



Durham E-Theses

Optical tweezer arrays of strontium atoms

HILL, MATTHEW

How to cite:

HILL, MATTHEW (2022) *Optical tweezer arrays of strontium atoms*, Durham theses, Durham University. Available at Durham E-Theses Online: <http://etheses.dur.ac.uk/14662/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Optical tweezer arrays of strontium atoms

Matthew Hill

A thesis presented for the degree of
Doctor of Philosophy



Quantum Light and Matter
Department of Physics
Durham University
United Kingdom
October 2022

Abstract

This thesis presents the development of a strontium optical tweezer array apparatus for precision measurements, towards a spin squeezed tweezer array clock. We demonstrate fast single atom detection of strontium atoms in a tweezer in just 200 μs , with fidelity 0.989(6). With the addition of Sisyphus cooling we measure comparable fidelity though with a longer 30 ms exposure time, and can resolve zero, one, and two atoms all with fidelity > 0.8 . Holographic magic wavelength tweezer arrays of strontium atoms are constructed, demonstrating loaded arrays of over 20 tweezers. High magnification imaging is used to show site-resolved imaging of atoms in tweezers spaced by as little as 6 μm , despite aberrations in the design of the objective lens collecting the fluorescence.

The precision measurement and cancellation of the electric field at the location of the atoms is demonstrated using Rydberg atoms as sensors, measuring the stray electric field to be $39.8 \pm 1.5 \text{ mVcm}^{-1}$, before cancelling it to a residual field of $1.9_{-1.9}^{+2.2} \text{ mVcm}^{-1}$. The demonstrated residual electric field produces an environment compatible with precision measurements, where the residual electric field would not be limiting on any atomic clock, while also opening a pathway to precision measurements of Rydberg states. Our clock laser is introduced and characterised against a GPS-referenced difference frequency comb, showing stability compatible with the use as a local oscillator for an atomic clock. We consider plans for the implementation of this clock laser system towards making clock measurements.

Contents

Abstract	ii
Contents	iii
List of figures	viii
List of tables	x
Nomenclature	xi
Declaration	xvi
Acknowledgements	xvii
1 Introduction	1
1.1 General introduction	1
1.2 Thesis outline	3
1.3 Publications	4
2 Background	5
2.1 Strontium	7
2.2 Optical tweezers	10
2.2.1 Trapping in far off-resonance potentials	11
2.2.2 Optical tweezers vs optical lattices	12
2.2.3 Choice of trapping wavelengths	15
2.3 Atomic clocks	17
2.3.1 Components of an atomic clock	19

2.3.2	From microwave to optical	20
2.3.3	Atoms, ions and isotopes	23
2.3.4	Quantum projection noise	25
2.4	Rydberg atoms	27
2.4.1	Rydberg energy levels	28
2.4.2	Rydberg properties	29
2.4.3	DC Stark effect	29
2.4.4	Rydberg blockade and dressing	30
2.5	Summary	32
3	The cold strontium experiment	33
3.1	Preparing cold atoms	33
3.1.1	Oven, Zeeman slower & main chamber	35
3.1.2	Blue MOT	37
3.1.3	Red MOT	39
3.1.4	Loading a tweezer	42
3.2	Laser systems	43
3.2.1	461 nm lasers	43
3.2.2	689 nm lasers	50
3.2.3	Dipole trapping lasers (532/813 nm)	54
3.2.4	Rydberg excitation laser (317/319 nm)	55
3.2.5	Clock laser (698 nm)	59
3.2.6	Other lasers	62
3.3	<i>In-vacuo</i> equipment	64
3.3.1	<i>In-vacuo</i> lenses	65
3.3.2	Magnetic field coils	66
3.3.3	Electrodes	67
3.3.4	Microchannel plate (MCP)	68
3.4	Imaging	68
3.4.1	SPAD array	69
3.4.2	CCDs	71
3.5	Summary	72

4	Number-resolved imaging in a tweezer	73
4.1	Imaging atoms	74
4.1.1	461 nm probe	75
4.1.2	SPAD imaging	75
4.1.3	Photon collection	77
4.2	Sisyphus cooling	78
4.2.1	Motivation	78
4.2.2	Method	79
4.3	Histograms	82
4.3.1	Composite Poisson distribution	82
4.3.2	Assumptions	85
4.4	Fidelity	86
4.4.1	Thresholds	86
4.4.2	Limiting infidelity and loss	87
4.4.3	Selective fidelity	89
4.5	Results	89
4.5.1	Imaging with continuous Sisyphus cooling	90
4.5.2	Fast single atom imaging	94
4.5.3	Enhanced fidelity imaging	96
4.6	Summary	99
5	Magic wavelength tweezer arrays	101
5.1	Magic wavelength tweezers	101
5.1.1	Benefits and drawbacks	102
5.1.2	Making a tweezer	104
5.2	Generating tweezer arrays	106
5.2.1	Methods to make arrays	106
5.2.2	Basic principles of the SLM	109
5.2.3	Phase modulation	110
5.3	High magnification imaging	111
5.3.1	Motivation of high magnification	112
5.3.2	Double telescope	114

5.3.3	Aberration and corrective optics	116
5.4	Loaded tweezer arrays	119
5.4.1	Loading from a BB MOT	120
5.4.2	Gallery of tweezer arrays	123
5.4.3	Loading schemes	123
5.5	Future imaging protocol	124
5.6	Summary	127
6	Rydberg electrometry	128
6.1	Measuring the DC Stark shift	129
6.1.1	Choosing a Rydberg state	129
6.1.2	Experimental sequence	130
6.2	Electric field control	132
6.3	Stark maps	134
6.3.1	Electric field cancellation	134
6.3.2	Comparison with ARC	137
6.4	Summary	145
7	Conclusions and outlook	148
	Bibliography	151
	Appendix	175
A	Experimental methods	175
A.1	461 nm beat lock components	175
B	Number-resolved imaging	177
B.1	Sisyphus cooling optimisation	177
B.2	Histograms including loss	181
B.3	Selective fidelity	184
C	Magic wavelength tweezers	186
C.1	813 nm tweezer preparation	186

C.1.1	Initial alignment	186
C.1.2	Loading and imaging a tweezer	187
C.1.3	Optimisation of tweezer waist	188
C.2	Phase masks for the SLM	192
C.2.1	SLM calibration	193
C.2.2	Types of phase masks	194
C.2.3	Gerchberg-Saxton algorithm	200
C.3	Corrective optics for imaging	202
C.3.1	Aberration correction	202
C.3.2	Theoretical imaging performance	204
D	Rydberg electrometry	207
D.1	Finding Rydberg states	207
D.1.1	UV laser optimisation	207
D.1.2	Depletion in a BB MOT	208
E	Towards clock measurements	210
E.1	Clock stability	210
E.1.1	Frequency comb	211
E.1.2	GPS reference	213
E.1.3	Beat measurement	213
E.1.4	Overlapping Allan deviation (OADEV)	215
E.1.5	Laser linewidth	219
E.2	Clock excitation	221
E.2.1	Magnetically induced spectroscopy (MIS)	221
E.2.2	Systematic shifts	222
E.3	Summary	226

List of figures

2.1	Strontium energy level diagram	8
2.2	Ground state polarisability	15
2.3	Differential AC Stark shift	18
3.1	Trapping and imaging setup	34
3.2	Vacuum system layout	36
3.3	Photograph of the Zeeman slowed atoms	36
3.4	Photograph of the blue MOT	38
3.5	SPAD image of a BB MOT (low magnification)	40
3.6	SPAD image of a NB MOT (low magnification)	40
3.7	Spectrum of lasers used	44
3.8	461 nm beat lock	46
3.9	461 nm beat lock error signal	47
3.10	461 nm beat lock stability	47
3.11	Layout for the 461 nm lasers	48
3.12	Layout for the 689 nm laser	50
3.13	In-loop error signal for the 689 nm laser	53
3.14	Long term cavity drift of the 689 nm laser	53
3.15	Rydberg excitation laser layout	56
3.16	In-loop error signal for the 638 nm laser	58
3.17	Layout for the clock laser	60
3.18	Cavity ring down measurement of the clock laser	61
3.19	In-loop error signal for the clock laser	61
3.20	Background count rate for the SPAD array	70

4.1	532 nm tweezer and imaging setup	74
4.2	Single shot image of a single atom	76
4.3	Sisyphus cooling	80
4.4	Histogram showing example of number-resolved detection	84
4.5	Histogram showing number-resolved imaging of atoms, with Sisyphus cooling and a weak probe beam	91
4.6	Number-resolved imaging losses	92
4.7	Histogram showing single atom imaging with a strong probe beam and no cooling	95
4.8	Histogram showing single atom imaging with an interleaved sequence of strong probe pulses separated by Sisyphus cooling	98
5.1	813 nm tweezer array setup	105
5.2	Pair of tweezers at low magnification	113
5.3	High magnification imaging with corrective optics	114
5.4	SPAD image of a narrowband MOT (high magnification)	115
5.5	Single tweezer at high magnification	118
5.6	Pair of tweezers at high magnification	118
5.7	Gallery of tweezer arrays	122
6.1	Depletion from a released tweezer	131
6.2	Electrode configurations	133
6.3	Stark map in the $\pm\hat{z}$ direction	135
6.4	Stark map in the $\pm\hat{y}''$ direction	136
6.5	Stark map in the $\pm\hat{x}''$ direction	136
6.6	Simulated Stark map for $5s79d\ ^3D_1 m_j = 0$ up to $\pm 200\text{ mVcm}^{-1}$	138
6.7	Simulated Stark map for $5s79d\ ^3D_1 m_j = 1$ up to $\pm 200\text{ mVcm}^{-1}$	138
6.8	Mixing of the Rydberg states for the $5s79d\ ^3D_1 m_j = 1$ state	140
6.9	Mixing of the Rydberg states for the $5s79d\ ^3D_1 m_j = 0$ state	140
6.10	Stark map in the $\pm\hat{z}$ direction with residuals	142
6.11	Stark map comparing experimental data with ARC	143
B.1	Effectiveness of Sisyphus cooling vs Sisyphus frequency	179

B.2	Effectiveness of Sisyphus cooling vs probe scattering rate	180
C.1	Trap frequency measurement	190
C.2	Trap frequency vs power	190
C.3	Tweezer waist optimisation (vertical alignment)	191
C.4	Tweezer waist optimisation (horizontal alignment)	191
C.5	Phase masks for translations	196
C.6	Gallery of phase masks for tweezer arrays	199
C.7	Phase mask correcting for SLM surface	200
C.8	Gerchberg-Saxton algorithm	201
E.1	Frequency comb	212
E.2	Beat frequency of the clock laser and frequency comb	216
E.3	Residuals of the beat between the clock laser and comb	216
E.4	Overlapping Allan deviation (OADEV) of the clock laser	218

List of tables

3.1	Configuration A of the 461 nm laser system	45
3.2	Configuration B of the 461 nm laser system	49
4.1	Comparison of different imaging methods	100
6.1	Electric field calibration with ARC	144

Nomenclature

Wavelengths

319 (317) nm - $5s5p^3P_1 \rightarrow 5snl^3L_J$ ($5s5p^3P_0 \rightarrow 5snl^3L_J$), where n is large and L is typically S or D.

408 nm - $5s^2S_{\frac{1}{2}} \rightarrow 5p^2P_{\frac{3}{2}}$ (auto-ionisation)

413 nm - $5s5p^1P_1 \rightarrow 5snl^1L_J$ where n is large and L is typically S or D.

461 nm - $5s^2^1S_0 \rightarrow 5s5p^1P_1$

532 nm - Tweezer wavelength λ_{trap}

638 (634) nm - SHG to 319 (317) nm

679 nm - $5s5p^3P_0 \rightarrow 5s6s^3S_1$

688 nm - $5s5p^3P_1 \rightarrow 5s6s^3S_1$

689 nm - $5s^2^1S_0 \rightarrow 5s5p^3P_1$

698 nm - $5s^2^1S_0 \rightarrow 5s5p^3P_0$

707 nm - $5s5p^3P_2 \rightarrow 5s6s^3S_1$

813 nm - Tweezer wavelength λ_{trap} (magic for $5s^2^1S_0 \rightarrow 5s5p^3P_0$)

922 nm - SHG to 461 nm

1079 (1070) nm - SFG with 1549-1565 nm to 638 (634) nm

1549-1565 nm - SFG with 1079 (1070) nm to 638 (634) nm

Abbreviations

AC - Alternating current

ADEV - Allan deviation

AFG - Arbitrary function generator

AOM - Acousto-optic modulator

APC - Angled physical contact

AR - Anti-reflection

ARC - Alkali(ne) Rydberg Calculator (simulation software [1])

BB - Broadband

BBR - Black-body radiation

CAD - Computer aided design

CCD - Charge-coupled device

CEO - Carrier envelope offset

CMOS - Complementary metal-oxide-semiconductor

Comb - Shortened name for frequency comb (see also DFC and OFC)

DBS - Diffractive beam splitter

DC - Direct current

DExTer - Durham Experimental Terminal (experimental control software [2])

DFC - Difference frequency comb

DFG - Difference frequency generation

DL - Diode laser

DMD - Digital micromirror device

EASLM - Electrically-addressed spatial light modulator

ECDL - External-cavity diode laser

EDFA - Er-doped fiber amplifier

EMCCD - Electron multiplying charge-coupled device

EOM - Electro-optic modulator

FALC - Fast analog linewidth control

FFT - Fast Fourier transform

FSR - Free spectral range

GNSS - Global navigation satellite system

GPS - Global positioning system

HWP - Half waveplate

IR - Infrared

LCOS - Liquid-crystal on silicon

LIAD - Light induced atomic desorption

LPI - Local position invariance

LUT - Look up table

MCP - Microchannel plate detector

MIS - Magnetically induced spectroscopy

MOT - Magneto-optical trap

MTS - Modulation transfer spectroscopy

NA - Numerical aperture

NB - Narrowband

NBK-7 - Common type of optical glass

NDD - Non-demolition detection

NPBS - Non-polarising beam splitter

NPL - National Physical Laboratory

OADEV - Overlapping Allan deviation.

OASLM - Optically-addressed spatial light modulator

OD - Optical depth

OFC - Optical frequency comb

PBS - Polarising beam splitter

PC - Physical contact

PCF - Photonic crystal fiber

PD - Photodiode

PDH - Pound–Drever–Hall

PPLN - Periodically-poled lithium niobate

PPSLT - Periodically-poled stoichiometric lithium tantalate

QED - Quantum electrodynamics

QPN - Quantum projection noise

QWP - Quarter waveplate

RF - Radiofrequency

RMS - Root mean square

SFG - Sum frequency generation

SHG - Second harmonic generation

SI - International System of Units

SLM - Spatial light modulator

SNR - Signal to noise ratio

SolsTis - Ti:Sa laser made by M Squared Lasers

SPAD - Single photon avalanche diode

Sprout - 532 nm laser (Lighthouse Photonics, Sprout-G 15 W)

TA - Tapered amplifier

TEM₀₀ - The transverse electromagnetic 00 mode (Gaussian beam)

Ti:Sa - Titanium-sapphire laser

ULE - Ultra-low expansion

UV - Ultraviolet

VCO - Voltage-controlled oscillator

YDFA - Yb-doped fiber amplifier

Declaration

I confirm that no part of the material offered has previously been submitted by myself for a degree in this or any other University. Where material has been generated through joint work, the work of others has been indicated. The work in this thesis was supported by a studentship from EPSRC. The data presented in this thesis is available from the author, upon reasonable request.

Copyright © 2022 by Matthew Hill.

The copyright of this thesis rests with the author. No quotations from it should be published without the author's prior written consent and information derived from it should be acknowledged.

Acknowledgements

I would like to offer my thanks to everyone who supported me through the course of my PhD. Firstly I must thank the group of people who I have worked alongside in the strontium project. My supervisors Matt Jones and Charles Adams provided guidance and support whenever I needed it. Niamh Jackson trained me across the many areas required in a cold strontium experiment, giving me a great start to my PhD, and was brilliant to work alongside in the work detecting single atoms. Zhongxiao Xu was a great help around the lab, and Frédéric Leroux took care of optical design. Mitch Walker and Liam Gallagher helped take the final results of my PhD in the optimisation of the imaging system for tweezer arrays, and made the lab a pleasant place to work, I look forward to continuing working with them.

Former members of the strontium project and others around QLM and the department have also been a great support, whom I generally thank. Dani Pizzey has been a fantastic help whenever there was a problem with a vacuum system. Ryan Hanley constructed the first tweezer in our lab, and both him and Alistair Bounds have been great at answering my many questions. The mechanical workshop team were great in producing key components for the project quickly, and always took the time to explain things and help refine designs. Outside of Durham I also thank the research group of Antoine Browaeys at the Institut d'Optique for providing advice and kindly sharing their code for operating the SLM.

Finally more personal thanks to my family, in particular my parents, for their support throughout my time at universities and all the steps beforehand along the way. I must also thank my girlfriend Katie for being there for me throughout my time in Durham, and because I am expecting an acknowledgement in her PhD thesis so it would be rude not to thank her first.

Chapter 1

Introduction

1.1 General introduction

Atomic clocks have shown incredible advancement in recent times, where three species of optical atomic clocks (Sr [3–5], Yb [6] and Al^+ [7]) have now passed below the $\sim 2 \times 10^{-18}$ fractional frequency systematic uncertainty level, corresponding to an uncertainty below a second in the age of the universe. State of the art measurements with Sr clocks are now able to observe shifts on the 7.6×10^{-21} fractional frequency precision level, resolving gravitational redshift on the sub millimetre scale [8]. These extremely high precision measurements offer applications in testing fundamental physics [9–14], such as gravitational wave sensing, searching for variation of fundamental constants and imposing constraints upon dark matter. Other applications include global navigation and communication, relativistic geodesy [15], geophysical measurements [16], and redefining the second [17–19].

State of the art atomic clocks are increasingly approaching the QPN (quantum projection noise) [20] limit, with ion clocks operating at the QPN limit [7]. Continuing to improve clocks beyond this limit to improve their performance will be key to higher quality measurements for their applications. The QPN limit applies for uncorrelated atoms, where to increase precision then either further correlations between the atoms must be introduced, or the probe time increased. Methods to generate these correlations include, but are not limited to, Rydberg dressing [21, 22]

and cavity QED (quantum electrodynamics) measurements [23–31] to prepare spin-squeezed states [32]. Squeezing reduces the uncertainty in one direction, at the expense of increasing it in another, permitting the uncertainty to be reduced below the QPN limit. In this thesis we work towards the implementation of a spin-squeezed atomic clock, using Rydberg-dressed interactions in an array of Sr atoms [21].

To implement a spin-squeezed atomic clock based upon the scheme described in the work of Gil et al. [21], we desire control of the spacing and geometry of the array of atoms, as well as the ability to accurately measure the state of individual atoms. Atoms can be trapped on the micron scale using either optical lattices or optical tweezers, where magic wavelength trapping [33, 34] avoids unnecessary perturbations to the clock transition frequency. In this work we use optical tweezer arrays due to the greater control of the spacing between atoms, and the availability of arbitrary geometries [35, 36]. Tweezer arrays can be used to produce arrays of single atoms in 1D [37, 38], 2D [39–41] and 3D [35, 36], where 3D arrays allow a maximum number of sites within the Rydberg-dressed volume for stronger spin-squeezing [21]. Tweezer arrays also give more control of the spacing between atoms than lattices, such that single atoms in each site can be resolved during measurements, making them an ideal platform for our experiments.

Single atom control in optical tweezers has recently been extended from monovalent atoms [42] to divalent atoms such as Sr [43–48] and Yb [49–51], allowing the toolbox of optical tweezers to be brought to optical atomic clock experiments. Loading of single atoms into tweezer arrays to higher fill fractions is also under continuing development, which is of importance for maximising the number of occupied sites for spin-squeezing protocols. Tweezer arrays with low number of unoccupied trapping sites have been demonstrated through greater than 50% loading methods [50, 52, 53], and shuttling of atoms to rearrange into smaller filled arrays [37, 38, 54–58].

Tweezer arrays are not only a promising platform for improving atomic clocks, they are useful for a range of other experiments with arrays of atoms and molecules [59] including quantum simulation [38, 60–62], quantum computation [41], and quantum chemistry [63, 64]. Tweezer arrays of Sr also have applications in the precision measurement of Rydberg states [1, 65–68] for tests of fundamental physics [69, 70].

1.2 Thesis outline

The work in this thesis creates a strontium tweezer array platform towards the implementation of a spin-squeezed optical atomic clock. The key parts of this thesis are single atom control in tweezers, the construction of arbitrary tweezer arrays with site resolved imaging, the use of Rydberg states for electrometry, and preparation towards an optical atomic clock. The thesis is structured as follows:

- Chapter 1 introduces the work undertaken in this thesis, and is where you are now, welcome to chapter 1.
- Chapter 2 discusses the QPN limit and the spin-squeezing protocol which we would like to use to beat it. It then introduces details of strontium and its transitions needed throughout this work. We then provide background on the key components towards developing a spin-squeezed tweezer array clock: optical tweezers, atomic clocks, and Rydberg atoms. The use of dipole trapping is introduced in the context of optical tweezers and lattices, as well as considering magic wavelengths. Atomic clocks are introduced, including the reasons to work with optical clock transitions like the one found in strontium, and the specifics of making a bosonic strontium clock. Rydberg atoms are then introduced in both the contexts of Rydberg-dressing for the generation of spin-squeezing, and as precision sensors of the electric field.
- Chapter 3 describes the preparation of the cold atomic sample in optical tweezers. The experimental system is detailed, including laser systems, vacuum systems and detectors.
- Chapters 4 and 5 cover trapping strontium atoms in tweezers. Chapter 4 describes imaging and cooling of atoms in a single 532 nm tweezer, demonstrating number-resolved imaging of atoms. Chapter 5 then generalises the system to a holographic array of tweezers at the magic wavelength, producing arbitrary tweezer arrays of atoms as a platform for future measurements.
- Chapter 6 uses Rydberg atoms as electric field sensors to measure and cancel

the stray electric field at the position of the tweezers, to produce an environment compatible with precision measurement for both Rydberg and clock experiments.

- Chapter 7 summarises the conclusions of the thesis, and provides outlook on where we can go next with the experimental platform developed in this work.
- Appendix E characterises our clock laser system combined with the frequency comb, with discussions of the imminent use of the clock laser on the atoms.

1.3 Publications

Below is a publication arising from this work, forming the basis of chapter 4:

N. C. Jackson, R. K. Hanley, M. Hill, F. Leroux, C. S. Adams and M. P. A. Jones, *Number-resolved imaging of ^{88}Sr atoms in a long working distance optical tweezer*, *SciPost Phys.* **8**, 038 (2020). [43]

A second publication containing the key results of chapter 5, chapter 6, and appendix E is planned to be prepared after submission of this thesis.

Chapter 2

Background

Optical atomic clocks are increasingly becoming limited by QPN [20]. The QPN limit for a clock of N_{at} uncorrelated atoms is given by

$$\sigma(\tau) = \frac{\beta}{Q\sqrt{N_{\text{at}}}} \sqrt{\frac{T_c}{\tau}}, \quad (2.0.1)$$

where Q is the quality factor¹ of the clock transition, T_c is the clock cycle time before correction, and β is a constant of order unity arising from the specifics of the locking method used [10, 71]. As increasing N_{at} and T_c further to improve the instability becomes less practical, it is desirable to get around this limit.

The limit arises from the uncertainty relation, where the many uncorrelated two level atoms behave like a total spin \hat{J} , the sum of the effective spin of each atom. This spin \hat{J} can be described on a Bloch sphere, with uncertainties in the directions orthogonal to the mean spin $\langle \hat{J} \rangle$. The uncertainty relation is then

$$\Delta \hat{J}_{\perp,1} \Delta \hat{J}_{\perp,2} \geq \langle \hat{J} \rangle / 2, \quad (2.0.2)$$

where the uncertainty in each orthogonal directions $\Delta \hat{J}_{\perp,1}$ and $\Delta \hat{J}_{\perp,2}$ are equal for uncorrelated atoms.

To exceed the QPN limit, correlations should be added between the atoms such that the uncertainty in the direction of the measurement is reduced at the expense of the other, squeezing the circular distribution of uncertainty into an ellipse. With

¹ $Q = \frac{\Delta\nu}{\gamma}$ is the ratio of the clock transition frequency $\Delta\nu$ and its linewidth γ .

quantum correlations between the atoms, the new absolute limit becomes the Heisenberg limit [26], allowing performance up to a factor of $\sqrt{N_{\text{at}}}$ better than the QPN limit. We aim towards implementing a spin-squeezed clock array of strontium atoms to demonstrate atomic clock performance beyond the QPN limit, based upon the proposal by Gil et al. [21].

Atoms in an array are initially uncorrelated, as the interactions when in the ground or excited states of the clock transition are negligible at long distances. In order to generate correlations between the atoms, state-dependent long range interactions are required. The scheme described in Gil et al. [21] relies on off-resonantly coupling the excited clock state to a Rydberg state (Rydberg dressing the excited clock state with a small amount of a Rydberg state), leading to tunable long-range interactions suitable for generating considerable squeezing in just a few microseconds.

The atom is initially excited with a $\frac{\pi}{2}$ pulse to an equal superposition of the ground and excited clock states. The application of the Rydberg dressing light leads to a one-axis twisting of the uncertainty ellipse, reducing uncertainty in one direction. The rest of a spin-echo sequence then follows to remove undesired linear spin rotations and broadening from inhomogeneous detunings. [21]

Rydberg dressed interactions have dressed blockade radii on the scale of several microns, meaning the spacing of the individual atoms must be on the scale of a few microns to achieve sufficient squeezing. A tweezer array is a great apparatus in which to implement the array of single atoms, with trapping regions and spacings on the scale of a few microns. The strength of the squeezing also relies on the density of atoms within the dressed blockade radius, where generalising the tweezer array to 3D can yield stronger squeezing with shorter dressing times [21]. Tweezers also permit single site-resolved readout of the individual states of each atom in the array, with arrays in arbitrary 3D geometries [35].

To implement this scheme we then need four main ingredients: an atomic sample with suitable properties including an optical clock transition, a tweezer array with which to trap the atoms, the construction of an atomic clock, and the addition of Rydberg interactions between the atoms. We will now discuss each of these aspects,

introducing the underlying concepts in more detail.

2.1 Strontium

Throughout this thesis we exclusively use the ^{88}Sr isotope of strontium as our atomic sample. Strontium has a high quality factor (Q), due to the optical transition frequency and narrow linewidth, allowing low fractional frequency uncertainties to be reached. Strontium is one of the three species which have demonstrated fractional frequency uncertainty below a second in the age of the universe [3–5], and has the lowest fractional frequency precision recorded with an atomic clock [8]. It is conveniently trappable in optical tweezer arrays, and has suitable and accessible Rydberg states to create near-isotropic Rydberg dressed interactions for spin-squeezing [21]. Strontium is a highly suitable element with which to implement a spin-squeezed tweezer array clock, where we now introduce the properties of ^{88}Sr more generally.

The key property of strontium is that it is a divalent atom, having electronic structure $[\text{Kr}] 5s^2$. The two valence electrons are the defining characteristic of the electronic structure of strontium, giving rise to both spin singlet and spin triplet electronic states. The electronic structure has a wealth of useful properties, as does that of many other alkaline-earth(like) elements which share a similar structure. An energy level diagram including key transitions of strontium is shown in figure 2.1.

There are three key transitions from the ground state, $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$, $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$, and $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$. The linewidths of these three transitions vary by many orders of magnitude. With both spin singlet and spin triplet states, we first consider the dipole allowed transition within the singlet manifold. The $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition is a broad ($\Gamma = 2\pi \times 30.2$ MHz) transition from the ground state, and is very useful for both cooling and imaging due to the high scattering rates achievable. We use this transition slowing, cooling, and trapping atoms, as well as all of our fluorescence imaging of atoms.

