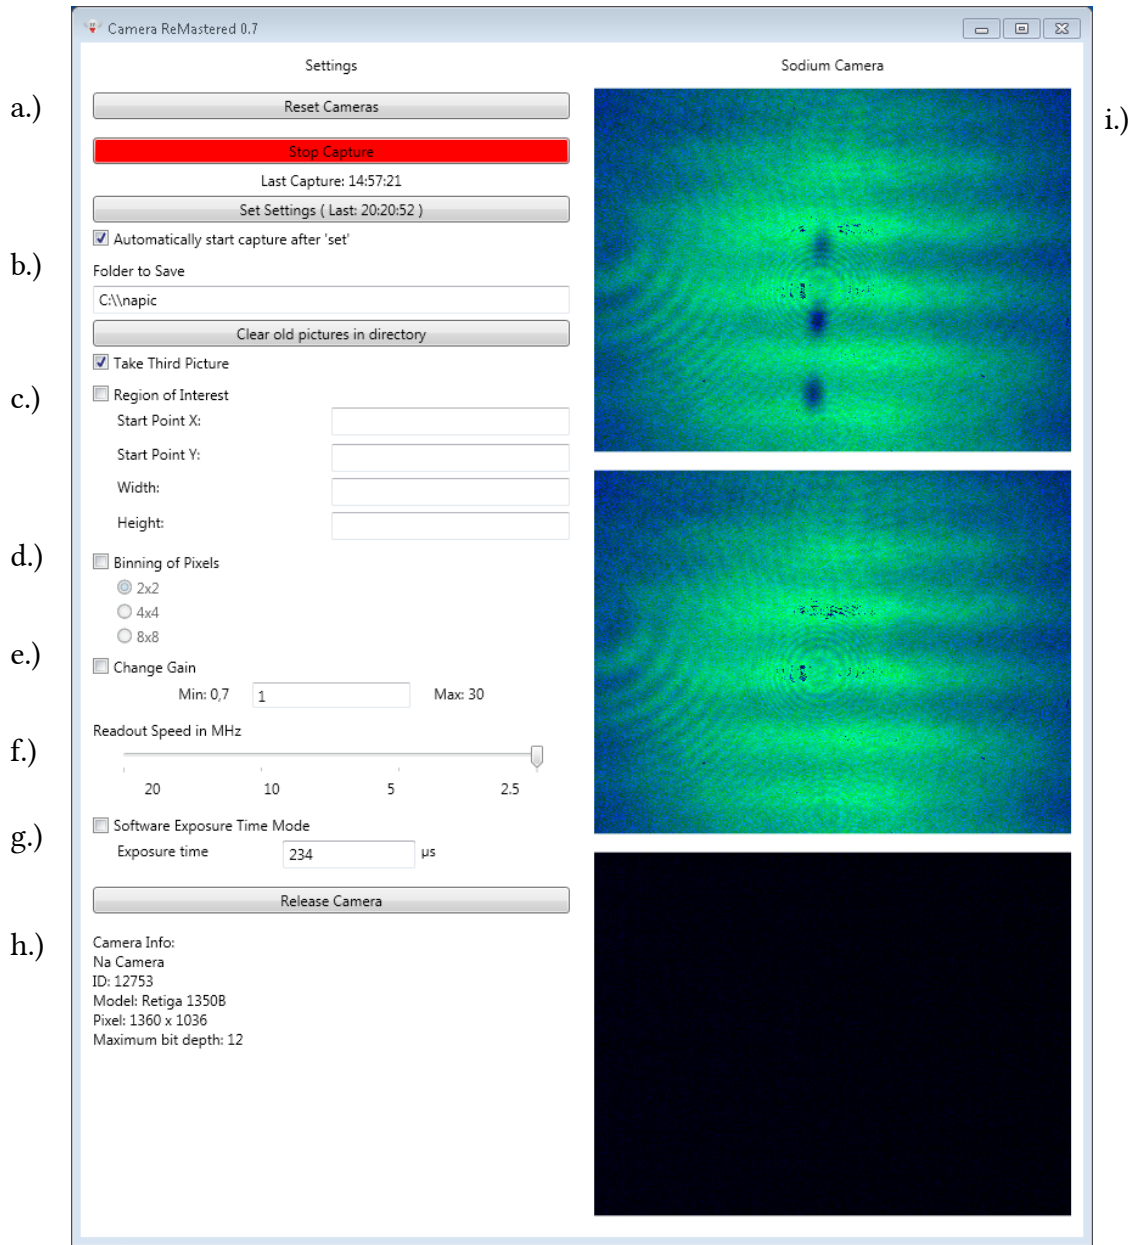
Figure 3.5: The graphical user interface (GUI) of the camera software written during this thesis. It is used to initialize and control the camera, set settings and preview the pictures taken. a.) Buttons to initialize and reset camera, start and stop the capture and set settings. b.) Folder selection for saved pictures and clear folder button c.) Region of interest, d.) Binning and e.) Gainsettings f.) readout speed slider g.) Software exposure settings h.) Information on the camera connected, including serial number to identify Sodium and Lithium camera i.) preview of last pictures taken, from top to bottom: Atom picture, reference picture, dark picture. On the atom picture one can see the absorption of three small clouds, which are the different magnetic sub-states of the F=1 ground state manifold of Sodium after a Stern-Gerlach separation The different fields are explained in more detail in the text.

Camera button, which then changes to *Reset Camera*. An error message box pops up if no camera is connected or the last connection crashed. In this case one should check the connection cable and turn off and on the camera manually. It then reads out the cameras internal information, to identify the camera as the Sodium or the Lithium camera. According to that information it automatically loads the last saved settings for that camera and fills it into the input fields of the GUI, like the save path (see b)). It also sets the name on top of the right side (see i)).

With the *Set Settings* button one can send the current settings that are seen on the GUI to the camera. Once this is done for the first time the *Start Capture* is put out of it's initial disabled status and can be clicked. It starts the preparation of the picture buffers and puts the picture objects into a queue. The active capture is indicated by the button turning into a red *Stop Capture* button. A new set settings call will disable the capture again, but there is the option of automatically starting the capture procedure each time one has set new settings. This is done by a checkbox below *Set Settings*.

b) After the completion of the last picture queued it's callback function will save all pictures into a specified save folder, which are accessible from the network to fetch the pictures. One has to take care though to write the folder path with a double backslash (\\) rather than with a single one. The Matlab script on the evaluation PC will take the last three pictures of in the folder, regardless of the content. So there is a button to delete the content of this folder in order to make sure the script doesn't mess up the acquisition of the newest pictures. There is also the possibility to not take three, but two pictures by omitting the dark picture. To do so one has to uncheck the *Take Third Picture* check box.

c) Region of interest (ROI) is a hardware limitation on the image size and should not be confused with the region of interest used for rescaling. The hardware ROI is commanding the camera to only read out a rectangle with a size described by *Width* and *Hidth* starting at the point defined by *Start Point Y* and *Start Point X*. This can reduce the image size and therefore improve the transfer speed from camera to the PC memory. One needs to keep in mind that the rescaling algorithm of equation (3.7) also uses a different ROI, which again decreases the final computed image size. The achievable size reduction therefore is limited, but can still be significant.

d) One other possibility to reduce the image size is to reduce the resolution of the image. Usually this is done software wise through after processing on the PC, but this doesn't increase the transfer time, since the whole picture still has to be sent over the limiting FireWire connection. The binning option of this camera however reduces the resolution on the hardware side of the capture process. Instead of reading out each single pixel of the CCD, the camera can electronically connect several pixels and read their electron count out as if coming from only one pixel. This super-pixel is always a square and its size can be chosen with the radio buttons in the GUI between 2×2, 4×4 and 8×8 pixels. This leads to an addition of the

values of the connected pixels, which can be seen as an increased gain factor (4-, 16-, or 64-fold respectively) to the pixel's image count. It also reduces the number of readout procedures per image, which usually introduce a lot of noise. This can also be described as a higher light sensitivity leading to better signal to noise. With a 4×4 binning set, this technique reduces the resolution of a full image from 1392×1040 pixels (file size 2.76 MB) to 348×260 pixels (file size of 0.086 MB). When using the ROI setting (see item c)) one has to adapt the values that describe the chosen rectangle on the image: Since the ROI size is defined through pixels one has to think of the picture after the binning, e. g. change *Starting Point X* from 400 (no binning) to 100 (4×4 binning). Binning also acts as a averaging filter compared to the full resolution picture, so the statistical noise is reduced. However this is no real advantage, since this can also be done post processing.
One of course has to weight the advantages of the increase in transfer speed to the big loss of information due to the reduced resolution of the picture.

e) Adjusting the gain of the camera sets the conversion factor from electrons read from the CCD chip (see section 3.2) to the 12-bit digital value assigned to the pixel in the Analogue-to-Digital-Converter (ADC). Electronic saturation on a pixel is reached with a count of around 18,000 electrons (Linear Full Well). The factory settings g_0 are adjusted to enable the best full dynamical range of 12 bit (0 to 4095) below this value. With gain g set to 1 (standard value) this results in a conversion of $g_0 = 4.5$ electrons to one digital integer step. Every other gain factor g changes this conversion accordingly:

$$N_{e^-} g_0 \cdot g = V_{\text{digital}} \tag{3.8}$$

Here N_{e^-} denotes the electrons sent to the ADC from one pixel and V_{digital} the digital pixel value saved in the image.

The readout noise in the ADC is not only dependent on the incoming electron number N_{e^-} but also on an internal intrinsic read noise. So one can reduce the relative noise by adjusting the gain of the signal to use the full dynamical range, therefore increasing the signal to noise ratio. For an exact atom number calculation, with the absorption imaging technique described earlier, one has to include this factor as well.

f) The camera's readout speed $f_{readout}$ defines how fast the set of pixels is read from the CCD chip. It will take a time t_{read} to read out a certain number of pixels N_{pixel}.

$$t_{read} = \frac{N_{pixel}}{f_{readout}} \tag{3.9}$$

The camera is capable of using four different readout speed settings from $2.5\,\text{MHz}$ to $20\,\text{MHz}$. For a full image with $1392 \times 1040 \approx 1.45 \times 10^6$ pixels the readout process will take $579\,\text{ms}$ in slowest and $72\,\text{ms}$ in the fastest settings. A faster readout speed generally leads to an increase in readout noise, depending on the capabilities of the ADC. In our case this noise is about 8 electrons per readout and not separately specified for different speeds.

g) As already described in section 3.2 the camera can switch between a hardware and software exposure time mode. In both a TTL signal sets the start of capture, but the duration of exposure is either set by signal high time (hardware exposure) or an internally set time (software exposure). The exposure time for the second mode can be set in microsecond steps, but one has to keep in mind that the imaging light pulse arrives only $125\,\mu s$ after the trigger signal (see figure 3.3).

h) On shut down the program will call a routine that dissolves the computer internal binding of the camera, e. g. the camera software object. If this is not done correctly, future communication attempts with the camera will fail, since it is still bound to a now non-existing structure on the PC. For debugging reasons one can do this manually by using the *Release Camera* button. However closing the program and reopening it will have the same effect.
On every initialisation of a communication with the camera some basic information will be asked and displayed at the bottom left of the GUI. The serial number is displayed next to the tag *ID* and is then used to distinguish between the lithium and the sodium camera. It also loads accordingly the last saved settings of that camera and puts its values in the GUI.

i) If the program can identify the serial number of the camera, it displays the name on top, e. g. *Sodium Camera*. The rest of the GUI's right side is used for previewing all three different pictures taken in one sequence.
The topmost picture is the atom picture, where in the shot shown one can see the absorption of three separate atom clouds. This is done by using a Stern-Gerlach sequence to separate the different magnetic sub-states of the sodium hyperfine-splitting [37]. In the middle the reference picture without atoms is displayed and on the bottom the dark picture without any light is shown.
To display a 12 bit greyscale picture one needs at least a 16 bit greyscale image format. All files are therefore saved in this format in a TIFF file. However there are some issues with the .NET 4.0 framework this program is working in and displaying the 16 bit greyscale format. Therefore we used a red-green-blue format which uses only 5 bits per colour channel. The colour scale of all pictures is fixed! Pixels that have values with only 10 bit values (e. g. 0 to 1023) only use blue and green colours. In an area with higher intensity, where the pixel values go over 1023 one can see a separation line. There only one bit of the red channel is added and the blue and green channel values start from zero again. It is no artefact of the camera nor its image, but only of the colour-scale used. Since the purpose of these images in the GUI are for preview only this was deemed acceptable.

We chose the settings implemented in the software according to our specific set-up and with the intention on improving the signal-to-noise ratio of our imaging. There are other settings the software development kit can access we don't need, like post processing options or the addition of a global offset. Others are connected to colour imaging, which is not possible with our cameras.

To address both cameras at once, they were connect to the same PC, which ran two instances of the same software. On initialisation it will detect the cameras not yet bound and try to connect to first the lithium and then the sodium camera. If there are any problems at the initialisation, e. g. only one camera is recognised, one should close all software instances and manually restart both camera. The cause for not recognizing a camera is usually not correctly terminated software instances or the removal of a FireWire cable during operation.

3.4 Adjustment and Optimisation of Lithium Imaging

We wanted to optimise our imaging to the best extend as possible with the software settings described in the previous section. Here we will only focus on the images calculated with the algorithm described in section 3.1. For this we use the two characteristics described in the following:

Noise on Empty Regions

For the definition of this quantity, a region R of our picture was chosen with no atoms in it. Given a correctly adjusted rescaling algorithm and no noise sources, one would expect an even distribution over the whole picture with value zero, which would correspond to zero absorbing atoms in this region. However this is not the case, since statistical noise and other systematics have to be accounted for. We define the noise σ_R in an imaging region as the deviation of the measured value V_p from the expected value $V_{exp,p} = 0$ with the p denoting one of N pixels in the region R, we examine:

$$\sigma_R^2 = \sum_{p \in R} \left(\frac{V_p - V_{exp,p}}{N-1} \right)^2 \tag{3.10}$$

This definition is also valid for an atomic distribution in our picture, however it is hard to know the exact distribution of the atoms beforehand. Therefore we will focus on examining this characteristic in empty regions, knowing the noise there is directly correlated to the regions where atoms are.

We discussed already most of the possible systematic errors of the absorption imaging method in section 3.1. There we stated that shortening the delay between atom and reference picture should reduce any deviations in the calculated image. There are however some noise sources only dependent on the process of taking and reading out the image itself. These noise sources can be summarised in [23]:

- **Readout noise**
  All components of the readout electronics can produce some noise, from the electron transfer to the analogue-to-digital converter itself. Generally this noise is characterised by the manufacturer, in our case around 8 electrons per readout and pixel. It depends however also on the readout speed, and is generally reduced by

a slower readout. Binning can also help to reduce this noise, since several pixels are read out at once, but with the loss of resolution.

- **Thermal noise**
  The CCD chip consists essentially of an array of photo diodes, in which electrons can be exited to the conducting band with residual thermal energy. To reduce this noise one has to cool the CCD chip itself, which is not feasible with our camera.

- **Photon noise**
  With a fixed light intensity the probability to detect a photon number N is given in a Poissonian distribution. For large photon numbers, e. g. high intensities, this distribution can be approximated with a Gaussian distribution with standard deviation of $\sqrt{N}$ centred around a mean value $\bar{N}$. This intrinsic variation on the photon count cannot be reduced easily, and gives us a lower limit of our imaging noise.

Signal-to-Noise

More important than the pure noise characteristic is its relation to the signal obtained. e. g. for us this signal could be the photon atom number N_{atom} we calculate. That means we would like to increase the laser's absorption efficiency of the atom while still reducing the noise δN_{atom} on our image. In that case we would define the signal-to-noise ratio (SNR) as:

$$\text{SNR} = \frac{N_{\text{atom}}}{\delta N_{\text{atom}}} \tag{3.11}$$

We would also like to reduce other imaging errors, like off focus imaging, leading to a blurring of our atomic cloud absorption picture or the possible air motion in the imaging laser's beam path. A known atomic cloud distribution, like the Thomas-Fermi profile of a BEC, could be used as a reference. A wrong focus leads to broadening of this profile on the calculated atom picture, so we would try to minimize it.

3.4.1 Reducing the Imaging Noise

With the new software interface, we evaluated the effect of different camera settings on the noise. For that, σ_R was calculated in an empty region R without any atoms. To do so, the optical trap was turned off a few seconds before the imaging to let all atoms escape. The lasers were still frequency locked and the standard sequence was used to replicate accurate experimental conditions. We also used the existing algorithm for the calculation of atom counts, which is in part described in section 3.1. This algorithm and its scaling was calibrated for the previous hardware settings, such that the atom counts and therefore the deviations should not be viewed in absolute units, but in relative changes. For every data point in the following measurements 10 pictures were taken. To also test the possible improvements for shorter delay time, the time to capture and transfer a
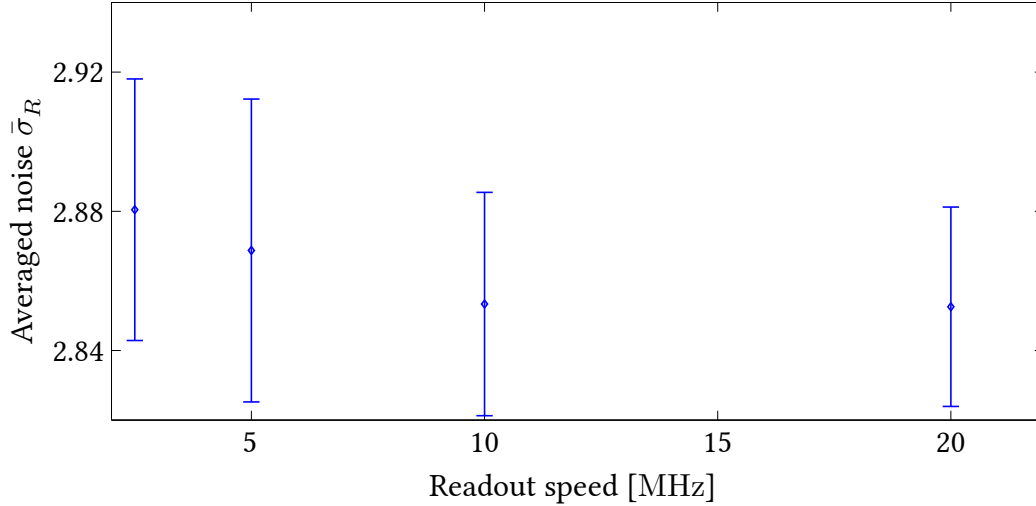
Figure 3.6: Measurement of the impact of a faster pixel readout speed on the noise in empty regions. The points are averaged values over the variance σ_R of 10 iterations. We could not observe a worse image noise for a faster readout, even though one would expect this.

picture to the PC had to be reduced. This was achieved by selecting a region of interest (ROI) of 200 x 200 pixels in the centre of our the CCD chip, which reduced the size of the image transferred from $2.76\,\mathrm{MB}$ to $0.078\,\mathrm{MB}$. At the same time, fewer pixels have to be read out on the camera chip, which also leads to an reduction of total capture time by a factor of 35. The intention hereby was to show the limits of the camera and not to prepare doing experiments with this setting, since the pictures are too small to use for more advanced post processing methods.

