

University of Nottingham
UK | CHINA | MALAYSIA

Interfacing cold atoms with optical waveguides for quantum technologies

Thesis submitted to the University of Nottingham for the degree of
Doctor of Philosophy, December 19, 2025.

Matt Overton

Supervised by

Lucia Hackermüller

School of Physics and Astronomy

University of Nottingham

Abstract

Cold atoms are very useful for a wide range of applications in quantum technologies. One property of neutral atoms that is of particular interest is their strong, coherent interactions with light. Interfacing cold atoms with integrated photonics is an area of key interest, to which we offer a solution utilising micro-scale through-holes in an optical waveguide. We trap about 400 atoms in our interaction region, achieving transmissions as low as 0.6%, corresponding to an optical depth of 1700/cm. Here I present 2 experiments done with cold atoms in this atom-photon interface.

The first is a slow light experiment utilising electromagnetically induced transparency (EIT). We achieve a photon group velocity reduction of $\sim 10^5$, a velocity of 2000 m/s. This phenomenon can also be used for photon memory, with applications in quantum computing and quantum communications. Work is currently being done to achieve this goal.

The second experiment utilises Bayesian statistics to improve our estimation efficiency. We experimentally observe a 5-fold precision enhancement in our atom number estimates using the same number of measurements. Equivalently, we could achieve the same precision as our standard method using 1/3 as many measurements. This technique is not only useful for this experiment, but can be widely applied to many areas of quantum technologies.

Acknowledgements

There are many people I'd like to thank for their support up to and throughout my PhD. My family, friends, and partners for being with me whenever I needed them. Emory for giving me the help I need even when I am unable to ask for it. Dr. Lucia Hackermüller in her role as my supervisor offering advice and direction to my work. Dr. Nathan Cooper with all his experimental help and oversight. The members of the my research group: Dr. David Johnson, Ben Hopton, Danielle Baldolini, Alexander Abbey, Ayelen Paez, for all the helpful discussions and assistance. The wider Nottingham cold atom group: Dr. Thomas Fernholz, Dr. Vilius Atkočius, Dr. Josh New, Dr. Rhys Morrison, Dr. Freya Owen, Dr. Ali Akbar, Roshan Giji, Ellis Hewins, for the entertaining lunch time discussions. Dr. Chris Overton, Dr. James Thompson, Dr. Richard Hussey, for the advice when I was considering starting this PhD. Jan Rogozinski for sparking and inspiring my love of science.

Declaration of Authorship

I, Matthew James Overton, declare that this thesis titled "Interfacing cold atoms with optical waveguides for quantum technologies" and the work presented in it are my own. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University.
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated.
- Where I have consulted the published work of others, this is always clearly attributed.
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work.
- I have acknowledged all main sources of help.
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself.

Contents

Abstract **i**

Chapter 1 Introduction **1**

1.1 Overview 3

1.2 Publications 5

Chapter 2 Background theory **7**

2.1 Electron structure 7

2.1.1 Caesium 11

2.2 Laser cooling and trapping 12

2.2.1 Magneto-optical trap 17

2.2.2 Dipole force 19

Chapter 3 Experimental techniques **21**

3.1 Laser locking 22

3.1.1 Spectroscopy 22

3.1.2 Frequency lock 24

3.1.3 Phase lock 26

3.2 Imaging 35

3.2.1 Absorption imaging 36

3.2.2 Absorption in the fibre junction 38

3.3 Temperature measurement 41

Chapter 4 Setup and experimental overview **43**

4.1 Apparatus 44

4.1.1 Lasers 44

4.1.2 Magnetic fields 47

4.1.3 Fibre junction 48
4.1.4 Experimental sequence 49
4.2 Magneto-optical trap 52
4.3 Dipole trap 55
4.4 Fibre absorption 63
Chapter 5 Electromagnetically induced transparency 72
5.1 EIT theory 73
5.1.1 Susceptibility 76
5.2 Experimental results 81
5.2.1 Slow light 84
5.2.2 Next steps 86
Chapter 6 Bayesian measurement strategies 87
6.1 Bayesian methods 88
6.1.1 Optimal estimators 90
6.1.2 Adaptive strategies 92
6.2 Our system 95
6.2.1 Maximum ignorance prior 97
6.2.2 Optimal estimator 98
6.2.3 Computational calculation 100
6.3 Results 102
Chapter 7 Conclusion 107
7.1 Outlook and next steps 108
Bibliography 111
Appendices 128
Appendix A Publications 129
Appendix B Optics layout 152
Appendix C Phase lock circuits 156
Appendix D Adaptive Bayesian code 158

List of Figures

1.1 Simplified diagram of the experiment 2

2.1 Caesium hyperfine structure 11
2.2 Zeeman splitting in a 1D MOT 18

3.1 Doppler-free spectroscopy optics 22
3.2 Laser spectroscopy 24
3.3 Offset lock circuit 25
3.4 Schematic of the optical phase lock 27
3.5 Phase frequency detector timing diagram 29
3.6 Feedback response strength 31
3.7 Unlocked beat note 33
3.8 Phase locked beat note 33
3.9 Zoomed in phase locked beat note 34

4.1 Laser frequency scheme 45
4.2 Image and diagram of the fibre junction 49
4.3 Experimental sequence 50
4.4 MOT images 52
4.5 Big MOT image 53
4.6 MOT temperature 54
4.7 Dipole beam waist analysis 56
4.8 Images of the dipole trap 57
4.9 Optical pumping 58
4.10 Dipole trap temperature 60

4.11 Lifetime of atoms in the dipole trap 61
4.12 Dipole trap aligned through the fibre 62
4.13 Fibre absorption spectrum 65
4.14 Lifetime in the fibre junction 66
4.15 Temperature in fibre junction 68
4.16 Absorption spectrum on the $|F = 4\rangle \rightarrow |F' = 4\rangle$ transition. . 70

5.1 Energy level diagrams for different EIT schemes. 73
5.2 Expected EIT transmittance for no dephasing. 78
5.3 Expected EIT transmittance with the effects of dephasing. . 78
5.4 EIT transmittance curves 82
5.5 Real susceptibility estimate 84
5.6 Slow light 85

6.1 Demonstration of the adaptive strategy 101
6.2 Atom number estimates for different measurement strategies 103
6.3 Noise-to-signal ratio of different measurement strategies . . . 106

B.1 Optical layout of the main optics table. 153
B.2 Optical layout of the probe and control lasers. 154
B.3 Optical layout around the vacuum chamber 155

List of Tables

4.1 Laser system overview 46

4.2 Transmittance vs power and exposure time 68

5.1 Transparency curve fitting parameters 81

6.1 Measurement strategy comparison 105

Abbreviations

AOM acousto-optical modulator.

ECDL external-cavity diode laser.

EIT electromagnetically induced transparency.

MLE maximum likelihood estimator.

MOT magneto-optical trap.

NSR noise-to-signal ratio.

PDF probability density function.

PFD phase frequency detector.

SPCM single photon counting module.

STIRAP stimulated Raman adiabatic passage.

VCO voltage controlled oscillator.

Chapter 1

Introduction

Cold atoms are a useful platform for many quantum technological applications, from sensing of gravity [1,2], acceleration [3,4], and magnetic fields [5,6], to quantum communication [7–9] and quantum computing [10–12]. However, cold atom systems tend to be large and utilise free space optics, making them hard to integrate into distributed systems, such as a quantum communication network [7,13]. For this reason, miniaturisation of these systems is a large area of active research, exploring techniques such as atom chips [14,15], advanced manufacturing [16–18], and the integration of cold atom systems with photonic circuits [19–26]. Investigating one realisation of how cold atoms can be integrated with photonic circuits constitutes the bulk of this thesis.

Several different approaches to this have been explored. One technique involves utilising the large evanescent field of an optical nanofibre (an optical fibre that has been stretched so its diameter is $\sim 300\,\text{nm}$) [27–32]. Other groups have explored transporting/trapping atoms in hollow core fibres [22,33–37]. Atoms have also been interfaced with waveguides, utilising either micro-trenches or evanescent fields [25,38–41]. Here, we build upon

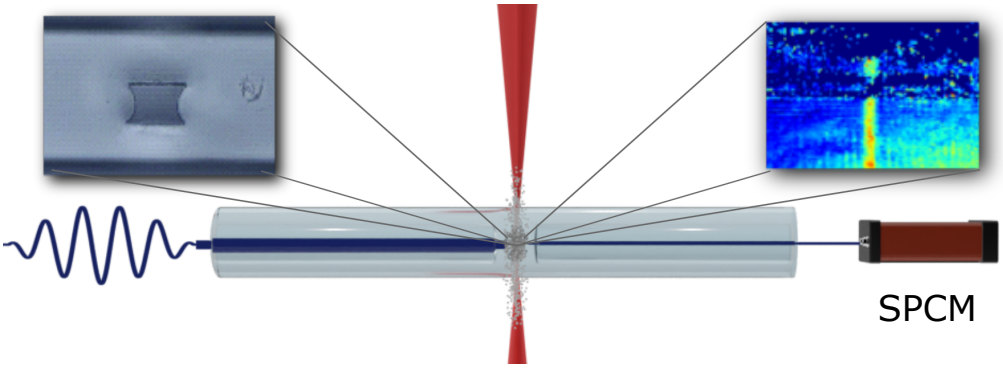

Figure 1.1: Explanatory diagram of the experimental principle. Atoms are loaded into a transverse hole through and optical fibre using an optical dipole trap. These atoms are probed by light travelling along that fibre, detected by a single photon counting module (SPCM). Left inset shows a microscope image of the fibre junction, right inset shown an atomic absorption image of atoms loaded into the junction. Figure is adapted from Ref. [46].

a technique developed earlier in the research group, utilising laser-drilled through-holes in an optical fibre [42–45].

For our system, we have a single-mode fibre running through a vacuum chamber. This fibre has a 30 µm hole drilled through it, transverse to the axis of the fibre core. Atoms can be trapped in this hole using an optical dipole trap, achieving high densities and a very large optical depth per unit length. While we demonstrate this system in an optical fibre, the core principle is very flexible and could be expanded to waveguide chips. The techniques could even be expanded to include multiple nodes on a single chip.

One key area of interest for quantum computing and quantum communication is in photon memory [7,47–49], mapping photonic quantum information into a stationary medium so that it can be retrieved again at a later desired time. There are several ways to implement a quantum memory, such as waveguide cavities [50], off-resonance cascaded absorption [51], and Raman transitions [52,53]. Here we develop a technique using electromag-

netically induced transparency (EIT), which has been demonstrated before in cold and warm atomic vapours [54–58]. It has also been demonstrated in both nanofibre [59,60] and hollow-core fibre systems [61–63]. We successfully demonstrate slow light in our fibre junction, achieving a photon group velocity of 2000 m/s. Work is still being done to further improve upon this and demonstrate photon storage.

Noise is a common obstacle to overcome in many quantum technology applications [12,18,64–66]. As an example, photon shot noise is a major limiting factor in the experiment detailed in this thesis. So it is important to be able to rapidly and efficiently estimate parameters from a noisy signal. Bayesian statistics offers a powerful toolbox for getting these estimates [67–70]. Here, we propose and demonstrate an adaptive Bayesian strategy where we optimally change how we perform a measurement while the measurement is being performed. This is demonstrated in our fibre junction however, the technique is generally applicable to a wide range of experiments, expanding upon previous work so it is not only applicable to phase estimation [71–74], but also allows estimates of location, scale, and weight parameters [46,75–78].

1.1 Overview

This thesis covers all the work done with this atom-fibre junction over the duration of my PhD. Most of the work on this experiment was done by myself; however, at various times, assistance was given by Dr. Nathan Cooper, Dr. David Johnson, Daniele Baldolini, and Alexander Abbey. The work done builds upon the work of previous researchers in the group, who built large parts of the experimental setup, notably Dr. Elisa Da Ros [42],

Dr. David Johnson [43], and Dr. Nathan Cooper. Additionally, the work on the adaptive Bayesian estimation strategy was done in collaboration with Dr. Jesús Rubio Jiménez, who developed the statistical method.

The thesis is structured as follows:

Chapter 2: Background theory. This chapter introduces the main theoretical concepts used in any cold atom experiment. It covers the hyperfine structure of alkali metals, with a focus a caesium, as that is the species used in this experiment. We then cover 2-level atom-photon interactions, to explain the principles of laser cooling and trapping.

Chapter 3: Experimental principles. This chapter covers the ideas and theory behind various experimental techniques utilised in this experiment. We cover atomic spectroscopy and various methods for laser frequency stabilisation. Particular detail is given to optical phase-locking, covering a phase-lock I designed and built during my PhD. We also cover the principles behind how we image atoms in various parts of our system, and how we can extract information about the temperature of the trapped atoms from these images.

Chapter 4: Experimental overview. This chapter details the experimental apparatus used in this experiment. It also covers results from earlier stages of the experiment and how they need to be optimised to enable the slow light experiment.

Chapter 5: Electromagnetically induced transparency. The theory of 3-level atom-photon interactions is covered to explain the physical origin of this phenomenon and how it leads to slow and stored light. Then our experimental results are presented, showing the transparency peaks and the reduction in group velocity.

Chapter 6: Bayesian measurement strategies. This chapter gives a quick introduction to Bayesian statistics and demonstrates the derivation of the adaptive protocol and how it is applied to our system. We then give an overview of how we perform these computations quickly and show our experimental results.

Chapter 7: Conclusion. Finally, this chapter concludes this thesis, giving an overview of the results, and discusses various technical improvements and outlines potential directions for future experiments with the system.

1.2 Publications

Over the course of my PhD, I have co-authored the following papers:

- M. Overton, J. Rubio, N. Cooper, D. Baldolini, D. Johnson, J. Anders, and L. Hackermüller. Substantial precision enhancements via adaptive symmetry-informed Bayesian metrology. *arXiv preprint arXiv:2410.10615, 2024.*
- M. Overton, A. Abbey, D. Johnson, D. Baldolini, N. Cooper, L. Hackermüller. Slow light in an atom-photon fibre junction. *In Preparation.*
- N. Cooper, D. Johnson, B. Hopton, M. Overton, D. Stupple, A. Bratu, E. Wilson, J. Robinson, L. Coles, M. Papastavrou, L. Hackermüller. Exploiting complex 3D-printed surface structures for portable quantum technologies *In Preparation.*
- D. Baldolini, N. Cooper, D. Johnson, A. Perez, M. Overton, L. Hackermüller. Shortcut to adiabaticity in ^{6}Li Feshbach molecule conversion. *In Preparation.*

- B. Hopton, D. Johnson, N. Cooper, M. Overton, A. Abbey, L. Hackermüller. Multi-frequency light for enhanced MOT formation. *In Preparation.*

The first 2 papers constitute the main results of this thesis, and are further expanded on in Chapters 5 & 6.

Chapter 2

Background theory

Laser cooling and trapping techniques are essential for all cold atom experiments. The development of these techniques has enabled this wide field of physics, offering significant improvements in metrology [79,80], sensing [1,3], and the realisation of novel technologies like quantum computing [11,12]. In this chapter, I give an introduction to some of the theory that is core to this field. I start by covering the fine and hyperfine structure of atoms (Sec. 2.1), before covering 2 level atom-light interactions and how they lead to optical cooling and trapping (Sec. 2.2).

2.1 Electron structure

The electron structure of atoms is important to understand for cold atom physics. The lasers used in these experiments need to target not only individual hyperfine transitions, but also individual transitions between the Zeeman sub-levels. So an understanding of this structure is important. The vast majority of cold atom experiments (this one included) use alkali metals, which have a hydrogenic structure (i.e. 1 valence electron). This

means multi-electron effects do not need to be considered.

Electrons in an atom are arranged in shells and orbitals determined by a principal quantum number n and orbital angular momentum ℓ (however other quantum numbers such as the spin of the electron and nucleus also come into play). At low-level approximation, the energy of these levels depends only on n. However, as will be discussed, there is splitting due to the total electron angular momentum and the spin of the electron and nucleus.

Fine Structure

The fine structure comes from the coupling of the electron's orbital angular momentum with its own spin [81–83]. Due to the motion of the electron through the nucleus's electric field, the electron experiences an effective magnetic field. The interaction between this magnetic field and the electron's intrinsic magnetic moment from its spin shifts the binding energy of the electron.

The effective magnetic field (from the electron's reference frame) is given by

$$\mathbf{B} = -\frac{Ze^2}{4\pi\epsilon_0}\frac{1}{m_e c^2 e r^3}\mathbf{L} \tag{2.1}$$

where $\mathbf{L}$ is the electron's orbital angular momentum. The electron's magnetic moment is given by

$$\mu_{\mathbf{s}} = -\mu_B g_s \frac{\mathbf{S}}{\hbar} \tag{2.2}$$

where μ_B is the Bohr magnetron, g_s is the electron g-factor, and $\mathbf{S}$ is the electron spin operator. The Hamiltonian for the fine structure (after

transforming back into the lab frame) is then given by

$$H_{fs} = -\mu_{\mathbf{s}} \cdot \mathbf{B} = \frac{Ze^2}{4\pi\epsilon_o} \frac{g_s - 1}{2m_e^2 c^2 r^3} \mathbf{L} \cdot \mathbf{S} \,. \tag{2.3}$$

To evaluate this, we introduce a composite quantum number

$$\mathbf{J} = \mathbf{L} + \mathbf{S} \tag{2.4}$$

which has a value $|L - S| \leq J \leq L + S$. This allows the Hamiltonian to be expressed in terms of the standard operators as

$$\mathbf{L} \cdot \mathbf{S} = \frac{1}{2} \left(\mathbf{J}^2 - \mathbf{L}^2 - \mathbf{S}^2 \right) \,. \tag{2.5}$$

This fine structure leads to the splitting of different L and S states of the electron due to the electron's total angular momentum J.

It is worth noting that there is another energy shift that also arises from the electron's angular momentum. This shift is due to relativistic effects on the electron's mass. This relativistic term plays a significant role in the energy levels of hydrogen; however, for larger atoms, this effect is less significant, and so we do not go into detail with them here.

Hyperfine structure

A further splitting of the energy levels comes from the interaction between the total electron angular momentum **J** and the nuclear angular momentum **I**. The intrinsic magnetic moment of the nucleus interacts with the magnetic field induced by the electron's motion. This shifts and splits the fine structure states based on the relative orientation of these fields. This leads to what is called the hyperfine structure [81–83].

The nuclear magnetic moment is

$$\mu_{\mathbf{I}} = \mu_B g_I \frac{\mathbf{I}}{\hbar} \tag{2.6}$$

where $\mathbf{I}$ is the nuclear spin operator, and g_I is the nuclear spin g-factor. The magnetic field induced by the electron's motion is

$$\mathbf{B} = -b\mathbf{J} \tag{2.7}$$

where b is a positive constant. So, the Hamiltonian for the hyperfine splitting is then given by

$$H_{hfs} = -\mu_{\mathbf{I}} \cdot \mathbf{B} = A_{hfs} \frac{\mathbf{I} \cdot \mathbf{J}}{\hbar^2} \tag{2.8}$$

where $A_{hfs} = -\mu_B g_I b\hbar$ is the magnetic dipole hyperfine constant.

This can be treated in a similar way to the fine structure splitting, by defining a new total atomic angular momentum quantum number

$$\mathbf{F} = \mathbf{I} + \mathbf{J} \tag{2.9}$$

such that

$$\mathbf{I} \cdot \mathbf{J} = \mathbf{F}^2 - \mathbf{I}^2 - \mathbf{J}^2 \,. \tag{2.10}$$

The energy shift from this hyperfine splitting is

$$\Delta E_{hfs} = \frac{1}{2} A_{hfs} \left[F(F+1) - I(I+1) - J(J+1) \right] \,. \tag{2.11}$$

There are higher-order corrections to this model that account for the electric quadrupole and magnetic octupole interactions. However, these are much smaller and not relevant for the physics at play and as such are not

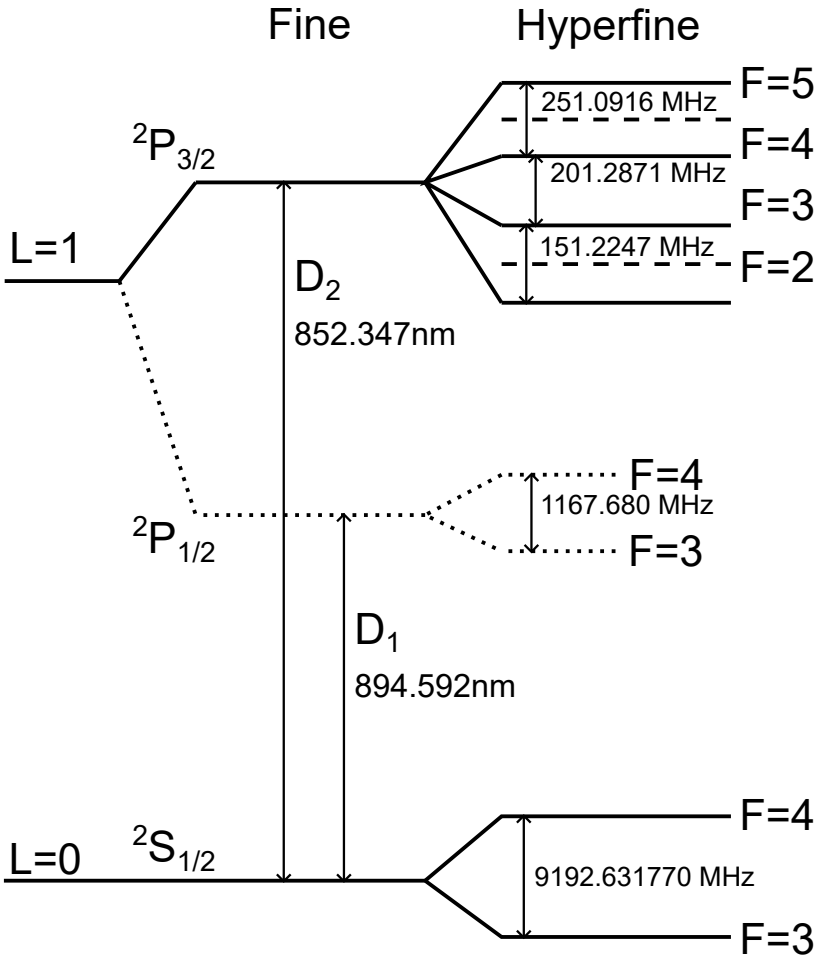

Figure 2.1: A diagram showing the fine and hyperfine splitting on the D line of caesium-133. The data are taken from [84]. In this experiment, only the caesium D_2 line is used. The dashed lines in the hyperfine structure show the cross-over peaks that the lasers are locked to.

included in this thesis.

2.1.1 Caesium

Caesium-133 is the main atomic species used for this work, for this reason its specific electron structure will be covered in more depth. Caesium has a nuclear spin of $I = 7/2$. Its ground state ($6^2S_{1/2}$) has $L = 0$ so $J = 1/2$ and it is split into 2 hyperfine states with $F = 3$ and $F = 4$. This hyperfine splitting has a frequency difference of exactly 9.192 631 770 GHz and is used as the current definition of the second [84]. The excited state with $L = 1$ is

split into 2 fine states with $J = 1/2$ ($6^2P_{1/2}$) and $J = 3/2$ ($6^2P_{3/2}$). These are then split into 2 ($F = 3, 4$) and 4 ($F = 2, 3, 4, 5$) hyperfine states, respectively.

Transitions from the ground state to this excited state ($L = 0 \rightarrow L = 1$) are D line transitions, with $6^2S_{1/2} \rightarrow 6^2P_{1/2}$ being D_1 and $6^2S_{1/2} \rightarrow 6^2P_{3/2}$ being D_2. All transitions used in this work are on the ^{133}Cs D_2 line. A diagram of the different hyperfine states on the caesium D line can be seen in Figure 2.1.

2.2 Laser cooling and trapping

One can consider the optical force experienced by an atom as having 2 main components: the dipole force, related to the gradient of the beam intensity; and the scattering force, a resonant effect to do with the absorption and emission of photons. For laser cooling, the main component to consider is the scattering force. The dipole force will be explored more in section 2.2.2.

When an atom absorbs a photon of frequency ω, an electron is excited to a higher energy state. As the photon has momentum, the atom also receives a change in momentum of $\hbar k$ in the direction of the photon propagation (where k is the wave-number of the incident photon) [81,85]. When the excited electron later decays back to the ground state, another photon is released. This emission will also cause a change in the atom's momentum of the same magnitude, but not necessarily in the same direction. Emission of light from an atom can occur in 2 ways: stimulated emission and spontaneous emission.

Stimulated emission is when the excited electron interacts with the mode

of an incident light field. When the excited state decays, the photon is emitted into that same field. Spontaneous emission is independent of any external field, the photon is emitted into a random mode/direction. The rate of spontaneous emission depends only on the population of the excited state. As stimulated emission is a reversible process (it is the reverse of absorption), it produces a conservative force, so is irrelevant for the process of laser cooling. Spontaneous emission, on the other hand, is irreversible (and hence dissipative) and plays a key role in laser cooling [85].

The photon generated from a spontaneous emission event propagates in a random direction, so the net change in momentum over many emissions is 0. However, the net change in momentum from absorption does not average to 0. Therefore, over many absorption-emission cycles, the atoms experience a scattering force acting in the direction of the incoming photons [42,85,86].

This force of photon absorption is given by

$$\mathbf{F_{sc}} = \hbar \mathbf{k} R_{ab} \tag{2.12}$$

where R_{ab} is the rate at which photons are absorbed. In the steady state, the rate of absorption equals the rate of emission $R_{ab} = R_{sp} = \Gamma \rho_{ee}$. Here, Γ is a rate factor, and ρ_{ee} is the population of the excited state. Now the scattering force can be written as

$$\mathbf{F_{sc}} = \hbar \mathbf{k} \Gamma \rho_{ee} = \frac{\hbar \Gamma}{2} \frac{s_0}{1 + s_0 + (2\delta/\Gamma)^2} \mathbf{k} \tag{2.13}$$

where $\delta = \omega - \omega_0$ is the detuning of the light from the resonant frequency of the transition, and $s_0 = 2(|\Omega|/\Gamma)^2 = I/I_{sat}$ is the on-resonance saturation parameter (where Ω, I, I_{sat} are the Rabi frequency, intensity, and saturation intensity respectively).

This equation is valid for atoms at rest. To expand it for moving atoms, one has to account for the Doppler shift ω_D caused by the atom's velocity $\mathbf{v}$. This means that the detuning is given by

$$\delta' = \delta + \omega_D = \delta - \mathbf{k} \cdot \mathbf{v} \tag{2.14}$$

which adds a velocity dependence to Eq. 2.13.

Take the case where the incident light is red-detuned ($\delta < 0$). For atoms moving towards the laser source, the frequency will be shifted upwards towards resonance, thus increasing the force experienced by these atoms. For atoms moving away from the laser source, the frequency will be shifted away from resonance, and the force experienced will decrease. The net result of this is a force that opposes motion towards the source of the beam. So, 2 counter-propagating beams (or a standing wave) generate a force that opposes all motion in that axis.

This technique is called optical molasses [85,86]. The force here is simply the summation of the forces of the 2 beams,

$$\mathbf{F}_{OM} = \mathbf{F}_{sc}(\delta - \mathbf{k} \cdot \mathbf{v}) - \mathbf{F}_{sc}(\delta + \mathbf{k} \cdot \mathbf{v}). \tag{2.15}$$

In the low velocity limit ($\mathbf{k} \cdot \mathbf{v} \ll \Gamma$), the scattering force can be approximated as

$$\mathbf{F}_{sc}(\delta - \mathbf{k} \cdot \mathbf{v}) \approx \mathbf{F}_{sc}(\delta) - \mathbf{k} \cdot \mathbf{v} \, \frac{\mathrm{d}\mathbf{F}_{sc}(\delta)}{\mathrm{d}\delta}. \tag{2.16}$$

If we assume our laser is directed along the z axis, this leads to the optical molasses force being given by [85]

$$\mathbf{F}_{OM} \approx \frac{8\hbar s_0 \delta \mathbf{k}}{\Gamma\big(1 + s_0 + (2\delta/\Gamma)^2\big)^2} \, \mathbf{k} \cdot \mathbf{v} = -\beta(\delta) v_z \, \hat{\mathbf{z}} \, , \tag{2.17}$$

where v_z is the velocity of the atoms in the z-axis. By having 6 beams in 3 orthogonal axes, the optical molasses force opposes motion in all directions, slowing (and hence cooling) the atoms.

Doppler Limit

Based purely on equation 2.17, it would seem like the atoms would quickly be slowed to $v = 0$, and hence $T = 0$K. In reality, while the temperature can get very low, there is also some heating caused by the light. This heating originates from the fact that the atom can only make discrete momentum steps.

For each absorption and emission event, the atom's momentum changes by $\hbar k$, which corresponds to a recoil energy of $E_r = \hbar^2 k^2/2m = \hbar\omega_r$ (where m is the mass of the atom). This means the average frequency of each absorption process is $\omega_{abs} = \omega_0 + \omega_r$, and each emission is $\omega_{em} = \omega_0 - \omega_r$. So, for each scattering event, the light field loses (and the atom gains) an average energy of $2\hbar\omega_r$. A single beam causes a scattering rate of $R_{ab} = \Gamma\rho_{ee}$ (as described by Eqs. 2.12 & 2.13), so 2 beams would have a rate of $2\Gamma\rho_{ee}$. This leads to a heating rate of

$$\frac{\mathrm{d}E_{heat}}{\mathrm{d}t} = 4\hbar\omega_r\Gamma\rho_{ee}\,. \tag{2.18}$$

The optical molasses force has a cooling rate of

$$\frac{\mathrm{d}E_{cool}}{\mathrm{d}t} = F_{OM}v = -\beta(\delta)v^2\,. \tag{2.19}$$

The Doppler limit is reached when the two rates are at equilibrium, which

happens when the velocity equals

$$v_{eq}^2 = \frac{\hbar\omega_r\Gamma\rho_{ee}}{\beta(\delta)}\,, \tag{2.20}$$

which corresponds to a temperature

$$k_B T = m v_{eq}^2 = \frac{\hbar\Gamma}{4}\left(\frac{2|\delta|}{\Gamma} + \frac{\Gamma}{2|\delta|}\right). \tag{2.21}$$

The equation has a minimum when $|\delta| = \Gamma/2$. This minimum is called the Doppler temperature, and is given by [86]

$$T_D = \frac{\hbar\Gamma}{2k_B}\,. \tag{2.22}$$

This temperature limit does not depend on the wavelength, atomic mass, or laser intensity, only on the rate (linewidth) of the transition. The Doppler temperature is 240 µK for ^{23}Na and 124 µK for ^{133}Cs [85].

Sub-Doppler cooling

The first precise temperature measurements of a cooled cloud of sodium showed that the temperature was 40 µK, well below the Doppler cooling limit [87]. The theory of Doppler cooling approximates the atoms as a 2-level system. To understand sub-Doppler cooling, one needs to account for the m_F splitting of the ground state, and how those sub-levels are affected by the light field.

Sisyphus cooling [85] involves the atoms constantly climbing a potential hill in their ground state, before the electron is excited and pumped into the lower state. This means the energy of the photon emitted is greater than the energy absorbed. These potential hills arise from the polarisation of

the light (circular polarisation causes the m_F levels to split). Two counter-propagating linear beams create a polarisation gradient that shifts between σ^+ to linear to σ^- and so on; this gradient causes the m_F splitting to shift sinusoidally, which creates these potential hills.

Another mechanism of sub-Doppler cooling is $\sigma^+ - \sigma^-$ polarisation gradient cooling [85]. If the 2 counter-propagating beams are circularly polarised (as is the case in a MOT), the polarisation gradient generated is linear everywhere, but changes direction like a corkscrew. This doesn't generate any splitting in the sub-levels (and therefore the Sisyphus mechanism described above doesn't work), but imbalances in the m_F levels affect how strongly the atoms interact with the light. Atoms at rest are pumped to the $m_F = 0$ sub-level, but atoms travelling towards the σ^+ (σ^-) beam get pumped into positive (negative) m_F states. Atoms in a positive m_F state interact with the σ^+ beam much more strongly than the σ^- (with the inverse true for $m_F < 0$). This creates an additional damping force, but one that doesn't depend on the Doppler shift of the light, just the imbalance in the m_F states as the atom moves.

2.2.1 Magneto-optical trap

The scheme described above can cool the atoms, however, it does not confine them. The atoms are able to drift outside the laser beams and be lost. The most common way to both cool and trap the atoms is to use counter-propagating beams (like for optical molasses) in a non-uniform magnetic field to create what is known as a magneto-optical trap (MOT). The scattering force from the light slows the atoms as before, but now the magnetic field adds a spatially dependent splitting to the atomic energy levels such that the scattering force also confines the atoms [42,85].

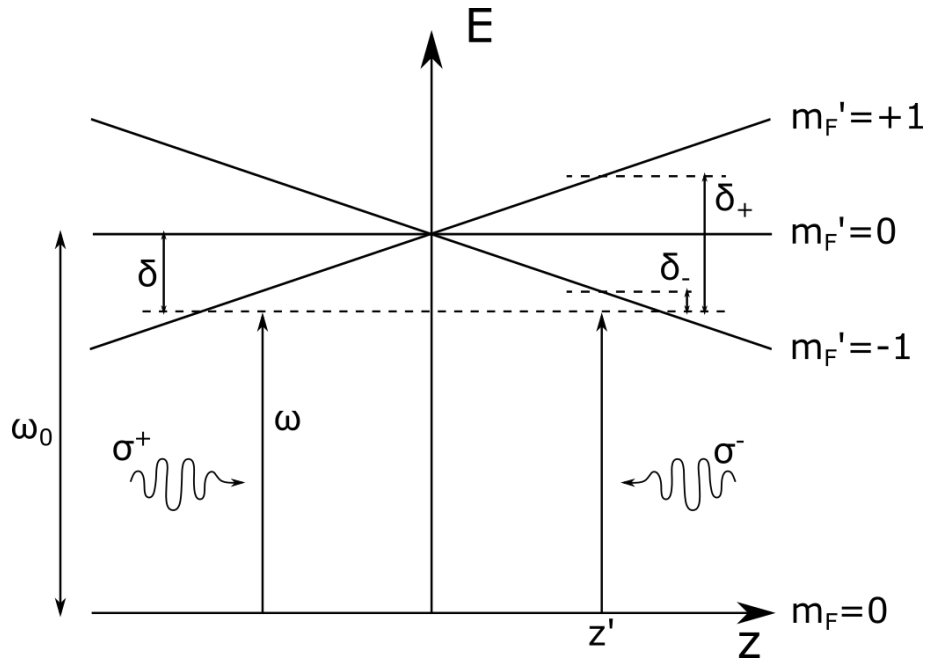

Figure 2.2: Energy level diagram for a 1D MOT. For an atom at $z' > 0$, $\delta_- < \delta_+$ so the $\Delta m = -1$ transition is closer to the frequency of the laser, meaning more σ^- light will be scattered by the atom (the inverse is true for $z' < 0$). This leads to a net force that pushes the atoms to the centre of the trap.

The non-uniform field is generated using a pair of coils in the reverse Helmholtz configuration. This generates an approximately linear magnetic field gradient around the centre of the coils (with $\mathbf{B} = 0$ at the centre). The Zeeman effect causes the hyperfine F states to split into separate m_F sublevels, whose energy varies linearly with position. Figure 2.2 shows a 1D simplification of this for a $F = 0 \rightarrow F' = 1$ transition.

In a weak linearly inhomogeneous field $B_z(z) = Az$ (such that you see the linear Zeeman effect), the hyperfine states split according to the equation

$$\Delta E = g_F \mu_B B_z m_F \propto m_F z \tag{2.23}$$

where g_F is the g-value, and μ_B is the Bohr magneton [83]. For an atom at $z = z' > 0$, $B_z > 0$ so the $m_F{}' = +1\,(-1)$ excited state is shifted up (down). Two counter-propagating beams with opposite circular polarisa-

tion are incident on the atoms. If these beams are red detuned relative to ω_0, then the $\Delta m_F = -1$ transition is shifted closer to resonance, while the $\Delta m_F = +1$ transition is pushed further out. By choosing the polarisation of the beams such that a σ^- beam travels in the $-\hat{\mathbf{z}}$ direction (and σ^+ in the $+\hat{\mathbf{z}}$ direction), then the atom scatters more light from the σ^- beam and experiences a force pushing it towards $z = 0$. For an atom at $z = z' < 0$, so $B_z < 0$, the $m_F = \pm 1$ states are flipped, more light is scattered from the σ^+ beam and the atom experiences a force in the $+\hat{\mathbf{z}}$ direction.

As the light is detuned below resonance, like for optical molasses, a MOT works to both trap and cool the atoms. Similarly to optical molasses, the MOT can easily be extended to 3D space by using 3 orthogonal pairs of counter-propagating beams. Additionally, while a $F = 0 \rightarrow F' = 1$ is quite rare in real atoms, this technique works for any transition where $F \rightarrow F' = F + 1$ [85].

2.2.2 Dipole force

The other type of optical trap that is relevant to this project is the optical dipole trap. The oscillating electric field of the laser induces an oscillating electric dipole moment in the atom. In a non-homogeneous laser field, the interaction between the field and the dipole moment varies with position, hence producing a potential and a force, called the dipole force [85]. If the frequency of this laser field is tuned below resonance ($\delta < 0$), then this force attracts atoms towards higher intensity areas. If $\delta > 0$ however, then the force is repulsive and moves atoms away for high intensities.

The dipole potential is given by [88]

$$U_{dip}(r,z) = -\frac{1}{2}\langle \mathbf{pE} \rangle = -\frac{3\pi c^2}{2\omega_0^3}\left(\frac{\Gamma}{\omega_0 - \omega} + \frac{\Gamma}{\omega_0 + \omega}\right) I(r,z)\,, \tag{2.24}$$

where $\mathbf{p}$ is the atom dipole operator, $\mathbf{E}$ is the electric field of the laser, and $I(r,z)$ is the intensity of the laser. The dipole force is then given by the gradient of this potential

$$\mathbf{F}_{dip}(r,z) = -\nabla U_{dip}(r,z) \propto \nabla I(r,z)\,. \tag{2.25}$$

The simplest form a dipole trap can take is just a strongly focused Gaussian beam detuned far below resonance. This is the type used in this experiment.

In a Gaussian beam, the intensity varies as

$$I(r,z) = I_0 \frac{w_0^2}{w(z)^2} e^{-2r^2/w(z)^2} \tag{2.26}$$

where $w(z)$ describes the radius of the beam and w_0 (the beam waist) describes the radius of the beam at the focus. Here, z is the propagation axis of the laser, and r is the radial distance from the central axis of the beam. As the light is far off resonance for any atomic transitions, the effects of excitation and scatter can be ignored, and the force is given purely by the dipole potential in Eq. 2.24. This potential has a typical trap depth of a few mK. However, as the trapping force is conservative, it is not able to cool the atoms. So it is necessary to cool the atoms (e.g. using a MOT) before loading them into an optical dipole trap.

Chapter 3

Experimental techniques

Cold atom experiments are highly technical experiments comprising of many different components. With the age and history of the field, there are many techniques that have been developed and adapted specifically for using within cold atom physics. In this chapter, I aim to cover many of these techniques and present the theory behind them.

The start of this chapter covers laser locking (i.e. laser frequency stabilisation), describing several methods of doing this (Sec. 3.1). Particular detail is given to a high-frequency optical phase lock I developed for the EIT experiment (Sec. 3.1.3). Then a description of how these lasers are used to image the atoms is given, covering standard absorption imaging of a MOT using a CCD camera (Sec. 3.2.1) and also how we image atoms trapped in our fibre junction (Sec. 3.2.2). Finally, I cover how we use these images to get an estimate for the temperature of our atom cloud (Sec. 3.3). Diagrams of all the optics of this experiment are included in Appendix B.

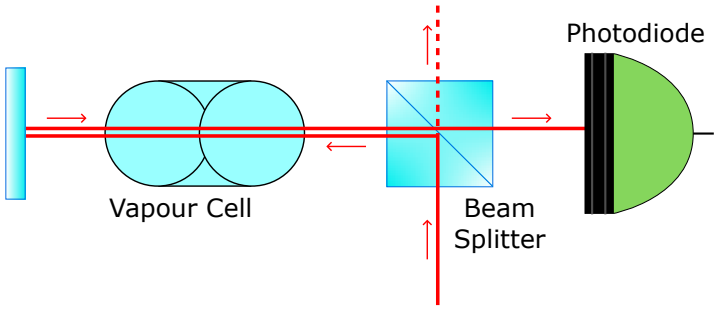

Figure 3.1: Optics layout for Doppler-free saturated absorption spectroscopy.

3.1 Laser locking

Laser locking and frequency stabilisation are an essential part of any cold atom experiment. The hyperfine states in alkali metals typically have a spacing of $\sim$ 10MHz $-$ 1GHz and a linewidth of $\sim$ 5MHz (the transition diagram for caesium can be seen in Fig. 2.1). In order to controllably address a single hyperfine transition, we need precise frequency stabilisation of lasers with a typical linewidth of $\sim$ 1MHz.

In this section, I detail several methods used to stabilise lasers. Starting with using atomic spectroscopy to lock to an atomic transition (Sec. 3.1.1), then describing 2 methods to stabilise a laser by fixing the frequency difference relative to another laser (Secs. 3.1.2 & 3.1.3). All the locking techniques described here are used in this experiment, as detailed in Sec. 4.1.1.

3.1.1 Spectroscopy

For caesium vapour at room temperature, the Doppler broadening of the transitions is far larger than the natural linewidth or the hyperfine separation of the excited state. However, by retro-reflecting the beam back onto itself, sub-Doppler features are revealed, and individual hyperfine transitions can be identified [89]. The experimental setup required for this can

be seen in Fig. 3.1.

This retro-reflection means there are effectively 2 beams with the same frequency ω, but opposite propagation directions. When this frequency is far from resonance $\delta = \omega - \omega_0 \ll \Gamma$, each beam is resonant with 2 entirely separate velocity classes of atoms, so the absorption of 1 beam is unaffected by the absorption of the other beam. However, there are 2 special cases to this situation. The first is when the frequency is on resonance with a hyperfine transition ($\delta = 0$), both beams then only interact with atoms with a 1D velocity $v = 0$. If one beam saturates the transition for these atoms, less light will be absorbed in the other beam, and there will be a small peak in the intensity of light detected at the photodiode.

The more significant special case is when the frequency is exactly halfway between 2 different hyperfine transitions $\delta = \omega - \omega_{01} = -(\omega - \omega_{02})$. In this case, the light from 1 beam is resonant with atoms moving in both the same and opposite directions $v = \pm c\frac{\delta}{\omega}$. Like for the resonant case, the counter-propagating beam is addressing the same atoms, so, due to saturation, less light is absorbed. However, as the number of atoms being addressed is now larger, the height of the intensity peak is more significant than it is for the resonant case. These peaks are called the "cross-over peaks."

An example of what the spectroscopy signals for the caesium D_2 line look like can be seen in Fig. 3.2. In order to stabilise a laser from this signal, the frequency of the laser is modulated and the spectroscopy signal (along with the sign of this modulation) is fed into a lock-in amplifier. This generates an error-signal that is a derivative of the spectroscopy signal (turning the peaks into slopes). The error-signal is then fed back to the laser via a PI card, which uses the slope of this error-signal to stabilise the laser frequency to one of the peaks in the spectroscopy.

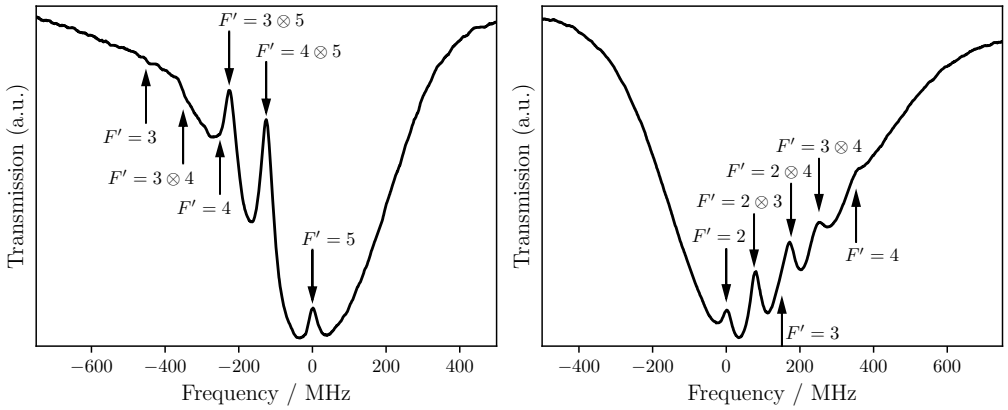

Figure 3.2: Doppler-free saturated absorption spectra for Caesium vapour at room temperature. (Left) Spectrum of the transitions from the $F = 4$ ground state. (Right) Spectrum of the transitions from the $F = 3$ ground state.

3.1.2 Frequency lock

A frequency lock (also called an offset lock) is a way to stabilise the frequency of one laser by fixing the frequency difference relative to another laser (which is typically spectroscopy locked). There are several methods to do this, the electronic scheme we use is based on the design of U. Schünemann et al. (1999) [90]. The 2 lasers (denoted reference and slave) are overlapped and interference between them creates intensity fluctuations with frequency $\omega_{beat} = |\omega_{ref} - \omega_{slave}|$. These intensity fluctuations can be detected on a fast photodiode, where the intensity detected is

$$I(t) = I_1 + I_2 + 2\sqrt{I_1 I_2}\cos(\omega_{beat} t)\,. \tag{3.1}$$

This is called a beat note. One has to be careful to ensure that the 2 lasers have identical polarisation and that the bandwidth of the photodiode is greater than the frequency of the beat note.

This photodiode signal is then mixed with the output of a voltage controlled oscillator (VCO). A VCO outputs an radio frequency (RF) signal with a frequency ω_{vco} proportional to a DC voltage input. These 2 signals are

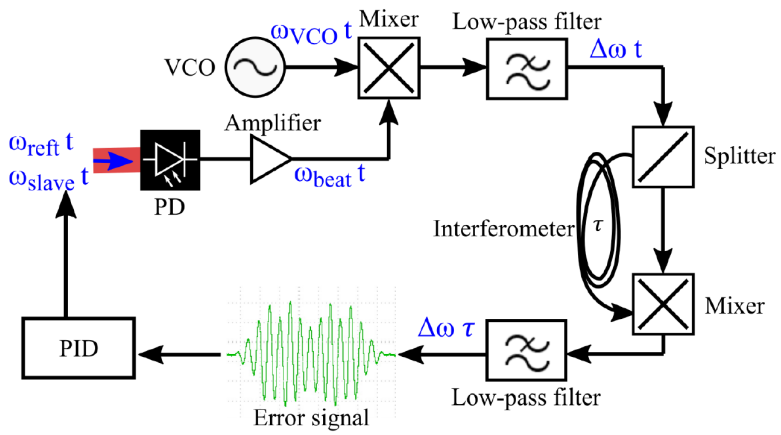

Figure 3.3: Schematic of the offset locking circuit used in this experiment. Figure adapted from Ref. 42.

combined at a mixer used for frequency down conversion. The output of an ideal mixer would be

$$\cos(\omega_{beat}t)\cos(\omega_{vco}t) = \frac{\cos\left((\omega_{beat}+\omega_{vco})t\right)+\cos\left((\omega_{beat}-\omega_{vco})t\right)}{2}, \quad (3.2)$$

then, using a low-pass filter, we can select only the $\Delta\omega = \omega_{beat} - \omega_{vco}$ component.

This signal is then sent through an interferometer where one path is much longer than the other. Due to this difference in length ($\sim$10 m), a delay τ and hence phase difference $\varphi = \Delta\omega\tau$ is built up between the 2 paths. Then, by combining them at a mixer

$$\cos\left(\Delta\omega t\right)\cos\left(\Delta\omega(t+\tau)\right) = \frac{\cos(\Delta\omega(2t+\tau))+\cos(\Delta\omega\tau)}{2} \quad (3.3)$$

and filtering out the high frequency components, one is left with just the $\cos(\Delta\omega\tau)$ term. As this now depends only on the frequency and path delay (and not time), this can be used to derive an error signal and a slope of that error signal can be used to lock the slave laser such that the value of $\Delta\omega$ is fixed.

This setup allows for continuous control over the frequency of the slave laser by varying the voltage input of the VCO. Increasing the VCO input voltage increases ω_{vco} then, as $\Delta\omega$ and ω_{ref} are fixed, ω_{slave} must change to account for this. This makes offset locks particularly useful for our imaging and cooling beams, where precise control of the frequency is desired.

3.1.3 Phase lock

Lasers do not output a perfect sine wave. Spontaneous emission in the amplification medium (that is then further amplified) leads to phase noise, instantaneous changes in the laser's phase $\phi(t) = \omega t + \varphi(t)$ (where $\phi(t)$ is the instantaneous phase of the laser, and $\varphi(t)$ is the phase noise). When detecting a beat note between 2 lasers, the instantaneous frequency difference between them is

$$\Delta\omega = \frac{\mathrm{d}}{\mathrm{d}t}\left[\omega_1 t + \varphi_1(t) - \omega_2 t - \varphi_2(t)\right] = \omega_1 - \omega_2 + \dot{\varphi}_1(t) - \dot{\varphi}_2(t)\,. \quad (3.4)$$

The frequency stabilisation methods described above fix the values of ω_1 and ω_2 relative to each other, but due to the phase noise, there can still be instantaneous changes in $\Delta\omega$. This frequency difference can have a significant impact on coherent effects (such as EIT, which is covered in Chap. 5 of this thesis) [91–93].

The amplitude of a laser's phase noise is a property of how the laser is built, and therefore cannot be reduced easily. However, with fast enough feedback, the phase noise of one laser can be tied to the phase noise of another such that $\Delta\varphi = \varphi_1(t) - \varphi_2(t)$ is constant. So $\Delta\dot{\varphi}(t) = 0$ and $\Delta\omega$ is constant. This is called a phase lock.

