LUND UNIVERSITY
Faculty of Science

Development of narrow-bandwidth filters for the suppression of scattered light for optical and ultrasound analysis of tissue

Philip Dalsbecker

Thesis submitted for the degree of Master of Science
Project duration: 4 months

Supervised by Lars Rippe, assistant supervisors Stefan Kröll and Adam Nilsson

Department of Physics
Division of Atomic Physics
May 2015

Abstract

Ultrasound-modulated optical tomography (UOT) is a proposed diagnostic technique with potential as a means of discovering cancerous tumours. The technique combines ultrasound and laser radiation to probe tissue in a non-invasive, harmless way. An important aspect in UOT is the ability to distinguish the probing light from background scattered light generated by the technique, which limits its resolution and usefulness. In this thesis work, spectrally hole-burned Pr^{3+}:Y_2SiO_5 crystals are used as spectral filters to block out the background light. Thus, the resolution depends on the filters, which this thesis work has aimed to improve.

As part of this work, simulations and experiments were performed to increase the ratio between transmission inside the spectral hole and attenuation outside it for such hole-burned crystals. This was done both by optimizing the pulse shape of the hole-burning pulses and by burning and probing the structure with different polarizations of light. As light polarized along the different axes of the biaxial crystal have different absorption coefficients, the response to the burning pulses varies with polarization.

The simulations showed that burning the hole with a weakly absorbed polarization and probing it with a strongly absorbed one produces sharper and more uniform spectral holes than burning and probing with the same polarization does. A set of optimal pulse parameters were also found. In the experiments, an attenuation of approximately 53 dB was found when comparing filtered light at the centre frequency of the spectral hole to filtered light with a 2 MHz higher frequency. The results corresponded to an absorption coefficient outside the spectral hole of approximately 20 cm^{-1} for both polarizations tested, falling short of the given literature value of 47 cm^{-1}. A number of potential reasons are given, as well as several suggestions for future improvements. Still, the attenuation achieved improved upon previous results as presented by Zhang *et al.*[1] (2012), which was the main objective of the thesis.

Acknowledgements

I wish to thank my supervisor, Assistant Professor Lars Rippe, for the opportunity to take on this project, which turned out to be quite a learning opportunity for me, and for general advice and answers.

I would also like to thank PhD students Adam Nilsson and Qian Li for their invaluable help during the experimental phase of this thesis work, and for answers and help with regards to the data analysis.

Furthermore, I wish to thank PhD student Axel Thuresson, whose code files and guidance in understanding them made my work so much easier than it could have been.

I also wish to give thanks to Professor Stefan Kröll for offering me this project and for offering advice during the process.

Lastly, I wish to thank my parents, Britt-Marie and Thomas Dalsbecker, for their unending support, love and guidance. It is thanks to you that I have made it this far.

Contents

Acronyms and abbreviations

UOT – Ultrasound-modulated optical tomography

AOM – Acousto-optic modulator

PMT – Photomultiplier tube

YSO – Yttrium orthosilicate, Y_2SiO_5

PEEK – Polyether ether ketone

FWHM – Full width at half maximum

Ref. – Reference

Eq. – Equation

Fig. – Figure

At.% – atomic percent, i.e., percent of the total number of atoms

Introduction

This thesis deals with the creation of spectral filters, intended for future medical use, through a technique called spectral hole-burning. Spectral hole-burning has numerous applications (see, for example, Refs. 2-3), but its usefulness in medicine has only fairly recently come to light[4]. The medical application in question, ultrasound-modulated optical tomography (UOT), has been researched for some time[5], but the methods vary and not all rely on spectrally hole-burned filters.

Common for all the UOT techniques is the combined use of ultrasound and laser light in order to probe tissue or other biological matter for irregularities[6]. The particular method considered here is mostly the same as that described by Zhang *et al.*[1] Laser light and ultrasound pulses are made to collide inside a certain region of the tissue, Doppler shifting the frequency of the light passing that region. This makes it discernible from the light which did not pass through that region.

The technique is, among other things, intended to be used for cancer detection, but has not yet been implemented in medicine. So far, it is only being tested on tissue phantoms, substances with light and sound scattering properties similar to those of human tissue. The more the technique is improved, the more likely it becomes that the technique can at some point be used in medicine as a harmless, non-invasive form of cancer detection.

For this reason, this thesis work aims to improve one important aspect of such a UOT setup, in the type of detection used by Zhang *et al.*[1]: the spectral filter. It is meant to let through the light coming from the region of tissue being probed while blocking out the laser light which has been scattered through the surrounding tissue. This is vital in the type of setup described here, as it decides the detection limit and reliability of the measurements.

In this Master's thesis, the creation of such filters through spectral hole-burning in crystals of praseodymium-doped yttrium orthosilicate (Pr^{3+}:Y_2SiO_5) is investigated theoretically and experimentally. The theoretical part was performed through simulations, which were used to optimize the shape of the spectral holes, making holes with sharp edges in frequency space. Such holes allow for high transmission of the probing light and high absorption of the scattered light without requiring a large frequency difference between the two. This is an important property, as the frequency difference is decided by the ultrasound.

The simulations were also used to investigate whether the hole shape could be improved using different polarizations of light with respect to the crystal axes. With the simulation results known, experiments were performed in order to determine how good spectral filters that could really be made. Thereby, the reliability of the simulations could also be checked. Knowing the results of the experiments, a number of possible improvements of the procedure were suggested for future work with simulations and experiments on this kind of system.

This thesis work aims to improve upon the results from the similar setup as reported by Zhang *et al.*[1], where an attenuation of 30 dB was achieved. This corresponds to three orders of magnitude of difference in light transmission between the inside and outside of the spectral hole. The aim is fulfilled through the use and modification of the above mentioned simulations and through the experiments based on said simulations. The final objective of this thesis is to produce results which show an improvement of suppression compared to the results reported by Zhang *et al.* It must be clarified, however, that those results were obtained using scattered light rather than the collimated light used in this thesis work. The use of

collimated light makes a high contrast between the inside and outside of the spectral hole easier to achieve, which must be remembered when comparing the results.

All that being said, the thesis does not deal with, study or explain the mathematics involved in the simulations, nor with the creation of code to study the phenomena at hand. The code which is used to simulate the spectral hole-burning was supplied by Axel Thuresson[7] and has been modified as part of this thesis work. However, that is the only manipulation of code that has featured in this thesis work.

Furthermore, this thesis work does not involve any considerations regarding the usefulness of slow light in UOT using spectrally hole-burned filters, as brought up by Zhang *et al.*[1] The effect is simply acknowledged when observed, with a reference given for further reading. Nor does this thesis explain in detail the effects of apparent polarization rotation, which may arise when light passes through the type of biaxial crystal used in the experiments here. The effect is merely mentioned as it is used in calibration for the experiments.

Lastly, no calculations are performed with regards to the attenuation necessary for a viable UOT setup. The thesis aims merely to improve the performance of the filters, under the assumption that 30 dB was not enough, seeing as UOT has yet to be implemented in the medical discipline.

All that being said, the thesis does provide the necessary theoretical background needed in order to understand the concepts discussed within it. It covers to some extent such phenomena as spectral broadening and coherent light-matter interaction as well as a more detailed explanation of the UOT technique. Hopefully, the reader will find this to be sufficient in order to understand the thesis as a whole.

Theory

This section is intended to provide the necessary knowledge of the relevant theory needed to understand the main body of the thesis. Here follow the relevant features of the crystals and the dopant ions, the basics of persistent spectral hole-burning and coherent light-matter interaction, the Maxwell-Bloch equations and a brief summary of the UOT technique.

Spectral broadening

The following section draws most of its information from the book Fundamentals of Photonics[8], which is a very good source of further information for the interested reader.

In dealing with light-matter interaction, a quite important effect to consider is that of spectral broadening. The mechanism of spectral broadening in general is a consequence of several underlying physical principles, which affect the absorption profiles of atoms.

Several processes are known to affect the width of these absorption lines. Lifetime broadening, for example, in which the lifetimes and decay rates for transitions from one level to several others influence the uncertainty in the energies involved in these transitions. This, in turn, widens the absorption profiles of the atoms. Collision broadening, on the other hand, results from energy transfer and phase shifts that take place when atoms collide. In a homogeneous ensemble of atoms, these broadening effects affect all atoms equally, hence the term homogeneous broadening. Their absorption profiles can simply be added together to form a homogeneous absorption profile for the ensemble. However, there are other effects at play as well, affecting different atoms in different ways.

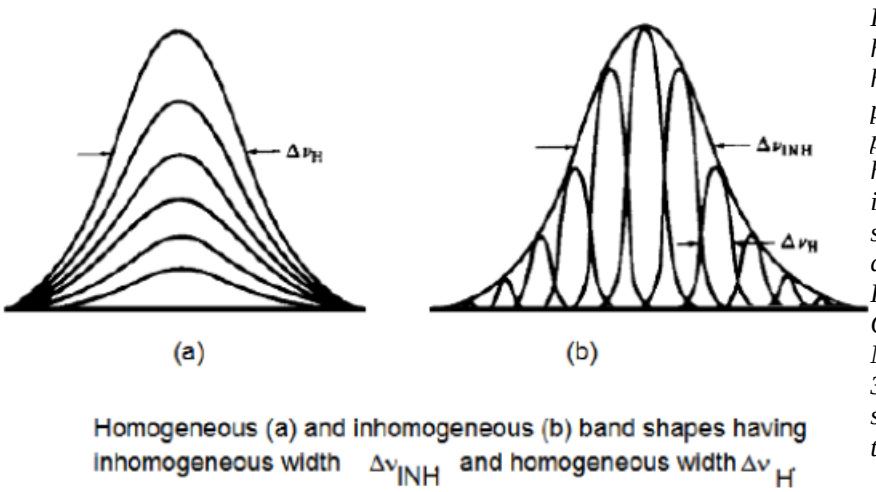

Figure 1: Examples of homogeneous (a) and inhomogeneous (b) spectral profiles, where in (b) the place and role of the atoms' homogeneous profiles in the inhomogeneous one is shown. Picture taken directly from Simons & Hayashi[9] under the Creative Commons Attribution-Noncommercial-Share Alike 3.0 United States License; see Ref. 26 for a summary thereof.

One example of this is Doppler broadening. As the atoms vibrate and move around with a certain kinetic energy, they will in general be moving in different directions at different times. The velocity of an atom relative to an incoming photon will determine which frequency the photon has in the atom's reference frame due to the relativistic Doppler effect. As such, an observer at rest in the atom's reference frame would see a different frequency photon than an observer at rest with the laboratory reference frame. Since the atoms generally move around in all different directions, some experience the photon frequency as lower while others experience a higher frequency than the one that is supplied by the light source. Thus, the sum of all the atoms' individual absorption profiles is broadened, as several of the individual profiles have been shifted in frequency. Since the effect differs between atoms in a homogeneous ensemble, this type of process is called inhomogeneous broadening; see Fig. 1 for an illustration of homogeneous and inhomogeneous absorption profiles.

In the crystals dealt with in this thesis, described in greater detail below, the atoms in question are doped in ionic form into a transparent crystal, which is cooled down to approximately 2 K or -271 °C. Thereby, very few collision and thermal effects are expected, but there is still a significant linewidth broadening. This stems from the interactions of the dopant ions and the crystal lattice, as the lattice field is not entirely homogeneous. Therefore, the absorption frequencies of the different ions become spread out. By summing together all these profiles, the inhomogeneously broadened shape takes form. There, each frequency within the profile corresponds to the absorption frequency of a given set of the ions, but not all of them. This is important in the context of spectral hole-burning, which is discussed below.

Spectral hole-burning

The experiments were performed using yttrium orthosilicate crystals (Y_2SiO_5, abbreviated YSO) doped with praseodymium ions (Pr^{3+}) at the dopant concentration of 0.05 at.% relative to the number of yttrium ions. The crystals keep the locations of the ions fixed in space, forming a cluster where the ions can interact with the incident light[10].

The interaction of the ions with the crystal lattice atoms as well as with each other cause different degrees of energy level splitting. The Stark splitting[10] is of the order of THz, with the hyperfine level separation[12] being a few to tens of MHz. The Stark levels are too far apart to be of any concern to the work performed as part of this thesis work, but the hyperfine levels play a very important role. Fig. 2 shows a schematic sketch of the part of the energy level system relevant to the work described here. While the level spacing between the hyperfine states of each level is generally constant, inhomogeneous broadening of the absorption profile causes the spacing between the upper and lower levels to be different for different absorbers. For one absorber, a given frequency might correspond to a transition from the ±3/2 level of the ground state to the ±1/2 level of the excited one, whereas for another absorber the same frequency instead corresponds to the energy gap between the ±5/2 levels of the two states.

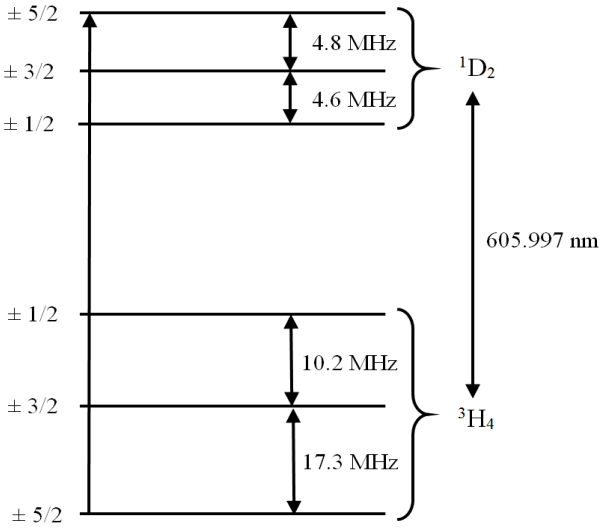

Figure 2: A schematic, not-to-scale sketch of the energy levels relevant to the work described in this thesis. $^3H_4$ is the ground state and $^1D_2$ is an excited state. Different hyperfine levels of the two states matter for different absorbers, due to inhomogeneous broadening. Adapted with permission, see References.

As the undoped form of the crystal is transparent to a range of wavelengths around 605.997 nm (494.725 THz), light in that range can be used to pump the atomic population of absorbers from the ground state $^3H_4$ to the excited state $^1D_2$, having the lifetime of 164±5 µs according to Equall *et al.*[11]

From the excited state levels, the atomic population is allowed to decay into either of the hyperfine levels of the ground state. The alternative ground state levels for a given absorber collectively form a reservoir of atomic population, where the atoms have a certain probability of ending up once they decay from the excited state. The hyperfine levels of the ground state have a relaxation time of approximately[13] 100 s for the praseodymium dopant. If the level to which it then relaxes not is the one it was excited from, its absorption frequency changes and it no longer absorbs the radiation that excited it the first time. This is the beginning of a

spectral hole, which deepens as more absorbers become unable to absorb a certain frequency or range of frequencies. The hole lasts until the absorbers have all relaxed back to the levels from which they originated. An example is shown in Fig. 3, where the inhomogeneous absorption spectrum of a spectral hole is shown in a simplified, exaggerated manner after a spectral hole has been burned over a range of frequencies $\Delta\nu_{burn}$.

It bears remembering that, as seen in the figure, the atomic population which responds to the incident light is not the entire available population. Incident laser radiation will generally excite several different groups of absorbing ions, since the laser wavelength corresponds to transitions between different hyperfine levels for different ion groups. This tends to create multiple holes, because as the ions are removed from the ground state, they can no longer absorb at any of the frequencies that would normally excite them from the ground state. These auxiliary holes in the absorption spectrum are referred to as side-holes. It is important to be aware of where these side-holes appear, to avoid them letting through frequencies intended to be blocked.

