University of Southern Denmark

Master Thesis

Towards Rydberg quantum optics with ultra-cold Yb atoms

Mogens Henrik From

June 11, 2021

supervised by
Sebastian Hofferberth, University of Southern Denmark

Abstract

This thesis proposes a design for a versatile DDS based frequency source, which is build and characterized, confirming that the device proves promising for a wide range of experimental setups. In particular, the device is used to implement Sawtooth Wave Adiabatic Passage (SWAP) cooling in a green intercombination magneto-optical trap (MOT) with ytterbium atoms. It is confirmed that SWAP cooling can be used to realise smaller atomic clouds, and a preliminary analysis confirms the SWAP cooling effects in a MOT. Furthermore, the new RF source is used to implement time-averaged dipole traps, for splitting clouds of ultracold atoms. The stability and precision of the setup are confirmed in a preliminary analysis where an array of beams is created with an interlaced frequency signal through an AOM. The device is integrated with an existing computer control system, which is improved for use in a brand new lab, and a new monitoring system for experimental equipment and lab climate is developed. In the scope of this thesis, the experiment is being reconstructed, requiring disassembly of the vacuum chamber. During this the differential pumping tube in the vacuum chamber is replaced, and the vacuum chamber is reassembled, pumped down, and baked, achieving a final pressure of 10^{-11} mBar. The new RF source provides a promising solution for future experiments, not only for realisation of SWAP cooling and structured arrays of optical traps, but in any case where a versatile, precise, and stable RF source is needed. As the vacuum chamber is reconstructed, and the new RF source developed, the experiment is ready for the realisation of ultra-cold ytterbium atoms in a SWAP MOT in the near future.

Resumé

I denne afhandling bliver en DDS baseret frekvens generator designet og bygget. Apparatet karakteriseres, og det bekræftes at systemet er lovende for et stort udvalg af eksperimentelle opsætninger. Specifikt bliver systemet brugt til at implementere Sawtooth Wave Adiabatic Passage (SWAP) køling i en grøn magneto-optisk fælde (MOT) med ytterbium atomer. Det bekræftes at SWAP køling kan bruges til at realisere mindre atomare skyer, og en indledende analyse bekræfter at SWAP kølingseffekterne i en MOT. Frekvensgeneratoren bruges til at implementere tidsmidlede dipol-fælder til at dele skyer med ultrakolde atomer. Stabiliteten og præcisionen af systemet bliver bekræftet af en indledende analyse hvor rækker af stråler bliver genereret med sammenflettede frekvenser igennem en AOM. Enheden er blevet integreret med det eksisterende computer kontrol system, der er blevet forbedret til brugen i et helt nyt laboratorie, og et nyt overvågningssystem til eksperimentelt udstyr og laboratorieklimaet er blevet udviklet. Som en del af denne afhandling bliver eksperimentet genopbygget, hvilket kræver arbejde i vakuum-kammeret. Som en del af dette bliver det differentielle pumpe-rør i kammeret skiftet, og kammeret pumpes og bages til vakuum, med et endeligt tryk på $\sim 10^{-11}$ mBar. Den nye RF generator er en lovende løsning til fremtidige eksperimenter, ikke kun til at realisere SWAP køling og strukturerede rækker med optiske fælder, men også i tilfælde hvor alsidige, præcise, og stabile RF kilder er nødvendige. Da vakuum-kammeret er genopsat, og den nye RF kilde er udviklet, er eksperimentet nu klar til at realisere ultra-kolde ytterbiumatomer i en SWAP MOT i den nærmeste fremtid.

Contents

Introduction

The field of nonlinear quantum optics has the core goal of investigating effective interactions between individual photons by exploiting strong nonlinearities in optical materials[1]. This field provides new insights on an elementary level of quantum physics and could be the basis for new technologies within optics-based high-speed information processing[2]. The field has already proven fruitful paving the way for quantum computation[3] and quantum simulation[4], but also through the creation and investigation of exotic quantum many-body states of light, such as photon crystals and photon liquids[5, 6, 7, 8].

Even though individual photons do in principle interact in free space, the scattering cross-section is by far negligible, and interactions between individual photons are thus not generally observed. The solution is then to instead investigate effective photon-photon interaction mediated through other media. However, the atom-photon cross-section is also very weak, and as a consequence nonlinear effects are not readily observed. Thus, to observe strong nonlinear effects mediated by atom-light interactions, the atom-light cross-section must be enhanced. This is the main issue in the field, and can generally be solved by two approaches: The first idea is to increase the nonlinearities by increasing the light intensities, and the second is to instead use Rydberg atoms[9, 10] to exploit long-range dipole-dipole interactions and electromagnetically induced transparency (EIT) to create effective photon-photon interactions[11]. The first has been implemented by confined light temporally in cavities[12] or waveguides[13], however, this thesis is mostly concerned about experiments exploiting the nonlinearities in Rydberg atoms.

Within the field of Rydberg physics, there are two main paths: The Rydberg blockade, where a nearby Rydberg excitation is prohibited due to an energy shift of the atoms Rydberg levels[14], or EIT experiments, where a two-photon transition renders an initially opaque medium transparent through destructive interference of the transition amplitudes[15]. Both approaches are very interesting, as the Rydberg blockade can lead to coherent manipulation of an ensemble of atoms by a single photon[16, 17], while EIT has been used to effectively slow down photons significantly through excitation to a dark-state polariton, which is a quasi-particle in EIT that exhibits both photon and matter like properties[18, 19].

The common factor for all Rydberg experiments is that they need an experimental platform, which except for the case of vapour-cell experiments, is most likely a cloud of cold atoms. Thus, the preparation of cold atomic ensembles is crucial in the realization of experiments. The basis for cooling atoms came with Maxwell proposing the theory of electromagnetism, in which radiation pressure was a consequence. In 1903 the force was first demonstrated[20, 21], but it was not until after the development of the laser in the 1960s[22], that further improvements were made. In the mid 1970s, laser cooling was introduced separately by Hänsch and Schawlow[23], and Wineland and Dehmelt[24], based on scattering forces. Furthermore, the dipole force was recognized. In 1968 it was proposed it could be used to trap atoms, and the first experimental realization was in the early 1970s when Ashkin managed to trap and levitate a small glass bead between focused laser beams[25].

Thus, the advent of laser paved the way for cooling and trapping of atoms, and within 10 years laser cooling schemes were proposed, developed, and even experimentally realised. Already in the late 1970s, the first laser cooling experiments were realized, trapping Mg and Ba ions[26, 27]. After this, the first magneto-optical trap was realised in 1987[28], which to this day provides a common starting point in most laser cooling experiments. Since then, several further improvements to cooling have been proposed, surpassing the Doppler limit of simple scattering based cooling techniques. Among these are Sisyphus cooling[29], Raman sideband cooling[30], and sawtooth wave adiabatic passage cooling (SWAP) techniques[31].

In this project, we design and build a versatile RF source for the experimental realisation of SWAP cooling in a core-shell MOT[32] in the ytterbium based experiment in the group of Sebastian Hofferberth at the University of Bonn. Furthermore, the device is tested thoroughly, ensuring that it fulfils the requirements of the experiment, such that it can be implemented in the existing experiment, and improve the loading rate and atom number of the narrow-line MOT. Furthermore, the proposed setup allows for the creation of optical tweezers, which provide a promising path for the creation of novel experimental platforms for modern cold-atom experiments and quantum simulation.

Even though the device is developed with the particular experiment in mind, it is purposely built in a versatile manner, and the design is readily usable in a wide range of other experiments. Both experiments implementing SWAP cooling, and in particular setups relying on individually controllable optical tweezers, or arrays thereof. The final result is a very versatile RF source that is designed specifically for SWAP cooling and moving optical tweezers, but with a large control of the output signals, the device could readily be used for any other application which requires precisely controlled RF signals.

The thesis is structured in 6 separate chapters, going through an introduction and motivation for laser cooling and dipole trapping, then designing, building, and testing the device, before interfacing with the existing experiment, and finally moving and reconstructing the experiment in a new lab.. The first chapter 1 provides an overview of the theory behind laser cooling. In particular, optical molasses and the magneto-optical trap, as well as a brief introduction to dipole trapping. Then, the second chapter 2 discuss the design and development of the RF source, motivating the design choices before the device is assembled, and finally interfacing it with the existing main experiment control. Chapter 3 contains a thorough characterization, to confirm that the experimental requirements are fulfilled. Chapter 4 presents results after the implementation, and confirms the performance of the device as part of the experiment.

Then Chapter 5 briefly describes the reconstruction of the experiment in a brand new lab, focussing in particular on the development of a lab monitoring system, and rebuilding of the vacuum chamber. Finally, Chapter 6 briefly summarize the thesis, and provides an outlook for possible further work in the future.

Chapter 1

Rydberg quantum optics with ultracold ytterbium

1.1 Why use ytterbium?

In the previous decades, alkali atoms have proven very popular in cold atom experiments, due to their simple electron structure, resembling the structure of hydrogen, with only a single valence electron. In recent years, however, an increasing number of experiments are choosing the alkaline-earth-like atoms instead[33, 34, 35, 36, 37], due to their more complex electron structure, with transitions with a wide range of linewidths from MHz to the very-narrow few Hz clock transitions. These could prove interesting for realising new phenomena in the field of Rydberg physics. In particular, ytterbium stands out as an interesting candidate for Rydberg physics, and its use was pioneered in Kyoto with the realisation of degenerate Bose and Fermi gases of Yb[38, 39]. The most obvious difference between the alkalis and the alkaline-earth-like elements is the number of valence electrons. The alkaline-earth-like elements have two valence electrons, which allows for the creation of both singlet and triplet states. The extra valence electron also provides the benefit of being able to excite one electron to the Rydberg state, while the remaining electron creates an optically active core[40], and the second valence electron furthermore makes it possible to simultaneously trap ground and Rydberg states[41].

One of the reasons for the previous favouring of the alkalis is the more complex level structure, and the typically shorter transition wavelengths, of the alkaline-earth-like elements, which render them more difficult to trap and cool. However, with the development of improved cooling techniques, as discussed in this chapter, it is now possible to sufficiently cool them, and exploit their structure for new regimes in quantum optics. In particular, for ytterbium, the presence of both a narrow and a broad line provides clear advantages for cooling[42], e.g. through the realisation of a core-shell magneto-optical trap, as we take advantage of in this project.

This chapter gives first a brief overview of the core properties of ytterbium and the electronic structure. Then, the basics of Doppler cooling and magneto-optical traps are discussed, before dipole trapping is introduced, motivating the following chapter.

1.1.1 Core properties

Yb is a rare-earth element with atomic number 70 and is thus part of the lanthanide series. The full electronic structure of ytterbium is $[Xe]4f^{14}6s^2$, which is, at least for our interest, dominated by the two valence electrons in the outer $6s$ shell. These two valence electrons make ytterbium very similar to an alkaline-earth atom, and Yb is, therefore, part of the alkaline-earth-metal-like atoms.

An important property of Yb is the large number of stable isotopes, of both fermionic and bosonic type. There are 7 stable isotopes of Yb, 2 of which are fermionic (^{171}Yb and ^{173}Yb), and 5 of which are bosonic (^{168}Yb, ^{170}Yb, ^{172}Yb, ^{174}Yb, and ^{176}Yb). Another analogy of ytterbium to the alkaline-earth atoms is the total angular momentum of the ground state, which interestingly is $J = 0$. The bosonic isotopes furthermore have a nuclear spin of $I = 0$, while the two fermionic isotopes ^{171}Yb and ^{173}Yb have nuclear spins of $I = 1/2$ and $I = 5/2$, respectively.

The element itself is soft and malleable and has a melting point of 824°C, and a boiling point of 1196°C[]. Furthermore, it has a very low vapour pressure on the order of 10^{-21} mbar, which is more than 10 orders of magnitude lower than Rb, a widely used element in cold-atoms experiments, with a vapour pressure of 10^{-7} mBar[43]. This very low vapour pressure requires precautions when ytterbium is used in vacuum, as ytterbium atoms generally stick to any surface inside the vacuum chamber. Furthermore, ytterbium need to be heated strongly to get a sufficient amount of atoms. This in combination with the low vapour pressure means that MOT cannot simply be loaded from background vapour. Several groups have reported issues with degraded optical access in their chambers due to Yb depositions on the glass surfaces[44, 45].

1.1.2 Electronic structure

The level structure of ytterbium is shown in Figure 1.1, where each level is denoted $^{2S+1}L_J$. Here S is the electronic spin angular momentum, L the electronic orbital angular momentum, and $J = L + S$ the total electronic angular momentum. With the $4f$ shell fully occupied, ytterbium is dominated by two valence electrons from the $6s$ shell, which are arranged in a spin singlet with $S = 0$, or a triplet state with $S = 1$, creating two distinct manifolds. The transitions within each manifold are generally broad and within the visible spectrum, while the so-called *intercombination* transitions going between manifolds are very narrow, as they are electric dipole forbidden. They are, however, still coupled slightly due to spin-orbit interaction between the $^1P_1$ and the $^3P_1$ states.

The ground state is the $^1S_0$ singlet state, which has an antiparallel alignment of the two valence electrons, and thus possesses zero total angular momentum ($J = 0$). This implies that for isotopes with non-zero nuclear spin, there is no hyperfine structure present in the ground state[44]. This lack of substructure becomes important when laser cooling ytterbium. Sub-Doppler cooling techniques can be applied on atomic species where the ground state consists of multiple hyperfine states[47], but this is not possible in isotopes where there is no substructure in the ground state. Even the fermionic isotopes are very insensitive to magnetic fields and lack fine structure in the ground state. Therefore, decay to a lower hyperfine state during cooling is not an issue, and ytterbium can be cooled without repumping light, which significantly simplifies the laser setup when compared to cooling alkali atoms. Further details and the theoretical background of laser cooling will be discussed in Section 1.2. Even though this magnetic insensitivity seems to solve many issues, it does have a number of other effects: It is not possible to magnetically trap ground-state atoms (unless extreme magnetic gradients are used[46]), and Magnetic Stern-Gerlach separation of spins states is therefore also not realistic. This, however, is not an issue for our setup.

For this experiment, we focus mainly on two specific transitions: The broad $^1S_0 \rightarrow {}^1P_1$ singlet transition, and the narrow $^1S_0 \rightarrow {}^3P_1$ triplet transition, both of which are used for cooling. The $^1S_0 \rightarrow {}^1P_1$ transition has a wavelength of 398.9nm (near-violet blue in the optical spectrum), and a broad linewidth of $\Gamma = 2\pi \cdot 29.13(3)$MHz[48]. Due to the large linewidth, this transition is exploited for the initial trapping in a magneto-optical trap as well as absorption

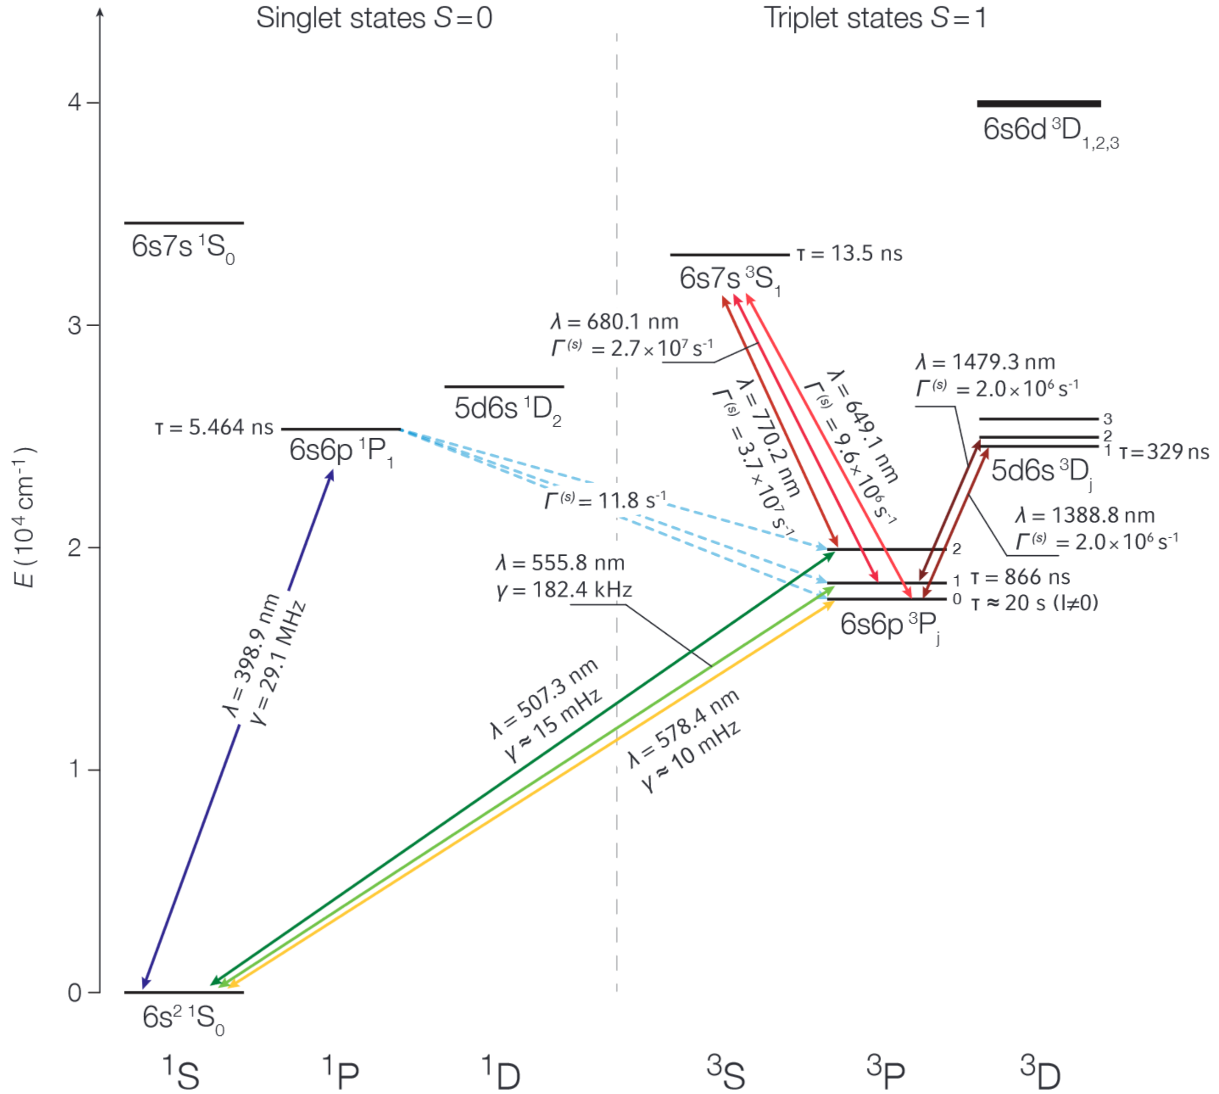

Figure 1.1: *Level diagram of ^{174}Yb, with energy levels and transitions mainly relevant for cooling. Solid lines represent optical transitions, while dashed lines are multi-photon decay channels. Illustration is from [46].*

imaging. The transition is almost closed, with a very weak radiative decay channel from $^1P_1$ through the triplet states $5d6s\,^3D_{1,2}$, and then to the $^3P_{0,1,3}$ states. The decay from each of the D states have lifetimes of respectively 329ns and 460ns[49, 50]. This weak decay channel sets a limit on the number of atoms trapped in the MOT, and their lifetime, which can be avoided by pumping the atoms back to the cooling cycle[51]. That said, the decay is weak enough that efficient cooling is still possible without repumping light[46, 44, 47].

The narrow $^1S_0 \rightarrow {}^3P_1$ transition is in the green range of the optical spectrum, with a wavelength of 555.8nm, and a natural linewidth of $\Gamma = 2\pi \cdot 182.4(4)$kHz. The transition is essentially closed, with a very weak and negligible magnetic dipole transition decay channel from $^3P_1$ state to the $^3P_0$ state[52]. Although this decay channel is present, it is weak enough to be negligible in our setup. The very narrow linewidth ensures that this transition can be used for a very low Doppler limit at $T_D = \hbar\Gamma/2k_B = 4.4\mu$K, which is the reason we use this line to further cool the atoms in a 3D MOT, after the initial cooling using the broader blue singlet transition.

1.1.3 Ytterbium as a platform for Rydberg physics

The electronic structure of ytterbium is not only convenient for the cooling setup, but also provides several paths to Rydberg excitations of the atoms. In particular, there are both ultraviolet single-photon transitions[53], and the two-photon transition with an intermediate state we utilise here, with furthermore allows for the realisation of EIT. More complex multi-photon transitions also exist, but we will not discuss those any further. Specifically for ytterbium, the two-photon transition is particularly interesting, as both transitions are similar in wavelength, and in the visible spectrum, which is both convenient when working with the experiment, and allows for long coherence times, as we discuss below.

The introduction describes the creation of strong optical nonlinearities are the goal of this experiment, and one way of achieving this is by mapping the interaction between atomic Rydberg states onto a few-photon field. Such a mapping is possible because of the extreme properties of atoms in Rydberg states where the atoms are excited to a state of high principal quantum number[9]: Rydberg atoms are large in size, they have very long lifetimes, and large polarizabilities. The large polarizability which goes as n^7 with principal quantum number n gives rise to long-range interactions between Rydberg atoms, as the presence of one atom in a Rydberg state will shift the energy levels of Rydberg states of atoms in the vicinity[54]. These energy shifts can for instance be exploited for facilitation of Rydberg excitations[55] or creation of Rydberg atoms, and are the main mechanism behind the so-called Rydberg blockade where one existing Rydberg excitation suppresses nearby excitations.

The single excitation is shared among the atoms within a certain volume and thus the blockade effects allow large atomic ensembles to behave collectively, and interactions with e.g. the driving field are strongly enhanced to reach even single-photon level where the Rydberg-Rydberg interactions are mapped onto the photons. This has so far allowed for a number of technical applications[1].

It is clear that Rydberg atoms provide a promising platform for realising strong nonlinear effects, however the atomic ensemble will not remain in the collective state indefinitely. With time, the state will lose coherence due to thermal motion of the individual atoms, and coupling to light. The collective Rydberg excitation of N atoms is described as

$$|R\rangle = \frac{1}{\sqrt{N}} \sum_{j=1}^{N} \exp(i\,[k_c - k_p]\,r_j)\,|g_1, \ldots, r_j, \ldots, g_N\rangle\,, \tag{1.1}$$

with counter-propagating probe and control beam k_p and k_c, respectively, and the position r_j of atom j[56]. The state $|g_1, \ldots, r_j, \ldots, g_N\rangle$ has all atoms in the ground stage $|g\rangle$, apart for the j'th atom which is in the Rydberg state. For this state, and mean velocity v of atoms in the ensemble, the scale of the coherence time is set by the phase factor

$$\frac{t}{\tau_{\text{coh}}} (k_c - k_p)\, r_j = 2\pi \left(\frac{1}{\lambda_c} - \frac{1}{\lambda_p} \right) vt. \tag{1.2}$$

This makes it clear that there are two main causes to short coherence times: The thermal motion of the individual atoms, and thus the temperature of the cloud, and the difference in wavelength between the probe and control transitions[56]. The thermal motion can obviously be reduced by cooling the cloud further, although the larger issue is the difference in wavelength of the transitions.

This behaviour is one of the reasons ytterbium is particularly interesting for Rydberg physics, as it not only provides a narrow linewidth cooling transition to lower the temperature of the ensemble, but also provides a two-photon transition to the Rydberg state, where both transitions are similar in wavelength, only separated by about 5 nm. Both of these factors are leading to the large upper limit of the coherence time scale, which for ytterbium near the narrow-line Doppler limit of $T_D = 4.4\,\mu$K is $\tau_{\text{c,Yb}} = 80\,\mu$s, while the coherence time scale for rubidium is instead $\tau_{\text{c,Rb}} = 2\,\mu$s, which is significantly shorter.

Many tools for single photon operations, such as photon sources, gates, single photon absorbers, and single photon transistors[57, 58, 59, 60, 61], are limited by decoherence. This is one of the main reasons why ytterbium is a promising choice for future experiments with nonlinear quantum optics in Rydberg atoms.

Apart from the long coherence times, ytterbium is interesting due to the large number of stable isotopes of both fermionic and bosonic types. in particular, fermionic isotopes could provide a novel platform for modern experiments, due to the lack of collisions between identical fermions[62]. Thus, ytterbium could offer a platform for Rydberg-EIT experiments with fermionic species. Furthermore, as have already been mentioned, the lack of hyperfine structure in the ground state simplifies the cooling setup, as no repumping light is needed, and the combination of a broad and narrow cooling transition allows for a combination of traps with large capture ranges, and still a low Doppler limit. Thus, ytterbium atoms allow for cooling to low temperatures while keeping sufficient atom numbers.

1.2 Green 3D MOT and SWAP Cooling

1.2.1 Doppler cooling and optical molasses

In order to realise experiments with cold or ultra-cold atomic gasses, it is obvious that we need a way to cool the atoms first. This can be done in several ways, but perhaps the simplest idea is Doppler cooling. This cooling technique was first proposed in 1975 simultaneously by Wineland and Dehmelt[63], and by Hänsch and Schawlow[23], and the principles behind can be described with an intuitive picture.

The principle of Doppler cooling is to irradiate laser light onto an ensemble of atoms with a velocity distribution. The Doppler shift of the light in the atomic rest frame depends on the velocity of the atom, causing the light to be more or less on resonance with the individual atoms. Absorption of a photon and thus momentum transfer to an atom, therefore, depends on the atomic velocity. Hence, atoms with velocity within a certain region will be addressed by the light, and experience the resulting momentum transfer. This velocity-dependent momentum

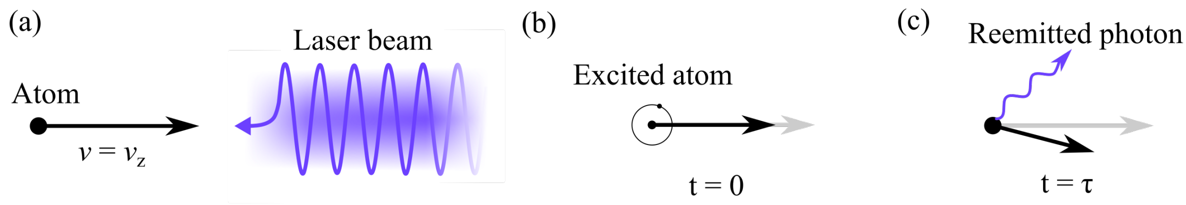

Figure 1.2: *Illustration of Doppler cooling setup, where (a) the atom first absorbs a photon from the counter-propagating laser beam. The atom is moving with a velocity v_z opposite to the propagation direction of a laser beam, causing a Doppler shift of the laser frequency in the rest frame of the atom. (b) The absorption of a photon then excites the atom for some period τ after which (c) the atom de-excites and spontaneously emits a photon in a random direction. The black arrow represents the momentum of the atom in each step, compared to the gray arrow showing the momentum from previous step.*

transfer only affects atoms until the velocity is outside the addressed region, and the velocity distribution, therefore, tends to this lower limit of the region. That is, the velocity-dependent momentum transfer narrows the velocity distribution, which effectively means that the atoms are colder in the relevant rest frame.

Consider an atom moving with some velocity v_x in the $+x$-direction, and a laser beam with a frequency ν_L propagating in the $-x$ direction, as shown in Figure 1.2(a). If the laser frequency is resonant with the atomic transition, then the atom can absorb a photon from the laser beam. The absorption process transfers the momentum of the photon to the atom. This momentum 'kick' will be in the $-x$-direction, thus slowing down the atom. The now-excited atom eventually decays back to the ground state, and then re-emit the photon. During the re-emission, the atom will again gain some momentum, but the direction of re-emission, and thus the momentum transfer, is random (Figure 1.2). The atom is then back in the ground state and a new photon can then be absorbed, repeating the process. As the absorbed photon will always be 'pushing against' the atom, but the photon is re-emitted in a random direction, there will be a net change in momentum in the x direction of

$$\Delta p_x = -\frac{h}{\lambda}, \tag{1.3}$$

with the photon momentum $\frac{h}{\lambda}$. This follows directly from conservation of momentum, and the fact that the absorbed photon is always in the $-x$-direction, while the emitted photon is in a random direction. Thus over time, the net momentum contributions from photon re-emission will average to zero.

In reality, a laser beam with the exact transition frequency of the atom would not be on resonance, due to the Doppler shift. In the rest frame of the atom, the laser source is moving towards the atom, thus leading to a Doppler shift of the laser frequency, such that the laser in the lab frame needs to be slightly red-detuned. If we write the detuning as $\nu_L = \nu_A + \delta$, with the atom transition frequency ν_A and detuning δ, then the Doppler shift is given by[64]

$$\nu_L' = \nu_L \left(1 + \frac{v_x}{c}\right) = (\nu_A + \delta)\left(1 + \frac{v_x}{c}\right) \approx \nu_A + \delta + \nu_A \frac{v_x}{c}, \tag{1.4}$$

where the expansion assumes that $v_x \ll c$. With this, we find that $\nu_L' = \nu_A$ when

$$\delta = -\nu_A \frac{v_x}{c} = -\frac{v_x}{\lambda}, \tag{1.5}$$

with the transition wavelength $\nu_A/c = \lambda$. This then provides a measure of the red-detuning needed before the Doppler cooling works for a particular velocity range.

This result implies that a friction-like force will be applied on the resonant atoms as they propagate against the laser, thus, in general, leading to the atoms slowing down. The idea presented above is intuitively true but is a bit of an over-simplification. To calculate the actual force on the atoms, and derive a lower limit of the achievable temperature, we introduce a new system: Consider a laser beam with detuning $\Delta = 2\pi\delta$ in angular frequency units, and optical intensity I, and an atom propagating with a velocity v_A with respect to the laser. The force on the atom would then be given by the momentum change of each absorption-emission cycle, multiplied by the number of such cycles, such that

$$F_A = -\hbar k R\left(I, \Delta\right), \tag{1.6}$$

with the wave vector $k = 2\pi/\lambda$ of the photon, and the net absorption rate $R\left(I, \Delta\right)$. The net rate should be given by the absorption rate, minus the spontaneous emission rate. We define the natural linewidth $\gamma = 1/\tau$, the saturation intensity I_s of the transition, and we can then write the rate as [65]

$$R\left(I, \Delta\right) = \frac{\gamma}{2}\left(\frac{I/I_s}{1 + I/I_s + \left[2\left(\Delta + kv_A\right)/\gamma\right]^2}\right). \tag{1.7}$$

It is clear that at high intensities, the fraction will vanish, leading to $R = \gamma/2$, and that at low intensities the intensity term in the denominator can be neglected, leading to a linear intensity dependence of the fraction.

The technique as described until now works well with large laser detunings, but as the atoms are slowing down, the detuning required for further cooling is also lowered, and will at some point become comparable to the linewidth. This is an issue, as a detuning close to the linewidth will increase the probability that atoms at rest, or co-propagating atoms absorb a photon from the laser, leading to a momentum gain in the propagation direction of the atoms. This means that as the atoms cool down, the laser will effectively start to accelerate the atoms again. In order to prevent this, and achieve low temperatures, we add another laser counter-propagating to the first one. In this case, the atoms will experience a separate force from each of the laser beams, leading to a net force of

$$F_A = F_+ + F_-, \tag{1.8}$$

where $F_\pm$ is the force from the laser beam propagating in the $\pm$ direction, respectively. When the velocity is high, the detuning is also high, and the atoms will only be affected by one of the laser beams as $F_+ \gg F_-$, or opposite, depending on the atom direction. However, when $|v_A|$ becomes small, and the detuning approaches the linewidth, a more careful analysis is required to find the net force on the atom.

