



LUND UNIVERSITY

Quantum computing with naturally trapped sub-nanometre-spaced ions

Rippe, Lars

2006

[Link to publication](#)

Citation for published version (APA):

Rippe, L. (2006). *Quantum computing with naturally trapped sub-nanometre-spaced ions*. [Doctoral Thesis (compilation), Atomic Physics]. Division of Atomic Physics, Department of Physics, Faculty of Engineering, LTH, Lund University.

Total number of authors:

1

General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117
221 00 Lund
+46 46-222 00 00

QUANTUM COMPUTING WITH NATURALLY TRAPPED SUB-NANOMETRE-SPACED IONS

Lars Rippe

Doctoral Thesis
2006



LUND UNIVERSITY

QUANTUM COMPUTING WITH NATURALLY TRAPPED SUB-NANOMETRE-
SPACED IONS

© 2006 Lars Rippe
All rights reserved
Printed in Sweden by Media-Tryck, Lund, 2006

Division of Atomic Physics
Department of Physics
Faculty of Engineering, LTH
Lund University
P.O. Box 118
SE-221 00 Lund
Sweden
<http://www.atom.fysik.lth.se>

ISSN 0281-2762
Lund Reports on Atomic Physics, LRAP-365

ISBN 13: 978-91-628-6907-6
ISBN 10: 91-628-6907-8

TO MY WIFE'S SURPRISE

ABSTRACT

The main aim of this work, was to lay the foundations for the experimental realisation of a quantum mechanical controlled NOT gate in rare-earth-metal-ion-doped crystals.

Small amounts of rare-earth elements, added during the growth of some inorganic crystals, will become substituted into the crystal lattice as trivalent ions. The trivalent rare-earth-metal ions between cerium, with atomic number 58, and ytterbium, with atomic number 70, have a partly filled 4f shell, which does not extend spatially outside the full 5s and 5p shells. The 4f vacancies make electronic *inner* shell transitions possible between spectroscopic 4f terms. Some of these optical transitions have coherence times of the order of milliseconds, when the crystals are cooled down to ~ 4 K. There are several reasons for these extraordinary coherence times, which are approximately 8 orders of magnitude greater than those typical for electronic transitions in solids. The most important one is the cage-like shield which the outer 5s and 5p shells provide for the 4f electrons. Furthermore, since these ions are naturally trapped inside the crystal lattice there is no Doppler broadening of the line-width. The coherence properties of these optical transitions is one of the features that makes these materials attractive for use as a solid-state platform for quantum computing, using these ions as qubits. Another appealing characteristic is the fact that different ions have different optical resonance frequencies, which means that ions belonging to different qubits, which only have nm separation, can still be addressed separately by using different laser frequencies. Since the inter-ion spacing is so small, it is possible to make two ions interact strongly, although they are well shielded, through a permanent dipole-dipole interaction. This interaction can be turned on and off by switching between two different ways of encoding the qubit, a most useful feature. When the qubit is represented as a superposition between two ground state hyperfine levels, the interaction is turned off. The interaction is turned on selectively by transferring this superposition to the optical transition with a π -pulse, for the specific ions that are to interact.

This thesis describes how peaks of ions, absorbing on a single transition, residing in spectral pits with no other ions, have been isolated. It is shown how these ions can be coherently transferred between hyperfine levels via the optically excited state, how the interaction between such peaks of ions representing qubits can be turned on and off, and how subgroups of ions with strong interaction can be distilled out. All the work described here has been performed using the ensemble approach.

The ensemble approach will, however, be difficult to scale up to large numbers of qubits. A method employing a single ion in each qubit, using a specialised ion for readout, has therefore also been proposed.

The rare-earth-metal-ion-based quantum computing experiments require a laser with coherence properties which at least match those of the material. To this end a stabilisation system was developed for a dye laser. This system uses a transient spectral hole in a rare-earth-metal-ion-doped crystal, of the same kind that is used in the experiments, as frequency reference, and is to the authors knowledge the first demonstration of locking a dye laser to a spectral hole. This system provides a line-width of 1 kHz on a 10 μ s timescale and a frequency drift below 1 kHz/s.

POPULÄRVETENSKAPLIG SAMMANFATTNING

Kvantdatorer har den unika egenskapen att de kan utföra samma beräkning för många olika startvärden samtidigt. Informationen i en vanlig dator är lagrad i ett minne som består av bitar, som kan ha värdena 1 eller 0. När ett program körs beror slutresultatet av startvärdena hos dessa bitar. Minnet i en kvantdator lagrar istället kvantbitar, som även de kan ha värdena 1 eller 0. Men kvantbitarna kan även befinna sig i ett "kanske-tillstånd", en så kallad superposition, där de kanske har värdet 0 och kanske 1. Om några av kvantbitarna i ett kvantdatorminne från början sätts i ett "kanske-tillstånd", utförs alla beräkningar, vilka motsvarar de olika möjliga kombinationerna av "kanske-bitarna", samtidigt när programmet körs. När man sedan ser på resultatet av programmet får man slumpvis ett av de möjliga svaren, vilket kan tyckas vara ineffektivt. Det visar sig, emellertid, att vissa problem som tar orimligt lång tid att beräkna med en vanlig dator, kan lösas effektivt av en kvantdator med specialskrivna program.

I detta arbetet har jag jobbat med kvantbitar, som är baserade på joner som hör till gruppen sällsynta jordartsmetaller. Jonerna sitter mycket nära varandra i genomskinliga kristaller. Dessa joner har den ovanliga egenskapen att kunna komma ihåg "kanske-tillståndet" för en kvantbit länge, trots att "kanske-tillstånd" normalt sett är mycket kortlivade. Programmen i våra kvantdatorer består av pulser av laserljus, med vilka vi belyser jonerna. Olika joner påverkas av olika frekvenser hos laserljuset, på liknande sätt som en radiomottagare inställd på en kanal bara tar emot den kanalen. Genom att byta frekvens på ljuset kan olika grupper av joner påverkas separat.

I arbetet, som denna avhandling beskriver, har en teknik utvecklats för att med hjälp av ljuspulser plocka bort alla joner inom ett frekvensintervall och sedan lyfta tillbaka joner med bara en välbestämd frekvens, som skall utgöra en kvantbit i vår kvantdator. Denna kvantbit skiftas sedan mellan värde 1 och 0 upprepade gånger, en så kallad en-bitars grind. Om man upprepar denna

procedur för två olika frekvensintervaller, så kan man skapa två kvantbitar. För att kunna skapa kvantgrindar med två bitar, så måste en kvantbit kunna påverka en annan. I vårt fall kan en jon, som tillhör en av kvantbitarna, ändra vilken ljusfrekvens som en annan jon annars skulle ha påverkats av. Ändringen kan liknas vid att radions kanalinställning ändras lite, varvid man inte längre kan ta emot radiokanalen. I detta arbete har vi visat detta fenomen.

Det är endast de joner som råkar sitta nära varandra, som har tillräckligt stor påverkan för att kunna användas som kvantbitar. Vi har experimentellt visat hur endast de joner, som påverkas mycket, kan väljas ut.

Ovanstående kvantdatorschema fungerar för två kvantbitar, men skall man ha fler kvantbitar blir det för få joner kvar i ett närliggande område som alla påverkar varandra, för att man skall kunna detektera dem. En metod som gör det möjligt att ha fler kvantbitar beskrivs i denna avhandling. I denna metod har man endast en jon per kvantbit och utläsningen sker med hjälp av en speciell utläsningsjon, som har dåliga kvantbitsegenskaper, men är lätt att detektera.

Frekvensen och fasen på de laserpulser som används, till att manipulera jonerna, måste vara mycket stabil. Jag beskriver därför, hur vi har stabiliserat en laser genom att jämföra dess frekvens med frekvensen hos samma typ av joner, som kvantbitarna består av, och sedan korrigera frekvensfelen.

LIST OF PUBLICATIONS

This thesis is based on the following papers, which will be referred to by their roman numerals in the text.

- I Initial experiments concerning quantum information processing in rare-earth-ion doped crystals**
M. Nilsson, L. Rippe, N. Ohlsson, T. Christiansson and S. Kröll.
Physica Scripta **T102**, 178 (2002).
- II Holeburning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids, demonstrated in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$**
M. Nilsson, L. Rippe, R. Klieber, D. Suter and S. Kröll.
Phys. Rev. B **70**, 214116 (2004).
- III Experimental demonstration of efficient and selective population transfer and qubit distillation in a rare-earth-metal-ion-doped crystal**
L. Rippe, M. Nilsson, R. Klieber, D. Suter and S. Kröll.,
Phys. Rev. A **71**, 062328 (2005).
- IV Scalable designs for quantum computing with rare-earth-ion-doped crystals**
J. H. Wesenberg, K. Moelmer, L. Rippe, and S. Kröll.
Submitted to Phys. Rev. A, arXiv:quant-ph/0601141, (2006).
- V Laser stabilization using spectral hole burning**
L. Rippe, B. Julsgaard, A. Walther and S. Kröll.
Submitted to Phys. Rev. A, arXiv:quant-ph/0611056, (2006).

VI Mode-hop-free electro-optically tuned diode laser

L. Levin (Levin is the authors former surname).

Opt. Lett. **27**, 237 (2002).

ABBREVIATIONS

AD	Analogue to digital (converter)
AOM	Acousto-optic modulator
DPSS	Diode-pumped solid-state...
DDS	Direct digital synthesiser
DSP	Digital signal processing
FID	Free induction decay
FWHM	Full width at half maximum
E-field	Electric field
EOM	Electro-optic modulator
FET	Field effect transistor
FFT	Fast-Fourier transform
MRI	Magnetic resonance imaging
NMR	Nuclear magnetic resonance
PI	Proportional-integral (controller)
PID	Proportional-integral-derivative (controller)
RAM	Residual amplitude modulation
REQC	Rare-earth-metal-ion-doped-crystal-based quantum computing
RF	Radio frequency
RMS	Root mean square
SFDR	Spurious-free dynamic range
ULE	Ultra-low thermal expansion (glass)

CONTENTS

1	Introduction	1
1.1	Bits, qubits and quantum computation	2
1.2	Naturally trapped ions as qubits	3
1.3	Light-controlled cubits	3
1.4	The aim and outline of this thesis	4
2	Naturally trapped ions	5
2.1	Long coherence times in a solid material	6
2.2	Homogeneous line-width	8
2.3	Hyperfine levels	8
2.4	Inhomogeneous broadening	9
2.5	Transition strengths	9
2.6	Ion-ion interactions	10
3	Spectral hole burning	13
3.1	Hole-burning mechanisms	13
3.2	Isolating a peak of ions in a single hyperfine state	15
4	Coherent manipulation	17
4.1	Two-level systems	17
4.1.1	Field-induced electron displacement	17
4.1.2	Time varying fields	19
4.2	The Bloch sphere	20
4.3	The Bloch equations	21
4.3.1	Including decay and decoherence	23
4.3.2	Bloch equations in different forms	23
4.3.3	Pulse area	23
4.4	Maxwell-Bloch equations	25
4.4.1	Steady-state solutions	26
4.5	Free induction decay	27
4.6	The complex hyperbolic secant pulse	30
4.6.1	Theoretical transfer efficiency	31
4.6.2	Pulses with broader inversion	31
5	Quantum computing with Rare-earth-metal-ion-doped crystals	33
5.1	Introduction	33
5.2	What is needed to construct a quantum computer?	34
5.2.1	Qubits	34
5.2.2	Initialisation to a pure state	35
5.2.3	One-bit gates	36
5.2.4	Two-qubit interactions	36
5.2.5	Two-qubit CNOT gates	38
5.2.6	Readout	39
5.3	Single-ion REQC	39

6	Equipment	43
6.1	The dye laser	43
6.2	Cooled diode lasers	44
6.3	Acousto-optic modulators	45
6.4	Cryostats	46
7	Laser stabilisation	47
7.1	Locking to cavities	48
7.1.1	Cavity field build up	48
7.1.2	Two mirrors, a cavity	50
7.1.3	Pound-Drever-Hall locking	51
7.2	Locking to holes	55
7.2.1	Constant frequency error	56
7.2.2	Phase jump	57
7.2.3	Sign sensitive error detection	58
7.2.4	Locking to holes with different storage mechanisms	59
7.2.5	Hole drift	59
7.2.6	Advantages with spectral hole locking	60
7.3	System design	61
8	Laser stabilisation electronics	63
8.1	General design considerations	66
8.1.1	Short delay	66
8.1.2	Low-noise electronics	66
8.1.3	Avoiding ground loops	66
8.1.4	RF shielding	66
8.1.5	Avoiding signal reflections	67
8.1.6	Circuit board design considerations	67
8.2	Electro-optic modulator	68
8.2.1	Phase modulator, B1800	68
8.2.2	Intracavity EOM, B2300	69
8.3	Error photodetector, B2200	70
8.4	Phase detector, B1400	70
8.5	EOM low voltage generator, B300	70
8.6	EOM high-voltage generator, B400	71
8.7	Suspend regulation controller, B500	73
8.8	Operation software and hardware	74
8.8.1	Digital signal processing card	74
8.8.2	Realignment cameras	75
8.8.3	Break-out box, B100	75
8.9	Modifications to dye laser detector unit, B700	76
8.10	Phase shifter, B2000	76
8.11	Diplexer filters	76
8.12	Amplitude stabilisation system	78
8.13	Table of systems, components and descriptions	80
	Epilogue	85
	Comments on the Papers	87
	Acknowledgements	91
	References	93
	Schematics	100

Papers

I	Initial experiments concerning quantum information processing in rare-earth-ion doped crystals	117
II	Holeburning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids, demonstrated in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$	127
III	Experimental demonstration of efficient and selective population transfer and qubit distillation in a rare-earth-metal-ion-doped crystal	141
IV	Scalable designs for quantum computing with rare-earth-ion-doped crystals	155
V	Laser stabilization using spectral hole burning	165
VI	Mode-hop-free electro-optically tuned diode laser	193

INTRODUCTION

Nature itself uses impressive structures which are only a few atoms wide, such as the information carrying DNA spiral which is ~ 2 nm across. Over times chemists have learnt how to create specific molecules by heating, cooling changing the pH and altering the pressure in reactions, and so on. Driven by the quest for faster computers, where hand-held mobile phones today have the same computer power as a mainframe computer had 30 years ago [1], the size of the structure fabricated by man is shrinking. The “wires” in the latest integrated circuits are less than a thousand atoms wide, and even smaller structures are being fabricated in the field of nanotechnology. As the structures shrink towards the size of a few atoms, the quantum aspects of matter will become more and more apparent. To be able to utilise the quantum aspect to our advantage, it is important increase our knowledge about them and how they can be manipulated in a controlled way. The field of quantum computing provides an organised way of understanding and learning to coherently manipulate these quantum properties, and this is important in itself.

The concept of quantum computers became widely known following Peter Shor’s famous paper about quantum factorisation of large numbers [2, 3] in 1994, and Lov K. Grover’s quantum search algorithm [4, 5] in 1996. There are, however, many technical difficulties to be overcome before we can construct a quantum computer, which utilising these algorithms can beat a classical computer, since this requires a many-qubit quantum computer with very a low error rate.

Quantum computing research is however useful at a much earlier stage than when it reaches the full implementation of Shor’s and Grover’s algorithm. As mentioned above, it provides a organised way of gaining knowledge on how quantum mechanical systems can be coherently manipulated. A quantum computer with only a few bits can be used to purify and relay entanglement,

a key concept for long-distance quantum cryptography. Furthermore, even a fairly small quantum computer can simulate quantum mechanics more efficiently than a classical computer, and this may be important when designing structures of a size where the quantum mechanical properties make themselves known. The possibility of simulating any quantum system efficiently with a quantum computer was suggested by Richard P. Feynman in 1982 [6], and proven in 1996 by Seth Lloyd [7].

1.1 Bits, qubits and quantum computation

In a normal computer, the smallest piece of information is called a bit, and the bit can have either the value zero or one. When a computation is carried out, deterministic decisions are made based upon the values of the bits, and given the same input the computer will always give the same result. The equivalent to the bit in a quantum computer is called a qubit, and a qubit can also have the value zero or one. If the qubits in a quantum computer are initialised to the same sequence of zeros and ones as for a classical computer, and the same program is executed, the quantum computer will come to exactly the same result. Why then go through all the trouble of building a quantum computer? The answer lies in the fact that qubits can not only be initialised to the values zero or one, they can also be put in a *mixture* of one and zero, a so-called superposition. If one of the input qubits is put into such a superposition state before computation is started, the quantum computer will carry out both computations corresponding to the case when the qubit has the value zero and one, simultaneously. It is not until you measure the result of the computation that it is decided which of the possible answers will be presented, and at this measurement the input qubit, which was originally in a superposition state, collapses to reflect the value, zero or one, that would correspond to the result which was measured. If several of the input qubits are in superposition states before the computation is started, the quantum computer will calculate the results for *all combinations* of input states for the qubits in superposition simultaneously, but again, the readout will only result in one answer, together with the corresponding input values. Let us consider a calculation in which many of the combinations of input give the same result, i.e. some outputs are much more *probable* than others. Since the quantum computer is probabilistic, it is much more likely that we will get the result to which many of the input states corresponded. When we read out a result that corresponds to multiple questions, the input will collapse to a superposition of those input states. One could say that, in some sense, the quantum computer does not have a preferred direction of calculation, and this can be used to solve some problems more efficiently than any known

algorithm on a normal computer [3, 5–9].

1.2 Naturally trapped ions as qubits

When some transparent inorganic crystals are grown, for example, yttrium silicate, Y_2SiO_5 , small amounts of added rare-earth elements will form trivalent ions. When these crystals are cooled down to $\sim 4\text{ K}$ a few of the optical transitions in some of the rare-earth-metal ions will have unique coherent properties, with coherence times of the order of milliseconds. One reason for these unique coherence properties is that the transitions are between different spectroscopic 4f terms, and the 4f electrons responsible for these transitions are not spatially located at the very outmost part of the ion. On the contrary, they are shielded by full 5s and 5p electron shells, which act as a cage protecting the 4f electrons. Furthermore, there is no Doppler broadening¹, since the ions are trapped by the crystal lattice, and a line-width as low as 50 Hz has been measured for an optical electronic transition in these materials [10].

It is possible to drive these transitions with laser light, and thereby transfer the ions between different spectroscopic levels.

Although the line-width of an individual ion is very narrow, different ions have different resonance frequencies, and together they form an absorption line-width of the order of gigahertz. Since different ions have different resonance frequencies, it is possible to individually manipulate ions which are located much closer together than the smallest spot to which a laser can be focused, by changing the frequency of the laser light.

Groups of ions located very close together can affect each other through dipole-dipole interactions. Which ions that interact can be selected by tuning a laser to the resonance frequency of the appropriate ions.

1.3 Light-controlled cubits

Manipulating ion qubits, that is, transferring them between the states corresponding to zero, one and any superposition state, is done with pulses of laser light. These pulses of light are, in some way, equivalent to the program which is read into the processor in a normal computer. Continuing this analogy, designing the shape of the pulses can be said to be equivalent to writing the computer program. However, it is a great technical challenge to design a laser system in which the light is coherent for a time that is as long

¹If the ions were moving while emitting, for example the atoms in a gas, the frequency recorded by an observer would appear different depending on whether the atoms were moving towards or away from the observer. In a gas at room temperature, this effect is approximately 2 GHz.

as the coherence time of our qubits, and where the pulses which imitate the desired modulation can be created with high enough accuracy.

1.4 The aim and outline of this thesis

The aim of this work, was to explore the possibilities of creating a basic quantum gate in a rare-earth-ion-doped crystal. In Paper [I](#) we describe the initial attempts of creating a qubit structure and also a measurement of the dipole-dipole interaction using the widely tunable laser described in Paper [VI](#). In Paper [II](#) the techniques creating qubit structures are refined and different properties of the material are measured, some of which were not possible to measure with previous techniques. Paper [III](#) then continues to show how the qubit structures can be manipulated in a coherent way, and show how a qubit containing only ions that interacted strongly with another qubit can be distilled out. As the experiments are being refined, the demands on the stability of the phase and frequency of the laser, used to manipulate the qubits, increase. I therefore continued with a project stabilising a dye-laser to a spectral hole, created in a crystal similar to those used for the quantum computing experiments. This laser, together with an elaborate theory of stabilisation to a spectral hole, is described in Paper [V](#). The scheme for quantum computing in rare-earth-ion-doped-crystals has limited possibilities of scaling, and Paper [IV](#) describes two different ways of altering the scheme to make it scalable.

I find that obtaining good conceptual model of the basics is often difficult and takes a long time. Once this is done however, understanding how things work is fairly easy. For me, the conceptual models are also necessary in order to come up with new ideas, like for example the method to achieve mode-hop-free-tuning described in Paper [VI](#). In this thesis a large effort is devoted to explaining some such basic concepts and the aim of this is to simplify the understanding of the papers describing my research.

The lines and shapes in figures has been referred to with the colour they appear with in the electronic version of the thesis. Where necessary, the lines have also been coded with letters or different line styles.

NATURALLY TRAPPED IONS

Quantum computing requires quantum mechanical systems with phase memory, also called coherence time, which is long compared with the time required to carry out one operation. For most quantum systems in gases and solids, the phase memory is rapidly destroyed by uncontrolled interactions with surrounding atoms or ions. There are however, several different ways of achieving long coherence times. One method which has been successfully used and has produced many impressive results in the field of quantum computing is that employed in ion traps [11–15], where undesired interactions are avoided by removing the surrounding ions and atoms. This is achieved by keeping the ions in an electromagnetic trap under high vacuum. Another method is to use a gas consisting of atoms where the quantum property under consideration is not changed when the atoms collide with the paraffin-covered walls of their container [16]. A quantum system in solids that often has very good coherence properties, is the spin of the nucleus of an atom or ion. This has been exploited extensively in nuclear magnetic resonance (NMR) spectroscopy and magnetic resonance imaging (MRI). The problem with NMR quantum information experiments, is related to the difficulty in creating pure starting states.

In this chapter the materials, in which the work described in this thesis is based on, is presented. It is a crystal that contains ions where not only the nuclear spin states, but also the optically excited states, have good coherence properties. Therefore, it is possible to use lasers to manipulate the hyperfine states of ions via optical excitation to the electronically excited states, without destroying the coherence. Accurate coherent state preparation requires lasers with correspondingly good coherent properties, as discussed in Chapter 4, and since no such laser was commercially available within the wavelength range needed, one was constructed, as described in Chapters 7 and 8. Optical manipulation allows spa-

tial subgroups of ions to be addressed by frequency and allows pure states to be created by optical pumping.

2.1 Long coherence times in a solid material

Rare-earth elements doped into some inorganic crystals, which are cooled down to liquid helium temperatures can, in some cases, exhibit extraordinary coherence properties. Let us very briefly consider a few reasons for these long coherence times, starting with a short description of the rare-earth elements.

The rare-earth elements include scandium and yttrium, which will not be further considered since they do not have the necessary properties, and the 15 lanthanides which start with lanthanum, with atomic number 57, and end with lutetium, with atomic number 71. When these elements are doped into some inorganic crystals they will form trivalent ions. These ions all have full 5s and 5p shells, which act as a shield protecting the inner 4f electrons. The 4f shell can have a maximum of 14 electrons. Lanthanum³⁺ has no 4f electron, and as one progresses along the lanthanide ions there is one electron more in the 4f shell for each lanthanide, finishing with lutetium³⁺ which has a full 4f shell with 14 electrons. Apart from lanthanum, which has no 4f electrons, and lutetium which has a full shell, the 4f electrons of the other rare-earth-metal ions can make inner shell 4f-4f transitions.

Figure 2.1, shows the energy levels for the rare-earth-metal ions doped into LaCl_3 , starting with cerium³⁺ which has one electron, up to ytterbium³⁺ with 13 electrons. Since the energy levels are mainly decided by spin-orbit interactions and not crystal-field interactions, the energy levels will be similar in other crystal hosts. The reason for the good coherence properties of 4f-4f transitions in the rare-earth elements is directly connected with the outer-lying 5s and 5p electrons. They act as a cage, shielding the 4f electrons from perturbations from the surroundings, which means that these elements have free-ion-like coherence properties, even when they are inside a crystal. This in itself is not sufficient to achieve long coherence times. Interaction with phonons will destroy the coherence properties, but the phonon density can be lowered by cooling the crystal to between 1.5 K and 4 K. In the present work liquid helium was used to cool the crystals, but there are now several closed-cycle cryostats for these temperatures available on the market. One problem of closed-cycle cryostats, however, can be vibrations. This problem is being addressed by the manufacturers, and the levels of vibration are continuously being reduced.

The main part of this work was performed in the crystal hosts Y_2SiO_5 (yttrium silicate or YSO) and $\text{Y}_3\text{Al}_5\text{O}_{12}$ (yttrium aluminium garnet or YAG).

G. F. Imbusch and R. Kopelman




Figure 2.1. Energy levels for triply ionised rare-earth elements doped into LaCl₃ crystals, measured by Deike et al. [17]. The positions of the energy levels for other hosts are similar, since the main contribution is from spin-orbit interaction and not crystal-field interactions. Figure reproduced from [18].




Figure 2.2. Panel **a** shows the hyperfine structure for the $^3H_4 \rightarrow ^1D_2$ transition in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$, and **b** the relative oscillator strength for the optical transitions between these hyperfine levels.

The coherence properties of rare-earth-metal-ion-doped crystals have been extensively studied using both hole-burning spectroscopy and coherent transient methods. An excellent review of this subject is given by Macfarlane and Shelby [19].