This section goes into the design and testing of a new phase lock that

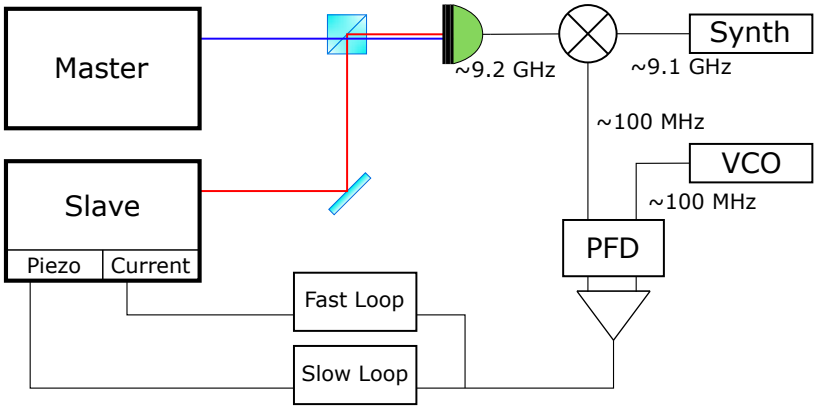

Figure 3.4: Schematic diagram of the optical phase lock. The high-frequency beat note is down-converted at a frequency mixer using a frequency signal from a synthesiser. This down-converted signal is then compared to another reference at a phase frequency detector (PFD) to produce an error signal. This error signal is fed back into the laser through 2 feedback loops with different response speeds.

I built for this experiment. It starts by giving an overview of how the phase lock works, before going into more detail about the phase frequency detector, loop filters, and the performance of the lock. The design used is based on the designs in Müller (2010) [92] and Hockel (2009) [91]. A block diagram of the phase lock can be seen in Fig. 3.4 (more detailed diagrams of individual parts are included in App. C).

The 2 lasers that are phase locked in this experiment address the $F = 3 \rightarrow F' = 4$ and $F = 4 \rightarrow F' = 4$ transitions (denoted "slave" and "master" respectively, like for a frequency lock). The ground-state hyperfine splitting in caesium means they have a frequency difference of 9.192 631 770 GHz [84]. The beat note between these 2 lasers has to be detected with a high-speed, high-bandwidth photodiode[1]. Due to the difficulty of working with GHz signals, the electrical signal generated at the photodiode is down-converted to the $\sim$ 100MHz range by combining the photodiode signal with a $\sim$ 9GHz signal from a synthesiser[2] at a frequency mixer.

[1]Hamamatsu G4176

[2]Windfreak SynthUSB II with 2 Mini-Circuits frequency doublers (ZX90-2-36-S+ and ZX90-2-50S+)

This down-converted signal is then compared to a reference signal at a phase frequency detector (PFD)[3]. The signal outputted by the PFD goes through loop filters and is fed back to the "slave" laser. This feedback works to match the phase and frequency of the down-converted signal to the reference signal. Like for a frequency lock, the frequencies of the "master" laser and this reference signal are fixed. Additionally, the reference signal has much lower phase noise than either of the lasers. So the feedback to the "slave" laser stabilises the frequency and causes there to be a fixed phase relation between the "slave" and "master" lasers.

Phase frequency detectors

There are several different designs of phase detector often based on either analogue mixers or digital flip-flops [92]. The type of phase frequency detector (PFD) used here is a digital three state sequential PFD. It was selected due to its large operating bandwidth and as it leads to a relatively simple design for the lock. This type of PFD has 2 inputs, 4 outputs, and 3 output states. The outputs are: up (U), down (D), not up (NU) and not down (ND) (the "not" outputs are the inverse of the other outputs). The 3 possible output states are

$$\begin{aligned} +\text{VE} &: U = 1, NU = 0, D = 0, ND = 1 \\ 0 &: U = 0, NU = 1, D = 0, ND = 1 \\ -\text{VE} &: U = 0, NU = 1, D = 1, ND = 0\,. \end{aligned} \tag{3.5}$$

The PFD takes 2 inputs, the (down-converted) beat-note from the photo-diode, and a reference oscillator that the beat-note's phase and frequency are compared to. The convention in analogue electronics denotes these the

[3]Analog Devices 105809-HMC439QS16G

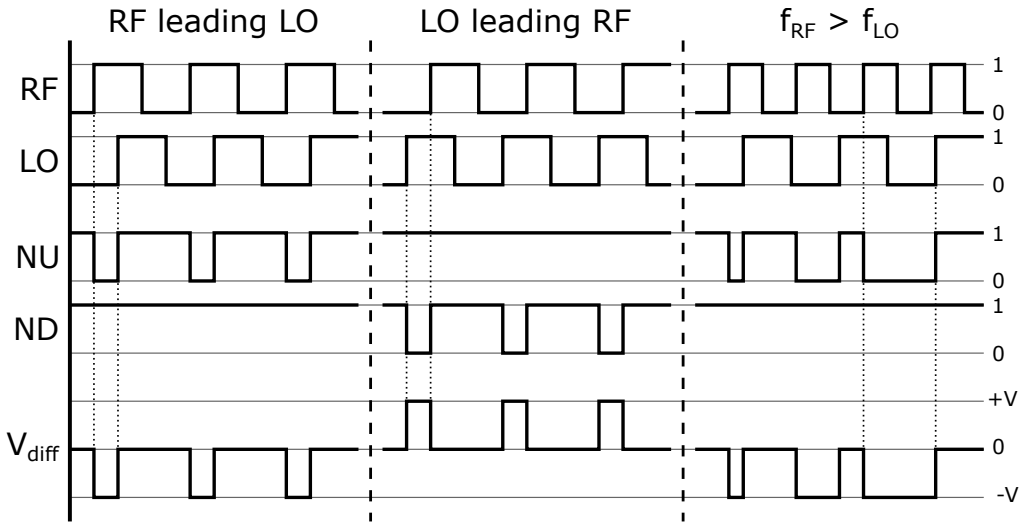

Figure 3.5: Timing diagram for the phase frequency detector (PFD), showing the "not" outputs and V_{diff} for 3 different cases.

"radio frequency"(RF) and "local oscillator" (LO) respectively (in phase locking, the RF might also be referred to as the VCO). A positive edge on the RF changes the output state into the next "higher" state, unless it is already in +VE. A positive edge on the LO does the opposite, going "down" unless it is already in -VE.

The 2 "not" outputs of the PFD are compared and subtracted from each other using a differential amplifier circuit (see Appendix C). So the overall output is $V_{diff} = V_{NU} - V_{ND}$ (where V_{NU} and V_{ND} are the voltages of the NU and ND outputs, respectively). V_{diff} is then the input of the loop filters that provide the feedback to the laser.

A timing diagram showing how the PFD (and V_{diff}) responds to various differences in phase and frequency can be seen in Fig. 3.5. If RF is leading LO, the PFD alternates between the +VE and 0 states (i.e. NU alternates between 0 and 1, while ND is constantly 1), so V_{diff} alternates between 0 and $-V_{out}$ and the mean input into the feedback $\overline{V}_{diff} < 0$. If LO is leading RF, the PFD alternates between -VE and 0, and $\overline{V}_{diff} > 0$. So far, this has assumed $f_{RF} = f_{LO}$. If $f_{RF} < f_{LO}$, rising edges from LO will be

arriving more frequently than from RF, so LO will always be "leading" RF. As a result, V_{diff} will be alternating between 0 and $+V_{out}$, with increasing durations of $+V_{out}$ pulses (with the inverse true if $f_{RF} > f_{LO}$). All of this can be combined to provide feedback such that the frequency and phase of the RF input matches the frequency and phase of the LO.

Loop filter design

Loop filters are how the output of the PFD is feedback to the laser. In this experiment, the laser being locked is an external-cavity diode lasers (ECDLs). The frequency of ECDL lasers can be controlled in 3 ways: the position of a grating (and hence the wavelength of back reflected light and the length of the external cavity), the current through the laser diode, and the temperature of the laser diode. These 3 methods to control the frequency all have different stabilities and different response times. Modulating the laser's temperature has a very slow response time ($\sim 10\text{s} - 1\text{min}$), therefore, external modulation of the temperature is not used as a way to stabilise the frequency. The position of the grating is controlled using a piezoelectric crystal, and has a good response bandwidth of about $0.1\text{Hz} - 1\text{kHz}$. This is how we typically scan and stabilise the frequency in the other locking schemes discussed above. However, this does not provide fast enough feedback for a phase lock, as the bandwidth of the feedback needs to be on the same order as the bandwidth of the phase noise (i.e. the linewidth of the laser, $\sim 1\text{MHz}$). So the phase lock feedback needs to be via current modulation. This can be done via 2 potential paths: the modulation input of the laser current driver, or injecting an RF signal into the laser current using a bias-tee.

Figure 3.6 shows the strength of the frequency response of these differ-

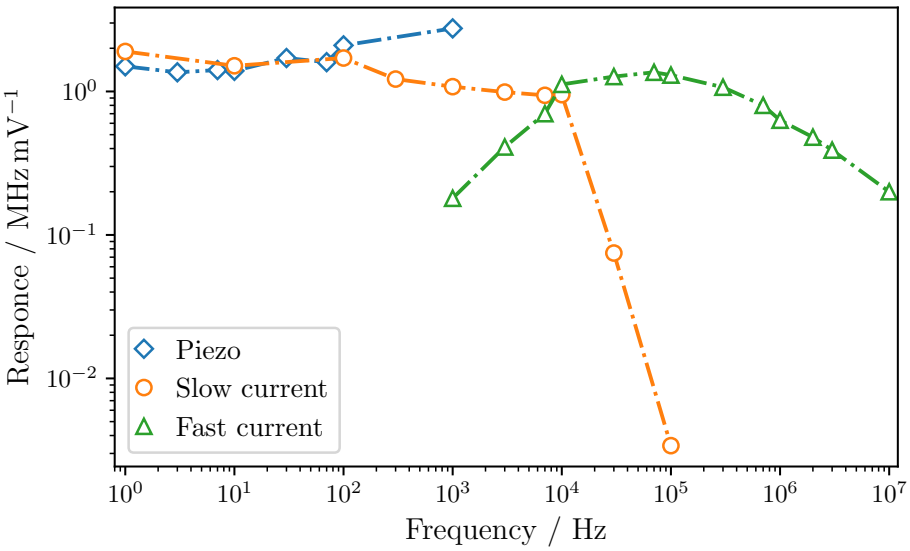

Figure 3.6: Measured strengths of the laser frequency response for different feedback paths. Three different feedback paths are used that alter the cavity (piezo) or current through the diode (fast and slow).

ent feedback paths at different frequencies. These response strengths were measured by looking at how the frequency of the beat note varied in response to a known control signal at different frequencies. The piezo offers very strong feedback at low frequencies, the bias-tee has good feedback at the frequencies required for the phase lock, and the current driver bridges the gap between these 2 paths. For the phase lock, it is important to combine these different feedback paths together for the full lock. The piezo offers long-term stability, and works over a wider frequency range than the current modulation, so can bring the frequency closer to the reference frequency, allowing the current modulation feedback to work.

Each feedback path needs its own loop filter. This controls the strength, speed and phase of the feedback sent to each path. The loop filters I designed are based on the ones described in Müller (2010) [92]. Before use, I simulated them using LTspice to ensure their bandwidths match the bandwidths of the different feedback paths, before building and testing them on prototype board. One change I made compared to Müller [92] was

to add a buffer onto the slow and medium paths to better isolate them from the other paths. These buffers could also be changed to an inverter using a toggle switch, allowing the phase of the feedback to be easily flipped if necessary.

After testing, it turned out that the medium feedback path was unnecessary, offering no obvious difference in the performance of the lock. Additionally, with the other electronic components present earlier in the system, the filter for the fast path only required a potentiometer to control the amplitude of the feedback. No additional phase-shifting or frequency-limiting components were required. The full circuit for the loop filters (and the differential amplifier) is available in App. C, the gain, phase, and sign of each path has been correctly tuned for the laser and lock point we are using.

Phase lock performance

Following implementation, the performance of the phase lock is analysed using the Fourier transform of the beat note between the probe and control lasers. The beat note is measured using an oscilloscope[4] after being down converted to the 100 MHz range by the synthesiser. We obtained the frequency spectrum of this beat note by fast Fourier transform (FFT) on the signal, collecting data at a rate of 1 GSa/s over a period of 1-2s. A better signal could have been measured with a spectrum analyser, however, we did not have this equipment available in our lab.

The spectrum of the beat note with no frequency or phase feedback can be seen in Fig. 3.7. This gives us a measure of the linewidth of the beat note, and the best it can be with only frequency feedback. The data are fit to a Voigt profile (the red line), allowing us to estimate the linewidth (full-width

[4]Teledyne T3DSO1204

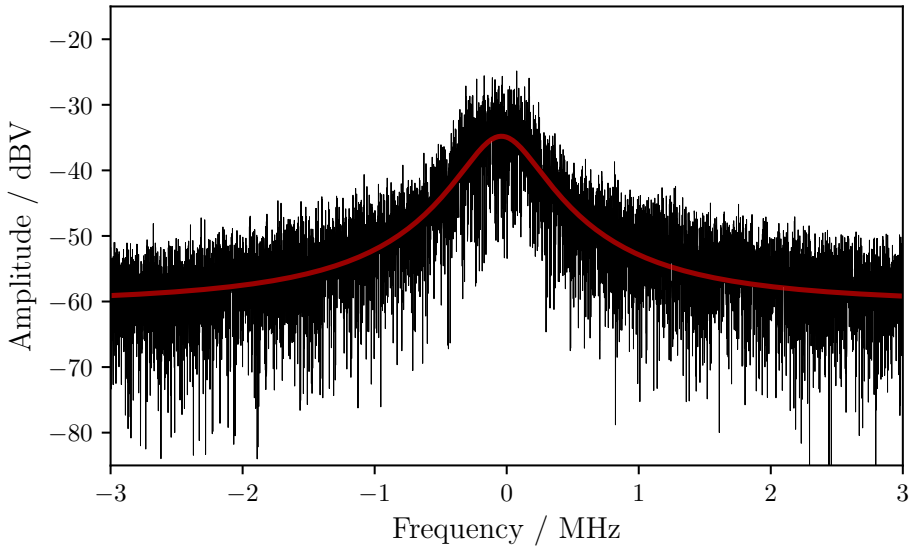

Figure 3.7: Fourier transform of the unlocked beat note between the 2 lasers. The red line shows a Voigt profile fit to the spectrum. This spectrum has a width of about 300 kHz but, while unlocked, nothing stabilises it to this frequency and it would drift over the course of seconds.

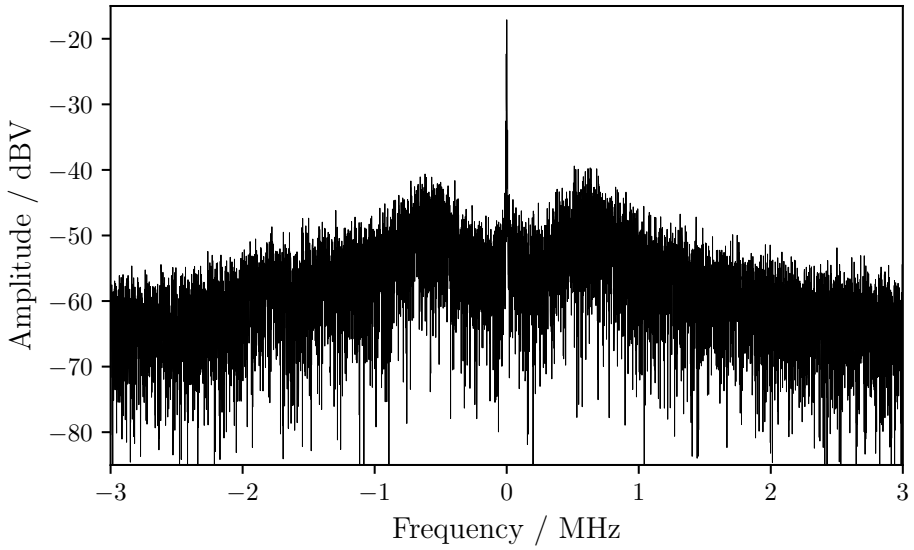

Figure 3.8: Fourier transform of the phase locked beat note with the same span as Fig. 3.7.

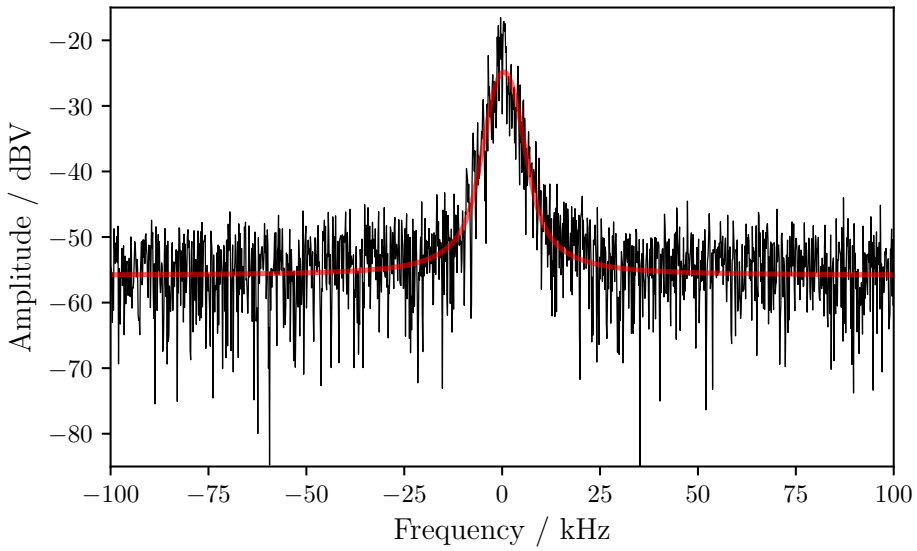

Figure 3.9: Fourier transform of the phase locked beat note with a finer span, so the central peak is clearer. Red line shows a Voigt profile fit to the spectrum. This peak has a linewidth of 3.2 kHz

half maxima) from the Lorentzian linewidth and the Gaussian standard deviation of the fit [92,94]. This gives us a linewidth of (301 ± 7) kHz.

The spectrum of the beat note with full phase lock feedback can be seen in Figs. 3.8 & 3.9. Both figures use the same data, but Fig. 3.9 shows a smaller frequency range to focus on the central peak of Fig. 3.8. The central peak is also fit to a Voigt profile, and has a linewidth of (3.2 ± 0.3) kHz, 2 orders of magnitude smaller than the unlocked beat note.

Although this is a significant improvement, it is a long way off state-of-the-art commercial products and phase locks in the literature, which can achieve beat linewidths of $<$ 10Hz [91,95,96]. The main reason for the comparatively low performance of this phase lock is balancing cost and required performance. This improvement by 2 orders of magnitude, while not state of the art, is sufficient that phase noise is no longer a limiting factor in the experiment.

The main limiting factor in the performance of this lock are the 2 local

oscillators. While the phase noise of electric oscillators is much better than optical oscillators, neither the synthesiser nor the VCO are of particularly high quality. Additionally, their phase noise is uncorrelated, so noise from both contributes to the phase noise of the lock. The best way to improve this would be to have an additional, high-quality reference oscillator (e.g. a 10MHz GPS clock). The local oscillators could then both be phase locked to that reference, reducing these additional sources of phase noise.

Another key area to improve would be the amplifiers in the photodiode circuit. The photodiode used produces a small signal, so a large number of amplifiers had to be used to make the signal strong enough for the PFD. The use of higher quality, higher gain amplifiers (particularly for the 1st amplifier in the circuit) would both help reduce the noise floor of the signal and reduce the propagation delay through the system, increasing the bandwidth.

While there is potential for further improvements in the future, this phase lock has given a significant improvement in the frequency stability of the 2 lasers used for EIT. This lock is now a key component in that experiment (as described in Chapter 5).

3.2 Imaging

Imaging is a key part of any cold atom experiment. Images of various types are used for diagnostics, calibration, and optimisation, as well as being one of the main ways to collect key experimental results. In this section, I describe the various imaging techniques used in this experiment.

3.2.1 Absorption imaging

A common method for imaging cold atoms is through absorption imaging. The atom cloud is illuminated with near-resonant light, which is detected on a CCD camera[5]. This image is then compared to a reference image with no atoms present. By looking at the difference between the 2 images, one can infer how many atoms are present by how much of the imaging light was absorbed.

Light travelling through the atoms is absorbed following the equation

$$\frac{\mathrm{d}I}{\mathrm{d}z} = -a(\omega)I \tag{3.6}$$

which solves to

$$I = I_0 e^{-a(\omega)z} \tag{3.7}$$

where I is the intensity of the laser, z is the distance travelled in the direction of propagation, and $a(\omega)$ is the absorption coefficient. The absorption coefficient is related to the atomic density by

$$a(\omega) = n\sigma_{sc}(\omega) = n\frac{\sigma_0}{1 + 4(\omega - \omega_0)^2/\Gamma^2 + I/I_{sat}} \tag{3.8}$$

where n is the atom number density of the atomic cloud, σ_{sc} is the scattering cross section (which has a Lorentzian frequency response), and $\sigma_0 = \hbar\omega_0\Gamma/2I_{sat}$ is the on-resonance low intensity cross section, ω_0 is the resonant frequency, and I_{sat} is the saturation intensity. The scattering rate Γ also describes the linewidth of this Lorentzian.

[5]Allied Vision, Guppy F-038B

Considering the light detected by a single pixel of the CCD,

$$I_{at}(x,y) = I_0(x,y)e^{-n_z(x,y)\sigma_{sc}} \tag{3.9}$$

where

$$n_z(x,y) = \int_0^z n(x,y,z)dz \tag{3.10}$$

is the column density of the atom cloud in the z direction at that position. Here, $n_z\sigma_{sc} = \ln(I_0/I_{at})$ is commonly known as the optical depth OD. From Eq. 3.9, one can find $n_z(x,y)$ from the intensity of light detected by the CCD using

$$n_z(x,y) = -\frac{1}{\sigma_{sc}(\omega)} \ln\left(\frac{I_{at}(x,y)}{I_0(x,y)}\right) = \frac{\mathrm{OD}(x,y)}{\sigma_{sc}(\omega)} \tag{3.11}$$

so the total number of atoms imaged is then

$$N = \int n_z(x,y)dxdy\,. \tag{3.12}$$

The ratio I_{at}/I_0 is calculated by comparing 3 separate images. For the first image, with intensity I_{at}, resonant light illuminates the atoms, attenuating some of the signal. For the second, I_{ref}, taken about 100 ms later, resonant light illuminates the chamber, but no atoms have been loaded (and 100 ms is long enough that no atoms remain from the original image). The final image, I_{dark}, is taken without any imaging light pulse, so only stray background light illuminates the camera. The optical density of the atom cloud at each pixel can then be calculated using

$$\mathrm{OD}(x,y) = \ln\left(\frac{I_{ref} - I_{dark}}{I_{at} - I_{dark}}\right)\,. \tag{3.13}$$

The total atom number can then be found using the optical depth on each

pixel, summing over every pixel while taking account of the pixel size and the magnification of the imaging system.

3.2.2 Absorption in the fibre junction

Important for this thesis is the ability to probe the atoms in the junction via photons travelling through the fibre. Measuring the number of atoms trapped in this fibre junction works on the same principle as absorption imaging. Resonant light illuminates the atoms, and by calculating how much the signal is attenuated by comparing it to a reference with no atoms, one can find the optical depth and hence the number/density of atoms in the junction. The main difference is that rather than the light being detected by a CCD and summing over different pixel values, the light is detected at a single point using a single photon counting module (SPCM)[6].

The core of the fibre is very narrow ($\sim$ 5µm), so light in the interaction region can reach very high intensities with very low optical powers. To remain in the low saturation limit, very low optical powers are used ($\sim$ 100pW). A standard photo-diode would struggle to detect this signal, which is why a SPCM is used. More details about the exact fibre system are given in Sec. 4.1.3.

Single photon counters

A SPCM outputs a binary logic pulse every time it detects a photon (a typical SPCM has a detection efficiency of $\sim 30\%$). These pulses are counted using a binary counter that is then read out using an Arduino. The number of photon detection events within a certain time window is then recorded,

[6]Excelitas Technologies SPCM-AQRH-14

and that gives the intensity signal of the imaging pulse.

The logic pulse that occurs after a detection event has a finite duration (typically $\sim$ 10ns). Additionally, after this pulse, the detector will have a short "dead time" before it resets and can detect further photons (typically 22ns according to specification). If a photon arrives at the detector during either of these periods, it will not be detected. This leads to what is called linearity correction. The manufacturer of the SPCM specifies correction factors where, for a given rate of logic output pulses, the true number of photons arriving in that period is larger by the specified factor [97].

An SPCM has much higher sensitivity to background light than a normal photodiode. As such, a number of measures are taken to reduce the background light and improve the signal-to-noise ratio. The detector is physically shielded from other sources of light, with the output of the fibre being the only significant path for light to reach the detector. Additionally, a 950 nm short-pass filter and a laser-line filter[7] are used so that only light of the appropriate wavelength for our signal can be detected. Finally, the SPCM is gated, so that photons are only detected in a narrow detection window around the imaging pulse (typically $\sim 10\,\mu s$). Even with these measures, the detector will have a non-zero dark count rate, so it is necessary to do experimental cycles with no imaging light to calibrate the number of dark counts detected.

Probing atoms in the fibre junction

As is the case for absorption imaging, the optical depth is calculated using the ratio of intensities of probe pulses with and without atoms present (taking into account any systematic background counts). With the optical

[7] Thorlabs FESH0900, Thorlabs FB850-10, Semrock LL01-852-12.5

powers used here, noise is limited by the photon-shot noise on both these pulses. This means a single set of probe pulses does not give a good enough signal-to-noise, so it is necessary to repeat the experimental cycle many times and average over several sets of probe pulses. Typically, we average over 15-50 experimental cycles for a single data point.

When probed with resonant light ($\sigma_{sc}(\omega = \omega_0) = \sigma_0$), the number of atoms trapped in the interaction region of the fibre junction can then be calculated using

$$N = -\frac{A}{\sigma_0} \ln\left(\frac{\overline{n}_{at} - \overline{n}_{d1}}{\overline{n}_{ref} - \overline{n}_{d2}}\right) \tag{3.14}$$

where A is the area of the probe beam, $\overline{n}_{at}$ is the mean number of photons detected for a probe pulse with atoms present, $\overline{n}_{ref}$ is the mean number of photons detected for a probe pulse with no atoms present, and $\overline{n}_{d1(2)}$ are the mean number of "dark counts" on the 1st and 2nd probe pulse respectively. The error associated with N is calculated using the standard error on the mean of n_{at} and n_{ref} and standard propagation of errors.

When probing atoms in the fibre, it is sometimes useful to take an absorption spectrum measurement. The transmission through the fibre is measured for an experimental run of ~ 30 measurements. This is recorded, and then another experimental run is done with a different probe frequency. Repeating many times for different frequencies, one can construct a spectrum of how the transmission through the fibre varies with the frequency of the probe light. Assuming we're in the low saturation limit, the transmission as a function of frequency is then

$$T(\omega) = \frac{\overline{n}_{at}(\omega) - \overline{n}_{d1}}{\overline{n}_{ref} - \overline{n}_{d2}} = \exp\left[-\frac{\mathrm{OD}_{res}}{1 + 4(\omega - \omega_0)^2/\Gamma^2}\right] \tag{3.15}$$

where OD_{res} is the on-resonance optical depth, from which one can calculate the number of atoms in the fibre junction. Fitting this curve also

allows the linewidth Γ to be determined to check if there is any broadening in the system due to intensity or collisions. The other advantage this measurement offers is that for high optical depth (high atom numbers), it can be difficult to get an accurate atom number estimate from just an on-resonance measurement. Probing the full spectrum allows for the optical depth to be determined more accurately, giving a better atom number estimate. This topic of accurate measurement strategies is explored in more detail in Chapter 6.

3.3 Temperature measurement

The temperature of a cold atom cloud is determined by observing how quickly the cloud expands and, from that, determining the thermal velocity of the trapped atoms. There are several techniques one can use to determine this. The method used in this experiment involves taking absorption images with varying "drop times" (time between the trap turning off and the image being taken). This allows us to determine how quickly the cloud expands and hence determine the temperature of the trapped atoms. This technique is commonly known as a time-of-flight temperature measurement.

In the model presented here, we assume classical ballistic atoms and neglect the effects of collisions between atoms. These assumptions are valid for the temperature and density regimes achieved in this experiment. The energy of atoms confined in a harmonic trap is

$$E(\mathbf{r}, \mathbf{p}) = \frac{1}{2} m \omega_T^2 \mathbf{r}^2 + \frac{\mathbf{p}^2}{2m} \tag{3.16}$$

where m is the mass of the atom, ω_T is the trap frequency, and the trap frequency is the same in every dimension. The distribution of atoms in this

trap is given by the Maxwell-Boltzmann distribution

$$f(\mathbf{r},\mathbf{p}) \sim e^{E(\mathbf{r},\mathbf{p})/k_B T}\,. \tag{3.17}$$

Combining this with Eq. 3.16 and normalising, we get

$$f(\mathbf{r},\mathbf{p}) = \frac{1}{(2\pi\sigma_r^2)^{3/2}} e^{-\frac{\mathbf{r}^2}{2\sigma_r^2}}\ \frac{1}{(2\pi\sigma_p^2)^{3/2}} e^{-\frac{\mathbf{p}^2}{2\sigma_p^2}} \tag{3.18}$$

where $\sigma_r^2 = \frac{k_B T}{m\omega_T^2}$ and $\sigma_p^2 = m k_B T$. If we let the atoms fall from the trap, after time t each atom will have a new position $\mathbf{r}'(t) = \mathbf{r} + \mathbf{p}t/m$. The phase space distribution can now be expressed as

$$f(\mathbf{r}',\mathbf{p}) = \frac{1}{(2\pi\sigma_r^2)^{3/2}} e^{-\frac{|\mathbf{r}'(t)-\mathbf{p}t/m|^2}{2\sigma_r^2}}\ \frac{1}{(2\pi\sigma_p^2)^{3/2}} e^{-\frac{\mathbf{p}^2}{2\sigma_p^2}}\,. \tag{3.19}$$

Integrating this over all momenta, we get

$$f(\mathbf{r}'(t)) = \frac{1}{(2\pi\sigma_t(t)^2)^{3/2}} e^{-\frac{\mathbf{r}'(t)}{2\sigma_t(t)^2}} \tag{3.20}$$

where $\sigma_t(t)^2 = \frac{t^2}{m^2}\sigma_p^2 + \sigma_r^2$ [42,98]. This new spatial distribution is a normalised Gaussian that remains a normalised Gaussian as it evolves with time. So the temperature of the atoms can be determined by fitting the cloud to a Gaussian and seeing how the standard deviation (i.e. size) of the cloud expands with time, fitting that to

$$\sigma_t(t) = \sqrt{\frac{k_B T}{m} t^2 + \sigma_r^2}\,. \tag{3.21}$$

Chapter 4

Setup and experimental overview

The aim of this experiment is to integrate cold atoms with integrated optics. To that end, we have an atom-waveguide junction that consists of a transverse hole through an optical fibre, within which we trap the atoms so they interact with the guided mode of the fibre. This experiment has several stages, all of which need to be optimised both individually and in how they affect the later stages. First, atoms must be loaded into a magneto-optical trap (MOT) to cool them down. Once cooled, they can be loaded into an optical dipole trap. That dipole trap has to be carefully aligned through the fibre junction to maximise the number of atoms in the interaction region. Once all that is optimised, only then can we start performing the EIT and Bayesian experiments.

I start this chapter by giving an overview of all the equipment, notably all the lasers and the atomic transitions they address (Sec. 4.1). I then cover experimental optimisation results from the MOT (Sec. 4.2), dipole trap (Sec. 4.3), and the fibre junction (Sec. 4.4).

4.1 Apparatus

In this section, I give a more detailed description of the specific equipment used for the experiments in this thesis. This experiment builds upon the work of previous members of the group, notably Dr. Elisa Da Ros [42], and Dr. David Johnson [43]. Many parts of the experiment were initially built by them and are described in more detail in their theses. Diagrams of the optics layouts around the lasers and the vacuum chamber are included in appendix B.

4.1.1 Lasers

There are 7 different lasers used in this experiment for a variety of purposes, including: frequency references, cooling, imaging, pumping, and trapping. 4 of the 7 lasers are homebuilt external-cavity diode lasers (ECDLs) with output powers of $\sim$ 10mW. These ECDLs are set up in a Littrow configuration. These lasers all output 852 nm light and are frequency stabilised to optical transitions using various methods as described in Sec. 3.1. Also operating at 852 nm, we use a Toptica TApro with an output power $>$ 700mW, which is used to generate the cooling light for the MOT. The 2 remaining lasers operate at 1064 nm with 1 acting as the "seed" for the 2nd, which has an output power of $\sim$ 30W. This far detuned, high power laser is used to create an optical dipole trap (see Sec. 2.2.2).

Most of the lasers are setup on a separate optical table to the main experimental chamber, in order to better isolate them from sources of vibration in order to achieve better stability. Light from all the lasers (on both the experimental and optics table) are brought over to the experiment chamber through polarisation maintaining optical fibres. These are used not

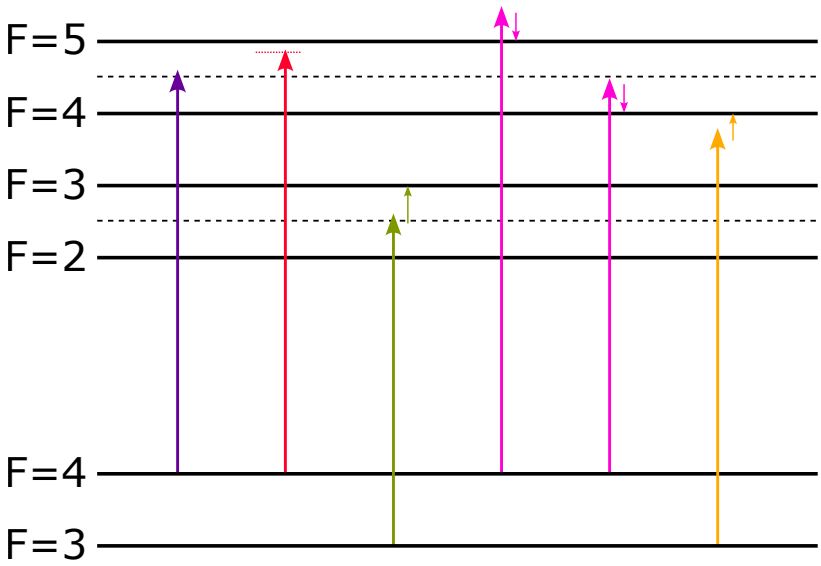

Figure 4.1: Energy level diagram showing the transitions that the resonant lasers address. Large arrows represent the frequency the laser is stabilised to, and the small arrows represent how the laser frequency is shifted by AOMs. The colour of the arrow represents which laser it is. (Purple) Reference, (Red) Cooler, (Green) Repumper, (Pink) Probe, (Orange) Control.

only to easily cover the distance between different components, but are used to "clean" the mode of the light out of the fibre and to isolate any alignment around the chamber from alignment nearer the lasers. A combination of shutters and acousto-optical modulators (AOMs) are used to turn the beams on/off at various points in the experimental sequence. Shutters (most of which are home-built using old hard drives) can withstand higher light intensities than AOMs and are easier to align. However, shutters have a much slower response time than AOMs (typically $\sim$ 3ms for a shutter and $\sim$ 10ns for an AOM). AOMs also have the additional effect of shifting the laser frequency by a fixed amount dependent on the frequency at which it is driven.

A scheme of which caesium transitions we address and which lasers address them is shown in Fig. 4.1. Details of what each of these resonant lasers do and what method we use to stabilise their frequency are summarised in Table 4.1. Schematics of the full optics layouts around the lasers and

Name	Function	Stabilisation technique
Reference	Frequency reference for Cooler and Probe lasers	Doppler-free absorption spectroscopy
Cooler	Cooling light for MOT	Frequency offset locked to Reference
Repumper	Repumper light for MOT	Doppler-free absorption spectroscopy
Probe	Imaging light and EIT probe	Frequency offset locked to Reference
Control	EIT control	Phase locked to Probe

Table 4.1: Summary of what the different lasers are used for and how their frequency is stabilised. This table only describes the lasers that resonantly address transitions; the high-power lasers used for the dipole trap are not included. All lasers other than the Cooler are homebuilt ECDLs.

around the chamber are available in appendix B.

The "reference" laser (a homebuilt ECDL) is spectroscopy locked to the $|F = 4\rangle \rightarrow |F' = 4 \otimes 5\rangle$ crossover peak. This laser is used as a frequency reference for 2 other lasers. The "cooler" laser (Toptica TApro) is used to generate the cooling light used in the MOT, with some light also used for optical pumping. It is frequency locked to the reference, which allows the frequency to be scanned from on resonance with the $|F = 4\rangle \rightarrow |F' = 5\rangle$ transition to $-80\,\mathrm{MHz}$ detuned from that transition. The "repumper" laser is spectroscopy locked to the $|F = 3\rangle \rightarrow |F' = 2 \otimes 3\rangle$ crossover peak. The light from this laser is shifted up to resonance with the $|F = 3\rangle \rightarrow |F' = 3\rangle$ by an AOM and is combined with the cooler light to aid in forming the MOT. While the $|F = 4\rangle \rightarrow |F' = 5\rangle$ transition is a cycling transition, due to quadrupole transitions, collisions, and off-resonant scattering of the $|F = 4\rangle \rightarrow |F' = 4\rangle$ transition, there is still a small chance of the atoms ending up in the $|F = 3\rangle$ ground state. The repumper pumps atoms out of this state back into the $|F = 4\rangle$ ground state so they can be addressed by the cooler light and are not lost from the trap.

The "probe" laser is the main laser used for imaging (both normal absorption images and imaging in the fibre). It is also used as the probe light for EIT (see Chap. 5). The probe is frequency locked to the reference laser and makes a double pass through an AOM, which allows the frequency of the laser to be tuned in 2 separate ways (by changing the frequency of the AOM driver or the frequency of the frequency lock VCO). There are also 2 different lock points used for the probe laser, depending on whether we want to address the $|F = 4\rangle \rightarrow |F' = 5\rangle$ or $|F = 4\rangle \rightarrow |F' = 4\rangle$ transitions. The "control" laser is phase locked to the probe laser (see Sec. 3.1.3), the exact frequency of this can be changed by adjusting the frequency of the synthesiser. This laser is used to generate the control light for EIT. The use of a phase lock here minimises the instantaneous frequency difference between the 2 light fields, which is important for effects like EIT. The existence of the double-pass AOM in the probe path allows the frequency of the probe light to be changed without affecting the frequency of the control laser. Whereas shifting the probe frequency using the frequency lock VCO would shift both the probe and control.

4.1.2 Magnetic fields

Magnetic fields play an important role in any cold atom experiment. The presence of a magnetic field gradient is essential for the formation of a magneto-optical trap (MOT). In this experiment, we have 4 sets of magnetic coils, 1 in anti-Helmholtz configuration to create the field gradient for the MOT, and then 3 orthogonal Helmholtz pairs used to counteract any background magnetic fields and change the position of the field centre. One of the offset coils is wound around the same block as the anti-Helmholtz pair, this is denoted as the "Z coil". The remaining coils produce fields

on the 2 remaining orthogonal axes. One of them is aligned to produce a vertical field, while the other produces a horizontal field that is orthogonal to the Z coil field and in axis with the fibre. These coils are denoted "V coil" and "H coil" respectively. The coil axes are labelled in Fig. B.3.

The anti-Helmholtz pair used for the MOT are powered by 2 separate power supplies, both running at 10A, producing a field gradient of about 10 G/cm. The currents through the offset coils are initially set such that the MOT is formed away from the fibre. Then, shortly before the dipole trap is loaded, the current in the offset coils is ramped so the field centre (and hence MOT) is "pushed" towards the fibre junction, to improve the loading efficiency. During this "push" phase, the current in the anti-Helmholtz coils is ramped down and the frequency of the cooling beams is reduced, so the push phase also acts as a molasses phase, allowing us to cool the cloud below the Doppler limit. The coils and magnetic fields are described in more detail in the thesis of Dr. Elisa da Ros [42], however, it should be noted that the chamber has been rotated since that thesis was written. This rotation was done so that the dipole trap is now horizontal and not against gravity.

4.1.3 Fibre junction

Central to the experiments in this thesis is the atom-waveguide junction, which consists of a transverse hole through an optical fibre. Previously, our group has conducted research into how to optimally shape these waveguide holes to maximise transmission across them. The optimal shape was found to be a 3D parabolically convex surface [42,44,45]. However, due to manufacturing limitations, 3D curved surfaces are hard to achieve. The design of the junction used in this thesis is a 2D circularly convex surface. Images of the fibre can be seen in Fig. 4.2.

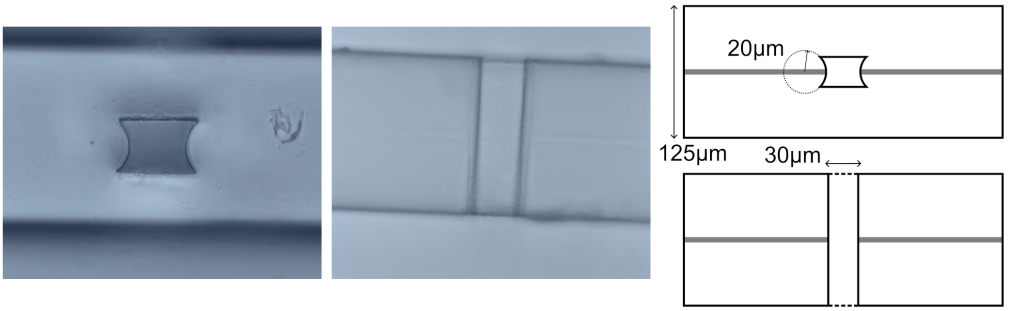

Figure 4.2: Images of the fibre used in this experiment. These images show both top and side views of the fibre junction. The diagram to the side shows the dimensions of the junction (to scale). The diameter of the fibre cladding is 125 µm, the core (light grey line) is 5 µm. The junction is 30 µm across, the faces of the junction have circular curvature with a radius of curvature of 20 µm. Images of the fibre junction were taken by Workshop of Photonics.

The fibre is a Thorlabs 780HP, which was drilled by Workshop of Photonics in Lithuania. The fibre has a mean mode diameter of 5 µm at 850 nm. The junction is 30 µm across (at closest approach), and the surfaces have a radius of curvature of 20 µm. The fibre is mounted on a glass slide with a 16 mm diameter hole centred where the junction is. This is to minimise any disturbance the glass slide may have on the MOT. The ends of the fibre exit the vacuum chamber via teflon feed-throughs. The overall measured transmittance through the fibre is $\sim 30\%$, which gives a transmittance of $\sim 40\%$ across the fibre junction. This is lower than the $\sim 70\%$ predicted by simulations [45], largely due to the surface roughness of the junction

4.1.4 Experimental sequence

Precise control and timing of many different components is essential for an experiment like this. To implement the experimental sequence, we use a National Instruments PXIe-8130 controller. This is a field-programmable gate array (FPGA) with 1 µs timing resolution, 32 digital and 12 analogue output channels. The FPGA is controlled using a custom piece of Python

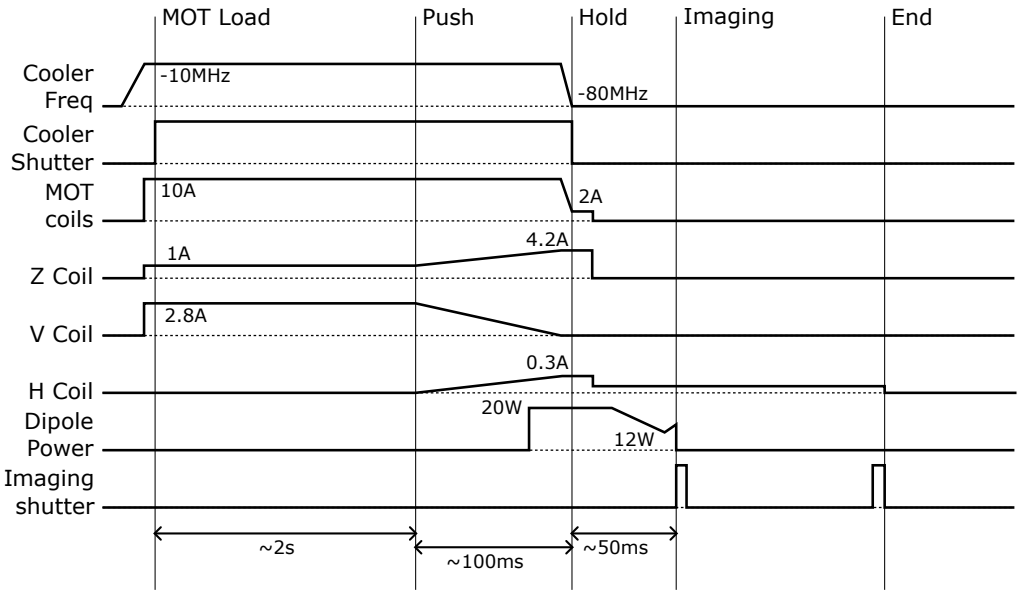

Figure 4.3: Diagram illustrating a simplified version of the sequence of the experiment.

software developed by previous members of the research group. While the timing resolution is only 1 µs, the consistency is much better, with a timing drift between runs of ±5ns.

The digital output channels are mostly used for controlling components such as shutters and AOM drivers, while the analogue channels are used to control the currents through magnetic field coils and laser frequencies. For the slow light experiments (see Sec. 5.1.1), the timing resolution of 1 µs is not sufficient, however, the precision is. So, digital output pulses from the National Instruments box are sometimes used to control pulse/function generators, which are capable of creating shorter pulses of arbitrary shape.

An example of some of the timings of the experimental sequence can be seen in Fig. 4.3. This illustrates a simplified version of a typical experimental sequence. A MOT is loaded, pushed closer to the fibre junction, cooled, then loaded into a dipole trap before being imaged. During MOT loading, the cooler (and repumper) beams are on, the cooler is about 10 MHz red detuned, and the MOT coils have a high current flowing through them.

In the "push" phase, the current in the 3 sets of offset coils is ramped to move the trap centre closer to the fibre. They are ramped relatively slowly ($\sim$ 100ms) to avoid losing atoms from the trap. After the MOT has moved, the current in the MOT coils and the frequency of the cooler are both ramped down to implement a short molasses phase, cooling the atoms below the Doppler limit. then, 20 ms before the end of the push phase, the dipole laser AOM is turned on and atoms start to be loaded into the dipole trap. At the end of the push phase, the cooler beams turn off while the dipole beam remains on. This means atoms that were in the MOT but not in the dipole trap fall away from the fibre. During this phase, the power in the dipole beam is also ramped to achieve evaporative cooling and to compress the atoms towards the focus of the trap.

During the imaging phase, the dipole beam is turned off, and 3 µs later, a probe pulse illuminates the atoms. The time between the trap being turned off and an image being taken is kept to a minimum to avoid atoms from the trap expanding too much before being imaged. The exact duration of the imaging phase, as well as the duration of individual imaging pulses, depends on the type of image being taken and the axis it is being taken on. In some cases, the MOT and dipole beams may be turned on again during the imaging phase to ensure the later reference images are as similar to the "with atoms" image as possible.

Figure 4.3 gives a good illustration of what a typical sequence looks like. However, it should be noted that there are many components in the real experiment that are not included in this illustration (e.g. camera triggers, repumper beam, and probe frequency, to name a few). The exact details of the sequence may also change depending on what is currently being measured. For example, in order to optimise a MOT, the atoms would not be loaded into the dipole trap, and there would not be a long hold phase.

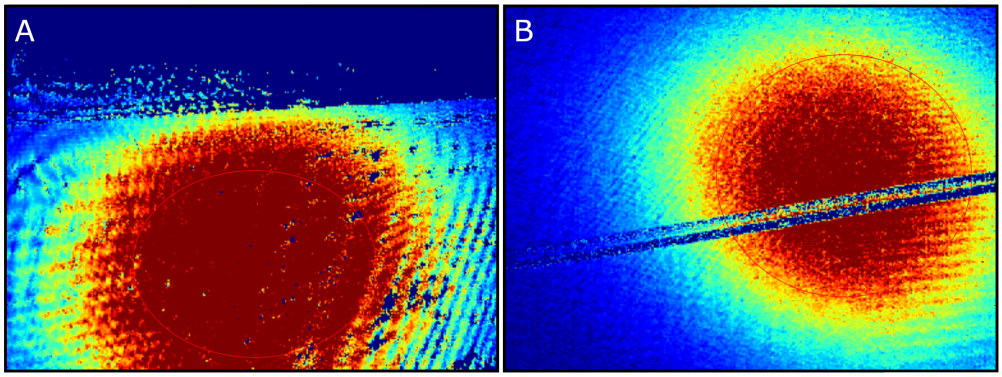

Figure 4.4: False colour images of the MOT taken on the (A) vertical and (B) horizontal imaging axes. Red represents areas with high optical depth (large numbers of atoms), while blue represents areas with very few caesium atoms. The fibre can be seen cutting across the MOT on the horizontal axis, and the large dark patch on the vertical axis is the fibre and the glass slide the fibre is mounted to.

4.2 Magneto-optical trap

The formation of a magneto-optical trap (MOT) is the first stage of the experimental procedure. The effectiveness of our MOT loading will affect the loading of the dipole trap, and hence the loading of atoms into the fibre junction. Total atom number, atom density, temperature, shape, position, and stability are all important parameters to optimise when optimising MOT loading. The parameter space to optimise over is also large, covering laser beam powers, beam alignment, beam power balance, laser frequency, magnetic field gradient, magnetic field centre, and repumper power. In this experiment, the beam alignment is a particularly sensitive parameter due to the glass chip the fibre is mounted on blocking part of the MOT beams. As well as affecting the maximum size of MOT that can be formed, poor alignment of the beams can also limit how close the MOT can get to the fibre junction.

While all the parameters listed above are important, the ones we particularly focus on optimising here are temperature, density, and position.

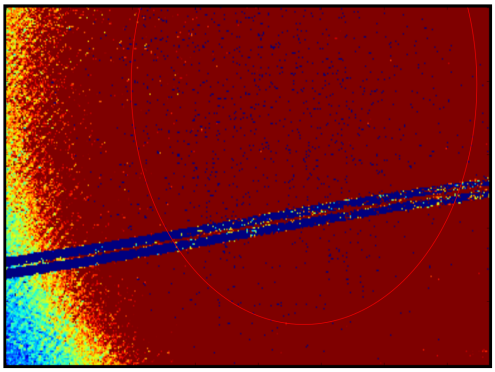

Figure 4.5: Image of the MOT taken on the horizontal axis. The atom cloud is too large for the camera to properly resolve it due to the limited field of view. It should be noted that one cannot determine how close this atom cloud is to the fibre based only on a horizontal axis image.

