

Chapter 3

Experimental Apparatus: the femtosecond Ti:S setup and the MOT

This chapter discusses the two primary components of the experimental setup, the femtosecond laser and the MOT. I will first present the idea for the Ti:S laser stabilization and then describe one of the most commonly used self-referencing schemes. I will conclude by briefly discussing the stabilization of the cw lasers and the setup used for cooling and trapping in the MOT. Chapter 4 will present the integration of the femtosecond laser and the MOT.

3.1 The femtosecond Ti:S laser and its stabilization setup

Recent advances in controlling the pulse repetition period and the carrier-envelope phase of mode-locked lasers have ushered in a series of novel applications using their frequency combs [42]. In the introduction I presented some of the fundamental aspects of mode-locked lasers. In this section, I will review some of the basic concepts behind carrier-envelope phase stabilization and I will explain how the two dynamic laser parameters are optically detected and stabilized in the experiment. Stabilization of the two degrees of freedom is a necessary step for using these lasers for spectroscopic applications.

3.1.1 Stabilizing the two degrees of freedom of mode-locked lasers

For the experiment, we used a home-built titanium-doped sapphire (Ti:S) laser pumped by a 5 W single-frequency, diode-pumped, frequency-doubled Nd:YVO₄ laser (Coherent Verdi) operating at 532 nm. The Ti:S laser generates a 100-MHz pulse train with pulse widths on the order of 15 fs and energies of about 5 nJ per pulse, using Kerr lens mode-locking (KLM). A diagram of our KLM laser is shown in Fig. 3.1. The configuration used is a standing-wave, folded geometry, with the cavity delimited by two flat mirrors: the laser end mirror (EM) and the partially reflective output coupler (OC). The nonlinear gain element for the laser cavity is the Ti:S crystal. The passive mode-locking mechanism employed here is the nonlinear Kerr Effect, which gives an increase in the index of refraction of the crystal as the optical intensity is increased (self-phase modulation, self-focusing, see for example [43]). The Ti:S acts like a nonlinear lens, slightly focusing the intracavity beam, but it also spreads the pulses temporally through normal dispersion (the long wavelengths travel faster than the short ones). As a consequence, a sequence of two prisms is used to compensate for the group velocity dispersion (GVD) in the crystal: the first prism spatially disperses the pulse, leading the long wavelengths to traverse more glass in the second prism than the short wavelengths. The beam is still horizontally dispersed after the second prism, though collimated. Upon reflection by the laser cavity EM, the pulse travels again through the pair of prisms, thus canceling the spatial dispersion.

As mentioned before, f_r determines the mode spacing and is inversely proportional to the cavity length L , $f_r = \frac{v_g}{L}$, while f_0 sets the absolute comb position and is proportional to the intracavity dispersion, $f_0 = \frac{\omega_c}{2\pi} (1 - \frac{v_g}{v_p})$. They are both influenced by the group velocity v_g , but there is no apparent f_0 dependence on the laser cavity length. Ideally one would like to control the two parameters independently. One adjustable parameter can be the cavity length, for f_r control; yet, an additional degree of




Figure 3.1: Optical layout of the KLM Ti:S laser used in the experiment; the end mirror (EM) and the output coupler (OC) are the two flat mirrors bounding the cavity, used for the f_0 and f_r stabilization, respectively. The two prisms P_1 and P_2 are used for dispersion compensation.

freedom is needed for changing v_g . The one employed in my laser was first suggested in 1999 by the Hänsch group [44] and uses small-angle rotations about the vertical axis of the laser EM. As previously stated, the different spectral components in the laser beam are horizontally spread out at the EM, leading to a linear relationship between the spatial coordinate and the wavelength. For small angle (β) changes, there is a linear path length change, i.e. phase shift, with frequency. This is effectively a group delay, linear with the angle. Under this assumption that the swivel mirror changes the group delay by a small amount $\sim \beta$, it can be shown [28] that f_0 is only controlled by β , while f_r depends on both β and L . Therefore, one can use the swivel angle β to control f_0 (this will also change f_r !), and the cavity length L to compensate for the subsequent change in f_r . Details of the experimental implementation of the v_g , and hence f_0 , control will be given in 3.1.5.