For a single atom propagating in the $+x$ direction, the net force will be given by

$$F_A = F_+\left(\Delta + kv_A\right) - F_-\left(\Delta - kv_A\right). \tag{1.9}$$

A simple expansion yields

$$F_\pm\left(\Delta \pm kv_A\right) = F\left(\Delta\right) \pm \frac{\mathrm{d}F\left(\Delta\right)}{\mathrm{d}\Delta} kv_A, \tag{1.10}$$

where Equations (1.6) and (1.7) already gives

$$F\left(\Delta\right) = -\hbar k \frac{\gamma}{2}\left(\frac{I/I_s}{1 + I/I_s + \left(2\Delta/\gamma\right)^2}\right). \tag{1.11}$$

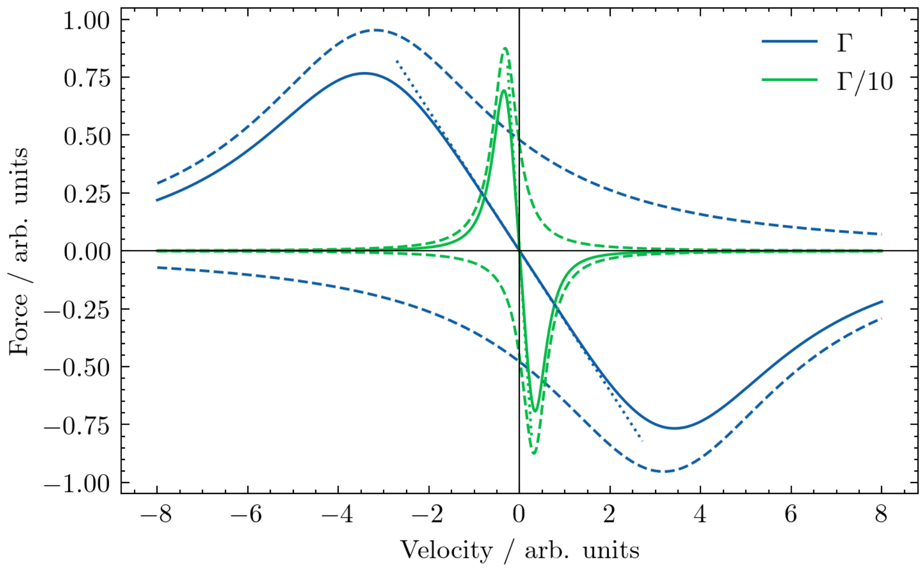

Figure 1.3: *Force versus velocity for $\Delta = \Gamma/2$ and with $I/I_s \ll 1$ for an optical molasses. The dashed lines are $F_\pm$ (F_+ on top, F_- below), the solid line is the total force. The dotted line shows the linear approximation of the force for $|kv_A| \ll |\Delta|$ or $|kv_A| \ll \Gamma$, ie. the low-velocity region. The plot is created for the case of a broad line Γ, and a narrower line at $\Gamma/10$.*

Thus, the derivative is

$$\frac{\mathrm{d}F(\Delta)}{\mathrm{d}\Delta} = \frac{4\hbar k\Delta}{\gamma} \frac{I/I_s}{\left(1 + I/I_s + 4\Delta^2/\gamma^2\right)^2}. \tag{1.12}$$

Combining everything then leads to the final result,

$$F_a = F(\Delta) + \frac{\mathrm{d}F(\Delta)}{\mathrm{d}\Delta} kv_A - \left(F(\Delta) - \frac{\mathrm{d}F(\Delta)}{\mathrm{d}\Delta} kv_A\right) \tag{1.13}$$

$$= 2\frac{\mathrm{d}F(\Delta)}{\mathrm{d}\Delta} kv_A \tag{1.14}$$

$$= \frac{8\hbar k^2\Delta}{\gamma} \frac{I/I_s}{\left(1 + I/I_s + 4\Delta^2/\gamma^2\right)^2} v_A. \tag{1.15}$$

This can be rewritten in the form of a damping force, $F = -\alpha v_A$, with the damping coefficient

$$\alpha = -\frac{8\hbar k^2\Delta}{\gamma} \frac{I/I_s}{\left(1 + I/I_s + 4\Delta^2/\gamma^2\right)^2}. \tag{1.16}$$

From this result, it is clear that depending on the sign of the detuning Δ, the atoms will either be accelerated or decelerated by the force. This behaviour is similar to the one of a particle moving in a viscous fluid, which is the reason that the setup with counter-propagating beams is also known as an optical molasses[65]. The forces from each beam, the combined force, and the linear regime from Equation (1.15) are shown in Figure 1.3.

The lower temperature limit achievable by this cooling setup is determined by the cooling effect of the damping force, compared to the heating effects[66]. The heating effect stems from the repeated absorption and emission of photons, and reflects the momentum in random directions that the spontaneous emissions will cause. That is, we cannot reach zero temperature due to the atoms' continuous interaction with the laser, and the resulting momentum diffusion in the cloud. Thus, the heating rate is given by

$$\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right)_{\text{heat}} = \frac{D_p}{m}, \qquad D_p = \frac{1}{2}\frac{\mathrm{d}\left\langle p^2\right\rangle}{\mathrm{d}t}, \tag{1.17}$$

where D_p is the momentum diffusion constant, which reflects the rate at which the atoms heat each other through the emitted photons and collisions. That is, the spontaneous emission is causing the heating, while collisions cause rethermalization of the cloud. The momentum diffusion constant is thus a reflection of the mean square of the momentum of the atoms, which is non-zero, even though the average velocity of the atoms is zero. This is because the random momentum kicks from the spontaneous emissions vanish when averaged but still lead to some recoil of the atoms. In fact, the recoil caused by the random momentum contributions can be described as the atoms performing a random walk in the cloud. In a random walk with a large number of steps, N, the average momentum will be zero, but the mean square is given by[67]

$$\left\langle p^2\right\rangle = 2N\left(\hbar k\right)^2, \tag{1.18}$$

where $\hbar k$ stems from the momentum contribution of a single photon absorption. As there are two laser beams in the molasses, the number of interactions in time t is given by $N = 2Rt$, where R was defined in Equation (1.7). We thus have that the momentum diffusion constant is

$$D_p = \frac{1}{2}\frac{\mathrm{d}\left\langle p^2\right\rangle}{\mathrm{d}t} = 2\hbar^2 k^2 R. \tag{1.19}$$

The heating rate was defined in Equation (1.17), and the cooling rate is simply given by

$$\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right)_{\text{cool}} = F_A v_A = -\alpha v_A^2. \tag{1.20}$$

In the lower limit, there is an equilibrium between the heating and cooling rate, such that the total energy change is zero. That is,

$$\left(\frac{\mathrm{d}E}{\mathrm{d}t}\right)_{\text{heat}} + \left(\frac{\mathrm{d}E}{\mathrm{d}t}\right)_{\text{cool}} = 0 \Leftrightarrow \frac{D_p}{m} - \alpha v_A^2 = 0, \tag{1.21}$$

which is satisfied by

$$v_A^2 = \frac{D_p}{\alpha m}. \tag{1.22}$$

Atoms do not interact over the typical interatomic distance in a MOT, and thus the only degree of freedom to consider is the translational motion. The r.m.s thermal velocity is then found with

$$\frac{1}{2} m v_A = \frac{1}{2} k_B T, \tag{1.23}$$

by setting the kinetic of each velocity component to $\frac{k_B T}{2}$. Using this, we can find temperature corresponding to 1.22:

$$\frac{1}{2} k_B T = \frac{m}{2}\frac{D_p}{\alpha m} = \frac{D_p}{2\alpha}, \tag{1.24}$$

and the temperature limit we obtain is thus

$$T = \frac{D_p}{k_B \alpha}. \tag{1.25}$$

It is evident that this limit is minimized by optimising the ratio of D_p to α.

We can now substitute the definition of R (Equation (1.7)) into the momentum diffusion constant (Equation (1.19)), and add this, as well as the definition of α (Equation (1.16)), in Equation (1.25), leading to

$$T = \frac{\hbar\gamma}{8k_B} \frac{1 + I/I_s + 4\Delta^2/\gamma^2}{\Delta/\gamma}. \tag{1.26}$$

In the low temperature limit, that is, $I \ll I_s$, the intensity term vanishes, and by minimizing we get the detuning $\Delta = -\gamma/2$ which leads to the minimum temperature

$$T_{\min} = \frac{\hbar\gamma}{2k_B}. \tag{1.27}$$

This temperature is also known as the Doppler limit and puts a lower limit on the temperature we can achieve by simple Doppler cooling in an optical molasses. Some experiments have achieved sub-Doppler cooling, surpassing this limit[68], through methods e.g. like Sisyphus cooling[29], Raman sideband cooling[30], and cavity-cooling techniques[69]. It is worth noting that the sub-Doppler cooling techniques mentioned here require structure in the atomic ground state, and they are therefore not applicable for the bosonic isotopes of ytterbium.

1.2.2 3D MOT

There are two large issues with the cooling setup described until now. First of all, we have only accounted for 1 dimension, and secondly, it does not actually trap the atoms. The first issue is solved rather simply by adding more beams in the remaining dimensions. That is, instead of a pair of counter-propagating beams in the x-direction only, we add pairs in the y, and z-directions, as shown in Figure 1.4b. This will create an optical molasses that slows the atoms in all 3 spatial dimensions. The optical molasses, also in 3D, does provide viscous damping, but the force is not position dependent, and the atoms will eventually reach the surface of the interaction region, which is given by the size of the laser beams, and escape.

This is the second issue, which is a bit more involved, but it is still fairly simple to solve. A solution is to introduce a conservative force that keeps the atoms at a specific point in space. In other words, we need a force that depends on the position of the atoms, instead of only on the atom momentum. Such a position-dependent atom-light interaction can be introduced by applying a magnetic quadrupole field to the system. Due to the Zeeman effect, the magnetic field will cause a splitting of the m_F levels of each state, which would normally be degenerate. As the quadrupole field is spatially varying, the splitting of the levels will also be spatially depending. This means that that the levels will only be degenerate in the centre where the field is zero, with the splitting increasing with the distance to the centre. Furthermore, the m_F levels are sensitive to polarization of light, and the transition which is driven can thus be controlled by choosing the polarization of the incident beams. Specifically, choosing the circular polarisation of the MOT beams can determine which of the beams couple to each of the $\sigma^{\pm}$-transitions of the m_F levels. That is, in the case where the magnetic field is parallel to the beam. This allows us to set up the trap such that atoms will always experience a force in the direction towards the

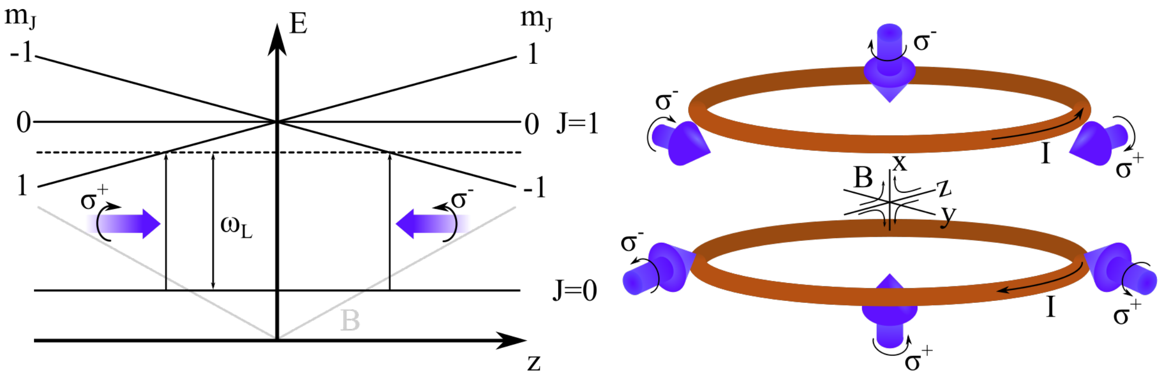

(a) *Sketch of energy against position in a 3D MOT. The energy levels of the $J = 1$ state are shifted by the magnetic field, making the resonance frequency position and polarisation dependent. The incoming beams are oppositely polarised, ensuring that the atom will only be resonant with a counter-propagation beam, raltive to the atom, ensuring that the atom is slowed down.*

(b) *Sketch of a typical MOT setup, with beam polarisation. 3 pairs of beams with opposite polarisation are crossed in the center point of the quadrupole field created by the magnetic coils in an anti-Helmholtz configuration.*

Figure 1.4: *Illustration of a 3D MOT setup.* ***(a)*** *shows the position-dependent splitting of energy levels due to the magnetic field, and the corresponding shift of resonance, and* ***(b)*** *is the experimental implementation of a 3D MOT, with three beams crossing in the center of the magnetic field.*

trap centre, creating a restoring force. The strength of the force at different distances to the centre depends on a combination of the magnetic field gradient and the detuning.

A complete three-dimensional Magneto-Optical Trap (MOT) setup then comprises six beams, pairwise counter-propagating, and all intersecting at right angles at the centre of the quadrupole magnetic field. Such a setup, with the correct polarisations, is sketched in Figure 1.4b. In this MOT setup, the magnetic field is providing the conservative force that pulls the atoms to the trap centre, while the optical molasses is providing a dissipative force that cools and captures the atoms in general. While the conservative force is absolutely crucial for the atoms to be trapped, the cooling effect is still provided by the molasses, and the cooling limit is, therefore, the same as derived in Equation (1.27)[70]. In turn, the capture velocity of the trap is also still determined primarily by the laser detuning and bandwidth.

1.2.3 Green vs. Blue MOT

It has been established that it is possible to trap atoms in 3D by using a combination of an optical molasses, and a quadrupole trap, in a combined setup called a MOT. However, the capture velocity depends on the detuning and bandwidth of the MOT beams, and the cooling limit depends on the linewidth of the atomic transition. We will ignore optimisation of the beams for now, and focus on the linewidth of the transitions.

In this experiment, we are working with ytterbium, which as described above, has two main transitions that are used for trapping. One is the broad $1S_0 - 1P_1$ transition with $\gamma/2\pi = 29.1$MHz, and the other is the narrow intercombination transition $1S_0 - 3P_1$ with linewidth $\gamma/2\pi = 182.3(3)$kHz. Given the Doppler limit $T_{\text{min}} = \frac{\hbar\gamma}{2k_B}$, it is very clear that the lower the linewidth, the lower the cooling limit. In fact, between these two transitions, the Doppler limit goes from $T_{(1S_0-1P_1)} \sim 10^{-4}$K to $T_{(1S_0-3P_1)} \sim 10^{-6}$K. It then seems that one should clearly use the narrow line transition, in order to reach lower temperatures and higher densities. However, the smaller linewidth also leads to a lower capture velocity such that faster atoms will not be captured. This can be seen in Figure 1.3, where the narrow-line force clearly has a steeper force gradient, while the velocity range is significantly narrower. This will ultimately lead to a cloud with fewer atoms. A number of solutions to this problem have been suggested, the first starts with a broad transition, and then switches to a narrow transition for further cooling once a desirable number of atoms have been captured[71]. This approach successfully improves both the amount of captured atoms and cooling, although even better methods have been suggested. The second improvement is a core-shell MOT, where the cooling beam consists of two segments: A narrow transition beam in the centre, surrounded by a broad transition beam on the mantle[32]. This scheme avoids the problem of losing atoms when switching between the narrow and broad MOT, by simply running both simultaneously, and directly loading from the broad to the narrow MOT.

The core-shell MOT successfully improves the loading rate and capture velocity of the MOT, while keeping the low Doppler limit of the narrow MOT, however, it is still limited by the Doppler cooling limit. Thus, we need another solution if we want to cool the atoms further than the Doppler limit, which for the narrow-line of ytterbium is 4.5μK. One approach is Sawtooth Wave Adiabatic Passage, which is what this project is focused on.

1.2.4 Sawtooth Wave Adiabatic Passage (SWAP) Cooling

We have established that Doppler cooling is a simple, and fairly effective cooling scheme, especially when combined with a magnetic field, thus forming a MOT. However, Doppler cooling is strongly dependent on spontaneous emission for getting the atoms back to the ground state, but at the same time, this spontaneous emission is the main contribution to the lower temperature limit achievable with this method. Thus, in order to achieve lower temperatures, we need a cooling scheme that either remove this dependence on spontaneous emission altogether or at least limits the heating effect of the spontaneous emission.

In this section, we will introduce the Sawtooth Wave Adiabatic Passage (SWAP) cooling scheme, which to some extent manages to get rid of the dependence of spontaneous emission, and thus achieve much lower cooling temperatures[72]. Here we will introduce the basic mechanism, and the setup required to realise SWAP cooling. In the next section, we will then discuss the design and building of an experimental setup that implements this cooling scheme.

The basic setup As the name suggests, the basic idea of SWAP cooling is to force a particle toward zero momentum through adiabatic passage. In order to describe the cooling system, we consider a simple two-level particle with internal states $|g\rangle$ (ground) and $|e\rangle$ (excited), which are separated by $\hbar\omega_a$. The particle has momentum p, and freedom of movement only in the z-direction, as we will only consider 1 dimension for simplicity. The particle is then interacting with the light from two counter-propagating laser beams, similar to the setup used for Doppler cooling. The crucial difference here is that the laser frequency, $\omega_L(t)$, is modulated in time to follow an asymmetric sawtooth waveform with frequency range Δ_s, and period T_s (See Figure 1.5). The sawtooth ramp is centred around the transition frequency ω_a and is modulated linearly from below to above ω_a.

With the system defined, we can consider the interaction of a moving particle with the laser beams. The interaction depends on the direction of the particle motion, as this defines the order in which the particle interacts with each of the beams. First, we consider a particle initially in state $|g\rangle$, with momentum $p > 0$. We assume that the laser frequency starts at ω_{min} and is

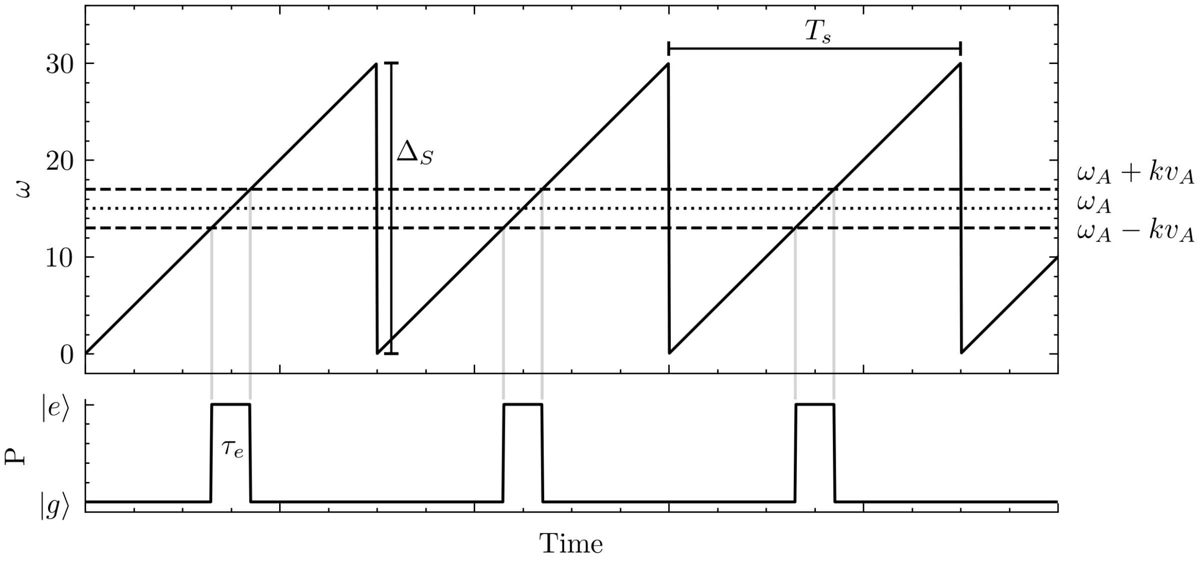

Figure 1.5: *Laser frequency as a function of time, where T_s denotes the period of the ramps, Δ_s the range of the frequency modulation, and τ_e the time in which the atom is excited. ω_a marks the centre frequency of the modulation, and dashed line above and below denote the resonance frequencies after the Doppler shift.*

swept linearly to ω_{max}. Depending on the propagation direction of an atom, the two beams are brought into resonance at different points in time due to the Doppler shift. An atom moving in the z-direction will see the counter-propagating beam first, and only later see the beam propagating in the $+z$-direction. The interaction with the counter-propagating beam will then adiabatically transfer the atom to state $|e\rangle$ with momentum $p - \hbar k$, where k is the wavenumber, via absorption. Then, when the frequency has swept far enough, the co-propagating beam will become resonant, and adiabatically transfer the particle to state $|g\rangle$, with momentum $p - 2\hbar k$ via stimulated emission, rather than spontaneous emission. In the case of an initial particle with $p < 0$, the process is the same, although in reverse, and the process is thus: $|g\rangle$, $p \to |e\rangle$, $p + \hbar k$ $\to |g\rangle$, $p + 2\hbar k$. In both cases, the particle is transferred closer to zero momentum, with a net removal of momentum of $2\hbar k$, and the particle is furthermore returned to its initial internal state, due to the stimulated emission. Each subsequent sweep would continue to remove momentum in steps of $2\hbar k$, and as the particles are left in their initial states, there are no apparent problems in repeating the cycle. In this idealised situation, the momentum of the particle would be expected to follow the trend shown in Figure 1.5.

Requirements We have just established that each cycle would remove two units of momentum from a particle in the system, if the particle starts in the ground $|g\rangle$ state. If the particle were to start in the excited $|e\rangle$ state instead, the process described above would be reversed, such that each interaction with the beams would increase the momentum of the particle, and the particles would thus be heated instead of cooled. It is therefore absolutely critical for this cooling scheme that the particles are in the ground state when each new sawtooth cycle starts. One way to ensure that this is the case is by making sure that the particles spend most time outside the two resonances, such that atoms left in $|e\rangle$ at the end of a frequency sweep have time to decay to $|g\rangle$ via spontaneous emission. The most obvious way to achieve this experimentally is by introducing a waiting time between each ramp in the sawtooth wave. However, we usually want to cool as quickly as possible, and arbitrary waiting times should thus be avoided. The compromise we use is to ensure that the frequency range of the sawtooth ramps fulfils[31]

$$\Delta_s > 4|kv|. \tag{1.28}$$

As the photon resonances are spaced by $2|kv|$, this condition ensures that at least half the time will be spent away from resonance, thus ensuring atoms are more likely in $|g\rangle$ rather than $|e\rangle$ when the next ramp starts. At the same time, this requirement enforces that the sweep range is large enough that both beams will reach resonance with the particle.

There are multiple other requirements for this scheme to work. First, we want to be certain that the particles are transferred adiabatically at each resonance. The probability for an adiabatic transition at each resonance is given by [73]

$$P_a = 1 - \exp\left(-\frac{\pi\Omega^2}{2\kappa}\right), \tag{1.29}$$

with Rabi frequency Ω of each laser beam, and $\kappa = \Delta_s/T_s$ the sweep rate of the laser frequency. In order to ensure a substantial probability of an adiabatic transition, we then have the condition that

$$\frac{\Omega^2}{\kappa} \geq 1. \tag{1.30}$$

When this condition is fulfilled, the system is said to be in the adiabatic regime, and any case that does not fulfil the condition is instead within the diabatic regime.

The last major requirement for the system is related to the decay time of the excited state. As the scheme relies on stimulated emission by the co-propagating beam, we need to ensure that there is a low probability of spontaneous decay before the second beam becomes resonant with the particle. Thus, we require that

$$\tau_e \ll \frac{1}{\gamma}, \tag{1.31}$$

where τ_e is the time from resonance with the first beam until the second beam becomes resonant, ie. the time the particle is in the excited state, and γ is the linewidth of the transition. In the case of the intercombination transition in Yb, with a linewidth of $\gamma = 182.3\,\mathrm{kHz}$, this leads to $\tau_e \ll 5.49\,\mathrm{ms}$, and thus the condition is very easily fulfilled when using narrow transitions. This requirement, together with the frequency range in Equation (1.28), sets the limits on the ramp. In intuitive terms, looking at Figure 1.5, what we want to ensure is that the slope of the ramp is large enough that the time between the first and second beam being on resonance is sufficiently short, while at the same time the frequency range should be large enough that the total duration of each cycle is long enough to allow any excited atoms to spontaneously decay. A large frequency range also comes with the advantage of addressing atoms with a wider ranger of velocities.

So far, we have only considered the SWAP cooling mechanism in a one-dimensional system. For a simple static-frequency MOT the one-dimensional setup simply generalise to three dimension. This, however, is not the case for SWAP cooling. The basic principles as described in this section are true in any number of dimensions, but when expanding to multiple dimensions, the crossing beams make the picture more complicated. Essentially, when there are multiple beams, unwanted cross-talk in the stimulated emission can occur between them. That is, instead of the atom getting excited by one beam, and de-excitation by the counter-propagating beam, the de-excitement could be caused by one of the perpendicular beams, thus not leading to the expected loss of momentum. Experimental results show that simply modulating the frequency of all beams at the same time, does lead to improved cooling when compared to a static-frequency MOT[74]. The stimulated emission between beam can be avoided by only illuminating one axis at a time, while switching between the axes. This approach leads to even better cooling, although at the cost of only cooling on axes at a time[75].

In the rest of this thesis, we will not further investigate the complexities of implementing SWAP cooling in three dimensions, and we use only the simplest approach of illuminating and sweeping all beams at once. The setup that will be discussed in Chapter 2 could, however, easily be extended to drive a setup with individually illuminated beams for more complex implementation of a SWAP cooling in a 3D MOT.

1.2.5 Dipole trap

Another way of trapping atoms is with optical dipole traps. Dipole traps use a very different approach to trapping when compared with for instance a MOT. A dipole trap does not rely on the radiation force, which is the main feature of a MOT, but instead takes advantage of dipole interactions. It turns out that a dielectric particle will be attracted or repelled by the high-intensity region of an electric field, depending on detuning. Atoms do not have an intrinsic electric dipole moment, but only a polarisability. The dipole trap thus relies on the transition dipole moment induced by the trap light, which is a second-order effect. This has the advantage that the contribution to heating from the dipole trap is minimal[76]. It is also important to note that, contrary to the optical molasses, the dipole trap is not really cooling the atoms, but is rather only trapping.

The first description of optical dipole trapping stems from the work of Ashkin on trapping and levitating transparent microscopic particles using light[77, 78]. The basic principle is that in an electric field E, a particle with polarisability α, has a dipole moment αE and the potential energy $-\alpha E^2$[79]. Thus for positive polarisability ($\alpha > 0$), the potential will decrease with higher intensity, and the particle is therefore attracted to high-intensity regions. The force itself is known simply as the *dipole force*, while the trap created from this is an *optical dipole trap* when the field is provided by a laser. For ground-state atoms, the polarisability is positive for laser frequencies below the resonance, and a trap can thus be created by using a red detuned and strongly focused laser beam (to create a high-intensity region).

Assuming that the dipole trap laser is tuned close to the resonance at ω_0, such that $|\Delta| \ll \omega_0$, we are in the regime where the rotating-wave approximation[80] can be applied. With this approximation, the dipole potential for a transition with decay rate Γ, and resonance frequency ω_0 is

$$U_{\text{dip}}(\mathbf{r}) = \frac{3\pi c^2}{2\omega_0^3} \frac{\Gamma}{\omega - \omega_0} I(\mathbf{r}), \tag{1.32}$$

with the laser frequency ω, intensity I, and the detuning defined as $\Delta = \omega - \omega_0$. From this, it is clear that the trap depth is scaling inversely with the detuning, and thus a deep trap can be realised with $\Delta \approx 0$. However, the spontaneous scattering rate is[76]

$$\Gamma_{\text{sc}}(\mathbf{r}) = \frac{3\pi c^2}{2\hbar\omega_0^3} \left(\frac{\Gamma}{\omega - \omega_0}\right)^2 I(\mathbf{r}), \tag{1.33}$$

which is scaling with $1/\Delta^2$, and the detuning $\Delta \approx 0$ is thus not desirable for a trap, due to the increased heating rate from scattering in this limit. Since a large trap depth is still very much desirable, the most common approach is to use a large detuning to keep the scattering rate low, and then use high-intensity laser fields to gain back trap depth. This approach is feasible as the scattering rate is scaling quadratically with detuning, while the scaling with the potential is only linear, such that any loss in depth from the detuning can be compensated by the linearly scaling intensity. It should be noted, however, that for the setup with a large detuning, the rotating wave approximation will no longer be valid, and all transitions will contribute to the potential.

From this potential, it is also clearly seen that the potential is negative only for a red detuned laser frequency ($\Delta < 0$), while a blue detuned laser frequency ($\Delta > 0$) would instead lead to a positive potential. In this simple model of trapping, the laser must be red detuned, as anything else would lead to repulsion of the atoms from the high-intensity region of the laser field. However, it should be noted that another approach to optical dipole trapping takes advantage of the repulsive potential of blue-detuned light. The repulsion can be used to create 'dark traps', where a tailored beam is used to create e.g. a ring potential around the atoms, trapping them inside a barrier of repulsive blue-detuned light. The atoms will then be resting in the dark region between beams, instead of the high-intensity region inside a beam. Note also that dipole traps are generally very shallow, and are therefore usually not used alone, but are instead loaded from e.g. a 3D MOT, that can ensure the initial capture and pre-cooling of the atoms.

A large advantage of the optical dipole trap is that it can be created by only a single laser beam. This allows for easy modifications of the trap in the form of e.g. moving the trap around or very rapid toggling of the trap potential, as there are no slow mechanical processes involved. This has many use cases, e.g. in the form of optical tweezers that can trap particles in a large range of sizes, from single atoms to particles of several micrometres, for instance in the form of bacteria[81]. Furthermore, with precise control of the dipole trap positions, optical tweezers can be used to move particles, and even Bose-Einstein condensates (BEC)[82], and can thus be used to create near arbitrary arrays of individual atoms, atomic ensembles, or BECs, providing a novel platform for experiments and quantum simulation.

It is then clear that optical tweezers, ie. dipole traps, is a very promising addition to cold-atoms experiments, as they allow precise control of the experimental platform. In particular, another experiment in the group of Sebastian Hofferberth has used an array of three dipole traps to implement cascaded Rydberg superatoms in three tightly confined cold atom ensembles[83]. However, their setup is currently very limited, and not easily extended. Thus, a versatile solution is needed to create these optical tweezers for use in experiments. The device proposed in Chapter 2 will allow the creation and manipulation of arrays of multiple optical tweezers, and could readily be implemented in the existing experiment in the future. Furthermore, the design provides a very general and versatile solution, that could readily be used in a wider range of experimental setups.

Chapter 2

A versatile FPGA controlled DDS frequency source

In the previous chapter, we established ytterbium as a promising choice for cold atom experiments, and in particular, realised that this choice of atomic species allows for improved cooling schemes. Here we are specifically taking advantage of the narrow $1S0 \rightarrow 3P1$ transition with a green 3D MOT, in a core-shell MOT [32], but with further improvements in the form of SWAP cooling. In order to implement this SWAP cooling scheme, we need a way to modulate the frequency of the laser fast and accurately. Frequency modulations of a laser beam are easily implemented with the use of an acousto-optical modulator (AOM), through which the frequency shift of the laser is defined by the RF frequency applied to the AOM. This chapter explains our approach to build a frequency source for an AOM is such a setup. First, we propose a design and discuss a versatile setup that can generate a large range of RF signals with high precision. This includes the basic methods for signal generation and reasoning as to why the specific solution we choose is advantageous. Then, the specific components are chosen and combined in a single-unit device, which is thoroughly characterized to confirm that the final solution fulfils the set requirements.

The RF source is made mainly for the implementation of SWAP cooling, which was discussed in Section 1.2.4, and the requirements are therefore: 1. High RF power such that an AOM can be driven, 2. continuous frequency tuning, and 3. fast triggering for reliable implementation in the experiment cycle. While the device is designed with SWAP cooling in mind, it is purposely designed to be versatile, such that it can also be used in a wide range of other experimental setups. For instance, the precise control of arrays of optical tweezers, or as a frequency reference for laser locking.

2.1 Hardware options

There are two options for the main hardware in the frequency source. One can use either an Arbitrary Waveform Generator (AWG), or a Direct Digital Synthesizer (DDS). In general the AWG is the more versatile choice, but comes at a high costs. The DDS based solutions are often cheaper, and for specific setups equally good. One common issue is the bandwidth with which the output frequency of the DDS can be changed. A DDS is generally not operating standalone, and additional hardware is required to control the output of the DDS. This additional hardware can be a major limiting factor for output frequency updates of the DDS. The specific limit depends on the choice of hardware, which will usually be some kind of processing unit,

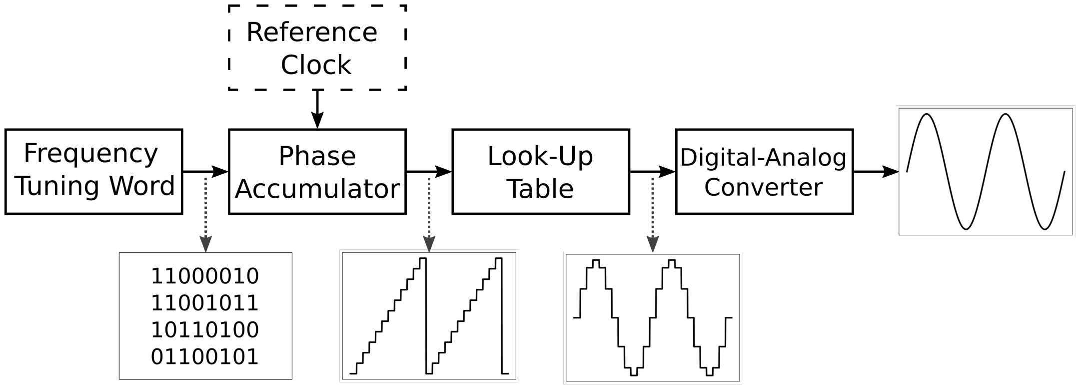

Figure 2.1: *Diagram of the components inside a Direct Digital Synthesizer. The boxes on dashed lines are showing the signal after each step. A frequency tuning word is set in a binary register, with is added to a phase-accumulator with every clock cycle. The phase accumulator creates a digital signal which, in this particular case of a sine signal output, resembles a sawtooth. This phase is then converted to amplitude by a lookup table, resulting in a digitized sine. Lastly, a digital to analogue converter converts the digitized sine to the final output signal.*

for instance a microcontroller device. Specifically microcontroller devices are in general easy to implement, but they are also particularly slow compared to other solutions. In this work, we have decided to base the RF source on a DDS controlled by a Field-Programmable Gate Array (FPGA). An FPGA is essentially a middle ground between a System-on-Chip (SoC) device, such as a microcontroller (MCU), and an application-specific integrated circuit (ASIC)[84]. The clock frequency of the FPGA device is running with a frequency on the same order as the clock frequency of the DDS itself, and does not e.g. limit the response time of triggers and time to update the RF signal, in the same ways as e.g. an MCU could have.