We started by scanning the readout frequency from $2.5\,\mathrm{MHz}$ to $20\,\mathrm{MHz}$, where $2.5\,\mathrm{MHz}$ was the previously used standard setting (see item f) in the previous section). The measured mean noise $\bar{\sigma}_R$ is plotted versus the readout speed in figure 3.6. Error bars are the standard deviation of σ_R for the ten different calculated images. An increased readout speed is usually connected to more electronic noise, which then can be seen on the calculated picture. We could however not see a significantly worsening of our noise with faster readout. Therefore the speed for the next measurements was set to the maximum value of $20\,\mathrm{MHz}$. Later we observed that for the whole picture with no ROI active only $10\,\mathrm{MHz}$ would work reliable, but we couldn't find any reason for that neither in the software nor the camera.

We expected to see the biggest reduction of noise by reducing the delay time t_{delay} between the atom picture and the reference picture. If t_{delay} is bigger than the usual time scale of a drift in frequency or fluctuations in the air, then the resulting diffraction pattern on the reference picture to the atom picture should be more or less randomly placed. Only if the time delay is shorter than the time scale of any fluctuations one
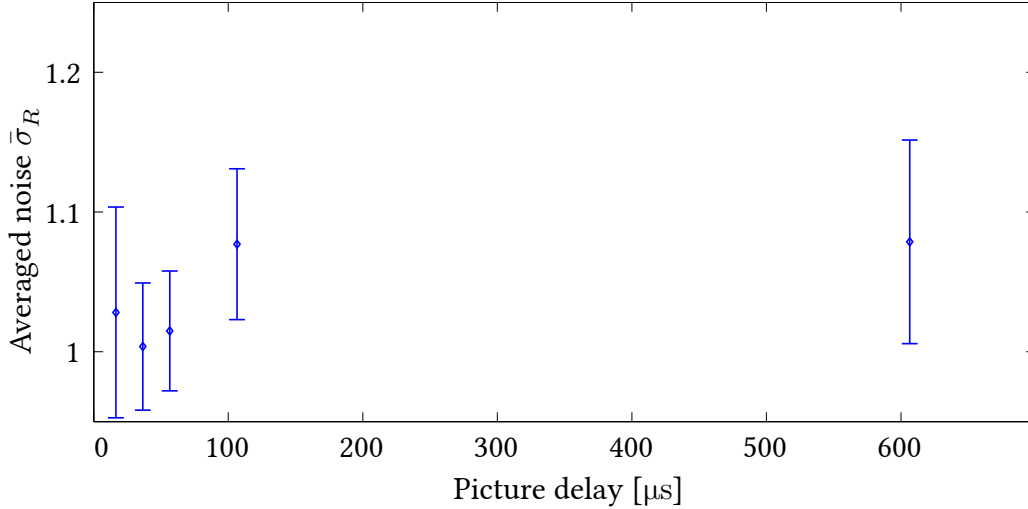
Figure 3.7: Change in noise on empty image regions dependent on the delay time between atom and reference picture. The points are averaged values over the variance σ_R^2 of 10 iterations. The previously used setting was $650\,\mathrm{ms}$ for this delay timer. To achieve times below $120\,\mathrm{ms}$ the image size for readout and transfer to the computer had to be reduced significantly. One can observe a small relative reduction of about 5 percent in noise starting at $56\,\mathrm{ms}$ and lower.

would see an improvement in the noise on the calculated picture. Assuming for example some $50\,\mathrm{Hz}$ noise on the laser control electronics from the house power line, then we would expect to also see oscillations on the laser frequency. This could lead to small changes in the beam path and the diffraction pattern, since the optical density is frequency dependent. This would therefore lead to an oscillation of our diffraction pattern with a period of $20\,\mathrm{ms}$. Therefore it would be ideal to go to delay times in the order of a few milliseconds. Using the described changes up to now (ROI of 200 x 200 pixels and readout speed of $20\,\mathrm{MHz}$) we could only reproducibly take two pictures with a delay of $32\,\mathrm{ms}$. To further reduce the image size a 4 x 4 binning of the picture was chosen, which lead to an effective reduction in image size and readout of a factor of 16. With that we could go down to $16\,\mathrm{ms}$, which is only half the time of the previous measurement. Since this scaling is not proportional it already shows that the size of the image including readout is not the limiting factor anymore. The communication between computer and camera, triggering and internal processes now limit the capture time of a single picture. The maximal frame rate of the camera specified by the manufacturer is 165 frames per second, which would lead to a delay time of about $6\,\mathrm{ms}$ [24]. This is however with not further specified binning and ROI and especially not with an external trigger signal as needed in our experiment.

With 4 x 4 binning, a ROI of 50 x 50 pixels and $20\,\mathrm{MHz}$ readout a scan of the delay time from $16\,\mathrm{ms}$ to $106\,\mathrm{ms}$ was performed, with an additional measurement at around $650\,\mathrm{ms}$ which was the previous setting used. The averaged noise of 10 different pictures $\bar{\sigma}_R$ is
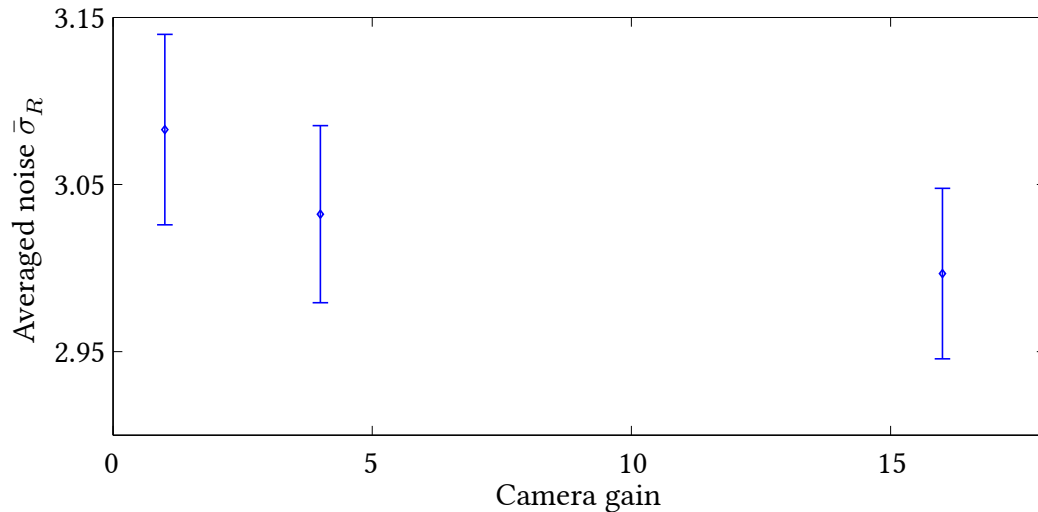
Figure 3.8: Plotted change of the noise against three different gain settings. Each point is an average of 10 iterations. One can see a small improvement in imaging noise by going to higher gain settings.

plotted against the delay time in figure 3.7. One can see that with a delay time of $56\,\mathrm{ms}$ a small reduction in noise is possible, but already with a delay of $100\,\mathrm{ms}$ this improvement is lost again. It takes a delay of around $120\,\mathrm{ms}$ to capture two full resolution pictures (no binning or ROI) which would be necessary for intensive post processing. The trade off for the improved noise characteristic is therefore very high: One would have to either reduce the resolution of the camera with binning, or only use part of the image with a ROI.

As a next step we wanted to see the effects of a change in the camera internal gain (see item e) in the previous section). By changing the gain factor of the ADC one will of course also amplify any noise that is already existing beforehand, like photon noise or thermal noise. But it can reduce the internal electronic noise of the ADC relative to the overall light intensity, provided one doesn't reach saturation. Therefore it should always be tried to use the full dynamical range of the 12 bit ADC. For this measurement the binning was turned off, so that we had a 200 x 200 pixel ROI, with a readout speed of $20\,\mathrm{MHz}$. The averaged noise over 10 iterations of each setting is plotted against the gain setting in figure 3.8. One can again see a small reduction of noise going from a gain of 1 to a gain of 16 of about 5 percent. This however is dependent on the initial intensity and should be set to use the full 12 bit dynamical range of the camera. We were however only able to set the gain to 16 since the previously lithium intensity was set to a very low value. Additionally one has to keep in mind that changing the gain has lots of implications on the exact calibration of the algorithm calculating atom numbers, which would have to be completely readjusted.

Finally we looked into the possible different exposure time settings and their effect on the image noise. This of course can only be adjusted sensibly with atoms present, since the detuning and the overall absorption is highly dependent on the light pulse time, as
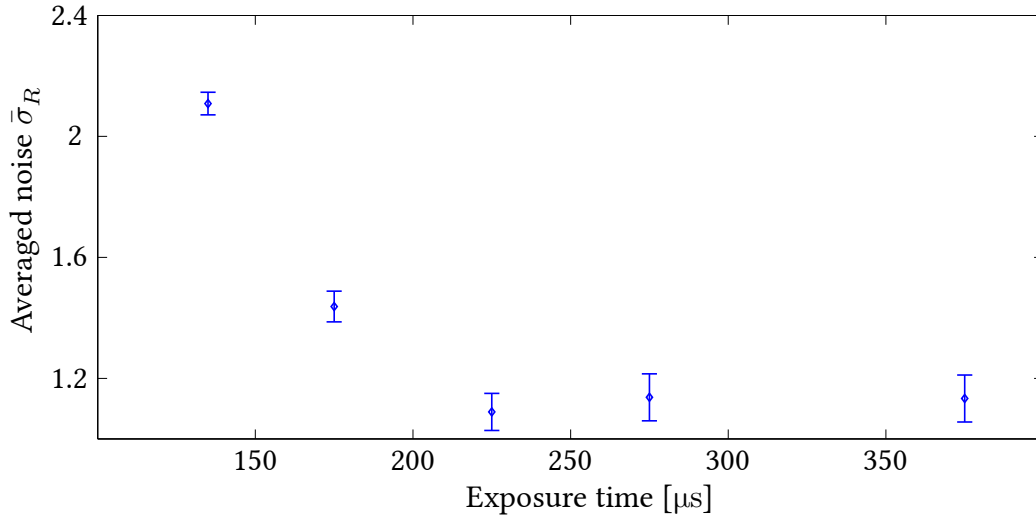
Figure 3.9: Graph showing the averaged noise versus the software exposure time cal-
culated out of 10 iterations for each data point. The time scale starts at the
trigger signal going high ($t = 0$). The light pulse is shone onto the camera
between $t = 125\,\mu s$ and $t = 250\,\mu s$. A longer exposure time means more
light onto the CCD chip. Therefore the raw signal of atom and reference pic-
ture is small relative to the electronic and thermal noise on these images. This
translates into a bigger noise on the calculated image as shown here. After
$250\,\mu s$ the noise is constant since the imaging intensity stays the same.

shown section 2.1. It is however possible to check the function of the camera setting sep-
arately without atoms and see if it would improve the noise characteristics. We again
used 4 x 4 binning, a 200 x 200 pixel ROI and $20\,MHz$ readout speed.

The software exposure setting was activated (see item g)) and the exposure time scanned,
starting from the trigger low to high transition. In figure 3.9 the results of the measure-
ment are shown. To correctly understand these pictures, one has to keep in mind that
the camera trigger signal goes high $125\,\mu s$ before the imaging light pulse of $125\,\mu s$ starts.
So after $250\,\mu s$ from the trigger the light should disappear again and no change in counts
on the image should appear. At $135\,\mu s$ the CCD chip was exposed to a small amount of
light, therefore the electronic and thermal noise play a more important role relative to
the light intensity. This then translates into image noise on the calculated picture shown
in the figure. With more exposure time the raw signal compared to the noise is increased
again, until the light pulse ends after $250\,\mu s$ and no big change is observed anymore.

Longer exposure times are nevertheless not always useful, since after a certain time the
atoms are accelerated with increasing detuning. Off resonant absorption is smaller, so
the atom calculation has to be adjusted precisely. It also leads to imaging disturbances,
comparable with focusing errors. We therefore have to not only look for a reduction in
noise, but also to its relation to the signal.

44

3.4.2 Optimising the Lithium Imaging

During the course of this thesis a lot of parameters were changed in the experimental sequence, in order to observe spin dynamics in the sodium lithium mixture. We switched from fermionic ^{6}Li to bosnic ^{7}Li and changed the coils inducing the Stern-Gerlach separation pulse for a faster separation of the magnetic sub states. Therefore we took the chance and optimised the lithium imaging again, while also using the settings tested in the previous section. The main parameters tuned were the exposure time, the imaging frequency detuning and the position of the camera relative to the magnifying optics.

As already discussed in the first chapter, a big problem with the absorption imaging of lithium is its small mass. Due to this and the photon recoil, the atoms are very fast blown out of resonance, as seen in figure 2.2. A heavier atom like sodium is not as easily affected, and can be exposed to light longer, while staying resonant. One has to choose the exposure time to a value where most atoms can absorb the incoming light with very little detuning. For a longer exposure the absorption rate changes and the algorithm calculating atom numbers is underestimating the atom number. But as seen in the previous measurement however, a longer exposure time can reduce noise on the calculated image. Therefore we scanned the exposure time t_{exposure} from $125\,\mathrm{\mu s}$ (beginning of the laser pulse) to $200\,\mathrm{\mu s}$, including a control shot with $375\,\mathrm{\mu s}$ (old settings). Within each setting 5 measurement iterations were performed to reduce experimental fluctuations, e. g. of the atom number. The atom pictures with no or only very few atoms were sorted out as measurement errors. All further evaluations was then done over the averaged picture of the remaining images for each setting. They are shown in figure 3.10. Here one can already observe qualitatively that after an initial rise, the peak intensity is again reduced. This is caused by the decreased absorption due to acceleration. The change in detuning also has a big effect on the susceptibility of an absorbing material [28]. Since the refraction index is directly coupled to it, one can observe a big change in the refraction index close to the resonance. Some of the light away from resonance therefore will be diffracted out of the parallel beam path. The cloud of atoms is acting then as a lens, leading to an effect similar to defocusing. This can be observed in the images, since the shape of the atoms gets blurred out. One can also see the previously measured decrease in noise in regions with no atoms for longer exposure time.

To quantify this effect more and to objectively choose the best settings we would like to calculate a signal-to-noise ratio. For this the mean noise in an empty region below the atoms (size 200 x 100 pixels) was measured. The definition of a *signal* quantity of our images is however not so straight forward. One could think of the total atom count which is linked to the absorption of light on the atoms. With stationary atoms, this value should stay constant even if the light frequency is not on resonance. If since the atoms are moveable, we saw already in figure 2.2 and equation (2.36) that the atoms are soon detuned from resonance and the absorption rate is not constant any more. So after a certain exposure time, the calculated atom numbers should be underestimated. As explained in the paragraph above one also has to include blurring out of the light shape due to off resonant diffraction which lead to wrong calculation of atom numbers.
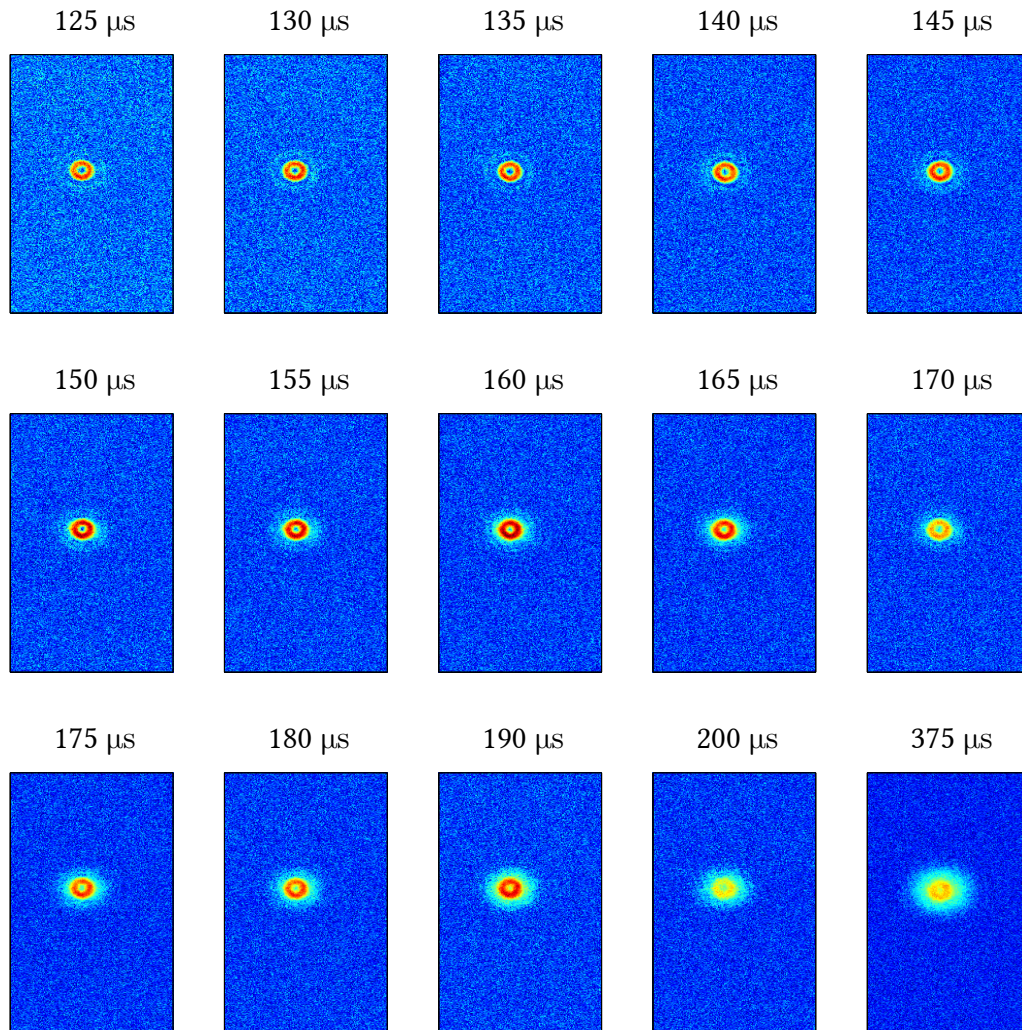
Figure 3.10: Measurement of the effects of different exposure times to the calculated atomic pictures. For each setting an averaged image out of a set of 5 pictures was created. All averaged images have a fixed colour scale. The ring form of the atomic cloud is a result of defocusing and a diffraction effect due to frequency detuning. These settings will be adjusted in the following. One can see the decrease in noise in regions with no atoms with increasing exposure time. At the same time the atomic cloud gets blurrier after a long exposure time, since the measured absorption in the centre regions is increased. The calculated peak intensity of the atoms also decreases.
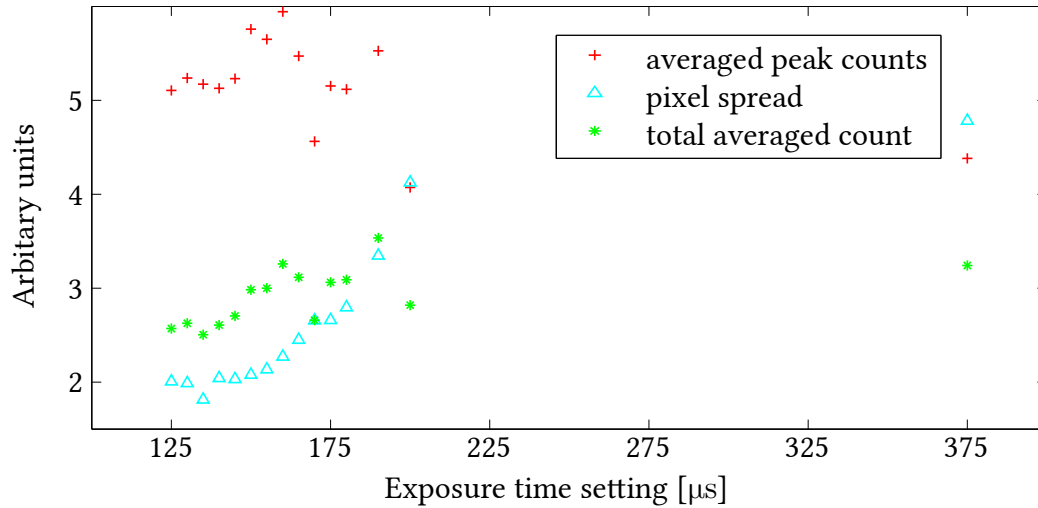
Figure 3.11: Different atom image properties, that could help to define the signal on the calculated atom pictures for different exposure time settings. The overall averaged atom count of all pixels is not useful since the main signal here is the sharpness of the absorption profile. For this we use the peak count, averaged over the 200 pixels with the highest value and the pixel spread, that describes the area relative to that the peak count that. In equation (3.12) we defined the signal as the ratio of peak count to pixel spread.