One might have expected that transitions between the singlet and triplet manifolds would be forbidden by the selection rules for pure LS coupling² ($\Delta S = 0$),

²Where L and S are the total orbital angular momentum and total spin angular momentum

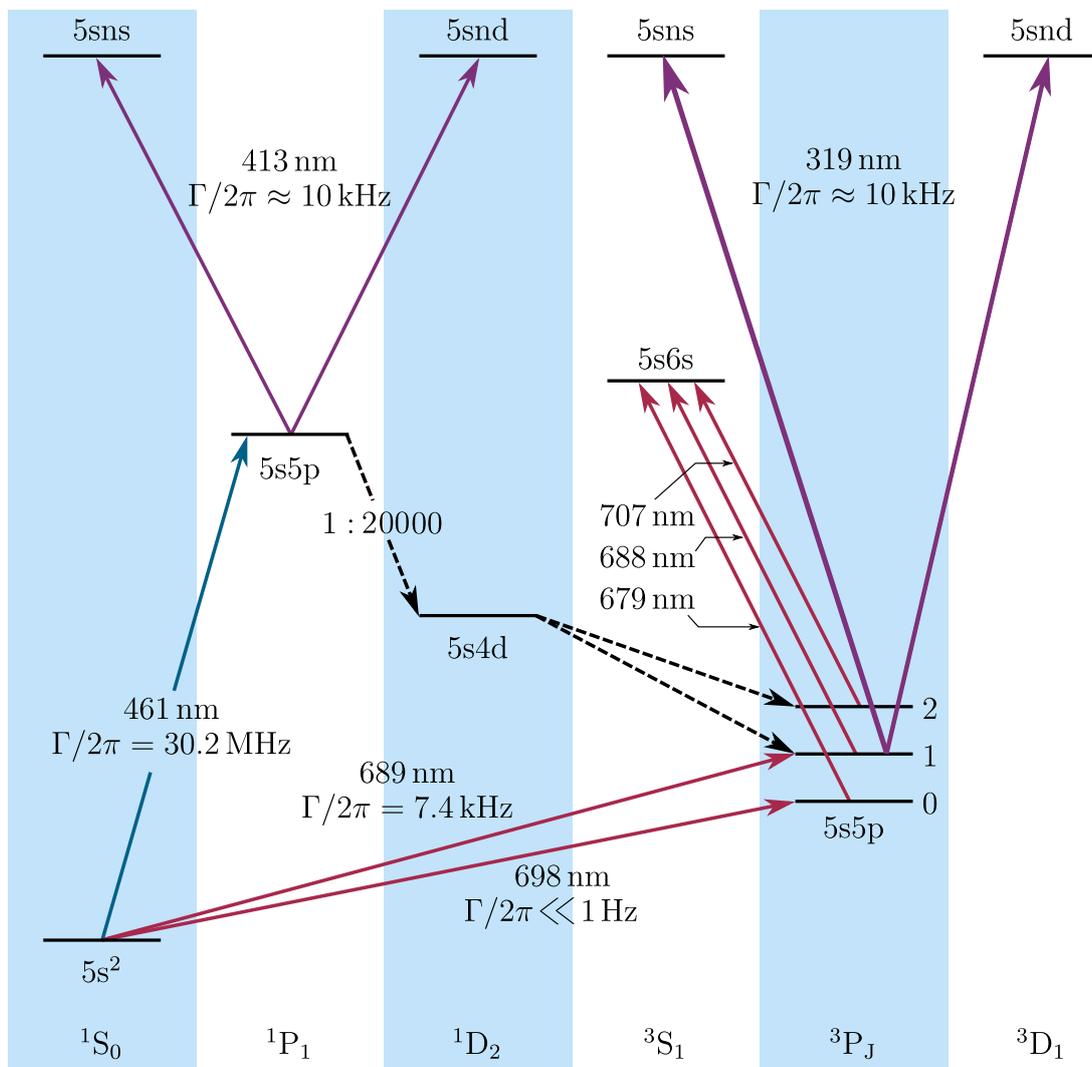


Figure 2.1: Energy level diagram for strontium. This figure is modified from a figure in the thesis of Ryan Hanley [72], which was based upon the original figure from the thesis of Danielle Pizzey [73].

though in fact some of these transitions are weakly allowed. The selection rules assume that we are in the LS coupling regime, where the total angular momentum $J = L + S$; this regime holds well for low mass atoms. In the case of heavy atoms we instead move to the jj -coupling regime, where the total angular momentum is better described by $J = j_1 + j_2$.³ With 88 nucleons, strontium's structure is reasonably well described with the LS coupling regime, but it has sufficient mass such there is some mixing of states with the same J . This mixes small amounts of the spin singlet states into the spin triplet states of the same J , acquiring a small electric dipole moment; this weakly allows singlet to triplet transitions ($\Delta S = 1$). In the case of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition, it is weakly allowed by mixing in a small amount of the $5s5p\ ^1P_1$ state [72, 74], resulting in a linewidth $\frac{\Gamma}{2\pi} = 7.4$ kHz. The narrow $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition is an important narrow transition for laser cooling to temperatures of < 1 μK .

There is an even narrower linewidth transition available in strontium from the ground state, which can be found by violating even more selection rules. The $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ 'clock' transition is doubly forbidden: firstly, as before, it is spin forbidden ($\Delta S = 1$), and secondly it is angular momentum forbidden ($J = 0 \rightarrow J = 0$) to all orders of the multipole expansion [10]. As a result, the linewidth of this transition $\Gamma/2\pi \ll 1$ Hz, where the exact linewidth depends upon parameters including the choice of isotope. We use the bosonic ^{88}Sr isotope, where there is no nuclear spin the weakly allow the transition, so the linewidth is instead set by the specifics of the excitation scheme used to allow the transition, where we plan to use magnetically induced spectroscopy (MIS, see section 2.3.3) [75]. This extremely narrow linewidth optical clock transition is key to the use of strontium in atomic clocks, as we discuss further in section 2.3.

The divalent nature of strontium also offers several other advantages. There is only a single ground state, the $5s^2\ ^1S_0$ state, making it simpler to work with. Furthermore, for the bosonic isotopes of strontium (such as ^{88}Sr) there is no nuclear spin ($I = 0$), meaning that there is no hyperfine structure to consider, simplifying

respectively.

³Where each of the two valence electrons has its own individual total angular momentum, j .

clock schemes and laser systems.

The availability of both spin singlet and spin triplet states extends to having both singlet [73, 76–78] and triplet [72, 79, 80] Rydberg states.⁴ These states offer a variety of attractive and repulsive interactions, where we discussed some of the uses of Rydberg atoms in section 2.4. In particular, the $5sns\ ^3S_1$ states offer near isotropic repulsive van der Waals interactions which will be useful for Rydberg dressing [21]. The presence of a second valence electron is also directly useful in detection of Rydberg atoms; the second electron can be used in an auto-ionisation transition (at 408 nm) to generate Sr^+ ions, which can then be detected by a microchannel plate detector (MCP).

One further set of transitions of interest are the three repumping transitions from each of the $5s5p\ ^3P_{0,1,2}$ states to the $5s6s\ ^3S_1$ state at 679, 688 and 707 nm, respectively. These offer the ability to move population around within the triplet manifold, where each of the three states could be the desired target depending upon the application.

The energy level structure of strontium gives a toolbox of atomic transitions with applications in laser cooling, atomic clocks, Rydberg physics, imaging, and detection. We have a suitable atom to use, and now move on to discussion of the other key components of tweezers, clocks and Rydberg states.

2.2 Optical tweezers

There are well-established techniques for the preparation of a cold sample of strontium atoms, through the use of magneto-optical traps (MOTs). Strontium atoms are cooled to temperatures of $< 1\ \mu\text{K}$, and localised to a region of a few hundred microns across. The atoms in a MOT have continuously changing positions and separations, and the atoms must scatter photons from the MOT beams to remain trapped. To overcome these limitations we must then load atoms into optical tweezers, where the atoms are trapped in a conservative far-off resonance potential that does not depend on scattering photons (which cause recoils of the atoms), and works

⁴In this work we exclusively use the spin triplet Rydberg states.

on much smaller length scales. The atoms in the tweezer array also then have controlled positions, with the geometry and separation of the array of atoms set by the tweezer array.

In this section we will introduce how trapping in far off-resonance potentials such as optical tweezers and optical lattices works, before contrasting the two approaches. The wavelengths of trapping light used will be discussed, including the the concept of magic wavelength tweezers which will be required for building an atomic clock in the tweezer array.

2.2.1 Trapping in far off-resonance potentials

Trapping atoms in far off-resonance optical potentials is a common experimental technique, where these potentials are called dipole traps [81]. Two common dipole trap configurations are optical tweezers [35, 42, 44, 45, 82] and optical lattices [8, 83, 84]. Here the light used is very far off-resonance ($|\Delta| \gg \Omega$)⁵ from the atomic transitions, such that there is very little scattering of light by the atoms. Instead the laser field acts to perturb the energy levels of the atoms through the AC Stark effect. The energy levels of the atom shift in the presence of the far off-resonant light in the form

$$U = \frac{-\text{Re}[\alpha]I}{2\epsilon_0 c}, \quad (2.2.1)$$

where $\text{Re}[\alpha]$ is the real part of the polarisability (α) of the atomic state (energy level) in question, I is the intensity of the laser field at the position of the atoms, and all other values are fundamental constants (ϵ_0 is the permittivity of free space, and c is the speed of light).

The conservative potential experienced by the atoms is then a function of α and I . The polarisability α depends on both the electronic state of the atom, and the wavelength of the off-resonant trapping light (λ_{trap}). The polarisability of state $|i\rangle$

⁵Where the magnitude of the detuning Δ is much greater than the Rabi frequency Ω .

is given by⁶

$$\alpha_i = 6\pi\epsilon_0 c^3 \sum_j \frac{A_{ij}}{\omega_{ij}^2(\omega_{ij}^2 - \omega_{trap}^2)}, \quad (2.2.2)$$

where $|j\rangle$ are other states of the atom, ω_{ij} and A_{ij} are the transition frequency and decay rate between states $|j\rangle$ and $|i\rangle$ [72], and $\omega_{trap} = \frac{2\pi c}{\lambda_{trap}}$. Assuming the laser field is spatially varying, I depends upon position. Where there is a spatially varying potential, there is a force acting along the potential gradient,

$$F(\vec{r}) = -\vec{\nabla}U(\vec{r}) = \frac{\text{Re}[\alpha]}{2\epsilon_0 c} \vec{\nabla}I(\vec{r}), \quad (2.2.3)$$

where quantities dependant upon the position (\vec{r}) are marked as such.

For positive values of the polarisability α , atoms are attracted to the local maxima of the light intensity. For negative values of α , atoms are repelled from the off-resonant laser light, and instead seek out a local minima of the light intensity. We will always use positive values of α for trapping the atoms, as is the typical case for trapping in tweezers and lattices.⁷ We will discuss choices of trapping wavelength in detail later in section 2.2.3. The trap depth of the ground state, U_0 , is given by equation 2.2.1 and substituting in the ground state polarisability for α and the maximum intensity $I = I_0$ of the laser field. Trap depths are often quoted in temperature units, where the value then refers to $\frac{U_0}{k_B}$.

2.2.2 Optical tweezers vs optical lattices

In the previous section we discussed how the AC Stark shift due to off-resonant laser light can be used to generate a trapping potential which depends upon the gradient of the laser field $\vec{\nabla}I(\vec{r})$. However as we are yet to define the form of $I(\vec{r})$, this tells us nothing about the size of the confinement region. For tight (micron scale) confinement, we contrast two options, optical tweezers and optical lattices.

⁶Higher order shifts such as the vector and tensor polarisabilities and the hyperpolarisability are left for appendix E.2.

⁷We note that it is also possible to trap when the polarisability is negative through the use of bottle beams traps [62, 85].

A basic optical tweezer is a tightly focused Gaussian beam produced by a high numerical aperture (NA) lens, which has an intensity distribution given by

$$I(r, z) = I_0 \left(\frac{w_0}{w(z)} \right)^2 e^{-2\left(\frac{r}{w(z)}\right)^2}, \quad (2.2.4)$$

where the waist⁸ of the beam $w(z)$ has size w_0 at the focus of the beam. The waist is a function of longitudinal position, given by

$$w(z) = w_0 \sqrt{1 + \left(\frac{z}{z_R} \right)^2}, \quad (2.2.5)$$

where z_R is the Rayleigh length of the tweezer,

$$z_R = \frac{\pi w_0^2}{\lambda_{\text{trap}}}. \quad (2.2.6)$$

Optical tweezers are tightly focused to sizes of order a few microns, providing three-dimensional spatial confinement. The two transverse directions are the most strongly confined, with a length scale dependent upon the waist w_0 . The longitudinal direction has much weaker confinement dependant upon the Rayleigh length z_R , though this is typically within an order of magnitude for the length scales in which we work. Optical tweezers have larger trapping frequencies, which are given in the radial direction by

$$\omega_r = \sqrt{\frac{4U_0}{mw_0^2}}, \quad (2.2.7)$$

and in the axial direction by

$$\omega_r = \sqrt{\frac{2U_0}{mz_R^2}}, \quad (2.2.8)$$

where m is the mass of a strontium atom.

Optical lattices typically use much larger beams, retro-reflected to produce a standing wave. In the case of the one dimensional optical lattice formed by a retro-reflected beam, a series of ‘pancakes’ are formed, where atoms can be trapped in discs separated by $\frac{\lambda_{\text{trap}}}{2}$.⁹ This confinement is smaller than the typical waists one would expect when working in an optical tweezer of the same trapping wavelength.

⁸The waist is the $\frac{1}{e^2}$ radius of the intensity of the beam.

⁹In other configurations this spacing can be controlled via the angle of the beams.

Optical lattices can be used in the one dimensional case [8, 83], but this approach generalises up to three dimensional lattices [84, 86, 87]. A variety of lattice geometries exist, including square/rectangular, triangular/hexagonal, honeycomb and Kagome lattices.

Optical lattices offer closer separations between adjacent trapping sites than optical tweezers. Retro-reflected lattices have separations of only $\frac{\lambda_{\text{trap}}}{2}$, whereas tweezers are Abbe diffraction limited to $\frac{\lambda_{\text{trap}}}{2NA}$ by the NA of the lens that focuses the Gaussian beam. High NA lenses that are placed even on the order of ~ 10 mm away from the atoms only have NAs of ~ 0.5 , so lattices have closer site spacing. 3D lattices can generate many more sites than are possible with tweezers, and have better 3D confinement as they do not share the long Rayleigh lengths of tweezers. The long Rayleigh length in tweezers this can be mitigated in some scenarios, for example by crossing two orthogonal propagating tweezers [88].

On the other hand, tweezer arrays offer far more tunability than is available in a lattice setup. Lattices, once constructed, exist in more of a fixed geometry, whereas tweezers can be placed at arbitrary and re-configurable separations and geometries [35, 56, 89]. Individual tweezers can have their trap depth changed [48], and be independently turned off or moved during an experimental sequence [37, 55, 56, 90]. The methods to produce arrays of tweezers will be discussed when required in section 5.2.

Both tweezers and lattices offer their own benefits and drawbacks, but for our application tweezers are the more appropriate choice. To implement a spin-squeezed protocol we will need to use Rydberg interactions which have characteristic length scales on the order of microns, where we need to be able to control the separations of the array, as well as the geometry. We also require the ability to measure the states of individual atoms in neighbouring trapping sites, where the spacing of the tweezer array can be chosen to allow good resolution. Alternatively tweezer arrays and lattices can be combined together, as has recently been demonstrated [22].

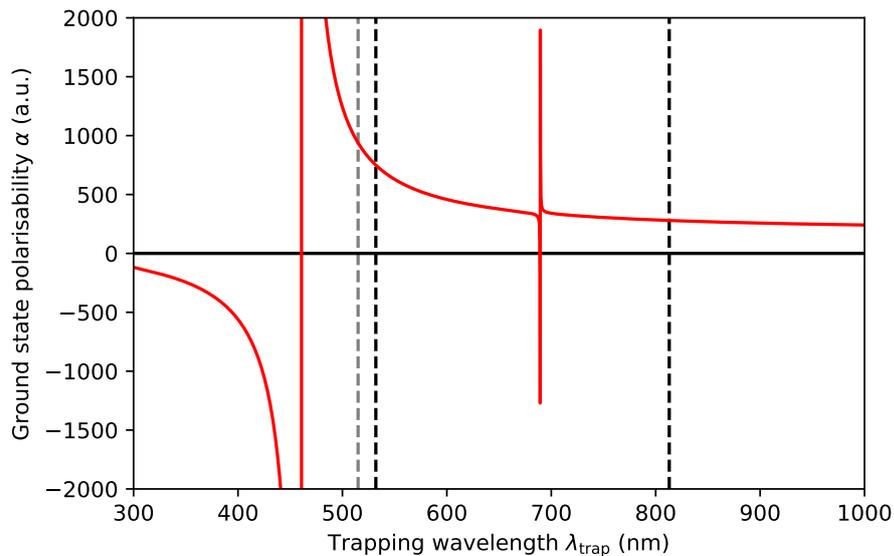


Figure 2.2: Polarisability of the ground state ($5s^2\ ^1S_0$) as a function of trapping wavelength is shown in red. Atomic units (a.u.) can be converted to SI units by multiplication with $4\pi\epsilon_0 a_0^3$ (where ϵ_0 is the permittivity of free space, and a_0 is the Bohr radius). The black dashed lines show the trapping wavelengths used in this thesis, 532 nm and 813 nm. The grey dashed line shows 515 nm, a trapping wavelength used by other research groups [44, 45] which acts similarly to our tweezer at 532 nm. At all of these trapping wavelengths the polarisability of the ground state is positive.

2.2.3 Choice of trapping wavelengths

To trap strontium in its ground state, one must choose a sensible trapping wavelength, λ_{trap} . Figure 2.2 shows that the polarisability (α) of the ground state is dominated by the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition at 461 nm, due to the large transition decay rate on that transition. In general the polarisability is positive if red-detuned from the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, and negative if blue-detuned from it. Trapping wavelengths in the green part of the spectrum, such as 515 nm and 532 nm, result in larger polarisabilities, whereas wavelengths in the infrared (IR), such as 813 nm, result in smaller but still positive polarisabilities. In this work we use 532 nm and 813 nm tweezers in chapters 4 and 5 respectively. Here we contrast the drawbacks and benefits of working with each system.

The first comparison to make is the difference in ground state polarisability. At 532 nm the ground state polarisability is ~ 2.7 times larger than that at 813 nm [72], leading to an increase in trap depth on the same scale for equal laser power and beam shape at each wavelength. The contrasting benefit for the 813 nm tweezer is that the trapping light is further off-resonant from the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, so the scattering of the trapping light from the atoms is reduced by an order of magnitude.

The shorter wavelength of 532 nm offers tighter confinement for reasons beyond just the polarisability. The diffraction limited tweezer waist available is smaller for shorter wavelengths, where the waist is proportional to the wavelength as given by the Abbe diffraction limit. The tighter confinement leads to an increase in tweezer intensity proportional to the inverse of the waist squared, so the diffraction limited size of the tweezer changes trap depth as $\frac{1}{\lambda^2}$, giving another factor of ~ 2.3 in trap depth for the 532 nm tweezer compared to the 813 nm tweezer.

Rydberg states of strontium (see section 2.4) are expected to be trapped by linearly polarised dipole traps of wavelengths $422\text{ nm} < \lambda_{\text{trap}} < 590\text{ nm}$, where the polarisability of the inner ion is larger in magnitude than the ponderomotive polarisability of the Rydberg electron [91]. This means that the common trapping wavelengths of 515 nm [44, 45] and 532 nm [43] are expected to be able to trap Rydberg atoms, whereas 813 nm tweezers do not trap Rydberg atoms, as has been experimentally demonstrated [91, 92]. Therefore when working with Rydberg atoms at $\lambda_{\text{trap}} = 813\text{ nm}$, the trap should either be turned off [92], or be shallow [22].¹⁰

So far, on the whole, it appears preferable to work with tweezers at shorter wavelengths, as we use in chapter 4, but there are also significant advantages to working with a longer wavelength tweezer. The driving reason for the use of 813 nm tweezers is that it is a magic wavelength for the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ clock transition [33, 34].¹¹ Each state of the atom has its own value of the polarisability, α , leading

¹⁰Alternatively complex trap geometries such as a bottle beam traps [62, 85] could be used.

¹¹There are other magic wavelengths, for example at 390 nm [93], though both the $5s^2\ ^1S_0$ and $5s5p\ ^3P_0$ states are repelled by the 390 nm light. Throughout this work we will refer to 813 nm as the magic wavelength.

to each state having a different AC Stark shift (see equation 2.2.1). The difference between the polarisabilities of two states α_i and α_j ¹² leads to a change in the energy difference between the two states, and thus changes the transition frequency. The energy difference between the two states is then given by

$$\Delta E = \frac{\text{Re}[\alpha_i(\lambda) - \alpha_j(\lambda)]I(\vec{r})}{2\epsilon_0 c}. \quad (2.2.9)$$

This effect is called the differential AC Stark shift. Where the light causing the shift has a spatial dependence (as it does in a tweezer $I(\vec{r})$), this leads to a different transition frequency as a function of position.

As we will see in section 2.3, when making an atomic clock it is important not to shift the clock transition away from its natural frequency where it can be avoided, and certainly not in a way that depends upon trap depth or position of the atoms! To avoid this we can choose the magic wavelength where $\alpha_i(\lambda) = \alpha_j(\lambda)$ and the differential AC Stark shift $\Delta E = 0$.¹³ Note this only is the magic wavelength for the $|i\rangle \rightarrow |j\rangle$ transition, in this case the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ transition. For all other transitions we expect to see a different, non-zero, differential AC Stark shift. The difference between working in a tweezer with non-zero and zero differential AC Stark shift is shown in figure 2.3, where at the magic wavelength of 813 nm the transition frequency does not depend upon position in the tweezer. Avoiding shifting the clock transition is the main motivation for the switch to 813 nm tweezers in chapter 5. There are other additional benefits to working in 813 nm tweezers, but these will be discussed later as required in section 5.1.1.

2.3 Atomic clocks

Strontium has an ultra-narrow $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ optical clock transition, making it well-suited to use in atomic clocks. In this section we introduce atomic clocks as a concept, and the essential components required to make one. We detail why

¹²Here state i is taken to be at a lower energy than state j , to provide the correct sign convention.

¹³At the magic wavelength, the differential AC Stark shift ‘vanishes’, and does not depend on trap depth or position of the atoms.

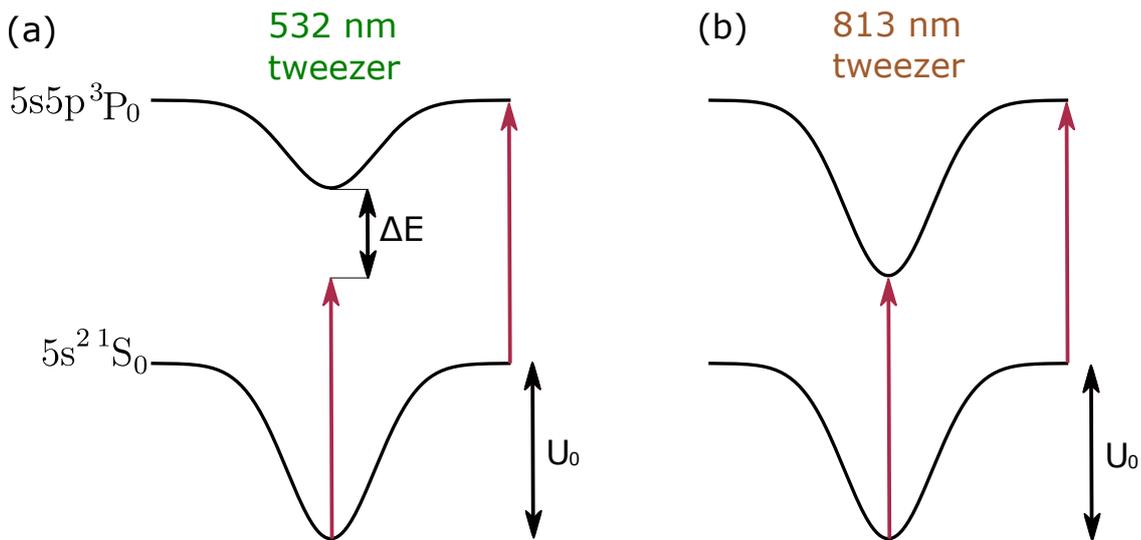


Figure 2.3: Diagram showing the effects of the differential AC Stark shift on the clock transition. The black curves show the tweezer potential for each state at each trapping wavelength, with the horizontal direction corresponding to a change in position, and the vertical direction corresponding to a change in energy. (a) A non-zero differential AC Stark shift on the clock transition for a tweezer at 532 nm, the transition frequency changes as a function of the position of the atoms in the tweezer. (b) A tweezer at the magic wavelength of 813 nm, where the differential AC Stark shift $\Delta E = 0$ and the clock transition frequency is unchanged with position. This figure is not to scale, in particular for the separation of the energy levels and the relative trap depths, and is for conceptual understanding only.

optical clock transitions such as the one found in strontium are far exceeding the performance of microwave clocks with which the second is currently defined, and the tools required to instead work with optical clock transitions. We then contrast different options for the implementation of optical atomic clocks, including between ^{88}Sr used in our experiment and ^{87}Sr which has shown the smallest fractional frequency measurement precision [8] in an atomic clock. We introduce MIS [75], an essential part of allowing ^{88}Sr to be used in an atomic clocks. Finally we discuss the QPN limit for clocks based upon uncorrelated atoms, motivating the generation of correlations between the atoms through Rydberg dressing.

2.3.1 Components of an atomic clock

There are three main ingredients in making a clock [10, 71, 94, 95] :

1. A local oscillator,
2. A frequency reference,
3. A counter.

The local oscillator produces the periodic signal, where each period of oscillation is one ‘tick’ of our clock. The frequency reference is something with which we keep the local oscillator at the correct frequency, so ‘ticking’ at the correct rate. Finally, the counter keeps track of the number of ‘ticks’ of the local oscillator, essentially keeping a record of the time. In an atomic clock, the local oscillator is a short-term stable microwave oscillator or laser, which is stabilised against the frequency reference of an atomic transition in some atomic species, where the oscillations are then counted using a frequency counter.

The key feature of an atomic clock is, as the name might suggest, the atoms. The atoms can be used as a frequency reference due to the energy level splitting between two atomic states $|i\rangle$ and $|j\rangle$ via $\Delta\nu = \frac{E_j - E_i}{h}$, where E_i and E_j are the energy levels of the two clock states, and h is Planck’s constant. The microwave oscillator/laser is then stabilised to this $|i\rangle \rightarrow |j\rangle$ transition in the atomic reference. Atoms act as an great frequency reference, because all the atoms of the same isotope

are identical. This means that aside from environmental perturbations, an atomic clock based on the energy level splitting between two states in an atom¹⁴ should have the same frequency as a clock based on the same states in another experiment elsewhere. This makes atoms a reproducible frequency reference wherever you are, subject to environmental perturbations shifting the clock transition frequency.

Atomic transitions suitable for use as frequency references in atomic clocks are those with large quality factors (Q). The quality factor is the ratio of the transition frequency and the linewidth of the transition,

$$Q = \frac{\Delta\nu}{\gamma}, \quad (2.3.1)$$

where $\gamma = \frac{\Gamma}{2\pi}$ is the transition linewidth, and Γ is the decay rate from the higher energy state. The narrower γ is relative to $\Delta\nu$, the better precision with which we can determine a frequency using the clock in a given time. References with large values of Q tend to make good clocks, and atomic clock transitions are reproducible references that come with very large quality factors.

2.3.2 From microwave to optical

The standard of modern timekeeping is defined by the atomic clock. The high level of precision of atomic clocks means that the SI unit¹⁵ of the second is currently defined by ¹³³Cs atomic clock based on a microwave transition. Caesium clocks have been a demonstrated method since 1955, where they first showed fractional frequency uncertainties on the scale of $\sim 1 \times 10^{-9}$ [96], beating the best achievable with astronomical observations [94]. An uncertainty of $\sim 1 \times 10^{-9}$ corresponds to an error of roughly one second in 30 years, so less than a second in my life so far.

Since the original Cs clock, our ability to make these measurements has vastly improved of course. State of the art caesium fountain clocks have fractional frequency uncertainties of $\sim 2.3 \times 10^{-16}$ [97, 98], corresponding to an error of one second in ~ 140 million years, or in other words an error of one second since the end

¹⁴Of the same isotope.

¹⁵International System of Units.

of the Jurassic period¹⁶ [99]. The accuracy to which experimental realisations of the second can be made significantly exceeds the measurements of other fundamental units, such that the definition of all but one of the other SI base units are based upon the definition of the second. Atomic clocks can demonstrate fractional uncertainties far below that achieved using caesium, and for that we now look to optical atomic clocks.

Clocks based on alkali elements like caesium utilise microwave transitions between the hyperfine split ground state levels, utilising the narrow linewidth γ of the transition. To make a better clock, we want to find a transition with a higher quality factor Q . In order to do this we either want to decrease the linewidth, or increase the transition frequency. Inspecting the atomic structures, the natural route is then to move to optical transitions (~ 100 s THz) instead of the much smaller (~ 9 GHz) hyperfine microwave transition, an improvement of over four orders of magnitude in transition frequency. To find optical transitions with narrow linewidths, we must look past monovalent atoms and instead to divalent atoms and also to ions, which offer linewidths ranging from order Hz to narrower than a mHz [10, 12]. Suitable divalent atoms include Sr [3, 4, 8, 83, 100–103], Yb [6] and Hg [104]. Suitable ions include Al⁺ [7], Sr⁺ [105, 106], Yb⁺ [107, 108],¹⁷ Hg⁺ [109], Ca⁺ [110], In⁺ [111] and some highly charged ions [112]. Beyond optical transitions in the atoms and ions discussed here, there is also outlook to even higher frequencies, considering nuclear transitions in Th [113, 114].