2.2 Direct Digital Synthesis

Direct Digital Synthesis is a fairly simple way to generate repeating waveforms. The method, which is shown in Figure 2.1, is as follows: First, a digital time-varying signal is generated, and then a digital to analogue converter (DAC) is used to create the final analogue signal. The digital signal is generated mainly by utilising a phase-accumulator, which is essentially just a counter, that will keep increasing up to a certain point, after which it wraps back to zero. A value is added to the phase-accumulator for each clock cycle in the device, and the value of the phase is used in a lookup table that converts the phase to an amplitude. In the final step, the amplitude is converted to a analogue signal by the DAC. The digital signal then depends on two parameters: The frequency of the internal clock, and the frequency tuning word (FTW), which is the value that gets added to the phase accumulator in each clock cycle. One of the beauties of this system is that the lookup table can in principle be chosen to create any form of periodic signal. Another advantage of a DDS is that the output signal will always be continuous in phase. That is, the analogue output signal will never show a jump in phase, even if when changing the FTW significantly between clock cycles. This is because the signals are generated only through modification of the instantaneous phase in the phase-accumulator.

In this setup, we will be using a DDS with a sinusoidal lookup table, in which case the repetitive angular phase has a range of 2π. This means that the amplitude is also cycling such that the value at phase 0 and 2π are the same, and the output frequency is therefore controlled by the rate at which the phase is changed. In more specific terms, the output frequency of the DDS is given by

$$\omega_{\text{out}} = \frac{M\omega_{\text{clk}}}{2^n}, \tag{2.1}$$

where M is the n-bit frequency tuning word, and ω_{clk} is the clock frequency of the DDS[85]. The clock frequency is typically fixed, and thus defines the range of possible output frequencies of the setup, while the bit-resolution of the phase accumulator defines the steps with which the system can change frequencies. The FTW can take any value from 1 to $2^n/2$, which corresponds to 1 full period of the output signal requiring 2^n or 2 clock cycles, respectively. The limit on the maximum value of the FTW, corresponding to 2 clock cycles for a period, is given by the Nyquist limit[86].

A simple sinusoidal signal is then generated with a DDS, by setting the FTW and running the DDS with any clock frequency. However, in many cases, a simple constant frequency is not sufficient, and signals with frequency ramps are needed. Fortunately, simple time-dependent frequency signals can also be created with a DDS setup, by modifying the FTW accordingly. For instance, a linear frequency ramp is created by increasing/decreasing the FTW with every clock cycle, such that the rate of change of the phase is increasing or decreasing continuously.

A DDS, therefore, provides a very stable and versatile platform for generating sinusoidal signals of static, or linearly changing frequencies. More complicated frequency sweeps are in principle also possible, and are limited by the implementation of the correct sequence of FTWs.

2.2.1 Controlling the DDS with an FPGA

We have established that the DDS is capable of generating a wide range of signals, by setting the corresponding FTW. However, as the DDS is a hardware device that is not working stand-alone, another device is needed to control the DDS, and specify the FTW. This task can be achieved in several ways, and here we choose to use a Field Programmable Gate Array (FPGA). Put it simply, an FPGA is an array of logic gates that can be connected in various ways through programming, in comparison to an ASIC where the connections are decided already at manufacturing. The main advantage of the FPGA is that it provides high operating speed and precision, similar to ASICs while being configurable as is known from MCUs. The disadvantage of an FPGA is that the programming is very low-level and that the devices usually have significantly less memory and functions compared to processors, e.g. the ATMEGA328P as known from the Arduino devices.

In this setup, we need high time-precision in the microsecond regime, as well as high reliability. The only feasible option is then FPGA devices, as the alternative microprocessors are changing the internal data-paths in real-time, and therefore lack predictable time-delays for computations. That is, the FPGA will perform exactly the same calculation each time, while a microprocessor will decide what to do on the fly, and thus might perform different calculations at different times, leading to different time delays on the μs scale. If, on the other hand, we required precision only on the ms scale, a microprocessor solution would typically be sufficient.

2.2.2 Our setup

Specifically, we use an AD9959 DDS Evaluation board from Analog Devices, and controlling it with the Artix-7 XC7A35T FPGA chip from Xilinx. Furthermore, we use the FTDI FT232H chip on an Adafruit breakout board to provide an interface between a desktop computer (through USB) and the FPGA device. The setup is illustrated in Figure 2.2. The FPGA is controlling

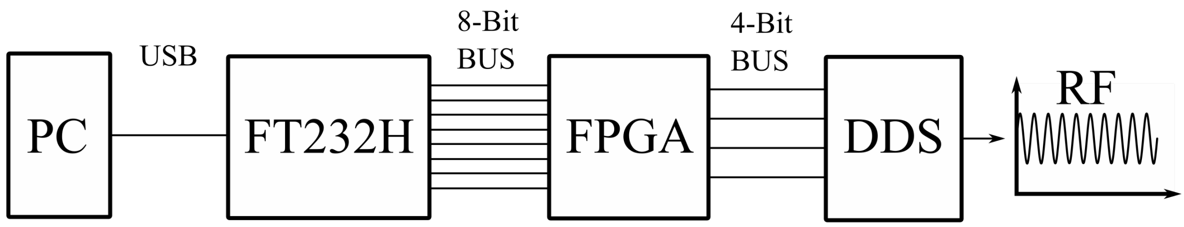

Figure 2.2: *Diagram of the communication setup from computer with high level code, to the DDS. A Python application on the computer generates the binary signal data, which are transferred to the FTDI interface through a USB connection. From here, the binary data is transferred to the FPGA memory using an 8-bit BUS. The data from the FPGA are transferred to the DDS in a 4-bit BUS, before it can finally be used to generate an RF signal.*

the DDS, but the FPGA still needs to be controlled by something else, which determines what output signal should be sent to the DDS. This part is done simply by USB communication with a desktop computer. The computer writes data directly to the memory of the FPGA, which is then read by the FPGA at runtime, and parsed and transferred to the DDS as needed. This means that this entire assembly can very easily be controlled through high-level scripts and application on a computer, without worrying about the specifics of FPGA programming.

The AD9959 DDS Evaluation board has a 32-bit resolution on the phase accumulator, and thus FTW, and the clock frequency is 500 MHz. The clock signal is provided by the FPGA, to ensure that the boards are in sync. In principle, the frequency could be increased, but that could lead to a potentially less stable system. The smallest possible frequency step of the setup is thus $\delta\omega = \frac{500\text{MHz}}{2^{32}} = 0.12$ Hz, which is much lower than the typical operation frequency of ~ 100 MHz, and thus allows sufficiently precise frequency tuning for this setup.

This setup is very versatile, as the FPGA can change the DDS settings 'on the fly', and thus allow for near-arbitrary output signals, at least on the FPGA clocks timescale which allows for a changed output every 2 μs. The main limitation is the memory of the FPGA, which in this particular case is 1800 Kb, corresponding to more than 10000 changes, meaning that realistically this limit will not be an issue for this setup. The basic operation of the system is as follows: The FPGA is programming the DDS, and sending the triggers to the DDS when the output signal needs to change. The DDS board has build-in modes for simple signals, e.g. a linear sweep, which can be activated by the FPGA without any further computations required. For more complicated signals, e.g. interlacing frequencies, or rapid jumps, the FPGA is controlling the DDS more directly during runtime.

This setup in itself is in principle good enough and is capable of outputting all the necessary signals. However, the output directly from the DDS board has very low power, on the order of 0 dBm, and is thus not suitable for driving an AOM. In order to prepare the signal for the AOM, it goes through a series of RF devices. First, the signal from the DDS is sent through a low-pass filter (built-in on the DDS board) to clean up high-frequency noise. From there, the signal goes to an analogue attenuator, which allows the user to control the amplitude of the output signal through some external signal. This could for instance be the experiment control, which will be discussed in Section 2.4.1. After the attenuator, the signal goes to a digital switch, such that the signal can easily be turned on/off with a TTL signal. Finally, the signal pass a preamplifier,

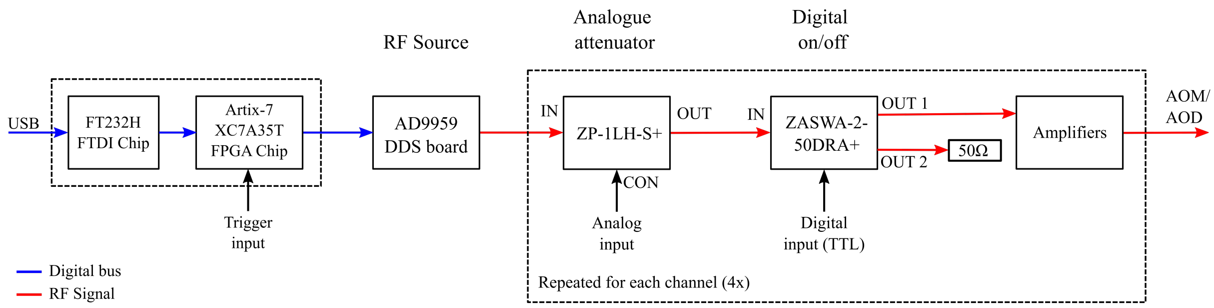

Figure 2.3: *Diagram of the components inside the FPGA-DDS based frequency source. From left to right, a computer provides the frequency data via a USB connection, which the FPGA reads and stores. When an external trigger is received, the FPGA sends the frequency data to the DDS, which then generates analog RF signals on 4 individually controlled channels. The signal on each channel then goes through an analogue attenuator, a TTL switch, and an array of amplifiers, before going to the experiment.*

before it enters the high power amplifier, which amplifies the signal by 2 W in order to drive the AOM. This array of devices is also presented in Figure 2.3. As the DDS board has 4 output channels, the entire setup is copied for each of the channels.

2.3 Assembling the device

The entire setup with DDS, FPGA, Amplifiers, etc., is assembled in a two-unit 19” rack case. Figure 2.4 shows an image of the finished assembly.

On the front panel, there are two micro-USB ports, one directly connected to the FPGA board, for programming, and one for transferring data to the FPGA memory via the FT232H USB interface. Furthermore, there is one column of BNC connectors for monitoring signals from the FPGA and for TTL inputs directly to the FPGA, and one column for fast TTL signals. The next connectors are for the front panel board, which handles both the attenuation and slow TTL signals. The board can either take external signals, which are then scaled accordingly to fit with the MiniCircuits attenuator and switch that is used inside the box, or the board can generate the voltages directly. Switches on the front panel are controlling if external or local voltages are used, and the locally generated voltages are also controlled from the front panel. This allows for easy switching between attenuation controlled by the experiment control setup, and manual control for alignment and debugging. The last row of connectors on the front panel is for the final output signals. The final components on the front panel are the switches for the high power amplifiers. These are important to avoid sending high power signals to devices that cannot handle them, or burning the amplifier with reflections from empty connectors.

The switches, attenuators, and amplifiers are all from MiniCircuits and are connected either by straight SMA connectors or with rigid SMA cables. Due to the number of different devices in the assembly, the box needs quite a few different voltages to power the devices. This is handled by a custom-built power division and distribution board which is designed for a ± 15 V input voltage. The board takes an input of $+15$V, GND, and -15 V, and produces ± 15 V, ± 5 V, $+3.3$ V, and $+1.8$ V.

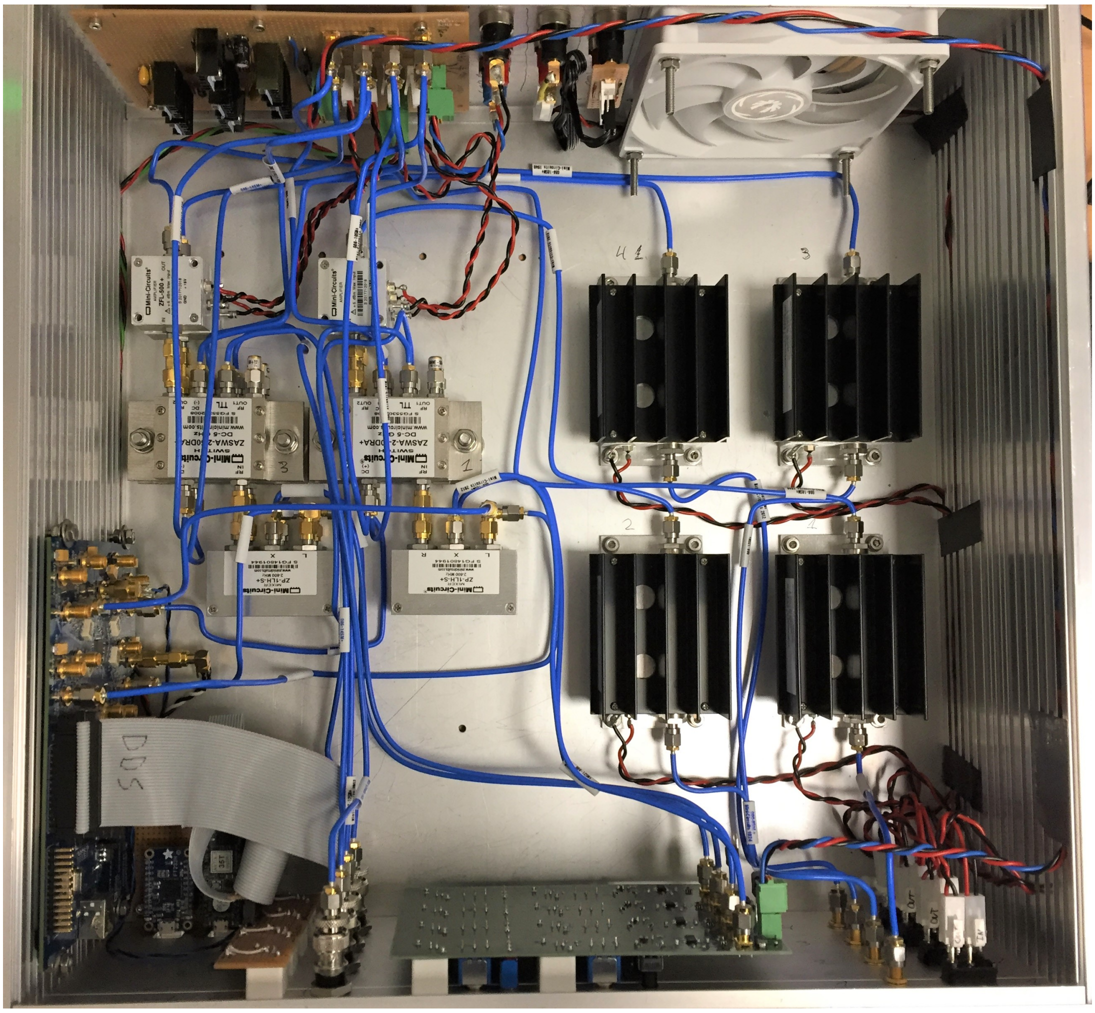

Figure 2.4: *Top view into the FPGA DDS device, with the front-panel in the bottom of the image. The power distribution board is placed top left, with the attenuators and switches below. In the lower left the DDS board is mounted on the side-panel, with the FPGA next to it. The black pieces on the right side are the high-power amplifiers, and the PCB on the bottom of the image is the front-panel board.*

2.4 Controlling the device

As mentioned above, the FPGA is controlled by data from a desktop computer, where the data can be generated by high-level scripts and applications. This requires that the FPGA is already programmed to take inputs from a computer. This programming is done using a Hardware Defined Language (HDL), which is not part of the scope of this work.

As the HDL programming is ready, the FPGA can be controlled simply by sending binary data directly from the computer, via a USB connection. The binary data is following a specific format, dictated by the HDL program. The first chunk of data is configuration data for the FPGA and DDS, and will generally not be changed. After the control chunk, each line of data corresponds to a specific output segment. That is, one segment could be a sawtooth wave with a specific set of parameters, a certain sweep, or a static signal. Each line contains 128 bytes and specifies all parameters for the particular segment. Each segment can specify a number of repetitions before automatically continuing to the next segment, or the FPGA can be triggered externally, forcing the change to the next segment. The reliability, stability, and speed of the FPGA responding to such an external trigger is investigated in the next section.

Thus, the task for the high-level application on the computer is to take a number of segments, convert them to the corresponding binary format, and then transfer it to the FPGA. In order to ease the use of this setup, a Graphical User Interface (GUI) application has been developed to control the FPGA, and in turn the DDS. The application is developed in Python 3.8, using PyQT5 for the graphical interface. The application is developed in two parts: One part to visualise and generate the signals, and one part that converts the data to the binary format and transfers the data to the FPGA. This segmentation is chosen for a number of reasons: First and foremost to prevent any bugs and issues in the main GUI application from affecting the communication with the FPGA. Secondly, it allows for the FPGA to be connected to one computer, and the GUI application running on another computer. Thus, the FPGA DDS box can be controlled remotely, and as will be described in Section 2.4.1, this allows easy interfacing the setup with the existing experiment control software. The application that generates and visualises the signals is a complete GUI, as shown in Figure 2.5, while the communication with the FPGA is a simpler terminal application. The software can then run in three different modes:

Local control In the case of local control, only the GUI application is running, and when the signals are ready to be transferred to the FPGA, the application will use the data-transferring script from the remote application. This mode is especially useful for testing and development when a single computer can be used to quickly and easily modify the output signal directly.

Remote control This mode is using the remote application on a computer that is connected to the FPGA, which listens for the data from the Control GUI that can be running on any other computer in the network. When the signals are as desired in the control GUI, the data is saved to a network drive, and a network ping is sent to the remote application, which will then read the data from the network drive, generate the binary data, and transfer it to the FPGA.

External lookup mode In this mode, the device is controlled by an external setup, from where the parameters and triggers are provided. In our experiment, the signals are controlled by the main experiment control software, which is also controlling all other lab devices. This is used when the setup is part of the experiment sequence, and the signal parameters must correspond to the specific timings and settings of all other devices in the sequence. This mode is very similar to the Remote Control mode, in the sense that the data are still transferred to the

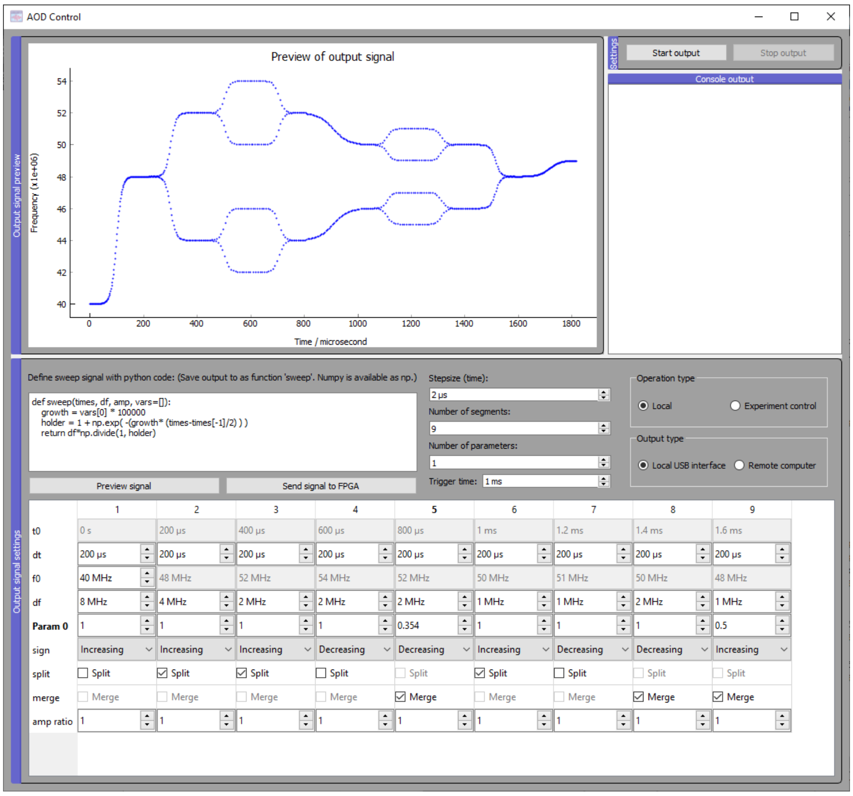

Figure 2.5: *Image of the GUI application used to control the interlaced multi-tone signals of the FPGA DDS box. The GUI window is split into three main sections. In the top left, there is a plot widget used to preview the output signal. To the right of the plot, there is a console output, used mainly for status messages and debugging information. In the lower part of the window is the main widget, used to specify the output signals through an input function, and a table of parameters.*

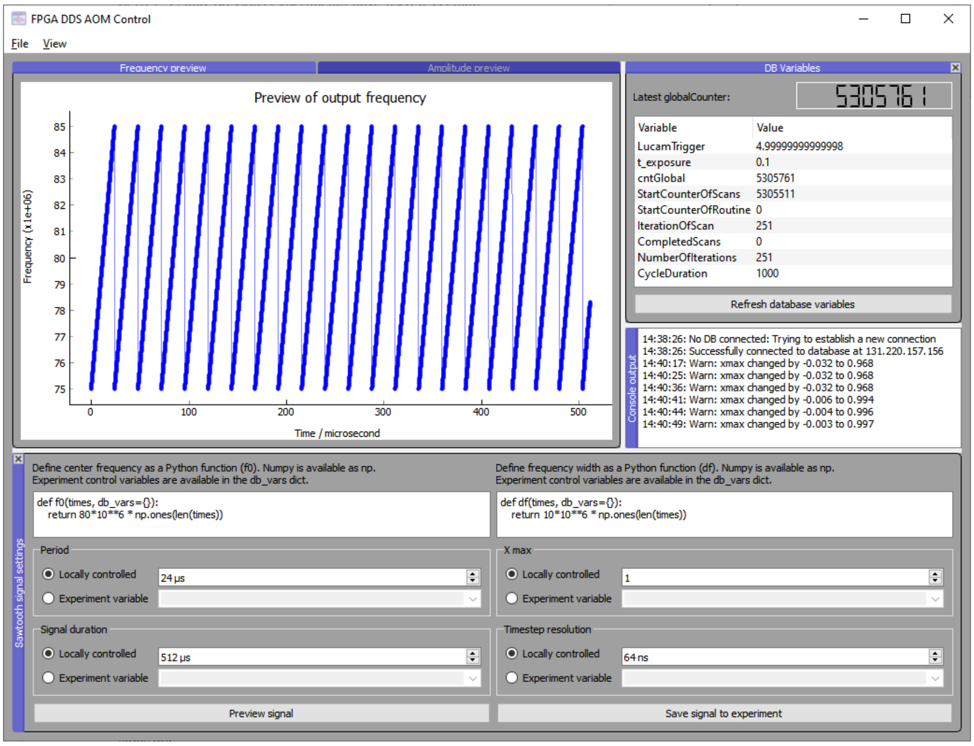

Figure 2.6: *Image of the GUI application used to control the segmented output signals of the FPGA DDS device. There are three main sections in the GUI: Top left is a preview of the output signal, top right is an overview of the parameters available in the database, and a console output. The lower section of the GUI is used to specify the signal parameters, and the functions used to generate the output signals.*

FPGA by the remote application, and the control GUI is still used to generate the data that are read by the remote application. The important difference is that the remote application reads variables from the main experiment software, and can use them as parameters in the signal.

2.4.1 Computer control setup for the experiment

In the external lookup mode, the DDS device reads parameters from a database, and then receives an external trigger to start the output. This mode is built mainly with the existing computer control setup in mind. This subsection will briefly describe the existing setup, and in particular how the new DDS device is interfaced with it.

Computer and device setup

In order to run the experiment, many different devices must be working together in sync. That is, each device needs to be configured properly, and start and stop at the correct time, relative to all other devices. A device could be a power supply, an RF source for the AOMs, a camera, or perhaps a shutter.

To control these different devices, an elaborate computer control setup is used. The setup is based on one main application (Cold Physics Experiments Control Software v1.4.4[87]) that controls most parameters of the sequence, determines which devices to run, and when the sequence should start and stop. The control application is also communicating directly with an ADwin Pro II device, which is used to send digital (32 channels) and analogue (16 channels) signals to other devices. The digital signals are mostly used to trigger other devices, while the analogue channels are used for e.g. attenuating RF signals for AOMs. The control software is working in unison with an SQL database that stores all parameters for each sequence. The database serves two main purposes: One is to provide an easy way for all devices to access the parameters when the sequence runs, and the second is to save parameters for future reference.

The entire control setup is spread out on several computers, both for practical purposes and in an effort to improve the stability of the system. The setup is currently using 5 different computers, 3 of which are essential for running the experiment, with the last two being used for data acquisition:

Control computer: The control computer is running the main control software, from where most of the experiment is controlled. This computer controls the starting time of each sequence, and communicates to the other computers with network pings whenever a new sequence is started.

Vault computer: The vault computer is hosting an SQL database that is used to store all experiment parameters for each sequence. When a new sequence starts, the control software will write data to the database, and all other computers and devices can then read the parameters from the database.

Apart from the database, the computer also hosts a network drive, that is used to store experimental data and results. This could for instance be images from an experimental sequence.

Devices computer: The devices computer is the busiest one of the setup. It communicates with all the experiment devices (power supplies, RF generators, pulse generators, etc.), that are usually connected through USB. Each device has its own application running on this computer (generally a Python 3.7 script), that handles the communication with the device. When the

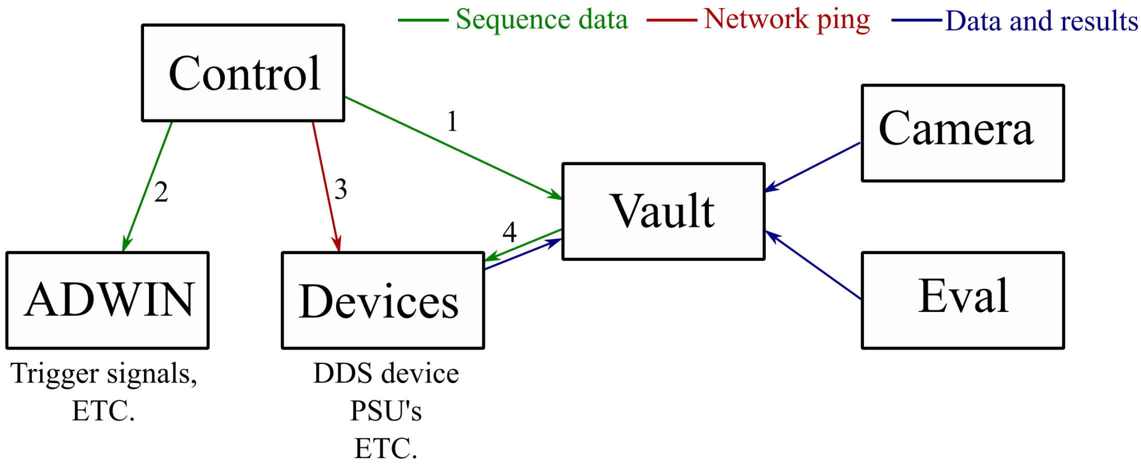

Figure 2.7: *Diagram of the computer control setup during a single experimental sequence. First, (1) the computer control saves parameters to the database, and (2) instructions are transmitted to the ADwin to configure experimental output. Then, (3) specific settings and a general network ping is sent to Devices computer, which configures experimental instruments. If needed, the devices will (4) read parameters from the database. The sequence is then ready to begin, and any data recorded by e.g. the camera and eval computers are saved on the Vault computer and in the database.*

sequence starts, the application will receive a network ping from the control software, and then the relevant variables can be read from the SQL database, before using this to configure the device for the particular sequence. As precise timings are difficult with software, the device scripts will usually only configure a device, which is then later triggered by a hardware trigger from the ADwin.

Eval computer: This computer is not a part of running the experiment sequence, but is critical for data acquisition, as it runs the Time Tagger, which is a high-speed, high-precision counting device, used to count triggers from the MCP or single-photon counters.

Camera computer: The camera computer is also not part of running the experiment, but is important for data acquisition, as it is managing the cameras. A number of cameras are used to image the MOT at different stages, and thus for instance perform time of flight measurements. The cameras are prepared similarly to the other devices and are also triggered externally in order to get a precise triggering time. When the data is acquired, it is saved to the network drive on the Vault computer.

Here we use the particular computer setup as described above, but the specifics could be adapted to fit any other experiment. The crucial part of the setup is the database, which is the backbone that ensures the relevant parameters and data are available to all devices when needed, and also ensures that the experimental sequences are properly saved and logged, for future reference.

The full experimental sequence is described in Figure 2.7. When the sequence starts, the control software will increment a global counter, and save all parameters to a new row in the SQL database on the Vault. After saving the data, it will run a Python script that sends network pings

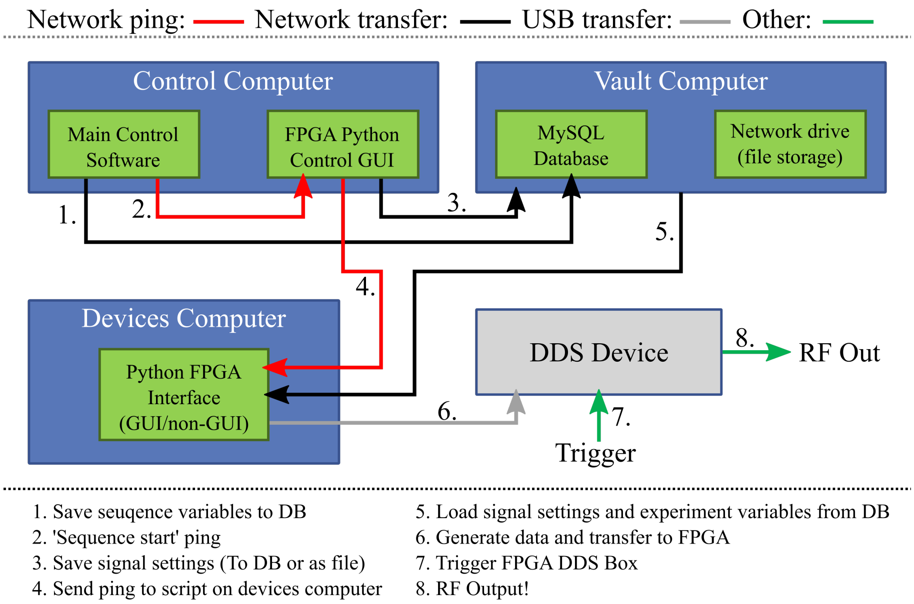

Figure 2.8: *Diagram of the communication between the computers in the experiment control, and the DDS device. The communication sequence is split into 8 distinct parts, as described in the text below the diagram.*

to all the device applications. Most device applications are running on the devices computer, and when they receive the network ping, they will read data from the SQL database, and use the relevant parameters to configure their device. From sequence start until all devices are ready and database entries have been written takes about 1 second. When the devices are all ready, the experiment sequence can begin and it is started by the control software that triggers the ADwin device. This will then trigger all the external devices used in the experiment sequence. The hardware triggers can be set with a time resolution of 20 μs. When the sequence is finished, the control software will increment the counter again, and start the next sequence. A typical sequence takes about 2 seconds, but the actual time depends on the specific measurement. Note that the control software is not aware whether the devices are ready when the ADwin is started, and issues can thus arise if a device is not ready within the set idle time of the sequence.

Interfacing the FPGA DDS frequency source

The FPGA DDS box described in the previous chapter is a prime example of one such experiment device. As described before, it can be controlled in several ways, where one method is interfacing with the experiment control setup. This subsection will discuss this interfacing. In this mode, the FPGA DDS box is connected to the devices computer with a USB connection, and the remote application of the FPGA DDS box is running on the computer, to handle the conversion to binary data, and the transferring of the data to the device. On the other end, on the control computer, a control GUI for the FPGA DDS box is running alongside the main experiment control software. The control GUI is used to choose the operation mode of the DDS, and select the type of signal (static frequency, sawtooth, segments). Apart from the basic settings, the control GUI also sets parameters for the output signal (amplitude, frequency, sweep rate, etc.), either by specifying values directly or by referring to variables from the experiment control software.

When this setup is running in the sequence, the control GUI will listen to the network ping

```python
import time

start = time.time()
task_to_measure()
stop = time.time()
duration = stop - start
```

Figure 2.9: *Example of a Python code snippet used to measure the duration of the tasks that are part of preparing the FPGA DDS device for each sequence.*

from the experiment control software. This is the case even if the control GUI is also running on the control computer. When the ping is received, the control GUI will save all parameters to a file in the network drive of the Vault computer, and then send a network signal to the remote application on the devices computer. As the remote application receives the signal, it will read the file from the drive on the Vault, and if any of the parameters are controlled by the experiment control software, the variables will be read from the SQL database, and substituted in, before the parameters are finally converted to the correct binary format, and transferred to the FPGA DDS box through the USB connection. This whole process is also described in Figure 2.8.