Total MOT atom number is less important than density due to how small a sample of the atoms end up trapped in the junction. A standard MOT would typically contain $\sim 10^7 - 10^8$ atoms, whereas we only end up probing ~ 300 atoms in the interaction region of the junction. Therefore, the MOT density and temperature are the most important parameters.

Images

Caesium atoms are produces in the chamber by dispensers running at about 4A. The atoms are cooled and trapped by MOT beams with a 25mm diameter. Images of the MOT after the push phase can be seen in Fig. 4.4. Images are taken on both horizontal and vertical axes to allow us to determine how close the MOT is to the fibre and the fibre junction. The magnification on both these axes is quite high to allow us to easily see the fibre hole on the horizontal axis, and to see between the fibre and the glass slide on the vertical axis. This high magnification can cause some issues when imaging an MOT that has not been pushed, due to the limited field of view. This can be seen in Fig. 4.5. However, the benefit of the improved resolution when doing fine alignment around the fibre junction outweighs

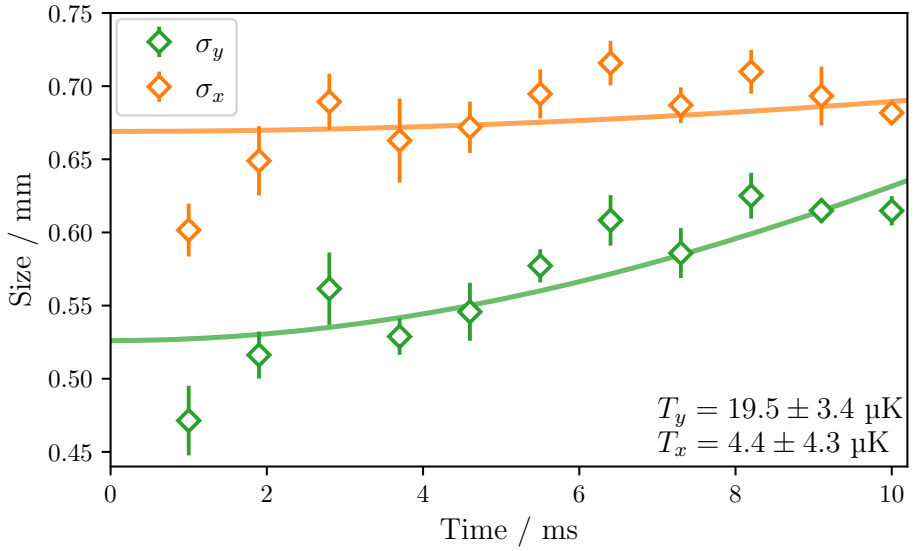

Figure 4.6: Time-of-flight temperature measurement of the MOT after the push and molasses phase. Green and orange data points represent the mean size of the atom cloud in the vertical and horizontal axes of the image respectively, error bars represent the standard deviation in the data set used. Solid lies are the fit to Eq. 3.21.

the impact of the field of view limitation.

After the push and cooling phase, our MOT typically contains $\sim 4 \times 10^7$ atoms, with a peak density of $\sim 3 \times 10^{10}\,\mathrm{cm}^3$.

Temperature

The temperature of the atom cloud was measured after the push phase using a time-of-flight measurement (see Sec. 3.3). At the end of the push phase, the trapping field is turned down (coil current drops to 2A), and the cooling light is detuned to $-80\,\mathrm{MHz}$. This is to implement a molasses stage to achieve sub-Doppler cooling. The Doppler limit for caesium is $125\,\mu\mathrm{K}$ [84]. Following molasses cooling, our atom cloud is significantly colder than this, getting down to $T = (19 \pm 3)\,\mu\mathrm{K}$. The time-of-flight measurement used to determine this temperature can be seen in Fig. 4.6. There is a large discrepancy in the temperature in the 2 axes, this is an

artefact of the imaging system. The glass slide that the fibre is mounted on obscures a section of the image, meaning the expansion of the MOT in that direction cannot be determined accurately.

4.3 Dipole trap

Loading atoms from the MOT into an optical dipole trap is the next step of the experimental procedure. This trapping field is generated using a 30W fibre laser with a wavelength of 1064nm, that gets focused down to a width of $\sim$10 µm. Atoms trapped at the focus of the dipole beam are ultimately the atoms that will be in the interaction region of the fibre junction. As with MOT loading there are many parameters that need to be optimised when optimising dipole trap loading. However, again, in order to optimise the atom number/optical depth in the fibre junction, the most important parameters for this experiment are temperature and density, with the total atom number mattering less.

Beam profile

In order to have a good dipole trap, it is important for the trapping beam to have a good beam profile. For a trap like this, the beam should have a clean Gaussian profile. Any unevenness in this profile can lead to additional heating of atoms in the trap, increasing collisions and reducing the density.

While working on the experiment, I noticed the temperature of atoms in the dipole trap was very high ($\sim$ 1mK). After checking the profile of the beam, it was clearly very uneven. This profile was improved by cleaning all the optics and improving the alignment through the lenses in the

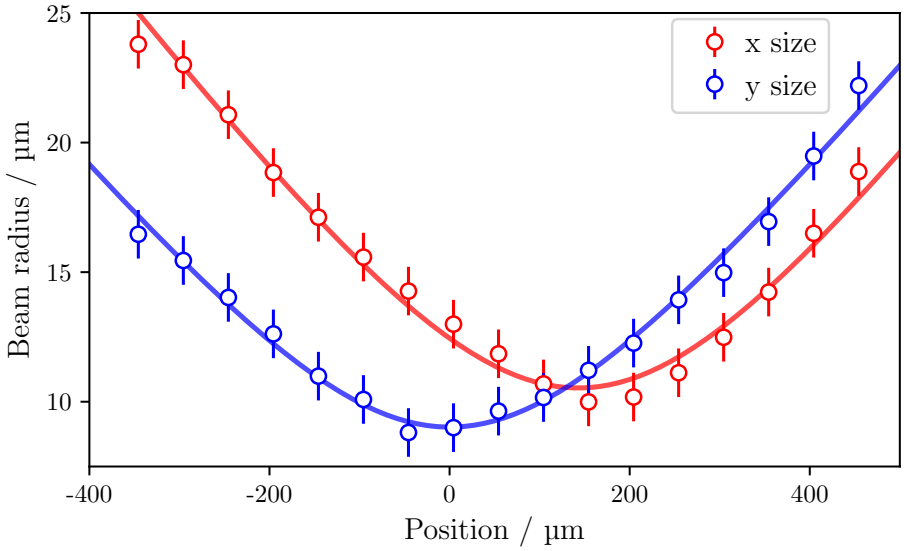

Figure 4.7: Measurement of the waist of the dipole beam around the focus of the trap. The error bars are based on the pixel size of the CCD camera.

telescope. This significantly improved the profile of the beam. The beam also displayed significant astigmatism, which was also improved by careful alignment through the telescope.

How the profile varies with position around the focus of the trap after optimisation is shown in Fig. 4.7. The beam still has a slight astigmatism, however, it is no longer very significant. The M^2 of the beam is 1.44 ± 0.04. Another important part of the profile to check is that it will cleanly fit through the fibre hole. The beam waist is $< 15\mu m$ for a range of $\sim 300\mu m$, so it can cleanly fit through the fibre hole.

This data was collected by creating an equivalent focus outside the chamber and (while the beam was at low power) imaging the profile on a CCD camera[1]. This data was then analysed using Python [99]. At the focus, the size of the beam was equivalent to the size of the pixels of the camera. This limitation of the pixel size was the main source of error for this profile analysis.

[1] Allied Vision Guppy F-038-B

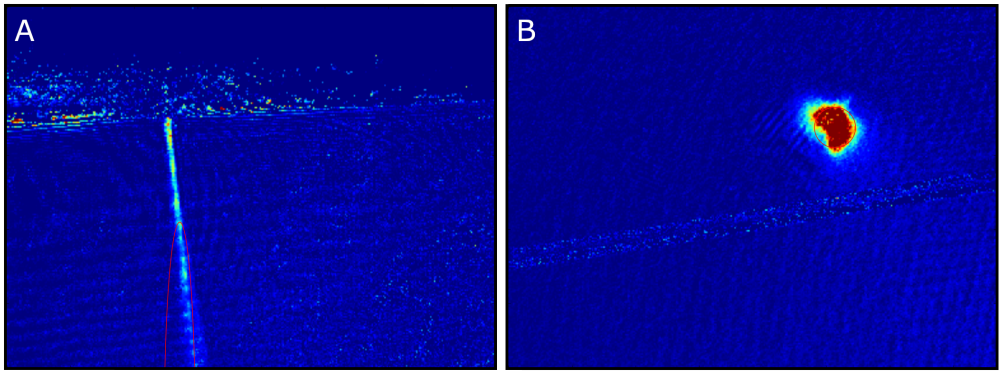

Figure 4.8: Images of the dipole trap on the (A) vertical and (B) horizontal imaging axes. Images are taken with 50 ms of dipole hold time, so all the MOT atoms that are untrapped have fallen out of the imaging region. The dipole trap in these images is not aligned through the fibre hole.

Loading

The dipole beam is turned on about 20 µs before the end of the push and molasses phase of the MOT loading. Atoms from the MOT will start to load into the dipole trap while they're still being pushed and will continue to be loaded after the MOT beams have turned off. Figure 4.8 shows images of the dipole trap with a long enough hold time such that none of the MOT atoms are visible. We typically load $\sim 5 \times 10^5$ atoms into the dipole trap with a density of $\sim 10^{12}\,\mathrm{cm}^{-3}$ at the focus.

The dipole beam starts at max power, to maximise trap depth and capture as many atoms as possible. Then, during the hold phase, the power is ramped down to achieve some evaporative cooling before being ramped back up to compress the atoms towards the focus, increasing the density in that region.

Pumping

Optical pumping in the Zeeman sub-levels is another important factor to optimise. Having all the atoms in the same sub-level reduces collisions

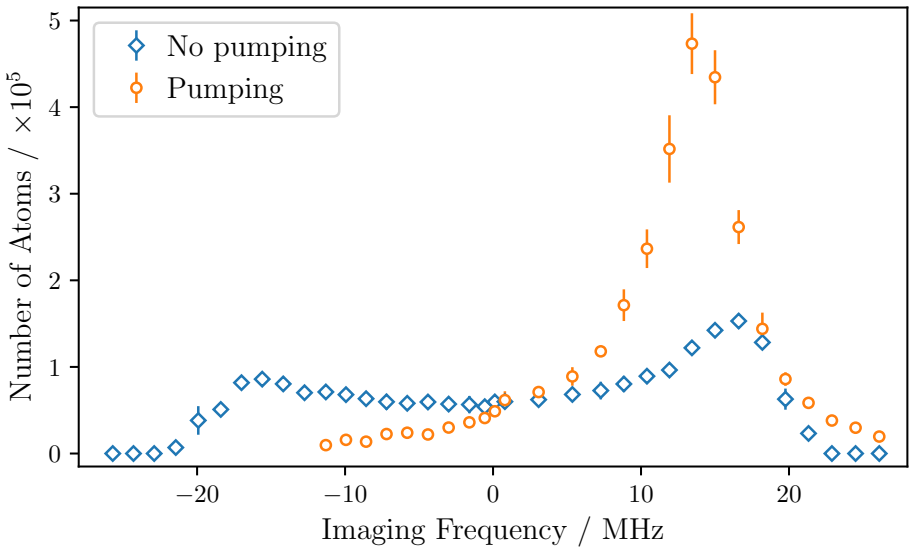

Figure 4.9: Plot of how the number of atoms imaged in the dipole trap varies with frequency, both before and after an optical pumping scheme was implemented. The atoms were imaged in a uniform magnetic field of about 8 G, leading to a breaking of the degeneracy of the Zeeman sub-levels. The state purity after optical pumping is clearly significantly better than it was before the optical pumping.

between trapped atoms, reducing heating from 2 and 3 body collision in the trap, allowing the trap to reach higher densities.

The typical way to determine what m_F state the atoms are in is to use a Stern-Gerlach experiment, where a large magnetic field gradient spatially separates the different states [42,85]. However, the limited field of view of the cameras makes this difficult. Additionally, this magnetic field gradient would require an additional coil that would be difficult to fit in around the rest of the experimental equipment. Instead, we utilise the frequency shift and degeneracy breaking of the sub-levels in a uniform magnetic field. This method requires a smaller total field than a Stern-Gerlach experiment, meaning the existing offset coils can be used. However, the sub-level distribution cannot be determined from a single image using this method, instead, it is necessary to combine many repeats at different imaging frequencies.

Due to the relatively small field used ($\sim 8\,\mathrm{G}$), we cannot accurately determine the exact population of each m_F state, as the linewidth of the transitions means they overlap. However, this is still a useful tool to diagnose the rough population distribution. A plot of how the number of atoms imaged in the dipole trap varies with frequency, both before and after an optical pumping scheme was implemented, can be seen in Fig. 4.9. Before the optical pumping was implemented, the atoms were clearly distributed across a wide range of the Zeeman sub-levels. After an optical pumping scheme was implemented, the state purity is much better, with the vast majority of atoms seeming to be in a single sublevel. This was mostly achieved by improving the polarisation of the MOT beams, which seemed to have drifted away from being correctly circularly polarised. Additionally, a pumping phase was implemented using a single beam of σ^+-polarised light from the cooler laser, in the presence of a small bias field.

Temperature

A plot showing a time-of-flight temperature measurement of the dipole trap can be seen in Fig. 4.10, demonstrating a temperature of $T = (83 \pm 4)\,\mu\mathrm{K}$. This temperature measurement was taken after improvements in the beam profile (but before optical pumping was implemented). Atoms trapped in the dipole trap are slightly hotter than they are cooled to in the MOT, however, the temperature is a lot better than the 1 mK observed before the beam profile was improved. The dipole beam can lead to additional heating due to power fluctuations, pointing instabilities, and imperfections in the beam profile [88]. There may also be heating from atoms in the trap colliding with the fibre.

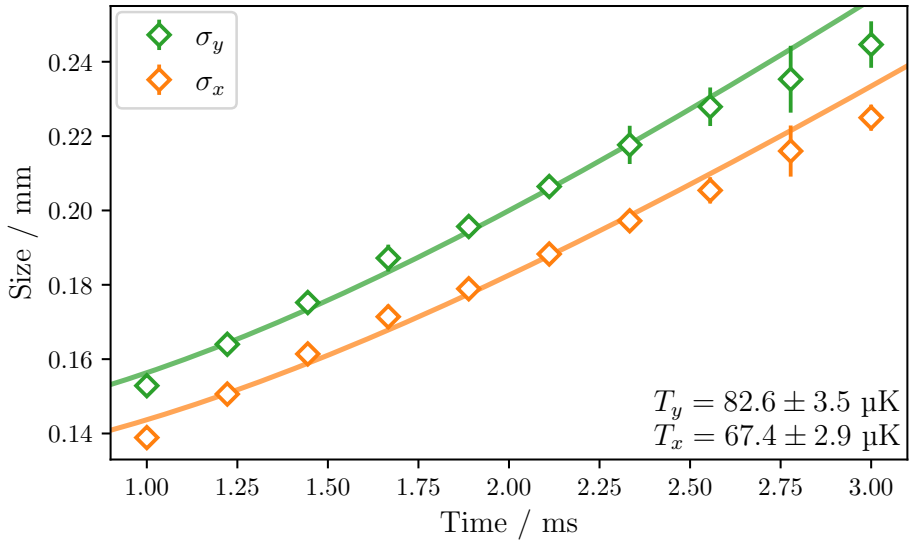

Figure 4.10: Time-of-flight temperature measurement of the dipole trap. Images were taken on the horizontal imaging axis. (Green) and (Red) data points represent the mean size of the atom cloud in the horizontal and vertical axes of the image, error bars represent the standard deviation in the data set used. Solid lies are the fit to Eq. 3.21.

Lifetime

A plot demonstrating the lifetime of atoms in our optical dipole trap can be seen in Fig. 4.11. Atoms were held in the trap for a varying amount of time between the MOT beams turning off, and the image being taken. The data are fit to an exponential decay, giving a lifetime of $\tau = (260 \pm 20)$ ms, with an inital atom number of $N = (4.6 \pm 0.4) \times 10^5$ atoms. This fit only models single-body losses (interactions between the trapped atoms and the background gas), and doesn't account for collisions between trapped atoms. This modelling of only single-body losses is not valid however, it still gives us a useful metric for how long one can keep atoms in the trap.

Alignment to Fibre

The final optimisation step of the dipole trap is aligning it to the fibre junction. There are several steps to this alignment: coarse alignment using

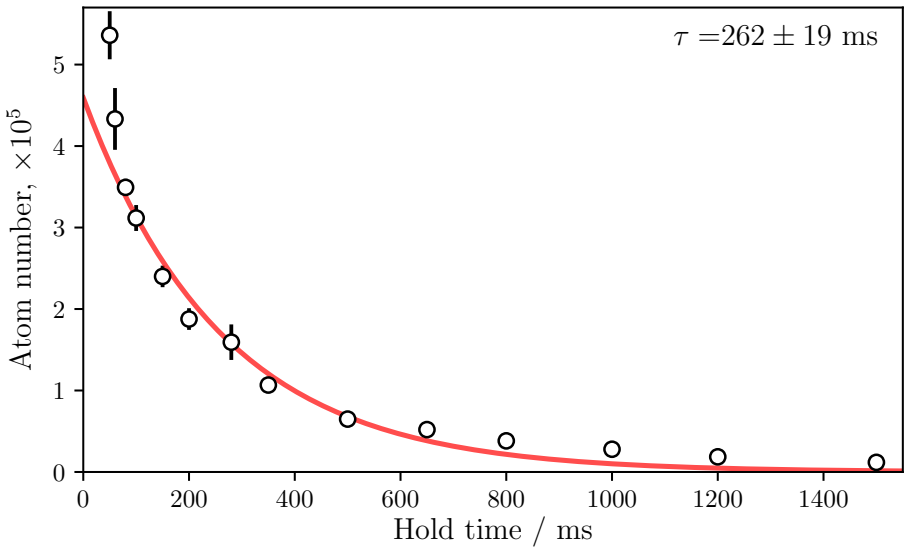

Figure 4.11: Plot of how the number of atoms imaged in the dipole trap decays for longer hold times. Black circles represent the average atom number across 3 images, with the error bars representing the standard deviation. The red curve represents a fit of the data to an exponential decay, modelling single-body losses. This fit gives a lifetime of $\tau = 260 \pm 20$ ms.

horizontal imaging, looking for atoms going through the fibre on the vertical imaging, and then using absorption through the fibre to finely optimise the alignment. This alignment requires a high degree of precision, due to the small size of the fibre hole. For this, the final mirror in the dipole beam path has been customised to use a high precision adjuster[2]. The adjuster used can be easily moved with a precision of ~ 0.5µm, allowing the angle of the mirror to be controlled to ~ 10µrad (at ~ 0.5mrad/rev).

The first step involves running the dipole laser at low power ($\sim$1 W) while recording at the beam on the horizontal imaging camera. By looking at distortions in the beam shape, one can determine when the beam is hitting the fibre. Then, by paying particular attention to fine features in the distortion, the location of the fibre hole can be determined, and the approximate range of motion for the beam alignment where it is still passing through the hole can be inferred.

[2]Newport 9853 gimbal mount with 2 Thorlabs DM22 differential Adjusters

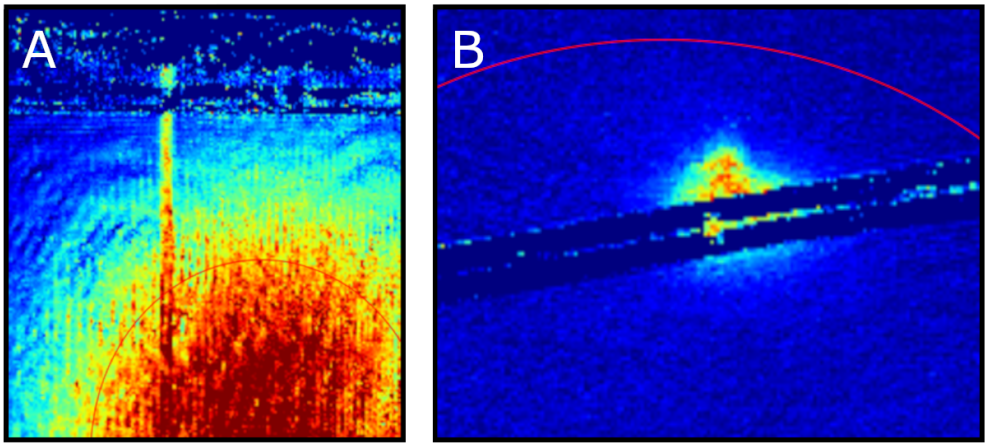

Figure 4.12: Cropped images of the dipole trap on the (A) vertical and (B) horizontal imaging axes. (A) was taken with 10 ms of hold time, so the MOT atoms are still visible, while (B) was taken with 50 ms hold time and the MOT atoms are not visible. The dipole trap in these images is aligned through the fibre hole. Atoms in the trap can be seen passing on both sides of the fibre in (A), implying the trapping field passes cleanly through the fibre hole.

Then we switch to absorption imaging on the vertical axis. There is a small gap between the fibre and the glass slide it is mounted to. By looking at the atoms imaged in this region, we can confirm whether the beam is well aligned. If the beam is far from the fibre, atoms will be clearly visible on both sides of the fibre. If the beam is aligned to the fibre, but not well aligned to the hole, there will not be any atoms from the dipole trap visible in this region. They will only be trapped before the fibre. There will also be some additional losses due to the beam clipping on the fibre. If the beam is well aligned through the fibre hole, then atoms will be visible on both sides of the fibre. Fig. 4.12 shows images of the dipole trap aligned to the fibre hole on both the vertical and horizontal axes. On the vertical images, atoms are clearly still trapped in the region between the fibre and the slide, confirming the beam is aligned through the fibre hole.

4.4 Fibre absorption

Interacting with atoms trapped in the fibre junction is the final stage of the experimental procedure. The atoms trapped in this junction will be used for the later experiments on EIT (Ch. 5) and Bayesian optimisation (Ch. 6). Having the largest number of atoms trapped in this region is the final parameter to optimise. This is optimised by ensuring the earlier experimental steps are optimised, and by fine-tuning the alignment of the final mirror in the dipole beam path (see Section above for more details).

Details of how the atoms in the fibre junction are measured are described in more detail in Sec. 3.2.2. The trap is turned off, before 2 probe pulses are sent through the fibre. The first about 2 µs after the trap is turned off, the second about 200 ms after the trap is turned off, so that transmittance of the atoms can be determined. The key equation is that the number of atoms in the interaction region is given by

$$N = \frac{A}{\sigma_0}\mathrm{OD}_{res} \tag{4.1}$$

where A is the area of the imaging beam, σ_0 is the scattering cross section of the transition, and OD_{res} is the on-resonance optical depth, which is related to the transmittance of the atoms by

$$T = \exp\left(-\frac{\mathrm{OD}_{res}}{1 + 4(\delta/\Gamma)^2}\right) \tag{4.2}$$

where $\delta = \omega - \omega_0$ is the detuning of the laser from the atomic resonance, and Γ is the linewidth of the transition.

When measuring the transmission on resonance, the lowest transmittance we have observed so far is $(0.59 \pm 0.30)\,\%$. This corresponds to an optical

depth of $OD_{res} = 5.1 \pm 0.5$, so there is a total of 436 ± 40 atoms in the interaction region. This corresponds to an optical depth per unit length of 1700/cm (over the length of the fibre junction) and an atomic density of $7.4 \times 10^{11}/\text{cm}^3$. This optical depth per unit length is higher than equivalent hollow-core fibre or tapered optical nanofibre systems [35,100] and is higher than what was observed previously in this system [42,44], showing the experimental changes implemented since have resulted in significant improvements.

The atom number and atom density estimates here have used the resonant cross-section for π-polarised light [84]. While this gives a reasonable estimate, there are a few unknown factors that this makes certain assumptions about. Notably, the unknown polarisation of the light in the junction, and the unknown Zeeman sub-level distribution of the atoms. The assumption of linear polarisation was made as that is the polarisation state coupled into the fibre, and (despite not being a polarisation maintaining fibre) the polarisation coming out the fibre is very stable and consistent. These assumptions do lead to a large systematic shift in the estimates of the atom number, however, ultimately the more important metric to consider is the optical depth, which can be determined with high precision and accuracy.

Absorption spectrum

It is often useful to investigate the transmittance as a function of the detuning. This allows the atom number to be determined with better accuracy and to determine the linewidth of the transition to check if there are any sources of broadening. A plot of this can be seen in Fig. 4.13. The red curve is a fit of the data to Eq. 4.2. From this fit the optical depth is determined to be $OD_{res} = 3.4 \pm 0.3$ and the linewidth of the transition is

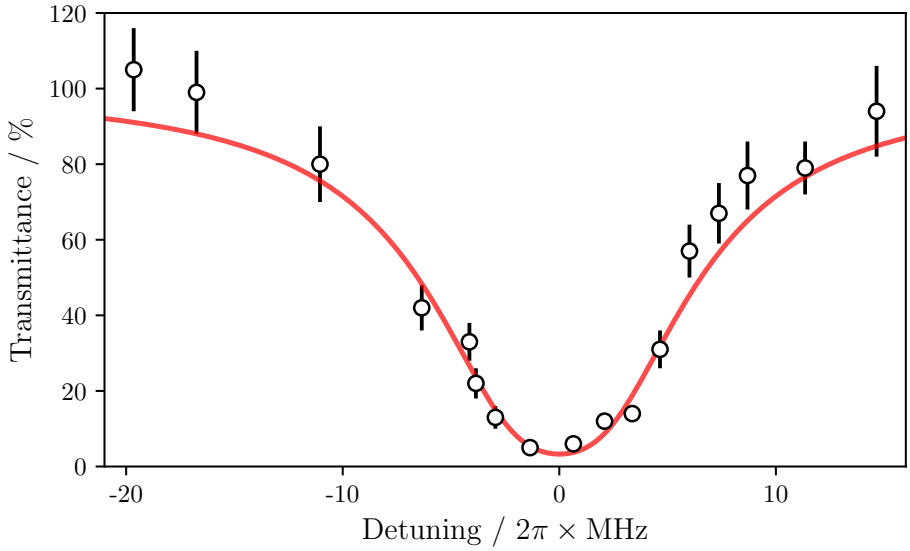

Figure 4.13: Transmittance through the fibre junction as a function of detuning. Imaging on the $|F = 4\rangle \rightarrow |F' = 5\rangle$ transition. Each data point represents the mean of 30 data point and the error bars are the standard deviation from that data set.

$\Gamma = 2\pi \times (6.60 \pm 0.56)$ MHz. For this measurement, the optical depth corresponds to an atom number of 290 ± 26 (assuming linearly polarised light). The linewidth here is slightly broadened from the natural linewidth [84], this is due to collisional broadening as the local background pressure near the fibre is higher than it is elsewhere in the chamber. This local pressure increase is due to the material the fibre coating is made from, and could be reduced by stripping away more of the fibre coating.

Lifetime

The lifetime of atoms in the dipole trap is an important parameter that can inform us of any heating of the atoms, from collisions or alignment. It can be determined by holding the atoms in the trap for varying lengths of time before imaging the atoms. Assuming losses from the trap are dominated by

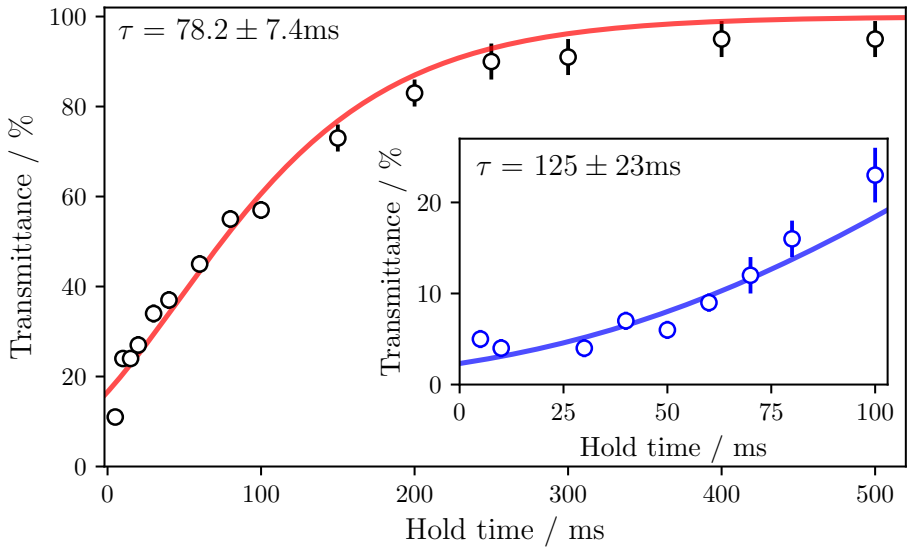

Figure 4.14: Transmittance through the fibre as a function of hold time (time between MOT end and the imaging pulse), to determine the lifetime of the trapped atoms. The main figure shows the lifetime before the optical pumping scheme was implemented and the inset shows the lifetime after it was implemented (see Sec. 4.3). Each data point represents the mean of 30 data points and the error bars are the standard deviation from that data set.

single-body losses, they can be described by an exponential decay as [101]

$$\frac{\mathrm{d}N}{\mathrm{d}t} = -\alpha N\,. \tag{4.3}$$

This means the transmittance as a function of time is given by

$$T(t) = \exp\left[-\mathrm{OD}_{res}\, e^{t/\tau}\right] \tag{4.4}$$

where τ is the lifetime of atoms in the trap.

Plots of how transmittance varies as a function of time can be seen in Fig. 4.14. The main body of the figure shows the lifetime of atoms before the optical pumping scheme was implemented, giving a lifetime of $\tau = (78 \pm 7)\,\mathrm{ms}$, with an initial optical depth of $\mathrm{OD}_{res} = 1.80 \pm 0.09$. The inset shows the lifetime after the optical pumping scheme was implemented, giving $\tau = (125 \pm 20)\,\mathrm{ms}$, with an initial optical depth of $\mathrm{OD}_{res} = 3.8 \pm$

0.3. Comparing this to the lifetime of similar nanofibre traps, our system outperforms many of them, which have lifetimes of $< 50\,\mathrm{ms}$ [30,32,102, 103]. However, with active sub-Doppler cooling of the trapped atoms, nanofibre traps can achieve much higher lifetimes of $> 1\,\mathrm{s}$ [104], limited by the background pressure.

It can be seen that the pumping has improved the lifetime of atoms in our trap, as expected. However, it's interesting to note that the lifetime measured here is significantly lower than that measured for the dipole trap when aligned away from the fibre (see Fig. 4.11). We suspect the reason for this is the dipole beam causing excess heating of the fibre. The "wings" of the high power laser are absorbed by the fibre, heating it up, causing material to be ejected from the surface of the fibre. The particle ejection increases the local background pressure, increasing the rate of single body collisions, and hence the rate of losses from the trap. This problem has also been exacerbated due to a small "burn" in the current system. This heating is why the data for the inset in Fig. 4.14 doesn't go to longer times.

Temperature

One important parameter to optimise when imaging atoms in the fibre junction is the exposure time of the imaging light. Due to the Poissonian nature of photon shot noise, a longer exposure time leads to lower relative noise on the photon counts. However, longer exposure time can also decrease the signal, due to optical pumping effects and from atoms moving out of the imaging region due to their thermal motion.

To investigate which of these was the dominant mechanism for the reduction in signal, we measured the absorption through the fibre for various exposure times and intensities. If pumping is the dominant mechanism,

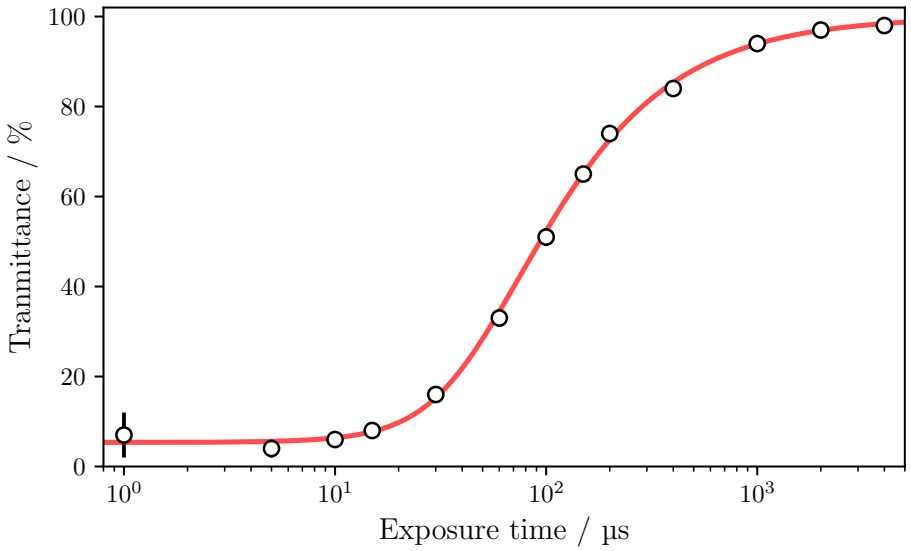

Figure 4.15: Transmittance through the fibre as a function of exposure time. The atom cloud is undergoing thermal expansion while the imaging pulse is interacting with the cloud. This allows the temperature of the cloud to be determined.

there should be a power dependence to the signal. However, thermal expansion will only depend on time, and not on the power of the probe.

A table of how the transmittance varies with both exposure time and power can be seen in Tab. 4.2. The data gathered showed no power dependence (in the region we were measuring), so thermal expansion is the dominant mechanism reducing the signal strength with longer exposure time. This means a plot of transmittance vs exposure time can be used to determine the temperature of atoms in the fibre junction. A plot of how this can be seen in Fig. 4.15.

	Exposure time		
Power	10 µs	30 µs	60 µs
1 pW	(5 ± 3) %	(11 ± 1) %	(30 ± 2) %
10 pW	(4 ± 1) %	(14 ± 1) %	(32 ± 1) %

Table 4.2: Table of how the transmittance varies with both power and exposure time. There is a strong dependence on exposure time, but very little power dependence. This suggests thermal expansion is the dominant mechanism reducing the transmittance with increasing exposure.

To determine the temperature from this plot, we expand upon the work done in Sec. 3.3. Assuming the atoms follow a Gaussian spatial distribution, the density of the cloud evolves as

$$n(r,t) = \frac{N}{\pi^{3/2}\left(\sigma_r^2 + \frac{\sigma_p^2}{m^2}\,t^2\right)^{3/2}}\ \exp\left[-\frac{r^2}{\sigma_r^2 + \frac{\sigma_p^2}{m^2}\,t^2}\right] \tag{4.5}$$

where N is the number of atoms in the cloud, r is the distance from the centre of the cloud, and $\sigma_p^2 = mk_BT$ where m is the mass of the atom (as in Sec. 3.3) [98,105,106]. This cloud is illuminated by a narrow beam, propagating in the $\hat{z}$ direction. Following Eq. 3.7, the transmittance through the cloud as a function of time is then [106]

$$T_{inst}(t) = \frac{I(t)}{I_0} = \exp\left[-\frac{\sigma_0\,N}{\pi\left(\sigma_r^2 + \frac{\sigma_p^2}{m^2}\,t^2\right)}\right]\,. \tag{4.6}$$

The equation above gives the instantaneous absorption at a particular fall time t (the time between the trap turning off and the image being taken). However, the data in Fig. 4.15 is a continuous cumulative measurement from $t = 0$ until the end of the exposure time $t = t_{exp}$. So that data follows the time average of Eq. 4.6,

$$T(t_{exp}) = \frac{1}{t_{exp}}\int_0^{t_{exp}} \exp\left[-\frac{\sigma_0\,N}{\pi\left(\sigma_r^2 + \frac{\sigma_p^2}{m^2}\,t^2\right)}\right]\,dt\,. \tag{4.7}$$

The red line in Fig. 4.15 is a fit of the data to this equation. Setting the total atom number as $N = 270$ (which is the atom number estimate from the data point with the highest absorption in that set), we can extract a temperature estimate of $T = (69 \pm 6)\,\mu\text{K}$.

This estimate is in reasonable agreement with the temperatures measured

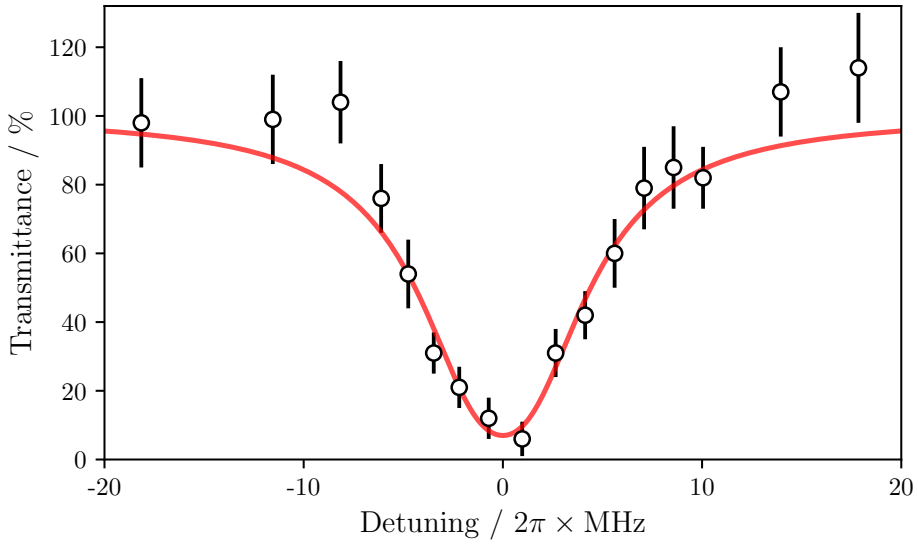

Figure 4.16: Transmittance through the fibre junction as a function of detuning. Imaging on the $|F = 4\rangle \rightarrow |F' = 4\rangle$ transition. Each data point represents the mean of 30 data points and the error bars are the standard deviation from that data set.

in Sec. 4.3. However, there are a few assumptions made that aren't truly valid. First, the atom number used only accounts for atoms initially in the imaging region, so the equation only models those atoms moving out of the interaction region and doesn't account for other atoms moving into it. Additionally, Eq. 4.5 assumes a spherically symmetric trap, which (as can be seen in Fig. 4.12(A)) is clearly not the case for this dipole trap. These assumptions suggest the estimate above is likely a lower bound on the true temperature of the atoms in the interaction region.

4-4 transition

It is also informative to image the atoms in the fibre junction on the $|F = 4\rangle \rightarrow |F' = 4\rangle$. This is the transition the probe addresses for the EIT experiments (see Ch. 5), so it is informative to check the behaviour of atoms on this transition with no control field present. The dependence of the transmittance to detuning on this transition can be seen in Fig. 4.16.

The on-resonance optical depth here is 2.7 ± 0.5, this optical depth gives an atom number estimate of 480 ± 90.

One thing to be cautious of when imaging on this transition is that, unlike the $|F = 4\rangle \rightarrow |F' = 5\rangle$ transition, it is not a cycling transition. So, if there is no repumper light, the atoms will be quickly pumped to be $|F = 3\rangle$ ground state during the imaging pulse, increasing the transmittance of long pulses.

Chapter 5

Electromagnetically induced transparency

Electromagnetically induced transparency (EIT) is a phenomenon observed when a 3-level system interacts with 2 light fields. The absorption of a resonant laser beam can be eliminated due to the presence of another laser field, making the medium transparent [107,108]. While this has applications as an all-optical switch [109,110], the property that generally attracts more interest is that it also leads to a slow light effect. In 1999, using EIT in an ultra-cold cloud of Sodium, a photon group velocity of 17 m/s was recorded [111]. Additionally, using stimulated Raman adiabatic passage (STIRAP), the light field can be transferred to a spin-wave in the atomic population. This process can be controllably reversed, leading to photon storage [112]. Photon storage times of up to 1 s have been observed in warm Caesium vapours [113]. Recent work has also been done to achieve storage efficiencies of $> 80\%$ [56,114–116].

Here I present the theory of 3-level atom-light interactions and how they lead to EIT (Sec. 5.1), and how this transparency leads to slow and stored

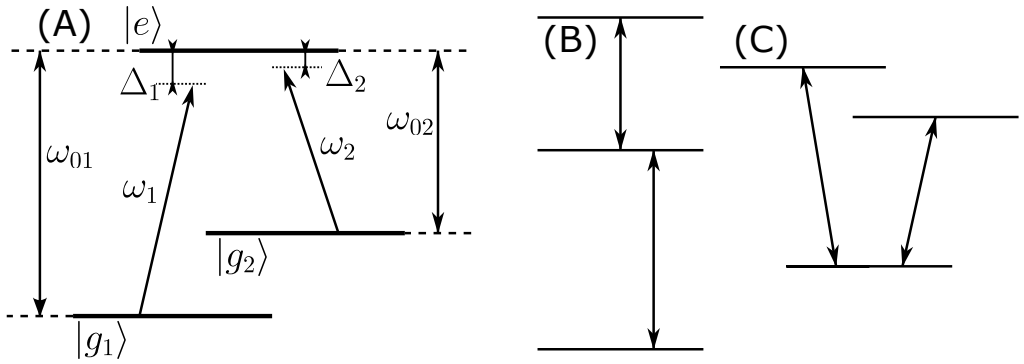

Figure 5.1: Generic energy level diagrams for an EIT system. Different arrangements of a 3-level structure are shown: (A) a Λ system, (B) a ladder system, (C) a V system. The frequencies of the energy levels and optical fields are shown on the λ structure.

light (Sec. 5.1.1). I then present my results of implementing EIT in a micro-scale hole through an optical fibre, as a method to integrate a cold atom photon memory with optical waveguides (Sec. 5.2).

5.1 EIT theory

Electromagnetically induced transparency (EIT) can be observed in any 3-level system (shown in Fig. 5.1). However, for this analysis, we will focus on a Lambda scheme, as that is the type of scheme used in this experiment (and the most common one for photon storage experiments). The Hamiltonian for the atom at rest is [42,82,107]

$$H_A = -\hbar\omega_{01} |g_1\rangle \langle g_1| - \hbar\omega_{02} |g_2\rangle \langle g_2| \tag{5.1}$$

where ω_{0i} is the frequency that couples the ground state $|g_i\rangle$ to the excited state $|e\rangle$, and the excited state energy is set to be 0. Two lasers interact

with these transitions with the electric fields

$$\begin{aligned}\mathbf{E_1}(t) &= \hat{\epsilon}_1 E_{01}\cos(\omega_1 t)\\ \mathbf{E_2}(t) &= \hat{\epsilon}_2 E_{02}\cos(\omega_2 t)\end{aligned} \tag{5.2}$$

where ω_i is their angular frequency, E_{0i} is the electric field amplitude, ϵ_i is the unit polarisation vector, and we have assumed the atom is small compared to the wavelength of the light (dipole approximation). It is convenient to express the total electric field as

$$\mathbf{E}(t) = \mathbf{E_1}(t) + \mathbf{E_2}(t) = \mathbf{E}^{(+)}(t) + \mathbf{E}^{(-)}(t) \tag{5.3}$$

where

$$\mathbf{E}^{(\pm)}(t) = \frac{1}{2}\left(\hat{\epsilon}_1 E_{01} e^{\mp i\omega_1 t} + \hat{\epsilon}_2 E_{02} e^{\mp i\omega_2 t}\right) . \tag{5.4}$$

The interaction between the atom and this electric field is given by

$$H_{AF} = -\mathbf{d}\cdot\mathbf{E} \tag{5.5}$$

where $\mathbf{d}$ is the atomic dipole operator given by

$$\begin{aligned}\mathbf{d} &= \mathbf{d}^{(+)} + \mathbf{d}^{(-)}\\ &= \Big[\langle g_1|\,\mathbf{d}\,|e\rangle\,\sigma_1 + \langle g_2|\,\mathbf{d}\,|e\rangle\,\sigma_2\Big] + \Big[\langle g_1|\,\mathbf{d}\,|e\rangle\,\sigma_1^\dagger + \langle g_2|\,\mathbf{d}\,|e\rangle\,\sigma_2^\dagger\Big]\end{aligned} \tag{5.6}$$

where $\sigma_i = |g_i\rangle\langle e|$.

Making the rotating wave approximation, and transforming into the rotating frame, these Hamiltonians can be expressed as [82,107]

$$\tilde{H}_A = \hbar\Delta_1 |g_1\rangle\langle g_1| + \hbar\Delta_2 |g_2\rangle\langle g_2| \tag{5.7}$$

and

$$\begin{aligned}\tilde{H}_{AF} &= -\mathbf{d}^{(+)} \cdot \mathbf{E}^{(-)} - \mathbf{d}^{(-)} \cdot \mathbf{E}^{(+)} \\ &= \frac{\hbar\Omega_1}{2}\left(\sigma_1 + \sigma_1^\dagger\right) + \frac{\hbar\Omega_2}{2}\left(\sigma_2 + \sigma_2^\dagger\right)\end{aligned} \tag{5.8}$$

where $\Delta_i = \omega_i - \omega_{0i}$ is the detuning between the laser frequency and the atomic resonance, and $\Omega_i = -\langle g_i|\, \hat{\epsilon}_i \cdot \mathbf{d}\, |e\rangle\, E_{0i}/\hbar$ is the Rabi frequency of the $E_i(t)$ beam [82]. The eigenstates of the total Hamiltonian $\tilde{H} = \tilde{H}_A + \tilde{H}_{AF}$ when both fields are resonant ($\Delta_1 = \Delta_2 = 0$) are

$$\begin{aligned}|0\rangle &= \frac{\Omega_2 |g_1\rangle - \Omega_1 |g_2\rangle}{\sqrt{\Omega_1^2 + \Omega_2^2}} \\ |-\rangle &= \frac{\Omega_1 |g_1\rangle + \Omega_2 |g_2\rangle}{\sqrt{2}\sqrt{\Omega_1^2 + \Omega_2^2}} - \frac{1}{\sqrt{2}}|e\rangle \\ |+\rangle &= -\frac{\Omega_1 |g_1\rangle + \Omega_2 |g_2\rangle}{\sqrt{2}\sqrt{\Omega_1^2 + \Omega_2^2}} + \frac{1}{\sqrt{2}}|e\rangle\ .\end{aligned} \tag{5.9}$$

The key eigenstate to notice is $|0\rangle$. As this eigenstate does not contain the excited state $|e\rangle$, there is no way for this state to absorb or emit a photon. As this $|0\rangle$ state does not interact with either laser field, it is a dark state. Atoms can be pumped into this dark state from the $|\pm\rangle$ states via spontaneous emission from $|e\rangle$, in a phenomenon called coherent population trapping [82,107]. Atoms interacting with these 2 light fields get pumped into a dark state and stop interacting with them.

In the case most relevant for EIT, both fields are on resonance with the atomic transitions ($\Delta_1 = \Delta_2 = 0$), and one field (the control) is strong compared to the other (the probe) ($\Omega_1 \gg \Omega_2$). Here, the eigenstates be-

come

$$\begin{aligned} |0\rangle &\approx |g_2\rangle \\ |-\rangle &\approx \frac{|g_1\rangle - |e\rangle}{\sqrt{2}} \\ |+\rangle &\approx \frac{|g_1\rangle + |e\rangle}{\sqrt{2}}\,. \end{aligned} \tag{5.10}$$

The dark state is a bare state of the atom, and that atomic state does not appear in either of the $|\pm\rangle$ states. So under EIT conditions, $|g_2\rangle$ does not interact with either field, making the medium transparent to the $\mathbf{E_2}$ field.

5.1.1 Susceptibility

In order to get a more thorough understanding of this transparency phenomenon, it is necessary to consider the effect it has on the susceptibility of the atomic medium. Here, we label $\mathbf{E_1}$ as a strong "control" field and $\mathbf{E_2}$ as a weak "probe" ($\Omega_2 \ll \Omega_1$). We also treat the atoms as initially occupying the $|g_2\rangle$ ground state. To calculate the linear susceptibility of the probe beam, we need to calculate the coherence between this ground state and the excited state $\tilde{\rho}_{eg_2}$. This can be found from the master equation

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}\tilde{\rho} = -\frac{i}{\hbar}\left[\tilde{H}_A + \tilde{H}_{AF}, \tilde{\rho}\right] &+ \frac{\Gamma_1}{2}\left[2\sigma_1\tilde{\rho}\sigma_1^\dagger - \sigma_1^\dagger\sigma_1\tilde{\rho} - \tilde{\rho}\sigma_1^\dagger\sigma_1\right] \\ &+ \frac{\Gamma_2}{2}\left[2\sigma_2\tilde{\rho}\sigma_2^\dagger - \sigma_2^\dagger\sigma_2\tilde{\rho} - \tilde{\rho}\sigma_2^\dagger\sigma_2\right] \\ &+ \frac{\gamma_g}{2}\left[2\sigma_g\tilde{\rho}\sigma_g^\dagger - \sigma_g^\dagger\sigma_g\tilde{\rho} - \tilde{\rho}\sigma_g^\dagger\sigma_g\right] \end{aligned} \tag{5.11}$$

where the 2nd and 3rd terms describe spontaneous emission from the excited state to the 2 ground states with rates Γ_1 and Γ_2, and the final term describes dephasing between the 2 ground states with rate γ_g and $\sigma_g = |g_2\rangle\langle g_2| - |g_1\rangle\langle g_1|$ [82]. This dephasing can arise from a number of sources, the most significant ones here being atomic collisions and phase

fluctuations in the lasers [107,117]. Solving Eq. 5.11 for $\tilde{\rho}_{eg_2}$ we get

$$\tilde{\rho}_{eg_2} = i\frac{\Omega_2}{2}\frac{\delta_{12} + i\gamma_g}{(i\Delta_2 - \Gamma_2/2)(\delta_{12} + i\gamma_g) - i(\Omega_1/2)^2} \tag{5.12}$$

where $\delta_{12} = \Delta_2 - \Delta_1$. From this, we can find the polarisation of the medium by the probe beam using

$$P_{probe} = \epsilon_0 \chi E_2 = n \langle g_2 | \hat{\epsilon}_2 \cdot \mathbf{d} | e \rangle \tilde{\rho}_{eg_2} \tag{5.13}$$

where n is the number density of the atomic medium, and χ is the linear susceptibility. So the linear susceptibility is given by [82]

$$\chi = -i\frac{n|\langle g_2 | \hat{\epsilon}_2 \cdot \mathbf{d} | e \rangle|^2}{\epsilon_0 \hbar}\frac{\delta_{12} + i\gamma_g}{(i\Delta_2 - \Gamma_2/2)(\delta_{12} + i\gamma_g) - i(\Omega_1/2)^2}\,. \tag{5.14}$$

Equation 5.14 fully describes the optical response of the medium to the probe beam. The real component of χ determines the refractive index of the medium, while the imaginary component determines the absorptive properties. In the case of two-photon resonance ($\delta_{12} = 0$) and no ground state decoherence ($\gamma_g = 0$), both real and imaginary components drop to 0, so the medium is fully transparent to the probe beam, as expected from the derivation of the states above. The transmission through the atom cloud is related to the complex susceptibility by

$$T = \exp\left(-\frac{k_0 z}{2}\,\mathrm{Im}\{\chi\}\right) \tag{5.15}$$

where k_0 is the resonant wave number of the probe beam, and z is the length of the interaction region [107].