Therefore the total atom number is not a good quality to use in the definition of a signal. The averaged atom pixel counts of the images are nevertheless plotted in figure 3.11 as additional information.

The discussion on the sum of all atom counts directs us to the in this case most interesting image property, the sharpness of the calculated shape. If the absorption properties of the cloud change over time then this leads to a change in the resulting momentary absorption profile. The resulting image on the chip then will be the integrated absorption shapes over the exposure time. If no change happens this overall picture will have the same absorption profile during the time of the exposure and we will see a sharp profile. This doesn't necessarily have to be the actual cloud shape, since we haven't calibrated the imaging frequency and the camera focus.

To define this *sharpness* we will use two quantities, the averaged peak count and the pixel spread: To reduce statistical influences on the maximal pixel value of a picture we will take the average of the 200 highest counting pixels, and will define this as the peak count. Any blurring out of the absorption shape will reduce this value. The pixel spread is measured relative to the peak values. It describes on how many pixels the biggest part of the signal with is distributed. For that we count the pixels with a value higher than half the averaged peak count and scale them for better display with constant factor. This pixel spread is also shown in figure 3.11 together with the averaged peak count. One can see that it is roughly constant up to an exposure time of $150\,\mu s$, then it rises as the image gets blurred out more and more.
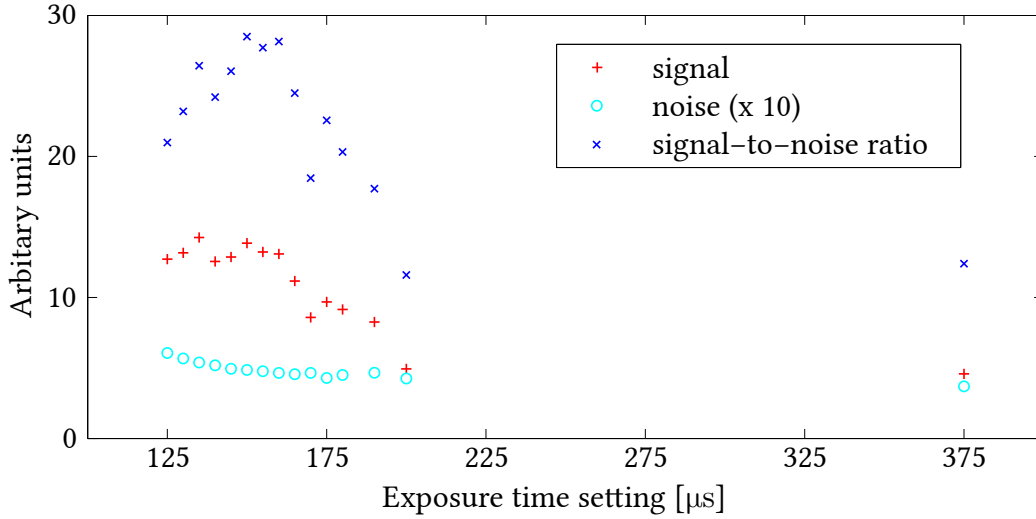
Figure 3.12: Signal-to-noise measurement for the exposure time. Plotted against the exposure time is the peak signal, the noise and the signal-to-noise ratio. One can see the already in figure 3.9 observed decrease in noise. The ratio signal-to-noise has its peak around $155\,\mu s$, which would be a good value to work with.

With these two values we can define our signal as:

$$\text{signal} = \frac{\text{average peak count}}{\text{pixel spread}} \tag{3.12}$$

Figure 3.12 shows the resulting signal, the noise calculated as described in the paragraph above and the resulting signal-to-noise for different exposure time settings. The overall maximum of the SNR is with an exposure time of $160\,\mu s$, close to the time of the signal maximum. The SNR alone is more than twice as high for in the peak region around $155\,\mu s$ than for the old setting of $375\,\mu s$, while the signal we defined is nearly tripled. Looking again at the pictures in figure 3.10 one can qualitatively see that the pictures around that exposure time setting are both sharp and with little noise. For all further measurements with lithium we switched to $160\,\mu s$ exposure time in the software exposure mode of the camera.

The now clearly visible ring shape can have two origins. One would be a defocused image, which we could improve by changing the distance of the camera to the magnifying lens stack. Another possibility is however that the imaging laser frequency is detuned. While the real part of the optical refraction index is unity on resonance, it is non-zero if not on resonance [28]. Assuming a spherical cloud this would have the same effect as a small lens on the otherwise parallel laser beam. Since these two effects can cancel out each other, we have to find a way to optimise one without using the shape of the cloud. An independently observable value is the pixel atom count summed over the whole picture. The calculated atom number is dependent on the absorbed light intensity which itself is dependent on the detuning of the laser to resonance. So by optimising on the atom count one can optimise the frequency and therefore its effect on the imaged
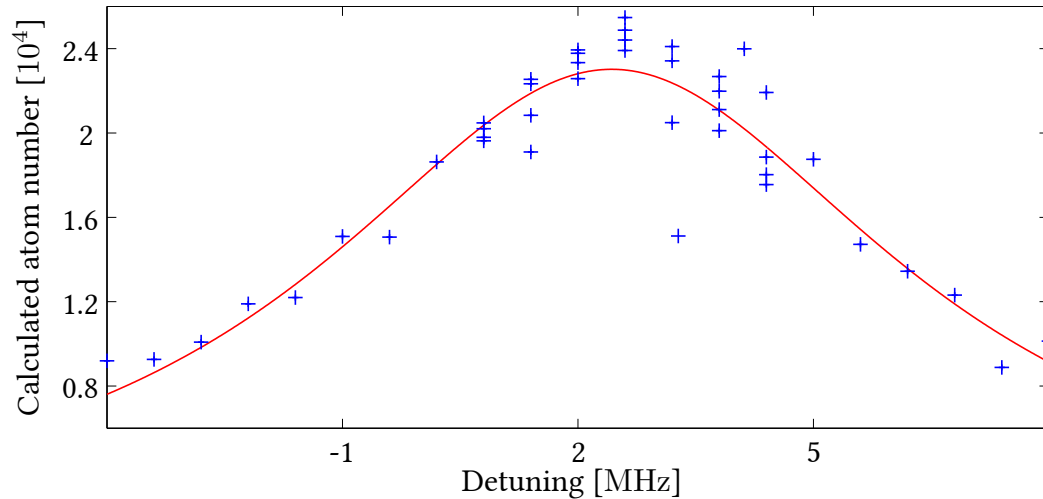
Figure 3.13: Plot of the sum of all calculated pixel atom counts in one picture against the beat lock frequency between the spectroscopy and the tapered amplifier laser. The calculated atom numbers are for the assumption to be on resonance, so there is no change in real atom number, but on the detection. The detuning frequency was determined by the beat lock frequency. It is not completely linear, but sufficiently enough to fit a Lorentzian to the data set, seen in red. As one can see the maximum is about $3\,\mathrm{MHz}$ away from resonance, which corresponds to around half the line width.

cloud shape.

In figure 3.13 we plotted the sum of the calculated atoms dependent on the imaging frequency. Fitting Lorentzian to the data we could see it's peak atom count and therefore the optimal detuning for our imaging to be at around $\delta \approx 3.6\,\mathrm{MHz}$ offset lock frequency. This corresponds to a detuning of around half the line width, which is reasonable in light of the results of figure 2.2(a) and a light exposure time of $160\,\mathrm{\mu s} - 125\,\mathrm{\mu s} = 35\,\mathrm{\mu s}$.

Using the motorised stage the lithium camera is mounted on we could change the position of the camera very accurately. This was used to scan the position compared relative to the initial one over a range of 20 mm. We could observe a shift in the position of the atoms on the images taken as well as a change in shape. The change in position results in not perfect positioning and alignment of the imaging system to the symmetry axis of the magnification objective. The shape of the cloud changed from the doughnut-like ring to a very smooth, positive distribution with a smaller width. To quantify this we could use a function build into the image algorithm. It creates two profiles of the atomic distribution by adding up all values in each pixel row and column respectively. Fitting a Gaussian distribution to this profile doesn't match perfectly, it is however good enough to get a quantified value of the width of the atomic distribution. Using these values to compare different settings figure 3.14 was created, showing a minimum in width around a displacement of around 10 mm. Any circular deformation disappeared and due to the concentration of the absorption onto fewer pixels, we got a very good signal to noise.
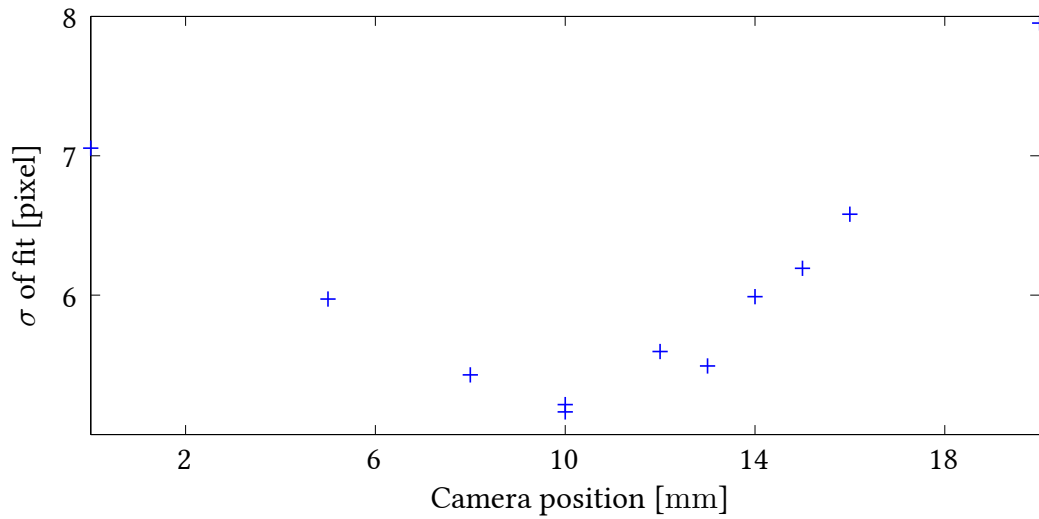
Figure 3.14: Measurement to determine the right camera position to be in focus. Plotted relative to the camera position is the width of a Gaussian fitted onto the column wise summed profile of the atomic picture in x direction. The minimum and therefore the sharpest image is at 10 mm. This can also be done for the y axis, but they both yield the same results.

This improvement made further measurements possible, so that the magnetic field stabilisation envisioned in this thesis could be constructed.

4 Stabilisation of the Magnetic Field

The experiments that are currently being done at our experiment are very sensitive to the magnetic field, such that we would like to have a field stability of less than a milligauss at the centre of the atom trap. In the scope of this work an active magnetic field stabilisation was designed tested and a prototype implemented in the experiment.

In this chapter, first coarse measurements of the magnetic field during the experimental sequence are described. Then the active stabilisation scheme is introduce,d which includes the tests done and the different configurations used, discussing advantages and disadvantages. Finally, the first results on the stabilisation will be presented.

4.1 Drifts and Noise of the Magnetic Field

The atoms we work with are both alkali metals with one valence electron. As described in section 2.2 this leads to a high sensitivity of the internal energy levels to magnetic fields. We use this in our sequence for example to magnetically trap atoms and also to separate different magnetic hyperfine states in the Stern-Gerlach pulse sequence, just before the imaging [37]. These magnetic fields are produced by stable current supplies and are switched on and off in a very controlled way, so that one can characterise them and adjust the experimental sequence in the desired manner.

There are however external fields that are not controlled by the sequence. They can influence experiments which are critical on the magnetic stability and an exact energy level splitting. We can separately discuss these fields by their frequency f:

- **Drifting fields:** $B(f << 1\,\text{Hz})$
  These fields are also called the *DC* fields, since similar to a DC current they have a mostly constant value. The earth magnetic field in Heidelberg for example has a value of around 0.45 Gauss [12]. Also permanent magnets like the large ion pumps in our lab give a contribution to the permanent magnetic field. There are also field sources that can't be assumed to be constant over time, like the superconducting 15 T electromagnet used in an other lab in our institute. Even with a constant background outside the lab, the DC fields may drift, due to the change of ferromagnetic magnetisation during our sequence.This can lead to a change over the course of a day due to the large fields we use in our magnetic trapping phase. The slow drift in these fields is dangerous for our measurements and frequent and time consuming recalibration is necessary.

- **Low frequency fields:** $B(1\,\text{Hz} < f < 1\,\text{kHz})$
  The grid electricity in Europe is normed to 230 V oscillating at 50 Hz. All electronics in our lab are powered with this AC line, even high power line cables are
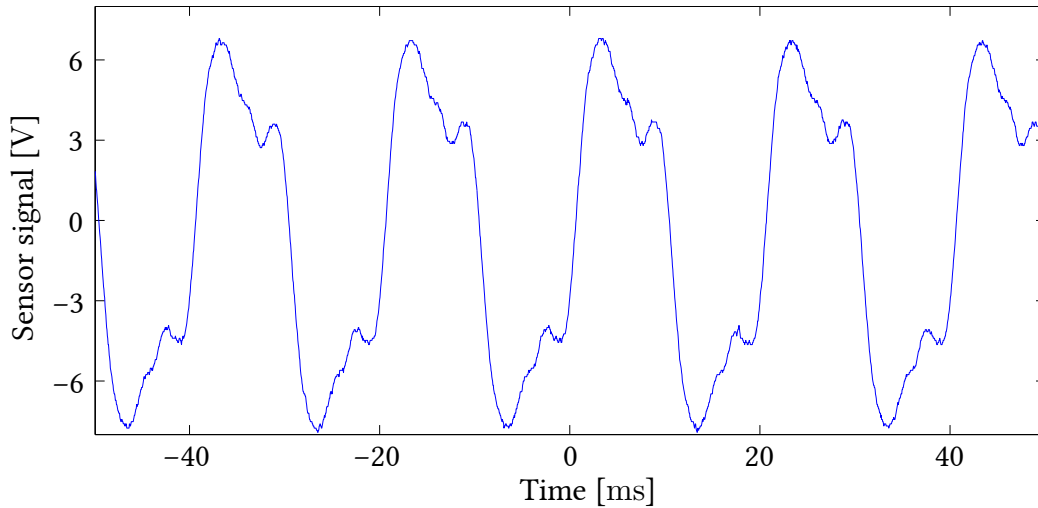
Figure 4.1: Amplified magnetic background oscillations in our lab. One can clearly see the 50 Hz with it's 20 ms oscillation period and higher harmonics at 100 Hz and 150 Hz. This was an averaged measurement using a flux-gate sensor and a self built amplification with filter described in section 4.3, triggered on the 50 Hz line. The shape and exact amplitude is dependent on the direction and the exact location in the lab.

just three phase shifted voltages with this oscillation frequency. Every cable carrying this 50 Hz current acts as an antenna, producing an oscillating field at this frequency and its higher harmonics. This 50 Hz noise can be seen everywhere in the lab and also affects our atoms. The rubidium BEC experiment in our research group measured this oscillation on their atoms [22]. The period of 20 ms is comparable to typical sequences in our experiment and can therefore potentially disrupt our measurements. Triggering the experimental sequence on the power line can give a better reproducibility of the measurements data, however it still doesn't get rid of the magnetic field change during the oscillations.

- **High frequency fields:** $B(1\,\mathrm{kHz} < f)$
  There is also higher frequency electromagnetic noise in our lab, with various sources. Switching power supplies are known to produce high frequency in their output current and voltage (see [17]). Depending on their construction the noise frequency is in the order of several tens of kilohertz. If a power supply like this drives for example an amplification stage, then one can detect these frequency components also on their output, which then radiates into the lab. The rubidium frequency reference we use also operates on 10 MHz and if not properly shielded emits electromagnetic noise.

Slow magnetic field drifts can be measured with trapped cold atoms in an initially well known state with a microwave frequency sweep.One can calculated the precise energy shift of the magnetic hyperfine sub states of the ground state with the Breit-Rabi

formula shown in equation (2.43). With a direct-digital-synthesizer (DDS) system implemented by Arno Trautmann [37], which uses a rubidium frequency reference one can control the microwave frequency very accurately. Probing then the atoms with a microwave frequency sweep then gives the transition probability of the population, which is maximal for resonant excitation. This can then be used to calculate the magnetic field inducing the energy shift. For fast oscillating magnetic field ($f > 50\,\text{Hz}$) this technique is not suitable any more, since the RAP ramp duration limits the possible probing time. To get rid of drifts and most of the other noise in the magnetic field one needs to search for other concepts.

The idea of an active magnetic field stabilisation is to suppress low frequency noise and cancel any drifts with a well controlled magnetic field. This can be done by a feedback control loop, as described in section 2.3.
To implement this technique one needs an accurate and fast sensor, which can measure the fluctuations in the background and a fast regulation that is capable of countering the measured changes. To get an indication on the magnetic field behaviour during our sequence we would first like to characterise this roughly with a low resolution sensor. This then can be used to define the more exact requirements on the sensor and the regulatory circuit.