There are many steps from the control software starts the sequence by sending the first ping, and until the device has been prepared and is ready to be triggered. It is therefore a concern if this process is fast enough to avoid slowing down the sequence. The typical sequence takes about 2 seconds, with 1 second dedicated to the preparation of devices. The actual preparation time needed is different between devices, but all devices currently used in the lab, are well within 1 second. Thus, to avoid slowing down the sequence unnecessarily, the device should be ready to receive a trigger at most 1 second after the sequence begins. There are many small processes before the device is ready, but the main time consumers are: Saving the data from the control GUI to the network drive, sending a network ping to the devices computer, reading the data from the network drive, converting to binary, and transferring to the FPGA. Some of these durations are well known, and the rest can be estimated or measured precisely enough to give a reasonable estimate of the total duration. The reading/writing of files to the network drive, the conversion to binary, and the transfer to the FPGA are all measured using the time module in Python (See Figure 2.9). The resulting data are shown in Table 2.1, and visualised in Figure 2.10. Note that the network pings from the control computer are not included here. The first ping is when the sequence starts, and the delay is therefore not affecting anything, while the second ping is between the control GUI and the remote application. This delay will increase the total duration, but should be less than 10 ms, which is not a significant compared to the total duration.

When the binary data have been transferred to the DDS device, the device is ready for the start of the actual sequence, and is waiting for an external trigger. From this point, depending on signal type, and how critical timing is, the device can operate in two different ways: Either by being pre-programmed to update signals (e.g. move to next segment) at pre-determined times after the initial trigger or by updating the signal on following external triggers. Updating on triggers that are pre-programmed, and thus internally controlled, is done with nanosecond precision, with no further delays or noise on the trigger time.

Contrary to the internal triggering, the external triggers will have some unknown delay between the trigger and the actual update of the output signal. This is due to the way the

Table 2.1: *Duration of each part from start of the sequence, until the FPGA DDS box is ready to output a signal. All delays are measured inside Python, and is averaged over 10 repetitions. The number in parentheses is the standard deviation in units of the last digit.*

Description	Duration / ms
Save data to vault	9.8(3)
Load variables from database	64.8(55)
Load data from vault	13.5(37)
Convert data to binary	9.99(2)
Transfer data to device	< 5

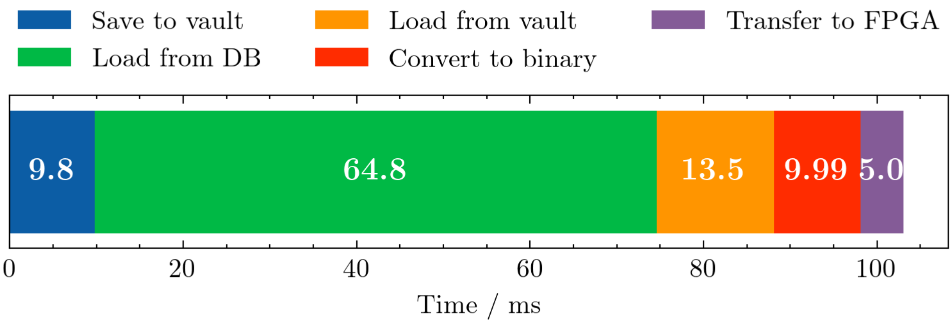

Figure 2.10: *Chart of the time it takes for each step from the sequence starts, until the FPGA DDS is ready to output a signal. All numbers are measured in the Python code, and averaged over 10 measurements. Standard deviations are provided in Table 2.1.*

FPGA is working and essentially reflects that the FPGA is only performing computations with a tick in internal clock. Thus, the time from the external trigger, until the next clock tick in the FPGA, will cause jitter on the update time. This will be further investigated in Chapter 3.

2.4.2 Segmented output signal

So far, the general control setup has been discussed. Further, the system can generate two different types of output signals: Firstly, the output can be a single frequency that can be kept fixed, be changed, or swept, and secondly, the output can be multiple frequencies that can be interlaced and individually changed. The first setup, i.e. the segmented output, is used to implement SWAP cooling as discussed in Section 1.2.4, while the second would be used for moving dipole traps as discussed in 1.2.5, and splitting atomic clouds. The two operation modes are fundamentally different on the FPGA level, and switching between them requires to re-flash (reprogram) the FPGA. This is rather easily achievable but is something that can only be done between experiments.

In the case of a segmented output, a number of static or sawtooth type signal are defined, and the system can go through them one by one. That is, a list is defined, and when the setup is triggered, the first segment will run, either for a predefined amount of time or until the next trigger is received. In principle, each segment could include more complicated waveforms, but the current setup only allows for static frequencies or repetitions of linear ramps. That said, with the current setup, many more complex waveforms could be well approximated by a number of linearly ramping segments. The segments are all controlled by rows in the FPGA memory, which sets the limit of the number of segments. The FPGA memory is 1800 Kbits, of which

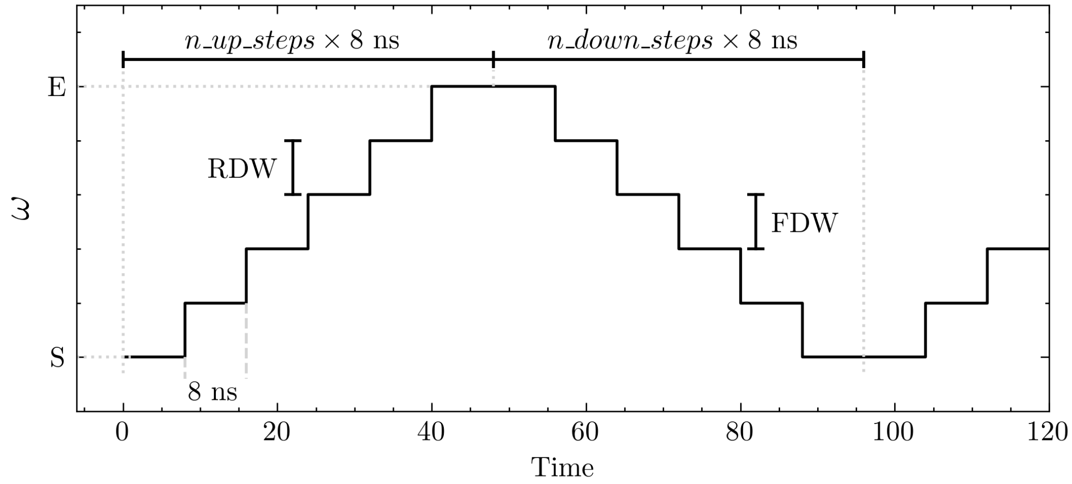

Figure 2.11: *The sawtooth waveform generated by the DDS, based on the parameters in Table A.1. Each single step is 8 ns long, and the parameters specify the change of frequency when ramping up and down, as well as the duration of the increasing and decreasing ramps.*

about 1.3 Kbits are used for configuration data, and therefore the number of segments is limited to 28000 lines of 64 bit. Each segment takes 3 rows for the parameters, and thus a total of 192 bit. The structure of each of the three rows is described in Table A.1 and together the rows specify the upper and lower frequency of the sweep, the rate of change of the frequency when increasing and decreasing, and the number of repetitions of full cycles (once cycle up and down in frequency) before continuing to the next segment. These parameters and their effect on an example waveform are shown in Figure 2.11.

2.4.3 Interlacing of multiple frequencies

The second mode of operation is developed to allow for moving dipole traps, and splitting and merging dipole traps. The idea is still to achieve this by sending the dipole trap beam through an AOM/AOD, which is driven with multiple frequencies. The DDS setup does not allow for true multi-tone signals, so the solution is to instead interlace multiple frequencies. That is, to rapidly switch between different frequencies, such that at any moment, only one single frequency is generated.

In this mode, the memory rows are also 128 bits long, although the structure is a bit different. The binary structure in this mode is shown in Table A.2. The first row is specifying a starting frequency, while the following rows specify the change of frequency relative to the final frequency of the previous row, the sign of the frequency change, the ramping speed, and whether the segment is splitting or merging beams, by starting or stopping the interlacing of frequencies. Based on the parameters, each segment will then go from the start to the end frequency, in increments of $2\,\mu$s steps, as is shown in Figure 2.12. That is, the ramps are not truly linear frequency sweeps, but rather a collection of static frequencies in $2\,\mu$s bins. This $2\,\mu$s period could in principle be shorter, but with the trade-off of creating more sidebands. With a single beam (no interlacing) each segment will simply increase (or decrease) slightly in frequency from the previous one until the desired level is reached. If a segment is to split a beam, and thus interlace two (or more) frequencies, each of the $2\,\mu$s segments will be dedicated to one of the interlaced frequencies. The FPGA will by itself calculate how the frequencies are changing for

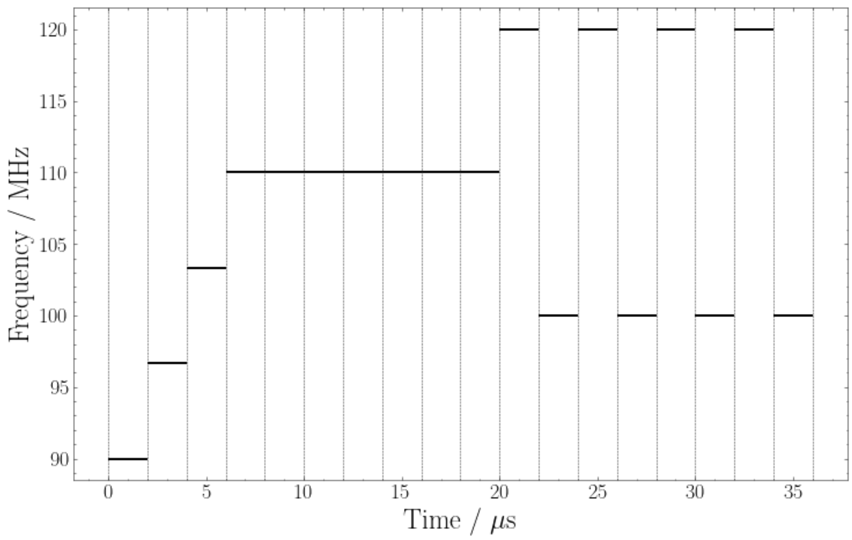

Figure 2.12: *Example of a signal with interlaced frequencies. The signal begins at* 90 *MHz, and sweeps the frequency to* 110 *MHz, before remaining static for a few cycles, after which the frequency interlacing is started, splitting the signal into one part with increasing frequency, and a part with decreasing frequency. After the two interlaced signals have reached* 120 *and* 100 *MHz, they remain static for the rest of the sequence.*

the interlacing, so we still only provide the frequency change, which is then used to increase one frequency, and decrease the other. The current setup is limited to at most 8 interlaced beams, but could in principle easily be increased, however this would increase the time required to cycle all frequencies ($t_{\text{total}} = n \cdot 2\,\mu\text{s}$), and if to many frequencies are added, the process would no longer be adiabatic. In order to keep the time between addressing each segment constant regardless of the number of interlaced beams, the segments are bunched, such that for 2 interlaced beams, the first 4 segments will address the first frequency, and the next 4 the second frequency. For 4 frequencies, two adjacent segments are addressing the same frequency, while for 8 frequencies, the individual segments are all addressing different frequencies.

Chapter 3

Characterisation of the DDS device

The final device should in principle be very stable and precise, due to the independence of higher level programs and hardware in the form of regular processors. The choice of an FPGA as the main logic unit in the device is expected to make the system very reliable altogether. Though, to confirm this expectation, a simple setup is created to test the precision and reliability of the system. In the following, we will investigate and test two features of the device: Firstly, the stability and precision of the external trigger, to confirm whether there are any delays on the external trigger, and if it can be used to reliably interface the device with an external controller like for instance the experimental control. The trigger analysis is done both for static and swept frequency signals, and for both the initial trigger as well as any consecutive triggers. Secondly, we investigate the creation of multi-frequency signals by interlacing, the timing and precision when switching frequencies, and briefly touch upon higher order effects like mixing.

3.1 Stability of external trigger

The main concern is the stability and delay on the triggering of the RF output signal, as that will directly affect the SWAP cooling during an experiment sequence, as discussed in Section 1.2.4. Thus, the test setup we implement takes an RF signal with multiple frequency changes, and measures the delay from the external trigger, until the output signal is updated. This is done both with the initial trigger, that turns on the signal, and for a triggered change of segment within the signal. The signal used for this case is shown in Figure 3.1.

The output signal of the DDS is connected to a Lecroy Wavepro 404HD oscilloscope, which with a samplerate of 10GS/s on 4 channels is fast enough to record the raw waveforms of the 70 to 100MHz signals we use in this case. Apart from the RF signal itself, two of the IO ports on the FPGA board are used to output monitoring signals to the scope. The first signal is a TTL envelope signal, that is high when the FPGA sequence is running, and is low otherwise. The other monitoring signal is an internal trigger signal, that is high whenever the frequency is increasing, and low when decreasing. As we are measuring static frequency segments in this case, the signal will momentarily go low when a new segments starts, but will otherwise always be high. The last signal that is recorded, is the external trigger, which is generated by a Rigol DX1022Z signal generator. The triggering signal is split in two, and goes to both the RF device, and the oscilloscope, where it is terminated with 10MΩ, in order to keep the voltage high enough for the RF device. All other signals are 50Ω terminated by the oscilloscope. An example of all the measurement of these 4 signals is shown in Figure 3.2.

This setup was used to measure 100 traces with the RF signal in Figure 3.1. Furthermore, another 100 traces with a frequency range of 7 MHz to 17 MHz were also measured. For the

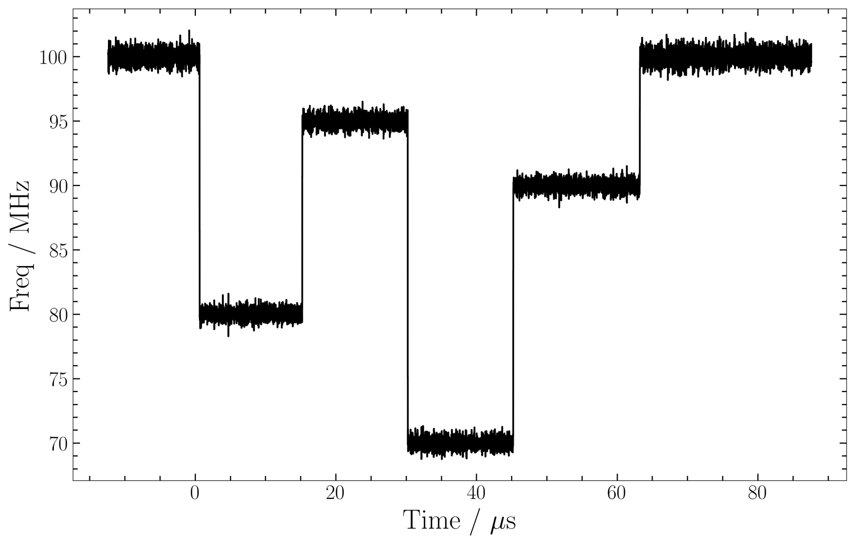

Figure 3.1: *The RF signal used for the trigger analysis. The signal starts at at the last frequency of the previous sequence, and then goes through 5 segments of different static frequencies. The last frequency is continuously outputted until the next sequence is started.*

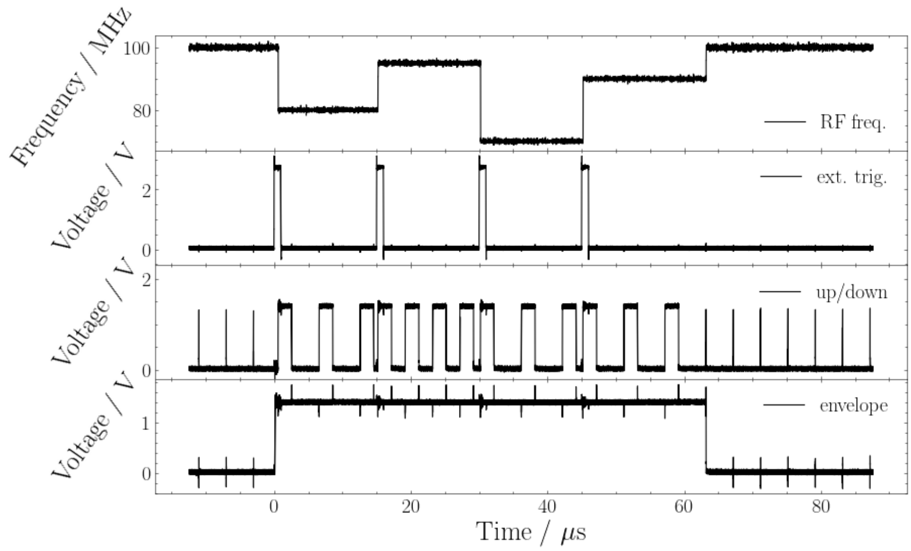

Figure 3.2: *Measured traces of RF frequency, the FPGA Envelope signal, the up/down internal trigger, and the external trigger. When the first external trigger is received at* $t = 0$ *the envelope signal goes high, and remains so until the last segment is finished. The up-down signal is mostly relevant when the frequency is swept, which is not the case here, but it is still visible that the signal is momentarily updated with each external trigger. Each segment runs until an external trigger is received, apart from the last segment that terminates by itself.*

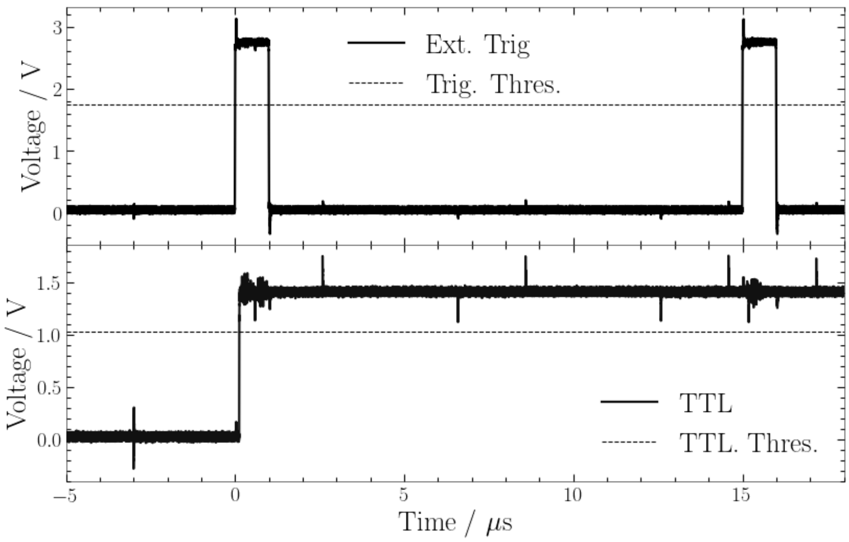

Figure 3.3: *Envelope signal from the FPGA, showing cross-talk between signal in the form of extra spikes on the signal every* $6\mu s$*. The dashed line is showing the threshold used for estimation of the trigger point. It is clear that the threshold is slightly higher than* 50% *of the high level of the trigger, due to the spikes from cross-talk in the device.*

lower frequency range, the relative frequency changes are the same, and the signal thus recreate Figure 3.1, although with the frequency axis scaled accordingly. Measurements with different frequency scales were carried out to determine if the order of magnitude of the output frequency affected the stability or delay on triggering the RF device. The results from both measurement sequences are similar, and the data are thus combined in the further analysis.

The first thing to analyze, is the time from the external trigger is received, until the FPGA registers the trigger. This time is found by comparing the external trigger, to the envelope signal from the FPGA. The analysis is performed in Python, and the trigger time on both signals is found by creating a threshold halfway between the maximum and minimum value of the signal, and using the time where the signal is crossing the threshold. This method has been compared to the time at which the oscilloscope registers the rising edge of the external trigger, and the method does add a delay of 1.3ns, although with no deviation between measurements (within the precision of the data). Also note that there are some cross-talk (See Figure 3.3) on the FPGA signals, which means that the maximum value is slightly higher than the 'high' value of the trigger, and this is likely the cause of the small delay on the determined time, as the threshold will then be slightly more than 50% of the signal. The time between the external trigger, and the envelope signal, determined by using this method, is summarized in a histogram in Figure 3.4a. It is clear that there is some jittering of the trigger delay, of 4.7ns. This jitter is actually expected from the device, as it is present due to the lacking synchronisation of the external trigger with the internal clock of the FPGA. That is, the FPGA is running at a clock of 250MHz, which corresponds to a clock tick every 4ns. The FPGA is essentially not doing anything between ticks, and can thus only register a high signal on the trigger once every clock cycle. This means that the external trigger will 'jitter' by up to 4ns, depending on when the external trigger is received, relative to the clock ticks. This, however, only accounts for 4ns, out of the 4.7ns in the data.

The last 0.7ns are not expected in the same way, and are a result of a slight design and measurement flaw in the system. The external trigger signal, and the envelope signal, are shown

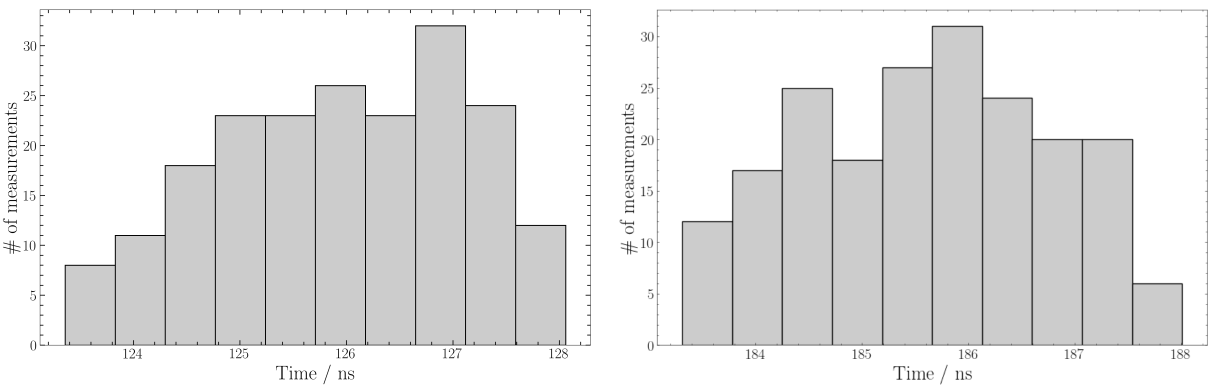

(a) *Distribution of the time from the external trigger until the FPGA TTL turns on. The histogram is based on* 200 *measurements of the delay, measured with the setup described in the main text. The delays are distributed from about* 124*ns to* 128*ns, and are not perfectly distributed evenly, however there is also no very clear tendency to a non-random distribution.*

(b) *Time until the RF signal is changed after an external trigger, while the device is already outputting a segment. The histogram is based on* 200 *measurements, which roughly evenly distributed in the range from just below* 184*ns to* 188*ns.*

Figure 3.4: *Distribution of the response time on an external trigger for* ***(a)*** *the initial trigger and* ***(b)*** *a second trigger updating the output signal to the next segment.*

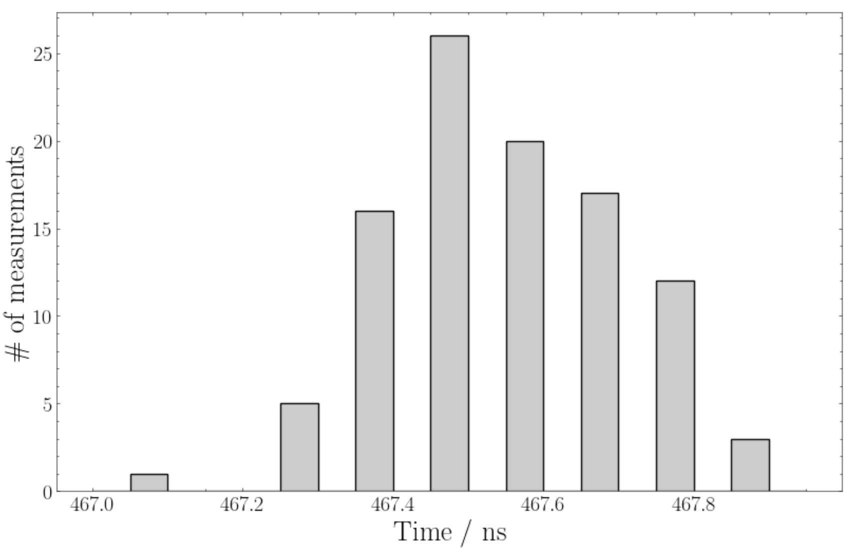

Figure 3.5: *Distribution of the time from the FPGA envelope signal, until the actual frequency change of the output signal from the RF device. The histogram is based on* 100 *measurements, and shows that the delay is distributed within a range of* 0.7*ns, which is very close to the measurement resolution of* 0.1*ns. Each bin is* 0.05*ns wide, and the gaps thus show that the times are constrained to steps of* 0.1*ns, which fit with the time resolution of the oscilloscope (at* 10*GS/s).*

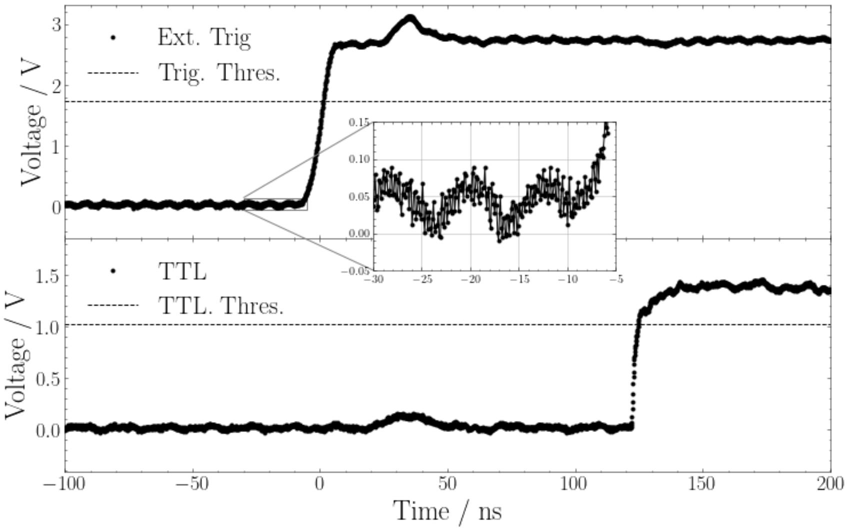

Figure 3.6: *An example of a trace of the external trigger signal, and the output TTL signal from the FPGA. It is clear that there is an oscillation on the* 0.1 *V level in both signals. The horisontal dashed line marks the threshold used to determine the trigger point, and the vertical dashed line marks the resulting trigger time.*

a much shorter timescale in Figure 3.6. Here, it is clear that there is a constant oscillation on both signals with an amplitude of 0.1V. This signal is likely a result of a fluctuation in the power load of the device, due to the RF output, which causes the overall voltage levels inside the device to fluctuate. Regardless of the origin, the oscillation is present, and can cause a slight jitter on the measured trigger time: The main trigger signal is changing from 0V to 3V in about 10ns, which corresponds to a slope of 3.3ns/V. Thus a fluctuation of 0.1V could then change the measure time by $0.1\text{V} \cdot 3.3\text{ns/V} = 0.33\text{ns}$. This jitter could happen on both the determination of trigger time on the external trigger signal, and on the envelope signal, leading to a total jitter of about 0.7ns.

The histogram in Figure 3.4a shows times on the level of 120ns, but the time before the actual output signal is actually about 450ns more. This further delay is the same with every single cycle, and is due to the calculations and data transfer needed to prepare the DDS for the change of signal, which is summarized in Figure 3.7. In detail, when the FPGA reads a high value on the external trigger, it waits for 124ns, to determine if the signal is a true trigger, or a false trigger in the form of noise on the signal. This is added to ensure that the system will not trigger falsely, even if the trigger signal becomes very noisy. When it is determined that the signal is a real trigger, the FPGA will proceed to read the parameters of the upcoming output signal from the FPGA memory. This process takes 44ns, as it needs to read 3 memory addresses, and each takes 2 clock cycles (of 4ns), plus an added 20ns in unused clock ticks to ensure stability of the system. When the parameters are loaded from the memory, the data (22 bytes) is transferred to the DDS. The transfer takes 352ns, as it requires 44 clock cycles, but the DDS communication clock is running at half the frequency of the FPGA. At this point the FPGA is done, and turns on the envelope TTL signal. The FPGA transfers data in to the buffer of the DDS, but the DDS needs to load it into registers before it can use the data. This process takes another 64ns, after which the RF signal will finally be changed. In total the delay is 580ns, plus the jitter of 4ns.

Note that the time required to load data from the FPGA memory, transfer to the DDS, and

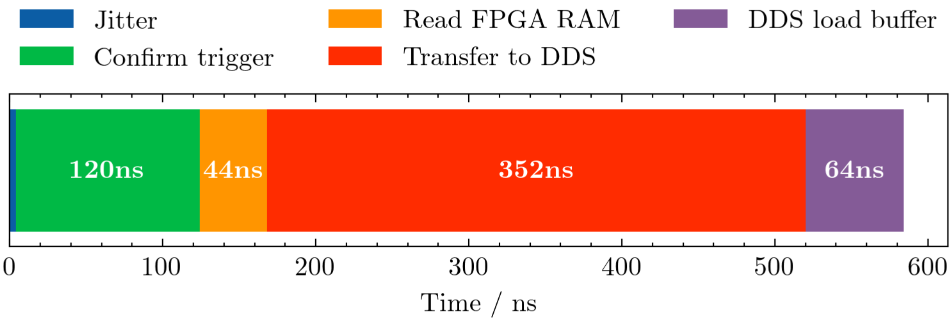

Figure 3.7: *Contributions from each of the delays as part of the process from external trigger, until the update of the output signal. First, there is a* 0 *to* 4.7*ns jitter, then a delay to confirm that the trigger is not false, before the data are loaded from memory and then transferred from the FPGA to the DDS. In the last step, the DDS reads the transferred data from the buffer, and outputs the new signal.*

load to DDS registers (460ns), is somewhat setting a lower limit on the duration of a segment, as this is the time required to prepare the next segment. That is, in principle a segment can be as short as 8ns, but it is not possible to proceed to the next segment before 460ns has passed. This limit could in principle be lowered by increasing the 4 wire SPI interface between the FPGA and DDS to an 8 wire interface, essentially halving the transfer time. This would, however, require a different choice of DDS board, e.g. the AD9910, which has an 18 wire SPI. Another approach would be to simplify the data structure, such that the amount data to transfer is lower. Although it should be noted that currently only 22 bytes are transferred to the DDS for each segment, so a simpler data structure does not seem to be a feasible improvement.

The analysis so far have only considered the initial trigger, turning on the sequence of RF output segments. However, the setup also allows for the use of an external trigger to continue to the next segment. In this case, the same jitter caused by the clock cycles of the FPGA is expected. The measurements used until now, also include external triggers to switch segment, and we thus use the same data set for this analysis. The external trigger is recorded directly on the oscilloscope, while the determination of the time at which the signal is updating is based on the up/down signal from the FPGA. That is, the up/down signal will go from low to high when the RF output is updated, and thus provides a useful point of reference. An example of measured traces, with the up/down signal, is visible in Figure 3.2.

Recording the time from the second external trigger, until the up-down signal is updated, produces the distribution in Figure 3.4b. It is clear that the distribution has the same width as previously found, which specifically is 4.7ns. This is expected following the same arguments as for the initial trigger. The important thing to note, however, is that the response time is distributed from 184ns to 188ns in this case, which is 64ns more than the previous case in Figure 3.4a. This is due to the fact that, opposed to the previous measurement, we are now including the time the DDS spends loading the buffer, although not the transferring of data from the FPGA to the DDS, as that is done as soon as a segment begins. That is, when the first segment begins, the data for the next segment is transferred to the DDS, precisely in order to avoid transferring the data later. This means that the only delay we expect in this situation is the waiting time to confirm the trigger (120ns), and the time the DDS takes to load the transferred data from its buffer to the registers (64ns). In total, the expected delay is thus $120\text{ns} + 64\text{ns} = 184\text{ns}$, plus the jitter of 4.7ns. The data thus shows that the delay times are behaving as expected, and it is confirmed that the delay on reacting to an external trigger is precise within 4.7ns.

We have determined that the fluctuation of the delay from the external trigger, until the FPGA registers the trigger, is less than 5ns, but it is still unclear if there is any fluctuation afterwards. When the FPGA gets the trigger, it transfers data to the FPGA, which then needs to load it from the buffer, before the an RF signal can be generated. This delay is measured by comparing the envelope signal of the FPGA, to the time where the signal is actually updated, which is determined by the up/down signal. The delay times from the 200 measurements leads to the distribution shown in Figure 3.5, where it is clear that there is no significant jitter. In fact, the standard deviation of the measurements is only 0.16 ns. Thus, without further analysis it is clear that there is no jittering from this part of the process. This was also expected, as the loading from the buffer will take the same amount of clock cycles every time, as long as the amount of data is kept constant, which is the case for us. Therefore, any jittering on this loading time would be caused by instabilities in the clock signal, which would generally not be expected on any relevant level.

It is now clear that there is a fluctuation on the response time to the external trigger due to the FPGA clock frequency, and low-amplitude oscillations on the TTL signal, but that the fluctuation is less than 5ns. This fluctuation is present on both the initial trigger and following segment triggers, and is in either case much lower than the precision required by the experiment which is running with a precision on the scale of microseconds. The delay itself is much larger, but is the same every single time, and is thus something that can be compensated for in the experiment sequence.

In terms of pure long-term stability, a thorough test has not been performed. Though, the system have been running continuously with a trigger every second for some weeks, without showing any problems. The output signal was not monitored actively, but only in short glimpses, and after some weeks the output was still as expected, and the triggering was still running fine. Thus, it is determined that for this part in particular, the setup is by far stable enough for implementation and use in the main experiment.