3.1.2 Stabilization of the laser repetition rate

The laser repetition frequency f_r is a radio frequency (rf) ~ 100 MHz and is measured directly using a fast photodiode. To reduce the increase in phase noise that is due to the large frequency multiplication factor up into the optical region ($\sim 10^6$), a

higher harmonic of the pulse repetition rate is used for stabilization. A portion of the pulse train is detected with a high-speed photodetector; the 10th harmonic of f_r at 1 GHz is filtered and amplified (to 0 dBm, suitable for the rf port of a double-balanced mixer) and then phase-locked to the 1 GHz provided by a stable microwave source. The error signal obtained by comparing against this reference frequency is integrated, filtered and amplified in a JILA-built loop filter (LF) and then used for active feedback to the laser, as shown in Fig. 3.2. Cavity length corrections are controlled with a small (~ 5 mm) tube piezoelectric transducer (PZT), on which the output coupler of the laser has been mounted. All the comb modes are moved together by this translating PZT which varies the laser cavity length.




Figure 3.2: Stabilization diagram used for phase-locking of the laser repetition rate f_r ; BPF stands for band-pass filter, DDS for direct digital synthesizer and LF for loop filter. More details are given in the text.

Several stable rf oscillators were tried as references for the repetition rate (I will address this issue further in Chapter 4) and in the end a custom-made Wenzel oscillator, based on a crystal oscillator that has a very good short-term frequency stability, was settled upon. The short-term stability of the rf source is very important, given the huge multiplication factor up into the optical domain. Unfortunately, the Wenzel tuning range is only 3.2 kHz around 1 GHz, corresponding to a 320 Hz tunability of

the repetition rate very close to 100 MHz. One would like to have more freedom than that in choosing a suitable f_r value, so an extra direct digital synthesizer (DDS) had to be later added to the stabilization setup. Its widely tunable output (\sim MHz) is mixed with the Wenzel and the repetition rate is forced to track the resulting frequency via the phase lock described above.

3.1.3 Self-referencing

In 1999 Telle and coworkers suggested several methods allowing for direct measurement and stabilization of the frequency offset of the fs comb, without using another optical standard [21]. These ‘self-referencing’ methods employ non-linear processes like second harmonic generation to compare the spectral extremes of the comb. In that work, it is shown that frequency-doubling of a number of modes from the infrared end of the spectrum and heterodyning them with the existing visible section of the spectrum will result in a beat frequency, equal to the offset of the fs comb from integer multiples of the repetition rate:

$$2\nu_N - \nu_{2N} = 2(Nf_r + f_0) - (2Nf_r + f_0) = f_0. \quad (3.1)$$

Thus, f_0 can be derived directly from the beat note, with no knowledge of the laser repetition rate. The requirement for this detection scheme to work is that the frequency comb contain on the blue side the second harmonic of the red side, i.e. that the laser optical bandwidth cover an optical octave. Along the same line of thought, the use of mode-locked lasers for absolute metrology of optical frequencies, where the frequency of every comb line is an integer multiple of the fixed comb spacing ($f_0 = 0$), was suggested in a NIST competence proposal from March 1999 [45].

Fortunately, the introduction in 2000 of special air-silica microstructure (MS) optical fibers [22], which broaden the output of the fs comb to span an entire optical octave solved the bandwidth requirement. These photonic crystal fibers have a very small core

diameter ($1.7 \mu\text{m}$), providing a very tight spatial confinement that generates very high intensities ($\sim \text{TW}/\text{cm}^2$). They have been engineered to exhibit zero GVD near 800 nm, which minimizes the temporal spread of the pulses propagating in the fiber over significant distances (several cm). The main non-linear effect leading to spectral broadening in the fiber is strong four-wave-mixing among the different frequency components in the input pulse spectrum, which generates new spectral components, while maintaining the periodicity of the incoming frequency comb. The spectral bandwidth at the output of the MS fiber can extend well beyond one optical octave e.g. 400 to 1200 nm.

