University of Sussex

A University of Sussex DPhil thesis

Available online via Sussex Research Online:

http://sro.sussex.ac.uk/

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details



A fibre-based single-photon source

Andrew G. Riley-Watson

Submitted for the degree of Doctor of Philosophy University of Sussex September 2012

Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree.

Signature:

Andrew Riley-Watson

UNIVERSITY OF SUSSEX

ANDREW G. RILEY-WATSON, DOCTOR OF PHILOSOPHY

A FIBRE-BASED SINGLE-PHOTON SOURCE

ABSTRACT

The controlled emission and absorption of single photons is an important enabling technology in the fields of quantum communication, cryptography and computing. We have realised a novel single photon source, based on a miniature 'endcap' ion trap with integrated optical fibres. To minimise distortion of the trapping field the fibres are tightly integrated and recessed within the two hollow cylindrical rf electrodes of the trap. This allows us to bring the fibres to within approximately 300 μ m of the trapped ion. With the fibres in place we are able to collect the ion's fluorescence using no further optics. In this thesis the ion trap is fully characterised, and the quantum nature of the light demonstrated in the results of a variant on the Hanbury Brown-Twiss photon-correlation experiment. The scheme will ultimately be extended to implement a coherent ion-photon interface through strong coupling cavity-QED. Towards this end, an ultra-high-finesse cavity has been designed and fabricated by laser-machining and coating surfaces with a range of radii of curvature on the end facets of the fibres. To improve the stability and precision manipulation of the cavity in-vacuum, an entirely new trap has been designed and built. Finally, the current status of the fibre cavity and the oulook for the experiment are presented.

Acknowledgements

The work of this thesis was carried out in the group of Professor Wolfgang Lange at the University of Sussex between October 2008 and September 2012. Adventures in experimental physics are rarely a solo enterprise, and over the last four years many people— post-docs, students and faculty members—have contributed to the success of our experiment. Perhaps the most rewarding part of the last four years has been meeting and working with so many interesting and talented people, all of whom deserve my thanks.

In particular, I would like to thank my supervisor Matthias Keller for his support and understanding. I will miss the conversations, and the mud, during our 10k lunchtime South Downs runs. Keep up the running—I'll be checking.

Thanks to Alex Wilson for introducing me to ion traps right in the beginning, and for all his hard work building up the beautiful little trap we used so successfully over the last few years. To Hiroki, who wrote most of the LabView programs we used and with whom I worked on the experiment for about three years—thanks for your hard work, your endless patience with my endless questions, and for being such a great ambassador for Japan. Have a fun!

Thanks to Alan Butler in our workshop for always having the time to tap the hole we forgot to make, for his amazing knowledge of vacuum technology, and for all the chats around the milling machine. Thanks also to John Knight, Peter, Bob, and all in the mechanical workshop for their advice, time, and for the many, many miniature parts we made you make for the new trap.

A shadow fell on our group in April of this year with the sudden death of Wolfgang Lange. It's hard to describe what the loss of such a brilliant and kind-hearted leader meant to us all. Our lives, our work, and the field of quantum optics in general, are the poorer for his loss. Wolfgang, *wir vermissen Dich so sehr*.

To all the past and current members of the group, thanks for your help over the years: Fedja, for designing the laser-machining setup we used, and helping me move house; Dan for the badminton, Nic, Kevin, Lizzy, Stephen, Anders and Peter; Michael for taking over from me, Atsushi for his work on the fibre cavity, and everyone else who, in my haste to get to the printers in time, I have omitted.

Thanks to Penny, and to my parents, who would have loved to have seen this.

Above all, I would like to thank Clare, for supporting me in every possible way on this long journey, and my wonderful children, for always being there.

Hamish, Elspeth and Ossian, this thesis is dedicated to you and the children of your generation. One day, one of you will find a *much* better way of doing this stuff.

"Extracting one molecule's signature [in spectral analysis] from the rest of the signatures is hard work, sort of like picking out the sound of your toddler's voice in a roomful of screaming children during playtime. It's hard, but you can do it."

> Neil De Grasse Tyson Death By Black Hole

Contents

List of Tables viii					
Li	st of l	Figures	xi		
1	Intr 1.1	oduction Motivation—Quantum Networking	1 3 4 5 5 6		
2	Rad 2.1 2.2 2.3 2.4 2.5 2.6	iofrequency Ion Traps Earnshaw's theorem The quadrupole potential Equations of motion of an ion in a quadrupole field The Pseudopotential approximation The Trap Depth Micromotion	7 7 8 11 13 14 15		
3	Lase 3.1 3.2	The quantum hardware—Calcium Calcium 3.1.1 The ⁴⁰ Ca ⁺ ion Doppler Cooling 3.2.1 The Density Matrix Formalism Doppler Cooling 3.2.2 The Optical Bloch Equations Doppler Cooling 3.2.3 Three-Level Model - repumping on the 850 nm and 854 nm transitions	17 17 17 18 19 20 23		
4	Cav 4.1 4.2	ity Quantum Electrodynamics Optical cavities 4.1.0.1 Cavity types and the conditions for 'stability' 4.1.1 Finesse Atom-Cavity Interaction	25 25 27 28 30 31 33 34 35		
5	Exp 5.1 5.2	erimental Setup Overview Fluorescence detection and computer control 5.2.1 Fluorescence detection	36 36 38 38		

			5.2.1.1 The PMT
			5.2.1.2 The EMCCD Camera
		5.2.2	Generation of calcium ions
	5.3	The en	dcap ion trap
	5.4	Simula	tions
		5.4.1	The electric field and rf potential
		5.4.2	The pseudo-potential
		5.4.3	The effect of a dc-bias on the pseudo-potential 46
		5.4.4	Secular frequencies
		5.4.5	The separation of the electrodes
		5.4.6	The recess of the outer electrodes
		5.4.7	The effect of the fibres on the pseudo-potential
		5.4.8	The electrode structure 53
			5.4.8.1 The solid angle subtended by the fibres
	5.5	RF dri	ve and electrical connections
		5.5.1	RF drive
			5.5.1.1 The autotransformer
	5.6	The Va	cuum System
	5.7	Magne	tic Field Control
		5.7.1	Experiment Control
			5.7.1.1 Controlling laser power—accousto-optic modulators 64
6	Cha	racteris	ation of the endcap trap 66
	6.1	Trappi	ng ions
	6.2	Micron	notion minimisation
		6.2.1	Coarse micromotion minimisation
		6.2.2	RF-Photon correlation
		6.2.3	Automatic 3D micromotion minimisation
			6.2.3.1 Sensitivities of the techniques
		6.2.4	Sources of stray fields after loading
	6.3	Magne	tic field compensation
	6.4	Secular	r frequency measurements
7	A fit	ore-coup	bled source of single-photons 86
	7.1	Theory	y of the measurements
		7.1.1	Classical theory of intensity fluctuations
			7.1.1.1 Second-Order Coherence
			7.1.1.2 The Hanbury Brown-Twiss Experiment
		7.1.2	Quantum theory of second-order coherence
		7.1.3	Photon Antibunching
			7.1.3.1 The second-order correlation function for a two-level atom 93
	7.2	Results	Continuous Single-Photon Generation
		7.2.1	Spectroscopy with fibre-based detection
		7.2.2	Saturation measurements of beam 1
		7.2.3	The correlation signal path
		7.2.4	Continuous Correlation Measurements
			7.2.4.1 Correlation measurements with a high intensity 100
			7.2.4.2 Correlation measurements with a low intensity
	7.3	Results	—Pulsed Single-Photon Generation
		7.3.1	Outline of the setup
		7.3.2	The sensitivity of the PMTs to 866 nm light
		7.3.3	The pulse sequence and signal path

9	9.1 Conc		145
	0.1 Cono	lusion	1/13
(Conclusion	n and outlook	143
	0.3.2		150
	832	The tran structure	137
C	0.5 A IIEV	The new vacuum design	130
c	02 1 100	8.2.1.4 Conclusion	130
		8.2.1.3 Possible cavity geometries	135
		8.2.1.2 Mode volume and the CQED Coupling parameter	133
		8.2.1.1 Cavity waist and clipping losses	130
	8.2.1	Cavity parameters	130
5	8.2 The f	ibre cavity	129
	8.1.5	Coating the fibres	128
		8.1.4.1 White-light interferometry	122
	8.1.4	Analysing the shot	. 121
	8.1.3	Shooting the fibres	119
	8.1.2	The fibres—etching and cleaving	116
	8.1.1	The CO_2 laser setup $\ldots \ldots \ldots$	114
8	8.1 Laser	machining the fibre facets	114
]	The Fibre	Fabry-Pérot cavity	114
	1.3.1		112
	7.3.0	Single Photon Efficiency	112
	7.3.3	Single Diston Pulse Shape	110
	7.5.4	Correlation Massurements	109
	724	7.3.3.1 The timing of the shelving pulse	109
		7.3.4	7.3.3.1 The timing of the shelving pulse

List of Tables

3.1	Table of transitions in ${}^{40}Ca^+$	18
5.1	Degrees of vacuum	60

List of Figures

1.1	An NP problem—prime factorisation	•	•	•	•	•	•	3
2.1 2.2	The hyperbolic electrodes of an ideal Paul trap	•			•		•	9 10
2.3	The 'saddle' potential							11
2.4	The stability diagram for an ideal Paul trap			•				13
2.5	The secular motion and the micromotion	•	•	•	•	•	•	16
3.1	Energy levels in the ${}^{40}Ca^+$ ion							18
3.2	A simplified three-level model of the ${}^{40}Ca^+$ ion							23
3.3	Dark resonance in the three-level model	•		•	•		•	24
4.1	Schematic drawing of a Fabry-Pérot cavity							26
4.2	The cavity stability diagram							28
4.3	Cavity transmission as a ratio of input intensity							30
4.4	An atom in a cavity							31
4.5	The 'dressed' atom-cavity state	•		•			•	34
5.1	Top view of the main optical setup							37
5.2	Alignment of the PMT							39
5.3	Energy levels in the 40 Ca atom							40
5.4	The Endcap Trap							42
5.5	Simulations: The electric field							44
5.6	Simulations: The instantaneous electric potential							44
5.7	Simulations: The pseudo-potential							45
5.8	Simulations: Radial and axial pseudo-potentials							45
5.9	Simulations: dc potentials							47
5.10) Simulations: Combined rf and dc potentials							47
5.11	1 Trap depths with a dc-bias							48
5.12	2 Extracting the secular frequency from the pseudo-potential	•						49
5.13	3 Simulations: The effect of the electrode separation							50
5.14	4 Simulations: The deviation from the parabolic potential							51
5.15	5 Simulations: The recess of the outer electrode	•		•			•	52
5.16	5 Simulations: The effect of the fibre recess			•			•	53
5.17	7 Cutaway view of the electrode structure			•			•	54
5.18	8 Simplified drawing of the trap mount			•			•	55
5.19	The solid angle subtended by the fibres at the ion						•	56
5.20) A series LCR circuit	•		•			•	57
5.21	1 A drawing of the autotransformer	•		•	•		•	58
5.22	2 A comparison between simulated and real LCR resonances	•		•	•		•	58
5.23	3 Electrical connections to the trap	•		•	•		•	60
5.24	4 The vacuum chamber							61

5.25	The chamber and vacuum flange assembly	52
5.26	A drawing of the trap surrounded by the magnetic field coils	64
5.27	Double-pass AOM setup	5
6.1	A single ion in the trap	8
6.2	Quantum jumps on the $D_{\frac{3}{2}} \leftrightarrow P_{\frac{3}{2}}$ transition	8
6.3	Spectroscopy using the cooling beams	9
6.4	Coarse micromotion compensation	1
6.5	The rf-photon correlation signal path	2
6.6	The phase flip in micromotion rf-photon correlations	5
6.7	Micromotion amplitudes	6
6.8	Planes in compensation-voltage space	17
6.9	The variation in the three micromotion compensation voltages over time 7	9
6.10	The calculated positions of origin of the stray charge	60
6.11	Optical pumping to dark states in Calcium	31
6.12	Magnetic field compensation	32
6.13	Scanning over the radial and axial secular frequencies	34
6.14	Secular frequencies: comparison with simulation	5
011 1		
7.1	The Hanbury Brown-Twiss experiment)1
7.2	The Hanbury Brown-Twiss experiment with photons	2
7.3	Bunched and antibunched photon streams)3
7.4	The $g^{(2)}(\tau)$ function calculated analytically	95
7.5	The $g^{(2)}(\tau)$ function by integration of the optical Bloch equations	96
7.6	The ultra-low scatter setup of 'beam 1'	8
7.7	Spectroscopy of a single ion using beam 1	8
7.8	Saturation power measurement using beam 1	9
7.9	The setup for correlation measurements)0
7.10	Correlation measurement with a high intensity)1
7.11	Correlation measurements at low intensity and detuning	13
7.12	The setup for the pulsed correlation measurements)4
7 13	Transmission from the bottom to the top fibre)5
7 14	PMT sensitivity to 866 nm light with Semrock filters)6
7 1 5	I aser pulse sequence for pulsed single photons	17
7.15	The pulsed single photon signal path	,, 18
7.10	The timing of the shelving and re-cooling pulses	10 10
7.19	Shalving afficiency massurements	10
7.10	Pulsed correlation measurements with hear 1	11
7.19	Single photon pulse shape	י. ר
7.20		2
8.1	The CO ₂ laser setup \ldots	5
8.2	The partially and fully etched fibre sections	7
8.3	Poorly-cleaved fibre facets	9
84	The laser-shutter pulse sequence 12	20
8 5	Shoting a fibre at various rotations	21
8.6	The white-light interferometer 12	2
87	Interferometric images of a fibre facet 12	2
8.8	Reconstruction of the fibre shape 12	6
8 Q	Interferometric image of a flat fibre	ט. דו
0.9 8 10	The distribution of radii of curvature	. 1 1 Q
0.1U Q 11	Drawings of the fibre transport assembly	0, 0
0.11	Drawings of the fibre transport assembly	ר, הו
0.12	r notographs of the note transport assentiony 12	9

8.13	Cavity waist—symmetric cavity 132
8.14	Cavity waist—half-symmetric cavities
8.15	Clipping losses and finesse
8.16	Mode volume in symmetric and half-symmetric cavities
8.17	Projected g-parameter in symmetric and half-symmetric cavities
8.18	Cavity decay rate
8.19	Some possible fibre-cavity geometries
8.20	The new bottom flange assembly
8.21	The new chamber
8.22	Vibration-isolation blocks
8.23	The 'cage'
8.24	The top and bottom electrode assemblies
8.25	Cross-section through the cage
8.26	Fibre strain-relief devices
8.27	The new trap

Chapter 1

Introduction

"...in the now current interpretation, ample information is forthcoming about a host of experiments that nobody has ever been or ever will be able to perform."

E. Schrödinger, British Journal of the Philosophy of Sciences, 3, 1952.

1.1 Motivation—Quantum Networking

In 1952, Erwin Schrödinger argued that we would never be able to perform experiments on single electrons or atoms[1]. Today, in the atomic physics laboratories of the world, these experiments have become routine. Although this is no longer news, the confinement of single atomic particles is now a prerequisite in several important fields of research, and trapped single ions are considered promising contenders for the hardware of a future quantum computer. The work presented here is motivated by the problem of efficient communication between these quantum computers by means of single photons. Although it is in spirit about communication, rather than computing, a brief introduction to the principles of quantum computing will be given in this section, to give an understanding of the problem this work is aimed at solving.

Why a quantum computer?

Moore's Law¹ states broadly that the number of transistors that can be placed on a chip will double approximately every 2 years. Although there is a whole field of study around how long into the future Moore's Law can be sustained, with opinions in recent years ranging from 10–20 years [2], to forever [3], we intuitively see that we will run into problems as transistors approach the size of the de Broglie wavelength of the electrons carrying charge through them—for example, tunnelling of electrons through potential barriers. If we want Moore's Law to hold far into the future it seems clear that some new processor architecture is required (that said, Weber et al. at the University of New South Wales have just reported observation of constant resistivity in phosphorous wires one atom high and four atoms wide embedded in a silicon crystal, suggesting that Ohm's law may hold down to even single atoms, given the right conditions [4]).

¹After Gordon Moore, co-founder of Intel Corporation, 1965.

Richard Feynman first discussed the idea of simulating quantum systems with computers in 1981 [5], noting that it does not seem possible to simulate quantum systems with a classical computer. David Deutsch's description of the 'Universal Quantum Computer' in 1985 [6] showed that a quantum computer might be able to efficiently solve problems which a classical computer could not. In 1986 in an effort to investigate whether, as we shrink transistors down to microscopic size, quantum mechanics places any limit on the amount of energy required per calculation 'step' in a computation, Feynman introduced the idea of a 'Quantum Mechanical Computer' [7].

Classes of computing problems

We classify computing problems according to the complexity of the resources used in obtaining their solutions. For example, the P *complexity class* contains the set of problems solvable by a deterministic Turing machine (such as a classical computer) in *polynomial time*. This means that the time taken to solve the problem scales with the input as a polynomial. By contrast, the NP complexity class contains the set of problems solvable by a *non*-deterministic Turing machine in polynomial time. This class of problem is extremely difficult for a classical computer to solve. There is a specific class of problems known as BQP—bounded error quantum polynomial time—which is solvable by a quantum computer, with high probability (within some error bounds), in polynomial time, and this is suspected to include some, but not all, NP problems [8].

The classic example of an NP/BQP problem is the prime factorisation of large numbers. That is, factoring a number into its constituent primes. The number of operations required to do this increases exponentially with the number, N, and this is dramatically demonstrated in figure 1.1².

In 1994, Peter Shor proposed a quantum algorithm to solve this problem [9] in polynomial time, and in 2001 the group of Isaac Cheung at IBM's Almaden Research Centre in California experimentally implemented Shor's algorithm using 7 quantum bits (qubits) to factorise the number 15 [10].

²The spikes in the data are due to the varying demands of the computer's CPU at the time the function was called. This graph was produced from a very simple prime-factoring algorithm—there are many more efficient algorithms available.



Figure 1.1: An NP problem—the time taken to factorise the numbers from 1 to 10000 into their constituent primes, in terms of Matlab function call time, using a simple prime-factoring algorithm. As a point of interest, the total time taken in terms of function calls, that is, the area under this graph, was around 12 hours. We can calculate from the fitted function that at this rate it would have taken 89 s to factorise 20,000, 527 s to factorise 30,000, 4 hours to factorise 50,000, 3 years to factorise 100,000, and 6.6×10^{68} years to factorise 1,000,000. The RSA public-key encryption algorithm is based on exactly this problem.

1.1.1 Qubits

Bits in a classical computer are represented by transistors. They are either on (1), or off (0). The processor performs millions of logical operations on one or two bits at a time. The qubits in a quantum computer can in principle be any quantum system with two distinguishable and accessible states, corresponding to the on and off of the classical transistor. Let us use Dirac notation from here on and call these two states $|a\rangle$ and $|b\rangle$. The crucial difference between quantum and classical bits is that quantum bits can exist in a superposition of the $|a\rangle$ and $|b\rangle$ states:

$$|\Psi\rangle = c_a |a\rangle + c_b |b\rangle, \tag{1.1}$$

where $|c_a|^2$ and $|c_b|^2$ represent the probability of observing state $|a\rangle$ or $|b\rangle$ when a measurement is performed. In such a superposition state then, the *information* carried by the qubit exists in the amplitudes c_a and c_b —one qubit can carry two pieces of information.

If we add one more qubit to our system, then the system can exist in a superposition of four states:

$$|\Psi\rangle = c_{aa}|aa\rangle + c_{ab}|ab\rangle + c_{ba}|ba\rangle + c_{bb}|bb\rangle.$$
(1.2)

We can see that N qubits are capable of storing 2^N pieces of information in the amplitudes of their states, and any operation performed on the qubit is simultaneously performed on all 2^N amplitudes.

There have been many proposals for qubit candidates, including (this is by no means an

exhaustive list):

- Trapped ions in an ion trap [11], where the qubits are two internal energy levels of ions.
- NMR, where nuclear spin holds the information [12].
- Neutral atoms, for example [13].
- The energy states of quantum dots [14].
- The spin states of donor impurities in, for example Silicon [15].
- Nitrogen-vacancies in diamond and similar defects in other materials, such as silicon carbide [16].

1.1.2 Why a Trapped Ion Quantum Computer?

The DiVincenzo criteria

In 1998, in a work inspired by Cirac and Zoller's 1994 proposal for a quantum computer using cold ions in a trap [11] David DiVincenzo, at IBM's Research Division, set out five criteria which he considered essential to a physical implementation of a quantum computer [17]. They are:

- 1. The quantum states of the qubits must be precisely enumerated, that is, we must know the state well, and know how to confine the qubit to that state. Also, the system must be scalable—we must be able to add qubits to the system³.
- 2. State preparation: we must be able to set the system into some simple state, for example $|a\rangle$.
- 3. Low decoherence: the environment should not perturb the quantum interference (coherence) of the system. DiVincenzo quotes a decoherence time of 1×10^6 times longer than the gate operation time.
- 4. Controlled unitary transformations: it must be possible to perform gate operations on the system.
- 5. Measurement: it must be possible to perform a measurement on individual qubits.

Vincenzo noted at the time that the Cirac-Zoller proposal for trapped ions appeared to satisfy all five of these criteria, for example:

- 1. Ions can be routinely cooled to their ground state using laser cooling.
- 2. The internal states of trapped ions are very well known and understood. We could imagine scaling such a device by simply adding more ions to the trap (although in practice this is not the case, because the ions are all coupled through their motion. Also, as the ion string grows it becomes harder to laser-cool the whole string to the ground state).

³Factorising a 200-digit number requires about 3,500 qubits [18] and in a theoretical work Hughes et al. found that an ion trap quantum computer would be decoherence-limited at around ten or so ions [19].

- 3. Long internal state coherence times.
- 4. Internal states are easily manipulated and systems for quantum logic operations are available.
- 5. Measurement can be carried out with high efficiency.

1.1.3 Quantum Networking

There are two current approaches to the problem of scaling mentioned in section 1.1.2. One solution is to build arrays of traps [20], in which ions are shuttled between different trapping regions, with each region only containing a small amount of ions. The group of David Wineland at NIST recently implemented such an array, trapping two ${}^{9}\text{Be}^{+}$ and two ${}^{24}\text{Mg}^{+}$ ions in a segmented trap [21], and ions have been successfully shuttled around 90° corners in a 'T-Junction' ion trap array [22]. However, our long-term interest lies more in an alternative idea; that of providing a communications channel between many remote ion trap devices. This has applications beyond the idea of scaling an ion trap quantum computer.

In 2000, in the prettily titled and very readable 'Desiderata for Quantum Communication' [23] DiVincenzo added the following two criteria to the list in section 1.1.2:

- 6. The ability to interconvert stationary and flying qubits.⁴
- 7. The ability faithfully to transmit flying qubits between specified locations.

The work presented here is based around a scheme designed to fulfill these final two criteria. The flying qubit in our case is a photon and the stationary qubit is a trapped ion.

1.1.3.1 Single-photon sources

If we restrict ourselves to the paradigm of the flying-qubit-as-photon, produced by a single quantum emitter (and the justification for such a restriction is that we require only one photon to be emitted for each 'request', so it is reasonable to use a single quantum emitter to produce it), then a prerequisite is the development of a single-photon source which fulfills the final two criteria.

The generation of single-photons has been an active research area in recent years and several different quantum emmitters have been tried. A review of many such schemes is given in [24]. The interest in these schemes lies in their utility as *enabling technologies* for several applications, aside from scaling quantum computers, some with extensive commercial relevance, such as:

• Key distribution/cryptography: the ability to transmit data between two points without any possibility of it being intercepted [25].⁵

⁴The phrase 'flying qubit' has become standard terminology in the field. In most schemes the qubit is encoded in a photon.

⁵Quantum cryptopgraphy is probably the most immediately obvious beneficiary of quantum networking technology developments. The link to quantum information transmission being that quantum cryptography schemes depend upon the reliable generation and transmission of a flying qubit (usually a photon—see, for example the BB84 protocol [26]). Ones and zeroes can be encoded in the polarisation of the photon in some basis; for example the vertical and horizontal polarisation for zero and one respectively. The basis can be chosen at random, for example a $\pm 45^{\circ}$ linear polarisation can encode a one/zero. In this way a 'key' can be transmitted between two parties and used to decrypt a subsequent transmission. Without knowing the basis, it would be impossible for an eavesdropper to know the key. For a recent review of this field, see [27].

- Authentication: the ability to transmit data between two points in a way that ensures it was not corrupted in passage [25, 27].
- A 'quantum repeater': to transfer the state of one ion in one trap to a remote ion in a distant trap—effectively to lengthen the distance over which quantum information can be transmitted [28].
- The distribution of 'quantum software' [29].
- An absolute standard of optical brightness—if the rate of single photon production, r, and their energy, is known, then the optical power is $\hbar\omega r$.
- 'Blind' quantum computing. In a paper published at the time of writing, Barz et al. [30] report a demonstration of a kind of quantum cloud computing, whereby a user only has to prepare photonic qubits and transmit them to a remote quantum computer where computations are performed.

We now move on to describe in detail the background theory necessary to understand the components of our implementation, that is:

- Radiofrequency ion traps
- Laser cooling
- Cavity Quantum Electrodynamics

1.1.4 This Thesis

In this thesis I will describe the realisation of a single-photon source based on a single calcium ion in a novel miniature ion trap. We take as our model the CQED-based scheme of Cirac and Zoller [31] for a quantum network composed of individual ion trap 'nodes' connected by optical fibres. The ultimate goal of our work is to produce a single-photon source suitable for use in such a scheme, that is, a single ion strongly coupled to a single cavity mode, as described in chapter 4. The work presented here consists of firstly the design and implementation of a miniature ion trap, the detection of an ion's fluorescence through optical fibres placed close to the ion, pulsed and continuous non-cavity single-photon generation, and finally the design and fabrication of a miniature optical cavity by laser-machining and coating the end facets of the fibres.

After covering the basic theory behind the trapping and laser-cooling of charged particles in chapters 2 and 3, and the principles of cavity-QED in chapter 4, I will describe the technical implementation of the experiment including materials and devices used (chapter 5). Chapter 6 presents a full characterisation of the ion trap, including comparison with simulations. Chapter 7 presents the main results of the experiment so far, that is, production and detection of single photons.

In Chapter 8 I describe the techniques and results of work done to laser-machine concave surfaces onto the $\sim 200 \ \mu m$ diameter end facets of optical fibres using a high-power CO₂ laser. Some predicted CQED parameters of the fibre cavity are calculated.

Chapter 9 presents a conclusion and the outlook for the experiment, including some details of the design and build of an entirely new ion trap to accommodate the new fibre cavity.

Chapter 2

Radiofrequency Ion Traps

"I think it is a sad situation in all our chemistry that we are unable to suspend the constituents of matter free."

G.C. Lichtenberg, quoted in Wolfgang Paul's Nobel Lecture, December 8, 1989.

In 1989, Wolfgang Paul, Hans G. Dehmelt and Norman F. Ramsey were awarded the Nobel Prize in Physics 'for contributions of importance for the development of atomic precision spectroscopy.' Paul and Dehmelt's contribution was the ion trap, that is, the 'suspension' of individual particles of matter for long periods of time, allowing access to unprecedented regimes of precision¹ in atomic spectroscopy measurements [32]. The first *single* ion was trapped in 1980 [33]. Ion traps are now commonly used in many fields, most commercially mass spectrometry—for a review of the history of this application of ion traps see [34]. Ion traps also represent the world's most precise frequency (and therefore, time) standards. In 2010, NIST² reported the worlds most precise clock, based on a single trapped aluminium ion [35, 36]. The most common types of ion trap today are the *Penning* trap [37], and the *radiofrequency* (rf), or Paul, trap. In this thesis we will concern ourselves only with rf ion traps, and this section aims to describe the essential theory behind their operation, with any necessary modifications relevant to our particular designs.

2.1 Earnshaw's theorem

To perfectly trap a single ion in space with an electric field alone we might reasonably think that the field lines must all point inwards, that is towards the equilibrium position of the particle. But there is a problem with this.

Consider the divergence of a vector field:

$$\nabla \cdot \mathbf{F} \equiv \lim_{V \to 0} \frac{\oint_S \mathbf{F} \cdot d\mathbf{a}}{V}.$$
(2.1)

¹Heisenberg's energy/time uncertainty relation tells us that over long periods of time we can make increasingly precise measurements of energy.

²America's National Institute for Standards and Technology.

That is, the divergence of the vector force field surrounding a point is the flux of the field through a closed surface surrounding a volume around that point as we shrink that volume to zero. In other words, it is the rate of creation or removal of electric flux at that point. If we want our electric field lines to terminate at the point where the particle is at equilibrium, such that if the particles moves away, it is pulled back, then that point must be a *sink* for the field—the divergence must be negative. Since we know from Maxwell's equations that the divergence of an electric field is zero in free space, and the curl of an electric field is zero, then we see that this is impossible.³ It turns out in fact that there is no way to trap a charged particle with electrostatic forces alone, and this is known as *Earnshaw's theorem*. We can, however, trap a charged particle with electrodynamic forces—an oscillating electric field, described by Wuerker et al. in 1958 as 'a closed form of W. Paul's and M. Raether's electric mass filter.' [39]

2.2 The quadrupole potential

If we want to confine a charged particle at the bottom of a potential well, from which it does not have sufficient energy to escape, then a simple solution would be to apply a potential in three dimensions which has quadratic dependence on the particle's coordinates x, y, and z:

$$\phi = \phi_0(\alpha x^2 + \beta y^2 + \gamma z^2), \qquad (2.2)$$

where ϕ_0 is the potential at the origin.

For this potential to obey Laplace's equation, we require

$$\alpha + \beta + \gamma = 0. \tag{2.3}$$

What we see from this is that while there could be a positive, confining potential in two of the three directions, there must then be a negative 'anti-trapping' potential in the third—we must always have a 'saddle' potential. This is illustrated in figure 2.3.

In a three-dimensional field, $\alpha = 1$, $\beta = 1$ and $\gamma = -2$ satisfy this constraint and the potential can be described by [41]

$$\phi = \phi_0 \left(\frac{r^2 - 2z^2}{2r_0^2} \right), \tag{2.4}$$

where *r* represents the vector $x^2 + y^2$ and r_0 is a factor depending on the geometry of the trap (indicated on figure 2.1. Figure 2.2(a) shows a cross-section of the equipotentials of such a quadrupole potential and the ideal electrode structure which achieves this along all directions is shown in figure 2.1. This consists of two hyperboloids of revolution about the z-axis (the 'endcap' electrodes) and a ring electrode with a hyperbolic cross-section, with the condition $r_0^2 = 2z_0^2$. This condition arises from the boundary conditions at the electrodes. We see that at z_0 the potential is $\phi_0(2z_0^2)$ and at r_0 it is $\phi_o(-r_0^2)$. Paul traps are often operated using a voltage of opposite sign on the endcaps and ring; in other words $\phi(r_0) = -\phi(z_0)$, and so we have $r_0^2 = 2z_0^2$. This condition is not necessary for stable trapping, and we do not adhere to it strictly in this work—see chapter 5.

³We often invoke Laplace's equation to show this: $\nabla \cdot \mathbf{E} = \nabla \cdot (-\nabla \phi) = -\nabla^2 \phi$. See, for example, [38].



Figure 2.1: An ideal Paul trap consisting of three electrodes. The top and bottom 'endcap' electrodes are infinite hyperbolas of revolution about the z-axis, shown in cross section. The East/West electrodes are hyperbolic cross-sections through a hypothetically infinite ring electrode. This geometry produces in principle a perfectly quadratic potential in all directions. A typical rf voltage arrangement is shown, whereby an oscillating voltage is applied to the endcaps while the ring is held at rf ground. The hyperbolic cross-sections are described by $\frac{r^2}{r_0^2} - \frac{z^2}{b_1^2} = 1$ for the ring and $-\frac{r^2}{z_0^2} + \frac{z^2}{b_2^2} = 1$ for the endcaps. In this plot $r_0^2 = 2z_0^2$, as described in the text, and also the two hyperbolas have the same asymptote, determined by the *b*-parameters, such that $b_2^2 = \frac{r_0^2 z_0^2}{b^2}$. See [40]. Dimensions are arbitrary (although the r_0 used is the real r_0 in our trap).

However, for a general analysis of ion trapping theory we will use this approximation.

Over the years, Paul traps have moved further and further away from the ideal hyperbolic electrode structure. Some alternative trap structures are mentioned in chapter 5, and the geometry of a trap can be varied greatly depending on the intended purpose. For example, in our work we are concerned with trapping single, laser-cooled ions. For this purpose we only require that the quadrupole nature of the potential be maintained close to the trap centre, as we know that using laser-cooling we can localise our ion there. Hyperbolic electrodes, therefore, are not necessary, and we can design our electrode geometry in a way that gives, for example, better laser access, or that can accommodate additional experimental apparatus. The type of trap used in our experiment consists of two cylindrical electrodes, and two concentric 'outer' electrodes, recessed some distance from the inner electrodes. This is known as an endcap trap, and is described in detail in chapter 5. The approximately quadrupole field generated by applying a voltage to one pair of electrodes (it doesn't matter which pair) can be determined by finite element analysis and the equipotentials

9

of the field are shown for an arbitrary voltage in cross-section in figure 2.2(b). We can see in this figure the potential starting to deviate from ideal quadrupole behaviour as we move far from the centre. The potential distribution in the endcap trap is discussed in far more detail in chapter 6.

In our trap an approximately identical rf voltage of the form

$$\phi_0 = U_{\rm dc} + V_0 \cos \Omega_{\rm rf} t \tag{2.5}$$

is applied to the top and bottom endcap electrodes, and the 'ring' electrode is grounded. U_{dc} is the dc offset of the rf voltage. We have not used a dc offset on our rf electrodes, but we leave the parameter in for the purpose of this general analysis. V_0 is the 0–peak amplitude of the rf voltage and Ω_{rf} is its angular frequency.



Figure 2.2: Figure (a) shows the ideal quadrupole potential in the Paul trap of figure 2.1, determined analytically, while figure (b) shows, for comparison, the approximate quadrupole potential near the centre of the endcap trap used in these experiments, calculated by finite element analysis. We see that, provided we confine our ion to the centre of the trap, we can assume a quadrupole potential. Note that the hyperbolic electrodes in figure (a) are not visible because the z_0 parameter of the ideal case differs from our z_0 .



Figure 2.3: The ideal 'saddle' potential surface at two points in a half-cycle of the rf showing how a charged particle is (a) trapped in the radial direction, and (b) trapped in the axial direction.