Optical clock transitions offer a far superior frequency reference to microwave transitions. The change of local oscillator to accommodate this is obvious, we now must use optical frequencies rather than microwaves, so we use a laser rather than a microwave oscillator. The more complex part is the need for the third major component of a clock, the counter. Counting at microwave frequencies with electronics is long established and works highly effectively, but we need to be able to count

¹⁶To provide context in terms of dinosaurs, even if a Tyrannosaurus rex had somehow managed to construct a Cs clock of this quality, it would have only picked up about half a second of error by now!

¹⁷Yb⁺ has two clock transitions.

at optical frequencies in the 100s of THz. Electronics cannot handle frequencies so high, as they work well in the microwave and radio regime. Infrared and optical frequencies are typically manipulated with optics, we are on the wrong side of the THz gap [115, 116]. The solution to this is to beat the clock laser (local oscillator) with a nearby optical frequency, such that the beat note is in the range of frequencies where typical frequency counters work. We then need a device that can produce a nearby laser frequency with suitable stability and well-determined frequency.

The breakthrough that allowed the use of these optical clocks was the invention of the optical frequency comb¹⁸ [117, 118]. Optical frequency combs use a modelocked laser to generate a series of laser frequencies, with well defined spacing and absolute frequency. Beating one of the laser frequencies from a frequency comb with a clock laser will typically produce a beat frequency in the 10s of MHz, which is well within the working range of standard frequency counters. The optical frequency comb will be discussed in more detail in chapter E, including the specifics of our system. The optical atomic clock now has all its constituent parts, a clock laser as the local oscillator, a narrow linewidth optical transition as the frequency reference, and a frequency comb used with a frequency counter as the counter.

Experimentally realised optical atomic clocks now far surpass the fractional uncertainties of microwave clocks by orders of magnitude. As mentioned earlier, for the definition of the second based on ^{133}Cs there are clocks operating at fractional uncertainties of $\sim 2.3 \times 10^{-16}$ [97, 98], corresponding to an error equivalent to a second since the end of the Jurassic period. Three species of optical clocks have now passed below the $\sim 2 \times 10^{-18}$ fraction frequency uncertainty level corresponding to below a second in the age of the universe, Sr [3, 4, 8, 100, 101], Yb [6] and Al^+ [7]. The best fractional frequency precision measured to date is between two uncorrelated regions of an ^{87}Sr sample, at 7.6×10^{-21} using ^{87}Sr [8], roughly equivalent to the wingbeat of a honeybee [119] in the age of the universe. This also demonstrates that strontium is an very attractive element to work with for atomic clocks, given the

¹⁸The optical frequency comb may be abbreviated to OFC, may alternatively be referred to as a femtosecond comb, and will often be shortened simply to either frequency comb or just ‘comb’ for convenience.

best measurement in the world was taken with ^{87}Sr , albeit we work with a different isotope, ^{88}Sr .

2.3.3 Atoms, ions and isotopes

Given the wide range of options to make optical atomic clocks, we now motivate the choice of ^{88}Sr in this thesis, and discuss the particulars of using it. It is worth noting to begin with that there is not necessarily a best candidate for making a new definition of the SI second with optical atomic clocks, proposals for making a new definition include taking weighted averages of different atomic clock transitions [17–19]. However for the sake of an individual platform you must make a choice of the species that you are working with.

As we wish to have multiple atoms within a small volume to demonstrate spin-squeezing via Rydberg dressing in a tweezer array, we need to work with atoms rather than ions. Ion clock experiments typically use the clock transition of a single ion in a Paul trap [7, 120, 121], though it is possible to use multiple ions [122]. Atoms such as strontium also have well-characterised Rydberg states, which will be important for them to be used in Rydberg dressing. As for optical clocks made with atoms, the typical remaining choices are Sr, Yb and Hg. Reviewing the properties of each of these three elements [10, 12], one can see that many of the relevant laser wavelengths¹⁹ for Hg are in the ultraviolet (UV), making it a more difficult environment to realise an experimental system in. Finer details in the comparison of Sr, Yb and Hg can be found in the literature [10]. From here we decide to make an atomic clock from strontium atoms, the final remaining choice being which isotope should be used.

The choice of strontium isotopes for use in atomic clocks is mainly between the fermionic isotope ^{87}Sr , and the bosonic isotope ^{88}Sr . Each has its benefits and drawbacks relative to the other. In terms of experimentally realised performance, ^{87}Sr has the superior measured fractional frequency uncertainty [4], much better than progress in ^{88}Sr [102]. ^{88}Sr has by far the highest natural abundance, > 10

¹⁹Including laser cooling, the clock transition, and the magic wavelength.

times more common than ^{87}Sr , making loading with ^{88}Sr faster. The key difference between the two isotopes is the nuclear spin. The bosonic isotopes like ^{88}Sr have nuclear spin $I = 0$, whereas ^{87}Sr has nuclear spin $I = \frac{9}{2}$. Having $I = 0$ in ^{88}Sr greatly simplifies the experimental setup; there is no hyperfine structure to worry about, so there is no need for the addition of a ‘stirring’ beam²⁰ [10]. The lack of hyperfine structure also removes terms from the uncertainty budget, which sets the linear Zeeman shift to zero, as well as setting the vector and tensor parts of the polarisability to zero which removes some of the shifts from the trapping light. It is worth noting however, that strategies already exist to counter these non-zero terms in ^{87}Sr , such as combining two clock servos with two samples of atoms in the $m_f = \pm\frac{9}{2}$ states, and using a sensible choice of the lattice polarisation along with an effective magic wavelength [10, 123].

The difference in nuclear spin (I) between ^{87}Sr and ^{88}Sr causes another major difference between the two isotopes beyond just hyperfine structure, it has a drastic impact on the linewidth of the clock transition. In ^{88}Sr the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ transition is heavily forbidden by selection rules²¹, so has a natural lifetime of thousands of years [124]. It is not practical to excite this transition directly with a laser when the linewidth is so exceptionally narrow. In ^{87}Sr things are very different, the hyperfine interaction leads to mixing of other states into the clock transition, primarily a small fraction of the $5s5p\ ^3P_1$ state [125]. This leads to a linewidth on the clock transition of ~ 1 mHz, with a lifetime > 100 seconds [126].

The drastic difference in clock linewidths for the two isotopes could naively lead to the conclusion that ^{87}Sr would be the only option, and that hyperfine interactions to mix in some of the $5s5p\ ^3P_1$ state into the $5s5p\ ^3P_0$ state are a necessity. There is, however, a way to also achieve a similar state mixing in ^{88}Sr to artificially broaden the linewidth to become a practically viable clock transition. This method is called

²⁰When using ^{87}Sr , a stirring beam is applied for spin-polarisation into a single magnetic sublevel prior to clock interrogation. It is also used to randomise the spin in the red MOT stage to improve the MOT lifetime. [10].

²¹The clock transition is spin forbidden ($\Delta S = 1$), and angular momentum forbidden ($J = 0 \rightarrow J = 0$) to all orders of the multipole expansion [10].

magnetically induced spectroscopy (MIS), where an external magnetic field is applied to mix in a small fraction of the $5s5p\ ^3P_1$ state [75].

MIS [75] was first demonstrated in ^{174}Yb [127, 128], before being demonstrated in ^{88}Sr [129]. The applied magnetic field allows a Rabi frequency on the clock transition, Ω_{clock} of

$$\frac{\Omega_{\text{clock}}}{2\pi} = \alpha\sqrt{IB}\cos\theta. \quad (2.3.2)$$

α is a constant depending on the element being used²², here $\alpha = 6.26 \frac{\text{Hz}}{\text{mT}\sqrt{\text{Wcm}^{-2}}}$ [10, 75]. I is the intensity of the clock laser light²³, B is the uniform external magnetic field used to allow the clock transition, and θ is the angle between the clock laser polarisation and the magnetic field direction. To maximise the Rabi frequency for the smallest possible shifts induced by the magnetic field and the clock laser, the polarisation of the clock laser should be parallel with the magnetic field direction. In order to drive the clock transition in ^{88}Sr , one typically uses magnetic fields of the order of $\sim 1\text{ mT}$ ²⁴ [10, 46, 47, 94]. The need to apply these magnetic fields and large probe power will add to the uncertainty budget, where the typical magnetic fields applied are much larger than the much smaller and well-characterised bias fields needed to resolve the hyperfine levels in ^{87}Sr [10]. The good news though, is that with suitable application of MIS, ^{88}Sr is a viable isotope for optical atomic clocks. MIS is considered in more detail in appendix E.2, including systematic shifts due to the applied magnetic field and probe light.

2.3.4 Quantum projection noise

When describing atomic clocks, we have not yet discussed the specifics of how the local oscillator (laser/microwave oscillator) is actually stabilised to the frequency reference (atomic transition). How this works brings us towards how one experimentally arrives at a key factor that limits atomic clock performance, QPN [20].

To stabilise the clock laser to the atomic reference, an error signal must be

²² α here is not a polarisability.

²³Apologies for using I as both the nuclear spin and clock laser intensity in the same section!

²⁴1 mT = 10 G.

produced such that when the laser drifts away from the atomic transition there is feedback to stabilise it. Measuring the clock transition directly is challenging due to how weak the transition is, so instead electron shelving is typically used. Light from the clock laser is applied to the atoms to prepare them in a superposition of the ground and clock states. Application of 461 nm light on the strong $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition projects each atom into either the ground or clock state in a quantum jump. Fluorescence from the atoms that were projected into the ground state²⁵ then provide a measure of the ground state atom number. The atoms in the ground state are then blown away by the 461 nm light, before atoms from the clock state are brought back to the ground state and are imaged. This gives a measure of both the ground state and clock state populations [10, 94, 130]. Alternatively, if cooling is used to compensate heating from the imaging light, then the atoms can be imaged multiple times without significant loss. After measuring the ground state population, atoms from the clock state can be repumped to the $5s5p\ ^3P_1$ state such that they can decay to the ground state, then the entire atomic population can be imaged. This gives a measure of the ground state and total populations, and has the benefit of the atomic sample being reusable [131].

The error signal must steer the laser frequency in the correct direction such that it kept near the atomic resonance. To do this, one must measure in a pair of sequences on either side of a resonance, such as a Rabi π -pulse lock or a Ramsey lock [10, 132]. The detection scheme should be set such that the measurements aim to occur around an equal superposition of the ground and clock states for both measurements, as that produces a signal of the greatest slope and sensitivity to shifts in the clock laser frequency. Information from these measurements can then be fed back to the laser to steer it towards the correct frequency.

The need to make a projective measurement provides a limit on the achievable clock performance; there will always be a variance in the measured fraction of atoms to be excited to the clock state. For N_{at} atoms measured, the excitation fraction F

²⁵Atoms projected to the ground state cycle on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition and fluoresce, atoms projected to the clock state do not interact with the 461 nm light and remain dark.

with mean fraction F_0 will have a variance

$$(\Delta F_{\text{QPN}})^2 = \frac{F_0(1 - F_0)}{N_{\text{at}}} \approx \frac{1}{4N_{\text{at}}}, \quad (2.3.3)$$

where the interrogation scheme has $F_0 \approx \frac{1}{2}$. This leads to a fundamental uncertainty limit (the QPN limit) based on a measurement of N_{at} uncorrelated atoms, which we reproduce from earlier in this chapter:

$$\sigma(\tau) = \frac{\beta}{Q\sqrt{N_{\text{at}}}} \sqrt{\frac{T_c}{\tau}}. \quad (2.0.1)$$

There are many other factors that limit clock performance, such as the Dick effect that arises from dead time between measurements [10, 71], but as clock experiments improve, QPN is fast approaching becoming the limiting factor, and is already a limit in some low N_{at} systems.

When quantum projection noise becomes a significant limiting factor, one must then consider what can be done to improve clocks further. Increasing the time of the clock laser interrogation is limited by the coherence time. The achievable number of atoms that can be interrogated is also limited, only provides benefit with a $\frac{1}{\sqrt{N_{\text{at}}}}$ scaling. We instead look to how to stop quantum projection noise being a limiting factor, and for that we have to go back to the concept that QPN is the limit for N_{at} *uncorrelated* atoms. We therefore want to generate correlations between the atoms, where for that we plan to use Rydberg dressing [21].

2.4 Rydberg atoms

Strontium atoms in tweezer sites separated by several microns do not typically interact with each other. In order to generate correlations between atoms in different tweezer sites to implement a spin-squeezing protocol, we require a method to introduce long-range interactions between atoms in different tweezers in a controllable way. These long-range interactions on the scale of several microns are readily achievable with the use of Rydberg states, where Rydberg dressing is used to generate the spin-squeezing in a tweezer array clock.

Rydberg atoms are atoms which have an electron in a highly excited electronic

state, where the principal quantum number (n) is large²⁶. The long-range interactions between atoms allow the generation of entanglement, as well as Rydberg-blockades and Rydberg dressing. In addition to this, Rydberg atoms have high sensitivity to the environment, where we can use them as sensors of electric fields for managing uncertainty in atomic clocks [134] and to enable precision measurements.

Other applications of Rydberg atoms include quantum simulation [38, 61, 135] and quantum computation [136]. The wide variety of applications of Rydberg states encourages the precision measurement of their energies and study of interactions between them. A tweezer array is a great platform with which to study this, due to the isolation of single atoms and the controllable separations between atoms. Improved precision measurement of Rydberg states [65, 66] would advance modelling of Rydberg atoms [1, 67, 68], with potential applications even including tests of fundamental physics [70]. Chapter 6 will consider electric field control through Rydberg electrometry, both for managing the systematic uncertainty in an atomic clock, and in the precision measurement of Rydberg states.

2.4.1 Rydberg energy levels

Unlike in hydrogen, the Rydberg energy levels of strontium cannot simply be given by a normal Rydberg energy series $E_n = \frac{-R_y}{n^2}$ [137], where the Rydberg constant R_y is defined in terms of fundamental constants [138]. This arises because the Rydberg valance electron can still penetrate the core shells for lower values of orbital angular (l), so they experience a modified potential. To account for this, the energy levels in strontium can be approximated by the Rydberg-Ritz formula [1, 66, 68, 139, 140],

$$E_n = \frac{-R_y}{(n - \delta)^2}, \quad (2.4.1)$$

²⁶In this thesis we primarily work with $n \sim 80$, and in general when discussing Rydbergs we are talking about $n \gtrsim 30$ [40, 80]. Lower values of n can be still considered Rydberg states though depending upon the system, such as demonstrated at $n \sim 13$ in Terahertz imaging of Cs [115, 116] and at $n \lesssim 10$ in cuprite excitons [133].

where the quantum defect δ is a function of the principal quantum number n and the particular Rydberg series,

$$\delta = \delta_0 + \frac{\delta_2}{(n - \delta_0)^2} + \frac{\delta_4}{(n - \delta_4)^4} + \dots, \quad (2.4.2)$$

where δ_0 , δ_2 and δ_4 depend upon the Rydberg series in question. To calculate the energy levels, the values for the quantum defect can be retrieved from literature, [1, 66, 68, 80, 140], giving us an estimate of the energies of the Rydberg states. A particularly useful tool for performing these calculations is ARC 3.0 [1],²⁷ which can be used to do a wide range of calculations on divalent Rydberg atoms. The calculated energy levels from the Rydberg-Ritz formula can be used to calculate the appropriate frequency for the lasers in the excitation scheme up to the Rydberg state.

2.4.2 Rydberg properties

Rydberg atoms have very different properties to atoms in lower electronic states, in particular having much stronger interactions with each other and the environment. Many of their properties scale with high powers of n , as is documented elsewhere in the literature [40, 80]. The key property for the goals of this thesis are the long-range van der Waals interactions between Rydberg atoms that scale as $V \propto n^{11}$, giving interactions between atoms separated by several microns. An additional use of Rydberg atoms is their sensitivity to the electric field environment, scaling the polarisability of the atoms as $\alpha \propto n^7$.

2.4.3 DC Stark effect

The application of an electric field $\vec{E} = E\hat{r}$ with amplitude E in direction \hat{r} modifies the Hamiltonian an atom with the addition of a Stark interaction term, perturbing the energy levels and mixing the bare atomic states²⁸. The DC Stark shift from an

²⁷Alkali(ne) Rydberg Calculator

²⁸The eigenstates of the total Hamiltonian at a particular field \vec{E} are linear superposition of bare atomic states [76].

applied electric field (E) is given by

$$\Delta_{\text{DC}} = \frac{-\alpha E^2}{2}, \quad (2.4.3)$$

in the same form as the AC Stark shift in equation 2.2.1. The difference here is that the electric field arises from the constant stray/applied electric field, rather than an oscillating electric field arising from laser light. Rydberg atoms have high sensitivity to the electric field environment through the DC Stark effect, due to their polarisabilities being very large ($\alpha \propto n^7$), making them ideal for use as electric field sensors. The energy levels of the Rydberg atoms are perturbed by the DC Stark effect, where the shift of the energy level can be used to characterise the electric field [71, 76, 134, 141] not only with high sensitivity, but measured at precisely the position of the atoms.

Measuring the electric field at the position of the atoms is extremely useful, as it is otherwise difficult to measure the precise electric field at particular positions inside of a vacuum chamber. Measuring the electric field using the atoms is also particularly useful in a tweezer array, where we can measure the electric field in very specific regions on the scale of microns, unlike the larger volumes used when performing electrometry experiments using atoms released from MOTs [134, 141].

2.4.4 Rydberg blockade and dressing

The long-range van der Waals interactions between Rydberg atoms are of the form

$$V = \frac{C_6}{r^6} \quad (2.4.4)$$

where the C_6 coefficient is a property of the Rydberg state and scales roughly as $C_6 \propto n^{11}$, and r is the separation between atoms. These long range interactions are useful for generating entanglement between atoms [92, 142–144], and for applications in quantum simulation [38, 61, 62].

Interactions between Rydberg atoms can be so strong that they form the Rydberg blockade [145, 146], where the presence of one Rydberg atom prevents the excitation of another atom to the Rydberg state if it is within the blockaded volume, i.e. within the blockade radius

$$R_b = \left(\frac{|C_6|}{\hbar\Omega} \right)^{\frac{1}{6}}, \quad (2.4.5)$$

where Ω is the Rabi frequency [79]. This occurs as the presence of the first Rydberg atom perturbs the Rydberg transition in the second atom by more than the Rabi frequency of the laser.

Along with the Rydberg blockade, there is also access to a regime called Rydberg dressing. Rydberg dressing is where a small fraction of a Rydberg state is mixed into a lower lying state in the form

$$|\tilde{e}\rangle \sim |e\rangle - \epsilon|r\rangle, \quad (2.4.6)$$

where $|e\rangle$ and $|r\rangle$ are the excited clock state and a Rydberg state, and the dressing fraction $\epsilon = \frac{\Omega}{2\Delta_d} \ll 1$. This brings a small amount of Rydberg character, and allows use of Rydberg properties without having to fully excite to the Rydberg state. An example of the advantage of dressed states include an extended lifetime compared to the Rydberg state, increasing proportional to $\frac{1}{\epsilon^2}$ and allowing sufficient coherence time to produce spin-squeezed states. Detailed discussion of Rydberg dressing can be found in previous theses from our laboratory [72, 79, 80]. In the Rydberg dressed regime one uses a larger detuning from the transition $|\Delta_d| \gg \Omega$, resulting in a dressed blockade radius

$$R_d = \left| \frac{C_6}{2\hbar\Delta_d} \right|^{\frac{1}{6}}, \quad (2.4.7)$$

though with a potential different to that of the van der Waal interaction [21, 72, 79, 80, 147] (see equation 2.4.4), instead in the form

$$\begin{aligned} V &= \frac{\hbar\Omega^4}{8|\Delta_d|^3 \left(1 + \left(\frac{r}{R_d} \right)^6 \right)} \\ &= \frac{V_0}{\left(1 + \left(\frac{r}{R_d} \right)^6 \right)}, \end{aligned} \quad (2.4.8)$$

for a pair of Rydberg dressed atoms [72, 79, 80], which is a long-range potential with a soft core, as can be seen in reference [21]. The length scale of these interactions scales as $R_d \propto |C_6|^{\frac{1}{6}} \propto n^{\frac{11}{6}}$, meaning that higher n leads to larger dressed blockade radii. Rydberg dressing with $5sns^3S_1$ states with high n offers near-isotropic long-range interactions on the scale of many microns, with which we can implement spin-squeezing in a tweezer array .

The collective energy shift for N atoms within the dressed blockade volume scales approximately with a factor of $\frac{N(N-1)}{2}$ compared to the pair of atoms case [72, 80]. Maximising the number of atoms within the dressed blockade volume is important to Rydberg dressing schemes for spin-squeezing [21], and is a key motivation for working with tweezer arrays that can be extended to 3D. Rydberg dressing has been demonstrated between pairs of atoms [148], in MOTs [149], optical lattices [150], and optical tweezer arrays [22].

2.5 Summary

We have introduced the key parts needed to implement a spin-squeezed tweezer array clock of strontium atoms: optical tweezers, atomic clocks and Rydbergs. Optical tweezers have been contrasted with optical lattices, as well as contrasting two trapping wavelengths for deep tweezers (see chapter 4) and magic wavelength trapping (see chapter 5). An overview of atomic clocks has been given, including the reasons to work with optical clock transitions such as the one found in strontium, as well as some specifics of what is required to prepare for making an atomic clock with ^{88}Sr (see appendix E). Rydberg atoms have been introduced as incredible sensors of electric field and as something that we want to make precision measurements of (see chapter 6), as well as in the context of using Rydberg dressing to create a spin-squeezed atomic clock.

Chapter 3

The cold strontium experiment

To create an array of strontium atoms for applications such as a spin-squeezed atomic clock, we first need to prepare an ultracold gas of strontium atoms and load the atoms into optical tweezers. A cold strontium experiment requires a large quantity of equipment and experimental techniques. Numerous laser systems are required, along with an ultra-high vacuum apparatus, electronics, detectors, and control software. The apparatus used in this project has been developed over several previous PhDs [72, 73, 76–80], has continued its development during my own time on the project, and will carry on further into the future. This chapter will introduce the key parts of the experiment from the perspective of the work undertaken during the course of this thesis. We first detail the preparation of a cold atomic sample of strontium into tweezers. We then describe the required laser systems, *in-vacuo* equipment, and imaging system in more detail.

3.1 Preparing cold atoms

This section introduces the preparation of cold atoms into optical tweezers. The optical tweezers are formed using a $\text{NA} = 0.26$ *in-vacuo* lens, focusing trapping light at either 532 nm or 813 nm, as shown in figure 3.1. The cold atomic sample must be prepared between the two *in-vacuo* lenses, such that the atoms can then be transferred into the optical tweezers. The same *in-vacuo* lens used to form the

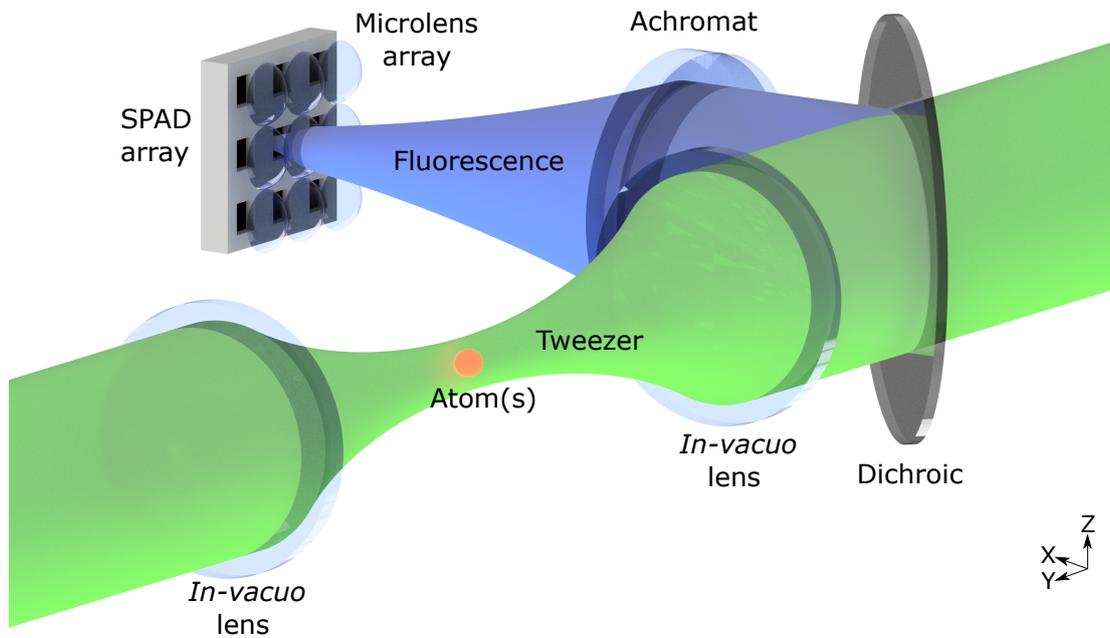


Figure 3.1: Strontium atoms are trapped in a optical tweezer(s) formed by a $\text{NA} = 0.26$ *in-vacuo* lens. Fluorescence is collected from the atoms via the same *in-vacuo* lens used to form the tweezer. The fluorescence is separated from the trapping light using a dichroic, and is then imaged onto the single photon avalanche diode (SPAD) array using a series of lenses. The 3D model in this figure is the work of Ryan Hanley. In chapter 4, a single 532 nm tweezer will be used, and the imaging system will be the $\times 10.8$ magnification imaging system shown here. In chapter 5 the same *in-vacuo* lens will be used to create an 813 nm tweezer array (shown in figure 5.1), where in section 5.3 we implement a higher magnification ($\times 48.6$) imaging system (shown in figure 5.3).

tweezer is also used to image the atoms unless otherwise stated. 461 nm fluorescence from the atoms is collected by the *in-vacuo* lens, before being separated from the trapping light using a dichroic¹, and imaged onto our SPAD array.

We now introduce the procedure for loading cold atoms, beginning with Zeeman slowing an atomic beam, followed by a blue MOT, broadband red MOT (BB MOT), narrowband red MOT (NB MOT), and then loading into optical tweezers.

3.1.1 Oven, Zeeman slower & main chamber

In order to produce a cold atomic vapour of strontium, we must work in a vacuum system. The vacuum system used in this thesis is the product of work of many previous members of our laboratory. Parts of the vacuum system date back over a decade to James Millen's PhD thesis [76], notably we still use the original Zeeman slower. More modern parts include the main vacuum chamber, which was upgraded in the thesis of Ryan Hanley [72]. The contents of the vacuum chamber have not been altered during the course of my thesis, however new connections have been made to the contents of the chamber using the feedthrough pins to various components. The overall layout of the vacuum system is shown in figure 3.2.

The experiment starts with a strontium oven heated to ~ 700 K [72], this is required to produce sufficient quantity of strontium vapour for our experiments. Capillary tubes inside the oven produce a well-collimated beam of strontium atoms, which are directed down the Zeeman slower tube and towards the main vacuum chamber.

Atoms from the oven travel at hundreds of metres per second, so must be rapidly slowed so that they can be trapped at the centre of the main chamber. To do this, the Zeeman slower is used [76]. A 461 nm beam of ~ 50 mW is sent into the Zeeman slower counter-propagating with the atomic beam. The light scatters on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, vastly reducing the velocity of the atoms. The 461 nm light is kept at the correct detuning relative to the changing Doppler shift of the atoms as they slow, because of the variable Zeeman shift along the \hat{x} axis counteracting

¹Thorlabs DMLP490L longpass dichroic mirror, 490 nm cut-on.

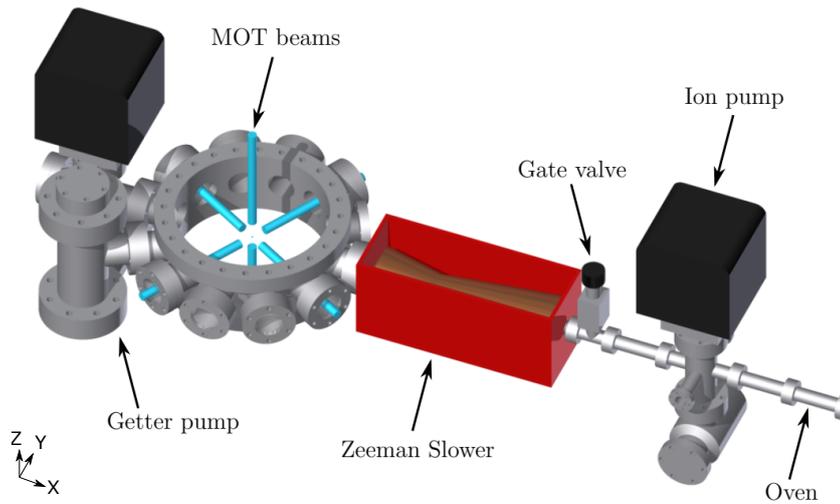


Figure 3.2: Layout of the vacuum system, including the oven, Zeeman slower and main chamber. This figure is modified from a figure from the thesis of Ryan Hanley [72], based upon the original figure from the thesis of Graham Lohead [77].