The first experimental implementation of the self-referencing technique using external broadening in a MS fiber, followed by doubling of the 1040 nm light and mixing with the fundamental at 520 nm, was reported by Jones and co-workers in 2000 [1]. Ultra-broadband mode-locked lasers whose direct output spans on optical octave, allowing for direct self-referencing and stabilization of the pulse train, have also been recently demonstrated [7].

3.1.4 The ν -to- 2ν prism interferometer

To achieve spectral broadening, $\sim 120 \text{ mW}$ of the fs laser output light ($\sim 500 \text{ mW}$) are coupled into a $\sim 7\text{-cm}$ piece of MS fiber. The output of the fiber is then sent into an ν -to- 2ν interferometer for self-comparison.

The configuration chosen for the interferometer is the so-called ‘prism-pair’ geometry, presented in Fig. 3.3. The sequence of SF10 prisms (8-cm prism separation) spatially separates the green (e.g. 532 nm) and red (e.g. 1064 nm) components of the spectrum, which are then reflected by two mirrors located very close to each other. The green-reflecting mirror is mounted on a translation stage, to easily adjust the temporal overlap between the red and green pulses at the detector. When the interferometer is well aligned and working, the distance between the end mirrors is $\sim 1 \text{ mm}$.

The beams retrace their path through the prism pair with a slight vertical offset

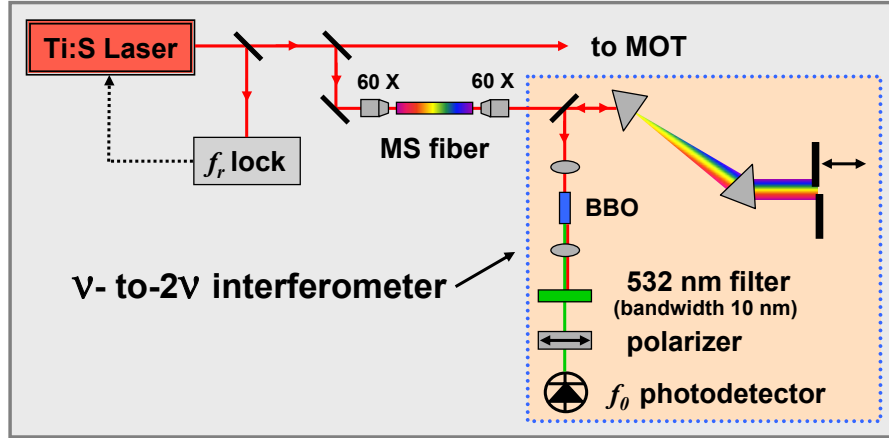


Figure 3.3: Diagram of the prism interferometer used for the detection of the comb offset frequency. The broad spectrum from the MS fiber is sent into a prism pair for spatial dispersion. The light near 532 nm is reflected by one end mirror, while the light near 1064 nm hits a second end mirror. Both beams are reflected back through the prism pair, recombined and then sent through a nonlinear BBO crystal tuned for efficient second harmonic generation at 1064 nm. The output is filtered, and the f_0 beat is detected on a PD.

in order to be picked off by a mirror positioned below the incoming beam. The pick-off mirror sends the recombined beams to a lens, which focuses them into a nonlinear crystal (1-mm BBO, type I phase-matching), angle-tuned for efficient second harmonic generation at 1064 nm.

Two types of green light are present at the output of the BBO crystal: the frequency-doubled beam obtained from the infrared part of the spectrum and the fundamental green coming out of the MS fiber; the polarization of the frequency-doubled green is rotated by 90 degrees with respect to the fundamental green. A 10-nm interference filter centered at 532 nm removes the extra, non-contributing spectral components, while a polarizer ensures that the polarizations of the two green beams are aligned along the same axis. Spatial overlap and mode-matching of the two beams are essential for a good signal-to-noise ratio (S/N) of the beat signal. A photodiode detects their difference frequency, the f_0 beat, along with the laser repetition rate. In reality, the detected beat

arises from mixing of thousands of comb lines, not just the two frequencies mentioned, giving a better S/N. In the experiment, the typical S/N of the offset beat is ~ 45 dB in a 100 kHz resolution bandwidth.