3.2 Sweeping segments

So far, the analysis of triggers have been based on segments of static frequency signals, but the main use case of the setup would be with sweeping signals. In the case of sweeping signals, the behaviour is exactly the same as the static segments, for the initial trigger. That is, there is a jitter of 4.7ns, and the delays from preparing and transferring data, as shown in Figure 3.7. Even though the behaviour on the initial trigger is the same, the system will behave in a slightly different manner when using an external trigger to start the next segment.

As the segments are sweeping, the frequency is naturally time-dependent, with a period dependent on the sweep rate and frequency range. It is not desirable to suddenly change the signal mid-sweep, if the setup is used for e.g. SWAP cooling, as this could lead to atoms left in the exited state. Thus, the setup is made such that a sweep will always finish, before continuing to the next segment, regardless of the trigger time. The process of using an external trigger to go to the next segment is therefore as follows: When the external trigger is received, the FPGA will wait for 120ns to verify the trigger signal, after which the FPGA will wait until the current sweep is finished, at which point it will tell the DDS to continue to the next segment. The data for the next segment would already have been transferred to the DDS before the trigger time, and is thus not a part of this process. It is expected that the 4.7ns jitter on the trigger time is still present in this case, as all relevant parts of the setup remain exactly the same.

3.3 Multitone signals by interlaced frequencies

The last operational mode of the DDS setup is interlacing frequencies, as a way of emulating multitones to create arrays of dipole traps. In principle arrays of dipole traps can be created in two ways: By driving the AOM/AOD with multiple frequencies at once, or by interlacing single-frequency signals, such that the input beam is quickly switching between different deflection angles. Each method has its own set of advantages and disadvantages, but as our RF source is based on direct digital synthesis, it can not create true multitone signals, and we are thus left with the interlacing as the only viable solution.

The basic idea is that the AOM is driven by a single frequency at a time, but that the frequency is rapidly switching between an array of values, such that the deflection angle of the beam is in turn also changing rapidly. An example of such a signal is shown in Figure 3.8b. This should create n time-averaged output beams, each with a deflection angle corresponding to one of the n input frequencies[88]. If we define τ as the duration for which each individual frequency is driven, that is, the time between each change in frequency, and thus also the time each beam angle is addressed, we can write the period of a whole cycle as $\tau_{\text{cycle}} = n\tau$. That is, the time in which all n frequencies are addressed once. In this setup, the intensity of each beam would then be reduced as a consequence of the time-averaging[89]. In our setup, the beam will be continuously switching between the n beam angles, addressing each angle for equal amounts of time, and the time-averaged optical trapping power is therefore the input beam power divided by the number of traps, n. Note however, that this is no different from the reduced intensity achieved by driving the AOM with a true multitone signal, which would also divide the input power equally between the outgoing beams, although continuously instead of by time-averaging. An important detail to note about the time-averaged traps is that there is a lower limit on the switching rate, if the dipole traps are to create an average potential. Specifically, the switching rate must be faster than the trapping frequency of the atoms, as the atom otherwise would essentially experience the trap turning on and of, instead of just an averaged potential with a lower trap depth.

One of the major issues of using interlaced single-frequency signals for arrays of dipole traps, is that the interlacing of frequencies is essentially similar to mixing the single frequency with a square wave. That is, the toggling on and off of each single frequency is similar to mixing a static signal with a square-wave signal with the period of the frequency toggling. This mixing of signals will introduce sidebands to the spectrum of the input RF signal, and might thus also create sidebands on the outgoing beams from the AOM. Figure 3.8a shows the spectrum of the signal in 3.8b which is switching between $\omega_1 = 160\,\text{MHz}$ and $\omega_2 = 200\,\text{MHz}$ at a switching frequency of $\omega_{\text{switch}} = 125\,\text{kHz}$. The figure clearly shows multiple orders of sidebands on the initial signal, although this also assumes perfect mixing, and no filtering of the signals, and the real-world case might not be as dominated by sidebands as this suggests. Although it is still very clear that the interlacing of signals can create sidebands, which is a potential issue to be aware of when implementing the setup in the experiment. Note also that higher order mixing terms should generally be neglected in this setup, as the AOM itself would act as a low-pass filter, thus removing any potential higher order terms.

3.4 Amplitude corrections

The goal of this setup is to accurately control the frequency of a laser, by passing through an AOM or AOD. The AOM/AOD takes care of modulating the laser frequency with an RF frequency, such that we only have to provide the RF. Even though we can simply drive the

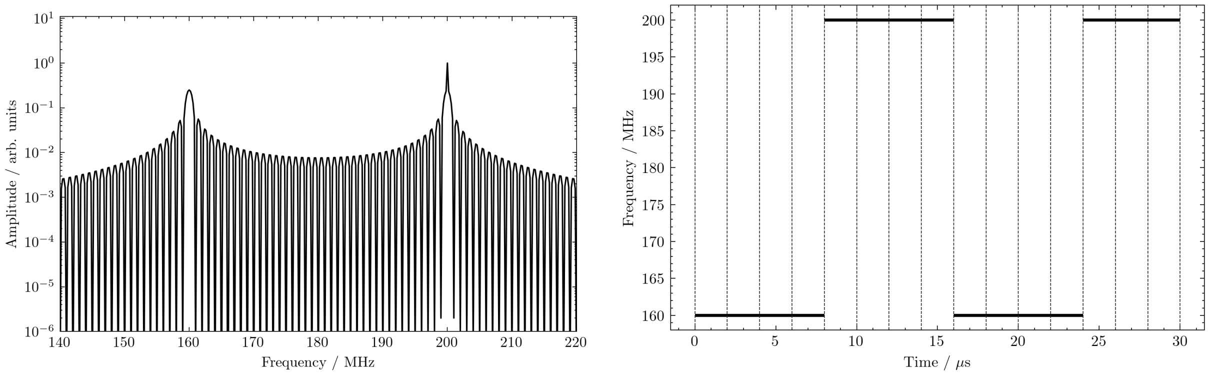

(a) *Plot of the frequency spectrum of a sine wave interlacing between two frequencies $\omega_1 = 160\,MHz$ and $\omega_2 = 200\,MHz$, at a switching frequency of $\omega_{switch} = 125\,kHz$. The mixing clearly creates multiple orders of sidebands on the original signals, which will in the end lead to sidebands on the laser beam*

(b) *The signal with interlacing frequencies which creates the waveform in **(a)**. The frequency is going between $160,MHz$ and $200\,MHz$ with a change every $8\,\mu s$.*

Figure 3.8: *Illustration of a **(b)** simple interlaced signal, and **(a)** the frequency spectrum when mixing with a square signal, which is the behaviour caused by interlacing.*

AOM/AOD and call it done, one issue that arise is that the intensity of the outgoing beam from the AOM/AOD is frequency dependent. This is less of an issue in an AOD as it is in an AOM, but the issue is present in both cases, and will cause the outgoing beam intensity to change throughout the frequency spectrum. Depending on the scale of this intensity change, it might not be an issue, but here we develop a system to correct for the intensity response of the deflector, such that the intensity of the outgoing beams remain constant within the relevant frequency region.

In the setup the FPGA is controlling the DDS, which generates the raw RF signals that goes through a number of attenuators and amplifiers before going to the AOM. In this setup, the amplitude can be controlled in two ways. The first, is by using the dedicated attenuators in the setup, which allows for control by an external signals (eg. from the experiment control), while the second way is to use the amplitude settings on the DDS board. This register would have to be set by the FPGA, and the register values would then be used to limit the amplitude by controlling the DAC. This 10-bit register allows for scaling the amplitude from zero to the maximal voltage in 1024 steps. Generally there would be no reason to limit the amplitude through the DDS, as it requires more from the software level, and the amplitude could easier be modulated through the analog attenuators. However, in this case, we chose to use the DDS amplitude settings for the corrections, as the device would then be more self-contained, meaning that no external signals are strictly required for operation. The idea is that the device is calibrated once (with the specific AOM and beam path), and then the intensity is stable, regardless of the analog attenuation. It should be noted however, that the attenuation input in principle allows to couple this device together with a PID loop or similar, if intensity stability is a major concern, even though it would most likely be excessive in most cases.

Specifically, the way we implement the calibration is by creating a look-up table (LUT) on the FPGA, which is used to find the correct amplitude value for each frequency. That is, in the non-calibrated case, the FPGA would tell the DDS to output a frequency at the maximal amplitude, then go on the next signal at full amplitude, and so on. With the calibration implemented,

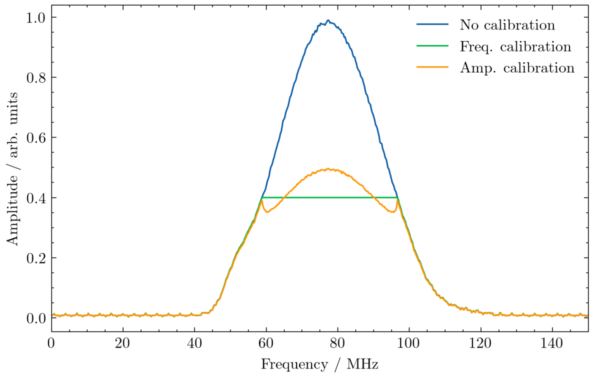

(a) *Oscilloscope trace of the measured light intensity after a double-pass AOM setup, without any amplitude calibration on the DDS device, and the theoretical traces after calibration, using both methods of calibration described in the main text. Note that the amplitude calibration does not create a flat RF amplitude from the DDS, in order to correct for the non-linear response in the AOM.*

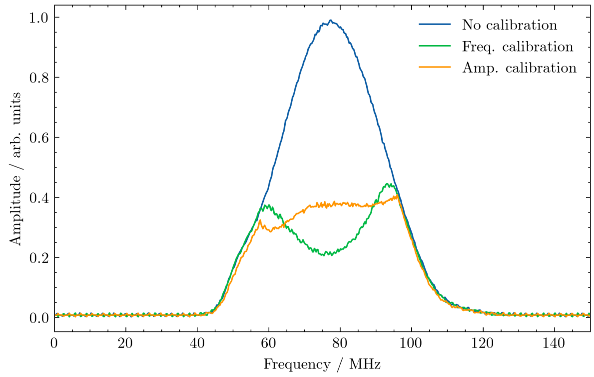

(b) *Oscilloscope trace showing the intensity of the beam as a function of AOM frequency before and after the calibration of the DDS device. The simple calibration is clearly undershooting, while the calibration methods that considers the non-linear intensity response is closer to being constant. The second order calibration still has two dips, around the point (80MHz) where the DDS amplitude vs. RF power calibration curve was measured, while being almost entirely flat in a small region around 80 MHz.*

Figure 3.9: *Traces of the beam intensity before and after calibrating the amplitude to achieve a flat intensity region. In **(a)** the un-calibrated measurement is compared with the theoretical prediction of the intensity after calibration, and **(b)** shows the beam intensities measured before and after calibration using both calibration methods.*

the FPGA takes the frequency, and use it to find a corresponding amplitude, which is then sent to the DDS together with the frequency value. This is done with 1024 amplitude steps, which for a large frequency range of 200MHz would allow for separate amplitude values for every $200\text{MHz}/2^{10} \approx 200\text{kHz}$, which should be plenty for most setups. The calibration data is measured once, and then loaded onto the FPGA, and is thus never considered by the higher-level Python applications.

In order to create the LUT, and test the calibration, a simple experimental setup is created. A laser beam is going through a double-pass AOM setup, meaning that the beam is passing through the AOM, and then gets reflected back through the AOM, after which the beam is sent through a fiber, and onto a fast photodiode. The advantage of this setup, is that by reflecting the beam back through the AOM, the final beam is not deflected by the frequency change, and a large frequency spectrum can thus be used without the beam moving away from the fiber tip. The fiber is used for convenience, and could in principle be removed from the setup. After the fiber the intensity is measured on a photodiode, which should have a fast response time, as the RF frequency for any realistic setup would be changing in a period of at most some 100μs. Note that the beam is passing through the AOM twice in this setup, and the calibration LUT should thus either be corrected or created with a different setup, if used for a situation where the beam is only passing the AOM once. This would for instance be the case for the interlaced signals. Here we use the measurements as is, and ignore any conversion between single-pass and double-pass setups.

The uncalibrated intensity response, meaning the response before the calibration is done, is somewhat Gaussian-like, with a peak near the centre of the AOM's frequency range, and

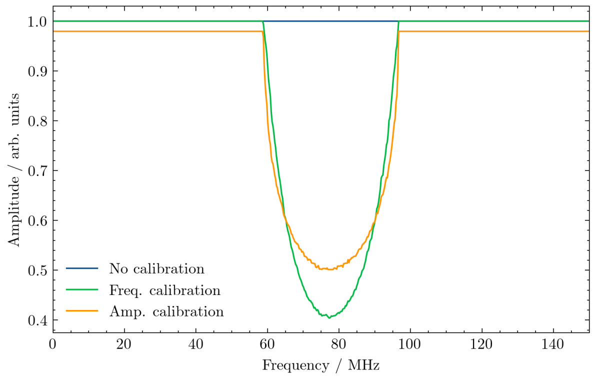

(a) *Shape of the look-up tables used to calibrate the DDS device for the two different calibration methods as described in the main text. The green trace is simply the inverse of the intensity value within the calibrated region, while the yellow lookup table is creating by inverting and multiplying by the corresponding intensity to amplitude conversion factor from **(b)**.*

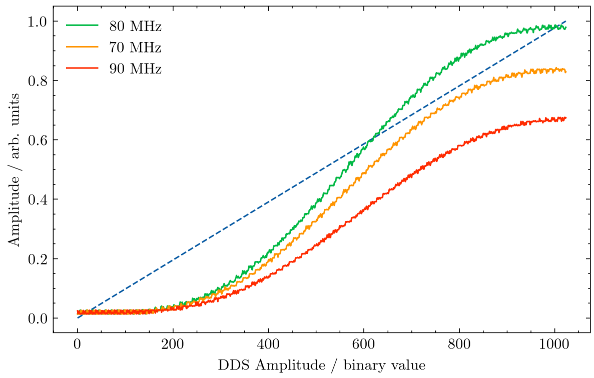

(b) *Measurement of the relation between output RF power from the device, and the set amplitude in the DDS. The dashed line is showing the simplest approximation, which assumes that the actual power and the amplitude are related only by a constant. The solid lines are measurements at different static frequencies, where* 80 MHz *is near the peak frequency, and* 70 *and* 90 MHz *are on each side of the peak.*

Figure 3.10: *Left figure shows the lookup tables used for each of the two calibration methods, while **(b)** is the relation between RF power and DDS amplitude register for different frequencies in the calibrated range.*

lower outgoing intensities at the borders of the range, as can be seen in Figure 3.9a. It is not possible to arbitrarily increase the amplitude with the calibration LUT, and the only means for correction is to decrease the amplitude of the peaks. Due to the shape of the uncalibrated intensity response, this means that a broader calibrated frequency range will cause a lower total amplitude of the system. To put it simply, the calibration is cutting of the peak of a mountain to create a plateau, and the broader the plateau needed, the lower we have to cut. This can to some extent be counteracted by using stronger amplifiers after the DDS, although this could also lead to a noisier signal, which might not be desirable. In our setup, it seems that the loss in amplitude caused by the calibration is not large enough to pose any further issues in the system.

We try to calibrate the intensity in with two different approaches, where the first is based purely on the measured intensity as a function of driving frequency, and the second approach takes into account that the output intensity as a function of DDS output amplitude is not strictly linear. Thus, we expect that the first approach will not be perfect, as e.g. a halving of the DDS amplitude setting would not necessarily lead to half the intensity value of the beam. Specifically, the relation between intensity and DDS amplitude setting is as shown in Figure 3.10b, which is more of an S-shape. The figure shows the relation measured for 70, 80, and 90 MHz, which corresponds to roughly the frequency of the intensity peak (in Figure 3.9), and 10MHz above and below. This clearly shows that the relation is not constant throughout the range, and a perfect calibration would thus require a measurement for this relation for every single frequency in the desired range. This could in principle be done, but is not feasible if the system is regularly calibrated (note that even slight drifts in the laser or alignment of elements could alter the calibration). Thus, we instead decide to do an approximate calibration where it is assumed that the intensity against DDS amplitude relation is constant throughout the range, specifically with the relation measured near the peak value at 80 MHz.

Using the setup described above, the intensity response before calibration is measured for the frequency range from 0 MHz to 250 MHz, and the recorded data is used to create a calibration LUT. First, the intensity relations are measured simply by sweeping the amplitude setting on the DDS, while keeping the output frequency constant. With the measurements done, we can generate the calibration LUTs. First, we generate the simplest calibration, where it is naively assumed that the relation between amplitude and intensity is linear. This leads to the naively predicted intensity vs. frequency measurement in Figure 3.9a. The figure furthermore shows what the assumption of a linear intensity vs. amplitude response predicts for the other calibration scheme that takes the nonlinear relation into account. The plot clearly shows that the first calibration is expected to be perfect, while the second calibration leads to some more complicated shape, even though it is still an improvement. The lookup tables corresponding to these predictions are shown in Figure 3.10a.

After generating the calibration data, we use the same setup to sweep the frequency from 0 to 250 MHz, and measure the resulting intensity on the photodiode, leading to the traces in Figure 3.9. Note that the frequency range of the figure is cropped, as the remaining range to 250 MHz is flat. From the figure, it is clear that the first, simplistic, calibration is over-compensating for the peak, and results in a significant drop in intensity in the centre region. The second calibration takes the amplitude vs. intensity relation into account, and produces a significantly better calibration, although still with clear peaks and valleys in the range. As mentioned above, the calibration could be further improved by using amplitude vs. intensity curves from more frequencies, and in this manner the calibration could be improved as much as desired, until the entire phase-space is used to generate the calibrations.

This amplitude calibration setup is not yet implemented in the operating mode with interlacing beams, however this will be done, and requires only further development of the HDL code. The preliminary code has already been developed[1] for use with the measurement described above, and thus only needs to be implemented with the existing firmware. This, however, was not prioritized for this project.

3.5 Limitations

In this chapter we have established that the device is triggering accurately, within 4 ns, that it is possible to calibrate the amplitudes to produce flat intensity responses, and that interlaced frequencies might cause higher-order effects that could lead to issues. It is furthermore evident that there is a limit on the lower limit on the duration of individual segments, due to the transfer speed between the FPGA and the DDS boards. There are also some long delays from external trigger until the output is updated, due to the transfer times in each step before the output. The choice of DDS itself also creates a limitation in terms of the possible output signals, by limiting how fast the frequencies can be updated. The last main limit is the size of the memory on the FPGA, which puts a limit on the amount of output segments, and thus indirectly on the possible complexity of the output signals.

All of these delays and limits could in principle be improved in future iterations of the design. In particular, the lower limit on the segment duration could be decreased by using a DDS with a larger BUS instead of the current 4-bit setup, and one could use a DDS with a much faster update speed on the frequency output. Even though improvements are possible, they would all require either better hardware, potentially leading to higher cost, or more sophisticated logic. For instance, the transfer times might be lowered by engineering a very memory-efficient binary

[1] By Mohammad Noaman

structure, although likely at the cost of readability, and the memory limit of the FPGA could be solved by utilising the larger already on-board memory.

The design presented and tested above, is chosen as the trade-off between price and complexity, and the requirements of the lab. Even though improvements could be implemented, the current designed is by far fulfilling the requirements of the current experimental setup, and it is furthermore expected that it will remain so in the future. Thus, we acknowledge that there are possible, and to some extent also obvious improvements, but we are confident that for our use-case the improvements would not add any benefit, and they are thus not implemented here. The setup, however, is general enough that individual improvements could be implemented as needed in the future, without disturbing the combined setup. That is, in particular the DDS board could be changed, without affecting the higher-level Python code and control setup, and conversely the entire control setup could be modified without modifying anything on the hardware level.

Chapter 4

Implementing in the experiment setup

The DDS based frequency source can of course be used to modulate the laser frequency to implement the SWAP cooling scheme discussed in Section 1.2.4, as it was build for, but that is not the only use case. The setup is very versatile, and also allows for sweeping the frequencies with more complicated shapes (by approximating linear segments), and even allows for effective multi-tone signals by interlacing multiple frequencies. That is, the RF can jump between different frequencies, effectively creating a signal with two, or more, separate frequency components.

The time-averaged effective multi-tones are in particular is useful in combination with optical traps, as it allows the creation of structured traps arrays. Simple structured arrays have already proven useful to e.g. observe BEC interference in a double-well setup[90]. In general optical traps are very flexible, and in particular relevant for tweezer experiments, where the ability to control and move traps precisely could pace the way for new experimental platforms.

4.1 Moving dipole traps

If a sweeping and interlaced RF signal is used to drive an AOM/AOD, the beam will move around, and jump between multiple positions. This by itself might not sound desirable, but by implementing this beam movement on a dipole trap laser, we can create movable dipole traps, as described in Section 1.2.5. This is useful as it allows for precise control of the position of the atomic cloud, and the relative position of multiple clouds. Furthermore, by starting with a single dipole trap, that is then changed to two interlaced dipole traps moving away from each other, one would be able to split or merge the atomic clouds as desired[91].

4.1.1 Experimental setup for testing

To test the setup, and confirm that the beam is indeed behaving as expected, the DDS based RF source is used to drive an AOM[1], and the outgoing light is captured on a Lumenera Lt365RM CCD camera. The complete setup is shown in Figure 4.1, and starts with a (399nm) laser source that goes through a waveplate and a polarising beam splitter, before it enters the AOM. The AOM is driven by the DDS source, and the outgoing light from the AOM is focused on to the CCD chip through a $f = 400$mm lens. In order to filter out ambient light, a long lens tube, and neutral density filters (ND= 5.0) were added directly in front of the camera sensor.

[1] Gooch and Housego 3200-129

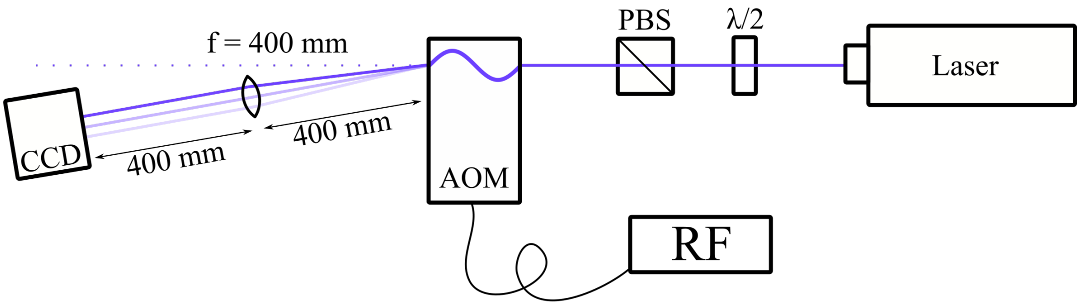

Figure 4.1: *Illustration of the experimental setup used for imaging the interlaced beams on a CCD. The blue (399nm) laser beam goes through a waveplate an a polarising beamsplitter (PBS), before going through the AOM. After the AOM, the negative first order is collected and focused/collimated in a* $f = 400$*mm lens, before the CCD camera is used to image the beam.*

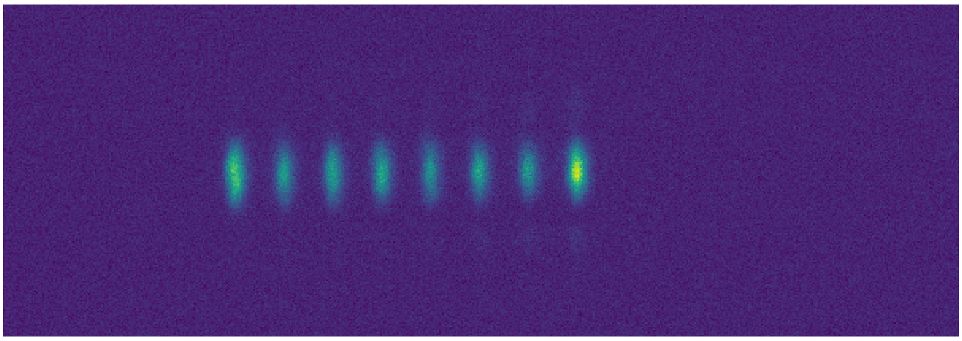

Figure 4.2: *A raw image of* 8 *interlaced beams on the image sensor. The image is exposed for* $20\mu s$*, during which the AOM (RF signal) addresses each beam for* $2\mu s$*, and then addresses two beams a second time. This is the reason that the left- and right-most beams are higher intensity than the rest.*

The CCD camera is based on the Sony ICX674 monochrome image sensor, which has a pixel size of $4.54 \times 4.54\mu$m, and resolution of 1936×1456 pixels. The image width is thus $4.54\mu\text{m} \cdot 1936 = 8.79\text{mm}$. In order to achieve the highest possible spatial resolution of the measurement, it is desirable that the moving beams span the entire camera sensor. As the beams are diverging after the AOM, this is in principle easily achieved by adjusting the distance from the AOM to the camera sensor. However, while the beams are diverging from each other, each of the individual beams are in principle collimated. Although the individual beams are already collimated, the setup would be improved by focusing the beams down. This is achieved by a lens in the path, that will not only focus each individual beam down, but will also change the relative angle of the beams, such that they become parallel. That is, if the lens is placed with the AOM in the focal point. The longest focal-length lens available in the lab was $f = 400$mm, and this was thus used for the measurements. With this lens, the distance to the AOM is not long enough to fully span the image sensors with the beams, but it is determined that the gain by focusing the individual beams make it up for this. In the end, the beams spans about 850 pixels, which is a bit less than half the sensor, and the width of each individual beam is 30 pixels, or $30 \cdot 4.54 = 136.2\mu$m.

The AOM we use has a resonance frequency at 200 MHz, but can operate in a large range around this. Operating it off resonance will, however, lead to less than peak efficiency, and the transmission intensity is expected to decrease as the frequency is moving away from resonance. This becomes especially obvious if the AOM is operated in a large frequency range, in which case the intensity would create a curve resembling a bell-shape, as was seen in Figure 3.9. This essentially shows the difference between an AOM and an AOD (acousto-optical deflector): The AOM has a resonance circuit at the input, which enables efficient diffraction at the resonance frequency, while the AOD has no such circuit, and thus deflect less at the peak, but with the gain of larger frequency ranges without large intensity changes. Thus, for purposes where the main focus is the beam deflection, an AOD would provide the better option. However, at the time of measurement, only an AOM was available, and should still provide for an ample setup for testing.

For the test, a signal that shows the capabilities and is still realistic in an experiment setting is used. The signal starts a static frequency of 180MHz, then sweeps some 10MHz, before the splitting starts by interlacing the signals. After the splitting the beam 3 times, leading to 8 parallel beams, they start to merge again, and the signal finish with one static beam. The total signal duration is 4.8ms, with a frequency span of 140MHz to 220MHz. The signal is shown in Figure 4.3. We thus have at most 8 beams, each with a width of 30 pixels, spanning 850 pixels in total, and as such each the beams are easily distinguishable on the sensor, and positional changes of $\sim 10\mu$m should be detected by the setup. An image of the fully-split signal with 8 interlaced beams on the camera sensor is shown in Figure 4.2.

The problem with this setup is that the lowest feasible exposure time of the camera is circa $100\,\mu$s, while the RF signal is switching through all interlaced beams with a period of $16\,\mu$s. Thus, the long exposure time of the camera would merge many cycles together in a single image, and therefore not lead to a satisfactory resolution in time. Furthermore, with a signal duration of 4.8 ms, and an exposure time of $100\,\mu$s, the entire signal would only produce 48 images. In order to achieve a higher resolution in time, despite the high exposure time, we take advantage of the fast TTL switch that is implemented in the DDS source. The idea is to only send light on to the camera sensor in the time window we want to image, by only turning on the TTL in a short period. The scheme is then as follows: First, the DDS source is triggered, starting the RF output, although not sending it to the AOM due to the low TTL signal. Then, the camera is triggered $40\,\mu$s before the TTL is turned on, as the camera requires some preparation time before it starts exposing the image. The TTL is then turned on for $20\,\mu$s, during which the beam goes through the AOM, and onto the camera sensor. This whole scheme, which is also shown in Figure 4.4, is then repeated such that an image is produced for every $20\,\mu$s segment of the full 4.8 ms signal.

The signal is loaded to the DDS source once, through the computer control setup as described in the previous section, and all the triggers are provided by the ADwin which is controlled by the Experiment control software. This allows for controlling the trigger times in multiples of $20\,\mu$s, and ensures that all triggers are coming at the correct time, relative to each other. Furthermore, the control software provides an automated way of increasing the trigger delay with each measurement, such that the many repetitions are easily performed.

4.1.2 Measured beam deflections

The measurement resulted in 250 images, each representing a time period $20\,\mu$s, as some padding was added to the signal in order to confirm that the tails are also behaving as expected. Each image is then cropped to the region of interest in the plane of the moving beams, and then projected down to a 1D array by adding all values in each column. We then have the intensity

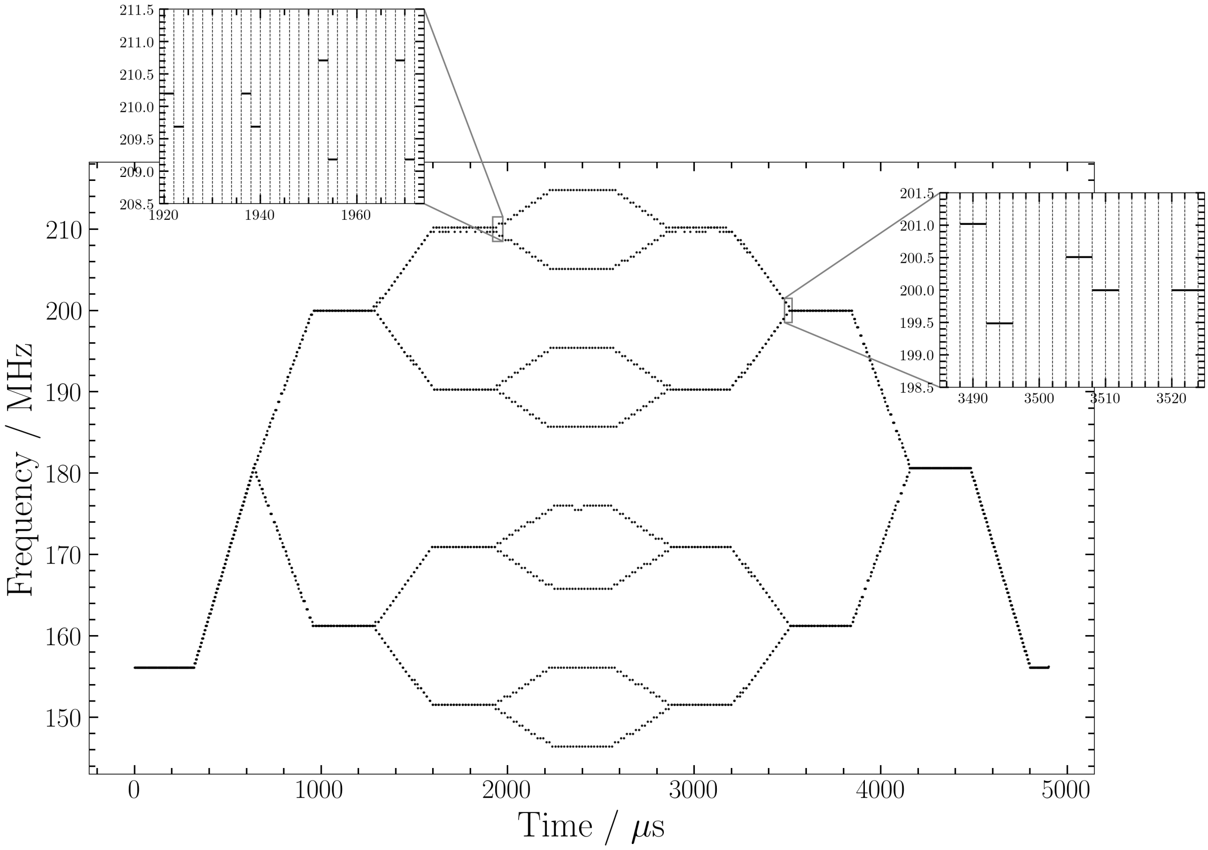

Figure 4.3: *Plot of the frequency generated for the moving beams by the DDS device as a function of time. The frequency is updated with a period of* $2\,\mu s$*, with a constant frequency in each segment. Each point in the plot marks the frequency of the segment outputted at any given time. The signal is starting with a single frequency at* $156\,MHz$*, which is increased to* $181\,MHz$*, before splitting into 2, 4, and finally 8 interlaced frequencies, and later merging together again. The upper left and right corners show zoomed-in plots of two places where the interlaced signals are splitting and merging. When splitting from 4 to 8 segments, each frequency is only addressed once every 8th segment, whereas when merging 4 to 2, each frequency is addressed 2 times for every 8 segments.*

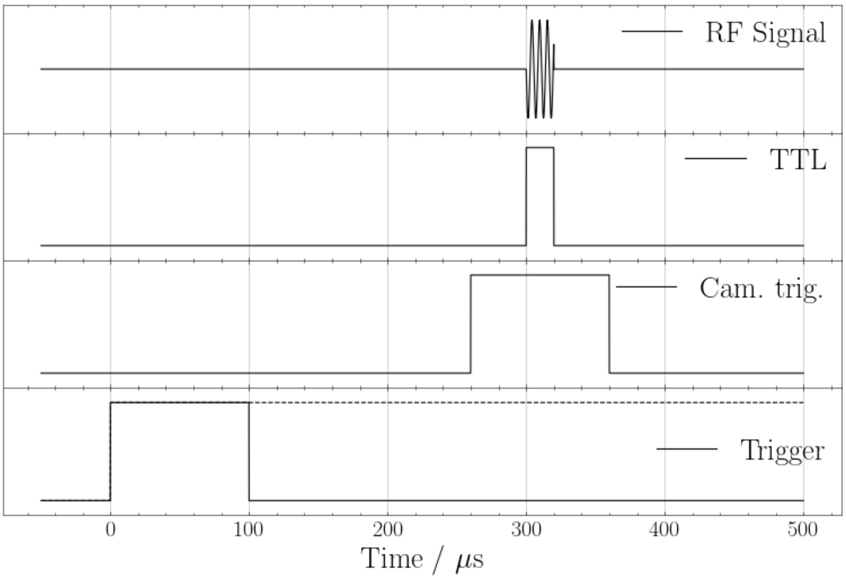

Figure 4.4: *Plot of the trigger and TTL signals used to image the moving beams from the AOM. The dashed line with the trigger signal marks the window where the sequence is running. At $t = 0$ the DDS device is triggered, and starts generating a signal. However, the TTL is off until the desired time, and there is therefore no output RF. The camera is triggered $40\,\mu s$ before the TTL is turned on for $20\,\mu s$, ensuring that the camera is indeed capturing while the RF signal is turned on by the TTL. This cycle is then repeated until the entire output signal has been recorded.*

distribution in the dimension of the moving beams (x-direction), allowing for further analysis.