2.3 Equations of motion of an ion in a quadrupole field

The electric field is defined by

$$\mathbf{E} = -\boldsymbol{\nabla}\phi. \tag{2.6}$$

So we can construct an equation of motion for an ion in the radial plane, that is, the **r**-direction of equation 2.4, and using the definition of the potential given in equation 2.5:

$$\mathbf{F} = m\mathbf{a}$$

$$\Rightarrow -e\frac{\partial}{\partial \mathbf{r}}\phi_{\mathbf{r}} = m\ddot{\mathbf{r}}$$

$$\Rightarrow \ddot{\mathbf{r}} + \frac{e}{m}\frac{\partial}{\partial \mathbf{r}}\phi_{r} = 0$$

$$\Rightarrow \ddot{\mathbf{r}} + \frac{e}{mr_{0}^{2}}\left(U_{dc} + V_{0}\cos\Omega t\right)r = 0.$$
(2.7)

And similarly for the axial plane:

$$\ddot{z} - \frac{2e}{mr_0^2} \left(U_{\rm dc} + V_0 \cos \Omega t \right) z = 0.$$
(2.8)

If we make the following substitutions in these equations:

$$a_z = -2a_r = \frac{8eU_{\rm dc}}{mr_0^2\Omega^2},\tag{2.9}$$

$$-q_z = 2q_r = \frac{4eV_0}{mr_0^2 \Omega^2},$$
 (2.10)

11

$$\zeta = \frac{\Omega t}{2},\tag{2.11}$$

then we can recast the equations of motion into Mathieu equations:

$$\frac{d^2r}{d\zeta^2} + (a_r - 2q\cos 2\zeta)r = 0, \qquad (2.12)$$

$$\frac{d^2z}{d\zeta^2} + (a_z - 2q_z \cos 2\zeta) z = 0.$$
(2.13)

Stable solutions, that is, solutions for which r, z is finite as $\zeta \to \infty^4$, are given by [41]

$$r, z(\zeta) = A \sum_{n = -\infty}^{n = +\infty} C_{2n} \cos(2n \pm \beta) \zeta + B \sum_{n = -\infty}^{n = +\infty} C_{2n} \sin(2n \pm \beta) \zeta, \qquad (2.14)$$

where the C_{2n} represents an amplitude of motion and $(2n \pm \beta \zeta)$ represents a frequency. Equating this latter with ωt and substituting equation 2.11, we have

$$\omega_z = \frac{(2n \pm \beta_z)\Omega}{2},\tag{2.15}$$

and

$$\omega_r = \frac{(2n \pm \beta_r)\Omega}{2} \tag{2.16}$$

where Ω is as usual the angular frequency of the rf drive. For the fundamental frequency of oscillation, n = 0 and $\omega_{r,z} = \beta \Omega/2$. Higher harmonics are clearly allowed by the equations, and indeed these sometimes show up in measurements of a trapped ion's oscillatory spectrum.

We can see that at $\beta = 0$, the first non-zero frequency term is Ω , and for $\beta = 1$ the first term is $\Omega/2$. We can now define a stability region in *a* and *q* space. For each value of β in the stability region, then, there is an associated fundamental frequency of oscillation.

This frequency is termed the **secular frequency**. The secular frequency is in contrast to a second, high-frequency oscillation driven directly by, and having a similar frequency to, the trap rf. This driven motion is termed the **micromotion**.

Solutions for $\beta = 0$ and $\beta = 1$ in terms of the *a* and *q* parameters are [39, 42]

$$a = -\frac{1}{2}q^2 + \frac{1}{128}q^4 - \frac{29}{2304}q^6 + \frac{68687}{18874368}q^8,$$
(2.17)

and

$$a = 1 - q - \frac{1}{8}q^2 + \frac{1}{64}q^3 - \frac{1}{1536}q^4 - \frac{11}{35864}q^5 \dots$$
(2.18)

Plotting a against q for the axial and radial directions gives us figure 2.4. The shaded area represents stable trapping values of a and a for an ideal Paul trap.

⁴Stable here means that the amplitude of the motion of the ion is restricted for all times.



Figure 2.4: The first stable trapping region for an axially symmetric ideal Paul trap. This figure is a plot of equations 2.17 and 2.18. That is, $a_z(q_z)$ for $\beta_z = 0, 1$, then $a_r(q_r) = -2a_z(q_z/2)$ for $\beta_r = 0, 1$, where equations 2.12 and 2.13 provide the relation between a_z, a_r, q_z , and q_r .

2.4 The Pseudopotential approximation

In this section we show how switching the potential on the electrodes of an ideal Paul trap generates an effective potential well in both the radial and axial directions at the same time. This is known as the **pseudopotential**.

Consider the equation of motion of the ion in the z-direction given by equation 2.13. The amplitude of the motion at a given time can written as a sum of the secular and micromotion amplitudes, Z and δ :

$$z = Z + \delta, \tag{2.19}$$

Now we make two assumptions; firstly that the micromotion amplitude δ is much smaller than the secular amplitude Z, and secondly that the micromotion amplitude varies much faster than the secular. With these assumptions, equation 2.13 can clearly be written as

$$\frac{\mathrm{d}^2\delta}{\mathrm{d}\zeta^2} + (a_z - 2q_z\cos 2\zeta)Z = 0. \tag{2.20}$$

We further assume that $a_z \ll q_z$ and that the secular amplitude is constant over one cycle of the rf oscillation. Now we can integrate equation 2.20 to give

$$\delta = -\frac{q_z Z}{2} \cos 2\zeta, \qquad (2.21)$$

which we can rewrite as

$$z = Z - \frac{q_z Z}{2} \cos 2\zeta. \tag{2.22}$$

Substituting this into equation 2.13 gives

$$\frac{d^2z}{d\zeta^2} = -a_z Z + \frac{a_z q_z Z}{2} \cos 2\zeta + 2q_z Z \cos 2\zeta - q_z^2 Z \cos 2\zeta.$$
(2.23)

We now consider the time average of the two motions over one cycle of the rf. Clearly the micromotion averages to zero. The secular motion averages to

$$\frac{1}{\pi} \int_{0}^{\pi} \frac{\mathrm{d}^2 z}{\mathrm{d}\zeta^2} \,\mathrm{d}\zeta. \tag{2.24}$$

Now substituting $d^2z/d\zeta^2$ from equation 2.23 we have

$$\frac{d^2 Z}{d\zeta^2} = -(a_z + \frac{q_z^2}{2})Z,$$
(2.25)

and using $\zeta = \Omega t/2$ we have

$$\frac{\mathrm{d}^2 Z}{\mathrm{d}t^2} = -\left(a_z + \frac{q_z^2}{2}\right)\frac{\Omega^2}{4}Z.$$
(2.26)

We can write this in another form

$$\frac{\mathrm{d}^2 Z}{\mathrm{d}t^2} = -\omega_z^2 Z,\tag{2.27}$$

which describes simple harmonic motion at angular frequency ω . This enables us to identify ω_z as

$$\omega_z = -\left(a_z + \frac{q_z}{2}\right)^{\frac{1}{2}} \frac{\Omega}{2}.$$
 (2.28)

If we recall that $\omega_0 = \beta \Omega/2$, we can write β as

$$\beta^2 = \left(a_z + \frac{q_z^2}{2}\right). \tag{2.29}$$

In an ideal trap, of the form $r_0^2 = 2z_0^2$ the radial secular frequency is then given by $\omega_z = 2\omega_r$. Substituting q_z in equation 2.28 and assuming that $a_z = 0$, that is, no DC bias, we have

$$\omega_z = \frac{2eV}{\sqrt{2}mr_0^2\Omega}.$$
(2.30)

2.5 The Trap Depth

We term the depth of the pseudopotential well the *trap depth*. The radial or axial depth of the well here means the depth at the centre of the trap relative to that at the top of the well.

Recall the definition of q_z from equation 2.10. Letting the DC bias be zero, we have from equation 2.26:

$$\frac{\mathrm{d}^2 Z}{\mathrm{d}t^2} = -\frac{q_z^2 \Omega^2}{8} Z \tag{2.31}$$

$$= -\frac{2e^2V^2}{m^2r_0^4\Omega^2}Z.$$
 (2.32)

If we now write this as

$$m\frac{\mathrm{d}^2 Z}{\mathrm{d}t^2} = -e\boldsymbol{\nabla}\Phi,\tag{2.33}$$

we can interpret the left side as the force produced by an electric field resulting from a *pseudo* potential Φ . Let us denote the depth of this potential between z = 0 and $z = z_0$ as Φ_z , so that $\nabla \Phi = \frac{d\Phi_z}{dZ}$. Then we have

$$\frac{\mathrm{d}\Phi_z}{\mathrm{d}Z} = \frac{2eV^2Z}{mr_0^4\Omega^2} = \boldsymbol{\nabla}\Phi.$$
(2.34)

Integrating this we have

$$D_z = \int_{z=0}^{z=z_0} \frac{\mathrm{d}\Phi_z}{\mathrm{d}Z} \,\mathrm{d}Z \tag{2.35}$$

$$=\frac{eV^2 z_0^2}{mr_0^4 \Omega^2}.$$
 (2.36)

2.6 Micromotion

At the trap centre, the electric field resulting from the rf drive is exactly zero. Equation 2.15 tells us that the ion oscillates with some amplitude around the trap centre, and so experiences a force resulting from the trap rf drive. This faster oscillation at the trap frequency is termed *intrinsic* micromotion, since it is characteristic of the system, and can be reduced by restricting the ion's secular excursions from the trap centre (by laser cooling—see chapter 3). There is another form of unwelcome micromotion, resulting from stray charges which can build up on dielectrics inside the trap structure. These can push the ion off the trap centre and so generate what is termed *excess* micromotion. We can reduce this by keeping dielectric surfaces well away from the ion, and again by pushing/pulling the ion to the trap centre with dc voltages. The methods we use for compensating for this excess micromotion will be described later.

We can describe the total amplitude of the ion's motion in terms of the secular motion and micromotion. Recall equation 2.21 where we describe the micromotion amplitude, given certain assumptions, and equation 2.27 which shows that the secular motion can be written in simple harmonic terms, $Z = A \sin(\omega t)$, with A the amplitude of the motion. We can then write the total amplitude z as

$$z = Z + \delta$$

$$= A \sin(\omega_z t) - \frac{q_z Z}{2} \cos \Omega t$$
(2.37)

$$=A\sin(\omega_z t)\left[1-\frac{q_z}{2}\cos\Omega t\right].$$
(2.38)

with *A* the amplitude of the secular motion. A plot of this motion for some typical parameters in our experiment is shown in figure 2.5.



Figure 2.5: The fast micromotion oscillation superimposed on the slower secular frequency, as described by equation 2.37. Some parameters from our experiment were used: $\Omega = 2\pi \times 15$ MHz, $\omega_r = 2\pi \times 2$ MHz, and q_z was chosen to be 0.4. The amplitude is arbitrary. A = 1 in this case. Of note is the fact that the micromotion amplitude increases as the ion moves away from the trap centre.

Chapter 3

Laser cooling

In 1975, Hänsch and Schawlow at Stanford University wrote: 'We wish to point out that if the laser radiation is essentially isotropic, but confined to frequencies on the lower half of the Dopplerbroadened absorption line of an atomic vapor, the gas can be cooled.' [43]. With this simple statement the field of laser cooling was launched. Wineland and Dehmelt at the University of Washington independently proposed a scheme for laser cooling in the same year [44].

Ions in a trap must have less kinetic energy than the energy of the potential barrier formed by the pseudoptential, and as we shall see in chapter 8 they must also be localised as much as possible to prevent variations in the coupling strength between the ion and optical cavity. Experimental laser cooling of trapped ions has been around since 1978, when Neuhauser et al. cooled a cloud of Barium ions in a miniature Paul trap [45]. The theory and practice of laser cooling for both neutral atoms and trapped ions has been extensively covered elsewhere—for a textbook see, for example [46], for a theoretical treatment of the laser cooling of a single ion see Cirac et al. [47] or for an extensive review of the field in 2003 with particular emphasis on trapped ions, see Eschner et al. [48]. In this section, therefore, we will restrict the discussion to the principles of the specific laser cooling technique we use in this work, following the straightforward approach given by Lett et al. in [49].

3.1 The quantum hardware—Calcium

3.1.1 The 40 Ca $^+$ ion

Singly-ionised calcium holds great promise as the hardware of quantum computing—its lowest five energy levels are already well-understood and accessible via commercially available or easily-built lasers and it has a forbidden transition to the ground state from the metastable $D_{3/2}$ state, giving long optical qubit lifetimes.

The transitions of interest for laser cooling are shown in figure 3.1. The main cooling transition is $P_{1/2} \leftrightarrow S_{1/2}$, with $P_{1/2} \leftrightarrow D_{3/2}$, $P_{3/2} \leftrightarrow D_{3/2}$, and $P_{3/2} \leftrightarrow D_{5/2}$ the so-called 'repumper' transitions, so named because when, as dictated by the branching ratio, an ion is 'shelved' in the metastable $D_{3/2}$ or $D_{5/2}$ states, it must be 'repumped' out of it to either the $P_{3/2}$ or $P_{1/2}$, from where it decays quickly to the ground state.

In addition to the transitions shown, there are several other transitions available within this



Figure 3.1: Partial energy level diagram of the ⁴⁰Ca⁺ ion showing the first five lowest-lying levels. The transitions we address in our experiments are marked, along with their wavelengths. Not to scale.

Transition	Wavelength	Rate/ 2π	Lifetime						
	(air)/nm								
	Electric Dipole								
$P_{1/2} \rightarrow S_{1/2}$	396.847	20.7 MHz	7.7(2) ns						
$P_{1/2} \rightarrow D_{3/2}$	866.214	1.7 MHz	94.3 ns						
$P_{3/2} \rightarrow S_{1/2}$	393.366	21.5 MHz	7.4(3) ns						
$P_{3/2} \rightarrow D_{3/2}$	849.802	0.2 MHz	901 ns						
$P_{3/2} \rightarrow D_{5/2}$	854.209	1.6 MHz	101 ns						
Electric Quadrupole									
$S_{1/2} \rightarrow D_{3/2}$	732.389	0.14 Hz	1.16 s						
$S_{1/2} \rightarrow D_{5/2}$	729.147	0.14 Hz	1.14 s						

Table 3.1: Electric dipole and quadrupole transition wavelengths, decay rates and natural lifetimes of the ⁴⁰Ca⁺ ion [50–52]. Uncertainties and accuracies differ because the various results originate from different papers. Details can be found in [50].

energy-level structure. A summary of all the transitions, their wavelengths and rates is shown below in table 3.1.

3.2 Doppler Cooling

In principle, Doppler laser cooling of a trapped ion is very straightforward—we tune a laser with wavevector **k** to a wavelength slightly longer ('redder') than that of an atomic transition, ω_{ab} , so that in the rest frame of an ion travelling towards the laser with velocity v, the laser frequency is Doppler-shifted by a factor of $1 \pm v/c$. The ion, therefore, sees 'bluer' light, that is, closer to

its resonance, and the probability of it absorbing a photon is higher. On absorbing a photon, the atom can re-emit via stimulated or spontaneous emission. In the case of stimulated emission, the photon is emitted back into the laser mode and so the momentum 'kick' received by the atom, $\rho = \hbar \mathbf{k}$ is cancelled out. Whilst the atom always absorbs a photon from the direction of the laser, its spontaneous emission occurs isotropically. The net momentum transfer to and from the ion by spontaneous emission therefore, is zero; however, there is a net transfer of momentum from the laser to the ion, since this always occurs in the direction of \mathbf{k} . The ion loses kinetic energy as a result.

For the work described in this thesis, laser cooling is performed on the $P_{1/2} \leftrightarrow S_{1/2}$ transition in ${}^{40}Ca^+$ (see figure 3.1 on page 18).

In order to understand the cooling mechanism we must first describe the laser-atom interaction. We start with a brief introduction to the density matrix formalism of describing quantum states, and then use it to obtain expressions for the force experienced by an atom in a laser field.

3.2.1 The Density Matrix Formalism

We consider the example of a two-level atom with a lower and upper energy states $|a\rangle$ and $|b\rangle$ respectively. In this case the orthonormal basis states spanning the Hilbert space are simply the states $|a\rangle$ and $|b\rangle$. Then the state vector of the system can be written as

$$|\psi\rangle = C_a |a\rangle + C_b |b\rangle, \qquad (3.1)$$

with $|C_{a,b}|^2$ the probabilities of finding the system is state $|a,b\rangle$.

An alternative and more versatile way to describe quantum states is by the **density operator**. For a pure state this is simply the projector operator, for example

$$\hat{\boldsymbol{\rho}} = |\boldsymbol{\psi}\rangle\langle\boldsymbol{\psi}| \tag{3.2}$$

and the matrix elements of this operator give us the **density matrix**. For the two-level system of equation 3.1 we have matrix elements

$$\begin{pmatrix} |C_0|^2 & C_0 C_1^* \\ C_1 C_0^* & |C_1|^2 \end{pmatrix}.$$
(3.3)

The diagonal terms are the probabilities of the atom being in state $|a\rangle$ or $|b\rangle$, while the offdiagonal terms represent coherent superpositions of the states. We can see that the trace $\text{Tr}\hat{\rho} = 1$ for a pure state.

The expectation value of the observable \hat{O} is given by

$$\langle \hat{O} \rangle = \operatorname{Tr}(\hat{\rho}\hat{O}),$$
 (3.4)

and the time evolution by

$$\frac{d}{dt}\rho = \frac{1}{i\hbar}[\hat{H},\hat{\rho}].$$
(3.5)

The density matrix has the advantage that it can also be used to describe a statistical mixture of

states. In this case the density operator is not simply the projector operator, rather it is

$$\hat{\rho} = \sum_{n} p_{n} |\psi_{n}\rangle \langle\psi_{n}|, \qquad (3.6)$$

that is, the sum of each state's density operator, weighted with its individual probability in the ensemble. The time evolution of the density operator for a mixed state is also given by equation 3.5, and its trace is also equal to 1.

The advantages of describing a state using the density matrix over the wavefunction or state vector are:

- It eliminates the arbitrary global phase factor that exists in a state vector description (the states e^{iθ} |ψ⟩ and |ψ⟩ are the same state. In the density matrix description this phase factor disappears).
- It can describe mixed, as well as pure, states.
- It gives a somewhat intuitive view of the states of the system. For example, in a two-level system the density matrix elements (1,1) and (2,2) represent the states |a⟩ and |b⟩.

3.2.2 The Optical Bloch Equations

We now consider how we may use the density matrix of a two-level system to make some specific calculations.

First, we consider the combined Hamiltonians of an atom and its interaction with a light field:

$$\hat{H} = \hat{H}_0 + \hat{H}_I,$$
 (3.7)

where the free atom Hamiltonian is given by

$$\hat{H}_0 = \hbar \omega_a |a\rangle \langle a| + \hbar \omega_b |b\rangle \langle b|, \qquad (3.8)$$

 $\hbar\omega_{a,b}$ being the respective energies of the states.

The dipole interaction Hamiltonian is given by

$$\hat{H}_I = -\mu \cdot \hat{E} \tag{3.9}$$

$$= -\left(|a\rangle\langle b| + |b\rangle\langle a|\right)\mu_{ab}E,\tag{3.10}$$

where μ_{ab} is the dipole matrix element of the transition, given by $\mu_{ab} = ex_{ab}$, if we assume the light is polarised in the *x*-direction.

The electric field of the light is given by

$$E = E_0 \cos \omega t = \frac{1}{2} E_0 \left(e^{-i\omega t} + e^{i\omega t} \right), \qquad (3.11)$$

and the Rabi frequency of the light driving the $a \rightarrow b$ transition is defined as

$$\Omega = \frac{\mu_{ab} E_0}{\hbar}.\tag{3.12}$$

We normally write the atomic raising and lowering operators $|a\rangle\langle b|$ and $|b\rangle\langle a|$ as $\hat{\sigma}_+$ and $\hat{\sigma}_-$ respectively, and so in the Heisenberg picture we can write the interaction Hamiltonian as

$$\hat{H}_{I} = -\frac{\hbar\Omega}{2} \left(\hat{\sigma}_{+} e^{i\omega_{ab}t} + \hat{\sigma}_{-} e^{-i\omega_{ab}t} \right) \left(e^{-i\omega t} + e^{i\omega t} \right).$$
(3.13)

Making the rotating wave approximation, we find

$$\hat{H}_{I} = -\frac{\hbar\Omega}{2} \left(\hat{\sigma}_{+} e^{-i\Delta t} + \hat{\sigma}_{-} e^{i\Delta t} \right).$$
(3.14)

where $\Delta = \omega - \omega_{ab}$.

We now make use of the time evolution of the density operator given in equation 3.5, adding a term to represent spontaneous decay [53]:

$$\frac{d}{dt}\hat{\rho} = \frac{1}{i\hbar}[\hat{H},\hat{\rho}] - \frac{\Gamma}{2}[\hat{\sigma}_{+}\hat{\sigma}_{-}\hat{\rho} - 2\hat{\sigma}_{-}\hat{\rho}\hat{\sigma}_{+} + \rho\hat{\sigma}_{+}\hat{\sigma}_{-}], \qquad (3.15)$$

with Γ the spontaneous decay rate of the transition.

An equation of motion of a density matrix is known as a **master equation**, and in this form is known as the **Lindblad form**. It can equivalently be written as [54]

$$\frac{d}{dt}\boldsymbol{\rho} = -\frac{i}{\hbar}[\hat{H},\hat{\rho}] + \frac{\Gamma}{2}(2|a\rangle\langle b|\hat{\rho}|b\rangle\langle a| - |b\rangle\langle |b\hat{\rho} - \hat{\rho}|b\rangle\langle b).$$
(3.16)

In general, a master equation can be written as

$$\frac{d\hat{\rho}}{dt} = \mathscr{L}\hat{\rho},\tag{3.17}$$

with \mathscr{L} the *Liouvillian superoperator*, which as we saw above can contain both pre-and postmultiplications using several other operators.

For a two-level atom, the time evolution of the elements of the density matrix are described by the **optical Bloch equations** [46]:

$$\frac{d}{dt}\hat{\rho}_{00} = -\Gamma\hat{\rho}_{00} + \frac{i\Omega}{2}\left(\tilde{\rho}_{10} - \tilde{\rho}_{01}\right),\tag{3.18}$$

$$\frac{d}{dt}\hat{\rho}_{11} = \Gamma\hat{\rho}_{00} - \frac{i\Omega}{2}\left(\tilde{\rho}_{10} - \tilde{\rho}_{01}\right),\tag{3.19}$$

$$\frac{d}{dt}\tilde{\rho}_{01} = \left(i\Delta - \frac{\Gamma}{2}\right)\tilde{\rho}_{01} + \frac{i\Omega}{2}\left(\hat{\rho}_{11} - \hat{\rho}_{00}\right),\tag{3.20}$$

where we have made the substitution $\tilde{\rho}_{01} = e^{i\omega t} \hat{\rho}_{01}$, and $\Delta = \omega - \omega_{ab}$, as before.

The steady-state solution to the optical Bloch equation for the excited state population is given by [46]

$$\hat{\rho}_{bb} = \frac{s_0/2}{1 + s_0 + (2\Delta/\Gamma)^2},\tag{3.21}$$

with $s_0 = 2\Omega^2/\Gamma^2$ the *on-resonance saturation parameter*.

The rate of transfer of momentum from the laser to the atom is determined by the steady state solution 3.21 and the rate of spontaneous decay Γ :

$$\gamma = \Gamma \rho_{bb}. \tag{3.22}$$

The rate of momentum transfer is the force experienced by the atom for one photon, multiplied by the scattering rate:

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = \hbar \mathbf{k} \cdot \boldsymbol{\gamma},\tag{3.23}$$

with the momentum ρ not to be confused with the density matrix $\hat{\rho}$.

So, using equation 3.21, we can now write the force on the atom as

$$\mathbf{F} = \hbar \mathbf{k} \cdot \frac{s_0/2}{1 + s_0 + (2\Delta/\Gamma)^2}.$$
(3.24)

In the rest frame of the atom, the laser frequency appears Doppler-shifted to ω_D

$$\omega_D = \left(1 \pm \frac{v}{c}\right) = \omega \pm kv, \qquad (3.25)$$

and the detuning Δ_D as seen by the atom becomes

$$\Delta_D = \Delta \pm kv. \tag{3.26}$$

Using this in equation 3.24, we have

$$\mathbf{F}(v) = \hbar \mathbf{k} \cdot \frac{s_0 \Gamma/2}{1 + s_0 + \left(\frac{2(\Delta \pm kv)}{\Gamma}\right)^2}$$
(3.27)

We can make a Taylor expansion of this force, around v = 0, and write the force as $\mathbf{F} = \mathbf{F}_0 + \beta v$, so

$$\mathbf{F}_0 = \frac{\hbar \mathbf{k} s_0 \Gamma/2}{1 + s_0 + (2\Delta/\Gamma)^2},\tag{3.28}$$

and

$$\boldsymbol{\beta} = -\frac{4\hbar |\mathbf{k}| \Delta}{\Gamma} \cdot \frac{s_0 \mathbf{k}}{(1 + s_0 + (2\Delta/\Gamma)^2)^2}.$$
(3.29)

We conclude from this that for small atomic velocities, the laser force is composed of a constant component, and a component that depends on the atomic velocity. The constant force is known as the *radiation pressure force*, and the velocity-dependent component is a viscous damping force. For cooling of neutral atoms, the radiation pressure force can be balanced out by a counter-propagating laser.

By balancing the cooling rate with the heating rate produced by the recoil of the atom from each emission event, it can be shown that the minimum temperature it is possible to reach is [49]

$$T_{\min} = \frac{\hbar\Gamma}{2k_{\rm B}},\tag{3.30}$$

and this occurs at a detuning of $\Delta = -\Gamma/2$. In calcium this corresponds to around 0.5 mK. Since

$$\frac{1}{2}mv^2 = \frac{1}{2}k_{\rm B}T,\tag{3.31}$$

we can write the minimum velocity as

$$v_{\min} = \sqrt{\frac{\hbar\Gamma}{2m}},\tag{3.32}$$

corresponding to around half a metre per second in calcium.

This argument, which applies to a neutral atom, also applies to the Doppler cooling of a single ion in a trap, although in a rf ion trap it is possible to cool a single ion using only one laser beam, angled such that it addresses every possible direction of motion of the ion. The ions are not free to continually move away from the laser, since the rf electric fields must always pull the ion back towards the trap centre.

3.2.3 Three-Level Model - repumping on the 850 nm and 854 nm transitions

Consider the simplified three-level diagram of the calcium ion shown in figure 3.2. It is quite straightforward to extend the master equation 3.16 to three levels, such as an atom with a metastable D-state, which we will denote $|m\rangle$, by adding an extra term :



Figure 3.2: A simplified three-level model of the ⁴⁰Ca⁺ ion.

$$\frac{\partial \hat{\rho}}{\partial t} = -\frac{i}{\hbar} [\hat{H}, \hat{\rho}] + \frac{\Gamma_b}{2} (2 \times |a\rangle \langle b|\hat{\rho}|b\rangle \langle a| - |b\rangle \langle b|\hat{\rho} - \hat{\rho}|b\rangle \langle b|)
+ \frac{\Gamma_m}{2} (2 \times |m\rangle \langle b|\hat{\rho}|b\rangle \langle m| - |b\rangle \langle b|\hat{\rho} - \hat{\rho}|b\rangle \langle b|).$$
(3.33)

Provided the Liouvillian superoperator of section 3.2.2 does not change in time¹, it is possible to write the solution of the master equation as the sum of a series of complex exponentials of the form $e^{s_j t}$ with each *s* an eigenvalue of the Liouvillian [55]. Once this is done, we can plot

¹As it would do with pulsed light, for example.

the expectation value for any element of the density matrix over time. Plotting the upper-state population as a function of the detuning of the 866 nm repumper (section 3.1.1) reveals a *dark resonance* at the point where the detunings of the 866 nm and 397 nm lasers are equal. The dark resonance is the signature of a *Raman transition* between the $D_{\frac{3}{2}}$ and $S_{\frac{1}{2}}$ state, which bypasses the $P_{\frac{1}{2}}$ state entirely. This is a situation which in our observation of continuous 397 nm ion fluorescence we would like to avoid. So rather than repumping on the $D_{\frac{3}{2}} \leftrightarrow P_{\frac{1}{2}}$ 866 nm transition, we implement an alternative setup, addressing the $D_{\frac{3}{2}} \leftrightarrow P_{\frac{3}{2}}$ transition with a 850 nm laser, and the $D_{\frac{5}{2}} \leftrightarrow P_{\frac{3}{2}}$ transition with a 854 nm laser (see figure 3.1 on page 18). This new setup entirely bypasses the $D_{\frac{3}{2}} \leftrightarrow P_{\frac{1}{2}}$ transition, avoiding any possibility of driving Raman transitions between those levels. Implementing this repumper setup allows us to treat the system as effectively two-level, and we will use this approximation in chapter 7, together with the two-level optical Bloch equations, to calculate the upper-state population.



Figure 3.3: A dark resonance occurs when the detunings of the 397 nm and 866 nm laser are equal, as shown in this three-level simulation. Here, the detuning of the 397 nm laser is fixed at $0.5 \cdot \Gamma$, and the detuning of the 866 nm laser is swept. We see the upper state population drop to zero as the 866 nm detuning passes $0.5 \cdot \Gamma$. This situation is avoided if we repump on the 854 nm and 850 nm transitions. In this way we decouple the metastable D-state from the P-state and create an effectively two-level system.
Chapter 4

Cavity Quantum Electrodynamics

Cavity quantum electrodynamics (CQED) describes how the radiative properties of atoms are modified by the proximity of boundaries to the electric field. It dates from E.M. Purcell's discovery in 1946 that the probability of nuclear magnetic moment transitions can be enhanced by coupling the system to a resonant electrical circuit [56]. In general terms this has become known as the *Purcell Effect* and refers to the effect on atomic lifetimes of modifying the boundary conditions on the vacuum around the system. An early review paper by Serge Haroche gives a very readable introduction to the field [57].

In the context of ion trap experiments, we refer to *optical* cavity QED. Optical because we place around the ion an optical cavity, usually consisting of two highly reflective mirrors. The confinement, high degree of localisation, and isolation from the environment available in an ion trap provides the ideal setting in which to test CQED theories and concepts.

In this chapter we first recap the general theory of optical cavities, before moving on to describe how simply placing a cavity around a radiating particle is able to dramatically alter the properties of the radiating atom. In general, detailed explanations of cavity-QED can be found in [58–60].

4.1 Optical cavities

We start by considering the simplest kind of cavity, that is, the *Fabry-Pérot* cavity. This consists of two mirrors placed parallel to each other a distance *L* apart. An electric field in the cavity must be zero at the boundaries, that is, the mirrors, and so a standing-wave is produced inside the cavity when the following condition is met:

$$L = n \frac{\lambda}{2}$$

= $n \frac{\pi c}{\omega}$, (4.1)

Rearranging equation 4.1 for ω gives us the frequencies of the standing waves allowed in a cavity of length *L*. These frequencies are termed axial, or longitudinal *modes*:

$$\omega = n \frac{\pi c}{L},\tag{4.2}$$



26

Figure 4.1: Schematic drawing of a Fabry-Pérot cavity formed of two concave mirrors M1 and M2, with radii of curvature R_1 and R_2 . The cavity losses due to the transmissivity of the mirrors, T_1 and T_2 are shown, with the cavity losses due to absorption and scattering by the mirrors represented by S_1 and S_2 .

and the free spectral range of the cavity is the frequency spacing of its longitudinal modes:

$$FSR = \frac{c}{2L}.$$
(4.3)

The spatial distribution of the electric field inside the cavity can be decomposed into a longitudinal mode function $\Phi(x, y, z)$ and two transverse mode functions $\Psi_n(x, z)$ and $\Psi_m(y, z)$:

$$\mathbf{E}_{nm} = (x, y, z) = E_0 \Psi_n(x, z) \Psi_m(y, z) \Phi(x, y, z), \tag{4.4}$$

with E_0 the amplitude of the electric field. The transverse electromagnetic (TEM) modes are labelled with the indices n and m

The transverse mode structure is defined by the *Hermite polynomials* H_m and H_n where each combination of the integers *m* and *n* define one particular transverse electric field distribution. A detailed derivation of these mode functions can be found in [61]. Let us assume that the cavity's axis is aligned along the *z*-direction, then the transverse mode structure is given by

$$\Psi_{m,n}(x,y,z) = \sqrt{\frac{w_0}{w(z)}} H_{m,n} \frac{\sqrt{2}(x,y)}{w(z)} \exp\left(-\frac{x^2 + y^2}{w^2}\right),\tag{4.5}$$

where w(z) is the radius of the mode at the axial point z, such that

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R}\right)^2},\tag{4.6}$$

where $z_R = \pi \omega_0 / \lambda$ is the *Rayleigh Range*, and represents the distance one must travel axially along the beam such that the area of the cross-section is doubled. w_0 is the beam waist.

The first Hermite polynomial is given by

$$H_0 = 1,$$
 (4.7)

so that equation 4.5 becomes

$$\Psi_{00}(x,y) = \frac{w_0}{w(z)} \exp\left(-\frac{x^2 + y^2}{w^2}\right).$$
(4.8)

This TEM_{00} mode is known as the **fundamental** or **Gaussian** mode. It has the smallest possible beam waist and the smallest divergence.

27

The longitudinal field mode function is given by [61]

$$\Phi(x, y, z) = \sin\left(kz - (1 + n + m)\arctan\frac{z}{z_{\rm R}} + \frac{k(x^2 + y^2)}{2R(z)}\right),\tag{4.9}$$

where the wavenumber $k = 2\pi/\lambda$ and R is the radius of curvature of the wavefront, such that

$$R(z) = [1 + (z_{\rm R}/z)^2].$$
(4.10)

The factor $\arctan(z_R)$ represents the relative phase difference between the TEM₀₀ Gaussian mode and a plane wave of the same frequency. The resonance condition is that the phase shift of a wave after one round trip of the cavity must be an integer multiple of 2π . The resonant frequencies of modes given by the transverse mode indices *m*, *n*, **and** the longitudinal mode index *q*, which must be a positive integer, are given by

$$\mathbf{v}_{nmq} = \mathbf{v}_{\text{FSR}} \left[q + \frac{1}{\pi} (1 + n + m) \arccos\left(1 - \frac{L}{R_i}\right) \right]. \tag{4.11}$$

Equation 4.11 tells us that in a confocal cavity, where $L = R_i$, transverse modes TEM_{nm} and TEM_{mn} are degenerate.

4.1.0.1 Cavity types and the conditions for 'stability'

There are several types of optical cavity, defined by the radii of curvature of the two mirrors, R_1 and R_2 , and the distance between them, *L*. A summary of some important cavity types is given here.