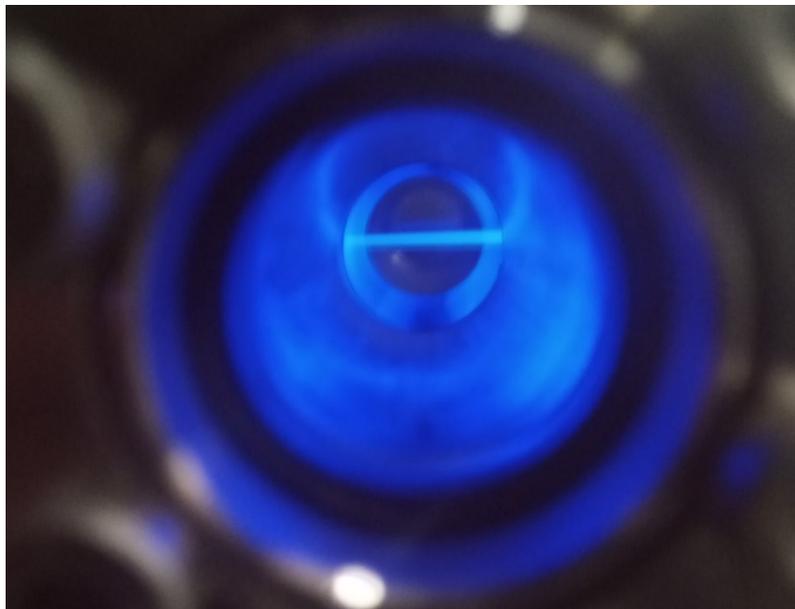


Figure 3.3: A photograph of the fluorescence from the atomic beam from the oven as it enters the beginning of the Zeeman slower. The beam is viewed through a DN16 viewport, where the image was taken with a regular mobile phone camera.

the shift. The variable Zeeman shift is produced by the different number of current carrying coils wrapped around the vacuum system along the length of the Zeeman slower. A photograph of the slowing beam of atoms is shown in figure 3.3.

The slowed atoms arrive in the main chamber, which is a DN200 pancake chamber with fourteen DN40 flanges, twelve horizontal flanges spaced every 30° around the edge, and a vertical flange at the centre of the top and bottom faces of the chamber. The magneto-optical traps (MOTs) are formed at the centre of the chamber using six orthogonal beams, four in the horizontal plane and two in the vertical. The MOTs can then trap the slowed atoms at the centre of the chamber.

The main chamber contains several pieces of equipment inside of the vacuum. These include *in-vacuo* lenses, MOT coils, electrodes and the MCP.² We will discuss these components in detail in section 3.3.

3.1.2 Blue MOT

After the Zeeman slower has drastically decreased the speed of the atoms, the next stage is to form a blue MOT. A blue MOT is a magneto-optical trap using 461 nm light on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1\ m_j = -1$ transition, which offers a high scattering rate to help rapidly cool the atoms. It comprises of three orthogonal counter-propagating circularly polarised laser beams, and a quadrupole magnetic field of $33\ \text{G cm}^{-1}$.³

Laser cooling in a MOT is well established, where details of preparation of MOTs of strontium can be found in references such as the thesis of Ryan Hanley [72]. The laser beams and magnetic field together provide both velocity and spatially dependant forces that trap and cool the atoms. The motion of the atoms in any direction brings the laser beam opposing the motion closer to resonance, providing a friction force against the motion of the atom. An atom displaced away from the centre of the trap is Zeeman shifted by the quadrupole magnetic field, bringing it

²For a detailed layout of these components see the thesis of Ryan Hanley [72].

³The quadrupole magnetic field is produce by the MOT coils inside the main chamber, which we will describe in section 3.3.2.

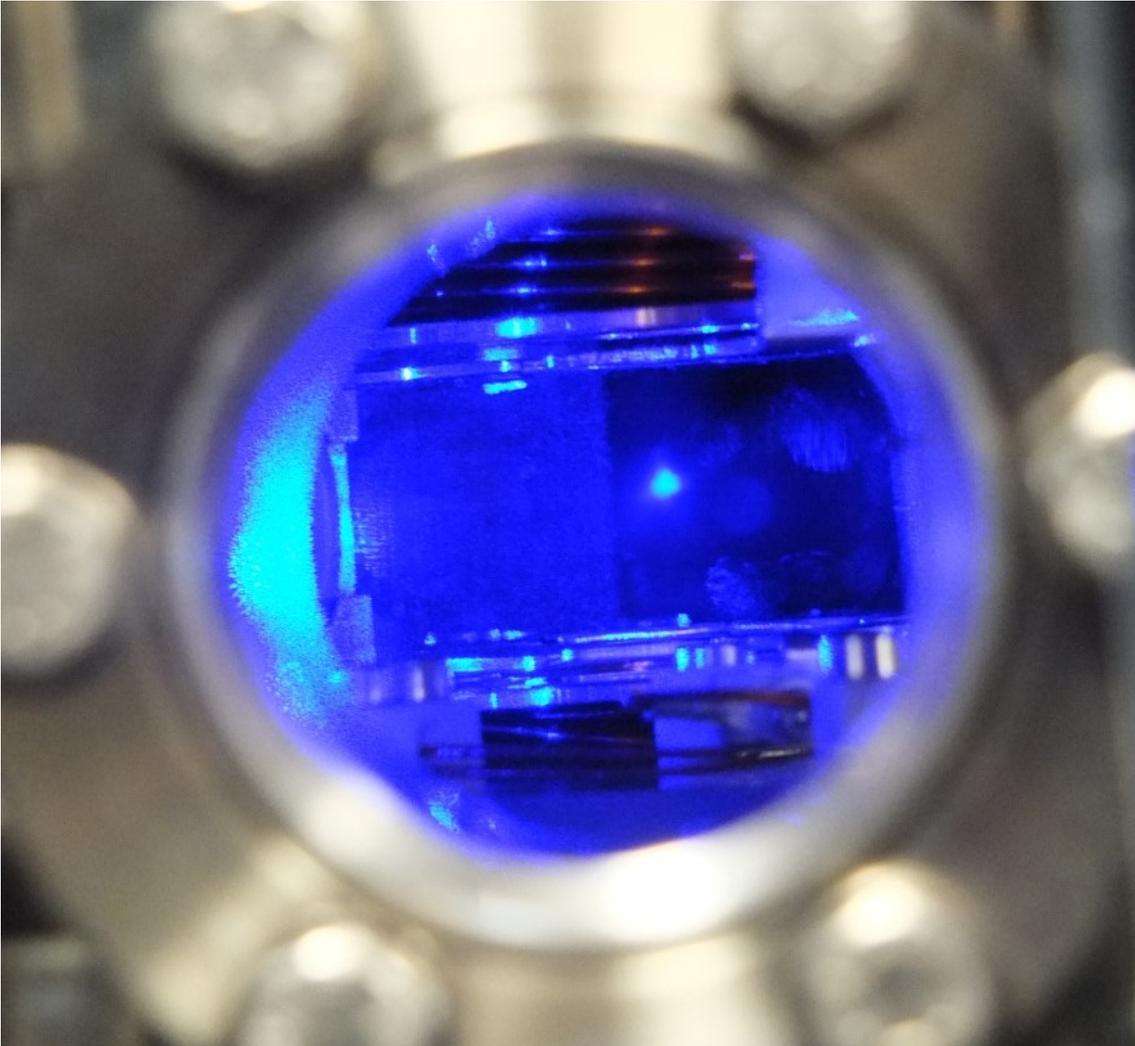


Figure 3.4: A photograph of the blue MOT taken using a regular camera through one of the viewports. The electrodes and MOT coils can be seen both above and below the atoms.

closer to resonance with the appropriate circularly polarised MOT beam, providing a restoring force to the centre of the trap.

It takes < 100 ms to fully load the blue MOT, cooling to a typical temperature of ~ 5 mK. This is above the Doppler temperature of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition of $T_D = \frac{\hbar\Gamma}{2k_B} = 0.7$ mK⁴, though ~ 5 mK is still sufficiently cold for the following stages of the experiment. The atomic cloud formed is visible to the naked eye through the viewports of the vacuum chamber, and appears as a bright blue ellipsoidal cloud at the centre of the chamber. A photo of a blue MOT taken with a regular camera⁵ is shown in figure 3.4.

3.1.3 Red MOT

To reach temperatures $\ll 1$ mK (much less than the trap depth of typical optical tweezer), we can no longer use the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition due to the large linewidth (and thus large Doppler temperature). After loading the blue MOT we instead switch to cooling on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition, which is over three orders of magnitude narrower. Here the Doppler temperature is a mere $0.18\ \mu\text{K}$, though as we will see for practical considerations we do not reach temperatures that low in the MOT.

It is reasonable to wonder why not start with a red MOT (a MOT based on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition) in the first place. Due to the much lower scattering rate on the transition, the red MOT can only capture atoms that have low kinetic energy, so that enough scattering events can occur before they escape the trapping region. This means for the experimental sequence that we transfer atoms from a blue MOT to a red MOT. In fact, the difference between using the blue and red transitions to cool is so great that we actually have two different stages of red MOT to enhance the probability of transferring atoms between them. The two stages of red MOT are the broadband MOT (BB MOT) and the narrowband MOT (NB MOT).

⁴Where $\frac{\Gamma}{2\pi}$ is the linewidth of the transition.

⁵Fujifilm X-S1.

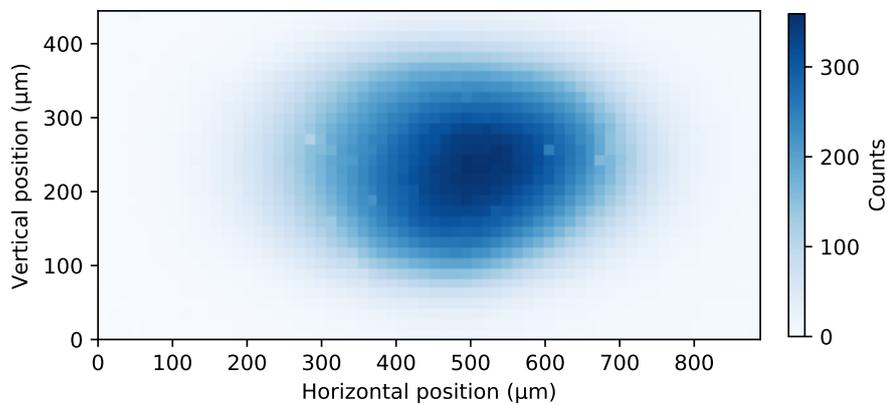


Figure 3.5: A typical averaged fluorescence image of 250 broadband red MOTs, taken at $\times 10.8$ magnification onto the SPAD array. A blue colourmap is used to represent the 461 nm fluorescence collected to image the MOT. Unlike the blue MOT, the red MOTs cannot be see directly by eye.

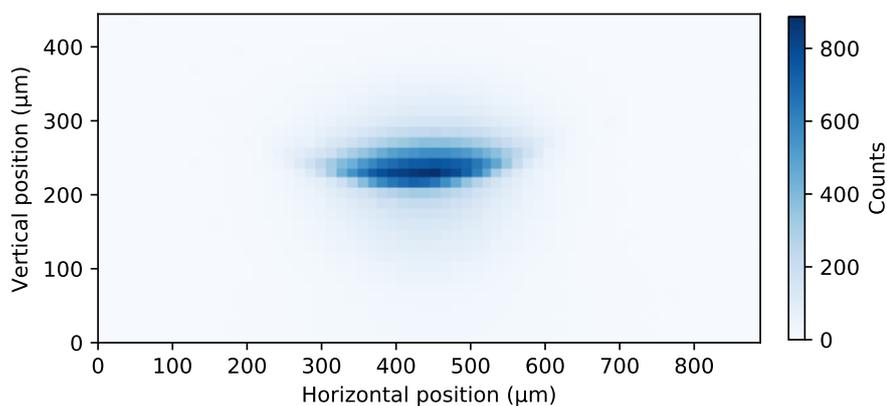


Figure 3.6: A typical averaged fluorescence image of 100 narrowband red MOTs, taken at $\times 10.8$ magnification onto the SPAD array. Each pixel corresponds to approximately $14 \mu\text{m}$ at the position of the atoms. The image is plotted against the full SPAD array to allow contrast with the higher magnification image of the narrowband MOT used in figure 5.4.

The BB MOT uses light that is artificially broadened to match more closely the temperature distribution of atoms from the blue MOT, which experience a variety of Doppler shifts. The light is broadened using the AOM used to switch the red MOT light⁶. The RF signal is modulated with frequency 50 kHz and deviation of 2 MHz to produce a comb of discrete frequencies. The BB MOT stage typically lasts for ~ 150 ms, reaching a temperature of ~ 20 μK . An averaged fluorescence image of many BB MOTs is shown in figure 3.5. The fluorescence images are taken with the $\times 10.8$ magnification imaging system to the SPAD array shown in figure 3.1, where a 52 μs pulse of the blue MOT beam light is used to cause the atoms to fluoresce.

We then transfer from the BB MOT to the NB MOT by switching to a single frequency driving the AOM. This removes the broadened laser light distribution, and allows us to cool further to typical temperature of ~ 1 μK , where the NB MOT stage takes ~ 100 ms. This is greater than the Doppler temperature of 0.18 μK , but is a reasonable final temperature to reach. The limiting temperature instead arises from the photon recoil energy of 689 nm photons, which have an energy equivalent to 0.23 μK . The NB MOT beams are typically set to a detuning of -280 kHz compared to the freespace atomic resonance, and are regularly corrected for the frequency drift of the cavity that the 689 nm laser is locked to.

An averaged fluorescence image of many NB MOTs is shown in figure 3.6. Once again this is taken with the $\times 10.8$ magnification imaging system to the SPAD array shown in figure 3.1, where a 52 μs pulse of the blue MOT beam light is used to cause the atoms to fluoresce. The cloud shape of the NB MOT no longer forms an ellipsoid like other MOTs, as gravity has become a significant force in the dynamics, and the atoms sit in a shell just above the resonance condition of the MOT light [151]. The cloud of atoms in the NB MOT is significantly smaller, colder and denser than the BB MOT, note the increased peak fluorescence for the same imaging parameters between figures 3.5 and 3.6.

To measure the temperature of the MOT, we perform a ballistic expansion of

⁶The red MOT light is switched into a fiber between the optical table with the laser and the optical table with the main chamber.

the cloud, where the change in the size of the cloud as a function of time tells us the temperature of the atoms on each axis. A ballistic expansion of the NB MOT taken along with the image in figure 3.6 yielded a horizontal temperature of $T_x = 0.83 \pm 0.03 \mu\text{K}$ and a vertical temperature of $T_z = 1.31 \pm 0.02 \mu\text{K}$.

3.1.4 Loading a tweezer

We now have a cold cloud of atoms in the NB MOT, with which we can load the atoms into the tweezer(s). The position of the NB MOT is set by applying an external magnetic field⁷ to shift the NB MOT to overlap with the tweezer(s).⁸ Good spatial overlap between the NB MOT and the the tweezer(s) is important to facilitate efficient loading.

The atoms are transferred to the tweezers by turning on the tweezer(s) during part of the NB MOT phase of the experiment, such that atoms can be cooled into the tweezers. To load many atoms into the tweezers the overlap time t_{load} typically lasts for half the NB MOT stage (50 ms). To control the number of atoms loaded into the tweezer, the time t_{load} can be varied. For loading on the single atom level as we will demonstrate in section 4.5, we use overlap times $t_{\text{load}} < 5$ ms.

The atoms are initially loaded into tweezers of trap depth $< 30 \mu\text{K}$, where at the end of the NB MOT stage the MOT is turned off, and the atoms not in the tweezer are allowed to fall away. We then adiabatically increase the trap depth of the tweezer over several hundred milliseconds to the desired trap depth for experiments. For example in chapter 4, the final trap depth of a 532 nm tweezer is 7.5 mK unless otherwise stated.

The reason for ramping the trap depth of the tweezers is discussed in detail in the thesis of Ryan Hanley [72]. The focus of an optical tweezer is not exactly a Gaussian beam, but is instead more like a sinc² function. The apodisation of the

⁷The external magnetic field is produced by the compensation coils, which we will describe in section 3.3.2.

⁸The position the tweezers are formed at is defined by the *in-vacuo* lens and the details of the trapping light sent to the lens. Whereas the position the MOT forms is defined by the magnetic field and the detuning of the MOT beams.

trapping light on the *in-vuaco* lens causes subsidiary intensity maxima, which can potentially trap atoms. Given the $\sim 1 \mu\text{K}$ temperature of the atoms from the NB MOT, we need to keep the trap depth of the subsidiary maxima on the $\lesssim 1 \mu\text{K}$ level when the NB MOT is present, in order to prevent atoms from loading into the subsidiary maxima. This is readily achieved for trap depths of $< 30 \mu\text{K}$. Once the NB MOT has fallen away, the trap depth can then be ramped to the required value. This leaves us with cold atoms prepared in optical tweezers, ready to do experiments with.

3.2 Laser systems

In order to prepare cold atoms, create tweezer arrays, and then manipulate the atoms in experiments, a variety of different laser systems are required. This section outlines the different laser systems in use, highlighting changes from the most recent theses from our laboratory [72, 79, 80]. Figure 3.7 shows the wavelengths used a spectrum. A summary list of relevant wavelengths can also be found in the nomenclature (see section starting on page xi), and key atomic transitions are shown in figure 2.1.

We begin by focusing on the key lasers for preparing a cold atomic sample in a tweezer. The 461 nm and 689 nm lasers to produce the MOTs of strontium and for imaging the atoms are described. We then give an overview of the 532 nm and 813 nm lasers used to make the tweezers in chapters 4 and 5, respectively. Two other ultrastable lasers are described: the UV laser for excitation to triplet Rydberg states in chapter 6, and the new 698 nm clock laser for driving the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ clock transition (used in appendix E). Other systems of the repump and auto-ionisation lasers, along with the wavemeter and frequency comb, are then briefly introduced.

3.2.1 461 nm lasers

The 461 nm lasers are used for addressing the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition (see figure 2.1), which is required for the Zeeman slower, blue MOT, and for fluorescence imaging of the atoms. Over the course of the work presented in this thesis, two different configurations of 461 nm laser systems will be used, using three dif-

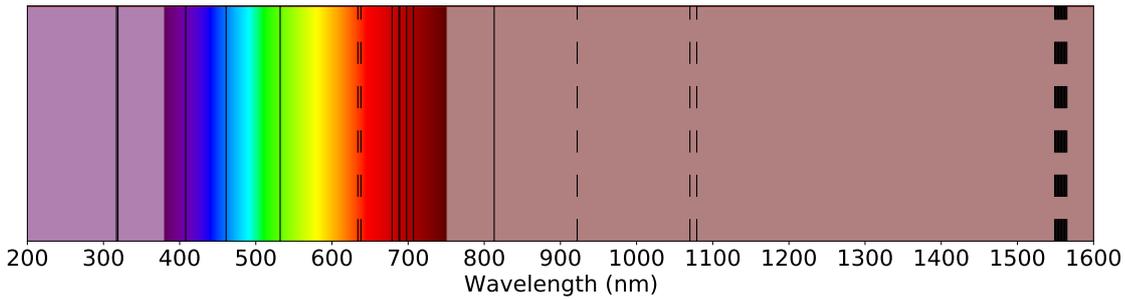


Figure 3.7: A spectrum showing the range of lasers used in our experiment. Solid lines show lasers that are actually used in the experiment, dashed lines show lasers that are later changed by second harmonic generation (SHG) or sum frequency generation (SFG) to produce a different wavelength.

ferent lasers. The configurations used for chapters 4 and 5 will be referred to as configurations A and B respectively.

Configuration A

This layout of the 461 nm system corresponds to that used in chapter 4. The configuration is described in detail in the thesis of Ryan Hanley [72], though with the addition of a widely tunable probe beam.

The main 461 nm laser is a frequency-doubled diode laser⁹. A 922 nm seed laser is amplified by a tapered amplifier (TA), before a nonlinear crystal in a ring cavity doubles the light to 461 nm. The laser output is shaped and fiber coupled to produce a higher quality beam profile, before being split off into different arms of the system for its different uses, where each arm is shifted by an acousto-optic modulator (AOM) before being sent to its application. The main laser is used for the Zeeman slowing beam, modulation transfer spectroscopy (MTS) locking to a reference cell [152], a fixed frequency probe beam that was not used in this work¹⁰, and to produce the blue MOT beams.

A pick off of the light from the main 461 nm laser is used stabilise another 461

⁹Toptica Photonics DL-SHG frequency-doubled system.

¹⁰We instead required a widely tunable probe beam for the work presented in chapter 4

Laser	Application	AOM /MHz	Detuning /MHz
TA SHG	Output	N/A	-240
TA SHG	Cell lock	+120×2	0
TA SHG	Beat lock	N/A	-240
TA SHG	Zeeman slower	-130×2	-500
TA SHG	MOT Beams	+200×1	-40
MOGLabs	Output	N/A	$\nu_{\text{Beat}} - 240$
MOGLabs	Beat lock	N/A	$\nu_{\text{Beat}} - 240$
MOGLabs	Probe beam	-120×1	$\nu_{\text{Beat}} - 360$

Table 3.1: 461 nm system configuration A. The beat lock operates over a very wide range, and is typically used in the $\nu_{\text{Beat}} = 300\text{-}800$ MHz range. This covers both the bare $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, and the two split AC Stark shifted $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ $|m_j| = 0$ and 1 transitions in a deep (7.5 mK) tweezer at $\lambda_{\text{trap}} = 532$ nm.

nm laser via a beat lock. The second 461 nm laser¹¹ is used as a probe beam of variable frequency and power, where the frequency is controlled via a beat lock, and the power by an AOM. The use of a second laser both increases the available power, and allows us to tune to different frequencies with ease. The variable probe is useful for exploring imaging in deep traps where there are large AC Stark shifts on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ imaging transition, this will be further detailed in section 4.1.1. The frequencies of each of the beams is shown in table 3.1.

Beat lock

The two 461 nm lasers are locked together to keep their relative frequencies stabilised, and so that the secondary 461 nm laser does not have to be locked to its own strontium reference cell. Initially the lasers were stabilised via a wavemeter lock, though this proved unreliable. The relative frequency drifted by ~ 20 MHz over timescales of several minutes, and occasionally the wavemeter would misread

¹¹MOGLabs ECDL.

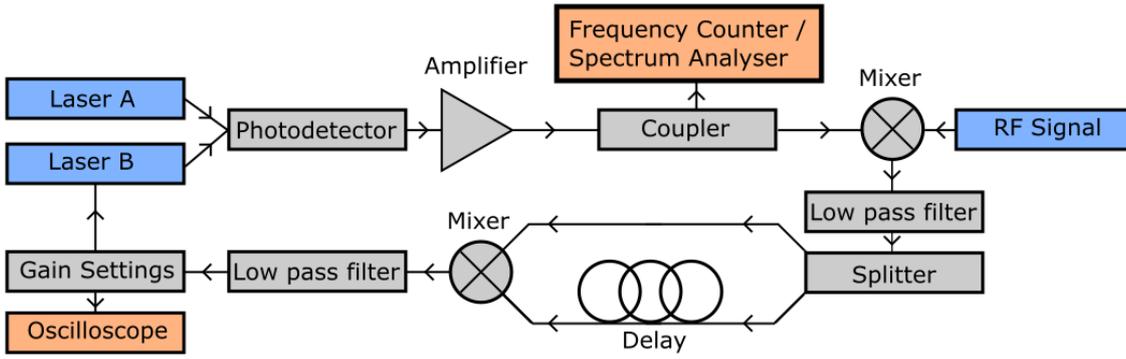


Figure 3.8: Design of the beat lock between the two 461 nm lasers. Input frequencies are coloured blue, analysis tools orange, and other electronics grey. Laser A is the laser that is MTS locked to a strontium reference cell, laser B is being locked by the beat lock. Components used in the beat lock are detailed in appendix A.1

the value, causing the laser to unlock. The stability was insufficient as it was comparable to the linewidth of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, therefore we instead constructed a beat lock between the two lasers.

The two lasers are overlapped on a non-polarising beam splitter (NPBS), before being reflected from a polarising beam splitter (PBS) to select the same polarisation. The light is then fiber coupled to a photodiode¹², where the signal from the photodiode is converted into an error signal using the locking electronics shown in figure 3.8 [153]. The error signal produced is shown in figure 3.9, where we lock to a zero crossing of the error signal. This locking offers sub MHz relative stability between the two lasers, as seen in figure 3.10, significantly narrower than the 30.2 MHz linewidth of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition. This stability is critical for the measurements in chapter 4, where variations in the probe frequency would have induced fluctuations in the expected fluorescence per atom and reduced the ability to resolve different atom numbers. The beat lock also allows us to set the relative frequency of the lasers by matching it to a computer controlled RF signal, allowing easy remote control of the probe frequency over a range of > 500 MHz in configuration A for use in chapter 4.

¹²Thorlabs DET025AL/M

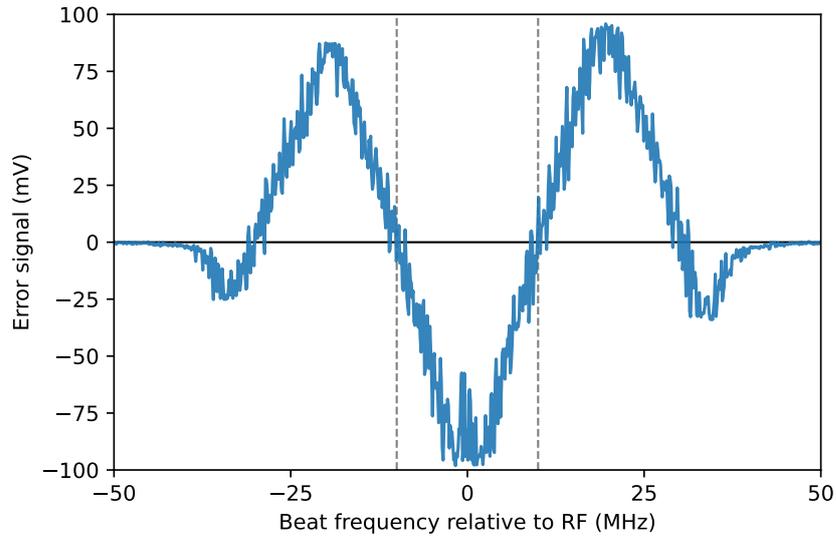


Figure 3.9: Error signal produced by the 461 nm beat lock. We lock to one of the steep zero crossings marked with the dashed lines. This has a ± 10 MHz offset in the locked beat frequency compared to the RF signal, so the correct zero crossing must be chosen.

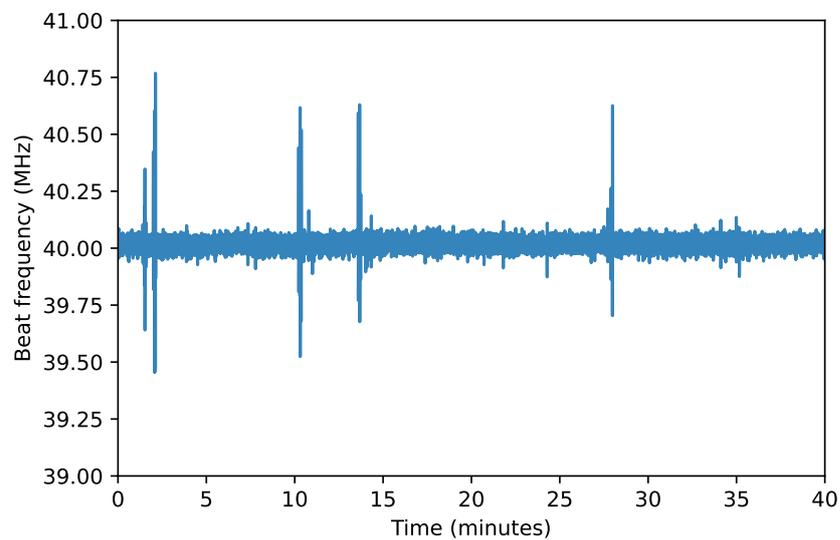


Figure 3.10: To assess the stability of the beat lock, a frequency counter is used to record the beat frequency every 100 ms. Here the RF signal in figure 3.8 was set to 50 MHz, and the lower frequency zero crossing was locked to such that the beat frequency between the two lasers was 40 MHz. The beat frequency is stable on the order of 0.1 MHz, with occasional spikes of < 1 MHz lasting < 1 s.