As each image is exposed for 20μs, but the RF is going through all beams in $16\,\mu$s, there will be some overlap, such that not all beams are exposed equally in each image. This means that some beams will have significantly higher peaks than others. We want to be able to see the movement of the beams, ignoring the intensity fluctuations for now, and thus all peaks are normalised to the same height. Furthermore, the noise floor is determined and subtracted from the data. With this, we can stack all images column by column, creating a plot with position on one axis, and time through the stack of images. This is shown in Figure 4.5.

The path of the beams through time are clearly visible in the figure. In fact, the traces bears a clear resemblance of the raw RF signal in Figure 4.3. This is to be expected, to some extent, as the deflection angle through the AOM is mostly linear with the input frequency. There are a few dark spots in the traces, especially when splitting/merging to and from 8 beams. This is due to the detection of peaks, that is fails to detect both peaks when the beams are too close, thus failing to normalise their intensities. The dark spots are thus not directly present in the data, but are rather artefacts of the analysis and plotting. It also appears in the image that the level of noise is larger when there are more beams present. This is due to the fact that the intensity of each beam is lower when the power is split into multiple beams, but the analysis is shifting all peaks to the same intensity. In other words, the signal-to-noise ratio is lower with more beams, and the noise is amplified together with the peaks, leading to more noise.

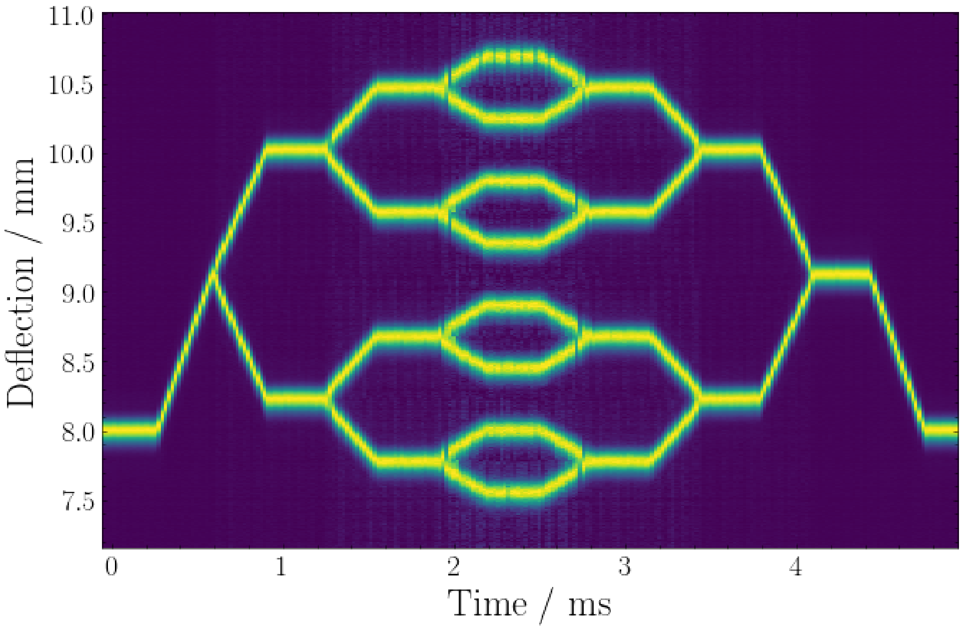

Figure 4.5: *Images of the interlaced beams on a CCD. The columns af each raw image are summed, such that the image is only 1D, and the 1D images are then stacked through time. The dark spots/lines present in the regions when splitting/merging 4 to 8 beams are artefacts from the analysis, and are thus not present in the experiment.*

4.1.3 Averaged intensity in the interlaced signal

In the previous section, the movement of the beams through the AOM was investigated, but another very interesting point is the intensities of the beams. Even though the images are showing multiple beams in one image, there is only a single beam present, which is rapidly changing the position. It is therefore expected that the intensity of each beam in the images with multiple beams, is lower than the intensity of a single beam. In particular, as the AOM is cycling through the beams at a constant rate, and thus addressing each beam for equal amounts of time, it is expected that the intensity is scaling inversely with the amount of beams, regarding the fact that some beams will be exposed more than others in the 20μs exposure time of the sensor.

Figure 4.6 shows the intensity distribution in 1D of an image with 8 beams, where the intensities are relative to the peak intensity of the single beam from the first $0.4\,\mu$s of the measurement. It is clearly seen that the first and last peaks are slightly higher than the rest, as they are exposed twice as long as the others. The average intensity of each beam is expected to be 1/10th of the single-beam intensity, as each beam is addressed for $2\,\mu$s, and the camera is exposed for $20\,\mu$s. The dashed line in the plot is marking the average peak height at 0.095, where the left and right most peaks are divided by 2 before averaging. The sum of the peak heights on the other hand is 1.02, which is only slightly higher than unity. Thus, the average is slightly lower than expected, while the sum is slightly higher, although within the precision of this very simple setup, these deviations are to be expected. Especially as the intensity of the beam from the AOM will be depending on angle, where the response is highest near the AOM center frequency at 200 MHz. Furthermore, the double exposure of the outermost beams will not necessarily result in a perfectly doubled peak height, and simply dividing by 2 could therefore lead to an under-estimated single-exposure intensity.

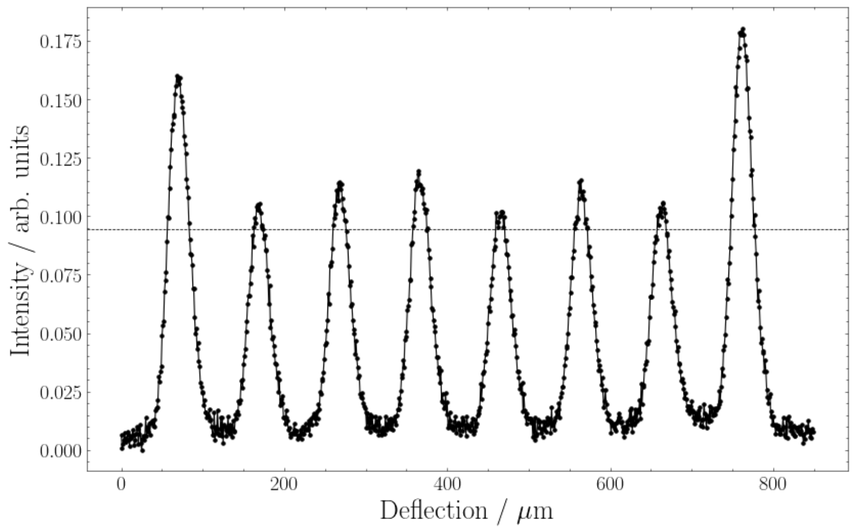

Figure 4.6: *1D intensity distribution of an image with 8 interlaced beams. The y-axis is the relative intensity compared to the peak intensity of a single beam on the image sensor. The outer peaks are significantly higher than the others, as they are exposed for twice as long, due to the nature of the measurement setup.*

All in all, it is clear that the measured intensity of the interlaced beams are scaling with the amount of time for which each beam is addressed. This behaviour also provides a good analogy of the time-averaged trapping potential such a setup will create when interlacing dipole trap beams. That is, if we imagine the atoms perceiving the laser in the same way as the camera sensor, they don't see the flickering of each individual beam, but rather see 8 stable beams of an accordingly lower intensity. This is of course only true if the interlacing is happening at a rate faster than the atoms can react. In more precise terms, the interlacing frequency should be faster than the trapping frequency of the atom.

With that, we conclude that the developed DDS source in indeed useful for implementing moving dipole traps, and even further use the beam movements to split and merge atomic clouds, and create 1-dimensional arrays of cold atomic clouds. The system is yet to be implemented in the experiment with actual atoms, but the proof-of-concept is promising, and shows that the system allows for precise control of the individual beams, and that the interlaced RF signals can be used to create arrays of time-averaged traps, based on a single trapping beam. For implementation in an experiment, the AOM should be replaced by an AOD to provide a larger range of angular displacement, but apart from this, the setup as described to this point is ready for direct implementation in the experiment.

4.2 Implementing SWAP cooling in a Green 3D MOT

One purpose of the DDS RF source is the moving dipole traps as described in the previous section, but the main reason to develop this setup is the implementation of SWAP cooling of ytterbium atoms in a green 3D MOT. The purpose of this section is to describe how well the setup is handling this use case, and how we tested the SWAP cooling in the green intercombination MOT. This was the first, qualitative, test of SWAP cooling before the experiment was disassembled for moving. Since then, the experiment have been under reconstruction, as will be discussed in Chapter 5. A more thorough analysis will be carried out once the experiment is reconstructed. That said, the preliminary testing acts as a proof of concept, and confirms the potentially large gain in MOT size and loading rates this SWAP cooling setup can offer.

The experimental setup, and in particular the vacuum system and the 2D/3D MOT scheme is described in [92]. In short, the Yb atoms are provided by dispensers in one chamber, and are loaded directly into a blue 2D MOT. From here, the atoms are pushed through a differential pumping tube, and into another chamber where they are loaded into a blue 3D MOT. While the atoms are in the blue 3D MOT, the significantly smaller green (ie. narrow line) MOT is turned on, and thus loading atoms from the blue MOT. After a initial loading for about 1 second, the blue MOT is turned off, and the green MOT is then ready for measurement. Here, we image the green MOT with absorption imaging, and use time-of-flight measurement to estimate the temperature of the MOT. Furthermore, the very initial measurement is a qualitative measurement showing the effect of different frequency modulations of the trapping beams. For all measurements, the trapping beam of the green 3D MOT is passing through an AOM to allow for control of the laser frequency. Note that the first negative order of the AOM is used, and a higher radio frequency is therefore corresponding to a lower laser frequency.

First, we want to investigate how much the shape of the frequency modulation affects the size of the MOT. For now, we only perform a qualitative measurement of the MOT by imaging the MOT fluorescence. The setup is as described above, where the frequency modulation of the 3D MOT beams is controlled by an AOM driven with the DDS setup. The baseline is provided by driving the AOM with a static frequency of 78.3MHz, which leads to the MOT in the top right of Figure 4.7. The MOT clouds from different SWAP frequency ramps are all shown in Figure 4.7. With a baseline for comparison, the next step is to implement the SWAP cooling. The first and simplest idea is to modulate the frequency with a sawtooth wave, meaning that the rate of change of frequency is the same when increasing and decreasing. We modulate the frequency between 74 MHz and 90 MHz with a period of 1 μs. This leads to the bottom left MOT in the figure, which is clearly brighter (and thus larger) than the static frequency MOT. This is already a clear improvement, but from the theory we know that ramp of increasing frequency will cool, while decreasing ramps are essentially forcing the atoms away from the trap. Thus, the next step is to remove the decreasing frequency ramp, by skewing the sawtooth such that frequency jump directly to the initial frequency after each ramp. Note that an increasing frequency ramp on the laser corresponds to a negative ramp on the RF, due to the used AOM order. The period of the sweep is still kept at 1μs, and the range is also the same. This setup creates the MOT shown in the bottom right, which again is significantly brighter than the MOT created with sawtooth modulation.

The last modulation we test is a signal which only ramps down in frequency, which is essentially opposite to the previous measurement. This should, from the theory, not improve cooling at all, but rather optimize heating, ie. pushing the atoms away from the trap centre. The modulation is a mirror of the previous one, meaning a period of 1μs with a modulation from 74 MHz to 90 MHz. The resulting MOT, or rather, lack thereof, is shown in the top left panel of Figure 4.7, where no MOT is to be seen. Thus, we have confirmed that ramps of decreasing laser frequency will indeed remove the MOT entirely, due to the heating effect.

A comparison of the fluorescence measurement is provided in Figure 4.8a, where the total fluorescence for each of the SWAP ramps are shown. Figure 4.8b shows the estimated size of the clouds. It is clear from the figures that the SWAP cooled MOT with a decreasing frequency ramp (which corresponds to an increasing laser frequency) is the brightest and smallest cloud. The normal non-SWAP cooled MOT provides that largest of the clouds, which is as expected. However, the normal MOT is brighter than the symmetric sawtooth SWAP cooled MOT. This would not obviously be expected, but as the measurement is only quantitative, the behaviour

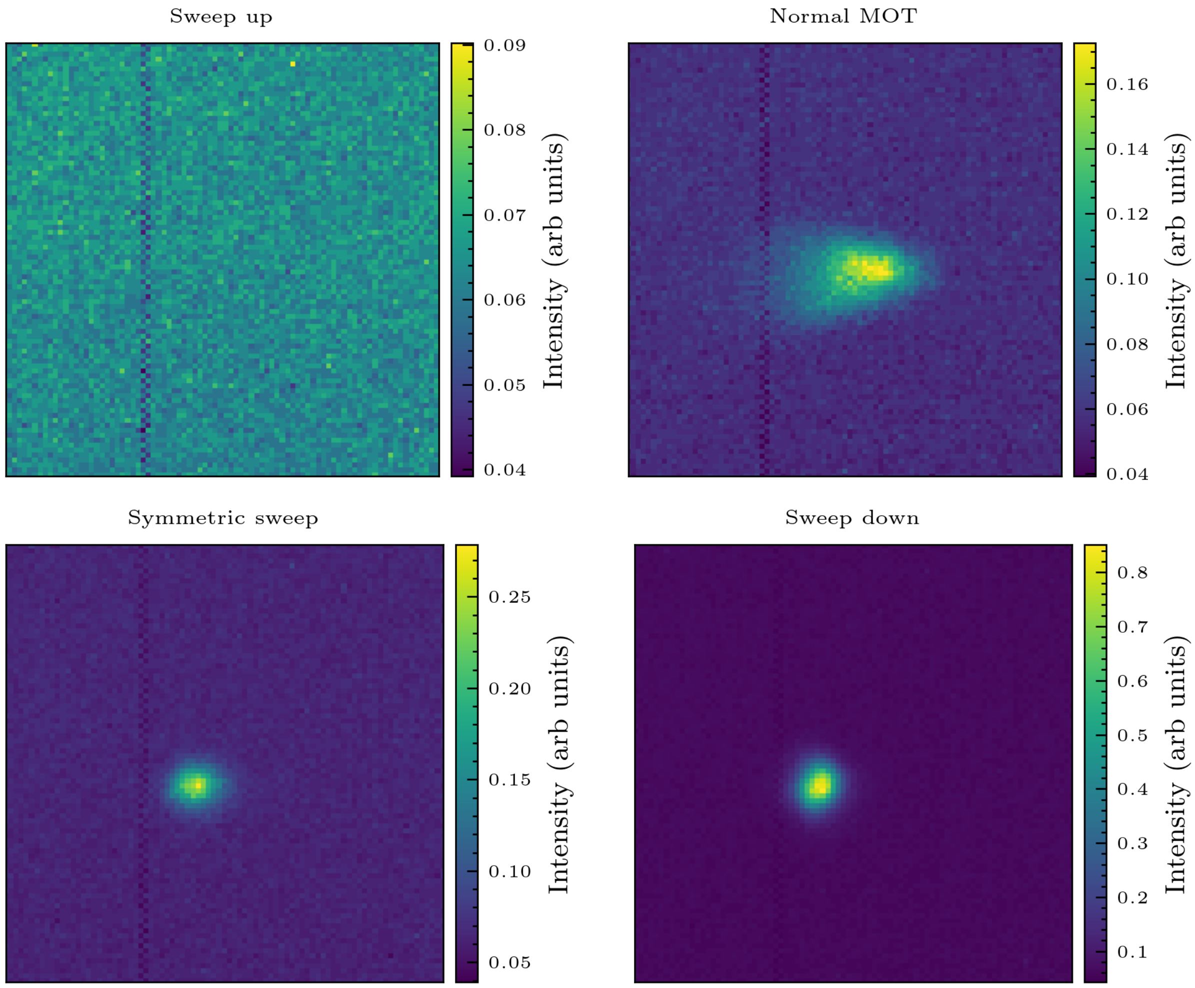

Figure 4.7: *Fluorescence images of the MOT after different SWAP ramps. The images are only qualitative, and therefore lack a length-scale. Top left: The MOT achieved by sweeping down the frequency, which as expected is empty. Top right is a MOT with a static frequency, and thus no SWAP cooling. The MOT is clearly visible, and very asymmetric. Bottom left image is with a symmetric sweep, which leads to a smaller, more symmetric cloud, with a higher peak intensity. Bottom right is expected to be the best, with the optimal ramp sweeping down. This creates a significantly brighter MOT than the others. Note that the colour-range in all images are the same, but the axis are scaled.*

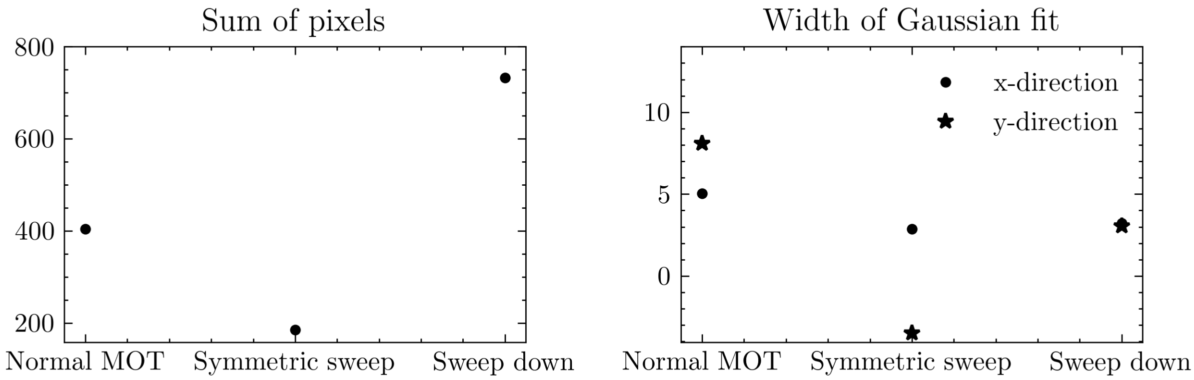

(a) *Plot of the total intensity in each of the fluorescence images of the clouds in SWAP MOTs with different frequency ramps. The intensities are in arbitrary units. The corresponding images of the clouds are shown in Figure 4.7.*

(b) *Size of the MOT clouds for different SWAP cooling frequency ramps. The normal MOT provides a baseline, and both SWAP MOTs are smaller than the non-SWAP case. It is not clear what caused the asymmetry of the MOT with a symmetric sweep. The size is in arbitrary units.*

Figure 4.8: *Total intensity **(a)** and cloud size **(b)** of the MOT clouds with different SWAP frequency ramps. The SWAP cooled MOTs tend to be smaller than the baseline, while, the tendency on intensity is not as clear.*

is not necessarily entirely representative, and further investigations are needed before anything can be concluded from this. That said, the data clearly shows that the SWAP cooling leads to a smaller cloud, and to some extend also a brighter cloud. Furthermore, is it shown that the best of the three cases is the down-ramping SWAP cooling, and that an up-ramping SWAP setup leads to an empty MOT, which confirms the expectations from theory.

At this point, we have established that the optimal frequency modulation is, as expected, increasing ramps where any decreasing of frequency is done in an instant. Strictly speaking, the optimal shape might not be a purely linear frequency ramp, although this analysis is something to explore when the experimental setup is reconstructed.

4.2.1 Temperature measurements of a SWAP cooled cloud

Using the linear frequency ramps, we carry out a further analysis of the achievable MOT temperatures by implementing the SWAP cooling scheme. This is done with absorption imaging of the MOT based on the setup described in the beginning of this section.

We image the cloud after a time of flight of $0, 1, 2$ and $3\,\text{ms}$, and estimate the size of the cloud in each image. The temperature is the determined by fitting the data to[93]

$$\sigma^2(t) = \sigma^2(0) + \frac{k_B T}{M} t^2, \tag{4.1}$$

where $\sigma(t)$ is the rms size of the cloud at time t after releasing the MOT, k_B is the Boltzmann constant, T is the temperature of the cloud, and M is the atomic mass. Thus, a linear fit of σ^2 against t^2 gives directly the temperature of the cloud. Note that the rms size is only measured in one dimension at a time, and this relation is thus used to estimate the temperature (ie. thermal motion) in each direction of the MOT. This measurement is done for two different ramp rates: First, we sweep the AOM frequency from $90\,\text{MHz}$ to $74\,\text{MHz}$ in $4\,\mu\text{s}$, and secondly we sweep the same range in a duration of $2\,\mu\text{s}$. The resulting linear fit is shown in Figure 4.9, where the $4\,\mu\text{s}$

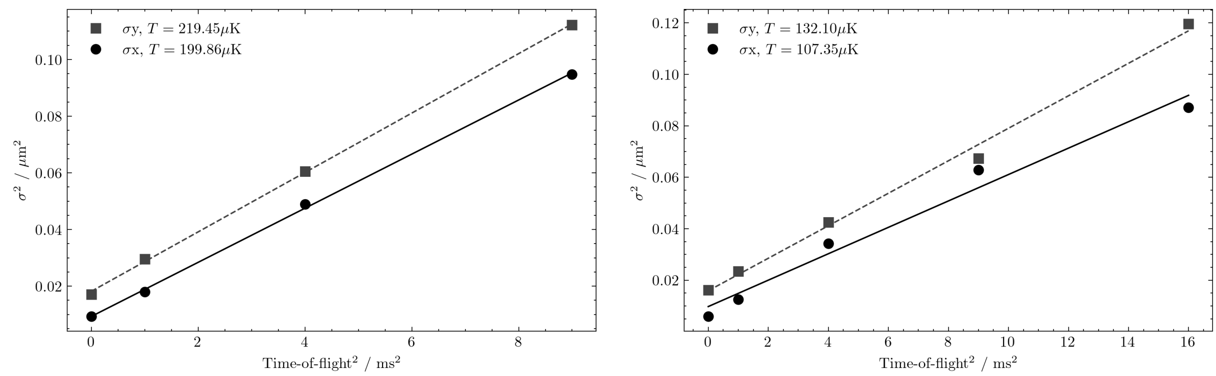

(a) *Estimated cloud size against time of flight for the* $2\,\mu s$ *ramp period.*

(b) *Estimated cloud size against time of flight for the* $4\,\mu s$ *ramp period.*

Figure 4.9: *Temperature estimation based on time of flight of the green MOT with different ramp rates of the laser frequency.*

leads to a temperature of 200 to 220 μK in the x and y-direction, respectively, while the faster sweep rate of 2 μs creates a MOT with a temperature of 107 to 132 μK in the x and y-direction. Thus, the temperature is nearly halved by doubling the ramp rate. This result is somewhat as expected, as a faster ramp will lead to a shorter cooling cycle (that is, stimulated absorption and emission cycle), allowing better cooling. This, however, is only true to some extent, as the atoms should have time to absorb the photon before the co-propagating cooling beam becomes resonant with the excited atoms. The temperature can therefore not be improved arbitrarily simply by increasing the ramping speed.

This shows that the temperature of the MOT depends directly on the ramping rates of the frequency, and this parameter should thus be optimised in order to obtain the best possible MOT. This optimisation is not carried out here, due to issues with the experiment, however it would be very interesting to determine the best ramp rates, and furthermore investigate how other types of ramps, e.g. Sigmoid-type etc., affect the loading rate and temperature limit of the MOT. From other similar experiments with different atomic species, it is expected that with proper optimisation of the SWAP parameters, MOT temperatures of $\sim 20\,\mu$K should be achievable[44, 72, 94]. Furthermore, it is expected that changing the frequency range and ramp rate as the MOT is loading might also provide fruitful to improve the MOT. That is, essentially narrow the frequency range as the temperatures of the atoms are decreasing, in order to remain resonant with the atoms, even as their velocity are changing significantly.

Chapter 5

Rebuilding the experiment in a new lab

The main goal of this thesis was the design and development of a versatile DDS based RF source for the implementation of SWAP cooling and arrays of optical tweezers. This part has been discussed in the previous chapter, where a device has been designed and assembled, characterized, and the preliminary steps to implementation in the experiment are discussed in Chapter 4. After the first steps, the experiment was entirely disassembled for moving, leaving further analysis impossible until the experimental setup is once again reconstructed in a new lab.

This chapter comprises some of the steps needed for the construction of a brand new lab, and the rebuilding of the experiment. In particular, we develop and discuss a versatile monitoring solution for the new lab, that allows for easy monitoring of not only the climate in the lab space, but also devices used in the experiments. Furthermore, we discuss the modifications performed on the experimental vacuum chamber as part of the reconstruction, and in relation to this, discuss the precautions and preparations needed when working with ultra-high vacuum (UHV) chambers.

5.1 Monitoring lab equipment and climate

For long term stability of the experiment, for documentation purposes, and to follow if anything are suddenly changing, it is important to monitor the lab, and the experiment devices. In particular, the temperature, pressure, and humidity of the lab are monitored, and the pressures of the experiment chamber and laser cavity are also monitored. More devices could easily be added in the future, but the aforementioned devices are the most critical.

In order to monitor these different devices, a general monitoring setup is needed. In general, the data to monitor is already available through some interface. For the experiment chamber, the pressure can be read through a serial port on the gauge controller, and similarly for the laser cavity pressure and temperature controllers. The lab temperatures are available through an HTTP socket from the climate regulators, and the pressure and humidity is measured with sensors attached to the GPIO of a Raspberry Pi. Thus, a general monitoring setup must be able to read data from virtually any kind of source and communication method, and combine the data in a unified structure, that also allows for easy visualisation.

For communication with such a wide array of devices, we use a very modular approach. A Python application is developed, that can run on any Linux or Windows computer, and can read

Figure 5.1: *Screen capture of the Grafana dashboard used to plot data from the monitoring system. In this figure the data shown is the pressure of the experiment chambers, where the peaks are caused by activation of one of the ytterbium dispensers.*

the data from a number of connected devices. The data from each device is transferred to an SQL database on the Monitoring computer, from where it can later be accessed for visualisation. The beauty of the system is that the application can run on many different computers, each monitoring their own subset of devices, while all the data is stored in one common database in the end. The application on one computer will be initiated with a list of connected devices, and will then read the logging interval from the database. It will then start a dedicated for thread for each of the devices, such that they can each log data a independent intervals, while still running in a single application. The thread for a device is running a specific measurement script developed for that particular device, such that differences in e.g. communication or data structure can be specified in a single place. As each device has a custom script, it does not change anything for the main application whether a device is communicating with a serial connection or through the network. This allows for connecting virtually any device to the setup by developing a suitable custom script for the particular device. It also means that one can make changes to the monitoring scripts of one device, without affecting all the other devices.

In the current setup, a single computer is used to monitor the pressure gauges of the experiment chamber and laser cavity, as well as the temperature of the laser cavity. The lab pressure and humidity is monitored by a Raspberry Pi, that is directly connected to the sensor through the GPIO ports. Thus, the monitoring script is running on two different devices, but one could easily add or remove both devices and computers to monitor them, as is needed in the lab.

5.1.1 Visualising the data

As described above, all data from the monitoring setup is stored directly in an SQL database. The python application doesn't provide any means of neither visualising the data nor sending alerts if anything seems wrong. Instead, the visualisation and alerting is done through Grafana, which is a web-dashboard application. Grafana is running on the Monitoring computer, and serves the dashboards in a web interface. When Grafana is connected to the SQL database, the dashboards with plots, gauges, etc., can be configured directly in the web interface. An example of a simple Grafana dashboard is shown in Figure 5.1.

While this has not been configured yet in the current setup, Grafana also provide an easy way to send out alerts when the data exceeds specific thresholds, or in other ways behaves

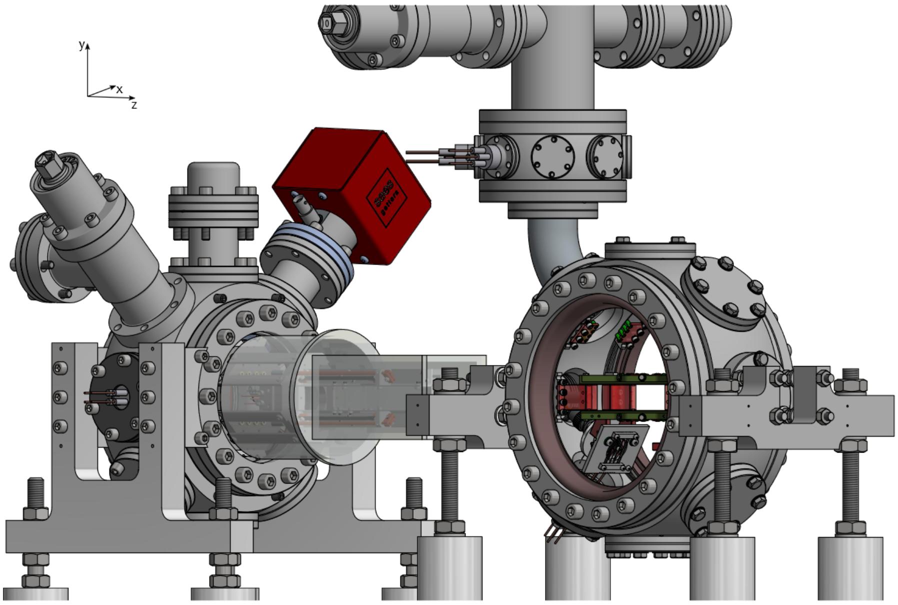

Figure 5.2: *CAD drawing of the experiment chambers. The 2D chamber on the left with the glass-cell used for the 2D MOT itself, with the main 3D chamber on the right with the large viewports. The chambers are connected with a differential pumping tube. The coils are not shown here, but the glass cell is surrounded with a cage of coils used for the 2D MOT, and the 3D MOT coils are placed above and below the 3D chamber. Smaller compensation coils are also present on all sides of the 3D chamber. Illustration is from [92].*

unexpectedly. This would be very useful e.g. in order to be alerted if the pressure in the experiment chamber suddenly starts to rise.

The setup is currently running, and working very well, but could still easily be expanded with more experiment devices. Particularly it would be useful to monitor the Uninterruptible Power Supply (UPS), and keep track of the battery lifetime, or log the flow rate of the cooling water used for the dipole trap laser and the MOT coils.

5.2 Assembling the vacuum chamber

In order to be able to use the proof-of-concept setup described in the previous section, a full experimental setup is needed, and in particular a vacuum chamber with the Yb source is needed for the realisation of cold atomic clouds. The specific design and creation of such a chamber will not be discussed here, as this work has already been done previously in [92], however during the reconstruction of the experiment, some vacuum work was needed, with the thorough cleaning and bake-out procedure that follows. A drawing of the chamber design is shown in Figure 5.2. Thus, this chapter will describe the cleaning process needed when preparing for vacuum, and the bake-out process afterwards, as well as the alignment procedure we went through to place the experimental chamber correctly on the table.

5.2.1 Preparing for vacuum

A large amount of care is needed when working with vacuum, as even small mistakes can render the entire chamber useless. Therefore, a lot of the time used for vacuum work is spent in preparation for working on the chamber itself. That is, the planning and cleaning phases are actually the longest. Every single tool, part, etc. that is needed for the vacuum work need to be ready in advance, and cleaned to go inside the chamber. It is also desirable to have the chamber open for as short a period as possible, and there is thus no time to fetch new parts and tools when the chamber is open.

The preparation therefore starts many weeks before the chamber is opened, making sure that all needed parts are ordered/received. This includes both any needed vacuum pieces, gaskets, tools, gloves, hairnets, aluminium foil, solvents for cleaning, and anything else needed for the process. Then, all work that is needed is carefully planned step-by-step, noting down exactly what tools and pieces are needed for each step. In case there are any doubts as to whether a step or procedure is actually possible, backup plans and corresponding steps are also discussed. There should be a thorough list of plan A, B, and maybe even C, with a complete list of the needed tools and equipment for each step of each plan. This list is then used to collect all tools and parts for cleaning, making sure there are spare tools and pieces (nuts and bolts etc.), in case any tools are dropped or otherwise contaminated during the work. All parts that needs cleaning are then cleaned with the procedure described below, before the actual vacuum work can begin.