4.2 First Sequence Characterisation

For these initial measurements a small and inexpensive magnetic field sensor like the Honeywell HMC5883L was deemed to be ideal. It is designed for use in low power, low cost applications, e.g. as a compass for mobile navigation. One of these chips including a breakout board (GY-273) can be bought online for 4 to 8 €. It is controllable via the I²C standard which makes it ideal to use with a low cost microcontroller like the Arduino. The three axes can be measured simultaneously and then are digitised with a 12 bit DAC. In table 4.1 some further specifications of the chip and the breakout board are listed [16].
The connection of this chip with an Arduino is relatively simple. The on board voltage regulators of the GY-273 accept an input voltage of 3 to 5 volts, so it can be powered over the 3.3 V and 5 V power supply pins on the Arduino. Apart from the ground connection, one only needs two cables, serial clock (SCL) and serial data (SDA), for the I²C communication. There exists another breakout board, the GY-271, which additionally adapts the signal levels of these channels to a better suited voltage, but for us the GY-273 was sufficient. The Arduino has special pins to be used for these two communication lines which then will be addressed with the `wire.h` library. All communication between the two devices is then done by the exchange of full bytes, which is always initiated by the master device, in our case the Arduino. It can order read or write operation with the chips identification number (first byte) on 13 chip internal 8 bit registers (second byte), used for configuration, operational mode, status and data output (third byte).
Sending for example the three bytes (written in hexadecimal):
`0x3C 0x00 0x70`

Table 4.1: Some of the specifications of the Honeywell HMC5883L [16] magneto-resistive sensor and it's breakout board GY-273 we used for initial measurements, like the one shown in figure 4.2. It is recommended to average over several measurement iterations, to reduce the electric readout noise.

Sensor type	3 - axis magneto-resistive sensor
Size chip	$3.0 \times 3.0 \times 0.9\,\text{mm}$ surface mount
Readout	12 bit ADC with 8 gain settings
Measurement range	$\pm 0.88\,\text{G}$ (min. gain)
	$\pm 8.1\,\text{G}$ (max. gain)
Digital resolution	$0.73\,\text{mG}$ per count (min. gain)
	$4.35\,\text{mG}$ per count (max. gain)
Output rate	$160\,\text{Hz}$ maximum
Noise floor (100 samples)	$2\,\text{mG}$ standard deviation
Supply voltage (GY-273)	$3.3\,\text{V}$ - $5\,\text{V}$

addresses the chip on a write operation with 0x3C, and writing the byte 0x70 into the first register 0x00. This operation sets the chip to do a normal measurement with internal average over 8 measurements in 15 Hz repetition speed.

For reading one first has to move the register pointer to the register one wants to read with a write operation, for example to the forth register 0x03:

0x3C 0x03

Then one can specify the number of bytes to read, with every operation moving the internal register pointer one register further. The order

0x3D 0x06

is therefore the request of a read operation to the chip 0x3D (last bit in the byte has changed to indicate a read operation) to send six bytes 0x06 back over the serial data channel. In our specific case this would result in the chip sending register 4 to 9 back to the Ardunio, which contain the x, y and z value of the latest measurements. The full code used including more documentation is added in the appendix of this thesis's online version.

The performance of this very low cost sensor was impressive, since we could detect changes of the magnetic field in the order of $1\,\text{mG}$. For this we used averaged values over several hundred measurements. This is since then used to detect the operation of the superconducting $15\,\text{T}$ electromagnet three floors above the lab. When it is working or changing its field output, the energy states are shifted, requiring a recalibration of the microwave preparation sequence.

We measured the magnetic field during our experimental sequence to gain additional information on the behaviour of the fields we wanted to regulate. In order to do so, a code was written that measured all three axes, stored the data on an Arduino and repeated the measurement a certain number of times. This was then read out later over the USB connection to the computer using a matlab script. Since the storage for live variables of the Ardiuno was limited we had to move from an Arduino Uno (2 kB SRAM) to an
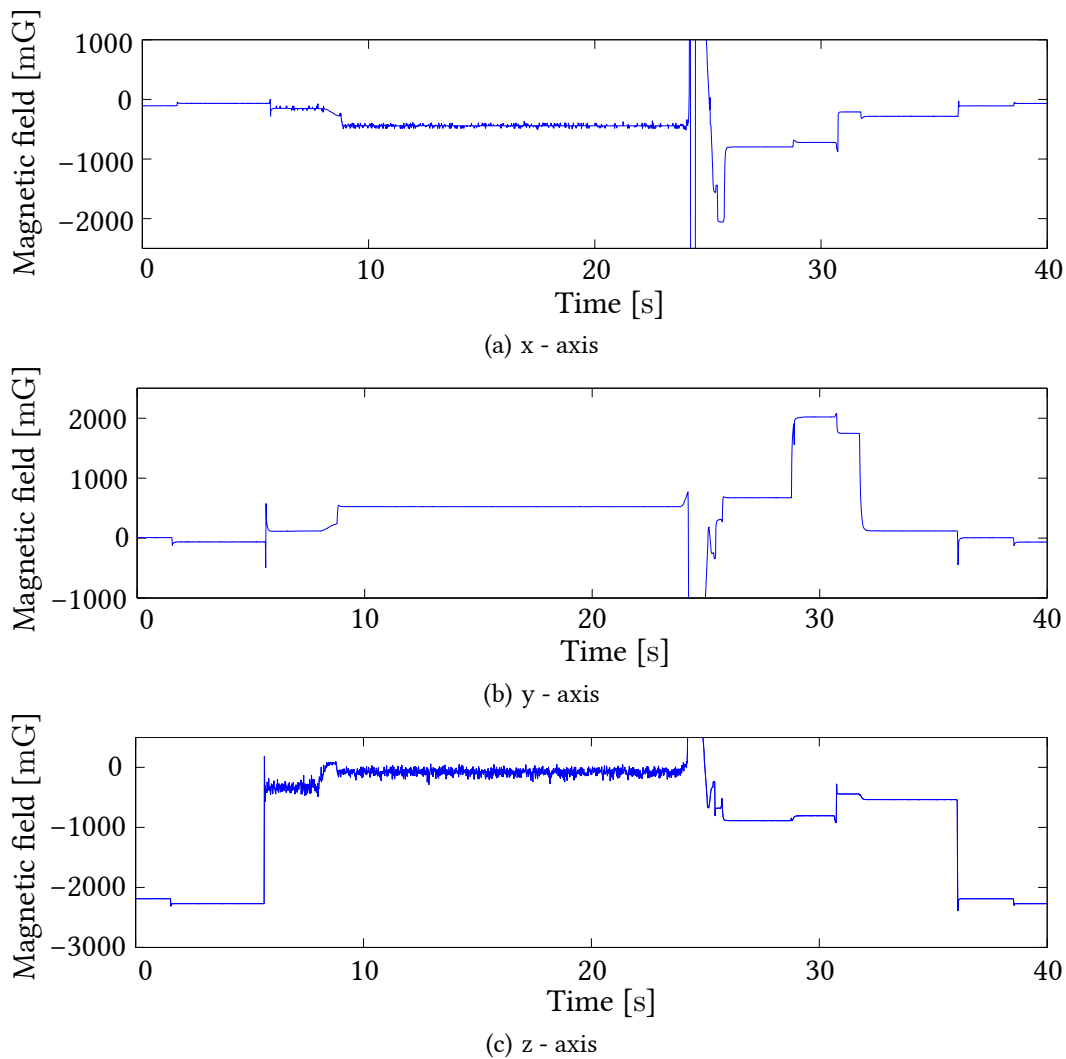
Figure 4.2: Measurement of the magnetic field 30 cm away from the atom trap using the
Honeywell HMC5883L magneto-resistive sensor [16]. We choose a setting
that enabled us to measure in a range of ± 4.0 G with a 12 bit resolution,
leading to 2.5 mG digital steps. All three axes were measured simultaneously
and stored in 13.3 ms time steps on the RAM of an Arduino Mega. With
that one whole experimental run could be observed, starting at $t = 0$ s. The
different phases of the experimental sequence can be identified, but one has
to remember that the absolute values the sensor sees are vastly different to
the field the atoms experience due to the large distance to the actual trapping
region. A nice feature to observe is the microwave for evaporation visible on
the z axis.

Arduino Mega (8 kB SRAM) and add pauses in between the measurements. With the careful adjustments of these breaks we perform a measurement of the magnetic field in all three axes during a sequence of 40 seconds with 3000 measured values for each axis. The resulting data can be seen in figure 4.2. Even though the sensor was placed 30 cm away from the atoms to avoid saturation, we could see and identify the different phases in our sequence:

Time [s]	Start of phase
1.5	NaLi - MOT
5.5	Magnetic trap (MT) and cleaning sweep
8	Slower field off
9	Evaporative cooling with microwave
24	Turning off MT and switching to dipole trap
26	Evaporative cooling in dipole trap
29	Experiment in the ultra cold
32	Transmitting new sequence and wait
36	Li-only - MOT

One can nicely see the microwave that is used for evaporative cooling during the cleaning sweep and the microwave cooling phases on the x and z axis. After turning off the current for the magnetic trap the field largely fluctuates due to the sudden current change in the coils and leads the sensor into saturation for a short time.

With this measurement and the previous tests on accuracy we could see that a digital measurement device would not be fast enough to detect the noise we wanted to eliminate (with $f = 50\,\text{Hz}$) due to the limitations in readout speed and data acquisition. In addition to that, this sensor was not accurate enough to resolve fluctuations with low magnitude nor detect the DC field with the required high precision. However, measuring the rise times of coils' magnetic field was also used to assure the experiments were done in a period of magnetic field stability. This is a prerequisite for accurate measurements of Feshbach resonances and other processes mainly dependent on the magnetic field.

4.3 The Flux Gate Sensor

Out of the previous tests with the HMC5883L magnetometer, we could rule out that a digital sensor would be suitable for a fast regulation and stabilisation below one milligauss. Analogue flux gate magnetic field sensors are very fast in their response while granting precise and exact measurements, provided the readout electronic is designed well enough [14, 27]. They are therefore ideal for our intended application.

The sensor we decided to use is the Bartington MAG-03MS1000 [5] three axis, low noise magnetic field sensor. It has a measurement bandwidth of 3 kHz in a range between $\pm10\,\text{G}$, giving out a not amplified voltage of $1\,\text{V/G}$. The electronics are optimised for low noise operation, resulting in a noise floor of less than $100\,\text{nG}/\sqrt{\text{Hz}}$ at 1 Hz.

Another experiment in our research group already uses one of these to stabilise the magnetic field around a value of $9.2\,\mathrm{G}$. With their set up they could achieve a stability of about $45\,\mu\mathrm{G}$ with a single axis stabilisation [22]. To be able to also achieve a high stability also at lower fields around 0 G we need to stabilise all three spatial directions at the same time. We still could use their experience and expertise to build up and enhance our system.

Operation Requirements and First Amplification

The sensor needs a well regulated and stable supply voltage of more than $\pm$(measurement output range + 2.5 V) = $\pm12.5\,\mathrm{V}$, which is preferably floating compared to other grounds in the experiment [6]. A separate box connecting the sensor via a shielded VGA cable is providing the needed supply voltages, and was used to feed signals of the three axes to separate BNC ports. Care was taken to connect the grounds of supply voltage and signal only directly at the signal, to avoid any influence on the signal. This box was connected to a linear DC power supply, so that no electric connection to other devices was established. We knew from the other experiment, that in their lab the dominant 50 Hz noise had an amplitude of about $2\,\mathrm{mG}$. This corresponds to a raw sensor signal oscillation of $2\,\mathrm{mV}$. We would like to amplify our signal to be able to detect and then regulate on that signal. This amplification then can additionally act as an impedance converter, since the sensor readout requires a high input impedance. Using a differential amplifier between signal ground and the signal would have the big advantage of decoupling the sensor ground completely from all other electronic grounds in the lab, leading to a electronically robust system, which can easily be shielded from outside fields.

This scheme has however the big disadvantage that one is limited in the total amplification possible: A differential amplifier with supply voltages of $\pm12\,\mathrm{V}$ is only able to put out up to $10\,\mathrm{V}$ [2]. For a measurement at $4\,\mathrm{G}$ one would expect a $4\,\mathrm{V}$ signal with a $2\,\mathrm{mV}$ amplitude oscillation on top. Using a gain factor of 2 would be possible, since the total signal is below $10\,\mathrm{V}$, but already an amplification by the factor of 3 is not possible any more. This disadvantage is mostly irrelevant if one aims at measuring fields close to 0 G. For our first test we however wanted to be able to access the whole $\pm10\,\mathrm{G}$ measurement range with an amplification of the AC magnetic field component between 1 to 1000 fold. This initial amplification device can be seen in figure 4.3. We build a digitally controlled voltage source, that could put out between 0 and $10\,\mathrm{V}$ in 1024 steps. This voltage was then subtracted form the sensor signal and amplified. The voltage source consisted of a 10 bit digital potentiometer AD5293 from Analog Devices and a REF01 $10\,\mathrm{V}$ reference, including voltage regulated power supplies. The digital potentiometer was then controlled and given commands by an Arduino. The big disadvantage of this scheme was that one had to connect the ground of the voltage source to the signal ground, and therefore breaking the electric isolation of the input signal. To reduce any negative effects we used the same floating power supply unit that drove the sensor for the voltage source as well and decoupled to the Ardunio with the use of opto-couplers. The amplification was done by a low cost instrumental amplifier AD622 with maximal gain of 1000 and a gain
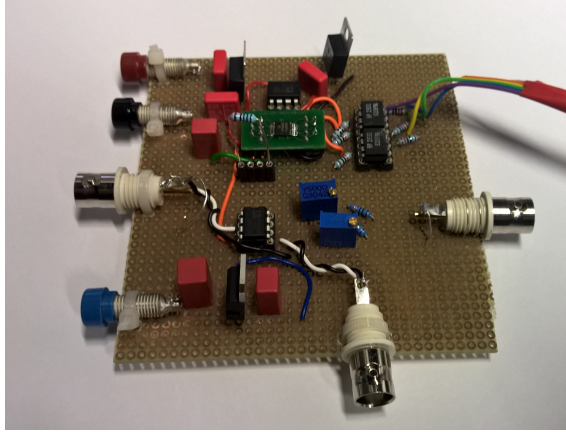
Figure 4.3: First amplification stage with a digitally controlled voltage source. The top-most DIP IC is an REF01 10 V voltage reference, which is then sent to the digital potentiometer below on the small separate board. The control SPI signal from an Arduino is connected through the opto-couplers to the right of both ICs. The Sensor signal, coming in from the left, then gets subtracted from this set voltage and amplified with the instrumentation amplifier in the lower centre of the board. The signal can then be directly read out (BNC connector at the bottom of picture) or sent through a notch filter described later in the text (BNC connector to the right).

bandwidth product (GBP) of $8\,\mathrm{MHz}$. This factor is a description of the deterioration of the amplification bandwidth with higher gains, e.g. with gain $G = 1000$ the bandwidth of the amplifier is $8\,\mathrm{MHz}/1000 = 8\,\mathrm{kHz}$. This is totally sufficient for the use with the $3\,\mathrm{kHz}$ bandwidth sensor. The gain of the amplifier is determined by a resistor R_G connected to it, with the gain $G = 1 + 50.5\,\mathrm{k\Omega}/R$. Using for example a $50\,\mathrm{k\Omega}$ potentiometer enables us to continuously tune the gain factor from 2 to 1000.

Figure 4.4 shows the raw sensor signal and the resulting amplified signal. It is important to note the different scales for the two signals. Comparing the initial signal amplitude of $2.5\,\mathrm{mV}$ to the resulting one of $2\,\mathrm{V}$ one can deduce the gain factor G to be around 800. Looking at the absolute values one will note that the raw signal is positive ($193\,\mathrm{mV}$ mean) whereas the amplified one is negative ($-3\,\mathrm{V}$ mean), therefore a bigger voltage than $192\,\mathrm{mV}$ had to be subtracted before the amplification. This is already showing the strength of this method: Without the subtraction and the same amplification the resulting signal should theoretically be as big as $154\,\mathrm{V}$, which exceeds by far the output range of any IC amplifier typically used in our experiment. Like this we can measure and regulate on a magnetic field with a system sensitivity of up to $800\,\mathrm{V/G}$. One can also see that there is a lot of high frequency noise that broadens the width of the amplified signal. Analysing the frequency spectrum one can detect a very strong frequency component of $15.78\,\mathrm{kHz}$ on the sensor signal. This is a lot higher than the bandwidth of the sensor and is an artefact of the measurement method. Each flux gate sensor has a fixed frequency of operation that has to be filtered out afterwards. For further information
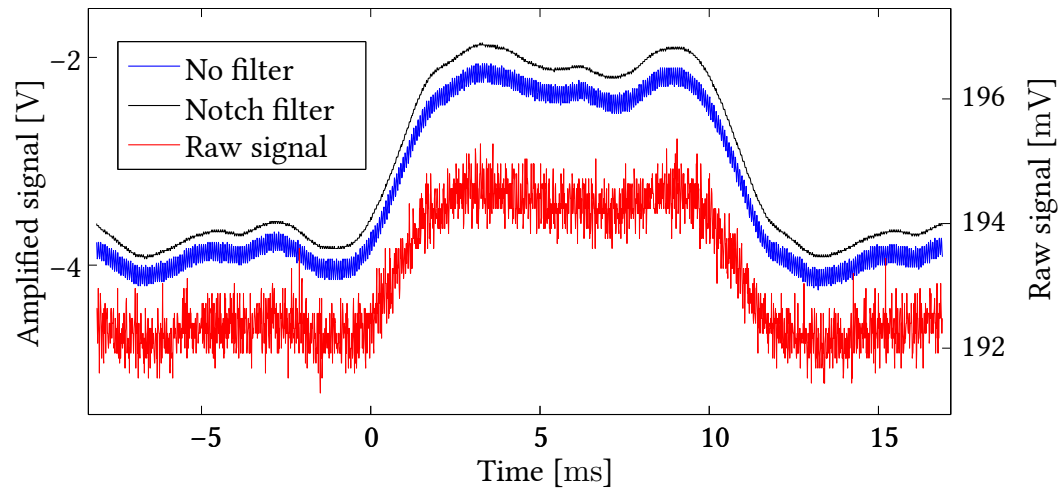
Figure 4.4: Amplification of the raw sensor signal with the circuit described in this section. Plotted in red is the raw signal before amplification, with an amplitude of about $2.5\,\text{mV}$, corresponding to a magnetic field of $2.5\,\text{mG}$. The blue data then shows the amplified signal, now spanning $2\,\text{V}$. The notch filtered signal in black was measured shortly after the two other data sets and one can observe the slight shift of the DC magnetic field.

on the operation of a flux gate sensor one can look into [27]. Usually this is done using low pass filters, but one can also use a notch filter, which is in principle a bandstop filter, blocking a certain frequency f_c.