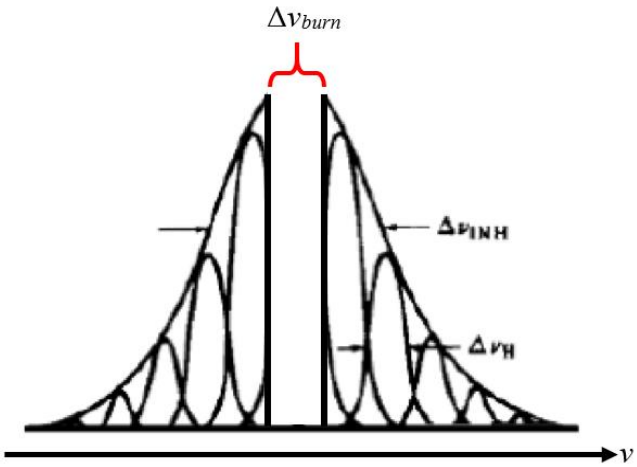

Figure 3: The inhomogeneous absorption profile as seen in Fig. 1, modified as though a spectral hole had been burned over the frequency interval $\Delta\nu_{burn}$. The original figure was taken directly from Simons & Hayashi[9] under the Creative Commons Attribution-Noncommercial-Share Alike 3.0 United States License, under which modifications were also performed, resulting in this figure. The reuse of this figure is permitted under that same license, summarized in Ref. 26.

Furthermore, as the ions are pumped away from the ground state, the states in which they end up can absorb light much more strongly as the population is increased there. This causes so-called "anti-holes" with increased absorption in the spectrum at the frequencies corresponding to transitions from the levels where the population is thusly increased. Creating holes in one part of the spectrum thus also means creating anti-holes elsewhere, as the population is simply moved from some levels to others through the interaction with the laser light.

Coherent light-matter interaction

The Maxwell-Bloch Equations

Here follows a brief outline of the necessary features from coherent light-matter interaction necessary to understand the physics and the terminology in the sections to follow. The modelling of the light-matter interaction described in this thesis takes a semi-classical approach, where the absorbers are quantized but the field is classical, i.e., continuous. This approach allows the interaction to be described by a set of equations known collectively as the Maxwell-Bloch equations (Eqs. 1-5). The part based on the Bloch equations treats the behaviour of the atomic system quantum mechanically, and the Maxwell part describe the classical light field itself. The equations presented here were used in the simulations on the form as provided by Thuresson, from whom the Maxwell-Bloch simulation code itself was also received. Therefore, only the basic features are given here; a full derivation can be found in Thuresson's thesis[7]. The equations are:

$$\frac{dr_x(z',\tau,\Delta)}{d\tau} = -\Delta r_y(z',\tau,\Delta) - \Omega_i(z',\tau) r_z(z',\tau,\Delta) - \frac{r_x(z',\tau,\Delta)}{T_2} \quad (1)$$

$$\frac{dr_y(z',\tau,\Delta)}{d\tau} = \Delta r_x(z',\tau,\Delta) + \Omega_r(z',\tau) r_z(z',\tau,\Delta) - \frac{r_y(z',\tau,\Delta)}{T_2} \quad (2)$$

$$\frac{dr_z(z',\tau,\Delta)}{d\tau} = \Omega_i(z',\tau) r_x(z',\tau,\Delta) - \Omega_r(z',\tau) r_y(z',\tau,\Delta) - \frac{1+r_z(z',\tau,\Delta)}{T_1} \quad (3)$$

$$\frac{d\Omega_r(z',\tau)}{dz'} = \frac{\alpha_0}{2\pi} \int_{-\infty}^{\infty} g(\Delta)\, r_y(z',\tau,\Delta) d\Delta \quad (4)$$

$$\frac{d\Omega_i(z',\tau)}{dz'} = -\frac{\alpha_0}{2\pi} \int_{-\infty}^{\infty} g(\Delta)\, r_x(z',\tau,\Delta) d\Delta \quad (5)$$

wherein, as in Thuresson's thesis:

- Δ represents the detuning, which is the difference between the pulse's centre frequency and the laser frequency currently being used, $\Delta = \omega_0 - \omega_L$..
- The Rabi frequency, relating to the strength of the interaction between light and matter, is defined as

  $$\Omega_{i\to j} = \frac{\vec{\mu}_{i\to j} \cdot \vec{E}(z,t)}{\hbar} e^{i\varphi(t)} \quad (6)$$

  where μ is the transition dipole moment, and indexes *i* and *j* denote the initial and final states, respectively. *E* is the electric field from which the complex oscillation due to the phase φ has been extracted. Ω_r and Ω_i are the real and imaginary components of the Rabi frequency, respectively.
- The parameters r_x, r_y and r_z are the components of the Bloch vector, which is explained below.
- The function $g(\Delta)$ represents the inhomogeneous absorption profile of the atomic system.
- The value of α_0 is a renormalization of the absorption coefficient α, dividing it with g for the values at the resonance frequency for both variables.
- T_1 is the lifetime of the excited state, and T_2 is the decoherence time. Their inverses are the rates of decay and loss of coherence, respectively.
- The retarded time is defined as $\tau(z,t) = t - zn/c$ with *t* as time, *n* as the refractive index and *c* as the light speed in vacuum. In that same variable substitution[7], the distance along the propagation axis, *z*, is replaced by the fully equivalent variable *z'*.

The above equations are integral in the code used to simulate the behaviour of the atomic system. The code solves the equations numerically using the fourth-order Runge-Kutta method to find how the interaction between light and matter affects the light pulse and the absorbers. To be able to understand the equations governing this interaction, the Bloch sphere is introduced as a fairly simple, yet descriptive, model.

The Bloch sphere

As the transition from ground state to excited state is driven by the incoming coherent laser radiation, the state of the absorber can be described as a coherent superposition of the two states. Both states have time-dependent oscillating probability amplitudes which, borrowing a bit of formalism from Nilsson[10], can be expressed as

$$|\psi\rangle = c_0(t)|0\rangle + c_1(t)e^{-i\omega_{01}t}|1\rangle \quad (7)$$

Here, c_0 and c_1 are the probability amplitudes of the respective states, 0 being the ground state and 1 the excited state. ω_{01} is the frequency corresponding to the energy difference between the two levels, expressed in rad/s. These probability amplitudes are normalized such that

$$|c_0(t)|^2 + |c_1(t)|^2 = 1 \quad (8)$$

As stated by Nilsson, the states can be described on this form – as pure states – if all interactions save for that between atom and light field are left out or accounted for in other ways, such as decoherence. In that case, the coefficients can be used to form the so-called Bloch vector, consisting of the real elements r_x, r_y and r_z as mentioned above and as seen in Fig. 4. The vector is made up of the coefficients according to the formulae

$$r_x(t) - ir_y(t) = 2 \cdot c_0 \cdot c_1^* \quad (9)$$

$$r_z = |c_1(t)|^2 - |c_0(t)|^2 \quad (10)$$

The first two elements describe the oscillation of the electric dipole moment and its two phase-dependent components r_x and r_y, 90° apart in phase from one another. The third component, r_z, is called inversion and describes the difference between the probabilities of finding the atom in the excited state (|1>) and finding it in the ground state (|0>). As seen above, these vector components play an important role in the Maxwell-Bloch equations.

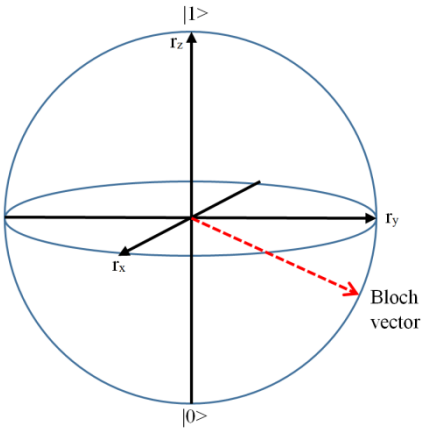

Figure 4: Simple schematic diagram showing the Bloch sphere and the eponymous vector at an arbitrary atomic state. Shown are also the vector components on the axes, with the two states in superposition (|0> and |1>) shown on the vertical axis.

This being said, the function and intricacies of the Maxwell-Bloch equations are not the main focus of this thesis, and so, there will be no further discussion regarding them here. Instead, this thesis work has relied on them as a tool to pursue improvement in spectral hole-burning with a specific application in mind, as explained in the following section.

Intended application

While spectral hole-burning has several possible areas of application[2-3], the intended application in this thesis work is a potential medical technique called Ultrasound-modulated Optical Tomography, or simply UOT. While there are more ways than one to perform the technique[6, 15], this section focuses on the one which relies on crystals which have been spectrally burned. The particular variation of the technique which is referred to here is the one described by Zhang *et al.*[1], illustrated in Fig. 5.

UOT is a potential technique for harmless probing deep into turbid media, such as biological tissue. One possible application of the technique could be to search living tissue for cancerous tumours. While the technique has not yet been tested *in vivo*, tests using fake tissue – so-called tissue phantoms – have been performed[1]. The UOT technique combines the use of light and ultrasound, which separately already have several biomedical applications. The benefit of this technique is the combination of the advantages with optical tomography, such as colour and texture sensitivity, with the superior spatial resolution of ultrasound techniques.

The basic operating principle of UOT is that a light pulse (green in the figure) and an ultrasound pulse (shown in blue) are sent into the tissue to be examined. Due to the large

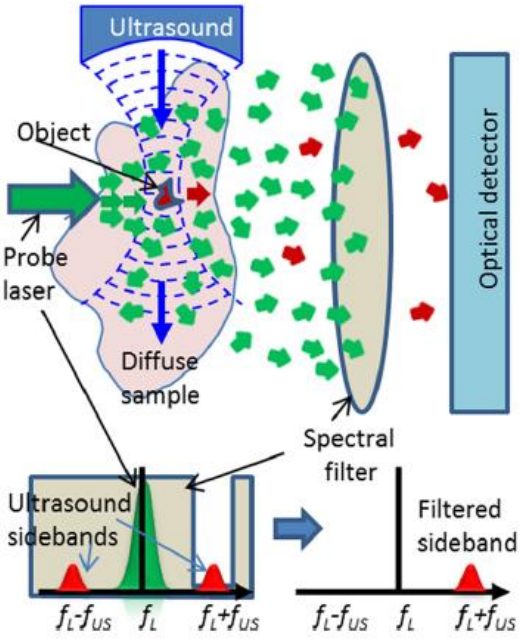

Figure 5: The working principles of UOT, as described in the text. The ultrasound pulse (blue) shifts the frequency of the probe laser light (green) where the two meet. The shifted light (red) then needs to be filtered out of the unshifted background light, which requires the use of a spectral filter. Reproduced with permission, see References.

difference in relative velocities of the sound and the light in the tissue, the sound pulse can be considered to be unmoving for the duration that the light is present in the tissue. As the ultrasound enters the tissue, the pulse causes relative shifts of pressure, density and refractive index of the tissue. Light passing through the region where the focus of the sound pulse is currently located experiences a shift in frequency due to the Doppler effect. This gives rise to so-called sidebands of light which are shifted up or down in frequency by an amount corresponding to the ultrasound frequency, as seen in the lower part of Fig. 5.

By gradually adjusting the intersection point and time delay between the sound and light pulses, the shifted light will originate from different points inside the tissue. However, as the light is strongly scattered by the tissue, the portion which interacts with the tightly focussed ultrasound pulse is quite small, and therefore, the shifted light is all but drowned out by the unshifted light. This is where the techniques diverge, offering different ways to detect the shifted light without it being obscured by the far stronger unshifted light.

The method considered here lets the output light pass through a spectrally hole-burned crystal, with a spectral hole corresponding to the frequency of one sideband. Thus, the main peak – the unshifted light – is blocked by the filter, as seen in Fig. 5. Since the frequency of the ultrasound used in UOT generally[6] lies in the range of 1 to 5 MHz, it becomes important that over less than 5 MHz, the spectral filter can go from maximal transmission to maximal attenuation of light. This is why this master's thesis work has focussed on burning spectral holes in such a way that their edges are sharp enough to meet this requirement.

Summary of the theory

- The absorption profiles of the individual ions are shifted and spread out in frequency due to the inhomogeneous lattice field. This broadens the total absorption profile.
- The ions' hyperfine levels are a few to tens of MHz apart, and have a relaxation time of about 100 s. By exciting these ions and letting them decay, they might end up in different hyperfine states. From these, they will not be excited by the same frequency of light anymore. This leaves a spectral hole at the old absorption frequency.
- The interaction between light and ions is here modelled using the Maxwell-Bloch equations, Eqs. 1-5.
- Crystals doped with these ions can, when hole-burned, function as very narrow-bandwidth filters. One application for such filters is ultrasound-modulated optical tomography, in which laser light and ultrasound are used to probe biological tissue.

The idea

While various spectral hole-burning experiments have been performed[16] since the mid-1970s, this particular thesis work deals specifically with the benefits of using light of different polarizations for burning and probing the spectral holes. The crystals used in this thesis work are anisotropic, having different optical properties depending on the direction and polarization of the incoming light. The crystals are biaxial, having three different orthogonal axes denoted[17] D1, D2 and b, as seen in Fig. 6. In this thesis work, only light which is incident along one of the axes b and D1 has been considered. When light impinges upon the crystal along one axis, the absorption of the light depends on its polarization relative to the other two axes, or rather, its polarization relative to the transition dipole moment. The orientation of this dipole moment vector is shown in the (D1, D2) plane in Fig. 7. Furthermore, the dipole moment vector is orthogonal to the b axis.

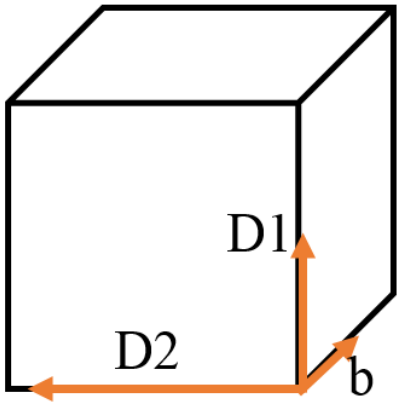

Figure 6: A simple depiction of the crystal and its three orthogonal axes, denoted D1, D2 and b. The crystal used in the experiments also has the square cuboid form seen here, where the D1 and D2 axes span a square surface.

As the polarization of incoming light is parallel to the electric field of the light wave, the relevance of the angle appears in the Rabi frequency of the transition, where (omitting the phase-dependent part of eq. 6):

$$\Omega_{i\to j} = \frac{\vec{\mu}_{i\to j}\cdot\vec{E}_0}{\hbar} = \frac{|\vec{\mu}_{i\to j}|\cdot|\vec{E}_0|\cdot cos(\theta)}{\hbar} \qquad (11)$$

Here, $\Omega_{i\to j}$ represents the Rabi frequency amplitude for the transition, E_0 is the electric field amplitude and $\hbar$ is the reduced Planck's constant, $1.055\cdot10^{-34}$ Js. The cosine function in the scalar product between the electric field and the transition dipole moment uses the angle θ between the electric field components and the dipole moment vector, as shown in Fig. 7 for an electric field along the D2 axis.