Cleaning procedure

One of the most important parts of preparing for vacuum work, if not the single most important, is cleaning. All parts that go inside the vacuum, and all tools to be used, need to be thoroughly cleaned before use. This is because even small patches of dirt and grease can contaminate the chamber, and evaporate for a long time, preventing the chamber from reaching the desired pressure. Following the literature, and general suggestions, it turns out that there is no standard agreed-upon by all process of cleaning for vacuum. This is therefore no claim that the following procedure is the best one, nor the only correct way of cleaning, but rather an explanation of what was done for this project, and the reasoning behind. It is safe to say that this procedure is based on the philosophy "Better safe than sorry", and is by far thorough enough for the vacuum we aim for.

The basic idea of the cleaning is to remove any organic material, and in particular remove any remnants of oil and grease. In general, stainless steel parts are already vacuum clean from the manufacturer, although if the parts have been machined in a workshop, they are generally covered in oil. Different types of materials and parts will require different cleaning procedures, but for this work only cleaning of stainless steel parts was needed, and we therefore describe only this process. The cleaning procedure involves a number of steps, starting with rough cleaning with a cloth, and going through solvents in an ultrasonic bath. The step-by-step process we use is:

1. **Rough cleaning** As a very first, preliminary step, the pieces should be cleaned from any visible dirt, dust, grease, etc. When this step is done, the part should look clean to the eye, and preferably not leave any residue to the touch, or if wiped with a cloth.

2. **Degreasing with soap** This step is to remove any grease and oil from the surface of the parts, and is done by submerging the parts in a ultrasonic bath with distilled water and an industrial degreaser. In this particular case we use Sonoswiss Cleaner T1 and T2 depending on the metal (T1 for hard metals, T2 for soft metals), which was diluted to a 5% solution by volume. To avoid using large amounts of degreaser, and solvents in the later steps, the ultrasonic bath is filled with tap water, and a glass-beaker with the parts

to be cleaned, and the cleaning fluid, is submerged into the baht with tap water.
For this step, the parts stay in the ultrasonic bath for 1 hour at a temperature of 70°C.

3. **Rinse of soap** As the parts have been cleaned in soap, the next step is to remove any residue of soap. This is done by first rinsing the parts in distilled water, until most of the soap is gone. At this point, the parts are submerged in distilled water, and submerged in the ultrasonic part for 10 minutes. After the ultrasonic bath, the distilled water should still look clean, and if the water is 'soapy' after the bath, the rinsing and bathing process should be repeated. At this points the parts should be void of most oil and grease, and it should therefore be avoided to touch the pieces with bare hands. Thus, from this point on, the pieces are only handled with clean gloves, and are only placed on clean surfaces (eg. clean aluminium foil).

4. **Solvent 1: Acetone** In this step the pieces go into an ultrasonic bath with acetone for 1 hour. The acetone will dissolve a wide range of contaminants, but might leave a residue itself. It is therefore used as the first solvent, as it will remove many contaminants, but will itself become a contaminant. Note that if the same ultrasonic bath is used for both the soap cleaning and the solvents, it must be thoroughly cooled in between the steps, as the acetone, with a boiling point of 56°C, will boil at this temperature.

5. **Solvent 2: Isopropyl Alcohol** Directly after leaving the acetone bath, and while still wet from acetone, the pieces are rinsed with Isopropyl Alcohol (IPA) to get rid of the acetone. The reason this must be done while still wet, is that if the acetone evaporates first, there is a risk that the dissolved contaminants will stick back to the surface. When the acetone has been rinsed of with IPA, they go into the ultrasonic bath with IPA for 15 minutes. This is again to remove any remnants of acetone, and the contaminants the acetone dissolved. IPA evaporates slower than acetone, but will in general is also not dissolving as many contaminants[95], and is therefore mainly used to clean the acetone, which acts as the main device to dissolve contaminants. The alcohols are in general not dissolving as many substances, but will evaporate without residue. In principle other alcohols like ethanol and methanol could be used instead of IPA, although IPA is convenient as it is evaporating slower than the other alcohols, and thus leaves less vapour in the room.

6. **Dry and store safely** After the IPA bath, the pieces are in principle clean, but need to dry. This is done easiest by shaking of most of the IPA, ans spreading out the pieces on a clean piece of foil, leaving them for a moment to let the IPA evaporate. If the pieces need to dry very fast, one can blow of the IPA with nitrogen, but it is necessary to be very careful that the nitrogen source is clean, as there can easily be oil or dirt in the nozzle or pipes. The safest is to just let the pieces air dry. When the pieces are dry, they are ready for use in vacuum, and can be wrapped in clean aluminium foil until they are needed.

In general, when going between steps, the pieces are placed on a new piece of clean aluminium foil, until ready to go into the next bath.

There are other sources [96] that suggest starting with the solvents, followed by a soap bath, but we choose not to do this as many pieces coming from the mechanical workshop are covered in oil, in which case it makes more sense to start by degreasing with soap. Cleaning a piece in solvents when it is covered in a film of oil seems like a waste, and the oil might even 'protect' any smaller contaminants on the surface from being dissolved by the solvents. That said, solvents should also be able to clean away the layer of oil, in which case soap before solvents might make sense. The conclusion is, as was also mentioned in the beginning of this section, that different

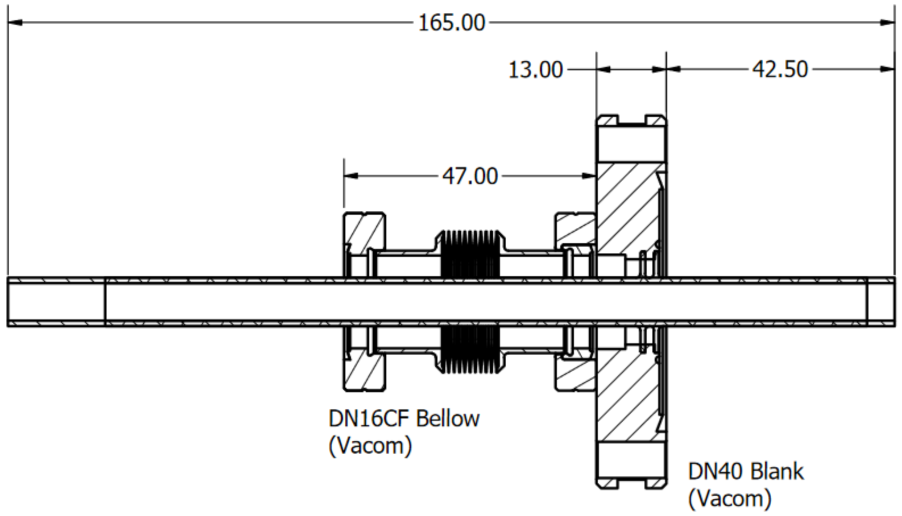

Figure 5.3: *Drawing of the differential pumping state, that connect the 3D chamber to the glass-cell, and in turn the 2D chamber. The differential tube is welded to a DN40 flange, on to which a DN16 bellow is attached. The DN40 flange is mounted to the 3D chamber, while the bellow connects to the glass-cell.*

groups have different procedures, all of which will most likely lead to pieces that are clean enough for vacuum.

5.2.2 Modifications inside the chamber

With the necessary cleaning and preparation procedure, we are ready to work on the vacuum chamber. As part of the reconstruction, there are two separate tasks in the vacuum chamber. Firstly, the differential pumping tube between the 2D and 3D chambers should be changed, and secondly, the ytterbium dispensers need replacement after being exposed to atmosphere during the disassembly and moving. Furthermore, the glass-cell was transported separately, and thus needs to be mounted back on the 2D chamber before the chambers chan be assembled again.

Replacing the differential pumping tube

The differential pumping tube is the interface between the 2D and 3D chambers, as shown in Figure 5.3. The tube is welded to a DN40 flange mounted on the 3D chamber, with a DN16 bellow on the outside, which connects to the end of the glass-cell on the 2D chamber. The bellow is a crucial part on the piece, at it provides a flexible connection between the two chambers, and importantly between the end of the sensitive glass-cell and the 3D chamber.

The tube is changed as it was determined that the old tube was unnecessarily wide, which leads to a lower possible pressure-difference between the two chambers. Thus, the new tube is narrower, and is expected to allow a larger difference in pressure. In particular, the goal of the differential pumping tube is to prevent the higher pressure in the 2D side during dispensing from causing a large pressure increase in the 3D chamber. The old tube is $L_{\text{old}} = 150\,\text{mm}$ long, and has an inside diameter of $d_{\text{old}} = 7\,\text{mm}$, while the new has a length of $L_{\text{new}} = 165\,\text{mm}$ and a diameter $d_{\text{new}} = 3\,\text{mm}$. For these parameters, the old and new pressure difference ratios are[92][97]

$$P_{\text{ratio,old}} \sim 300, \qquad P_{\text{ratio,new}} \sim 2500. \tag{5.1}$$

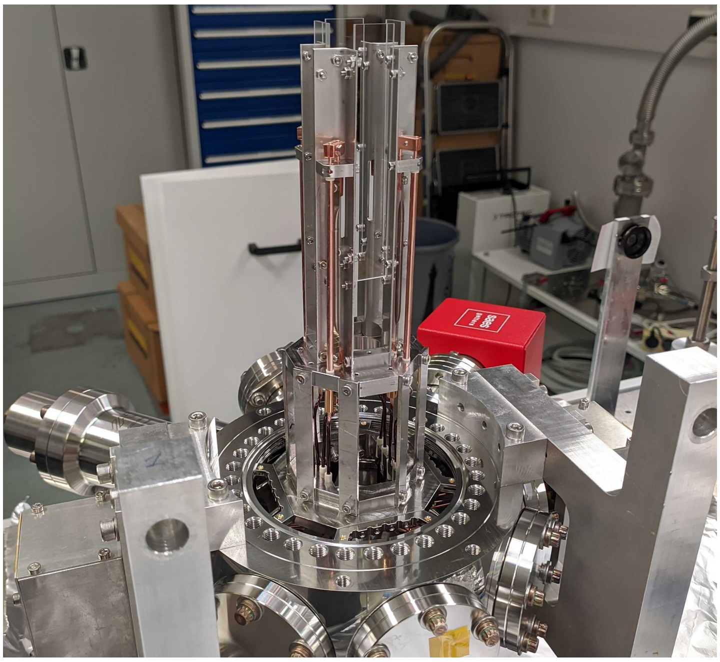

(a) *Image of the inside of the ytterbium dispenser mount on the 2D chamber, during replacement of the dispensers. The glass-cell is not present in the image, but would be mounted on the DN100 flange visible in the bottom of the image, and would cover the dispenser mounts.*

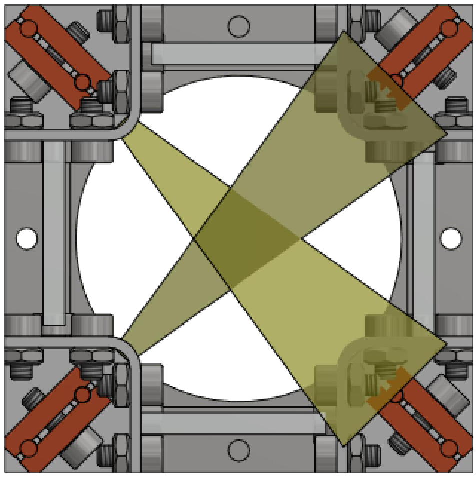

(b) *CAD drawing of the dispenser mount structure from the top, of the structure in **(a)**. The shaded triangles show the ytterbium beam from the dispensers. It shows that the alignment of the dispensers is important, as the beam should not reach the glass, as that could lead to deposition of ytterbium, and thus limited optical access.*

Figure 5.4: ***(a)*** *Image of the dispensers mounted on the 2D chamber, and* ***(b)*** *a CAD drawing of the dispenser mount, showing the expected ytterbium beam from the dispensers.*

The replacement of the tube is therefore expected to increase the maintainable pressure difference between the chambers by an order of magnitude.

Replacing the differential pumping stage is rather simple, while the chambers are separated. The DN40 flange is simply unmounted from the 3D chamber, and the new differential tube assembly is mounted. At all times being careful not to touch the tube itself, as it will go inside the chamber, and only handle the piece on the outer edge of the DN40 flange. The knife-edges of both the DN40 flange and the chamber itself are controlled before mounting, and then the DN40 flange is bolted to the 3D chamber with a new gasket in place. At this point, the next step is to connect the other end to the glass-cell, however, before this is possible, the glass-cell must be placed back on the 2D chamber.

Assembling the 2D chamber

The 2D chamber and the glass-cell are transported separately, with a DN100 nipple as a temporary cover on the 2D chamber. Thus, the steps for the 2D chamber are as follows: First, the DN100 nipple is removed, to expose the old dispensers. An image of the mounted dispensers is shown in Figure 5.4a. Then, the 4 dispensers need to be replaced, before the glass-cell is mounted on the chamber.

Each of the four dispensers are mounted on a copper plate in each end, with a nut and bolt. The cobber plates merge into a tube that goes through a hole on the structure, thus creating rotational freedom of the dispenser. The copper tubes furthermore connect to wires that in the end go to a feed-through in the chamber. To a replace a dispenser, the old one is unmounted by removing the bolts in both ends, and the new one is placed. The nuts and bolts are rather small, and all pieces are going inside the chamber, so extreme care is needed, and the work is helped along with a variety of clean pliers, and furthermore require frequent glove changes to ensure clean hands at all times. When the new dispenser is placed by tightening the bolt in each end, the dispenser is rotated until the dispenser-slit is in the pointing straight through the slit in the structure. A CAD drawing of the ytterbium beam from the dispensers is shown in Figure 5.4b. The alignment of the dispensers is important as the ytterbium should not be dispensed onto the structure, as that would be a waste, but at the same time it is important that the ytterbium cannot directly fly past the adjacent dispenser mount and onto the glass. This is particularly important due to the low vapour-pressure of ytterbium as discussed in Section 1.1.

After replacing all four dispensers, the glass-cell is mounted on the 2D chamber. The glass cell has a DN100 flange on one end, which merge into the glass piece, and ultimately back to a steel DN16 flange on the other end. The glass-cell is only slightly larger than the internal structure and dispenser mount on the 2D chamber, and is thus very carefully placed over the dispensers when mounted. The DN100 flange is the bolted to the 2D chamber, with a new gasket, and tightening the bolts evenly around the flange, ensuring even compression of the gasket.

With this, the 2D chamber is fully assembled, and ready to be mounted to the 3D chamber. However, the chambers are not easily moved after they are connected, and it should thus be ensured that the chambers are placed correctly on the table before they are connected.

5.2.3 Aligning a chamber on the experiment table

In principle, the chambers could be placed more or less carelessly on the table, without disturbing the experiment, but later placement and alignment of optics is significantly easier if the chambers are carefully placed in alignment with the optical table. Thus, a method for aligning the vacuum chambers is needed.

The alignment of the chambers is split into two separate steps. First, the placement and alignment of the 3D chamber, and secondly alignment of the 2D chamber, which is done when the chambers are connected. Thus, the first step is to align the 3D chamber. For convenience, the chamber is placed such that the MOT centre is above a specific hole in the table, and the z-axis going through both chambers is parallel to the length of the table. To help the alignment of the chamber, a beam-cross is placed in the MOT-centre, before the chamber is placed. The cross comprise 3 beams, placed far away to allow for plentiful working space, and is aligned using a combination of irises and carefully placed threads, as shown in Figure 5.5. With the beam-cross in place, the 3D chamber is placed such that the beams cross through the centre of the viewports in all directions. For the large viewports on the x-direction, we mount a plastic alignment piece on each viewport, which acts as an iris with a hole in the centre of the viewport. The pieces mounted on the chamber are shown in Figure 5.6. With these pieces, the placement on the table is fixed in one direction, and the height and tilt of the chamber is adjusted. Then, for the z-direction, the differential tube is used as an iris, and the chamber is placed such that the beam is not clipped by the tube. This locks both placement and rotation on the plane of the table, but furthermore locks the height and tilt of the chamber, due to the length of the tube. That is, the tube essentially acts as a pair of irises, with an iris at each end. There is

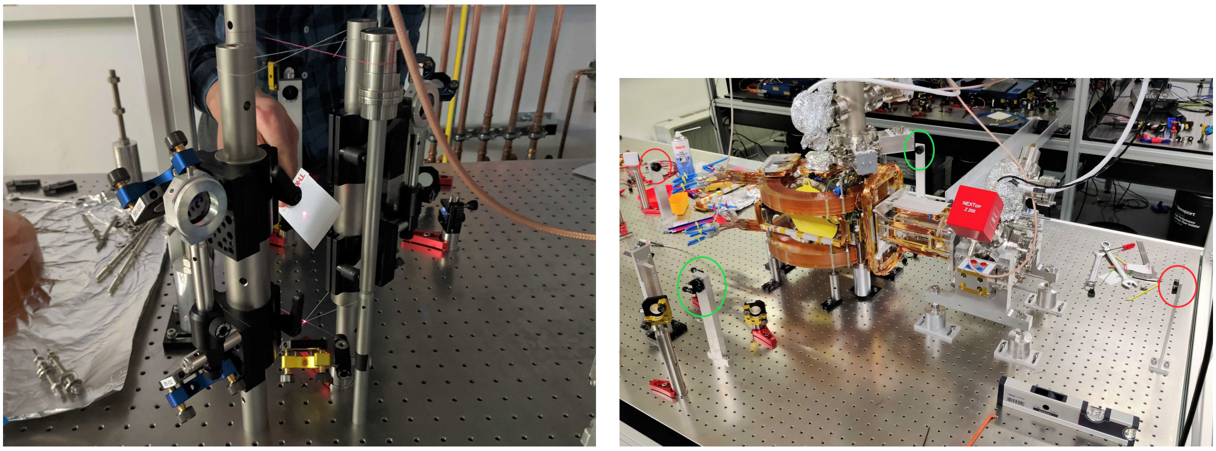

(a) *Image of the tower of optics mounts with crossing thread used for alignment of the beam along the y-directions in the beam-cross. All three beams of the cross are crossing in the red spot visible on the piece of paper.*

(b) *Image of the irises used for the alignment of the beams in the x- and z-directions in the beam-cross used for alignment of the 3D chamber. The z-axis irises are circled in red in the image, while the x-axis is marked with green. The image is taken after the chambers were aligned and connected.*

Figure 5.5: *Images of **(a)** the tower used to align the z-axis of the beam cross, and **(b)** the irises used for the alignment of the other axes.*

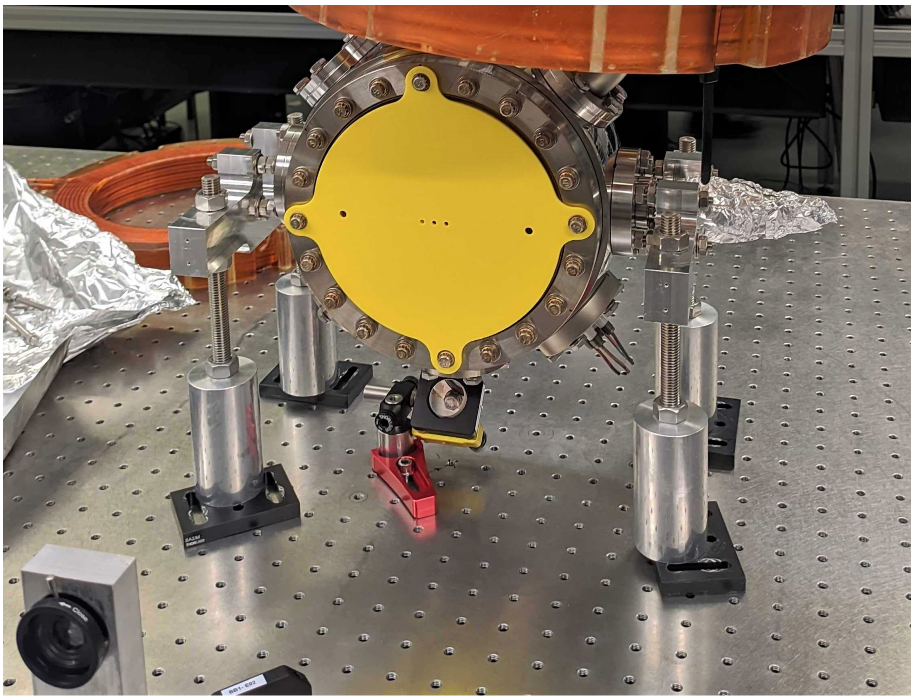

Figure 5.6: *Image of the 3D printed alignment pieces used to place the 3D chamber in the beam-cross. The plates are mounted on the bolts of the viewport, and provides a reference for the centre of the viewport in the form of a hole, thus acting like an iris.*

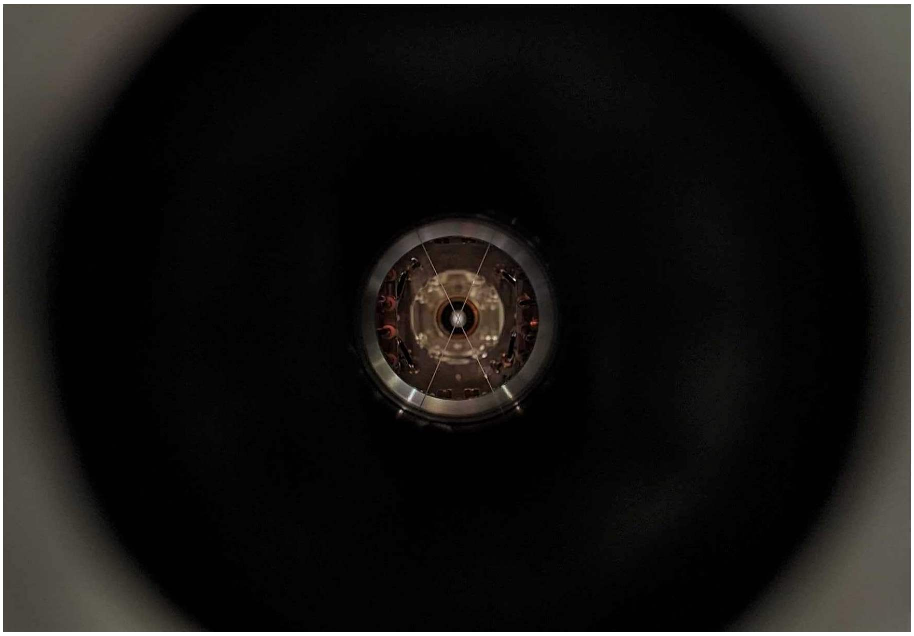

Figure 5.7: *Image of the differential tube taken from the rear viewport of the 2D chamber. This is used to align the 3D and 2D chambers to the z-axis, which is done visually by ensuring that the differential tube and irises all align, such that any clipping is symmetric.*

unfortunately no clear reference point on the chamber for the vertical beam, and the alignment is thus only confirmed by eye, and by creating a cross of threads between bolts on the viewports, to create an estimate of the centre-point of the viewport.

This provides the initial alignment of the chamber, but due to imperfections in the beam-cross, and due to the size of the laser-beam, the alignment is not perfect. To further improve the alignment, the cross-beams are turned off, and the chamber is visually aligned through the irises from the beam-cross, as well as the differential tube. An image of this approach in shown in Figure 5.7. The idea is to align through both irises, and the tube, ensuring that the clipping by the irises is symmetric around the tube.

After the visual alignment, the 2D chamber of connected to the DN16 bellow on the differential tube, and thus the 3D chamber. The 2D chamber is to some extent aligned to the z axis of the beam-cross, however only the rear viewport can provide a reference for the beam, and the entire glass-piece is placed by eye, ensuring that the differential tube is centred. This is done in a combination of looking through the glass-cell from the top and sides, and looking through the rear-viewport to align through the irises and the differential tube, as shown in Figure 5.7.

5.2.4 Pumping down to vacuum

After the work inside the chamber is finished, and the chamber is once again closed, it is time to pump down to vacuum. This is, similarly to all the other parts of vacuum work, done in a number of steps. In our case, the pumping is performed by three main devices: A roughing pump, a turbomolecular (turbo) pump, and NEG getter pumps. The roughing pump is connected to the exhaust of the turbo pump, and acts as the backing pump for the turbo, while the NEGs are placed inside the 2D and 3D chambers. The basic principle is that to use the roughing pump to lower the pressure enough for the turbo to turn on, and then use the turbo as the main pump until the pressure is low enough for the getters. Furthermore, as part of the pumping down procedure, it is also necessary to degas and activate the NEGs, gauge filaments, and Yb dispensers. The initial pump-down process we use is as follows.

1. **Roughing pump** The very first initial step is to turn on the roughing pump, and let the pressure decrease to $\sim$ 1mBar, at which point it is safe to turn on the turbo.

2. **Turbo pump** The turbo pump is turned on, and the pressure will drop very rapidly for a few minutes, until the turbo is at full speed, and the pressure is much lower. After 10 minutes of the turbo running, the pressure is $3.0 \cdot 10^{-5}$hPa, and after 2 hours the pressure has decreased further to $8.6 \cdot 10^{-6}$hPa. If at this point the pressure is decreasing at too slow of a rate, or the pressure is settling at a higher value than expected, the connection should be checked, and should be tightened further if needed, as there might be some leaks in the system. We let the turbo run for 18 hours, at which point the pressure has reached $3.0 \cdot 10^{-7}$hPa.

3. **Degassing and activation** At this point, the should only be slowly decreasing, if not already settled, and it is time to degas the ion gauge filaments. The degassing is necessary to remove any particles that have gotten stuck to the filament while exposed to the atmosphere. The degassing is done at this step such that the turbo can remove the degassed particles from the chamber entirely.

   There is an ion gauge in each of the 2 chambers, both of which have 2 filaments. Thus, there are 4 filaments that need degassing. There is a built-in degassing procedure in the ion gauge controller, which is used for this purpose, and takes 15 minutes each time. The 4 filaments are degassed one-by-one, with time for the pressure to settle between each filament. The degassing procedure increased the pressure to 1.0-$1.5 \cdot 10^{-5}$hPa in the turbo for each filament.

   After the degassing of the filaments, it is time to activate the NEG pumps. As with the gauges, there is an NEG in each of the chambers, that both need activation. The activation process itself takes 1 hour, but the NEG needs to heat up to 550°C before the activation starts, and the heating takes 30 to 60 minutes. Thus, the activation takes about 2 hours for each NEG. We start with the NEG in the 2D chamber, and during the activation the pressure peaked at $3.0 \cdot 10^{-5}$, after which it began to drop again. When the activation of the 2D side NEG finished, the pressure in the chamber is $1.2 \cdot 10^{-5}$hPa. Directly after this, the activation process for the NEG in the 3D chamber is started. The pressure increased to $7.4 \cdot 10^{-5}$hPa during activation.

4. **Turn on ion gauges** At this point, the pumps and filaments should be clean and ready to use, and it is thus safe to turn on the ion gauges. This allows for a much more accurate pressure measurement of the chambers, both because the gauges are placed inside the chamber, and because they are generally more precise than the gauge on the turbo pump. After the ion gauges are turned on, the pressures in the chambers are:

$$\text{2D chamber: } 1.0 \cdot 10^{-7}\,\text{mBar}, \text{ 3D chamber: } 3.2 \cdot 10^{-8}\,\text{mBar},$$
$$\text{turbo pump: } 5.4 \cdot 10^{-8}\,\text{hPa}.$$

At this point, the turbo should be left running for some days, after which the initial pumping down is finished. When the initial pumping is finished, the chamber should be tested for leaks. This is easily done by dropping some IPA on the joints, and monitor the pressure change. If there are any leaks, the pressure should increase rapidly as the IPA evaporates from the joint.

If there are any leaks, the joint should be tightened. If this doesn't help, there are two options: Accepting the leak, and the current pressure, or breaking vacuum to replace the gasket. In our setup, there were no obvious leaks, and the pressure after 4 days of pumping had reached:

$$\text{2D chamber: } 2.8 \cdot 10^{-8}\,\text{mBar},\ \text{3D chamber: } 3.6 \cdot 10^{-9}\,\text{mBar},$$
$$\text{turbo pump: } 1.4 \cdot 10^{-8}\,\text{hPa}.$$

If the pressure is settled at an acceptable level, the chamber is in principle ready for use after activation of the Yb dispensers. However, the pressure in the chamber can be further improved by baking of the entire chamber. That is, by heating up the entire chamber to increase the outgassing rate, and thus clean the chamber. In simple terms, the outgassing rate will increase with temperature, and thus the higher the temperature, the better the final vacuum will be.

In our setup, the BNC feed-through for the MCP inside the chamber cannot sustain temperatures higher than 100°C, which sets the upper limit for the bake-out temperature. Although, generally higher temperatures would be desirable. There are two main approaches to baking the chamber, either baking the whole chamber inside an oven, or by heating it with heating wires. As we do not have an oven that fits the chamber, and also have to bake on the table, our approach is to use heating tapes. One thing to be very careful with, is local heating with the heating wires, as any overlap between wires can cause points of very high heating.

The main concern when baking the chamber is to heat it uniformly, and slowly, to avoid any heat gradients across the chambers, and especially the glass cell. Large gradients could cause strain in the pieces that might lead to leaks, or ultimately cracks in the pieces. In order to achieve uniform heating, it is desirable to be able to control the heating very locally. The heating is therefore done with a number of shorter heating wires, that can be individually controlled, where each wire heats a small part of the chamber. The ability to control the temperature is important to avoid gradients, but being able to follow the temperatures is equally crucial, and we therefore place 15 temperature probes evenly spread out on the chamber.

Before the heating wires can be placed on the chamber, all viewports, and especially the glass cell is covered in a layer of aluminium foil for protection. The foil cannot rest directly on the glass, so for the glass cell a cage is built to hold the foil, and the main windows on the 3D chamber are covered with a piece of aluminium plate. On the smaller viewports, the foil is carefully placed on the bolts of the viewport, ensuring that the foil does not touch the windows. After the initial protective covering, the entirety of the chamber is covered in multiple layers of foil, before the heating wires are added. This is done in order to provide a layer where the heat from the wires can spread out, before reaching the chamber, hopefully causing more even heating. In this step, it is important to fill out any air gaps caused by e.g. flanges and cross pieces, as air-gaps can cause inhomogeneous heating, and cold/warm spots. The heating wires are then placed such that they cover the entire chamber, being careful never to cross the wires, and ensuring that each wire is touching the chamber throughout its length. On top of the heating wires, multiple layers of aluminium foil are added, in an effort to insulate, and even out the heating. At this point, the chamber itself is no longer visible, and what is left is a large ball of aluminium with wires coming out for the temperature probes, heating wires, and connections to the dispensers, that should be activated during the bake-out.

The bake-out procedure started at room temperature (22°C), and the heating wires are all turned on. Each wire is powered through a current regulator, such that we can control the heating rate. Initially all heating wires a turned on at the same low current, while the temperatures on the chamber are carefully monitored. Minor adjustments on the current of individual heating wires are performed as needed, to compensate for any developing gradients, while the temperate is increasing. The current control is manual, and the temperature increase

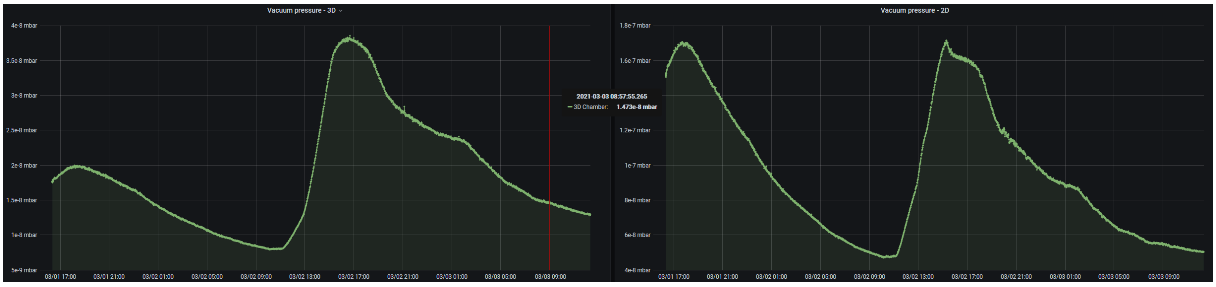

Figure 5.8: *Pressure of the 2D and 3D chambers during the last temperature increase of the chambers. The plots are from the Grafana frontend of the monitoring system, and show the pressure of the 3D chamber in the left panel, and the 2D chamber pressure in the right side. The large increase in pressure is caused by an increase in temperature during bakeout of the chambers.*

is therefore not perfectly constant. The current of the heating wires is increased every few hours throughout the day (remaining constant at night), and the temperature is in this way increased to 100°C over the course of 72 hours. At this point, the current controller are adjusted to make sure that the temperature remains stable, and the chamber is left at this temperature for 12 full days, before the temperature is lowered by over the course of 2 days, again ensuring that the cooling is even, and no temperature gradients develop.