- Plane-parallel: $R_1 = R_2 = \infty$.
- Symmetric cavities: $R_1 = R_2 = R$.
- Confocal symmetric cavities: $R_1 = R_2 = L$.
- Concentric cavities: $L = R_1 + R_2$.
- Half-symmetric cavities: $R_1 = L, R_2 = \infty$.

We can use ray transfer matrices to treat the passage of a ray through the cavity. If we stipulate that the beam must reproduce its original shape after one round-trip of the cavity, then, after some matrix algebra (see, for example [62, Ch. 8]), we arrive at the condition

$$0 < (1 - L/R_1)(1 - L/R_2) < 1.$$
(4.12)

We define the cavity stability parameters

$$g_1 = 1 - \frac{L}{R_1} \tag{4.13}$$

$$g_2 = 1 - \frac{L}{R_2}.$$
 (4.14)

So the stability condition becomes

$$0 < g_1 g_2 < 1. \tag{4.15}$$

This, then, is the condition for the existence of a low-loss standing wave in a cavity¹. We must bear this condition in mind when deciding on cavity parameters. A representation of the stability criteria is shown in figure 4.2, along with the position of three stable cavity types. Any cavity that exists in the parameter space represented by the grey-shaded area will be stable, and we can see that stable symmetrical cavities lie along the $g_1 = g_2$ line.



Figure 4.2: The cavity stability diagram, defined by the lines $g_1g_2 = 1$.

4.1.1 Finesse

The parameters most often used to characterise optical cavities are the free spectral range of equation 4.3, and the spectral width of the TEM₀₀ cavity resonance. We measure the latter as the full-width at half-maximum δv in frequency units of the cavity resonance, or alternatively as the half-width at half-maximum κ in angular frequency units. In this form, and in the field of cavity-QED we usually term κ the *cavity field decay rate*, or the *cavity linewidth*. The relationship between δv and κ is therefore

¹This has nothing to do with the stability of the cavity *per se*, rather it is the condition for the cavity to support low-loss Gaussian-Hermite modes. 'Stability', however, is the common term for this.

$$\kappa = \frac{2\pi\delta v}{2}.\tag{4.16}$$

The *finesse* of the cavity \mathscr{F} is given by the ratio of the FSR to the cavity linewidth

$$\mathscr{F} = \frac{\mathrm{FSR}}{\delta v} = \frac{2\pi \mathrm{FSR}}{2\kappa},\tag{4.17}$$

or alternatively

$$\mathscr{F} = \frac{2\pi}{L_{\text{Tot}}},\tag{4.18}$$

with L_{Tot} the total round-trip loss, $= T_1 + T_2 + S_1 + S_2$.

The cavity decay rate κ depends on the total cavity losses, and the cavity length, as

$$\kappa = \frac{cL_{\text{Tot}}}{4L} = \frac{\pi c}{2L\mathscr{F}}.$$
(4.19)

The cavity finesse is entirely determined by the cavity losses, and so essentially describes the quality of the mirrors. When a photon hits a mirror, it is either transmitted or reflected and this does not depend on the length of the cavity. Finessse can be extremely high—Rempe et al. measured a finesse of 1.9×10^6 in 1992, with mirrors of transmissivity 2 parts per million (ppm) [63].

The finesse divided by π gives us the average number of times a photon crosses the cavity before it is lost, that is, the number of times the photon passes a trapped ion.

A typical cavity, which we use in our experiment, might have the following parameters:

- $T_1 = T_2 = 30$ ppm or 30×10^{-6}
- Cavity length $L \approx 300 \ \mu \text{m}$
- FSR = $c/2L \approx 500 GHz$
- Finesse $\mathscr{F} \approx 100,000$

With these numbers, figure 4.3 shows the influence of the absorption and scattering losses S_1, S_2 on the cavity transmission. This is an important consideration for the cavities considered in this thesis, and will be touched on in later sections.

29



30

Figure 4.3: Transmission as a ratio of input intensity for a cavity with the parameters given in the text, and for the case of no, 10 ppm, and 30 ppm absorption/scattering (A/S) losses. The figure of 30 ppm is a pessimistic estimate of losses in a cavity of the type which will be described later in this thesis.

4.2 Atom-Cavity Interaction

In the theory of quantum electrodynamics, which quantises the electromagnetic field, each mode of the field can be associated with a quantum harmonic oscillator and photons are interpreted as elementary excitations 'inside' these oscillators, or modes. Each oscillator may in principle contain from zero to an infinite number of excitations. The standard result of treating the simple harmonic oscillator quantum mechanically [64] gives the energy of the oscillator as

$$E_n = \left(n + \frac{1}{2}\right)\hbar\omega,\tag{4.20}$$

and this imposes a quantisation on the energy contained in the electric field in a cavity. The energy in any mode of the field, oscillating at angular frequency ω can be represented by an integer value *n* in equation 4.20. In this case there are *n* quanta of energy in the mode of the field².

The Hamiltonian of each of these oscillators, or modes, reads

$$\hat{H} = \hbar \omega \left(\hat{a}^{\dagger} \hat{a} + \frac{1}{2} \right), \tag{4.21}$$

where we define the photon *creation* and *annihilation* operators \hat{a}^{\dagger} and \hat{a} as respectively increasing and decreasing the number of photons in the mode by 1. Furthermore, we define the number operator

$$\hat{n} = \hat{a}^{\dagger} \hat{a}, \tag{4.22}$$

which has as its eigenstates the photon number states $|n\rangle$, or 'Fock States'. Each photon number

²Detailed accounts of the theory of quantum electrodynamics can be found in [59] and [60]

state represents an elementary excitation in the mode. The annihilation and creation operators obey the following relations:

$$\hat{i}^{\dagger}\hat{a}|n\rangle = n|n\rangle \tag{4.23}$$

$$\hat{a}, \hat{a}^{\dagger} = 1 \tag{4.24}$$

$$\hat{a}^{\dagger}|n\rangle = \sqrt{n+1}|n+1\rangle \tag{4.25}$$

$$\hat{a}|n\rangle = \sqrt{n}|n-1\rangle, \tag{4.26}$$

and any Fock state can be generated by applying the creation operator to the ground state $|0\rangle$ as

$$|n\rangle = \frac{a^{\dagger n}}{\sqrt{n!}}|0\rangle. \tag{4.27}$$

Equation 4.21 also reveals the energy offset of the vacuum state—we see that even if the eigenstates of the photon number operator are zero (no photons in the mode), the vacuum still possesses an energy $\hbar\omega/2$. It turns out that the expectation value of the electric field is zero in all photon number states. The average energy density, however, is non-zero for all states, including the vacuum state.

4.2.1 The Jaynes-Cummings Model

Figure 4.4: An atom in a cavity with the loss mechanisms indicated. Here, $S_1 + S_2$ are the combined absorption and scattering losses of both mirrors. Γ is the (spontaneous) emission of the atom into modes other than the cavity mode. *g* is the coherent atom-cavity mode coupling rate.

The coupling of a single two-level system, such as a trapped ion, with a single quantum harmonic oscillator (a single mode of the quantised cavity field) is described by the Jaynes-Cummings model [65, 66]. For this simple model, we assume that only one mode of the cavity is populated with

photons, and that this mode only interacts with the two levels of the atom, not any other levels. We also assume no dissipative processes, such as spontaneous emission or cavity losses.

32

A two-level system is often termed a *spin system* by analogy with a spin-1/2 system evolving between two spin eigenstates $|a\rangle$ and $|b\rangle$ created by a magnetic field oriented along a *z*-axis. We make use of the atomic raising and lowering operators, $\hat{\sigma}_+$ and $\hat{\sigma}_-$ of section 3.2.2 on page 20, to represent the transition of the atom from the lower to the upper, or excited state, and *vice versa*. The operators have the properties:

$$\sigma_{+}|a\rangle = |b\rangle \tag{4.28}$$

$$\sigma_{-}|b\rangle = |a\rangle \tag{4.29}$$

$$\sigma_+|b\rangle = 0 \tag{4.30}$$

$$\sigma_{-}|a\rangle = 0 \tag{4.31}$$

We write the Hamiltonian of the entire system as

$$\hat{H} = \hat{H}_a + \hat{H}_c + \hat{H}_{\text{Int}} \tag{4.32}$$

where \hat{H}_a , \hat{H}_c and \hat{H}_{Int} represent respectively the energies of the atom, the cavity mode, and the interaction.

The Hamiltonian for the atom is given by [59]:

$$\hat{H}_a = \frac{\hbar\omega}{2}\hat{\sigma}_+\hat{\sigma}_-,\tag{4.33}$$

where $\omega/2\pi$ is the atomic transition frequency.

The Hamiltonian for the cavity is $\hat{H}_c = \hbar \omega \hat{n}$, where we have set the vacuum energy to be zero for simplification—we are aware that there is actually a $\hbar \omega/2$ offset everywhere.

The electric field in a cavity can be written as [59]

$$\hat{E} = i\mathscr{E}_0 \left[\mathbf{f}(\mathbf{r})\hat{a} - \mathbf{f}^*(\mathbf{r})\hat{a}^\dagger \right], \qquad (4.34)$$

where $\mathbf{f}(\mathbf{r})$ is a function describing the electric field amplitude and polarisation across the mode (see the mode functions of equation 4.4). At the electric field maximum we normalise this expression so that $\mathbf{f}(\mathbf{r}) = 1$. The mode volume is

$$\mathscr{V} = \int |\mathbf{f}(\mathbf{r})|^2 d^3 \mathbf{r}, \qquad (4.35)$$

and the normalisation

$$\mathscr{E}_0 = \sqrt{\frac{\hbar\omega}{2\varepsilon_0\mathscr{V}}},\tag{4.36}$$

is the **r.m.s. electric field amplitude of the vacuum** for an angular frequency ω and mode volume \mathscr{V} . ε_0 is the usual permittivity of the vacuum.

The dipole interaction Hamiltonian is then written as

$$H_{\text{Int}} = -\mu_{ab} \left(\hat{\sigma}_{+} \hat{\sigma}_{-} \right) \cdot i \mathscr{E}_{0} \left(\hat{a} - \hat{a}^{\dagger} \right). \tag{4.37}$$

Here, μ_{ab} is the dipole matrix element of the atomic transition. If we expand this scalar product we find that the terms proportional to $\hat{\sigma}_{-}\hat{a}$ and $\hat{\sigma}_{+}\hat{a}^{\dagger}$ correspond to non-resonant processes (transition from $|b\rangle$ to $|a\rangle$ with the annihilation of one photon, and transition from $|a\rangle$ to $|b\rangle$ with the creation of one photon), and in the case that the cavity mode and atomic transition angular frequencies are close, then we can neglect these terms. This corresponds to the *rotating wave approximation*. The Hamiltonian then becomes, in a frame rotating with the laser frequency:

$$H_{\rm Int} = -i\hbar \frac{\Omega_0}{2} [\hat{a}\hat{\sigma}_+ - \hat{a}^{\dagger}\hat{\sigma}_-], \qquad (4.38)$$

where

$$\Omega_0 = \frac{2\mu_{ab}\mathcal{E}_0}{\hbar} \tag{4.39}$$

is the **vacuum Rabi frequency**, which can be thought of as proportional to the interaction energy of an atomic dipole with a classical field having an r.m.s. value of the vacuum field. Generally we quantify the strength of coupling between the atom and the single field mode as

$$g = \frac{\Omega_0}{2} = \frac{\mu_{ab}\mathcal{E}_0}{\hbar}.$$
(4.40)

This is the rate at which single excitations are coherently exchanged between the atom and the single field mode.

4.2.2 Dressed States

For the purposes of this experiment we consider a single ion and a single photon, so that the possible states are $|a,0\rangle$, $|b,0\rangle$, and $|a,1\rangle$, with the 0,1 referring to the number of photons in the cavity mode. For a resonant light field, the states $|a,1\rangle$ and $|b,0\rangle$ are degenerate. The interaction with the cavity causes a mixing of the states—the cavity field 'dresses' the bare states—such that the eigenstates of the coupled system become

$$\frac{1}{\sqrt{2}}(|a,1\rangle \pm |b,0\rangle). \tag{4.41}$$

The energy difference between the states is given by $2\hbar g = \hbar \Omega_0$. This is the so-called *vacuum Rabi* splitting [67].

In reality, there are dissipative processes in the system—spontaneous emission into modes other than the cavity mode and the other loss processes shown in figure 4.4. We now move on to quantify what we mean by a 'large' vacuum Rabi splitting, in terms of the size of the coherent coupling parameter *g* relative to these losses.

Ì



Figure 4.5: The 'dressed state' of the degenerate $|a,1\rangle$, $|b,0\rangle$ level at zero detuning resulting from the atom-cavity interaction is a splitting of magnitude $\hbar\Omega_0$.

4.2.3 Coupling Regimes

The proposal of Cirac and Zoller for a quantum network in [31] depends critically on the photon generation process being *reversible*, that is, a photon is emitted from an atom in a trap, transmitted along an optical fibre before entering an *identical* trap and mapping the state of the original atom onto the new atom. The key is the reversibility of the emission of a photon into the cavity mode, as compared to the irreversibility of the emission of the photon into the many alternative vacuum modes, or the loss of a photon through the mirror coating (the cavity decay rate). The irreversibility comes from the random coupling of the atom to the non-cavity modes—if a photon is emitted into one of these modes, we can't get it back. In the reversible regime, however, the atom emits a photon into the cavity mode, it bounces back and forth between the cavity mirrors and is re-absorbed by the atom, thus exactly reversing the emission process. Eventually the cavity photon must leak out of the cavity due to the finite reflectivity of the mirrors, but if it leaks out into one end of an optical fibre, the other end of which is connected to an identical cavity, then it can 'leak in' to the identical cavity mode and be absorbed by an identical ion, thus reversing the emission but in a different ion.

We characterise three coupling regimes by comparing the coupling rate g with the cavity loss rates to spontaneous emission Γ and the cavity decay rate κ of equation 4.19, which depends on the cavity loss rates and the length of the cavity.

If the reversible emission of photons into the cavity mode is faster than κ and Γ , the spontaneous emission rate, then we are in the strong coupling regime.³

The second regime of interest is the 'bad-cavity regime', in which $\kappa \gg g \gg \Gamma$. In this regime, decay of the photon through cavity losses occurs faster than the atom can re-absorb a cavity photon, and no vacuum Rabi splitting would be observed. However, emission into the cavity mode still dominates spontaneous emission, and as such this regime could satisfy the requirements of a quantum communication device.

If both the spontaneous emission rate and the cavity losses are greater than *g*, then we are in the 'weak-coupling regime.'

Experimentally, the vacuum Rabi splitting could be observed in the transmission spectrum of a cavity locked to the atomic resonance as we sweep the frequency of the input laser field to the cavity. We would observe two peaks in the spectrum, as the laser sweeps over the two frequencies of the split. To observe this, we need the split to be as large as possible, and so from equation

³One point worth mentioning regarding an *N*-atom case is that the splitting of the dressed states becomes $2g\sqrt{N}$ and so the conditions for strong coupling are relaxed by a factor of \sqrt{N} .

4.40, we require \mathscr{E}_0 , or the dipole matrix element of the transion to be large. At optical frequencies, this amounts to increasing \mathscr{E}_0 , which in turn requires the cavity mode volume \mathscr{V} to be decreased. The reduction of the cavity mode volume is the motivation for the design of the experimental apparatus described in chapter 5. We must also work to minimise loss mechanisms, which in practice means making the quality of the mirror coatings as good as possible. By using specialist coating companies⁴, it is possible to achieve mirrors with transmissivities of a few ppm. However, since we want to transmit our photon, it must leak out of the cavity eventually, so the transmissivity must be carefully chosen so that a high coupling rate is achievable, given what we know about the spontaneous emission rate of the atom and the possible cavity volumes, whilst still allowing us to extract a cavity photon. This process, along with calculations of potential cavity lengths, will be covered in chapter 8. In general, along with highly-reflective mirror coatings, cavity lengths of a few hundred microns are required.

Cavity parameters as a function of cavity-QED parameters are discussed in chapter 8.

4.2.4 Experimental realisation of strong coupling

A 2006 review of many significant experiments in both the weak and strong coupling regimes of CQED concentrating mostly on single atoms can be found in [68].

Strong coupling was achieved in the microwave domain in 1987 by Rempe et al. in Garching. A beam of Rubidium atoms was strongly coupled to the field in a superconducting microwave cavity. Rydberg atoms such as these have been used because of their very large dipole matrix elements, and long atomic lifetimes. One disadvantage to using Rydberg atoms is that it is harder to detect microwave photons and the experiments have relied on detecting atoms exiting the cavity.

In general, the coupling parameter *g* varies throughout the cavity, as the electric field of the cavity standing wave varies. For quantum information processing, we would like the coupling to be constant and predictable, and this is one reason for pursuing strong coupling in the trapped ion context, where we can localise an ion precisely in the centre of a node of the standing wave, where the coupling is strongest. Strong coupling has been realised using trapped ion Coulomb crystals by the group of Michael Drewsen in 2009 [69]. They were able to achieve strong coupling with crystals of over 500 ions and a cavity length of 11.8 mm.

Strong coupling has been achieved with other systems, such as a single quantum dot (2004) [70], and a micromechanical resonator (2009) [71], and the field is extremely active. An interesting new approach just suggested by Chang and Kimble [72] proposes a cavity formed of a lattice of atoms, with the cavity atom an impurity in the lattice.

To date, nobody has achieved strong coupling with single trapped ions, although the boundary has been reached [73].

³⁵

⁴Advanced Thin Films, Boulder, CO.

Chapter 5

Experimental Setup

We discussed in chapter 1 the idea of implementing a trapped-ion/single-photon system, and we saw in chapter 4, how placing a very small optical cavity around the ion can produce the coupling between the trapped ion and the cavity light field necessary for reversible single-photon production. With the idea of a miniature optical cavity firmly in mind then we now move on to describe our experimental setup.

Central to the experiment is a miniature ion trap of a unique design, such that an optical cavity consisting of the end facets of two optical fibres can be introduced within around 200 μ m of the ion, critically with no degradation of the trapping potential.

This chapter describes the overall physical setup of the experiment. Firstly, the production of the ions through photoionisation of calcium atoms is described, and a description of the lasers used to address the ions is given. The ion trap itself is described, starting with the critical trapping region and moving outwards through the supporting construction to the attendant rf electronics and vacuum systems. Finally, fluorescence detection and computer control systems are described. Where the setup differs from this basic layout, it will be described in the relevant chapter.

5.1 Overview

Figure 5.1 shows a top view schematic of the main components of the optical setup. At the centre of the figure is the octagonal vacuum chamber housing the trap—see section 5.6. Also inside the vacuum is a resistively-heated oven, filled with calcium powder. A current of around 1.8 A is passed through the oven, causing evaporation of calcium atoms into the trapping region at the centre of the trap. The direction of the atomic flux is indicated with an arrow. Lasers are directed through the chamber windows into the central region where a two-stage photoionisation process takes place followed by Doppler cooling of the resultant ions. The cooling lasers¹ cause the ions to emit fluorescence which we detect with a free-space photomultiplier tube (PMT)², and an electron-multiplying charge-coupled-device camera (EMCCD). For some experiments other detection schemes were used, and these will be described where they arise.

¹Labelled beams 1, 4, 6 and 7 as the 1st, 4th, 6th and 7th beams in our 397 nm laser system.

² 'Free-space' here means the PMT is outside the vacuum—the photons are detected after passing through the chamber windows to the atmosphere.



Figure 5.1: A top view of the main optical setup.

5.2 Fluorescence detection and computer control

5.2.1 Fluorescence detection

Fluorescence from a trapped ion is collected by two methods in our setup. The eight available windows and excellent optical access of the trap means that we are able to 'look' from more than one direction at the same time. We use a photomultiplier tube (PMT)³ to count photons, and an electron multiplying charge coupled device (EMCCD) camera⁴ to image the ions. Referring again to the optical table schematic shown in figure 5.1, let us refer to the window through which the EMCCD camera looks as the front of the trap. The lens system used to collect light from the ion is practically the same for the PMT and the EMCCD camera. A Nikon Measurescope objective lens of ten times magnification is mounted on an translation stage allowing movement along the optical axis and placed immediately in front of the front trap window and also the PMT window. Attached to these lenses are 22 cm pieces of light-tight tubing. Inserted in each tube is a Semrock FF01-395/11-25 bandpass filter. This ensures that only light of the fluorescence wavelength makes it to the detector (the filter is not perfect—for an examination of its transmission of 866 nm light, for example, see section 7.3.2 on page 104).

5.2.1.1 The PMT

On the PMT side of the trap, the tubing terminates in a small hole mounted in an adjustable xyz-stage (visible in figure 5.2(a)). This allows us to focus more easily on the central trap area. Once this is aligned a further pinhole is inserted, to block unwanted scattered light, followed by the Semrock filter. Aligning the whole PMT setup then follows this procedure:

- Look at the pinhole mount with a CCD camera positioned very close to the mount (figure 5.2(a). Focus the CCD on the mount.
- Focus the image of the trap electrodes on the plane of the mount (figure 5.2(b)) by translating the objective lens back and forth. Now, the electrodes are imaged onto the plane of the mount (as can be seen in figure 5.2(b).
- Move the mount hole using the *xyz*-stage so that the image of the electrodes is in the centre of the hole. That is, the centre of the trap is in the centre of the hole.
- Remove the camera, place a pinhole (150 μm in this case) and the Semrock filter into the mount, and attach the PMT.

This procedure guarantees that the PMT can see the central trapping region, however, since moving the pinhole by 1 mm corresponds to a moving a distance of 100 μ m inside the trap it's clear that if the pinhole is even half a millimetre away from the optimal point on the outside, it could be 50 μ m away on the inside and as a result the PMT may entirely miss the ion's fluorescence. Consequently a considerable amount of effort has been expended in systematically searching for ions by stepping the pinhole over the central trap region. We have found that over timescales of

³Hamamatsu H7360-02.

⁴Andor Luca S.

weeks the pinhole position drifts and has to be realigned. As a result we have chosen not to use pinholes smaller than 150 μ m, although doing so would have resulted in lower background light.

Some important characteristics of this PMT are its sensitivity to our wavelength of light, its dark count rate, and its dead time ⁵. The two-photon resolution of this PMT is quoted as 18 ns. This corresponds to a possible count rate of around 55 MHz, which is well above the rates we measure from a single ion. This feature, therefore, is not a concern. The dark count rate is specified to be $60-300 \text{ s}^{-1}$. Turning to sensitivity, this is specified as $4.1 \times 10^{-5} \text{s}^{-1} \text{pW}^{-1}$ at 400 nm, and so we can define its *quantum efficiency* as:

$$Q_{\text{eff}} = \frac{\text{count sensitivity at 397 nm in number of photons per second per pW}}{\text{real number of photons in 1 second per pW}}$$
$$= \frac{4.1 \times 10^5 \text{ s}^{-1} \text{pW}^{-1}}{2.0 \times 10^6 \text{ s}^{-1} \text{pW}^{-1}}$$
$$= 20\% \tag{5.1}$$



Figure 5.2: The alignment setup for the PMT. (a) - an image of the trap is focused onto a plane formed by a metal pinhole held in place on a *xyz*-stage, and (b) - the magnified image of the trap can be seen through the hole.

The influence of the solid angle subtended by the window (more precisely the brass clamp aperture—see section 5.4.8.1 dominates the collection efficiency of the setup, with the quantum efficiency of the PMT a distant second. One suggestion for improvement of collection efficiency would therefore be bigger windows (which means a bigger chamber). In order to take advantage of bigger windows however, we would also need an objective lens with a bigger numerical aperture, so the question comes down to machining costs and equipment costs.

5.2.1.2 The EMCCD Camera

For imaging trapped ions we have chosen to use an EMCCD camera. The significant advantage of EMCCD over mainstream CCD technology is that signals are amplified before reaching an amplifier, that is, on the CCD chip itself. This means that extremely low signal levels can be amplified above the noise of the amplifier.

⁵Dead time is the length of time the detector is 'dead' after detecting one photon and before becoming ready to detect another one.

As with the PMT, a Nikon measurescope objective lens of ten times magnification produces an image at the plane of the camera chip. A Semrock filter, as used in the PMT setup, is placed in the tubing just before the camera to filter unwanted light. It was found that, especially at high EM gain, spots of scattered 866 nm light were visible, so a small mirror, coated to reflect 866 nm light ('B'-coating), was placed in the tubing before the camera and reflects any stray 866 nm light back into the chamber. It was found that with this objective lens, one pixel on the camera's display corresponded to approximately 1 μ m in the plane of the ion.

The camera has been most useful during the times when the position of the PMT has drifted over time. Locating an ion first with the camera then optimising the PMT position has proved to be a relatively quick means of recovering from some misalignment or rearrangement of the experiment.

5.2.2 Generation of calcium ions



Figure 5.3: Energy levels and transition wavelengths in neutral calcium showing two possible photoionisation schemes. In scheme 1 intense 272 nm light is used in a three-stage process including decay on the 672 nm transition. Scheme 2, which we use, is a two-stage process on the 423 nm and 389 nm transitions. Energy levels are not to scale.

Neutral calcium has an atomic number of 20, and is shown in the yellow column in figure 5.3, with some of its neighbouring alkaline earth elements. Like these elements, for example beryllium, magnesium or strontium, calcium has two valence electrons, and preferentially loses one or two electrons to become singly or doubly ionised. The most abundant isotope—⁴⁰Ca—has 20 protons and 20 neutrons in its nucleus, and consequently no nuclear spin and no hyperfine structure. This isotope represents 96.94% of naturally occurring calcium.[74]

Calcium atoms are generated in the trap oven as described in section 5.3. There are several ways of generating ions from an atomic source in an ion trap. Electron bombardment, as its name suggests, bombards the calcium atoms with high-energy electrons. A more efficient alternative is the two-stage resonance-enhanced photoionisation process first demonstrated in 1999 by Kjaergaard et al.[75]. This is the method used in our work. The $4s^2 {}^1S_0 \leftrightarrow 4s4p {}^1P_1$ transition shown in figure 5.3 is excited by a 423 nm laser at 45° to the atomic beam. Exact wavelengths of all the lasers we use are shown in appendix A. The electron is then promoted to the continuum by a 375 nm laser. The exact wavelength of this laser is not critical, provided it is less than around 389 nm, and in fact an incoherent source has been used for this purpose by Lucas et al. in 2003 [76]. The two lasers are overlapped and focused to a spot size of approximately 200 μ m at the trap centre.

Photoionisation has several benefits over electron bombardment ionisation, some of which are:

- The isotope shifts in calcium are larger than the 423 nm laser linewidth, meaning we can selectively ionise only ⁴⁰Ca.
- Atoms of other elements that may exist in the trap will not be (resonantly) ionised.
- It is around five orders of magnitude more efficient than electron bombardment [77].
- As a result of the increased efficiency, the atomic beam flux can be decreased.
- An electron beam can charge up dielectrics inside the trap.

Details of the lasers used and their configuration can be found in appendix A.

5.3 The endcap ion trap

The ion trap used in the work presented here is pictured in figure 5.4. It is a miniature threedimensional Paul trap [32] based on the *endcap* design originally described by Schrama et al. [78] in 1993, and implemented since by, for example, Roberts et al. [79] and Wilson [80]. Our trap differs from this plan in that the central rf electrodes are hollow, rather than solid, tubes. This unique design allows the insertion of a fibre-based optical cavity into the actual trap electrodes.

A useful way of thinking of this design is as an extension of the *ring* trap design [81], which consists of a ring, to which an rf voltage is applied, and two endcaps with hyperbolic surfaces of revolution—in the endcap trap we can imagine the rf ring being split in two and pulled back over the two endcaps. These two endcaps become the rf electrodes in the new trap, and the two halves of the ring become ground electrodes. Alternatively, we can imagine the more extreme ring trap demonstrated in [82], which consisted only of a ring, surrounded at a distance with metal plates to which a dc voltage could be applied. We could imagine the ring in such a trap being split in two



Figure 5.4: The Endcap Trap.

and pulled apart to form the two hollow electrodes, and then just adding ground electrodes around them.

The effect of the outer electrodes is dramatically shown in the simulations of section 5.4.6 on page 51, where we see that the trap depth is enhanced by their presence.

The geometry of the endcap trap means that the potential produced is not the ideal quadrupole potential of equation 2.2 on page 8. For a non-ideal trap we must use a more general expression for the potential distribution [81]:

$$\Phi = \sum_{l=0}^{\infty} A_{2l} (x^2 + y^2 + z^2)^l H_{2l}(x, y, z),$$
(5.2)

where H_{2l} are spherical harmonics, and A_2 are expansion coefficients.

Near the centre of the trap, the potential expansion is dominated by the quadrupole potential, and so can be approximated by equation 2.2. Further from the centre, however, higher-order terms in the potential become significant and a trapped ion's motion at this point would no longer be

entirely harmonic. Since in this work we are only concerned with single laser-cooled ions held very close to the centre of the trap, we will neglect this concern and assume an approximately quadrupole potential distribution.

The influence of the higher-order potential terms causes another concern in the endcap trap—we must drive the trap with a higher rf amplitude to achieve the same secular frequencies as in an ideal trap. This increase in voltage can be characterised as the 'voltage loss factor', which we can think of as a loss of power to higher-order potential terms. Schrama et al. found the voltage loss factor of a trap geometry consisting only of two cylinders to be 3.5, and of an endcap trap consisting of two cylinders and two grounded electrodes to be 1.7 [78].

The relatively minor drawback of the voltage loss factor is outweighed in this case by the improved optical access to the trap. We are able to address a trapped ion with lasers at angles of up to 20% above and below the radial plane.⁶

5.4 Simulations

A number of simulations of the trap geometry and parameters have been carried out using the finiteelement program Femlab, controlled from a Matlab front-end. These have enabled us to predict the effect on the trapping potential of parameters such as the distance the fibres are recessed within the inner electrodes, the inner electrode separation and the rf amplitude. From these simulations we can extract the trap depth, as defined in chapter 2.5 and the secular frequencies of the trap under various conditions. These simulations were used to define the physical dimensions used in the trap.

5.4.1 The electric field and rf potential

Figure 5.5 shows a plot of selected electric field lines⁷ in the central area of interest within the endcap trap. These are the lines along which the resultant force is directed. If we imagine the direction of this force changing with every rf half-cycle, then we can easily see how a charged particle is trapped at the centre.

Figure 5.6 shows a snapshot of the electric potential resulting from the field. In the crosssections we can see the quadratic nature of the potential quite clearly.

⁶That is, perpendicular to the axial plane of the electrodes.

⁷Of course there are electric field lines everywhere in the trap, but for clarity the figure is restricted to the central trapping region.



Figure 5.5: Selected simulated electric field lines in a cross-section through the endcap trap. These indicate the direction of the force experienced by the positively charged ion.



Figure 5.6: A simulation of the instantaneous electric potential in the trap resulting from an rf amplitude of 200 V. Figure (a) shows a surface plot of the potential while figure (b) shows a cross-section through the radial and axial potential.

5.4.2 The pseudo-potential

The pseudo-potential is simulated for a range of rf amplitude of 50–400 V. Figure 5.7(a) shows a surface plot, while figures 5.7(b), 5.8(a) and 5.8(b) show cross-sections through the radial and axial pseudo-potentials. As expected, the depth of the pseudo-potential well increases with increasing rf amplitude.

44



Figure 5.7: (a) The pseudo-potential for an rf amplitude of 200 V and $\Omega/2\pi = 15$ MHz, and (b) axial and radial cross-sections through the pseudo-potential.



Figure 5.8: Cross-sections through the radial and axial pseudo-potential for rf amplitudes between 50 V and 400 V and a frequency $\Omega/2\pi = 15$ MHz, showing the increase in the depth of the potential well at higher amplitudes.

5.4.3 The effect of a dc-bias on the pseudo-potential

We now investigate the effect that a dc bias on the outer ground electrodes has on the simulated trapping potentials and secular frequencies. Later we will do this experimentally, and compare the measured secular frequencies with these simulations.

We refer to the dc-bias on the outer ground electrode as an *a*-value. The values of *a* are simply numbers used by the LabView control software, but they correspond to the real voltages on the electrodes via an empirically-determined linear relationship. We chose to investigate five *a*-values, with the corresponding dc voltages:

- a = -2: -7.46 V
- a = -1: -3.76 V
- a = 0: -0.06 V
- a = 1: 3.64 V
- a = 2: 7.34 V

We assume a *w*-value of zero. This is reasonable because although the real *w* compensation voltage varied throughout the day as the stray charges within the trap varied, this is balanced by the stray charge to give the effectively 'ideal' w = 0 situation which we model in the simulation. The voltages corresponding to our five *a*-values are shown in figure 5.9. Next, we will show the effect of adding these potentials to the existing pseudo-potential in the trap. Figure 5.10 shows as an example the combined rf and dc axial and radial pseudo-potentials for an rf voltages of 100 V. The simulations were run for 8 voltages between 50 V and 400 V, and the trap depths calculated from the combined potential. These are shown in figure 5.11. The 'missing points' indicate where the potential was anti-trapping.

The simulations in this section were performed with identical dc voltages on upper and lower electrodes, but for a more accurate simulation, the small gain and offset differences in the amplifiers supplying the upper and lower electrodes should be taken into account.



Figure 5.9: The radial and axial dc potentials corresponding to *a*-values of -2, -1, 0, 1 and 2.



Figure 5.10: Radial and axial cross-sections of the combined rf and dc potentials for an rf-amplitude of 200 V.

47



Figure 5.11: Trap depths extracted from combined dc and rf potentials.

5.4.4

We can extract the secular frequency from the combined pseudo-potentials in the following way:

Consider the equation of motion of an ion:

$$m\frac{\mathrm{d}^2 x}{\mathrm{d}t^2} = -e\frac{\mathrm{d}\phi}{\mathrm{d}x},\tag{5.3}$$

and considering a parabolic potential $\phi = ax^2 + b$, we have

$$\frac{\mathrm{d}\phi}{\mathrm{d}x} = 2ax$$
$$\Rightarrow \ddot{x} = -\frac{2eax}{m}.$$
(5.4)

We know from simple harmonic motion theory that $\ddot{x} = -\omega^2 x$, so substituting this in equation 5.4 we have

$$\omega = \sqrt{\frac{2ea}{m}}.$$
(5.5)

Therefore, if we know the *a*-parameter of a parabolic confining potential, we can determine the secular frequency, given the other constants (the electronic charge and the atomic mass of calcium in this case). The central region of the pseudo-potential in the endcap trap is approximately parabolic, so we first extract that region and fit the equation $y = a(x - c)^2 + b$ to it, as shown in figure 5.12. The fit to the equation then, for this particular example gives the equation $y = 174.8x^2 + 1.3$, with *c* negligibly small. Using equation 5.5 we obtain a radial secular frequency of 4.62 MHz. The secular frequencies calculated for all rf and dc voltages are shown in figure 6.14(a).