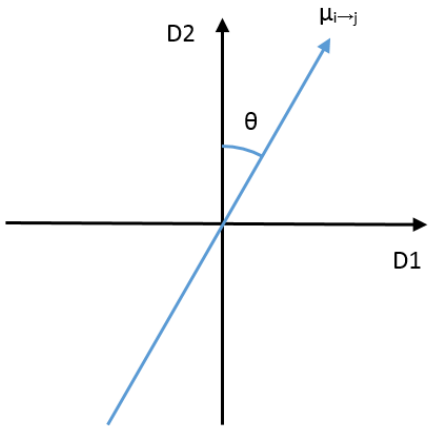

Figure 7: A schematic diagram of the location of the two axes of the D1, D2 plane in relation to the transition dipole moment for the transition driven by the absorption of light. The angle θ is measured relative to $\mu_{i\to j}$, and the value is positive for angles going clockwise from $\mu_{i\to j}$.

In the model used as basis for the simulations, it is assumed that light is sent in with a polarization perfectly parallel to one of the axes, as was done in the experiments. Light along the D2 axis has a much higher Rabi frequency, and thus interacts and is absorbed much more strongly by the crystal. While Fig. 7 is not drawn to represent it, the angle between the D2 axis and $\mu_{i\to j}$ is (-)15.5°, making the angle between D1 and $\mu_{i\to j}$ 74.5° as the two axes are orthogonal. If light is incident with a polarization which has components in both the D1 and D2 directions of the crystal, birefringent effects will cause a turning of the polarization. This was used in the experimental setup to allow for determination of the crystal axes, as only light which is polarized along these axes is stable with respect to the birefringence.

The idea behind this Master's thesis, suggested by Rippe[18], is to create spectral holes with the weakly absorbed D1 polarization, and then probe them as filters using the more highly

absorbed D2 polarization. Using the low-absorption polarization, the holes are more evenly burned, reducing the broadening of the hole in the front part of the crystal.

When a light pulse that interacts strongly with the absorbers reaches the crystal, it affects primarily the absorbers in the front end of the crystal until these have been mostly pumped away. The next pulse does the same at a slightly deeper point in the crystal, and so it goes on. This means that when the absorbers at the end of the crystal have burned a spectral hole that has the appropriate shape, most of the light pulses will already have passed through the hole burnt through the rest of the crystal. However, the absorbers at the frequency edges of the spectral hole have absorption profiles which, due to broadening, reach into the spectral hole. Given enough pulses which are not absorbed inside the spectral hole, these absorbers are pumped away into their respective auxiliary states, broadening the hole at the front end. This results in an absorption profile with edges that are not very steep, and thus not very useful for UOT. Fig. 8 contains an illustration of this effect, although the crystal is not quite burned through in that figure.

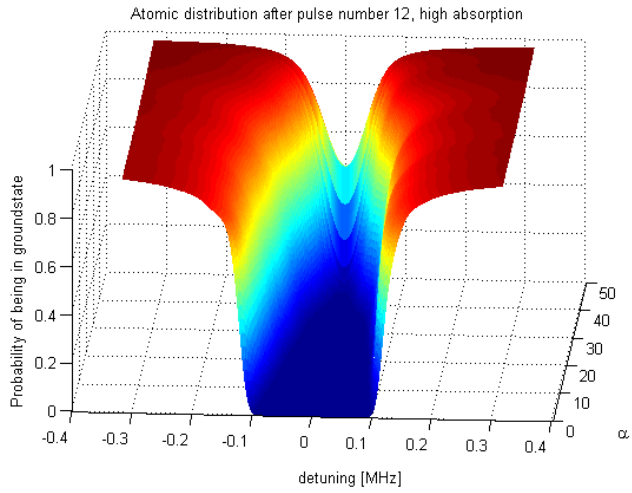

Figure 8: An illustration of the problem with burning using the D2 polarization. The horizontal axis shows the detuning, which translates as the frequency with respect to the central frequency of the spectral hole. The vertical axis shows the probability of the absorbers being in the ground state and thus able to absorb light as a function of the frequency (or detuning) of the light. The z-axis shows the depth into the crystal, illustrating how the spectral hole shape changes over the crystal's length. As the light is just starting to burn through the back of the crystal, it has already broadened the front end quite significantly. This makes such a spectral hole unsuitable for UOT, as explained in the text above.

Using the less strongly absorbed polarization (D1), an incoming light pulse instead pumps away a small portion of absorbers over a greater depth of the crystal. Thus, when the hole has been created, roughly the same amount of light will have passed each point along the light's path, and the shape of the spectral hole is much more even. As the light constantly has absorbers to interact with along the full length of the crystal until the hole has been burned, the risk of absorption by absorbers outside the hole is considerably smaller, resulting in a sharper absorption profile.

However, if the hole is probed using the D1 polarization, the suppression of light outside the spectral hole becomes too weak. Therefore, the hole is burned using D1 light and probed with D2 light. As the atomic population in the spectral hole has already been pumped away, the transmission of light through the hole is unaffected by the change of polarization, but the absorption is far greater outside the hole for D2 light. Using this idea, a series of simulations and experiments were performed, which are described in the appropriate sections below.

Summary of the idea

- The crystal used in this thesis has three different axes with different probabilities of interacting with incoming light.
- The polarization of light with respect to these axes decides the absorption coefficient.
- This thesis explores the importance of the polarization in spectral hole-burning.

Simulations

The atomic system

The idea outlined above was the foundation for the simulations performed in this thesis work. The simulations were performed in MATLAB, using a code originally developed by Axel Thuresson as part of his thesis[7], which was modified as part of this thesis work. The code simulates a two-level system, where after a pulse has passed through the crystal, a set fraction of the excited absorbers are transferred to an auxiliary state from which they do not return. A simple illustration of this level system can be found in Fig. 9. The interaction between the absorbers and the incoming light pulse is governed by the Maxwell-Bloch equations as described in the Theory section, where the equations were solved numerically using the fourth order Runge-Kutta method[7].

The code contains the necessary parameters to simulate the behaviour of the absorbers and their interaction with the incoming light field. It models both the properties of the absorbers (such as coherence time, lifetime and absorption coefficient) and those of the light pulses (such as spectral width, duration and amplitude). A low-absorption polarization setting was added in order to compare the resulting spectral holes with the two different effective absorption coefficients. In order to display what the spectral holes would look like when probed with the high-absorption polarization in both cases, the resulting graphs from low-absorption burning were all scaled up so that the apparent maximal optical depths were the same for both polarizations.

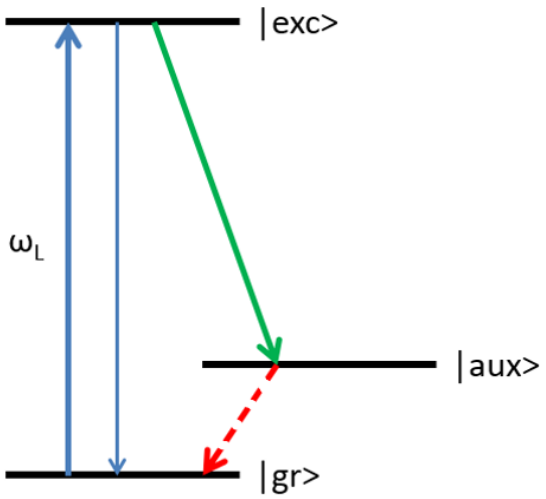

Figure 9: A representation of the energy levels as they are treated by the program. The excitation takes place from the ground state to the excited state, and after the light pulse has passed, 50% of the absorbers in the excited state return to the ground state and the rest disappear into the auxiliary state. The red, dashed arrow indicates the relaxation which would normally happen but which the program does not account for.

In addition to the plots which were pre-programmed in the code, a number of other relevant graphs were added. These display data such as the total optical depth experienced by the incident light pulse and the transmission of light through the crystal as functions of the detuning of the light.

Limitations of the program

It is important to mention one significant drawback to the simulation program. In order to achieve very deep or very wide spectral holes, the light pulses may sometimes need to be rather lengthy. The program, however, is not designed for such pulses for the purpose of spectral hole-burning. As stated above, and as illustrated in Fig. 9, the code simulates a two-level system, consisting of an excited state and a ground state, but also having an auxiliary reservoir state to which absorbers can be permanently transferred. The two-level code, though, does not consider this reservoir continuously as the light pulse traverses through the system; instead, it is only after a light pulse has passed through the system that the simulation transfers 50% of the absorbers which are currently excited to the reservoir state. For pulses lasting only a fairly short time compared to the lifetime, this has very little noticeable effect, but for pulses as long as or longer than the lifetime of the excited state (set to 164 μs) the decay from the excited state to the ground state becomes significant. Increasing the lifetime of

the excited state instead limits the decay to the ground state which would normally happen during the passage of the pulse for longer pulse durations. Thus, such a method is too uncertain, as depending on the new setting for the lifetime and on the pulse duration, it may give a result that is either better or worse than the real result would be.

Modifying the pulses

In order to achieve the highest efficiency of each given pulse, a specific pulse shape was used, namely that of a hyperbolic secant (*sech*). It was shown by Silver *et al*.[19] in 1985 that a magnetic pulse with the shape of a complex *sech* function was able to cause a highly selective, narrow-spectrum inversion of nuclear spins in a system. As the authors themselves say in their paper, nuclear magnetic resonance shares a number of interesting properties with coherent optics. It has since been seen[13] that the efficient inversion of a population by using that same particular pulse form is another such similarity.

A complex *sech* pulse[10] has the form

$$\Omega(t) = (sech[\beta(t - t_0)])^{1+i\mu} \tag{12}$$

where the exponent contains the product of a real number μ and the imaginary unit i. However, the imaginary part of the exponent on the *sech* expression is equivalent to introducing a time-dependent sweeping of the pulse frequency on the form

$$\Delta\nu(t) = \frac{\mu\beta}{2\pi} tanh\big(\beta(t - t_0)\big) \tag{13}$$

using the hyperbolic tangent (*tanh*) function. Summarizing this, the pulse behaviour is described as follows:

$$\Omega_{Complex}(t) = \Omega_0 \left(sech\big(\beta(t - t_0)\big)\right)^{1+i\mu} \Leftrightarrow \begin{cases} \Omega_{Real}(t) = \Omega_0 sech\big(\beta(t - t_0)\big) \\ \nu(t) = \nu_0 + \frac{\mu\beta}{2\pi} tanh\big(\beta(t - t_0)\big) \end{cases} \tag{14}$$

wherein:

- $\Omega_{Complex}$ and Ω_{Real} are the complex and real Rabi frequencies, respectively (both expressed in Hz)
- Ω_0 is the maximal amplitude of the Rabi frequency (also in Hz), which needs to be greater than $\mu\beta/2\pi$ for optimal transfer efficiency[13]
- μ is a real constant which needs to be greater than 2 for the inversion to be independent of the pulse amplitude[19]
- β is connected to the full width at half maximum (FWHM) of the pulse's intensity as

$$\beta \approx \frac{1.76}{t_{FWHM}} \tag{15}$$

- ν is the frequency (in Hz) of the light making up the pulse as a function of time t (in seconds) with respect to the time t_0, which corresponds to the midpoint of a pulse without an additional linear sweep component.
- ν_0 is the central frequency around which the *tanh* frequency sweep is centred.

The shape and frequency of such a pulse are shown in Figs. 10a and 10b, respectively. To create spectral holes with nearly maximal transmission inside the holes, very steep edges and as little transmission as possible outside the holes, the pulses were modified through a trial-and-error process. The parameters being modified were:

- The Rabi frequency amplitude Ω_0, measured in Hz, corresponding to the height of the peak in Fig. 10a.
- The FWHM of the pulse's intensity in time, t_{FWHM}, measured in seconds and scaling with the width of the pulse in Fig. 10a.
- The scan width f_{scan}, measured in Hz and representing the addition of a linear part at the top of the pulse in Fig. 10a and to the middle of the *tanh* shape in Fig. 10b.
- The width of the selected spectral region to be inverted, $\Delta\nu$, measured in Hz, corresponding to the frequency span of the graph in Fig. 10b.

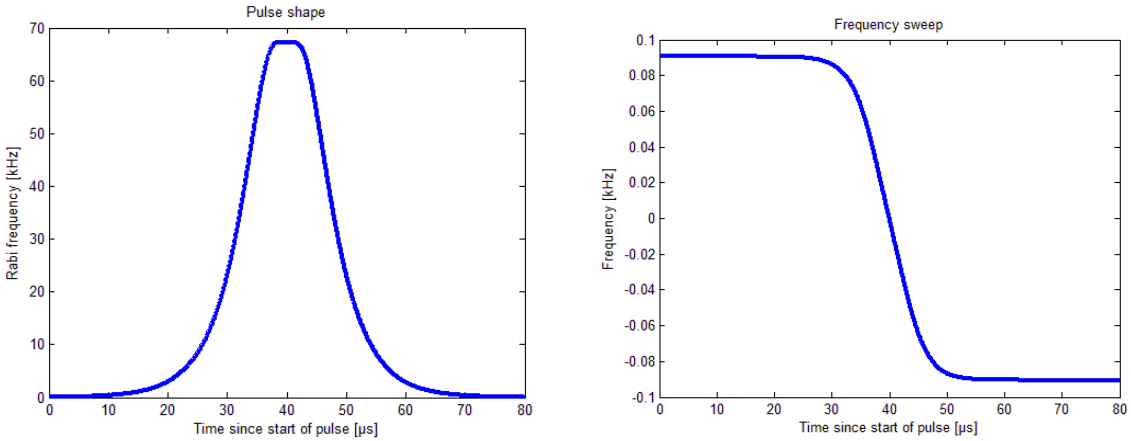

Figure 10: (a, left) The pulse shape of a modified hyperbolic secant pulse, and (b, right) its corresponding modified hyperbolic tangent frequency sweep. The modifications correspond to the addition of the f_{scan} component to both plots, seen as the flat top part of Fig. 10a and the straight middle part of Fig. 10b.

Since μ is constrained to a value greater than or equal to 2, it is useful to formulate this constraint in terms of the four parameters. As shown by Nilsson *et al.*[13], the spectral width of the population inversion region when using *sech* pulses is in total

$$\Delta\nu = \frac{2\mu\beta}{2\pi} \tag{16}$$

Using Eq. 15, it is possible to rewrite μ as a function of $\Delta\nu$ and t_{FWHM} as follows:

$$\mu = \frac{\pi \cdot t_{FWHM} \cdot \Delta\nu}{1.76} = t_{FWHM} \cdot \Delta\nu \cdot 1.785 \tag{17}$$

which together with the demand that $\mu \geq 2$ gives the constraint that

$$t_{FWHM} \cdot \Delta\nu \geq 1.12 \tag{18}$$

Furthermore, for optimal transfer efficiency, it is required for Ω_0 to be greater than $\mu\beta/2\pi$, which can be rewritten as

$$\Omega_0 \geq \frac{2\mu\beta}{2\pi} \cdot \frac{1}{2} = \frac{\Delta\nu}{2} \tag{19}$$

Apart from these constraints, theory does not limit the parameters any further, but care has to be taken nonetheless to make sure that the parameters do not exceed the physical limitations of the system which is to produce the pulses. Specifically, the maximal output Rabi frequency of the laser system was limited to a given value, as further described in the following section.

Performing the simulations

Two different experimental scenarios were envisioned when the simulations were performed. First, the case when the burn and probe pulses were sent to the crystal from the same direction was considered, in which case a spectral hole was burned with a small diameter, 0.1 mm. From an applications viewpoint, this could potentially be disadvantageous as it permits light to pass through the filter only at a small spot, meaning that less light can be filtered. As described above in the Intended application section, UOT already struggles with the fact that very little frequency-shifted light leaves the tissue. Thus, it is also worth considering an alternative.