2.2 Homogeneous line-width

The homogeneous line-width, Γ_h , is the line-width of the single ion, and for a given transition it is the same for all ions in the crystal. The line-width, i.e. how accurately the frequency of the emitter can be determined, is ultimately given by the amount of time the emitter radiates before losing its phase. The homogeneous line-width can be calculated from the coherence time, T_2 , by:

$$\Gamma_h = \frac{1}{\pi T_2}. \quad (2.1)$$

The maximum coherence time is twice the lifetime, T_1 , of the radiating state. The lifetime of the 4f electrons can be over 10 ms [20], which is 6 orders of magnitude longer than the typical lifetime for other electronic states. This allows for very long coherence times and up to 6.4 ms has been reported in $\text{Er}^{3+}:\text{Y}_2\text{SiO}_5$ [10].

2.3 Hyperfine levels

For rare-earth elements with a nuclear spin greater than 1/2, there will be hyperfine splitting in both the electronic ground and excited state, even without an external field. Such hyperfine levels can be very long-lived. For example $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$, has a hyperfine lifetime of several days at sufficiently low temperatures, and $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ approximately 100 s at 4 K. The hyperfine splitting in these two crystals is of the order of tens of megahertz. The hyperfine levels are numbered $|\pm 1/2\rangle$, $|\pm 3/2\rangle$ and so on, according to I_z [19], even though this might not be a good quantum number in the crystal. As an example, Figure 2.2a shows the hyperfine level structure for $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$, which was studied in Paper II. If an external magnetic field is applied, the degeneracy between the states with positive and negative I_z projection will be broken, and each level will be split into two. This also means that elements with a nuclear spin of 1/2 will exhibit hyperfine splitting in the presence of a magnetic field.

The coherence time of the hyperfine transitions is important for the quantum computing scheme described in Chapter 5, since the quantum state is encoded in the hyperfine states. It has been shown that the coherence time of the hyperfine transitions can be prolonged by decoupling the effects of spin flips in the host, by applying a carefully chosen magnetic field, and values up to 860 ms have been achieved in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ [21, 22]. By applying

so-called “bang-bang” pulses, which swaps the states back and forth, the coherence time has been extended even further, up to 30 s [21].

2.4 Inhomogeneous broadening

The resonance frequency of one particular ion inside the crystal can be very well defined, resulting in a very narrow line-width, and this is referred to as the homogeneous line-width, Γ_h , as discussed in Section 2.2. For example, $\Gamma_h = 122$ Hz has been reported in $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$ [23]. But even the most carefully grown crystal is not perfect, and different ions are located in slightly different surroundings, as schematically illustrated in Figure 2.3a. The consequence of this is that the resonance frequency of each particular ion is shifted randomly, as shown in Figure 2.3b, leading to an absorption line which is much broader than the homogeneous line-width. This is referred to as the inhomogeneous line-width, Γ_{inh} .

The inhomogeneous line-width varies from material to material, with typical values ranging from 0.5 GHz to 200 GHz [10]. It also varies with dopant concentration. Let us again consider the crystal mentioned above, $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$, site 1. At very low concentrations it has a residual dopant-concentration-independent line-width of 0.63 GHz, which then increases linearly with increasing dopant concentration at (21 GHz/unitconcentration in molar %) leading to a 150 GHz line-width at a concentration of 7 % [24]. Similar data, although with a residual line-width of 1.8 GHz, have also been reported by others [25]. Other factors that can affect the line-width are growth method, crystal quality and external static pressure on the crystal. The latter gave rise to a spectacular phenomenon in our frequency locking-crystal, which has a diameter of 19 mm, and is mounted in a copper holder, which causes stress on the crystal. Before it enters the crystal, as the dye laser is scanned over the inhomogeneous profile, the beam is completely even. But when it is viewed after passing the crystal, it looks as if a dark cloud moves across it as the laser is scanned. This was interpreted as both broadening and shifting of the inhomogeneous line profile occurs, depending of the local pressure in different parts of the crystal. To decrease this effect we decreased the pressure exerted by the copper holder on the crystal.

2.5 Transition strengths

The transition strength defines how easy it is to drive a transition. The transitions in rare-earth-metal ions are weak compared with allowed electronic transitions. Typical values for the oscillator strength of these transitions range from $f = 1.3 \times 10^{-8}$ for $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$ [24] to $f = 8 \times 10^{-6}$ for $\text{Nd}^{3+}:\text{YVO}_4$ [10]. The




Figure 2.3. Each ion has a very narrow line-width, called the homogeneous line-width, Γ_h . However, the surroundings in the crystal, around each particular ion, vary slightly, as depicted in a. This will shift the resonance frequency of each ion differently, leading to a broader total absorption profile, which is called the inhomogeneous line-width, Γ_{inh} , shown in b.




Figure 2.4. Panels **a** and **b** show two dopant ions located close to each other. Adopting the terminology of quantum computing the left ion is called the control and the ion on the right is called the target. Panel **c** schematically illustrates how the E-field surrounding the control ion changes as it is excited, as the ions have different static dipole moment in their ground and excited states. As a consequence, the resonant frequency of the target ion, shown in panel **d**, will be Stark shifted $\Delta\nu$ by the change in the E-field.

latter material was found when searching for rare-earth materials with high oscillator strengths. A high transition strength means, for example, that a crystal with a given dopant concentration and thickness will have a comparably high absorption, and that the transition can be driven in a shorter time with a given amount of laser power. On the other hand, the high transition strength also leads to greater spontaneous emission and thus a shorter lifetime of the upper state. It can be shown that if the coherence time is limited by the lifetime of the upper state, a larger number of operations can be carried out during the coherence time, with low transition strength and long upper state lifetime, for a given laser power [26]. This means that for the quantum computing schemes a low transition strength is favourable, as long as the laser power is limited.

As discussed in Section 2.3, and exemplified in Figure 2.2a, many of the rare-earth-metal ions have hyperfine structure. The different hyperfine transitions have different oscillator strengths, but this can in many cases be difficult to measure, since the individual transitions are “buried” under the inhomogeneous broadening. In Paper II a technique using advanced hole burning is described, through which it is possible to isolate and measure the properties of individual hyperfine levels. As an example, the relative oscillator strength measured for the different hyperfine transitions of the $^3H_4 \rightarrow ^1D_2$ transition in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ is shown in Figure 2.2b. (The table has been slightly altered compared with the one in Paper II, so that both the row and column sums are equal to one.)

A very useful article when working with these quantities is *Einstein coefficients, cross-sections, f values, dipole moments, and all that*, by Robert Hilborn [27].

2.6 Ion-ion interactions

In the last section of this chapter we will consider the electric dipole-dipole interaction in these materials. The dopant ions, in for example $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$, have a static dipole moment. This means that there is a static electric dipole field surrounding each ion. Furthermore, this electric dipole moment is different for an ion in a ground state, $|g\rangle$, and in the optically excited state, $|e\rangle$. This means that when an ion is excited, the E-field in its surroundings will change. In Figure 2.4a an ion is depicted in its ground state. We now adopt the terminology that will be used in Chapter 5 concerning quantum computing with rare-earth-metal-ion-doped crystals, calling this the control ion. In panel b we see another ion which is located very close to the first one, denoted the target ion. If the control ion is transferred to its optically excited state, the electric field (E-field) around it will change, as is illustrated in c.

When the E-field around the target ion changes, this will cause a Stark shift of the target ion's transition frequency which will be shifted $\Delta\nu$.

The quantum computing scheme our group is working on, uses electric dipole-dipole interaction to create quantum gates with multiple ions [28]. Studies of the dipole-dipole interaction are therefore described in Papers I and III.

SPECTRAL HOLE BURNING

In most materials and for most transitions, the absorption is distributed over a frequency range which is much broader than the frequency width of each individual absorber, because the centre frequency of each individual absorber varies. This is referred to as an inhomogeneously broadened transition. Some inhomogeneously broadened materials exhibit a spectral hole-burning mechanism. If such a material, is illuminated with a light source which is spectrally narrow compared with the inhomogeneous line-width, it will be bleached and show stronger transmission in the spectral region that is illuminated. This is referred to as the spectral-hole. The lifetime of the spectral hole depends on the physical mechanism responsible for the hole burning mechanism. If the frequency of the laser is scanned across the spectral hole within the lifetime of the hole, a dip in the absorption will be seen, i.e. an increase in transmission, as the laser frequency passes the spectral hole.

A general overview of science and applications of persistent spectral hole burning has been written by Moerner et al. [29], and two good reviews on hole burning in rare-earth-metal-ion-doped crystals have been written by Macfarlane and Shelby [19] and by Macfarlane [30].

In Paper II an advanced form of spectral hole burning is described which allowed single hyperfine transitions to be isolated.

3.1 Hole-burning mechanisms

When a spectral hole is created, some of the absorbers that have their resonance frequency at the hole frequency, for some reason stop absorbing. A number of different processes can be responsible for this, and we will consider three that are relevant to the present work. The first one includes only two levels, the ground state, $|g\rangle$, and the optically excited state, $|e\rangle$, as depicted in Figure 3.1a.




Figure 3.1. Three hole-burning situations, with different storage mechanisms. In **a** the ion is stored in the optically excited state, in **b** in an intermediate reservoir state, and finally in **c** storage takes place in another ground state hyperfine level.




Figure 3.2. Hyperfine levels for $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$. The discussion will be simplified by considering only one hyperfine level in the optically excited state, as indicated by the crosses.

Some of the population at the hole-burning laser frequency, indicated by an arrow, will be transferred to the excited state, leading to a decrease in absorption at this frequency. A hole of this type, has for example, been used to frequency lock a laser [31]. A second storage mechanism is depicted in panel **b**, where there is an intermediate reservoir state, $|r\rangle$, with a longer lifetime than the excited state between the excited state and the ground state. After laser excitation from the ground state to the excited state, there is some probability of decay to the intermediate state where the ion is trapped. This is the case, for example, in $\text{Tm}^{3+}:\text{Y}_3\text{Al}_5\text{O}_{12}$, where the metastable 3F_4 level acts as a reservoir with a lifetime of ~ 10 ms [32]. The third storage mechanism can occur when the ground state has several hyperfine levels as illustrated in panel **c**, where, after excitation to the optically excited state, the ions has some probability of relaxing to a different hyperfine state. This storage mechanism is present in some of the crystals studied in this work, namely $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$, which has a hyperfine lifetime of ~ 100 s, and $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$ with a hyperfine lifetime of up to several days. The hyperfine splitting in these two crystals is of the order of tens of megahertz, which is much less than their inhomogeneous broadening. This means that if the ion decays to another hyperfine level, as is depicted in panel **c**, there will be an *increase* in absorption corresponding to the resonance frequency from that hyperfine level. This is often referred to as an *anti-hole*. If there are multiple hyperfine levels in the optically excited state, there will also be side holes with decreased absorption, just as for the centre hole. The structure of holes and anti-holes is often used to reveal information about hyperfine structure.

3.2 Isolating a peak of ions in a single hyperfine state

Let us consider $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$, which has three hyperfine levels in both the ground and the excited state. To simplify the discussion we will ignore the hyperfine structure of the excited state, as shown in Figure 3.2. If a laser illuminates one frequency within the inhomogeneous broadening, for a time much longer than the lifetime of the excited state, ions with a transition frequency equal to the laser frequency will be excited. If they decay back to the same ground state hyperfine level, they will be excited again, until they decay to a different ground state hyperfine level, where they are no longer resonant with the exciting laser, and will therefore remain there. In this way it is possible to create a spectral hole with almost no absorption. If the laser is instead scanned forward and back in frequency, as indicated in Figure 3.3a, the same mechanism will create a wide pit rather than a narrow hole. The pit can of course not be infinitely wide, as the ion must reside in one of the three ground state hyperfine levels. If the laser is then turned on with a shifted frequency, as illustrated in Figure 3.3b, this will pump the ions away from this frequency, and they will eventually revert to the other two ground state hyperfine levels. If the frequency is chosen such that one of these hyperfine levels is located inside the pit, in the case of $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$, for example, 17.3 MHz away, a peak will appear, and all the ions in this peak will be in the same ground state hyperfine level.

For clarity this description was somewhat simplified. A thorough description of this technique, which was developed by our group, can be found in Paper I.

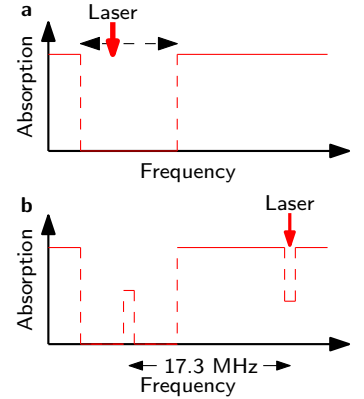


Figure 3.3. If a laser is scanned forward and back within a frequency interval, it is possible to create a wide pit with very little absorption, **a**. If the frequency distance is chosen correctly, it is possible to move back a peak of ions, all in the same ground state hyperfine level, **b**.

COHERENT MANIPULATION

In our research much of the equipment used is very tangible, such as mirrors, holders, optical tables, vacuum equipment and lasers. When it comes to the light pulses used to program manipulate the ions, detectors and oscilloscopes are necessary, but this is still fairly straightforward. However, it is very difficult to visualise the effect that the light pulses have on the atoms in the cooled samples. You can not see this effect. I have found it useful to develop an intuitive model of the effect the laser field has on the atoms. To this end I found the Bloch sphere and the Bloch equations to be very useful. This chapter starts with an illustrative description of the interaction between light and atoms, followed by a description of how the state of the ion can be described using the Bloch sphere and finally the Bloch equation, which provides an accurate tool for predicting the effect of light on the state will be presented.

4.1 Two-level systems

This discussion is in no way intended to be quantum mechanically stringent. A rigorous quantum mechanical treatment, starting with the light-atom Hamiltonian using the density matrix formalism can, for example, be found in [33] or [34].

4.1.1 Field-induced electron displacement

Imagine the electron as a spherical cloud around the nucleus, as in Figure 4.1a. This cloud is referred to as a wave function, and the square of its absolute value gives the probability distribution for the electron, i.e. there is a higher probability of finding the electron where the cloud is dense. Quantum mechanics tells us that if this is a base state, the probability distribution will not vary with time, which simply means that whenever we measure, the probability of finding the electron somewhere is the same. Let us call this




Figure 4.1. Panels **a**, and **c** show the wave function for the ground state, $|g\rangle$, and the excited state, $|e\rangle$ of an atom. When an E -field is applied, the electron cloud will be displaced as shown in **b**. This displacement can be described as a superposition of $|g\rangle$ and $|e\rangle$, as depicted in **d** and **e**.




Figure 4.2. An atom in a homogeneous E -field, \mathbf{E} . The electron and nuclear will be slightly displaced a distance, \mathbf{a} , creating a dipole with dipole moment, \mathbf{p} . The dipole will in turn induce a E -field \mathbf{E}_d .

cloud $|g\rangle$, denoting the ground state. If an electric field (E -field) is applied, the electron will be affected since it is charged, and the cloud will be displaced, as depicted in Figure 4.1**b**. The charge displacement will in turn create a dipole, with a dipole moment, \mathbf{p} , see Figure 4.2. The amount of dipole moment per unit volume is called the polarisation of the material, \mathbf{P} , not to be confused with the polarisation of light. The charge displacement can be described quantum mechanically as a mixture of many base states. In this discussion we will only consider the effect of the two base states: $|g\rangle$, which is depicted in Figure 4.1**a**, and an excited state, $|e\rangle$, shown in Figure 4.1**c** where the top part is negative and the bottom part is positive. This has no effect on the probability distribution, as the wave function is squared when this is calculated. A system with two levels, or base states, $|g\rangle$ and $|e\rangle$ is often referred to as a two-level system. We will now add these two states together, to form a superposition of the two states, in a way which will allow us to somewhat mimic the charge displacement seen in Figure 4.1**b**. This is done including a time dependence which will be used later. We can write this superposition as:

$$|\psi\rangle = \alpha |e\rangle + e^{i\omega t} \beta |g\rangle, \quad (4.1)$$

where α and β are related to the composition of the superposition through $|\alpha|^2$ and $|\beta|^2$ which give the probability of finding the system in level $|e\rangle$ and $|g\rangle$, respectively. They have to obey the relation:

$$|\alpha|^2 + |\beta|^2 = 1, \quad (4.2)$$

which simply means that the probability of finding the electron somewhere adds up to one. The term $e^{i\omega t}$ takes into account the phase difference between the two states, where ω is the frequency of the laser, which is resonant with the transition, in our case often around 500 THz ($= 500 \times 10^{12}$ Hz). Let us insert some values for the coefficients in Equation (4.1), and see what the result is:

$$\alpha = \sqrt{\frac{1}{5}}, \quad \beta = \sqrt{\frac{4}{5}} \quad \text{and} \quad \omega t = 0 + n \times 2\pi \Rightarrow$$

$$|\psi\rangle = \sqrt{\frac{1}{5}}|e\rangle + \sqrt{\frac{4}{5}}|g\rangle, \quad (4.3)$$

n is an integer, which means that $e^{i\omega t} = 1$ for all n . As can be seen in Figure 4.1d this gives a downward displacement of the electron cloud since the lower part of the wave functions have the same sign and therefore interfere constructively, while the top halves have opposite signs and will consequently partly cancel each other out. On the other hand, setting $\omega t = \pi + n \times 2\pi$ means that $e^{i\omega t} = -1$ and

$$|\psi\rangle = \sqrt{\frac{1}{5}}|e\rangle - \sqrt{\frac{4}{5}}|g\rangle, \quad (4.4)$$

which means that the ground state wave function will have changed sign. This leads to constructive interference for the upper half, and destructive for the lower half which corresponds to a displacement upwards, as can be seen in Figure 4.1e. This means that there is a material polarisation that oscillates with angular frequency ω , as time progresses. To recapitulate: an E-field causes an electron cloud displacement; this displacement gives rise to a polarisation in the material and finally, in the two-level case, the displacement and thus the polarisation will oscillate at the resonance frequency of the transition.

4.1.2 Time varying fields

Next we will consider what happens when light with a frequency that is below, on and above resonance interacts with the two-level system, in a conceptual way. A model that can accurately predict the interaction between light and a two-level system is introduced in Section 4.3. Since we are now describing displacement in terms of adding part of the excited state wave function, we could say that the population is partly transferred. It should first be pointed out that the period of the light at a frequency of 500 THz is only 2 fs




Figure 4.3. Phase, dashed red, and amplitude, solid black, for an oscillator driven below, (a), on, (b), and above its resonance frequency, (c).




Figure 4.4. At resonance, the electron displacement (black line) is 90° after the driving E-field (red arrows). This is the most efficient way of transferring energy from the field to the electron. During the time the field is positive the electron is pushed downwards (a), and vice versa (b).

($= 2 \times 10^{-15}$ s), and only a very small part of the population can be transferred during this time.

When the driving field oscillates at a frequency that is far below the resonance frequency of the two-level system, ω , the polarisation will always want to go faster than the driving field, and the nett effect is a displacement (and polarisation) that has the same phase as the driving field, as we show in Section 4.3. This is shown in Figure 4.3a.

Two things take place at resonance. Firstly, the polarisation will be 90° after the driving field, which is the most efficient way of transferring energy, as illustrated in Figure 4.4. Secondly, the phase relationship between the polarisation and the driving field is the same from cycle to cycle, which means that they can interfere positively for a long time and a larger and larger part will be transferred to the excited state. This is of course exactly the same as in any other oscillator driven at resonance. According to Equation (4.1), the oscillations will become large as β starts to increase, as shown in Figure 4.3b. However, this picture then breaks down, as there is nothing that explains what will happen when β grows increases further.

When the driving field is far above resonance, the polarisation lags behind, and will approach 180° behind in phase. Again there is no polarisation build-up, and the amplitude is low (see Figure 4.3c).

After this simplified description we will raise the level of abstraction slightly, to obtain a model that can actually be used to calculate the behaviour of the two-level system. This is called the Bloch model.

4.2 The Bloch sphere

We will now consider a description of a two-level system which, allows us to visualise the state of the system, describe how it will develop with time and the effect light has on the system. Recalling that $|\alpha|^2 + |\beta|^2 = 1$ from Equation (4.2), Equation (4.1) can be rewritten as

$$|\psi\rangle = \cos \frac{\theta}{2} |e\rangle + e^{i\phi} \sin \frac{\theta}{2} |g\rangle, \quad (4.5)$$

ignoring the global phase factor¹. θ and ϕ are a real numbers, and are the co-ordinates in a normal spherical co-ordinate system, as shown in Figure 4.5, where the pure state of the system is represented by a vector pointing on the unit sphere, known as the Bloch sphere². The two base states $|g\rangle$ and $|e\rangle$ are now represented

¹If α or β is complex there can also be a global phase factor, but as explained in a standard textbook on introductory quantum mechanics or quantum information [35], this has no observable effect, and can be dropped.

²Different labels for the coordinates are widely used: x, y and z ; u, v and w ; and r_1, r_2 and r_3 . Here u, v and w will be used.

by vectors pointing to the south and north pole of the Bloch sphere respectively. A vector pointing straight at the equator corresponds to a 50-50 superposition of $|g\rangle$ and $|e\rangle$. In Section 4.1.1 the phase between the two base states governed the displacement up and down of the electron cloud. This phase is now represented by ϕ , which for a 50-50 superposition specifies the position on the equator. So far, we have only discussed states which lie on the surface of the unit sphere. This represents states that are fully known, i.e. pure states. A state that is not fully known is called a mixed state. These will be discussed further in Section 4.3. These states can also be described with the Bloch sphere, and are simply represented by a vector pointing somewhere inside the sphere, i.e. which has a length less than one.

In the next section we describe state changes in the Bloch sphere under different conditions. This is done in Cartesian coordinates, where the vector describing the state in the abstract $(\hat{\mathbf{e}}_u, \hat{\mathbf{e}}_v, \hat{\mathbf{e}}_w)$ space is:

$$\mathbf{R} = (u, v, w), \quad (4.6)$$

and is referred to as the Bloch vector.

4.3 The Bloch equations

The Bloch equations describe how the state vector moves inside the Bloch sphere. They were introduced in 1946 by Felix Bloch [36] to describe the movement of nuclear spins acted upon by radio-frequency electromagnetic fields; what today is called nuclear magnetic resonance (NMR). The Bloch sphere was later suggested for use to describe optical resonances by R. P. Feynman, F. L. Vernon and R. W. Hellwarth in 1957 [37].

The pictorial description in Section 4.1.1 suggested, that, depending on the differences between the wave functions, different transitions would have different susceptibility of being put into oscillation by an E-field. The factor describing this susceptibility is called the dipole moment, μ .³ Using μ it is possible to define the (angular) Rabi frequency, Ω , in rad/s, as

$$\Omega(t) = \frac{\mu \mathcal{E}(t)}{\hbar}, \quad (4.7)$$

a quantity that describes how fast the field makes the Bloch vector spin. Here \mathcal{E} is defined as a complex E-field, so that phase variations can be included, and μ is the dipole moment along \mathcal{E} . The real and imaginary parts of Ω are separated out, for inclusion in the real Bloch equations as

$$\Omega = \Omega_{\text{re}} + i \Omega_{\text{im}}. \quad (4.8)$$

³If one wishes to control the Bloch vector it is important to know the value of μ , and this was therefore measured for the $^3H_4 - ^1D_2$ hyperfine transitions in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ (Paper II).




Figure 4.5. In the Bloch sphere, the state of the two-level pure system is represented as a vector pointing to the surface of the unit sphere. A mixed state is described by a vector pointing inside the sphere. A vector pointing at the north pole of the Bloch sphere corresponds to a system in the excited state and on the south pole the ground state and a vector pointing in between represents a mixture between the two states.




Figure 4.6. The torque vector decomposed into its three basic components.

For a single-frequency field, with only amplitude variations, $\Omega_{\text{im}} = 0$ may be set equal to zero to simplify the equations. We wish to include the contributions from atoms and ions with different resonance frequencies. u , v and w are therefore described as functions of both time, t , and detuning, Δ (rad/s), which is the frequency difference between the resonance frequency of the atom or ion and an arbitrarily chosen reference frequency.

The full movement of the Bloch vector, \mathbf{R} , can now be described by

$$\frac{\partial \mathbf{R}}{\partial t} = \mathbf{\Omega}_T \times \mathbf{R} \quad (4.9)$$

which simply states that \mathbf{R} rotates around the “torque” vector, $\mathbf{\Omega}_T$, defined by

$$\mathbf{\Omega}_T = (-\Omega_{\text{re}}, \Omega_{\text{im}}, \Delta). \quad (4.10)$$

The frequency of rotation is given by the length of the vector,

$$\Omega_G = \sqrt{\Omega_{\text{re}}^2 + \Omega_{\text{im}}^2 + \Delta^2} \quad (4.11)$$

which is often called the generalised Rabi frequency. The rotation is now decomposed into three different components, as shown in Figure 4.6. The following definitions are made for convenience:

$$a\mathbf{\Omega}_{\text{re}} = (-\Omega_{\text{re}}, 0, 0) \quad (4.12)$$

$$\mathbf{\Omega}_{\text{im}} = (0, \Omega_{\text{im}}, 0) \quad (4.13)$$

$$\mathbf{\Delta} = (0, 0, \Delta). \quad (4.14)$$

Figure 4.6a describes rotation around w with a magnitude given by Δ . This means that if we consider an atom or ion that is detuned from the arbitrarily chosen reference frequency, it will rotate around the w -axis at a constant angular speed, which seems completely reasonable. Figure 4.6b and c show the rotation vectors for the real and imaginary parts of the Rabi frequency. Figure 4.7 shows the rotation around a torque vector including all three components.