The prism pair geometry is advantageous to use compared to other configurations, e.g. a Mach-Zender interferometer, because most of its elements are common-mode (the only differential path is the prism pair region), which makes it easier to align and tweak.

3.1.5 Stabilization of the laser offset frequency

The output of the ν -to- 2ν PD is sent through two 50 MHz low-pass filters (LPF), to reduce all the unnecessary rf frequencies in the spectrum (otherwise, they would saturate the subsequent amplifiers), then a small portion of the signal is sent via a directional coupler to an rf spectrum analyzer for monitoring of the beat note. The main signal is amplified to ~ -5 dBm and then frequency-prescaled (divided by 64) before its phase is compared to a stable rf oscillator such as a DDS. Unlike the f_r lock which employs an analog phase detector, the stabilization circuit used for f_0 employs a digital phase detector (DPD), which provides both a frequency and a phase lock. This allows for a larger capture and locking range than the double-balanced mixer case and increases the stability of the lock. Again, the error signal derived by comparing against the reference frequency is integrated, filtered and amplified in a JILA LF (see Fig. 3.4) and then sent to the laser for frequency correction. Note that the local oscillators used for the f_r and f_0 stabilization are all referenced to the same local commercial cesium clock.

As announced earlier in section 3.1.1, f_0 control is carried out by fine-tuning the laser end mirror angle using a twister PZT. The twister PZT used to tilt the end mirror is a 0.5 inch diameter cylinder, about 0.5 inch long, that has been modified such that the outer electrical surface is split to allow for the two outer sides to have different voltages with respect to the inner surface of the PZT. When applying opposite-sign voltages to

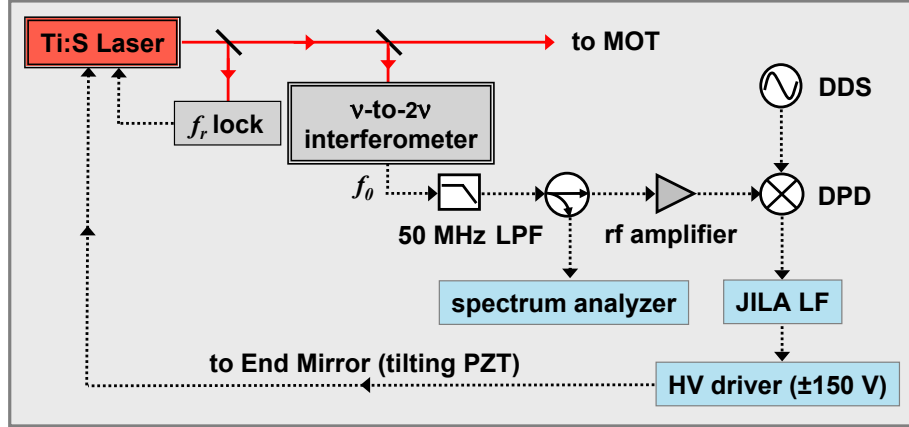


Figure 3.4: Stabilization diagram used for phase-locking of the laser offset frequency f_0 ; LPF stands for low-pass filter, DPD for direct phase detector, DDS for direct digital synthesizer and LF for loop filter. The details are given in the text.

the two sides, one contracts, while the other expands, with the net effect of an angle change for the end mirror mounted on the PZT. As mentioned before, this is effectively a group delay, which depends on the voltage sent to the PZT. Changing the voltage in the 0 to ± 150 V range allows for scanning of the offset frequency beat over its full range of ~ 50 MHz.

Another method of adjusting and stabilizing f_0 is employing an acousto-optic modulator (AOM) to control the pump laser power. Modulating the pump laser power actually affects several laser parameters. Systematic studies of intensity-related dynamics in mode-locked Ti:S lasers have been recently undertaken [46,47] to understand the effects on f_r and f_0 .

To summarize, f_r and f_0 stabilization and referencing to an atomic clock result in knowledge of the absolute frequencies of the $\sim 10^6$ comb lines, with an accuracy only limited by the frequency reference used. The stabilized comb can now be employed as a precision frequency ruler for any optical frequency in the visible to near-infrared range.