The second selected approach was therefore to widen the beam diameter by a factor 100 from the original 0.1 mm to 10 mm, which would permit more light to pass through the spectral hole. In return, though, it would put a far narrower constraint on the attainable intensity and thus limit the maximal Rabi frequency, Ω_0. Since[13] $I \propto E^2 \propto \Omega^2$($I$ being the intensity of the light) and since $I \propto 1/r^2$ for a circular spot of radius r, it follows that $E \propto \Omega \propto 1/r$. For the second set of simulations, the pulse parameters thus had to be optimized separately, as the maximal Rabi frequency was constrained to one hundredth of the maximal value for the first set.

The wider beam enables the full width of the crystal to be hit by the burning beam, allowing the whole crystal to function as a spectral filter. This makes the use of the crystal in UOT more practical, although burning the spectral hole takes many more pulses. This stems from the fact that the criterion $\Omega_0 \geq \Delta\nu/2$ needs to be violated as the Rabi frequency amplitude is decreased. Furthermore, after burning the full volume of the crystal, probing can be done both along the line of the burning beam and orthogonal to it.

For both situations, the burning simulations were essentially limited by the maximal output power of the laser. For the small-radius burning with the high-absorption polarization, this was given as corresponding to a maximal Rabi frequency of 1.5 MHz. Thereby, the 100 times wider radius beam of the same polarization was limited to 15 kHz. The low-absorption polarization's upper limit was then calculated from the corresponding high-absorption value by the following means:

It is known that the maximal intensity is the same whichever polarization is incident for a given spot size, as the only difference is the angle of the electric field to the transition dipole moment vector. This also means that the electric field, proportional to the square root of the intensity, is of the same absolute value in both cases. As the case where $E_1 = -E_2$ bears no physical significance here, E_1 is set as equal to E_2. Thus, as $E = \frac{\Omega}{\mu_{i\to j} cos\theta}$, it is clear that

$$\frac{\Omega_1}{\mu_{i\to j} cos\theta_1} = \frac{\Omega_2}{\mu_{i\to j} cos\theta_2} \Leftrightarrow \frac{\Omega_1}{\Omega_2} = \frac{cos\theta_1}{cos\theta_2} \tag{20}$$

From this, knowing the two angles and one Rabi frequency, it was found that the maximal Rabi frequency for the low-absorption polarization is 416 kHz, with 4.16 kHz as the maximum for the wider beam.

For both radius categories, numerous simulations were performed in order to find the optimal settings for burning sharp spectral holes. To find out whether the D1 burning really does produce sharper spectral holes, the same optimizations were also performed for the use of D2-polarized light for burning. The results of the optimizations are found in Table 1 below.

Summary of the simulation procedure

- Using MATLAB code[7] to solve the Maxwell-Bloch equations numerically, the ions were modelled as a two-level system as they interacted with incoming light.
- In order to burn efficient and sharp spectral holes, the pulse form of the incoming light was chosen to be that of a hyperbolic secant. The pulses were also frequency-swept following the shape of a hyperbolic tangent; see Eq. 14.
- In order to optimize the sharpness of the spectral hole, the burning pulses were altered until, through trial and error, the optimal pulse parameters were found.
- The parameters being adjusted were the spectral and temporal widths of the pulses, the Rabi frequency (see Eq. 11) and a linear widening parameter.
- Both a narrow (0.1 mm) and a wide (10 mm) diameter beam were simulated. Furthermore, for each beam size, polarizations along both the D1 and the D2 axis were tested.

Results of the simulations

Here follow the results of the simulations which formed the basis of the burning pulses used in the experimental part, listed in Table 1.

One significant point that stands out quite clearly in the table is the correlation between the radius category and the necessary number of pulses. As the wider radius constrained the possible Rabi amplitude, this had to be compensated for by increasing the number of pulses by quite a bit. The simulations showed that this kind of compensation was indeed a viable strategy, as will be shown below.

Polarization	Diameter category	Rabi amplitude (kHz)	FWHM of pulse in time (μs)	Frequency sweep width (3) (kHz)	Inversion region width (kHz)	Number of pulses needed
High-absorption (D2)	Small	100	8.415	50	182	20
	Wide	1	16.83 (1)	50	91 (2)	>30000 (4)
Low-absorption (D1)	Small	67	8.415	50	133	10
	Wide	0.67	16.83 (1)	50	66.5 (2)	10000

Table 1: The optimal parameters used to create spectral holes in the simulations. The footnotes are explained in the text below. The small diameter category means that the beam diameter was 0.1 mm, whereas the wide beam had a diameter of 10 mm.

The limitations of the program also came into play in a quite important way in the optimization process. By increasing the values marked (1) in the table while decreasing those labelled (2) by the same factor, a more efficient burning could be achieved, requiring fewer pulses while sacrificing some of the hole's width. However, doing so would extend the pulse duration beyond what the program can reliably handle, as discussed above in Limitations of the program. This limitation would not exist in the experimental case, however, allowing for this modification to be implemented if needed.

The frequency width, labelled (3), was kept at a lower value for most of the simulations, but was raised to 50 kHz to be closer to the experimental case. In the actual experiments, an even higher value was used to make sure that the spectral width of the hole was greater than that of the light pulse used to probe it.

Lastly, the label (4) on the pulse number indicates the highest number of pulses used in simulations of the given kind. For the wide-beam D2 burning, 30000 pulses were insufficient in producing enough transmission. Due to time constraints, no lengthier simulations were performed. The remaining parameters were determined using a thinner crystal, reducing the simulation time. Therefore, their accuracy is less certain than that of the other settings.

Below are the simulated optimal results for the four different incident beam categories, corresponding to the four rows of Table 1, displayed graphically in a few different forms. Irrespective of which polarization the holes were burned with, the results have been scaled to match the shape they would have when probed using the D2 polarization.

The graphs supplied by the code display the data in a number of different ways. In order to see the overall shape of the spectral hole, the atomic population after a given number of pulses was plotted against the depth into the crystal and the detuning. An example of this, using the pulses optimized for the narrow beam and the D2 polarization, is shown in Fig. 11a. To further clarify the difference in shape of the hole profile over the depth of the crystal, five "slices" were cut out of Fig. 11a at five different depths. These are found in Fig. 11b.

Both types of graphs were primarily used in the process of optimizing the shapes and sharpness of the burned holes, as they show the way that the structure changes as a function of the depth into the crystal. As can be seen in both figures, it tends to narrow as the depth increases, showing the absorption in the sides of the hole at lower depths as discussed above.

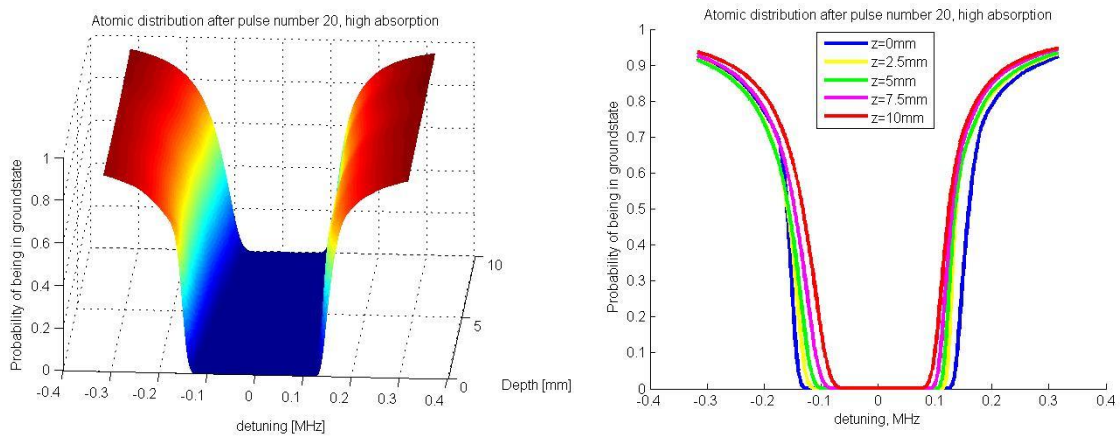

Figure 11: The optimal hole shape for the narrow-beam (0.1 mm) high-absorption (D2) burning, shown as though probed using D2 light. In both, the horizontal axis shows the detuning, the difference between the pulse's centre frequency and the sweeping laser frequency. The vertical axis shows the probability of finding an absorber in the ground state and thus able to absorb light at that detuning. (a, left) The relationship between detuning and probability of being in the ground state as it changes over the crystal depth, shown on the z axis of the graph. (b, right) The same relationship as in Fig. 11a, but cut out at five specific depths in the 10 mm crystal.

However, the plots in Fig. 11 show fairly little of how the crystal actually interacts with the incoming light. For that reason, the effective absorption was extracted from the data. For every possible "slice" of the data in Fig. 11a, the value on the probability axis was multiplied with the absorption factor per slice. All such slices were then summed together to yield the total absorption, or *integrated αL*. Here, α is the absorption coefficient in cm^{-1} and L is the crystal length in cm. This kind of graph shows the absorption experienced by D2-polarized light as it enters the crystal after the structure has been burned. Such a data set of absorption values was collected for each of the optimized cases, both the wide and narrow beam for both of the polarizations D1 and D2, and plotted together in Fig. 12 as though probed with D2 light.

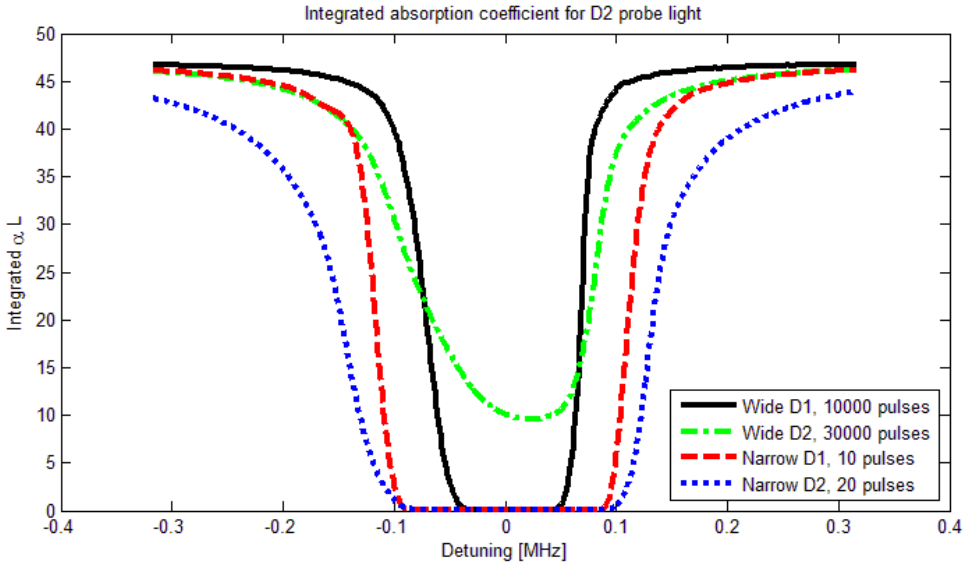

Figure 12: The integrated absorption profile for all four optimized burning pulses. This graph shows the absorption experienced by a pulse of D2-polarized light as a function of its detuning. The four lines correspond to the type of pulse (diameter type, pulse number and polarization) used to burn the hole, as shown in the legend.

The 30000 pulses used when burning with the wide D2 beam were unable to burn the structure all the way through and down, which is why its curve never reaches zero. The other three all manage to deplete the population in the centre of the hole, bringing the effective absorption to zero, but the difference lies in the sharpness of the edges. Specifically, the D2 narrow-beam burning comes up short in comparison to the D1 beams of both radii.

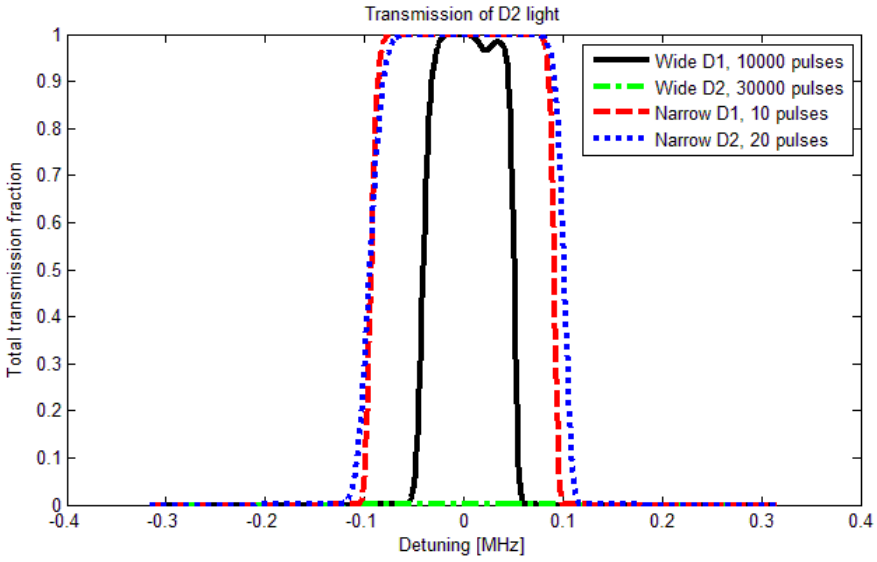

Figure 13: The fraction of the incoming probe light intensity that makes it through the burned structure, as a function of its detuning. The probe light is D2-polarized, and the four lines correspond to the four types of pulses used to burn the holes. The legend shows which line corresponds to which diameter, polarization and pulse number in order to create the transmission profile shown in the graph.

An even clearer representation of the crystal's usefulness as a filter can be found by looking at the transmission of light rather than the absorption. The Beer-Lambert law gives the intensity transmission through a lossy medium:

$$I = I_0 e^{-\alpha L} \Leftrightarrow \frac{I}{I_0} = e^{-\alpha L} \tag{21}$$

Knowing the integrated absorption αL, it is easy to find the ratio between the light intensity entering the crystal (I_0) and the intensity passing through it (I). From the data in Fig. 12, this information was extracted and plotted in Fig. 13.

The transmission fraction alone shows whether the crystals have the wanted combination of high transmission within the hole and low transmission outside it. Looking at Fig. 13, it is clear that all except for the wide-beam D2 burning pulses achieve more or less perfect transmission inside the spectral hole. However, the scale of the graph makes it difficult to tell what degree of suppression the hole profiles offer outside the region that has been burned down. Seeing as an attenuation of unwanted light of a factor greater than 10^3 is desired, it is important that this kind of information can be read out from the graph. For this reason, the data in Fig. 13 was also plotted with the transmission in a base-10 logarithmic scale as Fig. 14, making it proportional to the integrated absorption shown in Fig. 12.