Figure 4.7. The Bloch vector, $\mathbf{R}(t)$, rotating around the torque vector $\mathbf{\Omega}_T$

4.3.1 Including decay and decoherence

The next step is to include the influence of the limited upper-state lifetime, called T_1 , and a limited coherence time, denoted T_2 . The Bloch equations can now be written:

$$\frac{\partial \mathbf{R}}{\partial t} = \Omega_T \times \mathbf{R} - \left(\frac{1}{T_2} u, \frac{1}{T_2} v, \frac{1}{T_1} (w - w_0) \right), \quad (4.15)$$

where w_0 is the equilibrium state. In the cases considered here, most states decay to the ground state, $|g\rangle$, and consequently $w_0 = -1$. To describe it with words, if $w_0 = -1$, T_1 makes the Bloch vector relax downwards in the Bloch sphere towards $w = -1$, and T_2 relaxes the state in the uv -plane towards the w -axis. This also explains why T_1 and T_2 are often denoted longitudinal and transverse relaxation times, respectively.

4.3.2 Bloch equations in different forms

The Bloch equations can be written in different forms, which are useful for different purposes. If the cross product is calculated, and the three different directions are separated, Equation (4.15) can be written as

$$\frac{\partial u}{\partial t} = -\frac{1}{T_2} u - \Delta v + \Omega_{\text{im}} w, \quad (4.16)$$

$$\frac{\partial v}{\partial t} = \Delta u - \frac{1}{T_2} v + \Omega_{\text{re}} w, \quad (4.17)$$

$$\frac{\partial w}{\partial t} = -\Omega_{\text{im}} u - \Omega_{\text{re}} v - \frac{1}{T_1} (w_0 + w). \quad (4.18)$$

For calculations it is often useful to write them in complex form,

$$\frac{\partial}{\partial t} (u - iv) = -\left(\frac{1}{T_2} + i\Delta\right)(u - iv) - i\Omega w, \quad (4.19)$$

$$\frac{\partial}{\partial t} w = \frac{i}{2} [\Omega(u + iv) - \Omega^*(u - iv)] - \frac{1}{T_1} (1 + w). \quad (4.20)$$

4.3.3 Pulse area

Considering the case when the resonance frequency of the atoms and the driving laser field both coincide with the chosen reference frequency, and their transition is being driven with a single real frequency. This is basically the case illustrated in Figure 4.6b. If the Rabi frequency is constant, the angle through which the Bloch vector is turned is then given by:

$$\Theta = \Omega_{\text{re}} \times t. \quad (4.21)$$




Figure 4.8. Evolution of a Bloch vector, starting in $|g\rangle$, subjected to pulses with different pulse areas. These pulses are often referred to as $\frac{\pi}{2}$ -pulses (a), π -pulses (b) and 2π -pulses (c).

Θ in this context is often called the pulse area. For a field with a slowly varying envelope the pulse area is given by

$$\Theta = \int_0^t \Omega_{\text{re}}(t') dt'. \quad (4.22)$$

In Figure 4.8 the evolution of the Bloch vector for pulses according to the above with different pulse areas is displayed. These pulses are commonly used in coherent transients⁴ and are often referred to as $\frac{\pi}{2}$ -pulses, π -pulses and 2π -pulses. The $\frac{\pi}{2}$ -pulse starting in the ground state results in a 50-50 superposition. The π -pulse creates a so-called inverted population as it transfers the population from the ground state to the excited state.

Gaussian pulse

The Gaussian pulse is commonly used since it is narrow both in time and frequency. Therefore, it will be used in an example of calculating the pulse area. The intensity envelope of the Gaussian pulse is given by

$$I(t) = I_0 \exp\left(\frac{-t^2}{\sigma^2}\right) \quad (4.23)$$

where I_0 is the intensity on the peak of the pulse, t = time and

$$\sigma^2 = \frac{t_{\text{FWHM}}^2}{4 \ln 2}, \quad (4.24)$$

where t_{FWHM} is the pulse length measured at the full width half maximum of the intensity. As

$$I \propto \mathcal{E}^2 \propto \Omega^2, \quad (4.25)$$

⁴Coherent transients is a collective name for several techniques such as spatial-spectral holography, time-domain spectral hole burning and photon echoes

where E is the E-field and Ω the Rabi frequency, Eq. (4.23) for Ω (in rad/s) can be written as:

$$\Omega(t) = \Omega_0 \sqrt{\exp\left(\frac{-t^2}{\sigma^2}\right)} = \Omega_0 \exp\left(\frac{-t^2}{2\sigma^2}\right). \quad (4.26)$$

The total angle, in the Bloch sphere, through which the ions rotate, i.e. the pulse area in radians is:

$$\Theta = \int_{-\infty}^{\infty} \Omega(t) dt = \int_{-\infty}^{\infty} \Omega_0 \exp\left(\frac{-t^2}{2\sigma^2}\right) dt. \quad (4.27)$$

This can be calculated from the standard integral:

$$\int_{-\infty}^{\infty} \exp(-x^2) dx = \sqrt{\pi}. \quad (4.28)$$

Combining the result from Equations (4.27) and (4.24) allows the pulse area for a Gaussian pulse to be expressed as:

$$\Theta = \Omega_0 \sqrt{2\pi} \cdot \sigma = \Omega_0 \sqrt{\frac{\pi}{2 \ln 2}} \cdot t_{\text{FWHM}}. \quad (4.29)$$

where t_{FWHM} refers to the intensity FWHM.

4.4 Maxwell-Bloch equations

So far we have only considered the effect of light on the state of the system. But, as has been show above a state in a superposition corresponds to an oscillating electron cloud, and oscillating charges radiate electromagnetic waves. The effect on the electromagnetic field due to the atomic response can be taken into account effectively, if the Bloch equations are expanded, to the so-called Maxwell-Bloch equations. Since the system is radiating, the electromagnetic field will be altered as it travels through the sample, and it is necessary to make Ω dependent on the position along the pulse propagation axis, which we define as being the z direction, i.e. $\Omega(z, t)$. But if Ω changes, it will not have the same effect on u , v and w , and they must also be dependent on the distance traveled through the sample, i.e. $u(z, t, \Delta)$, $v(z, t, \Delta)$ and $w(z, t, \Delta)$. It is of course also possible to take into account different radial distributions, but here we will simply discuss plane waves. The Maxwell-Bloch equations can be expressed as follows:

$$\left(\frac{\partial}{\partial z} + \frac{n}{c} \frac{\partial}{\partial t}\right) \Omega_{\text{re}} = \frac{\alpha_0}{2\pi} \int_{-\infty}^{\infty} g(\Delta) v d\Delta, \quad (4.30)$$

$$\left(\frac{\partial}{\partial z} + \frac{n}{c} \frac{\partial}{\partial t}\right) \Omega_{\text{im}} = \frac{\alpha_0}{2\pi} \int_{-\infty}^{\infty} g(\Delta) u d\Delta, \quad (4.31)$$




Figure 4.9. The frequency of the driving laser field is below the resonance frequency of the atoms (Δ).




Figure 4.10. The driving frequency of the laser is below the resonance frequency of the atoms. The Bloch vector is spiralling around the torque vector.




Figure 4.11. The frequency of the driving laser is above the resonance frequency of the atom. The Bloch vector is spiralling around the torque vector

where n is the index of refraction for the non-absorbing background atoms, c is the speed of light, α_0 is the absorption coefficient and $g(\Delta)$ is a function that takes into account variations in the absorption as a function of detuning, and is normalised so that it is equal to one where α_0 is measured. The term $\frac{n}{c} \frac{\partial}{\partial t}$ only contributes when the changes are fast, on a timescale of $\frac{L}{c}$, where L is the length of the sample. This is never the case in our experiments, and the term will therefore be ignored. Let us again consider the example in Figure 4.8a. The Bloch vector in this example points straight along the negative v direction, $\mathbf{R} = (0, -1, 0)$. If we only considered the atoms near resonance and insert this into Equation (4.30) a decrease in Ω_{re} is observed. This corresponds to absorption, which is what we would expect. Finally, the Maxwell-Bloch equation is given in its complex form.

$$\left(\frac{\partial}{\partial z} + \frac{n}{c} \frac{\partial}{\partial t} \right) \Omega = \frac{i\alpha_0}{2\pi} \int_{-\infty}^{\infty} g(\Delta) (u - iv) d\Delta \quad (4.32)$$

4.4.1 Steady-state solutions

The succeeding sections consider how the Bloch vector is transferred from the south-pole of the Bloch sphere to the steady-state solution when a weak E-field is suddenly applied. The following conditions are thus assumed: a weak driving field, corresponding to a low Rabi frequency, and in this context this means that the Bloch vector is turned considerably less than a quarter of a turn in the lifetime of the upper state, on resonance. This conditions ensures that the tip of the Bloch vector remains close to the south pole in the Bloch sphere. If this demand is not met effects such as Rabi oscillations can occur. Furthermore only the case where Ω_{re} is constant and positive, and $\Omega_{\text{im}} = 0$, will be considered. The steady-state solution can, of course, easily be analytically calculated by calculating the steady-state solution to Equations (4.16) to (4.18) with a constant driving Rabi frequency.

On resonance

Since only Ω_{re} is being considered, it can be seen in Figure 4.6b that the Bloch vector will move towards negative values of v , where, under the current restrictions, it will come into steady state with the decay. According to equation Equation (4.30) this will decrease the value of Ω_{re} , which is again just absorption.

Far below resonance

Let us assume that the driving field is at $\Delta = 0$, which follows from the assumption that the driving field is real. The resonance frequency of the transition is located at a positive frequency as shown in Figure 4.9. From Figures 4.6a, b and 4.7 it can be

deduced that the Bloch vector will rotate as shown in Figure 4.10. Note, that even though the Rabi frequency is low compared to the inverse of the decay ($1/T_1$), the generalised Rabi frequency, given by Equation (4.11), can still be high, and this dictates how fast the Bloch vector spins around the torque vector. When the field is initially turned on the Bloch vector will spin around the circle, but the decay will make it spiral towards the centre of the circle, which is the equilibrium state according to the Bloch equations, which means that the nett polarisation will be pointing in the positive u direction. If we think of the decay as a discrete process, the decay would at random bring the state of the atoms to the south pole of the Bloch sphere, from where they would start rotating around the circle, and summing there individual contributions to the nett polarisation yields a vector pointing in the centre of the circle. According to Equation (4.31) this will cause a positive Ω_{im} field, which is equivalent to saying that in this case the phase of the field emitted by the atoms is 90° before the phase of the incoming field. When the incoming field and the field emitted by the atoms are added, a phase shift results such that the phase of the light leads after passing the atoms, compared to the case when no atoms are present.

Far above resonance

The same reasoning applies here as in Section 4.4.1, but in this case the situation is as illustrated in Figure 4.11. In this case is the field emitted by the atoms 90° after the driving field. When the incoming field and the field emitted by the atoms are added, this again results is a phase shift but this time such that the phase of the light lags after passing the atoms, compared to the case when no atoms are present.

4.5 Free induction decay

If a crystal that contains inhomogeneously broadened absorbers, having long coherence times, is illuminated, and the light is suddenly turned off, the ions will keep on emitting coherent light. This phenomenon is called free induction decay, FID. Section 4.4.1 discussed the Bloch vector reaches its steady state value, where especially Figures 4.10 and 4.11 showed the rotation of the Bloch vector when exposed to a off-resonant field. Let us study this in more detail. Figure 4.12 shows the region close to the bottom of the Bloch sphere from above, when $w \approx -1$. The lines show the transient response of the tip of the Bloch vector, starting in the ground state. The cross corresponds to the steady-state solution. If the driving field has a considerably lower frequency than the resonance of the system, the spiral will look as in panel **a**. If the driving field is slightly beneath or above resonance it will have the




Figure 4.12. Transient response of the tip of the Bloch vector for different resonance frequencies of the transition. For details see the text.

appearance shown in **b** and **c** respectively. Finally if the frequency of the field is above the the resonance frequency of the transition it will behave as shown in Figure 4.12. The steady-state values, in the case of small Rabi frequencies compared to the inverse coherence time, i.e. $\Omega T_2 \ll 1$, can easily be solved analytically by setting the time derivatives equal to zero and $w = -1$ in Equation (4.19). As previously, we assume that $\Omega = \Omega_{\text{re}}$ for simplicity, which gives:

$$u(\Delta) = \frac{\Delta \Omega_{\text{re}}}{\frac{1}{T_2^2} + \Delta^2}, \quad v(\Delta) = -\frac{\frac{\Omega_{\text{re}}}{T_2}}{\frac{1}{T_2^2} + \Delta^2}. \quad (4.33)$$

These functions are parametrically plotted, as functions of Δ in Figure 4.13. As can be seen in this figure, and can easily be shown from Equation (4.33), the steady-state values forms a circle. If u and v are instead plotted in two separate graphs as functions of Δ directly, it is easier to see how the polarisation varies with detuning.

Let us return to the problem of what happens when the field is suddenly turned off. We will first look at two points, corresponding to the tips of the Bloch vectors, with equal absolute value of

frequency detuning from the field, one above and one below, as in Figure 4.14. These points are stationary in this steady state solution, but they have different Δ which means that as soon as the field is turned off they will start rotating around w , as indicated by the arrows. Since their detuning is the same but with opposite signs, they will rotate with equal speed but in opposite directions. Consequently, the total u -polarisation will be zero, while this is not the case for the v -polarisation. This symmetry can be expanded to include all values of detuning, all ions can be grouped into pairs as was done here. Since we have concluded that there will be no net u -polarisation, let us concentrate on the v -polarisation. When the field is turned off, the Bloch vectors start to rotate at an angular speed determined by the detuning. They rotate from the negative v half-plane into the positive half-plane, and eventually back again. The field emitted by the ions can be calculated with Equation (4.30). This is shown in Paper V, for the case when the field is suddenly turned off. If we are in the linear regime, which simplified means that the polarisation is proportional to the pulse area at that particular frequency, or that the tip of the Bloch vector is very close to the bottom of the Bloch sphere and thus there is no inversion, i.e. $w = -1$, the calculation in Paper V shows that no light will be emitted. The reason that there still is free induction decay is that the system is *not* linear. This can be easily described in the Bloch sphere. For small deviations from the south-pole the polarisation is proportional to the angle between the Bloch vector and the z -axis. But for larger angles, as will be the case close to resonance, this does not hold, and the polarisation will grow less than linearly. There could also be another reason for the polarisation at the laser frequency to grow less than linearly. If there is a reservoir state, the ions could have been moved there, which is an effect that is proportional to the inversion, and not to the polarisation.

The larger the detuning, the faster the ions rotate on the Bloch sphere, which means that the ions with the highest detuning are the first to be rotated into the positive half-plane. Since their resonance frequencies are far away from the frequency of the driving field, their individual polarisation is of course very small. But the frequency interval that contains ions rotated into the positive half-plane is also very wide, which means that the integrated polarisation from these ions still can be considerable. The net effect of this is that there will be, an initially very wide, and then shrinking interval of ions moving into the positive half plane, and the limits for this interval changes from high detuning towards small. When there are comparably fewer ions close to the centre, this wave will cause the integral of v to be positive. Thus we can conclude from this discussion that once the driving field is turned off, the FID signal has the same phase as the former driving field. will cause the emitted light to be in phase with the former driving field.




Figure 4.13. Steady-state values for the tips of Bloch vectors, when the Rabi frequency of the driving field is small compared to the inverse of the coherence time, $\Omega T_2 \ll 1$. In this case $\Omega_{re} \times T_2 = 0.1$. The detuning, Δ changes by 3.1 krad/s, which is equal to 1 kHz, between each cross.




Figure 4.14. The crosses illustrate how the tip of the Bloch vectors of two groups of ions, with equal detuning in frequency, but with opposite signs, from the driving field, moves if the field is suddenly turned off. The polarisation in the u direction will cancel out, since one is positive and one is negative, but there will be a net v polarisation.




Figure 4.15. The amplitude envelope of a complex hyperbolic secant pulse, is a real hyperbolic secant pulse as shown above.




Figure 4.16. The frequency sweep of a complex hyperbolic secant pulse, is a real hyperbolic tangent sweep, as shown above.

Another observation can also be mentioned here. In Sections 4.4 and 4.4.1 we concluded that the light emitted by the ions, when irradiated by the field, was opposite in phase to the incoming light, which is simply one way to describe an absorption process. When the field is turned off, the ions start to emit light with the same phase as the incoming beam had before the light was switched off. Fast phenomena generally requires a large frequency bandwidth. So what is it that defines how fast the phase flips? We concluded above that the emission in phase with the former driving field was caused by the Bloch vector being turned into the positive half-plane, and the rate at which this happens for any ion is given by its frequency detuning. It is not until there exists ions with a certain frequency detuning that the emitted field can swap phase. In our case, this means that it is the width of the whole inhomogeneously broadened profile that dictates how fast this happens; the broader the profile, the faster the switch.

The decay time for the FID emission is given by the width of the hole. Finally, in practice, the pulses used in free induction decay experiments are not infinitely long, typically experiments are done with pulse lengths of about 1 μs to 30 μs .

4.6 The complex hyperbolic secant pulse

Since the systems in which this work was carried out are inhomogeneously broadened, it is important that the pulses used have the same effect on ions with slightly different detuning, on the same time as ions outside a certain frequency interval should be unaffected. This section describes the complex hyperbolic secant pulses that fulfils both those requirements. The complex pulse is given by:

$$\begin{aligned}\Omega(t) &= \Omega_0 \{\text{sech}[\beta(t - t_0)]\}^{1+i\mu} = \\ &= \Omega_0 \text{sech}[\beta(t - t_0)] e^{+i\mu \ln\{\text{sech}[\beta(t-t_0)]\}},\end{aligned}\quad (4.34)$$

where β is related to the width of a pulse, μ is a real constant and Ω_0 is the Rabi frequency at the centre of the pulse. As we see from the real part of Equation (4.34) the pulse envelope is given by,

$$\Omega(t) = \Omega_0 \text{sech}[\beta(t - t_0)]. \quad (4.35)$$

The pulse envelope is shown in Figure 4.15. The time derivative of the phase in Equation (4.34) gives the instantaneous angular frequency,

$$\frac{d}{dt} \left(-\mu \ln \{\text{sech}[\beta(t - t_0)]\} \right) = \mu\beta \tanh[\beta(t - t_0)], \quad (4.36)$$

in rad/s, using the convention that the phase of $e^{-i\phi(t)}$ is $\phi(t)$. This is shown in Figure 4.16. The effect on the state vector in the Bloch sphere can be seen in Figure 4.17.

As $I \propto \Omega^2$, Equation (4.37) we have

$$I(t) = I_0 \operatorname{sech}^2 [\beta(t - t_0)], \quad (4.37)$$

where I_0 is the intensity that gives the Rabi frequency Ω_0 . Calculating the full width half maximum of intensity, t_{FWHM} , as a function of β yields:

$$t_{\text{FWHM}} = \frac{2}{\beta} \operatorname{arcsech} \sqrt{\frac{0.5 I_0}{I_0}} \approx \frac{1.76}{\beta}, \quad (4.38)$$

and the frequency width of the sweep (Hz), which is given by:

$$\nu_{\text{width}} = \frac{\mu \beta}{\pi}. \quad (4.39)$$

4.6.1 Theoretical transfer efficiency

The transfer efficiency for complex hyperbolic secant pulses can be calculated analytically for the case of an infinitely long pulse, with infinite T_1 and T_2 [38]. The transfer efficiency is given by

$$\begin{aligned} w = & w_0 \tanh \left[\pi \left(\frac{\Delta}{2\beta} + \frac{\mu}{2} \right) \right] \tanh \left[\pi \left(\frac{\Delta}{2\beta} - \frac{\mu}{2} \right) \right] \\ & + \cos \left\{ \pi \left[\left(\frac{\Omega_0}{\beta} \right)^2 - \mu^2 \right]^{1/2} \right\} \\ & \times \operatorname{sech} \left[\pi \left(\frac{\Delta}{2\beta} + \frac{\mu}{2} \right) \right] \operatorname{sech} \left[\pi \left(\frac{\Delta}{2\beta} - \frac{\mu}{2} \right) \right] \end{aligned} \quad (4.40)$$

where w is the inversion component of the Bloch sphere, Δ is the detuning (rad/s), β is related to the width of the pulse, μ is a real constant and Ω_0 is the maximum Rabi frequency. w_0 is the initial value of w before the pulse, in our case when the atom is in the ground state, $|g\rangle$, $w_0 = -1$. For this equation to be valid and for the complex hyperbolic secant pulse to work the following two requirements, as also noted by Silver et al. [38], must be fulfilled,

$$\mu > 2, \quad \nu_{\text{width}} \leq 2\Omega_0. \quad (4.41)$$

4.6.2 Pulses with broader inversion

The frequency inversion width for a complex hyperbolic secant pulse is limited according to $\nu_{\text{width}} \leq 2\Omega_0$, see Section 4.6.1. However, it is possible to create pulse shapes which are optimised for different purposes, for example for inverting broader regions [39–45].




Figure 4.17. The paths taken by a few different ions with different detuning from the ground state to the excited state driven by a complex hyperbolic secant pulse.

QUANTUM COMPUTING WITH RARE-EARTH-METAL-ION-DOPED CRYSTALS

A great deal there is a lot of experimental work is currently being carried out aimed at constructing a very basic quantum computer, (see, for example, references in [35, 46–49]. In this chapter I will discuss how a quantum gate can be constructed in rare-earth-metal-ion-doped crystals.

Much of this work has been carried out as part of the European research projects, REQC hardware, “Development of quantum computer hardware based on rare-earth-ion-doped inorganic crystals”, and the project following it, ESQUIRE, “Experimental realisation of quantum gates and development of scalable quantum computer schemes in rare-earth-ion-doped inorganic crystals”. The participants in these projects were: Lund, LTH [50], Aarhus, QUANTOP [51], Orsay, LAC [52], Paris, LCAES [53] and Caen, CIRIL [54]. Many impressive results in this particular field have been achieved in Matthew Sellar’s group in Canberra, Australia [21, 22, 25, 55, 56], and work is also being initiated by Mingzhen Tian’s group in Bozeman, USA [57].

A general introduction to the field can be found in the widely used text book by Nielsen and Chuang [35], the textbook by Stolze and Suter [58] or online sources [47–49].

5.1 Introduction

The long-term goal of our research is, of course, to build a quantum computer with many bits and very high fidelity. A good first goal towards this is constructing a good fidelity controlled NOT quan-




Figure 5.1. Classical XOR gate.




Figure 5.2. A qubit is normally stored in two of the hyperfine levels of the ground state, and the optically excited state is an intermediate state that can be used during the operations which are driven with laser light.




Figure 5.3. A quantum system is initially in a superposition between the two ground states, as shown in the left side of the figure. When a π -pulse is applied on the $|1\rangle \rightarrow |e\rangle$ transition this superposition is transferred to the $|0\rangle$ and $|e\rangle$ states, as shown on the right-hand side.

tum logical gate, or CNOT gate, and much of our experimental work has been devoted to this.

Let us start by considering the classical case. The smallest unit of information in a traditional computer is represented by either the high voltage or a low voltage inside the processor. To make it less complicated to write, the low voltage is simply called 0, and high voltage 1. Figure 5.1 shows the symbol and truth table for the classical equivalent of a CNOT gate which is called XOR. The function can be described as follows, where we designate the first bit C, for control, and the second bit, T, for target. If the control bit is 0, nothing should happen to the target bit, and if the control bit is 1 the value of the target bit should be changed. We will now see step by step what is needed to construct a CNOT gate, and how this is implemented in the rare-earth-metal-ion scheme.

5.2 What is needed to construct a quantum computer?

The very famous Turing machine, which can simulate any computation that can be made on a computer, shows that at a fundamental level all computers may be regarded as equivalent [59, 60]. A similar concept that is very useful for quantum computers is the result that if we are restricted to only a two-qubit CNOT gate and a set of single qubit gates, it is actually still possible to carry out any quantum computation with these basic building blocks [61, 62].

Let us try to formulate some criteria for what is needed to actually construct a rare-earth-metal-ion-doped-crystal-based quantum computer, based on the well-known DiVincenzo criteria [63].

5.2.1 Qubits

A quantum mechanical bit, qubit, can just as its classical counterpart have the values 0 or 1, but the extra power of the quantum computer originates from the possibility for qubits to be in coherent mixtures, superpositions, between these two values. A good way to describe and visualise a qubit, is again to use the Bloch sphere, which was introduced in Section 4.2. When the Bloch sphere is used to represent a qubit, we can denote the south pole of the Bloch sphere $|1\rangle$ and the north pole $|0\rangle$. States that are fully known, so-called pure states, are represented by points on the surface of the sphere, and states which are not fully known, mixed states, are represented by points inside the sphere.

In Chapter 4 we used the Bloch sphere to describe a two-level system where the two base states were a ground state and an optically excited state. In the REQC scheme the qubit representation can be viewed as being moved between two different two-level systems, which enables the interaction between different qubits to be

turned on and off. This is explained in more detail in Section 5.2.5. The Bloch sphere only describes two-level systems, and to predict the behaviour of our full system, simulations solving the equations of motion in the space, spanned by all the states, have been carried out [64], but for the discussions here we will consider a simpler case, as shown in Figure 5.2.