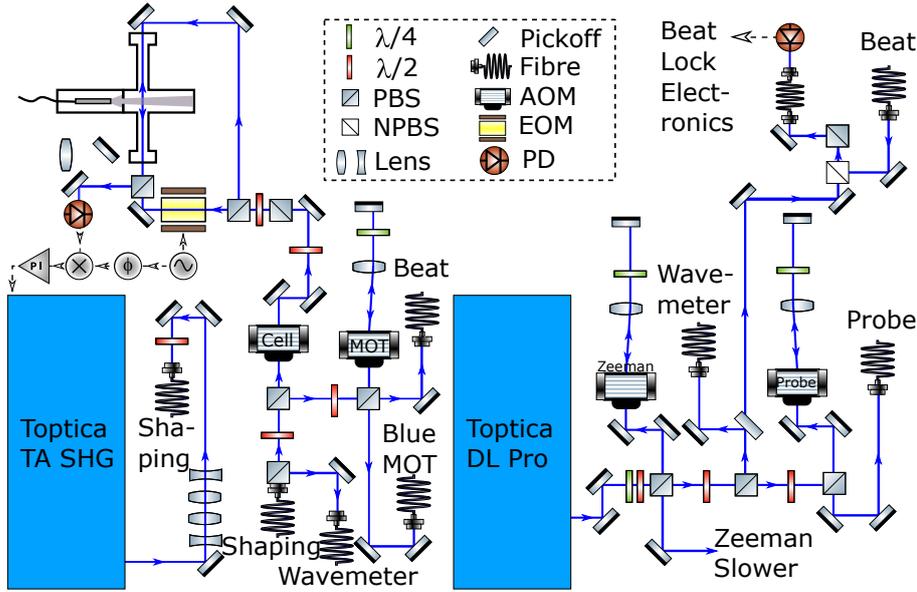


Figure 3.11: The optics layout for the 461 nm lasers in configuration B.

Configuration B

During the course of this work, we added a spatial light modulator (SLM) to produce tweezer arrays at $\lambda_{\text{trap}} = 813$ nm. In order to fit this near the vacuum chamber, a rearrangement of the lasers on the main optical table was required. We took this opportunity to free up space taken by the 461 nm laser system, at the same time improving it and increasing the available power. The reader may wish to skip the details of configuration B until they reach chapter 5, as it is not implemented until that point.

When rearranging the 461 nm laser system, we once again decided to beat lock two lasers together. However in this iteration, we had different priorities. We replace the MOGLabs ECDL with a Toptica DL Pro for high power output (> 100 mW) and better stability. This configuration is used for the results in chapter 5, where the probe beam does not require such a wide tunability due to the traps at 813 nm having a lower differential AC Stark shift on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition and working at lower trap depths. This meant the the second laser can take on more of the tasks, and there is greater power available to each of the 461 nm beams.

The Toptica TA SHG laser was moved to a secondary optical table, due to the large space it takes up. The Toptica TA SHG laser is used for the blue MOT beams,

Laser	Application	AOM /MHz	Detuning /MHz
TA SHG	Output	N/A	+200
TA SHG	Cell lock	-200×1	0
TA SHG	Beat lock	N/A	+200
TA SHG	MOT Beams	-120×2	-40
DL Pro	Output	N/A	-265
DL Pro	Beat lock	N/A	-265
DL Pro	Zeeman slower	-105×2	-475
DL Pro	Probe beam	+122.5×2	-20

Table 3.2: 461 nm system configuration B, where the beat lock operates at $\nu_{\text{Beat}} = 465$ MHz. The probe beam is tunable around the AC Stark shifted $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition in a shallow 813 nm tweezer.

and is MTS locked to the strontium reference cell. The Toptica DL Pro laser (which is much smaller than the Toptica TA SHG laser) was placed on the main optical table, where is it used for the Zeeman slower and the probe beam. The detailed layouts of the two lasers are given in figure 3.11, with the frequencies in each beam summarised in table 3.2.

Using the Toptica DL Pro laser for both the Zeeman slower and probe beam is ideal as we require the ability to tune the probe frequency. We have significant tuning range on both double pass AOMs¹³ used in the probe and Zeeman paths, as well as the beat lock frequency, allowing us to widely change the probe beam frequency without affecting the Zeeman slowing light. In addition to this, the Zeeman slower is effective over a range of frequencies, so could be adjusted for the benefit of the probe beam if necessary, without significantly impacting loading of atoms.¹⁴

¹³AA Opto Electronic MT130-B50-A1.5-461.

¹⁴Though this should not be necessary.

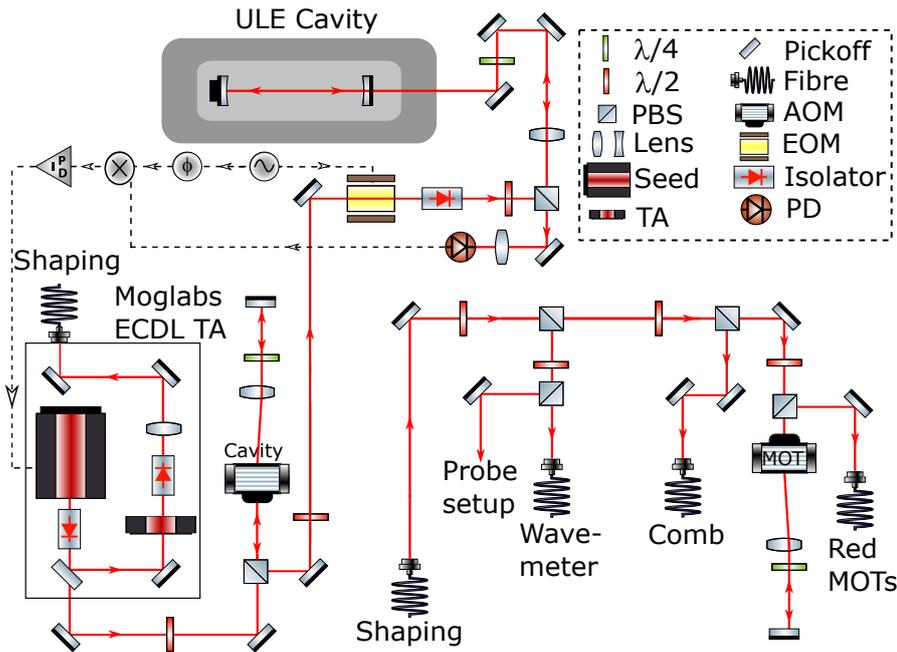


Figure 3.12: The optics layout for the 689 nm laser.

3.2.2 689 nm lasers

The 689 nm laser addresses the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition (see figure 2.1), and has uses including formation of both BB and NB MOTs, Sisyphus cooling in tweezers, and as the first stage in some Rydberg excitation schemes. For the results in chapter 4, two home built 689 nm lasers were used, as described in previous theses from our group [72, 80]. For results in chapter 5 this system was replaced with a commercial laser system with a TA¹⁵, which shall be described here and is shown in figure 3.12.

New 689 nm laser

A 30 mW diode laser is used as the seed laser for a TA, with output powers of up to 600 mW¹⁶. The output is then fiber coupled and split into different arms of the

¹⁵MOGLabs ECDL TA.

¹⁶Note that the full power of the 689 nm system is not required, and we do not want to unnecessarily power broaden the transitions or shorten the lifetime of the TA chip, so we operate the 689 nm system at lower power of ~ 200 mW with the TA current at 0.8 A.

experiment for different applications. The applications include red MOT¹⁷ beams and probe beam¹⁸, where both use the +1 order of their respective double pass ~ 80 MHz¹⁹. Other light is sent to the wavemeter and frequency comb (see section 3.2.6) such that the wavelength of the 689 nm light can be accurately measured.

Unlike the 461 nm laser which operates on a broad (30.2 MHz) transition, the 689 nm laser is used on a narrow (7.4 kHz) transition, and must be frequency stabilised more carefully to reduce the laser linewidth. To achieve this, a small (5%) pick-off of the seed laser light is sent to an ultra-low expansion (ULE) cavity for Pound-Drever-Hall (PDH) locking as shown in figure 3.12.

PDH locking

Pound-Drever-Hall (PDH) locking is a commonly used locking method in our experiment, and is used in three of our laser systems at 689, 638 and 698 nm (see sections 3.2.4 and 3.2.5 for locking of the 638 and 698 nm lasers respectively). This section aims to introduce PDH locking using the case study of the 689 nm laser. For a more thorough explanation of PDH locking I would recommend reference [154], where specific examples of applying PDH locking can be found in previous PhD theses from our project [73, 80]. PDH locking stabilises the wavelength of the laser against the length of a stable cavity. The ultra-low expansion (ULE) cavity used for the 689 nm laser has a finesse of $\mathcal{F} = 236,000$ [72], and a free spectral range (FSR) of 1.5 GHz.

To achieve PDH locking, the light picked off from the seed laser (see figure 3.12) to go to the cavity is phase modulated by adding an electro-optic modulator (EOM) before the cavity, where here we modulate at 10.110 MHz. The laser light is then mode matched into the cavity to favour transmission in the transverse electromagnetic 00 (TEM₀₀) mode. When a Gaussian beam is transmitted through the cavity,

¹⁷The same light path is used for both the BB and NB MOTs, where the only difference is the frequency and amplitude of the RF signal driving the AOM that switches the light.

¹⁸The 689 nm probe beam setup has not yet been constructed.

¹⁹Exact frequency depends upon whether a BB or NB MOT is being formed, or the required frequency of the probing light.

then the wavelength of the laser has matched a TEM_{00} mode. This occurs when the laser wavelength matches a condition set by the length of the cavity. The reflected signal from the input face of the cavity is focused onto a photodiode, where the optical signal from the cavity is converted to an electronic one. The signal is still modulated at 10.110 MHz, so we demodulate it by mixing it with 10.110 MHz RF of the correct phase. This signal forms the typical PDH error signal shape [154], and has a sharp zero crossing that can be locked to by processing the signal with electronics²⁰ and feeding it back into the laser. As the laser frequency moves and the cavity reflection signal changes, the locking electronics act to correct the laser frequency. This suppression of the fluctuations in laser frequency reduce the linewidth of the laser.

The settings of the locking electronics are crucial to achieving a narrow linewidth and a reliable laser lock. The electronics need to counteract low frequency noise on the laser, where frequencies below the bandwidth of the locking electronics are well controlled by the lock. The lock should be configured such that it has a bandwidth large enough to cover the frequency range of the dominant noise spectrum of the laser [154]. To optimise the locking electronics, the in-loop error signal is observed (see figure 3.13). The in-loop error signal is viewed by picking off some of the signal from the reflection photodiode²¹ and observing it on a spectrum analyser. When the laser is locked the in-loop error signal is composed of a sharp spike at the EOM modulation frequency, and two wide noise sidebands on either side. A good lock will have the gain balanced such that the sidebands are far from the carrier frequency (large servo bandwidth), but not so far that the loop starts to oscillate and the sidebands significantly increase in amplitude [73, 154]. The distance the sidebands are from the carrier for these optimised locking parameters indicates the servo bandwidth, which tells us below what frequency the lock is removing noise from the laser effectively. Figure 3.13 shows a servo bandwidth of 0.8 MHz.

With PDH locking we now have a 689 nm laser with a narrow linewidth locked

²⁰MOGLabs Fast Servo Controller

²¹Preferably with a directional coupler, though alternatively a splitter or T piece could be used if required.

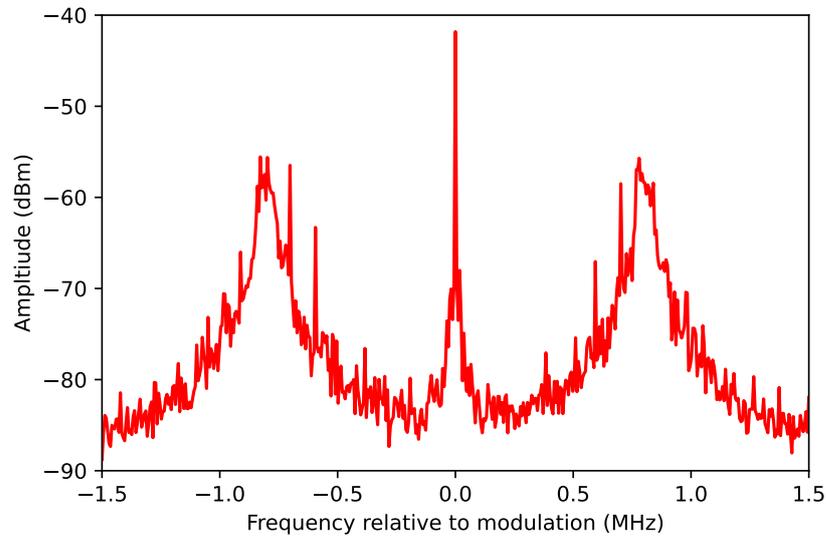


Figure 3.13: The locked in-loop error signal for the 689 nm laser when the signal is picked off after the photodiode with a splitter. The servo bandwidth achieved here is 0.8 MHz.

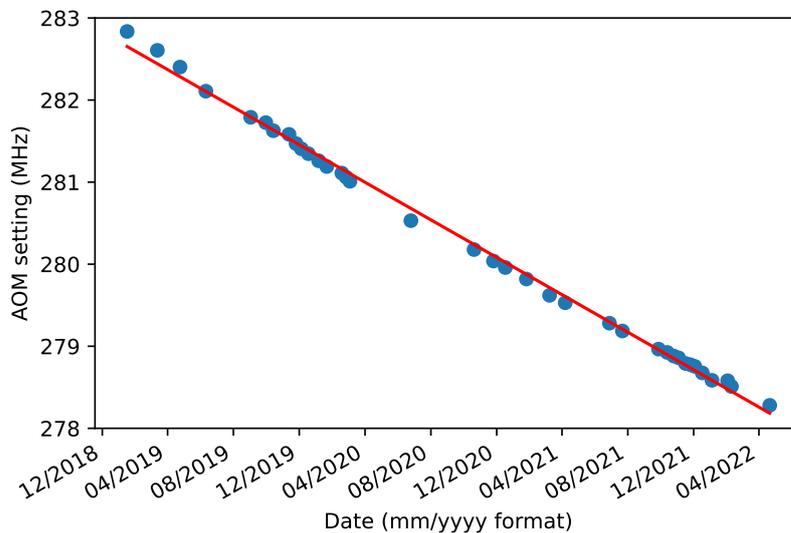


Figure 3.14: The long term change in the frequency of the double pass AOM between the 689 nm seed laser and the ULE cavity used to compensate for the cavity drift. The cavity drifts by 7.50 ± 0.07 kHz per day, which must be compensated for periodically to maintain a suitable detuning for the red MOT.

to a reference. This alone does not mean that the laser is at the correct frequency. Unfortunately the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition falls half way between two of the TEM_{00} modes, which are spaced by the 1.5 GHz FSR. To resolve this, the seed light to the cavity is first double passed through a 280 MHz AOM in the -1 order. The combined shift of the cavity path and MOT beam AOMs bring the laser frequency to the target, in the case of the MOT beams slightly red-detuned of resonance.²²

Another consideration that must be made is the stability of the cavity itself. PDH locking stabilises the laser wavelength to the length of the cavity. The ULE cavity is specifically designed to provide a stable length reference, however the length of the cavity still varies in time. For the cavity used with the 689 nm system, the frequency drifts by 7.50 ± 0.07 kHz per day, and the AOM frequency in the cavity path must be adjusted periodically to compensate. Figure 3.14 shows some of the corrections made to maintain the red MOT detuning over several years, revealing the long term linear drift of the cavity.

3.2.3 Dipole trapping lasers (532/813 nm)

The laser systems introduced so far have corresponded to atomic transitions of strontium. Here we introduce two lasers that are far from any transition wavelengths, used as far off resonance optical tweezers for trapping atoms. We use two different wavelengths of tweezers throughout this work, at 532 nm and 813 nm (as discussed in section 2.2.3).

The results in chapter 4 use a diode-pumped solid state laser called the Sprout²³, producing 15 W²⁴ of light at 532 nm. Some of the 532 nm light is taken from a pick off to make the single deep tweezer. The wavelength of the Sprout laser is not tunable, but the exact frequency does not matter as it does not need to be at an atomic transition or a magic wavelength.

²²Typically -280 kHz detuned for a NB MOT.

²³Lighthouse Photonics, Sprout-G 15 W.

²⁴Alternatively could be described as a 20 microhorsepower laser, though that unit system may not catch on for lasers.

The majority of the 532 nm power is used to pump the titanium-sapphire (Ti:Sa) laser²⁵. The Ti:Sa has a very widely tunable frequency (100s of nm). We use the Ti:Sa at the magic wavelength of 813 nm [33, 34], where the ground and excited states of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ transition have the same AC Stark shift. At 813 nm the Ti:Sa can produce a power of ~ 2.7 W.

Both 532 and 813 nm laser systems are switched by 80 MHz AOMs into high-power optical fibres to the main experiment, where after the fiber the beams are expanded to the appropriate size to focus with the *in-vacuo* lens to make tweezer(s). The AOM allows the control of the power coupled into the fiber, allowing the power and therefore trap depth of the tweezer array to be dynamically controlled. The control of the tweezer power is important for adiabatically ramping the trap depth, as discussed in section 3.1.4.

The specifics of the alignment of the 532 nm system are detailed in the thesis of Ryan Hanley [72]. The specifics of the alignment of the 813 nm system involve detailed discussion of the spatial light modulator (SLM), so are left for chapter 5.

3.2.4 Rydberg excitation laser (317/319 nm)

One of the most complex laser system in our experiment is the ultraviolet (UV) Rydberg excitation laser. The UV laser operates at either 317 or 319 nm, exciting to a chosen Rydberg state ($5snl\ ^3L_J$, where n is large) from either the $5s5p\ ^3P_0$ or $5s5p\ ^3P_1$ states, respectively. In this work we exclusively use the UV laser at 319 nm for coupling to the triplet Rydberg states via the $5s5p\ ^3P_1$ state. However, using the laser at 317 nm will be important for Rydberg dressing the clock transition in future work toward spin-squeezing an atomic clock. The UV laser requires a narrow linewidth to drive the narrow Rydberg transitions, a wide tunability to allow us to select from a vast range of triplet Rydberg states, and as much power available as possible to allow high Rabi frequencies.

To achieve this we use the setup described in previous work from our group

²⁵M Squared SolsTiS.

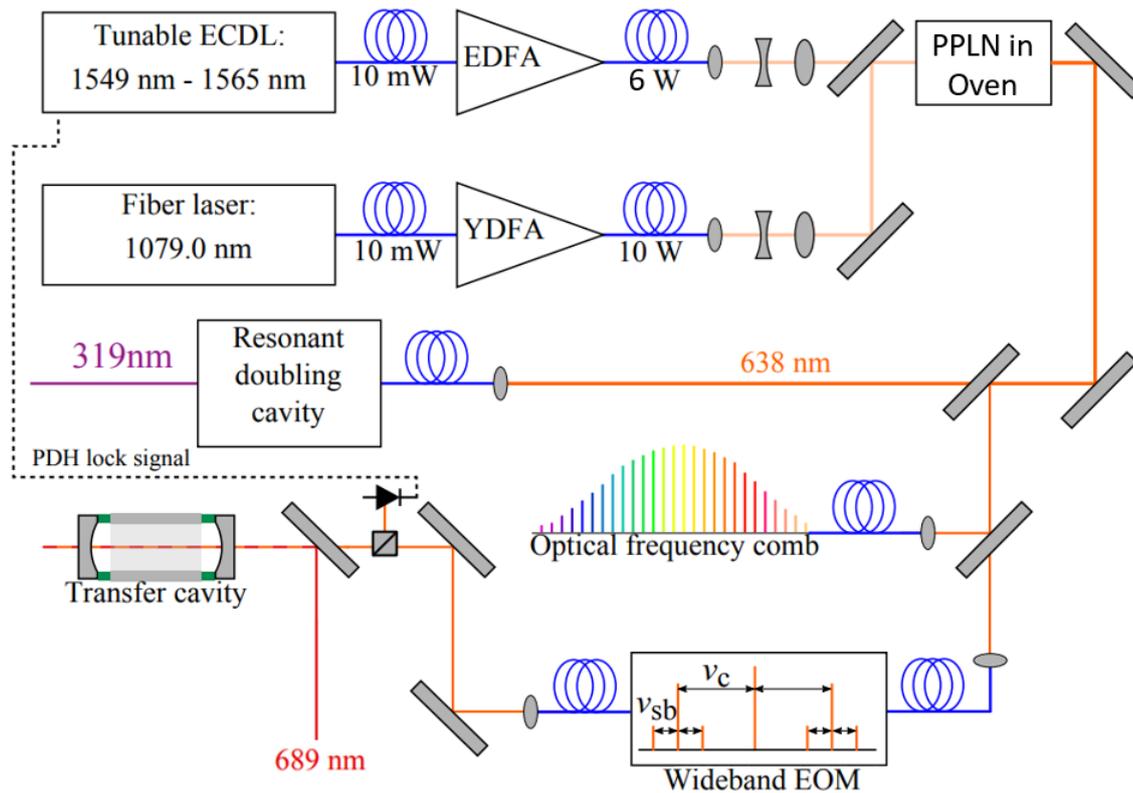


Figure 3.15: Layout of how the Rydberg excitation laser produces large quantities of UV, here in case of producing light at 319 nm. It is locked to the same cavity as the 689 nm system, though instead using electronic sideband locking. This figure is modified from a figure in the thesis of Alistair Bounds [79], other versions of this figure can be found in references [80, 155].

[79, 80, 155], aside for changes of some components²⁶. The experimental setup is shown in figure 3.15, which we will now recap. We start with two infrared seed lasers. The first is a tunable ECDL²⁷ offering a wide range of wavelengths from 1549-1565 nm, this laser is the source of the UV laser's tunability between the Rydberg states. The second IR seed laser are interchangeable 1070 and 1079 nm lasers²⁸. The former is used when intending to couple to the Rydberg state from the $5s5p\ ^3P_0$ state, and the latter when intending to couple from the $5s5p\ ^3P_1$ state. For all measurements in this thesis we use the 1079 nm laser to drive $5s5p\ ^3P_1 \rightarrow 5snl\ ^3L_J$ transitions.

The two seed lasers are both then amplified using fiber amplifiers²⁹ to increase the powers to 6 W and 10 W for the 1549-1565 nm and 1070/1079 nm beams respectively. These beams are overlapped and focused in a temperature controlled PPLN³⁰ crystal, frequency doubling the light to produce 4 W of 634 or 638 nm light. This light is then fibered to the main optical table using a photonic crystal fiber (PCF)³¹. The 634 or 638 nm light is then frequency doubled in a bowtie cavity³² to produce many 100s of mW³³ of UV at 317 or 319 nm.

The frequency of the laser is stabilised using the same ULE cavity that is used in the 689 nm system, where the 634/638 nm light is electronic sideband locked [156]. The light to the cavity is picked off from the 634/638 nm laser, where the ULE cavity also has high reflectivity. The PDH signal feeds back to the 1549-1565 nm laser in order to stabilise the 634/638 nm (and therefore also the UV) frequency.

A key difference between locking of the Rydberg excitation laser compared to the

²⁶The changes include the summing crystal, the 1549-1565 nm fiber amplifier, and the RF source for driving the EOM.

²⁷Toptica ECDL.

²⁸NP Photonics 'The Rock'.

²⁹Nufern and Keopsys.

³⁰Periodically-poled lithium niobate, instead of the previous PPLST (periodically-poled stoichiometric lithium tantalate).

³¹NKT Photonics, LMA-PM-15, 10.0 m FC-APC / FC-APC, 3mm PVC.

³²Toptica frequency doubling unit.

³³The maximum power of the laser system has not been characterised since the upgrades.

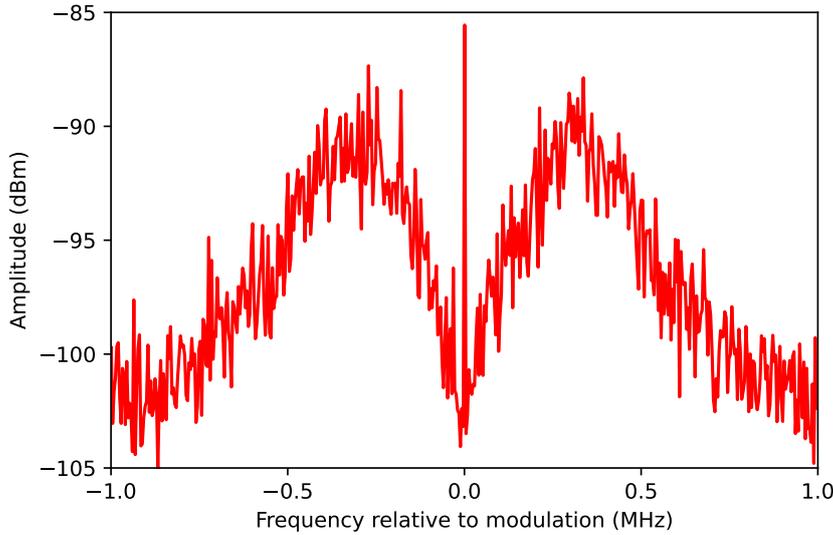


Figure 3.16: The locked in-loop error signal for the 638 nm laser when the signal is picked off after the photodiode with a coupler. The servo bandwidth achieved here is 0.3 MHz.

689 nm laser, is that we want to be able to widely tune it to a variety of transitions, rather than just a single one. We therefore need to be able to vary the detuning of the laser compared to the cavity modes by the FSR of the cavity (1.5 GHz). To achieve this, we instead use a modified PDH locking scheme known as electronic sideband locking [80, 156]. The EOM is driven with a RF source³⁴ at a carrier frequency of 100s of MHz, with sidebands at 8.333 MHz. Varying the carrier frequency varies the detuning of the laser from the cavity mode, allowing the laser to be locked at the frequency of any of the desired Rydberg states, as well as allowing us to scan the UV across transitions. Other than the presence of a controllable scanning carrier frequency to shift the laser, and the additional complexity of working with multiple lasers in order to produce UV, the locking otherwise behaves similarly to that of the 689 nm laser. The in-loop error signal for locking the 638 nm laser is shown in figure 3.16.

The produced UV light is then sent to the chamber in an anodized aluminium laser safety enclosure. Along the path, the light to the chamber passes through

³⁴Keysight N5171B.

a half-waveplate to set the polarisation, and is then switched by both an AOM³⁵ and a shutter. The UV light passes through uncoated fused silica viewports into the chamber, unlike the AR coated NBK-7 glass viewports used for all of the other beampaths, which would strongly absorb UV light [72]. There is the option to add beam shaping optics, such as to produce a beam with uniform intensity [72, 79, 80, 149], though we do not use any at present.

3.2.5 Clock laser (698 nm)

As the local oscillator for an atomic clock, we require a narrow linewidth laser at a wavelength of 698 nm corresponding to the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ clock transition. The 698 nm laser will be referred to as the clock laser in order to prevent the wavelength from getting confused with the 638, 688 and 689 nm lasers that also form parts of our apparatus.

Due to the narrowness of the clock transition ($\ll 1$ Hz) we require a much narrower linewidth laser than the kHz linewidth lasers used in the 638 nm (see section 3.2.4) and 689 nm (see section 3.2.2) systems. To achieve this we use a new laser to our laboratory that was produced by Menlo Systems. The linewidth of this laser will be discussed in greater detail throughout this chapter, where the test report of the laser indicates a linewidth < 1 Hz on timescale of 1 second.

The clock laser system is shown in figure 3.17. An external-cavity diode laser (ECDL)³⁶ produces ~ 27 mW of 698 nm light. 3 mW of the light is picked off to the cavity section of figure 3.17 to be sent for laser locking. The picked off light is shifted into the -1 order of a 286 MHz³⁷ AOM in order ensure that the light from the laser output can be tuned to the atomic transition frequency when locked to the ULE cavity. The -1 order is fiber coupled to the cavity breadboard, which is mounted on an active anti-vibration stage³⁸ and surrounded by an acoustic isolation

³⁵AA Opto Electronic MQ110-A3-UV

³⁶Moglabs Cateye ECDL

³⁷Driving RF is produced by a VCO (voltage controlled oscillator) internal to the Menlo unit. All RF signals within the clock laser setup use a 10 MHz reference from GPS.

³⁸Accurion halcyonics_i4.