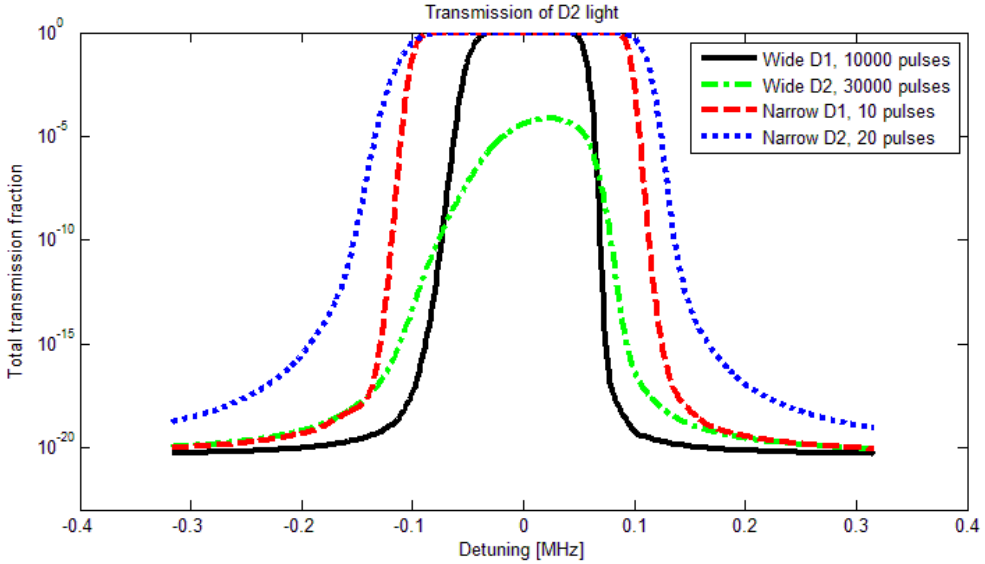

Figure 14: The transmission fraction of incoming D2 probe light through a hole burned using the settings listed in the legends. The data is the same as in Fig. 13, but the transmission fraction axis has now been made base-10 logarithmic in order to show the magnitude of the suppression outside the holes.

As can be seen in Fig. 14, the crystals offer a very high degree of suppression of light outside the spectral hole. The difference in sharpness is also quite clear, and it turns out that the two D1 burning regimes are those that fared the best. Even the narrow-beam D2 burning pulse form did manage to reach a suppression of about 18 orders of magnitude, however, which is quite promising for the filters' ability to surpass the previous experimental result of three orders of magnitude, as reported by Zhang *et al.*

Before moving on from the simulations, however, one practical consideration needs to be brought up. For the probe light to experience the filtering effect shown in the above graph, it has to have the right polarization. While this is easy to provide in an experimental setup, it would be more difficult in a clinical case, as the scattered light exiting the biological tissue would have lost its polarization purity due to scattering.

The solution[20] considered here is to let the emerging light pass through two hole-burned filters which are rotated 90° to each other. As light with any polarization can be divided into two

orthogonal components, one along D1 and one along D2, this means that the component which passes the first filter aligned with its D1 axis would then reach the second filter aligned with the D2 axis of that crystal, and vice versa.

To investigate the suppression and transmission of such a filter setup, further simulations were performed. The two filters were simulated separately, and their transmission fraction results were subsequently multiplied together. The transmission fraction of the two crossed filters is shown in Fig. 15 in the same logarithmic scale as Fig. 14. For clarity, it should be added that both 10 mm crystals used would first have been burned separately using the beam size and polarization shown in the legend, whereupon one would have been rotated 90° with respect to the other and placed right after it.

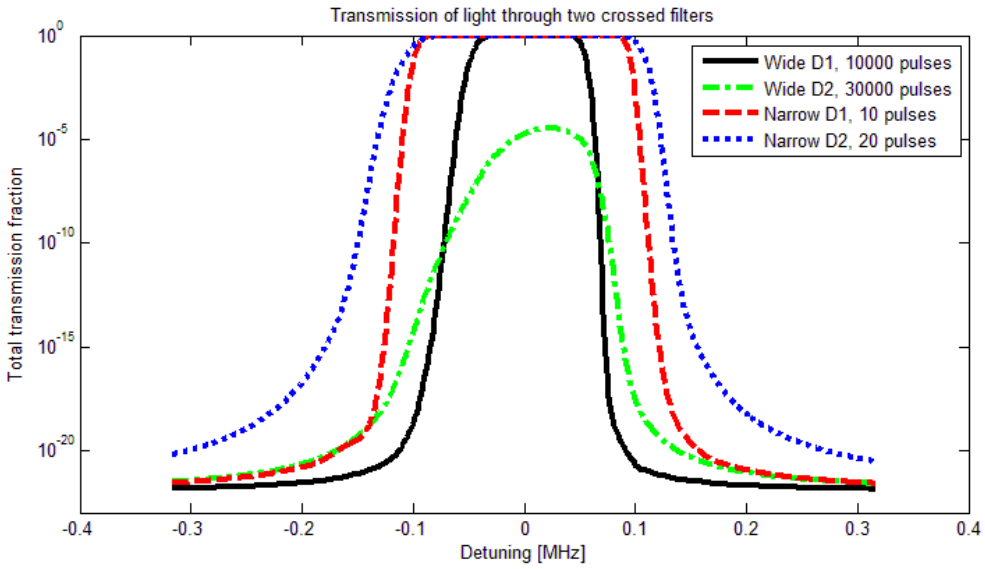

Figure 15: The transmission fraction of probe light of any polarization passing through two spectrally hole-burned filters, where one is rotated 90° with respect to the other. The legend shows which profile corresponds to which type of burning beam. The burning was simulated as performed separately on the two crystals before one was rotated and they were put together.

Fig. 15 shows that, with the two filters crossed, even the less-than-optimal narrow D2 beam can achieve 20 orders of magnitude of suppression. Still, the hypothesis that D1 burning should produce sharper edges holds true in the simulations, as both D1 beams produce sharper edges to their spectral holes than the narrow-beam D2 burning does.

Summary of the simulation results

- The optimal parameters for the different beam sizes and polarization are found in Table 1.
- According to the simulations, burning with the D1 polarization gives sharper spectral holes than the D2 one.
- Both polarizations manage to offer at least 18 orders of magnitude of suppression in the simulations.

This concludes the description of the simulations. They served mainly to find the optimal parameters for the spectral hole-burning and to estimate the quality of such a spectral filter. With that done, all that remained was to test the process experimentally.

Experiments

Experimental setup

Setup layout

This section describes the setup used to perform the experiments described in the Experiments section below. A schematic drawing of the setup is found in Fig. 16.

The laser light entering the system (shown in orange in the figure) was given off by a modified dye laser system[21, 22], giving a wavelength of 605.977 nm (494.725 THz). It entered the setup via an optic fibre at the point denoted 1. The beam was led through a zoom beam expander (2) and onto an Acousto-Optic Modulator (AOM, 4a) with the use of a mirror (3). When turned on, the AOM split the incoming beam into a horizontal diffraction pattern. As only two components were wanted (the zeroth and negative first order components), an iris (5a) was used to block out all other beams. For this reason, only two beams are shown leaving the AOM in the figure.

The negative first order beam was then, when the AOM was turned on, deflected onto a mirror followed by a beamsplitter (6a), which exhibited an 8% reflectivity for the incoming D1 light. The reflected component of the light was focussed using a lens (7) and a mirror onto a Thorlabs detector (15), which would serve as a reference signal. This would be used to compare the light intensity before the cryostat with the measurements after it, allowing for calibration, comparison and interpretation of results.

The component of the beam which was transmitted through the beamsplitter continued via another mirror onto a half-wave plate (8). The light had up to this point been very nearly vertically polarized, thanks to both the laser setup and the polarization-preserving optic fibre leading it to the setup. The half-wave plate was adjusted to turn the incoming polarization by 90°, making it nearly horizontal. The light was then led through a polarizer (9a), ensuring even further that the light entering the cryostat was as purely horizontal as possible.

The cryostat itself is denoted 10 in the figure, a cylindrical container containing a sample chamber that could be cooled down to roughly 2 K using liquid helium. If necessary, a magnetic field could also be applied to the sample chamber thanks to its superconducting magnet system. The cryostat had four windows, seen as smaller ellipses within the cryostat in the figure, through which light could enter and exit. The crystal (11) was lowered into the middle of the cryostat, suspended by a polyether ether ketone (PEEK) holder (see Fig. 17 and the appendix) attached to a metal rod, which was used to adjust the height and orientation of the crystal inside the cryostat.

After passing through the cryostat, the light was sent through another polarizer (9c), which was set either to the same or the opposite orientation compared to the previous one, depending on which polarization component was to be examined. Having passed this, the light reached another beamsplitter (6b) with 89% transmission for D1 light and 97% for D2, where the reflected part was focussed onto another Thorlabs detector (13) using a mirror (3) and a lens (7). The transmitted light was instead led towards a photomultiplier tube (PMT, 16), protected by a shutter. The PMT was used as a complement to the Thorlabs detector (13), as the two covered different intensity ranges.

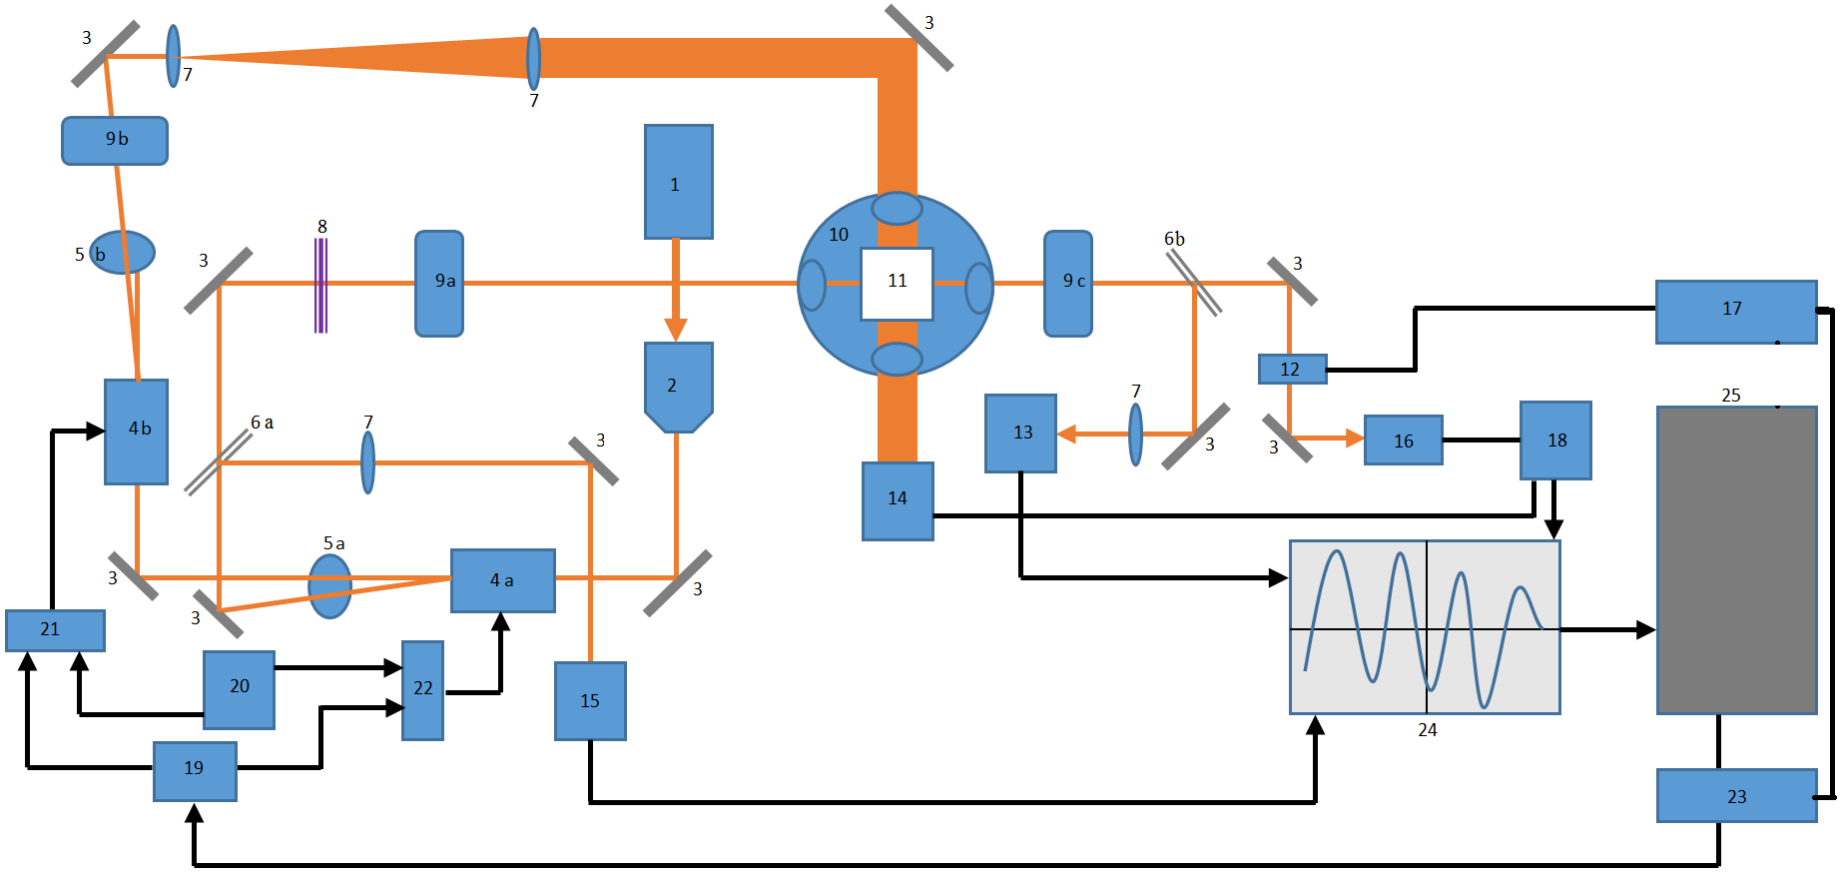

Figure 16: A schematic drawing of the experimental setup. The orange lines indicate the laser light, and the black lines represent electrical cables connecting the various components. 1) Optical fibre collimator, where the laser light enters. 2) Zoom beam expander. 3) Mirror. 4a,b) Acousto-optic modulators (AOMs). 5a,b) Irises. 6a,b) Beamsplitters. 7) Lens. 8) Half-wave plate. 9a-c) Polarizers. 10) Cryostat. 11) Crystal. 12) Shutter. 13) Absorption detector, Thorlabs. 14) Auxiliary detector. 15) Reference detector, Thorlabs. 16) Photomultiplier tube. 17) Shutter control box. 18) Transimpedance amplifier. 19) Switch. 20) Power supply. 21, 22) Radiofrequency amplifiers. 23) Arbitrary waveform generator. 24) Oscilloscope. 25) Computer. The components are described in more detail in the text.

The light from the zeroth-order beam of the first AOM (4a) which passed through the iris (5a) was reflected onto a second AOM (denoted 4b). There, the negative first order light was deflected by the same amount as was done by the first AOM. The iris that followed (5b) let through only that light, blocking all other orders. The deflected light was led through a polarizer (9b), ensuring that the polarization passing through was orthogonal to that of the other beam. Next, the light was led through a lens system (both lenses denoted 7), where the first lens (f = 20 mm) expanded the beam, magnifying it greatly, before the second lens (f = 750 mm) collimated the beam again. In total, the beam was expanded by a factor of 37.5 by the lenses, making it wide enough to hit the entire side of the crystal. After passing through the cryostat, the light from the expanded beam reached a detector (14).

The black lines in the figure denote the cables connecting the electronics. Each one of the two AOMs was connected to an amplifier (22 and 21, respectively). Both were connected to a power supply (20), and received their signal via a switch (19). The signal, determining which AOM was to be active and thus which one of the two paths the light should take, was decided through the computer system (25) which controlled the Arbitrary Waveform Generator (23) that directed the laser system.