Figure 4.5 shows several different designs with different advantages and disadvantages. The quality factor Q of a notch filter is the quantity describing the width of the blocking dip in the frequency spectrum. Defining the bandwidth Δf of the filter as the width where the signal is attenuated by $3\,\text{dB}$, one can calculate the quality factor as $Q = f_c/\Delta f$. Q also determines the phase response in the frequency domain, where a higher Q leads to a smaller region in which the phase of the input signal is changed significantly. Notch filters are usually designed using capacitors and inductances to determine the filter frequency like shown in figure 4.5(a). These filters can perform with an excellent quality factor, e.g. a very sharp frequency response. One needs big inductances for filtering higher frequencies with this design. [17] offers another passive solution without inductors, a twin-T filter design that only uses resistors and capacitances and which is therefore easy to assemble and reproduce, as displayed in figure 4.5(b). The quality factor in this design is dependent on the exact matching of all components, but can't exceed 1/4. Especially conductors have to be carefully chosen and should only differ from the chosen values within one percent. The filter frequency is then still fixed. To change it is only possible by switching half of all parts simultaneously. With some modifications to the filter one can adjust the resonance frequency with a single potentiometer as shown in figure 4.5(c). Q is still limited to at maximum 1/4. The only way to improve this is to *bootstrap* the filter, by adding an artificial ground with two operational amplifiers, see figure 4.5(d). With that one can increase the quality up to a certain degree and therefore
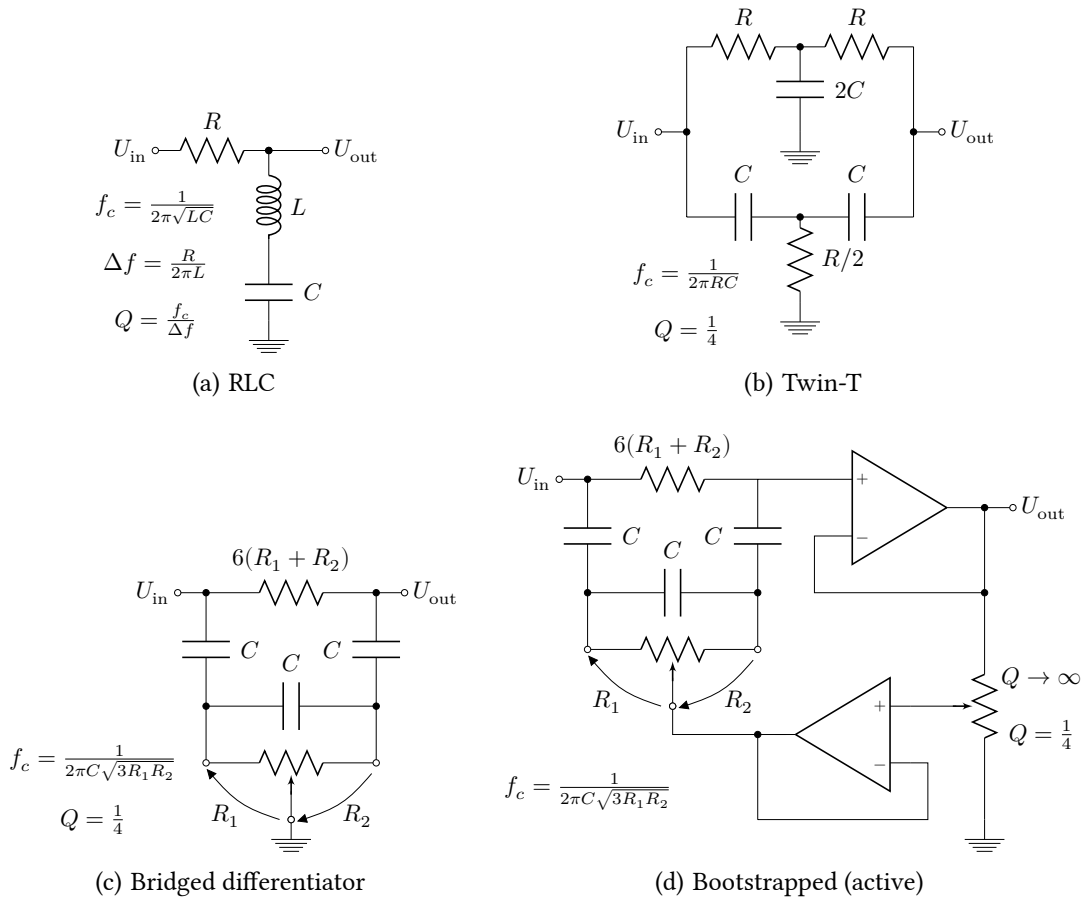
(a) RLC

$f_c = \frac{1}{2\pi\sqrt{LC}}$

$\Delta f = \frac{R}{2\pi L}$

$Q = \frac{f_c}{\Delta f}$

(b) Twin-T

$f_c = \frac{1}{2\pi RC}$

$Q = \frac{1}{4}$

(c) Bridged differentiator

$f_c = \frac{1}{2\pi C\sqrt{3R_1 R_2}}$

$Q = \frac{1}{4}$

(d) Bootstrapped (active)

$f_c = \frac{1}{2\pi C\sqrt{3R_1 R_2}}$

$Q \to \infty$

$Q = \frac{1}{4}$

Figure 4.5: Different notch filter designs used in this thesis, which can be found in the book of Horowitz and Hill [17]. The most basic design 4.5(a) is using an inductance and has a very sharp and deep transmission dip in their frequency response. Getting big and precise inductors may not be desirable, but there is also another passive circuit layout using only resistors and capacitors, shown in 4.5(b). Its disadvantage is the requirement of precise component matching (< 1%), a fixed centre attenuation frequency, while having a rather bad Q factor. The improved circuit design, called bridged differential notch filter, of figure 4.5(c) allows you to change the centre frequency f_c in a certain range, while maintaining the not ideal Q factor value of 1/4. Bootstrapping the ground in figure 4.5(d), to create an artificial signal ground level for the filter finally enables to tune the quality factor. $Q \to \infty$ is however only a theoretical value, limited by the operational amplifier performance.

optimising the phase response of our filtered signal. The notch filter used in figure 4.4 was a twin-T design with not adjustable f_c. It could still reduce the $15.78\,\text{kHz}$ noise significantly, while not changing the phase response of the frequencies interesting for us. For figure 4.1 a tunable twin-T filter was used. Later measurements were done with the bootstrapped, active version.

4.4 Setting up an Active Magnetic Field Regulation

With the sensor being able to detect and characterise the magnetic field noise in our lab, we could turn to the realisation of an active stabilisation. As discussed in section 2.3, for a feedback control system one needs several components:
The actual value we want to regulate is the magnetic field, sensed by the atoms. It is however not possible to directly access this in situ. Therefore the fluxgate sensor described in the previous section is used. It is not possible to place it into the glass cell where the atoms are, but it will be mounted as close as possible. The field gradient of the sensor to the atom trap has to be evaluated with separate measurements.
The other difference to the theoretical measurements is that the voltage signal of the sensor can not always be directly converted into a real magnetic field value, e.g. the $15.78\,\text{kHz}$ noise on the sensor signal seen in figure 4.4 is not originating of a field fluctuation of this frequency, but is an artefact of the measurement. One also has to consider that the sensor core can be magnetised if a field much larger than the saturating field is applied to the sensor. This effect can be reduced by continuous operation and should not play a big role in our system [34]. The amplification of the signal and subtraction of a set value can be done as explained in the previous section, we will however present an alternative way later in this section.
In both set ups, a PID controller PCB designed by Moritz Höfer was used. In addition to the separately adjustable proportional, integral and differential gain values K_P, K_I and K_D, this design features an initial amplified subtraction of either an external set value or an on board reference value. We then will use this control voltage to drive a voltage controlled current source connected to a coil pair. This will give us a feedback controlled regulation loop, as was discussed in section 2.3.

Prototype Regulation

The set up on our test bench is shown in figure 4.6. The sensor was placed in between a coil pair in Helmholtz configuration, with the sensing part centred in the symmetry axis. The signal then gets subtracted by a digitally controlled voltage as described in the previous section. The then amplified signal is sent through a twin-T notch filter to eliminate the $15.78\,\text{kHz}$ noise on it. The PID controller board uses an instrumentation amplifier INA128 that creates the error signal of our regulation. The set signal for the subtraction is either set to zero volt with a $50\,\Omega$ termination or a set voltage produced by a function generator. So there are two separate differential amplification stages, the first one set by the Arduino, which can only regulate in 2 mV steps, and the second one

coil pair   supply/readout   function generator

sensor   differential amplifier   current source   PID controller
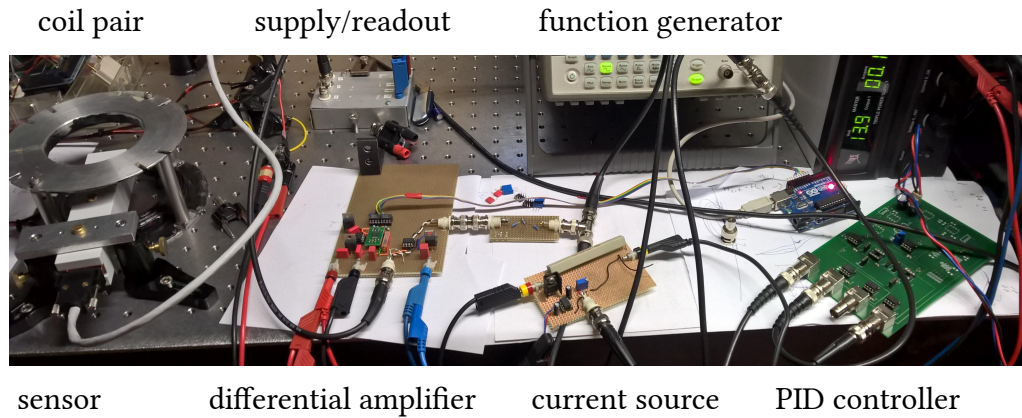
Figure 4.6: The prototype of the active field stabilisation, set up on a test bench. On the left the flux gate sensor is placed between a Helmholtz coil pair. The readout box sends the signal to the differential amplifier described in the previous section. After passing a twin-T notch filter the signal is sent to a PID controller that creates an error signal relative to the set signal of a function generator. The control voltage then is sent to a simple current source, driving the coils and therefore closing the feedback loop. Placed between the function generator and the PID board is the blue Arduino microcontroller, that sets the offset of the differential amplification stage.

that can then adjust for any value in between, depending on the amplification of the first stage.

One has to take care however not to push the instrumentation amplifiers to their output limits. This is set by the supply voltage, which is about $\pm 14\,\text{V}$, since their operation behaviour is not good any more [2]. A good rule of thumb to set the amplification gain was to not let the amplitude of the amplified signal be greater than $10\,\text{V}$. The PID controller itself gave out a control voltage, that was then sent to a simple current source, which used a Darlington transistor. The current then was driving the coils, hence closing the feedback loop.

There were however some drawbacks of this system: for example that the voltage initially subtracted from the signal could only have positive values, since the digital potentiometer was only using a positive voltage reference. The sensor had to be positioned so that the earth magnetic field was positive in its measurement axis. Additionally, the transistor limited the current direction one could control, so the coil connections had to be chosen appropriately. To adjust the control circuit for the active stabilisation we used the methods described in section 2.3. A more than ten-fold noise suppression with good field control could be achieved.

We could also use the parts of this system to test a more or less passive noise reduction with a *feed-forward* system. For that one uses the function generator to put out a sine wave at the frequency one wants to suppress. Then it is connected directly to the current source, driving a magnetic field at that frequency. By controlling the amplitude and the relative phase one then can effectively cancel this frequency component in the
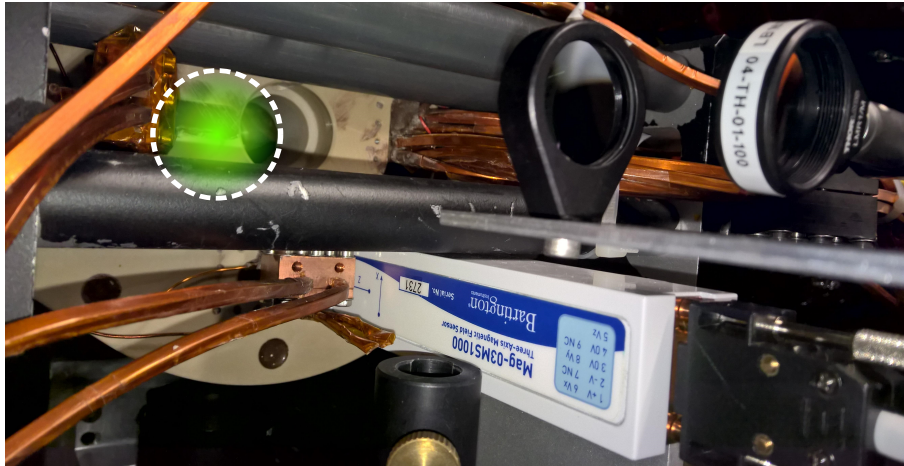
Figure 4.7: Positioning of the magnetic field sensor in the main experiment. Indicated in the top left of the image is the trapping region, which is visually blocked by the large coils for the magnetic trap.

background signal. The optimised measurement is plotted in figure 4.11 and will also be discussed in detail later in section 4.5.

Moving to the Main Experiment

While it was important to be able to test the whole system including all different sub-parts, in the end only measurements on the real system can hold significance for us. The placement of the sensor was chosen such that the optical access to the experiment was not reduced, while still positioned as close as possible to the atom trapping region. The final mounting of the sensor can be seen in figure 4.7. The design of the adjustable, non-magnetic holder was done in the CAD software Solidworks and machined by the internal workshop. The drawings are included in the appendix. It is placed next to the coils for the magnetic trap, but underneath the beam paths of MOT and optical dipole trap. With that position we can be as close as about $8\,\mathrm{cm}$ away from the atoms. Initial concerns about the magnetisation of the sensor during the time of the magnetic trap phase were resolved after discussions with a Bartington support team member [34]. Additionally some test measurements without an active regulation during the trapping. No drifts in the mean magnetic field directly after the trapping phase that could indicate a slow demagnetisation could be observed. It is important however to let the sensor run also in this phase, since the flux gate working principle then partially shields the core from magnetisation.
While the principle of the feedback loop stayed the same there were some changes in the individual parts of the system:

- **Initial amplification:**
  In the prototype design the control system was intended to stabilise the filed in between $\pm 10\,\mathrm{G}$. Therefore the subtraction of a first rough set value had to be
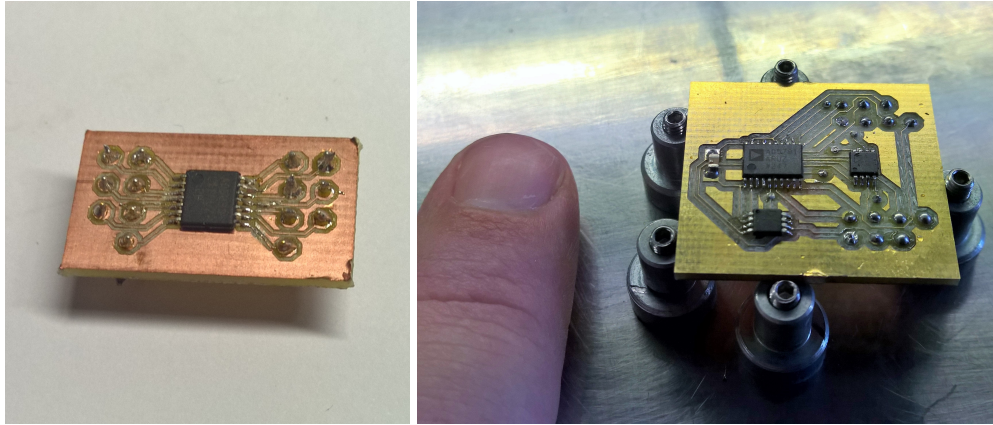
(a) AD5293 adapter board　　　　　　　　(b) AD5791 adapter board

Figure 4.8: Comparison of the 10 bit digital potentiometer AD5293 and the 20 bit DAC AD5791 board designed in this thesis side by side. They were produced with the Kirchhoff institutes's PCB milling machine and the ICs soldered onto it by hand. Both were designed to be plugged into special slots on the PID board and use the same pin configuration. The latest version of the DAC board design is included in the appendix.

done during the initial amplification stage. It had the disadvantage to connect somehow the signal ground to the rest of the circuit, introducing the possibility of noise transmission. The focus of our experiments however shifted more towards a stabilisation around zero magnetic field. So we could introduce an initial differential amplification of signal relative to signal ground, therefore making it (ground-wise) floating. Since we still had to be able to adjust the stabilisation value in some field range of about ± 1 G, the initial amplification of the instrumentation amplifier INA128 was chosen to be small ($R_{\mathrm{gain}} = 5.6 \,..\, 1\,\mathrm{k}\Omega$):

$$G_{\mathrm{INA128}}(R_{\mathrm{gain}}) = 1 + \frac{50\,\mathrm{k}\Omega}{R_{\mathrm{gain}}}$$

$$G_{\mathrm{INA128}}(5.6\,\mathrm{k}\Omega) \approx 9.93$$
$$G_{\mathrm{INA128}}(1.0\,\mathrm{k}\Omega) = 51$$

We used the bigger amplification of $G = 51$ in the main magnetisation axis of previous experiments, e.g. the Feshbach measurements, which was the z - axis. After the observation of stability issues with the simultaneous stabilisation of all three axes, we set the two other gain settings to $G \approx 9.9$. We will discuss the issues of the cross axis talk in the next section.

- **Notch filter:**
  The not activated, bridged differential circuit (see figure 4.5(c)) was ideal to adjust the filter frequency f_c. There was however some slight phase shift of for example

150 Hz oscillations due to the circuit. With the bootlegged filter (see figure 4.5(d)) Q was adjustable again and a high Q was chosen while remaining stable. The components where chosen to match within 3 percent accuracy to the calculated values, to increase performance.

- **PID controller board:**
  The PID board has an instrumentation amplifier for the creation of the error signal by subtraction from either an external signal given in via BNC connector or an internal adjustable voltage. To do so a jumper has to be switched, to connect one of the inputs with a 10 bit digital potentiometer, connected to $\pm$ 10 V from a voltage reference. It is the same chip used in the previous section and can also be controlled over SPI with an Arduino. The step size of this voltage source is then $20\,\mathrm{V}/1024 \approx 20\,\mathrm{mV}$. Including the initial amplification factor of about 10 this corresponds to a step size in the possible magnetic field stabilisation of $2\,\mathrm{mG}$.
  The error signal for the regulation is then produced by differential amplification of the two input signals with an INA128 and a gain of $G = 6$. The different parameters for the PID regulation then can be adjusted by potentiometers, summed and put onto a operational amplifier used as buffer.