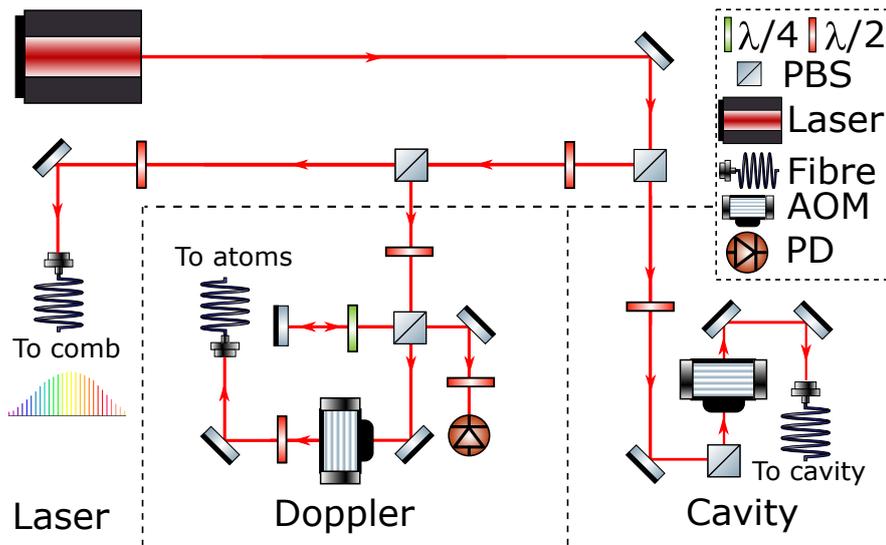


Figure 3.17: The optics layout for the 698 nm clock laser. The laser is locked to an ultrastable cavity through the cavity section, and has two outputs to the frequency comb and to the atoms in the laser and Doppler sections, respectively. The light to the atoms is Doppler noise cancelled.

box. An EOM then adds 20 MHz sidebands, where 63 μW of light arrives at the cavity input to achieve high quality locking. The ULE cavity here has a finesse of $\mathcal{F} = 230,000$, as shown in figure 3.18. The in-loop error signal of the PDH locking is shown in figure 3.19.

The rest of the light produced by the laser is then split between laser and Doppler sections of figure 3.17. The laser section sends light to beat with the frequency comb (see section E.1) via a fiber³⁹ with APC⁴⁰ connectors to measure the frequency of the clock laser. The beat of the clock laser with the frequency comb is discussed in detail in appendix E.1. The Doppler section is used to sent light to the atoms, where the Doppler AOM first shifts the light into the -1 order of a 40 MHz AOM,

³⁹For the results in this thesis the clock laser and frequency comb were connected by a 70 m long fiber link. This frequency comb has since been moved to the same laboratory as the strontium experiment, and is now connected by a 5 m long fiber link.

⁴⁰Angled physical contact (APC) is a connector type where the end of fiber is polished at an angle (typically 8°) to prevent back-reflections into the fiber from the output.

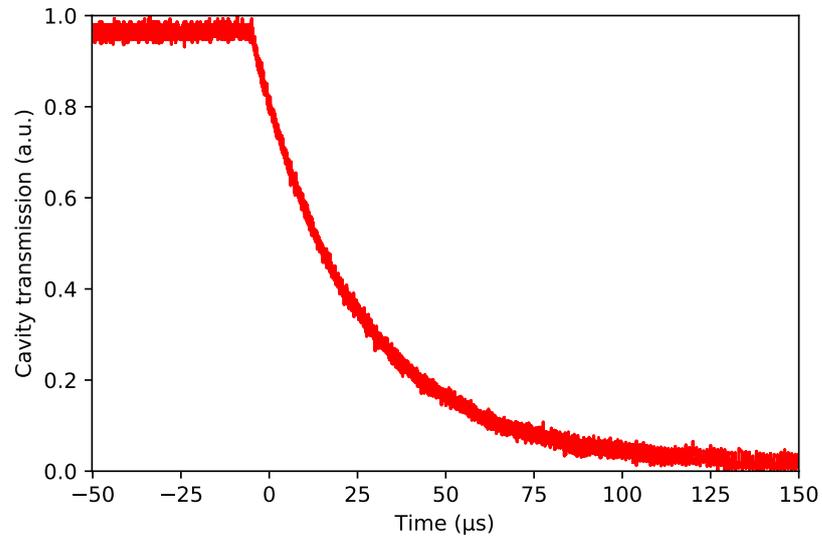


Figure 3.18: The ring down measurement of the ULE cavity used for the clock laser. Fitting the transmission curve shows a ringdown time of $29.61 \mu\text{s}$, corresponding to a finesse of $\mathcal{F} = 230,000$. The data for the cavity ringdown was kindly provided by Menlo systems.

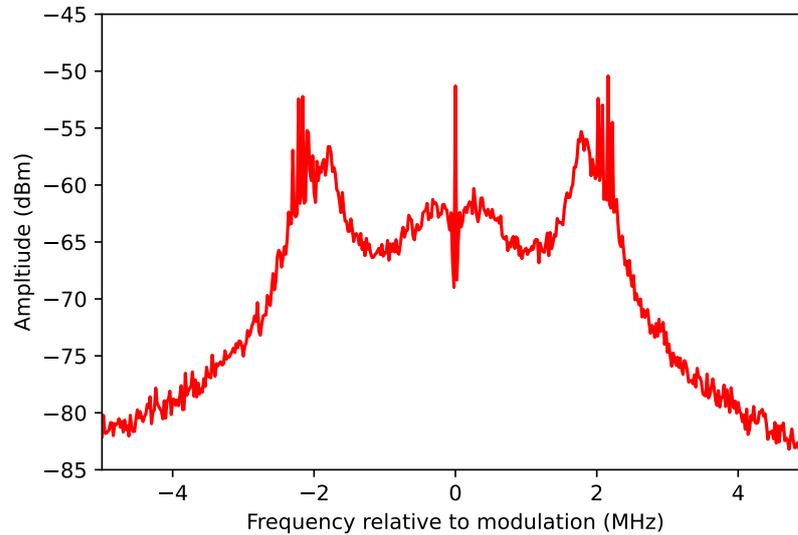


Figure 3.19: In loop error signal of the clock laser, where the frequency axis has been corrected relative to the 20 MHz modulation frequency. The lock is extremely stable, with the laser staying locked to the TEM_{00} mode for weeks at a time.

before coupling into an APC to PC⁴¹ connector fiber. The light then arrives next to the main chamber, the light is switched with an 80 MHz AOM in the +1 order, and the light is beam shaped before being sent to the atoms (not shown in diagram). The clock beam to the atoms can then be shaped for uniformity across the array, or focused to drive the transition at high Rabi frequencies.

The Doppler section uses a fiber with a PC end such that Doppler noise accumulated down the fiber path can be actively cancelled [157, 158]. The back reflected light from the end of the fiber is compared with a short free space optical reference path shown in the Doppler section of figure 3.17. The two paths are overlapped and fiber coupled to a photodiode, where the beat frequency between the two paths is measured. This beat frequency is then used to automatically adjust the frequency of the Doppler AOM away from 40 MHz to cancel the noise that is accumulated in the fiber.

3.2.6 Other lasers

Beyond cooling and imaging with 461 and 689 nm light, forming optical tweezers with 532 and 813 nm light, and Rydberg and clock excitation with 319 and 698 nm light, there are a few more laser systems in our laboratory worth mentioning. We now provide a brief overview of the other laser systems.

Repump and auto-ionisation lasers

There are four home-built ECDLs that have not been used in the work presented in this thesis, but will be important going forwards. These lasers are the repump lasers at 679, 688 and 707 nm (see figure 2.1), and the auto-ionisation laser at 408 nm.

There are several circumstances where it is necessary to move atoms between the states within the $5s5p\ ^3P_{0,1,2}$ manifold, and multiple available schemes to do so [79, 159, 160]. The repump lasers can be used to prevent buildup in the $5s5p\ ^3P_2$

⁴¹Physical contact (PC) connectors are not polished at an angle, so there is a back reflection from the output face of the fiber.

and $5s5p\ ^3P_0$ states when making a MOT⁴² [79], to incoherently pump to the clock state [92] and for state-resolved detection [131]. We have three home-built ECDLs, where each can be stabilised to the wavemeter (see section 3.2.6).

One of the benefits of strontium being a divalent atom, is that the second valence electron can be used to trigger ion production. When the first electron is excited to a high lying Rydberg state, the other electron is still left in the 5s orbital. This electron sees a potential like that of a Sr^+ ion, where there is a transition at 408 nm to excite the second electron up to the 5p state. The atom can then auto-ionise, as the two valence electrons have more than the ionisation energy between them. To drive this transition we have a homebuilt ECDL at 408 nm. The transition is quite broad, so frequency stability can be handled using a wavemeter (see section 3.2.6).

Wavemeter and frequency comb

An important part of all the laser systems is the ability to check the wavelength of the laser against a reference, such that the lasers can be accurately set and measured. For coarse measurements we measure the wavelengths using wavemeters, and for precise measurements we use an optical frequency comb.

All of our laser systems can be checked using a wavemeter⁴³ to measure the frequencies with a precision of 10 MHz and an accuracy of 60 MHz. This is a useful tool for checking that your laser wavelengths are at the correct frequency and are single mode. The wavemeter can be calibrated to the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition, by setting the 689 nm laser to the resonance condition destroying a NB MOT.

The 638 nm, 689 nm and 698 nm (clock) laser systems mentioned above are all PDH locked to high finesse cavities, this makes the lasers very stable and they

⁴²In this work we do not use repump lasers during the MOTs, as we do not need the increased atom number.

⁴³All of our lasers can use the same HighFinesse WS7-60 (VIS/Standard) wavemeter, with one exception. The 1549-1565 nm laser is not within the working range of the wavemeter, so instead can be coarsely set using a Bristol Instruments 671 Series Laser Wavelength Meter. We use this when significantly changing principal quantum number, before fine-tuning with the 638 nm light on the WS7.

operate with narrow linewidths. Given that lasers in our lab have linewidths of order kHz from the PDH locking, a wavemeter alone is not an accurate enough tool to assess the stability of these lasers, and cannot make precision measurements of the atomic transitions. To enable precision measurement of atomic transitions, and to assess the frequencies and linewidths of our lasers, we also have the facility of an optical frequency comb⁴⁴ [161], that works at 638, 689 and 698 nm. The frequency comb enables high precision measurements of the laser frequency, by beating each laser with a nearby ‘tooth’ of the frequency comb. Furthermore, the comb also allows us to constrain the linewidths. The frequency comb is a whole laser system in itself, where further explanation shall be left until it is required in appendix E.1. Note that when working with the comb we instead measure the 638, 689 and 698 nm with a higher precision wavemeter⁴⁵, so that we can identify which particular comb ‘tooth’ the laser is beating with.

3.3 *In-vacuo* equipment

To prepare, control and detect cold atoms, we also require a series of tools inside the vacuum chamber. An *in-vacuo* lens with $NA = 0.26$ is used to both create tweezers, and as part of the imaging system to collect the atomic fluorescence. We need control of both the magnetic and electric fields inside the vacuum chamber. MOT coils are used for both the quadrupole field used to create MOTs, and generating uniform magnetic fields to allow the clock transition through MIS. Electrodes give control of the electric field at the atoms, which allows the cancellation of undesirable DC Stark shifts of the clock transition and also of Rydberg states. As an additional detection method, there is also a microchannel plate (MCP) to allow the detection of ions in the vacuum chamber.

⁴⁴Toptica Difference Frequency Comb. This may also be referred to simply as the comb.

⁴⁵HighFinesse WSU-30 (IR1).

3.3.1 *In-vacuo* lenses

A pair of custom aspheric *in-vacuo* lenses are configured as shown in figure 3.1. These lenses can be used for both the forming of optical tweezers, and the collection of atomic fluorescence for imaging. The chamber-facing sides of the lenses are anti-reflection (AR) coated at all relevant wavelengths for trapping and imaging to minimise losses. The atom facing sides of the lenses are indium tin oxide (ITO) coated⁴⁶ to provide a conductive coating that reduces stray electric fields caused by the lenses [162].

The lenses are at a long working distance of 37 mm from the atoms, with a 20 mm clear aperture, and have numerical aperture $NA = 0.26$. The lenses can be used to form tweezers at both $\lambda_{\text{trap}} = 532$ nm and 813 nm, for which we achieve trapping waists of 1.28(1) μm [72] and 2.25(1) μm , respectively. The working distance used here is much larger than other similar experiments, particularly those based on glass cells with high NA lenses outside of the vacuum system, where the distance between the atoms and the nearest surface may be only a few mm. The long working distance was chosen here primarily for the benefit of reducing stray fields from the surfaces of the lenses. Large unwanted stray fields would cause significant problems for atomic clocks [163] and for experiments with Rydberg atoms [162, 164].

Even with the conductive ITO coating, adsorbates on the lenses [165], build up of charges [163], and work function differences between materials⁴⁷, can cause large stray electric fields. Longer working distances will minimise these effects, for example the DC Stark shift from small patch potentials will scale with the distance d between the surface and the atoms as d^{-4} . This problem is also more notable in strontium in that the method LIAD⁴⁸ [166] where regularly applied UV light is used to remove rubidium atoms from surfaces [38] has not been demonstrated to work for strontium or ytterbium. The application of UV light can also lead to charging

⁴⁶Coatings applied by Diamond Coatings.

⁴⁷There is a difference in work function between ITO and steel which can cause a stray field, though this stray field should be relatively stable.

⁴⁸Light induced atomic desorption.

of the ITO coating [167, 168].

The long working distance also offers better access, allowing a line of sight between the atomic cloud and the MCP so that ions produced by the experiment can be detected, and allowing orthogonal MOT beams in the horizontal plane for efficient trapping. The disadvantage of the long working distance is the reduced numerical aperture of the lenses, which limits our possible collection efficiency of fluorescence from the atoms.

In all experiments described in this work we use a single in-vacuo lens for both generating the tweezer traps, and collecting the atomic fluorescence (unless otherwise stated). The second in-vacuo lens is for recollimating the trapping light so that it can be imaged after the chamber for analysis of the trap depths of each tweezer when we are forming arrays.

3.3.2 Magnetic field coils

There are two sets of coils used to apply magnetic fields to the atoms in the main vacuum chamber, the MOT coils and the compensation coils.

The MOT coils are found inside the chamber, above and below where the atomic cloud forms. The MOT coils are used to produce either the quadrupole field required for trapping atoms in MOTs using the anti-Helmholtz configuration, or to produce uniform magnetic fields in \hat{z} direction using the Helmholtz configuration. The large uniform magnetic fields in the Helmholtz configuration are required for exciting the clock transition in bosonic strontium through MIS (see section 2.3.3). The type of magnetic field produced is controlled using a relay that can reverse the direction of the current in one of the coils, which changes the field between anti-Helmholtz and Helmholtz configurations [72].

The compensation coils are three pairs of coils surrounding the outside of the main chamber. Each pair can be used to apply a small uniform magnetic field along each of the lab axes. These coils can be used to cancel stray magnetic fields in the experiment. The compensation coils are more often used to change the location of the field zero when forming MOTs. This means that we can shift the location of the MOTs in 3D, most notably for optimising the overlap of a MOT with the optical

tweezers as mentioned in section 3.1.4. This is very useful, as the position where an optimised optical tweezer forms depends upon the *in-vacuo* lens, and the position the MOT naturally forms depends on the MOT coils. These two positions may not be close enough to each other to transfer atoms between the MOT and tweezer well, so the compensation coils provide an approach to better overlap them.

The power supplies used for the MOT coils⁴⁹, the \hat{x} and \hat{y} compensation coils⁵⁰ and the \hat{z} compensation coil⁵¹ are all stable current sources, with setting accuracy of $0.3\% \pm 0.5$ mA, and low RMS output noise.

3.3.3 Electrodes

There are twelve electrodes inside the vacuum chamber for applying electric fields to the atoms in any direction. They are positioned between the MOT coils in two sets. Six electrodes sit atop the bottom MOT coil, and six below the top MOT coil. The electrodes are in a hexagonal split-ring geometry [72], and are at 15° to our lab axes so that four of them line up with the Microchannel plate (MCP)⁵²

The electrodes are connected via feedthrough pins to a multichannel high precision low noise voltage source.⁵³ The voltage source has twelve outputs, one to each of our electrodes, which are connected via a shielded cable⁵⁴ directly to the vacuum feed-through pins. When using the voltage source our electrodes can be independently set to voltages within the range ± 300 mV, where each output shares a common ground in order to try to avoid stray voltages. The outputs to each electrode have a typical offset error of ± 25 μ V, with short term fluctuations on the scale of ~ 20 μ V. The twelve output channels are well isolated from each other with,

⁴⁹TTi QL564P.

⁵⁰TTi QL355TP.

⁵¹TTi QL564P.

⁵²The MCP is at 45° to the \hat{x} lab direction to avoid obstructing the optical axes, where for convenience the electrodes are also rotated such that a minimal number of electrodes are needed to direct ions to the MCP, as detailed in the thesis of Ryan Hanley [72].

⁵³Stahl Electronics, BS 1-12-0.3.

⁵⁴Mil-C-26482, 2.4 m cable.

typical cross-talk between the channels to each electrode of 10 ppm. The voltage source offers high precision control of electric fields, which is of great importance for both Rydberg experiments [162, 164] and for atomic clocks [163].

The achievable electric field strengths with the ± 300 mV voltage source are ~ 100 mVcm⁻¹, where this conversion arises from the configuration of the electrodes, which we will discuss further in section 6.2. Preliminary results before purchasing the precision voltage source indicated that the stray electric field is less than half the achievable range of voltage source, where the voltage source was chosen to be able to comfortably cancel the stray field, while not having the higher noise levels of voltage sources with larger output ranges.

3.3.4 Microchannel plate (MCP)

To enable the detections of any ions produced in Rydberg experiments, we have a microchannel plate (MCP)⁵⁵ in the main chamber. It is mounted on the inside edge of the chamber, halfway between two viewports such that it doesn't block any optical access. The MCP in the chamber has not been necessary for any results in this thesis, and initial testing showed excess ringing in the detected signals. They are however introduced here as due to their relevance in the electric field control, and the potential for single ion detection in future work.

3.4 Imaging

There are several cameras on the optical table around the main chamber to act as detectors for both atomic fluorescence and to image trapping light. The SPAD array in particular is a critical component to our experiments, as it is used as the primary detection method for imaging atoms.

⁵⁵Hamamatsu F4655-14.

3.4.1 SPAD array

A key piece of equipment to the measurements in chapters 4 and 5 is our main camera. We use a SPAD array⁵⁶ as the main detector for imaging 461 nm fluorescence from atoms in the vacuum chamber. The fluorescence is collected via the *in-vacuo* lens (see section 3.3.1), before being magnified onto the SPAD array. The choice of the SPAD array as our detector is a key factor in how we choose to implement our imaging protocols as it has different properties compared to more typical detectors such as CCDs⁵⁷, EMCCDs⁵⁸ and CMOS⁵⁹ cameras.

The SPAD array is comprised of 2048 pixels, in a 64×32 grid. The pixels have a pitch of $150 \mu\text{m} \times 150 \mu\text{m}$, each containing a $30 \mu\text{m}$ diameter circular active detection area. The fill-factor of each pixel is therefore only $\frac{\pi}{100} = 3.14\%$, so to prevent the majority of the fluorescence being lost to the region outside of the active area, the camera comes with a fitted microlens array (depicted in figure 3.1) that increases the effective fill factor to $\frac{\pi}{4} = 78\%$. The quantum efficiency of the camera peaks for blue light, where at our imaging wavelength of 461 nm the performance is in the region of 45 – 50%. Combining the effective fill factor and quantum efficiency together gives a photon detection efficiency of $\sim 38\%$. The microlens does restrict the acceptance half angle of the SPAD to 4%, above which the collection efficiency will fall, though this can be avoided with reasonable setup of imaging optics [169].

The SPAD array offers frame rates as high as 96 kHz, recording photon counts in a series of consecutive frames with negligible inter-frame dead time, enabling monitoring of fluorescence as a function of time on very short ($\sim 10 \mu\text{s}$) timescales. There is negligible readout noise on each frame, and a typical dark count rate of ~ 100 counts per second [169]. There is some variation over the array of the dark count rate as can be seen in figure 3.20, where there are a few ‘hot’ pixels with

⁵⁶Micro photon devices MPD SPC3

⁵⁷Charge-coupled device

⁵⁸Electron multiplying CCD

⁵⁹Complementary metal-oxide-semiconductor

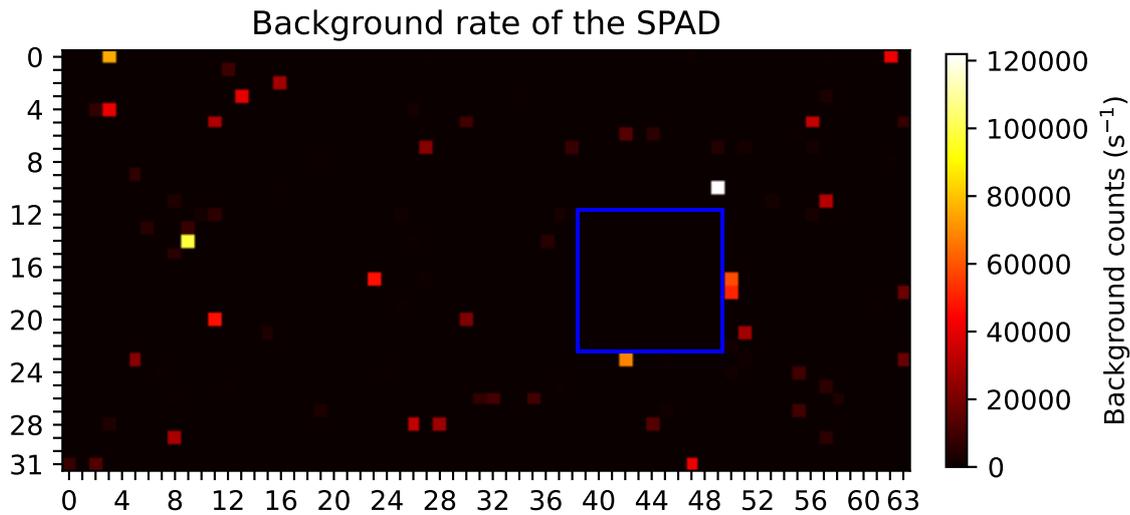


Figure 3.20: Background count rate for the SPAD array. An 11×11 pixel region is highlighted in blue to show an area of the SPAD with a dark count rate that does not exceed 140 s^{-1} .

higher dark counts [72], though these only account for $< 5\%$ of the pixels⁶⁰. The high dark counts on some pixels do make those individual pixels unsuitable for single atom imaging, however there are large regions of the SPAD array that can be used instead, so the hot pixels can be avoided. On our particular SPAD array there is an 11×11 pixel region highlighted in figure 3.20, where the typical dark count rate does not exceed 140 counts per second. This would be an ideal region of the sensor to image a 6×6 square tweezer array with spacing of two SPAD pixels.

The SPAD is protected by a series of filters to reduce stray light from wavelengths other than 461 nm reaching it. For the results of chapter 4 using a 532 nm tweezer, the SPAD is protected by a bandpass filter around 460 nm⁶¹, and a notch filter to block 532 nm⁶². For the results of chapter 5 using 813 nm tweezers, the SPAD is protected by two bandpass filters around 460 nm⁶³. The optical depth (OD) is a measure of what fractions of photons are allowed through, on a logarithmic scale of

⁶⁰98 of the 2048 pixels have a dark count rate of > 200 counts per second (double the typical background rate).

⁶¹Edmund optics 88010.

⁶²Thorlabs NF533-17.

⁶³Edmund optics 88010 $\times 2$.

base 10. The protection when using the 532 nm tweezer was OD4⁶⁴ for most stray light, and OD10⁶⁵ for stray tweezer light. The protection when using the 813 nm tweezer was OD8⁶⁶ for stray light.

To make the most of the SPAD array when imaging the atoms, one should make use of the low dark count rate. As the dark counts are effectively linear with time, this means that one should take the shortest exposures possible. The SPAD lends itself to fast imaging methods, though if a longer duration is required then multiple frames or gated operation can be used to avoid excessive build-up of dark counts.

The 150 μm pixel size is much larger than the size of typical detectors such as CCDs which normally have pixels that are only a few microns in size. Contrasting the size of SPAD pixels against the typical size of optical tweezers and their separations being only a few microns, one can see to resolve multiple tweezers the SPAD must image through a high magnification imaging system. This will not be necessary in chapter 4 where we only use a single tweezer, but will be considered in detail in chapter 5 where we image tweezer arrays. For the images shown in chapters 3 and 4 we image onto the SPAD with a $\times 10.8$ magnification imaging system formed of the *in-vacuo* lens and a 400 mm achromatic doublet [72], as shown in figures 3.1 and 4.1. For the images in chapter 5 we begin with the same imaging system, before moving to the upgraded $\times 48.6$ magnification imaging system which shall be described in detail there and is shown in figure 5.3.

3.4.2 CCDs

A more conventional kind of detector to use than a SPAD array is a CCD⁶⁷. We have two key CCDs in our experiment, used for the imaging of atoms and of trapping light. These roles are taken on by the Pixelfly⁶⁸ camera and the DMK⁶⁹ camera,

⁶⁴OD4 = 99.99% blocked.

⁶⁵OD10 = 99.99999999% blocked.

⁶⁶OD8 = 99.999999% blocked.

⁶⁷Charge coupled device.

⁶⁸PCO Pixelfly QE.

⁶⁹DMK 33UX273.

respectively.

As a secondary detector for imaging atomic fluorescence, we have a Pixelfly camera. The Pixelfly offers a slightly superior quantum efficiency of 55 – 60% around 461 nm, and does not have the associated losses of the microlens array, but comes with the drawback of having readout noise⁷⁰. The Pixelfly has much smaller pixel size ($6.45 \mu\text{m} \times 6.45 \mu\text{m}$) than the SPAD, enabling it to work well in lower magnification systems and take on a different role. It can be used to image atoms on different viewports (where all optics used are outside of the vacuum system), for example when aligning the tweezers (see section C.1.2). It also becomes integrated into the permanent setup in section 5.3.2, as an alternative low-magnification imaging method for studying the MOTs.

In chapter 5 we will require the ability to image the recollimated trapping light used to make the tweezer array. This will be used to allow us to equalise the trap depths of the tweezers as part of an adapted [36, 39, 40] Gerchberg-Saxton algorithm [170]. For this we use the DMK camera, which has the smallest pixel size ($3.45 \mu\text{m} \times 3.45 \mu\text{m}$), to help in resolving the trapping light used to make adjacent tweezers in arrays.

3.5 Summary

We have given an overview of the complex experimental apparatus required to study cold strontium in tweezers, with application from Rydbergs to clocks. The many laser systems have been described, particularly where they have changed from previous work, except where details of the setup have been left for where they are required in later chapters. The vacuum apparatus has been introduced, including the *in-vacuo* electric and magnetic field control and lenses. We also introduce the cameras, most notably the SPAD array with which we will typically image our atoms. The preparation of the cold atomic sample loaded into tweezers has been detailed, and from here we are now ready to discuss experiments using strontium in tweezers in chapters 4, 5 and 6.

⁷⁰The typical RMS readout noise is $7e^-$.

Chapter 4

Number-resolved imaging in a tweezer

This chapter presents our results detecting single and multiple atoms in a single 532 nm tweezer. The key data presented in this chapter can also be found in one of our publications [43]. Single atom control of divalent atoms in optical tweezers has so far been demonstrated with both strontium [43–48] and ytterbium [49–51], where here we demonstrate number-resolved imaging of strontium.