The shutter (12) was controlled by a home-made control unit (17), controlled through the waveform generator so that the shutter's motion could be synchronized with the incoming light pulses. Two of the detectors, denoted 14 and 16, required amplification of their signals, and were therefore connected to a transimpedance amplifier (18). This amplifier, along with the reference detector (15), was connected to the oscilloscope setup (24). From this, the data was gathered in the computer system and saved for analysis.

Designing the crystal holder

For the crystal to be held in place inside a cryostat, a crystal holder needed to be designed and crafted. The crystal and its surrounding components were to be cooled using liquid helium, and the temperature difference of roughly 290 K risked exposing the components to stress due to thermal contraction. For this reason, the material for the holder was chosen to be polyether ether ketone, commonly abbreviated PEEK. For a temperature decrease from circa 300 K to 2 K, PEEK compresses only about 1.25% (calculated from a figure in Ref. 23).

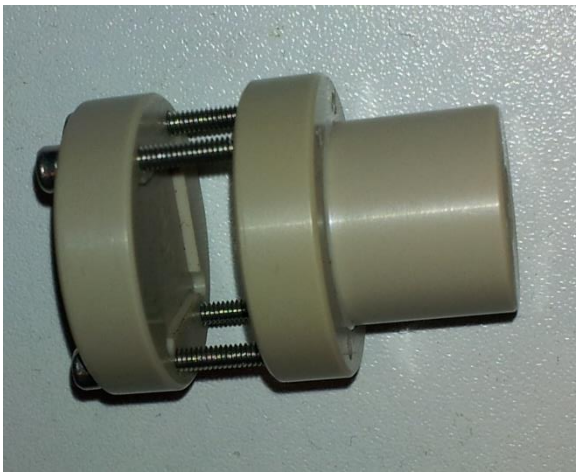

Figure 17: A photography of the assembled crystal holder. Not visible in the picture is the indentation on the bottom side of the upper component, which has the same square shape with rounded corners as the one seen on the bottom part. The appendix contains the SolidWorks drawings used to craft the holder.

Due to data regarding the thermal expansion coefficient of YSO being hard to find, the crystals were assumed not to shrink at all when cooled down. Therefore, sufficient margins were necessary so that the holder would not subject the crystal to stress when shrinking. For this reason, and to allow light to reach the crystal through any window of the cryostat, a design consisting of two parts was chosen, connected with four screws as shown in Fig. 17. The design of the two PEEK components is given in the appendix for further clarity.

The screws, too, needed to be chosen such that their thermal expansion coefficient ensured that they would shrink very little when cooled. For this reason, stainless steel 304 was chosen, exhibiting a contraction of less than 0.3% over the given temperature range[24]. Given the design, the only important requirement would therefore be that the screws were long enough that the crystal could be given some margin. The screws were ordered from Westfield Fasteners, Ltd., and the PEEK parts were crafted by BEMA Slip AB.

Testing the components

Before the setup was assembled, several components needed to be tested. Some were tested to see that they worked, others to check that they were good enough for the experiments, and still others to learn how to operate them.

The main concern about the experimental components was the purity of the polarization and the conservation of that purity. The purity would set the limit of the attainable ratio, and no conventional polarizer of today can offer a purity ratio of 18 orders of magnitude. A ratio of four to six orders of magnitude would be more realistic, and it would be relatively easy to tell whether the results were limited only by that purity. Leakage of unwanted polarizations into the burning pulses would be relatively insignificant, as that would only slightly change the hole profile. However, D1- or b-polarized light leaking through in the D2-polarized probe light would be a bigger problem, since the D1 polarization has a transmission of roughly 3% outside the hole according to simulations. Thus, conservation of the purity of the polarization was vital.

For this reason, a number of components in the setup needed testing in order to see what purity they could retain.

- The polarizers to be used (labelled 9a-c in Fig. 16) were tested first. Using a helium-neon laser, the transmission through the polarizers when aligned and when crossed was compared. The transmitted power differed by a factor of 10^5, meaning that the polarizers would decrease the power of a blocked polarization by that factor.

- The AOMs were also tested, as even when no voltage was applied, they scattered some light to the sides of the zero-order spot. As can be seen in Fig. 16, this would allow light from one of the two beams to impinge upon the crystal while the AOMs are set in such a way that light should only be reaching the crystal from the other beam. It was found that the scattered light in the direction of the negative first order had a power of $2.5 \cdot 10^{-5}$ times that of the zero-order beam.

    - While surprisingly good for an AOM, it was not entirely satisfactory for the experiment. However, simply choosing which order of which AOM was to be used as probe beam allowed for the potential leakage of light to be brought below one part in 10^5, which was the limit of the polarizers.

- Another concern for the polarization purity was raised regarding the cryostat windows[20]. Due to the temperature difference between the cryostat's inside and the room temperature on the outside, and/or due to the pressure difference that causes, the windows may affect the polarization of incoming light. This was checked by placing one polarizer in front of the cryostat and one behind it, comparing the maximal and minimal transmissions of light by turning the second polarizer.

- It was found that the polarization effects depend quite a lot upon whether or not the incoming light hits the windows at a 90° angle. The measurements proved quite fickle, and varied by fairly large factors from one day to the next. Ratios as high as 228000:1 were found when the cryostat was cooled down and the crystal removed, but with little to no changes to the setup, ratios that high were at other times unattainable.

- Nonetheless, it was found that a polarization purity slightly better than five orders of magnitude could be reliably achieved if the beam impinging on the windows was perpendicular to them. Perpendicular to the best possible extent, that is, as the windows were not entirely parallel with one another to begin with.

Two more things needed to be tested before the experiments could begin. The crystal to be used was cut as a square cuboid, having two square sides with 10 mm sides and four rectangular ones of length 6 mm. However, it was not certain that the three perpendicular crystal axes were in fact situated perpendicular to the faces of the crystal, nor which axis was which out of D1, D2 and b. Still, both questions could – in theory – be rather easily answered. Fig. 18 shows the crystal, photographed in the process of determining the axes.

* By shining the laser light onto the crystal with a perfectly vertical polarization, the crystal's effect on the polarization was checked. If the polarization of the incoming light were along one of the crystal axes, the birefringence would not cause the polarization to rotate. Using two crossed polarizers, it was checked whether the transmission of light through the second polarizer increased when the crystal was inserted, indicating a turned polarization.

* At the same time, the fluorescence given off by the crystal as the beam passed through it would reveal which axis was which. This would be visible as the fluorescence would be greater for light along an axis with a high absorption coefficient, as this light interacts more strongly with the absorbers. Light along the b axis would, as the crystal is perfectly transparent to it, cause no fluorescence at all. As the tests were performed at room temperature, no persistent spectral holes were created which could otherwise have altered the outcomes of both tests.

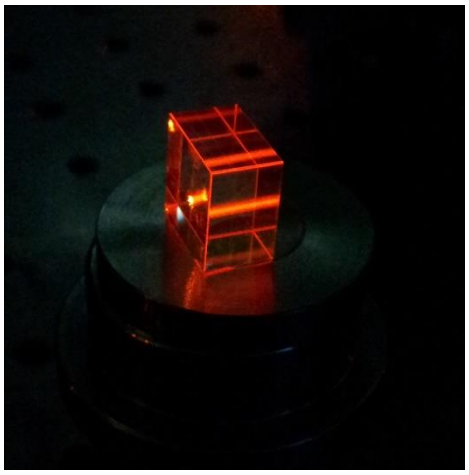

Figure 18: A photo of the crystal used in the experiments, as it was being tested to determine the orientation of the axes. The fluorescence due to the laser beam is clearly visible.

It was found, by turning the crystal around the propagation axis of the light, that the axes were not entirely perpendicular to the faces of the crystal. All three axes were found to be tilted a few degrees with respect to the light's propagation direction, respectively. Furthermore, being able to adjust the rotation in only one plane at a time, the visible fluorescence was insufficient to determine which axis was which out of D1 and b. These had to be distinguished later as a part of the experiment.

Performing the experiments

Having set up all the components according to Fig. 16, the components were aligned to make sure that the light entered the cryostat windows as close to perpendicularly as possible. In the end, achieving this perfectly proved impossible as the cryostat windows were not in

themselves perfectly parallel to one another. Further alignment problems arose as the crystal, carefully placed into its holder in such a way that it should be normal to the incoming light, became misaligned as the rod holding it in place was subject to thermal contraction, bending it slightly. This also resulted in the axis measurements previously performed being incorrect, and new measurements were performed with the crystal on location inside the cryostat.

Measurements of the polarization turning were repeatedly performed as described above. This method was later complemented by also studying the crystal's absorption profile. When the absorption of the horizontally polarized light was maximal, the polarization was parallel with the D2 axis. A similar calibration was also performed for the D1-polarized light in the other beam. It was also through measurements of the absorption that the crystal axes of the smaller crystal could unambiguously be told apart. The polarization of the incoming light as well as the crystal's orientation were adjusted in order to see the absorption along all three axes.

Scanning the laser frequency, the inhomogeneous absorption profile was mapped out and the absorption maximum was found. The apparent absorption coefficient was calculated for a number of frequencies in the profile and, transmission values were also taken offline, i.e., suitably far from the absorption profile. These values would later enable the calibration for the intensity losses in the system. Further calibration data was also collected by means of power measurements at various points in the system and the corresponding detector signals.

The experiments included burning along D2 using a narrow beam with a 1 mm diameter, as well as burning along D1 with a wider beam of about 10 mm diameter. The wide beam would hit the crystal on one side, burning a spectral hole through all of it, with the probe beam coming from another side. The narrow beam was used both for the D2 burning and as probe beam for all experiments.

Each run of the experiments consisted of the structure being burned with the appropriate beam and settings, whereupon the structure was probed with at the center and the edge of the hole. The measurement at the edge was done with a 2 MHz higher frequency than the measurement at the centre of the hole.

It was found, using the sideways beam line's detector (14 in Fig. 16) that the D1 burning took so long to complete that the ions started relaxing to the original, absorbing ground-state level, ruining the spectral hole. To remedy this, a magnetic field of 10 mT was applied to the crystal inside the cryostat, which helped lengthen the relaxation time of the hyperfine levels, thus improving the results. Further adjustments to the D1 burning regime include an increase in the scan width from 50 kHz to 500 kHz, an increased number of pulses from 10000 to 50000, and an estimated effective Rabi frequency of 1.8 kHz rather than the simulated 0.67 kHz. The values used are summarized in Table 2.

The uncertainty in the used Rabi frequencies bears mentioning. The Rabi frequency depends on the output of the laser and the polarization with respect to the dipole moment, as seen in Eq. 11. Furthermore, the electric field suffers losses when passing through the various components, making it harder to know the actual Rabi frequency reaching the crystal. Previous users of the laser and cryostat system had left behind a conversion table, but since the components in the setup and the beam diameter were different for these experiments, the Rabi frequencies used here are merely estimates.

A strategy[18] was used to improve the filtering properties of the crystal. As the burning of spectral holes moves the population around between a small number of energy levels as seen in Fig. 2, it is possible to figure out at what frequencies anti-holes are likely to be created for a

given burning frequency. Then, burning at the frequencies where anti-holes would appear should instead make an anti-hole appear near the original, intended burning frequency.

If an anti-hole is created near or around a hole, the contrast between their respective absorptions should be greater than that achieved by just burning a hole in the structure, even though the anti-hole is only populated by certain groups of ions. For that reason, before each run of burning the structure, the ranges from -45 MHz to -22 MHz and +8 MHz to +32 MHz with respect to the centre of the hole were burned with a linearly sweeping frequency.

Once the hole had been burned in the structure, it was probed with the D2 beam. The probing was done using pulses that lasted 3 μs, with an estimated effective Rabi frequency of 25 kHz.

The procedure was mostly the same for the D2 burning as for the D1 case. The magnetic field was kept at 10 mT and the frequency scan was increased from 50 kHz to 500 kHz. The number of pulses was increased from 20 to 1000, and the estimated Rabi frequency for the burning matched the intended 100 kHz from the simulations, as seen in Table 2.

Polarization	Diameter category	Rabi amplitude (kHz)	FWHM of pulse in time (μs)	Frequency sweep width (kHz)	Inversion region width (kHz)	Number of pulses needed
High-absorption (D2)	Small	100	8.415	500 (50)	182	1000 (20)
Low-absorption (D1)	Wide	1.8 (0.67)	16.83	500 (50)	66.5	50000 (1000)

Table 2: A summary of the settings used in the experiments. The values in parentheses are the optimal values found through the simulations, where these deviate from the ones actually used. The used Rabi frequencies are merely estimates.

Summary of setup and preparations

- The laser beam, with a wavelength of 605.977 nm (494.725 THz), was split into two possible paths using acousto-optic modulators. Thus, the beam could quickly be switched between the two paths.
- The two beam paths had different widths and polarizations, and reached the crystal from different directions.
- The components in the setup were able to offer a purity ratio of 10^5:1 between the two polarizations D1 and D2.
- The crystalline axes were found to be tilted by a few degrees with respect to the axes of the crystal cuboid itself.
- In the actual experiments, two of the four alternatives for burning from Table 1 were used. The narrow (1 mm diameter) beam burned with D2 light from the front, whereas the wide (10 mm diameter) beam burned with D1 from the side. Both settings used the narrow D2 beam for probing.
- There are some uncertainties regarding the true values of the Rabi frequencies used in the experiments.

Results of the experiments

Using the data points collected, converted into watts using the calibration data, a ratio of intensities before and after the crystal could be calculated. Using the Beer-Lambert law (eq. 21), the absorption coefficient α for the 6 mm crystal could be calculated from each data set. Below are listed the results from the two measurements, where the data is given with one standard deviation of accuracy. It must be clarified that the power shown in the graphs below is the power measured by the detectors after the cryostat, having gone through the cryostat windows and the remainder of the setup before reaching the detectors.

Perpendicular burning, D1

In Fig. 19, the results of the D1 burning are shown as probed with D2 light of the highest purity possible in the setup, roughly $1{:}10^{-5}$. The top graph shows the light powers of the two polarizations after passing through the crystal 2 MHz outside the centre of the spectral hole, whereas the lower plot shows the two components having passed through the middle of the hole. The probe pulses had the same intensity in both experiments.

The presence of D1 light can still be seen in the probing outside the hole, where it is of about the same order of magnitude as the D2 light. This is because the D1 light is far less strongly absorbed by the crystal, and so, the two become comparable when light of this purity passes through outside the spectral hole. The lower plot, showing the probing inside the hole, is still almost purely D2 light.