The qubit is normally stored as a coherent superposition between only two of the hyperfine levels of the ground state, and this two-level system can then be described in the Bloch sphere. If a π -pulse is applied, for example between $|1\rangle$ and the $|e\rangle$ state, the part of the wave function that was in $|1\rangle$ will be transferred to $|e\rangle$, and after the π -pulse there is no population left in state $|1\rangle$ and the Bloch sphere now can be used to describe the two-level system consisting of $|0\rangle$ and $|e\rangle$, as indicated in Figure 5.3. The state of one such qubit, based on a single Pr^{3+} ion, would be very difficult, if not impossible, to detect directly. We have therefore, so far, been working experimentally with ensembles of such ions, but an approach for single-ion detection is described in Section 5.3. In Section 3.2 and Papers II and III it is described how peaks containing many ions can be created, manipulated and measured. Strictly speaking, each individual ion in such a peak is a qubit belonging to one instance of many quantum computers running in parallel, but the whole peak is usually referred to as a qubit. The very heart of the REQC scheme is the use of frequency addressing, which builds on the inhomogeneous broadening, which means that different ions have different resonance frequencies, and it is therefore possible, as we shall see, to make qubits that can be addressed separately by choosing an appropriate frequency of the laser light.

One might think that having many quantum computers running in parallel would cause problems, since, for example, if several intermediate results of Shor's algorithm are added bitwise, the information is lost. It has however, been shown that if only the full calculation is carried out in the quantum computer this is not a problem (See [65] or Section 7.7.3 in Nielsen and Chuang's book [35]).

5.2.2 Initialisation to a pure state

Before computation can begin, it must be possible to initialise the qubit to a known state. In the REQC scheme, this is equivalent to isolating a peak of ions in a single hyperfine state, which was discussed in Section 3.2. This can be done very effectively, as is shown in Paper II.

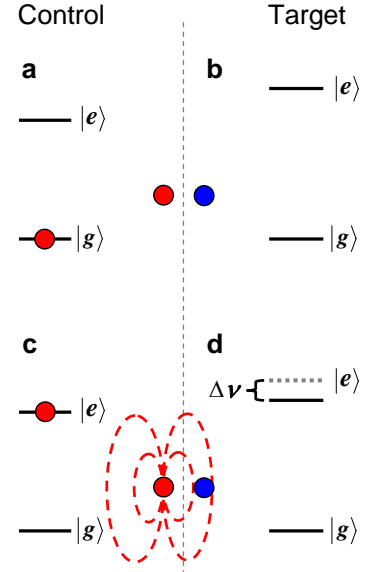


Figure 5.4. Two ions a few nanometres apart which have different transition frequency, and can therefore be manipulated independently. We call the first one the control, a, and the second one the target, b. The ions have permanent dipole moments which are different in their ground and optically excited states. If the control ion is transferred to its optically excited state, the electric field (E-field) around it will therefore change, c. When the E-field surrounding the target ion changes, its optical resonance frequency will be Stark shifted, d.




Figure 5.5. Two ions with different optical resonance frequencies, spatially located very close to each other, one denoted the control and the other the target, **a**. When the control ion is transferred to its optically excited state, which is indicated by its disappearance, the resonance frequency of the target ion is shifted, **b**.




Figure 5.6. Consider two peaks, each consisting of many ions, one containing control ions, and the other target ions, **a**. Since the ions are randomly located, some target ions will happen to be close to a control ion and some will not. When the ions in the control peak are transferred to their excited state, some target ions will experience a large shift in resonance frequency, while others will experience almost none, **b**.

5.2.3 One-bit gates

The next requirement, is the possibility to perform single qubit gates to arbitrary input states. In REQC this requires some special attention, since ions within a peak belonging to different instances of quantum computers, have slightly different resonance frequencies, because of the inhomogeneous broadening within the peak. One way to create single-qubit gates with inhomogeneously broadened peaks employs laser pulses with two frequencies simultaneously [64]. An experimental system to realise such pulses has been constructed, and these pulses will hopefully be experimentally tested very soon. One class of single-qubit pulses has, however, already been experimentally demonstrated (see Paper III). The pulses used in these experiments are called complex hyperbolic secant pulses and were described in Section 4.6. They have the ability to completely transfer the probability amplitude for all instances in the inhomogeneously broadened peak coherently from one state to another.

5.2.4 Two-qubit interactions

To enable two-qubit gates, there has to be some interaction between the two ions constituting the two qubits. Such an interaction exists, namely the dipole ion-ion interaction which was introduced in Section 2.6, and it can be used to perform two qubit gates between pairs of ions which lie close enough together. The ions have different static dipole moments in the ground state and in the excited state. This simply means that the E-field in the vicinity the ion changes when it is excited. In Figure 5.4a an ion is illustrated in its ground state, denoted the control. A few nanometres from the control ion, is another ion, the target, shown in panel b. The control and the target ions have different resonance frequencies and can therefore be manipulated independently of each other. When the control ion is transferred to its optically excited state its static dipole moment changes, which results in a change in the E-field surrounding it, as indicated in panel c. When the local E-field surrounding the target ion changes, its optical transition frequency will change, as depicted in d. This is illustrated slightly differently in Figure 5.5, where resonance frequencies of the control and the target bit are indicated in panel a. A π -pulse applied on the optical transition will transfer the control ion into its excited state, which is indicated in panel b by the disappearance of the peak, accompanied by the subsequent change in resonance frequency of the target ion. Let us consider a situation that better resembles the one in the crystals studied, where there are many ions with slightly different resonance frequencies which together make up a peak, as illustrated in Figure 5.6a. The spatial positions of ions are random; the ions in the target bit may or may not

be close to an ion from the control peak. This means that when the control peak is excited, as shown in the left side of **b**, some ions that happen to be close to a control ion will experience a large change in E-field and thus get a large frequency shift, while others further away from any control ion will not experience any change in E-field, and their resonance frequency will be unchanged. This situation is been depicted in **b**.

We now wish to remove the ions in the target bit that interact too weakly with the control bit, i.e. the ions that only experience a small frequency shift when the control ions are excited. Before doing this, let us first consider a few of the properties of one of the REQC materials, $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$. Some of these properties were revealed in the studies leading up to **II** and Papers **III**. The relevant properties are summarised in Figure 5.7, where suitable uses for the different hyperfine levels are indicated. The ground state levels $\pm 1/2$ and $\pm 3/2$ are denoted $|0\rangle$ $|1\rangle$, due to the high oscillator strength of the cross-over transitions via excited $\pm 1/2$ and $\pm 3/2$ states, which are here denoted $|e1\rangle$ and $|e2\rangle$. As mentioned above, we wish to remove the target ions with only small frequency shifts, and the state denoted $|\text{aux}\rangle$, meaning auxiliary, will be used as a reservoir for these ions. We have seen some evidence that the branching ratios of the excited states seem to follow the relative oscillator strengths. These observations were made in other experiments, but no quantitative investigations have been made. Of special interest is the high oscillator strength, and thus the high probability of decay via this route for the $\pm 5/2$ to $\pm 5/2$ transition. This means that if the "undesired" ions can be transferred to the excited state $\pm 5/2$, they will have a very high probability of relaxing to the $|\text{aux}\rangle$ state. This is why the excited $\pm 5/2$ state is denoted $|\text{aux funnel}\rangle$.

As discussed above, the REQC scheme uses frequency addressing, with pulses, such as the complex hyperbolic secant pulses, which are highly frequency selective, and will only affect ions within a very specific frequency interval. Consider Figure 5.6 again which shows the broadening of the target peak when the control ions were excited. If a frequency-selective pulse is applied to the target ion, which only excites the unshifted population to the $|\text{aux funnel}\rangle$ state, and then another pulse that returns the control ion to the ground state, the target ions in the $|\text{aux funnel}\rangle$ state will mainly decay to the $|\text{aux}\rangle$ state. In this way most of the unwanted ions will be removed. This procedure is schematically depicted in Figure 5.8, and was experimentally demonstrated in Paper **III**, although in this paper the unwanted ions were moved to the other computational base state. The procedure can be repeated until the desired distillation purification is reached.




Figure 5.7. Properties of $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$. The different hyperfine levels are denoted in accordance with the intended usage in the REQC scheme. The table gives the relative oscillator strengths of the transitions. The table has been slightly altered compared with that in Paper **II**, so that both the row and column sums are equal to one.




Figure 5.8. Schematic description of the distillation process which selects only the target ions that experience a large frequency shift when the control qubit is excited. The process starts with both control and target qubits in their $|0\rangle$ state, **a**. The control qubit is then transferred to the optically excited state, upon which the resonance frequencies are shifted more or less, depending on how close each ion is to a control ion, **b**. A frequency-selective pulse then transfers the ions with insufficient shift to the $|\text{aux funnel}\rangle$ state, **c**. After this, the control qubit is transferred back to the $|0\rangle$ state, upon which the ions that were earlier transferred to the $|\text{aux funnel}\rangle$ state decay to the $|\text{aux}\rangle$ state. The remaining ions in the target qubits are now all controlled by the control bit, **d**.

5.2.5 Two-qubit CNOT gates

Figure 5.9 shows the symbol for the CNOT gate, together with its matrix representation. In the computational basis $|C, T\rangle$ nothing happens to the target bit if the control bit is $|0\rangle$ and if the control bit is $|1\rangle$ the value of the target bit is swapped.

I will first describe how to make a NOT gate with REQC, and then extend it to a CNOT gate [28]. As the hyperfine transitions in these crystals can be driven with RF fields, one can ask why not simply perform a NOT gate between the computational $|0\rangle$ and $|1\rangle$ states with an RF π -pulse, as indicated in Figure 5.10a. The answer is that as all ions have nearly the same RF transition frequencies for the $|0\rangle$ to $|1\rangle$ transition, this would perform a NOT gate on all qubits, and even worse on all other ions in the crystal. However, since the transition frequencies to the optically excited state are different for different qubits, this can be used to perform a NOT gate to a selected qubit with three optical π -pulses, as shown in Figure 5.10b. If we start with a superposition, $\alpha|0\rangle + \beta|1\rangle$, the first pulse transfers the population of state $|0\rangle$ to state $|e\rangle$ and we will have $\alpha|e\rangle + \beta|1\rangle$. The second pulse will swap the populations

between states $|1\rangle$ and $|e\rangle$ leaving us with $\beta|e\rangle + \alpha|1\rangle$, and finally the third pulse will return the population from state $|e\rangle$ to state $|0\rangle$, and the final result will be $\beta|0\rangle + \alpha|1\rangle$, i.e. the populations of the two states have been swapped.

In Section 5.2.4 it was shown how only the target ions that experienced a large shift in frequency, when the control ions were excited, were selected. We can now construct a REQC CNOT gate by combining this with the NOT gate described above.

Let us start with the case when the control ion is in the state $|0\rangle$, see Figure 5.11, top row. Pulse 1 excites the control ion to state $|e\rangle$, which then changes the local E-field, **a**, which in turn shifts the resonance frequency of the target ions, **b**. When the NOT pulses, pulse 2 to 4, are applied to the target, they are no longer resonant, and will have no effect, **c**. A final pulse, not shown, on the $|0\rangle$ to $|e\rangle$ transition will return the ions from the excited state. In the lower row the same pulse sequence is applied but the control ion is initially in state $|1\rangle$. The first pulse on the $|0\rangle$ to $|e\rangle$ transition now has no effect, **a**, the resonant frequency of the target ions is not shifted, **b**, and pulses 2 to 4 perform the NOT gate on the target ions, **c**. The final pulse, not shown, on the $|0\rangle$ to $|e\rangle$ transition, again has no effect.

5.2.6 Readout

At the end of the computation the result must be read out. Readout in other bases than the computational one is achieved by a rotation into the computational basis before the readout. It is the average value given by all the parallel quantum computers that have been running that is read out. There are several different methods of doing this. In the first one, the laser frequency is simply scanned across the qubit peaks, while the absorption is measured. In the second method, the laser is again scanned across the features, but this time the fluorescence is measured. A third method, is to put the qubit into a superposition with a brief $\pi/2$ -pulse, change the laser frequency to a value where there are no absorbing ions, and monitor the beating between the free induction decay (FID) from the qubit ions and the laser.

5.3 Single-ion REQC

My belief is that it is possible to construct a two-bit REQC CNOT gate with the ensemble approach, using many ions in each qubit. But since we simply choose the positions of the frequency channels beforehand, and then remove the quantum computers that are not represented in all qubits, this scales poorly if the probability of the quantum computer being represented in a frequency interval corresponding to the qubit is low. If we instead have a quantum




Figure 5.9. Symbol and matrix representation of a quantum CNOT gate.




Figure 5.10. One might imagine that a NOT gate could be performed with a RF by π -pulse, **a**, but that would perform a NOT gate on all ions in the crystal. Since the ions in the different qubits have different resonance frequencies, a qubit specific NOT gate can be performed using the optical transition, **b**.

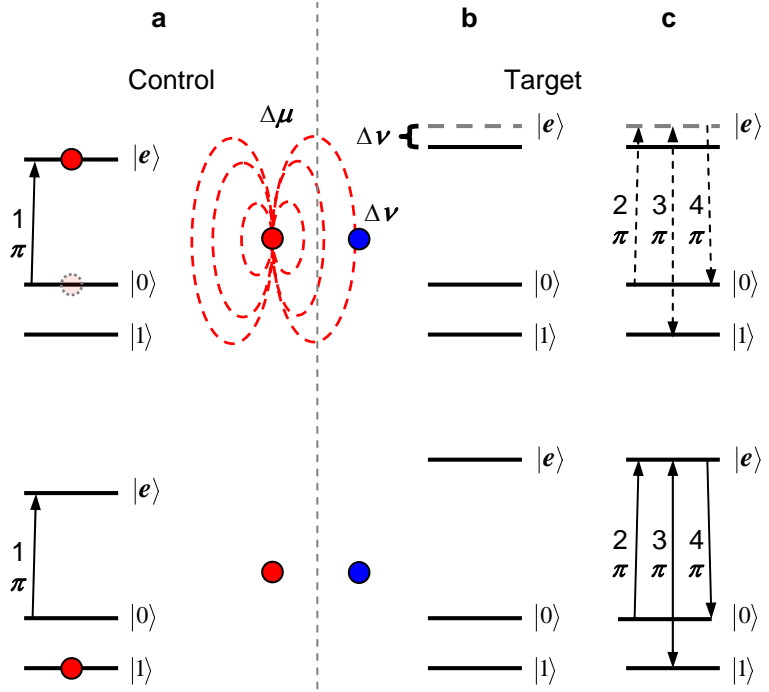


Figure 5.11. Principle of the REQC CNOT gate. The upper row shows how the local E-field changes as a result of the control ion being excited, **a**. The transition frequencies of the target ions are shifted, **b**, and the subsequent NOT sequence on the target has no effect. The last pulse, not shown, on the $|0\rangle$ to $|e\rangle$ transition of the control ion, returns it to state $|0\rangle$. In the lower row exactly the same pulse sequence is applied, but the control ion starts in the $|1\rangle$ state, and is subsequently not excited by the first π -pulse on the $|0\rangle$ to $|e\rangle$ transition, **a**. The resonance frequencies of the target ions are not changed, **b** and the three NOT pulses on the target ions swap the population between states $|0\rangle$ and $|1\rangle$, **c**.

computer based on a single ion for each qubit, we can use a search strategy to find the frequencies at which there are qubits. There are two problems with the single-instance approach. First, it is not possible to focus the laser down to a sufficiently small spot, and secondly I can not conceive of a way of reading out the state of a single Pr^{3+} ion. It might, however, be possible to solve both these problems if we implant an ion of a different species, that has a transition that can be cycled many times for the readout procedure.

The basic concept is illustrated in Figure 5.12. In reality, there is of course a continuous random distribution of ions, but let us here describe it as clusters of possible quantum computers, as the star symbols in panel **a** indicate. The large fuzzy circular area

shows the spot on which the laser is focused. Inside this spot there are many ions in any given qubit frequency channel. The trick is now to dope in a second species, which we will call a readout ion (indicated by the small dot in panel **a**), at such a low concentration that, on average, only one readout ion is detected within the focal spot. Two issues must be pointed out. First, the excitation wavelength for the readout species is completely different from the qubit ions, which, on the one hand means that two lasers have to be used but, on the other, the readout and the qubit ions can be manipulated complete separately. Second, the spectral profile of the readout ions is also inhomogeneously broadened, which means that there can actually be several readout ions within the focal spot, as long as there is only one readout ion per frequency channel¹. The readout ion should not participate in the actual computation, and the long lifetime required for the qubit is neither necessary nor desirable. The readout ion should preferably have three levels, as shown in Figure 5.13. The single ion can be detected by illuminating the $|g_r\rangle$ to $|e_r\rangle$ transition, while measuring the fluorescence of the $|e_r\rangle$ to $|i_r\rangle$ transition. To obtain a good fluorescence signal it should be possible to cycle the ion very rapidly, and the states $|e_r\rangle$ and $|i_r\rangle$ should therefore be short lived.

The readout procedure uses the same dipole blockade effect as the CNOT, but in a slightly different way. Figure 5.12b shows an enlargement of the surroundings of the readout ion. As can be seen there is a qubit ion, q1, very close to the readout ion. Figure 5.11 showed how the resonance frequency of the target ion was shifted when the control ion was originally in state $|0\rangle$, when the control ion was conditionally excited by pulse 1, and the subsequent NOT pulses were therefore only resonant if the control ion was originally in state $|1\rangle$. To read out the state of q1 out, a π -pulse is applied to its $|0\rangle$ to $|e\rangle$ transition, and the resonance frequency of the readout ion is shifted condition on q1 originally residing in state $|0\rangle$. In Figure 5.13 the conditional shift of the readout is indicated by $\Delta\nu$. If it is shifted, the laser at the original frequency of the $|g_r\rangle$ to $|e_r\rangle$ transition will no longer be resonant and the fluorescence signal will disappear. This means that by monitoring the fluorescence signal from the readout ion it can be decided whether q1 was in state $|0\rangle$ or $|1\rangle$.

To find q1 continuous monitoring of the readout ion is needed, while the qubit laser is scanned. When the fluorescence signal from the readout ion disappears, the qubit ion has been found. This can be repeated for several ions, which either directly quenches the light from the readout ion or, more probably, quenches the readout ion via q1. It is only necessary for one of the qubit ions




Figure 5.12. An ion from a different element is doped into the crystal at such a low concentration that with the laser tuned to an arbitrary frequency within the absorption line of this ion, only one ion will, on the average, be excited within the focal spot, **a**. An enlargement of the volume around the readout ion shows the qubit ions surrounding the read out ion, **b**.

¹It can be confirmed that there is only one ion per channel, by measuring the coincidence of arrival times in two detectors, one at each exit of a beam splitter.




Figure 5.13. A good readout ion should have three levels, where the laser excites the $|g_r\rangle$ to $|e_r\rangle$ transition, while the fluorescence is detected from the $|e_r\rangle$ to $|i_r\rangle$ transition. The levels $|e_r\rangle$ and $|i_r\rangle$ should be short lived so that many photons can be collected per unit time, to give a strong detection signal. If a qubit close to the readout ion is excited, the E-field surrounding it will change, since it has different static dipole moments in the ground state and excited state. This will frequency shift state $|e_r\rangle$ an amount $\Delta\nu$, and thereby quench the laser excitation and consequently the detected fluorescence.

to be close enough to the readout ion, since the state of the other qubit ions can be sequentially shifted to q1 before they are read out.

EQUIPMENT

Making nature do things it has never done before is not always easy. We want to make atoms "dance" in a quantum mechanically coherent way. In order to persuade them to do this we have had to combine different technological systems in a well engineered way, and some systems that could not be purchased we have had to build ourselves.

The equipment closely related to the stabilised laser is described in Chapter 8.

6.1 The dye laser

For many experiments, and for the laser stabilisation work described in Chapter 7, we use a modified Coherent 699-21 dye laser. The stabilised laser is pumped by a 6 W, diode-pumped, solid-state, DPSS, Nd : YVO₄ laser, (Coherent Verdi V6). The dye laser consists of a butterfly laser cavity, where the amplifying medium consists of dye molecules in a liquid solvent. The liquid dye solution is ejected through a nozzle, and lasing takes place in a free-flowing dye stream. Variations in the thickness of this dye stream are the cause of the comparably large, and rapid phase variations that this laser type exhibits. Tuning over a frequency range of 30 GHz is accomplished by turning a galvo controlled glass plate, mounted in the beam inside the cavity in Brewster angle. Frequency errors with a bandwidth of around 10 Hz to 1 kHz are corrected by one of the cavity mirrors mounted, which is mounted on a piezoelectric stack. One of the most significant benefits of a dye laser is that the laser frequency can be changed across a very large range. By using different dyes it is possible to tune the laser from 540 nm to 1040 nm with a pump source of 532 nm. Furthermore, we have considerable experience of this particular type of dye laser. One disadvantage is that they usually require

a great deal of attention to ensure that they function properly. The following changes were therefore made to the laser, and to the environment that surrounds it, to ensure the it worked reliably.

- The original dye jet nozzle was replaced by a Radiant Dyes RDSN 02 nozzle.
- The original nozzle holder and pump mirror adjustment unit were replaced by a Radiant Dyes RDU 10 unit.
- The laser table was moved into a class 100 clean room from Terra Universal. The clean room has six MAC 10 IQ fan filter units made by Envirco. These fans were chosen because they have a very low sound level. A sketch of the clean room can be seen in Figure 6.1. The clean room was chosen so large that it is possible to walk and work inside the room on three sides of the optical table. The walls are made out of 60 mil antistatic PVC vinyl stripes, which allows passage anywhere without tearing. This turned out to be very convenient.
- A solid state cooler with PID controller was installed, keeping the dye at 12.0°C. Cooler: Supercool LA-160-24-02 air-liquid modified to liquid-liquid with two RPS-73620 liquid heat sinks and four L-TURBO-8 turbolators. The PID controller is a supercool PR-59, with a temperature sensor fitted just where the dye leaves the cooler. The cooler is placed after the dye pump. All the tubes and parts was insulated with Armaflex insulation.
- The ion laser which was used as pump laser in the first experiments, was replaced by a Coherent Verdi V6 pump laser was used.

6.2 Cooled diode lasers

Before stabilising the dye laser, we investigated other possible laser sources, with smaller inherent line-widths. It is not possible to use a diode laser at room temperature, since there are no diode lasers with lasing wavelengths in the gap from around 450 nm to 635 nm. Matt Sellars suggested to us to try cooling down a 635 nm laser diode and in that way lowering its lasing wavelength to the desired 605 nm, which is the wavelength of the $^3H_4 \rightarrow ^1D_2$ transition in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$. Cooled laser diodes have been used by spectroscopists to reach wavelengths that are not reached at room temperature before [66, 67]. We tried simply dipping a 635 nm laser diode, into liquid nitrogen, but this only lowered the lasing wavelength to 612 nm, which is not sufficient. Therefore the idea to cool a 635 nm laser diode was abandoned.

The single instance computing efforts uses a cerium ion for readout. The required wavelength in this case is 370.6 nm. No laser diode currently on the market can reach this wavelength at room temperature. For this purpose, a wavelength-selective Nichia NDHU110APAE3 laser diode was purchased from Toptica Photonics. Our particular laser can be run at 370.4 nm at 13°C mounted in an external cavity, and still deliver more than 3 mW.

Acousto-optic modulators (AOM) can be used to amplitude and frequency modulate laser beams. They can also be used to scan the angle of a beam, a feature which I have not explicitly utilised. The modulator consists of a piece of crystal (or glass), with a piezoelectric transducer mounted on one side. The transducer is driven with an oscillating electric RF-signal, and ejects a sound-wave into the crystal. The acoustic wave can be pictured as planes with alternately high and low pressure, moving through the crystal. The pressure changes causes also the index of refraction to alternate in the same way. This results in a Bragg grating, which is travelling through the crystal, and it is possible to deflect light off this grating. Since the Bragg grating is moving, the light which is deflected will be Doppler-shifted. The frequency shift is equal




Figure 6.2. AOM based, double-pass, phase, frequency and amplitude modulation setup in a bow tie arrangement.

to $m \times F$ where m is the order of diffraction, and F is the frequency of the sound wave. By changing the sound wave frequency, it is possible to very rapidly change the frequency and phase of the laser light. However, since the alteration of the frequency is accompanied with a change of the wavelength of the sound wave, and thereby the grating spacing, the angle of the deflected beam will change simultaneously with its frequency. This can be compensated either by using two AOMs in series, or by passing the beam through the same AOM twice, as shown in Figure 6.2, referred to as a double-pass AOM. Modifying the amplitude of the electric RF-signal will of course alter the intensity of the deflected beam. I have extensively used AOMs to modulate the frequency, phase and amplitude of optical beams in the different experiments described in this work.

6.4 Cryostats

A cryostat is an apparatus for maintaining a very low temperature. In the work described in this thesis several different types of liquid helium cryostats have been used, all with optical access to the sample chamber where the rare-earth-metal-ion-doped crystal is mounted. The crystal is either contained in vacuum on a cooling finger, sitting in cold helium gas, or submerged in liquid helium. The experiments are most often performed at 4 K, the temperature of liquid helium at atmospheric pressure. In bath cryostats helium gas bubbles might be generated in the liquid helium, especially when high laser beam intensities are being used, which distort the beam. This can be circumvented by lowering the pressure in the sample chamber below the λ -point, where the liquid helium becomes superfluid and exhibits super-heat-conductivity, causing the bubbles to disappear.

Close-cycle cryostats can be operated without the need for a supply of liquid helium. These cryostats do however, suffer from vibrations, a problem which currently is being actively addressed by the manufacturers.