As the temperature is increased to 100°C, the pressure of the chamber is also increasing. Figure 5.8 shows the pressure of the chambers through time. At 13:00 on the 2. of March, the current was increased for the last temperature increasing, going from circa 90°C to 100°C in the chambers. It is clear from the figure that the pressure is increasing with the temperature, but after a few hours, when the temperature stabilise, the pressure starts to drop more or less exponentially.

During the 12 days of baking, the ytterbium dispensers need to be activated. In order to prevent the ytterbium source from getting exposed to the atmosphere, the Yb is sealed behind an Indium seal inside the dispenser unit. Thus, before the dispensers can be used for the experiment, the Indium seal must be broken. This is done simply by running current through the dispenser, which will heat the unit, and melt the Indium. When the indium seal breaks, the inert gas inside is also released, creating a rapid increase in pressure. An example of this pressure increase is shown in Figure 5.9. There are 4 dispensers in total, placed in the 2D chamber, an they are activated one-by-one, allowing time for the pressure to settle before activating the next dispenser. The activation done by sending current through the dispenser to heat melt the Indium seal, but this is the same procedure as is used to actually dispense Yb atoms, which we would rather avoid at this stage. Therefore, the activation process is done carefully, only increasing the current slightly, until the seal bursts. After the initial sudden pressure increase, the seal is broken, and the dispenser is left at that current for a few minutes to ensure that the seal is entirely melted. At this point, the current is turned of, and the pressure is allowed time to settle before the next dispenser is turned on. The activation takes about 1 hour, but could be done faster by being more aggressive with the current increases.

In addition to the activation of the dispensers, the NEGs are also activated once again, even though they were already activated in the initial pumping down. This was done as less particles will stick to the inside surfaces during bakeout, and thus a second activation could lead to a better final pressure. Also, the philosophy is that it can't hurt to activate an extra time, to be certain that any contaminants are evaporated off.

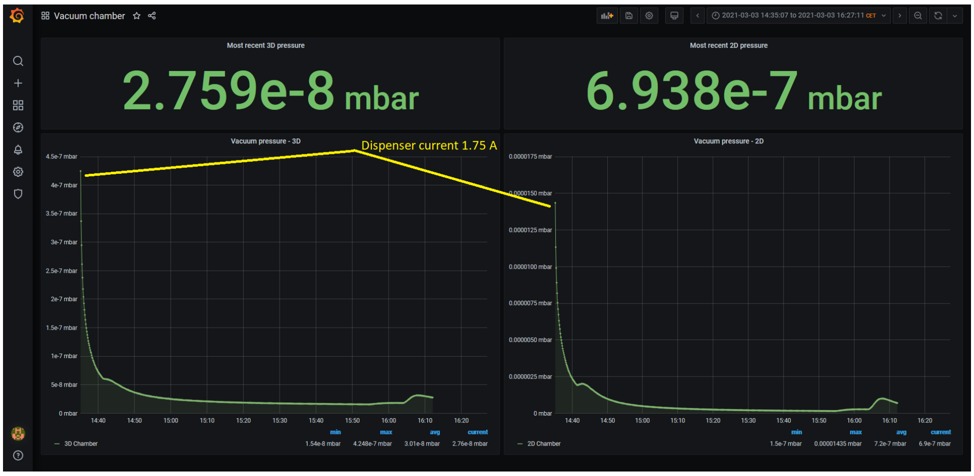

Figure 5.9: *Spike in pressure caused by the activation of one of the ytterbium dispensers inside the 2D chamber. When the dispenser is activated, the gas inside is released, causing a large increase in pressure.*

Apart from the sudden pressure increases caused by the dispenser and NEG activations, multiple virtual leaks were observed during the bake-out. A virtual leak is essentially when a small gas-pocket that have been trapped somewhere (eg. in a bellow), which gets released, causing a sudden change in pressure in the chamber. Virtual leaks can cause both very small and rather large pressure changes, although they shouldn't change the pressure by orders of magnitude by any means. We observed multiple virtual leaks throughout, but an example is shown in Figure 5.10.

With the chamber back at room temperature, and the turbo-pump disconnected, the chamber is at its final pressure, ready to use for the experiment. After bake-out, and with the NEGs running, we achieved the final pressure:

$$\text{2D CHAMBER: } 7.93 \cdot 10^{-11} \text{ MBAR}, \text{ 3D CHAMBER: } 7.74 \cdot 10^{-11} \text{ MBAR}.$$

When the dispensers are turned on, the pressure will increase slightly, but in general this final pressure is very satisfactory, and should provide a good vessel for the experiment.

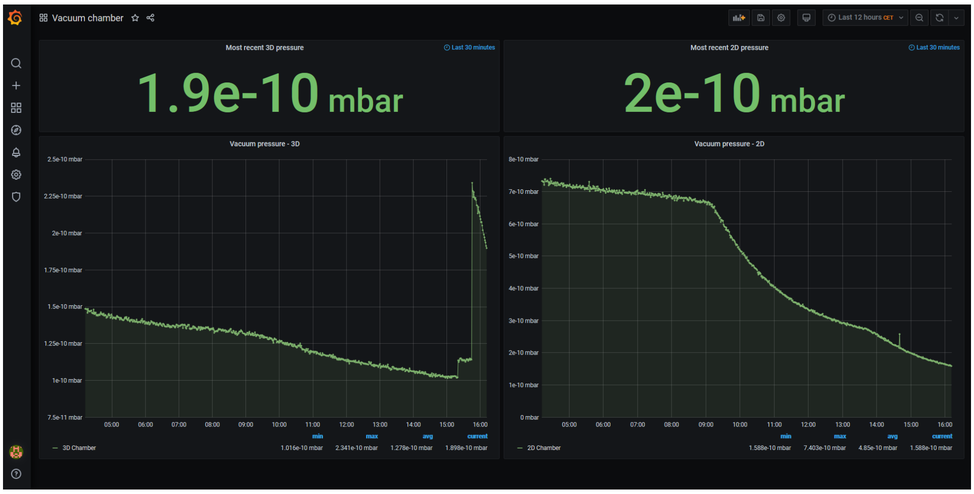

Figure 5.10: *An example of 2 virtual leaks observed during the baking. On the 3D side, the first is a minor leak, causing a 'step up' in pressure, before the rest of what is presumably the same pocket gets released. On the 2D side, a smaller virtual leak is seen in the form of a small spike in pressure. Note that the pressure is only logged every* 60 *seconds, so the actual peak of the leaks might have been higher than what the data is showing.*

Chapter 6

Summary

In this thesis, a new design for a versatile RF source is proposed, and the device is built and characterized. In particular, the device is tested with two separate operations in mind: Implementation of a SWAP cooling scheme, and realisation of movable arrays of dipole traps. The device is based on a DDS and an FPGA, due to their high reliability and precision. Both modes of operation are confirmed to be well within the requirements of the experiment in terms of stability, reliability, and timing. The new RF source is performing as expected, and is by far fulfilling the requirements set by the experiment. This is thus a very promising design for future experimental realisations of SWAP cooling, moving dipole traps, or other applications that require a versatile and precise frequency source. The RF source is constructed, and implemented in the experiment, however it still remains to use it in a real experimental sequence.

A preliminary measurement of SWAP cooling in a 3D MOT provides a proof of concept, and confirms that the cooling scheme and design provides a large potential for future experiments. The preliminary analysis implements a green SWAP MOT, which show that the ramp rate used in the SWAP scheme directly change the cloud temperature, in particular two ramps are implemented, leading to temperatures of $\sim 200\,\mu$K and $\sim 100\,\mu$K. While the device is very promising, a more detailed characterization of the SWAP cooling scheme remain to be carried out. In particular, SWAP cooling with a green 3D MOT should be thoroughly tested, to confirm how the free parameters affect the cooling efficiency, and ultimately optimise the trap as desired.

The device is furthermore used to create an array of time-averaged beams by interlacing static frequencies in the driving signal for an AOM. In this mode, the nonlinear amplitude response of an AOM/AOD is investigated, and a method to compensate for this nonlinear behaviour is developed, such that a flat intensity response can be achieved across a wide frequency range. The moving dipole traps remain to be implemented with ultracold atoms, and in particular the amplitude calibration scheme is still to be correctly implemented in the device.

The device is controlled by a computer control setup, through a number of custom Python applications. The applications have been developed as specifically for this purpose as part of this thesis, and they interface the hardware device with the existing control setup. Through this setup, the output signals are controlled with high precision, both in terms of frequency value and timing of the frequency outputs. Apart from the integration with existing control setup, a stand-alone application is also developed, and the final system is thus very flexible. In general, the control scheme and all related applications are developed in a very modular structure, such that the setup is easily modified for implementation in other experimental setups.

While the current design is very versatile, it does also have some limitations, mostly in the hardware. For instance, the bus-width of the DDS is limiting the speed at which the Rf signal can be updated, and the FPGA memory limits the number of segments, and in extent the possible complexity of the RF signal. Many of these could in principle be improved in future implementations, for instance by using the AD9910 DDS board with a wider BUS, however for the moment it is not determined to be fruitful, as the current experimental setup is not directly affected by the limitations of the current device.

During this thesis, the experiment was disassembled and moved to a brand new lab. Moving the experiment to a new lab required a brand new infrastructure around the new experiment, and the development of several new systems. This task was done in the scope of this thesis. In particular, a new monitoring system for equipment in the lab, as well as the lab climate is developed. Apart from logging data from the devices, the system also provides a notification system that alerts us in case of critical events or failures in the lab.

The monitoring system is also specifically used to monitor the vacuum chamber, and proved particularly useful in the pump-down process, after the vacuum work which was required during reconstruction. In terms of vacuum work, we describe how to prepare for working with UHV chambers, and specifically discuss the cleaning process. In the scope of this thesis in particular, we replace the differential pumping tube in the chamber during reconstruction. After reassembling the chamber, we carry out the bake-out process that is required to reach the desired pressures. After assembly and bake-out, the vacuum chamber successfully reaches the desired pressure of $\sim 10^{-11}$ mBar.

The newly developed RF device is promising for new experiments, or improvements of existing systems, and in particular it is expected that the design developed here will be implemented as a frequency source for steerable optical tweezers in another experiment in the group. Furthermore, it is expected that the SWAP cooling scheme can be implemented and thoroughly characterised in the ytterbium experiment in the near future.

Bibliography

[1] O Firstenberg, C S Adams, and S Hofferberth. Nonlinear quantum optics mediated by Rydberg interactions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(15):152003, jun 2016.

[2] H. John Caulfield and Shlomi Dolev. Why future supercomputing requires optics. *Nature Photonics*, 4(5):261–263, may 2010.

[3] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, jun 1982.

[4] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information.* Cambridge University Press, 2009.

[5] P. Bienias, S. Choi, O. Firstenberg, M. F. Maghrebi, M. Gullans, M. D. Lukin, A. V. Gorshkov, and H. P. Büchler. Scattering resonances and bound states for strongly interacting Rydberg polaritons. *Physical Review A*, 90(5):053804, nov 2014.

[6] Johannes Otterbach, Matthias Moos, Dominik Muth, and Michael Fleischhauer. Wigner crystallization of single photons in cold Rydberg ensembles. *Physical Review Letters*, 111(11):113001, sep 2013.

[7] Matthias Moos, Michael Höning, Razmik Unanyan, and Michael Fleischhauer. Many-body physics of Rydberg dark-state polaritons in the strongly interacting regime. *Physical Review A*, 92(5):053846, nov 2015.

[8] Ofer Firstenberg, Thibault Peyronel, Qi-Yu Liang, Alexey V. Gorshkov, Mikhail D. Lukin, and Vladan Vuletić. Attractive photons in a quantum nonlinear medium. *Nature*, 502(7469):71–75, sep 2013.

[9] Thomas F. Gallagher. *Rydberg Atoms.* Cambridge University Press, sep 1994.

[10] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with Rydberg atoms. *Reviews of Modern Physics*, 82(3):2313–2363, aug 2010.

[11] Inbal Friedler, David Petrosyan, Michael Fleischhauer, and Gershon Kurizki. Long-range interactions and entanglement of slow single-photon pulses. *Physical Review A*, 72(4):043803, oct 2005.

[12] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe. Single-atom single-photon quantum interface. *Science*, 317(5837):488–490, jul 2007.

[13] Danny O'Shea, Christian Junge, Jürgen Volz, and Arno Rauschenbeutel. Fiber-optical switch controlled by a single atom. *Physical Review Letters*, 111(19):193601, nov 2013.

[14] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and P. Zoller. Dipole blockade and quantum information processing in mesoscopic atomic ensembles. *Physical Review Letters*, 87(3):037901, jun 2001.

[15] Michael Fleischhauer, Atac Imamoglu, and Jonathan P. Marangos. Electromagnetically induced transparency: Optics in coherent media. *Reviews of Modern Physics*, 77(2):633–673, jul 2005.

[16] Y. O. Dudin and A. Kuzmich. Strongly interacting Rydberg excitations of a cold atomic gas. *Science*, 336(6083):887–889, apr 2012.

[17] Asaf Paris-Mandoki, Christoph Braun, Jan Kumlin, Christoph Tresp, Ivan Mirgorodskiy, Florian Christaller, Hans Peter Büchler, and Sebastian Hofferberth. Free-space quantum electrodynamics with a single Rydberg superatom. *Physical Review X*, 7(4):041010, oct 2017.

[18] M. Fleischhauer and M. D. Lukin. Dark-state polaritons in electromagnetically induced transparency. *Physical Review Letters*, 84(22):5094–5097, may 2000.

[19] Lene Vestergaard Hau, S. E. Harris, Zachary Dutton, and Cyrus H. Behroozi. Light speed reduction to 17 metres per second in an ultracold atomic gas. *Nature*, 397(6720):594–598, feb 1999.

[20] E. F. Nichols and G. F. Hull. The pressure due to radiation. *Proceedings of the American Academy of Arts and Sciences*, 38(20):559, 1903.

[21] Peter Lebedew. Untersuchungen über die druckkräfte des lichtes. *Annalen der Physik*, 311(11):433–458, 1901.

[22] T. H. Maiman. Stimulated optical radiation in ruby. *Nature*, 187(4736):493–494, aug 1960.

[23] T.W. Hänsch and A.L. Schawlow. Cooling of gases by laser radiation. *Optics Communications*, 13(1):68–69, jan 1975.

[24] Hans G. Dehmelt. Proposed 10^{14} $\nu/\Delta\nu$ laser fluorescence spectroscopy on $T1^{+}$ monoion oscillator ii. *Bull. Am. Phys. Soc.*, 20:60, 1975.

[25] William D. Phillips. Nobel lecture: Laser cooling and trapping of neutral atoms. *Reviews of Modern Physics*, 70(3):721–741, jul 1998.

[26] D. J. Wineland, R. E. Drullinger, and F. L. Walls. Radiation-pressure cooling of bound resonant absorbers. *Physical Review Letters*, 40(25):1639–1642, jun 1978.

[27] W. Neuhauser, M. Hohenstatt, P. Toschek, and H. Dehmelt. Optical-sideband cooling of visible atom cloud confined in parabolic well. *Physical Review Letters*, 41(4):233–236, jul 1978.

[28] Charles S. Adams and Erling Riis. Laser cooling and manipulation of neutral particles. In *The New Optics*. Cambridge University Press, 2001.

[29] J. Dalibard and C. Cohen-Tannoudji. Dressed-atom approach to atomic motion in laser light: the dipole force revisited. *Journal of the Optical Society of America B*, 2(11):1707, nov 1985.

[30] D. J. Heinzen and D. J. Wineland. Quantum-limited cooling and detection of radio-frequency oscillations by laser-cooled ions. *Physical Review A*, 42(5):2977–2994, sep 1990.

[31] John P. Bartolotta, Matthew A. Norcia, Julia R. K. Cline, James K. Thompson, and Murray J. Holland. Laser cooling by sawtooth-wave adiabatic passage. *Physical Review A*, 98(2), aug 2018.

[32] Jeongwon Lee, Jae Hoon Lee, Jiho Noh, and Jongchul Mun. Core-shell magneto-optical trap for alkaline-earth-metal-like atoms. *Physical Review A*, 91(5), may 2015.

[33] C. Gaul, B. J. DeSalvo, J. A. Aman, F. B. Dunning, T. C. Killian, and T. Pohl. Resonant Rydberg dressing of alkaline-earth atoms via electromagnetically induced transparency. *Physical Review Letters*, 116(24):243001, jun 2016.

[34] F B Dunning, T C Killian, S Yoshida, and J Burgdörfer. Recent advances in Rydberg physics using alkaline-earth atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(11):112003, may 2016.

[35] J. Millen, G. Lochead, and M. P. A. Jones. Two-electron excitation of an interacting cold Rydberg gas. *Physical Review Letters*, 105(21):213004, nov 2010.

[36] D. P. Sadler, E. M. Bridge, D. Boddy, A. D. Bounds, N. C. Keegan, G. Lochead, M. P. A. Jones, and B. Olmos. Radiation trapping in a dense cold Rydberg gas. *Physical Review A*, 95(1):013839, jan 2017.

[37] Ryan K Hanley, Alistair D Bounds, Paul Huillery, Niamh C Keegan, Riccardo Faoro, Elizabeth M Bridge, Kevin J Weatherill, and Matthew P A Jones. Probing interactions of thermal Sr Rydberg atoms using simultaneous optical and ion detection. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 50(11):115002, may 2017.

[38] Asaf Paris-Mandoki, Hannes Gorniaczyk, Christoph Tresp, Ivan Mirgorodskiy, and Sebastian Hofferberth. Tailoring Rydberg interactions via Förster resonances: state combinations, hopping and angular dependence. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 49(16):164001, jul 2016.

[39] Y. Takasu, K. Honda, K. Komori, T. Kuwamoto, M. Kumakura, Y. Takahashi, and T. Yabuzaki. High-density trapping of cold ytterbium atoms by an optical dipole force. *Physical Review Letters*, 90(2):023003, jan 2003.

[40] Santosh Narayanan Pisharody. *Electron Dynamics in Double Rydberg Wavepackets*. PhD thesis, Indian Institue of Technology, 1997.

[41] R Mukherjee, J Millen, R Nath, M P A Jones, and T Pohl. Many-body physics with alkaline-earth Rydberg lattices. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(18):184010, sep 2011.

[42] Ryuta Yamamoto, Jun Kobayashi, Takuma Kuno, Kohei Kato, and Yoshiro Takahashi. An ytterbium quantum gas microscope with narrow-line laser cooling. *New Journal of Physics*, 18(2):023016, feb 2016.

[43] Daniel A. Steck. Rubidium 87 d line data, November 2019.

[44] Sören Dörscher. *Creation of ytterbium quantum gases with a compact 2D-/3D-MOT setup*. PhD thesis, Universität Hamburg, 2013.

[45] Reina Maruyama. *Optical Trapping of Ytterbium Atoms*. PhD thesis, University of Washington, 2003.

[46] Francesco Scazza. *Probing SU(N)-symmetric orbital interactions with ytterbium Fermi gases in optical lattices*. phdthesis, LMU München: Faculty of Physics, 2015.

[47] R. Maruyama, R. H. Wynar, M. V. Romalis, A. Andalkar, M. D. Swallows, C. E. Pearson, and E. N. Fortson. Investigation of sub-Doppler cooling in an ytterbium magneto-optical trap. *Physical Review A*, 68(1):011403, jul 2003.

[48] Y. Takasu, K. Komori, K. Honda, M. Kumakura, T. Yabuzaki, and Y. Takahashi. Photoassociation spectroscopy of laser-cooled ytterbium atoms. *Physical Review Letters*, 93(12):123202, sep 2004.

[49] K. Beloy, J. A. Sherman, N. D. Lemke, N. Hinkley, C. W. Oates, and A. D. Ludlow. Determination of the5d6s3d1state lifetime and blackbody-radiation clock shift in Yb. *Physical Review A*, 86(5):051404, nov 2012.

[50] C. J. Bowers, D. Budker, E. D. Commins, D. DeMille, S. J. Freedman, A.-T. Nguyen, S.-Q. Shang, and M. Zolotorev. Experimental investigation of excited-state lifetimes in atomic ytterbium. *Physical Review A*, 53(5):3103–3109, may 1996.

[51] Jun Woo Cho, Han gyeol Lee, Sangkyung Lee, Jaewook Ahn, Won-Kyu Lee, Dai-Hyuk Yu, Sun Kyung Lee, and Chang Yong Park. Optical repumping of triplet-pstates enhances magneto-optical trapping of ytterbium atoms. *Physical Review A*, 85(3):035401, mar 2012.

[52] J Migdalek and W E Baylis. Relativistic transition probabilities and lifetimes of low-lying levels in ytterbium. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 24(4):L99–L102, feb 1991.

[53] H. Lehec, A. Zuliani, W. Maineult, E. Luc-Koenig, P. Pillet, P. Cheinet, F. Niyaz, and T. F. Gallagher. Laser and microwave spectroscopy of even-parity Rydberg states of neutral ytterbium and multichannel-quantum-defect-theory analysis. *Physical Review A*, 98(6):062506, dec 2018.

[54] H. Gorniaczyk, C. Tresp, P. Bienias, A. Paris-Mandoki, W. Li, I. Mirgorodskiy, H. P. Büchler, I. Lesanovsky, and S. Hofferberth. Enhancement of Rydberg-mediated single-photon nonlinearities by electrically tuned Förster resonances. *Nature Communications*, 7(1), aug 2016.

[55] Matteo Marcuzzi, Jiří Minář, Daniel Barredo, Sylvain de Léséleuc, Henning Labuhn, Thierry Lahaye, Antoine Browaeys, Emanuele Levi, and Igor Lesanovsky. Facilitation dynamics and localization phenomena in Rydberg lattice gases with position disorder. *Physical Review Letters*, 118(6):063606, feb 2017.

[56] Emanuele Distante, Pau Farrera, Auxiliadora Padrón-Brito, David Paredes-Barato, Georg Heinze, and Hugues de Riedmatten. Storing single photons emitted by a quantum memory on a highly excited Rydberg state. *Nature Communications*, 8(1), jan 2017.

[57] Y. O. Dudin, L. Li, F. Bariani, and A. Kuzmich. Observation of coherent many-body Rabi oscillations. *Nature Physics*, 8(11):790–794, sep 2012.

[58] T. G. Tiecke, J. D. Thompson, N. P. de Leon, L. R. Liu, V. Vuletić, and M. D. Lukin. Nanophotonic quantum phase switch with a single atom. *Nature*, 508(7495):241–244, apr 2014.

[59] H. Gorniaczyk, C. Tresp, J. Schmidt, H. Fedder, and S. Hofferberth. Single-photon transistor mediated by interstate Rydberg interactions. *Physical Review Letters*, 113(5):053601, jul 2014.

[60] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. D. Yavuz, T. G. Walker, and M. Saffman. Observation of Rydberg blockade between two atoms. *Nature Physics*, 5(2):110–114, jan 2009.

[61] C. Tresp, C. Zimmer, I. Mirgorodskiy, H. Gorniaczyk, A. Paris-Mandoki, and S. Hofferberth. Single-photon absorber based on strongly interacting Rydberg atoms. *Physical Review Letters*, 117(22):223001, nov 2016.

[62] Stefano Giorgini, Lev P. Pitaevskii, and Sandro Stringari. Theory of ultracold atomic Fermi gases. *Reviews of Modern Physics*, 80(4):1215–1274, oct 2008.

[63] H. Wineland, D. J.; Dehmelt. Proposed $10^{14}\delta\nu/\nu$ laser fluorescence spectroscopy on TI^+ mono-ion oscillator iii. *Bulletin of the American Physical Society*, 20(637), 1975.

[64] Christopher J. Foot. *Atomic Physics*. PAPERBACKSHOP UK IMPORT, January 2005.

[65] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook. Optical molasses. *Journal of the Optical Society of America B*, 6(11):2084, nov 1989.

[66] Mark Fox. *Quantum Optics*. Oxford University Press, 2006.

[67] Wilhelm Brenig. *Statistical Theory of Heat*. Springer Berlin Heidelberg, 1989.

[68] Paul D. Lett, Richard N. Watts, Christoph I. Westbrook, William D. Phillips, Phillip L. Gould, and Harold J. Metcalf. Observation of atoms laser cooled below the Doppler limit. *Physical Review Letters*, 61(2):169–172, jul 1988.

[69] Peter Horak, Gerald Hechenblaikner, Klaus M. Gheri, Herwig Stecher, and Helmut Ritsch. Cavity-induced atom cooling in the strong coupling regime. *Physical Review Letters*, 79(25):4974–4977, dec 1997.

[70] A. M. Steane, M. Chowdhury, and C. J. Foot. Radiation force in the magneto-optical trap. *Journal of the Optical Society of America B*, 9(12):2142, dec 1992.

[71] Hidetoshi Katori, Tetsuya Ido, Yoshitomo Isoya, and Makoto Kuwata-Gonokami. Magneto-optical trapping and cooling of strontium atoms down to the photon recoil temperature. *Physical Review Letters*, 82(6):1116–1119, feb 1999.

[72] John P. Bartolotta and Murray J. Holland. Sawtooth-wave adiabatic passage in a magneto-optical trap. *Physical Review A*, 101(5), may 2020.

[73] Clarence Zener. Non-adiabatic crossing of energy levels. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 137(833):696–702, sep 1932.

[74] Juan A. Muniz, Matthew A. Norcia, Julia R. K. Cline, and James K. Thompson. A robust narrow-line magneto-optical trap using adiabatic transfer. June 2018.

[75] S. Snigirev, A. J. Park, A. Heinz, I. Bloch, and S. Blatt. Fast and dense magneto-optical traps for strontium. *Physical Review A*, 99(6):063421, jun 2019.

[76] Rudolf Grimm, Matthias Weidemüller, and Yurii B. Ovchinnikov. Optical dipole traps for neutral atoms. In *Advances In Atomic, Molecular, and Optical Physics*, pages 95–170. Elsevier, 2000.

[77] A. Ashkin. Acceleration and trapping of particles by radiation pressure. *Physical Review Letters*, 24(4):156–159, jan 1970.

[78] A. Ashkin. Atomic-beam deflection by resonance-radiation pressure. *Physical Review Letters*, 25(19):1321–1324, nov 1970.

[79] H. J. Metcalf and P. van der Straten. Laser cooling and trapping of atoms. *Journal of the Optical Society of America B*, 20(5):887, may 2003.

[80] L. Allen and Joseph H. Eberly. *Optical Resonance and Two-Level Atoms*. DOVER PUBN INC, December 1987.

[81] A Ashkin and J. Dziedzic. Optical trapping and manipulation of viruses and bacteria. *Science*, 235(4795):1517–1520, mar 1987.

[82] T. L. Gustavson, A. P. Chikkatur, A. E. Leanhardt, A. Görlitz, S. Gupta, D. E. Pritchard, and W. Ketterle. Transport of Bose-Einstein condensates with optical tweezers. *Physical Review Letters*, 88(2):020401, dec 2001.

[83] Nina Stiesdal, Hannes Busche, Kevin Kleinbeck, Jan Kumlin, Mikkel G. Hansen, Hans Peter Büchler, and Sebastian Hofferberth. Controlled multi-photon subtraction with cascaded Rydberg superatoms as single-photon absorbers. March 2021.

[84] Ian Kuon and Jonathan Rose. Measuring the gap between FPGAs and ASICs. In *Proceedings of the internation symposium on Field programmable gate arrays - FPGA'06*. ACM Press, 2006.

[85] Eva Murphy and Colm Slattery. Ask the application engineer—33 all about direct digital synthesis. *Analog Dialogue*, 38(3):8–12, 2004.

[86] Ulf Grenander. *Probability and Statistics: The Harald Cramér Volume*. Almqvist and Wiksell, Wiley, 1959.

[87] Ghareeb Falazi. Cold physics experiments control software (cpecs), November 2020.

[88] Justin E. Molloy. Chapter 12 optical chopsticks: Digital synthesis of multiple optical traps. In *Methods in Cell Biology*, pages 205–216. Elsevier, 1997.

[89] S. K. Schnelle, E. D. van Ooijen, M. J. Davis, N. R. Heckenberg, and H. Rubinsztein-Dunlop. Versatile two-dimensional potentials for ultra-cold atoms. *Optics Express*, 16(3):1405, 2008.

[90] M. R. Andrews. Observation of interference between two Bose condensates. *Science*, 275(5300):637–641, jan 1997.

[91] K. O. Roberts, T. McKellar, J. Fekete, A. Rakonjac, A. B. Deb, and N. Kjærgaard. Steerable optical tweezers for ultracold atom studies. *Optics Letters*, 39(7):2012, mar 2014.

[92] Philipp Lunt. Design and construction of a new ultracold ytterbium experiment for Rydberg physics. mathesis, University of Southern Denmark, January 2019.

[93] S. Pradhan and B. N. Jagatap. Measurement of temperature of laser cooled atoms by one-dimensional expansion in a magneto-optical trap. *Review of Scientific Instruments*, 79(1):013101, jan 2008.

[94] Johann Gan, M. E. Pantalon, and F. Robicheaux. Simulations of sawtooth-wave adiabatic passage with losses. *Physical Review A*, 101(1):013422, jan 2020.

[95] The Quantum Optics Group Birnbaum, K. M. Ultra-high vacuum chambers. Technical report, California Institute of Technology, Norman Bridge Laboratory of Physics. CA 91125. Tech. Rep. 12-33, 2005.

[96] A. Roth. *Vacuum Technology*. Elsevier Science and Techn., December 1990.

[97] Pfeiffer Vacuum GmbH. The vacuum technology book volume ii - know how book, April 2013.

Acknowledgement

This thesis would not have been possible without the immense help I received from others. My acknowledgement of all the people who contributed directly or indirectly to this work could fill numerous pages! In particular, I would like to mention

- Sebastian Hofferberth for guidance and supervision, and especially for the opportunity to work on this project, despite of the entire experiment moving to Bonn. By now, I've had the unique opportunity of setting up an experiment from scratch twice, and I learn a lot about the lasers, and the many experimental methods every time. As with my previous projects with you, I've learned a lot, and really appreciate that you let me do yet another very technical project.
- The group members, both current and former, and especially Noaman, Rafael, Thilina, Nina, Hannes, Mikkel, Simon, Philipp, This group initially got me into experimental physics, when the ytterbium lab came was build in Odense, and the self-cooked lunch made me stay! It's been a pleasure to be a part of the group throughout the years, even though I've been coming and going from time to time.
- Mohammad Noaman for always being ready to explain all the picky details about th FPGA setup, and answer all my questions in the lab - even if I did not catch the answer the first few times. This thesis would not have been possible without your help with the FPGA programming, and insight in the inner working of the hardware.
- Nina Stiesdal for the immense help during the final days of the thesis, for always being happy, and for keeping me active outside the lab with running in Bonn. You are also always willing to help, no matter how stupid my questions are, and you have saved me multiple times, and can always point towards the right answer if something is unclear!
- Philipp Lunt for convincing me that experimental qunatum optics is the right path. The long evenings in Odense of initially constructing the experiment, the following beer, and the many duels in table-soccer have all been a large part of my introduction to the field.
- My fellow students, and in particular Signe Wind, Frederik Kamper Jørgensen, and Peter Græns Larsen for the near-weekly sessions of DnD, socialising, and complaining about our projects. Complaining about our problems in unison make the problems seem less significant, and tlaking to you provided nice, and needed, breaks from the thesis.

Appendix A

Tables of FPGA memory structure

Table A.1: *Structure of the binary data used for the SWAP cooling signals. The parameters are split into 3 different 64-bit long 'words', that are stored in the FPGA memory. The first word specifies the initial and final frequencies for the sawtooth ramps, the second word specify the slope with RDW/FDW, and output channel. The last word specify the number of steps in the increasing and decreasing ramps, and the number of full cycles in the segment. The use of these parameters in a waveform is shown in Figure 2.11.*

Binary word 1	
63:32	31:0
S	E
Start frequency, ie. lower frequency of the sawtooth ramp	End frequency, ie. upper frequency of the sawtooth ramp

Binary word 2					
63:40	39:16	15:12	11:9	8	7:0
RDW[31:8]	FDW[31:8]	channel_select	reserved	stop_increment	profile_pin_setup
Rising Delta Word specifies the frequency change for each incrementing step	Falling Delta Word specifies frequency when decreasing	Selects the output channel (0, 1, 2, 3)		Specifies if this is the last segment	Not directly used for the signal. Controls the up/down signal etc.

Binary word 3		
63:48	47:32	31:0
n_up_steps	n_dn_steps	n_ramps
Duration in which the frequency is increased, in units of 4ns.	Duration in which the frequency is decreased, in units of 4ns.	Number of up and down cycles before moving to next segment (unless triggered externally before this time)

Table A.2: *Binary words for specifying interlaced signals. The data is specified in 64-bit binary rows. The first row specifies the starting frequency in absolute terms, as well as the number of following segments. Each following row then specifies a frequency change relative to the last frequency, the duration of the frequency change (ie., the ramp rate), and whether the signals are interlaced.*

First 64-bit row			
63:32	31:24	23:16	15:0
Starting frequency	Number of memory lines to read (ie. number of frequency updates/sweeps)	Not used	Not used

Following rows					
63:32	31:24	23:16	15:8	7:1	0
Change of frequency relative to previous segment	Unused	Number of segments to perform the frequency change. Each segment is 4ns.	Split mode: Specifies how many signal are interlaced at this point	Unused	Sign of sweep. Positive when splitting, and negative when merging together again.