Figure 5.12: Extracting the secular frequency from the pseudo-potential. The pseudo-potential is shown for an rf amplitude of 400 V and dc potential a = 2, with the central region extracted and fitted to a parabola. See the text for further details.

5.4.5 The separation of the electrodes

We now consider the effect of the separation of the electrodes on the trapping potential. We saw in chapter 1 that the fibre separation, that is, the cavity, should be as small as possible, but we must also have an acceptably high trapping potential⁸.

First, we simulate a range of separations of the whole (inner and outer) electrode structure. That is, we move both electrodes up and down together with the inner electrode protruding a fixed distance 0.42 mm from the outer. The simulation is shown in figure 5.13 for separations between 50 μ and 1 mm. We find that at very small separations the radial confinement is very strong, with the axial almost non-existent. We can understand this if we consider that at small separations the electric field lines are almost entirely radial at the ion.

The radial potential only starts to approximate a parabola at around a separation of 300 μ m. Below this point the potential becomes more like a square well, and so we cannot say that the secular frequencies calculated from a parabolic fit in this region are reliable. This point is illustrated in the sum of deviations plot of figure 5.14. Each point is the sum of the deviations of the fit from the real data for that electrode separation. A high sum of deviations indicates a poor fit. The trap depths, on the other hand, are not calculated from this fit and are a reasonable measure of the potential well depth.



Figure 5.13: The secular frequency and trap depth as the separation of the inner electrodes is varied over a wide range of possible (and practically impossible) geometries. See the text for further explanation.

⁸We have found that with trap depths of several eV we are able to trap for several hours.



Figure 5.14: Figure (a) shows the highly non-parabolic radial pseudo-potential at an electrode separation of 50 μ m, and (b) shows the sum of deviations taken from a parabolic fit to the central pseudo-potential region for each electrode separation (see figure 5.12). The high values for the radial fits for very small electrode separations indicate the region in which the potential is no longer parabolic.

5.4.6 The recess of the outer electrodes

We now fix the inner electrode separation and examine the effect of moving the outer ground electrodes relative to the plane of the inner electrode face. Again, we examine the anharmonicity of the potential at each position in terms of the fit to a parabola. The results are shown in figure 5.15. The sum of deviations plot shows that the radial potential becomes less parabolic as we approach zero recess. However, trap depth increases, the closer the outer electrode comes to the inner. We can understand the increase in the trap depth by considering the electric field lines, more of which act on the ion as the outer electrodes approach the inner. Based on these models, a reasonable compromise between optical access and trap depths was made and an inner electrode separation of 446 μ m was chosen with the outer electrodes recessed by 420 μ m. Our main reason for recessing the outer electrodes from the inner at all is to improve optical access—the maximum angle available to a laser increases as we move the outer electrodes further back.



Figure 5.15: The variation in secular frequency and trap depth as we move the outer ground electrodes up and down in terms of their recess from the plane of the inner electrode face. Therefore, a negative value means the inner electrode is recessed inside the outer electrode. In this simulation we move from a practically impossible region (where the inner electrodes are entirely enclosed by the outer electrodes, meaning we have no way, in the current setup, of injecting an atomic beam to the trapping region, or of introducing a cooling laser beam), to a far-recessed region. We see that the potential more closely approximates a parabola as the outer electrode is moved further from the inner, but that the deviation is not large relative to the trap depth.

5.4.7 The effect of the fibres on the pseudo-potential

We now consider the recess of the fibres from the plane of the inner electrode face. We find that as we move the fibre closer to the face, the axial trap depth starts to increase by a very small amount. We intuitively see that the fibre in this case is behaving as a 'weak metal'. We considered, therefore, that the possibility of stray charge accumulating on the fibre facet was the most significant factor in our placement of the fibre, and decided to recess it by approximately 50 μ m inside the inner electrode.

52



Figure 5.16: The fibre recess inside the inner electrode is plotted for a variety of realistic distances. We see that the axial trap depth is affected by this parameter. In the experiment we chose to use a recess of approximately 50 μ m as a compromise between trap depth and ease of insertion, although the main reason for recessing the fibres at all is to minimise the effect of stray charges on their surfaces. These simulations were run with an rf amplitude of 200 V and no *a*-value.

5.4.8 The electrode structure

The critical central electrode region (hereafter referred to as 'the trap') is shown in close-up in figure 5.17.

The trap is cylindrically symmetric and consists of two stainless steel tubes of outer diameter 0.457 mm and inner diameter 0.254 mm.⁹ An rf voltage of amplitude approximately 200 V and frequency 15 MHz is applied to these electrodes. They are fixed inside two hollow ground electrodes of outer diameter 1.22 mm and inner diameter 1.02 mm.¹⁰ A ceramic spacer¹¹ insulates the inner from the outer electrodes. The spacer is glued in place with UHV-compatible glue. The ceramic and glue are recessed so as to prevent a direct line-of-sight to the trapped ion at the centre of the trap. The inner electrodes protrude a distance of approximately 0.45 mm from the outer electrodes. See chapter 6 for the simulations of the trapping potential which allowed us to choose these dimensions and parameters.

A capacitor (which can be seen in figure 5.23 connects the electrode structure to the trap mount structure, holding the outer electrodes at rf ground.

The unique feature of this trap is that we can introduce fibre-optic cables as close to the ion as we like, limited practically only by the inter-electrode separation. Inside each of the top and bottom electrodes is a Thorlabs BFH48-200 multimode fibre, designed for ultra-violet to visible light. The fibre has a cladding diameter of 230 μ m and so fits snugly inside the electrode. The fibres are recessed by a distance of approximately 50 μ m from the end of the electrode in order to prevent any degradation of the trapping potentials¹². The distance from the fibre to the ion is 277 μ m. The top fibre is guided out of the electrode structure by a stainless steel hook attached to a mount on the bottom flange. Both fibres pass down through the vacuum system underneath the trap

⁹Coopers standard size stainless steel tubing, gauge 26.

¹⁰Coopers Special Size Stainless Steel Tubing.

¹¹Frialit-Degussit Technical Ceramics.

¹²By the accumulation of stray charge on the dielectric surface.



Figure 5.17: Cutaway view of the electrode structure showing the inner rf electrodes, outer ground electrodes and recessed fibres. All dimensions are in millimetres. Ceramic spacers separate the ground and inner electrodes.

and are connected to feedthroughs on CF40 flanges. See section 5.6 for more details on this.

Calcium atoms are generated in a resistively-heated oven made of a 20 mm long tantalum tube of outer diameter 1.02 mm and wall thickness 0.94 mm. From here, a vapour of calcium atoms passes into the collimator tube, which is a 6 mm-long stainless steel tube of outer diameter 0.51 mm and inner diameter 0.254 mm.¹³ This is positioned at a distance of approximately 2 mm from the centre of the trap. It has been found that a current of approximately 1.7 A provides a sufficient flux of atoms.

An alternative method of loading calcium atoms into the photoionisation region would be the all-optical method used by, for example, Sheridan et al. [83]. In this scheme calcium atoms are ablated from a target within the trap by a high-power laser. The resultant vapour is photoionised in the usual way at the trap centre. Photoablation can also be used to generate ions without the photoionisation step, if the laser intensity is high enough [84]. This method has the advantage of being quickly switchable. When enough atoms have been loaded, the laser can be stopped

¹³Coopers Needleworks standard size stainless steel tubing: gauge 25.

immediately, in contrast to a hot oven which takes some time to cool down, filling the trap with unwanted atoms all the while.

The electrode mounting structure is shown in figure 5.18. It is attached to a custom-made CF63 flange, through which pass electrical feedthroughs to the atmosphere. These carry all the electrical connections noted on figure 5.23.



Figure 5.18: A simplified drawing of the trap mount showing the electrodes, oven collimator and oven. The top fibre is guided by the stainless steel hook.

All electrical connections to the trap pass through capacitative feedthroughs to atmosphere underneath the custom CF63 flange which the trap sits on. Two copper wires carry the oven current in and out. There are four dc connections to micromotion-compensation electrodes, which are used to compensate for stray charge in the trap. These electrodes are the upper and lower rf ground electrodes, a side electrode—simply a short wire pointing approximately horizontally at the trap centre (hidden behind the trap in figure 5.23, and the oven. All these connections pass out underneath the trap and into the metal box housing the resonator, from where they leave by various BNC and other connectors.

The dc voltages to the compensation electrodes are passively filtered by low-pass RC filters before entering the trap. The feedthrough to vacuum also acts as a low-pass filter.

5.4.8.1 The solid angle subtended by the fibres

One of the factors limiting the ability of the fibres to collect fluorescence from a trapped ion is simply how much of the fluorescence the fibre surface can 'see' from its position. We can quantify this in the solid angle subtended by the fibres at the ion.

In general, the solid angle which an area on the surface of a sphere subtends at the centre of the sphere is given by

$$\Omega = \frac{A}{r^2},\tag{5.6}$$

where A is the area of the surface and r is radius of the sphere.

The solid angle subtended by a conical surface is given by

$$\Omega = 2\pi (1 - \cos \theta). \tag{5.7}$$

If we imagine a cone of light which the fibre is able to accept from an ion at the centre of the trap (figure 5.19), we can calculate the solid angle subtended by the fibre at the ion.

The distance between the ion and the fibre surface is approximately 270 μ m (approximately because the fibre recess is only approximately known), and the radius of the fibre (core) is 100 μ m, and so the half-angle subtended by the fibre at the ion is 20.3°. Then the solid angle is, from equation 5.7, $\Omega = 0.39$ steradians. For both fibres, we can double this: $\Omega = 0.78$ steradians. This is about 6% of the full 4π solid angle.

By a similar argument we can quantify the limitations of the chamber windows on fluorescence detection by the PMT or camera by calculating the solid angle subtended by the windows at the ion. This is constrained by the brass clamps outside the windows. The radius of the clamp aperture is 13 mm, and the radial distance from the ion to the window is the radius of the inscribed circle of the octagonal chamber plus the thickness of the window, that is (38.5+6) mm, so the half-angle subtended by the window at the ion is 18.7°. The solid angle, then, is 0.33, or 2.6% of the full 4π solid angle.



Figure 5.19: The solid angle subtended by the fibres at the ion.

5.5 **RF drive and electrical connections**

5.5.1 RF drive

Any system for generating the rf voltage for an ion trap should ideally produce a single frequency with minimal noise. Harmonics of the rf frequency and noise can cause heating of the ion which leads to low trapping lifetimes. A resonant circuit is used to filter out unwanted frequencies.

The oscillating trapping potentials are generated by a Rhode and Schwarz SMG signal generator. From there the signal passes to a Mini Circuits LZY-1 +44 dB rf amplifier. The signal then passes to an autotransformer which steps up and filters the signal before it enters the trap.

5.5.1.1 The autotransformer

The components of the trap taken together have a total capacitance and resistance. If we place an inductance in series with this we have a series LCR circuit (figure 5.20), and we can use the properties of such a circuit at resonance to our advantage.

LCR circuits have a complex impedance which depends on the driving frequency, given by:

$$Z = R + i\omega L + \frac{1}{i\omega C},\tag{5.8}$$

where ω is the angular frequency of the voltage source. At resonance, $\omega_0 = 1/\sqrt{LC}$; the impedances of *C* and *L* cancel each other because the voltages across them are equal and opposite. So we can say that at resonance the impedance is entirely real, that is, resistive.



Figure 5.20: A series LCR circuit.

The inductance that makes the LCR circuit with our trap is an autotransformer. This not only steps up the rf amplitude, but in theory ensures only a single narrow frequency band is passed to the trap (the resonant frequency of the circuit). An autotransformer is a transformer with just one single winding—a single coil of wire forms both the primary and secondary winding. It has the advantage of being simpler and more compact than a two-winding transformer, but there is no electrical insulation between the windings. The autotransformer used in our ion trap consists of 30 turns of enamelled copper wire wound around a teflon cylinder of diameter 28.8 mm and length 156 mm. The first 2 turns form the primary coil and the last 28 turns form the secondary coil. One end of the coil is connected to the common ground and the other end to the load—that is, the trap. The rf input is connected as shown in figure 5.21. The ratio of primary to secondary voltage is given by the ratio of number of turns on the coils, and so we expect an amplification of a factor of 14.

To firstly estimate the resonant frequency of our coil-trap system, we can simulate an LCR circuit using a two-turn transformer. Calculating the inductance of each coil using the formula:

$$L = \frac{\mu_0 N^2 A}{l},\tag{5.9}$$

where A is the (circular) area enclosed by the coil, μ_0 is the vacuum permability, N is the number of turns, and l is the length of the coil. This gives 0.02 μ H and 4.9 μ H respectively for the primary and secondary coils. We measure the capacitance of the trap approximately, between the rf feedthrough to vacuum and the vacuum chamber ground using an LCR meter¹⁴, as 40 pF. Modelling the circuit¹⁵

¹⁴ISO-TECH LCR 819.

¹⁵Using the circuit simulator QUCS.



Figure 5.21: A drawing of the autotransformer used in our trap with its three electrical connections marked.

with an rf amplitude of 1 V shows a resonance at $v \approx 13$ MHz, with a FWHM $\Delta v \approx 0.35$ MHz. The results of the simulation are shown in figure 5.22(a). We quantify the narrowness of the resonance with the *quality factor* or *Q*-factor of the resonance, given by:

$$Q = \frac{v}{\Delta v},\tag{5.10}$$

bearing in mind that this equation applies to *power*, not voltage, so instead of using the full width at *half* maximum, we must use the full width at $1/\sqrt{2} \times$ maximum, which in this case gives $Q \approx 56$. We would now like to compare the real resonance of the trap/resonator setup with the simulation.



Figure 5.22: (a) The resonance of a simulated LCR circuit and (b) the resonance of the real trap, with Lorentzian fits. The Q-value of the experimental resonance is determined from the fit to be 41, while from the simulated resonance it is 56.

To measure the resonance of the circuit, two methods are employed—firstly we use a single coil of wire around the end of the autotransformer core as an rf pickup to observe the changing rf amplitude delivered by the transformer as we vary the frequency. This data is shown in figure 5.22(b). The data shows two peaks, at around 15 MHz and 19 MHz. The second peak may result

from the influence of the monitor coil on the circuit.

The second method used is to place an rf reflectometer in the path of the rf signal to the trap. For no power to be reflected back from the trap, the reactance of the trap must be zero, that is, the impedance must be entirely resistive—the resonance condition. Also, that resistance must equal the output resistance of the rf source. The power absorbed by an LCR circuit is Lorentzian in shape, with a full width at half maximum (FWHM) $\gamma = v/Q$.¹⁶ Fitting this data gives the centre of the resonance to be at 15.4 MHz with a FWHM of 0.45 MHz. This gives a quality factor of $Q \approx 34$.

The two methods both give a good approximate idea of the resonance condition of the trap, but the first method is possibly more useful for our purposes, since it is more representative of the quantity - the voltage supplied by the coil - which we actually want to use. The second method tells us essentially when the impedance of the circuit matches 50 Ω . Neither method necessarily provides an accurate measurement of the resonance frequency of the trap, since both are influenced by connecting wires and neighbouring components.

Since we only see a dip in the reflected power around 15 MHz, it seems reasonable to assume that the 15 MHz peak in the monitor coil signal represents the 'real' resonance. To determine the quality of this resonance, therefore, we fit to that part of the data and obtain a centre resonance of approximately 15.03 MHz with a width at $1/\sqrt{2} \times$ the amplitude of 0.37 MHz. This gives a quality factor $Q \approx 41$.

There are a number of reasons why the quality of our resonator might differ from the expected value. Chief amongst these may be imperfect impedance matching between the rf source and the trap. We also do not know the coupling between the coils of the autotransformer.

Suggested improvements to the resonator design would be to to adjust the coupling between the coils of the transformer, perhaps by using a two-coil device, so that the impedance-matching condition is met.

The electrical connections to the trap are shown in figure 5.23. In addition to the rf connections and the oven, there are four dc 'compensation' electrodes shown. These voltages are used for micromotion-compensation purposes, and will be discussed in detail in section 6.2 on page 70.

 $^{^{16}}$ The shape is Lorentzian only for high-Q resonators, that is, resonators with narrow resonances.

Experimental Setup



Figure 5.23: The electrical connections to the trap.

5.6 The Vacuum System

The term *Ultra-High-Vacuum* is not strictly defined by any regulatory body, nor is it used consistently. However, for the purposes of our work we will use the definitions given by the National Physical Laboratory.[85].

The following table summarises the various 'degrees of vacuum' in millibars.¹⁷

Degree of vacuum	Pressure Range/mbar
Low vacuum	1000–30
Medium vacuum	$30-1 \times 10^{-3}$
High vacuum	$1 \times 10^{-3} - 1 \times 10^{-6}$
Very high vacuum	$1 \times 10^{-6} - 1 \times 10^{-9}$
Ultra-high vacuum (UHV)	$1 \times 10^{-9} - 1 \times 10^{-12}$
Extreme ultra-high vacuum (EHV or XHV)	$< 1 \times 10^{-12}$ mbar

Table 5.1: Degrees of vacuum

By these definitions then, our trap is held in ultra-high-vacuum. For the duration of most of the work described here, the pressure was $1 \times 10^{-10} - 1 \times 10^{-11}$ mbar.

The octagonal vacuum chamber, pictured in figure 5.24, is 77 mm in diameter (inscribed circle),

 $^{^{17}}$ A note on pressures—1 bar = 1×10^5 Pascals = 750 Torrs. Atomospheric pressure is approximately 1 bar. We will generally use millibars in this work.


Figure 5.24: A drawing of the trap enclosed in its vacuum chamber.

and sits 140 mm above the optical table. It is welded to a CF63 flange, while the trap itself is attached to a second CF63 flange below this. It has 8 anti-reflection-coated windows, and one top window, held in place with an indium-wire seal and brass clamp. Three vacuum pumps bring the chamber from atmospheric pressure to $\sim 10^{-10}$ mbar. All pumps are located under the table. First, a roughing pump, so called because it produces a 'rough' vacuum—approximately 1×10^{-3} mbar. Then, a turbomolecular pump, or 'turbo' pump takes over. This consists of a spinning turbine which sucks gas from its inlet to the exhaust with successive spinning blades pushing the molecules down and compressing them eventually to the roughing pressure. Finally, an ion pump reduces the pressure to around 1×10^{-10} mbar and maintains it. Ion pumps work by creating clouds of electrons by a very high voltage ($\sim 1 \text{ kV}$) then trapping them using a very high magnetic field (~ 1000 Gauss). The clouds of electrons ionise gas molecules passing through the assembly by collision and the resultant positive ion is accelerated towards a cathode. It is then removed from the vacuum by reacting chemically with the material of the cathode. In our setup, once UHV has been achieved, the roughing and turbo pumps are switched off and only the ion pump maintains the vacuum.

We have also used a titanium sublimation pump to reduce pressure when we have been unable to reduce pressure using just the ion pump. This kind of pump is very simple—a current of tens of Amps is passed through a titanium filament, causing sublimation. The chamber walls become coated with the titanium produced. Since titanium is very reactive, it forms chemical bonds with gas molecules in the chamber, causing a drop in pressure.

Another issue relevant to pressure is that of rf power dissipated in the trap. This causes heating which causes the pressure to rise. Using an rf power meter, we have found that it is possible to dissipate approximately 0.5 W of rf power in the trap before the pressure increases significantly.

Underneath the CF63 trap flange a tube travels down through the metal rf autotransformer box, through a hole in a removable breadboard attached to the optical table, and is welded to a custom CF150 flange which is bolted to the underside of the table. Figure 5.25 shows from top to bottom the chamber, the resonator box, the breadboard and optical table, the custom CF150 flange bolted to the underside of the table, and finally the reducing cross leading down to the turbo pump beneath. In principle, the breadboard and optics could be entirely removed from the optical table, giving full access to the trap and chamber, for maintenance or bake-out work. In practice, however, some of the optics extend onto the surrounding optical table, so this has not been possible. An improvement for the next iteration of this trap would be to make sure that everything is included on the one removable breadboard.



Figure 5.25: The chamber and vacuum flange assembly. Beneath this is the turbo pump.

The whole chamber and vacuum structure is baked for several days at around 100°C to remove any water molecules, cleaning agent residues, solvents and any other molecules that will inevitably be present on and in the metal surfaces inside the chamber and flanges. The temperature is restricted to 100°C because of the indium seals on the windows, but this has not had any detrimental effect on our vacuum.

5.7 Magnetic Field Control

For reasons that will be explained in detail later, we would like to ensure that the magnetic field surrounding the trap is of a known magnitude and direction. To this end we must first compensate for the Earth's ≈ 0.5 G (50 μ T) magnetic field and then create a bias field of known magnitude and direction.

We surround the trap with three pairs of coils, as shown in figure 5.26, their axes orthogonal to each other. Above and below the trap are two circular coils of diameter 138 mm, separated by approximately 110 mm. This provides an approximately homogeneous magnetic field along the (cylindrical) axis of the trap. To provide an homogeneous field we would use the Helmholtz configuration, in which we separate the coils by their radius. In our case the set of coils had already been built for another experiment, and so we had to balance the potential inhomogeneity of the field produced against the time required to machine and build a new set of coils. Since we were only concerned with the homogeneity of the field in a very small central area of the trap, we decided to use the existing coils, and characterised them as follows. For two concentric circular coils a distance d apart and with N turns of wire, the magentic field on their axes at the midpoint between them is given by

$$B = \mu_0 N I \frac{R^2}{\left[R^2 + (d/2)^2\right]^{3/2}},$$
(5.11)

and the magnetic field produced on the axis of two rectangular coil of side lengths l and w at their midpoint is given, using the Biot-Savart law, by

$$B = \frac{\mu_0 NI}{4\pi} \left[\frac{l}{\sqrt{\frac{l^2}{2} + d^2} \times \sqrt{\frac{l^2}{4} + d^2}} + \frac{w}{\sqrt{\frac{w^2}{2} + d^2} \times \sqrt{\frac{w^2}{4} + d^2}} \right].$$
 (5.12)

The magnetic field produced was measured at the centre of each pair, and found to be 22 G/A for the circular pair and 6 G/A for the rectangular pair. Using the above equations, there are approximately 250 turns of wire on the circular coils and 300 turns on the rectangular coils.

The two radial dimensions are covered by two pairs of rectangular coils, which again provide approximately homogeneous fields in the trapping region.



Figure 5.26: A drawing of the trap surrounded by the magnetic field coils.

5.7.1 Experiment Control

Throughout the experiments reported in this thesis control has been provided by LabView running on Windows XP. LabView vi's have mainly been written by Dr. Hiroki Takahashi, and control most of the devices we have used. This includes switching the oven on and off, switching the shutter for the photoionisation lasers, reading, recording and cross-correlating signals from the PMTs, automatically compensating for micromotion and measuring secular frequencies (chapter 6).

5.7.1.1 Controlling laser power—accousto-optic modulators

The 397 nm light from the Toptica system and the light from the 866 nm laser is split into several 'arms' using a series of half-wave plates and polarising beam splitter (PBS) cubes. A PBS essentially splits linearly-polarised light into its orthogonal components, transmitting one and reflecting the other. By altering the input polarisation one can alter the power passed to that arm of the setup. By this method we divide the 866 nm and 397 nm laser power into the amount of beams we require for our experiment, and for other experiments in the lab.

After power distribution, each 397 nm, 866 nm and 423 nm beam passes through an accoustooptic modulator (AOM) in the double-pass configuration.[86, 87]. A typical arm of the 866 nm laser setup is shown in figure 5.27. The efficiency of diffraction by the AOM into the first order depends on the amplitude of the rf delivered to the crystal, and this allows us to control the power of the beams at the trap (by simply varying the amplitude of the rf sent from the to the AOM), and to shutter the beams (by setting the rf amplitude to zero). By varying the *frequency* of the rf delivered to the AOM, one can vary the amount by which the diffracted light is shifted in frequency, and so exercise control over the laser wavelength without tuning the main laser diode, and without affecting the coupling of the laser to the fibre. This feature was not used for the work presented here.



Figure 5.27: A typical arm of the 866 nm laser AOM setup showing the main components and the beam path.

Chapter 6

Characterisation of the endcap trap

Every ion trap has a unique set of parameters, or characteristics, that define its operation. These depend mainly on elements of the trap geometry, for example the size and shape of the electrodes, the proximity of dielectrics to the ion, or the rf supply. In this chapter a characterisation of the endcap trap is presented, starting with some comments on trapping single ions, then moving on to the problem of micromotion. We end with a thorough investigation of the secular frequencies of the ions in the trap and the trap depth, with a comparison to the expected values from simulations.

The format of this chapter roughly follows the day-to-day procedure we followed in preparing a single ion for the more advanced experiments to be described in later chapters. In outline, then, this consists of:

- 1. Trapping a single ion.
- 2. Micromotion compensation.
- 3. Defining the magnetic field.
- 4. Performing laser spectroscopy on the trapped ion to ensure laser detuning is optimal.

The measurements of secular frequencies and trap depth are not a part of the preparation routine for any particular experiment, but are included as figures of interest in describing the trap. We find in the literature these figures often quoted in descriptions of ion traps, and their use gives the reader a good general idea of the trapping potential of a particular trap.

6.1 Trapping ions

The procedure for trapping a single ion is as follows. First, the cooling and photoionisation lasers are switched on and allowed a warm-up period of anything from one to several hours. Next, the AOMs are switched on manually, and the voltages sent to their VCOs are adjusted using a LabView vi until suitable laser powers are reached. Typical powers are:

- 1. 397 nm cooling lasers: $\sim 20 \ \mu$ W each.
- 2. Repumper lasers: $\sim 400 \ \mu$ W.

3. Photoionisation lasers: $\sim 400 \ \mu W$ each.

The laser frequencies are set and locked using the methods described in section ?? on page ??. Often, the temperature of the laser diodes will need to be adjusted to bring the wavelength close enough for fine adjustment. For optimal Doppler cooling, we saw in chapter 3 that the detuning should be $\Delta = -\Gamma/2$. We know that Γ for the $S_{1/2} \leftrightarrow P_{1/2}$ transition is around 20 MHz (section 3.1.1, page 17), and this translates to a detuning of approximately 0.01 pm¹. In practice, however, the wavemeter only has an stated accuracy of ± 60 MHz, so when we need optimal detuning, we determine it by examining spectroscopic lineshapes.

Once the lasers are ready, the calcium oven is switched on using a LabView vi signal to a current power supply. A current of around 1.8 A for around 45 seconds usually provides sufficient atomic flux across the trapping region to trap ions. The geomtry of the trap means that we rarely trap more than two ions at a time, but should this occur we eject unwanted ions from the trap by briefly heating them. This is done by blocking the repumper laser for a second or two. Cooling is drastically reduced and ions far from the trap centre, that is, hotter ions, are preferentially heated out of the trap. By this method we can usually reduce the number of ions to one very quickly. Figure 6.1 shows a single ion in the trap with the top electrode visible. We are able to determine the coordinates from knowing the magnification of the optics in front of the camera and the resolution per pixel of the camera's CCD chip.

In figure 6.2 we implement the alternative repumper method described in section 3.33, but using an extremely low-power 850 nm laser on the $D_{\frac{3}{2}} \leftrightarrow P_{\frac{3}{2}}$ transition. The sudden drops in fluorescence reveal the failure of the weak 850 nm laser to pump the ion out of the metastable $D_{\frac{3}{2}}$ state, that is, the abrupt transition of the electron to and from the $D_{\frac{3}{2}}$ state. These fluorescence spectra demonstrate nicely the phenomenon of quantum jumps.

67

¹A note on detunings: We measure on our wavemeter the fundamental frequency of light from the Toptica frequencydoubling diode laser system, not the doubled frequency. We have experimentally determined the resonance of the $S_{1/2} \leftrightarrow P_{1/2}$ transition (as will be shown by spectroscopy later) to lie around a fundamental wavelength of 793.91835 nm. In day-to-day ion trapping, where the cooling is not critical, we operate at a wavelength of around 793.91840 nm, that is, a detuning of 0.00005 nm. Since $v = c/\lambda$, we calculate the frequency detuning as $\Delta v = -(c/\lambda^2) \times \Delta \lambda$, which amounts to approximately 71 MHz. This means that the frequency-doubled light is detuned by twice this, or 142 MHz.



Figure 6.1: A single ion in the trap. The upper electrode is visible bathed in scattered 397 nm laser light, with the lower electrode slightly out of picture. The origin of coordinates has been set to the ion's position.



Figure 6.2: Quantum jumps on the $D_{\frac{3}{2}} \leftrightarrow P_{\frac{3}{2}}$ transition. Here we use an extremely low power 850 nm repumper at 0.5 μ W. The low repumping efficiency on this transition causes the ion to become occasionally shelved in the metastable $D_{\frac{3}{2}}$ state, where it stays until pumped out by a 850 nm photon (or naturally decaying after around a second).

Figure 6.3 shows scans over the $S_{\frac{1}{2}} \rightarrow P_{\frac{1}{2}}$ 397 nm resonance using each of the three principle cooling lasers indicated in figure 5.1 on page 37. Spectroscopy using the low-scatter beam is shown in section 7.2.1 on page 97.



Figure 6.3: Spectroscopy of the $S_{\frac{1}{2}} \rightarrow P_{\frac{1}{2}}$ cooling transition in a single ion using the three main cooling beams: 'beam 4', 'beam 6' and 'beam 7', named for the numbering system of the AOM branches in the Toptica 397 nm laser setup. Micromotion compensation had been optimised. The lineshapes are fitted to a half-Lorentzian curve. The signal-to-background ratios and fluorescence at a detuning of 30 MHz are as shown. The fluorescence was measured at the free-space side PMT.

6.2 Micromotion minimisation

The ability to localise a trapped ion very precisely at the rf zero-point in a trap is an important prerequisite to many cavity-QED experiments (see chapter 4). However, stray electric fields resulting from build-ups of charge on dielectrics inside the trap structure can displace an ion away from the trap centre. The effect of this is that the ion experiences a driven motion due to the ac electric field at the new position. The further from the trap centre the ion moves, the greater is the amplitude of the micromotion. This driven motion is termed *excess micromotion*, in contrast to the *intrinsic* micromotion induced as the ion's secular motion moves it back and forth through the rf zero-point. Unlike intrinsic micromotion, excess micromotion cannot be reduced by laser cooling. The Doppler shift in the ion's resonance frequency, causing a spread of frequencies $\Delta \omega = 2\omega_0/c$, is an unwanted extra factor in the experiments which will be described later, and in our work has been significantly reduced by the two methods described in this section.

Compensation of stray fields in three dimensions is achieved using the four dc compensation electrodes shown in figure 5.23 on page 60. In the radial plane, dc is applied to the oven collimator and to a thin wire positioned approximately 90° radially from it (the 'side' electrode)². In the axial plane, dc is applied to the top and bottom ground electrodes (see section 5.4.8 on page 53). Signals are sent from a LabView vi to a National Instruments data acquisition (DAQ) device³. This generates analogue voltages which are sent on to the compensation electrodes.

We introduce four parameters: u, v, w and a^4 to represent these voltages in our LabView control software, and we will denote the dc voltages sent from the DAQ as:

- A0 = u: Oven
- A1 = v: Side
- $A2 = \frac{w+a}{2}$: Upper ground electrode.
- $A3 = \frac{-w+a}{2}$: Lower ground electrode.

In addition, before reaching the electrodes the signals pass through a high-voltage amplifier, so there is a known gain and offset in the real compensation voltages realtive to those sent from the DAQ.

The effect of u and v is to push the ion along the direction of the oven and side electrode respectively, while w has the effect of moving the ion up or down (equal and opposite voltages are applied to top and bottom). a provides an dc offset to the upper and lower ground electrodes. We can think a as similar to the a-value described in equation 2.9 on page 11, although it differs in that here the dc is on the ground electrodes, rather than the rf electrodes.

These, then, are the tools with which we move the ion to the rf-zero point using the following methods.

²Note that the side electrode has a component in the axial direction. It was not possible make it entirely radial, for optical access reasons.

³NI-USB 6229.

⁴Not to be confused with the a-parameter of chapter 2.

6.2.1 Coarse micromotion minimisation

The principle of this method is that if the ion is sitting at some position far from the trap rf zero-point, as illustrated in figure 6.4 with the parabola being the trap pseudo-potential, then if we reduce the steepness (or 'stiffness') of the trapping potential (recall from chapter 2 that we can do this by altering the rf amplitude), then the ion will 'drop' to a new position. The objective is to move the ion to a position such that changing the rf amplitude causes no observable movement. We watch the ion with the EMCCD camera, and mark its position on the screen. We then decrease the trap stiffness. If the ion moves, we push it back to its original position with the compensation voltages. We then decrease the trap stiffness and repeat the process until the ion's position is independent of the rf amplitude. We see from figure 6.4 that this happens at the rf zero-point. This method is not sufficient, however, since we cannot use it to observe the micromotion in the direction of the EMCCD camera. Plus, we cannot use it for particularly sensitive adjustments.



Figure 6.4: Coarse micromotion compensation. If the ion is at some point far from the trap rf zero point, then if we suddenly decrease the rf amplitude ('open' or weaken the trapping potential), then the ion will shift to a new position, and this movement will be obvious on a CCD camera.

6.2.2 **RF-Photon correlation**

In order to more precisely minimise the micromotion, we look at the correlations between the arrival times of photons from the PMT positioned at the side of the trap, and a point in the phase of the trap rf [88]. If we consider micromotion along the direction of one red-detuned cooling laser first, we can see that the probability that the ion will scatter a photon increases as the ion moves towards the laser, and decreases as it moves away. Thus there is a correlation between the phase of the rf and the ion's fluorescence and this manifests as a sinusoidal variation in fluorescence. The objective of this method, then, is to detect this variation, and reduce the amplitude of the sinusoid as much as possible. The device used to measure the times between photon arrivals and the rf signal is an ACAM AM-GP2 time-to-digital converter.

We found that using the main trap signal generator to supply both the trap rf signal and the

ACAM correlation signal introduced some noise to the trap which may have caused poor trapping lifetimes. As a result we implemented the rf signal routing shown in figure 6.5.

72



Figure 6.5: The rf-photon correlation signal path. To avoid splitting the signal from the trap rf signal generator the following system is used. We take the 10 MHz timebase from the trap rf drive and amplify it with an inverting amplifier. This drives a signal generator which generates a 15 MHz rf signal in phase with the trap drive. This is then divided by four and the final signal sent to the ACAM card. The inverting amplifier is used because the amplitude of the 10 MHz timebase output of the Rhode & Schwartz is not enough to drive the DS345 signal generator. By dividing the rf signal by four we give ourselves four times as much correlation data as we would have otherwise had.

A typical series of rf-photon correlation spectra is shown in figure 6.6. These data are acquired in the following manner: acquire 20,000 data points (a data point is one $\Delta t = t_{\rm rf} - t_{\rm photon}$), plot a histogram of the time differences (these will never exceed ~0.26 µs because the rf frequency is ~15 MHz and we observe correlations over four rf cycles), and fit a sine to this.