- **Digital-to-analogue converter (DAC):**
  As mentioned above, with the previously used digital potentiometer we could only set voltage steps resulting in $2\,\mathrm{mG}$ resolution of our field stabilisation. This is sufficient for the initial stabilisation, but could be an issue in the future: For a stabilisation that can control the field up to a few hundred $\mu\mathrm{G}$ a smaller step size would be needed. For ease of use and minimal invasive changes in the overall system, a new adapter board was designed and milled. It had the same pin configuration as the 10 bit DAC to be able to switch between the new and the old solution quickly. Therefore we were bound to a SPI interface, a $\pm 10\,\mathrm{V}$ reference input voltage and a dual power supply. There are several ICs that would fit, with 14 bit (AD5446), 16 bit (AD5542), 18 bit (AD5781) and 20 bit (AD5791) resolution. The design for all of those ICs would have to be done completely new anyway, so we choose to use the 20 bit DAC from Analog Devices [36]. In combination with a $\pm 10\,\mathrm{V}$ reference voltage source, it can set values in steps of $20\,\mathrm{V}/(2^{20}) = 20\,\mathrm{V}/1048576 \approx 20\,\mu\mathrm{V}$. This chip can change it's output with speeds up $100\,\mathrm{kHz}$, is therefore fast and precise enough even for the most challenging tasks in our experiment. We used the PCB design software EAGLE to create the adapter board, also housing three ultra precise operational amplifiers (AD8676 and AD8675), used as input and output buffers. Figure 4.8 shows the 10 bit digital potentiometer next to the 20 bit DAC. The pin configuration stays the same, however it is very important to check if the power supply is capable of operating the chip. Some resistors on the PID board had to be replaced and bridged by a conductor compared to the original layout, otherwise the supply voltage would be less than the input reference voltage, destroying the DAC. The latest circuit design drawings and the Arduino code to address the 20 bit DAC are attached to this thesis.
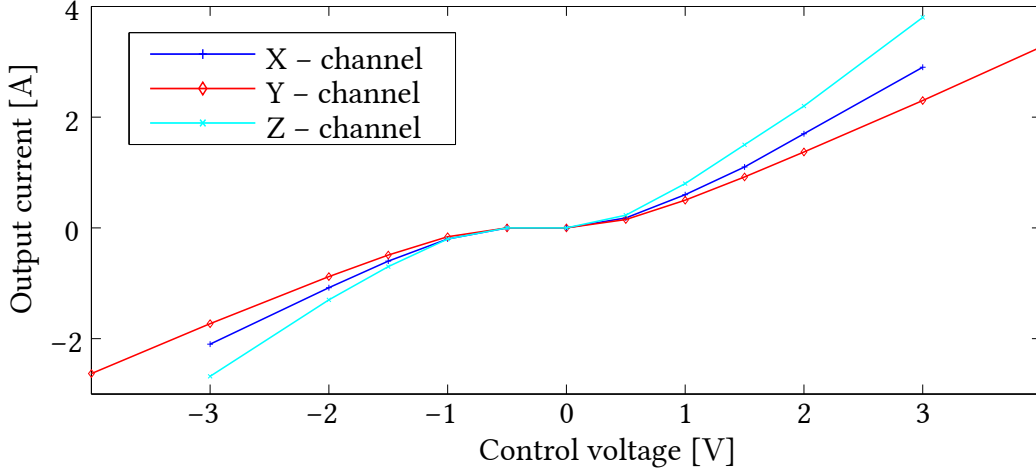
Figure 4.9: Characterisation of the current source used. Plotted are the three independent channels consisting of MOSFETs switched in parallel. The current source is bi-directional, meaning it can supply positive and negative currents. Close to zero one can see the non-linearity in the response of the MOSFETs. This region should be avoided during the stabilisation.

- **MOSFET transistor array:**
  The control signal then goes directly into a current source, an array of MOSFET transistors, we call the *passbank*. It was designed by Jürgen Schölles and build by Fabián Olivares. Four bi-directional channels with 5 power MOSFETs in parallel for each direction can be controlled individually and give out more than $\pm 10\,\mathrm{A}$ of current. The output curve, with input control voltage against output current for all three channels we used is plotted in figure 4.9.

- **Offset coils:**
  The coils where already build in and used in the experimental sequence to create small offset fields when needed. All three square coil pairs are centred around the trapping region of the atoms and described in table 4.2.

coil axis	number of windings	dimensions [m × m]	separation distance [m]
X	14	1.212×0.604	0.89
Y	26	0.825×0.604	1.23
Z	9	0.825×1.172	0.63

Table 4.2: Dimensions of the offset coils used for the stabilisation.

They were characterised in combination with the current source channels they were connected to, which is plotted in figure 4.10. These coils were designed to
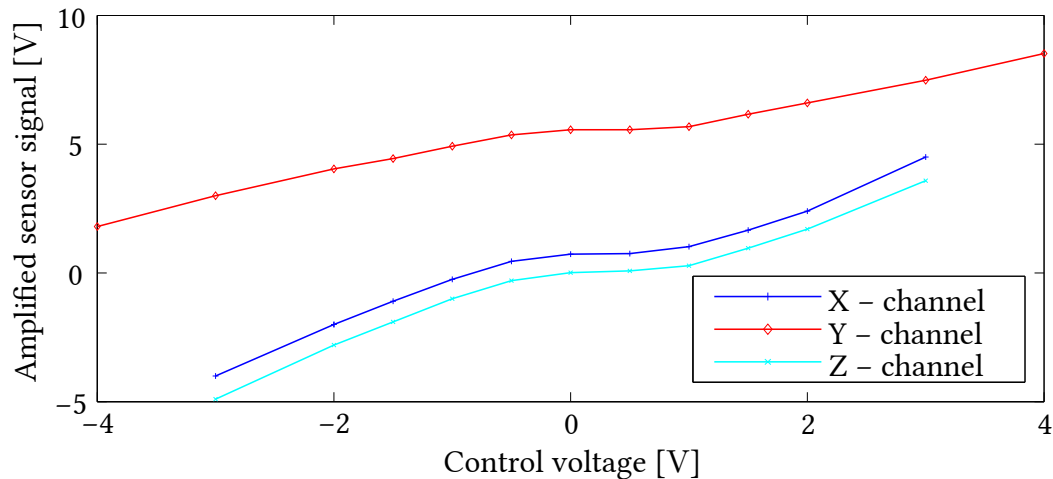
Figure 4.10: Characterisation of the coils used for the stabilisation. For comparability the same sensor gain of $\approx 9.9 \cdot$ V/G was used for all three sensor axes. With no current flowing a field of about 0.5 G is measured in y direction, which is about the earth magnetic field's magnitude [12].

give a very homogeneous magnetic field at the centre of the atom trap, which is unfortunately not precisely where the sensor is positioned. Hence the field at our sensor is not completely homogeneous and a field applied along one axis can change the field in the other two. We will refer to this effect as crosstalk.

- **Relay box:**
  Since both the current source and the coils were already in use in the previous sequence, some kind of switch between the old offset field, set by the sequence control, and the new active regulation had to be implemented. A relay box was build, that switches the whole set of input signals for the MOSFET array on the digital high of a TTL signal. With that no changes to the original sequence had to be made and the active stabilisation can be switched on or off timed with the experimental sequence control.

After getting all components of the regulation working again the feedback loop could be tested and adjusted.

4.5 Results and Adjustments

With the system described, we could gain good control over the behaviour of the magnetic field at the place of our sensor. With using a feed forward signal on our current source, as described in section 4.4 we could nearly completely eliminate any background field oscillations of the set frequency. For that we triggered a function generator to the power line 50 Hz and adjusted its own 50 Hz output phase and amplitude such that the oscillations in this signal were extinguished. This was possible because the 50 Hz noise
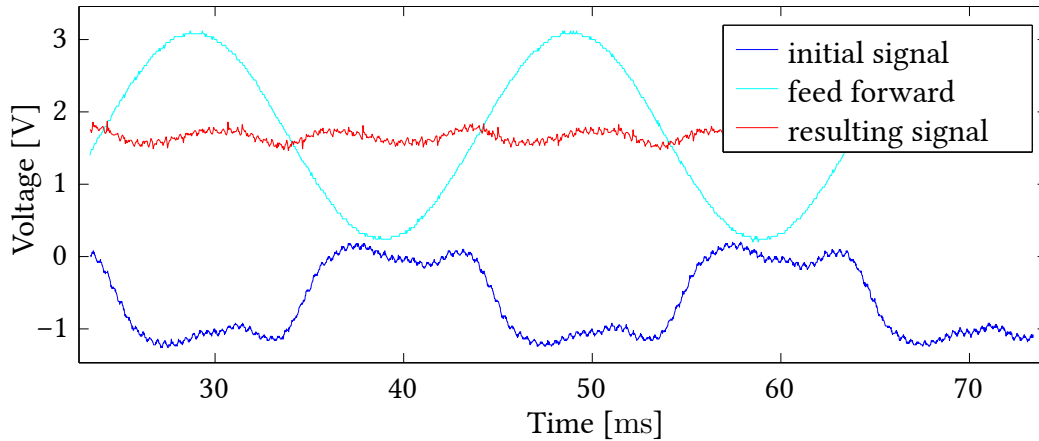
Figure 4.11: Feed forward suppression of 50 Hz background magnetic field oscillations measured with the first test system. Plotted in blue is the initial signal which is already subtracted by an offset and amplified. In cyan the feed forward signal is shown, which is used to control the current in the coils. Since it was not possible to drive the current in two directions, an offset had to be included, which can also be seen in the resulting sensor signal in red. The noise in this plot is due to aliasing of the 15.78 kHz which was not optimised at that time.

stems from the line power, so even higher harmonics will always have a fixed phase relation. The result of optimising these settings can be seen in figure 4.11. With nearly no 50 Hz background present after the feed forward current control, the amplitude of the background oscillations has been reduced by a factor of 4! What is left is mostly 150 Hz noise and other higher harmonics of the 50 Hz background. With adding and adjusting these frequencies into the feed forward signal a further reduction could be possible. This method is robust since it doesn't rely on active feedback, but rather on the passive stability of the system, which is given by the triggering to the power line. It needs however some time to adjust and has to be checked in certain intervals of time for any change in the background oscillations. As seen in our measurement one has to at least additionally take the 150 Hz background also into account. For three axes, with two frequencies each this results in an optimisation of 12 parameters, which would have to be done repeatedly in a real experiment.

While feed forward can effectively suppress the background oscillations it cannot be used to set the magnetic field to a certain value, e.g. zero. With *passive* stabilisation like monitoring and adjusting the current flow through the coils, one can change the magnetic field generated by the coils to a certain value. This doesn't react to any slow drifts of the background DC field nor has day to day stability. That's where a sensor based stabilisation is ideal. By using the feedback system described in this chapter one can adjust the regulatory feedback control with only the PID factors in the controller. We already explained the adjustments possible in section 2.3.
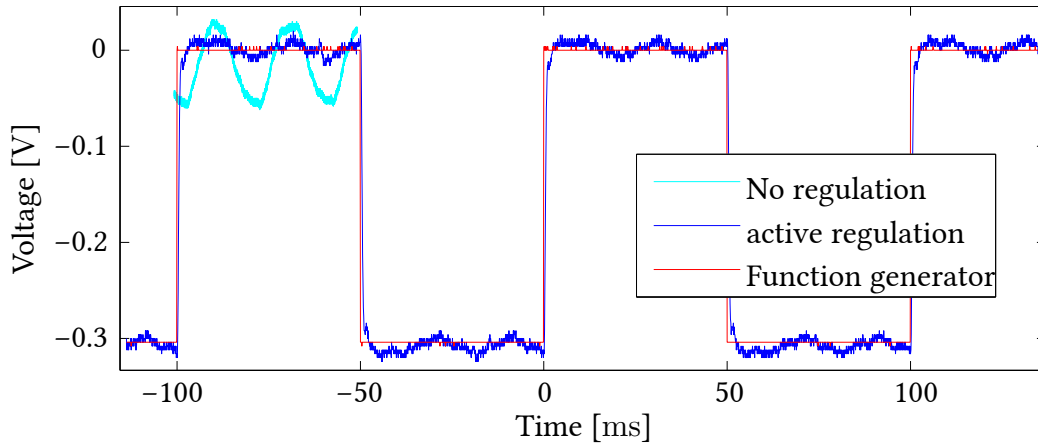
Figure 4.12: Active feedback control of the magnetic field in z direction with a long integration time T_I. The initial PID control settings were chosen according to the Ziegler-Nichols method, see section 2.3. In cyan the background oscillation without active stabilisation are added on top of one step for comparison. With the active regulation set to the function generator signal in red the sensor signal in blue exhibits less fluctuations.

Choosing for example a long integration time T_I compared to the typical noise frequencies and a suitably large proportional gain K_P will lead the regulated parameter, e.g. the magnetic field, to reach a set mean value, while suppressing the background oscillations only slightly. This setting is demonstrated in figure 4.12. The background oscillations are suppressed by a factor of three, however still clearly visible in the regulated signal. More importantly one can regulate the mean field accurately to a set value. Since the regulated value, the magnetic field, doesn't have an overshoot but rather approaches it smoothly we can say that this is a very stable regulation. The measurement was done in the z axis, which is the atoms' main magnetisation axis in our experiments. For this and the following measurements the other two axes where set to similar regulation parameters, with a not changing set value of zero. This granted us needed stability of the whole regulation due to magnetic field cross talk of the coil axes.

In order to optimise the field response and to minimize any oscillations of the background fields one has to use a shorter integration time T_I. This leads to high overshoots, which can be reduced by setting an appropriate value for the differential time T_D. There exist several methods on how to optimise the system response, however most of them are either suited for following a small signal change around a fixed value or for fast changes in the set value, the whole system has to react to. We would like to have something in between, being able to switch on the regulation in a very short time of few milliseconds and then following and cancelling accurately the background oscillations. In figure 4.12 we showed a stable and slow regulation, while in figure 4.13 we optimised the controller to follow and suppress the background oscillations. The regulation behaviour is very sensitive, leading to a big overshoot, which however reaches the desired value after only
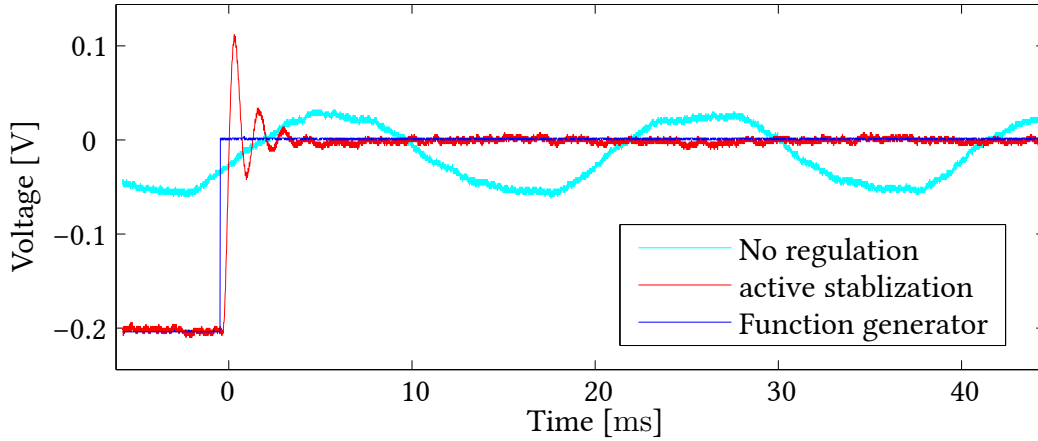
Figure 4.13: Short integration time T_I and a significant differential part in the active feedback control of the magnetic field in z direction. The settings are chosen to maximize the suppression of background fluctuations while still being able to switch very fast to a desired mean value.

a few short oscillations. Here we have a reduction in the noise amplitude of a factor of eight while the signal settled to the new set value after around 25 ms. While the background noise is suppressed we have to accept that the other factors like overshoot and settling time are not ideal.

Magnetic Field Measurements with Atoms

We could show that the stabilisation of the magnetic field worked at the place of the sensor, but the goal is to do so at the place of the atoms in the optical trap. Ideally the field there is the same as at the sensor, and all changes with the offset coils are homogeneously translated there too. It can be checked if that is true by probing the magnetic field sensitive hyperfine level splitting with a microwave. For that a frequency sweep of the microwave gets sweep discrete intervals is performed.

While a slight difference had to be expected we were surprised that by regulating the sensor output in the x and y axis to zero, the atoms experienced a total magnetic field of about 320 mG. By adjusting the set values of two axes active stabilisation (x and y) and a passive current stabilisation in the z axis, we could achieve a total field 3.7(28) mG measured with the atoms. The parameters chosen are shown in table 4.3:

Table 4.3: Parameters of the active regulation used to achieve 3.7(28) mG total field at the atoms.

coil axis	set value (active) [V]	current (passive) [A]
X	1.625	-
Y	-2.54	-
Z	-	1.456

We can also calculate the field at the place of the sensor for these values, but since there is crosstalk of the coils producing the field, summing them corresponding to the single axis characterisation would not give an accurate value for the difference between the sensor and the place of the atoms. This value $\sqrt{B_X^2 + B_Y^2 + B_Z^2} = 338\,\mathrm{mG}$ however is in qualitative agreement with the previous measurement.

Technical and Conceptional Problems

It was already mentioned that two of the three axes had to be set to a slow, but stable regulation. This was due to the big crosstalk of the different regulation axes. By setting a certain value for the y axis, one automatically could see a significant change in the other two directions. Optimising this is not easy, since there are some permanent magnets (ion pump magnets) and magnetizable objects (optical table) close to the atoms. The possible positioning in the experiment is additionally very limited. If now the regulation of two axes is coupled strongly enough and the time scales of the regulation comparable to the overshooting frequency then the regulation becomes very unstable. While in most cases the regulation of one axis adds the signal it wants to suppress to the other axis and vice versa, it also often leads to the regulation not being able to settle at the desired field at all. Hence only one axis can be regulated with a fast response shown in figure 4.13, which gives the best background noise suppression. This is sufficient for non zero total magnetic fields, since the vector addition of the spacial field components is quadratic. Therefore the noise on the smaller axes has less and less influence on the total field's noise with bigger total magnetic field. We want to stabilise however to zero magnetic field, where all three field components add equally to the total noise. This limits the total suppression of background field fluctuations possible with this system.