Plots showing how the transmission of the atomic medium changes with various Rabi frequencies, detunings, and dephasing rates can be seen in

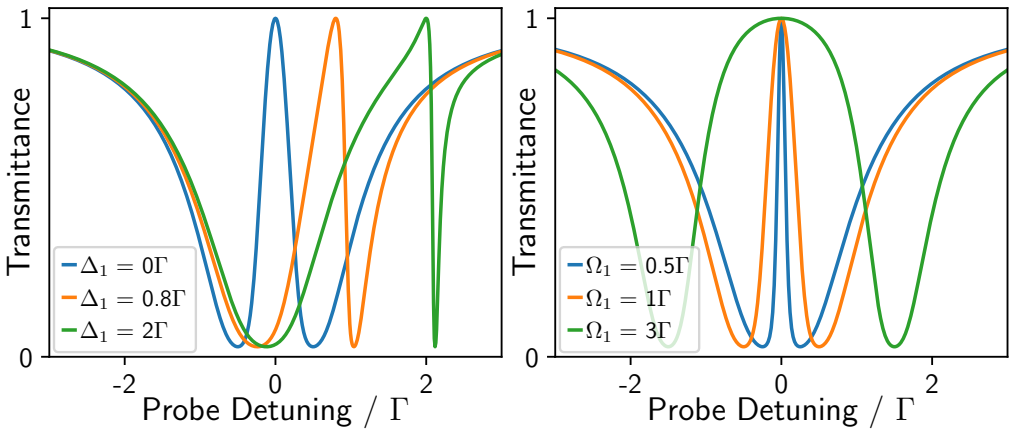

Figure 5.2: Plots showing the expected response of an EIT system to probe frequency, based on Eq. 5.15. These are plotted assuming $\gamma_g = 0$. Left=: How the response changes for different control detunings Δ_1, with $\Omega_1 = \Gamma$. Right: How the response changes for different control Rabi frequencies Ω_1, with the control on resonance ($\Delta_1 = 0$).

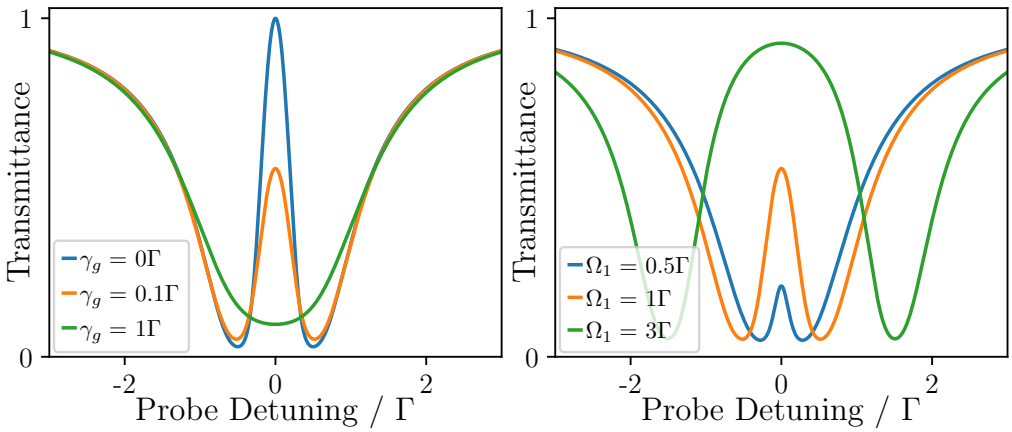

Figure 5.3: Plots showing the expected response of an EIT system to probe frequency, based on Eq. 5.15, with non-zero dephasing γ_g. Left: How the response changes for different dephasing rates γ_g, with $\Omega_1 = 1\,\Gamma$. Right: How the response changes for different control Rabi Frequencies Ω_1, with $\gamma_g = 0.1\Gamma$. Both plots have a resonant control, $\Delta_1 = 0$.

Figs 5.2 & 5.3. Two key points to notice are that the transparency occurs at the two-photon resonance and that lowering the control intensity makes the transparency window narrower. It is also important to note the effect dephasing has in reducing the height of the transparency. The EIT transparency transmittance is only 1 when $\Omega_1 \gg \gamma_g$, and if $\Omega_1 \approx \gamma_g$ the transparency disappears entirely.

As an aside, it is interesting to note that in the case of very high intensities ($\Omega_1 \gg \Gamma$) the transparency can actually be attributed to another effect called Autler-Townes splitting [107,118]. The AC Stark effect from the control field causes the excited state to split into a doublet state with a frequency difference of Ω_1 between the 2 dressed states. The absorption spectrum of the probe beam is then the sum of 2 separate Lorentzians, each with a linewidth $\Gamma/2$. Which is the behaviour of Eq. 5.14 in the limit of large Ω_1. The key difference between Autler-Townes splitting and EIT is that when $\gamma_g = 0$ and $\Delta = 0$ EIT predicts perfect transparency regardless of the laser intensities, whereas Autler-Townes splitting does not. So the narrow transparency observed for EIT cannot be explained only through the AC Stark effect, instead, one has to consider the coherences between the excitation pathways.

Slow light

One of the most interesting aspects of EIT is how it can lead to slow light and photon storage. To understand where this slow light phenomenon comes from, it is important to consider the effects of the real part of the susceptibility. The real component of the refractive index is given by [82, 107]

$$\mathrm{Re}\{n\} = \sqrt{1 + \mathrm{Re}\{\chi\}} \approx 1 + \frac{\mathrm{Re}\{\chi\}}{2}\,. \tag{5.16}$$

As $\mathrm{Re}\{\chi\} \ll 1$ the refractive index is near unity, and the phase velocity of light through the medium is unaffected. However, the transparency peak in $\mathrm{Im}\{\chi\}$ corresponds to a steep positive gradient in $\mathrm{Re}\{\chi\}$, which is linked to a strong reduction of the group velocity of the probe pulse

$$v_g = \frac{\mathrm{d}\omega_2}{\mathrm{d}k} = \frac{c}{n + \omega_2(\mathrm{d}n/\mathrm{d}\omega_2)}\,. \tag{5.17}$$

This allows the group velocity of the probe pulse to be slowed by several orders of magnitude, with one group famously slowing light to 17 m/s through a sodium Bose-Einstein condensate [111].

This slow light effect is associated with what is known as a dark state polariton, a quasiparticle that is a coherent mixture of an electromagnetic wave and an atomic spin excitation wave. Before the probe beam is turned on, the control pumps the atoms fully into the $|0\rangle = |g_2\rangle$ dark state. When the probe laser turns on, Ω_2 increases and some of the dark-state atoms are transferred to the $|g_1\rangle$ atomic state through stimulated Raman adiabatic passage (STIRAP). This polariton is how light is "stored" using EIT. While the probe pulse is propagating through the medium, the intensity of the control beam is ramped down until $\Omega_1 = 0$. If this is done adiabatically, the dark-state polariton shifts from an electromagnetic wave-atom spin combination to purely being a spin wave, and the velocity of the spin wave is reduced to 0. Then, by ramping the control intensity back up, this process is reversed and the electromagnetic wave is "retrieved" from the medium [107,113].

5.2 Experimental results

To achieve EIT in the fibre junction, both the probe and control light are coupled co-propagating into the fibre (see Fig. B.3 for an optics diagram). The 2 beams are combined at a polarising beam splitter (so both fields have orthogonal linear polarisation going into the fibre), and then the control field is filtered out before the SPCM using a linear polariser. When testing EIT, it is important to start by looking for the characteristic transparency peak. The correct frequencies can be found by setting the probe to resonance (which can be found using an absorption profile as in Sec. 4.4) while scanning the control frequency and looking for an increase in the transmission. Once a suitable control frequency has been found, one can then produce the characteristic transmission profiles by keeping the control frequency fixed while scanning the probe frequency. Some of the profiles observed in this experiment can be seen in Fig. 5.4. These curves have been fit to the transmittance in Eq. 5.15.

These fits allow us to extract information on the control detuning, control intensity, and the dephasing rate. All the fitting parameters are listed in Table 5.1, the Rabi frequencies are additionally included in Fig. 5.4. Panels (a)–(c) were all taken close together in time, with approximately the same atom density. Panels (a) and (b) show the position of the trans-

Panel	(a)	(b)	(c)	(d)
$\Delta_1\ /2\pi \times$ MHz	-2.2 ± 0.3	0.7 ± 0.3	-0.2 ± 1.3	2.3 ± 1.4
$\Omega_1\ /2\pi \times$ MHz	6.8 ± 0.3	6.5 ± 0.6	3.6 ± 1.0	7.3 ± 1.0
$\Gamma_2\ /2\pi \times$ MHz	7.4 ± 0.6	6.9 ± 0.8	7.0 ± 3.2	15.0 ± 3.2
$\gamma_{12}\ /2\pi \times$ MHz	1.5 ± 0.2	1.7 ± 0.2	0.83 ± 0.30	0.39 ± 0.31
$n\ / \times 10^{11}\,\mathrm{cm}^{-3}$	7.8 ± 0.3	7.0 ± 0.4	7.6 ± 1.0	44.6 ± 6.0

Table 5.1: Table of the fitting parameters for the 4 curves in Figure 5.4. The parameters are control detuning Δ_1, control Rabi frequency Ω_1, linewidth Γ_2, ground-state dephasing rate γ_{12}, and atom number density n.

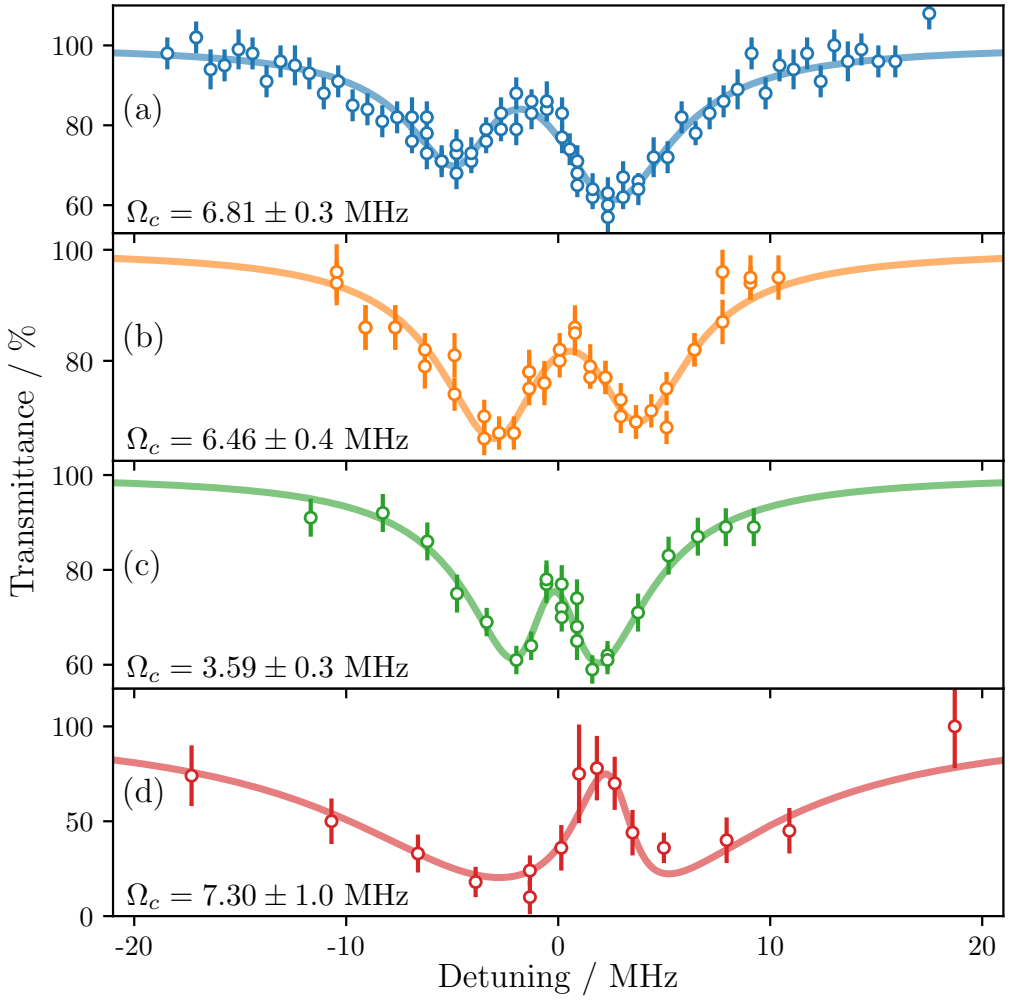

Figure 5.4: Scans of probe detuning vs transmittance with EIT. Scans are done at different control intensities and frequencies. Panel (d) has a much higher atom number than the other panels. A transparency peak can be clearly seen in the centre of the transition. Solid lines show a fit to Eq. 5.15. Fitting parameters are all listed in Table 5.1.

parency peak shift as the control frequency is changed (from -2.2MHz to +0.7MHz) while the intensity remains constant. Panels (b) and (c) show that the transparency window gets narrower as the control intensity is reduced while the frequency is kept the same. The dephasing rate is also reduced between these, this was done by minimising the magnetic fields in the fibre junction. Panel (d) was taken much later than the other 3, after several significant improvements had been made to the experiment. First, improvements to the dipole beam profile gave a 6-fold improvement in the atom density. Additionally, the dephasing has been reduced by making the probe and control beams co-propagating [119,120], while keeping the magnetic fields minimised (the beams were counter-propagating in the other 3 plots). The fitted value for the linewidth of panels (a) to (c) are all slightly broadened from caesium's natural linewidth, this is for the same reason given in Sec. 4.4 (background pressure near the fibre leads to some collisional broadening). The linewidth of (d) is much more significantly broadened. This is also due to collisional broadening, however for that one a small "burn mark" on the fibre means light from the dipole laser causes significant heating of the fibre. As the fibre heats up, material is ejected from the surface of the fibre, increasing the background pressure further, leading to a higher rate of collisions, and hence more collisional broadening.

From these fits, we can calculate the real component of the susceptibility, from which we can estimate the group velocity. An example of this can be seen in Fig. 5.5, which is based on the fit in Fig. 5.4(d). A steep positive gradient can be seen at the same frequencies as the transparency peak. Putting this gradient into Eq. 5.17, we get a group velocity of 310 m/s, which implies a delay time of 97 ns for photons travelling through the fibre junction.

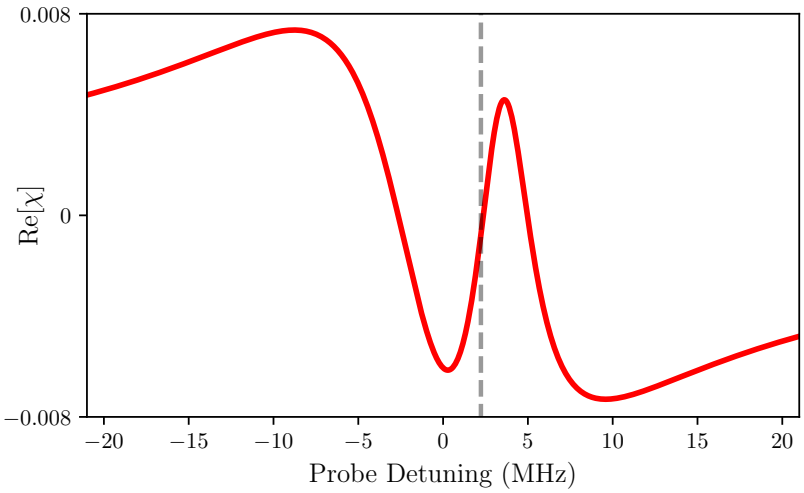

Figure 5.5: Curve showing an estimate of the real component of the susceptibility, based on the fitting parameters for Fig. 5.4(d). From this, we can estimate the group velocity by finding the gradient. The grey dashed line gives a visual representation of the frequency of the transparency peak.

5.2.1 Slow light

Once we have a good transparency peak, we can start to look for a reduction in the group velocity directly. This is done by measuring the arrival times of photons on the SPCM. While the length of the SPCM pulses is relatively long (10 ns), the timing resolution of the rising edge of this pulse is 350 ps [97]. The output of the SPCM is connected to an oscilloscope[1] which samples the channel at 2 GSa/s, the data from this is then sent to a computer that runs code to find the locations of the peaks. The oscilloscope is triggered off the same sequence pulse as the probe to minimise any timing drift between them. A histogram is formed from the arrival times and, by fitting this histogram to a Gaussian function, the timing of the pulse can be found.

A plot of this can be seen in Fig. 5.6. The grey curve shows photon detection events under conditions without EIT, while the orange curve shows detection events with EIT. The data without EIT was collected with no

[1] Rohde & Schwarz RTC1002 - 350MHz

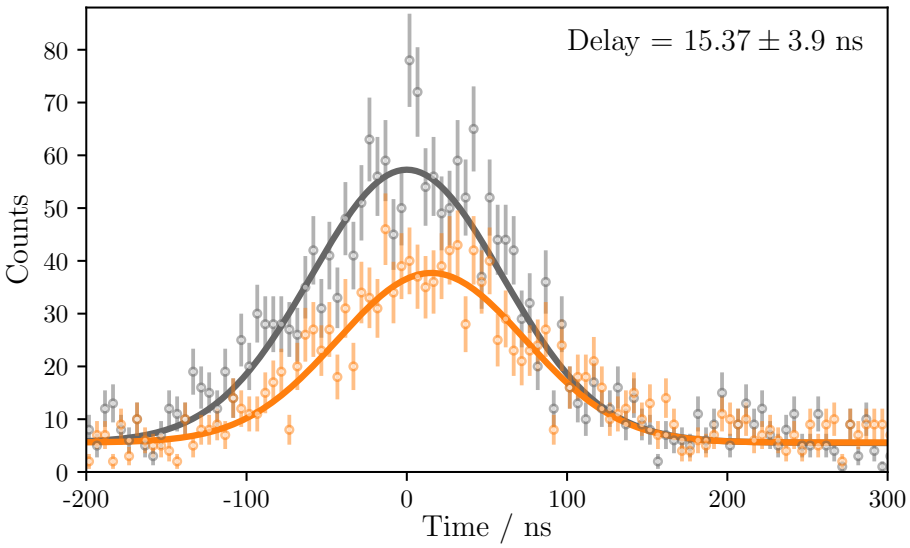

Figure 5.6: Plot of photon detection events against time, with $t = 0$ set to the centre of the grey peak. The 2 curves represent detection events with no EIT (grey), and detection events with EIT (orange). Both data sets are fit to a Voigt profile. The centre of the orange curve occurs (15.4 ± 3.9) ns after the grey, demonstrating a reduction of the photon group velocity. Errorbars represent the square root of the number of counts in that time bin.

atoms present, and calibration was done to ensure this was equal to the arrival time when there are atoms present but the probe is far detuned. The time of the peak centre was found from the Voigt fit with an error of about ± 3 ns on each peak. The time difference between these peaks (with an error found by adding the 2 errors in quadrature) was found to be (15.4 ± 3.9) ns. The fibre junction is 30 µm across, so setting that as our length, this delay corresponds to a group velocity of about 2000 m/s, over 5 orders of magnitude lower than the speed of light in vacuum.

This data was collected with an atom density of $(1.2 \pm 0.1) \times 10^{11}/\text{cm}^3$, a dephasing rate of $2\pi \times (0.6 \pm 0.4)$ MHz and a control Rabi frequency of $2\pi \times (6.5 \pm 0.6)$ MHz. This gives an expected group velocity of 1300 m/s (corresponding to a delay of 23 ns). Considering the error on the dephasing rate, this is in good agreement with the observed delay time. This data was deliberately collected with quite a broad transparency peak, in order

to make initial observation easier.

5.2.2 Next steps

This delay (while impressive) is not enough for us to be able to observe photon storage. For storage, we need to quickly ramp down the power of the control beam while the probe is contained in the atom sample. This process would take a time on the order of $\sim 10\,\mathrm{ns}$, so our total delay would need to be large compared to this time. Work is currently being done on reducing the group velocity to $\sim 300\,\mathrm{m/s}$, which would give us a delay of $\sim 100\,\mathrm{ns}$. With this delay, using a function generator with ns scale timing resolution, the control power would be able to be ramped down and photon storage could be observed, using the same method used to measure the reduction in group velocity. These conditions are definitely possible to reach in this experiment (see the expected group velocity for Fig. 5.4(d)), the main difficulty is in keeping the experiment stable while in this state, as the slow light data can take a few days to collect.

Chapter 6

Bayesian measurement strategies

Rapid and efficient estimation of unknown parameters is essential to quantum technologies (and almost any area of science) [12,18,64–66]. Bayesian estimation theory offers a powerful toolbox for this by using a probabilistic description of the experimental platform [67–70]. Its effectiveness has been demonstrated before in atomic, molecular, and optical experiments, though primarily for phase estimation [71,74,78]. Here, I present a systematic, symmetry-informed adaptive strategy for experimental sensing that enables more rapid information acquisition, or equivalently, higher estimation precision. The core idea is using Bayesian statistics to form an optimal estimator and to optimise the choice of a control parameter to minimise a loss function.

The work presented in this chapter was done in collaboration with Dr. Jesús Rubio of the University of Surrey and Prof. Janet Anders of the University of Exeter. This statistical strategy was developed by Dr. Rubio [76,77], with input on the atomic model from Dr. Nathan Cooper and me. The

experiment and analysis of the results were done by me. The results in this chapter are summarised in the paper "Substantial precision enhancements via adaptive symmetry-informed Bayesian metrology" (2025) [46] (included in App. A). The chapter starts by giving an overview of the basics of Bayesian statistics (Sec. 6.1), before covering the derivation of this adaptive strategy for a general system (sec. 6.1.2). This strategy is then applied to our system (Sec. 6.2) and the experimental results are shown in Sec. 6.3.

6.1 Bayesian methods

Bayesian statistics is built around Bayes' theorem for conditional probabilities [68]

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}\,, \tag{6.1}$$

where E and H are 2 arbitrary events (sometimes called event and hypothesis, respectively). However, rather than just describing probabilities, Bayesian statistics is used to describe a statistical probability density function (PDF).

Suppose ψ is a possible value for an unknown parameter Ψ that one desires to know the value of (the hypothesis). The parameter Ψ cannot be measured directly, but its value can be determined by a measurement with outcome x. The chance of obtaining a certain value of x is determined by a likelihood function $p(x|\psi)$. If the measurement is repeated k times, the data collected is $\boldsymbol{x} = (x_1, x_2, ..., x_k)$, and the total likelihood is

$$p(\boldsymbol{x}|\psi) = \prod_{i=1}^{k} p(x_i|\psi)\,. \tag{6.2}$$

The typical way to find an estimate of Ψ is to find the value of ψ that

maximises this likelihood function. This is called a maximum likelihood estimator (MLE). However, this method is only an efficient estimator in the regime of large sample sizes [121,122].

This is where Bayesian statistical methods become useful. Rather than looking at the PDF of the likelihood to collect data x given hypothesis ψ, we consider the probability of ψ to have a certain value given the data we've collected $p(\psi|x)$. This is called the posterior and can be found using Bayes' theorem [68]

$$p(\psi|\boldsymbol{x}) \propto p(\boldsymbol{x}|\psi)p(\psi) \tag{6.3}$$

where $p(\psi)$ represents the prior knowledge of the value of ψ before the measurement was performed. Here, the denominator of Eq. 6.1 has been excluded, as it is simply a scaling factor. The RHS of the equation is still a PDF and so should integrate to 1.

The posterior represents the knowledge we possess about Ψ after inferring the information from measurements $\boldsymbol{x}$. Therefore, the posterior can naturally be used to construct an estimator $\widetilde{\psi}$. This estimator can take many forms, one possibility would be to find the mean of ψ with respect to the posterior PDF

$$\widetilde{\psi} = \int \mathrm{d}\psi \, p(\psi|\boldsymbol{x})\psi \, . \tag{6.4}$$

However, this is not the only form an estimator can take. A discussion of the choice of what estimator to use is covered in Sec. 6.1.1.

All Bayesian estimation strategies depend explicitly on the form of a prior PDF. The idea of having a prior distribution may seem like it will inherently bias the results inferred. However, this idea of a prior distribution is already naturally used, just without being explicitly stated. When performing a measurement, there is typically an expectation of the order of magnitude

the result will be. For example, when measuring the intensity of a laser beam, the detector must be suitable for the expected intensity (a thermal power sensor measures on a very different order of magnitude to a single photon counter). This prior distribution for the intensity informs how the experiment is constructed. However, it is important to note that a poor choice of prior can dominate over any data collected; if the prior used is a Dirac delta function about a particular value of ψ, the posterior will be the same Dirac delta function. So it is important to choose an appropriate prior that accurately represents the existing knowledge of Ψ before measurement, without dominating over the measurement [67,68,122]. The choice of priors is discussed in more detail in Sec. 6.2.1.

It should be noted that in this first section, I treat all parameters and data as being continuous variables to ensure consistent notation. The logic and arguments presented still hold for discrete variables by simply exchanging any integral over the discrete variable for a sum. This becomes apparent in the next section (Sec. 6.2) where the data from our experiment, photon counts, are discrete.

6.1.1 Optimal estimators

When choosing an estimator, it can be chosen such that it minimises a type of error relevant to the system. An estimator that minimises an error like this is called an optimal estimator [76,77]. These errors are described by a loss function $\mathcal{L}(\widetilde{\psi}, \psi)$ that quantifies the information loss incurred by an estimate $\widetilde{\psi}$ should the hypothesis ψ be correct. The loss associated with measurements $\boldsymbol{x}$ is then

$$\mathcal{L}(\boldsymbol{x}) = \int \mathrm{d}\psi \, p(\psi|\boldsymbol{x}) \mathcal{L}(\widetilde{\psi}(\boldsymbol{x}), \psi) \, . \tag{6.5}$$

The form of $\mathcal{L}(\widetilde{\psi}, \psi)$ depends on the type of parameter being determined by $\widetilde{\psi}$. When determining a phase, trigonometric losses $\mathcal{L}(\widetilde{\psi}, \psi) = 4\sin^2[(\widetilde{\psi} - \psi)/2]$ are the most relevant [123]. Another common loss function, relevant for location parameters, is the mean square error $\mathcal{L}(\widetilde{\psi}, \psi) = (\widetilde{\psi} - \psi)^2$. The estimator that minimises this error fulfils the equation

$$\frac{\partial}{\partial \widetilde{\psi}} \int \mathrm{d}\psi \, p(\psi|\boldsymbol{x})(\widetilde{\psi}(\boldsymbol{x}) - \psi)^2 = 0 \,. \tag{6.6}$$

The estimator that fulfils this criteria, and is hence an optimal estimator, is

$$\widetilde{\psi}^{\mathrm{opt}}(\boldsymbol{x}) = \int \mathrm{d}\psi \, p(\psi|\boldsymbol{x})\psi \,, \tag{6.7}$$

which is the mean of ψ with respect to $p(\psi|\boldsymbol{x})$, as in Eq. 6.4 [122].

The choice of loss function (and hence the optimal estimator) depends on the parameter class. For a location parameter (as above), the loss function is quantified by how "far" the estimator is on a linear scale. For a scale parameter, the loss function is instead quantified by the distance on a logarithmic scale. It is possible to construct a more generalised optimal estimator that applies to classes of parameters that can be mapped into location parameters [76,77]. If Ψ is such a parameter, a function $f(\psi)$ will exist such that

$$f(\psi') = f(\psi) + c \tag{6.8}$$

where ψ and ψ' are 2 different but equally valid hypotheses about Ψ, related by a transformation $\psi' = g(\psi)$. This mapping $f(\psi)$ now behaves like a location parameter, so the loss can be described by a quadratic error

$$\mathcal{L}_f\big(\widetilde{\psi}(\boldsymbol{x}), \psi\big) = \Big[f\big(\widetilde{\psi}(\boldsymbol{x})\big) - f(\psi)\Big]^2 \,. \tag{6.9}$$

The optimal estimator for Ψ is then [46]

$$\widetilde{\psi}_f^{\text{opt}}(\boldsymbol{x}) = f^{-1}\left[\int \mathrm{d}\psi\, p(\psi|\boldsymbol{x}) f(\psi)\right] \tag{6.10}$$

and the uncertainty associated with this estimator is

$$\Delta\widetilde{\psi}_f^{\text{opt}}(\boldsymbol{x}) = \frac{\sqrt{\mathcal{L}_f(\boldsymbol{x})}}{\left|f'\big(\widetilde{\psi}_f^{\text{opt}}(\boldsymbol{x})\big)\right|} \tag{6.11}$$

where

$$\mathcal{L}_f(\boldsymbol{x}) = \int \mathrm{d}\psi\, p(\psi|\boldsymbol{x}) f(\psi)^2 - f\left[\widetilde{\psi}_f^{\text{opt}}(\boldsymbol{x})\right]^2 \tag{6.12}$$

is the loss function and results from substituting Eq. 6.9 into Eq. 6.5.

The function f accounts for the symmetries in the system being modelled. These symmetries lead to the association of different parameter types with certain transformations g under which the state of knowledge is unchanged. Phase parameters have circular symmetry $\psi' = \psi + 2\pi$, whereas scales are invariant under $\psi' = \gamma\psi$. The form of these symmetries becomes apparent when forming a maximum ignorance prior [46,68], which take the form

$$p_{\text{MI}}(\psi) \propto f'(\psi)\,. \tag{6.13}$$

The idea of a maximum ignorance prior is discussed more in Sec. 6.2.1.

6.1.2 Adaptive strategies

Let us build upon the simple model presented at the start of this chapter. Now we desire to estimate the value of a parameter Ψ using a measurement with outcome x where the outcome is also related to a known control parameter y. So the likelihood function is now $p(x|\psi, y)$ (where ψ is a

hypothesis), and after k measurements, the total likelihood is

$$p(\boldsymbol{x}|\psi,\boldsymbol{y}) = \prod_{i=1}^{k} p(x_i|\psi,y_i) \tag{6.14}$$

where $\boldsymbol{y} = (y_1, y_2, ..., y_k)$. All the equations above still generally hold and the optimal estimator is given by

$$\widetilde{\psi}^{\mathrm{opt}}_{\boldsymbol{y},f}(\boldsymbol{x}) = f^{-1}\left[\int \mathrm{d}\psi\, p(\psi|\boldsymbol{x},\boldsymbol{y}) f(\psi)\right] \tag{6.15}$$

and the loss function is

$$\mathcal{L}_{\boldsymbol{y},f}(\boldsymbol{x}) = \int \mathrm{d}\psi\, p(\psi|\boldsymbol{x},\boldsymbol{y}) f(\psi)^2 - f\left[\widetilde{\psi}^{\mathrm{opt}}_{\boldsymbol{y},f}(\boldsymbol{x})\right]^2 . \tag{6.16}$$

We desire to have a way to determine what value of y will yield the most information about Ψ. To do this, we need to find the value of y that will minimise the expected loss (Eq. 6.16) for any new data. The average loss is defined as

$$\bar{\mathcal{L}}_{\boldsymbol{y}} = \int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}|\boldsymbol{y}) \mathcal{L}_{\boldsymbol{y}}(\boldsymbol{x}) \tag{6.17}$$

where $p(\boldsymbol{x}|\boldsymbol{y}) = \int \mathrm{d}\psi\, p(\psi) p(\boldsymbol{x}|\psi,\boldsymbol{y})$. Substituting in Eq. 6.16 we get

$$\bar{\mathcal{L}}_{\boldsymbol{y},f} = \int \mathrm{d}\psi\, p(\psi|\boldsymbol{y}) f(\psi)^2 - \int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}|\boldsymbol{y}) f\left[\widetilde{\psi}^{\mathrm{opt}}_{\boldsymbol{y},f}(\boldsymbol{x})\right]^2 . \tag{6.18}$$

As ψ and y are independent, the value of y does not inform the value of ψ, so $p(\psi|\boldsymbol{y}) = p(\psi)$. The first term, therefore, only depends on the prior information. This means minimising the average loss amounts to maximising a gain function [46]

$$\mathcal{G}_{\boldsymbol{y},f} = \int \mathrm{d}\boldsymbol{x}\, p(\boldsymbol{x}|\boldsymbol{y}) f\left[\widetilde{\psi}^{\mathrm{opt}}_{\boldsymbol{y},f}(\boldsymbol{x})\right]^2 . \tag{6.19}$$

This allows y to be selected such that it maximises this gain, and hence minimises the expected loss and uncertainty.

The control parameter y can be optimised using the gain $\mathcal{G}_{y,f}$ before any measurements have been taken, using only the prior distribution $p(\psi)$, to give an *a priori* optimised value. Alternatively, it can be optimised adaptively, where the distribution is updated after every measurement, and the new gain (and hence optimum y) is found. For this adaptive method, the optimum control for the k-th measurement y_k is found by maximising

$$\mathcal{G}_{y_k,f} = \int \mathrm{d}x\, p(x|y_k) f\left[\widetilde{\psi}^{\mathrm{opt}}_{y_k,f}(x)\right]^2 \tag{6.20}$$

where $p(x|y_k) = \int \mathrm{d}\psi\, p_{k-1}(\psi) p(x|\psi, y_k)$ and $\widetilde{\psi}^{\mathrm{opt}}_{y_k,f}(x)$ is defined as in Eq. 6.15 with $p(\psi|x, y_k) \propto p_{k-1}(\psi) p(x|\psi, y_k)$. Here $p_{k-1}(\psi)$ is the updated prior, given by

$$p_{k-1}(\psi) = \begin{cases} p(\psi) & \text{if } k = 1 \\ p(\psi) \prod_{i=1}^{k-1} p(n_i|\psi, y_i) & \text{if } k > 1 \end{cases}. \tag{6.21}$$

We have thus devised a method to find the optimal value for a control parameter by maximising a gain function $\mathcal{G}_{y,f}$. This gain can be maximised after each individual measurement to adaptively optimise the next measurement, improving the final estimate.

To implement this adaptive measurement protocol, one should apply the following steps:

1. Calculate gain $\mathcal{G}_{y_1,f}$ with $p_0(\psi) = p(\psi)$, the prior distribution. Find the control parameter y_1 which maximises this.

2. Perform a measurement with control y_1 and collect the data x_1.

3. For the second trial, calculate the gain $\mathcal{G}_{y_2,f}$ with the updated prior

$p_1(\psi) \equiv p(\psi|x_1, y_1) \propto p_0(\psi)p(x_1|\psi, y_1)$. Find the new optimal control y_2 that maximises this gain.

4. Perform a measurement with y_2 and collect the data x_2.

5. For the third trial, calculate the gain $\mathcal{G}_{y_3,f}$ with $p_2(\psi) \propto p_1(\psi)p(x_2|\psi, y_2)$. Find y_3 which maximises this. Perform a measurement and collect the data x_3.

6. Repeat this process of finding y_k with $p_{k-1}(\psi) \propto p_{k-2}(\psi)p(x_{k-1}|\psi, y_{k-1})$ until the desired precision has been reach or the desired number of measurements have been performed.

This process will generate the datasets $\boldsymbol{x}$ and $\boldsymbol{y}$, which can be used to estimate Ψ using Eqs. 6.15 & 6.16.

6.2 Our system

This adaptive protocol described above was applied to the fibre experiment used throughout this thesis. When doing experimental optimisation, it is important to be able to quickly and efficiently determine how many atoms are trapped in the interaction region of the fibre junction. However, due to the low powers used and photon shot noise, we typically need to do many repeats, limiting the speed of experimental progress. Additionally, if the atoms are only imaged with resonant light, it can be hard to distinguish between high atom numbers. So the Bayesian scheme above was implemented with the imaging frequency ν as a control parameter, photon counts n are the measurement, and the atom number is the parameter we want to estimate.

In this system, there are 2 unknown parameters that need to be estimated. These are the expected photon counts with no atoms present Φ, and the on-resonance optical depth of the medium Θ, which is related to the number of atoms by $\mathcal{N} = \kappa\Theta$, where κ is a constant determined by the cross section of the transition and the area of the imaging beam. The light passing through the fibre is attenuated by the atoms such that the expected photon counts detected is

$$\bar{n}_\nu(\Phi, \Theta) = \Phi e^{-\zeta_\nu \Theta} \tag{6.22}$$

where

$$\zeta_\nu = \frac{\Gamma^2}{\Gamma^2 + 4(\nu - \nu_r)^2} \tag{6.23}$$

and Γ and ν_r are the linewidth and resonant frequency of the transition, respectively. More details of how atoms are imaged in the fibre are given in Sec. 3.2.2. Strictly, the linewidth and resonant frequency are also unknown parameters, however, they can be determined by a calibration run beforehand and their value does not change significantly between runs.

The detected photon counts follow a Poissonian distribution, so the likelihood function is

$$p(n|\phi, \theta, \nu) = p(n|\bar{n}_\nu(\phi, \theta)) = \frac{\bar{n}_\nu(\phi, \theta)^n e^{-\bar{n}_\nu(\phi,\theta)}}{n!} \,. \tag{6.24}$$

Here $\phi \in [\phi_{min}, \phi_{max}]$ and $\theta \in [\theta_{min}, \theta_{max}]$ are our hypotheses about Φ and Θ respectively. As with our typical imaging scheme (see Sec. 3.2), the data are collected in pairs measuring photon counts with atoms in the junction n_a and photon counts without atoms in the junction n_b. These both follow the Poissonian distribution in Eq. 6.24; however, for n_b, θ is set to 0. So when the prior is updated following Bayes' theorem, it is updated such that

$$p(\phi, \theta|n_a, n_b, \nu) \propto p(\phi, \theta) p(n_a|\phi, \theta, \nu) p(n_b|\phi, 0, \nu) \,. \tag{6.25}$$

6.2.1 Maximum ignorance prior

We now need to find a prior distribution for ϕ and θ. We have above defined some limits where we know the true values Φ and Θ lie within. But beyond that, we want our prior to bias the data as little as possible. For this, we construct what is known as a maximum ignorance prior [68,124].

Suppose 2 observers run the experiment at the same frequency ν. They could detect comparable average photon counts, but have different hypotheses due to different prior information [46,68]. That is $\bar{n}_\nu(\phi', \theta') = \bar{n}_\nu(\phi, \theta)$, so

$$\frac{\phi'}{\phi} = \exp\left[\zeta_\nu(\theta' - \theta)\right] . \tag{6.26}$$

This equation leads to the transformations

$$\phi' = \gamma\phi\,, \quad \zeta_\nu\theta' = \zeta_\nu\theta + \log(\gamma)\,, \tag{6.27}$$

where γ is an arbitrary positive constant, under which Eq. 6.24 is invariant. If both observers are equally ignorant, then their priors must respect this invariance

$$p_{\rm MI}(\phi', \theta')\,\mathrm{d}\phi'\,\mathrm{d}\theta' = p_{\rm MI}(\phi, \theta)\,\mathrm{d}\phi\,\mathrm{d}\theta\ . \tag{6.28}$$

This implies

$$\gamma\, p_{\rm MI}\left(\gamma\phi, \theta + \frac{\log(\gamma)}{\zeta_\nu}\right) = p_{\rm MI}(\phi, \theta)\,, \tag{6.29}$$

which, by taking the derivative with respect to γ on both sides, leads to the partial differential equation

$$\zeta_\nu p_{\rm MI}(a, b) + \zeta_\nu a\frac{\partial}{\partial a}p_{\rm MI}(a, b) + \frac{\partial}{\partial b}p_{\rm MI}(a, b) = 0\,. \tag{6.30}$$

Solving this, we find the prior has the form

$$p_{\mathrm{MI}}(\phi,\theta) = e^{-\zeta_\nu \theta} h\Big(\phi e^{-\zeta_\nu \theta}\Big) \tag{6.31}$$

where h is an arbitrary function. Enforcing this invariance alone does not fully determine the maximum ignorance prior. In order to define the prior, we also enforce the condition that the laser frequency ν does not directly inform the values of Φ and Θ, i.e. $\frac{\mathrm{d}\phi}{\mathrm{d}\nu} = \frac{\mathrm{d}\theta}{\mathrm{d}\nu} = 0$. We can achieve this by setting the derivative of Eq. 6.31 with respect to ν to zero

$$\frac{\mathrm{d}}{\mathrm{d}\nu} p_{\mathrm{MI}}(\phi,\theta) = \frac{\mathrm{d}}{\mathrm{d}\nu}\left[z h(z)\right] = \frac{\mathrm{d}z}{\mathrm{d}\nu}\left[h(z) + z h'(z)\right] = 0\,. \tag{6.32}$$

Provided that $\nu \neq \nu_r$ so $\frac{\mathrm{d}\zeta_\nu}{\mathrm{d}\nu} \neq 0$, the solution is $h(z) \propto 1/z$. Hence, the prior representing maximum ignorance about Φ and Θ is

$$p_{\mathrm{MI}}(\phi,\theta) \propto \frac{1}{\phi}\,. \tag{6.33}$$

Normalising this over our hypothesis range $[\phi_{min}, \phi_{max}] \times [\theta_{min}, \theta_{max}]$ we get [46]

$$p_{\mathrm{MI}}(\phi,\theta) = \left[(\theta_{max} - \theta_{min}) \log\left(\frac{\phi_{max}}{\phi_{min}}\right) \phi\right]^{-1}\,. \tag{6.34}$$

6.2.2 Optimal estimator

As we now know the maximum ignorance prior, we can construct the exact form of our optimal estimator using the relation in Eq. 6.13. In this system, there are 2 unknown parameters, the expected photon counts Φ and the optical depth Θ. However, ultimately we want to determine the expected atom number $\mathcal{N}$, for which we only need to determine Θ. This means Φ is a nuisance parameter, for which we do not need to know the numerical

value.

As we only need to form an estimate of Θ, the symmetry function f we need for our estimate (see Sec. 6.1.1) can be determined by marginalising our prior over ϕ,

$$p_{\mathrm{MI}}(\theta) = \int_{\phi_{min}}^{\phi_{max}} \mathrm{d}\phi \, p_{\mathrm{MI}}(\phi, \theta) \propto 1 \,. \tag{6.35}$$

This leads to the symmetry function

$$f(\theta) = c_1\theta + c_2 \tag{6.36}$$

where $c_{1(2)}$ are arbitrary constants, so we set $f(\theta) = \theta$. This allows us to determine the optimal atom number estimator and its uncertainty using Eqs. 6.15 & 6.11 and $\mathcal{N} = \kappa\Theta$, giving

$$\begin{aligned} \widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n}) &= \kappa \int_{\theta_{min}}^{\theta_{max}} \mathrm{d}\theta \, p(\theta|\boldsymbol{n}, \boldsymbol{\nu})\theta \,, \\ \Delta\widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n})^2 &= \kappa^2 \int_{\theta_{min}}^{\theta_{max}} \mathrm{d}\theta \, p(\theta|\boldsymbol{n}, \boldsymbol{\nu})\theta^2 - \widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n})^2 \,. \end{aligned} \tag{6.37}$$

Here,

$$p(\theta|\boldsymbol{n}, \boldsymbol{\nu}) = \int_{\phi_{min}}^{\phi_{max}} \mathrm{d}\phi \, p(\theta, \phi|\boldsymbol{n}, \boldsymbol{\nu}) \tag{6.38}$$

is the marginalised posterior.

The gain function Eq. 6.20 can also be determined as

$$\mathcal{G}_{\nu_k} = \sum_n p(n|\nu_k)\widetilde{\vartheta}_{\nu_k}(n)^2 \tag{6.39}$$

where

$$p(n|\nu_k) = \iint \mathrm{d}\theta \, \mathrm{d}\phi \, p_{k-1}(\phi, \theta) p(n|\phi, \theta, \nu_k) \tag{6.40}$$

and

$$\widetilde{\vartheta}_{\nu_k}(n) = \iint \mathrm{d}\theta\, \mathrm{d}\phi\, p_{k-1}(\phi,\theta) p(n|\phi,\theta,\nu_k)\theta\,. \tag{6.41}$$

Here

$$p_{k-1}(\phi,\theta) = \begin{cases} p(\phi,\theta) & \text{if } k=1 \\ p(\phi,\theta)\prod_{i=1}^{k-1} p(n_{a,i}|\phi,\theta,\nu_i)p(n_{b,i}|\phi,0,\nu_i) & \text{if } k>1 \end{cases} \tag{6.42}$$

and all the logic set out in Sec. 6.1.2 still holds.

6.2.3 Computational calculation

One challenge in implementing adaptive strategies like this is in efficient computational calculations. These integrals can be quite slow to calculate numerically. If each measurement takes 5 s, but it takes 30 s to calculate the gain and find the optimum frequency between measurements, then the improvement in measurement efficiency can easily be offset by the reduction in measurement speed. We were able to overcome these computational challenges, achieving sufficient speed with some fairly basic optimisation in Python 3. In this experiment, each measurement takes $5-7$s, and we were able to calculate the gain between each measurement in <1s. Most of this computational optimisation work was done by Daniele Baldolini and Dr. Nathan Cooper. It is likely that further optimisation of this is possible, through more sophisticated techniques, and through the use of a faster programming language. A copy of the code used is included in Appendix D.

To speed up the calculation, several standard techniques were used, such as caching of function outputs and control of error tolerances. Additionally, the probability distributions were adaptively truncated as information

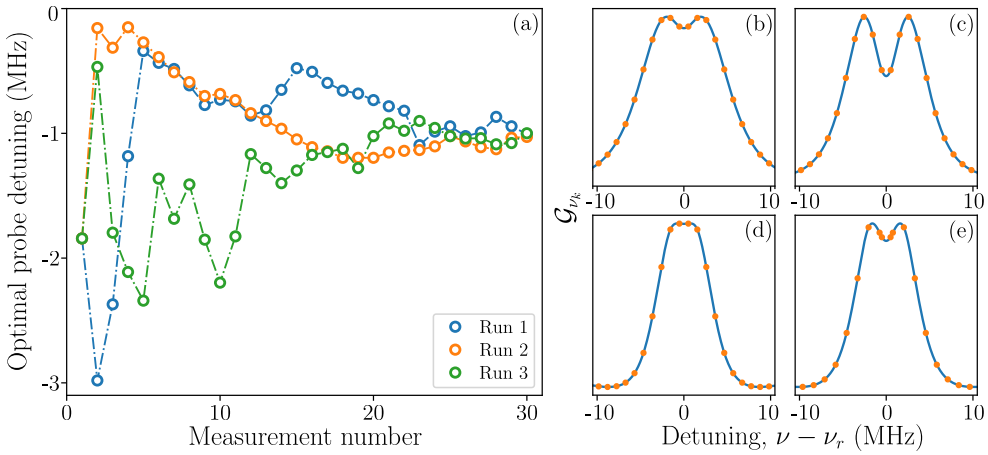

Figure 6.1: (a) Examples of how the optimal frequency changes during the adaptive strategy. Three consecutive experimental runs are shown with an atom number of $N \sim 250$ atoms in the interaction region. (b-e) examples of the gain $\mathcal{G}_{\nu_k}$ as a function of frequency. Each sub-plot shows the gain for a different current atom number estimate $\widetilde{N}$ after a different number of measurements k. Orange dots show explicitly calculated values, while the blue line is cubic interpolation between those points. (b) *a priori*, k=1; (c) $\widetilde{N} \sim 340$, k=4; (d) $\widetilde{N} \sim 200$, k=4; (e) $\widetilde{N} \sim 250$, k=30.

about Φ and Θ was gained. Finally, the gain $\mathcal{G}_{\nu_k}$ (Eq. 6.39) was only calculated explicitly for 10 frequencies, and the optimal frequency was chosen based on cubic interpolation between those points (see side panels in Fig. 6.1). Only frequencies below ν_r were used, as the function is symmetric about resonance.

An example of how the gain and optimum frequency vary during an experimental run can be seen in Fig. 6.1. All measurements start at the same frequency (based on the prior distribution), before initially diverging based on early data. Then (as the expected atom number is the same for each run) they eventually settle to approximately the same final optimum frequency.

The plots on the right show examples of how the gain function varies with frequency. Each sub-plot shows the calculation with a different current atom number estimate $\widetilde{N}$ after a different number of measurements k. It is

clear that the function is symmetric about ν_r. The behaviour demonstrated matches the behaviour we would intuitively expect from the system. When the atom number estimate is low, measuring with resonant light is optimal (see Fig. 6.1(d), where the optimum detuning is 0.75 MHz). As the atom number increases, the optimum frequency gets further from resonance (see Fig. 6.1(c), where the optimum detuning is 2.5 MHz).

6.3 Results

The effectiveness of this adaptive Bayesian strategy was tested by utilising it in the fibre experiment, and comparing its estimates to 4 other estimation strategies. For this, we take measurement runs with $k = 30$ measurements, determining an estimate and its error. These measurement runs are then repeated 12-18 times for each measurement/estimation strategy, allowing us to determine the mean and variance in the estimates.