6.2.3 Automatic 3D micromotion minimisation

The electric field at a position \mathbf{x} is given by (apart from a constant)

$$\mathbf{E}(\mathbf{x}) = (x, y, -2z),\tag{6.1}$$

with *C* a constant. Consider the micromotion along the direction of one cooling laser, pointing in the direction $(k) = (k_1, k_2, k_3)$, then that micromiton is minimised when there is no component of the electric field along that direction. That is,

$$\mathbf{E}(\mathbf{x}) \cdot \mathbf{k} = \mathbf{0}.\tag{6.2}$$

We can write this as

$$k_1 x + k_2 y - 2k_3 z = 0. ag{6.3}$$

This is the equation of the plane

$$\mathbf{x} \cdot \mathbf{n} = \mathbf{0},\tag{6.4}$$

with

$$\mathbf{n} = (k_1, k_2, -2k_3). \tag{6.5}$$

The distance an ion is pushed is proportional to the voltage applied to the electrode⁵. Consider the real-space directions \mathbf{u} , \mathbf{v} and \mathbf{w} as defined by the directions of the compensation electrodes, with a voltage determined by u, v, and w. Then the ion's displacement is given by

$$\mathbf{x} = A_u(u - u_0)\mathbf{u} + A_v(v - v_0)\mathbf{v} + A_w(w - w_0)\mathbf{w},$$
(6.6)

where A_u , A_v and A_w are the gains of the amplifiers used to supply the final voltages. u_0 , v_0 and w_0 represent stray electric fields in the trap. Substituting this into equation 6.4, we have

$$A_u(u-u_0)\mathbf{u}\cdot\mathbf{n} + A_v(v-v_0)\mathbf{v}\cdot\mathbf{n} + A_w(w-w_0)\mathbf{w}\cdot\mathbf{n} = 0.$$
(6.7)

Our scheme for automatic 3D micromotion minimisation uses a LabView vi and the above theory, but operates entirely in *uvw*-parameter-space ('*uvw*-space'). The program operates in the following way:

- 1. Start with laser 1. By moving around in *uvw*-space, find a set of voltage triplets defining a micromotion minimum plane P_1 . Let \mathbf{n}_1 be the normal vector to the plane in *uvw*-space (figure 6.8(a)). From equation 6.5 we see that once we know this vector, provided the direction of the laser never changes, it is fixed. In principle, then, we only need to empirically determine this direction once, and we can then use it again and again. The position of the plane P_1 along this vector however can, and does, change, depending on stray charges inside the trap. So if we move along a direction parallel to the normal vector, call it \mathbf{d}_1 , we must eventually encounter the plane P_1 (figure 6.8(b)). The program does this by moving 'up' or 'down' a direction parallel to the normal vector, acquiring 20,000 data points in the manner described above, for each set of *uvw*-values, and attempting to reduce the amplitude of the sine fit for each set. The number of *uvw*-points attempted can be set at will, but generally 5 were used. It then performs a linear fit on the 5 points and moves to the zero-micromotion point in *uvw*-space. A typical run of this program for one laser is shown in figure 6.7.
- 2. Repeat step 1 with laser 2 to determine \mathbf{n}_2 .
- 3. Find the intersection of P_1 and P_2 . This is the line **L** (figure 6.8(c)). Since this lies on the planes P_1 and P_2 , it must be perpendicular to both \mathbf{n}_1 and \mathbf{n}_2 , so we can write $\mathbf{L} \parallel \mathbf{n}_1 \times \mathbf{n}_2$. If we move perpendicular to L, then we must eventually hit it, but we must also remain on one

73

2

⁵If a uniform electric field is created then $\mathbf{F} = e\mathbf{E}$. A greater electric field means a proportionally greater force, therefore acceleration, and since $\phi = \int \mathbf{E} d\mathbf{r}$, then distance travelled is proportional to potential difference.

of the planes P_1 , or P_2 , so we can say that the direction to move, \mathbf{d}_2 , must satisfy:

$$\mathbf{d}_2 \perp \mathbf{n}_1 \tag{6.8}$$

$$\mathbf{d}_2 \perp L \tag{6.9}$$

$$\Rightarrow \mathbf{d}_2 \perp \mathbf{n}_1 \times \mathbf{n}_2 \tag{6.10}$$

$$\Rightarrow \mathbf{d}_2 \parallel \mathbf{n}_1 \times (\mathbf{n}_1 \times \mathbf{n}_2) \tag{6.11}$$

4. The global micromotion minimum must lie somewhere along the line L. The third direction, then, is just parallel to L (figure 6.8(d)).

We can summarise the three directions then, as:

$$\mathbf{d}_1 = \mathbf{n}_1 \tag{6.12}$$

$$\mathbf{d}_2 = \mathbf{n}_1 \times (\mathbf{n}_1 \times \mathbf{n}_2) \tag{6.13}$$

$$\mathbf{d}_3 = \mathbf{n}_1 \times \mathbf{n}_2. \tag{6.14}$$



Figure 6.6: Moving in uvw-space. RF-photon correlation spectra for six different sets of u,v, w compensation voltage parameters. The sine fit gives the amplitude and phase of the micromotion for those parameters. The flip in the phase of the sine as the micromotion passes over the zero micromotion point is evident in (a) and (e). See the text for details.



Figure 6.7: Automatic 3D micromotion minimisation. A plot of micromotion amplitudes taken from the sine fit in figure 6.6. The 's-parameter' of the axis refers to the distance, in *uvw*-space, along the normal vector to the micromotion minimum plane. The distance along this direction that the program has access to, in its attempts to find the micromotion zero-point, can be set by the operator.

6.2.3.1 Sensitivities of the techniques

The first technique, that of monitoring the spatial motion of the ion as the trap stiffness is changed, is less sensitive than that of rf-photon correlations. We find that after we have done what we can to eliminate the ion's spatial dependence on the rf amplitude, the rf-photon technique always shows a measurable micromotion in the correlation spectra. This makes sense if we consider that we cannot see the ion's motion in the direction of the camera, and also that the technique cannot detect phase differences in the rf applied to the top and bottom electrodes (considering the ion's equation of motion, we see that its average position would not change even if the phase of the rf did). The relative merits of the techniques are discussed in much greater detail in [89].



77

Figure 6.8: Mapping the parameter space of the dc compensation voltages. (a): A real micromotion plane fitted to, in this case, 9 *uvw*-triplets, each representing an empirically determined point of minimum micromotion. This is performed for each laser. Once we know the normal for each laser we follow the procedure outlined in the text and shown in the following three figures. (b): A theoretical micromotion minimum plane for one laser. To find this plane we move along its normal vector until we hit it. (c): A plane for laser 2 intersects plane 1. Now we need only move along the intersection line to find the global minimum. (d): The intersection of the planes of all three lasers. See the text for more details.

6.2.4 Sources of stray fields after loading

We would like to know whether build-up of stray charge on the fibres is a significant factor in the displacement of the ion from the trap rf-minimum. By tracking the changing micromotion compensation voltages over time, we can determine from which direction the stray charge originates.

78

We consider first the micromotion compensation in the axial direction, and let the steady-state compensation voltage be w_0 . Immediately after loading the trap, a stray field appears, which decays exponentially and consequently requires an exponentially decreasing compensation voltage. The top and bottom electrodes will require not only a different compensation voltage, but a voltage of opposite sign. The total voltage required is then

$$w(t) = w_0 + w_+ e^{-\gamma_+ t} - w_- e^{-\gamma_- t}.$$
(6.15)

The parameters are all correlated, and so can't be obtained by simply fitting to this function. To allow us to extract useful fit parameters we make the following definitions:

$$w_{\pm} = \bar{w} \pm \Delta w \tag{6.16}$$

$$\gamma_{\pm} = \bar{\gamma} \pm \delta. \tag{6.17}$$

We substitute these definitions into equation 6.15 and in the approximation that $\delta t \ll 1$ we have

$$w(t) \approx w_0 + 2[\bar{w}\delta t + \Delta w]e^{-\bar{\gamma}t}.$$
(6.18)

We see from this that if w_+ and w_- were the same, then $\Delta w = 0$, and if the decay rates were the same, then $\delta = 0$ and $w(t) = w_0$. However, if there is a difference in the quantity of charge on the top and bottom electrodes ($\Delta w \neq 0$), or a difference in the decay rate ($\delta \neq 0$), then exponentially changing compensation voltages must be applied. A fit to equation 6.18 cannot give us \bar{w} , and so we cannot know the absolute amount of charge on the electrodes. However, we are more interested in the direction of the stray fields, so we proceed towards that end.

There are four independent fitting parameters, making the fitting function

$$w(t) = p_1 + [p_2 t + p_3]e^{-t/p_4}.$$
(6.19)

For this model, we assume that the decay parameter γ is the same for all micromotion compensation components, making three fit parameters for each compensation electrode, p_1, p_2, p_3 and one global parameter p_4 .

We tracked the change in the three micromotion compensation voltages over time. Figure 6.9 shows one such measurement, together with a fit to the model described in this section. In order to extract the direction of the stray field source we first make a 3D finite-element calculation of the electric field at the trap centre in response to voltages applied in the direction of the four dc electrodes, which are the wire electrode in the +x-direction, the oven in the +y-direction, the top ground electrode in the +z-direction, and the bottom ground electrode in the -z-direction. This gives the following result:

$$\begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix} = \begin{pmatrix} 0.0 & 0.201 & 0.0 \\ 0.684 & 0.0 & -2.45 \times 10^{-3} \\ 0.0 & -3.513 \times 10^{-3} & 0.375 \end{pmatrix} \begin{pmatrix} u \\ v \\ w \end{pmatrix} \text{ mm}^{-1}.$$
(6.20)



79

Figure 6.9: A log of the variation in the three micromotion compensation voltages over time, together with a fit to equation 6.19

At time t = 0, we have a steady-state field, given by u_0 , v_0 and w_0 , and also the decaying compensation field given by Δu , Δv , and Δz . We can use the latter, with equations 6.19 and 6.20 to calculate the angles of the source (θ_0, ϕ_0). At longer times, the source is given using $\bar{u}, \bar{v}, \bar{w}$, which is given by parameter p_2 in equation 6.19.

The long-term measurements shown in figure 6.9 were repeated to give six datasets, and the angles calculated for each set. The azimuthal angle ϕ is consistently in the direction of the oven, with only 3 degrees variation, whilst the angle θ varies slightly more, but is always in the direction of the surface of the inner electrodes. None of the datasets give an angle within that subtended by the fibre. If we assume that the source lies on the end facet of the inner electrode, then we can calculate the position of the source, and this is shown in figure 6.10.



80

Figure 6.10: The calculated positions of origin of the stray charge, based on the model in this section.

6.3 Magnetic field compensation

Many of the measurements reported in this thesis have some dependence on the magnetic field at the ion, and so it is desirable that the field is of a known magnitude and direction. We have used the set of coils described in 5.7 on page 63 to first null the Earth's $\approx 50 \ \mu$ T field, and then provide a known bias field. In addition, at very low magnetic fields, the axis of atomic polarisation of the ion is defined by the polarisation vector of the laser—the electric dipole moment of the ion is aligned with the polarisation of the laser. The result is shown in figure 6.11: On the 397 nm S_{1/2} \leftrightarrow P_{1/2} cooling transition we see that linearly polarised light can always induce a $\Delta m_j = 0$ transition, so we need only make sure our light is linearly polarised. On the D_{3/2} \leftrightarrow P_{1/2} repumper transition, however, we see that whatever polarisation of light we use, the P_{1/2} state can always decay to a state from which it cannot be pumped. That is, the $m_j = \pm 1/2$ state for σ_{\pm} -polarised light and the $m_j = \pm 3/2$ state for π -polarised light. This phenomenon is termed 'optical pumping' and the resultant state, where it is an angular momentum eigenstate as shown here, is termed a 'dark state'.

Light of an arbitrary polarisation can always be described as a supersposition of π and σ polarisations, and so for any polarisation of light there will always be some corresponding superposition of D_{3/2} sublevels which will produce a dark state in the absence of a magentic field.

If we introduce a magnetic field, the energy of each sublevel changes, and since the time dependence changes with the energy, we find that the dark state does not persist—even if the ion, at some time, is in some superposition of states that would have constituted a dark state, the time evolution of each individual state means that the superposition is quickly destroyed. The theory of this destabilisation of the dark states by shifting the energies of the sublevels is laid out in [90], and this is the method we use in this work.

The method of shifting the energy levels only works when the laser is in a superposition of polarisations. If its polarisation is a pure state, that is, π , σ^{\pm} , then the dark state is one single sublevel, not a superposition of m_i states, and its time-dependence has no effect.



Figure 6.11: Optical pumping to dark states in ⁴⁰Ca⁺. In very low magnetic fields the axis of atomic polarisation of the ion is aligned along the laser's polarisation. In (a) this is not a problem provided our light is linearly polarised, but if it is circularly polarised it will eventually land in a 'dark state', and fluorescence will stop. Of more concern is the repumper transition shown in (b)—here, any polarisation of light will eventually put the ion into a dark state.

The procedure for nulling the background field is as follows:

- 1. Align the repumper polarisation in some specific direction, call it x.
- 2. Scan the magnetic field y and z components, (which simply means scanning the current in the coils), until a dip in fluorescence is noted. At the fluorescence minimum we can say we have nulled the field for the y and z components. The field now only has an x-component. Figures 6.12(a) and 6.12(b) show a typical iteration of this process.
- 3. Rotate the repumper polarisation by $\pi/2$. Now we only need to scan along the *x*-direction since we know the field only has this component. At the fluorescence minimum the field is minimised. Figure 6.12(c) shows this process, with a final fluorecence rate very close to the background count rate for the PMT used.
- 4. Note that we can never entirely null the field because the polarisation of the repumper is never entirely pure.

For our experiments, we perform several iterations of these scans before any critical measurement, to find the global minimum, but we find that it does not change significantly, day-to-day. We find that by nulling the field in three dimensions we are able to reduce the ion's fluorescence to approximately the background count rate for the free-space PMT located at the side of the trap, that is, a rate of 4–5 kHz. Once this is done we can easily apply a bias magnetic field of a known amplitude and direction.



Figure 6.12: Compensating the Earth's magnetic field in three dimensions. First we scan the current in one horizontal set of coils until the fluorescence minimum is found. At this point the field from the coils is approximately cancelling the Earth's field in that direction. We repeat for the second set of horizontal coils and then for the vertical coils until we have a global fluorescence minimum where the Earth's field is approximately compensated. Note that the values of the magnetic field in the figures are derived from the linear scaling of current and magnetic field, which we know for each coil.

6.4 Secular frequency measurements

To determine the secular frequency of a single trapped ion we apply an additional sinusoidal rf voltage to the compensation electrodes and scan its frequency over a range within which we expect the ion's secular frequency to lie. The principle of the method is that an additional electric field oscillating close to the resonant frequency of the ion in a particular direction will transfer energy to the ion, exciting the amplitude of its oscillations. This heating effect results in a reduced fluorescence. In our trap we have radial symmetry, and so we measure two clear secular frequencies: radial and axial⁶. We vary the amplitude of the trap rf drive—the *q*-parameter of equation 2.10 on page 11—and also the *a*-parameter of section 6.2 (recall that this gives a dc offset to the top and

⁶A splitting of the radial frequency into two components is also visible. We were able to directly observe the ion oscillating in different radial directions as the rf voltage was scanned over the two peaks. The breaking of the rotational symmetry of the trap here probably results from the influence of the oven collimator.

bottom outer ground electrodes). We compare the results with secular frequencies derived from finite-element simulations of the trap for a variety of rf amplitudes and similar dc offsets.

We used a TTi TGR2050 sythesised rf generator to provide the additional sinusoidal voltage, controlled by LabView via RS232.

First, the dc offset on the ground electrodes was set to as close to zero as possible. This corresponds to an *a*-value of zero. Then, the following steps were taken:

- 1. Connect the additional rf voltage to the side compensation electrode. This will preferentially excite the radial secular frequency.
- 2. Compensate micromotion in 3-D.
- 3. Scan the additional rf voltage over a wide frequency range to approximately locate the resonance. Try a variety of amplitudes—too much and the ion will be heated out of the trap, too little and the dip in resonance will not be visible.
- 4. Zoom in on the resonance for a more detailed scan.
- 5. Repeat the detailed scan at 6 different main trap rf amplitudes from -26 dBm to -21 dBm to see how the secular frequency shifts.
- 6. Repeat steps 2–5 for a = -2, -1, 1, 2.
- 7. Repeat steps 2–6 with the additional rf voltage applied to the upper ground electrode. This will excite mainly the axial secular frequency.

Figure 6.13 shows a typical measurement of the secular frequencies. In this figure we scan upwards over the resonance. This process was repeated for the range of parameters given above. Figure 6.14 presents the range of secular frequencies calculated from simulations alongside the experimental data. The experimental data have been fitted to equation 2.28 on page 14. The simulations cover a wider range of rf amplitudes than are available to us experimentally. In day-to-day operation of the trap, the rf amplitude was set at around 0.025 V and the *a*-value was zero. By comparing these figures it is possible to say that the rf amplitude in the endcap trap, for the signal generator level generally used, was approximately 200 V, and the corresponding *q*-value of the trap was approximately 0.5–0.7. We note also a similar functional dependency in the simulation and experimental results.



Figure 6.13: Resonantly exciting the ion's radial and axial secular frequencies with an additional rf voltage on (a): the side electrode, and (b): the upper electrode. We interpret the double radial peak as a splitting of the radial frequency into two components. See the text for details. For these plots the trap rf amplitude was 22 dBm at the first function generator, and the *a*-value was 1.



Figure 6.14: A comparison between (a) the secular frequencies derived from simulations and (b) the measured values. Error bars can be determined from the width of the resonances. Since these are on the order of 10 kHz they are not shown in this plot. See the text for further details.

Chapter 7

A fibre-coupled source of single-photons

This chapter presents the main results of the experiment, that is, the use of the endcap trap with integrated fibres as a single photon source. Two main experiments were performed; the first to excite a trapped ion with continuous light and observe the quantum nature of the fluorescence collected through the fibres; the second to excite the ion with a series of laser pulses, such that one photon is emitted for each pulse sequence. In both cases, the quantum nature of the light is strikingly shown in correlation spectra of photon arrival times from the fibres.

As discussed in chapter 1, the controlled creation and absorption of single photons enables schemes in quantum networking, cryptography, communication and computing. Prior to the development of quantum optics, single photons were produced using a laser beam attenuated to such an extent that the probability of more that one photon being emitted was negligible. This was unsatisfactory for two reasons—firstly, the probability of no photon being emitted was higher than that of one, and secondly the probability of two photons being emitted together was non-zero.

In the new field, then, of 'real' single-photon sources, we can identify two overarching schemes. The first—coherent emitter/photon interaction through strong coupling cavity QED—has been mentioned in the trapped-ion context in chapter 4. In this scheme, a single photon is emitted into a single cavity mode, providing a reversible coherent exchange of excitation between emitter and photon, as required by many quantum information processing proposals (for example, [31]). Although strong-coupling has been achieved with an ensemble of ~ 500 ions in a linear trap [69], to date it has never been reached with a single trapped ion. At present, work towards this is a highly active research area—see, for example [91]. A 2010 review of the theory and most important implementations of cavity-based single photon schemes, (not only ion trap schemes), can be found in [92]. Neutral atoms have also been used as a cavity-based single-photon source [93]. In this thesis however we are concerned only with ion trap schemes and we can identify the work by Keller et al. [54, 73] in 2004 to implement an ion-trap single photon scheme with calcium ions as a milestone in the field. The experiments in this thesis have several similarities to this work. The 2004 experiment was the first to couple a single trapped ion to an optical cavity, and furthermore the first to produce single photons from such a scheme. The ultimate goals of the experiments reported in this thesis are to be the first single ion coupled to a *fibre* cavity, and from it to produce single photons on demand. Further, due to the miniaturised cavity¹, to *strongly* couple the ion to the cavity field.

¹The cavity used in the 2004 work was 8 mm long.

In contrast to cavity-assisted schemes, the second distinct single-photon generation system is to use an emitter in free space. Some implementations of this include single ions, where this was first demonstrated by Diedrich et al. in 1987 [94], single molecules [95], nitrogen-vacancies in diamond (colour centres) [96] and quantum dots [97]. In the context of trapped ions, the atom emits into many different modes and as a result the amount of photons that can be captured by a collection system of some fixed numerical aperture is limited, although the efficiency of single-photon detection in such a scheme can approach that of the cavity-assisted scheme[98]. Review articles covering many single-photon generation schemes can be found in [99]. In this chapter the results of two non-cavity single-photon generation experiments are presented, as a prototype for future cavity-assisted schemes.

7.1 Theory of the measurements

7.1.1 Classical theory of intensity fluctuations

The aim of this chapter is to show the quantum nature of the light emitted by our single-photon source. That is, to prove the 'singleness' of the photons. We start, then, by outlining the classical theory of optical intensity fluctuations, and show that a quantum-mechanical reformulation of the theory predicts different results at the single-photon level—results which we can directly observe. The results of our experiment are no less than a proof of this quantum theory.

We consider first a chaotic light source. Such light is termed 'chaotic' because the contributions of the random phases emitted by the atoms of the source produces interference which results in random intensity peaks and troughs. See, for example, the random-walk model of collision-broadening in chaotic light, described in [59, Ch. 3]. It can be shown that the mean-square intensity of such a beam is given by

$$\langle \bar{I}(t)^2 \rangle = \left(2 - \frac{1}{\nu}\right) \bar{I}^2,\tag{7.1}$$

with v the number of radiating atoms and \overline{I} the mean intensity over long-times (times much greater than the coherence time), given by

$$\bar{I} = \frac{1}{2} \varepsilon_0 c E_0^2 v. \tag{7.2}$$

In this and hereafter the overbar denotes the *cycle average* intensity, that is, the intensity averaged over one cycle of the electric field, whilst the angle brackets denote the average of many values of the cycle-averaged intensity, taken over times much greater than the coherence time (the time-averaged intensity).

For a large number of atoms, equation 7.1 approximates to

$$\langle \bar{I}(t)^2 \rangle = 2\bar{I}^2. \tag{7.3}$$

By contrast, coherent light exhibits no intensity fluctuations:

$$\langle \bar{I}(t)^2 \rangle = \bar{I}^2. \tag{7.4}$$

7.1.1.1 Second-Order Coherence

If we now consider *pairs* of intensity fluctuations of the sort described in section 7.1.1, but with a fixed time delay τ between them. We define the **intensity correlation function**, $g^{(2)}(\tau)$, of the light as

$$g^{(2)}(\tau) = \langle \bar{I}(t)\bar{I}(t+\tau)\rangle, \tag{7.5}$$

and generally we use a normalised version:

$$g^{(2)}(\tau) = \frac{\langle \bar{I}(t)\bar{I}(t+\tau)\rangle}{\bar{I}^2}.$$
(7.6)

which we can write in terms of the electric field as

$$g^{(2)}(\tau) = \frac{\langle E^*(t)E^*(t+\tau)E(t+\tau)E(t)\rangle}{\langle E^*(t)E(t)\rangle^2}.$$
(7.7)

We now consider the values of $g^{(2)}(\tau)$ we might expect at $\tau = 0$ and $\tau > 0$. If we take a pair of measurements of intensity, at times t_1 and t_2 , then the following inequality must be true:

$$2\bar{I}(t_1)\bar{I}(t_2) \le \bar{I}(t_1)^2 + \bar{I}(t_2)^2.$$
(7.8)

Now, considering the statistical average of the two intensities, we find

$$\left[\frac{\bar{I}(t_1) + \bar{I}(t_2)}{2}\right]^2 \le \frac{\bar{I}(t_1)^2 + \bar{I}(t_2)^2}{2},\tag{7.9}$$

where we have replaced the cross terms in the left hand side of equation 7.9 with the inequality 7.8. So we have that

$$\langle \bar{I}(t) \rangle^2 \le \langle \bar{I}(t)^2 \rangle, \tag{7.10}$$

and since $\bar{I}^2 = \langle \bar{I}(t)^2 \rangle$, then from equation 7.6 we can see that for zero time delay, $\tau = 0$, classical light gives

$$g^{(2)}(0) = \frac{\langle \bar{I}(t)^2 \rangle}{\langle \bar{I}(t) \rangle^2}$$

$$\Rightarrow g^{(2)}(0) \ge 1.$$
(7.11)

In addition, we can write

$$\left(\sum_{i=1}^{N} \bar{I}(t_i)\bar{I}(t_i+\tau)\right)^2 \le \sum_{i=1}^{N} \bar{I}(t_i)^2 \sum_{i=1}^{N} \bar{I}(t_i+\tau)^2,$$
(7.12)

where the cross terms on the left hand side have again been replaced by the inequality 7.8. We can see that if the t_N are sufficiently large, then the two sums on the right hand side of equation 7.12 are equal. Taking a statistical average again, as in equation 7.9, we find

$$\langle \bar{I}(t)\bar{I}(t+\tau)\rangle \le \langle \bar{I}(t)^2\rangle,$$
(7.13)

and referring to equation 7.6, this means that

$$g^{(2)}(\tau) \le g^{(2)}(0). \tag{7.14}$$

It can be shown [59, Ch. 3] that for Gaussian-Lorentzian light²

$$g^{(2)}(\tau) = 1 + \exp\left[-\frac{2|\tau|}{\tau_0}\right],$$
 (7.15)

with τ_0 the radiative lifetime of the spectral transition, or the collision time, depending on the dominant broadening mechanism.

Similarly, for Gaussian-Gaussian light, such as Doppler broadened light:

$$g^{(2)}(\tau) = 1 + \exp\left[-\pi \left(\frac{\tau}{\tau_c}\right)^2\right],\tag{7.16}$$

with τ_c the coherence time.

It follows from equations 7.15 and 7.16 that

$$g^{(2)}(\tau) \to 1, \ \tau \gg \tau_C, \tag{7.17}$$

and

$$g^{(2)}(0) = 2, (7.18)$$

and these relations apply to all forms of chaotic light.

For coherent light, which exhibits no intensity fluctuations, we find from equation 7.6 that

$$g^{(2)}(\tau) = 1, \tag{7.19}$$

for all τ .

In equation 7.7 we have assumed a stationary beam of light³, with fixed linear polarisation, measured at a fixed point; however, it can be extended to cover non-stationary light fields measured at two points in the following way:

$$g^{(2)}(\mathbf{r}_{1},t_{1},\mathbf{r}_{2},t_{2}) = \frac{\langle E^{*}(\mathbf{r}_{1},t_{1})E^{*}(\mathbf{r}_{2},t_{2})E(\mathbf{r}_{2},t_{2})E(\mathbf{r}_{1},t_{1})\rangle}{\langle |E(\mathbf{r}_{1},t_{1})|^{2}\rangle\langle |E(\mathbf{r}_{2},t_{2})|^{2}\rangle},$$
(7.20)

where for non-stationary beams, the two averages in the denominator could be different. With the linearly-polarised, parallel beams we assume, we need only use one component of \mathbf{r} , say z. We

²Gaussian-Lorentzian means that the distribution of the electric field amplitude is Gaussian, while the spread of frequencies is Lorentzian, such as in homogeneously collision-broadened chaotic light. Conversely, Gaussian-Gaussian light means that the broadening process produces a Gaussian frequency spread, such as in inhomogeneous Doppler broadening.

³ 'Stationary' here means that the processes controlling the intensity fluctuations are constant in time, that is, the photon statistics do not change. It doesn't matter when we start measuring, the correlations only depend on the time difference $t_2 - t_1$. This contrasts with non-stationary sources, such as pulsed light, where the correlations would depend on exactly where in the pulse we start measuring.

can convert a measurement at two points (z_1, t_1) and (z_2, t_2) to a measurement at one point but two times, to produce a correlation as in equation 7.7, by redefining τ as

$$\tau = t_2 - t_1 - (z_2 - z_1)/c, \tag{7.21}$$

and we can then write

$$g^{(2)}(z_1, t_1, z_2, t_2) = g^{(2)}(\tau), \tag{7.22}$$

and all the relations for a single measurement point also apply to two separate measurement points.

7.1.1.2 The Hanbury Brown-Twiss Experiment

In 1956, two astronomers, Robert Hanbury Brown and Richard Twiss, published the results of an experiment to test the time-dependent intensity correlations in the chaotic light emitted by a mercury lamp [100–102]. They had previously used their technique to measure stellar diameters [103]. Figure 7.1 shows a simplified schematic of their setup. The principle of the experiment is that if the light falling on the two PMTs is coherent, then the output intensity fluctuations will be correlated. Hanbury Brown and Twiss used a zero time delay, but changed the distance between the PMTs. As the distance increased, the correlations dropped off, reflecting the decreasing spatial coherence of the two light sources. In our experiments we are more concerned with the temporal coherence of the light, that is, how the intensity correlations behave at difference time delays. For example, if we set $\tau = 0$, then the output of the experiment becomes

$$\langle I(t)I(t+\tau)\rangle = \langle I(t)^2\rangle, \tag{7.23}$$

which gives a non-zero result due to the random fluctuations of the light. If we make τ much greater than the coherence time then the intensity fluctuations will be uncorrelated and $\langle I(t)I(t+\tau)\rangle$ averages to zero. The drop-off in correlations in this case can be used to measure the coherence time of the light. The principle of observing intensity correlations as a function of time delay is the main theme of the experiments described in this chapter. Although Hanbury Brown and Twiss recorded analogue correlations between continuous photocurrents, later experiments recorded coincidences between discrete pulses in separate detectors [104].



Figure 7.1: The main components of the Hanbury Brown-Twiss experiment to measure time-dependent intensity correlations in light from a filtered mercury lamp. The photocurrents, $i_{1,2}$ (not to be confused with the intensity *I* discussed elsewhere in this chapter), measured by the photomultiplier tubes (PMTs) were AC-coupled, with the signal from one being delayed by an amount τ . The resultant fluctuations were multiplied and averaged. Hanbury Brown and Twiss used this apparatus to measure the spatial coherence of the light by placing PMT1 on a translation stage. This way the correlations could be observed as the distance between the paths was changed. They used this to measure the spatial coherence of the light.

7.1.2 Quantum theory of second-order coherence

The preceeding discussion of the classical degree of second-order coherence works well for the ordinary light intensities of the macroscopic world; however, for quantum correlations between single photons emitted from single ions we must formulate the theory in fully quantum-mechanical terms. This was first laid out in 1963 by Roy Glauber [105] and a summary of the main points is given here. Figure 7.2 shows a typical setup for investigating correlations in single photon arrival times, that is, a quantum version of the Hanbury Brown-Twiss experiment. The source light is split by a beam splitter. Photons arriving at detector D3 trigger a start event in the timer, while photons arriving in detector D4 trigger a stop event. After many events a histogram of delay times between start and stop events can be built up. In this case, since intensity is proportional to the number of counts detected, $g^{(2)}(\tau)$ can be written as

$$g^{(2)}(\tau) = \frac{\langle n_3(t)n_4(t+\tau)\rangle}{\langle n_3(t)\rangle\langle n_4(t+\tau)\rangle},\tag{7.24}$$

with $n_3(t)$ and $n_4(t)$ the number of photons arriving at detector 3 and 4 at time t.

We can write this in quantum-mechanical terms in the following way.

Recall the photon number operator \hat{n} of section 4.2 on page 30. Using equation 4.22, and re-ordering the operators into *normal order*,⁴ it is possible to write equation 7.24 in terms of the creation and annihilation operators, as

⁴Placing the operators with all the creation operators to the left and all the annihilation operators to the right is called *normal ordering*.



Figure 7.2: A typical Hanbury Brown-Twiss setup for measuring correlations in arrival times of photons from a source. In an ideal setup we would only need one detector which would measure each photon's arrival time accurately. In practice however detectors have a *dead time* after measuring an event, during which they cannot measure anything. To compensate for this, two detectors are used. This necessarily introduces the added complication of the beam splitter. The fields E1-E4 are the electric fields at the four ports of the beam splitter, with E1 the source, E2 the vacuum, E3 and E4 the beamsplitter outputs.

$$g^{(2)}(\tau) = \frac{\langle \hat{a}_{3}^{\mathsf{T}}(t) \hat{a}_{4}^{\mathsf{T}}(t+\tau) \hat{a}_{4}(t+\tau) \hat{a}_{3}(t) \rangle}{\langle \hat{a}_{3}^{\mathsf{T}}(t) \hat{a}_{3}(t) \rangle \langle \hat{a}_{4}^{\mathsf{T}}(t+\tau) \hat{a}_{4}(t+\tau)}.$$
(7.25)

Setting $\tau = 0$, and for a photon number state $|n\rangle$ this equation can be simplified to [106, Ch. 8]:

$$g^{(t)}(\tau) = \frac{\langle \hat{n}_1(\hat{n}_1 - 1) \rangle}{\langle \hat{n}_1 \rangle^2}.$$
(7.26)

This equation holds for any input field, but in the case that the input field is a photon number state $|n\rangle$ then we have

$$g^{(2)}(0) = \frac{n(n-1)}{n^2}.$$
(7.27)

So, for a source emitting single photons we obtain the non-classical result, $g^{(2)}(0) = 0$.

It can be shown that the degree of second-order coherence for stationary chaotic light equations 7.15 and 7.16—and for coherent light—equation 7.19—is the same in the quantum picture as in the classical picture. In the classical picture this describes the intensity fluctuations of chaotic light, while in the quantum picture it describes the surprising phenomenon of *photon bunching*, that is, the tendency of photons to clump together. For a more detailed explanation of the quantum mechanics of the photon bunching effect, see [107].

7.1.3 Photon Antibunching

Light in which the inequality

$$g^{(2)}(0) < g^{(2)}(\tau) \tag{7.28}$$

holds, is known as *antibunched* light and implies that the probability of two photons occuring close to each other is less than the probability of them occuring at greater separations. This kind of light is in violation of the classical theory of intensity fluctuations and can be considered a signature of the quantum nature of light.

Antibunched light often exhibits sub-Poissonian photon statistics⁵, another signature of nonclassical light, but the two phenomena are distinct and can occur separately [108, 109]. In contrast to inequality 7.28, sub-Poissonian statistics satisfy

$$g^{(2)}(0) < 1. (7.29)$$

Antibunching was first observed in the resonance fluorescence of sodium atoms by Kimble et al. in 1977 [110], and later from a single sodium atom in 1978 [111]. The sub-Poissonian nature of the radiation in the above was shown by Short in 1983 [112]. In the experiments reported in this thesis, antibunching in the resonance fluorescence from a single calcium atom is observed. In this case it is clear that after a single atom has emitted a photon it must be re-excited before it can emit a second. The resultant delay between photons is the cause of the antibunching.



Figure 7.3: Photon streams corresponding to chaotic light (bunched photons), and non-classical light (antibunched) photons.