LASER STABILISATION

The first laser was demonstrated in 1960 by T.H. Maiman [68]. Since then lasers have become an extremely versatile tool used in many different areas such as welding and cutting, storing and reading information from CD discs, measuring the distance to the moon and precision spectroscopy just to mention a few examples. Some of these applications, like precision metrology, rely on lasers with an extremely stable output frequency. Here the stable laser can be compared to a stable RF oscillator and radio transmitter, with the main difference that the frequency is much higher for the laser, around 500 THz. Since the frequency is much higher and thereby the wavelength much smaller, beams can be much more spatially confined than with radio waves.

If a laser is constructed to be single-mode and mechanically stable, like for example the non-planar ring oscillators [69], the line-width can be as low as 3 kHz over 100 ms without any active stabilisation [70]. These lasers can however only be made to work in certain, often fairly narrow, wavelength regions. For a quantum computer based on rare-earth-metal-ion-doped crystals the laser frequency is dictated by the transition frequencies for the material in question. For the transition $^3\text{H}_4 \rightarrow ^1\text{D}_2$ in Pr^{3+} at 605.977 nm = 494.726 THz, which was used in Papers II and III, we are not aware of any solid-state alternatives available at present, but work is in progress [71]. One possibility would be to adapt one of the laser types which are developed for the sodium laser guide star projects at 589.2 nm [72–74]. Another one, which turned out not to work for our purposes, is to frequency tune a 640 nm laser diode by cooling it, see Section 6.2 for more information.

We decided to use a modified Coherent 699-21 dye laser. For more details description of the modifications and the pros and cons with this laser type, see Section 6.1. The most important disadvantage dye laser as the 699, is the rather large and rapid phase fluctuations created by free-flowing dye jet, resulting in a

line-width of the order of 1 MHz. This means that the feedback loop stabilising the laser has to have a low response time, which is further discussed in Section 7.3 [75, 76].

There are many different methods of stabilising a laser. They are all based on measuring the phase or frequency, against some stable refer an element which can change the phase or frequency of the laser. It is almost always done by changing the optical length of the laser cavity in one way or an other. In 1983 the so-called Pound-Drever-Hall locking scheme was first demonstrated [75]. Since then, this method of locking to cavities, is the most widely used [76–84]. Before this the most frequent method was locking to the side of the transmission fringe of a cavity [85], which is also the method used as standard on the Coherent 699-21. Lately some interesting methods of using the phase shift between two different spatial modes of the resonant cavity have been suggested and used to stabilise lasers [86, 87]. Within the community working with coherent transients and spectral hole burning in rare-earth-metal-ion-doped crystals the Pound-Drever-Hall method had been used to lock lasers to a spectral hole burnt in these crystals, which is convenient since these crystals also are used for the experiments. Since it was first demonstrated in 1999 [88], it has become very popular both for locking to regenerative [89–92] and persistent [88, 93] holes. We have adopted this technique. There are however two major differences from previous work. In the Pr-doped material there is a semi permanent storage mechanism originating from multiple ground state hyperfine levels. For permanent holes it is difficult to get a good short time scale-phase stability, which is a necessary requirement for our quantum computing project. Therefore erasing radio fields, resonant with the hyperfine transitions, were added, which made the holes burnt into the material quasi-transient. The second difference is that instead of stabilising comparatively well behaved diode lasers we, for the first time, stabilised a noisy dye laser to a hole.

In this chapter we will mainly concentrate on describing the physics behind the frequency locking scheme. The technical details of how the locking system is constructed can be found in Chapter 8.

7.1 Locking to cavities

7.1.1 Cavity field build up

Let's start by looking at what happens when a an abruptly turned on light beam impinges on a high finesse cavity, schematically indicated by the line coming in from the left in Figure 7.1. The green dotted line is the spacer on which the two mirrors are mounted¹. To

¹The spacers holding the mirrors are often made of ultra-low thermal expansion glass, ULE glass, to minimise the change of resonance frequencies for

make the discussion easier to follow, let's give numbers to some of the cavity parameters: cavity length $L = 10$ cm, cavity mirror reflection $R = 99.99\%^2$, cavity free spectral range $\Delta\nu = 1.5$ GHz, cavity finesse $F = 30\,000$ and full width half maximum of a transmission fringe $\delta\nu = 50$ kHz. Some useful identities are:

$$\Delta\nu = \frac{c}{2L} \quad \delta\nu = \frac{\Delta\nu}{F} \quad F \approx \frac{2\pi}{1-\rho}, \quad (7.1)$$

where c is the speed of light and ρ is the fraction of the field left after one round-trip, in this discussion $1 - \rho \approx 2(1 - R)$. When the light is turned on, and hits the first cavity mirror, 99.99% of the light will to begin with be reflected back, as indicated by the dashed arrow in Figure 7.1.

Phases for reflection of dielectric mirror

Here we will make a small detour and consider the reflection from an imagined dielectric mirror consisting of only one layer with a thickness of $\lambda/4$ where λ is the wavelength of light inside the material of consideration. In Figure 7.2a we see the wave impinging from the right, 1. The transmitted and reflected E-field are given by reflection, Γ , and the transmission, τ , coefficients³

$$\Gamma = \frac{n_1 - n_2}{n_1 + n_2}, \quad \tau = \frac{2n_1}{n_1 + n_2}, \quad (7.2)$$

where n_1 and n_2 is the index of refraction of the material where the wave comes from, and goes into, respectively. Note, the sign of the transmission coefficient is always positive, while the reflection coefficient is negative when reflecting off the surface with higher index of refraction, and positive when reflecting off the surface towards lower index of refraction. In Figure 7.2b the light has entered the material, 2, and also been reflected on the first surface, 3. The E-field in the reflected wave has opposite sign, since it is reflected against a higher index of refraction. In panel c the wave has been reflected, 4, and transmitted, 5, against the second interface. In this reflection there is no change of sign since it occurs against the medium with lower index of refraction. In panel d finally, we noticed two things. Firstly, the wave which has been once through the slab, 6, and the one which were originally reflected from the first surface, 3, will interfere constructively. Secondly, the beam which was originally transmitted through the slab, 5, and the beam which was twice reflected, 7, interference destructively. Both these things means the same thing, a slab like

the cavity as a result of temperature changes.

²High reflection, low loss, mirrors can for example be purchased from Newport or Research electro-optics, REO

³See for example [94], Equations (8-140), (8-141) and (8-185).




Figure 7.1. Cavity used in laser locking. This geometry would typically be used when locking on the side of a transmission fringe. The green dotted line shows the spacer on which the mirrors are mounted.




Figure 7.2. A single $\lambda/4$ dielectric layer is used as an simple example of the dielectric mirror, to work out the relationships between the transmitted and the reflected waves. For details see the text.




Figure 7.3. Conceptual picture of a cavity on resonance. For details see the text.

7.1.2 Two mirrors, a cavity

Let's return to our multilayer dielectric mirror with 99.99 % reflectivity. If we only have one mirror 99.99 % will be reflected and that is it. As we shall see, adding another mirror with the same reflectivity at just the right distance, compared to the wavelength, can make a dramatic difference. Of the incoming beam, 1, in Figure 7.3, 99.99 % is reflected, 2, and only 0.01 % transmitted, 3, greatly enhanced for visibility. When another mirror is added the transmission of the first mirror is of course still only 0.01 %, and when this small fraction hits the second high reflectivity mirror it is almost completely reflected, 4. Let's now choose the distance between the mirrors as a multiple of $\lambda/2$, as in Figure 7.3, where the distance equals $2 \times \lambda/2$. When beam 4 hits mirror **a** two things should be noticed. Firstly, the beam which is transmitted, 5 has opposite sign a suppose to the originally reflected beam, 2, which means that these two beams will interfere destructively. Secondly, the part which is reflected, 6, is in phase with the beam already bouncing forward and back between the mirrors, 3.

Since the reflectivity is so high, one can say that the light bouncing between the mirrors is trapped, and will only slowly leak out. Further more, the part of 1 leaking in will continuously build up the circulating field. Assuming there is no losses, this build-up will continue until the field 2 and 4 completely cancels each other. At this stage the field inside the cavity will be 1×10^4 times higher than the field 1. The light leaking out through mirror **b**, field 7, will now be exactly as large as 1, just as it was completely transmitted. We have just described the transmission through a cavity on resonance.

Cavity light decay time

The time it takes for the intensity to build up inside the cavity, τ , or to decay if the light is suddenly turned off, is given by [95]

$$\tau = \frac{1}{2\pi\delta\nu}. \quad (7.3)$$

For a cavity with a line-width of $\delta\nu = 50$ kHz, the build-up, or ring down, time is $\tau \approx 3$ μ s.

If we again Figure 7.1 we can now say that it takes in order of 3 μ s for the intensity to build up inside the cavity and reach the detector. If the technique of locking on the side of the transmission fringe of the cavity is used [85], it takes quite a long time for the frequency error to change the intensity inside the cavity and before this happens, it cannot be detected by the detector after a cavity. It is of course possible to lower the finesse of a cavity, and thereby decreasing this time. But if this is done the line-width of a cavity will increase which means that the slope which you lock to won't be a sharp, leading to a smaller error signal for a given frequency error, which in turn means that the noise in the system will have a larger effect and will decrease the quality of the locking.

7.1.3 Pound-Drever-Hall locking

The problem with locking to the side of the transmission fringe is the delay before the error signal arrives when cavity with narrow line widths are used. One solution to this problem is the so-called Pound-Drever-Hall locking scheme, where instead the reflected light is used. For this to work, it has to be possible to separate the light going towards and returning from the cavity. This can be accomplished by inserting a polarising beam splitter cube and a quarter wave plate, as shown in Figure 7.4.

The light first passes through the beam splitter cube, up on which the light is turned circularly polarised by the $\lambda/4$ -plate before it enters the cavity. The light returning in the other direction is then turned back into linearly polarised light, on its second passage through the $\lambda/4$ -plate but now 90° turned in respect to the incoming light. This means that the light is now deflected towards the detector by the polarising beam splitter cube. This means we are actually dealing with circularly polarised light, but as circularly polarised light can be described as a superposition of two linearly polarised fields, we can conduct the discussion as if it was linear. In Figure 7.5 we have again drawn two of the fields discussed in Section 7.1.2. Field 1, is the direct reflection of the incoming beam on the first cavity mirror, **a**, field 2 is the leakage field from a huge field inside the cavity. These two fields are 180° out of phase, and will completely cancel each other, with the result being field 3.

Phase jump

What happens if we suddenly get a phase change on the incoming field? As we have said previously, the leakage field, 2, will not change, since it takes a long time to build up the new field inside the cavity. The reflection on mirror **a** will however change




Figure 7.4. The light coming out from the left side of the cavity is separated from incoming light with the polarising beam splitter cube shown to the left. The light first passes through the beam splitter cube, upon which its polarisation is turned 45° by the $\lambda/4$ -plate before it enters the cavity. The light returning in the other direction is then turned another 45° on its second passage through the $\lambda/4$ -plate, which means a total polarisation rotation of 90° . This means that the light is now deflected towards the detector by the polarising beam splitter cube.




Figure 7.5. The incoming beam reflection of the first mirror, 1, together with the leakage field from the cavity, 2. In the limit of no losses, these two fields will completely cancel each other resulting in field 3.




Figure 7.6. When a phase jump occurs on the incoming beam, the leakage field from the fields stored in the cavity will originally be unaffected by it, 2. The reflected field will however change instantly, 1, which means that these two fields do not any longer cancel each other and there will be a resulting field, 3. The two panels show positive and negative phase jumps.




Figure 7.7. Light slightly off resonance with the cavity. For details see text.

immediately, as can be seen in Figure 7.6, field 1. These two fields will no longer cancel and there will be a resulting field, 3. This can be detected as an increase in intensity which is proportional to the *phase error*. The two panels show jumps with positive and negative phase. As can be seen the phase of the resulting field, 3, is different in the two cases, but since the detector only measures intensity, this change cannot be detected. In both cases there will just be an increase in the intensity. We will get back to how this can be resolved. So with this method we can detect a phase error, but this error signal doesn't last for ever. If the error prevails, as the new phase leaks in to the cavity the new phase will be stored, and the error signal will disappear. This happens on the time scale of the build up time of cavity.

Constant frequency error

Next we will consider what happens if the frequency of the incoming beam is slightly off resonance with a cavity. Imagine that the frequency of the light is slightly higher than the resonance frequency for a cavity, and thereby the wavelength slightly shorter. This leads to that the wave is pulled a little bit for each round trip, as depicted in Figure 7.7a, where the red dotted line is the incoming light for reference, and the green lines are the returning field after one to five bounces on the other cavity mirror. For clarity it is only the waves travelling from right to left which are plotted. The waves become weaker for each bounce since a fraction of the field leaks out. If all these contributions are added we get the resultant field, which we can see as the blue dashed line in panel b. It is clear that the resulting field is phase shifted if compared to a field resonant with the cavity. A small part of this field will leak out from the end mirror, and this will of course also be phase shifted. In Figure 7.8a the on resonance case has been depicted, with the field inside and outside the cavity out of scale for clarity. The red dotted line is the incoming field, the blue dashed line is the intracavity field which is near in phase with the incoming field, and can be seen leaking out of a cavity 180° out of phase with the directly reflected in the coming light seen in dash-dotted black. This results in no field leaving the cavity to the left in the limit of no losses, solid green. Let's return to the frequency offset, and thereby phase shifted case just discussed. The cavity leakage is now also phase shifted, as shown in panel b. Adding the reflected field and then leakage field as in previous case, now gives the solid green line as a result, which is also phase shifted. This means that in this case we have the error signal which depends on the *frequency error*. Just as in the case of the phase jump in previous section, the phase will be different depending if the frequency is too high or too low, and this will not be resolved until next section. In the end of previous, it was concluded that the error signal

depending on the *phase error* died of as the field with a new phase was built at inside the cavity. And now we can conclude that the error signal which is dependent on the *frequency error* builds up on the same time scale.

Sign sensitive error detection

We now have an error signal which can measure both fast phase errors and detect if the laser is out of resonance with the cavity. Unfortunately the direction of the error is not known, which makes it impossible to know which way to change the phase or frequency to correct the error. The Pound-Drever-Hall method solves this by incorporating a phase modulator in the laser beam. This modulator dither the phase of the light forward and back, as shown in Figure 7.9, by applying a voltage over a material where the index of refraction depends on the field inside the material. The field after the modulator is

$$E = E_0 e^{i[\omega_0 t + m \sin(\omega_m t)]} \quad (7.4)$$

where ω_0 is the frequency of the light, ω_m is the modulation frequency, E_0 is the amplitude of the incoming E-field and m is the so-called modulation index, which is linked to the amplitude of the modulating field. It can be shown that if the modulation index is small this can be rewritten as

$$\begin{aligned} E &= E_0 e^{i[\omega_0 t + m \sin(\omega_m t)]} \\ &\approx E_0 \left[J_0(m) e^{i\omega_0 t} + J_1(m) e^{i(\omega_0 + \omega_m)t} - J_1(m) e^{i(\omega_0 - \omega_m)t} \right]. \end{aligned} \quad (7.5)$$

This corresponds to three fields at different frequencies. The first term at the original carrier frequency, and the next two are sidebands separated from the carrier with plus and minus the modulation frequency, as shown in Figure 7.10. $J_0(m)$ and $J_1(m)$ are Bessel functions of the first kind which determine the E-field amplitudes of the carrier and the sideband, respectively. Let's just consider the case when the modulation frequency is chosen so that the sidebands are located some distance from the transmission fringe of a cavity, or in the line of the description in first paragraph in Section 7.1.3 these fields will come out a of phase with themselves inside the cavity fast and interfere destructively and stop a high field from being built up inside the cavity. The cavity will due to this have no effect on the sidebands, which are effectively reflected from the first mirror surface of the cavity. We have already described how the light in the carrier that interacts with the cavity has a phase shift compared to the incoming light when there is a phase or frequency error. What we want to achieve now is to decide which way the phase is shifted. Let us therefore




Figure 7.8. The green solid line shows the total field which is returning to the left. In **a** the light frequency is precisely resonant with the cavity, and this field is zero. In **b** the light frequency is shifted slightly from the cavity resonance, and the cavity leakage field no longer cancels the cavity mirror reflection, resulting in a net returning field.




Figure 7.9. A sinusoidal electric field applied to an EOM inserted in a light beam will modulate the light.




Figure 7.10. Modulation sidebands




Figure 7.11. Panel a shows the central carrier, green curve with constant envelope, and the two lowest order sidebands added together, red curve with amplitude modulated envelope, and panel b shows the modulation signal.




Figure 7.12. Interference between carrier and sidebands. For details see text.

carry out some algebra on the two sidebands, E_s ,

$$\begin{aligned}
 E_s &= E_0 \left[J_1(m) e^{i(\omega_0 + \omega_m)t} - J_1(m) e^{i(\omega_0 - \omega_m)t} \right] \\
 &= E_0 J_1(m) e^{i\omega_0 t} 2i \frac{e^{i\omega_m t} - e^{-i\omega_m t}}{2i} \\
 &= E_0 J_1(m) e^{i\omega_0 t} 2i \sin(\omega_m t) \\
 &= 2E_0 J_1(m) e^{i(\omega_0 t + \frac{\pi}{2})} \times \sin(\omega_m t).
 \end{aligned} \tag{7.6}$$

Written like it is clear that these two fields can be described as a field with the same frequency as the carrier frequency, ω_0 , which is phase shifted 90° ($\pi/2$). On top of that the field is amplitude modulated by the term $\sin(\omega_m t)$, causing the amplitude to vary with twice the modulation frequency, ω_m . The sine term is, of course, positive every second half period and negative every other half period, which results in an alternating phase shift of 180° between each consecutive half cycle, on top of the 90° phase shift. The sidebands are shown in Figure 7.11a, red curve with amplitude modulated envelope, together with the central carrier, green line with constant envelope, where you can see that the green line changes between being 90° before and after between each half period. Panel b shows the phase of the driving voltage to the modulator for reference. A wave like this which switches between 90° before and 90° after every half cycle is actually the perfect tool to measure the phase of another wave, let's just call this the sideband field. Let's just see what happens if we phase shift the sideband field and interferes it with an unshifted wave. In Figure 7.12a, the red line with amplitude modulated envelope is the phase shifted sideband field, the green line with constant envelope is the unshifted wave, and it can clearly be seen how these two waves are in and out of phase between each period of the sideband field. This is even clearer in panel b where this downscaled sum has been drawn. In panel c the sideband field have been phase shifted in the other direction, and we will again get constructive and destructive interference successively, as is also seen when the two fields are added, panel d. But, if the phase of the interferences are compared to the driving voltage for the modulator, panel e, we see that the envelop, which is what the light detectors measure, is either in phase, or out of phase with the modulator driving voltage. And this phase we *can measure*, at a conveniently low frequency, as we choose the driving frequency for the modulator ourselves, in our case 50 MHz. The phase detection is practically done with a mixer and low pass filter.

So, when there is a phase jump or a frequency change in the light coming from the laser, this will of course be present after the modulator as well, both in the central carrier and in the side bands. The side bands are simply reflected from the first mirror surface of the cavity, which means that the phase shift is the same




Figure 7.13. The error signal, which is derived by mixing the beating signal from the detector with the modulation signal going to the modulator in the mixer, followed by low pass filtering in the low pass filter, LP. The phase shifter, PS, compensates for phase shifts in the setup, for example caused by time delays. The error signal goes into the regulation system, RS, which then feeds a control signal to the intracavity modulator, IM.

for the side bands after reflection. Due to the interaction with the cavity this is not the case for the carrier which was described previous, both for a frequency error and a phase jump, where the conclusion was that it was necessary in some way to measure the phase of the light, which we now have accomplished.

Conceptual cavity locking setup

We have now described how to get an error signal, but we still have to correct the frequency of the laser. If a modulator of a similar kind as the one which was used to create the phase modulation is inserted into the laser cavity, it is possible to change the optical length of a cavity by applying a voltage, and thereby change the resonance frequency of the laser. The whole setup can be seen in Figure 7.13.

7.2 Locking to holes

We will now repeat the same reasoning, as we did for cavities above, but this time for spectral holes. This time we would start with a constant frequency error.




Figure 7.14. A laser with a frequency in-between the frequency of two absorbing peaks of atoms.

7.2.1 Constant frequency error

We will start by considering what happens when a single frequency field impinges on a spectral hole in inhomogeneously broadened absorption profile like the ones described in Chapter 3. Let us recall the conclusions from Section 4.4.1. where we discussed what happens when a field which was either below or above in frequency, interacted off-resonantly with a transition, or peak of atoms. The conclusion was that when the field had a lower frequency than the resonance frequency of the atoms the fields phase would lag, and vice versa, if the field frequency was above the resonance frequency of the atom would cause a lag of the phase of the field. Let us now imagine that the frequency of the field is centered in between two equal peaks of absorbing atoms as shown in Figure 7.14. One of them alone will cause the phase of the laser to lead, and the other one by it self will cause the phase of the laser to lag, but when they are both present the phase shift caused by one will be cancelled by the phase shift caused by the other and the net effect would be no phase shift. What this little example serves to show is that, in steady state, if the absorption profile next to the laser frequency is symmetric there will be in no net phase shift. This means, for a laser which sits in the centre of a symmetric spectral hole, there will be no phase shift of the light. This is shown in Figure 7.15a, where the small insert shows the location of the laser frequency compared to the spectral hole. Here Ω_{re} is the incoming laser field, Ω_{atoms} is the field which is generated by the atoms and $\Omega_{\text{net}} = \Omega_{\text{re}} + \Omega_{\text{atoms}}$ is the field after the light has passed atoms. If the laser frequency is slightly higher then the centre of the hole, as in Figure 7.15b, there will be a excess of atoms at frequencies above the laser frequency compared to below the laser frequency, and the phase shift by the atoms above and below the laser frequency will no longer be cancelled, and the light which is sent out by the atoms is phase shifted compared to the laser light, which also causes a phase change in the transmitted field, as indicated in the figure. In Figure 7.15c there laser frequency has moved to lower frequencies, compared to the spectral hole, which causes the opposite phase shift. That means that we now have a phase shift which is depending on which side of the whole the laser is, in the same way as we had for the frequency error in the cavity case in Section 7.1.3.

The timescale on which the frequency error signal appears is inversely proportional to the hole width, which can be understood as follows. Ions with different detuning have different resonance frequencies, but if the field is suddenly turned on it will take some time before they come out of phase with a driving field. This means, that right after a field is turned on, the absorption comes from many atoms and the ions in the frequency range of a narrow is only responsible for a small part of the absorption. But as time




Figure 7.15. If the field with Rabi frequency Ω_{re} impinges on the hole burning sample, the ions in the sample will send out the field with a Rabi frequency of Ω_{atoms} . The total field, after the light has passed a sample, is given by $\Omega_{\text{net}} = \Omega_{\text{re}} + \Omega_{\text{atoms}}$. If the laser is centered on the spectral hole there will be no phase shift, **a**, if it is situated at a frequency above the centre of a whole there will be a phase shift as shown in **b** and if it is at the lower frequency the effect will be as shown in **c**

goes on the region which is still in phase with a driving field narrows and eventually it becomes as narrow as the hole and then the asymmetry of not being placed in the centre of a whole becomes large. Another way of saying this is that the frequency spectrum of a brief pulse is broad. This means that the time it takes for the frequency error signal from the hole to grow is inversely proportional to the hole width.

That just as in the cavity case a reference is needed to decide the sign of this phase shift. This is again accomplished with their Pound-Drever-Hall method, which we will get back to in Section 7.2.3.

7.2.2 Phase jump

In order to describe what happens when there is a rapid change of phase in the light, let us look back at Section 4.5, where we described free induction decay, where we concluded that if the light was turned off the ions continued to emit light with the same phase as the driving field had before, and that this field was emitted for a time inversely proportional to the width of the hole. If there is a sudden phase jump, the light emitted by the ions is much like the free induction decay keeping the old phase, and this is mixed with the phase of the light after the jump. The net effect of this is a phase in-between the old and the just as in the as we so in the cavity case. One can say that the ions remember the old phase, just as the cavity did. It is necessary that both the light emitted by the ions and the directly from the laser hits the detector, therefore should the absorption not be 100%. We will now analyse to the

problem of deciding which way the phase changed.

7.2.3 Sign sensitive error detection

In Sections 7.2.1 and 7.2.2 it was concluded that phase and frequency changes in light, passing through a hole burning material, could create the same phase shift response on the light as a cavity. So one would think that it is just straightforward to add the modulation, get some sidebands and from them derive an error signal, just as in the cavity case described in Section 7.1.3. This is almost true, but not quite. The difference is the following: in the cavity case the sidebands are not resonant with the cavity, and the sideband fields are therefore not stored in the cavity, but in the case of locking to inhomogeneously broadened transitions this is not the case, since the transitions are typically several gigahertz wide, which means that the sidebands will also be resonant with the ions, creating their own holes, and causing sideband fields to be stored as well.⁴

The situation is in some respect equivalent to choosing the free spectral range equal to the modulation frequency in the cavity case. This would mean that both the carrier and the sidebands would be resonant with the cavity, and they would therefore both be stored. Since the field storage in the cavity case is linear as a function of the incoming field, there would be no error signal in this case. The reason is that before, the phase shifts of the sidebands and the carrier were different, as the carrier was mixed with the stored field in the cavity.

In Section 4.5 I concluded that it was not until the system became non-linear that the free induction decay started to grow. It was also stated that the rate at which holes were created was proportional to the inversion and not to the polarisation, which causes the hole depth also to be non-linear in the field creating the hole. This means that the hole depth increases non-linearly with the field strength which means that also the free induction decay will increase non-linearly with the field strength. As long as the central hole grows faster than side holes, the free induction decay will be comparably larger from the centre hole, and it will be possible to maintain a locking signal. To get this nonlinear difference the centre field must be much larger than the side hole fields.