3.2 The MOT

The magneto-optical trap (MOT) is currently the starting point for many experiments with laser-cooled atoms. Cold atoms are used for precision measurements, atom interferometers, atomic clocks, to name just a few examples. Advances in laser cooling of alkali atoms [48] have also paved the way to producing the first Bose-Einstein condensates and degenerate Fermi gases [49–51]. In particular, vapor-cell traps [52] like the one used in this experiment, are simple and versatile tools for many experiments with cold atoms.

Our experiment was performed on a sample of laser-trapped and -cooled Rubidium 87 atoms (see Fig. 4.2 for a detailed ^{87}Rb level diagram, $I=3/2$). The $5S_{1/2} \rightarrow 5P_{3/2}$ cooling transitions can be easily accessed with commercially available laser diodes centered at 780 nm, making ^{87}Rb a convenient choice for laser cooling and trapping. We use two lasers to trap and cool the atoms. The cooling beams are ~ 10 MHz red-detuned from the $5S_{1/2}$ ($F=2$) \rightarrow $5P_{3/2}$ ($F'=3$) cycling transition. An additional laser, the so-called ‘repumper’, tuned to the $5S_{1/2}$ ($F=1$) \rightarrow $5P_{3/2}$ ($F'=2$) transition, is needed to return the atoms to the cooling cycle whenever they decay to the $F=1$ ground level, thus preventing them from falling dark due to optical pumping.

3.2.1 Diode lasers: description and stabilization

The laser diodes used are Sharp LT024MD (specified output power of 20 mW) for the repump laser (RL) and Sanyo DL7140-201 (specified output power of 70 mW) for the trap laser (TL). They are configured in an external cavity geometry, based on a design described in reference [53, 54], with ~ 3.5 cm extended cavity length. The lasers employ a diffraction grating (1200 lines/mm) aligned in the Littrow configuration [55] to couple the first diffraction order back into the laser. One PZT controls the horizontal tilt of the grating, permitting adjustments of the laser wavelength. A change in the PZT

length alters the grating angle and the external cavity length at the same time, allowing for continuous scans of typically ~ 5 GHz. Additionally, the grating is mounted on a fast disk PZT, which is servoed together with the laser diode current to stabilize the laser frequency.

Both diode lasers are locked to the specified atomic transitions using standard saturated absorption spectroscopy [53], which yields sub-Doppler hyperfine lines that are used to frequency-stabilize the laser (see Fig. 3.5). The TL is locked to the red side of the saturated absorption peak; a side lock is not ideal, since the locking set point is inevitably dependent on power broadening of the transition lines. We use acousto-optic modulators (AOMs) to create beams at various detunings and also to control the timing of the lasers in the experiment. The TL light is steered through two AOMs: the first AOM deflects a small amount of power (-1 order) into the saturated absorption spectrometer, for frequency stabilization. The second AOM acts as a shutter for the light, by directing the negative first order diffracted beam to the MOT cell. Thus, the MOT red detuning is set by the side lock along with the frequency difference of the two AOMs.

The RL is locked to the peak of the crossover transition between the $5S_{1/2}$ ($F=1$) \rightarrow $5P_{3/2}$ ($F'=1$) and the $5S_{1/2}$ ($F=1$) \rightarrow $5P_{3/2}$ ($F'=2$) transitions, using modulation sidebands at 4.6 MHz, written directly on the diode-laser current (the laser output current can be modulated rapidly by sending a voltage into the rf modulation input, this in turn will modulate the frequency of the laser). The signal from the RL saturated absorption spectrometer is detected with a homemade resonant photodetector (PD) and then compared in a double-balanced mixer with the local oscillator at 4.6 MHz, resulting in dispersion curves for all the features in the saturated absorption spectrum [56, 57]. The value of these dispersion curves changes sign at the peak, allowing for a lock at the zero-crossing of the error signal. Since the RL is offset peak-locked, an AOM is needed to shift the frequency into resonance with either the $5S_{1/2}$ ($F=1$) \rightarrow $5P_{3/2}$ ($F'=1$) or the




Figure 3.5: Saturated absorption spectra, showing the hyperfine structure of the $5S_{1/2} \rightarrow 5P_{3/2}$ transitions for the two Rb isotopes. The numbers labeling the hyperfine peaks correspond to F' states in the $5P_{3/2}$ manifold. Note that the spectra contain both true Doppler-free resonances (numbered peaks) and crossover resonances. The crossover resonances occur at frequencies $(\nu_a + \nu_b)/2$ for each pair of true resonances at frequency ν_a and ν_b , as exemplified for the bottom spectrum.