7.1.3.1 The second-order correlation function for a two-level atom

In the following sections we describe the theoretical $g^{(2)}(\tau)$ function for a two-level atom in two separate ways—firstly using an analytical expression, and secondly by numerically integrating the optical Bloch equations of section 3.2.2 on page 20. Carmichael and Walls [113] give an analytical expression for $g^{(2)}(\tau)$ in the case of light scattered at resonance from a two-level atom:

$$g^{(2)}(\tau) = 1 - \left(\cos\tilde{\Omega}\tau + \frac{3\Gamma}{4\tilde{\Omega}}\sin\tilde{\Omega}\tau\right)e^{-\frac{3}{4}\Gamma\tau},$$
(7.30)

where

⁵The arrival times of photons from a source is governed by probabilities related to the Poissonian distribution. Coherent light exhibits Poissonian photon statistics in which the square of its standard deviation $(\Delta n)^2$, equals its mean \bar{n} . Super-Poissonian and Sub-Poissonian statistics, on the other hand, have the properties $(\Delta n)^2 > \bar{n}$, and $(\Delta n)^2 < \bar{n}$ respectively.

$$\tilde{\Omega} = \sqrt{\Omega^2 + \left(\frac{\Gamma}{4}\right)},\tag{7.31}$$

with Ω the Rabi frequency⁶, and Γ the spontaneous decay rate of the transition. It is clear that the presence of spontaneous decay places a damping exponential envelope on the Rabi oscillations. Rabi oscillations depend on the evolution of the superposition between the two energy levels, and a spontaneous decay event abruptly interrupts the phase of the superposition. This occurs in a random way, at a rate determined by Γ , the decay rate. There may be other damping processes present which interrupt the coherence of the Rabi oscillations such as collisions with other atoms, but in the single-atom case we assume that spontaneous emission dominates. If we were to increase the electric field of the exciting light, we would increase the Rabi frequency, and if this became much greater than the damping processes, then we would in principle start to see clear Rabi oscillations.

So much for the situation where the exciting light is exactly on-resonance with the atomic transition. For an arbitrary detuning, a more general form is given by Dagenais and Mandel [111] in their work on photon correlations from a single sodium atom:

$$g^{(2)}(\tau) = 1 - \left(\frac{A + \frac{1}{3}}{B^2 + 9A^2}\right) \frac{\left(\frac{1}{9} - \frac{4}{3}A + 4A^2 + \Theta^2\right)\left(\frac{1}{2}\Omega^2/\beta^2 + 1 + \Theta^2\right)}{\left(\frac{\Omega^2}{\beta^2}\right)\left(\frac{1}{2}A - \frac{1}{3}\right)} e^{\left(-\frac{4}{3} + 2A\right)\beta\tau} - \frac{4\left(\frac{1}{2}\Omega^2/\beta^2 + 1 + \Theta^2\right)}{\Omega^2/\beta^2} \left(\frac{FH + GJ}{H^2 + J^2}\cos\left(B\beta\tau\right) - \frac{GH - FJ}{H^2 + J^2}\sin\left(B\beta\tau\right)\right) e^{\left(-\frac{4}{3} + A\right)\beta\tau}, \quad (7.32)$$

with the following definitions:

$$\begin{split} \beta &= \frac{\Gamma}{2} \\ \Theta &= \frac{(\omega_L - \omega_0)}{\beta} \\ A &= \frac{1}{2}(\eta_+ + \eta_-) \\ B &= \frac{1}{2}(\eta_+ + \eta_-) \\ \eta_{\pm} &= \frac{1}{3} \left\{ \frac{9\Omega^2}{2\beta^2} - 9\Omega^2 - 1 \pm \sqrt{\left(\frac{9\Omega^2}{2\beta^2} - 9\Omega^2 - 1\right)^2 + \left(\frac{3\Omega^2}{\beta^2} + 3\Omega^2 - 1\right)^3} \right\}^{\frac{1}{3}} \\ F &= \frac{2}{27} + \frac{1}{3}A + \frac{2}{3}\Theta^2 - A\Theta^2 - A^3 + 3AB^2 \\ G &= -\frac{1}{3}B + B\Theta^2 - B^3 + 3A^2B \\ H &= 8AB^2 + \frac{8}{3}B^2 \\ J &= 8AB + 6A^2B - 2B^3. \end{split}$$
(7.33)

⁶See [59] or [106]. Recall that the Rabi frequency, giving the rate of population oscillation between the two electronic energy levels, depends on the electric field amplitude of the exciting light as $\Omega = |2\mu E_0/\hbar|$, with μ the dipole matrix element describing the strength of coupling between the atomic transition and the light.



(a) The $g^{(2)}(\tau)$ function for three intensities of light as represented by the Rabi frequency, Ω . Γ is the natural linewidth of the transition.



(b) The $g^{(2)}(\tau)$ function for three detunings, shown in terms of the natural linewidth of the transition. The light intensity is fixed at $\Omega = \Gamma$ in this case.

Figure 7.4: The $g^{(2)}(\tau)$ function as given by equation 7.32, shown for (a) three intensities of light, as represented by the Rabi frequency Ω and the natural linewidth Γ , and (b) three different detunings from the resonance frequency of the transition, shown again in terms of the natural linewidth Γ .

The effect of the light intensity and the detuning from resonance on the $g^{(2)}(\tau)$ function is shown in figure 7.4.

The optical Bloch equations presented in section 3.2.2 can be integrated numerically to show the time evolution of the upper state. To close this section, figure 7.5 shows the solution for the population of the upper and lower states of the optical Bloch equations for a two-level atom, by numerical integration between 0 and 100 ns. From this the normalised $g^{(2)}(\tau)$ function is shown for varying Rabi frequencies and detunings. Note that the normalisation prevents us from seeing the expected drop in amplitude of the Rabi oscillations with increasing detuning. These solutions are consistent with the plots obtained from the analytical expression 7.32.



Figure 7.5: The normalised $g^{(2)}(\tau)$ function calculated by numerical integration of the optical Bloch equations for a two-level atom. Shown as the Rabi frequency Ω and then the detuning Δ is varied.

7.2 Results—Continuous Single-Photon Generation

Previous implementations of the HBT photon-correlation experiment have used a setup similar to that shown in figure 7.2. See, for example, [91], for a typical setup—the light to be investigated is sent to one of two detectors using a beam splitter, or mirror. In our case, however, this splitting is not necessary, since photons from our light source can enter either the fibre in the top electrode, or the fibre in the bottom electrode. Since the ion must be re-excited before emitting again, light collected from the fibres should be anti-correlated. In addition, there is no need for further collection optics, since the ends of the fibres can be directly coupled to PMTs.
7.2.1 Spectroscopy with fibre-based detection

If we recall that the frequency of Rabi oscillations increases with detuning whilst the amplitude decreases, then it is apparent that the detuning must be accurately determined, to ensure that Rabi oscillations are visible. To this end the scanning cavity lock mentioned in section **??** on page **??** was used to scan the laser frequency over the resonance and so allowed us to choose an accurate detuning. 'Beam 1', indicated in figure 5.1 on page 37 was used. In order to reduce light scattered from surfaces inside the trap the beam was first 'cleaned' using a spatial filter and then focused using a diffraction-limited, custom-made lens⁷. The setup is shown in figure 7.6. Using this setup and with very low beam powers (hundreds of nanowatts) we were able to achieve count rates of up to 55 kHz from a single ion, combining both fibres, against background rates of under 1 kHz. Figure 7.7 shows a typical lineshape scan of the resonance using this beam and collecting the fluorescence through the fibres. The signal-to-background is shown for a detuning of 30 MHz.

By way of completeness, signal-to-backgrounds for the three main cooling beams are shown in figure 6.3 on page 69; however, these should not be compared with the low-scatter beam. Beam 1 is measured through the fibres, while the other cooling beams are measured at the free-space side PMT. Count rates for the main cooling beams through the fibres were extremely high (\sim MHz). To summarise the discussion of this beam, its principal advantages over the other cooling beams were:

- Spatial filtering to remove high-frequency noise from the beam.
- Tight focusing at the trap centre ($\sim 6 \,\mu m$ waist).
- High signal-to-background rate.
- Very low scatter rate from surfaces inside the trap.

⁷Manufactured by Bernhard Halle Nachfl. GmbH.



Figure 7.6: The final ultra-low-scatter setup for the principal beam used in the correlation experiments— 'beam 1'. See the text for more details.



Figure 7.7: Beam 1—the *y*-axis shows combined fluorescence from the top and bottom fibres. The halfwidth-at-half-maximum and the fluorescence at a detuning of 30 MHz are as shown. The scan was taken using the scanning cavity lock (see chapter 5). The power in the beam was 390 nW, and the 850 nm and 854 nm repumpers were used.

7.2.2 Saturation measurements of beam 1

In order to ascertain the optimum power to use in beam 1 we took a series of measurements of the half-width at half-maximum of the fluorescence spectrum. The object being to find the *saturation power* of the transition. Figure 7.8 shows a series of half-width at half-maximum linewidth measurements of beam 1. The data are fitted to the equation

$$\gamma' = \gamma \sqrt{\left(1 + 2\frac{P}{P_{\text{Sat}}}\right)},\tag{7.34}$$

where γ' is the linewidth of the transition, broadened by power *P*, γ is the natural linewidth of the transition⁸, and *P*_{Sat} is the saturation power of the transition.



Figure 7.8: Linewidth measurements of beam 1 with the fitted saturation power shown. For these measurements the 866 nm repumper was used, at a power of 520 μ W.

7.2.3 The correlation signal path

The two multimode fibres inside the trap⁹ are connected to the atmosphere via vacuum feedthroughs¹⁰ to further optical fibres¹¹, then to two PMTs¹². The general setup for a correlation measurement is shown in figure 7.9.

The photon signals from the PMTs are received by a FAST 7072T time-to-digital converter, which measures the the delay between a start and stop event at its inputs. This device has two channels and we set these up such that each start event is also a stop event, and *vice versa*. This means that every photon is correlated against all other photons. We set the 7072T to record time differences up to 200 ns, with a *conversion range* of 2048. This means that the 200 ns range is divided up into 2048 bins. We can specify the *time resolution* then, of the device in this state, as $(200 \times 10^{-9}/2048) \approx 98$ ps.

We require that the arrival times of photons recorded by the PMTs be as close as possible to the real photon arrival time (or at least to have a consistent offset); however, there is always some variation in recorded arrival times relative to the real photon time. This is quantified in the 'jitter'

⁸20.6 MHz/2, taking a value from [50].

⁹Thorlabs BFH48-200, core diameter 200 μ m, cladding diameter 230 μ m.

¹⁰Core diameter 400 μ m.

¹¹Edmund Optics NT58-457, core diameter 600 μ m, cladding diameter 660 μ m.

¹²In this particular experiment, these are Hamamatsu H5773's.



Figure 7.9: The setup for correlation measurements. The two fibres are coupled directly into the PMTs. The PMT signal passes to various signal processing elements and from there to a computer which records the arrival times of each photon event. The arrival times are processed later to obtain correlation spectra.

of the PMTs. Jitter is an undesirable effect whereby the time difference between a photon event and the resultant PMT pulse differs, or 'jitters', randomly, within some standard deviation. We can see the reason why we would care about jitter if we consider the arrival times of photons around $g^{(2)}(0)$. If the arrival times vary randomly, then the dip at zero time delay is 'smoothed out'. The effect of a number of dips, offset from each other by random amounts of time, at around the zero point, would have the effect of making the resultant dip seem higher.

The jitter of the PMTs was determined by observing time differences between picosecond laser pulses, 25 ns apart,¹³ and the resultant PMT pulses. A new PMT—Hamamatsu H5773, was found to have the lowest jitter, with a standard deviation of photon arrival times relative to the nearest picosecond pulse of 0.1 ns. Our previous PMTs—Hamamatsu H7360-02, had a standard deviation of 1.2 ns.

7.2.4 Continuous Correlation Measurements

7.2.4.1 Correlation measurements with a high intensity

The following continuous correlation measurements were made using the alternative repumper scheme of section 3.2.3 on page 23. A delay was inserted into one channel of the timing module, to

¹³Using a PicoQuant PDL 800-B picosecond pulsed diode laser.

give us access to negative correlations. Plots have been shifted so that the centre of the plot is at zero time. At high intensity, Rabi oscillations should be visible in the correlation spectrum, but it was found that the extra random correlations from the scattered background obscured any clear oscillations. An example of this is shown in figure 7.10. We subsequently concentrated on getting the dip at zero time as close to zero as possible, and these measurements are shown in the following sections. The fits to the following data are performed using a two-level model with the method discussed in section 3.2.3 on page 23.



Figure 7.10: Correlation measurements taken with a high intensity (770 nW) in beam 1. We hoped to see clear Rabi oscillations by using a high intensity beam, as indicated by the fit, but a combination of the high detuning and unwanted ion/background and background/background correlations have obscured them. The detuning in this measurement was fixed and the Rabi frequency is fitted to the two-level model of section 3.2.3.

7.2.4.2 Correlation measurements with a low intensity

In order to improve the quality of the single-photon source we now concentrate on getting the background scatter rate as low as possible. In practice this means using as low an intensity in beam 1 as possible, and keeping the laboratory almost entirely dark for the duration of the measurement. This reduces the unwanted ion/background and background/background photon correlations which are responsible for the non-zero value of $g^{(2)}(0)$. Figure 7.11 shows a 40 minute correlation measurement. For this measurement, we determined the signal-to-background ratios in the two fibre channels to be $SBR_1 = 75$ and $SBR_2 = 26$, respectively, with the difference accounted for by a small angle in beam 1 which caused slightly different amounts of light to be scattered into the upper and lower fibres. We can now estimate the value of $g^{(2)}(0)$ from the signal-to-background ratios, and compare it to the real value.

Estimating the value of $g^{(2)}(0)$

Let us temporarily describe a photon emanating from the ion as a 'real' photon event, and all other photons as 'background' events.

Let us say we have definitely received a correlation event at time zero. There are three physical ways this could happen:

- 1. A real photon in the first PMT, and a background photon in the second.
- 2. A real photon in the second PMT, and a background photon in the first.
- 3. A background photon in both PMTs.

If we say that a correlation event has definitely happened, and we know that both events cannot be real photons, then we can assign P_1 to the probability that a real event has occurred in the first PMT, and $1 - P_1$ to the probability that a background event has occurred in the second PMT, and *vice-versa*.

So we can write the probability of a correlation event at $\tau = 0$ as

$$g^{(2)}(0) = P_1(1 - P_2) + P_2(1 - P_1) + (1 - P_1)(1 - P_2),$$
(7.35)

where $P_1(1-P_2)$ is the probability of receiving a real photon in detector 1 and a background photon in detector 2; $P_2(1-P_1)$ is the probability of receiving a real photon in detector 2 and a background photon in detector 1, and $(1-P_1)(1-P_2)$ is the probability of receiving two background photons together.

If we write the probabilities P_1 and P_2 in terms of signal-to-background: $S_1 = P_1/(1-P_1)$ and $S_2 = P_2/(1-P_2)$, then we can write

$$g^{(2)}(0) = \frac{S_1 + S_2 + 1}{(S_1 + 1)(S_2 + 1)}.$$
(7.36)

If we use the signal-to-backgrounds, $SBR_1 = 75$ and $SBR_2 = 26$ of this measurement in equation 7.36, then we obtain a value of $g^{(2)}(0) \approx 0.05$, which agrees with the value of $g^{(2)}(0) = 0.05$ obtained by normalising the *y*-axis of figure 7.11, and reading off the minimum.

A note on the two-level fit

It is clear that the two-level fit deviates somewhat from the real behaviour of the correlations. One reason for this is that the two-level model does not take into account the full Zeeman structure of the levels. In essence, there are two Rabi frequencies at work in our effective two-level system here, because the S, and P-levels are both split into $m = \pm 1/2$ magnetic sub-levels. The two Rabi frequencies result from the differing Clebsch-Gordan coefficients for the different transitions between the S and P sub-levels, and this results in a deviation from the two-level fit.

The results of this experiment can also be found in [114].



Figure 7.11: Correlation measurements taken at a low beam intensity and very low fixed detuning. Again, the fit is to the two-level model of section 3.2.3. The fitted Rabi frequency, Ω , is as shown. The power in beam 1 was 160 nW. The minimum of the *y*-axis is 4.

7.3 Results—Pulsed Single-Photon Generation

Our second implementation of a single-photon setup was a scheme to generate single photons at regular intervals, defined by a controllable laser pulse sequence. In principle the scheme is a simple one, consisting of the following steps:

- 1. 'Shelve' the ion in the metastable D-state using a pulse of 397 nm light of pre-determined duration.
- 2. Inject a pulse of 866 nm light through one fibre—the ion emits a single photon and returns to the ground state.
- Detect the single photon, with some detection efficiency, via PMTs connected to the other fibre.
- 4. Re-cool the ion with 397 nm and 866 nm light for a period, before repeating the sequence.

7.3.1 Outline of the setup

Recall from section 3.2.3 on page 23 that we used an 850 nm and 854 nm repumper setup to avoid the dark resonance condition in continuous fluorescence that occurs when the detunings of the 397 nm and 866 nm lasers are equal. In the pulsed case, only the 866 nm laser is on during the measurement period; that is, the single-photon emission phase of the pulse sequence. Therefore, there is no possibility of a dark resonance and we can simplify the laser pulse sequence by only using an 866 nm repumper.

For this experiment, we switched to a new timing card, a FAST Comtec P7888. The first PMT is coupled directly to the bottom trap fibre using a custom-made light-tight aluminium cap. In our

efforts to create as far as possible an all-fibre system, we decided to inject 866 nm light and detect the resultant single photons through the fibres. Particular care was needed therefore to minimise 866 nm repumper light transmitted through the trap. The setup is illustrated in figure 7.12.



Figure 7.12: The setup for the pulsed correlation measurements.

7.3.2 The sensitivity of the PMTs to 866 nm light

The sensitivity of the PMT to 400 nm light is quoted in its specifications¹⁴, but we would also like to know its sensitivity to 866 nm light. Since we would be not only triggering, but also detecting, a single photon through the fibres, we particularly needed to know how much, if any, of the 866 nm trigger light would make it through the trap and into the single-photon-detection setup. We determined this by injecting 866 nm light into one fibre and extracting it through the other. This light was then coupled directly to the PMT and the measured counts recorded. We placed one, and then two, Semrock FF01-395 11 bandpass filters in a custom-designed light-tight cap attached to the PMT lens. The results are shown in figures 7.14(b) and 7.14(c). The background light offset that was present during the measurement has been subtracted from this measurement. From the linear fits shown we know that the sensitivity of the PMT is:

• No filter: $2.65 \times 10^{-3} \text{ s}^{-1} \text{pW}^{-1}$

- One filter: $4.52 \times 10^{-5} \text{ s}^{-1} \text{pW}^{-1}$
- Two filters: $2.57 \times 10^{-5} \text{ s}^{-1} \text{pW}^{-1}$



Figure 7.13: (a) Here we calibrate the AOM voltage used for the 866 nm beam entering the trap via the bottom fibre to the real beam power. The fourth-degree polynomial fit serves to guide the eye. The calibration is used in all subsequent measurements of this beam's power. In (b) we measure how much of the power entering the bottom fibre makes it out the top fibre. From a linear fit to the data we determine the transmission efficiency to be 3.45%. In both plots the background count rate parameter of the fit has been subtracted.

Based on this data, we decided to use two Semrock filters in the PMTs shown in figure 7.12.



Figure 7.14: (a) The PMT's sensitivity to 866 nm light with no filter. (b) The effect of one, and (c) two Semrock filters. Sensitivities calculated from the linear fits are 2650 Hz/ μ W 45.2 Hz/ μ W and 25.7 Hz/ μ W, respectively. Background light has been subtracted. Figure (d) shows the transmission of the filter used at approximately the wavelengths of interest. Transmission data provided by Semrock.

7.3.3 The pulse sequence and signal path

The laser pulse sequence was provided by an FPGA¹⁵, controlled via USB by a graphical Python user interface. Figure 7.15 shows the pulse sequence used with details of the pulse durations.

Each timestamp recorded by the P7888 card consists of 32 bits, or 4 bytes, and we want to ideally record 1 photon or 4 bytes per 4 μ s cycle. This is $250 \times 10^3 \times 4$ bytes $s^{-1} = 1$ Mb s^{-1} . This is 3.6 Gb per hour, just to record the single photons. To record *all* photon timestamps in the pulse sequence, then, would be restrictive, in terms of storage, and processing time. For this reason we gate the PMTs using a signal from the FPGA and so only record photons arriving within a specified time window. This window is labelled 'gate' in figure 7.15. The gating process is controlled by the signal path shown in figure 7.16. NIM level conversion modules are used to convert the TTL signal output by the PMTs and FPGA to the NIM signal required by the P7888 card.

¹⁵Field Programmable Gate Array. This FPGA is a Xilinx Spartan 3, device:XC3S400, package: PQ208.



Figure 7.15: The pulse sequence used in the pulsed single photon experiment. In this case the 397 nm 'shelving' pulse is on for 2500 ns before switching off, ensuring that the ion is shelved in the metastable D-state with high probabillity. At 2500 ns, the 866 nm repumper pulse switches on, generating the single photon. The third channel is the gate 'window', that is, the period during which we actually count photons. It should be noted that a delay of around 500 ns was present in the repumper pulse, introduced mainly by the AOM used, so in fact the gate window and repumper pulse start at roughly the same time. This pulse sequence was used for the correlation measurements shown in figure 7.19. The repetition rate is 250 kHz. The reference pulse shown in the lower plot is used later in the single photon pulse shaping experiment.



Figure 7.16: The signal path used for the pulsed photon correlation measurements. The NIM module on the left is a level adaptor. The P7888 counting card accepts NIM-level signals, so we first convert the PMT and the FPGA gate TTL signals to NIM. The module in the middle is not used in this measurement. The module on the right is a coincidence unit. It outputs a signal when a coincidence is satisfied between the FPGA gate signal and the PMT signal. In this way we only record counts inside the gate window.

7.3.3.1 The timing of the shelving pulse

In order to determine the effect of the position of the 397 nm shelving pulse in relation to the measurement window we looked at count rate in the window as a function of shelving pulse timing. We found that count rates approached a steady background after around 3000 ns. We also tested the effect of a 're-cooling' pulse placed towards the end of a pulse sequence, but found that this appeared to make little difference. We used the 'shelving efficiency' data in figure 7.18(a) to decide on a beam power of 15 μ W for our initial measurements.



Figure 7.17: Background count rates in the measurement window as we move the relative timing of the shelving pulse, and then the re-cooling pulse.



Figure 7.18: (a) We look at the correlation counts received at the P7888 card in a 3500 ns period after the onset of the shelving pulse. (b) We quantify the shelving efficiency in the time constant of an exponential fit to the data in (a). This quantifies the decline of the P-state population as the ion becomes shelved in the metastable D-state. At longer times we see less probability of finding the ion in the P-state, as expected.

7.3.4 Subtracting the background

Denoting a photon emitted from an ion as a 'real' photon, and all others as background events, we can formulate a more accurate method of background subtraction.

Let us now denote the total number of counts at the first PMT as $P_1(t)$, and at the second PMT as $P_2(t)$. We denote real photon events as $p_1(t)$ and $p_2(t)$, and the contribution from background or dark count events as $d_1(t)$ and $d_2(t)$. Then the total counts registered at the first PMT are given by $P_1(t) = p_1(t) + d_1(t)$, and at the second by $P_2(t) = p_2(t) = d_2(t)$. We want to know the cross-correlation of real photon events only, that is, $\langle p_1(t)p_2(t+\tau)\rangle$, and this is given by

$$\langle p_1(t)p_2(t+\tau) \rangle = \langle (P_1(t) - d_1(t+\tau))(P_2(t) - d_2(t+\tau)) \rangle$$

= $\langle P_1(t)P_2(t+\tau) \rangle - \langle d_1(t)P_2(t+\tau) \rangle - \langle P_1(t)d_2(t+\tau) \rangle + \langle d_1(t)d_2(t+\tau) \rangle.$ (7.37)

The background-contribution terms in equation 7.37 can be determined if we take a measurement of the background under the same circumstances used for the real measurement, that is, the same laser power, the same ambient lighting conditions, but with no ion, and then construct a histogram of cross-correlations of the background measurement with the real measurement. This will give us a histogram of

$$\langle d_1(t)P_2(t+\tau)\rangle - \langle P_1(t)d_2(t+\tau)\rangle + \langle d_1(t)d_2(t+\tau)\rangle$$
(7.38)

which we can then subtract from the main measurement. This technique is used to subtract the background from the measurement shown in section 7.3.5.

7.3.5 Correlation Measurements

Our first measurements demonstrated clear antibunching around the zero time delay, but a non-zero peak at zero time delay. An ideal single photon source must in principle be free of two-photon contamination. Our effort, therefore, moved in the direction of getting scattered background light to a minimum. We used the ultra-low scatter beam—beam 1—which we had employed for the continuous correlation measurements of section 7.2.4. Due to the scatter in the tail of our shelving pulse, we were unable to exceed 250 kHz in this experiment (considerable time and effort went into trying to reduce this, by reducing the fall time of the AOM used while monitoring correlations in the tail of the pulse). We were, however, able to reduce background correlations around zero time considerably. Figure 7.19(a) shows the final result. In figure 7.19 the background has been subtracted using the method of section 7.3.4. The mean number of correlations around the zero delay peak, that is, between -1.5 μ s and 1.5 μ s, (chosen arbitrarily), is -0.12. Moreover, if we normalise the plot using the average peak height, we obtain a value for the correlation function at zero time delay of 0.005. This indicates a high-quality single photon source, almost entirely free of multiple-photon events.



(a) Correlation spectra obtained using beam 1 and the beamsplitter setup illustrated in figure 7.12. The pulse sequence shown in figure 7.15 was used, with a repetition rate of 250 kHz. This measurement was taken over a period of 28 minutes. The rate on channel 1 was 147 Hz, and on channel 2 170 Hz. No background has been subtracted from this plot. The bin size was 100 ns.



(b) The correlation measurement of (a) with a background measurement taken immediately afterwards subtracted. The background was taken over 25 minutes, with beam 1 on, and had rates of 7.14 on channel 1 and 8.80 on channel 2. The background was shuffled, and then removed using the method detailed in section 7.3.4. The random nature of the background subtraction results in some small negativities. The bin size again was 100 ns. See the text for further details.

Figure 7.19

7.3.6 Single Photon Pulse Shape

In this section we reconstruct the shape of a single photon pulse by correlating photon arrival times with the reference pulse of the pulse sequence (channel four of figure 7.15). To reduce file sizes, we divided the 250 kHz reference pulse by 100, to 2500 Hz. Our count rates are much smaller than this, so this doesn't cause any problems. The pulse shape of single photons has been studied in detail in various atom-cavity contexts; see, for example [54, 73, 115, 116]. In creating a single-photon pulse shape, the principle is to repeat the experiment (that is, the generation of one single photon) many times, and record the arrival time in each case. When we record a photon's arrival time, that is, the time we receive a detector click, what we are actually detecting is the square modulus of its electric field. This is essentially the square modulus of the photon's wavefunction, which is the probability of the photon occurring at that time. If the photons are *indistinguishable*, that is, *identical*, then the distribution of arrival times can be said to represent the temporal shape of a single photon. That the photons should be identical is quite generally a requirement for quantum information processing applications. This is not the case in our experiment—we do not prepare our ion in a specific Zeeman sub-level, such that the output polarisation is defined, and so our photons are emitted with random polarisations. We can say that the pulse shape shown in figure 7.20 is a combination of pulse shapes resulting from the differing Rabi dynamics of each transition between Zeeman sub-levels, governed by the Clebsch-Gordan coefficients of the transitions.



Figure 7.20: The single photon pulse shape reconstructed by correlating the arrival times of photons with respect to the reference pulse. The \sim 500 ns delay before the photon pulse is attributed to the 866 nm AOM.

7.3.7 Single Photon Efficiency

For the measurement of figure 7.19, we repeat the pulse sequence at a rate of 250 kHz, and we ultimately detect rates of 147 Hz and 170 Hz in channel 1 and 2, respectively. Adding the channels, we detect 317 counts/s. If we assume that one photon is created for every pulse sequence, then we have a detection efficiency of 0.13%. We know, however, that the quantum efficiency of the PMTs

is 20%, so the efficiency before the PMTs works out to be 0.63%. If we subtract the background rates in channel 1 and 2 quoted in figure 7.19(b), we get a revised efficiency of 0.12% and 0.60% after and before the PMTs, respectively.

Starting with 250,000 photons per second, the various losses contributing to the final detection in this measurement are:

- 1. Fibre capture for one fibre (from the solid angle calculation of section 5.4.8.1 on page 55 = 0.03. 7500 photons captured.
- 2. Dichroic mirror = 0.86. 6450 photons passed.
- 3. Semrock filter = 0.93^2 (there are two). 5580 photons passed.

We attribute the subsequent drop to the \sim 1580 photons which finally make it to the PMTs to losses in the propagation between the inner and outer trap fibres, and losses in the beamsplitter.

Chapter 8

The Fibre Fabry-Pérot cavity

8.1 Laser machining the fibre facets

In 1994, it was found that normal ground glass optics could be polished by the application of a pulse of CO₂ laser light [117]. Glass, and in the case of optical fibres, fused silica, absorbs the $\sim 10 \ \mu m$ CO₂ radiation strongly, and a very thin layer of melted material is produced. The surface tension of the melt layer and the flowing material underneath it produces a smoothing effect, so that imperfections on a scale of the depth of the melt layer are removed. It can also cause a global contraction, producing a convex bulge. The technique has been used to repair damage in fused silica optics [118], and much earlier, in 1975, to produce a lens on the end of a fibre [119]. At much higher powers, material is evaporated away from the surface, and a depression is formed on the fibre facet. Our technique of laser machining uses the same technique as that of Hunger et al., of the Max-Planck Institut für Quantenoptik in Munich [120]. This section describes the process we used to machine depressions of specific radii of curvature on the end of multimode fibres. A variety of radii were manufactured, of around 200–600 μ m, based on calculations of the range of most suitable radii of curvature for a fibre cavity in our trap.

8.1.1 The CO₂ laser setup

The laser itself is a Synrad 48-1KWL water-cooled CO_2 laser, outputting up to 10 W of light. The setup is shown in figure 8.1. CO_2 lasers are widely commercially available in the high powers required to melt silica glass. Special optical elements are required for these powers; for example, instead of a normal optical isolator, the arrangement (marked 'isolator' in figure 8.1) is used. This consists of a polarisation-selective mirror, which passes vertical linear polarisation, followed by a reflective polarisation rotator, which acts like a quarter-wave plate for this wavelength, but only when it is tilted at 45° to the plane of the optical table. The light is now circularly polarised. Any reflections that make it back are converted to linear polarisation, but at 90° to the input polarisation, so the polarisation-selective mirror rejects the returning polarisation.

The isolator helps to avoid unwanted laser power fluctuations. After the isolator setup, the beam passes through a telescope, the purpose of which is to provide a waist (calculated to be 80 μ m) at the shutter. The shutter, visible between the two telescope lenses in figure 8.1 is actually a very



Figure 8.1: The CO₂ laser setup.

small loudspeaker with a piece of stainless steel glued to it¹. In 'closed' mode the piece of metal reflects the beam at 90° into a power meter. In 'shoot' mode, a short electrical pulse is sent to the speaker, causing it to move down and up, allowing the beam to pass for the duration of the electrical pulse. After this the beam hits a final focusing lens very close to the fibre tip. It is focused (to a waist of a calculated 20 μ m at the fibre) and aligned onto the fibre, which is held in a v-groove in a stainless steel platform mounted on a *xyz* translation stage. The second telescope lens is mounted on a *z*-stage, allowing us to alter the final focusing at the fibre tip.

For viewing the fibre tip we remove the final focusing lens. This is mounted on a magnetic stage so it can be easily lifted off and on. We replace it with a small mirror which reflects the image of the fibre tip into the vertically-mounted microscope. At the fibre end of the microscope is a $10 \times$

¹A cement glue which can withstand temperatures up to 800°C was used.

MPL APO objective lens. The microscope image is illuminated by a high-intensity white-light illuminator², and the images are recorded digitally by a Moticam 2500 USB camera mounted to the microscope.

8.1.2 The fibres—etching and cleaving

The fibres are Oxford Electronics HPSIR200CB multimode copper-coated fibres. Copper coated fibres were chosen for their low-outgassing properties under UHV.

The new copper-coated fibres have a core diameter of 200 μ m, cladding diameter 212 μ m and copper coating diameter 280 μ m. Since the electrodes (see figure 5.17 on page 54) are of inner diameter 254 μ m, we had to chemically etch away at least 13 μ m of copper coating around the fibres in order to fit them into the electrodes. The process is as follows:

- Etching crystals³ are dissolved in a beaker in the ratio 1.1 kg to 5 l.
- The beaker is placed on a magnetic stirrer⁴ and stirred at 200 RPM. It is heated to 50° C.
- When the etching crystals are fully dissolved, four fibres are lowered a distance of 6 cm into the solution for 2.5 minutes. This removes some of the copper coating. This etched section is the part that will fit tightly into the electrode.
- The fibre-holder is then raised by 16 mm and the remaining 4.4 cm of fibre is allowed to completely etch down to the bare fibre. This fully-etched section will be cleaved off—described below⁵.

As an aside, an obvious way to supply the rf to the trap electrodes is to use the copper coating. The skin depth of copper at 15 MHz is 17 μ m. The copper coating is 34 μ m thick around the cladding, and after partial etching, 24 μ m. So, this at least would not be an obstacle to such a scheme. In fact, the skin depth would only become an issue for these fibres at frequencies of less than 10 MHz.

²Edmund Optics MI-150.

³Farnell Fine Etch Crystals.

⁴Corning PC-420D.

⁵The cleaving process requires a certain length of bare fibre to operate.

 $250 \ \mu m$



(a) The fibres in the etching solution.

(b) The partially and fully-etched sections.



(c) The etched fibre being inserted into an electrode

Figure 8.2

Cleaving a fibre refers to the method of cutting the bare end of a fibre, such that a face is created perpendicular to the longitudinal axis of the fibre. This is not done by cutting in the conventional sense. Rather, a very sharp blade makes a tiny crack at one point in the fibre circumference. A tension is then applied longitudinally to the fibre, causing the crack to propagate through the fibre. The tension must be carefully selected-too low and the fibre may slip perpedicular to the blade,

damaging the blade and sometimes not cleaving the fibre. Too-low tension can also result in an excessive intrusion of the blade into the fibre—this is visible as a large chip in the fibre. On the other hand, if the tension is too high the fibre can 'spring back' from the cut and the resultant airborne specks of silica glass can land on the new fibre face. A too-high tension also results in a rough texture around the fibre facet known as 'hackle', caused by the fibre tearing apart too quickly. The hackle on a fibre facet is visible in the cleaved fibre in figure 8.3(a). The tension must be set to minimise hackle and blade intrusion. We found that for our fibres a tension of 5 N was optimal, although the cleaving process was not entirely predictable and sometimes more than one attempt would be required to obtain an acceptable cleave. There are other more obvious hazards to the newly-cleaved facet—water droplets and general dust can very quickly accumulate. Figure 8.3(c) shows water droplets on the facet of a newly-cleaved fibre.