The loading of the atoms into a tweezer has already been detailed in section 3.1.4, where we overlap a NB MOT with a shallow tweezer, allow the NB MOT to fall away, and then ramp the tweezer to its final trap depth. We begin this chapter discussing the necessary apparatus used to image the atoms that have been loaded into the tweezer, and what fluorescence we would expect to detect. We then discuss the specifics of the Sisyphus cooling techniques used to control the temperature of the atoms under imaging light. Next we introduce a model with which we can generate expected histograms of atom fluorescence, before defining the fidelity with which we can observe single and multiple atoms within the tweezer from our data. Finally we present our results imaging single and multiple atoms in a tweezer, using three different imaging schemes including fast single atom detection in just 200 μs ,

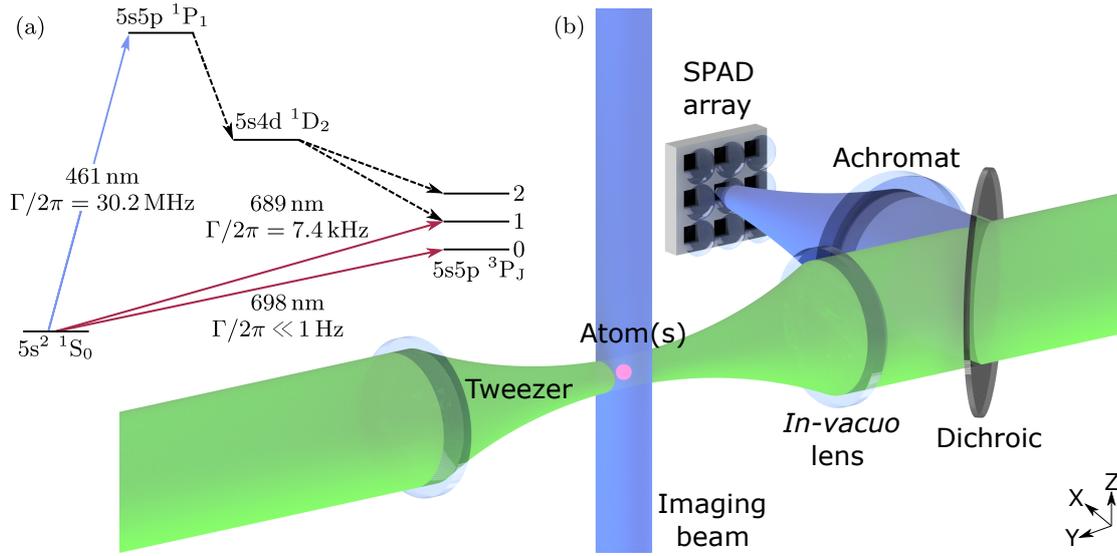


Figure 4.1: (a) Part of the strontium energy level diagram showing key transitions. (b) Strontium atom(s) are trapped in a single 532 nm optical tweezer formed by a $\text{NA} = 0.26$ *in-vacuo* lens. The atom(s) are imaged by a 461 nm probe beam, with the fluorescence collected through the same *in-vacuo* lens used to create the tweezer. The fluorescence is separated by the dichroic, and imaged onto the SPAD array via the 400 mm achromat and the microlens array. This figure is modified from a figure in our publication [43], where the 3D model of the system is the work of Ryan Hanley.

and optimised imaging with single atom fidelity of $0.998(2)^1$.

4.1 Imaging atoms

In order to measure the atom number in the tweezer, we apply a probe beam at 461 nm (on the $5s^2 \ ^1S_0 \rightarrow 5s5p \ ^1P_1$ transition) and collect the fluorescence. The fluorescence is collected using the imaging system shown in figure 4.1, where a $\times 10.8$ magnification imaging system formed of the *in-vacuo* lens and a 400 mm achromatic doublet [72]. The number of detected photons will be correlated to the atom number in the tweezer, so the counts detected can be used to identify the atom number up

¹ $99.8 \pm 0.2\%$.

to a confidence level of the fidelity.

4.1.1 461 nm probe

Atoms are imaged on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition with a 461 nm probe beam in the vertical ($-\hat{z}$) direction onto the atoms. The waist of the beam at the atoms is 0.4 mm, though both the power and detuning of the beam are variable parameters in the following experiments². The probe beam is linearly polarised along the (\hat{x}) direction, which is orthogonal to both the direction of propagation (\hat{y}) and polarisation (\hat{z}) of the tweezer. Therefore we image on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1\ |m_j| = 1$ transition, where the dipole radiation pattern of the atomic fluorescence is oriented such that the maximum possible fraction of the emitted light is collected by the *in-vacuo* lens.

Strontium atoms in a 532 nm tweezer experience a large differential AC Stark shift on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, resulting in a splitting of the $|m_j|$ levels [72]. The splitting of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition is particularly acute as we work in a deep tweezer where $\frac{U_0}{k_B} = 7.5$ mK, and the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1\ |m_j| = 1$ ($|m_j| = 0$) transitions shift by 26 ± 1 (56 ± 1) MHz/mK respectively [43]. The trap depth is 7.5 mK throughout this chapter, where other trap depths of 1.4 mK and 3.0 mK will be considered in appendix B.1 when studying the Sisyphus cooling mechanism in figure B.1(a-b). In this chapter detunings of the 461 nm probe are given relative to the empirically measured centre of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1\ |m_j| = 1$ transition in the tweezer for trap depth used.

4.1.2 SPAD imaging

For the data presented in this chapter a single pixel of the SPAD array was used to collect data for the fluorescence from the atom(s) in the tweezer. The light from the atoms was carefully focused to maximise the fluorescence gathered onto this one pixel, which has been chosen for having a low dark count rate and being near the centre of the SPAD array. A single shot image of a single atom is shown in figure

²These parameters will be specified where appropriate, where we either use a weak probe with a scattering rate of $0.14\ \mu\text{s}^{-1}$ or a strong probe with a scattering rate of $11.5\ \mu\text{s}^{-1}$.

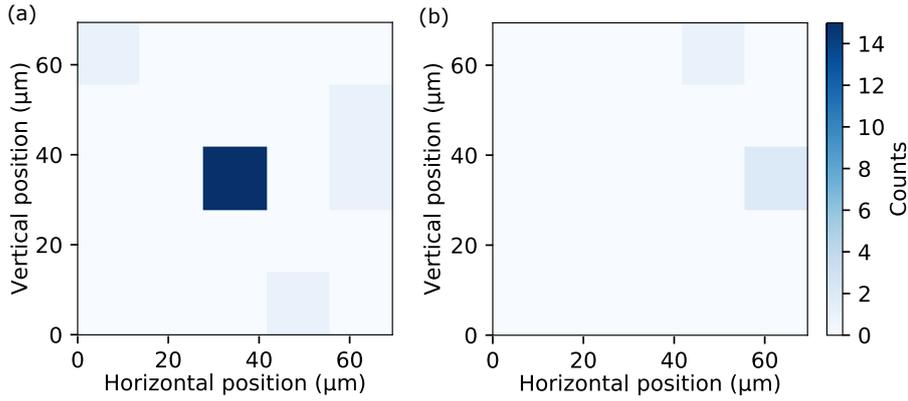


Figure 4.2: (a) Single shot image of a single atom with fidelity > 0.99 for a 1 ms exposure time. (b) Single shot image of an empty trap with a 1 ms exposure time. We will return to this data later in section 4.5.2 once we have introduced more of the framework used in acquiring and analysing this data.

4.2(a), with a corresponding image of an empty tweezer in figure 4.2(b). The data behind identifying the single shot image as a single atom will be discussed later in section 4.5.2.

Imaging with $\times 10.8$ magnification onto the large $150\ \mu\text{m}$ diameter SPAD pixels³ (see figure 4.1) means that each SPAD pixel corresponds to $14\ \mu\text{m}$ at the plane of the atoms. Given the small tweezer waist of $1.28(1)\ \mu\text{m}$ [72], one would expect that almost all of the fluorescence collected by the SPAD array would be collected on a single pixel.

Unfortunately we find that when comparing atomic fluorescence images with background images only $\sim 90\%$ of the light above the background level is collected by the target pixel, with spilling of light onto the eight nearest neighbours. This spilling of the fluorescence remains after a series of optimisations of the imaging system, including replacement of a strained dichroic, checking that the fluorescence is within the acceptance angle of the SPAD, and translating the relative positions of the SPAD array in all three dimensions.

We believe that this inability to focus a greater fraction of the fluorescence onto a single pixel is due to aberrations caused by the design of the *in-vacuo* lens

³As detailed in section 3.4.1, the $150\ \mu\text{m}$ diameter microlenses collect the light for each pixel.

at 461 nm. The high order aspheric coefficients of the lens strongly affect the shape towards the edge of the clear aperture of the *in-vacuo* lens, in a way that is beneficial for focusing Gaussian beams, but detrimental to collecting nearly-uniform emitted 461 nm fluorescence. For the results in this chapter the loss of fluorescence to neighbouring pixels is acceptable, but this will become problematic at higher magnifications when resolving nearby tweezers in chapter 5. The problem of the aberrations and the implemented solution to this will be discussed in more detail in section 5.3.

In order to keep the background counts in our imaging system low, we only use a single SPAD pixel for imaging our atoms. While this does sacrifice some of the emitted light, it is not worth including neighbouring pixels as the commensurate increase in background counts from using many SPAD pixels would reduce discrimination between empty and loaded tweezers.

4.1.3 Photon collection

An understanding of the imaging system is important to planning a sensible experiment. The efficiency of the imaging system used here is detailed in the thesis of Ryan Hanley [72], where there are two main losses associated with the imaging system. The primary loss is that the relatively low numerical aperture ($\text{NA} = 0.26$) *in-vacuo* lens used to collect the light can only collect a small fraction of the light as it does not make up much of the solid angle that the atoms emit into. The secondary loss is that each SPAD has a collection efficiency of $\sim 38\%$. There are also other smaller losses associated with the coating of optics etc. These factors together mean that we expect to collect just over 1 in every 200 photons. When imaging on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition there is a 1:20,000 branching ratio [45] into the strongly anti-trapped⁴ $5s4d\ ^1D_2$ state (see figure 2.1), where the atom will be lost from the tweezer. This puts a limit onto how many photons we can collect per atom ($\ll 100$) without there being a significant risk of losing the atom.

⁴This state is strongly anti-trapped in a 532 nm tweezer, where the atoms will be repelled by the tweezer light. In chapter 5 in an 813 nm this state will be trapped.

4.2 Sisyphus cooling

During imaging, the application of the 461 nm probe light will heat the atoms within the tweezer potential. We now consider the application of laser cooling during imaging to counteract this heating, which is important for both maintaining a cold atomic sample, and for accurate single atom imaging.

4.2.1 Motivation

Imaging in the tweezer with a 461 nm probe beam causes the atoms to be heated. To counteract this, we can apply Sisyphus cooling. In this chapter we work in very deep ($\frac{U_0}{k_B} = 7.5$ mK) tweezers, where for any reasonable imaging scheme we do not expect to be able to heat the atoms out of the tweezer, as loss via the $5s4d\ ^1D_2$ state would become too significant first⁵. The reasons to Sisyphus cool in such deep tweezers are to maintain a cold atomic sample in the tweezer, and to prevent the rate of atomic fluorescence from varying in time.

Keeping the atom(s) in the tweezer cold is important for their use afterwards, a cold single atom is more useful for subsequent experiments than a hot one. It also will be important in the ability to identify a single atom in the first place, where atoms heated in the trap potential will experience a different fluorescence rate under the 461 nm probe beam. This variable scattering rate arises due to the tweezer not being magic on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1\ |m_j| = 1$ imaging transition. We have set the detuning of the 461 nm probe beam relative to the AC Stark shifted transition near the bottom of the trap potential (the region occupied by the atoms). The 461 nm probe beam can heat the atoms through photon recoil, allowing them to climb the trap potential, which will decrease the AC Stark shift that they will experience. The change in Stark shift reduces the frequency of the light which would be resonant with these atoms, and effectively blue-detunes the probe beam from its original position, changing the rate at which photons are scattered.

⁵Heating out of the tweezer must be considered in section 5.5, where the trap depth is much shallower ($\frac{U_0}{k_B} = 0.4$ mK).

The change in scattering rate is problematic for repeatedly identifying the atom number, where we would like to be able to say that a particular number of counts will likely correspond to a particular atom number. If the scattering rate has changed, the number of detected counts in the second image will change, making misidentifying the atom number more likely. With the appropriate application of Sisyphus cooling the scattering rate of 461 nm light should be constant, and the identification of atom numbers more reliable.

4.2.2 Method

There are two recently demonstrated approaches to cooling strontium in tweezers using 689 nm light on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition; resolved-sideband cooling [44], and Sisyphus cooling [45, 131]. To cool in our tweezers we use Sisyphus cooling, of which there are two different mechanisms depending upon whether we are working in a 532 nm tweezer or 813 nm tweezer (in chapters 4 and 5 respectively), which are shown in figure 4.3. Sisyphus cooling the atoms in the tweezer utilises the fact that the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition is not magic in either a 532 nm or 813 nm tweezer⁶; there is a differential AC Stark shift⁷ on the transition.

In a 532 nm tweezer the shift of the ground state ($5s^2\ ^1S_0$) is greater than that of the excited state ($5s5p\ ^3P_1$). 689 nm light is applied that is red-detuned of the resonance condition at the centre of the trap. This 689 nm light is then on resonance with hotter atoms that are higher up the trap potential and experience a smaller differential AC Stark shift. As atoms are heated by the 461 nm imaging light, they climb the trap potential until the 689 nm light approaches resonance with the atoms. The atom is then excited to the $5s5p\ ^3P_1$ state, where it then moves in the trap potential, and then randomly de-excites. The atom is most likely to de-excite closer to the centre of the tweezer than where it was excited⁸, where the different

⁶There is a magic wavelength at 914 nm for the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1\ |m_j| = 1$ transition [72, 91, 171].

⁷See equation 2.2.9.

⁸Most of the energetically accessible region of the tweezer is towards the centre, the de-excitation probability is not spatially dependant.

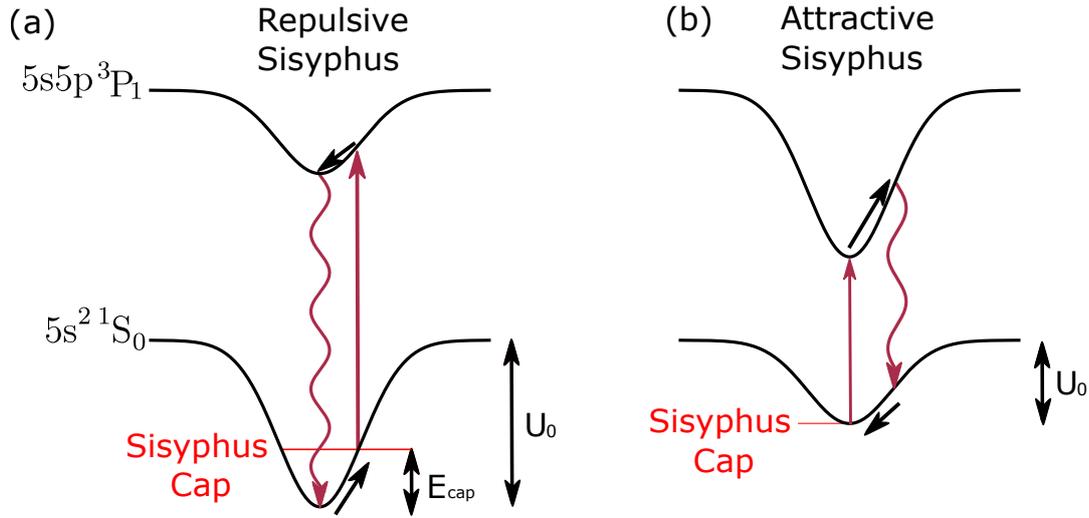


Figure 4.3: Diagram showing the two regimes of Sisyphus cooling. The differential AC Stark shift on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition is used such that 689 nm light can selectively excite atoms with energy E_{cap} . The excited atoms then move within the trap potential, typically de-exciting with a photon of more energy than that which was initially absorbed, decreasing the atoms kinetic energy and thus cooling it. (a) Repulsive Sisyphus cooling which we use in 532 nm tweezers, where the ground state has greater polarisability than the excited state. Atoms which heat to the Sisyphus cap at energy E_{cap} enter the Sisyphus cooling mechanism and are then cooled. The mechanism acts as repeller in energy space [45], pushing atoms below the Sisyphus cap to lower temperatures. (b) Attractive Sisyphus cooling which can be used in 813 nm tweezers (see chapter 5), where the ground state has lesser polarisability than the excited state. Here the detuning is set to resonance near to the bottom of the trap, where the Sisyphus cooling mechanism acts as an attractor in energy space [131], drawing atoms towards the bottom of the potential. Note that in both (a) and (b) the $5s5p\ ^3P_1\ |m_j| = 0$ and 1 levels have different polarisabilities, so the excited state of each subfigure refers to choosing a particular $|m_j|$ level. At 532 nm both $|m_j|$ levels have smaller polarisability than the ground state [72, 91] so either would be a repulsive Sisyphus mechanism, whereas at 813 nm the $|m_j| = 0$ (1) level has a smaller (larger) polarisability than the ground state [91] so operate on the repulsive (attractive) mechanism. For Sisyphus cooling at 532 nm we will use $|m_j| = 0$ on the repulsive mechanism, at 813 nm we will use $|m_j| = 1$ on the attractive mechanism. Please note that this figure is not to scale.

differential AC stark shifts means that the atom will emit a photon of greater energy than it absorbed. This removes energy from the atom, effectively cooling it. The detuning of the 689 nm light sets how high the atoms can climb up the trap potential before the cooling light acts upon it, setting a temperature limit in the tweezer, the ‘Sisyphus cap’. The optimal frequency for the cooling light will be highly dependant on both the tweezer wavelength and the trap depth of the tweezer. The mechanism discussed here is repulsive Sisyphus cooling [45], shown in figure 4.3(a), which we will use here in this chapter. Repulsive Sisyphus cooling works in a 532 nm tweezer as the ground state has a greater trap depth than the excited state.

To apply Sisyphus cooling to the atoms we use the red MOT beams shifted to a different frequency, though alternatively a single beam [131] of suitable polarisation could be used. The $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ $|m_j| = 0$ and $|m_j| = 1$ states experience significantly different differential AC Stark shifts in a 532 nm tweezer. The $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ $|m_j| = 0$ transition is near-magic [72], so small detunings of the cooling light are needed to set a Sisyphus cap, and this is the transition which we use cool for the 532 nm tweezer. The study and optimisation of Sisyphus cooling parameters is detailed in appendix B.1, where the effects of Sisyphus cooling are studied as functions of trap depth, Sisyphus detuning and 461 nm probe scattering rate.

Sisyphus cooling can also be done in 813 nm tweezers in the converse situation where the ground state is shallower than the excited state. Sisyphus cooling in the 813 nm tweezers can use an attractive Sisyphus cooling mechanism [131], shown in figure 4.3(b). The Sisyphus cap is set near the bottom of the trap potential, where once again the emitted 689 nm photon is likely to have more energy than the absorbed one⁹. We will discuss attractive Sisyphus cooling further in section 5.5, where we will want to use the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ $|m_j| = 1$ transition.

⁹The Sisyphus light is set to be resonant at the bottom of the trapping potential where the energy of the absorbed photon is minimal. If the atom is anywhere else in the trapping potential when de-excitation occurs, more energy will be emitted than was absorbed, thus cooling the atom.

4.3 Histograms

To determine the atom number from the detected fluorescence on the SPAD array, multiple repeats of the experiment should be taken, building up a histogram of the distribution of detected counts. As we will see, the histogram will form peaks around counts corresponding to each atom number, and thus instances where we detect those numbers of counts can be characterised as having a particular atom number. To better understand how this works, we now describe the underlying probabilities with which we expect each number of counts, and describe how we can simulate the expected counts.

4.3.1 Composite Poisson distribution

When repeating the experiment many times to identify atom number from photon counts, each time a fresh tweezer is loaded with a random atom number. The trap is imaged to record a number of detected photons, then the tweezer is released and the experimental cycle repeats. The photon count measured each time will be random, and depend on three distributions: the number of atoms loaded into the trap, the number of scattered photons detected per atom, and the number of dark counts detected by the SPAD. Each of these can be modelled by a Poisson distribution¹⁰, $\Pi(\mu, n)$. Therefore the detected counts will be formed by a composite Poisson distribution, and will result in histograms with peaks in photon counts corresponding to different atom numbers.¹¹ This method is based upon the work described in references [72, 172].

The probability to load a particular number of atoms, N_{at} , is given by

$$P_{\text{at}}(N_{\text{at}}) = \Pi(\overline{N}_{\text{at}}, N_{\text{at}}) \quad (4.3.1)$$

where \overline{N}_{at} is the mean number of atoms loaded into the tweezer for the chosen loading parameters.

¹⁰Here we assume Poissonian loading of atoms, i.e. the absence of light-assisted collisions.

¹¹For an example of this, see actual experimental data of atomic fluorescence from the tweezer in figure 4.4.

The probability of detecting n_s counts on the SPAD from light scattered by the atom(s) when there are N_{at} atoms in the trap is given by

$$P_s(n_s)|_{N_{\text{at}}} = \Pi(\alpha N_{\text{at}}, n_s) \quad (4.3.2)$$

where α is the average number of detected photons per atom.

However, as we will run the experiment many times and there will be a distribution of different atom numbers loaded into the tweezer, we must sum over all possible values of N_{at} to find the general expression of $P_s(n_s)$. This is what will give rise to the multiple peaks in the distribution, which will be later used to identify different atom numbers from the detected counts. Combining the two distributions to find the distribution of detected counts from scattered photons yields

$$\begin{aligned} P_s(n_s) &= \sum_{N_{\text{at}}=0}^{\infty} (P_{\text{at}}(N_{\text{at}}) \times P_s(n_s)|_{N_{\text{at}}}) \\ &= \sum_{N_{\text{at}}=0}^{\infty} \Pi(\overline{N}_{\text{at}}, N_{\text{at}}) \Pi(\alpha N_{\text{at}}, n_s). \end{aligned} \quad (4.3.3)$$

In addition to the counts due to the atom(s) scattering photons described by equation 4.3.3, the detected counts will also include background (dark) counts. We will be able to identify the mean dark count rate, but will not know which particular counts correspond to an atomic fluorescence signal, and which are part of the background. The probability to detect n_b background counts is given by

$$P_b(n_b) = \Pi(\overline{n}_b, n_b) \quad (4.3.4)$$

where \overline{n}_b is the mean number of dark counts per exposure. Due to the low background of the SPAD compared to an EMCCD this will be a relatively low number of counts.

The probability to then detect n_{tot} counts on the SPAD is simply the sum of the scattered counts and the background counts.

$$n_{\text{tot}} = n_s + n_b \quad (4.3.5)$$

The probability of n_{tot} counts on the detector then depends on both $P_s(n_s)$ and $P_b(n_b)$ for all possible combinations of n_s and n_b that satisfy equation 4.3.5. The

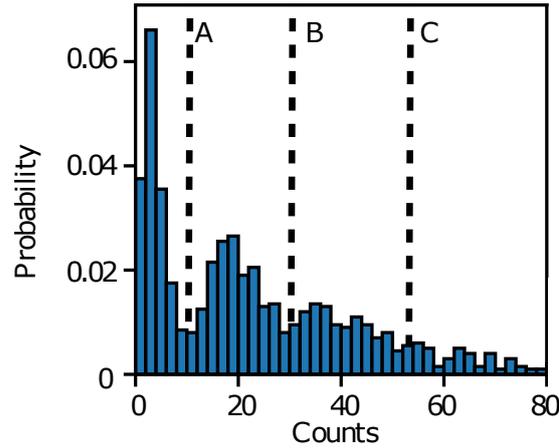


Figure 4.4: Actual experimental data of atomic fluorescence collected from a tweezer loaded with low atom number to illustrate what kind of distribution in photon counts we are going to see. There is a large narrow peak around $\bar{n}_b \approx 3$ corresponding to an empty tweezer, with broader peaks at $\alpha + \bar{n}_b \approx 20$ and $2\alpha + \bar{n}_b \approx 37$ corresponding to one and two atoms respectively. There are three dashed lines marking potential thresholds where we might expect the probable atom number to change. For example we could suggest that between dashed lines A and B likely corresponds to one atom, which will be discussed in greater detail in section 4.4. We will return to this data later in section 4.5.1 as figure 4.5(a) once we have introduced more of the framework used in acquiring and analysing this data.

probability of detecting n_{tot} counts is therefore given by

$$\begin{aligned}
 P_{\text{tot}}(n_{\text{tot}}) &= \sum_{k=0}^{n_{\text{tot}}} P_b(k) P_s(n_{\text{tot}} - k) \\
 &= \sum_{k=0}^{n_{\text{tot}}} \sum_{N_{\text{at}}=0}^{\infty} \Pi(\bar{n}_b, k) \Pi(\bar{N}_{\text{at}}, N_{\text{at}}) \Pi(\alpha N_{\text{at}}, n_{\text{tot}} - k).
 \end{aligned} \tag{4.3.6}$$

With equation 4.3.6 we arrive at the ability to model a histogram for many repetitions of the experiment. The shape of the histogram will form a series of bumps around likely photon counts such as \bar{n}_b , $\alpha + \bar{n}_b$, $2\alpha + \bar{n}_b$, etc. For an example to illustrate this point, see the peaks in the distribution of actual atomic fluorescence collected from the tweezer in figure 4.4.

4.3.2 Assumptions

There are underlying assumptions behind equation 4.3.6 which we must now consider. The first is that it is assumed that a Poissonian atom number is loaded, which means that there would be no light-assisted collisions causing two-body loss resulting in the atom number being reduced to zero or one (no sub-Poissonian loading) [42]. The absence or presence of light-assisted collisions is something that can be tested with our experimental data, where in the data presented in section 4.5 we see a distribution of atom numbers consistent with Poissonian loading.

The second assumption was explicitly made during equation 4.3.2, that each atom has an average fluorescence rate of α per image. This is a reasonable assumption that can be expected to be true, though if the image is comprised of multiple frames which are then compared, this must be approached with caution. Care must be taken because the fluorescence rate as a function of time within the image may vary due to heating effects, as will be discussed further in section 4.5.2.

The third assumption is that there is no loss of atoms during imaging, which is an assumption that we know to be false. As mentioned in section 4.1.3, there is a one-body loss mechanism when the imaging light is on. This will result in the atom number decreasing at random points during some imaging frames, and will lead to a filling between the peaks¹² in the histogram. The filling between peaks will reduce the fidelity with which atom numbers can be distinguished, and raise the possibility that the measured atom number will no longer be correct by the end of the measurement. Therefore the model must be modified to incorporate this one-body loss mechanism and the effects that it will have on the distribution of detected counts. A modified model for simulating histograms including loss is developed in appendix B.2, and will be used to compare with the final results in section 4.5. Due to the loss of atoms during imaging, the optimal imaging time becomes a balance between collecting sufficient photons to resolve the atom number, while scattering as few photons as needed to manage the probability of losing the atom.

¹²For example if we lose the atom half way through the frame, we will get fluorescence that corresponds to roughly 1.5 atoms.

4.4 Fidelity

We have now discussed what we expect to see in the detected photon distribution from a tweezer with a Poisson distribution of loaded atom number, and discussed the effects of loss to the $5s4d\ ^1D_2$ state. This leaves us ready to consider how we can determine the atom number from the fluorescence, and the confidence with which we can do that.

4.4.1 Thresholds

Once sufficient data is taken to produce histograms of the atom fluorescence, the number of detected counts can then be converted into the predicted atom number present in the tweezer by applying thresholds in the distribution. In our case we apply up to three thresholds, at A , B and C counts, placed between the peaks of each histogram to determine atom number, as shown in figure 4.4. For n_{tot} counts detected in a particular run of the experiment the atom number is then

$$N_{\text{at}} = \begin{cases} 0 \text{ atoms} & n_{\text{tot}} < A \\ 1 \text{ atoms} & A < n_{\text{tot}} < B \\ 2 \text{ atoms} & B < n_{\text{tot}} < C \\ m (> 2) \text{ atoms} & C < n_{\text{tot}} \end{cases} \quad (4.4.1)$$

where we always define the thresholds to be at non-integer counts to make clear any ambiguity where the counts equal the threshold. The case of ‘m’ (many) atoms is used for atom numbers > 2 where we cannot resolve peaks with the data to be presented later in section 4.5.

To make conversion of SPAD counts to atom number meaningful, the confidence in the measurement must be quantified. One could approach this by applying a model to the data, fitting each peak of the histogram applying the thresholds between each fitted peak to minimise the fraction of the tail of each distribution crossing the threshold. The suitability of this approach is reduced in the presence of the significant one-body loss mechanism filling the parts of the distribution between

the peaks, so *a priori* we cannot be certain of using this method. We therefore choose to develop upon the approach detailed in reference [44]¹³, comparing successive imaging frames to determine the fidelity with which the atom number is measured. The optimal position of each threshold is found by locating the position which minimises the number of instances where counts in successive frames cross the thresholds between frames. Once positions are optimised, the remaining cases where the threshold is crossed in successive frames indicates some kind of error.

The approach of comparing successive images is well suited to our imaging system, utilising the negligible readout noise of the SPAD array. Furthermore, throughout the data presented later in this chapter, the single atom exposures are divided into many frames. This allows us to choose optimal exposure times *a posteriori* by merging frames together, and allows us calculate the fidelities for multiple different exposure times from a single dataset.