An additional issue discovered during the optimisation of the regulation parameters was connected to the subsequent use of two instrumentation amplifiers. Both can amplify small signals precisely with a set gain, but their operation with bigger, already amplified signals has strong constrains. Due to the internal construction of the INA128 the linear input range is limited to around $1.4\,\mathrm{V}$ below the positive and $1.7\,\mathrm{V}$ below the negative power supply voltage [2]. If this threshold is crossed with one input, then the internal operation of the instrumentation amplifier is compromised. This leads to the output of the maximal negative or positive voltage. The behaviour is not easily predictable due to the complex internal structure of the instrumentation amplifier. In the usual operation one knows the possible, mostly small input ranges and adapts the supply voltage so that this never happens. If two amplifiers are however set in series with a big initial amplification it is easily possible that the first amplifier's output is already at the maximal output voltage. So if the second amplifier uses the same power supply then the signal is out of the input range. We observed several times the breakdown of the regulation of even just one axis if the initial amplification was too high. The PID control output was set to the maximum due to this wrong input and the resulting sensor signal grew even bigger, producing a lock down of the feedback circuit. Due to that reason all initial amplification was set to a very small value mentioned in the previous section.

Initial thoughts for doing the subtraction *after* the initial amplification was because it

was expected to regulate the field around the zero measurement value of the sensor. As was discovered later this was not the case since there is a difference in total field of about $320\,\text{mG}$ form the sensors position to the atom trap location. This in combination with the issue of the input range was leading to even less possible initial amplification. The other reason to do the subtraction after an amplification was due to the possible controllable step voltage of $2\,\text{mV}$ with the 10 bit digital potentiometer. But with the use of the 20 bit DAC one can now put out and use for subtraction $40\,\mu\text{V}$ steps, which would correspond to $40\,\mu\text{G}$ in the not amplified sensor signal. Therefore the subtraction of the set signal should happen before or with the initial amplification.

5 Conclusion and Outlook

In the framework of this thesis several additions to the experiment were tested and implemented. We would like to use the summary of the results to point out future, possible improvements to the system.

Imaging System

The new camera software allows for an easy usage and good control over all camera parameters relevant to the experiment. A new graphical user interface was presented and its internal function explained. By systematically checking imaging parameters we could achieve significant improvements in image quality and signal strength for lithium absorption imaging. This lead to an improved detection capability of small population changes, which is now not limiting the experiment anymore.

The high hopes of achieving very low delay times between the atom and the reference picture were however not fulfilled. Here we showed that even with pushing the camera to its limits and reducing resolution and size, we could not get close to time scales of only a few milliseconds. There are however cameras and techniques that can reach this regime. One can use for example frame-transfer CCD cameras, where half of the CCD chip is covered with a mask, thus making the pixels there insensitive to light. After the not concealed CCD pixels have been exposed to light and accumulated electrons in their potential wells, they get collectively shifted as one frame into the masked region, where they can be read out. This shifting process is done in less than a millisecond [1], whereas the readout itself may take much longer. With good shutter timing, one can already expose the sensitive region for a second picture, while the other frame, now stored under the mask, gets read out. It is therefore possible to capture two images with a time separation of around one millisecond. Since any shift in the fringe pattern of the initial images is then too slow to play a relevant role for the atom number calculation, the resulting, calculated pictures would have an even better quality, while maintaining maximal resolution and size.

Another project that started during the course of this thesis is to exchange the commercial microscope objective below the glass cell with a customised lens stack. This project, addressed by Alexander Mil, will aim at reaching diffraction limited imaging at a high numerical aperture.

Magnetic Field stabilisation

An active magnetic field stabilisation was built and tested and we could gain good control over the magnetic field at the position of the sensor. We achieved an at least four-fold reduction in background oscillations in each direction, while being able to set specific magnetic fields to a very good accuracy. The big gradients of the magnetic field close to the atom trap make it however difficult to translate this result to a field stabilisation at the position of the atoms. While the prototype was designed to stabilise to set signals around zero, we discovered that we had to include the absolute field difference of about $320\,\mathrm{mG}$ from the sensor to the atoms into our regulation scheme. The two stage amplification with the set signal subtraction after the initial amplification was therefore a limiting component. The 20 bit DAC circuit, designed during this thesis, was however not included in the initial planing. Using this can give us the possibilities for an even better suited design. With the new digitally controlled voltage step size of $40\,\mu\mathrm{V}$ we have the possibility to directly subtract the set signal within the initial amplification step. A second amplification stage is hence not needed anymore. This can give us back the feasible amplification factors of $G = 100$ to 1000 and a more stable regulation. A wanted feature would be to include a digital controlled potentiometer as the gain resistor R_G, enabling us to digitally control the amplification of the signal. With this added to the circuit one could ramp up the sensitivity of the sensor during the turn on phase of the stabilisation, allowing for a very controlled and smooth transition from unregulated to regulated fields.

Since the difference in sensor to atom trap magnetic field is so big, cancelling the oscillations seen by the sensor don't necessarily reduce the field fluctuation at the atom position. We should therefore also use a scheme to reduce the background oscillations with a more or less passive stabilisation: The feed forward noise cancellation was very promising in their ability to suppress the background oscillations at a specific frequency. Using a Ramsay spectroscopy sequence with microwave pulses and triggering the whole experiment to the $50\,\mathrm{Hz}$ power line would allow for a measurement of these oscillations directly with the atoms. This then could be used to optimise the feed forward signal which could be added to the controller voltage. Hence one could actively stabilise to an absolute magnetic field reference, the sensor, while still passively reducing the fluctuations.

A Appendix

A.1 Closed-Loop Transfer Function

Open Loop Regulation

For understanding the regulatory behaviour of a system it is interesting to look at the transfer function of the controller and the system-plant combined, giving the so called open loop gain O:

$$O(s) = \frac{\mathcal{L}\{a(t)\}}{\mathcal{L}\{e(t)\}} = H(s) \cdot G(s) \tag{A.1}$$

For example thinking of a controller using only PI parts and a system-plant modelled as a low pass with some actuation, as in equation (2.51) we can calculate:

$$O(s) = H(s) \cdot G(s) = \frac{A}{1 + T_L s} \cdot K_P \left(1 + \frac{1}{T_I s}\right) \tag{A.2}$$

$$= \left(\frac{1 + T_I s}{1 + T_L s}\right) (A \cdot K_P) \left(\frac{1}{T_I s}\right) \quad \text{with:} \quad T_I = \frac{K_P}{K_I} \tag{A.3}$$

Loop Regulation

We would like to calculate the frequency dependant reaction of the actual value $a(t)$ to a change in the set value $p(t)$:

Connecting the previous results and assuming that the sensor is ideal and instantaneous $r(t) = a(t)$ we can derive the following:

$$\mathcal{L}\{a(t)\} = H(s) \cdot \mathcal{L}\{u(t)\} = H(s)G(s) \cdot \mathcal{L}\{e(t)\} \tag{A.4}$$

$$\text{with:} \quad \mathcal{L}\{e(t)\} = \mathcal{L}\{p(t)\} - \mathcal{L}\{a(t)\} \tag{A.5}$$

$$\Rightarrow \mathcal{L}\{a(t)\} = H(s)G(s) \left[\mathcal{L}\{p(t)\} - \mathcal{L}\{a(t)\}\right] \tag{A.6}$$

Now we can define the so called *closed-loop transfer function* $C(s)$:

$$C(s) = \frac{\mathcal{L}\{a(t)\}}{\mathcal{L}\{p(t)\}} = \frac{H(s)G(s)}{1 + H(s)G(s)} \tag{A.7}$$

This completely describes the behaviour of our system and can be used to model most regulation in our lab. The usually desired outcome is that $C(s) \approx 1$, meaning that the system directly follows any change in set value!

A.2 Circuit Board Designs
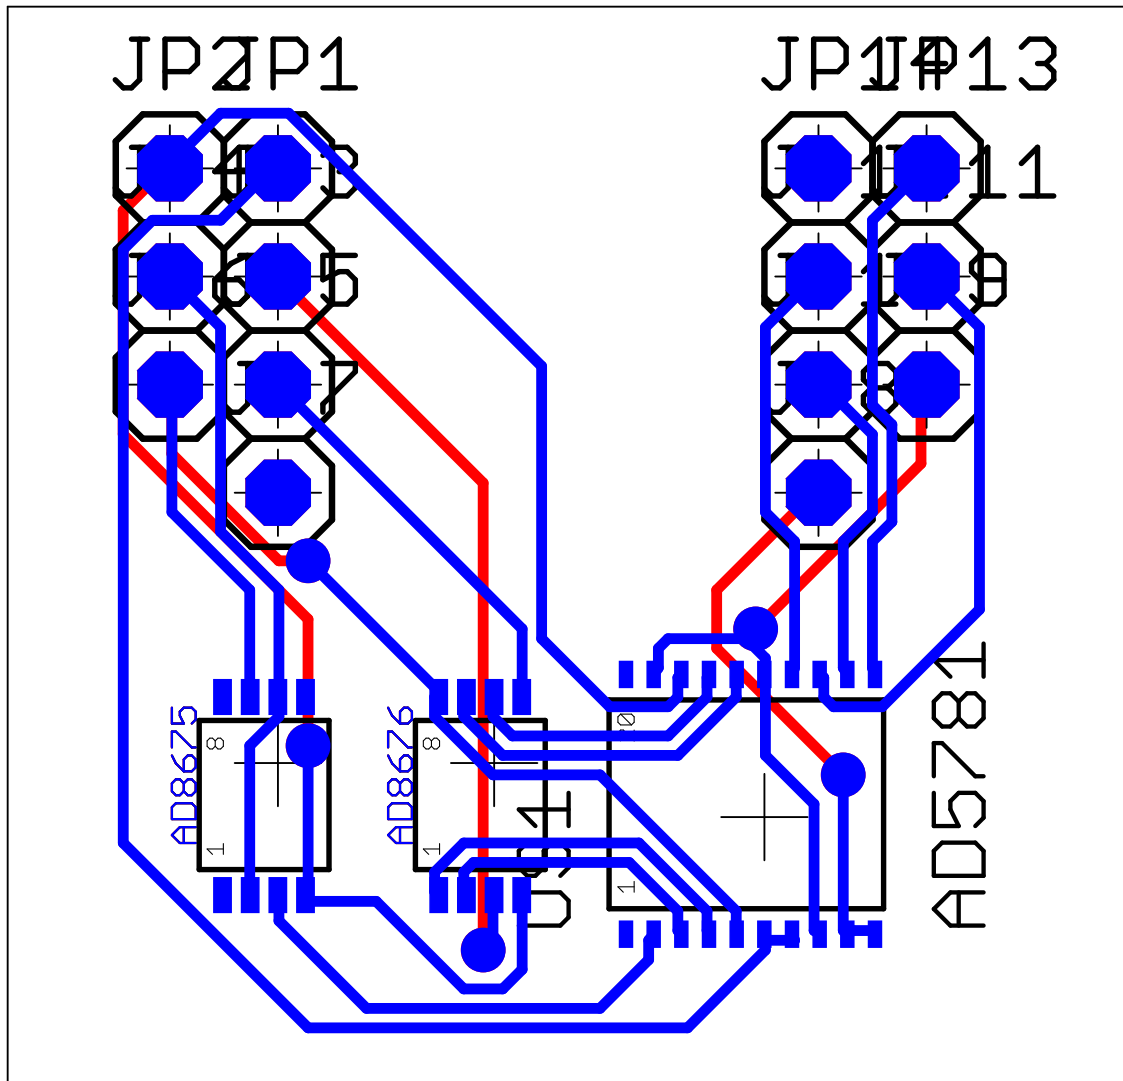
Figure A.1: Circuit design of the 18 Bit digital-to-analogue converter adapter board for the AD5781. The 20 bit DAC AD5791 is interchangeable with the AD5781 and was used in this work. The top layer is shown in blue, whereas the bottom layer is shown in red. The AD8676 dual operational amplifier is used as reference voltage input buffer, whereas the AD8676 is buffering the output voltage. It was created using the EAGLE PCB design software and used with the institute internal milling machine.

A.3 Technical Drawings
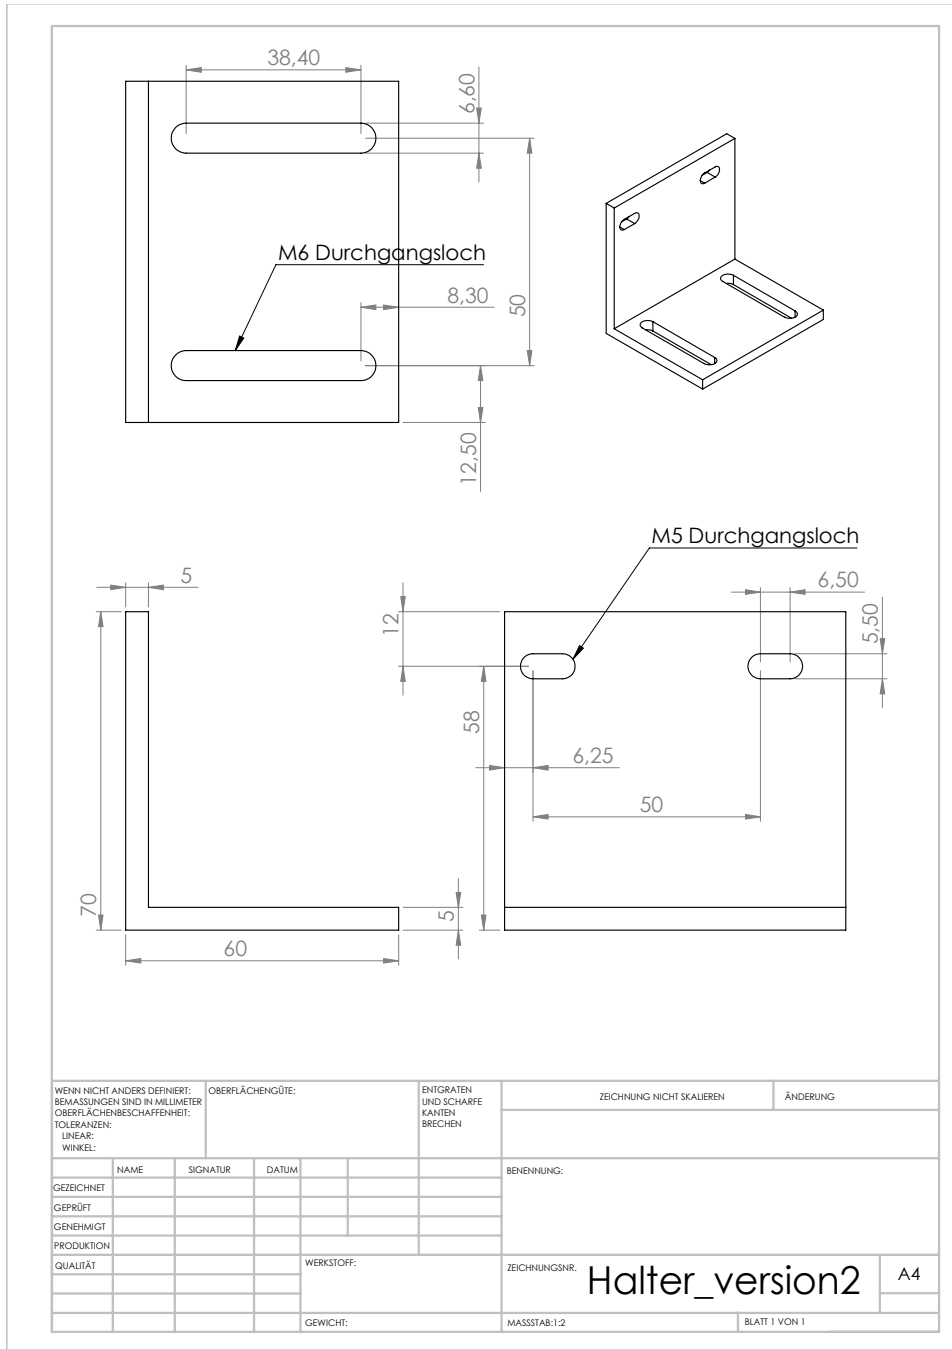
Figure A.2: Technical drawing for the sensor holder for the positioning close to the atoms. It's made out of aluminium and was machined by the institute workshop.

A.4 Arduino Codes

Magnetometer Readout

```
/*
 GY-273 module with Honeywell HMC5883L 3-Axis Digital Compass IC
 */

// math library for rounding
#include <math.h>

// include the I2C library
#include <Wire.h>

// include the library for the LCD screen
#include <LiquidCrystal.h>

/* lcd screen initialisation
 * LCD RS pin to digital pin 48
 * LCD Enable pin to digital pin 46
 * LCD D4 pin to digital pin 47
 * LCD D5 pin to digital pin 49
 * LCD D6 pin to digital pin 51
 * LCD D7 pin to digital pin 53
 * LCD R/W pin to ground
 */
LiquidCrystal lcd(48, 46, 47, 49, 51, 53);

//_____
// ------Global variables------
//----------------------------------------------------------------
// -----------------------------fast adjustment HMC5883L-----------------------
// Gain setting choosen
byte sample = 3;        // Sampled average: 1 ; 2 ; 4 ; 8
//(sample = 3 => 8 values averaged)
byte rate = 6;           // rates in Hertz:
//0.75 ; 1.5 ; 3 ; 7.5 ; 15 ; 30 ; 75; (rate = 3 => rate = 7.5Hz
byte pos_bias = 0;    // Settings for positive bias coils (1 for on)
byte neg_bias = 0;     // Settings for negative bias coils (1 for on)
byte single = 0;       // single meassurement if set to 1
byte gain = 3;         // global gain setting => can be overwritten in the
//configuration function , multiplicative factor defined below
float gain_array[] = {
  0.73, 0.92, 1.22, 1.52, 2.27, 2.56, 3.03, 4.35};
// amplification for gain setting: 100 LSB correspond to the gain in mG
//----------------------------------------------------------------
//----------------------------------------------------------------

// Possibility to use a digital averaging over the values
//to get a lower statistical variation
int N_averaging = 50;
int wait = 0; // Wait time between every sample , set to 0 if not wanted
```

```
// self calbiration on or off
boolean selfcal = false; // currently NOT WORKING

// We declare our variables for the different values of the magnetic field
float x = 0;
float y = 0;
float z = 0;
float total;

// declaring the sensitivity adjustments for each sensor (default = 100)
float x_m = 1.00;
float y_m = 1.00;
float z_m = 1.00;

// declaring the offsets of each axis (default = 0)
float x_o = 00;
float y_o = 00;
float z_o = 00;

// helping variable to use for getting the bytes
float x_uc = 0;
float y_uc = 0;
float z_uc = 0;

float x_1;
float y_1;
float z_1;

float sum_x = 0;
float sum_y = 0;
float sum_z = 0;

int x_show;
int y_show;
int z_show;


byte x1;
byte x2;
byte y1;
byte y2;
byte z1;
byte z2;

// value to be transmitted
int x_unmod;

// check the bnc input
boolean digital_input;

// input charakter, defines communication with the arduino
```

```
char input;

// array length for transfer to PC is defined here
const int ArraySize = 750;   // 2625 -> 35s
int DataArray[ArraySize];

//_____
void setup()
{
  // start the lcd interface with 16 characters and two lines
  lcd.begin(16, 2);
  // start the I2C bus service
  Wire.begin();
  // set the initial values for the configuration of the chip
  adjust_mag();

  // Start Communication with PC and MatLab
  Serial.begin(115200);
  // Serial.println("testworks");

  pinMode(30, INPUT);    // signal to turn on
  pinMode(31, OUTPUT);   // GND
  digitalWrite(31, LOW);

  Serial.println(0);   // startup is done and signals can be recieved
}


//_____

// Main Programm
void loop()
{
  // initalize the internal variables for digital averaging, first sum up,
  // then devide through the no of values
  sum_x = 0;
  sum_y = 0;
  sum_z = 0;

  // digital sampling loop, reads out a number of values, then averages
  // over them to get a lower statistical uncertainty (we assume it is
  // better than the internal averaging)
  for (int i = 0; i < N_averaging; i++) {
    // data_ready();
    get_data();       // requests set of data, including scaling
    sum_x = sum_x + x;
    sum_y = sum_y + y;
    sum_z = sum_z + z;
    delay(wait); // delay the cycle to get let the data be more readable
  }

  // set the x y z values on the average
  x_1 = sum_x / N_averaging;
```