In fact, however, the quality of the cleave is not of critical importance for our purposes, since most imperfections will be smoothed out by the shooting process later on.

We used a Nyfors Autocleaver LDF (large-diameter fibre) to cleave the long bare-fibre end of the etched fibres. The machine needs the length of bare fibre to 'hold onto' when it is applying tension.



(a) Hackle on a cleaved fibre facet.



(c) Water droplets on a cleaved fibre facet.

Figure 8.3: Three examples of fibre facets immediately after cleaving. In (a), hackle is visible around the top of the facet. This is caused by excess tension on the fibre during cleaving, causing it to tear apart too quickly. In (b) we see tiny specks of glass on the fibre facet. Again due to excess tension during cleaving. In (c), water droplets have immediately accumulated on the fibre facet. In all the fibres the chip from the cleaving blade is visible. The darker ring around the core is the cladding.

8.1.3 Shooting the fibres

In order to make sure the laser shot is aligned on the fibre tip we first shoot a piece of clear plastic with one high-power shot. By moving the microscope on a xyz-stage we position the centre of the shot at the centre of the computer screen. We then place a fibre on its xyz-platform, and move it so that its image is in the centre of the screen. Now we know that the fibre is at the position of the shot, and we can proceed.

There is some mechanical ringing when the loudspeaker shutter is switched, lasting for around 20 ms. To circumvent this, the laser switch-on is delayed by 20 ms from the shutter opening. The

119

laser and shutter are controlled by a single LabView vi, from which we also control the length of the laser pulse. Typically, we used a shot of 60 ms for our fibres, and a laser power of 1.2 W.



Figure 8.4: The laser-shutter pulse sequence. To avoid any influence of the mechanical movement of the shutter as it opens and closes, the laser is switched on at a predefined delay from the shutter.

In contrast to the single-shot process used in [120], we shot the fibre many times, rotating it by various degrees with each shot⁶. This method gave us not only fine control over the final radius of curvature of the shot, but also smoothed out any misalignment of the beam relative to the fibre face, and any ellipticity or astigmatism of the beam, ensuring that the depression in the central area was as symmetrical as possible. We also shot several flat fibres. For this purpose, we shot the fibre once, at low power, to clean and polish the surface. This gave the approximately flat surface shown in figure 8.9(a). Figure 8.5 shows a typical fibre-shooting process consisting of 37 shots. The number of shots was varied depending on the radius of curvature we required. The rotation is done by hand—the fibre is clamped several centimetres from the tip in a v-groove mounted to a manual rotator. The fibre is kept secured to its platform with two small magnets. We shoot the fibre at one angle, then remove the magnets and rotate to the next angle where we replace the magnets and shoot again. This manual procedure causes the fibre tip to misalign, with respect to the CO₂ beam, and so it must be optically inspected and returned to the correct position after each rotation.

⁶Hunger et al. shoot structures often smaller than ours; that is, the 1/e diameter of their single shot was 10–45 μ m, as compared to our average diameter, which was 89 μ m. They also used very short cavities, around a factor of ten smaller than the lengths we would probably consider. It seems that ellipticity concerns were less important for them in this parameter regime.



(a) After one shot.

(b) After four shots at 0° , 90° , 180° , 270° .



(c) After a further 4×4 shots at 45° , 135° , 225° , 315° . (d) After a further 4×4 shots at 22.5° , 112.5° , 202.5° , 292.5° .

Figure 8.5: Shooting the fibre shown in figure 8.3(c) 37 times at various different rotations. The dark ring around the fibre is not the cladding, rather it is a kind of 'lip' around the rim of the fibre—a chamfer, caused by melting. Note that the single shot of figure (a) has removed all the water droplets visible in figure 8.3(c).

8.1.4 Analysing the shot

In order to calculate the various parameters of a cavity, we need to know the radius of curvature and quality of the mirrors. This section describes the process of determining the radius of curvature of the surface using white-light interferometry, and gives the distribution of radii of curvature we were able to produce in a batch of 48 fibres. The methods used to extract a detailed phase map of the surface are detailed in [121–124], with [124] providing a general description of the process.

8.1.4.1 White-light interferometry

The first step in mapping the height-variation over the fibre surface is to produce an interferometric image of the facet. We use *white-light interferometry* for this purpose. A white-light interferometer is defined by the following conditions—the position of the zeroth-order interference fringe, and the spacing of the fringes, is independent of wavelength. It differs from a laser interferometer in that the path lengths of the two interfering beams need to be very accurately matched. This is not the case with lasers, because of their much longer coherence lengths. With short coherence-length white light, high-contrast fringes only appear when the paths are close-to-perfectly matched. With laser light, in a sense they appear too easily—unwanted laser reflections, for example from optics or glass surfaces in the imaging system, can produce fringe patterns almost indistinguishable from the pattern resulting from the sample surface. When the sample is highly-reflective, this may not be a problem, but in our case the sample is essentially just glass, similar to many other optical elements in the imaging path. It would be very difficult, therefore, with a laser interferometer, to distinguish the correct fringe pattern from the many incorrect ones.



Figure 8.6: The interferometer setup for analysing the topography of the fibre surface. White light enters a 50/50 beamsplitter placed in front of the fibre. Half the light passes to the fibre (the grey beam), where it is reflected back to the beamsplitter. Half is reflected downwards to the mirror (the red beam), where it is reflected back up. Half of this beam interferes with half of the beam reflected from the fibre facet, producing an interference pattern at the imaging system above. The dotted line indicates a reflected beam, which is lost.



Figure 8.7: (a) Interference fringes generated using the white-light interferometry setup of figure 8.6 and the fibre shown in figure 8.5. Figure (b) shows the same image with the central area of interest cropped out. Figure (c) shows a cross-section of the fringe amplitude at an angle of 0° through the image. Finally, figure (d) shows the fringe amplitude normalised and ready for processing.

The process starts, then, with an interferometric image with as high a contrast between fringes as possible. Interferometric images of the fibre shown in figure 8.5 are shown in figure 8.7. The red layer only is extracted, and the image is cropped to 320 pixels square The fringe amplitude is then extracted. In principle, we know that the spatial distance between fringes is $\lambda/2$, but we cannot know anything about the spatial variation in between peaks. In addition, the fringe amplitude is subject to fluctuations resulting from background 'noise'. Instead of using the amplitude, then, we use the phase of the pattern. We first remove noise, by applying a Gaussian filter, then extract the phase by simply taking the arccosine of the normalised fringe amplitude. This is shown in figure 8.8(a). We know that the spatial depth between phase peaks is $\lambda/2$, but we would like to know the depth at other points in the phase. In addition, we cannot know from the phase map whether an adjacent peak is higher or lower. The phase map of 8.8(a) is a *wrapped* phase map, in the sense that it repeats every 2π . We require a continuous phase map, and this is accomplished using the phase-unwrapping technique described in [122]. When we have a continuous phase map of the surface, the depth is then very simply calculated from the unwrapped phase, ϕ , and the wavelength, λ , of red light⁷ as

$$D = \frac{\lambda}{2} \times \frac{\phi}{2\pi}.$$
(8.1)

The analysis assumes that the shot and the depression formed have a Gaussian shape, and we show here how the radius of curvature of the shot can be determined. Each point on a Gaussian curve has some curvature. If we imagine a circle at the bottom of the Gaussian, then the slope of the tangent to the circle at the bottom must equal the slope of the Gaussian at that point, and the second derivative of the circle and Gaussian must likewise be the same here.

Recall the equation of a circle:

$$(x - x_0)^2 + (y - y_0)^2 = r^2.$$
(8.2)

Then the first derivative is, by implicit differentiation:

$$\frac{dy}{dx} = -\frac{x - x_0}{y - y_0}.$$
(8.3)

And the second derivative is:

$$\frac{d^2 y}{dx^2} = \frac{-1 - y^2}{y - y_0}.$$
(8.4)

The radius of a circle is:

$$r = \sqrt{(x - x_0)^2 + (y - y_0)^2},$$
(8.5)

where x_0, y_0 are the coordinates of the centre of the circle.

From 8.4 we have $y - y_0$, and from 8.3 we have $x - x_0$, so the radius is given by

$$r = \sqrt{\frac{-(1 + \frac{dy^2}{dx})^3}{\frac{d^2y^2}{dx^2}}}.$$
(8.6)

This equation gives the general radius of curvature at any point in a curve described by the function *y*. So we can use it with a Gaussian function:

$$y = ae^{\frac{-(x-b)^2}{2c^2}} + d.$$
 (8.7)

Recall that here b is the mean, c is the standard deviation, and a is the amplitude. Then the first derivative is:

$$\frac{dy}{dx} = -a\frac{x-b}{c^2}e^{\frac{-(x-b)^2}{2c^2}},$$
(8.8)

and the second:

⁷Which we take to be 611 nm. This number comes from a calibration of the red layer in the .TIFF images we use, using a commercial mirror of known curvature.

$$\frac{\mathrm{d}^2 y}{\mathrm{d}x^2} = \frac{-a}{c^2} e^{\frac{-(x-b)^2}{2c^2}} + a \left(\frac{x-b}{c^2}\right)^2 e^{\frac{-(x-b)^2}{2c^2}}.$$
(8.9)

At the bottom of the Gaussian (right in the middle of it), x = b, so $\frac{dy}{dx} = 0$ and $\frac{d^2y}{dx^2} = -\frac{a}{c^2}$, and:

$$r = \frac{c^2}{a},\tag{8.10}$$

and this result is used to obtain the x and y radii of curvature from a fit of the 2-D Gaussian of equation 8.11 to the reconstructed shape of each fibre. In general, we take the geometric mean of the two radii for our final result. This is often used when taking the mean of values intended to be multiplied together. This is the case with the 2-D Gaussian function, equation 8.11, which is the product of two exponential functions.

$$g = A \exp\left[-a\left((x - x_0)^2 - b(y - y_0)^2 + c(x - x_0)(y - y_0)\right)\right] + d,$$
(8.11)

with in this case a, b, and c equal to

$$a = \left(\frac{\cos^2(\theta)}{2s_x^2} + \frac{\sin^2(\theta)}{2s_y^2}\right)$$
(8.12)

$$b = \left(\frac{\sin^2(\theta)}{2s_x^2} + \frac{\cos^2(\theta)}{2s_y^2}\right)$$
(8.13)

$$c = 2\cos(\theta)\sin(\theta)\left(\frac{1}{2s_x^2} - \frac{1}{2s_y^2}\right).$$
(8.14)

The parameter θ refers to the rotation of the elliptical 2-D Gaussian relative to the *xy*-axis. The parameters s_x and s_y control the width of the Gaussian in the *x*, *y*-directions, and *d* is a constant.

In addition to a variety of curved fibres, we also shot several flat fibres. The reason for this will become clear in later sections.



Figure 8.8: Figure (a) shows the phase extracted from the fringe amplitude. We know that adjacent peaks are spatially $\lambda/2$ apart, but the phase requires further processing before it can be used to map the entire surface. See the text for details. Figure (b) shows the final, calculated depth through a cross-section of the shot, while the final two figures show the shape of the shot, constructed from a series of 12 cross-sections, and a 2-D Gaussian fit to the shape. The final radius of curvature of this fibre was found from the fit to be: $R_x = 340 \ \mu m$ and $R_y = 355 \ \mu m$. The FWHM was $39.8\mu m$ and $40.7 \ \mu m$ in the x and y-directions, respectively.



(a) After one shot.

(b) Interferometric image of the flat fibre.

Figure 8.9: (a) A fibre after one single shot to polish the surface, and (b) the rather beautiful interferometric image of the 'flat' surface of the fibre. If we recall that moving between fringes corresponds to a difference in depth of $\lambda/2$, then we see that in the central area of the facet, there is very little variation. The greatest variation comes from the area of cleave chip. We are most concerned with the central area in our cavities, and so this fibre is considered essentially flat for our purposes.

8.1.5 Coating the fibres

48 fibres were produced in total, of which 40 were curved, 6 were flat, and 2 were plain cleaved fibres. They were placed in two separate containers of 24 each. The fibres were coated by AT Films, of Boulder, Colorado, using the technique of ion-beam sputtering. Alternate layers of Ta_2O_5 and SiO_2 are deposited in a vacuum chamber⁸ using a reactive process. A 1-inch diameter piece of glass was coated with the fibres as a 'witness' piece, with which to characterise the coating. The transmission measured by AT Films is shown in figure 8.10(b). The coating is specified to have a transmissivity of ~30 ppm. We also had some half-inch mirror substrates coated as part of the same run.

The assembly of parts required to transport them safely was of critical importance. The technical details of the transport assembly are shown in figure 8.11.



Figure 8.10: (a) The distribution of the radii of curvature of our fibres, and (b) the transmission of the coating as measured by a 1-inch 'witness' piece of glass, coated with the fibres. This data was supplied by AT Films.



(a) The fibre chuck.

(b) The complete transport assembly.

Figure 8.11: (a) A drawing of the fibre chuck. This is a small aluminium cylinder, 25 mm long and 6 mm in diameter. The fibre rests in the notch cut along its axis, protruding by around half a millimetre. At the other end of the cylinder a teflon cylinder of diameter 3 mm and length 5 mm is inserted into a recess in the aluminium, and a M1.2 grub screw is inserted radially. This has the effect of 'closing' the teflon on the fibre. A metal shim is pressed down into the notch, holding the fibre in the notch along its length. The handle of the shim is visible protruding in figure (a). Figure (b) shows the full transport assembly for each batch of 24 fibres. The 24 cylinders with fibres are inserted into holes in an aluminium block and secured by screws from the side. The block is screwed to an aluminium disc. The discs are designed to fit into hard plastic boxes without any movement during transport. Strain reflief bars are provided on the disc for the fibre bunches. which are then wrapped around the inside of the plastic box. The two boxes are then sealed in heat-shrunk plastic before being packaged in layers of bubble-wrap and sent to the coating company. The design of the fibre chuck and transport assembly is based on a similar design by the group of Jakob Reichel at ENS in Paris.



Figure 8.12: (a) A fibre in its aluminium cylinder and (b) the cylinders, block and disc secured in their plastic box.

8.2 The fibre cavity

This section expands on the basic optical cavity material laid out in chapter 4, extending it to the principles of the kind of cavity we are constructing using the coated fibre facets. Some of the cavity and coupling parameters we hope to achieve will be discussed.

129

The current state of the art in fibre Fabry-Pérot cavities (FFPCs) is the work by Hunger et al. in 2010 [125] and owes much to previous work in the field. The idea of placing mirrors on the tips of fibre-optic cables is not new, but the method of creating the mirror surface has varied over time. For example, mirrors have been bonded to the end facets of short lengths of fibre with epoxy resin, to form an interferometer *within* the fibre [126]. Layers of dielectric coating have been deposited by vacuum evaporation [127]. The mirrors have been transferred from another substrate onto the fibre tip, to form an open cavity; that is, the cavity consists of the coated ends of two fibres, with the cavity light field in free space rather than inside the fibre [128]. The finesse of the cavities has increased over the years, from, for example 20 in 1982 [127], 300 in 1985 [126] and 500 in 1986 [129]. In 2006, Steinmetz et al. reported an open FFPC with a finesse of 1000 [128]; they suggested that with the increasingly high-quality coatings becoming available, the only remaining barrier to ultra-high finesses is the roughness of the fibre facet. This issue has, in effect, been solved by the CO₂ machining technology discussed in section 8.1. With losses due to the roughness of the substrate extremely small[125], the path to finesses on the order of 100,000 is cleared.

8.2.1 Cavity parameters

The mirror coating on our fibres is of transmissivity $T_{1,2} \approx 30 \times 10^{-6}$ at 866 nm. The finesse of the cavity is given by equation 4.18 on page 29. If we neglect the unknown scattering and absorption losses, then the maximum finesse we can expect is 104,000.

We consider the stability of the cavity types discussed in section 4.1.0.1 when selecting a potentical cavity geometry for the fibre cavity.

Plane-parallel cavities are not considered at all, due to their alignment problems—the mirrors must be very-nearly parallel to avoid the cavity field 'walking' off the cavity axis into free space (out of the cavity).

Consider the concentric case. $L = R_1 + R_2$. From the stability criterion, 4.13 on page 28, we have that $g_1g_2 = 1$, which is right at the limit of stability. It's easy to see that a small change in parameters could make the cavity unstable, and so this type of cavity is avoided.

Similarly for the confocal case, where $g_1g_2 = 0$, again at the edge of the stability region. If R_1 and R_2 are not exactly equal, the cavity will not be stable, and so this cavity is also avoided.

The half-symmetric, or 'near-hemispheric' cavity is a possible candidate. In this cavity one mirror is planar and one is curved. The waist is at the planar mirror and the spot size at the spherical mirror can be adjusted by changing the cavity length. We will consider in this section the various CQED parameters of symmetric (that is, $R_1 = R_2$) and half-symmetric cavities.

8.2.1.1 Cavity waist and clipping losses

There are two reasons why we might be interested in the waist of the cavity mode. For CQED purposes, to optimise the coupling of the ion to the field we would like the waist to be as small as possible at the ion. We are also interested in the spot size at the mirrors, because if this exceeds our mirror diameter then we incur clipping losses. In this section we calculate the cavity waist and spot size at the mirrors, and determine whether clipping losses are likely to pose a problem for the mirrors we have machined.

Fortunately, simple expressions for the cavity waist and spot size are available for the symmetric and half-symmetric cavities of interest. For a symmetric cavity, the waist is at the centre of the cavity, and is given by [130, Ch. 19]:

$$w_0^2 = \frac{L\lambda}{\pi} \sqrt{\frac{1+g}{4(1-g)}},$$
(8.15)

with $g = g_1 = g_2$. The spot size at the mirrors are the same, and given by

$$w^2 = \frac{L\lambda}{\pi} \sqrt{\frac{1}{1-g^2}}.$$
(8.16)

Figure 8.13 shows the variation of cavity waist with radius of curvature, in a symmetric cavity of several realistic lengths. This compares with half-symmetric cavities, where the position of the mode waist is on the left mirror. In this case, $g_1 = 1$ and $g = g_2$. The waist is given by

$$w_0^2 = \frac{L\lambda}{\pi} \sqrt{\frac{g}{1-g}},\tag{8.17}$$

and the spot size at the other mirror by

$$w_1^2 = \frac{L\lambda}{\pi} \sqrt{\frac{1}{g(1-g)}}.$$
 (8.18)

The variation of waist with cavity length and radius of curvature for a half-symmetric cavity is shown in figure 8.14. The cavity lengths have been chosen as candidate lengths based on our simulations of trap depth vs. electrode separation (see figure 5.13 on page 50). The cavity should be as short as possible, within the limits imposed by the trapping potential. If we consider a cavity of around 350 μ m, with mirrors of curvature around 300 μ m, we see that the waist in a symmetric cavity is around 7 μ m, whereas in a half-symmetric cavity it is around 6 μ m. Replacing the curved mirror with a planar one has made little difference. However, the position of the waist at the planar mirror and not in the centre of the cavity has implications for the coupling strength, which, since the ion is trapped at the centre of the cavity, would be weaker in a half-symmetric cavity. We will see shortly if this is an acceptable sacrifice.



Figure 8.13: The variation in (a) cavity waist at the centre of a symmetric cavity with radius of curvature of the two identical mirrors. Notice how the waist becomes vanishingly small as we approach the concentric region wheren L = 2R. (b) shows the variation in spot size on the mirrors. Plots are truncated at the limit of the cavity stability region.



Figure 8.14: The variation in (a) cavity waist at the planar mirror of the half-symmetric cavity with radius of curvature, shown for four cavity lengths. (b) shows the variation of the spot size at the curved mirror for the same parameters. Plots are truncated at the limit of the stability region.

Clipping losses

The Gaussian intensity distribution is given by

$$I(r,z) = I_0 \left(\frac{w_0}{w_z}\right)^2 \exp\left(\frac{-2r^2}{w^2(z)}\right),\tag{8.19}$$

with *r* the radial distance from the beam axis, w_0 the beam waist and w(z) the 'spot size', that is, the radial distance from the beam axis at which the intensity has dropped to 1/e of its initial value. We would like to know how much of the Gaussian distribution 'spills over' the edge of our finite-diameter mirrors. We can get an idea of this by integrating from D/2 to infinity and from 0 to 2π as follows, remembering to normalise:
$$L_{\text{Clip}} = \frac{\int_{0}^{2\pi} \mathrm{d}\theta \int_{D/2}^{\infty} e^{\frac{-2r^{2}}{w_{m}^{2}}} r \mathrm{d}r}{\int_{0}^{2\pi} \mathrm{d}\theta \int_{0}^{\infty} e^{\frac{-2r^{2}}{w_{m}^{2}}} r \mathrm{d}r}$$
(8.20)

$$=e^{\frac{-D^2}{2w_m^2}}.$$
 (8.21)

with w_m the spot size at the curved mirror and D the mirror diameter. The average diameter of our mirrors from one canister is 171 μ m, and so performing the integral using this diameter gives the results shown in figure 8.15(a). It seems that for spot sizes on the order of tens of microns, such as shown in figure 8.14(b), clipping losses are extremely small. Following a similar argument to [125], we can set a limit on clipping losses, and consequently consider the finesse achievable for each spot size. If we state that clipping losses must contribute 10% of total losses, then we have from equation 4.18 that

$$L_{\text{Clip}} = \frac{2\pi}{10F}$$

$$\Rightarrow F = \frac{2\pi}{10L_{\text{Clip}}}.$$
(8.22)



Figure 8.15: (a) Clipping loses as a function of spot size on the curved mirror for an average mirror diameter of 171 μ m. We see that for spot sizes of tens of microns clipping losses are very small. (b) The finesse achievable if we fix clipping losses at 10% of total losses and vary the spot size at the curved mirror.

So we see that clipping losses are not a significant barrier to achieving high finesse.

8.2.1.2 Mode volume and the CQED Coupling parameter

The cavity mode volume of equation 4.35 on page 32, is given by spatially integrating the full cavity mode function of equation 4.4 over the cavity length. This gives us, for the TEM_{00} mode

$$\mathscr{V} = \frac{\pi}{4} w_0^2 L. \tag{8.23}$$

The mode volume is shown in terms of cubic wavelengths, for symmetric and half-symmetric cavities in figure 8.16.



Figure 8.16: The mode volume in (a) symmetric and (b) half-symmetric cavities, given in terms of cubic wavelengths at 866.45250 nm. The radius of curvature in (b) refers to the curved mirror.

We can write the rate of coherent exchange of single atomic and photonic excitations of equation 4.40 on page 33, for a two-level atom in free space at the maximum field intensity as [125]

$$g = \sqrt{\frac{3\lambda^2 c\Gamma}{8\pi\mathcal{V}}},\tag{8.24}$$

with Γ the FWHM linewidth of the upper state.

Let us assume that the ion is trapped somewhere along the cavity axis, along the z-direction of equation 4.8 on page 27, such that (x, y) = 0. Further, we assume that the ion is trapped at a longitudinal antinode, such that the longitudinal mode function of equation 4.9 has the value of unity. Then, to calculate g at an arbitrary z-position, such as the middle of a half-symmetric cavity, we must multiply equation 8.25 by the mode functions $\Psi_n \Psi_m$ of equation 4.8, which amounts to multiplying by $w_0/w(z)$. So, we have

$$g = \frac{w_0}{w(z)} \sqrt{\frac{3\lambda^2 c\Gamma}{8\pi\mathscr{V}}}.$$
(8.25)

The variation of the coupling parameter with radius of curvature of the mirrors is shown for several cavity lengths in figure 8.17. As a point of interest, some predicted values of κ —the cavity linewidth, are shown in figure 8.18.

134



Figure 8.17: The projected coupling parameter *g* for (a) an ion at the centre of a symmetric cavity, and (b) for an ion at the centre of a half-symmetric cavity. Plots are truncated at cavity stability limits.



Figure 8.18: The variation in the cavity decay rate κ with total cavity losses, that is, the transmission, absorption and scattering losses of both mirrors, shown for several cavity lengths.

8.2.1.3 Possible cavity geometries

Let us assume we have access to single-mode as well as multi-mode coated fibres. There are many ways we might choose to set up a cavity. Figure 8.19 shows several possible geometries. In all cases, the beam enters the cavity from the left, in the direction of the arrow, and exits on the right.

Of these, we consider (a) the most likely to be successful. The flat surface of the single-mode fibre matches the radius of curvature of the plane waves exiting it. Using a multi-mode fibre on the output channel gives a higher coupling efficiency than would a single-mode fibre [125], since it is fixed by the numerical aperture of the fibre. The fibres in our experiment have a numerical aperture of 0.22, giving an acceptance angle of 25°. The divergence angle of the cavity mode is given by $2\lambda/(\pi w_0)$ at distances much greater than the Rayleigh range $z_R = \pi w_0^2/\lambda$, which if we consider as

135

an example the case of a half-symmetric cavity of length 300 μ m and mirrors of radius 300 μ m, is 130 μ m. So we see that the divergence is much less than the acceptance angle of the fibre for this waist, and indeed the waist would need to be less than 1.3μ m for the divergence to match the fibre acceptance angle in this case. In addition, the multi-mode fibre makes the cavity much more forgiving of misalignments. It is, however, a first step towards eventually using two single-mode fibres.



Figure 8.19: Some possible fibre-cavity options.

8.2.1.4 Conclusion

In conclusion for a cavity of a reasonable length, governed by possible trapping potentials, and for mirrors of around 300–400 μ m we should be able to enter the strong coupling regime, that is, $g > \kappa, \Gamma$, with both symmetric and half-symmetric cavities. This assumes the coating, and therefore finesse, is as expected. From figure 8.17 it is clear that this will be more difficult with a half-symmetric cavity. The current status of the fibre cavity will be covered in chapter 9.

8.3 A new trap

The ion trap used for the work in this thesis successfully achieved its objective, that is, to test the principle of placing optical fibres within the electrodes of the trap, extremely close to the trapped ion. However, in several ways it is not suited to housing a high-finesse miniature optical cavity, principally

- Poor vibration isolation
- No room to accommodate in-vacuum precision cavity manipulation devices

With these added points in mind a new trap has been designed and built, starting with the electrode mounting structures, and working out finally to a new chamber and vacuum setup. The

trap itself, that is, the inner and outer electrode assembly, is identical to the previous trap. All other parts have been designed and built from scratch, with the stability and precision movement of the fibre cavity foremost in mind. In this section, a description of the principal parts of the trap will be given, along with the design decisions that produced them.

8.3.1 The new vacuum design

Figure 8.20 shows the new vacuum flange assembly which sits directly beneath the chamber. The design differss from the original, shown in figure 5.25 on page 62 in that we have added a further four CF40 flanges to the existing four-way reducing cross. This allows us to bring up to four rf connections into the vacuum on two rf feedthroughs⁹, along with a dedicated 19-pin dc-connection flange¹⁰ to carry all dc connections into the vacuum. Inside the vacuum a ribbon cable connects the 19-pin feedthrough to a 25-pin sub-D connector attached by a bracket to the trap. Two further CF40 flanges are taken up with the fibre feedthroughs¹¹, and a final flange is used for the ion gauge. In total, 6 of the 8 CF40 flanges on the cross will be in use. The customised CF160 flange is attached with M6 bolts to the underside of the optical table.



Figure 8.20: The new bottom flange assembly. See the text for details.

⁹Lewvac FHP5-50C2-40CF 5 kV 50 A copper 2-pin feedthroughs mounted on CF40 flanges.

¹⁰Lewvac C19-KIT40CF, supplied with UHV and air cables.

¹¹Lewvac FO-SM800-40CF single-mode fibre feedthroughs for 800–900 nm, mounted on CF40 flanges.



Figure 8.21: The new chamber sits on top of the vacuum assembly of figure 8.20. As with the old trap, the windows are sealed with indium wire and secured with brass clamps.

8.3.2 The trap structure

Vibration isolation

To isolate the cavity from vibration, the trap sits on top of two stainless steel blocks, marked in figure 8.22, which in turn sit on viton rubber feet. The top vibration block is hollow and contains the cavity manipulation devices, which will be described later.



Figure 8.22: The trap sits on top of two heavy stainless-steel blocks. The blocks sit on viton rubber blocks to damp vibration. Two studs pass loosely through the customised flange beneath into the vibration blocks, to prevent any gross movement of the blocks and trap.

The 'cage'

Figure 8.23 shows the 'cage' surrounding the electrodes. For micromotion compensation we have four dc electrodes—the top and bottom outer rf-ground electrodes, the oven collimator and the side dc electrode. For the purpose of moving the rf-minimum, that is, the trap centre, to coincide with the cavity mode we have two side rf electrodes in addition to the main rf inner electrodes. The oven is held in place by a MACOR support, chosen for its high heat tolerance.



Figure 8.23: The 'cage', which supports the top trap electrode structure.

The top and bottom electrode assemblies

Figure 8.24(a) shows the top electrode assembly. A ring piezo¹² allows us to move the cavity axially over a range of 4 μ m. The operation of this assembly can be more clearly seen in the cross-section of figure 8.25. The top electrode structure is held in the cone mount with a small grub screw. The dc connection to the top ground electrode enters through one of the bolts holding the assembly together.

Figure 8.24(b) shows the bottom electrode assembly. A manual *xy*-stage¹³ is attached to the bottom of the top vibration block. This has a travel of 5 mm, and is accessed with an allen key through the cutout in the vibration block visible in figure 8.22. On top of this sits a UHV-compatible *xy*-piezo stage¹⁴ with a travel of 100 μ m. This stage and the ring piezo provide the precision cavity-alignment control we require. Both stages and the bottom of the vibration block have a central bore of 9 mm diameter, and the bottom fibre passes up to the electrode through this, insulated by a teflon cylinder. The dc connection to the bottom ground electrode enters at a small screw in the bottom cone-shaped mount, which sits on top of the *xy*-piezo. The cone mount is insulated from the piezo with polyether ether keton (PEEK) spacers. The electrode structure is held in the cone mount with a small grub screw, as with the top assembly.

The fibre strain relief

The fibres come in through the vacuum feedthroughs of figure 8.20, and then up to the top and bottom of the trap. We provide strain-relief to the fibres in the form of the three brackets shown in figure 8.26. A copper shim connected to the rf supply makes contact with the copper fibres at the middle bracket, and the rf is carried to the inner electrodes by the copper coating. To prevent

¹²Piezomechanik HPCh 150/15-8/3.

¹³Custom-made for UHV by our workshop.

¹⁴Piezosystem Jena PXY100, custom-built for UHV.



141

Figure 8.24: (a) The top and (b) bottom electrode assemblies.



Figure 8.25: A cross-section through the cage structure of figure 8.23 showing how the ring piezo moves the top electrode up and down.

anything beneath this point receiving the rf voltage, we will chemically etch a small portion of the copper fibre at a point between the bottom and middle brackets, as marked in figure 8.26. The top two brackets are machined from PEEK, while the bottom bracket is aluminium.

To conclude this section, figure 8.27 shows the assembled trap with the electrical connections about to be attached to the sub-D connector.



Figure 8.26: Strain relief is provided for the fibres by the three brackets shown in this figure. The cylinders are customised Swagelock devices. See the text for more details.



Figure 8.27: The new trap, fully built and ready for the electrical connections.

Chapter 9

Conclusion and outlook

9.1 Conclusion

A unique miniature ion trap has been designed and built with the express purpose of accommodating two optical fibres within the electrodes of the trap. The electrodes consist of two opposing metal tubes of outer diameter 458 μ m and inner diameter 254 μ m, separated by 446 μ m, giving an ionelectrode distance of 223 μ m. With cleaved optical fibres recessed by 50 μ m inside the electrodes we are able to trap single ⁴⁰Ca⁺ ions for several hours, proving that the fibres cause no significant degradation of the trapping potential.

The trap has been fully characterised in simulation and experiment. Micromotion has been compensated in three dimensions, and the location of stray fields after trap loading has been calculated and determined to originate not from the fibres but from the electrodes. Continuous single photons have been generated and the phenomenon of photon antibunching has been observed in the correlations between the photons received at PMTs connected to the two fibres. The high quality of the single-photon source has been verified in the value of the normalised second-order correlation function $g^{(2)}(0) = 0.05 \pm 0.04$.

In addition, pulsed single photons have been produced using a periodic trigger pulse injected into one fibre, and antibunching has been observed in the photons emitted into the other fibre using a Hanbury Brown-Twiss photon-correlation setup. Again, the lack of two-photon contamination of the correlation spectrum around zero time delay has been verified in the absence of net positive photon cross-correlations. A value of $g^{(2)}(0) = 0.005$ at zero time was noted. The difference between the continuous and pulsed measurements around zero time delay can be accounted for by the much-reduced presence of 397 nm light during the measurement time window in the pulsed measurement.

Using a CO_2 laser, concave surfaces have been machined on the end facets of 48 multimode fibres. The radii of curvature of the surfaces has been determined in each case using white-light interferometry. The fibres have been treated with a high-reflectivity coating, specified to have a transmissivity of 30 ppm at the wavelength of interest. A new trap and vacuum system has been designed and built to accommodate the new fibre cavity and the precision positioning apparatus required to manipulate it in vacuum.

9.2 Outlook

The experiment is at a particularly interesting stage at present. Since the work of this thesis was completed, work on producing a high-finesse cavity using the coated multimode fibres has been ongoing. At present, cavities have been successfully built using a flat fibre as the input and a concave fibre as the output, as well as concave-concave cavities. In both cases, finesses of approximately 60,000 have been measured (sidebands on an 854 nm input laser¹ were used to measure the free spectral range of the cavity, with the length of the cavity determined using a microscope). The finesse remains at approximately this figure over a cavity length of 100–500 μ m in both geometries.

A problem has been identified with using a curved fibre for the input in that there appears to be some leakage of light out of the 'lip' of the end facet (this is the black ring visible on the facet in, for example, figure 8.5(d)). The laser shots appear to have left a chamfer around the edge of the fiber facet. Although this doesn't affect the finesse of the cavity, it does produce an offset on the cavity signal, such that the cavity transmission and more importantly the dressed state of section 4.2.2 on page 33 could potentially be obscured. It seems likely, therefore, that in the immediate future we will use a flat-concave cavity.

Turning to progress on the new trap; the trap, vacuum system and electrical feedthroughs have all been assembled and work has started on the process of inserting the very delicate fibre and electrode assemblies.