The first method ("on resonance", blue down triangles) is the standard method as described in Sec. 3.2.2. Resonant light images the atoms and the mean photon counts give an atom number estimate according to [42,82]

$$\widetilde{N}_{\mathrm{MLE}} = -\frac{\kappa}{\zeta_\nu} \log\left[\frac{\langle n_a \rangle}{\langle n_b \rangle}\right] \tag{6.43}$$

where $\langle n_{a(b)} \rangle$ are the mean photon counts with (without) atoms present, ζ_ν is defined in Eq. 6.23, and κ is a constant determined by the interaction cross section and the area of the imaging pulse. The error associated with this estimator results from the standard deviation in n_a and n_b and standard error propagation. This method corresponds to a maximum likelihood estimator (MLE). The second method is analogous, but with the frequency detuned 5 MHz above resonance ("detuned", orange up trian-

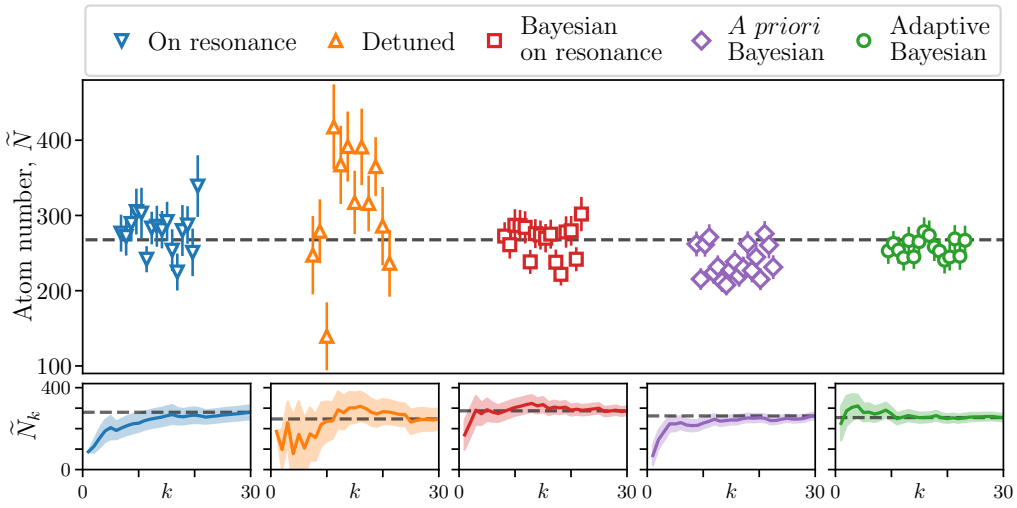

Figure 6.2: Atom number estimates $\widetilde{N}$ and their errors for different measurement strategies taken after 30 measurement shots. The lower sub-plots give a demonstration of how the estimates develop as the number of measurements k increases for each method, with the shaded region representing the error. The dashed lines indicate final estimates in the subplots and the overall average across all methods in the main figure, shown here to guide the eye. The small differences in the mean value obtained with each method are consistent with normal experimental drifts in loading efficiency

gles). This frequency was chosen arbitrarily, based on my own intuition of the experiment, before any Bayesian analysis had been done.

All the other methods are also Bayesian, utilising the optimal estimator in Eq. 6.37. However, each method employs a different strategy for determining the imaging frequency. For these methods, we set the hypothesis ranges to $5 \leq \phi \leq 20$ and $0 \leq \theta \leq 8$, and used the prior distribution from Eq. 6.34. The third method (first Bayesian method) again images the atoms with resonant light and calculates Eq. 6.37 with $\boldsymbol{\nu} = (\nu_r, \nu_r, ..., \nu_r)$ ("Bayesian on resonance", red squares). The data analysed using this method is the exact same data that was collected for the first method. The fourth method is as per the previous one, but with an *a priori* optimisation for the frequency ("*a priori* Bayesian", purple diamonds). The gain $\mathcal{G}_{\nu_0}$(Eq. 6.39) is calculated once at the start, before any data has been collected, and that frequency is used for the full experimental run. With our hypothesis

ranges, the *a priori* optimum frequency was 1.95 MHz below resonance. The fifth and final method is the full adaptive strategy outlined in this chapter ("adaptive Bayesian", green circles).

Figure 6.2 shows the atom number estimate and its error for all 5 measurement strategies. Each data point in the main body of the figure represents the estimate after $k = 30$ measurements, with the error bars representing the uncertainty in that specific estimate. The variance between estimates made using the adaptive Bayesian strategy is clearly smaller than the variance between estimates of any other strategy. We can quantify this using an empirical noise-to-signal ratio (NSR),

$$\mathrm{NSR} = \frac{\mathrm{Var}(\widetilde{N}_k)}{\langle \widetilde{N}_k \rangle^2} = \frac{\langle \widetilde{N}_k^2 \rangle - \langle \widetilde{N}_k \rangle^2}{\langle \widetilde{N}_k \rangle^2} \tag{6.44}$$

where $\langle \widetilde{N}_k \rangle$ is the arithmetic mean of all estimates made with the same strategy. Here, the subscript on $\widetilde{N}_k$ is used to represent the estimate after k measurements, rather than denoting the frequency dependence as it does in earlier equations.

This NSR for all the different strategies after $k = 30$ measurements is shown in Table 6.1. The detuned strategy is clearly the worst, with a ratio of 6%. While going detuned can improve the noise-to-signal ratio, going too far detuned reduces the signal too much leading to poor performance. The detuning I chose for this was clearly too far detuned for the optical depth being measured. The on resonance, Bayesian on resonance, and *a priori* Bayesian all show comparable performance. Though the Bayesian estimator does offer some advantage over the MLE, with "Bayesian on resonance" giving a NSR about 1/3 lower than "on resonance" from the exact same data. However, the adaptive strategy offers the best performance with a NSR of less than 0.2%. This means that this adaptive strategy offers an

	On resonance	Detuned	Bayesian on res.	*A priori* Bayesian	Adaptive Bayesian
NSR in %	0.89	6.10	0.62	0.77	0.19
Num. repeats	16	12	16	19	16
$\langle \widetilde{N}_k \rangle$	279	313	268	238	258
$\langle \Delta \widetilde{N}_k \rangle$	27	47	19	16	18
$k_{\min}$	13	17	12	12	8

Table 6.1: Comparison of atom number estimation strategies in terms of noise-to-signal ratio for $k = 30$ measurement shots, as well as mean minimal number of shots $k_{\min}$ after which the estimate stays within 10% of the final estimate, from the data shown in figure 6.2. Also included is the number of different estimates taken with each method, the empirical mean $\langle \widetilde{N}_k \rangle$ of all estimates for each method, and the mean error bar size $\langle \Delta \widetilde{N}_k \rangle$ of all estimates for each method.

approximately five-fold improvement over the standard method.

Figure 6.3 shows how the NSR varies as a function of the number of measurements included in each estimate k. The adaptive method outperforms the other methods even at low numbers of measurements. The NSR achieved by the standard method after $k = 30$ measurements is reached with only $k = 9$ by the adaptive technique. This means the same precision can be obtained with fewer measurements by using the adaptive technique, allowing experimental progress to be made faster.

The adaptive technique also converges on a final result faster than the other methods. Examples of how the estimates develop as the number of measurements increases are shown in the sub-plots of Fig. 6.2. For each experimental run, we recorded the number of measurements required before the current estimate settled within 10% of the final estimate (based on $k = 30$). We then found the mean of this for each method k_{min}, the results of which can be seen in Tab. 6.1. The adaptive strategy converges on the final estimate with 40% fewer measurements than the standard technique. This behaviour can be seen in the insets of Fig. 6.2, with the

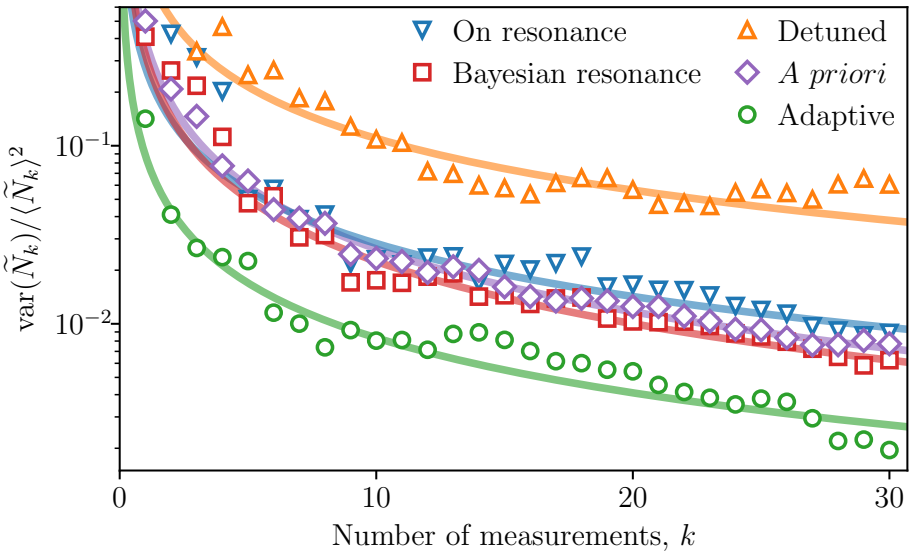

Figure 6.3: Noise-to-signal ratio (NSR) versus number of measurements k. The data points are obtained empirically using Eq. 6.44 by taking the variance between estimates after calculating the estimate after k measurements. Fitted curves act as guides to the eye.

adaptive reaching the final estimate after $\sim$ 10 measurements, while the other methods need 15-20. It should be noted that these insets only show a single run for each strategy, and so this behaviour could be due to chance. However, the values listed in Tab. 6.1 are the average across all estimates and are more reliable.

Chapter 7

Conclusion

This thesis demonstrates an innovative approach to integrating cold atoms with optical waveguides, offering an alternative to nanofibre and hollow core fibre systems. We have successfully exploited the system to demonstrate EIT and slow light, and are currently working to observe photon storage. We have also shown the utility of Bayesian estimation theory to quantum technologies, with adaptive measurement strategies offering significant improvements to the precision of our estimate.

Laser-drilled through holes in an optical waveguide offer a powerful method for interfacing cold atoms with optical waveguides. This method allows for high atom number densities in the interaction region, leading to strong atom-photon interactions. The improvements I made to the experiment by optimising the beam profile, introducing optical pumping, and reducing temperatures to $< 100\,\mu\text{K}$ allowed us to observe a transmittance of less than 1%. This implies an optical depth per unit length of 1700/cm, significantly higher than previously achieved in this system [44].

I was also able to directly observe a reduction in the photon group velocity

through the atoms under EIT conditions. We were able to measure a delay time of (15.4 ± 3.9) ns, corresponding to a reduction of the group velocity to 2000 m/s. Though further improvements are still necessary to conduct a photon storage experiment. A group velocity of 300 m/s would give us a delay of 100 ns, which would be enough time to implement a storage protocol. Since that slow light data was taken, improvements have been made to the experiment with the implementation of optical pumping, reducing the effects of collisional losses, allowing us to achieve higher atom number densities, leading to further reductions in the photon group velocity.

Finally, we demonstrated the application of a novel Bayesian measurement strategy to our experiment, enabling more efficient estimates. This is the first time a fully adaptive technique has been experimentally applied to the estimation of non-phase parameters. The strategy not only optimises the estimator, but also optimises a control parameter allowing for adaptive optimisation of how the measurement is being performed. Through this strategy, we achieved a 5-fold improvement in measurement precision when the same number of measurements had been performed. Alternatively, we could reach the same specified precision with only 1/3 the number of measurements. This strategy can be applied to a wide range of quantum technology experiments, from magnetometry and interferometry to read out in quantum computing.

7.1 Outlook and next steps

The immediate next steps for this fibre experiment are to continue the work on EIT, trying to achieve further reduction in the group velocity in order to be able to observe photon storage. The improvements that have been

made with the optical pumping and phase lock will allow us the achieve the factor of 5 improvement in the group velocity required to be able to implement the storage scheme.

Another improvement that needs to be made to the system is changing the fibre we are using in the experiment. The current fibre has been damaged by an accidental exposure to full power dipole light for an extended period. This exposure caused significant heating in the fibre, creating a small "burn" which absorbs light from the dipole beam more efficiently, heating the fibre more easily. While the fibre is still useable, this creates issues for the lifetime of atoms trapped in the fibre junction. Additionally, this heating can deform the fibre junction, changing the overall transmission through the fibre. It is therefore necessary to open up the vacuum chamber and replace the fibre inside. We have a number of spare fibres that are suitable, with flat edges to the interface and transmissions of about 20%. This transmission is lower than the current fibre used to be, but allows us to make the change immediately. Alternatively, we could get some new fibre junctions produced, which may have even greater transmissions than the current fibre, and may allow us to conduct some new experiments (such as a photonic circuit rather than fibre, or multiple junctions).

Outlook

There are several directions this experiment could be taken next. One possibility is to make further improvements to the lifetime of the trap (to $\sim 1\,\mathrm{s}$) to be able to implement an evaporative cooling stage and potentially form a Bose-Einstein condensate (BEC) in the fibre junction. One of the advantages of our system over similar nanofibre systems is our ability to have greater separation between the trapped atoms and the fibre walls,

which may allow us to get to colder temperatures than they have been able to achieve so far [104,125]. Forming a BEC in the fibre will allow us to get to even higher atom number densities ($\sim 10^{13}\,\mathrm{cm}^{-3}$) and hence stronger atom-photon interactions. However, care would need to be given to ensure the diameter of the BEC is larger than the mode field diameter of the fibre.

Other directions the experiment could be taken include: closing the fibre on itself to form a cavity; 4 wave mixing experiments; Rydberg excitations and how they interact with the fibre walls. The Bayesian work we did could also be taken further, by either directly applying the method to a different control parameter (such as imaging power or exposure time), or designing a sensing experiment that could be done in the fibre and use Bayesian statistics to fully optimise our measurement process.

Bibliography

[1] Y. Bidel, O. Carraz, R. Charrière, M. Cadoret, N. Zahzam, and A. Bresson, "Compact cold atom gravimeter for field applications," *Applied Physics Letters*, vol. 102, no. 14, 2013.

[2] J. E. Debs, P. A. Altin, T. H. Barter, D. Doering, G. R. Dennis, G. McDonald, R. P. Anderson, J. D. Close, and N. P. Robins, "Cold-atom gravimetry with a bose-einstein condensate," *Physical Review A*, vol. 84, no. 3, p. 033610, 2011.

[3] I. Dutta, D. Savoie, B. Fang, B. Venon, C. L. Garrido Alzar, R. Geiger, and A. Landragin, "Continuous cold-atom inertial sensor with 1 nrad/sec rotation stability," *Physical Review Letters*, vol. 116, no. 18, p. 183003, 2016.

[4] B. Canuel, F. Leduc, D. Holleville, A. Gauguet, J. Fils, A. Virdis, A. Clairon, N. Dimarcq, C. J. Bordé, A. Landragin, *et al.*, "Six-axis inertial sensor using cold-atom interferometry," *Physical Review Letters*, vol. 97, no. 1, p. 010402, 2006.

[5] N. Behbood, F. Martin Ciurana, G. Colangelo, M. Napolitano, M. W. Mitchell, and R. J. Sewell, "Real-time vector field tracking with a cold-atom magnetometer," *Applied Physics Letters*, vol. 102, no. 17, 2013.

[6] T. M. Tierney, N. Holmes, S. Mellor, J. D. López, G. Roberts, R. M. Hill, E. Boto, J. Leggett, V. Shah, M. J. Brookes, *et al.*, "Optically pumped magnetometers: From quantum origins to multi-channel magnetoencephalography," *NeuroImage*, vol. 199, pp. 598–608, 2019.

[7] S.-H. Wei, B. Jing, X.-Y. Zhang, J.-Y. Liao, C.-Z. Yuan, B.-Y. Fan, C. Lyu, D.-L. Zhou, Y. Wang, G.-W. Deng, *et al.*, "Towards real-world quantum networks: a review," *Laser & Photonics Reviews*, vol. 16, no. 3, p. 2100219, 2022.

[8] J. P. Covey, H. Weinfurter, and H. Bernien, "Quantum networks with neutral atom processing nodes," *npj Quantum Information*, vol. 9, no. 1, p. 90, 2023.

[9] L.-M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics," *Nature*, vol. 414, no. 6862, pp. 413–418, 2001.

[10] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, *et al.*, "Multi-qubit entanglement and algorithms on a neutral-atom quantum computer," *Nature*, vol. 604, no. 7906, pp. 457–462, 2022.

[11] A. Negretti, P. Treutlein, and T. Calarco, "Quantum computing implementations with neutral particles," *Quantum Information Processing*, vol. 10, pp. 721–753, 2011.

[12] K. Wintersperger, F. Dommert, T. Ehmer, A. Hoursanov, J. Klepsch, W. Mauerer, G. Reuber, T. Strohm, M. Yin, and S. Luber, "Neutral atom quantum computing hardware: performance and end-user perspective," *EPJ Quantum Technology*, vol. 10, no. 1, p. 32, 2023.

[13] H. J. Kimble, "The quantum internet," *Nature*, vol. 453, no. 7198, pp. 1023–1030, 2008.

[14] R. Folman, P. Krüger, D. Cassettari, B. Hessmo, T. Maier, and J. Schmiedmayer, "Controlling cold atoms using nanofabricated surfaces: Atom chips," *Physical Review Letters*, vol. 84, no. 20, p. 4749, 2000.

[15] P. Navez, S. Pandey, H. Mas, K. Poulios, T. Fernholz, and W. Von Klitzing, "Matter-wave interferometers using TAAP rings," *New Journal of Physics*, vol. 18, no. 7, p. 075014, 2016.

[16] C. C. Nshii, M. Vangeleyn, J. P. Cotter, P. F. Griffin, E. A. Hinds, C. N. Ironside, P. See, A. G. Sinclair, E. Riis, and A. S. Arnold, "A surface-patterned chip as a strong source of ultracold atoms for quantum technologies," *Nature Nanotechnology*, vol. 8, no. 5, pp. 321–324, 2013.

[17] S. H. Madkhaly, L. A. Coles, C. Morley, C. D. Colquhoun, T. M. Fromhold, N. Cooper, and L. Hackermüller, "Performance-optimized components for quantum technologies via additive manufacturing," *PRX Quantum*, vol. 2, no. 3, p. 030326, 2021.

[18] J. Lee, R. Ding, J. Christensen, R. R. Rosenthal, A. M. Ison, D. Gillund, D. Bossert, K. Fuerschbach, W. Kindel, and P. S. Finnegan, "A compact cold-atom interferometer with a high data-rate grating magneto-optical trap and a photonic-integrated-circuit-compatible laser system," *Nature Communications*, vol. 13, no. 1, 2022.

[19] A. Isichenko, N. Chauhan, D. Bose, J. Wang, P. D. Kunz, and D. J. Blumenthal, "Photonic integrated beam delivery for a rubidium

3D magneto-optical trap," *Nature Communications*, vol. 14, no. 1, p. 3080, 2023.

[20] D. J. Blumenthal, A. Isichenko, and N. Chauhan, "Enabling photonic integrated 3D magneto-optical traps for quantum sciences and applications," *Optica Quantum*, vol. 2, no. 6, pp. 444–457, 2024.

[21] J. Lee, G. Biedermann, J. Mudrick, E. A. Douglas, and Y.-Y. Jau, "Demonstration of a MOT in a sub-millimeter membrane hole," *Scientific Reports*, vol. 11, no. 1, p. 8807, 2021.

[22] A. P. Hilton, C. Perrella, F. Benabid, B. M. Sparkes, A. N. Luiten, and P. S. Light, "High-efficiency cold-atom transport into a waveguide trap," *Physical Review Applied*, vol. 10, no. 4, p. 044034, 2018.

[23] G. Lepert, M. Trupke, M. J. Hartmann, M. B. Plenio, and E. A. Hinds, "Arrays of waveguide-coupled optical cavities that interact strongly with atoms," *New Journal of Physics*, vol. 13, no. 11, p. 113002, 2011.

[24] J. Lee, D. H. Park, S. Mittal, M. Dagenais, and S. L. Rolston, "Integrated optical dipole trap for cold neutral atoms with an optical waveguide coupler," *New Journal of Physics*, vol. 15, no. 4, p. 043010, 2013.

[25] R. Ritter, N. Gruhler, H. Dobbertin, H. Kübler, S. Scheel, W. Pernice, T. Pfau, and R. Löw, "Coupling thermal atomic vapor to slot waveguides," *Physical Review X*, vol. 8, no. 2, p. 021032, 2018.

[26] W. Li, D. Brown, A. Vylegzhanin, Z. Shahrabifarahani, A. Raj, J. Du, and S. N. Chormaic, "Atom-light interactions using optical nanofibres - a perspective," *Journal of Physics Photonics*, vol. 6, no. 2, p. 021002–021002, 2024.

[27] G. Sagué, E. Vetsch, W. Alt, D. Meschede, and A. Rauschenbeutel, "Cold-atom physics using ultrathin optical fibers: Light-induced dipole forces and surface interactions," *Physical Review Letters*, vol. 99, no. 16, p. 163602, 2007.

[28] T. Nieddu, V. Gokhroo, and S. N. Chormaic, "Optical nanofibres and neutral atoms," *Journal of Optics*, vol. 18, no. 5, p. 053001, 2016.

[29] F. Le Kien, P. Schneeweiss, and A. Rauschenbeutel, "State-dependent potentials in a nanofiber-based two-color trap for cold atoms," *Physical Review A*, vol. 88, no. 3, p. 033840, 2013.

[30] E. Vetsch, S. T. Dawkins, R. Mitsch, D. Reitz, P. Schneeweiss, and A. Rauschenbeutel, "Nanofiber-based optical trapping of cold neutral atoms," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 18, no. 6, pp. 1763–1770, 2012.

[31] A. Nicolas, *Optical quantum memories with cold atomic ensembles: a free space implementation for multimode storage, or a nanofiber-based one for high collection efficiency.* PhD thesis, Université Pierre et Marie Curie-Paris VI, 2014.

[32] B. Gouraud, *Optical Nanofibers Interfacing Cold Atoms-A Tool for Quantum Optics.* PhD thesis, Université Pierre et Marie Curie-Paris VI, 2016.

[33] D. R. Häupl, D. Weller, R. Löw, and N. Y. Joly, "Spatially resolved spectroscopy of alkali metal vapour diffusing inside hollow-core photonic crystal fibres," *New Journal of Physics*, vol. 24, no. 11, p. 113017–113017, 2022.

[34] M. Bajcsy, S. Hofferberth, T. Peyronel, V. Balić, Q.-Y. Liang, A. S. Zibrov, V. Vuletić, and M. D. Lukin, "Laser-cooled atoms inside a

hollow-core photonic-crystal fiber," *Physical Review A*, vol. 83, no. 6, 2011.

[35] K. T. Kaczmarek, D. J. Saunders, M. R. Sprague, W. S. Kolthammer, A. Feizpour, P. M. Ledingham, B. Brecht, E. Poem, I. A. Walmsley, and J. Nunn, "Ultrahigh and persistent optical depths of cesium in kagomé-type hollow-core photonic crystal fibers," *Optics Letters*, vol. 40, no. 23, pp. 5582–5585, 2015.

[36] G. Epple, K. S. Kleinbach, T. G. Euser, N. Y. Joly, T. Pfau, P. S. J. Russell, and R. Löw, "Rydberg atoms in hollow-core photonic crystal fibres," *Nature Communications*, vol. 5, no. 1, 2014.

[37] M. Langbecker, M. Noaman, N. Kjærgaard, F. Benabid, and P. Windpassinger, "Rydberg excitation of cold atoms inside a hollow-core fiber," *Physical Review A*, vol. 96, no. 4, 2017.

[38] M. E. Kim, T.-H. Chang, B. M. Fields, C.-A. Chen, and C.-L. Hung, "Trapping single atoms on a nanophotonic circuit with configurable tweezer lattices," *Nature Communications*, vol. 10, no. 1, 2019.

[39] X. Zhou, H. Tamura, T.-H. Chang, and C.-L. Hung, "Trapped atoms and superradiance on an integrated nanophotonic microring circuit," *Physical Review X*, vol. 14, no. 3, 2024.

[40] A. Goban, C.-L. Hung, S. Yu, J. D. Hood, J. A. Muniz, J. H. Lee, M. J. Martin, A. McClung, K. S. Choi, D. E. Chang, O. Painter, and H. J. Kimble, "Atom–light interactions in photonic crystals," *Nature Communications*, vol. 5, no. 1, 2014.

[41] M. Kohnen, M. Succo, P. G. Petrov, R. A. Nyman, M. Trupke, and E. A. Hinds, "An array of integrated atom–photon junctions," *Nature Photonics*, vol. 5, no. 1, p. 35–38, 2010.

[42] E. Da Ros, *Integrating cold caesium atoms into optical waveguides*. PhD thesis, University of Nottingham, 2020.

[43] D. Johnson, *Novel exploitation of cold atom experimental techniques for the development of Quantum Technologies*. PhD thesis, University of Nottingham, 2025.

[44] E. Da Ros, N. Cooper, J. Nute, and L. Hackermüller, "Cold atoms in micromachined waveguides: A new platform for atom-photon interactions," *Physical Review Research*, vol. 2, no. 3, p. 033098, 2020.

[45] N. Cooper, E. Da Ros, C. Briddon, V. Naniyil, M. T. Greenaway, and L. Hackermüller, "Prospects for strongly coupled atom-photon quantum nodes," *Scientific Reports*, vol. 9, no. 1, pp. 1–7, 2019.

[46] M. Overton, J. Rubio, N. Cooper, D. Baldolini, D. Johnson, J. Anders, and L. Hackermüller, "Substantial precision enhancements via adaptive symmetry-informed Bayesian metrology," *arXiv preprint arXiv:2410.10615*, 2024.

[47] A. I. Lvovsky, B. C. Sanders, and W. Tittel, "Optical quantum memory," *Nature Photonics*, vol. 3, no. 12, p. 706–714, 2009.

[48] K. Heshami, D. England, P. C. Humphreys, P. J. Bustard, V. M. Acosta, J. Nunn, and B. J. Sussman, "Quantum memories: emerging applications and recent advances," *Journal of Modern Optics*, vol. 63, no. 20, p. 2005–2028, 2016.

[49] B. Jing, S. Wei, L. J. Zhang, D. Zhou, Y. He, X. Zou, W. Pan, H. Song, and L. Yan, "Approaching scalable quantum memory with integrated atomic devices," *Applied Physics Reviews*, vol. 11, no. 3, 2024.

[50] X.-X. Li, P. Zhou, Y.-H. Chen, and X. Zhang, "Integrated optical quantum memory controlled by the electro-optic effect," *Physical Review A*, vol. 108, no. 1, p. 012614, 2023.

[51] K. T. Kaczmarek, P. M. Ledingham, B. Brecht, S. E. Thomas, G. S. Thekkadath, O. Lazo-Arjona, J. H. D. Munns, E. Poem, A. Feizpour, D. J. Saunders, J. Nunn, and I. A. Walmsley, "High-speed noise-free optical quantum memory," *Physical Review A*, vol. 97, no. 4, 2018.

[52] K. F. Reim, J. Nunn, V. O. Lorenz, B. J. Sussman, K. C. Lee, N. K. Langford, D. Jaksch, and I. A. Walmsley, "Towards high-speed optical quantum memories," *Nature Photonics*, vol. 4, no. 4, pp. 218–221, 2010.

[53] J.-P. Dou, A.-L. Yang, M.-Y. Du, D. Lao, J. Gao, L.-F. Qiao, H. Li, X.-L. Pang, Z. Feng, H. Tang, *et al.*, "A broadband DLCZ quantum memory in room-temperature atoms," *Communications Physics*, vol. 1, no. 1, p. 55, 2018.

[54] P. Vernaz-Gris, K. Huang, M. Cao, A. S. Sheremet, and J. Laurat, "Highly-efficient quantum memory for polarization qubits in a spatially-multiplexed cold atomic ensemble," *Nature Communications*, vol. 9, no. 1, 2018.

[55] X. Bao, A. Reingruber, P. Dietrich, J. Rui, A. Dück, T. Straßel, L. Li, N.-L. Liu, B. Zhao, and J.-W. Pan, "Efficient and long-lived quantum memory with cold atoms inside a ring cavity," *Nature Physics*, vol. 8, no. 7, p. 517–521, 2012.

[56] Y. Cho, G. Campbell, J. L. Everett, J. Bernu, D. Higginbottom, M. Cao, J. Geng, N. P. Robins, P. K. Lam, and B. C. Buchler, "Highly efficient optical quantum memory with long coherence time in cold atoms," *Optica*, vol. 3, no. 1, p. 100–100, 2016.

[57] Q. Zhang and G. Huang, "Quantum memory of single-photon polarization qubits via double electromagnetically induced transparency," *Physical Review A*, vol. 104, no. 3, 2021.

[58] L. Esguerra, L. Meßner, E. Robertson, N. V. Ewald, M. Gündoğan, and J. Wolters, "Optimization and readout-noise analysis of a warm-vapor electromagnetically-induced-transparency memory on the Cs D_1 line," *Physical Review A*, vol. 107, no. 4, 2023.

[59] B. Gouraud, D. Maxein, A. Nicolas, O. Morin, and J. Laurat, "Demonstration of a memory for tightly guided light in an optical nanofiber," *Physical Review Letters*, vol. 114, no. 18, p. 180503, 2015.

[60] C. Sayrin, C. Clausen, B. Albrecht, P. Schneeweiß, and A. Rauschenbeutel, "Storage of fiber-guided light in a nanofiber-trapped ensemble of cold atoms," *Optica*, vol. 2, no. 4, p. 353–353, 2015.

[61] W. S. Leong, M. Xin, C. Huang, Z. Chen, and S.-Y. Lan, "Long light storage time in an optical fiber," *Physical Review Research*, vol. 2, no. 3, p. 033320, 2020.

[62] F. Blatt, L. S. Simeonov, T. Halfmann, and T. Peters, "Stationary light pulses and narrowband light storage in a laser-cooled ensemble loaded into a hollow-core fiber," *Physical Review A*, vol. 94, no. 4, p. 043833, 2016.

[63] T. Peters, T.-P. Wang, A. Neumann, L. S. Simeonov, and T. Halfmann, "Single-photon-level narrowband memory in a hollow-core photonic bandgap fiber," *Optics Express*, vol. 28, no. 4, pp. 5340–5354, 2020.

[64] Y. L. Len, T. Gefen, A. Retzker, and J. Kołodyński, "Quantum metrology with imperfect measurements," *Nature Communications*, vol. 13, no. 1, p. 6971, 2022.

[65] P. Horodecki, Ł. Rudnicki, and K. Życzkowski, "Five open problems in quantum information theory," *PRX Quantum*, vol. 3, no. 1, p. 010101, 2022.

[66] X. Deng, S. Li, Z.-J. Chen, Z. Ni, Y. Cai, J. Mai, L. Zhang, P. Zheng, H. Yu, C.-L. Zou, *et al.*, "Quantum-enhanced metrology with large Fock states," *Nature Physics*, pp. 1–7, 2024.

[67] P. M. Lee, *Bayesian statistics : an introduction*. Chichester, West Sussex: Wiley, 4th ed., 2012.

[68] E. T. Jaynes and G. L. Bretthorst, *Probability Theory : The Logic of Science*. Cambridge: Cambridge University Press, 2003.

[69] R. Demkowicz-Dobrzański, W. Górecki, and M. Guţă, "Multi-parameter estimation beyond quantum fisher information," *Journal of Physics A: Mathematical and Theoretical*, vol. 53, no. 36, p. 363001, 2020.

[70] J. Rubio and J. Dunningham, "Quantum metrology in the presence of limited data," *New Journal of Physics*, vol. 21, no. 4, p. 043037, 2019.

[71] S. Paesani, A. A. Gentile, R. Santagati, J. Wang, N. Wiebe, D. P. Tew, J. L. O'Brien, and M. G. Thompson, "Experimental Bayesian quantum phase estimation on a silicon photonic chip," *Physical Review Letters*, vol. 118, no. 10, p. 100503, 2017.

[72] A. Lumino, E. Polino, A. S. Rab, G. Milani, N. Spagnolo, N. Wiebe, and F. Sciarrino, "Experimental phase estimation enhanced by ma-

chine learning," *Physical Review Applied*, vol. 10, no. 4, p. 044033, 2018.

[73] L. Gerster, F. Martínez-García, P. Hrmo, M. W. van Mourik, B. Wilhelm, D. Vodola, M. Müller, R. Blatt, P. Schindler, and T. Monz, "Experimental Bayesian calibration of trapped-ion entangling operations," *PRX Quantum*, vol. 3, no. 2, p. 020350, 2022.

[74] M. Valeri, V. Cimini, S. Piacentini, F. Ceccarelli, E. Polino, F. Hoch, G. Bizzarri, G. Corrielli, N. Spagnolo, R. Osellame, *et al.*, "Experimental multiparameter quantum metrology in adaptive regime," *Physical Review Research*, vol. 5, no. 1, p. 013138, 2023.

[75] J. Rubio, J. Anders, and L. A. Correa, "Global quantum thermometry," *Physical Review Letters*, vol. 127, no. 19, p. 190402, 2021.

[76] J. Rubio, "Quantum scale estimation," *Quantum Science and Technology*, vol. 8, no. 1, p. 015009, 2022.

[77] J. Rubio, "First-principles construction of symmetry-informed quantum metrologies," *Physical Review A*, vol. 110, no. 3, p. L030401, 2024.

[78] J. Glatthard, J. Rubio, R. Sawant, T. Hewitt, G. Barontini, and L. A. Correa, "Optimal cold atom thermometry using adaptive Bayesian strategies," *PRX Quantum*, vol. 3, no. 4, p. 040330, 2022.

[79] S. Weyers, U. Hübner, R. Schröder, C. Tamm, and A. Bauch, "Uncertainty evaluation of the atomic caesium fountain CSF1 of the PTB," *Metrologia*, vol. 38, no. 4, p. 343, 2001.

[80] B. Fang, I. Dutta, P. Gillot, D. Savoie, J. Lautier, B. Cheng, C. L. Garrido Alzar, R. Geiger, S. Merlet, F. Pereira Dos Santos, *et al.*,

"Metrology with atom interferometry: inertial sensors from laboratory to field applications," in *Journal of Physics: Conference Series*, vol. 723, p. 012049, IOP Publishing, 2016.

[81] P. Van der Straten and H. Metcalf, *Atoms and molecules interacting with light: Atomic physics for the laser era*. Cambridge University Press, 2016.

[82] D. A. Steck, "Quantum and atom optics, revision 0.16.2." http://steck.us/alkalidata, November 2024. Accessed: 2025-01-10.

[83] G. K. Woodgate, *Elementary atomic structure*. Oxford science publications, Oxford: Clarendon Press, 2nd ed., 1980.

[84] D. A. Steck, "Cesium D line data, revision 2.3.3." http://steck.us/alkalidata, May 2024. Accessed: 2025-01-10.

[85] H. Metcalf and P. van der Straten, *Laser cooling and trapping*. Graduate texts in contemporary physics, New York: Springer, 1999.

[86] P. D. Lett, W. D. Phillips, S. L. Rolston, C. E. Tanner, R. N. Watts, and C. I. Westbrook, "Optical molasses," *JOSA B*, vol. 6, no. 11, pp. 2084–2107, 1989.

[87] P. D. Lett, R. N. Watts, C. I. Westbrook, W. D. Phillips, P. L. Gould, and H. J. Metcalf, "Observation of atoms laser cooled below the Doppler limit," *Physical Review Letters*, vol. 61, no. 2, p. 169, 1988.

[88] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, "Optical dipole traps for neutral atoms," in *Advances in Atomic, Molecular, and Optical Physics*, vol. 42, pp. 95–170, Elsevier, 2000.

[89] D. W. Preston, "Doppler-free saturated absorption: Laser spectroscopy," *American Journal of Physics*, vol. 64, no. 11, pp. 1432–1436, 1996.

[90] U. Schünemann, H. Engler, R. Grimm, M. Weidemüller, and M. Zielonkowski, "Simple scheme for tunable frequency offset locking of two lasers," *Review of Scientific Instruments*, vol. 70, no. 1, pp. 242–243, 1999.

[91] D. Höckel, M. Scholz, and O. Benson, "A robust phase-locked diode laser system for EIT experiments in cesium," *Applied Physics B*, vol. 94, pp. 429–435, 2009.

[92] M. Müller, "Realisation and characterization of a phase locked laser system for coherent spectroscopy of fiber-coupled cesium atoms," Master's thesis, Johannes Gutenberg-Universität Mainz - Institut für Physik, July 2010.

[93] R. Wynands and A. Nagel, "Precision spectroscopy with coherent dark states.," *Applied Physics B: Lasers & Optics*, vol. 68, no. 1, 1999.

[94] J. J. Olivero and R. Longbothum, "Empirical fits to the Voigt line width: A brief review," *Journal of Quantitative Spectroscopy and Radiative Transfer*, vol. 17, no. 2, pp. 233–236, 1977.

[95] Z. Xu, X. Zhang, K. Huang, and X. Lu, "A digital optical phase-locked loop for diode lasers based on field programmable gate array," *Review of Scientific Instruments*, vol. 83, no. 9, 2012.

[96] A. M. Marino and C. R. Stroud, "Phase-locked laser system for use in atomic coherence experiments," *Review of Scientific Instruments*, vol. 79, no. 1, 2008.

[97] Excelitas, *SPCM-AQRH series, datasheet operating manual*, 2023.

[98] A. Lambrecht, E. Giacobino, and S. Reynaud, "Atomic number fluctuations in a falling cold atom cloud," *Quantum and Semiclassical Optics: Journal of the European Optical Society Part B*, vol. 8, no. 3, p. 457, 1996.

[99] S. Prahl, "laserbeamsize: a python module for ISO 11146 analysis of laser beams," Nov. 2023.

[100] S. T. Dawkins, R. Mitsch, D. Reitz, E. Vetsch, and A. Rauschenbeutel, "Dispersive optical interface based on nanofiber-trapped atoms," *Physical Review Letters*, vol. 107, no. 24, p. 243601, 2011.

[101] J. Weiner, V. S. Bagnato, S. Zilio, and P. S. Julienne, "Experiments and theory in cold and ultracold collisions," *Reviews of Modern Physics*, vol. 71, no. 1, p. 1, 1999.

[102] D. Hümmer, P. Schneeweiss, A. Rauschenbeutel, and O. Romero-Isart, "Heating in nanophotonic traps for cold atoms," *Physical Review X*, vol. 9, no. 4, p. 041034, 2019.

[103] K. P. Nayak, M. Sadgrove, R. Yalla, F. Le Kien, and K. Hakuta, "Nanofiber quantum photonics," *Journal of Optics*, vol. 20, no. 7, p. 073001, 2018.

[104] Y. Meng, A. Dareau, P. Schneeweiss, and A. Rauschenbeutel, "Near-ground-state cooling of atoms optically trapped 300 nm away from a hot surface," *Physical Review X*, vol. 8, no. 3, p. 031054, 2018.

[105] R. Kumar, V. Gokhroo, V. B. Tiwari, and S. N. Chormaic, "Temperature measurement of cold atoms using transient absorption of a resonant probe through an optical nanofibre," *Journal of Optics*, vol. 18, no. 11, p. 115401, 2016.

[106] R. R. Silva, K. M. F. Magalhães, E. A. d. L. Henn, L. G. Marcassa, and V. S. Bagnato, "Temperature determination for magneto optical trapped atoms using a single parameter transient absorption," *Optics Communications*, vol. 265, no. 2, pp. 526–531, 2006.

[107] M. Fleischhauer, A. Imamoglu, and J. P. Marangos, "Electromagnetically induced transparency: Optics in coherent media," *Reviews of Modern Physics*, vol. 77, no. 2, pp. 633–673, 2005.

[108] S. E. Harris, J. E. Field, and A. Imamoğlu, "Nonlinear optical processes using electromagnetically induced transparency," *Physical Review Letters*, vol. 64, no. 10, p. 1107, 1990.

[109] M. Albert, A. Dantan, and M. Drewsen, "Cavity electromagnetically induced transparency and all-optical switching using ion coulomb crystals," *Nature Photonics*, vol. 5, no. 10, pp. 633–636, 2011.

[110] F. Le Kien and A. Rauschenbeutel, "Nanofiber-based all-optical switches," *Physical Review A*, vol. 93, no. 1, p. 013849, 2016.

[111] L. V. Hau, S. E. Harris, Z. Dutton, and C. H. Behroozi, "Light speed reduction to 17 metres per second in an ultracold atomic gas," *Nature*, vol. 397, no. 6720, pp. 594–598, 1999.

[112] M. Fleischhauer and M. D. Lukin, "Dark-state polaritons in electromagnetically induced transparency," *Physical Review Letters*, vol. 84, no. 22, p. 5094, 2000.

[113] O. Katz and O. Firstenberg, "Light storage for one second in room-temperature alkali vapor," *Nature Communications*, vol. 9, no. 1, p. 2074, 2018.

[114] Y.-C. Wei, B.-H. Wu, Y.-F. Hsiao, P.-J. Tsai, and Y.-C. Chen, "Broadband coherent optical memory based on electromagnetically

induced transparency," *Physical Review A*, vol. 102, no. 6, p. 063720, 2020.

[115] A. C. Leung, W. Lau, A. D. Tranter, K. V. Paul, M. Rambach, B. C. Buchler, P. K. Lam, A. G. White, and T. J. Weinhold, "Highly efficient storage of cavity spdc single photons in room temperature gradient echo memory," *APL Quantum*, vol. 1, no. 3, 2024.

[116] Y.-F. Hsiao, P.-J. Tsai, H.-S. Chen, S.-X. Lin, C.-C. Hung, C.-H. Lee, Y.-H. Chen, Y.-F. Chen, I. A. Yu, and Y.-C. Chen, "Highly efficient coherent optical memory based on electromagnetically induced transparency," *Physical review letters*, vol. 120, no. 18, p. 183602, 2018.

[117] G. S. Agarwal, "Quantum statistical theory of optical-resonance phenomena in fluctuating laser fields," *Physical Review A*, vol. 18, no. 4, p. 1490, 1978.

[118] S. H. Autler and C. H. Townes, "Stark effect in rapidly varying fields," *Physical Review*, vol. 100, no. 2, p. 703, 1955.

[119] T. Peters, B. Wittrock, F. Blatt, T. Halfmann, and L. P. Yatsenko, "Thermometry of ultracold atoms by electromagnetically induced transparency," *Physical Review A*, vol. 85, no. 6, p. 063416, 2012.

[120] B. Zhao, Y.-A. Chen, X.-H. Bao, T. Strassel, C.-S. Chuu, X.-M. Jin, J. Schmiedmayer, Z.-S. Yuan, S. Chen, and J.-W. Pan, "A millisecond quantum memory for scalable quantum networks," *Nature Physics*, vol. 5, no. 2, pp. 95–99, 2009.

[121] S. M. Kay, "Statistical signal processing: estimation theory," *Prentice Hall*, vol. 1, pp. Chapter–3, 1993.

[122] R. Demkowicz-Dobrzański, M. Jarzyna, and J. Kołodyński, "Quantum limits in optical interferometry," *Progress in Optics*, vol. 60, pp. 345–435, 2015.

[123] R. Demkowicz-Dobrzański, "Optimal phase estimation with arbitrary *a priori* knowledge," *Physical Review A*, vol. 83, no. 6, p. 061802, 2011.

[124] R. E. Kass and L. Wasserman, "The selection of prior distributions by formal rules," *Journal of the American Statistical Association*, vol. 91, no. 435, pp. 1343–1370, 1996.

[125] G. Kestler, K. Ton, D. Filin, C. Cheung, P. Schneeweiss, T. Hoinkes, J. Volz, M. Safronova, A. Rauschenbeutel, and J. Barreiro, "State-insensitive trapping of alkaline-earth atoms in a nanofiber-based optical dipole trap," *PRX Quantum*, vol. 4, no. 4, p. 040308, 2023.

Appendices

Appendix A

Publications

Over the course of my PhD I have so far been an author on 2 papers

- M. Overton, J. Rubio, N. Cooper, D. Baldolini, D. Johnson, J. Anders, and L. Hackermüller. Substantial precision enhancements via adaptive symmetry-informed Bayesian metrology. *arXiv preprint arXiv:2410.10615, 2024.*

- N. Cooper, D. Johnson, B. Hopton, M. Overton, D. Stupple, A. Bratu, E. Wilson, J. Robinson, L. Coles, M. Papastavrou, L. Hackermüller. Exploiting complex 3D-printed surface structures for portable quantum technologies *In Preparation.*

the first has been submitted for publication and is currently available on the arXiv, the second is currently in preparation. These papers are included here.

Substantial precision enhancements via adaptive symmetry-informed Bayesian metrology

Matt Overton,[1,*] Jesús Rubio,[2,*] Nathan Cooper,[1] Daniele Baldolini,[1]
David Johnson,[1] Janet Anders,[3,4,†] and Lucia Hackermüller[1]

[1] *School of Physics and Astronomy, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom*
[2] *School of Mathematics and Physics, University of Surrey, Guildford GU2 7XH, United Kingdom*
[3] *Department of Physics and Astronomy, University of Exeter, Stocker Road, Exeter EX4 4QL, United Kingdom*
[4] *Institute of Physics and Astronomy, University of Potsdam, 14476 Potsdam, Germany*

Precision measurements are essential for addressing major challenges, ranging from gravitational wave detection to healthcare diagnostics. While quantum sensing experiments are constantly improving to deliver greater precision, the in-depth optimisation of measurement procedures beyond phase estimation has been overlooked. Here we present a systematic strategy for parameter estimation that can be applied across a wide range of experimental platforms operating in the low data limit. Its strength is the inclusion of experimental control parameters, which allows optimisation within a Bayesian procedure and adaptive repetition. We provide general expressions for the optimal estimator and error for any parameter amenable to symmetry-informed strategies. We demonstrate the power of this strategy by applying it to atom number estimation in a quantum technology experiment. Our protocol results in a five-fold reduction in the fractional variance of the atom number estimate compared to a standard, unoptimized protocol. Equivalently, it achieves the target precision with a third of the data points previously required. The enhanced device performance and accelerated data collection achieved with this Bayesian optimised strategy will be essential for applications in quantum computing, communication, metrology, and the wider quantum technology sector.

Introduction

From advances in gravitational wave detection [1, 2] to applications such as brain pattern imaging [3], magnetic sensing [4, 5] or inertial navigation [6], quantum technologies are revolutionising fundamental and applied research. Rapid estimation of platform parameters, as well as precise state preparation and control, underpins this success, enabling, e.g., efficient interferometric measurements [7, 8] and high-fidelity readout within a limited coherence time in quantum computing [9, 10]. Estimation efficiency [11, 12] is critical for advancing quantum technologies [4, 13] and for testing the fundamental laws of nature [14–18].

Local estimation theory [19–22] is the most common framework for benchmarking measurement processes, where the rate of information acquisition is increased by optimising the (quantum) Fisher information [23–26] over probe states and measurement operators. To achieve the theoretically optimal sensitivity, the Cramér-Rao bound must be saturated, which typically requires an asymptotically large number of repetitions [19, 21] leading to large acquisition times and often information loss at the data processing stage [27–32] of realistic experiments.

Bayesian estimation theory offers an alternative [23, 33] in which optimality is achieved by minimising uncertainty based on a probabilistic description of the experimental platform, enabling maximal information acquisition for any sample size or parameter range and thereby accelerating data collection. Its effectiveness has been demonstrated within a limited number of atomic, molecular, and optical experiments and has primarily been applied to phase estimation [34–46]. Despite the clear advantages of the Bayesian approach and its well-established foundations [21, 23, 33, 47–52], experimental adoption remains limited, revealing a need for a clear and unified methodology for experimental implementation.

Here we present a systematic, symmetry-informed strategy for experimental sensing that enables more rapid information acquisition, or equivalently, higher estimation precision. This is achieved by setting up a Bayesian optimisation for the most likely value of the parameter Θ to be estimated, where the optimisation can now include experimental control parameters y, as illustrated in Fig. 1(a). Our strategy is applicable to experimental parameters Θ that are not phases, such as atom number [53], decay rates [54, 55], or relative weights [23, 55], and significantly expands beyond conventional phase estimation.

We demonstrate the effectiveness of this adaptive strategy in an example experiment, where a cold atomic cloud is trapped inside an optical fibre [56, 57], see Fig. 1(b), and atom number is the system parameter to be estimated. Controlling the probing laser frequency according to our adaptive symmetry-informed Bayesian protocol results in significant enhancements in information acquisition speed and estimation precision.

For other experiments, particularly those operating in the low data limit, our strategy informs how to set up the estimation problem, including symmetries associated with the system parameter, and finding the expression for the optimal estimator. This systematic approach for control parameter optimisation in response to initial measurement data provides significant benefit when probing challenging regimes, such as in searches for new fundamental physics with wide *a priori* bounds for unknown parameters [2, 58, 59]. Adaptive methodologies [60–63]

* The first two authors contributed equally to this work.
† janet@qipc.org

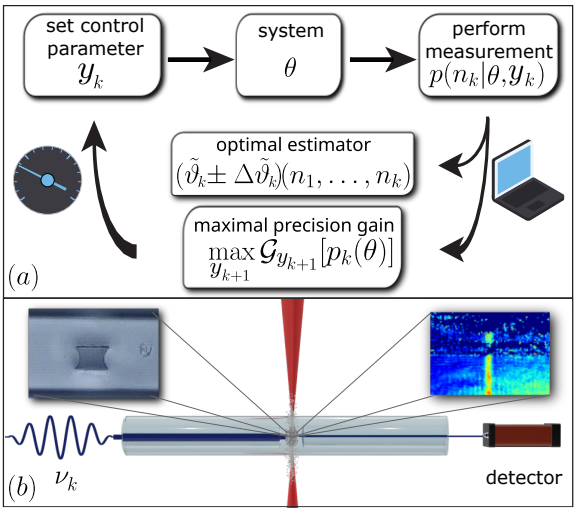

FIG. 1. (a) General adaptive, symmetry-informed, Bayesian estimation strategy. The information gained about a system parameter θ from a detected signal n_k is used to select a new optimal control parameter y_k which offers maximum information gain in the next measurement shot, $k+1$. This precision gain is invariant under the symmetries of the platform. (b) Application to an ensemble of cold Cs atoms confined in an optical dipole trap (red cones) within a microscopic hole (left inset) that intersects the core of an optical fibre. The atom number is probed via photon-absorption detection of light travelling along the fibre. The light's frequency ν_k is the control parameter which can be chosen optimally. The final estimate of atom number is sequentially improved as more photon-counts are recorded. The right inset shows an absorption image of the atoms passing through the optical fibre.

can therefore significantly boost the performance of a wide range of experiments and technologies in quantum sensing and quantum computing, while also facilitating the realisation of experiments that require empirical tuning. The proposed strategy provides a unified standard for optimal metrology in realistic experimental scenarios, facilitating adoption within the quantum technology sector.