In Paper V we thoroughly discuss the effects of the necessity to get into the non-linear regime and conclude that the optimal amount of power in the side bands should be lower than in the

⁴There are two reflections to be made here: firstly, gigahertz modulation and detection are by no means impossible, but it requires some special attention, secondly there are hole burning materials with an inhomogeneous broadening less than 200 MHz [96], which would of course lower the needed modulation frequency.

cavity case. In the cavity case the most effective modulation index is $m = 1.08$ [97], but if just taking the optimal signal strength of the locking signal into consideration, the optimal modulation index in the spectral hole locking case is $m = 0.56$.

7.2.4 Locking to holes with different storage mechanisms

The lowest Root Allen variance locking to spectral holes published so far, to the authors knowledge, was achieved in a transient system with a metastable state acting as reservoir [91]. In this material the holes created were transient and were refilled through population decay from a metastable state with a timescale of 10 ms. Hole locking to two level systems have also been demonstrated where the storage mechanism corresponded to transferring some population to the excited state [92]. Spectral hole locking in materials featuring permeant hole burning have also been shown [88, 93] but there are problems some problems attached to this approach. Since the ions have a limited lifetime in the upper state they will have Lorentzian line shapes, which give very broad wings. This means that if a laser is kept on a single frequency it will create a hole which becomes broader and broader, so called power broadening. This is not desirable since if you have a broad hole the phase memory becomes shorter, and the frequency error slope becomes flatter, which, due to signal to noise considerations, leads to that the locking becomes less accurate. To avoid this the power can be decreased, but this will also lead to a decreased signal-to-noise ratio. Therefore we prefer working with transient system where the ions by one mechanism or the other returns from the reservoir state filling up the hole after some time. Another issue which can occur is linear drift as we discuss in Section 7.2.5. The current quantum computing efforts in Lund, are based on the $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ and this crystal has storage based on multiple hyperfine levels, with a lifetime of approximately 2 min, a time which in this context can be considered to be permanent. To avoid the problems attached to locking to permeant materials the material was made, what we will from now on refer to as “quasi transient”, by applying radio fields resonant with the ground state hyperfine transitions on 10.2 MHz and 17.3 MHz. These radio fields simply drives the radio frequency transitions, and thereby effectively decreases there lifetime by shuffling the population around. This is described in more details in Paper V.

7.2.5 Hole drift

Spectral holes in a crystal with permanent storage mechanism can easily become so-called power broadened. Apart from the less accurate locking, already discussed in Section 7.2.4, another problem

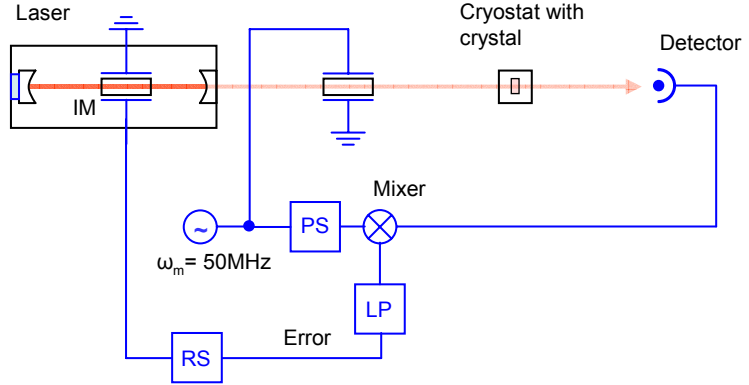


Figure 7.16. Conceptual setup for laser stabilisation to a hole. The designations used in the figure are: phase shifter, PS, low pass filter, PS, regulation system, RS and intracavity modulator, IM. For details see text.

arises. After a while the centre hole be burnt to the bottom and then it will start to become broader. As long as the holes gets deeper the phase shift in the transmitted light , as a result of a frequency error, will increase, but then it saturates. But the phase shift at the sideband frequencies keeps on growing. This can cause a sign flip in the locking signal which can cause a drift. This mechanism is thoroughly investigated in Paper [V](#).

7.2.6 Advantages with spectral hole locking

Both cavity and spectral hole locking have proven to be very useful methods, although cavity locking by far is the most widely used approach. There are however some distinct advantages with spectral hole locking. As the light bounces forward and back many thousands of times inside their high-finesse locking cavity, any movement of the mirrors will be enormously enhanced. Therefore the cavity has to be temperature controlled to mK accuracy and extremely thoroughly shielded from vibrations, see for example [\[84\]](#). This is not necessary in the spectral hole locking case, since it is the largely the coherence time of the ions which decides how stable the reference is . Another benefit is that no mode matching is needed. Since the laser mode has to be matched to the locking cavity mode there will unavoidably be back reflections towards the laser when a cavity is used. Therefore, it is customary to use an optical diode to protect the laser from back reflections. This is not necessary when locking to holes.

A disadvantage with spectral hole locking on the other hand is of course that the crystal has to be kept at liquid helium temper-

atures, which requires a cryostat.

7.3 System design

This section describes some general design considerations regarding laser stabilisation to holes. The implementation of the system is described in Chapter 8, and the detailed theory for spectral hole stabilisation and a characterisation of the system is described in Paper V. Figure 7.16 shows a spectral-hole-stabilisation setup conceptually. After the beam exits the laser, it passes through an EOM, driven by a 50 MHz oscillator, which creates the sidebands. Then it goes through the rare-earth-metal-ion-doped crystal and the spectral hole. If there is a frequency or phase error, this will cause a 50 MHz amplitude modulation in the optical beam, which is registered by the detector. The signal from the oscillator goes through a phase shifter, PS, which can compensate for the delays in the system, and is then mixed with the signal from the detector. Low-pass filtering gives the error signal. The error signal proceeds to the regulation system, RS, which controls the intracavity modulator, IM, which can alter the optical cavity length of the laser, and thereby change the frequency of the laser. This closes the loop. There are many considerations which has to be made, out of which a few important ones are discussed here.

The time delay from the occurrence of a phase drift, until the regulation system starts to correct the error should be short [75, 76]. Otherwise the phase will have time to make large excursions before the phase error is corrected. To meet this demand, the optical path should be made short, and the electronics in the fastest regulation loop should have a short delay. Good passive stability of the laser is also important, since this will lower the rate at which the laser drift, and consequently have impact on how much the phase has drifted before being corrected. The steps taken to improve the passive stability of the laser were described in Section 6.1.

It is also important to maximise the signal-to-noise ratio for the error signal, as the noise will be transferred to the light by the high gain regulation system. It is desirable to have high laser power, since this will increase the signal-to-noise ratio of the detector signal. At the same time, high intensity will create a deep and broad hole, which decreases the slope of the error signal. To compensate for the decrease in error slope, the gain in the regulation system has to be increased, which also amplifies the detection noise. With a large beam diameter both a high intensity and a steep error signal slope, can be obtained at once. Crystal and beam diameters of 19 mm where therefore used in the stabilisation system. Figure 7.17 shows a photograph of the laser stabilisation setup.



Figure 7.17. *Photograph of the laser stabilisation setup.*

LASER STABILISATION ELECTRONICS

Many mistakes can be made when designing a system, and each one of them takes a long time to isolate and solve, especially when one has little prior experience of building similar systems. When working on the electronic and optical design of the laser stabilisation system, I have therefore borrowed ideas and designs from many different places and people. I have written this chapter with this in mind, and have tried present the information in a way that it should be possible for others to borrow ideas from us. Our aim is also to make our HTML- based system documentation available online. All the parts of the laser stabilisation system have been numbered, and to make it easier to follow the descriptions in the chapter the principle of the numbering is described. All the electronic units have designations beginning with a B, e.g. box B300. The printed circuit boards inside a box are given the same number but beginning with a P and ending with a A, B, C... for each consecutive circuit board, e.g. P300A. The designation of cables begins with C, and the number chosen to be the same as a box in their vicinity, for example, C301. This will hopefully make it easier for someone with technical questions.

This chapter begins with an overview of the electronics, Figure 8.1, and the layout of the system, Figure 8.2. The reader will probably find it useful to refer to these figures. This is followed, in Section 8.1, by some general considerations that permeate the design, after which some of the key parts are discussed in more detail. Detailed circuit drawings of several of the sub-systems can be found in the Schematics section before the papers. The chapter ends with a fairly extensive list, Section 8.13, where all the elements in both Figure 8.1 and Figure 8.2 can be found, with references to the where in this chapter they are discussed and, in relevant cases, to the corresponding schematic.




Figure 8.2. Layout of the stabilisation set-up. Descriptions of the designations can be found in Section 8.13

8.1 General design considerations

A great deal of consideration has been dedicated to the details of the electronic and optical set-up for the locking system. In Section 7.3 some requirements were discussed in general terms, while this chapter concentrates on the technical implementation.

8.1.1 Short delay

It is important to keep the delay, i.e. the time from when a phase error occurs until it is corrected, short. To this end the optical path from the laser to the error detector should be as short as possible, since each metre the light travels corresponds to a delay of $t \approx 3.3$ ns. In our design this delay is 11.6 ns. This could be shortened, but this would make the experimental set-up more impractical to work with. The electrical delay should also be short, and this is done by placing the error detector close to the intracavity EOM (electro-optic modulator), keeping the cables short, making sure that the filters and other components have short delays, especially those used for the fast control loop.

8.1.2 Low-noise electronics

The noise in the electronics can be added to the light by the control system. Therefore, the electronic components chosen throughout the construction were often those with the lowest noise readily available on the market.

8.1.3 Avoiding ground loops

Ground loops can add excess noise and therefore efforts were made to avoid these throughout the whole system. As a consequence, many signals are transmitted differentially, the cable shields are typically only connected to ground at one end, many of the power supplies use floating potential, and isolated inputs with optocouplers are used for most digital signals. These precautions were only taken in the parts where the DC level is important, and not in the pure radio frequency (RF) circuitry.

8.1.4 RF shielding

The RF erasure system is basically a large transmitter, transmitting at 10 and 17 MHz. To prevent these radio frequencies from entering the other electronics, all critical systems were placed in RF-shielded boxes, cables with good shielding were employed and contacts that block the RF signals were used for all power supplies and DC signals. We also used coaxial cables with high RF shielding, the RG-174 was replaced by an LMR-100A and the RG-58 by

an LMR-195-UF, both replacements are manufactured by Times Microwave Systems. These cables provide more than 90 dB shielding effectiveness which is about 50 dB better than the cables they replaced.

8.1.5 Avoiding signal reflections

To enable test measurements on the locking system without reflections, all high-speed outputs were constructed with buffers that can drive $50\ \Omega$ loads. The buffers were built with the operational amplifier THS3121 from Texas Instruments, a very good circuit apart from its high input offset voltage, typically 2 mV. All buffers were therefore originally equipped with an auto-zero circuit, based on the Burr-Brown OPA277P, with a maximal offset voltage and offset drift of $20\ \mu\text{V}$ and $0.15\ \mu\text{V}/^\circ\text{C}$, respectively, which should null the THS3121's offset. Unfortunately, there was a design error in the auto-zero construction, and they had to be turned off. However, the offset errors could be compensated for in all critical cases, with manual offset adjusters in other parts of the circuitry.

Circuit board microstrip transmission lines

To minimise reflections as the signals reach and travel through the circuit boards, the necessary strips were constructed as microstrip transmission lines. This simply means that the strips have the same impedance as the coaxial cable and terminating resistance, in most cases $50\ \Omega$. Some strips on the circuit boards were made with $100\ \Omega$ impedance, since this gives both narrower strips and lower losses. For a wide circuit board, which in most cases is a good approximation, the characteristic impedance, Z_0 can be calculated as [98, 99]

$$Z_0 = \frac{87}{\sqrt{\varepsilon_r} + 1.41} \ln \left(\frac{5.98 \times H}{0.8 \times W + T} \right) \Omega, \quad (8.1)$$

where ε_r is dielectric constant of the circuit board material and the dimensions are defined according to Figure 8.3.

8.1.6 Circuit board design considerations

Most of the circuits were constructed using surface mounted components, to save board space, to avoid breaking the ground plane and because the supply of modern ICs circuits in hole-mounted versions are limited. The high-speed circuit-boards P300A, P400A and P500A were constructed using four-layer circuit boards, which makes it easier to maintain a solid ground plane, components can be mounted on both sides, and have large covering power planes. Below is a list of some general general considerations made in the construction.




Figure 8.3. Definitions for microstrip impedance calculations: W and T are the width and thickness of the trace, respectively, and H the height between the trace and reference plane.

- Appropriate decoupling close to the operational amplifiers¹.
- Opening up power and the ground plane where low capacitance is crucial, for example, under the negative feedback on the operational amplifiers.
- Apart from the above, use a completely covering ground plane.
- Either keep a signal transmitter and receiver very close together, or use suitably terminated microstrip transmission lines to transmit the signal.
- Use many, small vias² to connect planes with equal potential.

Figure 8.7 shows a photograph of one of the circuit boards which was constructed, namely P400A.

8.2 Electro-optic modulator

An electro-optical modulator, EOM, is a device where the laser beam passes through a crystal which change its index of refraction when a voltage is applied across it. To apply the voltage, electrodes, often made of gold, are usually attached to the crystal. EOMs can be used in several different ways, two of which we have used and are discussed below.

8.2.1 Phase modulator, B1800

When a sinusoidal voltage is applied to a EOM this will result in sidebands appearing in the laser light, as discussed in Section 7.1.3. The phase modulator used in the stabilisation set-up is a New Focus 4002-M MgO : LiNb₃ modulator. This modulator behaves electrically like a 23 pF capacitor, and can be modulated between 0 and 100 MHz, but the voltage required for a certain modulation is quite large. If the modulator capacitance is used as part of a resonant circuit, the voltage can be stepped up, in what is often referred to as a resonant tank [102, 103]. Our resonant tank design can be seen in Schematic P1700A.

When the EOM is used in the locking system it is important to know the modulation index, $\Delta\phi$, as a function of driving power

¹If, for example, an operational amplifier, is to give an alternating, high-frequency current, the supply current will also have a high frequency part. If the leads to the power supply are long, they will have a large inductance, preventing the high-frequency current, which will lead to a varying voltage at the amplifier. This can be avoided by putting a capacitor between ground and the power potential close to the amplifier. This is generally referred to as capacitive decoupling. More information on capacitive decoupling can be found in for example [100, 101].

²A via is a plated hole connecting different layers of a circuit board.

to the resonant tank, as we see in Paper V. A simple way to determine this is by beating the light after the EOM with unmodulated light. The problem with this approach is that the beat note of the unshifted carriers will have a frequency of zero, and the positive and negative sidebands will have the same frequency. If instead a acousto-optic modulator (AOM) is used as a frequency shifter, as shown in Figure 8.4, with suitably chosen frequencies, all the beat notes can be recorded. With the built-in fast Fourier transform (FFT) function of the oscilloscope, we can isolate each individual beat note, and then use windowed measurements to measure the area of each beat note separately. The voltage amplitude for each beat note is directly proportional to the E-field in that beat note. The crosses in Figure 8.5 show measured data from such an experiment, and the curves are a set of 4 Bessel functions which are simultaneously fitted to the data, using a single fitting parameter (the modulation index). By simultaneously recording the RF power that is fed to the resonant tank, a transfer coefficient can be obtained which relates the RF power applied to modulation index. In our case we obtained:

$$\Delta\phi = \frac{240 \times 10^{-9}}{\lambda} \times U_{\text{RMS}} = \frac{1.70 \times 10^{-6}}{\lambda} \times \sqrt{P}, \quad (8.2)$$

where λ is the wavelength of the light, U_{RMS} is the root mean square voltage driving the resonant tank and P is the corresponding power measured in 50 Ω .

8.2.2 Intracavity EOM, B2300

In Section 6.1 it was explained that the dye laser has two different elements for changing the cavity length, namely a scanning Brewster plate and a folding mirror mounted on the piezoelectric transducer. Neither of these can change the cavity length quickly enough to compensate for the fastest phase changes created by the thickness variations in the dye stream. An EOM inserted into the cavity will allow much faster changes [75, 77–81, 84, 104, 105]. An EOM intended for use inside a cavity must have low transmission losses, and be able to withstand high powers. The Linos PM25 EOM is suitable for this purpose. It has windows at the Brewster angle and two opposing Brewster-cut crystals, which cancel the walk-off [80, 106], leading to a transmission above 99 %. It can also withstand a power of 100 W. When this modulator is inserted into the cavity, it will normally continue lasing, and after a quick realignment, the power will typically have dropped less than 15 %. The bandwidth of the modulator is 100 MHz, the capacitance 30 pF and the connectors to the electrodes are two banana plugs. In Figure 8.8 the modulator can be seen inside the dye laser cavity.




Figure 8.4. Set-up used to measure the E-field in the EOM sidebands and carrier.




Figure 8.5. The crosses correspond to the measured E-field in the EOM sidebands and carrier. The lines are Bessel functions, as a function of modulation index, simultaneously fitted to the measured data (see text).




Figure 8.6. Phase detector output in a $50\ \Omega$ load.

8.3 Error photodetector, B2200

The error photodetector is designed to measure a small intensity variation, with a frequency of 50 MHz, riding on a large background when the system is locked. To achieve a good signal-to-noise ratio the transimpedance of the detector has to be high, and to stop the output from reaching the rail³ the detector should be AC-coupled before the transimpedance amplifier. A Thorlabs PDA10BS-AC was chosen as the error photodetector for our locking system. The detector was custom-made for us by Thorlabs, but is very similar to Thorlabs PDB110A-AC the main differences being that our detector has a lower transimpedance, $1.65 \times 10^4\ \text{V/A}$ at $50\ \Omega$ load, a lower conversion gain on the monitor outputs, $2\ \text{V/mW}$ and a DC saturation power of 10 mW. The AC saturation power was the same as the standard Thorlabs detector, 200 μW . This is a balanced detector, that has two inputs, the signal at one is subtracted from the signal on the other. Balanced detection makes it possible to remove amplitude fluctuations around 50 MHz. There is little amplitude fluctuation in the light emitted by the laser at these frequencies, but there is residual amplitude modulation, RAM, resulting from the electro-optic phase modulator, B1800, even though this was especially designed to have a low RAM. Amplitude fluctuations at low frequencies will anyway be suppressed by the low-pass filter after the mixer. The importance of using balanced detection has not been evaluated.

8.4 Phase detector, B1400

Our phase detector, is a typical double-balanced mixer design, Mini-Circuits LPD-2. This phase detector is designed to be used with a $500\ \Omega$ load on the phase detection output, I. We wanted to use a $50\ \Omega$ load instead, so that we could use $50\ \Omega$ cables, without reflection problems, between the mixer, filter and receiving operation amplifier in the EOM low-voltage generator. The main reason for separating the mixer and the operational amplifier was to avoid having RF circuits in the same box as the receiving circuitry. To characterise this we measured the output of the RPD-2 in a $50\ \Omega$ load, and the result can be seen in Figure 8.6.

8.5 EOM low voltage generator, B300

The task of the EOM low-voltage generator is to quickly send a correcting voltage to the intracavity EOM, B2300, when a phase error is detected. The intracavity EOM behaves electrically like a 30 pF capacitor, which should be quickly charged to a certain voltage, in order to correct for the error. Since only 30 cm of coaxial

³When an operational amplifier reaches its maximum output voltage.

cable would double the capacitance, and thus double the charge required to accomplish a certain voltage change, we decided to place the EOM low-voltage generator directly on top of the modulator, as can be seen in Figure 8.8. The first stage in the circuitry, (see Schematic P300A), consists of only one operational amplifier, with an effective input impedance of $50\ \Omega$, and an output range of $\pm 10\ \text{V}$, which directly drives one of the EOM electrodes⁴. The connection to the EOM was made with circuit-board-mounted banana plugs, the output of the operational amplifier was placed close to the banana plugs and the other circuit board layers were opened up beneath the trace going between the amplifier and the plug. Since the phase detector is sensitive to input offset, two offset adjustment circuits were included, one to compensate for the input bias current of the amplifier, and the other to counterbalance the input offset voltage. This offset adjustment was found to be difficult, and it would have been better to use OPA277-based automatic circuitry to compensate for the voltage offset, similar to that used for the PA85, (see Schematic P400A).

This stage is fast but its output range is only $\pm 10\ \text{V}$ and even though this is sufficient for correcting the fast phase errors, it would reach saturation before the laser's internal control system could take over. Since the total change in refractive index of the EOM is proportional to the differential voltage across it, a high voltage can be applied to the second EOM electrode, which allows for larger but slower voltage excursions [75]. It is more convenient to generate the high voltage in a separate unit, B400, since the heat sink for the high-voltage amplifier is quite large. The signal applied to the first EOM electrode, is buffered and passed to B400 via cable C302, and the high-voltage generated returns to B300 through cable C303. When the first EOM electrode is driven rapidly, the potential of the second electrode should not follow. This is accomplished with a $3\ \text{nF}$ high-voltage capacitor placed between the second electrode and ground. To minimise the loop through which the high-frequency current must travel, this capacitor is placed in B300, making the loop to the fast amplifier as small as possible. The principal role of the EOM low- and high-voltage generators can be seen in Figure 8.1.

8.6 EOM high-voltage generator, B400

In the EOM high-voltage driver the PI regulator is implemented in the high voltage circuit, as can be seen in Schematic P400A. The high-voltage regulator is implemented with an Apex PA85

⁴The amplification of this amplifier found to be unfortunately high, and it would have been better to use a composite or cascaded design. To remedy this and speed up the design of the system we plan to incorporate an RF amplifier after the error photodetector, B2000, which will allow us to decrease the amplification of B300 by approximately a factor 10.




Figure 8.7. The P400A for the high- voltage generator, one of the constructed 4-layer circuit boards.




Figure 8.8. The intracavity EOM, B2300, is the gray cylinder shown by the arrow, with one of the Brewster windows seen protruding on its right side. The black box on top of the modulator is the EOM low-voltage generator, B300. Cables can be seen leaving the back of the box connecting it to the modulator.

operational amplifier. If a low-voltage PI regulator, followed by a high-voltage amplifier had been used, the offset error in the low-voltage PI regulator would have been amplified around 40 times by the high-voltage amplifier. To further decrease the offset error, an automatic offset-adjusting circuitry, based on the Texas Instrument OPA277 was implemented. The high voltage is then passed back to B300 via cable C303, where it is applied to the second electrode.

The high voltage is also measured in two different ways. The first is fairly slow, but has high precision, see Schematic P400A. This design is based on Analog Devices AD629 [107] followed by a combined differential transmitter and LP filter used to give an accumulated error signal, which is the input error signal to the digital control system. The second involves the method used in a normal oscilloscope probe, with a capacitive and resistive voltage divider in parallel. This is not as accurate, but it has the advantage that it is fast. It is followed by a Analog Devices AD823, which is a FET input operational amplifier with low input capacitance, as buffer, so that the probe design does not have to be recalibrated with each new measurement load. As is noted on page 2, note 8 of the P400A schematics, the stray capacitance was slightly larger than expected, which has often been the case in our designs, and therefore C21 should be decreased at some time.

There are also two digital control inputs into the high-voltage generator. The first makes it possible to freeze the integration temporarily, and the second resets the integrator. These signals are further discussed in Section 8.7.

8.7 Suspend regulation controller, B500

The laser beam may be momentarily interrupted by, for example, a speck of dust falling through the beam inside the cavity, or an air bubble in the dye jet stream. If this happens, the integrators in the control system should be temporarily frozen, so that locking can resume with unchanged control signals when the light source returns. We therefore constructed a system that measures the intensity after the crystal, and compares it with a preset reference, see Schematic P500A. One advantage of measuring the intensity after the crystal is that not only sudden interruption of the laser beam causes the intensity to fall, but an extraordinarily fast and large frequency error, as might happen if there is a very large change in diameter of the dye stream, will cause a decrease in intensity. When the intensity falls below the preset value the *release high-voltage integrator* signal to the EOM high-voltage generator goes low, which temporarily freezes the integrator. Another signal is sent to the digital control system to halt the digital integrators.

Furthermore, this system sends the value of the transmitted

intensity to the digital control system. This signal is used when the absorption profile is scanned to determine the frequency lock point, see Section 8.8. There are two additional signals from the digital control system: reset high-voltage integration and forced high-voltage integration. The first is used to reset the high-voltage integrator in B400, before a locking sequence is initiated, and the second is used at the beginning of a new locking sequence, before the spectral hole is deep enough, and the transmitted intensity has therefore not reached the preset reference value, which would release the integrators, as described above.

This system was constructed and built based on our experience, and that of others, that dye lasers sometimes tend to blink, but after implementing the precautions discussed in Section 6.1, this system may not be absolutely necessary.

8.8 Operation software and hardware

The control PC, B1200, runs several different programs, the most important one of which is a Labview/Quview-based program that controls the digital signal processing (DSP) card, which is installed in, and runs asynchronously with, the computer. Several of the system functions can be run with this software. It is possible to control various functions of the amplitude stabilisation system, among them the set point. There is also a function that scans the dye laser frequency up to 30 GHz, without the external locking system being turned on, while the transmitted intensity is monitored with detector B2100. The operator can then choose where on the profile to engage the lock. The main task of this system is, however, to act as an interface to the frequency PID system running on the DSP card and allowing, among other things, the PI parameters to be set and adjusted: the outputs from the PI loop controls, the feed to the dye laser detector unit, B700, and the slow control of the external scan input of the dye laser, B2400. The details of the DSP card are given in Section 8.8.1.