Work has also begun preparing the CO_2 laser shooting setup for shooting copper-coated singlemode fibres². This includes automating the fibre-rotation process. Assuming success with a multimode fibre cavity, we will then test single-mode fibre cavities, probably in the configuration of single-mode input and multimode output. It is likely that the leakage problem described above will not occur in single-mode cavities, because the core is so small, and so far from the edge of the fibre.

¹An 854 nm laser was chosen simply because it had higher power than our 866 nm laser.

²IVG Fiber Cu800 copper-coated single-mode fibre with custom 200 μ m cladding.

Appendix A

The Lasers

As mentioned in chapter 5, the transitions of interest in the ${}^{40}Ca^+$ ion are accessible by commerciallyavailable or home-built solid-state lasers. Most of the lasers we use are constructed by us using commercial diodes and gratings. This section describes briefly the lasers and the methods used to stabilise their frequencies.

Our suite of lasers consists of home-built extended-cavity diode lasers (ECDL's) in the Littrow configuration [131], addressing transitions in ${}^{40}Ca^+$ at 866 nm, 850 nm, 854 nm, and in ${}^{40}Ca$ at 423 nm and 375 nm, alongside a Toptica SHG110 frequency-doubled 794 nm tapered-amplifier setup which is used to produce 397 nm light.

A small amount of light from each laser is picked off by a glass plate¹ and sent to a wavemeter.² We calibrate the wavemeter using the known wavelength of a commercial Helium-Neon laser³. A Labview virtual instrument (vi) reads the wavemeter at regular intervals and feeds back a voltage to the laser's horizontal grating alignment piezo. In this way the wavelengths are kept approximately within the linewidth of the relevant transitions. However, at a resolution of 60 MHz, a measured wavelength of 866.45250 nm could be as high as 866.45265 nm or as low as 866.45235 nm and for some experiments we must reduce and stabilise the linewidth still further. This is done using a fast scanning cavity lock—an absolute reference frequency is generated by locking a diode laser to a tunable cavity which is in turn locked to the D₁ transition in atomic caesium using polarisation spectroscopy. Our lasers are locked to the reference laser and an absolute stability of less than 100 kHz at 397 nm is achievable. The system, with performance measurements, is described in detail in [132].

Figure 5.1 on page 37 shows the arrangement of the lasers around the trap. 423 nm and 375 nm light is overlapped at a polarising beamsplitter and focused to a spot size of approximately 200 μ m at the trap centre. We have typically used powers of around 200 μ W for each of these beams. An electronically-controlled shutter is placed in the path of these beams, the control of which is integrated into an automatic ion-loading LabView vi.

A cooling laser at 397 nm and a repumper laser at 866 nm are overlapped at a dichroic mirror and focused to a spot size of around 100 μ m at the trap centre. These enter the trap at an angle of approximately 12° below the horizontal. Beams from the 850 nm and 854 nm lasers are overlapped

¹For example, a Thorlabs WG11050 broadband precision window, which reflects about 5% of incident light.

²High Finesse WS-7—resolution ~60 MHz.

³SIOS Messtechnik Gmbh model SL 02-1 at 633 nm.

using half-wave plates and a beamsplitter, and the combined beam is overlapped with a second 397 nm beam. These beams enter the trap at an angle of approximately 6° above the horizontal. A third 397 nm beam enters the trap at an angle of approximately 10° below the horizontal. Three cooling beams are required for micromotion-compensation purposes (this process is described in detail in section 6.2 on page 70 onwards). A fourth 397 nm laser ('beam 1') enters the trap horizontally. Considerable care has been taken to reduce the spot size of this beam and to reduce the amount of light scattered by the beam from elements of the trap. A spatial filter has been placed in the beam path, and the beam is focused into the trap using a custom-made lens. This beam was used where the experiment called for such low levels of scattered light, and its characteristics will be discussed in the relevant section.

Beams from all the lasers are transported to the trap area of the optical table using single-mode fibres.⁴ One advantage of using single-mode fibres are that the fibre acts as a spatial filter, ensuring the beam emerging has an approximately Gaussian shape. Plus of course it makes it very easy to transport the beam to wherever it may be needed.

⁴Typically Thorlabs P1-780PM for infra-red and Thorlabs PMS350HP for ultra-violet.

Bibliography

- [1] Schrodinger E. Are there quantum jumps? *Br J Philos Sci*, 3(10):109–123, 1952. doi: 10.1093/bjps/iii.10.109. Cited on 1
- [2] Manek Dubash. Moore's law is dead, says Gordon Moore. http://news.techworld.com/operating-systems/3477/moores-law-is-dead-says-gordonmoore/, 2005. Accessed on 5th Jan 2012. Cited on 1
- [3] Jeremy Geelan. Moore's law: 'We see no end in sight,' says Intel's Pat Gelsinger. http://java.sys-con.com/node/557154, 2008. Accessed on 5th Jan 2012. Cited on 1
- [4] B. Weber, S. Mahapatra, H. Ryu, S. Lee, A. Fuhrer, T. C. G. Reusch, D. L. Thompson, W. C. T. Lee, G. Klimeck, L. C. L. Hollenberg, and M. Y. Simmons. Ohms law survives to the atomic scale. *Science*, 335(6064):64–67, 2012. doi: 10.1126/science.1214319. Cited on 1
- [5] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6/7), May 1981. Cited on 2
- [6] David Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 400:97–117, Jul 1985. Cited on 2
- [7] Richard P. Feynman. Quantum mechanical computers. *Optics News*, 11(2):11–20, Feb 1985. doi: 10.1364/ON.11.2.000011. Cited on 2
- [8] Michael A. Nielsen and Isaac L. Chuang. Quantum computation and quantum information (Cambridge series on information and the Natural Sciences). Cambridge University Press, 1 edition, January 2004. ISBN 0521635039. Cited on 2
- [9] Peter W. Shor. Algorithms for quantum computation: discrete log and factoring. In Proceedings of the 35th Annual Symposium on Foundations of Computer Science, pages 124–134. IEEE Press, 1994. Cited on 2
- [10] Lieven M.K. Vandersypen, Matthias Steffen, Gregory Breyta, Constantino S. Yannoni, Mark H. Sherwood, and Isaac L. Chuang. Experimental realisation of Shor's quantum factoring algorithm using nuclear magnetic resonance. *Nature*, 414:883–887, October 2001. Cited on 2
- [11] J. I. Cirac and P. Zoller. Quantum computations with cold trapped ions. *Physical Review Letters*, 74(20), 1994. Cited on 4
- [12] Neil A. Gershenfeld and Isaac L. Chuang. Bulk spin-resonance quantum computation. *Science*, 275(17):350–356, 1997. Cited on 4
- [13] Q. A. Turchette, C. J. Hood, W. Lange, H. Mabuchi, and H. J. Kimble. Measurement of conditional phase shifts for quantum logic. *Phys. Rev. Lett.*, 75:4710–4713, Dec 1995. doi: 10.1103/PhysRevLett.75.4710. Cited on 4

- [14] Mark S Sherwin, Atac Imamoglu, and Thomas Montroy. Quantum computation with quantum dots and terahertz Cavity Quantum Electrodynamics. *Physical Review A*, 60(5):6, 1999. Cited on 4
- [15] John J.L. Morton, Dane R. McCamey, Mark A. Eriksson, and Stephen A. Lyon. Embracing the quantum limit in silicon computing. *Nature*, 479:345–353, 2011. Cited on 4
- [16] William F. Koehl, Bob B. Buckley, F. Joseph Heremans, Greg Calusine, and David D. Awschalom. Room temperature coherent control of defect spin qubits in silicon carbide. *Nature*, 479:84–87, 2011. Cited on 4
- [17] David P. Divincenzo and Daniel Loss. Quantum information is physical. Superlattices and Microstructures, 23(3/4), 1998. Cited on 4
- [18] David Beckman, Amalavoyal N. Chari, Srikrishna Devabhaktuni, and John Preskill. Efficient networks for quantum factoring. *Phys. Rev. A*, 54:1034–1063, Aug 1996. doi: 10.1103/ PhysRevA.54.1034. Cited on 4
- [19] Richard J. Hughes, Daniel F. V. James, Emanuel H. Knill, Raymond Laflamme, and Albert G. Petschek. Decoherence bounds on quantum computation with trapped ions. *Phys. Rev. Lett.*, 77:3240–3243, Oct 1996. doi: 10.1103/PhysRevLett.77.3240. Cited on 4
- [20] D. Kielpinski, C. Monroe, and D.J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417:709–711, June 2002. doi: 10.1038/nature00784. Cited on 5
- [21] Jonathan P. Home, David Hanneke, John D. Jost, Jason M. Amini, Dietrich Leibfried, and David J. Wineland. Complete methods set for scalable ion trap quantum information processing. *Science*, 325(5945):1227–1230, 2009. doi: 10.1126/science.1177077. Cited on 5
- [22] W. K. Hensinger, S. Olmschenk, D. Stick, D. Hucul, M. Yeo, M. Acton, L. Deslauriers, C. Monroe, and J. Rabchuk. T-junction ion trap array for two-dimensional ion shuttling, storage and manipulation. *Applied Physics Letters*, 88(3):034101, 2006. doi: 10.1063/1. 2164910. Cited on 5
- [23] D.P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48:771–784, 2000. Cited on 5
- [24] Mark Oxborrow and Alastair G Sinclair. Single-photon sources. *Contemporary Physics*, 46 (3):173–206, 2005. doi: 10.1080/00107510512331337936. Cited on 5
- [25] David P. DiVincenzo and Daniel Loss. Quantum computers and quantum coherence. *Journal of Magnetism and Magnetic Materials*, 200:202–218, 1999. Cited on 5, 6
- [26] C.H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of the IEEE International Conference on Computers, Systems and Signal Processing, Bangalore, India*, page 175. IEEE, New York, 1984. Cited on 5
- [27] Nicolas Gisin, Grégoire Ribordy, Wolfgang Tittel, and Hugo Zbinden. Quantum cryptopgraphy. *Reviews of Modern Physics*, 74:145–194, Mar 2002. Cited on 5, 6
- [28] H.-J. Briegel, W. Dür, J. I. Cirac, and P. Zoller. Quantum repeaters: The role of imperfect local operations in quantum communication. *Phys. Rev. Lett.*, 81:5932–5935, Dec 1998. doi: 10.1103/PhysRevLett.81.5932. Cited on 6
- [29] John Preskill. Plug-in quantum software. Nature, 402:357–358, 1999. Cited on 6

- [30] Stefanie Barz, Elham Kashefi, Anne Broadbent, Joseph F. Fitzsimons, Anton Zeilinger, and Philip Walther. Demonstration of blind quantum computing. *Science*, 335(6066):303–308, 2012. doi: 10.1126/science.1214707. Cited on 6
- [31] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Phys. Rev. Lett.*, 78:3221–3224, Apr 1997. doi: 10.1103/PhysRevLett.78.3221. Cited on 6, 34, 86
- [32] Wolfgang Paul. Electromagnetic traps for charged and neutral particles. *Rev. Mod. Phys.*, 62 (3):531–540, Jul 1990. doi: 10.1103/RevModPhys.62.531. Cited on 7, 41
- [33] W. Neuhauser, M. Hohenstatt, P. E. Toschek, and H. Dehmelt. Localized visible ba⁺ monoion oscillator. *Phys. Rev. A*, 22:1137–1140, Sep 1980. doi: 10.1103/PhysRevA.22.1137. Cited on 7
- [34] Raymond E and March. Quadrupole ion trap mass spectrometry: a view at the turn of the century. *International Journal of Mass Spectrometry*, 200(1-3):285 – 312, 2000. ISSN 1387-3806. doi: 10.1016/S1387-3806(00)00345-6. Cited on 7
- [35] NIST. Nist's second 'quantum logic clock' based on aluminium ion is now world's most precise clock. http://www.nist.gov/pml/div688/logicclock_020410.cfm, 2010. Cited on 7
- [36] C. W. Chou, D. B. Hume, J. C. J. Koelemeij, D. J. Wineland, and T. Rosenband. Frequency comparison of two high-accuracy al⁺ optical clocks. *Phys. Rev. Lett.*, 104:070802, Feb 2010. doi: 10.1103/PhysRevLett.104.070802. Cited on 7
- [37] F. M. Penning. Die glimmentladung bei niedrigem druck zwischen koaxialen zylindern in einem axialen magnetfeld. *Physica*, 3(9):873–894, November 1936. ISSN 00318914. doi: 10.1016/S0031-8914(36)80313-9. Cited on 7
- [38] David J. Griffiths. Introduction to Electrodynamics. Prentice Hall International, Inc., 1999. Cited on 8
- [39] R. F. Wuerker, H. Shelton, and R. V. Langmuir. Electrodynamic containment of charged particles. *Journal of Applied Physics*, 30(3):342–349, 1959. doi: 10.1063/1.1735165. Cited on 8, 12
- [40] Randall D. and Knight. The general form of the quadrupole ion trap potential. *International Journal of Mass Spectrometry and Ion Physics*, 51(1):127 131, 1983. ISSN 0020-7381. doi: 10.1016/0020-7381(83)85033-5. Cited on 9
- [41] P. Ghosh. Ion Traps. Clarendon Press, 1995. Cited on 8, 12
- [42] N.W. McLachlan. Theory and application of Mathieu functions. Dover Publications, 1964. Cited on 12
- [43] T. Hansch and A. Schawlow. Cooling of gases by laser radiation. *Optics Communications*, 13(1):68–69, January 1975. ISSN 00304018. doi: 10.1016/0030-4018(75)90159-5. Cited on 17
- [44] D.J. Wineland and H.G. Dehmelt. Proposed 10¹⁴ delta frequency resolution laser fluorescence spectroscopy on Ti⁺ mono-ion oscillator III. *B. Am. Phys. Soc.*, 20:637, January 1975. Cited on 17
- [45] W. Neuhauser, M. Hohenstatt, P. Toschek, and H. Dehmelt. Optical-sideband cooling of visible atom cloud confined in parabolic well. *Physical Review Letters*, 41(4):233–236, July 1978. doi: 10.1103/PhysRevLett.41.233. Cited on 17

- [46] Harold J. Metcalf and Peter van der Straten. *Laser Cooling and Trapping (Graduate Texts in Contemporary Physics)*. Springer, November 2001. ISBN 0387987282. Cited on 17, 21
- [47] J. I. Cirac, L. J. Garay, R. Blatt, A. S. Parkins, and P. Zoller. Laser cooling of trapped ions: The influence of micromotion. *Phys. Rev. A*, 49:421–432, Jan 1994. doi: 10.1103/PhysRevA. 49.421. Cited on 17
- [48] Jürgen Eschner, Giovanna Morigi, Ferdinand Schmidt-Kaler, and Rainer Blatt. Laser cooling of trapped ions. J. Opt. Soc. Am. B, 20(5):1003–1015, May 2003. doi: 10.1364/JOSAB.20. 001003. Cited on 17
- [49] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook. Optical molasses. J. Opt. Soc. Am. B, 6(11):2084–2107, Nov 1989. doi: 10.1364/JOSAB.6. 002084. Cited on 17, 22
- [50] Daniel F. V. James. Quantum dynamics of cold trapped ions, with application to quantum computation. *Applied Physics B*, 66(2):20, 1997. doi: 10.1007/s003400050373. Cited on 18, 99
- [51] Sy-Sang Liaw. *Ab initio* calculation of the lifetimes of 4 p and 3 d levels of Ca⁺. *Phys. Rev.* A, 51:R1723–R1726, Mar 1995. doi: 10.1103/PhysRevA.51.R1723. Cited on
- [52] Yu Ralchenko, A.E. Kramida, J Reader, National Institute of Standards NIST ASD Team (2011), and MD. Technology, Gaithersburg. Nist atomic spectroscopy database (ver. 4.1.0), [online]. http://physics.nist.gov/asd, 2011. Accessed on March 9th 2012. Cited on 18
- [53] Howard Carmichael. An open systems approach to quantum optics, volume 18. Springer-Verlag, 1993. Cited on 21
- [54] M Keller, B Lange, K Hayasaka, W Lange, and H Walther. A Calcium ion in a cavity as a controlled single-photon source. *New Journal of Physics*, 6(1):95, 2004. Cited on 21, 86, 112
- [55] Sze M Tan. A computational toolbox for quantum and atomic optics. *Journal of Optics B Quantum and Semiclassical Optics*, 1(4):424–432, 1999. Cited on 23
- [56] E. M. Purcell. Spontaneous emission probabilities at radio frequencies. In *Physical Review*, volume 69, pages 681+, 1946. Cited on 25
- [57] Haroche S. and Kleppner D. Cavity Quantum Electrodynamics. *Physics Today*, 42(1):24, 1989. Cited on 25
- [58] Serge Haroche and Jean-Michel Raimond. Exploring the Quantum: Atoms, Cavities, and Photons (Oxford Graduate Texts). Oxford University Press, USA, 1st ed edition, October 2006. ISBN 0198509146. Cited on 25
- [59] Rodney Loudon. The Quantum Theory of Light (Oxford Science Publications). Oxford University Press, USA, 3 edition, September 2000. ISBN 0198501765. Cited on 30, 32, 87, 89, 94
- [60] Claude Cohen-Tannoudji, Jacques Dupont-Roc, and Gilbert Grynberg. *Photons et atomes* : introduction à l'électrodynamique quantique. EDP Sciences : CNRS Ed., 2001. ISBN 286883535. Cited on 25, 30
- [61] H. Kogelnik and T. Li. Laser beams and resonators. Applied Optics, 5(10):1550–1567, 1966. Cited on 26, 27

- [62] Geoffrey Brooker. Modern Classical Optics (Oxford Master Series in Atomic, Optical and Laser Physics). Oxford University Press, USA, October 2003. ISBN 019859965X. Cited on 27
- [63] G. Rempe, R. J. Thompson, H. J. Kimble, and R. Lalezari. Measurement of ultralow losses in an optical interferometer. *Optics Letters*, 17:363–366, March 1992. Cited on 29
- [64] Paul Davies. *Quantum Mechanics*. Routledge & Kegan Paul, London and New York, 1987. Cited on 30
- [65] E. T. Jaynes and F. W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, January 1963. ISSN 0018-9219. doi: 10.1109/PROC.1963.1664. Cited on 31
- [66] Bruce W. Shore and Peter L. Knight. The Jaynes-Cummings Model. Journal of Modern Optics, 40(7):1195–1238, July 1993. doi: 10.1080/09500349314551321. Cited on 31
- [67] Marlan O. Scully and M. Suhail Zubairy. *Quantum Optics*. Cambridge University Press, 1 edition, September 1997. ISBN 0521435951. Cited on 33
- [68] Herbert Walther, Benjamin T H Varcoe, Berthold-Georg Englert, and Thomas Becker. Cavity Quantum Electrodynamics. *Reports on Progress in Physics*, 69(5):1325, 2006. Cited on 35
- [69] Peter F. Herskind, Aurelien Dantan, Joan P. Marler, Magnus Albert, and Michael Drewsen. Realization of collective strong coupling with ion Coulomb crystals in an optical cavity. *Nat Phys*, 5(7):494–498, June 2009. ISSN 1745-2473. doi: 10.1038/nphys1302. Cited on 35, 86
- [70] J. P. Reithmaier, G. Sek, A. Loffler, C. Hofmann, S. Kuhn, S. Reitzenstein, L. V. Keldysh, V. D. Kulakovskii, T. L. Reinecke, and A. Forchel. Strong coupling in a single quantum dot-semiconductor microcavity system. *Nature*, 432(7014):197–200, November 2004. ISSN 0028-0836. doi: 10.1038/nature02969. Cited on 35
- [71] Simon Groblacher, Klemens Hammerer, Michael R. Vanner, and Markus Aspelmeyer. Observation of strong coupling between a micromechanical resonator and an optical cavity field. *Nature*, 460(7256):724–727, August 2009. ISSN 0028-0836. doi: 10.1038/nature08171. Cited on 35
- [72] D. E. Chang, L. Jiang, A. V. Gorshkov, and H. J. Kimble. Cavity QED with atomic mirrors, February 2012. Cited on 35
- [73] Matthias Keller, Birgit Lange, Kazuhiro Hayasaka, Wolfgang Lange, and Herbert Walther. Continuous generation of single photons with controlled waveform in an ion-trap cavity system. *Nature*, 431(7012):1075–1078, October 2004. ISSN 0028-0836. doi: 10.1038/ nature02961. Cited on 35, 86, 112
- [74] National Physical Laboratory. Radioactive elements—table of nuclides. http: //www.kayelaby.npl.co.uk/atomic_and_nuclear_physics/4_6/4_6_1_part02_ 020_029.html, 2011. Accessed on 30th November 2011. Cited on 41
- [75] N. Kjaergaard, L. Hornekaer, A.M. Thommesen, Z. Videsen, and M.Drewsen. Isotope selective loading of an ion trap using resonance-enhanced two-photon ionization. *Appl. Phys. B*, 71:207–210, May 2000. doi: 10.1007/s003400000296. Cited on 41
- [76] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Nakayama, J.-P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. Isotope-selective photoionization for Calcium ion trapping. *Phys. Rev. A*, 69:012711, Jan 2004. doi: 10.1103/PhysRevA.69.012711. Cited on 41

- [77] S. Gulde, D. Rotter, P. Barton, F. Schmidt-Kaler, R. Blatt, and W. Hogervorst. Simple and efficient photo-ionization loading of ions for precision ion-trapping experiments. *Appl. Phys. B*, 73:861–863, May 2001. doi: 10.1007/s003400100749. Cited on 41
- [78] C.A. Schrama, E. Peik, W.W. Smith, and H. Walther. Novel miniature ion traps. *Optics Communications*, 101(1-2):32 36, 1993. ISSN 0030-4018. doi: DOI:10.1016/0030-4018(93) 90318-Y. Cited on 41, 43
- [79] M. Roberts, P. Taylor, G. P. Barwood, W. R. C. Rowley, and P. Gill. Observation of the 2S1/2 2F7/2 electric octupole transition in a single 171Yb+ ion. *Phys. Rev. A*, 62(2): 020501, Jul 2000. doi: 10.1103/PhysRevA.62.020501. Cited on 41
- [80] Mark Alexander Wilson. *Quantum state control of a single trapped Strontium ion*. PhD thesis, National Physical Laboratory/University of Strathclyde, 2001. Cited on 41
- [81] N. Yu and W. Nagourney. Analysis of Paul–Straubel trap and its variations. *Journal of Applied Physics*, 77(8):3623–3630, 1995. doi: 10.1063/1.358598. Cited on 41, 42
- [82] N. Yu, W. Nagourney, and H. Dehmelt. Demonstration of new Paul–Straubel trap for trapping single ions. *Journal of Applied Physics*, 69(6):3779–3781, 1991. doi: 10.1063/1.348478. Cited on 41
- [83] K. Sheridan, W. Lange, and M. Keller. All-optical ion generation for ion trap loading. *Applied Physics B: Lasers and Optics*, 104:755–761, September 2011. doi: 10.1007/ s00340-011-4563-7. Cited on 54
- [84] K Zimmermann, M V Okhapkin, O A Herrera-Sancho, and E Peik. Laser ablation loading of a radiofrequency ion trap. Technical Report arXiv:1112.1664, CERN, Dec 2011. Comments: submitted to Appl. Phys. B., special issue on ion trapping. Cited on 54
- [85] National Physical Laboratory. What do high vacuum and low vacuum mean? http://www.npl.co.uk/reference/faqs/ what-do-high-vacuum-and-low-vacuum-mean-(faq-pressure), 2011. Accessed on 28th November 2011. Cited on 60
- [86] Crystal Technology. Application note—modulator model 3000 series. http://www. crystaltechnology.com/docs/A0_Modulator3000_appnote.pdf, 2006. Accessed on 5th December 2011. Cited on 64
- [87] D. J. McCarron. A guide to acousto-optic modulators. http://massey.dur.ac.uk/ resources/slcornish/AOMGuide.pdf, 2007. Accessed on 5th December 2011. Cited on 64
- [88] K. Dholakia, G. Zs. K. Horvath, D. M. Segal, R. C. Thompson, D. M. Warrington, and D. C. Wilson. Photon-correlation detection of ion-oscillation frequencies in quadrupole ion traps. *Phys. Rev. A*, 47:441–448, Jan 1993. doi: 10.1103/PhysRevA.47.441. Cited on 71
- [89] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *Journal of Applied Physics*, 83(10):5025–5033, 1998. doi: 10.1063/1.367318. Cited on 76
- [90] D. J. Berkeland and M. G. Boshier. Destabilization of dark states and optical spectroscopy in Zeeman-degenerate atomic systems. *Phys. Rev. A*, 65:033413, Feb 2002. doi: 10.1103/ PhysRevA.65.033413. Cited on 80

- [91] H G Barros, A Stute, T E Northup, C Russo, P O Schmidt, and R Blatt. Deterministic single-photon source from a single ion. *New Journal of Physics*, 11(10):103004, 2009. Cited on 86, 96
- [92] A Kuhn and D Ljunggren. Cavity-based single-photon sources. *Contemporary Physics*, 51 (4):289–313, July 2010. ISSN 0010-7514. doi: 10.1080/00107511003602990. Cited on 86
- [93] M Hennrich, T Legero, A Kuhn, and G Rempe. Photon statistics of a non-stationary periodically driven single-photon source. *New Journal of Physics*, 6(1):86, 2004. Cited on 86
- [94] Frank Diedrich and Herbert Walther. Nonclassical radiation of a single stored ion. *Phys. Rev. Lett.*, 58:203–206, Jan 1987. doi: 10.1103/PhysRevLett.58.203. Cited on 87
- [95] B Lounis and W.E Moerner. Single photons on demand from a single molecule at room temperature. *Nature*, 407(6803):491–493, 2000. Cited on 87
- [96] Christian Kurtsiefer, Sonja Mayer, Patrick Zarda, and Harald Weinfurter. Stable solid-state source of single photons. *Phys. Rev. Lett.*, 85:290–293, Jul 2000. doi: 10.1103/PhysRevLett. 85.290. Cited on 87
- [97] P Michler, A Kiraz, C Becher, W V Schoenfeld, P M Petroff, L Zhang, E Hu, and A Imamoglu. A quantum dot single-photon turnstile device. *Science*, 290(5500):2282–5, 2000. Cited on 87
- [98] M. Almendros, J. Huwer, N. Piro, F. Rohde, C. Schuck, M. Hennrich, F. Dubin, and J. Eschner. Bandwidth-tunable single-photon source in an ion-trap quantum network. *Phys. Rev. Lett.*, 103:213601, Nov 2009. doi: 10.1103/PhysRevLett.103.213601. Cited on 87
- [99] Philippe Grangier, Barry Sanders, and Jelena Vuckovic. Focus on single photons on demand. *New Journal of Physics*, 6(1), 2004. Cited on 87
- [100] R Hanbury Brown and R Q Twiss. Correlation between photons in two coherent beams of light. *Nature*, 177(4497):27–29, 1956. Cited on 90
- [101] R. Hanbury Brown and R. Q. Twiss. Interferometry of the intensity fluctuations in light. I. Basic theory: The correlation between photons in coherent beams of radiation. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 242(1230): 300–324, 1957. doi: 10.1098/rspa.1957.0177. Cited on
- [102] R. Hanbury Brown and R. Q. Twiss. Interferometry of the intensity fluctuations in light II. An experimental test of the theory for partially coherent light. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 243(1234):291–319, 1958. doi: 10.1098/rspa.1958.0001. Cited on 90
- [103] R. Hanbury Brown and R.Q. Twiss. a test of a new type of stellar interferomenter on Sirius. *Nature*, 178:1046–1048, 1956. Cited on 90
- [104] G.A. Rebka and R.V. Pound. Time-correlated photons. *Nature*, 180:1035–1036, 1957. Cited on 90
- [105] Roy J. Glauber. The quantum theory of optical coherence. *Phys. Rev.*, 130:2529–2539, Jun 1963. doi: 10.1103/PhysRev.130.2529. Cited on 91
- [106] Mark Fox. Quantum Optics-An Introduction. OUP, 2007. Cited on 92, 94
- [107] Gordon Baym. The physics of Hanbury Brown–Twiss intensity interferometry: from stars to nuclear collisions. ACTAPHYS.POLON.B, 29:1839, 1998. Cited on 92

- [108] Surendra Singh. Antibunching, sub-Poissonian photon statistics and finite bandwidth effects in resonance fluorescence. *Optics Communications*, 44(4):254 – 258, 1983. ISSN 0030-4018. doi: 10.1016/0030-4018(83)90132-3. Cited on 93
- [109] X. T. Zou and L. Mandel. Photon-antibunching and sub-Poissonian photon statistics. *Phys. Rev. A*, 41:475–476, Jan 1990. doi: 10.1103/PhysRevA.41.475. Cited on 93
- [110] H. J. Kimble, M. Dagenais, and L. Mandel. Photon antibunching in resonance fluorescence. *Phys. Rev. Lett.*, 39:691–695, Sep 1977. doi: 10.1103/PhysRevLett.39.691. Cited on 93
- [111] M. Dagenais and L. Mandel. Investigation of two-time correlations in photon emissions from a single atom. *Phys. Rev. A*, 18:2217–2228, Nov 1978. doi: 10.1103/PhysRevA.18.2217. Cited on 93, 94
- [112] R. Short and L. Mandel. Observation of sub-Poissonian photon statistics. *Phys. Rev. Lett.*, 51:384–387, Aug 1983. doi: 10.1103/PhysRevLett.51.384. Cited on 93
- [113] H J Carmichael and D F Walls. Proposal for the measurement of the resonant Stark effect by photon correlation techniques. *Journal of Physics B: Atomic and Molecular Physics*, 9(4): L43, 1976. Cited on 93
- [114] Alex Wilson, Hiroki Takahashi, Andrew Riley-watson, Fedja Oru, Peter Blythe, Anders Mortensen, Daniel R Crick, Nicolas Seymour-smith, Elisabeth Brama, Matthias Keller, and et al. Fiber-coupled single ion as an efficient quantum light source. *Population English Edition*, quant-ph:0–3, 2011. URL http://arxiv.org/abs/1101.5877. Cited on 102
- [115] Genko S Vasilev, Daniel Ljunggren, and Axel Kuhn. Single photons made-to-measure. New Journal of Physics, 12(6):4, 2009. Cited on 112
- [116] Peter Nisbet, Jerome Dilley, and Axel Kuhn. Highly efficient source for indistinguishable photons of controlled shape. *New Journal of Physics*, 13(10):103036, 2011. Cited on 112
- [117] F Laguarta, N Lupon, and J Armengol. Optical glass polishing by controlled laser surfaceheat treatment. *Applied Optics*, 33(27):6508–6513, 1994. Cited on 114
- [118] E Mendez, K M Nowak, H J Baker, F J Villarreal, and D R Hall. Localized CO₂ laser damage repair of fused silica optics. *Applied Optics*, 45(21):5358–5367, 2006. Cited on 114
- [119] U C Paek and A L Weaver. Formation of a spherical lens at optical fiber ends with a CO₂ laser. Applied Optics, 14(2):294–298, 1975. Cited on 114
- [120] D Hunger, C Deutsch, R J Barbour, R J Warburton, and J Reichel. Laser micro-fabrication of concave, low-roughness features in silica. AIP Advances, 2(1):3–6, 2011. Cited on 114, 120
- [121] Qian Kemao and Seah Hock Soon. Sequential demodulation of a single fringe pattern guided by local frequencies. *Optics Letters*, 32(2):127–129, 2007. Cited on 121
- [122] Q Kemao. Two-dimensional windowed Fourier transform for fringe pattern analysis: Principles, applications and implementations. *Optics and Lasers in Engineering*, 45(2):304–317, 2007. Cited on 124
- [123] J M Bioucas-Dias and G Valadao. Phase unwrapping via graph cuts. IEEE Transactions on Image Processing, 16(3):698–709, 2007. Cited on
- [124] M Takeda, H Ina, and S Kobayashi. Fourier transform method of fringe pattern analysis for computer based topography and interferometry. *Journal of the Optical Society of America*, 72:156–160, 1982. Cited on 121

- [125] D Hunger, T Steinmetz, Y Colombe, C Deutsch, T W Hnsch, and J Reichel. A fiber Fabry-Pérot cavity with high finesse. *New Journal of Physics*, 12(6):065038, 2010. Cited on 130, 133, 134, 135
- [126] J Stone. Optical-fibre Fabry-Pérot interferometer with finesse of 300. *Elect. Lett.*, 21: 504–505, 1985. Cited on 130
- [127] T. Yoshino, K. Kurosawa, K. Itoh, and T. Ose. Fiber-optic Fabry-Pérot interferometer and its sensor applications. *Quantum Electronics, IEEE Journal of*, 18(10):1624–1633, 1982. Cited on 130
- [128] T. Steinmetz, Y. Colombe, D. Hunger, T. W. Hänsch, A. Balocchi, R. J. Warburton, and J. Reichel. Stable fiber-based Fabry-Pérot cavity. *Applied Physics Letters*, 89(11):11110, 2006. doi: 10.1063/1.2347892. Cited on 130
- [129] J Stone and D Marcuse. Ultrahigh finesse fiber Fabry-Pérot interferometers. IEEE J Lightwave Technol, 4(4):382–385, 1986. Cited on 130
- [130] A. Siegman. Lasers. University Science Books, 1st edition edition, October 1990. ISBN 0935702113. Cited on 131
- [131] Huanqian Loh, Yu-Ju Lin, Igor Teper, Marko Cetina, Jonathan Simon, James K. Thompson, and Vladan Vuletić. Influence of grating parameters on the linewidths of external-cavity diode lasers. *Appl. Opt.*, 45(36):9191–9197, Dec 2006. doi: 10.1364/AO.45.009191. Cited on 145
- [132] Nicolas Seymour-Smith, Peter Blythe, Matthias Keller, and Wolfgang Lange. Fast scanning cavity offset lock for laser frequency drift stabilization. *Review of Scientific Instruments*, 81 (7):075109, 2010. doi: 10.1063/1.3455830. Cited on 145

Colophon

This thesis was typeset in 11 point Times using the TeX Live distribution of LATEX a PC running Ubuntu 12.04 LTS. Emacs v23.3.1 with Auctex was used as an editor. All graphs were produced using MATLAB R2010a. Finite-element simulations were performed using Femlab v3 controlled by MATLAB. Drawings and diagrams were produced using Inkscape v0.48 and Open Office Draw. SolidWorks 2010-2011 was used for CAD, and 3D drawings. The text contains approximately 30,000 words.

"Then I started feeling bad for the photon, and I said maybe it wanted to continue but I got in its way. But then I said, no, those are probably happier photons than the one that slammed into the mountainside that will go unanalyzed and will not contribute to the depth of our understanding of the universe."

Neil DeGrasse Tyson Beyond Belief: Science, Reason, Religion and Survival, Salk Institute for Biological Studies, November 7, 2006