Results

Finding efficient estimates amounts to solving an optimisation problem [23, 47, 62, 64]. We first present a new and general adaptive symmetry-informed Bayesian strategy for the estimation of parameters that are not necessarily phases. It is applicable to all experiments where measurement precision depends on the appropriate selection of one or more control parameters.

Adaptive symmetry-informed estimation strategy

Let $\theta \in [\theta_{\min}, \theta_{\max}]$ be a *hypothesis* for an unknown *true* parameter Θ to be determined indirectly through a measurement with outcome n. This outcome is related with θ and possibly a control parameter y through a likelihood function $p(n|\theta, y)$. For k outcomes $\boldsymbol{n} = (n_1, \ldots, n_k)$ measured using control parameters $\boldsymbol{y} = (y_1, \ldots, y_k)$, the likelihood is

$$p(\boldsymbol{n}|\theta, \boldsymbol{y}) = \prod_{i=1}^{k} p(n_i|\theta, y_i). \tag{1}$$

Based on this, the default choice for an estimator for the unknown parameter is the maximum likelihood estimator (MLE)

$$\tilde{\theta}_{\boldsymbol{y}}(\boldsymbol{n}) = \underset{\theta}{\text{argmax}}[p(\boldsymbol{n}|\theta, \boldsymbol{y})]. \tag{2}$$

However, the MLE is known to be efficient only in the regime of large sample sizes [19, 21]. This motivates adopting a Bayesian approach, which is generally applicable. Here, a prior probability $p(\theta)$, which encodes the initial information, is *updated* via Bayes's theorem as

$$p(\theta) \mapsto p(\theta|\boldsymbol{n}, \boldsymbol{y}) \propto p(\theta) p(\boldsymbol{n}|\theta, \boldsymbol{y}) \tag{3}$$

upon recording $\boldsymbol{n}$, where

$$p(\theta|\boldsymbol{n}, \boldsymbol{y}) \propto p(\theta) \prod_{i=1}^{k} p(n_i|\theta, y_i) \tag{4}$$

is the posterior probability. This probability contains all the information available about the true Θ [23, 39].

Now, in addition to using Bayes's theorem, one needs to respect the symmetry of the problem. For example, while phases obey circular symmetry, $\theta' = \theta + 2\pi$ [65], temperature usually appears in ratios of energies, which are invariant under transformations $\theta' = \gamma\theta$, for positive γ [30]. The symmetry of the problem can be included in the estimation strategy by introducing a symmetry function f that weights the posterior probabilities (4), thereby enforcing the desired symmetry in the optimal estimation. For a given f, the symmetry-informed estimator for Θ is

$$\tilde{\vartheta}_{\boldsymbol{y},f}(\boldsymbol{n}) = f^{-1}\left[\int d\theta\, p(\theta|\boldsymbol{n}, \boldsymbol{y}) f(\theta)\right]. \tag{5}$$

This estimator is *optimal* under quadratic losses [52]. When $f(\theta) = \theta$, the estimator is simply the mean of the posterior probability, suitable for location parameters with translation symmetry $\theta' = \theta + \gamma$, with arbitrary γ, while the estimator functions will be different for other symmetries. As such, accounting for the symmetries of the context in which it is applied is an essential step in generalising this estimation strategy. Furthermore, the uncertainty associated with this estimator is

$$\Delta\tilde{\vartheta}_{\boldsymbol{y},f}(\boldsymbol{n}) = \frac{\sqrt{\mathcal{L}_{\boldsymbol{y},f}(\boldsymbol{n})}}{|f'[\tilde{\vartheta}_{\boldsymbol{y},f}(\boldsymbol{n})]|}, \tag{6}$$

where

$$\mathcal{L}_{\boldsymbol{y},f}(\boldsymbol{n}) = \int d\theta\, p(\theta|\boldsymbol{n}, \boldsymbol{y}) f(\theta)^2 - f[\tilde{\vartheta}_{\boldsymbol{y},f}(\boldsymbol{n})]^2 \tag{7}$$

is the information loss incurred by processing $(\boldsymbol{n}, \boldsymbol{y})$.

Eqs. (5) and (6) enable experimental symmetry-informed estimation [52]. This provides maximum information for any parameter range or sample size while respecting the symmetry of the problem through the choice of f. For example, the likelihood may be invariant under a symmetry group [66], or we may enforce invariance under all reparametrisations using information geometry [33, 67]. More generally, parameter types can be associated with transformations such as translations for shifts or rescalings for rates [23, 55], as discussed. These symmetries can be shown [52, 66] to determine maximum ignorance priors of the form

$$p_{\mathrm{MI}}(\theta) \propto f'(\theta), \tag{8}$$

from which the function f required to use Eqs. (5) and (6) can be calculated, once the left hand side of Eq. (8) has been derived.

Having taken care of symmetry considerations, the information loss (7) can now be *adaptively* reduced by selecting y_k for the k-th measurement to maximise the precision gain [39, 52]

$$\mathcal{G}_{y_k,f} = \sum_n p(n|y_k)\, f[\tilde{\vartheta}_{y_k,f}(n)]^2. \tag{9}$$

Here, $p(n|y_k) = \int d\theta\, p_{k-1}(\theta) p(n|\theta, y_k)$, with

$$p_{k-1}(\theta) \coloneqq \begin{cases} p(\theta) & \text{if } k = 1 \\ p(\theta) \prod_{i=1}^{k-1} p(n_i|\theta, y_i) & \text{for } k > 1, \end{cases} \tag{10}$$

and $\tilde{\vartheta}_{y_k,f}(n)$ is defined as in Eq. (5) using $p(\theta|n, y_k) \propto p_{k-1}(\theta) p(n|\theta, y_k)$.

Eqs. (5), (6) and (9) constitute our first result. Their derivation is provided in the Appendices. As we shall see next, they enable the experimental implementation of optimal estimators for any parameter for which a suitable f exists.

Application to experimental atom number metrology

Having derived the general equations for implementing adaptive symmetry-informed Bayesian metrology, we now apply them to a specific experiment to demonstrate their power. The example experiment is a cold atom setup [56], where Cs atoms are trapped within a hole cut into an optical fibre, see Fig. 1(b). This setup is a promising platform [56, 57] for the integration of cold-atom based sensors, quantum memories and multi-qubit gates within existing photonic circuit architectures [68]. The task is to determine the number of atoms N trapped within the fibre. Previously this required $30 - 100$ shots per estimate, rendering atom number estimation the main rate-limiting step. Requiring a large number of repetitions, or long measurement times, is typical e.g. for atom based systems such as atom-nanofibre experiments [69, 70], precision magnetometers [71] and atom interferometry [72]. Here we address this problem with an adaptive strategy based on Eq. (9) and an optimal atom number estimator $\widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n}) \pm \Delta\widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n})$ built upon Eqs. (5) and (6).

In the fibre setup, near-resonant light propagating through the fibre is absorbed by the atoms in the intersection at a rate proportional to [56]

$$\zeta_\nu = \frac{\Gamma^2}{\Gamma^2 + 4(\nu - \nu_{\mathrm{r}})^2}, \tag{11}$$

where ν is the laser frequency, and ν_{r} and Γ are the resonant frequency and natural linewidth of the atomic transition, respectively. Since the light intensity is low compared to the saturation intensity of the atomic transition, the optical power propagating through the fibre is attenuated such that the expected photon count is

$$\bar{n}_\nu(\Phi, \Theta) = \Phi\, \mathrm{e}^{-\zeta_\nu \Theta}. \tag{12}$$

Here, Φ is the expected photon number detected with no atoms present, $\exp(-\zeta_\nu\Theta)$ is the transmittance of the medium, and Θ its on-resonance optical depth. Both Φ and Θ are initially unknown. The atom number N we seek to determine is proportional to the optical depth Θ, i.e. $N = \kappa\,\Theta$ [56], with a constant κ determined by the mode area of the waveguide and the on-resonance absorption cross-section of the atoms. In such an absorption setup, the standard method of estimating N is then to count photons with and without atoms loaded into the hole ($n = n_{\mathrm{a}}$ and $n = n_{\mathrm{b}}$, respectively), and calculating the maximum likehood estimator $\frac{\kappa}{\zeta_\nu} \log \frac{\langle n_a \rangle}{\langle n_b \rangle}$, where ν is set on resonance with ν_r (i.e., $\zeta_\nu = 1$). This estimator arises from the fact that ordinary (unsqueezed) light exhibits a Poissonian number distribution, i.e. the photon count statistics reads

$$p(n|\Phi, \Theta, \nu) = P(n|\bar{n}_\nu(\Phi, \Theta)). \tag{13}$$

where $P(n|z) = z^n \mathrm{e}^{-z}/n!$.

Instead, we will now apply the new estimation protocol. Here, we consider the laser frequency ν as the control parameter y in Eq. (5). It can be tuned to be on resonance with the atomic transition ν_{r}, or differ from it. Hence, we will apply symmetry-informed estimation directly to the optical depth Θ and with Φ, the photon-number without atoms, featuring as an undetermined nuisance parameter. This is thus a multiparameter estimation problem [50].

Let $\theta \in [\theta_{\min}, \theta_{\max}]$ and $\varphi \in [\varphi_{\min}, \varphi_{\max}]$ be the hypotheses for the real Θ and Φ, respectively. When two observers run the experiment at frequency ν, their average photon counts are comparable, but their initial hypotheses can differ due to prior ignorance. That is, $\bar{n}_\nu(\varphi', \theta') = \bar{n}_\nu(\varphi, \theta)$, so that $\varphi'/\varphi = \exp[\zeta_\nu(\theta' - \theta)]$. Equating both sides to a positive γ reveals transformations

$$\varphi' = \gamma\varphi, \qquad \zeta_\nu\theta' = \zeta_\nu\theta + \log(\gamma) \tag{14}$$

under which Eq. (13) is invariant. These imply that the expected number of photons is a scale parameter, while the optical depth is a location parameter [23]. Then, if both observers are equally ignorant, and given that ν does not inform Φ and Θ, their priors must respect the aforementioned invariance as $p_{\mathrm{MI}}(\varphi', \theta')\, d\varphi'\, d\theta' =$

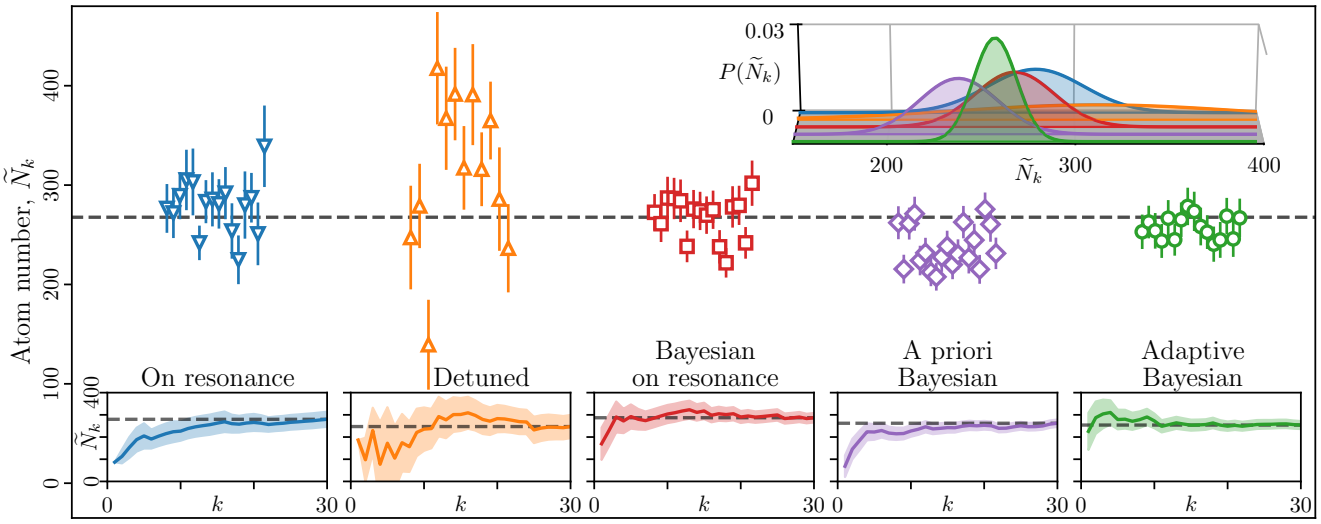

FIG. 2. Atom number estimates $\tilde{N}_k$ resulting from measurements using the standard (blue and orange) and Bayesian estimation strategies (red, purple and green) described in the text. Each point estimate employs $k = 30$ measurement shots. In the on-resonance and detuned cases, error bars are given by the standard error for each 30-shot measurement, and via Eq. (20b) otherwise. The lower insets show the progression of the estimate with increasing k, for a single measurement run, in each case. The shading denotes the aforementioned errors. The small differences in the mean value obtained with each method are consistent with normal experimental drifts in loading efficiency, owing to the temporal separation of the data collection using each methodology. The dashed lines indicate final estimates in the insets and overall average across all methods in the main figure, shown here to guide the eye. The inset in the top right corner gives a visual Gaussian representation, $P(\tilde{N}_k)$, of the mean and the standard deviation of the data presented in the main body of the figure for each method.

$p_{\mathrm{MI}}(\varphi,\theta)\,d\varphi\,d\theta$. As shown in the Appendices, this leads to the multiparameter prior

$$p_{\mathrm{MI}}(\varphi,\theta) = \left[(\theta_{\max}-\theta_{\min})\log\left(\frac{\varphi_{\max}}{\varphi_{\min}}\right)\varphi\right]^{-1}. \quad (15)$$

We can see that, despite having no *a priori* knowledge about the values of optical depth and expected number of photons beyond the initial ranges, this prior is *not* flat in φ. This is a consequence of the scale symmetry in Eq. (14) [23, 66], which the prior (15) respects.

Since Eq. (15) is a two-parameter prior, it gives rise to two different symmetry functions. For the expected number of photons φ, we marginalise over θ, i.e.,

$$p_{\mathrm{MI}}(\varphi) = \int_{\theta_{\min}}^{\theta_{\max}} d\theta\, p_{\mathrm{MI}}(\varphi,\theta) \propto \frac{1}{\varphi}. \quad (16)$$

Introducing this prior into Eq. (8) leads to the logarithmic symmetry function

$$f(\varphi) = c_1 \log\left(\frac{\varphi}{c_2}\right), \quad (17)$$

where c_1 and c_2 are free constants. For the optical depth θ, we instead marginalise over φ as

$$p_{\mathrm{MI}}(\theta) = \int_{\varphi_{\min}}^{\varphi_{\max}} d\varphi\, p_{\mathrm{MI}}(\varphi,\theta) \propto 1, \quad (18)$$

leading to the linear symmetry function

$$f(\theta) = c_1'\theta + c_2', \quad (19)$$

with constants c_1' and c_2'. Since our focus is θ, and φ is a nuisance parameter, only the linear function (19) is required in our particular case. We note, however, that the logarithmic function (17) has been shown useful in quantum thermometry, where the parameter of interest, temperature, is often a scale parameter [30, 39, 61]. In general, the function f can take many forms; for example, this methodology could be applied equally effectively to experiments that measure a conversion efficiency such as Ref. [73], which is a weight parameter and would thus require $f(z) = 2\,\mathrm{artanh}(2z-1)$ [52].

Once the relevant symmetry function (19) has been identified, and using the proportionality between N and Θ, Eqs. (5) and (6) straightforwardly deliver the optimal atom number estimator and its uncertainty

$$\widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n}) = \kappa \int_{\theta_{\min}}^{\theta_{\max}} d\theta\, p(\theta|\boldsymbol{n},\boldsymbol{\nu})\theta, \quad (20a)$$

$$\Delta\widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n})^2 = \kappa^2 \int_{\theta_{\min}}^{\theta_{\max}} d\theta\, p(\theta|\boldsymbol{n},\boldsymbol{\nu})\theta^2 - \widetilde{N}_{\boldsymbol{\nu}}(\boldsymbol{n})^2, \quad (20b)$$

where the constants in (19) were chosen to $c_1' = 1$ and $c_2' = 0$. Here,

$$p(\theta|\boldsymbol{n},\boldsymbol{\nu}) \propto \int_{\varphi_{\min}}^{\varphi_{\max}} \frac{d\varphi}{\varphi} \prod_{i=1}^{k} p(n_i|\varphi,\theta,\nu_i) \quad (21)$$

is the marginalised posterior. Eq. (20) respects the location symmetry associated with the optical depth θ and the scale symmetry associated with the nuisance parameter φ through the ignorance prior (8). This is our second result.

	On resonance	Detuned	On reso. Bayesian	*A priori* Bayesian	**Adaptive Bayesian**
NSR in %	0.89	6.10	0.62	0.77	**0.19**
$k_{\min}$	13	17	12	12	**8**

TABLE I. Comparison of atom number estimation strategies in terms of noise-to-signal ratio for $k = 30$ measurement shots, as well as mean minimal number of shots $k_{\min}$ after which the estimate stays within 10% of the final estimate, from the data shown in Fig. 2.

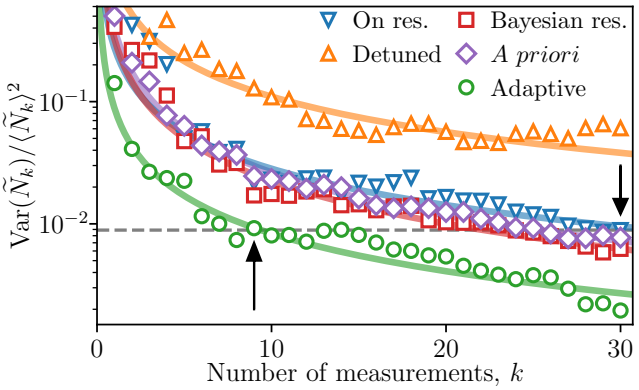

FIG. 3. Noise-to-signal ratio (NSR) versus shot number k for the on-resonance (blue inverted triangles), detuned (orange triangles), on-resonance Bayesian (red squares), *a priori* Bayesian (purple rhombuses), and adaptive Bayesian (green circles) strategies. The data points are obtained empirically from the distribution of m estimates after each shot number k. The dashed line indicates the NSR for on-resonance method with 30 shots (downward facing arrow), which is achieved by our adaptive protocol with only 9 shots (upward facing arrow).

Absorption measurements at atom–optical interface

We are now ready to employ the protocol experimentally. We measure the outcomes $\boldsymbol{n} = (n_{\mathrm{a},1}, n_{\mathrm{b},1}, \ldots, n_{\mathrm{a},k}, n_{\mathrm{b},k})$, with $1 \leq k \leq 30$, alternating photon counts with and without atoms and setting $p(n_i|\varphi, 0, \nu_i)$ in Eq. (21) for the latter. The laser frequency is used as an adaptive control parameter by optimising the precision gain (9) with $\tilde{\vartheta}_{\nu_k}(n) \propto \widetilde{N}_{\nu_k}(n)$, thus finding the frequencies $\boldsymbol{\nu} = (\nu_{1\mathrm{a}}, \nu_{1\mathrm{b}}, \ldots, \nu_{k\mathrm{a}}, \nu_{k\mathrm{b}})$, see Appendices Fig. 4. We choose prior ranges $5 \leq \varphi \leq 20$ and $0 \leq \theta \leq 8$. For the Cs D_2 $F = 4 \to F' = 5$ transition, the physical values are $\nu_{\mathrm{r}} = 351.721961\,\mathrm{THz}$, $\Gamma = 2\pi{\cdot}5.234\,\mathrm{MHz}$, and $\kappa = 84.9$ [56, 74]. All experimental parameters other than the control parameter are kept constant. Eqs. (9) and (20) are rapidly evaluated between shots by caching function results, managing tolerances, and adaptively truncating probabilities, see Appendices.

To quantify the effectiveness of the adaptive, symmetry-informed Bayesian strategy, we compare it to four alternative methods. The first ("on resonance") is the standard method [56] discussed above, which uses the optical depth estimate $\frac{\kappa}{\zeta_\nu} \log \frac{\langle n_a \rangle}{\langle n_b \rangle}$ and an error propagated from this formula. The second method ("detuned") is analogous, but with the laser frequency detuned below the atomic resonance, to $\nu - \nu_{\mathrm{r}} = 5\,\mathrm{MHz}$. The third method ("on resonance Bayesian") uses laser light again in resonance with the atomic transition, $\nu = \nu_{\mathrm{r}}$, but now adopting the Bayesian estimator in Eqs. (20). The fourth method ("*a priori* Bayesian") is similar to the third, but here the laser frequency is chosen to maximise the precision gain (9) for $k = 1$. This is an *a priori* optimisation as it involves no outcomes [27, 39], and here results in a detuning of $\nu - \nu_{\mathrm{r}} = 1.95\,\mathrm{MHz}$. Apart from the detuned case, these methods are ordered according to increasing level of optimisation, with the fifth method being our adaptive, symmetry-informed Bayesian strategy.

Fig. 2 shows the atom number estimates yielded by the on-resonance (blue inverted triangles), detuned (orange triangles), on-resonance Bayesian (red squares), *a priori* Bayesian (purple rhombuses), and adaptive Bayesian (green circles) methods. We repeat each method m times and quantify its performance by the empirical noise-to-signal ratio (NSR) $\widetilde{\mathrm{Var}}(\widetilde{N}_k)/\langle\widetilde{N}_k\rangle^2$ [39], where

$$\widetilde{\mathrm{Var}}(\widetilde{N}_k) = \langle\widetilde{N}_k^2\rangle - \langle\widetilde{N}_k\rangle^2, \quad \text{with} \quad \langle\widetilde{N}_k\rangle = \frac{1}{m}\sum_{j=1}^{m} \widetilde{N}_{k,j}, \tag{22}$$

and the k-th estimate during the j-th repetition is denoted as $\widetilde{N}_{k,j}$. The results, given in Tab. I, clearly show that the different strategies exhibit marked differences in their performance. While the two non-adaptive Bayesian protocols yield a NSR comparable to that for the on-resonance method, just under 1%, and all three outperform the detuned method by an order of magnitude, the adaptive protocol surpasses them all with a NSR of less than 0.2%. Notably, the latter outperforms the on-resonance method – which is the current standard for cold atom platforms – by a factor of 4.68. This approximately five-fold enhancement constitutes our third result.

Fig. 3 shows that these conclusions hold even for a low shot number. There, the aforementioned NSR is represented as a function of k. A crucial implication is that the NSR achieved by the standard technique (blue inverted triangles) using $k = 30$ shots is rendered by the adaptive protocol (green circles) with only $k = 9$ shots, i.e., about one third of the resources. This is our fourth result. Interestingly, Fig. 3 also shows that the *a priori* method (purple rhombuses), while simpler to implement, can be notably worse than an adaptive treatment despite already benefiting from use of the precision gain (9), which had been predicted in Ref. [39].

The adaptive protocol was also found to converge to a final result within a desired precision faster. I.e., when recording the number of measurement steps $k_{\min}$ needed before the current estimate settles within a fractional precision of $10\,\%$ of the final estimate based on 30 shots, and taking the mean of this value for the data sets shown in Fig. 2, we obtain the results listed in the last row of Tab. I. As can be seen, the adaptive protocol achieves the specified fractional precision with approximately 40% fewer shots (8 instead of 13) than the standard on-resonance method requires.

Discussion

In summary, we have developed a widely applicable symmetry-informed estimation strategy that will enable accelerated data collection across the quantum technology sector. We have demonstrated its application to a test platform, atom number estimation in a fibre experiment. The optimised protocol yielded substantial enhancements in the precision of the atom number estimate (five-fold increase), or alternatively, increased information acquisition speed (a third of data points), compared to current standards. We have further shown that real-time adaptive optimisation is feasible with standard computational resources even in multiparameter setups [28, 39]. In general, use of Eqs. (5), (6) and (9) expands beyond the application range of Bayesian phase estimation [65] to optimal Bayesian estimation of parameters with non-cyclical symmetry, such as temperature, atom number and state occupation.

Our approach is systematic and adaptable, making it applicable across a wide range of quantum technology experiments. It provides a rigorous estimation framework that performs *at least* on par with standard methods, while often delivering superior results, particularly in data-limited regimes. The achievable gain in precision and information acquisition speed makes it a powerful tool for quantum sensing, metrology, and computation. Given its generality and performance advantages, the presented systematic strategy should be adopted in experiments ranging from precision spectroscopy and magnetometry to quantum state preparation and fundamental physics tests.

At a fundamental level, the next theory step is implementing not only optimal estimators, but finding optimal states $\rho_y(\theta)$ and measurements $\mathcal{M}_y(n)$ with statistics $p(n|\theta, y) = \mathrm{Tr}[\mathcal{M}_y(n)\rho_y(\theta)]$. Optimal measurements are given by projecting onto the eigenspace of an operator $\mathcal{S}_{y,f}$ solution to the Lyapunov equation

$$\mathcal{S}_{y,f}\rho_{y,f,0} + \rho_{y,f,0}\mathcal{S}_{y,f} = 2\rho_{y,f,1}, \tag{23}$$

with $\rho_{y,f,l} = \int d\theta\, p(\theta)\rho_y(\theta)f(\theta)^l$ [52, 75], while optimal states can be found by combining this with techniques such as tensor networks and machine learning [51, 76–81].

Acknowledgments: JR gratefully thanks J. Boeyens, J. Glatthard, G. Barontini, V. Guarrera, M. Tsang and L. Correa for helpful discussions. JR acknowledges financial support from the Surrey Future Fellowship Programme. JA gratefully acknowledges funding from EPSRC (EP/R045577/1) and thanks the Royal Society for support. The experiment was supported by the grant 62420 from the John Templeton Foundation, the IUK project No.133086, EPSRC grants EP/T001046/1, EP/R024111/1, EP/M013294/1, EP/Y005139/1 and EP/Z533166/1, and by the European Commission grant ErBeStA (no.800942).

Data availability: All data upon which our conclusions are based are presented herein. All additional data related to the experimental system are available from the authors upon reasonable request.

For the purpose of open access, the authors have applied a 'Creative Commons Attribution' (CC BY) licence to any Author Accepted Manuscript version arising from this submission.

References

[1] J. M. Hogan and M. A. Kasevich, Atom-interferometric gravitational-wave detection using heterodyne laser links, Phys. Rev. A **94**, 033632 (2016).

[2] L. Badurina, E. Bentine, D. Blas, K. Bongs, D. Bortoletto, T. Bowcock, K. Bridges, W. Bowden, O. Buchmueller, C. Burrage, J. Coleman, G. Elertas, J. Ellis, C. Foot, V. Gibson, M. Haehnelt, T. Harte, S. Hedges, R. Hobson, M. Holynski, T. Jones, M. Langlois, S. Lellouch, M. Lewicki, R. Maiolino, P. Majewski, S. Malik, J. March-Russell, C. McCabe, D. Newbold, B. Sauer, U. Schneider, I. Shipsey, Y. Singh, M. Uchida, T. Valenzuela, M. v. d. Grinten, V. Vaskonen, J. Vossebeld, D. Weatherill, and I. Wilmut, AION: an atom interferometer observatory and network, J. Cosmol. Astropart. Phys. **2020** (05), 011–011.

[3] R. Zhang, W. Xiao, Y. Ding, Y. Feng, X. Peng, L. Shen, C. Sun, T. Wu, Y. Wu, Y. Yang, Z. Zheng, X. Zhang, J. Chen, and H. Guo, Recording brain activities in unshielded earth's field with optically pumped atomic magnetometers, Sci. Adv. **6**, eaba8792 (2020).

[4] C. L. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, Rev. Mod. Phys. **89**, 035002 (2017).

[5] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. D. Muñoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes, Moving magnetoencephalography towards real-world applications with a wearable system., Nature **555**, 657–66 (2018).

[6] F. A. Narducci, A. T. Black, and J. H. Burke, Advances toward fieldable atom interferometers, Adv. Phys. X **7**, 1946426 (2022).

[7] J. Lee, R. Ding, J. Christensen, R. R. Rosenthal, A. Ison, D. P. Gillund, D. Bossert, K. H. Fuerschbach, W. Kindel, P. S. Finnegan, J. R. Wendt, M. Gehl, A. Kodigala, H. McGuinness, C. A. Walker, S. A. Kemme, A. Lentine, G. Biedermann, and P. D. D. Schwindt, A compact cold-atom interferometer with a high data-rate grating magneto-optical trap and a photonic-integrated-circuit-compatible laser system, Nat. Commun. **13**, 5131 (2022).

[8] B. Adams, S. Kinge, K. Bongs, and Y. H. Lien, Magneto-optical trap performance for high-bandwidth applications, Phys. Rev. A **108**, 063111 (2023).

[9] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, M. Kwon, M. Ebert, J. Cherek, M. T. Lichtman, M. Gillette, J. Gilbert, D. Bowman, T. Ballance, C. Campbell, E. D. Dahl, O. Crawford, N. S. Blunt, B. Rogers, T. Noel, and M. Saffman, Multi-qubit entanglement and algorithms on a neutral-atom quantum computer, Nature **604**, 457 (2022).

[10] Wintersperger, Karen, Dommert, Florian, Ehmer, Thomas, Hoursanov, Andrey, Klepsch, Johannes, Mauerer, Wolfgang, Reuber, Georg, Strohm, Thomas, Yin, Ming, and Luber, Sebastian, Neutral atom quantum computing hardware: performance and end-user perspective, EPJ Quantum Technol. **10**, 32 (2023).

[11] X. Deng, S. Li, Z. Chen, Z. Ni, Y. Cai, J. Mai, L. Zhang, P. Zheng, H. Yu, C. Zou, S. Liu, F. Yan, Y. Xu, and D. Yu, Quantum-enhanced metrology with large fock states, Nat. Phys. **20**, 1874–1880 (2024).

[12] Y. Len, T. Gefen, A. Retzker, and J. Kołodyński, Quantum metrology with imperfect measurements, Nat. Commun. **13**, 6971 (2022).

[13] P. Horodecki, L. Rudnicki, and K. Życzkowski, Five open problems in quantum information theory, PRX Quantum **3**, 010101 (2022).

[14] S. G. Turyshev, U. E. Israelsson, M. Shao, N. Yu, A. Kusenko, E. L. Wright, C. W. F. Everitt, M. Kasevich, J. A. Lipa, J. C. Mester, R. D. Reasenberg, R. L. Walsworth, N. Ashby, H. Gould, and H. J. Paik, Space-based research in fundamental physics and quantum technologies, Int. J. Mod. Phys. D **16**, 1879 (2007).

[15] R. Kaltenbaek, A. Acin, L. Bacsardi, P. Bianco, P. Bouyer, E. Diamanti, C. Marquardt, Y. Omar, V. Pruneri, E. Rasel, *et al.*, Quantum technologies in space, Exp. Astron. **51**, 1677 (2021).

[16] A. Belenchia, M. Carlesso, Ömer Bayraktar, D. Dequal, I. Derkach, G. Gasbarri, W. Herr, Y. L. Li, M. Rademacher, J. Sidhu, D. K. Oi, S. T. Seidel, R. Kaltenbaek, C. Marquardt, H. Ulbricht, V. C. Usenko, L. Wörner, A. Xuereb, M. Paternostro, and A. Bassi, Quantum physics in space, Phys. Rep. **951**, 1 (2022).

[17] J. Ye and P. Zoller, Essay: Quantum sensing with atomic, molecular, and optical platforms for fundamental physics, Phys. Rev. Lett. **132**, 190001 (2024).

[18] S. D. Bass and M. Doser, Quantum sensing for particle physics, Nat. Rev. Phys. **6**, 329 (2024).

[19] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory* (Prentice Hall, 1993).

[20] M. G. Paris, Quantum estimation for quantum technology, Int. J. Quantum Inf. **7**, 125 (2009).

[21] R. Demkowicz-Dobrzański, M. Jarzyna, and J. Kołodyński, Quantum Limits in Optical Interferometry, Prog. Opt. **60**, 345 (2015).

[22] V. Montenegro, C. Mukhopadhyay, R. Yousefjani, S. Sarkar, U. Mishra, M. G. A. Paris, and A. Bayat, Review: Quantum metrology and sensing with many-body systems, arXiv:2408.15323 (2024).

[23] E. T. Jaynes, *Probability Theory: The Logic of Science* (Cambridge University Press, 2003).

[24] J. Liu, H. Yuan, X.-M. Lu, and X. Wang, Quantum Fisher information matrix and multiparameter estimation, J. Phys. A: Math. Theor. **53**, 023001 (2020).

[25] Y. L. Len, T. Gefen, A. Retzker, and J. Kołodyński, Quantum metrology with imperfect measurements, Nature Communications **13**, 6971 (2022).

[26] D. R. Arvidsson-Shukur, N. Yunger Halpern, H. V. Lepage, A. A. Lasek, C. H. Barnes, and S. Lloyd, Quantum advantage in postselected metrology, Nature communications **11**, 3775 (2020).

[27] J. Rubio and J. A. Dunningham, Quantum metrology in the presence of limited data, New J. Phys. **21**, 043037 (2019).

[28] M. Valeri, E. Polino, D. Poderini, I. Gianani, G. Corrielli, A. Crespi, R. Osellame, N. Spagnolo, and F. Sciarrino, Experimental adaptive Bayesian estimation of multiple phases with limited data, npj Quantum Inf. **6**, 92 (2020).

[29] S. Morelli, A. Usui, E. Agudelo, and N. Friis, Bayesian parameter estimation using Gaussian states and measurements, Quantum Sci. Technol. **6**, 025018 (2021).

[30] J. Rubio, J. Anders, and L. A. Correa, Global quantum thermometry, Phys. Rev. Lett. **127**, 190402 (2021).

[31] W. Salmon, S. Strelchuk, and D. Arvidsson-Shukur, Only Classical Parameterised States have Optimal Measurements under Least Squares Loss, Quantum **7**, 998 (2023).

[32] J. J. Meyer, S. Khatri, D. S. França, J. Eisert, and P. Faist, Quantum metrology in the finite-sample regime, arXiv:2307.06370 (2023).

[33] W. v. d. Linden, V. Dose, and U. v. Toussaint, *Bayesian Probability Theory: Applications in the Physical Sciences* (Cambridge University Press, 2014).

[34] D. Brivio, S. Cialdi, S. Vezzoli, B. T. Gebrehiwot, M. G. Genoni, S. Olivares, and M. G. A. Paris, Experimental estimation of one-parameter qubit gates in the presence of phase diffusion, Phys. Rev. A **81**, 012305 (2010).

[35] S. Paesani, A. A. Gentile, R. Santagati, J. Wang, N. Wiebe, D. P. Tew, J. L. O'Brien, and M. G. Thompson, Experimental Bayesian Quantum Phase Estimation on a Silicon Photonic Chip, Phys. Rev. Lett. **118**, 100503 (2017).

[36] A. Lumino, E. Polino, A. S. Rab, G. Milani, N. Spagnolo, N. Wiebe, and F. Sciarrino, Experimental phase estimation enhanced by machine learning, Phys. Rev. Appl. **10**, 044033 (2018).

[37] R. Kaubruegger, D. V. Vasilyev, M. Schulte, K. Hammerer, and P. Zoller, Quantum Variational Optimization of Ramsey Interferometry and Atomic Clocks, Phys. Rev. X **11**, 041045 (2021).

[38] I. Gianani, M. G. Genoni, and M. Barbieri, Assessing data postprocessing for quantum estimation, IEEE J. Sel. Top. Quantum Electron. **26**, 1 (2020).

[39] J. Glatthard, J. Rubio, R. Sawant, T. Hewitt, G. Barontini, and L. A. Correa, Optimal cold atom thermometry using adaptive bayesian strategies, PRX Quantum **3**, 040330 (2022).

[40] L. Gerster, F. Martínez-García, P. Hrmo, M. W. van Mourik, B. Wilhelm, D. Vodola, M. Müller, R. Blatt, P. Schindler, and T. Monz, Experimental Bayesian Calibration of Trapped-Ion Entangling Operations, PRX Quantum **3**, 020350 (2022).

[41] M. Valeri, V. Cimini, S. Piacentini, F. Ceccarelli, E. Polino, F. Hoch, G. Bizzarri, G. Corrielli, N. Spagnolo, R. Osellame, and F. Sciarrino, Experimental multiparameter quantum metrology in adaptive regime, Phys. Rev. Res. **5**, 013138 (2023).

[42] A. Oliveira, J. Rubio, C. Noble, J. Anderson, J. Anders, and A. Mulholland, Fluctuation relations to calculate protein redox potentials from molecular dynamics simulations, J. Chem. Theory Comput. **20**, 385 (2023).

[43] V. Cimini, E. Polino, M. Valeri, N. Spagnolo, and F. Sciarrino, Benchmarking Bayesian quantum estimation, Quantum Sci. Technol. **9**, 035035 (2024).

[44] F. Belliardo, V. Cimini, E. Polino, F. Hoch, B. Piccirillo, N. Spagnolo, V. Giovannetti, and F. Sciarrino, Optimizing quantum-enhanced Bayesian multiparameter estimation of phase and noise in practical sensors, Phys. Rev. Res. **6**, 023201 (2024).

[45] T. Hewitt, T. Bertheas, M. Jain, Y. Nishida, and G. Barontini, Controlling the interactions in a cold atom quantum impurity system, Quantum Sci. Technol. **9**, 035039 (2024).

[46] Z. Ma, C. Han, Z. Tan, H. He, S. Shi, X. Kang, J. Wu, J. Huang, B. Lu, and C. Lee, Adaptive cold-atom magnetometry mitigating the trade-off between sensitivity and dynamic range, Science Advances **11**, eadt3938 (2025).

[47] C. W. Helstrom, *Quantum Detection and Estimation The-*

ory (Academic Press, New York, 1976).
[48] G. D'Agostini, Bayesian inference in processing experimental data: principles and basic applications, Rep. Prog. Phys. **66**, 1383 (2003).
[49] K. Macieszczak, M. Fraas, and R. Demkowicz-Dobrzański, Bayesian quantum frequency estimation in presence of collective dephasing, New J. Phys. **16**, 113002 (2014).
[50] R. Demkowicz-Dobrzański, W. Górecki, and M. Guţă, Multi-parameter estimation beyond quantum Fisher information, J. Phys. A: Math. Theor. **53**, 363001 (2020).
[51] J. Bavaresco, P. Lipka-Bartosik, P. Sekatski, and M. Mehboudi, Designing optimal protocols in Bayesian quantum parameter estimation with higher-order operations, Phys. Rev. Res. **6**, 023305 (2024).
[52] J. Rubio, First-principles construction of symmetry-informed quantum metrologies, Phys. Rev. A **110**, L030401 (2024).
[53] L. Su, A. Douglas, M. Szurek, A. H. Hébert, A. Krahn, R. Groth, G. A. Phelps, O. Marković, and M. Greiner, Fast single atom imaging for optical lattice arrays., Nat. Commun. **16**, 1017 (2025).
[54] Q. Zhang, Y. Xia, W. Wu, C. Qiu, P. Chen, W. Li, H. Jing, and Y. Zhou, Observation of quantum strong Mpemba effect., Nat. Commun. **16**, 301 (2025).
[55] J. Rubio, Quantum scale estimation, Quantum Sci. Technol. **8**, 015009 (2022).
[56] E. D. Ros, N. Cooper, J. Nute, and L. Hackermueller, Cold atoms in micromachined waveguides: A new platform for atom-photon interactions, Phys. Rev. Research **2**, 033098 (2020).
[57] N. Cooper, E. D. Ros, C. Briddon, V. Naniyil, M. Greenaway, and L. Hackermueller, Prospects for strongly coupled atom-photon quantum nodes, Sci. Rep. **9**, 7798 (2019).
[58] B. Schrinski, P. Haslinger, J. Schmiedmayer, K. Hornberger, and S. Nimmrichter, Testing collapse models with bose-einstein-condensate interferometry, Phys. Rev. A **107**, 043320 (2023).
[59] K. Clements, B. Elder, L. Hackermueller, M. Fromhold, and C. Burrage, Detecting dark domain walls through their impact on particle trajectories in tailored ultrahigh vacuum environments, Phys. Rev. D **109**, 123023 (2024).
[60] E. Polino, M. Valeri, N. Spagnolo, and F. Sciarrino, Photonic quantum metrology, AVS Quantum Science **2**, 024703 (2020).
[61] M. Mehboudi, M. R. Jørgensen, S. Seah, J. B. Brask, J. Kołodyński, and M. Perarnau-Llobet, Fundamental Limits in Bayesian Thermometry and Attainability via Adaptive Strategies, Phys. Rev. Lett. **128**, 130502 (2022).
[62] S. Kurdziałek, W. Górecki, F. Albarelli, and R. Demkowicz-Dobrzański, Using adaptiveness and causal superpositions against noise in quantum metrology, Phys. Rev. Lett. **131**, 090801 (2023).
[63] J. G. Smith, C. H. W. Barnes, and D. R. M. Arvidsson-Shukur, Adaptive Bayesian quantum algorithm for phase estimation, Phys. Rev. A **109**, 042412 (2024).
[64] B. MacLellan, P. Roztocki, S. Czischek, and R. G. Melko, End-to-end variational quantum sensing, arXiv:2403.02394 (2024).
[65] R. Demkowicz-Dobrzański, Optimal phase estimation with arbitrary a priori knowledge, Phys. Rev. A **83**, 061802 (2011).
[66] R. E. Kass and L. Wasserman, The Selection of Prior Distributions by Formal Rules, J Am Stat Assoc **91**, 1343 (1996).
[67] S.-I. Amari, *Information Geometry and Its Applications*, 1st ed., Applied mathematical sciences (Springer, Tokyo, Japan, 2016).
[68] W. Bogaerts, D. Pérez, J. Capmany, D. A. B. Miller, J. Poon, D. Englund, F. Morichetti, and A. Melloni, Programmable photonic circuits, Nature **586**, 207 (2020).
[69] T. Nieddu, V. Gokhroo, and S. N. Chormaic, Optical nanofibres and neutral atoms, Journal of Optics **18** (2016).
[70] N. V. Corzo, J. Raskop, A. Chandra, A. S. Sheremet, B. Gouraud, and J. Laurat, Waveguide-coupled single collective excitation of atomic arrays, Nature **566**, 359 (2019).
[71] C. Troullinou, V. G. Lucivero, and M. W. Mitchell, Quantum-enhanced magnetometry at optimal number density, Phys. Rev. Lett. **131**, 133602 (2023).
[72] Z. Zhou, S. C. Carrasco, C. Sanner, V. S. Malinovsky, and R. Folman, Geometric phase amplification in a clock interferometer for enhanced metrology, arXiv: , 2405.10226 (2024).
[73] V. Naniyil, Y. Zhou, G. Simmonds, N. Cooper, W. Li, and L. Hackermueller, Observation of collectivity enhanced magnetoassociation of ^{6}Li in the quantum degenerate regime, New Journal of Physics **24**, 113005 (2022).
[74] D. A. Steck, Cesium D line data, revision 2.3.3, `http://steck.us/alkalidata` (2024), accessed: 2024-09-10.
[75] M. Tsang, Generalized conditional expectations for quantum retrodiction and smoothing, Phys. Rev. A **105**, 042213 (2022).
[76] I. S. M. Schuld and F. Petruccione, An introduction to quantum machine learning, Contemp. Phys. **56**, 172 (2015).
[77] L. J. Fiderer, J. Schuff, and D. Braun, Neural-Network Heuristics for Adaptive Bayesian Quantum Estimation, PRX Quantum **2**, 020303 (2021).
[78] D. Koutný, L. Motka, Z. c. v. Hradil, J. Řeháček, and L. L. Sánchez-Soto, Neural-network quantum state tomography, Phys. Rev. A **106**, 012409 (2022).
[79] V. Gebhart, R. Santagati, A. A. Gentile, E. M. Gauger, D. Craig, N. Ares, L. Banchi, F. Marquardt, L. Pezzè, and C. Bonato, Learning quantum systems, Nat. Rev. Phys. **5**, 141 (2023).
[80] E. Rinaldi, M. G. Lastre, S. G. Herreros, S. Ahmed, M. Khanahmadi, F. Nori, and C. S. Muñoz, Parameter estimation from quantum-jump data using neural networks, Quantum Sci. Technol. **9**, 035018 (2024).
[81] S. Kurdzialek, P. Dulian, J. Majsak, S. Chakraborty, and R. Demkowicz-Dobrzanski, Quantum metrology using quantum combs and tensor network formalism, arXiv:2403.04854 (2024).
[82] K. F. Riley, M. P. Hobson, and S. J. Bence, *Mathematical Methods for Physics and Engineering*, 3rd ed. (Cambridge University Press, Cambridge, 2006).
[83] M. Tsang, Operational meanings of a generalized conditional expectation in quantum metrology, Quantum **7**, 1162 (2023).
[84] W. Li, D. Brown, A. Vylegzhanin, Z. Shahrabifarahani, A. Raj, J. Du, and S. Nic Chormaic, Atom-light interactions using optical nanofibres—a perspective, Journal of Physics: Photonics **6**, 021002 (2024).
[85] C. Sayrin, C. Clausen, B. Albrecht, P. Schneeweiss, and A. Rauschenbeutel, Storage of fiber-guided light in a nanofiber-trapped ensemble of cold atoms, Optica **2**, 353 (2015).

Appendices

Bayesian inference with incorporated symmetries

A flexible approach to designing estimation strategies is minimising an information loss quantifier [23, 47]. Let $\mathcal{L}(\tilde{\theta},\theta)$ quantify the loss incurred by an estimate $\tilde{\theta}$ should the hypothesis θ be correct. The average loss is defined as

$$\bar{\mathcal{L}}_{\boldsymbol{y}} = \int d\boldsymbol{n}\, p(\boldsymbol{n}|\boldsymbol{y})\mathcal{L}_{\boldsymbol{y}}(\boldsymbol{n}), \tag{24}$$

where $p(\boldsymbol{n}|\boldsymbol{y}) = \int d\theta\, p(\theta)p(\boldsymbol{n}|\theta,\boldsymbol{y})$ is the evidence;

$$\mathcal{L}_{\boldsymbol{y}}(\boldsymbol{n}) = \int d\theta\, p(\theta|\boldsymbol{n},\boldsymbol{y})\mathcal{L}[\tilde{\theta}_{\boldsymbol{y}}(\boldsymbol{n}),\theta] \tag{25}$$

is the loss associated with the vector of outcomes $\boldsymbol{n}$, based on the posterior $p(\theta|\boldsymbol{n},\boldsymbol{y})$; $\boldsymbol{y}$ is a vector of control parameters; and $\tilde{\theta}_{\boldsymbol{y}}(\boldsymbol{n})$ is an estimator function, whose form we need to decide. To do so, we minimise Eq. (24) with respect to $\tilde{\theta}_{\boldsymbol{y}}(\boldsymbol{n})$ using variational calculus. This identifies the *optimal* estimator $\tilde{\vartheta}_{\boldsymbol{y}}(\boldsymbol{n})$, in the sense of leading to minimal information loss, as solution to [23, 82]

$$\int d\theta\, p(\theta|\boldsymbol{n},\boldsymbol{y})\, \partial_z \mathcal{L}(z,\theta)|_{z=\tilde{\vartheta}_{\boldsymbol{y}}(\boldsymbol{n})} = 0. \tag{26}$$

To solve Eq. (26), a loss function must be chosen. Phase estimation, for instance, deals with circular parameters and so it relies on trigonometric losses such as $\mathcal{L}(\tilde{\theta},\theta) = 4\sin^2[(\tilde{\theta}-\theta)/2]$ [65]. For parameters other than phases, the quadratic loss

$$\mathcal{L}_f(\tilde{\theta},\theta) = [f(\tilde{\theta}) - f(\theta)]^2 \tag{27}$$

can account for errors in a wide range of information-theoretic scenarios [75, 83]. Here, f, which is the symmetry function previously introduced, maps the hypothesis θ into a location (or shift) hypothesis $f(\theta)$ [52], in the sense that the aforementioned $\mathcal{L}_f$ is invariant under translations $f(\theta') = f(\theta) + c$, with arbitrary constant c [23, 33]. Those parameters for which f exists are thus referred to as location-isomorphic [52]. Use of Eq. (27) allows Eq. (26) to be straightforwardly solved, finding the estimator function (5). Under the quadratic loss criterion (27), this estimator is optimal for any parameter range, sample size, prior, and likelihood function.

To calculate the empirical information loss incurred by this estimator, $\mathcal{L}_{\boldsymbol{y},f}(\boldsymbol{n})$, given in Eq. (7), we insert Eqs. (27) and (5) into Eq. (25). By itself, $\mathcal{L}_{\boldsymbol{y},f}(\boldsymbol{n})$ does not offer a clear visualisation of the quality of our estimates, but it can be used to construct a Bayesian analogue of error bars. One approach is to note that, when $\tilde{\theta} \approx \theta$, the quadratic loss (27) can be Taylor-expanded and the square error written as $(\tilde{\theta}-\theta)^2 \approx \mathcal{L}_f(\tilde{\theta},\theta)/f'(\theta)^2$. Motivated by this, we choose the error bar (6) as a formal analogue of this idea and employ it as a symmetry-informed, dimensionally consistent standard deviation. Strictly speaking, error bars are most appropriate in the context of Gaussian-like distributions, but, for more general cases, they serve only as a heuristic and become less reliable with limited data. Nevertheless, we adopt this representation of Bayesian errors for consistency with current practice.

Finally, to derive the precision gain quantifier (9), we insert Eqs. (27) and (5) into the average loss (24). This renders the minimum average loss $\bar{\mathcal{L}}_{\boldsymbol{y},f,\min} = \int d\theta\, p(\theta) f(\theta)^2 - \mathcal{G}_{\boldsymbol{y},f}$. Since the first term is fully determined once the prior $p(\theta)$ and the symmetry function f are specified, reducing the average loss amounts to maximise the precision gain $\mathcal{G}_{\boldsymbol{y},f}$ over the control parameters $\boldsymbol{y}$.

Derivation of multiparameter ignorance priors

In this section we derive the multiparameter ignorance prior (8). This derivation provides a rationale that can be applied to other experimental platforms.

We seek the prior $p_{\rm MI}(\varphi,\theta|\nu)$ representing a meaningful notion of maximum ignorance about the unknown parameters of our platform, which may in principle be conditioned on ν. To find it, we impose two physically motivated conditions:

1. $p_{\rm MI}(\varphi',\theta'|\nu)d\varphi' d\theta' = p_{\rm MI}(\varphi,\theta|\nu)d\varphi d\theta$, which amounts to demand invariance under Eq. (14);
2. the laser frequency ν does not inform the value of the unknown parameters.