$5S_{1/2} (F=1) \rightarrow 5P_{3/2} (F'=2)$ transition. We opted for the latter and thus the positive first order deflected beam is sent to the MOT vapor cell. In order to operate the trap at zero repump detuning, the driving frequency for the RL AOM is equal to exactly half of the $F'=1 - F'=2$ $5P_{3/2}$ hyperfine interval (78.45 MHz). This AOM is a fast and convenient switch used to turn off the RL beam whenever needed.

3.2.2 MOT details

The main TL beam is spatially filtered by a $70 \mu\text{m}$ pinhole, which reduces the high-frequency spatial modes, and collimated to roughly a 1-cm beam diameter with a beam-enlarging telescope. It is then divided into three separate, intensity-balanced beams using $\lambda/2$ plates and polarizing beamsplitter cubes. The three beams are sent to the experiment glass cell using 1-inch diameter optics, in a standard retro-reflected MOT configuration [58]. The atoms are caught in the MOT using these three pairs of orthogonal, circularly-polarized beams in the typical $\sigma^+ - \sigma^-$ arrangement, and detuned ~ 10 MHz to the red of the atomic resonance. The RL beam is directed into the cell along one of the TL beams. We monitor the beams in the cell and the position of the trapped cloud with a low-cost CCD camera.

We load the MOT from the vapor produced by Rb dispensers ('getters') in a cylindrical glass cell. The vapor cell is continuously pumped by an ion pump and typically reaches a background pressure of 10^{-9} Torr. The source of Rb consists of three nichrome (Ni-Cr alloy) dispenser strips, connected in parallel and placed in a small side arm of the cell. They contain a Rb salt and release atoms through a chemical reaction when an electric current is passed through them and heats the nichrome. Only one dispenser is used at a time, and the operating current is usually 3.5 A.

The quadrupole magnetic field necessary to operate a MOT is provided by a pair of anti-Helmholtz coils. The two coils are wound from 'magnet wire', i.e. Cu wire (~ 1 mm diameter) coated with an insulating material, and are wrapped around the axis of

the cylindrical cell. Each coil is 7 cm in diameter and has 33 turns. The coils are spaced 5.5 cm apart and are supplied with ~ 4.5 A, generating a magnetic field gradient of ~ 10 G/cm.

Stray magnetic fields, due mainly to the Earth's field, can shift the zero-point of the MOT field. They are compensated by three pairs of orthogonal Helmholtz coils. The glass cell is surrounded by these pairs of shim coils, each run by a different current supply. In the early stages of the experiment, the beam alignment and ambient magnetic field compensation were verified by monitoring the isotropy of the molasses expansion. We were aiming for a slow and uniform expansion of the atomic cloud upon switching off the quadrupole field. At a later stage, after obtaining the two-photon signal as described in Chapter 4, we realized that we could use it to find the zero-point of the magnetic field.

Resonance fluorescence from the atoms is detected with a photodiode (PD), leading to an estimate of the atom number based on the PD current, the PD responsivity, the solid angle of the light collected by the imaging lens, the energy of a photon and the optical scattering rate. We typically trap 5×10^6 atoms in a 0.5 mm diameter MOT with a $1/e$ lifetime between 3 and 5 s.

After collecting the atoms in the MOT, we transfer them to optical molasses [59] for further cooling using polarization gradients (PGC). To experimentally implement PGC, the quadrupole field gradient is turned off, the TL is further red-detuned from resonance (~ 40 MHz) and the TL beam intensity is decreased. The resulting sub-Doppler cloud temperature established by PGC is ~ 15 μ K, providing an ideal initial condition for spectroscopy.

I have presented here the basic operation details of the MOT. In Chapter 4 I will describe how the trapped atoms are probed with the frequency comb of a femtosecond laser.