Other programs running on the computer allow the user to control the pump laser, B4000, the temperature controller for the dye and, as discussed in Section 8.8.2, realign the laser beam. The box that connects the DSP card with the rest of the system is discussed in Section 8.8.3

8.8.1 Digital signal processing card

The DSP card was obtained from Sheldon Instruments, and has model number SI-C6713DSP-PCI-256. This card is based on the floating-point DSP TMS320C6713 card from Texas instruments. A daughter module is mounted on top of the DSP card which has 64 single-ended or 32 differential 16-bit analogue inputs, 16 16-bit 180 kHz analogue outputs and 32 digital input/output lines.

The daughter module is denoted SI-MOD6800-250-DDS-16DAC. Connected to the inputs are four analogue-to-digital converters (AD converters), running at a maximum of 250 kHz which can be multiplexed. The multiplexer was found to give quite large, high-frequency spikes, and therefore we finally chose to use only four inputs in total, which enabled us to run the card with the multiplexer function turned off.

The main reason why we chose this card was the software that accompanied it. This software, Quview, makes it possible to program the card from within Labview, very similar to the way Labview is normally used. This approach has the same benefits as Labview, but as the DSP loop grows, it is difficult to keep track of which operation is the most time consuming, and thereby making it more difficult to optimise the code.

Another, similar DSP card that we considered was Toro made by Innovative Integration which has very impressive hardware specifications, for example, 90 dB SFDR. When our system was designed, this card had to be programmed with the Texas Instruments code composer studio, a rather expensive piece of software, considering that only one unit was to be programmed.

8.8.2 Realignment cameras

The realignment cameras, B3800 near field, and B3900 far field, are standard USB-type web cameras (IONetworks WatchPort V), which were mounted very firmly. They are used together with Labview beam analyser program written by us. The software measures the beam position and size, and allows the user to store and recall the earlier parameters and compare them with the current values. This makes it possible to realign the beam, both in direction and convergence, after the laser has been adjusted. In this way, total realignment of the whole set-up can be avoided when the beam path leaving the laser changes.

8.8.3 Break-out box, B100

The break-out box houses signal conditioning circuits that perform two things. Firstly, there are circuits that convert some incoming signals from differential to single-ended signals, a step that had to be taken to save on inputs, when it was found that only 4 inputs could be used simultaneously in this application. Secondly, some signals are filtered. We used two different prefabricated filters, D78L8L with a cut-off frequency of 5 kHz and D74L4L with a cut-off frequency of 10 kHz, both from Frequency Devices Inc. Figure 8.1 shows how the different signals are conditioned.

8.9 Modifications to dye laser detector unit, B700

The built-in locking system in the dye laser works by locking to the side of the transmission fringe of a cavity. The light transmitted through the cavity is measured with a photodiode, and the dye laser's control system strives to maintain this signal at the same level. To avoid the influence of amplitude fluctuations, the signal is normalised against a second reference photodiode. The control signal is added by simply feeding a small current into the signal coming from the reference photodiode. This is done with a differential input amplifier, which takes its power from the dye laser, followed by a 30 k Ω resistor. In this way, the circuits in our locking system and the dye laser locking system can be kept electrically separate.

8.10 Phase shifter, B2000

We used a variable-reactance, reflection phase shifter, although this type of phase shifter has some undesired features. Firstly, the output amplitude changes with the phase shift, secondly, at some phase shifts it creates a signal at twice the fundamental frequency. The first is mainly irritating, and not really a problem, since once the required phase shift has been determined it only needs to be adjusted very little, and the amplitude variations are small. The second problem was resolved by constructing a notch diplexer filter, B1100, blocking the 100 MHz signal. For further information on the filters constructed see Section 8.11. Furthermore, in this set-up we needed very precise control of the phase shift, and the control on our model was too crude. This could, however, easily be fixed by changing the potentiometer controlling the phase shift, from a single-turn version to a 10-turn version. Another way of creating the two phase-shifted signals could be to use a direct digital synthesiser, DDS, with two outputs with separate phase registers, for example Analog Devices AD9854. At the time of the construction of this system this circuit was not available, and we could not find any other suitable DDS. The unit can be bought mounted on an evaluation board, which simplifies the construction. Another benefit of a DDS is the precise, linear control of the phase shift. If necessary, narrow-band crystal filters can be used after the DDS to clean up the signal.

8.11 Diplexer filters

RF filters are often used to block specific frequencies. In a normal RF filter the frequencies that are blocked are reflected back towards the signal source. If the signal source is perfectly terminated, with

the same impedance as the coaxial cable or transmission line, the reflected signal will be absorbed in the termination resistor. However, some components do not have a perfect output impedance. In this case, it would be preferable for the filter to absorb the signal that where not transmitted, instead of reflecting it back towards the source. This is where diplexers prove useful.

Diplexers are filters with two outputs, and instead of reflecting the signal not passed to the first output, it is instead passed to a second output⁵. The most common use of diplexers is probably to combine or separate signals with different frequencies in multi-band broadcasting. In our case, we were only interested in the fact that the reflection is low, and therefore the second output was terminated with a 50 Ω resistor, internally in the diplexer. In some of the filters constructed it would not be possible, even in principle, to route the absorbed signal through a second port, since it may, for example, be partially absorbed in the internal resistance of a coil, but these filters are still referred to as diplexers. Three diplexers were constructed for the locking system. The filters were first simulated using a SPICE-compatible electronics simulator from Spectrum Software, called Micro-cap.

The first diplexer filter, B1000, acts as a low-pass filter on the IF output of the phase detector, B1400. It was most important to make this filter as a diplexer as reflections back into the mixer could affect its performance [109]. This is especially important at DC and on twice the modulation frequency. The signal at 100 MHz is very large and therefore an extra notch diplexer was used to remove this signal. The notch diplexer is a bridge-t (or bridge-tee) diplexer [110], which is the diplexer equivalent of a normal bridge-tee notch filter. In the left column of Figure 8.9 the behaviour of this filter is depicted. The red dashed line shows the filter behaviour simulated, using Micro-cap⁶, and the solid black line is the measured response using a network analyser. The design criteria for this filter were as follows: 3 dB point at 20 MHz which can be seen in panel **b**, a high attenuation at 100 MHz which amounted to 55 dB as seen in **a**, high return loss for all frequencies, which was achieved as seen in **c**, and since this filter is inside the regulation loop, perhaps the most important criterion is, a constant, low delay up to 20 MHz, as shown in **d**. For layout see Schematic P1000A.

The second diplexer filter, B900 (Schematic P900A), is placed after the error detector. This is a 3-stage notch diplexer, the design of which was altered by the author to achieve higher attenuation, with notch frequencies at 10.188, 17.310 and 100 MHz. The first two frequencies are used to block light from coherent effects originating from the erasure fields, and the third to block signals

⁵Some discussion about diplexers can be found on Todd Gale's home page [108].

⁶Another small programs which I found useful was Filter Free from Nuherz technologies.

appearing from the second sidebands, which cause offsets in the locking point [111]. We decided to use such a filter at the design stage, but have not actually measured the amplitude of the signals at any of these frequencies, although this would be simple. On the other hand, we have not tried to run the system without the filter, so we do not know what problems we have avoided. The behaviour of this diplexer filter is depicted in the right column in Figure 8.9. The most important design criteria were again low and constant delay, but here the interesting regime is ± 20 MHz centred around 50 MHz, since this is the signal measured by the phase detector. The measured and simulated delay can be seen in panel **h**. It is apparent here that the delay is longer for the measured case, but the difference is equivalent to the time taken for the signal to travel through about 20 cm of coaxial cable, and although the network analyser was compensated for the connection cables, it still takes the signal about this time to travel across the circuit board, and this was not taken into account in the simulations. The second design goal was high attenuation at the three notch frequencies, 35 dB, 41 dB and 59 dB, as can be seen in panel **e**.

The third diplexer, B1100 (Schematic P1100A), was designed to remove the undesirable 100 MHz signal from the phase shifter (see Section 8.10). It was constructed as a double bridge-t diplexer. The desired 50 MHz signal and the unwanted 100 MHz signal impinge on this filter, which should have low attenuation for the 50 MHz signal and absorb most of the 100 MHz signal. The attenuation at 50 MHz was 0.5 dB, identical the calculated value 0.5 dB, and the measured return loss was 45 dB. The attenuation at 100 MHz was 52 dB, and the calculated value was 62 dB, while the measured return loss was 27 dB. The observant reader may wonder why we did not use the same design as the 100 MHz notch diplexer in B900, which is easier to tune and contains only one stage, but with similar performance. The answer is that this was the first filter made and we had not yet conceived the design used in B900.

8.12 Amplitude stabilisation system

It has been mentioned previously that the free-flowing dye stream causes frequency instability, but it also leads to comparatively large amplitude fluctuations. When the laser is running well, the variations are small enough to for the frequency locking system to work well, but they are too large for the experiments that the laser was intended for. Therefore, we constructed a laser amplitude stabiliser. The basic idea is to measure the amplitude after the dye laser, and correct it by changing the input pump power. Since we wanted a fairly large bandwidth, we chose to change the pump power with an AOM located in the pump beam (B2800), see Fig-

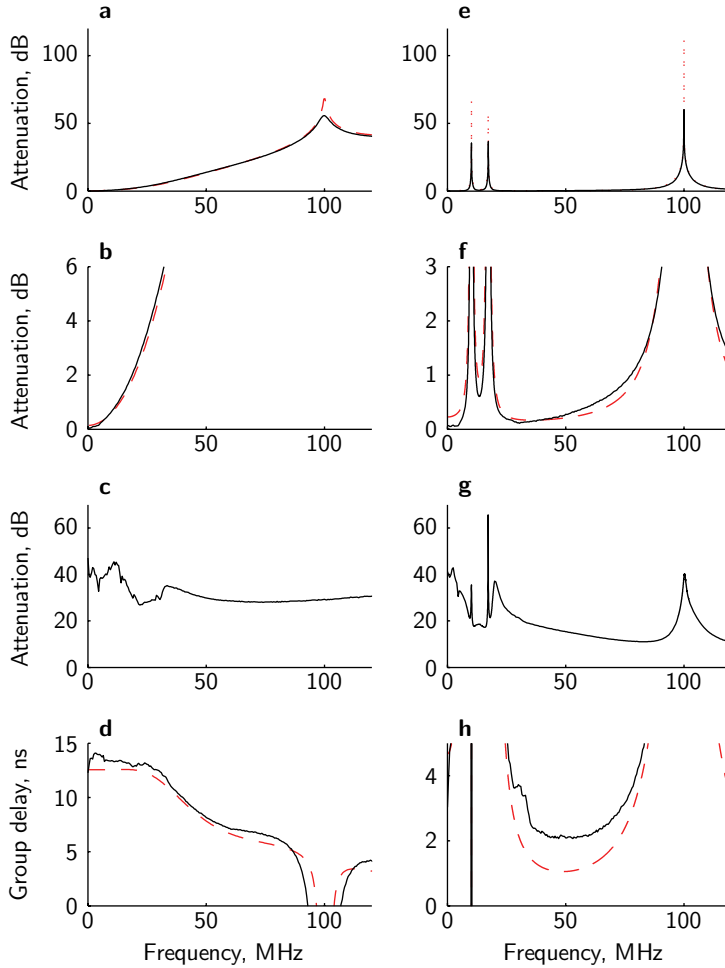


Figure 8.9. The left column shows the response of the 20 MHz low-pass and 100 MHz notch diplexer, B1000, and the right column shows the behaviour of the 3-stage diplexer notch filter, B900, with notches on 10.188, 17.310 and 100 MHz. The dashed red lines show the calculated response, and the solid black lines correspond to measured behaviour. The first and second row show the attenuation of the transmitted signal on different scales, the third row shows the return loss and the last row depicts the group delay.




Figure 8.10. Impedance converter B4200.

ure 8.2. When the AOM is turned off, no power is dumped and the dye laser gives maximum power. One of the benefits of this method is that most of the spatial beam variation introduced by the AOM is presumably removed by the dye laser cavity, which in this respect acts as a mode cleaning-filter. The idea of stabilising the laser in this way was suggested to us by Mike Jefferson [105], in the coffee room at the Hole Burning conference in Bozeman, Montana, 2003.

The stabilisation detector and most of the electronics are all contained in one unit, B200, shown in Figure 8.1. The output current from the photodiode (Hamamatsu S1223), was added to the current from a stable current source with the opposite sign in a summing junction, which simply means that they are connected to each other, before the transimpedance amplifier. The transimpedance is $10\text{ k}\Omega$ and the current source consists of a stable voltage reference and a resistor. It is also possible to use a reference voltage from the control computer, B1200. The controller consists of a normal operational amplifier PI controller with added anti-windup circuitry. To control the power to the AOM, a mixer is used as a current-controlled attenuator, B1400. The power to the AOM is then boosted with a 5 W amplifier, B2900.

The system's ability to suppress amplitude variations has not been analysed as a function of frequency, but the remaining root mean square fluctuation in our setup is better than 0.1% RMS. The bandwidth of the system is about 1 MHz .

8.13 Table of systems, components and descriptions

Systems and components

B100, Break-out box, see Section 8.8.3.

B200, Amplitude stabilisation, see Section 8.12 for further information.

B300, EOM low-voltage generator, see Schematic P300A and Section 8.5 for details.

B400, EOM high-voltage generator, see Schematic P400A and Section 8.6 for further information.

B500, Suspend regulation controller, see Schematic P500A and Section 8.7 for details.

B600, Dye laser, modified Coherent 699-21, control box.

B700, Dye laser, modified Coherent 699-21, detector unit, see Section 8.9.

- B800**, Stabilised power supply for low-voltage generator B300, not shown in any figures.
- B900**, 3-stage diplexer notch filter, 10.188, 17.310 and 100 MHz, see Schematic [P900A](#) and Section [8.11](#) for further information.
- B1000**, 20 MHz low-pass and 100 MHz notch diplexers with enable, see Schematic [P1000A](#) and Section [8.11](#) for further information.
- B1100**, 2-stage 100 MHz notch diplexer filter, see Schematic [P1100A](#) and Section [8.11](#) for further information.
- B1200**, Digital laser control PC, see Section [8.8](#) for further information.
- B1300**, Power supply for B500, Thurlby Thandar Instruments EB2025T, linear, not shown in any figures.
- B1400**, Phase detector, Mini-Circuits RPD-2, see Schematic [P1400A](#).
- B1500**, Power supply for B100, Thurlby Thandar Instruments EB2025T, linear, not shown in any figures.
- B1600**, Dual output 50 MHz crystal oscillator, Wenzel Sprinter, 500-14426, SC-cut crystal and SMA connectors.
- B1700**, 50 MHz resonant tank for phase modulator B1800, see Section [8.2.1](#) and Schematic [P1700A](#) for further information.
- B1800**, 50 MHz electro-optical phase modulator, New Focus 4002-M, see Section [8.2.1](#) for more information.
- B1900**, Phase modulator amplifier, Mini-Circuits ZHL-2010, amplification +20 dB, max output power +26 dBm.
- B2000**, 50 MHz phase shifter, Tele-Tech Corp, 360°, +10 dBm max input, see also Section [8.10](#).
- B2100**, Transmitted intensity detector, home built, photodiode Hamamatsu S1223, transimpedance 10 k Ω .
- B2200**, Error photodetector, Thorlabs PDA10BS-AC. See Section [8.3](#) for further information.
- B2300**, Intracavity EOM, Linos PM25 VIS APD 400-650 nm model number 50 2075, see also Section [8.2.2](#).
- B2400**, Dye laser, modified Coherent 699-21, see also Section [6.1](#).

- B2500**, Power supply for B800, Thurlby Thandar Instruments, EB2025T, linear, not shown in any figures.
- B2600**, ± 18 V power supply for B400, HQ power PS23023, linear, not shown in any figures.
- B2700**, ± 210 V power supply for B400, made from two XP Power LNB-0.12 which are the same as International power IHB200-0.12, modified according to [112], not shown in any figures.
- B2800**, Amplitude stabilisation AOM, Brimrose, FQM-80-1-.532/WC. The transmission of this AOM is 98 %, power handling 100 W/mm² and it can be used in the pump beam. See also Section 8.12
- B2900**, 5 W amplifier, Mini-Circuits ZHL-5W-1.
- B3000**, Power supply for B2900, HQ Power PS3003, linear, not shown in any figures
- B3100**, Power supply for B1600, not shown in any figures
- B3200**, Power supply for B2000, not shown in any figures
- B3300**, Power supply for B2200, not shown in any figures
- B3400**, Directional coupler, Mini-Circuits ZX30-20-4
- B3500**, Directional coupler, Mini-Circuits ZX30-20-4
- B3600**, Amplifier, Mini-Circuits ZFL-500HLN
- B3700**, Directional coupler, Mini-Circuits ZMDC-20-3
- B3800**, Realignment camera, close, IONetworks WatchPort V. For further information see Section 8.8.2
- B3900**, Realignment camera, far, IONetworks WatchPort V. For further information see Section 8.8.2
- B4000**, Pump laser, Coherent Verdi V6, see also Section 6.1
- B4100**, Voltage-controlled attenuator (mixer), Mini-Circuits ZAD-3
- B4200**, Impedance converter, see Figure 8.10
- B4300**, 80 MHz voltage-controlled oscillator, Mini-Circuits ZOS-100, followed by step variable attenuator 0 to -11 dB with step size 0.1 dB, JFW Industries, 50DR-060
- B4400**, Power supply for B200, not shown in any figures

B4500, Continuous flow helium cryostat with locking crystal, Oxford Instruments, Optistat CFV . The crystal is surrounded by two RF coils

B4600, Resonant circuits for two-tone RF erasure. This is described in Paper [V](#)

B4700, Waveform generator for RF erasure, Thurlby Thandar Instruments, TGA1244. The used RF pattern is discussed in Paper [V](#)

Key to Figure [8.1](#)

CB, Current booster

DCB1, DC block, Mini-Circuits BLK-89

DC1, Directional coupler, Mini-Circuits ZMDC-20-3

DC2, Directional coupler, Mini-Circuits ZX30-20-4

DM, Digital monitor output

GS, Ground sense

HV, High voltage

I, Integrator regulator, can also have a proportional part

LP, Low-pass filter, in some cases combined with notch filter

M, Monitor output

N, Notch filter, sometimes notches at more than one frequency

Ref, Voltage reference, variable

RN, Resistor network, see Figure [8.10](#)

RT, Resonant tank, see Section [8.2.1](#) and Schematic [P1700A](#) for further information

Rt, Transimpedance

-1 dB, -6 dB,..., Attenuators, one or several Mini-Circuits VAT-1, VAT-6,..., unless otherwise stated.

a In $50\ \Omega$. In high-impedance loads the voltage will be twice as high.

b Reset high-voltage integrator

c Forced high-voltage integration

- d** Transmitted intensity
- e** Suspend regulation
- f** Release high-voltage integrator
- g** AOM control signal
- h** External disable
- i** Intensity minus set point
- j** External set point

Key to Figure 8.2

- 1** High-power mirror from CVI
- 2** Periscope, from 135 mm coming from the laser to 75.4 mm which is the beam height for the rest of the set-up. Made with 1.5 inch post and and two stable post clamps, Thorlabs P200/M, C1501/M, and two mirror mounts with mirrors
- 3** Zoom beam expander, Sill Optics S6ASS2075/067 450-650 nm.
- 4** The camera locked at the finely ground back of a Thorlabs BB1-E02 mirror, through an absorptive OD 1 ND filter
- 5** Beam pick-off, Thorlabs BSF10-A1, picks off about 5 % of the light for the frequency locking system
- 6** Mirror with polished back, CVI BBD1-PM-1037-C
- 7** Beam splitter, Thorlabs BSW07
- 8** Beam shrinker, to make the beam fit through the EOM, B1800
- 9** α -BBO polariser, Casix PGL8308
- 10** Beam expander, to expand the beam to 19 mm, which is the size of the crystal
- 11** Iris diaphragm
- 12** Achromatic $\lambda/2$ -wave plate, Photop PAWP-I-0105 460-700 nm
- 13** Lens
- 14** Circular continuously variable ND filter
- 15** Beam splitter, Thorlabs BSW13

EPILOGUE

The focus of my PhD studies has been very clear: to create a CNOT in a rare-earth-metal-ion-doped crystal, and the work leading to this has been extremely interesting. During the course of this process many puzzling questions have arisen, such as: why can't we burn a hole which is so deep that the fluorescence stops? and where do all these other peaks come from when we try to burn back a single peak in a pit? These questions have been answered, one by one.

The demands on the equipment required for the CNOT gate are much higher than those for photon echo experiments, that was previously carried out in this research group. The experimental work has therefore been combined with the design of a new system, with which it is possible to create laser light with sufficient frequency and amplitude stability, and to carve out the precise pulses necessary for the CNOT experiment. This work is now completed, and we strongly believe that a CNOT gate experiment based on the ensemble approach will be performed very soon.

Since our own stabilised laser system has only recently been finished, we ran the experiments leading up to Papers [II](#) and [III](#) together with Robert Klieber, using Professor Dieter Suter's (Dortmund, Germany) stabilised dye-laser with our modulator and waveform generator set-up. These experiments were planned for over one year ahead, buying modulators and waveform generators that would allow us to perform coherent light pulses. We only had three weeks to do the experiments and prior to this we had not tried to generate a complex hyperbolic secant pulse, had not transferred populations between states in an efficient way using such pulses and had not measured the properties of single hyperfine levels in these systems, and to the best of the author's knowledge neither had anybody else. As we only had a limited time to perform the experiments, we had made a detailed 16-point list of exactly which experiments we wanted to carry out, and had programmed the necessary pulse sequences, in advance. The first thing we did was to break the modulator which was essential for the experiments. This set us back for a couple of days, but after a week when we got started on point number one on the list, everything just worked.

Sitting in a laboratory with a long list of experiments that no one had ever been near trying before, checking off point after point was an nice feeling, and when we left for home we had successfully all the 16 experiments completed. For those of you who are not experimentalists I will tell you this is not a very common occasion. This give us some hope that the CNOT experiment could be run without too many problems after having worked towards it for 2 and a half years.

So what will follow the CNOT experiment? We have an idea for addressing a single quantum computer, with a single ion per qubit, using a specialised read-out ion. Working with single ions will again require new technology, and it will be very interesting to see what lies ahead of us in this area of research.

COMMENTS ON THE PAPERS

I Initial experiments concerning quantum information processing in rare-earth-ion doped crystals

The first experiments, aiming at creating a peak of ions, are described in this paper. This work was carried out in $\text{Eu}^{3+}:\text{Y}_2\text{SiO}_5$. The dipole-dipole interaction between ions was also investigated in $\text{Tm}^{3+}:\text{YAG}$.

I participated in the planning and took part in all the experiments.

II Holeburning techniques for isolation and study of individual hyperfine transitions in inhomogeneously broadened solids, demonstrated in $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$

The transitions from hyperfine levels are, in many rare-earth-metal-ion-doped crystals, hidden under the broader inhomogeneous line, which makes it difficult to measure their individual properties. This paper describes how a narrow peak of ions, all absorbing on the same transition, is isolated in a pit with no other ions. All the nine different hyperfine transitions are isolated and their relative transition strengths can thereby be measured.

I participated in the planning and took part in all the experiments.

III Experimental demonstration of efficient and selective population transfer and qubit distillation in a rare-earth-metal-ion-doped crystal

This paper describes how peaks of ions, which forms the qubits in the rare-earth-metal-ion-doped quantum computing scheme, are manipulated coherently. A peak of ions is first created in one of the hyperfine level of the ground state. Using two consecutive complex hyperbolic secant pulses the peak is transferred to another hyperfine levels of the ground state via the optically excited state with a transfer efficiency of 93 %. It was shown how the dipole-dipole interaction broadens an ion peak when another peak is excited, and this was used to distill out only those ions with large frequency shifts.

I participated both in the planning and in the experimental work. I also evaluated all the data and wrote the manuscript.

IV Scalable designs for quantum computing with rare-earth-ion-doped crystals

This paper points out and explains the difficulties related to scaling of the previous rare-earth element based quantum computing schemes. Two different strategies are presented for how to modify the schemes in order to solve this problem. It is also shown that fast gates require strong interaction between the qubits.

I came up with the idea to use a special read-out ion to design a quantum computer with a single ion in each qubit.

V Laser stabilization using spectral hole burning

This paper presents a dye laser that is frequency and phase stabilised against a spectral hole in a $\text{Pr}^{3+}:\text{Y}_2\text{SiO}_5$ crystal. An elaborate analytical theory was also developed describing hole stabilisation, including a criteria that must be fulfilled to avoid large drifts.

I planned, designed and built large parts of the laser system. I came up with the idea to repopulate the hyperfine levels using RF fields. I contributed to the theory described in the paper, and participated in the experiments.

VI Mode-hop-free electro-optically tuned diode laser

This paper describes an external cavity diode laser, that can be tuned over 50 GHz without mode-hop. This is achieved by an intra-cavity electro-optic modulator which changes the optical cavity length, and deflects the beam simultaneously. This laser was used in Paper [I](#), and it, and copies of it, has been used in several other experiments [\[113–115\]](#).

I came up with the idea of how to achieve mod-hop-free tuning and did the optical design for the laser. I drew all the mechanical components in Pro/Engineer and manufactured most of them in an automatic milling machine. I assembled, adjusted and tested the laser. I wrote the manuscript.

ACKNOWLEDGEMENTS

The main reason that I started to do research in this group, is that I thought my tutor Stefan Kröll is such a nice person to work for, and during my time here, this opinion has become firmer founded. I think Stefan creates a extremely nice atmosphere to do research in and keeps on pushing on the right issues and for this I thank you Stefan.

I would also like to express great gratitude to my co-supervisors, Krishna Rupavatharam and Brian Julsgaard, for your way of always taking time to answer my questions, over a cup of coffee or two.