80

```
y_1 = sum_y / N_averaging;
z_1 = sum_z / N_averaging;

//get the correction factors for our scaling
total = sqrt((x_1 * x_m)*(x_1 * x_m) + (y_1 * y_m)*(y_1 * y_m) +
(z_1 * z_m)*(z_1 * z_m));
x_show = round(x_1 * x_m);
y_show = round(y_1 * y_m);
z_show = round(z_1 * z_m);
show_data(x_show, y_show, z_show); // updates the LCD

// check if pc sent a request to transfer the data collected
if (Serial.available()) {
  input = Serial.read();

  // loop until other input comes
  while (input == 'w') {

    // now get all data for the array to transfer
    lcd.setCursor(0,0);
    lcd.write("Data Collection ");
    lcd.setCursor(0,1);
    lcd.write("                ");

    // now wait for the signal of on a digital input
    boolean digital_input = true;
    while (digital_input) {
      digital_input = digitalRead(30);
    }

    // for determining the length of the data taken
    long int tick = micros();

    //Serial.println("start");
    //for (int i = 0; i < ArraySize; i++) {

    // continue as long as atom beam sutter is open
    while (!digital_input) {

      get_data();

      // directly print it to matlab
      Serial.print(x);
      Serial.print(' ');
      Serial.print(y);
      Serial.print(' ');
      Serial.print(z);
      Serial.print("; ");
      digital_input = digitalRead(30);
      delayMicroseconds(10221);

    }
```

```
      long int tack = micros ();
      Serial.println ();

      long int delta = tack − tick;
      //Serial.println(delta);

      //check if data accqisiton was aborted
      if (Serial.available()) {
        input = Serial.read();
      }
    } // end of while loop

    // set back to normal meassurement mode
    adjust_mag ();
  }
}

//_____

// pdates the LCD screen with the newes values of x,y,z and total

void show_data(int x, int y, int z){

  lcd.setCursor(0,0);
  lcd.write("x:␣␣␣␣␣␣");
  lcd.setCursor(3,0);
  lcd.print(x);

  lcd.setCursor(8,0);
  lcd.write("y:␣␣␣␣␣␣");
  lcd.setCursor(11,0);
  lcd.print(y);

  lcd.setCursor(0,1);
  lcd.write("z:␣␣␣␣␣␣");
  lcd.setCursor(3,1);
  lcd.print(z);

  // get the absolute value

  lcd.setCursor(8,1);
  lcd.print("T:␣");
  lcd.print(total);
  lcd.print("␣␣␣");
}

//_____

// This function writes a byte into the 3 configuration registers
// Register A: sample rate, averages, bias fields on (positive/negative)
// Register B: gain
// Register M(ode): single or continuous meassurement
```

```
void adjust_mag (){

  byte RegA  = (sample << 5)| (rate <<2)| (neg_bias <<1) | pos_bias;

  byte RegB = gain << 5;

  byte RegM =  single;

  // set the configuration register A
  Wire.beginTransmission(0x1E);   // the adress is 0x1E for the magnetometer
  Wire.write(B00000000);          // set to write first register
  Wire.write(RegA);               // -> configuration register A
  Wire.endTransmission();

  // set the configuration register B
  Wire.beginTransmission(0x1E);   // the adress is 0x1E for the magnetometer
  Wire.write(B00000001);          // set to write second register
  Wire.write(RegB);
  Wire.endTransmission();

  // set the mode register
  Wire.beginTransmission(0x1E);  // the adress is 0x1E for the magnetometer
  Wire.write(B00000010);          // set to write second register
  Wire.write(RegM);        // 1.bit HighSpeed I2C 0-> off 7.-8.bit operating
  // mode 00 -> set to continous meassurement mode
  Wire.endTransmission();
}

// _____

// This function calls the data registers and reads them out. First it has
// to set the device pointer to the right address (data X A => position 3)
// and then requests 6 bytes (1.X 2.Z 3.Y)
// All values are then written on the global variables x,y,z and total
void get_data (){

  // wait until data is ready (wait in ms -> optional, default = 0)
  data_ready();

  // get the data -> go to the first data register
  Wire.beginTransmission(0x1E);
  Wire.write(B00000011);          // register 3 -> X   5 -> Z      7 -> Y
  Wire.endTransmission();

  // Request 6 byte from the address of the magnetometer. After every
  // read process, the register is shifted by one -> one time read and
  // get all of it
  Wire.requestFrom(0x1E, 6);
  x1 = Wire.read();
  x2 = Wire.read();
  z1 = Wire.read();
  z2 = Wire.read();
```

```
  y1 = Wire.read();
  y2 = Wire.read();

  // uc stands for uncorrected -> different scaling
  x_uc = (x1 << 8) | x2;
  z_uc = (z1 << 8) | z2;
  y_uc = (y1 << 8) | y2;

  // scaling with two factors: the total scaling due to different gain
  // settings tot_m, and the different scaling for each axis x_m and
  // adjusts for an offset x_o

  x = (x_uc * gain_array[gain]) + x_o;
  y = (y_uc * gain_array[gain]) + y_o;
  z = (z_uc * gain_array[gain]) + z_o;

  return;
}

//_____

// This function calls the data registers and reads them out.
// First it has to set the device pointer to the right address
// (data X A => position 3) and then requests 6 bytes (1.X 2.Z 3.Y)
// All values are then written on the global variables x,y,z and total
void get_data_single(){

  // wait until data is ready (wait in ms -> optional, default = 0)
  data_ready();

  // get the data -> go to the first data register
  Wire.beginTransmission(0x1E);
  Wire.write(B00000011);          // register 3 -> X    5 -> Z      7 -> Y
  Wire.endTransmission();

  // Request 6 byte from the magnetometer. After every read process,
  // the register is shifted by one -> one time read and get all of it
  Wire.requestFrom(0x1E, 2);
  x1 = Wire.read();
  x2 = Wire.read();
  //z1 = Wire.read();
  //z2 = Wire.read();
  //y1 = Wire.read();
  //y2 = Wire.read();
  x_unmod = (x1 << 8) | x2;

  // delayMicroseconds(6000);

  return;
}

//_____
```

```
// function to check if new data available with the a defined loop wait time
void data_ready() {
  boolean Nready = true;   // this is set to false when data is ready
  while (Nready)
  {
    // check if there is new data -> go to status register
    Wire.beginTransmission(0x1E);
    Wire.write(B00001001);          // register 9
    Wire.endTransmission();

    // request 1 byte from the magnometer device
    Wire.requestFrom(0x1E, 1);
    byte status = Wire.read();
    if (status & (1 << 0))
    {
      Nready = false;
    }
  }
  return;
}
```

10 Bit Digital Potentiometer with Opto-Couplers

```
/*
   Digital Poti Control, following implementation by Tom Igo
   and Heather Dewey-Hagborg
   Adaptions by Arno Trautmann, 2012-2013
   Modifications by Marcell Gall, 2015
   10 bit digital poti AD5293, with SPI over optocouplers
*/

// inslude the SPI library:
#include <SPI.h>
#include <math.h>

// set pin 10 as the slave select for the digital pot:
const int slaveSelectPin = 10;

int input;

void jumpto_old(int jval){
    Serial.print("Jump to (old): ");
    Serial.println(jval);
    digitalWrite(slaveSelectPin,LOW);
    SPI.transfer(6146 >> 8);  //command bits to update calibration, allow change
    //of wiper position, performance mode, for faster switch to normal
    SPI.transfer(6146);
    digitalWrite(slaveSelectPin,HIGH);
    digitalWrite(slaveSelectPin,LOW);
    SPI.transfer(jval+1024 >> 8); // 1024 is setting one command bit
    //to go to the write mode
    SPI.transfer(jval+1024);
    digitalWrite(slaveSelectPin,HIGH);
}


void jumpto(int jval){
    Serial.print("Jump to: ");
    Serial.println(jval);
    digitalWrite(slaveSelectPin,HIGH);
    byte Abyte = 6146 >> 8;
    byte Bbyte = 6146;
    SPI.transfer(~Abyte);
    SPI.transfer(~Bbyte);
    digitalWrite(slaveSelectPin,LOW);
    digitalWrite(slaveSelectPin,HIGH);
    byte Cbyte = jval+1024 >>8;
    byte Dbyte = jval+1024;
    SPI.transfer(~Cbyte);
    SPI.transfer(~Dbyte);
    digitalWrite(slaveSelectPin,LOW);
}
```

```
void setup() {
    // set the slaveSelectPin as an output:
    pinMode (slaveSelectPin , OUTPUT);
    digitalWrite(slaveSelectPin ,LOW);
    // initialize SPI:
    SPI.begin();
    SPI.setBitOrder(MSBFIRST);
    SPI.setDataMode(SPI_MODE3); // is set in function jumpto
    SPI.setClockDivider(SPI_CLOCK_DIV128);

    Serial.begin(57600);

    Serial.println("Initialized to 0 Volt offset");

    jumpto(0); // 2V -> 2G
}

void loop() {

    // to change the offset value, send a new value
    if (Serial.available()) {
      input = Serial.parseInt();
      if (input < 1024) {
        jumpto(input);
      }
      Serial.println("Write values between 0 to 1023 to change value");
      }
}
```

20 DAC Serial Control

```
/*
  Digital Poti Control, following implementation by Tom Igoe
 and Heather Dewey-Hagborg
 Adaptions by Arno Trautmann, 2012-2013
 Adaptions by Marcell Gall, 2015
 Adapted to control the AD5791 - 20-bit DAC with Arduino Mega
 */

// inslude the SPI library:
#include <SPI.h>
#include <math.h>

// set pin 10 as the slave select for the digital pot: (mega 53)
const int slaveSelectPin = 53;

long input;

byte Abyte, Bbyte, Cbyte;

void jumpto(long jval){
  Serial.print("Jump to: ");
  Serial.println(jval);
  digitalWrite(slaveSelectPin,LOW);
  Abyte = jval+1024 >>16;
  Bbyte = jval+1024 >>8;
  Cbyte = jval+1024;

  Abyte = B00010000 ^ Abyte; // add the address of the DAC register (0001)
to the 20 bit word

  SPI.transfer(Abyte);
  SPI.transfer(Bbyte);
  SPI.transfer(Cbyte);
  digitalWrite(slaveSelectPin,HIGH);
}

void setup() {
  // set the slaveSelectPin as an output:
  pinMode(slaveSelectPin, OUTPUT);
  digitalWrite(slaveSelectPin,HIGH);
  // initialize SPI:
  SPI.begin();
  SPI.setBitOrder(MSBFIRST);
  SPI.setDataMode(SPI_MODE1);
  SPI.setClockDivider(SPI_CLOCK_DIV128);

  Serial.begin(57600);

  // set DAC in control-register:
  digitalWrite(slaveSelectPin,LOW);
```

```
    SPI.transfer(B00100000);
    SPI.transfer(B00000000);
    SPI.transfer(B00010010); // second to last bit is set
    //  in configuration with external condensator
    digitalWrite(slaveSelectPin,HIGH);

    Serial.println("Initialized to 0 Volt offset");

    jumpto(524288); // 2V -> 2G
}

void loop() {
    // to change the offset value, send a new value
    if (Serial.available()) {
        input = Serial.parseInt();
        if (input < 1048576) {
            jumpto(input);
        }
        Serial.println("Write values between 0 to 2^20-1 to change value");
    }
}
```

B Bibliography

[1] Princeton instruments catalog of high performance digital ccd cameras.

[2] Analog Devices Inc., One Technology Way, P.O. Box 9106, Norwood, MA 02062-9106, U.S.A. *Data Sheet AD5791*, 2013.

[3] Kiam Heong Ang, Gregory Chong, and Yun Li. Pid control system analysis, design, and technology. *Control Systems Technology, IEEE Transactions on*, 13(4):559–576, 2005.

[4] E Arimondo, M Inguscio, and P Violino. Experimental determinations of the hyperfine structure in the alkali atoms. *Reviews of Modern Physics*, 49(1):31, 1977.

[5] Bartington Instruments Ltd., Bartington Instruments Limited, 5, 10 and 11 Thorney Leys Business Park, Witney, Oxford, OX28 4GE, England. *Mag-03 Three-Axis Magnetic Field Sensors*, 2015.

[6] Bartington Instruments Ltd., Bartington Instruments Limited, 5, 10 and 11 Thorney Leys Business Park, Witney, Oxford, OX28 4GE, England. *Operation Manual for Mag-03 Three-Axis Magnetic Field Sensors*, 2015.

[7] Willard S Boyle and George E Smith. Charge coupled semiconductor devices. *Bell System Technical Journal*, 49(4):587–593, 1970.

[8] Anton Braun. Grundlagen der regelungstechnik. *Carl Hanser-Verlag*, 2005.

[9] Gregory Breit and II Rabi. Measurement of nuclear spin. *Physical Review*, 38(11):2082, 1931.

[10] Microsoft Corp. End user licence agreement microsoft visual studio community 2015. https://www.visualstudio.com/support/legal/mt171547, November 2015. [Online; posted 27-11-2015].

[11] Kendall B Davis, M-O Mewes, Michael R Andrews, NJ Van Druten, DS Durfee, DM Kurn, and Wolfgang Ketterle. Bose-einstein condensation in a gas of sodium atoms. *Physical review letters*, 75(22):3969, 1995.

[12] National Centers for Environmental Information. The world magnetic model - online calculator. http://www.ngdc.noaa.gov/geomag/WMM/calculators.shtml, May 2015. [Online; used 28-11-2015].

[13] Michael E Gehm. Properties of 6li. 2003.

[14] William A Geyger. New type of flux-gate magnetometer. *Journal of Applied Physics*, 33(3):1280–1281, 1962.

[15] Hermann Haken and Hans C Wolf. *Atom-und Quantenphysik: Einführung in die experimentellen und theoretischen Grundlagen.* Springer-Verlag, 2013.

[16] Honeywell International Inc., 12001 Highway 55, Plymouth, MN 55441, USA. *3-Axis Digital Compass IC HMC5883L*, 2013.

[17] Paul Horowitz and Winfield Hill. *The art of electronics.* Cambridge Univ. Press, 2015.

[18] James R Janesick, Tom Elliott, Stewart Collins, Morley M Blouke, and Jack Freeman. Scientific charge-coupled devices. *Optical Engineering*, 26(8):268692–268692, 1987.

[19] DS Jin, JR Ensher, MR Matthews, CE Wieman, and EA Cornell. Collective excitations of a bose-einstein condensate in a dilute gas. *Physical Review Letters*, 77(3):420, 1996.

[20] Jan Krieger. Zeeman-slower und experimentsteuerung fuer das nali-experiment. Diploma thesis, 2008.

[21] Harold J Metcalf and Peter Straten. *Laser cooling and trapping.* Springer, 1999.

[22] Wolfgang Muessel. *Scalable Spin Squeezing for Quantum-Enhanced Magnetometry with Bose-Einstein Condensates.* PhD thesis, Heidelberg University, Kirchhoff-Institut for Physics, 2014.

[23] Timo Bastian Ottenstein. A new objective for high resolution imaging of bose-einstein condensates. Diploma thesis, 2006.

[24] QImaging, 4401 Still Creek Drive, Suite 100, Willingdon Business Park, Burnaby, BC, Canada V5C 6G9. *Retiga EXi Fast 1394 Data Sheet*, 2003.

[25] Tobias Rentrop. unpublished work.

[26] Marc Repp. Aufbau einer vakuumapparatur fuer experimente mit ultrakalten fermionischen und bosonischen quantengasen. Diploma thesis, 2007.

[27] Pavel Ripka. Review of fluxgate sensors. *Sensors and Actuators A: Physical*, 33(3):129–141, 1992.

[28] Bahaa EA Saleh and Malvin Carl Teich. *Fundamentals of Photonics.* Wiley, 2007.

[29] Raphael Scelle. *Dynamics and Motional Coherence of Fermions Immersed in a Bose Gas.* PhD thesis, Heidelberg University, Kirchhoff-Institut for Physics, 2013.

[30] T Schuster, R Scelle, A Trautmann, S Knoop, MK Oberthaler, MM Haverhals, MR Goosen, SJJMF Kokkelmans, and E Tiemann. Feshbach spectroscopy and scattering properties of ultracold li+ na mixtures. *Physical Review A*, 85(4):042721, 2012.

[31] Tobias Schuster. *Feshbach Resonances and Periodic Potentials in Ultracold Bose-Fermi Mixtures*. PhD thesis, Heidelberg University, Kirchhoff-Institut for Physics, 2012.

[32] Marlan O Scully and M Suhail Zubairy. *Quantum optics*. Cambridge university press, 1997.

[33] RJ Sewell, M Koschorreck, M Napolitano, B Dubost, N Behbood, and MW Mitchell. Magnetic sensitivity beyond the projection noise limit by spin squeezing. *Physical review letters*, 109(25):253605, 2012.

[34] Tom Simcox. personal correspondance.

[35] Daniel A Steck. Sodium d line data. *Report, Los Alamos National Laboratory, Los Alamos*, 124, 2000.

[36] Texas Instruments Inc., Post Office Box 655303, Dallas, Texas 75265, U.S.A. *INA12x Precision, Low Power Instrumentation Amplifiers*, 2015.

[37] Arno Trautmann. Feshbach spectroscopy of sodium and sodium-lithium mixtures. Diploma thesis, 2011.

[38] Valentin Volchkov. Cold lithium atoms for future polaron experiments. Diploma thesis, 2009.

[39] Jens Wagner. Anleitung fuer das fortgeschrittenenprakikum elektronik. website of advanced practical courses in the department of physics, 2014.

[40] Stefan Weis. Setup of a laser system for ultracold sodium - towards a degenerate gas of ultracold fermions. Diploma thesis, 2007.

[41] Jinghua Zhong. Pid controller tuning: a short tutorial. *class lesson, Purdue University*, 2006.

Acknowledgement

Erklärung:

Ich versichere, dass ich diese Arbeit selbstständig verfasst habe und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, den 4. Dezember 2015 . . . .