Condition 1 implies the functional equation

$$\gamma\, p_{\rm MI}\left[\gamma\varphi, \theta + \frac{\log(\gamma)}{\zeta_\nu}\middle|\nu\right] = p_{\rm MI}(\varphi,\theta|\nu). \tag{28}$$

By taking the derivative with respect to γ on both sides, we can transform it into the partial differential equation

$$\zeta_\nu\, p_{\rm MI}(x,y|\nu) + \zeta_\nu x\, \partial_x p_{\rm MI}(x,y|\nu) + \partial_y\, p_{\rm MI}(x,y|\nu) = 0. \tag{29}$$

Solving for $p_{\rm MI}(x,y|\nu)$, we find

$$p_{\rm MI}(\varphi,\theta|\nu) = \mathrm{e}^{-\zeta_\nu\theta} h(\varphi\, \mathrm{e}^{-\zeta_\nu\theta}), \tag{30}$$

where h is an arbitrary function. As such, and unlike in many single-parameter scenarios [23, 33], condition 1 does not fully determine the ignorance prior we seek.

Condition 2, which allows us to drop the dependency of the prior on the laser frequency as $p_{\rm MI}(\varphi,\theta|\nu) \to p_{\rm MI}(\varphi,\theta)$ [23], suffices to address this problem. This condition is already satisfied for $\theta = 0$, since $p_{\rm MI}(\varphi,0) = h(\varphi)$. For $\theta \neq 0$, we can enforce the aforementioned independence by setting the derivative of Eq. (30) with respect to ν to zero. Provided that $\nu \neq \nu_{\rm r}$, so that $d\zeta_\nu/d\nu \neq 0$, we find

$$h(z) + z\, h'(z) = 0, \tag{31}$$

whose solution is $h(z) \propto 1/z$. Therefore, the prior representing maximum ignorance about the unknown parameters of our cold atom platform is

$$p_{\rm MI}(\varphi,\theta) \propto \frac{1}{\varphi}. \tag{32}$$

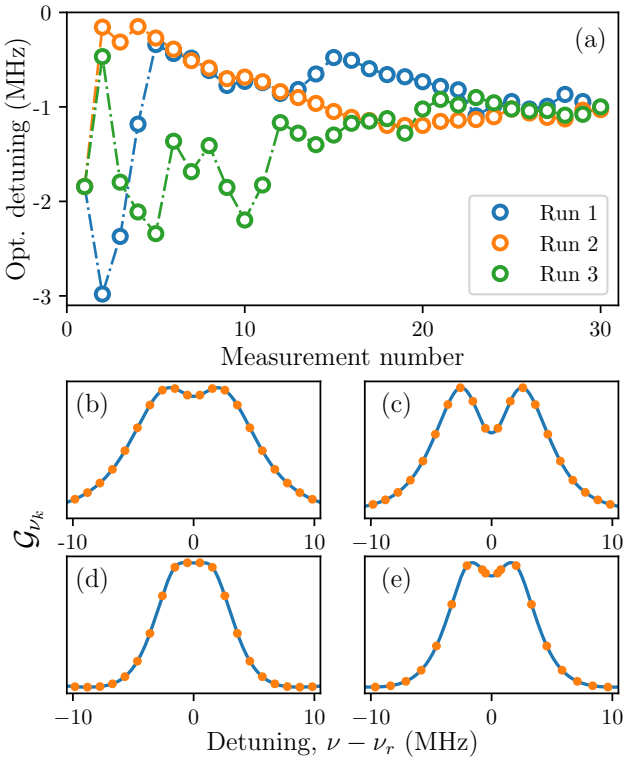

FIG. 4. (a) Example of how the optimal measurement frequency ν changes when the adaptive Bayesian feedback is applied. Three consecutive experimental runs are shown, where the sample has been prepared with an atom number of $N \sim 250$. (b-e) Examples of the gain function $\mathcal{G}_{\nu_k}$ used to find the optimal frequency [Eq. (9)]. Examples are shown for: (b) the *a priori* case, $k = 1$; (c) high current atom number estimate, $N \sim 340$ and $k = 4$; (d) low current atom number estimate, $N \sim 200$ and $k = 4$; (e) at the end of a run, $N \sim 250$ and $k = 30$. Orange dots show the calculated values, while the blue line is a cubic interpolation.

The prior (8) used for the adaptive protocol—and, for consistency, for the other two Bayesian methods—is recovered by normalising Eq. (32) over the prior range $[\varphi_{\min}, \varphi_{\max}] \times [\theta_{\min}, \theta_{\max}]$.

Experimental implementation

Atoms are loaded into the fibre-atom junction via the following procedure. A cloud of $\sim 10^8$ ^{133}Cs atoms is cooled in a magneto-optical trap and transferred into a red-detuned optical dipole trap formed using 30 W of light at 1064 nm, focussed to a waist of 7 μm. Approximately 10^5 atoms are thus captured in an elongated dipole trap that overlaps with the fibre junction, perpendicular to the fibre core. The number of atoms in this region of overlap is determined by sending a 10 μs pulse of light through the fibre, at a frequency that is close to resonance with the atomic D2 line transition, and measuring the fraction of the light that is transmitted using a single-photon counter. A detector dark count rate of $n \sim 1$ is measured independently and subtracted from both signals.

The small mode area of the intersection ($\sim 20\,\mu\text{m}^2$) together with the saturability of the atomic medium limits the applied optical power to $\sim$ pW, while expansion of the atomic cloud restricts the exposure duration. Total photon counts are thus limited to a few tens of photons per measurement step and photon-counting statistics is the limiting factor on measurement precision. Typically, averages of 30 - 100 measurement shots are recorded, each shot taking approximately 6 s. These requirements are comparable to, e.g., Refs. [71] and [84, 85].

Tab. II below expands the results for all methods in the main text.

	m	$\langle \widetilde{N}_k \rangle$	$[\mathrm{Var}(\widetilde{N}_k)]^{\frac{1}{2}}$	NSR	$\langle \Delta \widetilde{N}_k \rangle$
On resonance	16	279	26	0.89 %	27
Detuned	12	313	77	6.10 %	47
On resonance Bayesian	16	268	21	0.62 %	19
A priori Bayesian	19	238	21	0.77 %	16
Adaptive Bayesian	16	258	11	0.19 %	18

TABLE II. Overview of parameters of Fig. 2 in the main text, including: number of repeats m; mean of the individual estimates, $\langle \widetilde{N}_k \rangle$; the correspondent standard deviation, $[\mathrm{Var}(\widetilde{N}_k)]^{\frac{1}{2}}$; empirical noise-to-signal ratio; and mean of the individual error bars of each estimate $\widetilde{N}_k$ from m repeats, $\langle \Delta \widetilde{N}_k \rangle$.

To implement the adaptive Bayesian strategy, the precision gain quantifier $\mathcal{G}_{\nu_k}$ in Eq. (9) needs to be maximised and the optimal frequency calculated in real time and adapted physically sufficiently fast between the measurement shots. Optimal frequency computation was achieved at sufficient speed in Python 3 on a standard desktop computer. Standard techniques, such as caching of function outputs, control of error tolerances and interpolation were employed to reduce evaluation time. In addition, adaptive truncation of the probability distributions for θ, $\bar{n}_b$ and n (the on-resonance optical depth hypothesis, expected photon counts without atoms and photon counts recorded from a specific shot, respectively) was used. Truncation ranges were set such that, using the posterior probability distribution in light of all previous measurements, we integrated only over regions where the probability densities were greater than 1 % of their maximum values, with sums over possible n truncated at the same probability threshold, but under the worst-case assumption for θ and $\bar{n}_b$, i.e., the sum was only truncated once this threshold was passed for all possible values of θ and $\bar{n}_b$ remaining within the truncation window.

The laser frequency was adjusted via a voltage-controlled oscillator within an optical beat lock that stabilises the laser frequency. An example of how the gain function and hence optimal laser frequency changes over the course of a measurement run when the adaptive protocol is applied is shown in Fig. 4.

Measurement process and example

We now provide additional details of the measurement process and example data that help to illustrate

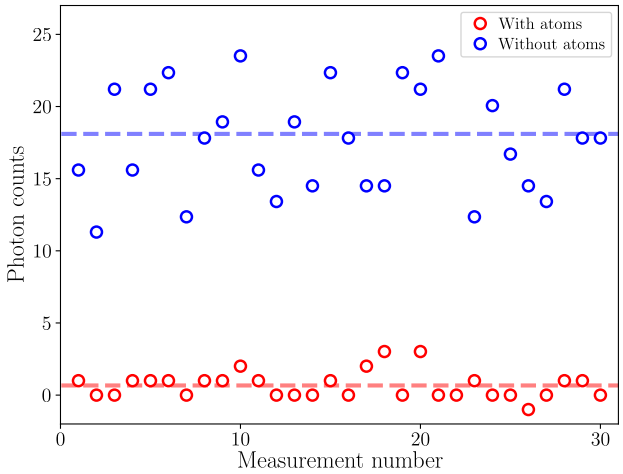

FIG. 5. Example of an experimental run consisting of 30 measurements, with atoms present in the hole (red circles) and without (blue circles), using resonant light.

the typical conditions under which the experiment operates. Fig. 5 shows an example of the data resulting from the standard measurement process, where the mean number of atoms present in the interaction region was ~ 250 and the laser light was applied at the resonant frequency ν_{r}. Thirty measurement shots were taken ($k = 30$) with atoms present (red circles), yielding photon counts $\boldsymbol{n}_a = (n_{a,1}, \ldots, n_{a,k})$, and the same number ($l = 30$) without atoms present (blue circles), yielding photon counts $\boldsymbol{n}_b = (n_{b,1}, \ldots, n_{b,l})$. Negative photon counts arise due to subtraction of a background dark count rate of ~ 1 count per shot, measured separately, from the data.

Data taken on-resonance like this can be analysed either via standard maximum likelihood estimation, or via Bayesian analysis. For the example in Fig. 5, $k = 30$ measurements are taken using resonant light with empirical means $\langle n_a \rangle = 0.7 \pm 0.2$ and $\langle n_b \rangle = 18.1 \pm 0.8$. When analysed using the standard protocol, this leads to an atom number estimate of $\tilde{N}_k = 283 \pm 22$. When analysed using the Bayesian protocol, as outlined in the main text, it gives an atom number of $\tilde{N}_k = 277 \pm 19$. Since these example data were collected with the frequency fixed on resonance, the other protocols do not apply.

Assessment of expected gain for different experimental settings

While an adaptive approach is the most robust and resource-effective strategy by construction, one may sometimes consider a preliminary estimation of the expected gain before implementation in a new experimental setting. A rigorous quantification requires a full determination of the expected information loss as a function of the number of measurements k, integrated over all possible hypotheses θ and measurement outcomes n. For a quick assessment, local estimation theory provides a practical means to evaluate whether different hypotheses θ within the *a priori* distribution $p_0(\theta)$ necessitate significantly different values of the control parameter y for optimal measurement. Specifically, it allows one to check whether applying the optimal control parameter value(s) associated with one plausible hypothesis θ_a would, according to local estimation theory, perform substantially worse under the assumption of another possible hypothesis θ_b than using the control parameter value(s) optimal to θ_b.

Exploiting complex 3D-printed surface structures for portable quantum technologies

N. Cooper,[1,*] D. Johnson,[1] B. Hopton,[1] M. Overton,[1] D. Stupple,[2] A. Bratu,[2]
E. Wilson,[2] J. Robinson,[2] L. Coles,[3] M. Papastavrou,[3] and L. Hackermueller[1]

[1] *School of Physics and Astronomy, University of Nottingham, University Park, Nottingham, NG7 2RD, UK*
[2] *Torr Scientific Ltd, Pebsham Lane, Bexhill-on-Sea, TN40 2RZ*
[3] *Metamorphic Additive Manufacturing Ltd, Riverside Chambers, Full Street Derby, DE1 3AF, UK*

(Dated: June 27, 2025)

Portable quantum technologies require robust, lightweight and compact systems. For techniques dependent upon high-vacuum environments, such as atom interferometers and atomic clocks, 3D-printed components can augment functionality by using fine-scale surface patterns to tailor in-vacuum gas propagation dynamics. These patterns can be integrated directly into additively manufactured components; they can be applied in close proximity to key experimental regions, can contribute to weight-reduction and can also enhance the stability of the geometry. We demonstrate intricate surface patterning of 3D-printed vacuum components to increase the rate at which gas particles collide with the surface. By applying a non-evaporable getter coating for use as a surface pump, we demonstrate that the patterned part of the surface pumps gas particles 3.8 times faster than an equivalent flat surface. Numerical simulations show good agreement with this result and predict larger increases in pumping rate, up to a factor of 10, for realistic surface structures. Our work has direct applications in enabling passively-pumped portable quantum technologies, but also demonstrates the broader potential of 3D-printed surface patterning to provide enhanced control over high-vacuum gas dynamics.

INTRODUCTION

A rapidly growing range of portable atom-based quantum technologies (QT) are exploiting ultra-high vacuum (UHV) conditions to enable sensing and timing applications with unprecedented performance [1–7]. A significant challenge for field-employment is to reduce the size, weight and power consumption (SWAP) of vacuum equipment, especially pumps. While much progress in the miniaturization of experimental hardware such as optics [8–12], vacuum vessels [13, 14], magnetic field generation [15], control electronics [16] and integrated atom traps [17] has been demonstrated, vacuum pumps have developed more slowly and are now a major limiting factor. Solving this problem will facilitate the deployment of portable QT for both industry and research. In addition to terrestrial QT applications, a range of exciting proposals to deploy space-based QT [18] have been put forward for Earth monitoring [19, 20], astronomical observations [21] and precision tests of fundamental physics [22–24]; low mass, power consumption and heat dissipation are essential for such deployments. Understanding and exploiting the benefits offered by 3D-printed surfaces—in particular in combination with 3D-printed chambers or hardware elements—can help overcome many of these challenges.

Passive pumping via non-evaporable getters (NEGs) [25–27] offers a promising route to unlock the full potential of portable QT [28, 29]. However, substantial increases in pumping rate and pump capacity are required, particularly for some "problem" particle species with low pumping efficiency [30].

By considering the physics of a diffuse gas interacting with a patterned surface, it is possible to engineer complex surfaces that enhance the performance of high-vacuum systems for QT; specifically, we show that additive manufacturing (AM, also known as 3D-printing) can create intricate surface structures that substantially increase the effectiveness of passive NEG pumps. Exploiting AM, these surfaces can be embedded directly within the walls of vacuum components, aiding compactness and facilitating integration into target systems. We demonstrate experimentally that appropriately structured parts of a 3D-printed, NEG-coated surface can pump 3.8 times faster than equivalent flat regions.

We present numerical simulations, validated by our experimental results, which suggest up to a 10-fold increase in pumping speed with readily-achievable surface geometries. The resulting improvement in pumping rate through complex geometries can exceed the corresponding increase in total surface area. Furthermore, this enhancement can be realized without any increase in the projected surface area—which serves as a metric for the spatial footprint of the surface pump within the vacuum system. This approach to modeling the behavior of a rarefied gas interacting with complex surface patterning thus reveals optimal structures for NEG surface pumping, while also paving the way for AM-enabled surface texture engineering to be applied more broadly to meet the needs of emerging QT in high-vacuum environments; this could include, for example, the use of optimized surface patterns to control channel conductances, particle dwell times or nozzle output distributions.

We produce AM surface structures in titanium alloy via additive manufacturing and they are confirmed as UHV compatible. Photographs of the surface structures

tested are shown in figure 2(a) and their exact geometries are described below. Previous work has not only demonstrated AM of UHV components [31], but has also shown that AM allows conformal application of intricate structures over complex 3D surfaces [8, 14]. Combined with the approach described herein, this establishes a clear pathway to integrated, structured pumping surfaces, suitable to cover a large fraction of the internal surface area of a vacuum system, within nearly any UHV component—see figure 1.

Beyond QT, passive surface pumps play a role in maintaining UHV environments within a wide range of technologies and devices, from electron microscopy [32] and mass spectrometry [33] to prototype fusion reactors [34] and particle accelerators [35]. All such applications may benefit from improved pumping performance and the technology introduced here.

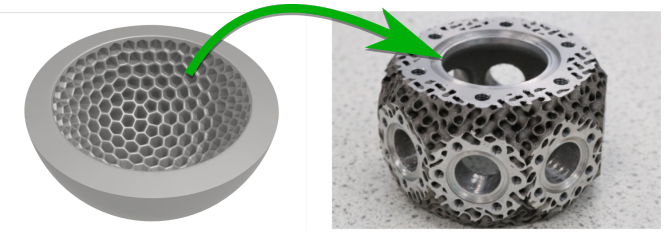

FIG. 1. Additively manufactured vacuum apparatus with internal surface structuring: the surface structures could be conformally mapped onto the interior surfaces of chambers and components. Herein it is shown that, with an NEG coating applied to such a surface, this can enable large improvements in surface pump performance. Similar methods will likely enable tailored control of gas flow within high-vacuum apparatus. The AM chamber photographed here is described fully in [14].

THEORY AND SIMULATIONS

Monte Carlo simulation methods, based on ballistic particle propagation in the long mean-free-path limit, were used to model the interactions of incident gas particles with structured surfaces. This approach allows rapid investigation of the properties of a wide range of surface structures, enabling optimization of surface patterns for specific applications. The employment of a Monte Carlo approach based on sequential simulation of individual particles makes the simulation process highly general; any surface pattern with either a repeated unit cell can be simulated. The generality of the method facilitates analysis of complex surface geometries, allowing surface patterns to be engineered to forms optimal for specific applications.

Here we target the application of optimizing pumping performance for surfaces coated with NEG material. Small-scale surface patterning increases the surface area available and thus yields a corresponding increase in total pumping capacity. However, these structures can also increase the pumping rate by factors significantly larger than the increase in surface area. The pumping rate is limited by the finite probability for an incident gas particle to be absorbed, via either binding or embedding, upon collision with the surface. This probability, P_s, depends on the particle species and getter material, but often $P_s \ll 1$ [30].

Complex surface structures can enhance pumping efficiency for any $P_s < 1$. To simulate particle interactions with complex surfaces, we consider particles incident upon both a flat and a structured surface, as depicted in figure 2(a). A particle incident upon the flat surface impacts the surface only once, yielding overall pumping probability P_s. However, a particle incident upon the structured surface, if not bound by the first collision, has a chance to re-impact the surface *before* passing back through the virtual surface marking the top-plane of the surface structure, and can undergo multiple collisions in this way. This yields an effective pumping probability $P_e > P_s$.

To enable evaluation of P_e, we define $n \geq 1$ as the number of surface collisions that would be undergone by an incident, non-sticking particle with $P_s = 0$. The effective pumping probability P_e for a real particle, with $P_s \neq 0$, is then given by

$$(P_e|n) = 1 - (1 - P_s)^n. \tag{1}$$

In the limit of $P_s \ll 1$, $(P_e|n)$ is equal to nP_s and thus $P_e = \langle n \rangle P_s$, resulting in an upper limit on the value of P_e for real particles, with $\langle n \rangle$ representing the maximum factor of improvement offered by the surface patterning.

Appropriate surface structures thus enable increased pumping rate within the same projected area. Labeling the probability that a particle with $P_s = 0$ undergoes n collisions as $P(n)$, we can combine the probability distribution for $P(n)$ with (1) to obtain P_e for any value of P_s:

$$P_e = \sum_{n=1}^{\infty} (P_e|n) P(n). \tag{2}$$

However, obtaining $P(n)$ is non-trivial and depends on the surface geometry. To do this, we developed a Monte Carlo simulation that models particle incidence upon the structured surface and particle propagation within the structured surface. Particles are sampled randomly from an isotropic background gas and propagated, in the long mean free path limit appropriate to high-vacuum environments, until they return through the top plane.

By default we model re-emission of particles that impact the surface using the standard cosine law [36], but we also consider isotropic re-emission into the available half-space and find similar results (see figure 3(b)). AM build processes create surface roughness on a scale

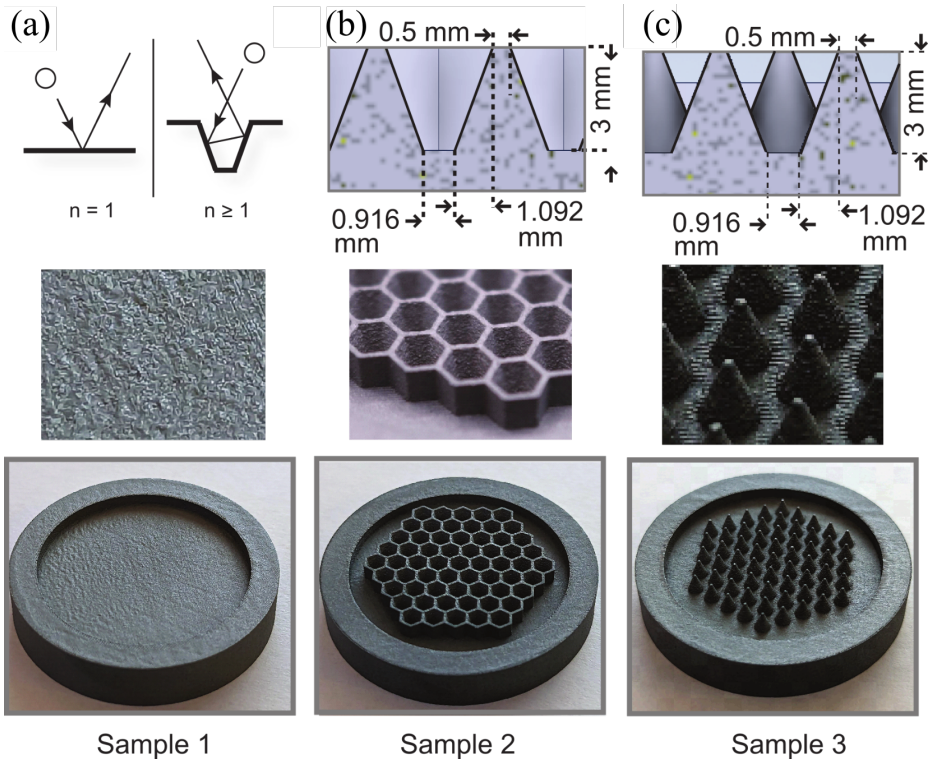

FIG. 2. The 3D-printed surface structures experimentally tested for pumping performance, including a flat reference surface. (a) Diagram illustrating how tailored surface structures enable multiple collisions; photographs of flat reference surface. (b) Array of tapered, truncated hexagonal pockets. (c) Hexagonally-packed array of conical protrusions.

that is small compared to the intentional feature size but large compared to the length scale of an atom or molecule. In the context of scattering from structured surfaces, this manifests as a partial randomization of the local surface orientation during particle emission, leading to increasingly isotropic emission. By considering the two extreme cases—a pure cosine distribution and fully isotropic emission—we establish bounds within which the actual system performance is expected to fall.

Particles are propagated in this way until they re-emerge from the top-plane of the surface. The number of collisions is then recorded and, by averaging over large numbers of particles (10^5 particles per data point for the data shown in figure 3(b)), the probability distribution $P(n)$ can be numerically approximated from the results. Equation (2) can then be applied to determine the increase in capture probability for an incident particle and thus the factor by which the pumping rate is improved, P_e/P_s.

Simulation results

Simulations were first conducted for the two surface geometries depicted in figure 2: an array of tapered, truncated, hexagonal pockets (structure/sample 2) and a hexagonally-packed array of conical protrusions (structure/sample 3). The exact, realised geometry of the printed structures is shown in the figure. In the case of structure 2, we simulate an idealised form in which the pockets are fully tessellated without interstitial flat regions on the top surface, as such regions occur in our sample only for historical and technical reasons. Figure 3(a) depicts the key parameters of the simulation, including the surface angle θ, which is measured between the inclined regions of the structured surface and the normal to the top plane of the surface. Other, more complex forms were also considered and simulated, as illustrated in figure 4.

Our simulations can predict the pumping rate improvement expected for any finite value of P_s, but we display here the theoretically optimum case where $P_s \ll 1$. This is a good approximation for many realistic situations [30] and is supported by our experimental results reported in the subsequent sections. This approximation inherently becomes a better estimate for the long-term pressure behavior of systems where NEG surface pumping is the main pumping mechanism, since species with low P_s will tend to become the dominant constituents of the residual background gas.

The simulation results for the geometries in figure 2 are shown in figure 3(b). They indicate that large in-

creases in pumping speed are plausible using fine-scale surface texturing. For example, with $\theta = 10°$ (an angle that remains plausible for sputter coating [37]), fully-tesselated hexagonal pockets are predicted to enhance pumping rate by a factor of between 5.3 and 5.9. They also show the strong dependence of the pumping efficiency on the surface angle θ. Figure 3 shows the results for either a standard cosine-law for re-emission of particles from the surface, or isotropic re-emission into the available half-space. Circles indicate experimental data points for structure 2 (upper point, blue) and upper and lower bounds for structure 3 (middle and lower points, green). The simulations assume complete coverage of an infinite surface.

Simulation results for $\theta < 10°$ indicate that the maximum pumping rate increase achievable with these structures is around a factor of 8, achieved for surface angles $\lesssim 1°$. This is relevant because certain printable materials—such as titanium—can function directly as getter materials when activated at sufficient temperature [38, 39], potentially allowing for steeper surface angles in future applications.

In addition, a broad range of patterns were simulated whose cross-sectional shape was that of a regular polygon with 3-8 sides. The relative pumping rate enhancement dependent on θ is similar for all these shapes with only minor dependence on the number of sides of the polygon and the truncation parameter ($t/h \leq 0.2$). The results are given in the Supplemental Material.

Beyond simple polygonal forms, simulations do predict improvements in performance for some more complex forms. To demonstrate the ability of the simulations to tackle complex surfaces and almost arbitrary geometries, we simulated the effect of intricate pocket shapes based upon geometry-inspired artwork ("Escher tilings"), one of which was found to yield significant improvements in pumping performance. This pocket shape is given by the height map shown in figure 4(a), and the corresponding maximum improvement in pumping rate is shown in figure 4(b). Because the form simulated involves a range of different surface angles, it cannot be compared directly with the simpler, polygonal pocket forms for any specific surface angle. However, in this case the *maximum* possible pumping rate improvement is considerably higher than that for simple polygonal pocket forms, at a factor 9/10 for isotropic/cosine-distributed scattering respectively. This motivates further investigation of complex, engineered surface patterns for enhancing QT devices and demonstrates the utility of these simulations in analyzing atom-surface interactions.

SAMPLE DESIGN AND FABRICATION

The small-scale surface structures were parametrically designed using the computational tool Grasshopper for Rhinoceros 3D, enabling a scalable approach to the creation of tessellating features for enhanced passive pumping. Although this method was initially applied to planar geometries, it was also found that the method is capable of conformally integrating with more complex, curved surfaces within larger components. This design methodology offers a versatile framework, ensuring that structured surfaces can be seamlessly applied to a wide range of geometrical configurations.

Hexagonal pockets (as recesses) and conical protrusions (as raised features) were selected for exploration due to their efficient packing arrangements, which can easily be implemented in both planar and complex surfaces. The cone structure offers advantages in terms of outgassing and avoiding trapped gas pockets. For the hexagonal pattern, the interconnected nature of the hexagonal pockets enhances the structural integrity during printing, it therefore offers advantages in terms of printability, especially on overhanging features of larger components. Truncation of the structures was employed to avoid build imperfections at sharp features; in the case of recessed features this limits the potential for the formation of trapped gas pockets and thus ensures UHV performance.

AM enables the creation of sufficiently small and numerous features for compact, integrated application within 3D-printed components (e.g. internal tubes, shields or vacuum parts). Alternatively it permits conformal application over complex internal surfaces of complete vacuum chambers (see Fig. 1) [14].

While various metals and UHV-compatible glass have been printed successfully [14, 31, 40], Ti-6Al-4V Grade 5 is particularly suitable for quantum technology applications, owing to its high strength-to-weight ratio and its relatively low electrical conductivity, which is important in mitigating eddy currents during atom-based sensing experiments. All experimental trials were therefore conducted using Ti-6Al-4V Grade 5 and fabricated through laser powder bed fusion. This method provides the necessary precision (< 100 μm resolution) and material properties to realize these designs at the required scale, demonstrating the suitability of additive manufacturing for the production of functional surfaces in advanced applications.

Following printing and post-processing, the samples (the hexagonal pocket array, hexagonally-packed cone array and a flat reference sample) were sputter coated with V-Zr-Ti NEG material. A composition range was identified from the literature [41] that could be activated at low temperatures around 200°C, which is V 40% ± 20%, Zr 40% ± 20% and Ti 20% ± 20% (all figures atomic % and all intervals absolute % points).

The sputter coating was performed with a dual DC-magnetron system with one target of V-Zr and one of Ti-V composition, with both magnetrons running simultaneously at appropriate intensity to achieve the desired

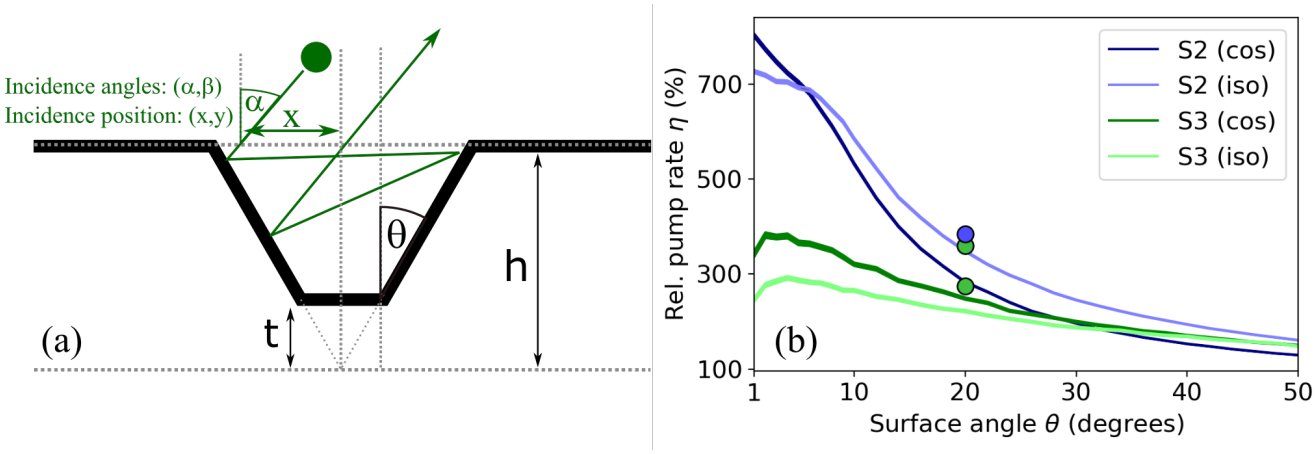

FIG. 3. (a) Sketch illustrating key simulation parameters; the truncation ratio is given by t/h. (b) The effect of surface angle θ on pumping efficiency, for $P_s \ll 1$, for both experimentally tested samples: hexagonal pocket arrays (structure 2) and hexagonally-packed cone arrays (structure 3). Simulations (blue lines for structure 2, green lines for structure 3), for either a cosine-law or isotropic scattering for re-emission of particles from the surface, employed 10^5 particles per point. The linewidth represents 1σ uncertainty bounds due to numerical error. Circles indicate experimental datapoints for structure 2 (upper point, blue) and upper and lower bounds for structure 3 (middle and lower points, green). Errors on experimental data points are ±10% (smaller than the marker size). Pockets are assumed fully-tesselated in structure 2.

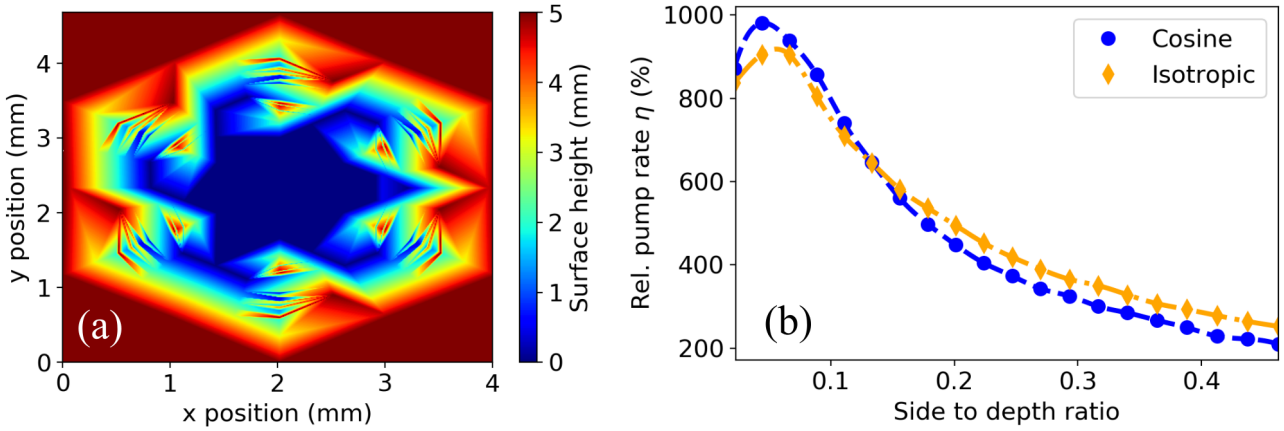

FIG. 4. (a) Height map of intricately-structured pocket with superior pumping performance. Total depth is set at 5 mm. (b) Simulated pumping rate increase resulting from the structure displayed in (a) as a function of the ratio between the side length of the hexagon and the total pocket depth, when subject to a uniform stretch along the direction of the normal to the surface top plane. Numerical errors are smaller than the marker size. Dashed lines are cubic interpolations between data. Note that although a direct comparison with simpler geometries is not possible due to the absence of a clearly defined surface angle θ, the *maximum* pumping rate improvement reaches a factor of 9/10 for isotropic/cosine-distributed re-emission, exceeding the respective factors of 7/8 found for polygonal pocket forms.

composition. The system was calibrated to achieve a thickness close to 2 μm. Energy-dispersive X-ray (EDX) analysis confirmed that the V, Zr and Ti were present in acceptable proportions and that each element was well-dispersed throughout the coating.

EXPERIMENT AND RESULTS

AM titanium samples were produced as 47.9 mm diameter discs so that, with a machined Conflat® knife edge, the samples could be incorporated into a standard DN40 rotatable vacuum flange and attached to corresponding vacuum ports. Samples with AM structures 1–3 were produced (see figure 2) and NEG coated. An additional flat 3D-printed sample was also manufactured for use as a control sample, without NEG coating. Due to initial

concerns over potential build imperfections—ultimately revealed to be unnecessary by the quality of the resulting samples—a 500 μm flat space was allowed between adjacent pockets in sample 2; this differs from the theoretical optimum version of structure 2 used for simulation purposes, but it is easy to reconcile the two by allowing for the flat area fraction of the real array (see below).

Visual examination of all samples revealed no significant build errors or distortions and the samples were sputter coated with V-Zr-Ti getter material following the procedure outlined above. The entire interior surface of the AM flange, a total area of $A_R = 1134\,\mathrm{mm}^2$, was coated with NEG material. The structured surface was restricted to an approximately hexagonal region of total area $A_S = 726\,\mathrm{mm}^2$. Within the hexagonal pocket array, flat areas of 500 μm width separate adjacent pockets, leaving a remaining $A_P = 508\,\mathrm{mm}^2$ of the area that actually corresponds to the interior of the pockets.

Initial pumping tests were performed using a residual gas analyzer (RGA, Stanford Research Systems RGA100 with electron multiplier), connected to a test rig containing a pressure gauge (Arun Microelectronics AIG17G), a turbomolecular pump that could be isolated via an all-metal bonnet valve and the sample under test. Standard baking and activation procedures were applied—see Supplemental Material. The bonnet valve was then closed and the changes in partial pressure for H_2, starting at the time of valving-off the turbomolecular pump, were recorded by the RGA and are displayed in figure 5. Partial pressures for CO_2 and N_2/CO were found to be negligible by comparison; data for this is shown in the Supplemental Material. The results clearly confirm that the surface structures of samples 2 and 3 substantially decrease the observed pressure rise, and that hydrogen, which has small P_s for the NEG coating used [30], is the dominant background gas species.

To quantitatively assess the pumping performance increase, tests were conducted using the apparatus shown in figure 6(a). The chamber, with an internal volume of 0.20 l, was evacuated and baked, and the NEG-coating activated as described in the Supplemental Material. The angle valve seen in figure 6(a) was then closed and the pressure as a function of time monitored using the pressure gauge (Pfeiffer IKR 270 Compact Cold Cathode Gauge). An initial pressure spike occurred due to the disturbance caused by physically closing the valve. This was followed in each case by a gradual decline in the measured pressure. The raw data is shown in figure 6(b). A small decrease in pressure is also observed for an uncoated sample (grey line in Fig. 6(b)). This can most likely be explained by contamination of the system with NEG material from previous tests, or by a pumping contribution from the ion gauge. By performing a differential measurement, comparing structured NEG-coated samples to a flat NEG-coated sample, the influence of all such sources of spurious pumping or outgassing can be negated.

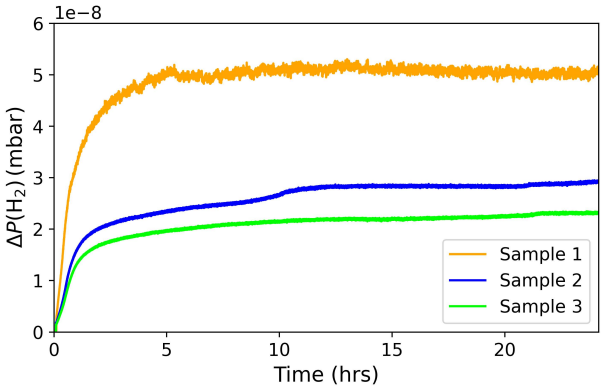

FIG. 5. Partial pressure increase for H_2 from RGA tests after closing of the gate valve for three different samples. The initial partial pressures were $P_{H_2} = 1.9 \times 10^{-10}$ mbar for sample 1, $P_{H_2} = 1.2 \times 10^{-10}$ mbar for sample 2 and $P_{H_2} = 1.0 \times 10^{-10}$ mbar for sample 3.

In order to determine a pumping rate coefficient γ (see equation (3) below) under comparable conditions, each data set was truncated at the point at which the measured pressure reached 10^{-5} mbar during the decline following the initial spike; this ensured repeatable and consistent starting conditions. The rate of pressure decrease dp/dt was then determined in each case as a function of pressure, using linear interpolation, where necessary, to enable a consistent set of data points across all samples. To eliminate the influence of spurious pumping/contamination sources, the rate of pressure drop for the uncoated sample was subtracted.

The residual rate of pressure drop attributable to the NEG coating was then plotted against pressure in each case—see figure 6(c). The net rate of pressure drop under constant pumping conditions is given by

$$\frac{dp}{dt} = C - \gamma p, \tag{3}$$

where C is a constant rate of pressure increase resulting from leaks and outgassing, p is the current pressure in the system and γ is the pumping rate coefficient, which represents the 'strength of pumping' [42, 43]. In figure 6(c), $-dp/dt$ is plotted against p and the gradient corresponds to the pumping rate coefficient γ_i for each sample. The ratio of the fitted gradients for the different samples therefore gives the ratio of the pumping rate coefficients of the different surfaces.

The results are summarized in table I and show that sample 3 offered an NEG-based pumping rate coefficient greater than that of sample 1 by a factor of $\gamma_3/\gamma_1 = 3.0 \pm 0.1$, while structure 2 increased it by a factor of $\gamma_2/\gamma_1 = 2.3 \pm 0.1$.

In order to quantify the improvement introduced by

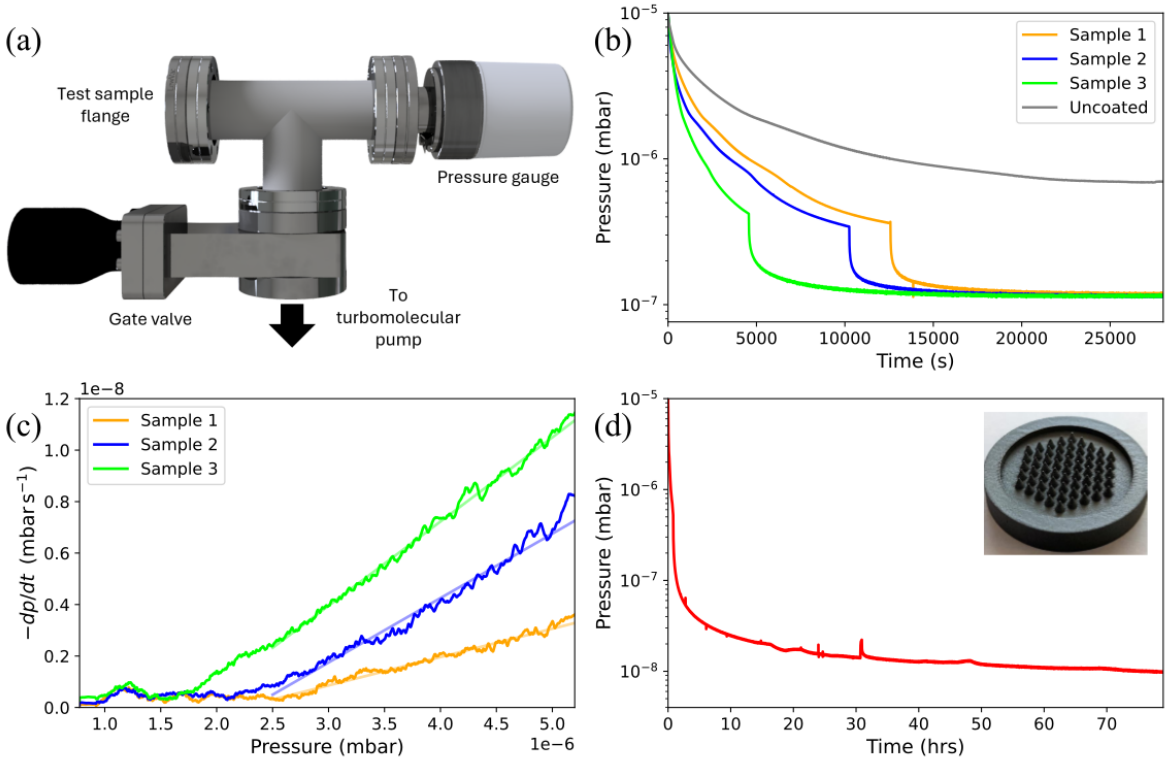

FIG. 6. (a) Setup used to test the pumping performance of the additively manufactured and NEG-coated structured samples. The closure of the gate valve allows pumping rates to be inferred from pressure data recorded by the vacuum gauge. (b) Measured chamber pressure against time for each test sample including a printed, uncoated sample. (c) Net pumping rates $-dp/dt$ for each NEG coated sample as a function of pressure. Linear fits have been used to calculate the pumping rate coefficients for the pumped volume, which are $\gamma_1 = 1.10(2) \times 10^{-3}\,\mathrm{s}^{-1}$, $\gamma_2 = 2.50(3) \times 10^{-3}\,\mathrm{s}^{-1}$ and $\gamma_3 = 3.27(3) \times 10^{-3}\,\mathrm{s}^{-1}$. (d) Long term pressure data for sample 3 (pictured) showing that pumping performance is maintained over > 3 days and a final pressure of 9.7×10^{-9} mbar is reached.

the 3D-printed geometry and to compare to simulations, a useful measure is the pumping rate coefficient per unit projected area of the structured surface, denoted $\gamma_{i,s}$ for sample i. To calculate $\gamma_{i,s}$, allowance must be made for the flat parts of the sample, as well as for the fact that particles can enter through the vertical sides of the array in structure 3 (see Supplemental Material). This second consideration results in the determination of upper and lower bounds to $\gamma_{3,s}$, depending on what assumptions are made about pumping by the sides of the array.

These calculations yield enhancements factors, $\eta_i = \gamma_{i,s}/\gamma_{1,s}$, of the pumping rate coefficient per unit area given by $\eta_2 = 3.0 \pm 0.1$ and upper and lower limits of $\eta_{3,\mathrm{max}} = 3.6 \pm 0.1$ and $\eta_{3,\mathrm{min}} = 2.7 \pm 0.1$ for structures 2 and 3 respectively. These factors of improvement show that the method offers substantial benefits even with relatively conservative surface designs. It is interesting to note that η exceeds the ratio by which the total surface area is increased.

For the purposes of comparison with simulation, η_{2^*}—the improvement in pumping rate coefficient per unit projected area for the hexagonal pockets themselves, excluding the interstitial flat regions—is found to be 3.8 ± 0.1. The results are displayed in figure 3 and include both bounds for structure 3.

Simulation and experiment show good correspondence, with the experimental data displaying a slightly larger pumping rate enhancement than expected. This can be explained by the influence of small-scale surface roughness caused by the AM build process, which was not included in the simulation.

Overall, the experimental results confirm the basic premiss of our theory and show that it provides a good description of actual pumping performance. In particular, our simulations are suitable estimates to predict the performance of AM surface structures of known geometry. This further supports the simulation results shown in figures 3 and 4, which indicate that much larger improvements in pumping performance are possible. For example, figure 4 predicts $7.5 < \eta < 8$ for pockets of the form shown with a side to depth ratio of 0.1; given that standard laser powder bed fusion manufacturing offers a

AM sample	Rel. SA	γ_i/γ_1	η_i
1	1	1	1
2	2.3	2.3(1)	3.0(1)
2^*	2.9	NA	$3.8(1)^*$
3	1.9	3.0(1)	2.7(1) (min) 3.6(1) (max)

TABLE I. Comparison of the experimentally tested samples; relative total surface area per unit projected area of the structured surface, Rel. SA, measured pumping rate coefficient enhancement γ_i/γ_1 for each sample flange and corresponding relative pumping rate coefficient per unit structured area, $\eta_i = \gamma_{i,s}/\gamma_{1,s}$, for each sample. Sample 2^*, together with the related η_2^* inferred from the measured γ_2, represents the theoretical optimum form of sample 2 without interstitial flat regions.

resolution $< 100\,\mu$m, this can be achieved with absolute pocket depths of a few mm. At the optimal side to depth ratio of 0.028, simulations predict $9.1 < \eta < 9.8$ for this pocket shape.

Figure 6(d) shows longer-term pressure data taken with sample 3, indicating that this sample can pump a volume of 0.20 l for > 70 hours of operation to a pressure $< 10^{-8}$ mbar. As no degradation of performance was observed within the window of observation, this data provides only a lower limit on pump longevity, which may in fact be far longer. In general, the increase in total surface area when applying a NEG coating to a structured surface, rather than a flat one, increases the total pump capacity and thus extends operational longevity.

CONCLUSIONS

We have demonstrated that additively manufactured, complex surface structures can be employed to manipulate gases in ultra-high vacuum chambers. In particular, we have shown that structures 3D-printed in ^{64}Ti are UHV compatible and enable a tailored enhancement in the number of surface scattering events induced by the structured elements. Outgassing tests (see Supplemental Material) did not measure any adverse effects of the structure or build process on vacuum performance. Applications include controlling particle propagation pathways and dwell times in UHV devices, surface pumping with high efficiency. Experimentally, we have shown a 3.8-fold increase in pumping rate per unit area by adding a purpose-designed surface texture to an NEG-coated surface. We have developed and presented Monte Carlo simulations that enable the prediction of the performance of complex AM surface geometries. Comparing theory and experiment, we find good agreement, while the experimental results slightly exceed the predictions of Monte Carlo simulations. These simulations reveal a range of interesting surface structures that offer even greater improvements in pumping rate, predicting up to a 10-fold increase in pumping rate for some more complex forms.

Promising directions for further development include adapting the technique to provide optimum performance from components that control gas flow via impact with a surface, such as collimation nozzles [44, 45] and cold plates [46]. Conformal application over a large fraction of a vacuum system's interior surface, coupled to direct activation of the titanium build material as an NEG pump, is also a very interesting avenue of future research, with the potential to enable extreme pumping performance and underpin portable UHV and even XHV systems.

The capabilities provided by such engineered AM surface structures are relevant to numerous technologies and devices across QT and beyond. QT applications will include next-generation portable gravimeters [47, 48], cold-atom magnetometers [49, 50], atomic clocks [51, 52] and quantum accelerometers [53]. Application areas outside QT include mass spectrometers [33], particle accelerators [54] and electron microscopes, especially where the use of methods such as environmental transmission electron microscopy or the study of novel materials imposes additional pumping and gas-flow requirements [55, 56]. The passive nature of such surfaces and the ease with which AM can shape components for integration into compact, lightweight systems makes this approach ideally suited to field-deployable [57] or space-based [58] devices.

ACKNOWLEDGEMENTS

This work has been supported by the UK Engineering and Physical Sciences Research council (EPSRC) grants EP/T001046/1, EP/R024111/1, EP/Y005139/1, EP/Z533166/1, the Innovate UK projects 10143484 (QAPTCHA) and 10031462 (QTEAM), the grant 62420 from the John Templeton Foundation, and the EPSRC Impact Acceleration Account EP/X525765/1.

* Contact author: nathan.cooper@nottingham.ac.uk

[1] V. Menoret, P. Vermeulen, N. Le Moigne, S. Bonvalot, P. Bouyer, A. Landragin, and B. Desruelle, Gravity measurements below 10^{-9} g with a transportable absolute quantum gravimeter, Sci. Rep. **8**, 12300 (2018).

[2] Y. Bidel, O. Carraz, R. Charriere, M. Cadoret, N. Zahzam, and A. Bresson, Compact cold atom gravimeter for field applications, Appl. Phys. Lett. **102**, 144107 (2013).

[3] M. Limes, E. Foley, T. Kornack, S. Caliga, S. McBride, A. Braun, W. Lee, V. Lucivero, and M. Romalis, Portable magnetometry for detection of biomagnetism in ambient environments, Phys. Rev. Appl. **14**, 011002 (2020).

[4] X.-B. Deng, Y.-Y. Xu, X.-C. Duan, and Z.-K. Hu, Precisely mapping the absolute magnetic field in vacuum by an optical ramsey atom interferometer, Phys. Rev. Appl. **15**, 054062 (2021).