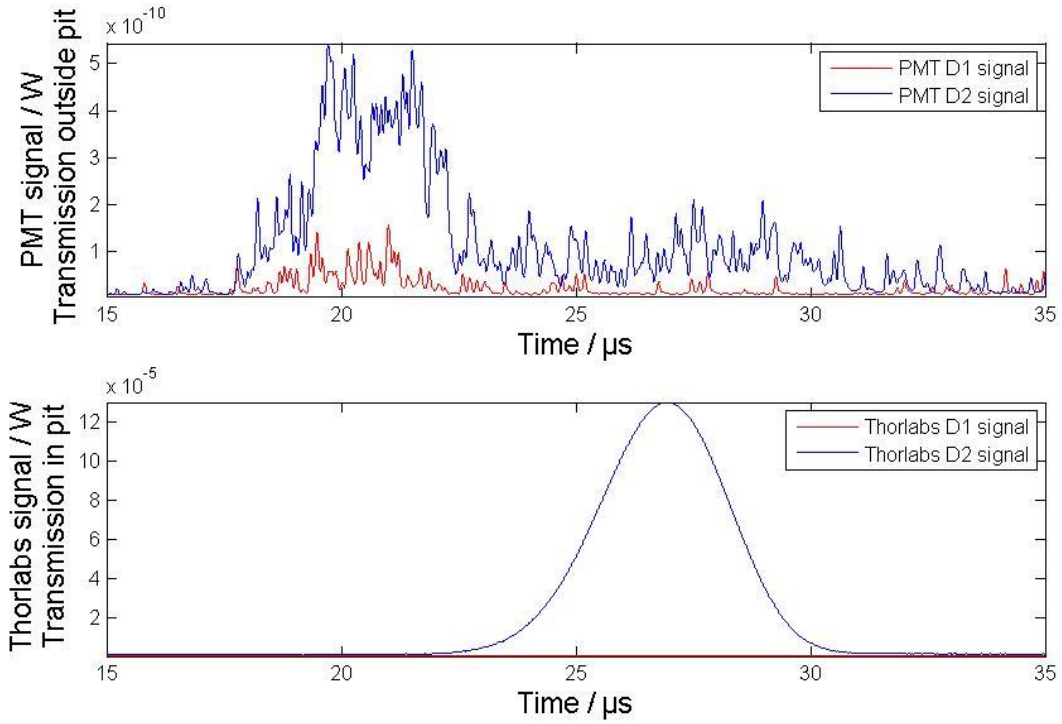

Figure 19: The results from the D1 burning experiment. The upper graph shows the signal detected by the PMT, taken 2 MHz outside the centre of the spectral hole. The lower plot shows the signal from the second Thorlabs detector, set to register the light passing through the spectral hole. In both cases, the probing used D2 light.

An interesting feature that appears in the graphs is the delay of the light which passed through the spectral hole. This is the slow light effect; light that passes through a region where absorption and refraction change rapidly as a function of frequency gains a slower group velocity. More on slow light can be found in, for example, the paper by Milonni[25].

Interestingly, the slow light effect is visible not only in the probing inside the hole but to some extent even in the PMT data, where the D2 light has a lagging tail that coincides fairly well with the bulk of the Thorlabs signal. This is most likely due to some of the probe pulse leaking through the spectral hole. While the lagging tail can be seen in Fig. 19, it is even clearer in Fig. 20 in the results from the D2 burning, due to reduced noise.

Using the data, the effective absorption coefficients could be calculated inside and outside the spectral hole. The transmission fractions, too, were calculated to determine the attenuation. For the hole burned with D1 light and probed with D2 light, the results are found in Table 3.

Parameter	**Value inside spectral hole**	**Value outside spectral hole**
Absorption coefficient	0.35 ± 0.1 cm^{-1}	21 ± 0.2 cm^{-1}
Transmission fraction	0.81 ± 0.07	$3.7 \cdot 10^{-6} \pm 0.4 \cdot 10^{-6}$
Ratio of transmission fractions	$4.6 \cdot 10^{-6} \pm 0.6 \cdot 10^{-6}$	

Table 3: The results of the experiment in which a hole was burned using the D1 polarization with a wide beam and subsequently probed using the D2 polarization with a narrow beam.

These results will be discussed in greater detail in the section titled Analysis and discussion.

Parallel burning, D2

The results of the narrow-beam D2 burning are shown in Fig. 20. Due to a different setting on the oscilloscope, this readout appears a lot less noisy than the one in Fig. 19. This is due to the data in Fig. 20 being an average of a greater number of measurements than the data in Fig. 19, making it slightly more reliable.

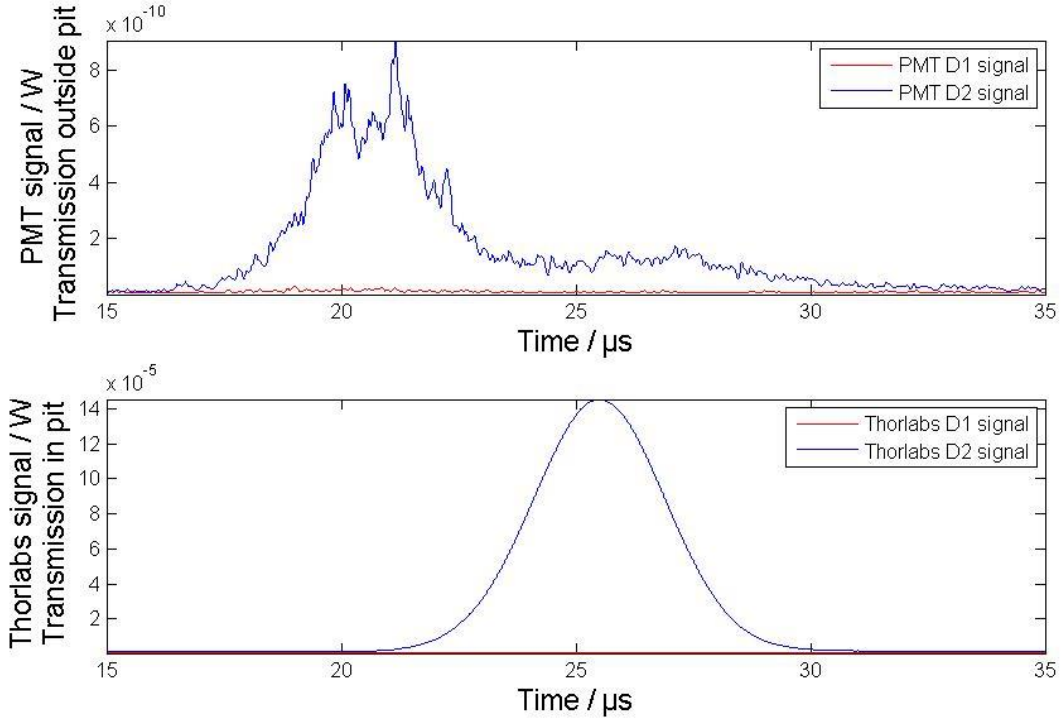

Figure 20: The detector output data from the narrow-beam experiment, with a spectral hole burned and probed using D2 light. As in Fig. 19, the top plot shows the light passing through outside the hole, while the lower one shows the power of light passing through the hole.

The slow light effects are as clearly visible here as they are in Fig. 19, with the Thorlabs signal and part of the PMT signal being delayed, the PMT indicating that part of the probe pulse leaks through the spectral hole in this case as well. The Thorlabs pulse seems a bit less delayed in this measurement, suggesting that the spectral profile does not vary as sharply in this experiment as in the previous one. The absorption coefficients and transmission fractions were calculated in the same way as before, with the results summarized in Table 4.

Parameter	**Value inside spectral hole**	**Value outside spectral hole**
Absorption coefficient	-0.084 ± 0.134 cm^{-1}	20 ± 0.2 cm^{-1}
Transmission fraction	1.1 ± 0.1	$4.8 \cdot 10^{-6} \pm 0.5 \cdot 10^{-6}$
Ratio of transmission fractions	$4.5 \cdot 10^{-6} \pm 0.6 \cdot 10^{-6}$	

Table 4: The results of the experiment using D2 light with a narrow beam for burning and probing.

As before, a more thorough discussion of these results follow in the section Analysis and discussion. It must be said already, though, that there is no way that the crystal should have a negative absorption in these experiments. This is most likely an artifact of the calibration, as the error margin of the result is large enough to allow the value to be greater than zero.

Summary of the experimental results

- In both experiments, a delay was observed in the light passing through the spectral hole. This is the slow light effect[25].
- This effect also appeared in the data from the probing outside the hole, showing that the probe pulses had some overlap with the hole.
- Both experiments showed an effective absorption coefficient of approximately 20 cm^{-1} outside the hole and close to 0 inside it.
- In both cases, the ratio of transmission fractions was about $4.5 \cdot 10^{-6}$.

Analysis and discussion

Quite a few facts stick out in Tables 3 and 4 above. First and foremost, the absorption coefficients need discussing. In the case of the D1 burning, the value inside the pit shows that the structure was burned down quite well, having an absorption coefficient fairly close to 0. The measured absorption coefficient outside the pit, however, is less than half of the value that was provided for the simulations, which was 47. Thus, the difference in suppression between inside and outside is more than 14 orders of magnitude lower than in the simulations.

It should be borne in mind, though, that the simulations were performed assuming a 10 mm thick crystal. Such a crystal with an absorption coefficient of 47 cm^{-1} and a perfectly pure polarization gives a transmission fraction on the order of 10^{-21} via the Beer-Lambert law (eq. 21). However, a crystal of thickness 6 mm would give a transmission fraction on the order of 10^{-13}. This is what the experimental results should be compared to, and as such, the deviation is explained by the differing absorption coefficient. As for why the simulations were performed using the properties of a 10 mm crystal, this was the value provided at the start of the simulations. Later changes led to a 6 mm crystal being used instead.

Still, despite the difference in magnitude between experiments and simulations, the standard deviations of the results are not very large, and the absorption values outside the pit match each other quite well between the two experiments. This is to be expected, as the values were taken outside of the range which was to be burned, meaning that the differences in hole shape should have no effect on those absorption values. The only thing that should have happened to the spectrum that far out is that the absorption should have been increased by the creation of an anti-hole around the hole, as described above.

As such, it is strange that the absorption is that low, but one clue can be found in the top graphs in Figs. 19 and 20. In both measurements, there is a leakage of D2 light through the spectral hole, as seen on the lagging tail to the right of the D2 component. This implies that a part of the probe pulse spectrum overlaps with a part of the spectral hole, letting more light pass through than it would if the probe pulse frequency had been farther from the hole. Thus, the transmission appears greater, and thereby the absorption coefficient seems smaller than it really was. This could perhaps be rectified or improved upon using probe pulses which are spectrally narrower than the current ones, meaning that they would have to last longer than 3 µs. However, they must not be long-lasting enough that they start to affect the shape of the profile by burning it.

Another reason for the low absorption might be that the Rabi frequency was different in the experiments compared to the simulations. The D1 burning had an estimated Rabi frequency which was too high compared to the value found in the simulations, resulting in a hole that would not be nearly as sharp as the one with the intended Rabi frequency. This may have reduced the absorption around the hole, lowering the absorption of the probe pulse. Still, the actual Rabi frequency to reach the crystal is a mere estimate, but for a future experiment, Rabi frequencies calibrated for the actual setup and beam would be an advantage.

Furthermore, the actual Rabi frequency for the D2 burning was most likely too low, since the pulse number needed an increase by a factor 50 to burn a hole. This increase may have affected the hole shape, causing a decrease of apparent absorption. The D1 burning required a greater number of pulses as well, implying that either the losses were far greater than anticipated, or some other factor limited the burning efficiency. One possible candidate for that is what is known as closed transitions.

It has previously been found[13] that transitions between the ±5/2 sublevels of the ground and excited states have a very high relative oscillator strength. This means that atoms excited to the ±5/2 excited state level are far more likely to decay to the ±5/2 level of the ground state than to any of its other levels. Thus, atoms in the ±5/2 ground state level that get excited to the ±5/2 excited state level by the incoming light are likely to decay back to where they came from, forming a closed transition. These, which the simulations do not take into consideration, might make the burning of a spectral hole more difficult and require more pulses, as in the experiments.

The simulations of the wider beam burning showed that a decrease in burning efficiency can be compensated with an increased pulse number without ruining the hole shape. However, the only successful simulations using thousands of pulses were those burning with the D1 polarization. Since the D2 photons used in the narrow-beam burning here are strongly interacting, hole broadening in the front end of the crystal would be a risk, as the absorbers in the front are excited and removed early on. Thus, the reduced absorption in the D2 burning may well be consequence of this hole broadening.

Another point that sticks out about the results is the negative absorption coefficient in the pit for the D2 burning. It is not realistic for the crystal to exhibit negative absorption given the experimental situation, and as can be seen, the standard deviation is large enough to allow the true value to still be positive. The reason that the value appears negative is the same that also makes the corresponding transmission fraction seem larger than one. Since there are losses in light intensity when the light passes through the cryostat and crystal off-resonance, a calibration was performed to estimate the size of the losses. These losses were overestimated to some degree, making it look as though there was more light after the crystal than before it. All the same, a physically acceptable value lies within one standard deviation of the results.

As seen in the tables, both the absorption coefficients are somewhat larger for the D1 burning, with the transmission fractions being correspondingly smaller. This could mean that the total burning was slightly more efficient in the D1, although it is hard to tell whether this was because of the polarization or just the combination of the different Rabi frequency and pulse number. Another explanation is that it can be a consequence of the D2 hole being slightly wider, causing a greater overlap between the probe pulse and the hole.

Looking at the total transmission fraction, which compares the transmission fractions outside and inside the pit, it can be seen in Tables 3 and 4 that the values are quite similar between the two experiments. Going by the simulations, this is only to be expected as both the D1 and D2 burnings achieved approximately the same degree of suppression. The deviations from the theorized values are as apparent here as in the absorption coefficients, and for the same reasons. In addition, it is worth considering whether the given absorption coefficients were in fact correct for the crystal used in the experiments. If not, that would further explain the discrepancy between the observed and simulated transmission fractions, as the given and measured alpha coefficients differed by a factor of about 2.

Either way, the leakage of light through the spectral hole is bound to have had a detrimental effect on the resulting perceived absorption coefficients. Still, looking at Figs. 19 and 20, it is clear that since the lagging part which went through the spectral hole is significantly smaller than the part which passed through outside the hole, the leakage cannot alone be responsible for the large difference in absorption compared to the given literature values.

Another consideration is that the polarization alignment may have been off with respect to the crystal itself. The inclination of the crystal axes was determined before the experiments but

later needed readjustment as the rod holding the crystal in the cryostat changed due to thermal contractions. It is not certain that the polarizations reaching the crystal were perfectly aligned with the crystal axes. The adjustment was done with good accuracy using the detectors, but nonetheless, a slight error may have occurred. Furthermore, the cryostat windows did have a discernible effect on the polarization and its purity, which must not be left out of the discussion.

Still, despite these flaws in the results, the total transmission fraction reaches an attenuation of more than five orders of magnitude, found as the inverse of the ratio of transmission fractions given in the tables. This is to be considered with respect to the main objective of this thesis work, which was to improve upon the result previously attained by Zhang *et al.*[1] As seen in their paper, they achieved an attenuation of 30 dB, corresponding to a power attenuation of 10^3 or a total transmission fraction of 10^{-3}. As such, the objective has been fulfilled as the total transmission fractions achieved here are approximately $4.5 \cdot 10^{-6}$ and the corresponding power attenuations are on the order of $2 \cdot 10^5$ (ca. 53 dB) for both burning polarizations. Again, the results obtained by Zhang *et al.* stemmed from scattered light rather than the collimated light used here, but nonetheless, an improvement of more than two orders of magnitude is still a very good result. Furthermore, this work relied on a slightly smaller frequency separation of the shifted and unshifted light than the results reported by Zhang *et al.*, specifically 2 MHz compared to the 2.3 MHz ultrasound frequency they used.

Looking at the two burning directions, they appear to have fared more or less equally well. In order to be more certain of whether the hole sharpness is better on one than the other, a full readout scan of the spectral profile would have been necessary, requiring a different detector than the ones available for this setup. Specifically, a detector capable of detecting very faint light as well as reasonably bright pulses would have been necessary, as the intensity obviously varies by more than five orders of magnitude over the spectrum. Thus, it is not possible to tell from the experimental data whether the hole shape is any better for the D1 burning than for the D2 case, as the simulations imply. However, there is also no particular reason to doubt the simulation results in this respect, and as such, further investigation is warranted.