For introducing me to atomic physics, through the course in atomic and molecular spectroscopy, which is taught in a most vivid way, I thank Sune Svanberg.

I would like to thank my fellow PhD students in the Quantum Information group Andreas Walther and Atia Amari, for your friendliness and help, and all the diploma workers with whom I have had the pleasure to work with, in our group.

I would to thank the former group members Nicklas Ohlsson and Mattias Nilsson, for all the interesting discussions about physics and other things.

During the construction of the diode laser I had much help from Ulf Gustafsson and Jonas Sandsten, and this was greatly appreciated.

Mike Jefferson and Pete Sellin are thanked for generously sharing their knowledge about laser stabilisation with me.

I thank our electronic engineers, first Åke Bergquist and later Thomas Wendel for all the help with electronic construction, sometimes on the most awkward of hours. I also thank Björn Lundberg, for the circuit layout of some of the electronic designs, Göran Jönsson, for helping out with RF-measurements and Bertil Hermansson for all the discussions and for making sure that our computers work.

I am also grateful to all the people in the mechanical workshop, for the help manufacturing many of the mechanical components.

During our visits to University of Dortmund, Robert Klieber was always ready to help us out, keeping the liquid helium running

and making sure that the laser always worked, and for this I thank you.

Anders Persson is appreciated for always keeping the door open to answer questions about optics, for never hesitating to lend me equipment and for creating a friendly atmosphere at our department, especially for the newcomers.

I thank Johan Mauritsson for all the discussions in the *early rises club*, giving me a glimpse into short pulses interaction with noble gases, for arranging the nice wine tastings at the department and elsewhere, and for being a good friend.

I would like to thank all my colleges at the department, for making it such a nice place to work at, especially Tomas Svensson how also took the photographs of this thesis, Per Johnsson who also helped out with all my questions about L^AT_EX, and Mats Andersson for all the discussions about optronics.

I thank Leif Magnusson and Struan Gray, the most friendly liquid helium suppliers one could ask for.

Elsie, for always taking the time to say a few friendly words and letting me in, in the mornings, when I forgot my keys.

I am also grateful to all the people that keep all the administrative things running at the department, Henrik Steen which pays our bills and is never late to share a word, Minna Ramkull our present secretary which fixes everything with a smile and Marie Holmdahl-Svensson and Britt-Marie Hansson, which is no longer working at the department. I want to thank Anneli Nilsson at the department for Combustion Physics for helping out with different matters. I also thank Laila Lewin, how unfortunately is no longer with us.

I would like to thank Anna and Bengt Rippe, for all their help with my daughters, and for all the nice meals I have been invited to share with you, during my PhD time.

I thank my parents for their fantastic support during my studies, for helping out with Freja and Matilda during my work with this thesis, and for creating a resort for my family. I would also like to thank my sister, Sus Levin, for dropping Freja off at school, lending me here car when my bike breaks down, and all the other things you help me out with, but most of all for being a great friend.

I would like to thank my two daughters, Freja and Matilda, you are the most fantastic thing that has ever happened to me, and you always bring things into perspective.

And finally I would like to thank Catarina, for helping me out with this thesis and your endless love and support. I love you.

REFERENCES

1. wikipedia. Supercomputer (2006). URL <http://en.wikipedia.org/wiki/Supercomputer>.
2. P. W. Shor. *Algorithms for quantum computation: discrete logarithms and factoring*. Proceedings of the IEEE 35 th Annual Symposium on Foundations of Computer Science pages 124–134 (1994).
3. wikipedia. Shor’s factorisation algorithm (2006). URL http://en.wikipedia.org/wiki/Shor%27s_algorithm.
4. L. K. Grover. *A fast quantum mechanical algorithm for database search*. Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing pages 211–219 (1996).
5. wikipedia. Grover’s search algorithm (2006). URL http://en.wikipedia.org/wiki/Grover%27s_algorithm.
6. R. P. Feynman. *Simulating Physics with Computers*. International Journal Of Theoretical Physics **21**, 467–488 (1982).
7. S. Lloyd. *Universal quantum simulators*. Science **273**, 1073–1078 (1996).
8. P. W. Shor. *Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer*. Siam J. Comput. **26**, 1484–1509 (1997).
9. L. K. Grover. *Quantum mechanics helps in searching for a needle in a haystack*. Phys. Rev. Lett. **79**, 325–328 (1997).
10. Y. Sun, C. W. Thiel, R. L. Cone, R. W. Equall and R. L. Hutcheson. *Recent progress in developing new rare earth materials for hole burning and coherent transient applications*. J.Lumin. **98**, 281–287 (2002).
11. J. I. Cirac and P. Zoller. *Quantum computations with cold trapped ions*. Phys. Rev. Lett. **74**, 4091–4094 (1995).
12. D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King and D. M. Meekhof. *Experimental issues in coherent quantum-state manipulation of trapped atomic ions*. J. Res. Natl. Inst. Stan. **103**, 259–328 (1998).
13. F. Schmidt-Kaler, H. Haffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner and R. Blatt. *Realization of the Cirac-Zoller controlled-NOT quantum gate*. Nature **422**, 408–411 (2003).

14. H. Haffner, W. Hansel, C. F. Roos, J. Benhelm, D. Chek-al kar, M. Chwalla, T. Korber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Guhne, W. Dur and R. Blatt. *Scalable multiparticle entanglement of trapped ions*. Nature **438**, 643–646 (2005).
15. D. Leibfried, E. Knill, S. Seidelin, J. Britton, R. B. Blakestad, J. Chiaverini, D. B. Hume, W. M. Itano, J. D. Jost, C. Langer, R. Ozeri, R. Reichle and D. J. Wineland. *Creation of a six-atom 'Schrödinger cat' state*. Nature **438**, 639–642 (2005).
16. B. Julsgaard, J. Sherson, J. I. Cirac, J. Fiurasek and E. S. Polzik. *Experimental demonstration of quantum memory for light*. Nature **432**, 482–486 (2004).
17. G. H. Dieke and H. M. Crosswhite. *The spectra of the doubly and triply ionized rare earths*. Appl.Opt. **2**, 675–686 (1963).
18. W. M. Yen, Selzer P. M., G. F. Imbush, Kopelman R., Holstein T., Lyo S. K., R. Orbach, Kuber D. L., Weber M. J. and Francis A. H. *Laser spectroscopy of solids*. Springer-Verlag Berlin (1981).
19. R. M. Macfarlane and R. M. Shelby. *Coherent transient and holeburning spectroscopy of rare earth ions in solids*. In A.A. Kaplyanskii and R.M. Macfarlane, editors, *Modern problems in condensed matter sciences*. North-Holland Amsterdam (1987).
20. T. Böttger, C. W. Thiel, Y. Sun and R. L. Cone. *Optical decoherence and spectral diffusion at $1.5\mu\text{m}$ in $\text{Er}^{3+} : \text{Y}_2\text{SiO}_5$ versus magnetic field, temperature, and Er^{3+} concentration*. Phys. Rev. B **73**, 075101 (2006).
21. E. Fraval, M. J. Sellars and J. J. Longdell. *Dynamic decoherence control of a solid-state nuclear-quadrupole qubit*. Phys. Rev. Lett. **95**, 030506 (2005).
22. E. Fraval, M. J. Sellars and J. J. Longdell. *Method of extending hyperfine coherence times in $\text{Pr}^{3+} : \text{Y}_2\text{SiO}_5$* . Phys. Rev. Lett. **92**, 077601 (2004).
23. R. W. Equall, Y. Sun, R. L. Cone and R. M. Macfarlane. *Ultraslow optical dephasing in $\text{Eu}^{3+} : \text{Y}_2\text{SiO}_5$* . Phys. Rev. Lett. **72**, 2179–2181 (1994).
24. F. Könz, Y. Sun, C. W. Thiel, R. L. Cone, R. W. Equall, R. L. Hutcheson and R. M. Macfarlane. *Temperature and concentration dependence of optical dephasing, spectral-hole lifetime, and anisotropic absorption in $\text{Eu}^{3+} : \text{Y}_2\text{SiO}_5$* . Phys. Rev. B **68**, 085109– (2003).
25. M. J. Sellars, E. Fraval and J. J. Longdell. *Investigation of static electric dipole-dipole coupling induced optical inhomogeneous broadening in $\text{Eu}^{3+} : \text{Y}_2\text{SiO}_5$* . J.Lumin. **107**, 150–154 (2004).
26. Mattias Nilsson. *Coherent interactions in rare-earth-ion-doped crystals for applications in quantum information science*. PhD thesis Division of Atomic Physics, LTH (2004).
27. Robert. C. Hilborn. *Einstein coefficients, cross-sections, f values, dipole moments, and all that*. Am.J.Phys. **50**, 982–986 (1982).
28. N. Ohlsson, R. K. Mohan and S. Kröll. *Quantum computer hardware based on rare-earth-ion-doped inorganic crystals*. Opt. Commun. **201**, 71–77 (2002).

29. W. E. Moerner, K. K. Rebane, L. A. Rebane, D. Haarer, R. M. Macfarlane, R. M. Shelby, J. M. Hayes, R. Jankowiak, G. J. Small, A. J. Sievers, W. Lentz and G. C. Bjorklund. *Persistent spectral hole-burning: Science and applications*. Number 44. Springer-Verlag Berlin (1988).
30. R. M. Macfarlane. *High-resolution laser spectroscopy of rare-earth doped insulators: a personal perspective*. J.Lumin. **100**, 1–20 (2002).
31. T. Böttger, Y. Sun, G. J. Pryde, G. Reinemer and R. L. Cone. *Diode laser frequency stabilization to transient spectral holes and spectral diffusion in $Er^{3+} : Y_2SiO_5$ at 1536 nm*. J. Lumin. **94**, 565–568 (2001).
32. R. M. Macfarlane. *Photon-Echo Measurements On the Trivalent Thulium Ion*. Opt.Lett. **18**, 1958–1960 (1993).
33. P. W. Milonni and J. H. Eberly. *Lasers*. John Wiley & Sons New York (1988).
34. F. Jonsson. *Lecture notes on nonlinear optics* (2006). URL <http://www.laserphysics.kth.se/nlopt/nlo2003.pdf>.
35. M. A. Nielsen and I. L. Chuang. *Quantum computation and quantum information*. Cambridge University Press (2000).
36. F. Bloch. *Nuclear induction*. Phys. Rev. **70**, 460–474 (1946).
37. R. P. Feynman, F. L. Vernon and R. W. Hellwarth. *Geometrical representation of the schrodinger equation for solving maser problems*. J. Appl. Phys. **28**, 49–52 (1957).
38. M. S. Silver, R. I. Joseph and D. I. Hoult. *Selective spin inversion in nuclear-magnetic resonance and coherent optics through an exact solution of the Bloch Riccati equation*. Phys. Rev. A **31**, R2753–R2755 (1985).
39. A. Tannus and M. Garwood. *Adiabatic pulses*. NMR Biomed. **10**, 423–434 (1997).
40. E. Kupce and R. Freeman. *Adiabatic pulses for wide-band inversion and broad-band decoupling*. J. Magn. Reson., Ser A **115**, 273–276 (1995).
41. K. E. Cano, M. A. Smith and A. J. Shaka. *Adjustable, broadband, selective excitation with uniform phase*. J. Magn. Reson. **155**, 131–139 (2002).
42. L. Mitschang, B. Ittermann, F. Schubert and H. Rinneberg. *Design of a constant adiabaticity pulse for selective population inversion*. J. Magn. Reson. **168**, 103–109 (2004).
43. D. Rosenfeld, S. L. Panfil and Y. Zur. *Optimization of adiabatic selective pulses*. J. Magn. Reson. **126**, 221–228 (1997).
44. E. Kupce and R. Freeman. *Optimized adiabatic pulses for wideband spin inversion*. J. Magn. Reson., Ser A **118**, 299–303 (1996).
45. E. Kupce and R. Freeman. *Stretched adiabatic pulses for broadband spin inversion*. J. Magn. Reson., Ser A **117**, 246–256 (1995).
46. qubit.org (2006). URL <http://www.qubit.org/>.
47. quantiki (2006). URL <http://www.quantiki.org/>.
48. qwiki (2006). URL <http://qwiki.caltech.edu>.

49. wikipedia. Quantum computing (2006). URL http://en.wikipedia.org/wiki/Quantum_computing.
50. Atomic Physics, LTH (2006). URL <http://www.atom.fysik.lth.se/QI/>.
51. Danish National Research Foundation Center for Quantum Optics (2004). URL <http://www.phys.au.dk/quantop/>.
52. Laboratoire Aim Cotton (2004). URL http://www.lac.u-psud.fr/LAC/us_index.html.
53. Laboratoire de Chimie Applique de l'Etat solide (2004). URL http://www.enscp.fr/en_recherche_index.html.
54. Centre Interdisciplinaire de Recherche Ions Lasers (2004). URL http://www.ganil.fr/ciril/index_gb.htm.
55. J. J. Longdell and M. J. Sellars. *Experimental demonstration of quantum-state tomography and qubit-qubit interactions for rare-earth-metal-ion-based solid-state qubits*. Phys. Rev. A **69**, 032307 (2004).
56. J. J. Longdell, M. J. Sellars and N. B. Manson. *Demonstration of conditional quantum phase shift between ions in a solid*. Phys. Rev. Lett. **93**, 130503 (2004).
57. M. Z. Tian, Z. W. Barber, J. A. Fischer and W. R. Babbitt. *Geometric manipulation of the quantum states of two-level atoms*. Phys. Rev. A **69**, 050301(R) (2004).
58. J. Stolze and D. Suter. *Quantum Computing: A short course from theory to experiment*. John Wiley & Sons (2004).
59. A. Church. *An unsolvable problem of elementary number theory*. Amer. J. Math. **58**, 345–363 (1936).
60. A. M. Turing. *On computable numbers, with an application to the entscheidungsproblem*. Proc. London Math. Soc. **42**, 230– (1936).
61. D. P. DiVincenzo. *Two-bit gates are universal for quantum computation*. Phys. Rev. A **51**, 1015–1022 (1995).
62. A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin and H. Weinfurter. *Elementary gates for quantum computation*. Phys. Rev. A **52**, 3457–3467 (1995).
63. D. P. DiVincenzo. *Quantum computation*. Science **270**, 255–261 (1995).
64. I. Roos and K. Mølmer. *Quantum computing with an inhomogeneously broadened ensemble of ions: Suppression of errors from detuning variations by specially adapted pulses and coherent population trapping*. Phys. Rev. A **69**, 022321– (2004).
65. N. A. Gershenfeld and I. L. Chuang. *Bulk spin-resonance quantum computation*. Science **275**, 350–356 (1997).
66. H. R. Barry, B. Bakowski, L. Corner, T. Freegarde, O. T. W. Hawkins, G. Hancock, R. M. J. Jacobs, R. Peverall and G. A. D. Ritchie. *OH detection by absorption of frequency-doubled diode laser radiation at 308 nm*. Chem. Phys. Lett. **319**, 125–130 (2000).
67. S. Urabe, M. Watanabe, H. Imajo, K. Hayasaka, U. Tanaka and R. Ohmukai. *Observation of Doppler sidebands of a laser-cooled Ca⁺ ion by using a low-temperature-operated laser diode*. APPLIED PHYSICS B-LASERS AND OPTICS **67**, 223–227 (1998).


-
68. T. H. Maiman. *Stimulated optical radiation in ruby*. Nature **187**, 493–(1960).
 69. T. J. Kane and R. L. Byer. *Monolithic, unidirectional single-mode nd-yag ring laser*. Opt. Lett. **10**, 65–67 (1985).
 70. T. J. Kane, A. C. Nilsson and R. L. Byer. *Frequency stability and off-set locking of a laser-diode-pumped Nd-YAG monolithic nonplanar ring oscillator*. Opt. Lett. **12**, 175–177 (1987).
 71. S. Bahbah, F. Bretenaker and C. Drag. *Single-frequency quasi-continuous red radiation generated by a green-pumped singly resonant optical parametric oscillator*. Opt. Lett. **31**, 1283–1285 (2006).
 72. J. C. Bienfang, C. A. Denman, B. W. Grime, P. D. Hillman, G. T. Moore and J. M. Telle. *20 W of continuous-wave sodium D-2 resonance radiation from sum-frequency generation with injection-locked lasers*. Opt.Lett. **28**, 2219–2221 (2003).
 73. C. A. Denman, P. D. Hillman, G. T. Moore, J. M. Telle, J. D. Drummond and A. L. Tuffli. *20 W CW 589 nm sodium beacon excitation source for adaptive optical telescope applications*. OPTICAL MATERIALS **26**, 507–513 (2004).
 74. wikipedia. Laser guide star (2006). URL http://en.wikipedia.org/wiki/Laser_guide_star.
 75. R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley and H. Ward. *Laser phase and frequency stabilization using an optical- resonator*. Appl. Phys. B **31**, 97–105 (1983).
 76. M. Zhu and J. L. Hall. *Stabilization of optical-phase frequency of a laser system - application to a commercial dye-laser with an external stabilizer*. J. Opt. Soc. Am. B **10**, 802–816 (1993).
 77. J. Hough, D. Hils, M. D. Rayman, L. S. Ma, L. Hollberg and J. L. Hall. *Dye-laser frequency stabilization using optical resonators*. Appl. Phys. B **33**, 179–185 (1984).
 78. J. Helmcke, J. J. Snyder, A. Morinaga, F. Mensing and M. Glaser. *New ultrahigh resolution dye-laser spectrometer utilizing a non-tunable reference resonator*. Appl. Phys. B **43**, 85–91 (1987).
 79. M. Houssin, M. Jardino, B. Gely and M. Desaintfuscien. *Design and performance of a few-kilohertz-linewidth dye-laser stabilized by reflection in an optical-resonator*. Opt. Lett. **13**, 823–825 (1988).
 80. R. Kallenbach, C. Zimmermann, D. H. McIntyre, T. W. Hansch and R. G. Devoe. *A blue-dye laser with sub-kilohertz stability*. Opt. Commun. **70**, 56–60 (1989).
 81. I. Steiner, V. Enders, F. Elsner, W. Neuhauser, P. E. Toschek, R. Blatt and J. Helmcke. *A dye ring-laser spectrometer for precision spectroscopy*. Appl. Phys. B **49**, 251–256 (1989).
 82. M. Zhu and J. L. Hall. *Short and long term stability of optical oscillators*. In *Proceedings of the 1992 IEEE Frequency Control Symposium* pages 44–55 Pennsylvania (1992). Hershey.
 83. M. Zhu and J. L. Hall. *Frequency stabilization of tunable lasers* chapter 5, pages 103–136. Academic press (1997).

84. B. C. Young, F. C. Cruz, W. M. Itano and J. C. Bergquist. *Visible lasers with subhertz linewidths*. Phys. Rev. Lett. **82**, 3799–3802 (1999).
85. R. L. Barger, M. S. Sorem and J. L. Hall. *Frequency stabilization of a CW dye laser*. Appl. Phys. Lett. **22**, 573–575 (1973).
86. M. Gray, D. Shaddock and McClelland D. Tilt Locking for laser stabilization. Technical report Department of Physics Australian National University Canberra (2000).
87. D. A. Shaddock, M. B. Gray and D. E. McClelland. *Frequency locking a laser to an optical cavity by use of spatial mode interference*. Opt. Lett. **24**, 1499–1501 (1999).
88. P. B. Sellin, N. M. Strickland, J. L. Carlsten and R. L. Cone. *Programmable frequency reference for subkilohertz laser stabilization by use of persistent spectral hole burning*. Opt. Lett. **24**, 1038–1040 (1999).
89. G. J. Pryde, T. Böttger, R. L. Cone and R. C.C. Ward. *Semiconductor lasers stabilized to spectral holes in rare earth crystals to a part in 10^{13} and their application to devices and spectroscopy*. J. Lumin. **98**, 309–315 (2002).
90. K. D. Merkel, R. D. Peters, P. B. Sellin, K. S. Repasky and W. R. Babbitt. *Accumulated programming of a complex spectral grating*. Opt. Lett. **25**, 1627–1629 (2000).
91. N. M. Strickland, P. B. Sellin, Y. Sun, J. L. Carlsten and R. L. Cone. *Laser frequency stabilization using regenerative spectral hole burning*. Phys. Rev. B **62**, 1473–1476 (2000).
92. P. B. Sellin, N. M. Strickland, T. Böttger, J. L. Carlsten and R. L. Cone. *Laser stabilization at 1536 nm using regenerative spectral hole burning*. Phys. Rev. B **63**, 155111 (2001).
93. T. Böttger, G. J. Pryde and R. L. Cone. *Programmable laser frequency stabilization at 1523 nm by use of persistent spectral hole burning*. Opt. Lett. **28**, 200–202 (2003).
94. D. K. Cheng. *Field and wave electromagnetics*. Addison-Wesley publishing company (1989).
95. D. Z. Anderson, J. C. Frisch and C. S. Masser. *Mirror reflectometer based on optical cavity decay time*. Applied Optics **23**, 1238–1245 (1984).
96. J.P.D. Martin, M.J. Sellars, P. Tuthill, N.B. Manson, G. Pryde and T. Dyke. *Resolved isotopic energy shift and hole burning in $\text{EuCl}_3\cdot 6\text{H}_2\text{O}$* . J. Lumin. **78**, 19–24 (1998).
97. T. Day, E. K. Gustafson and R. L. Byer. *Sub-hertz relative frequency stabilization of 2-diode laser-pumped Nd:YAG lasers locked to a Fabry-Perot-interferometer*. IEEE Journal of Quantum Electronics **28**, 1106–1117 (1992).
98. wikipedia. Microstrip (2006). URL <http://en.wikipedia.org/wiki/Microstrip>.
99. H. A. Wheeler. *Transmission-line properties of a strip on a dielectric sheet on a plane*. IEEE Trans. Microwave Theory Tech. **25**, 631–647 (1977).
100. B. Carter. How (not) to decouple high-speed operational amplifiers, Application report SLOA069. Technical report Texas instrument (2001). URL <http://focus.ti.com/lit/an/sloa069/sloa069.pdf>.

101. P Brokaw. An IC amplifier user's guide to decoupling, grounding, and making things go right for a change. Technical report Analog devices (2006). URL http://www.analog.com/UploadedFiles/Application_Notes/135208865AN-202.pdf.
102. Peter Sellin. Personal communication (2006).
103. B. Böttger. *Laser frequency stabilization to spectral hole burning frequency references in erbium-doped crystals: material and device optimization*. PhD thesis Montana state university (2002).
104. M. J. Sellars. *Ultra-high resolution laser spectroscopy of rare earth doped solids*. PhD thesis Australian National University (1995).
105. Michael Jefferson. Personal communication (2006).
106. G. E. Francois and F. M. Librecht. *ADP 45 degrees x-cut 4-crystal light modulator*. APPLIED OPTICS **11**, 472–& (1972).
107. M. Gerstanhaber and C. Tran. *Difference amplifier measures high voltages*. EDN **April**, 78–79 (2005). URL www.edn.com/contents/images/42805di.pdf.
108. T. Gale. Diplexers Topics (2006). URL <http://www.qrp.pops.net/dip2.htm>.
109. Nic Hamilton. *Improving direct conversion receiver design*. Radio communications **April**, 39–44 (1991).
110. T. Gale, H. Strickner and W. Hayward. Diplexer Supplemental Page (2006). URL http://www.qrp.pops.net/DIP_SUP.htm.
111. C. Salomon, D. Hils and J. L. Hall. *Laser stabilisation at the millihertz level*. J. Opt. Soc. Am. B **5**, 1576–1587 (1988).
112. J. Steele. *Modify power supply to enhance performance*. Electronic design analog application issue **June**, 57–59 (1994).
113. N. Ohlsson, M. Nilsson, R. K. Mohan and S. Kröll. *Long-time storage mechanism for Tm:YAG in magnetic field*. Opt.Lett. **28**, 450–452 (2003).
114. I. Lorgere, L. Menager, V. Lavielle, J. L. Le Gouet, D. Dolfi, S. Tonda and J. P. Huignard. *Demonstration of a radio-frequency spectrum analyser based on spectral hole burning*. J. Modern Opt. **49**, 2459–2475 (2002).
115. V. Lavielle, I. Lorgere, J. L. Le Gouet, S. Tonda and D. Dolfi. *Wide-band versatile radio-frequency spectrum analyzer*. Opt. Lett. **28**, 384–386 (2003).

SCHEMATICS

P300A	EOM low voltage generator, B300	101
P400A	EOM high voltage generator, B400	102
P500A	Suspend regulation controller, B500	106
P900A	3-stage diplexer notch filter 10.188, 17.310 and 100 MHz, B900	109
P1000A	20MHz low pass and 100 MHz notch diplexer filters with enable, B1000	110
P1100A	2-stage 100 MHz notch diplexer filter, B1100	111
P1400A	Phase detector, B1400	112
P1700A	50 MHz resonant tank, B1700, for phase modulator B1800	113



Box B-400, Farnell 91 65940 Heatsink, Farnell 523185	*NM = optional component, not mounted		Charged by	Date Charged September 28, 2006	Time 11:15
			3	Thursday	4

	1	2
To adjust CMRR:	<p>Apply the same signal to center and casing of X1.</p> <p>Connect 50 Ω between input and ground.</p> <p>Apply signals from DC and up and adjust \odot for minimum output</p>	<p>Short the center and casing of X1.</p> <p>Adjust the gain of the input voltmeter</p> <p>This should be done after proper warm-up.</p>
Offset adjust:	<p>Without R73,R74 mounted, ground \odot and adjust \odot for 0 output</p>	