[5] Y. B. Kale, A. Singh, M. Gellesch, J. M. Jones, D. Morris, M. Aldous, K. Bongs, and Y. Singh, Field deployable atomics package for an optical lattice clock, Quantum Science and Technology **7**, 045004 (2022).

[6] K. W. Martin, G. Phelps, N. D. Lemke, M. S. Bigelow, B. Stuhl, M. Wojcik, M. Holt, I. Coddington, M. W. Bishop, and J. H. Burke, Compact optical atomic clock based on a two-photon transition in rubidium, Phys. Rev. Appl. **9**, 014019 (2018).

[7] B. Tennstedt, N. Weddig, and S. Schön, Improved inertial navigation with cold atom interferometry, Gyroscopy and Navigation **12**, 294 (2021).

[8] S. Madkhaly, L. Coles, C. Morley, C. Colquhoun, M. Fromhold, N. Cooper, and L. Hackermüller, Performance-optimized components for quantum technologies via additive manufacturing, PRX Quantum **2**, 030326 (2021).

[9] S. Madkhaly, N. Cooper, L. Coles, and L. Hackermüller, High-performance, additively-manufactured atomic spectroscopy apparatus for portable quantum technologies, Optics Express **30**, 25753 (2022).

[10] A. Strangfeld, B. Wiegand, J. Kluge, M. Schoch, and M. Krutzik, Compact plug and play optical frequency reference device based on doppler-free spectroscopy of rubidium vapor, Optics Express **30**, 12039 (2022).

[11] S. Dyer, K. Gallacher, U. Hawley, A. Bregazzi, P. Griffin, A. Arnold, D. Paul, E. Riis, and J. McGilligan, Chip-scale packages for a tunable wavelength reference and laser cooling platform, Phys. Rev. Appl. **19**, 044015 (2023).

[12] F. Wang, N. Cooper, Y. He, B. Hopton, D. Johnson, P. Zhao, C. J. Tuck, R. Hague, T. M. Fromhold, R. D. Wildman, L. Turyanska, and L. Hackermüller, Additive manufacturing of functionalised atomic vapour cells for next-generation quantum technologies, Quantum Science and Technology **10**, 015019 (2024).

[13] K. Weng, B. Wu, J. Lin, Y. Zhou, B. Cheng, and Q. Lin, Compact magneto-optical trap with a quartz vacuum chamber for miniature gravimeters, J. Opt. Soc. Am. B **37**, 1637 (2020).

[14] N. Cooper, L. Coles, S. Everton, I. Maskery, R. Campion, S. Madkhaly, C. Morley, J. O'Shea, W. Evans, R. Saint, P. Krüger, F. Oručević, C. Tuck, R. Wildman, T. Fromhold, and L. Hackermüller, Additively manufactured ultra-high vacuum chamber for portable quantum technologies, Additive Manufacturing **40**, 101898 (2021).

[15] R. Saint, W. Evans, Y. Zhou, T. Barrett, T. M. Fromhold, E. Saleh, I. Maskery, C. Tuck, R. Wildman, F. Oručević, and P. Krüger, 3D-printed components for quantum devices, Scientific Reports **8**, 8368 (2018).

[16] P. Ruksasakchai, C. E. Cowdell, L. Sanchez, M. Weyland, and M. F. Andersen, A microcontroller based self-locking laser system, Review of Scientific Instruments **93**, 043008 (2022).

[17] L. Chen, C.-J. Huang, X.-B. Xu, Y.-C. Zhang, D.-Q. Ma, Z.-T. Lu, Z.-B. Wang, G.-J. Chen, J.-Z. Zhang, H. X. Tang, C.-H. Dong, W. Liu, G.-Y. Xiang, G.-C. Guo, and C.-L. Zou, Planar-integrated magneto-optical trap, Phys. Rev. Appl. **17**, 034031 (2022).

[18] A. Belenchia, M. Carlesso, Ömer Bayraktar, D. Dequal, I. Derkach, G. Gasbarri, W. Herr, Y. L. Li, M. Rademacher, J. Sidhu, D. K. Oi, S. T. Seidel, R. Kaltenbaek, C. Marquardt, H. Ulbricht, V. C. Usenko, L. Wörner, A. Xuereb, M. Paternostro, and A. Bassi, Quantum physics in space, Physics Reports **951**, 1 (2022), quantum Physics in Space.

[19] F. Migliaccio, M. Reguzzoni, K. Batsukh, G. M. Tino, G. Rosi, F. Sorrentino, C. Braitenberg, T. Pivetta, D. F. Barbolla, and S. Zoffoli, MOCASS: A satellite mission concept using cold atom interferometry for measuring the earth gravity field, Surveys in Geophysics **40**, 1029 (2019).

[20] R. Haagmans, C. Siemes, L. Massotti, O. Carraz, and P. Silvestrin, ESA's next-generation gravity mission concepts, Rendiconti Lincei. Scienze Fisiche e Naturali **31**, 15 (2020).

[21] F. Müller, O. Carraz, P. Visser, and O. Witasse, Cold atom gravimetry for planetary missions, Planetary and Space Science **194**, 105110 (2020).

[22] M. He, X. Chen, J. Fang, Q. Chen, H. Sun, Y. Wang, J. Zhong, L. Zhou, C. He, J. Li, D. Zhang, G. Ge, W. Wang, Y. Zhou, X. Li, X. Zhang, L. Qin, Z. Chen, R. Xu, Y. Wang, Z. Xiong, J. Jiang, Z. Cai, K. Li, G. Zheng, W. Peng, J. Wang, and M. Zhan, The space cold atom interferometer for testing the equivalence principle in the China space station, npj Microgravity **9**, 58 (2023).

[23] F. Sorrentino, K. Bongs, P. Bouyer, L. Cacciapuoti, M. de Angelis, H. Dittus, W. Ertmer, A. Giorgini, J. Hartwig, M. Hauth, S. Herrmann, M. Inguscio, E. Kajari, T. T. Könemann, C. Lämmerzahl, A. Landragin, G. Modugno, F. Pereira dos Santos, A. Peters, M. Prevedelli, E. M. Rasel, W. P. Schleich, M. Schmidt, A. Senger, K. Sengstock, G. Stern, G. M. Tino, and R. Walser, A compact atom interferometer for future space missions, Microgravity Science and Technology **22**, 551 (2010).

[24] L. Badurina, E. Bentine, D. Blas, K. Bongs, D. Bortoletto, T. Bowcock, K. Bridges, W. Bowden, O. Buchmueller, C. Burrage, J. Coleman, G. Elertas, J. Ellis, C. Foot, V. Gibson, M. Haehnelt, T. Harte, S. Hedges, R. Hobson, M. Holynski, T. Jones, M. Langlois, S. Lellouch, M. Lewicki, R. Maiolino, P. Majewski, S. Malik, J. March-Russell, C. McCabe, D. Newbold, B. Sauer, U. Schneider, I. Shipsey, Y. Singh, M. Uchida, T. Valenzuela, M. v. d. Grinten, V. Vaskonen, J. Vossebeld, D. Weatherill, and I. Wilmut, AION: an atom interferometer observatory and network, Journal of Cosmology and Astroparticle Physics **2020** (05), 011 (2020).

[25] E. Sartori, M. Siragusa, P. Sonato, F. Siviero, M. Mura, E. Maccallini, A. Ferrara, P. Manini, S. Hanke, and C. Day, Development of non evaporable getter pumps for large hydrogen throughput and capacity in high vacuum regimes, Vacuum **214**, 112198 (2023).

[26] Y. Li, D. Hess, R. Kersevan, and N. Mistry, Design and pumping characteristics of a compact titanium–vanadium non-evaporable getter pump, Journal of Vacuum Science & Technology A **16**, 1139 (1998).

[27] H. Kodama, S. Ohno, M. Tanaka, M. Tanaka, K. K. Okudaira, K. Mase, and T. Kikuchi, Low-cost, high-performance nonevaporable getter pumps using nonevaporable getter pills, J. Vac. Sci. Technol. A **34**, 051601 (2016).

[28] R. Boudot, J. P. McGilligan, K. R. Moore, V. Maurice, G. D. Martinez, A. Hansen, E. de Clercq, and J. Kitching, Enhanced observation time of magneto-optical traps using micro-machined non-evaporable getter pumps, Scientific Reports **10**, 16590 (2020).

[29] B. J. Little, G. W. Hoth, J. Christensen, C. Walker, D. J. De Smet, G. W. Biedermann, J. Lee, and P. D. Schwindt, A passively pumped vacuum package sustaining cold atoms for more than 200 days, AVS Quantum Sci. **3**, 035001 (2021).

[30] O. Seify, O. B. Malyshev, R. Širvinskaitė, R. Valizadeh, and A. Hannah, Optimisation of the RGA location in the evaluation of NEG coating pumping properties, Vacuum **195**, 207 (2019).

[31] J. Vovrosh, G. Voulazeris, P. G. Petrov, J. Zou, Y. Gaber, L. Benn, D. Woolger, M. M. Attallah, V. Boyer, K. Bongs, and M. Holynski, Additive manufacturing of magnetic shielding and ultra-high vacuum flange for cold atom sensors, Scientific Reports **8**, 2023 (2018).

[32] Q. Zhang, M. Tanaka, M. Takeguchi, and K. Furuya, Analytical UHV transmission electron microscopy studies of electronic structure changes between as-deposited Mn and Mn silicide on Si(111) surface, Surface Science **507**, 453 (2002).

[33] W. Zhang, X. Xia, Y. Zhang, T. Peng, and Q. Yang, A novel sample preparation method for ultra-high vacuum (UHV) secondary ion mass spectrometry (SIMS) analysis, J. Anal. At. Spectrom. **33**, 1559 (2018).

[34] E. Sartori, M. Siragusa, G. Berton, C. Cavallini, S. Dal Bello, M. Fadone, L. Grando, D. Marcuzzi, D. Rizzetto, G. Serianni, P. Sonato, M. Zaupa, F. Dinh, A. Ferrara, E. Maccallini, M. Mura, F. Siviero, and V. Toigo, Design of a large nonevaporable getter pump for the full size ITER beam source prototype, Journal of Vacuum Science & Technology B **41**, 034202 (2023).

[35] C. Benvenuti, A new pumping approach for the large electron positron collider (LEP), Nuclear Instruments and Methods in Physics Research **205**, 391 (1983).

[36] J. Greenwood, The correct and incorrect generation of a cosine distribution of scattered particles for Monte-Carlo modelling of vacuum systems, Vacuum **67**, 217 (2002).

[37] W. Zhou, J. Brock, M. Khan, and K. Eid, Oblique angle deposition-induced anisotropy in Co_2FeAl films, Journal of Magnetism and Magnetic Materials **456**, 353 (2018).

[38] J. Kamiya, K. Takano, K. Wada, and H. Yuza, Some methods of making titanium vacuum chamber act as getter pump for UHV/XHV, Joint accelerator conferences website (JACoW) **IPAC2021**, WEPAB337 (2021).

[39] J. Kamiya, K. Takano, H. Yuza, and K. Wada, Evaluation of titanium vacuum chamber as getter pump, e-Journal of Surface Science and Nanotechnology **20**, 107 (2022).

[40] F. Wang, N. Cooper, D. Johnson, B. Hopton, T. Fromhold, R. Hague, A. Murray, R. McMullen, L. Turyanska, and L. Hackermueller, Additive manufacturing for advanced quantum technologies, arxiv:2503.11570 (2025).

[41] A. Prodromides, C. Scheuerlein, and M. Taborelli, Lowering the activation temperature of TiZrV non-evaporable getter films, Vacuum **60**, 35 (2001).

[42] L. Detian, C. Yongjun, F. Yan, X. Zhenhua, and Z. Lan, Very low gas flow measurements for UHV/XHV and leak calibration, Metrologia **50**, 15 (2012).

[43] L. Peksa, T. Gronych, M. Jeřáb, P. Řepa, M. Vičar, Z. Krajíček, D. Pražák, F. Staněk, and J. Tesař, Dynamic extension – a new principle for a vacuum primary standard for UHV-range (106–1010 Pa) calibration of hot cathode ionisation gauges, Vacuum **85**, 1047 (2011).

[44] N. Cooper, E. Da Ros, J. Nute, D. Baldolini, P. Jouve, L. Hackermüller, and M. Langer, Collimated dual species oven source and its characterisation via spatially resolved fluorescence spectroscopy, Journal of Physics D: Applied Physics **51**, 105602 (2018).

[45] R. Hahn, T. Battard, O. Boucher, Y. J. Picard, H. Lignier, D. Comparat, N.-A. Keriel, C. Lopez, E. Oswald, M. Reveillard, and M. Viteau, Comparative analysis of recirculating and collimating cesium ovens, Review of Scientific Instruments **93**, 043302 (2022).

[46] E. W. Streed, A. P. Chikkatur, T. L. Gustavson, M. Boyd, Y. Torii, D. Schneble, G. K. Campbell, D. E. Pritchard, and W. Ketterle, Large atom number Bose-Einstein condensate machines, Review of Scientific Instruments **77**, 023106 (2006).

[47] C.-Y. Li, J.-B. Long, M.-Q. Huang, B. Chen, Y.-M. Yang, X. Jiang, C.-F. Xiang, Z.-L. Ma, D.-Q. He, L.-K. Chen, and S. Chen, Continuous gravity measurement with a portable atom gravimeter, Phys. Rev. A **108**, 032811 (2023).

[48] T. Zhang, L.-L. Chen, Y.-B. Shu, W.-J. Xu, Y. Cheng, Q. Luo, Z.-K. Hu, and M.-K. Zhou, Ultrahigh-sensitivity Bragg atom gravimeter and its application in testing Lorentz violation, Phys. Rev. Appl. **20**, 014067 (2023).

[49] A. Fregosi, C. Gabbanini, S. Gozzini, L. Lenci, C. Marinelli, and A. Fioretti, Magnetic induction imaging with a cold-atom radio frequency magnetometer, Applied Physics Letters **117**, 144102 (2020).

[50] H. A. Taylor, C. C. Bounds, A. Tritt, and L. Turner, Unambiguous measurement in an unshielded microscale magnetometer with sensitivity below 1 $\mathrm{pT}/\sqrt{\mathrm{Hz}}$, Phys. Rev. Appl. **22**, 024047 (2024).

[51] J. Grotti, S. Koller, S. Vogt, S. Häfner, U. Sterr, C. Lisdat, H. Denker, C. Voigt, L. Timmen, A. Rolland, F. N. Baynes, H. S. Margolis, M. Zampaolo, P. Thoumany, M. Pizzocaro, B. Rauf, F. Bregolin, A. Tampellini, P. Barbieri, M. Zucco, G. A. Costanzo, C. Clivati, F. Levi, and D. Calonico, Geodesy and metrology with a transportable optical clock, Nature Physics **14**, 437 (2018).

[52] Q. Hao, S. Yang, J. Ruan, P. Yun, and S. Zhang, Integrated pulsed optically pumped Rb atomic clock with frequency stability of 10^{-15}, Phys. Rev. Appl. **21**, 024003 (2024).

[53] F. Narducci, A. Black, and J. Burke, Advances toward fieldable atom interferometers, Advances in Physics: X **7** (2022).

[54] H. Hähnel and U. Ratzinger, First 3D printed IH-type linac structure—proof-of-concept for additive manufacturing of linac RF cavities, Instruments **6** (2022).

[55] R. Yang, L. Mei, Y. Fan, Q. Zhang, H.-G. Liao, J. Yang, J. Li, and Z. Zeng, Fabrication of liquid cell for in situ transmission electron microscopy of electrochemical processes, Nature Protocols **18**, 555 (2023).

[56] Y. Lin, M. Zhou, X. Tai, H. Li, X. Han, and J. Yu, Analytical transmission electron microscopy for emerging advanced materials, Matter **4**, 2309 (2021).

[57] S. E. Crawford, R. A. Shugayev, H. P. Paudel, P. Lu, M. Syamlal, P. R. Ohodnicki, B. Chorpening, R. Gentry,

and Y. Duan, Quantum sensing for energy applications: Review and perspective, Advanced Quantum Technologies **4**, 2100049 (2021).

[58] R. Kaltenbaek, A. Acin, L. Bacsardi, P. Bianco, P. Bouyer, E. Diamanti, C. Marquardt, Y. Omar, V. Pruneri, E. Rasel, B. Sang, S. Seidel, H. Ulbricht, R. Ursin, P. Villoresi, M. van den Bossche, W. von Klitzing, H. Zbinden, M. Paternostro, and A. Bassi, Quantum technologies in space, Experimental Astronomy **51**, 1677 (2021).

Appendix B

Optics layout

This chapter contains diagrams of most of the optics used in this experiment, showing the beam paths of each different laser. Optics around the lasers are used for frequency stabilisation and splitting and directing laser power into various optical fibres. These fibres are used to transport the prepared laser power to the main experimental chamber.

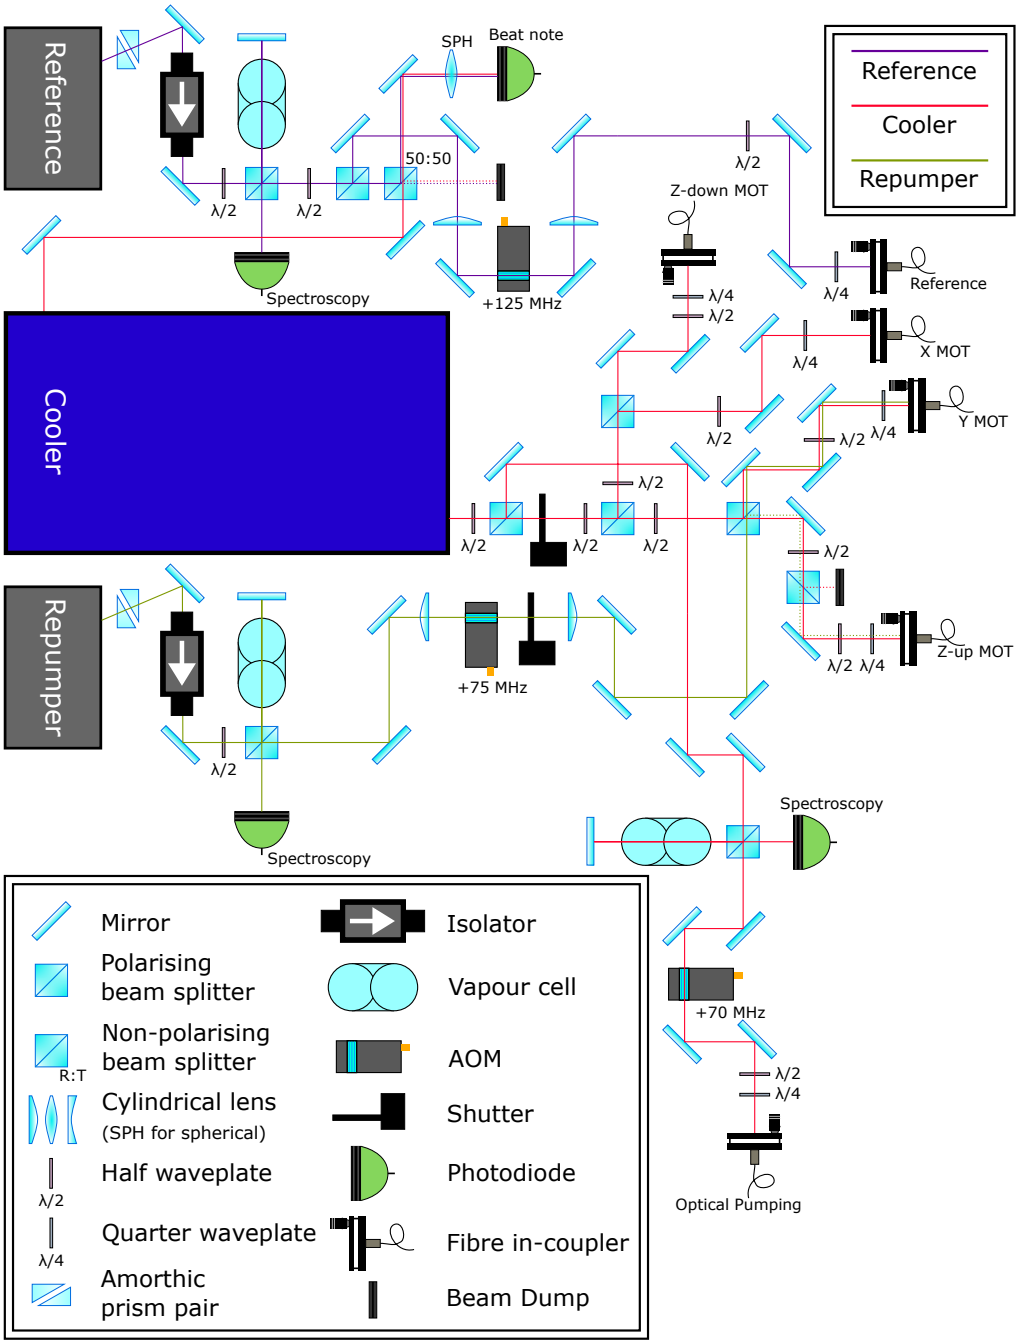

Figure B.1: Optical layout of the main optics table.

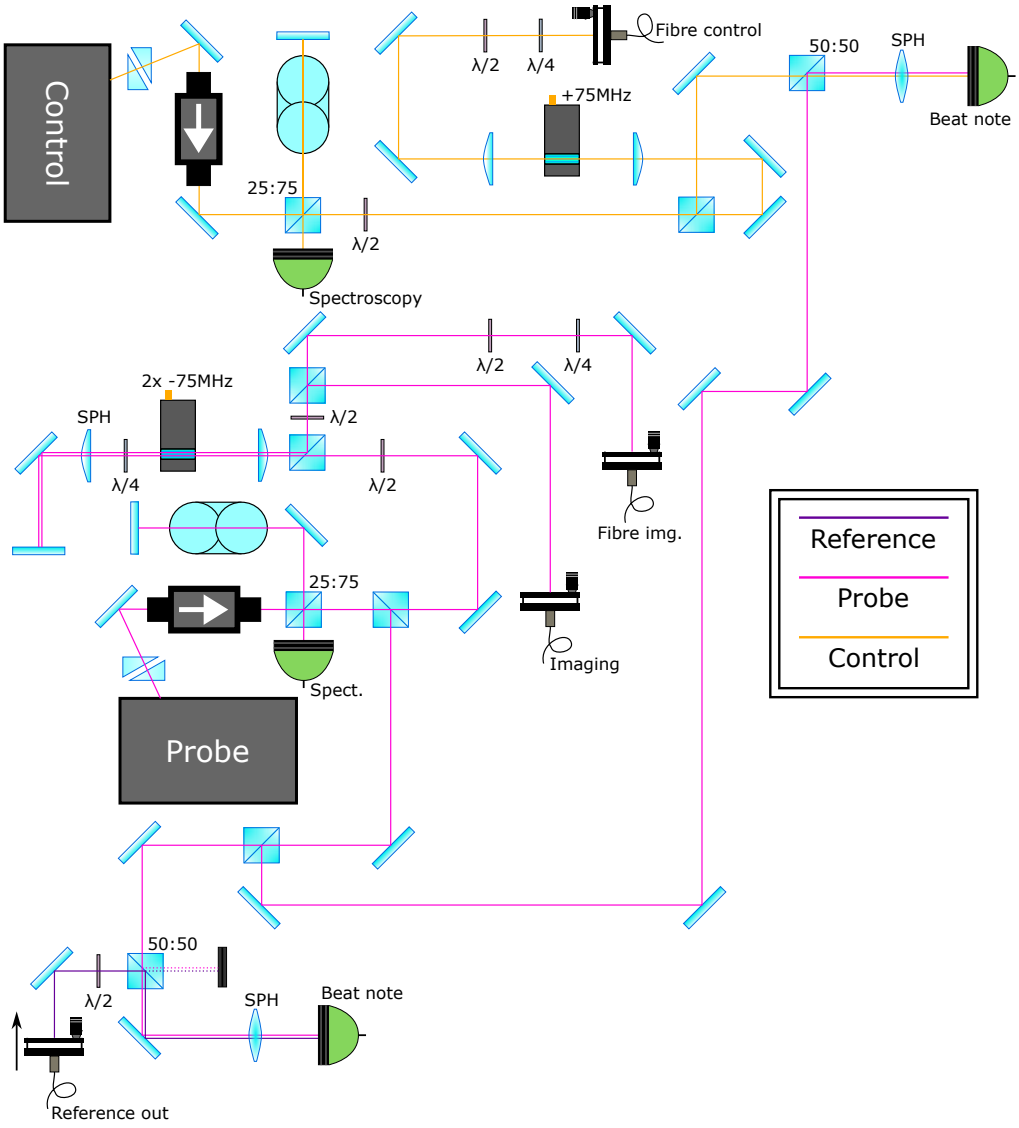

Figure B.2: Optical layout of the probe and control lasers.

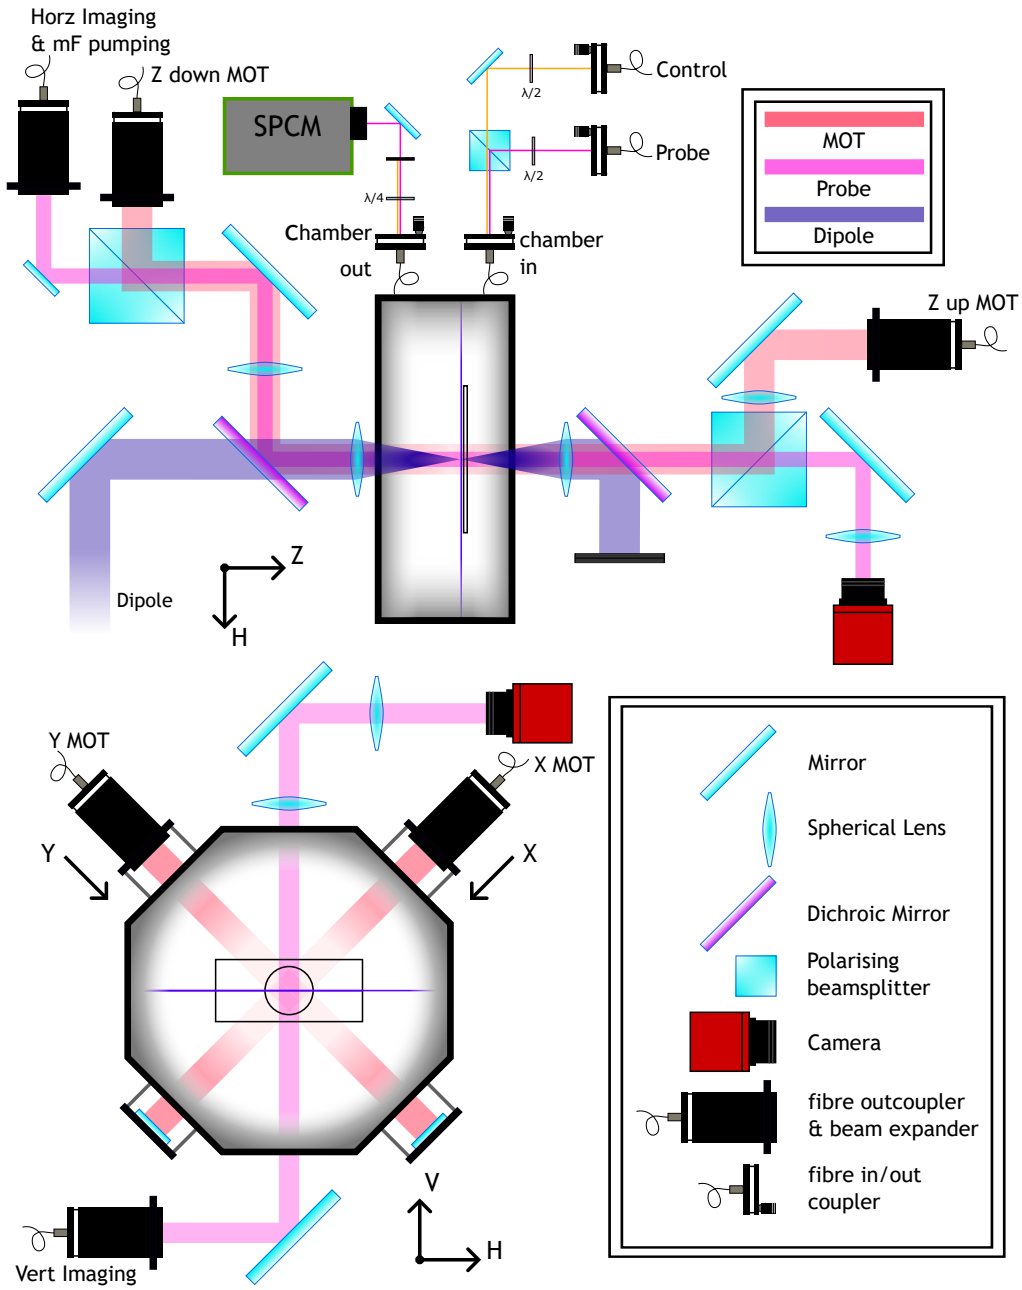

Figure B.3: Optical layout of all optics around the vacuum chamber. Both horizontal (bottom) and birds eye (top) views are shown. The grey rectangle/octagon shows the location of the chamber. The blue line running through the chamber represents the fibre, and the black rectangle represents the glass slide the fibre is mounted on.. Axis labels are shown at the bottom of both diagrams demonstrating the axes of the 3 offset coils (Z, H and V), the X and Y MOT beam axes are also shown.

Appendix C

Phase lock circuits

This chapter contains a circuit diagram of part of the optical phase lock I built for the experiment. The circuits shown here take the outputs from the phase frequency detector (PFD) and turn them into the feedback given to the laser. The lower circuit is the differential amplifier which subtracts the 2 PFD outputs from each other. The top circuit shows the loop filters which take the output of this differential amplifier and provide the locking signal for the 2 feedback paths.

The components for the amplification and down-conversion of the 9.2 GHz beat-note are not shown.

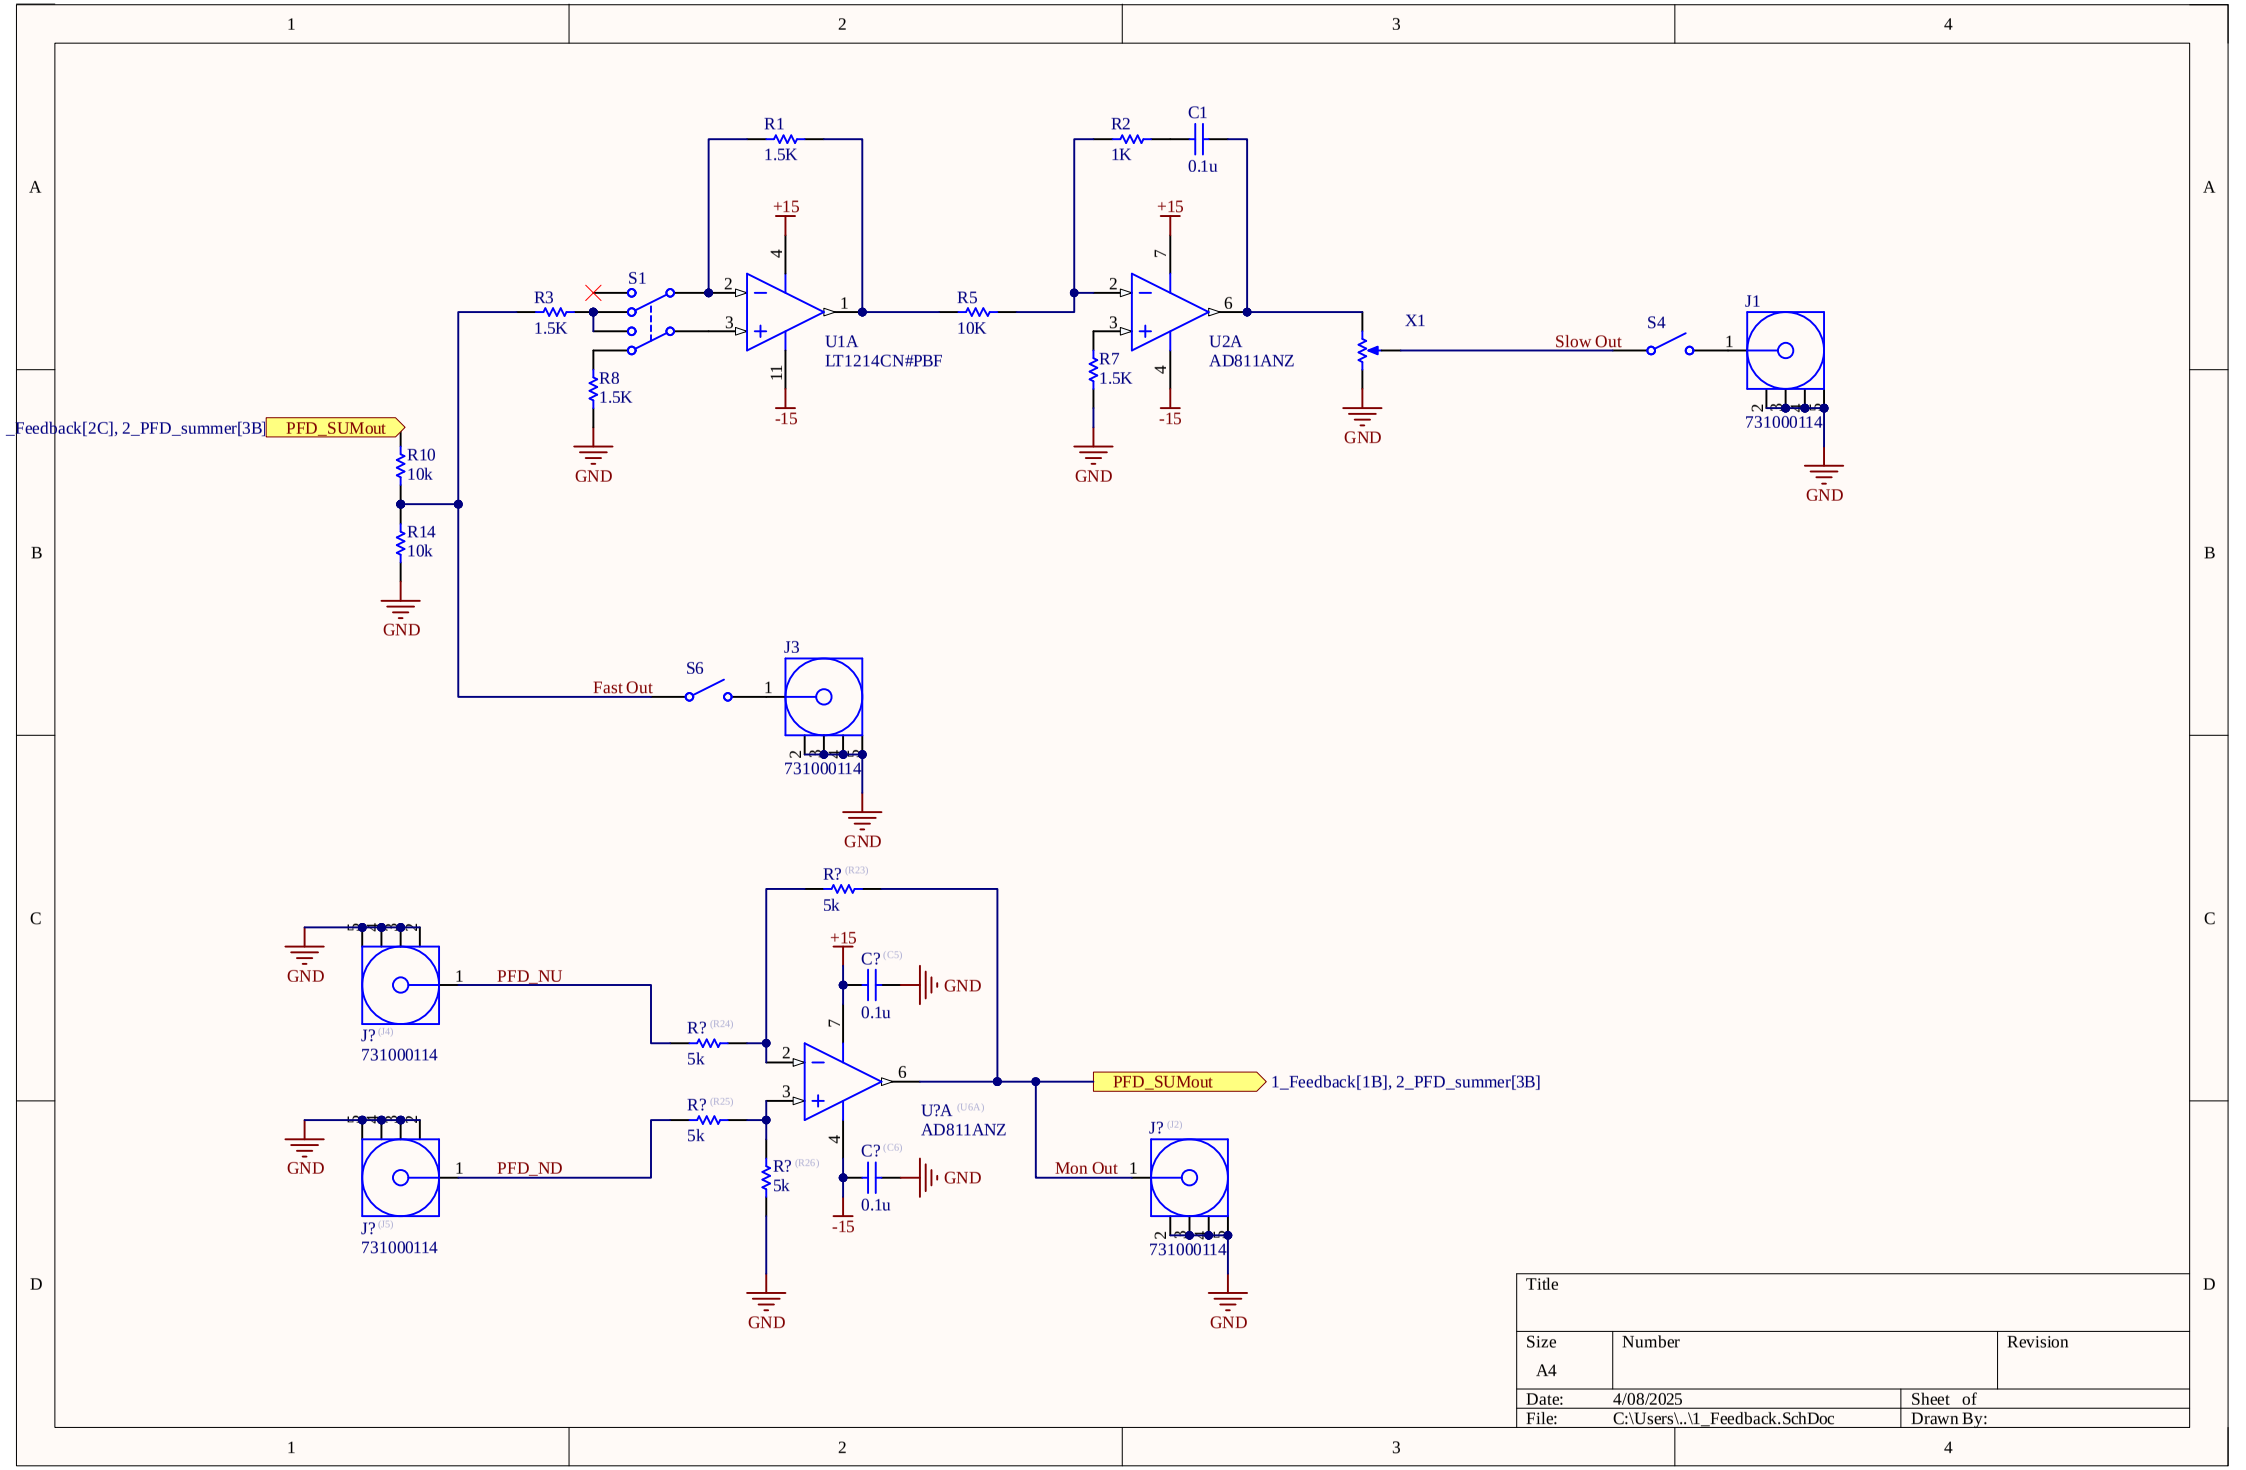

Title		
Size A4	Number	Revision
Date: 4/08/2025		Sheet of
File: C:\Users\..\1_Feedback.SchDoc		Drawn By:

Appendix D

Adaptive Bayesian code

This chapters contains the code used to evaluate the optimal frequency using the adaptive strategy outlined in Chap. 6. The code calculates the probability density function (PDF) and gains for 10 frequencies and uses cubic interpolation to find the frequency that maximises this gain. Caching of function outputs and adaptive truncation of the probability distributions are used to speed up the code.

```
import numpy as np
from scipy.integrate import nquad
from scipy.integrate import quad
from scipy.special import factorial
from scipy.interpolate import interp1d
import matplotlib.pyplot as plt
from functools import lru_cache
from tqdm import tqdm

'''
Constants Definition
'''

GAMMA = 5.234 #MHz
F_R = 75.9 #MHz
#Integration limits
THETA1_MIN = THETA1_MIN_i = 0
THETA1_MAX = THETA1_MAX_i = 8
THETA2_MIN = THETA2_MIN_i = 5
THETA2_MAX = THETA2_MAX_i = 20

options={'epsrel':1e-1,'limit':100}
# for double integration intervals, standard 50 was raising
# a warning
step = 0
max_steps = 30
f_best = []
n_a = [] # photon counts without atoms
```

```
n_b = [] # photon counts with atoms
initial_freq = 65
final_freq = F_R+0.5
freq_steps = 10

cutoff_threshold=0.01
ptnum=20

VCO_ref  = [0.00,0.25,0.50,0.75,1.00,1.10,1.20,1.25,1.30,
            1.40,1.50,1.60,1.70,1.75,1.80,1.90,2.00,2.05,
            2.10,2.15,2.20,2.25,2.30,2.35,2.40,2.45,2.50,
            2.75,3.00,3.25,3.50,3.75,4.00,4.25,4.50,4.75,
            5.00]
freq_ref = [56.1,57.4,59.0,64.7,64.7,64.7,64.7,64.9,65.1,
            69.4,71.1,71.6,71.7,71.8,71.9,72.2,72.8,73.1,
            73.5,74.0,74.4,74.9,75.4,76.0,76.4,76.8,77.2,
            78.8,80.4,82.1,83.8,85.4,87.1,88.8,90.4,92.0,
            93.6]
'''
Defining functions
'''
@lru_cache(maxsize=None)
def zeta(f):
    return GAMMA**2/(GAMMA**2 +4*(f - F_R)**2)

@lru_cache(maxsize=None)
def T(theta1, f):
    return np.exp(-zeta(f) * theta1)
```

```
57
58  @lru_cache(maxsize=None)
59  def p_n_given_theta_f(n, f, theta1, theta2):
60      p_n = (np.exp(-theta2 * T(theta1, f)) *
61             (theta2 * T(theta1, f))**n)/ factorial(n)
62      return p_n
63
64  @lru_cache(maxsize=None)
65  def p_n_given_f(n, f):
66      def integrand(theta1, theta2):
67          return p_theta(step, theta1, theta2) *
68                 p_n_given_theta_f(n, f, theta1, theta2)
69      return nquad(integrand, [[THETA1_MIN, THETA1_MAX],
70                   [THETA2_MIN, THETA2_MAX]], opts=[options,
71                                                    options])[0]
72
73  @lru_cache(maxsize=None)
74  def p_n_given_theta(n, theta2):
75      p_n = (np.exp(-theta2) * (theta2)**n)/ factorial(n)
76      return p_n
77
78  @lru_cache(maxsize=None)
79  def p_theta(step, theta1, theta2):
80      if step == 0:
81          return 1/( (THETA1_MAX-THETA1_MIN) *
82                 np.log10(theta2)/np.log10(THETA2_MAX/
83                                           THETA2_MIN))
84      else:
```

```
        return p_theta(step - 1, theta1, theta2) *
               p_n_given_theta(n_a[step-1], theta2) *
               p_n_given_theta_f(n_b[step-1], f_best[step -1],
                                 theta1, theta2)

def p_theta_given_n_f(n, f, theta1, theta2, marginal):
    return (p_theta(step, theta1, theta2) *
            p_n_given_theta_f(n, f, theta1, theta2))/marginal

def theta_tilde(n,f, marginal):
    def integrand(theta1, theta2):
        return theta1 * p_theta_given_n_f(n, f, theta1,
                                          theta2, marginal)
    return nquad(integrand, [[THETA1_MIN, THETA1_MAX],
                 [THETA2_MIN, THETA2_MAX]], opts=[options,
                                                  options])[0]

def clear_caches():
    p_theta.cache_clear()
    #p_n_given_theta.cache_clear()
    p_n_given_f.cache_clear()
    #p_n_given_theta_f.cache_clear()
    #T.cache_clear()
    #zeta.cache_clear()

def Gain(f, truncate):
    g_gamma = 0
    truncate = truncate
```

```
113     for n in np.arange(truncate):
114         marginal = p_n_given_f(n, f)
115         g_gamma += marginal * theta_tilde(n,f, marginal)**2
116     return g_gamma
117 
118 def p_theta_1(step,theta1):
119     def integrand(theta2):
120         return p_theta(step,theta1,theta2)
121     return quad(integrand, THETA2_MIN, THETA2_MAX)[0]
122 
123 def p_theta_2(step,theta2):
124     def integrand(theta1):
125         return p_theta(step,theta1,theta2)
126     return quad(integrand, THETA1_MIN, THETA1_MAX)[0]
127 
128 #%%Starting optimisation protocol
129 
130 print('Starting optimisation protocol')
131 f_best = []
132 n_a = [] # photon counts without atoms
133 n_b = [] # photon counts with atoms
134 od = []
135 od_err =[]
136 truncate_n = int(THETA2_MAX+3*THETA2_MAX**0.5)
137 # Where to truncate the infinite sum of eq 9
138 truncate_n_l = int(THETA2_MIN-3*THETA2_MIN**0.5)
139 if truncate_n_l < 0:
140     truncate_n_l=0
```

```
with tqdm(total=freq_steps, desc="Progress") as pbar:
    for step in np.arange(max_steps):
        print('\n' + str(step))
        clear_caches()
        gains = []
        freqs = np.linspace(initial_freq, final_freq,
                            num = freq_steps)
        for freq in freqs:
            gains.append(Gain(freq, truncate_n))
            pbar.update(1)
        pbar.refresh()
        ffunc=interp1d(freqs,gains,kind='cubic')
        testpoints=np.linspace(np.min(freqs),np.max(freqs),
                               5000)
        gvals=ffunc(testpoints)
        print('\nfreq: %g' %(testpoints[np.argmax(gvals)]))
        print('vco: %0.3g'
              %(np.interp(testpoints[np.argmax(gvals)],
                          freq_ref,VCO_ref)))
        f_best.append(testpoints[np.argmax(gvals)])
        #saves and feeds new best freq in the loop
        mu=step
        print('(Give a non-numeric input to end run.)')
        try:
            n_b.append(float(input("Input photon count " +
                                   "with atoms :")))
        except ValueError:
            print('\n** value error - exiting loop')
```

```
169             break
170         n_a.append(float(input("Input photon count " +
171                                "without atoms :")))
172 
173         def integrand_p_theta(theta1, theta2):
174             return p_theta(mu+1,theta1,theta2)
175         od_norm=nquad(integrand_p_theta,
176                       [[THETA1_MIN, THETA1_MAX],
177                        [THETA2_MIN, THETA2_MAX]],
178                       opts=[options,options])[0]
179 
180         def integrand_p_theta_theta(theta1, theta2):
181             return theta1*p_theta(mu+1,theta1,theta2)
182         od_top=nquad(integrand_p_theta_theta,
183                      [[THETA1_MIN, THETA1_MAX],
184                       [THETA2_MIN, THETA2_MAX]],
185                      opts=[options,options])[0]
186 
187         def integrand_p_theta_theta2(theta1,theta2):
188             return (theta1**2)*p_theta(mu+1,theta1,theta2)
189         od_err1=nquad(integrand_p_theta_theta2,
190                       [[THETA1_MIN, THETA1_MAX],
191                        [THETA2_MIN, THETA2_MAX]],
192                       opts=[options,options])[0]
193 
194         od.append(od_top/od_norm)
195         od_err.append(np.sqrt(od_err1/od_norm-od[-1]**2))
196         print('\nOptical depth estimate : %g \u00B1 %g'
```

```
            %(od[-1],od_err[-1]))

        testpoints_min1=np.linspace(THETA1_MIN,THETA1_MAX,
                                    ptnum+1)
        testres_min1=np.zeros(len(testpoints_min1))
        for k in range(len(testpoints_min1)):
            testres_min1[k]=p_theta_1(mu+1,
                                      testpoints_min1[k])
        maxres_min1=np.max(testres_min1)
        relvals_min1=testres_min1/maxres_min1
        for i in range(len(relvals_min1)):
            if relvals_min1[i]>cutoff_threshold:
                THETA1_MIN=testpoints_min1[i]
                break
        for j in range(len(relvals_min1)):
            if relvals_min1[-j-1]>cutoff_threshold:
                THETA1_MAX=testpoints_min1[-j-1]
                break
        testpoints_min2=np.linspace(THETA2_MIN,THETA2_MAX,
                                    ptnum+1)
        testres_min2=np.zeros(len(testpoints_min2))
        for k in range(len(testpoints_min2)):
            testres_min2[k]=p_theta_2(mu+1,
                                      testpoints_min2[k])
        maxres_min2=np.max(testres_min2)
        relvals_min2=testres_min2/maxres_min2
        for i in range(len(relvals_min2)):
            if relvals_min2[i]>cutoff_threshold:
```

```
                THETA2_MIN=testpoints_min2[i]
                break
        for j in range(len(relvals_min2)):
            if relvals_min2[-j-1]>cutoff_threshold:
                THETA2_MAX=testpoints_min2[-j-1]
                break

        if THETA1_MIN < THETA1_MIN_i:
            THETA1_MIN=THETA1_MIN_i
        if THETA1_MAX > THETA1_MAX_i:
            THETA1_MAX=THETA1_MAX_i
        if THETA2_MIN < THETA2_MIN_i:
            THETA2_MIN=THETA2_MIN_i
        if THETA2_MAX > THETA2_MAX_i:
            THETA2_MAX=THETA2_MAX_i

        pbar.reset()
'''
ploting results
'''

od = np.array(od)
od_err = np.array(od_err)
print('  ----  CODE ENDED  ----  ')
```