With regards to the usefulness of the filters for UOT, it is difficult to determine as the technique has yet to be implemented in use on human tissue. Experiments have been performed on tissue phantoms[1], but no estimates have been found in literature how great the attenuation would need to be in order to use UOT *in vivo* on human tissue. Nonetheless, a continued improvement in filter efficiency can only lead closer to a future implementation of UOT within the medical discipline.

Outlook

Looking towards the future, a number of changes could be made in both the simulations and the experiments in order to improve the outcome and its reliability. Since the experiments are based on the simulations, these are the best starting point.

For this kind of application, simulations that reflect all the relevant physical properties are essential. Therefore, a script which does not just deals with two atomic levels is a very good starting point. A three-level simulator would be a minimum requirement, as the excited absorbers could decay to either the ground state or the reservoir state during longer pulses as opposed to just the ground state, as in the current code. Furthermore, it would be even more desirable to have a code which includes all the relevant sublevels of the ground and excited states. Then, such phenomena as side-holes and anti-holes could be investigated theoretically when designing pulses and burning/resetting strategies.

Being able to base the experiments on more complete and reliable simulations, the results could quite possibly be improved without changes in the setup or performance. In addition to that, there are quite a few improvements which could be made to the setup and procedure. For one thing, the Rabi frequency would need to be properly calibrated and, better yet, more exactly determined. A setup with fewer risks of losses would also be beneficial. Furthermore, once higher precision is attained in the burning and probing, it could be an advantage for both processes if polarizers with a higher purity were implemented. Another potential improvement with respect to the polarization would be the use of cryostat windows which do not alter the polarization, irrespective of the angle of incidence of the light and the mechanical and thermal stress the windows are subjected to.

Once the system had been accurately modelled and the experimental errors had been ironed out, the experimental setup could be rearranged to allow for D1 and D2 burning from both the projected directions, in order to determine whether and to what extent D1 burning really is better than D2 burning for making sharp spectral holes.

It is recommended to still keep a polarizer after the cryostat, as was done in the experiments described above. Without it, the attenuation is limited by the absorption coefficient of the D1 component, as whatever component of D1 light passes through the polarizer for the probe beam will make it through the crystal relatively unaffected. As such, the light leaving the crystal might end up being mostly D1 light, which is why another polarizer is necessary unless the polarizer in the probe beam has a very high purity.

Further possible improvements include a crystal with well-determined axes, preferably aligned with the macroscopic cutting axes of the crystal's cuboidal shape; more measurement series, to increase the certainty of the measured results; a detector which is able to detect reliably within a range of intensities spanning as much as the projected transmission fraction ratio, to allow for a full spectral readout, if such a detector is even possible; and a rod inside the cryostat which contracts and bends as little as possible when subjected to cryogenic temperatures.

As can be seen, numerous improvements are possible which would, most probably, increase the quality and reliability of the results, making spectrally hole-burned filters all the more viable for UOT. It is the author's sincere hope that this research will go on, as it is interesting not only from a scientific viewpoint but also because of its potential in health care and medicine.

Appendix

This appendix contains the drawings from which the crystal holder, shown in Fig. 17, was crafted. The holder was designed as part of this thesis work, with the holder itself being crafted according to said designs by BEMA Slip AB.

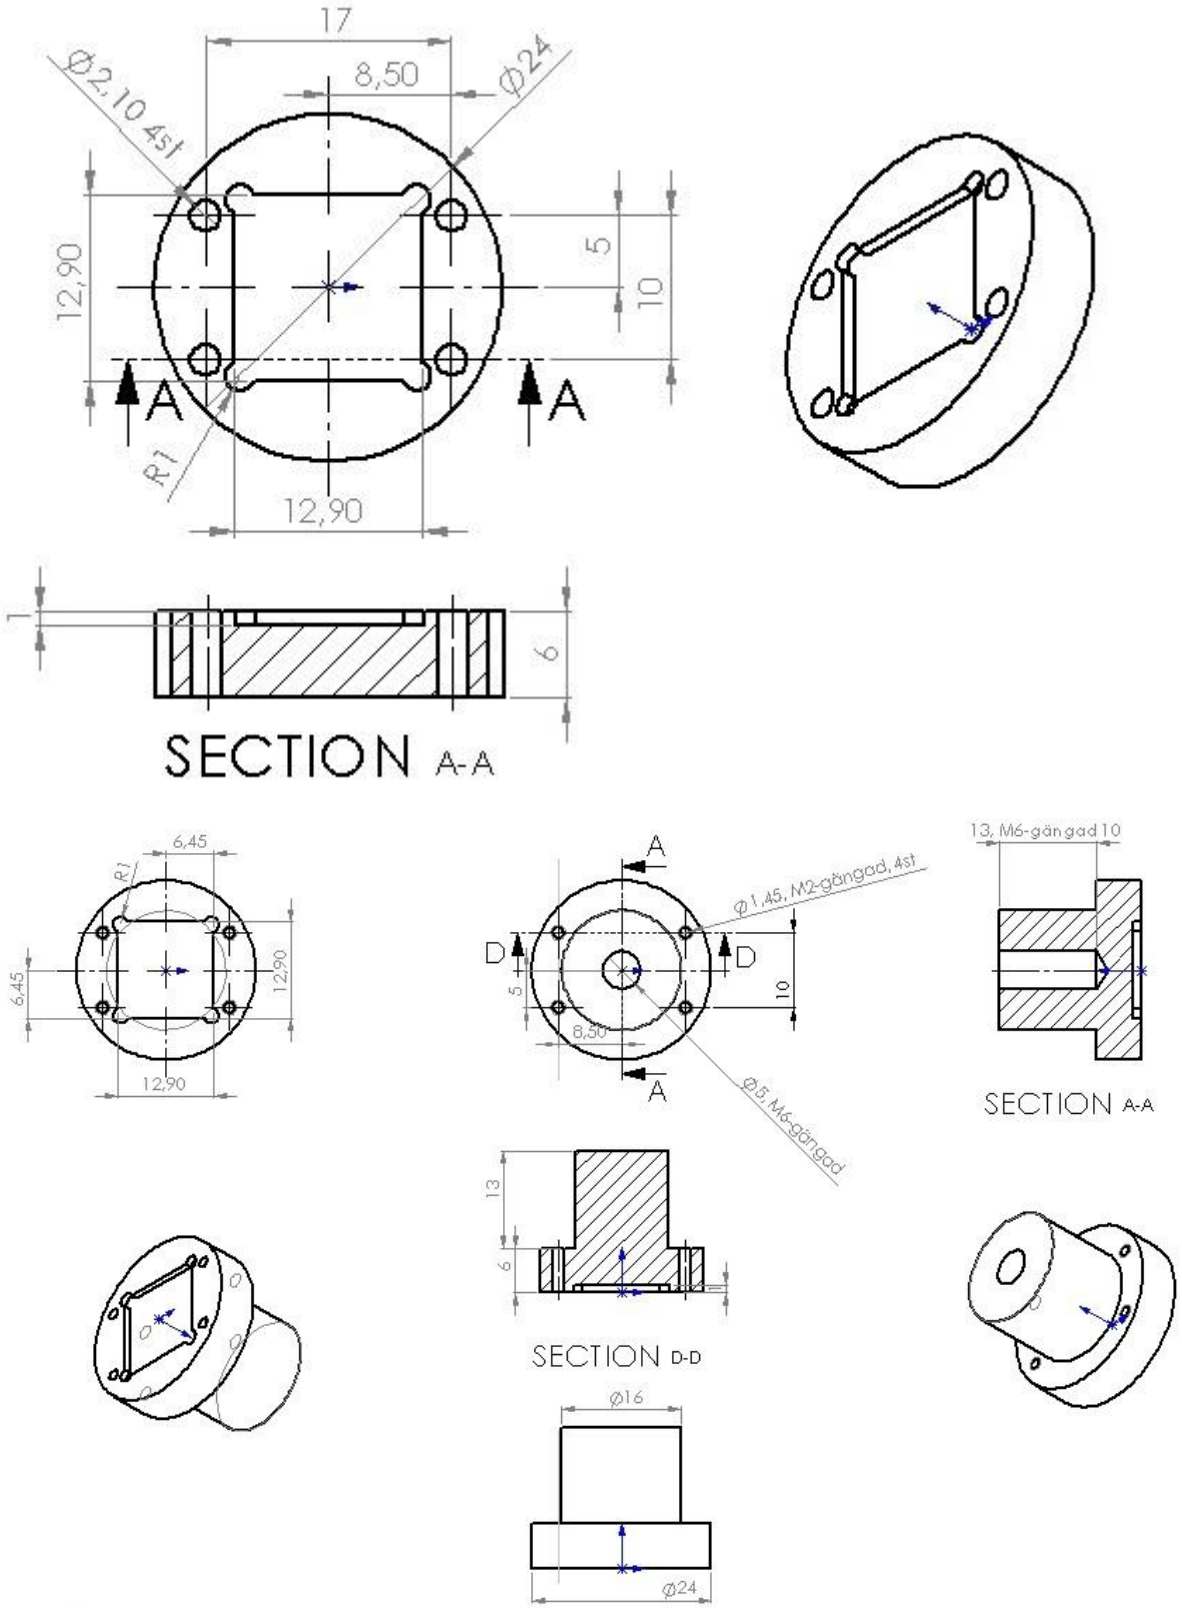

References

1. Zhang H, Sabooni M, Rippe L, Kim C, Kröll S, Wang LV, Hemmer PR. Slow light for deep tissue imaging with ultrasound modulation. Appl. Phys. Lett. [Internet]. 2012 Mar 26;100:131102:1-5.

2. Reddy NRS, Lyle PA, Small GJ. Applications of spectral hole burning spectroscopies to antenna and reaction center complexes. Photosynthesis Research. 1992;31:167-194.

3. Masumoto Y. Persistent spectral-hole-burning in semiconductor quantum dots and its application to spectroscopy. Jpn. J. Appl. Phys. 1999 Jan;(Pt 1):38(1B):570-576.

4. Li Y, Hemmer P, Kim C, Zhang H, Wang LV. Detection of ultrasound-modulated diffuse photons using spectral-hole burning. Opt. Expr. 2008 Sep 5;16(19): 13 p.

5. Marks FA, Tomlinson HW, Brooksby GW. A comprehensive approach to breast cancer detection using light: photon localization by ultrasound modulation and tissue characterization by spectral discrimination. Proc. SPIE. 1993 Sep 14;1888: p. 500-510.

6. Elson DS, Li R, Dunsby C, Eckersley R, Tang M-X. Ultrasound-mediated optical tomography: a review of current methods. Interface focus. 2011;1:632-648.

7. Thuresson A. Numerical simulations of highly efficient quantum memories [master's thesis]. [Lund]: Lund University; 2011. 83 p.

8. Saleh BEA, Teich MC. Fundamentals of Photonics. 2nd ed. Hoboken, New Jersey: John Wiley & Sons, Inc.; 2007. 1202 p.

9. Simons J, Hayashi T. 6.2: Experimental Probes of Electronic Structure [Internet]. Davis (CA): University of California, Davis. Publication date unknown [accessed 2015 Apr 28]. Available from: http://chemwiki.ucdavis.edu/Textbook_Maps/Theoretical_Chemistry_Textbook_Maps/Simon's_%22Advanced_Theoretical_Chemistry%22/6%3A_Electronic_Structures/6.2%3A_Experimental_Probes_of_Electronic_Structure

10. Nilsson M. Coherent interactions in rare-earth-ion-doped crystals for applications in quantum information science [doctoral thesis]. [Lund]: Lund Institute of Technology; 2004. 208 p.

11. Equall RW, Cone RL, Macfarlane RM. Homogeneous broadening and hyperfine structure of optical transitions in Pr^{3+}:Y_2SiO_5. Phys Rev B. 1995 Aug 1;52(6):3963-9.

12. Wang HH, Li AJ, Du DM, Fan YF, Wang L, Kang ZH, Jiang Y, Wu JH, Gao JY. All-optical routing by light storage in a Pr^{3+}:Y_2SiO_5 crystal. Appl Phys Lett. 2008;93(221112):1-3.

13. Nilsson M, Rippe L, Kröll S, Klieber R, Suter D. Hole-burning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids demonstrated in Pr^{3+}:Y_2SiO_5. Phys Rev B. 2004;70(214116):1-11.

14. Scully MO, Zubairy MS. Quantum optics. 1st edition, 6th printing. Cambridge: Cambridge University Press; 1997. 630 p.

15. Lai P, Xu X, Wang LV. Ultrasound-modulated optical tomography at new depth. J. Biomed. Opt. 2012 Jun;17(6):066006:1-6

16. Szabo A. Observation of hole burning and cross relaxation effects in ruby. Phys Rev B. 1975 Jun 1;11(11):4512-20.

17. Li C, Wyon C, Moncorgé R. Spectroscopic properties and fluorescence dynamics of Er^{3+} and Yb^{3+} in Y_2SiO_5. IEEE J. Quant. Elect.; 1992 Apr:28(4):1209-1221.

18. Rippe, Lars (Division of Atomic Physics, Department of Physics, Lund University, Lund, Sweden). Conversation with: Philip Dalsbecker (Division of Atomic Physics, Department of Physics, Lund University, Lund, Sweden). 2014 Aug-Sep.

19. Silver MS, Joseph RI, Hoult DI. Selective spin inversion in nuclear magnetic resonance and coherent optics through an exact solution of the Bloch-Riccati equation. Phys Rev A. 1985 Apr;31(4):2753-5.

20. Kröll, Stefan (Division of Atomic Physics, Department of Physics, Lund University, Lund, Sweden). Conversation with: Philip Dalsbecker (Division of Atomic Physics, Department of Physics, Lund University, Lund, Sweden). 2015 Jan-Feb.

21. Rippe L, Julsgaard B, Walther A, Kröll S. Laser stabilization using spectral hole burning. Submitted to Phys. Rev. A. Arxiv. Available from http://arxiv.org/abs/quant-ph/0611056. [Accessed: 25th April 2015].

22. Karlsson J. Investigations of cerium as a readout mechanism for a single instance qubit register in a crystal (preliminary title) [doctoral thesis]. [Lund]: Lund Institute of Technology; Forthcoming 2015

23. Baschek G, Hartwig G, Zahradnik F. Effects of water absorption in polymers at low and high temperatures. Polymers. 1999;40:3433-3441.

24. Van Schiver SW. Helium cryogenics. Online edition. New York: Springer Science+Business Media, LLC; 2012. 470 p.

25. Milonni PW. Controlling the speed of light pulses. J. Phys. B: At. Mol. Opt. Phys. 2002;35:R31–R56.

26. Creative Commons [Internet]. [place unknown]: creativecommons.org; Publication date unknown [accessed 2015 Jul 22]. Available from: https://creativecommons.org/licenses/by-nc-sa/3.0/us/legalcode

Fig. 2 adapted with permission from Wang HH, Li AJ, Du DM, Fan YF, Wang L, Kang ZH, Jiang Y, Wu JH, Gao JY. All-optical routing by light storage in a Pr^{3+}:Y_2SiO_5 crystal. Appl Phys Lett. 2008;93(221112):1-3. Copyright 2008, AIP Publishing LLC. DOI: http://dx.doi.org/10.1063/1.3041645

Fig. 5 reproduced with permission from Zhang H, Sabooni M, Rippe L, Kim C, Kröll S, Wang LV, Hemmer PR. Slow light for deep tissue imaging with ultrasound modulation. Appl. Phys. Lett. [Internet]. 2012 Mar 26;100:131102:1-5. Copyright 2012, AIP Publishing LLC. DOI: http://dx.doi.org/10.1063/1.3696307