4.4.2 Limiting infidelity and loss

We categorise misidentification of atom number into two distinct types of error: infidelity errors and loss errors. An infidelity error is when the atom number is incorrectly identified, a loss error is when the atom number has decreased during the imaging frame. Evidence of each will be seen by the detected counts crossing a threshold between subsequent frames. An increase in atom number can be confidently identified as an infidelity error, as the atom number in our tweezer cannot increase during imaging as there is no MOT present to re-load atoms from. However, a decrease in atom number could be either due to an infidelity error or a loss error. Therefore we cannot always be sure in an individual case which type of error occurred, and we must define how we derive each type of error from the data to carefully exclude contributions from the other kind of error. The separation of the errors into these two distinct forms is a worthwhile pursuit as not only are they results of

¹³The approach detailed here only works for cases where the tweezer is either empty or contains a single atom, i.e. in the regime of light-assisted collisions. We generalise this approach in the regime where multiple atoms can be loaded into the same tweezer.

different processes, but also that the one-body loss channel via the $5s4d\ ^1D_2$ state can be closed in an 813 nm tweezer, where the states are trapped and repumping is available. This means that in later work in 813 nm tweezers it is the infidelity error rate which will be the key value.

First we calculate the fidelity¹⁴ one would expect if we were in the regime of light-assisted collisions. This quantifies the limiting performance of our imaging system, where events of $n_{\text{tot}} > B$ are excluded, and defines the limiting single atom fidelity F_1 .

$$F_1 = 1 - \frac{P(01)}{P(0) + P(1)} \quad (4.4.2)$$

where $P(a)$ is the probability of loading a atoms, and $P(ab)$ is the probability of identifying a atoms in the first frame and b atoms in the second. The values of $P(a)$ and $P(ab)$ are extracted directly from the data. This limiting fidelity F_1 is useful for comparison between experiments, since it is independent of the mean atom number $\overline{N_{\text{at}}}$ in each dataset¹⁵. This definition of fidelity utilises the $P(01)$ term, which includes contributions of both misidentifying a single atom as an empty trap in the first frame, and misidentifying an empty trap as a single atom in the second frame.

The associated single atom loss rate is then defined by

$$L_1 = \frac{P(10) - P(01)}{P(1)} \quad (4.4.3)$$

where the difference between the rates of crossing the threshold in each direction identifies the fraction of the cases of $P(10)$ that are due to loss errors rather than infidelity errors.

Error on fidelities and losses are found by varying the positions of the thresholds. Moving the thresholds from their optimised positions by ± 1 counts results in deviations in the values of fidelity and loss, the size of which we take to be the associated errors for each value. This definition of error is chosen due to the position of

¹⁴Fidelity throughout this work is defined as being between 0 and 1, rather than as a percentage.

Discussion uses both fidelity (correctly identified atom number) and infidelity (misidentified atom number), where fidelity = 1 – infidelity and *vice versa*.

¹⁵Which is roughly set by adjusting t_{load} , but does not have fine control.

thresholds being the dominant source of uncertainty in the changes of atom number, as the thresholds are used to define the atom number.

4.4.3 Selective fidelity

To evaluate the accuracy with which the atom number N_{at} can be determined from the measured photon count distribution without parity projection (multi-atom loading regime), we define the selective fidelity F'_1 for single-atom ($N_{\text{at}} = 1$) detection. This selective fidelity is the fidelity which is achieved in each set of data for the mean atom number \bar{N}_{at} loaded, and represents the fidelity with which we could use the loaded single atoms. The selective fidelity is defined in appendix B.3, taking the form

$$F'_1(\bar{N}_{\text{at}}) \approx 1 - P(01) - \frac{P(12)}{1 - L_1}. \quad (4.4.4)$$

4.5 Results

While exploring how best to enable high fidelity detection of low numbers of atoms in our single 532 nm tweezer, multiple imaging schemes were trialled. The imaging approach needs to find the best way to image using our SPAD array. We need to leverage the negligible readout noise of the SPAD, while bearing in mind the lower photon collection efficiency.¹⁶

Given the loss to the $5s4d\ ^1D_2$ discussed in section 4.1.3, we expect to be able to collect $\ll 100$ photons per atom before there is significant chance the atom will be lost. The limiting performance is the same regardless of the imaging technique, though it is possible to do worse at single atom detection fidelity. Ways in which the fidelity can be made worse are through either the atoms being heated away from resonance, or by imaging for long times and increasing the number of background counts.

¹⁶The SPAD array has a lower detection efficiency than some EMCCDs. The long working distance *in-vacuo* lens also reduces the collection efficiency due to its relatively small NA of 0.26, compared to other systems which use $\text{NA} \geq 0.5$.

The three different imaging schemes focus on control of just heating, followed by control of just the background, and then finally controlling both issues in the same imaging protocol. Firstly in section 4.5.1 we use continuous Sisyphus cooling with a weak probe beam for effective control of the temperature. Secondly in section 4.5.2 a strong probe beam is used for fast imaging without any cooling, to collect with minimal background counts. Finally the benefits of both approaches are combined in section 4.5.3, where both the background and heating are minimised together.

Each set of data contains ~ 1000 repetitions of the experiment. The 461 nm probe beam settings will be detailed at the start of each subsection, including detuning from differential AC Stark shifted resonance, scattering rate, and saturation parameter $S = \frac{I}{I_{\text{sat}}}$. Atom number in each case is controlled to be in the single atom regime by carefully setting the overlap time of the NB MOT and the tweezer, typically this requires $t_{\text{load}} < 5$ ms. The results from each method are summarised at the end of this chapter in table 4.1.

4.5.1 Imaging with continuous Sisyphus cooling

The first approach to image single atoms is to use a weak 461 nm probe beam and Sisyphus cooling continuously together throughout the imaging sequence. The Sisyphus cooling beams are set to their optimal detuning from the bare resonance of $\Delta_{689}/2\pi = 7.7$ MHz. The probe beam is set -20 MHz off resonance of the AC Stark shifted imaging transition at the trap centre, with a scattering rate of $0.14 \mu\text{s}^{-1}$ ($S = 0.004$), well within the range of scattering rates ($< 0.30 \mu\text{s}^{-1}$) where the Sisyphus cooling is maximally effective as shown in figure B.2.

The data were taken in a series of fifteen consecutive 10 ms frames, allowing a variety of possible composite exposure times to be analysed.¹⁷ The histogram in figure 4.5(a) shows the number of counts obtained for a cumulative imaging time $t_{\text{exp}} = 30$ ms formed from the first three frames. While the loading rate is somewhat controlled by variation of t_{load} , the empirically chosen value here resulted in a signif-

¹⁷The negligible readout noise of the SPAD allows us to decide *a posteriori* the optimal imaging duration.

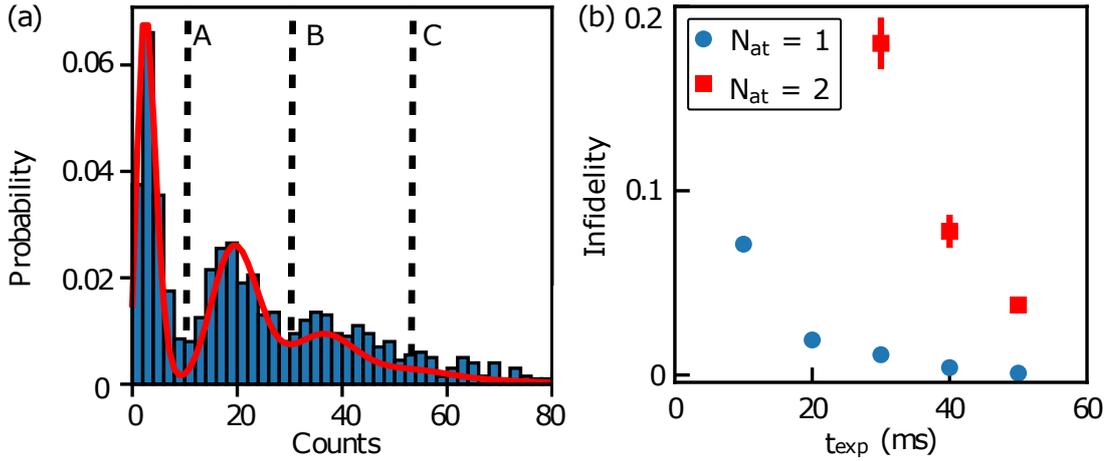


Figure 4.5: (a) A histogram of the number of detected counts for $t_{\text{exp}} = 30$ ms with a bin width of 2 counts is shown in blue, compared to the model shown in red. The optimised thresholds used to calculate the infidelities and loss from the data are indicated by the vertical dashed lines. (b) Limiting infidelity for one atom (blue circles) and selective fidelity for two atoms (red squares) for different exposure times. This figure is modified from a figure in our publication [43].

icant fraction of occurrences with multi-atom loading in the tweezer¹⁸. Three peaks are clearly visible, corresponding to $N_{\text{at}} = 0, 1, 2$ atoms in the tweezer. Not only can we clearly see from the data that light-assisted collisions are not dominant in our experiment, this also demonstrates that different atom numbers are resolvable. The data in the histogram is consistent with Poissonian loading of $\overline{N_{\text{at}}} = 1.2$ where $\alpha = 17$ photons are collected per atom per 30 ms exposure.

The values of fidelity and loss are calculated for histograms of a variety of different composite exposure times in figure 4.5(b). We see that the resolution between the peaks at different atom number improves with increased exposure time, drastically improving the fidelities in both the single and two atom cases. Ideally one would therefore choose long exposures to set the fidelity as required, however as shown in figure 4.6(a), an increased exposure time corresponds to a greater fraction of instances where atoms are lost from the tweezer. This is due to the increased

¹⁸To set the atom number more precisely one should use light-assisted collisions, though that precludes the ability to observe the 2 atom case.

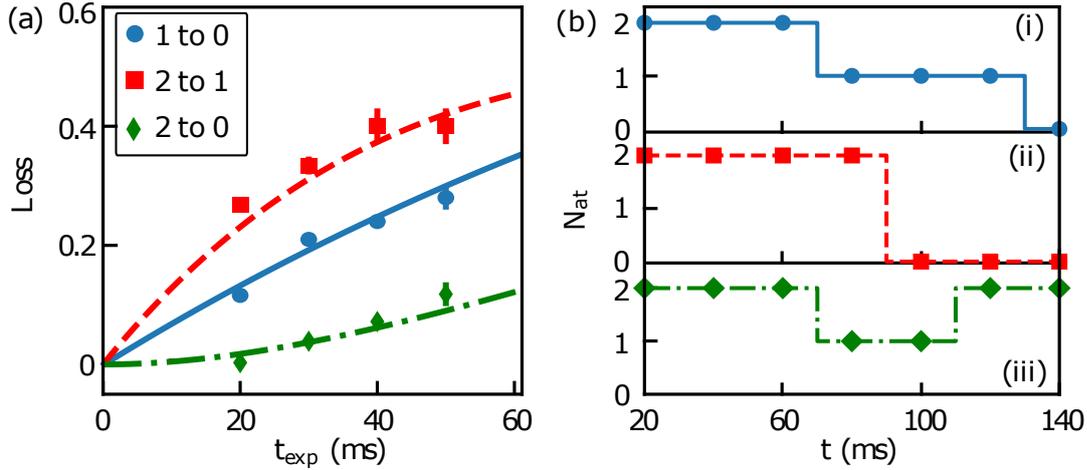


Figure 4.6: (a) Measured atom loss as a function of exposure time for N_{at} : $1 \rightarrow 0$ (blue circles), $2 \rightarrow 1$ (red squares) and $2 \rightarrow 0$ (green diamonds). Predicted losses via the $5s4d^1D_2$ state are plotted as lines against the data. (b) Trajectories of N_{at} using 20 ms frames, including examples of decay from (i) $2 \rightarrow 1 \rightarrow 0$ atoms (ii) $2 \rightarrow 0$ atoms and (iii) where an error occurs are shown. This figure is modified from a figure in our publication [43].

number of 461 nm photons being scattered during the image, and more loss to the anti-trapped $5s4d^1D_2$ state. Therefore as a compromise between these two considerations we choose an optimal composite exposure of 30 ms as in figure 4.5(a), where the single atom limiting (selective) fidelity F_1 (F'_1) = 0.989(3) (0.970(2)), and the associated single atom loss $L_1 = 0.211(8)$. The two atom selective fidelity and loss are $F'_2 = 0.83(2)$ and $L_2 = 0.373(11)$ respectively.

The losses associated with the data are shown in figure 4.6(a), we plot the losses from $1 \rightarrow 0$, $2 \rightarrow 1$ and $2 \rightarrow 0$ atoms as a function of exposure time. Three curves are added for the expected loss rates due to loss via the $5s4d^1D_2$ state. The model curves are not only consistent with our expected one-body loss rate for exposure times > 20 ms, but also in the $2 \rightarrow 0$ case show that we see no evidence of an additional two-body loss mechanism (i.e. we see no indication of light-assisted collisions being present).

The losses during imaging frames are further shown in figure 4.6(b), where three representative runs of the dataset have their detected atom number tracked for

seven consecutive 20 ms frames. In trace (i) we initially start with two atoms, before decaying to one atom then an empty trap in two discrete steps. In trace (ii) we initially start with two atoms, that by chance are both lost in the same frame, resulting in a sharp step to an empty trap. In trace (iii) we see an infidelity event, where there are two probable descriptions¹⁹. Either that there were two atoms in the tweezer for the full imaging duration, but they were misidentified as a single atom towards the middle of the imaging sequence, or there were initially two atoms of which one was lost, then this single atom was misidentified as being two atoms towards the end of the imaging sequence²⁰.

Imaging with a weak probe and continuous Sisyphus cooling provides high fidelity imaging of single atoms, as well as reasonable resolution to detect two atoms in the same tweezer. The key issue in the current tweezer system is that to resolve with the chosen fidelity, the loss terms are significant. These losses can be later addressed by moving to an 813 nm tweezer where the loss channel is closed, and the high fidelity should remain, making the fidelity values the key figure going forwards. The continuous Sisyphus cooling approach also ensures that the temperature of the atom(s) is kept under control during imaging, potentially an important aspect for their subsequent use in a 532 nm tweezer, and necessary for working in shallower tweezers.

The drawbacks of the continuous Sisyphus imaging approach centre around the need for a weak probe beam that does not overwhelm the cooling. This not only means we have to wait ~ 30 ms to measure that we have a single atom, but also leads to a higher background peak which reduces the fidelity. The centre of the background peak scales linearly with exposure time, because we do not suffer from readout noise but instead have a relatively high dark count rate (approximately 0.1

¹⁹Remember that reloading a second atom into the tweezer is not possible as there is no MOT present during the imaging sequence.

²⁰The descriptions of what occurred for specific runs of the experiment are of course not completely certain due to the finite infidelities in the assessment of atom number in each frame. However, with many events similar to each of these occurring, the general assessment of what generally happened in a class of events should be accurate.

counts per ms). We expect that this mean counts in the background peak could be approximately halved by doubling the probe scattering rate to the limit of where Sisyphus cooling is effective ($\sim 0.30 \mu\text{s}^{-1}$), but ideally we would like to decrease the background by more than a factor of two to maximise the fidelity even further. This leads us onto the following two imaging methods, where through changing the imaging sequence to include a strong 461 nm probe beam we significantly reduce the number of dark counts.

4.5.2 Fast single atom imaging

The second method with which we imaged single atoms seeks to minimise the number of dark counts. Here we work with a strong probe beam on resonance with the AC Stark shifted imaging transition, with a scattering rate of $11.5 \mu\text{s}^{-1}$ ($S = 0.14$). The drastically increased scattering rate mean that we require much shorter imaging times to resolve a single atom, but precludes the possibility of effective continuous Sisyphus cooling, as the scattering rate is over an order of magnitude higher where the cooling is maximally effective.

In this data the mean atom number loaded was much lower, $\overline{N_{\text{at}}} = 0.5$; the tweezer predominantly contained zero or one atoms. Images were taken in a series of ten consecutive $100 \mu\text{s}$ frames, where a histogram for a combination of the first two frames into a $200 \mu\text{s}$ frame is shown in figure 4.7(a). Here the mean number of background counts is $\overline{n_{\text{b}}} = 0.02$, resulting in a very sharp peak for empty traps²¹. The low likelihood of there being multiple background counts in the same frame greatly contributes to the resolution between the empty trap and single atom peaks, as it is highly likely for multiple counts that there must have been fluorescence from an atom. This imaging approach yields single atom limiting (selective) fidelity of F_1 (F'_1) = 0.989(6) (0.979(6)) for $200 \mu\text{s}$ frames. The small number of two atom events precludes reliable analysis of the two atom case.

The fidelity and loss as a function of imaging time is shown in figure 4.7(b),

²¹Note the large axis break in figure 4.7(a), the majority of the images of unoccupied tweezers result in 0 counts on the SPAD pixel!

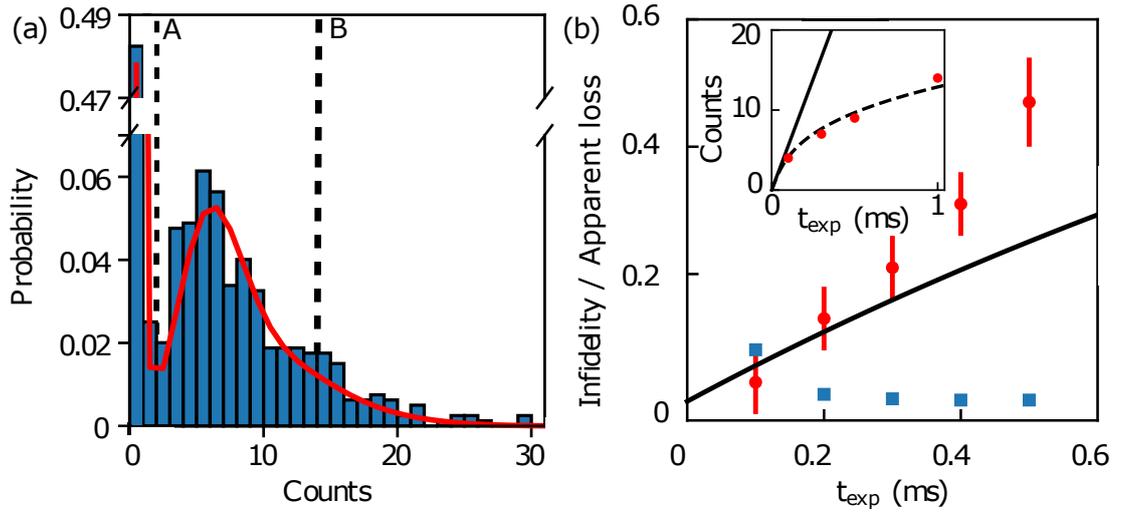


Figure 4.7: (a) A histogram of the number of detected counts for $t_{\text{exp}} = 200 \mu\text{s}$ is shown in blue, compared to the model shown in red. The optimised thresholds used to calculate the infidelities and loss from the data are indicated by the vertical dashed lines. (b) Limiting infidelity (blue squares) and apparent loss (red circles) for different exposure times. The expected loss via the $5s4d^1D_2$ state is shown (black line), the apparent loss rate increases beyond this line at longer exposure times where the effect of heating is more significant. The inset shows the detected counts (red circles) compared to predictions with (dashed line) and without (solid line) heating effects moving the probe beam off resonance and reducing the scattering rate. This figure is modified from a figure in our publication [43].

where for longer imaging times better fidelity is achieved at the cost of higher loss. For the chosen imaging time the loss is $L_1 = 0.13(5)$. The expected loss rate is shown as a black curve in the figure, where at longer times the measured loss significantly exceeds the theory. We attribute this measured loss to come from two components, firstly the genuine one body loss to the $5s4d^1D_2$ state, and secondly apparent loss due to probe induced heating. Single shot images of a single atom and of an empty trap for 1 ms exposure times are shown in figure 4.2, with fidelity $F_1 > 0.99$.²²

²²We do not have sufficiently long data to calculate fidelity for a 1 ms exposure, though the fidelity for 0.5 ms exposure time gives limiting fidelity of 0.998(1), and fidelity typically increases with exposure time. The longer exposure time is chosen to show the extremely low background.

In the absence of any Sisyphus cooling, probe induced heating drives the atom up the trapping potential, changing the experienced AC Stark shift and moving the atom out of resonance with the probe beam. This leads to a reduction in the collected fluorescence signal at longer imaging times, where the reduced fluorescence looks like loss. No atoms are lost through the heating, but the reduced scattering rate can cause them to emit below threshold levels in the second composite frame which is used to calculate fidelity. In the inset of figure 4.7(b) the average counts per atom as a function of imaging time is recorded as the red data points. The solid line corresponds to the expected rate of fluorescence, and the dashed line shows the rate from a simple heating model. The heating model assumes that each scattered 461 nm photon drives the atom up the trap potential by its recoil energy ($\frac{E_r}{k_B} = 0.514\mu\text{K}$), and represents the maximal reduction in fluorescence possible due to heating.

Imaging with a strong probe beam offers a greatly reduced number of background counts, utilising the negligible readout noise of the SPAD. However the heating of the atoms both reduces the fluorescence and increases the apparent loss rate. The strong probe beam offers high fidelity single atom imaging in just 200 μs , yielding comparable fidelity with exposures two orders of magnitude shorter than in section 4.5.1. While this approach offers single atom detection in much shorter timescales, the associated heating is undesirable. Ideally to maximise fidelity one would combine the two approaches to have low background imaging but without excessive heating, which shall be the topic of the section 4.5.3.

4.5.3 Enhanced fidelity imaging

Our final approach to imaging single atoms seeks to combine the benefits of the previous two, the temperature control of section 4.5.1 with the low background of section 4.5.2. To do this we require a long time of Sisyphus cooling, but a short exposure time. Therefore we apply a series of short, strong imaging pulses, separated by long periods of cooling.

The probe pulses share the same power and detuning as in section 4.5.2, the probe is on resonance with a scattering rate of $11.5\ \mu\text{s}^{-1}$ ($S = 0.14$). The difference

here is that the probe is pulse on for $41.6 \mu\text{s}$ ²³ with a 4.16 ms repetition time (1% duty cycle). Continuous Sisyphus cooling is applied throughout the sequence, but is only effective when the probe is off. Here the Sisyphus cooling is set at $\Delta_{689}/2\pi = 6.7$ MHz, 1 MHz lower than in the continuous Sisyphus cooling case. This lower value of 6.7 MHz was found to be effective for this set of pulsed imaging parameters. In this imaging method, the probe heats the atom(s) for a finite pulse time before the cooling can act, so a higher temperature cap (lower Sisyphus detuning Δ_{689}) reduces the chance of atom(s) heating beyond the Sisyphus cap during the imaging pulses.

The imaging cycle consists of fifty of the $41.6 \mu\text{s}$ pulses, each of which is carefully centred within a $52 \mu\text{s}$ SPAD frame²⁴ to collect the fluorescence. This slightly longer exposure time ensures that valuable scattered photons from the atom are still collected within the exposure time, for only a slight increase in background counts above the minimum required. Histograms are formed by combining the imaging frames corresponding to the probe being on, which for this particular set of parameters is every 80th frame.

One may wonder at the choice of very particular imaging times. The imaging frame duration has been chosen to be in the regime where a single pulse does not cause too much heating, while also being a multiple of the $10.4 \mu\text{s}$ hardware readout time of the SPAD [173] to ensure accurate timing. High quality timing is critical for this imaging method, as such the frames where the probe beam is off have also been analysed to ensure that there are no unexpected counts, confirming that the experiment was properly synchronised.

In figure 4.8(a) we combine the frames containing the first 20 pulses to form a histogram. We see a well resolved low-loading ($\overline{N_{\text{at}}} = 0.6$) histogram. This approach offers a narrow background peak with $\overline{n_{\text{b}}} \sim 0.1$ from the 1.04 ms of exposure time, along with a high number of photons detected per atom. For 20 pulses this approach yields single atom limiting (selective) fidelity of F_1 (F'_1) = 0.998(2) (0.991(2)) with a

²³This corresponds to an exposure four times longer than the shortest possible frame taken with the SPAD array.

²⁴This corresponds to an exposure five times longer than the shortest possible frame taken with the SPAD array.

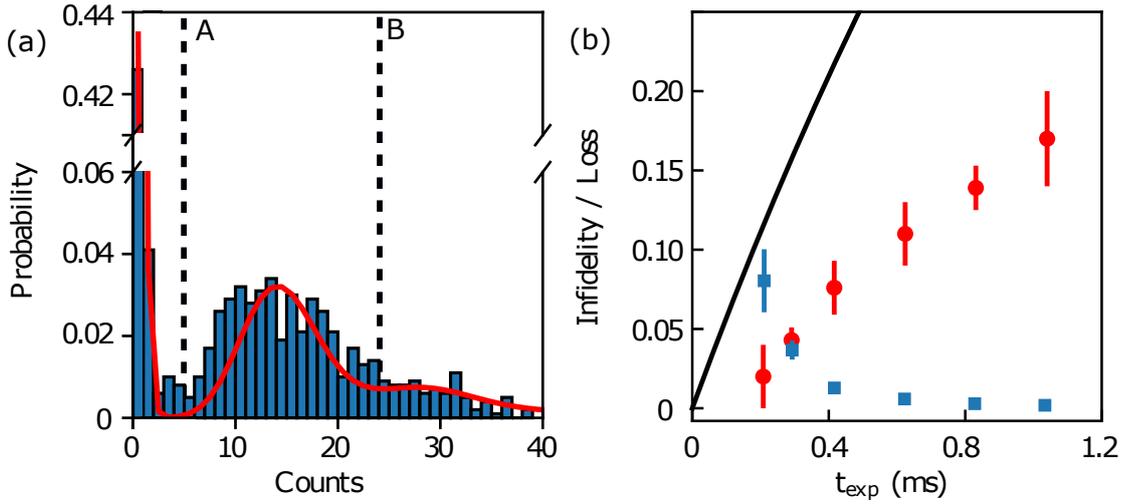


Figure 4.8: (a) A histogram of the number of detected counts for 832 μs of probe on time is shown in blue compared to the model shown in red. The optimised thresholds used to calculate the infidelities and loss from the data are indicated by the vertical dashed lines. (b) Limiting infidelity (blue squares), and measured loss (red circles), and expected losses (black line) are shown for different probe on times. This figure is modified from a figure in our publication [43].

corresponding loss $L_1 = 0.139(10)$. This method clearly offers a better fidelity than either of the previous approaches, where each approach is contrasted in table 4.1. As in section 4.5.2, the small number of two atom events in this dataset once again precludes reliable analysis of the two atom case.

The loss (L_1) as a function of imaging time is shown in figure 4.8(b), where we can see that the loss rate is actually below that predicted for our imaging parameters. We find that both the loss and the number of photons detected per atom are below the expected values by approximately the same factor. This indicates that while the interleaved sequence of probe pulses and cooling offers significant advantages, the cooling is not entirely effective and there is notable heating off resonance. This is likely due to a combination of single probe beam pushing atoms up the trap potential during pulses, and the 1 MHz higher Sisyphus cap permitting higher temperatures within the tweezer.

The one key drawback retained by this approach is that the sequence of probing and cooling requires 83.2 ms, though through optimisation this time could poten-

tially be reduced while still achieving the same fidelity. This could be done by increasing the average scattering rate of the light to the point where the Sisyphus cooling begins to be overwhelmed, though this may not be at the same average scattering rate of $\sim 0.30 \mu\text{s}^{-1}$ limit shown in figure B.2.

It is likely that a more exhaustive study of imaging and cooling parameters could better control the heating mechanism to offer improved fidelity for shorter imaging sequences. However with this type of experimental sequence having a large parameter space to explore, and that the optimum parameters will depend upon the tweezer wavelength used, we leave such a study for further work. The results here are more than sufficient to demonstrate the usefulness of the imaging technique, where for our research goals it is now more pertinent to move towards arrays of 813 nm tweezers in chapter 5.

4.6 Summary

In this chapter we have demonstrated single atom and number-resolved imaging in a single deep (7.5 mK) tweezer at $\lambda_{\text{trap}} = 532 \text{ nm}$, through a variety of imaging techniques. We introduced a model to simulate histograms in atomic fluorescence where the $5s4d^1D_2$ loss channel is open, along with generalised fidelity calculations for the multiple atom regime. We demonstrated the increased effectiveness of Sisyphus cooling in deeper tweezers, working in both continuous and pulsed modes of operation.

The imaging techniques used in this chapter are summarised in table 4.1. In the first case using a weak probe beam with Sisyphus cooling we demonstrated number-resolved detection of atoms in a single tweezer, with two atoms resolved with a fidelity of 0.83(2). In the second case a strong probe beam without cooling was used to enable fast single atom detection in just 200 μs with a fidelity of 0.989(6), thanks to the extremely low dark count rate of our SPAD. Finally we demonstrated an interleaved sequence of strong probe pulses with Sisyphus cooling, utilising the negligible readout noise of the SPAD to enhance fidelity further, showing single atom imaging fidelity of 0.998(2).

Method	t_{exp} (ms)	\overline{N}_{at}	F_1	$F'_1(\overline{N}_{\text{at}})$	L_1	$F'_2(\overline{N}_{\text{at}})$	L_2
Sec 4.5.1	30	1.2	0.989(3)	0.970(2)	0.211(8)	0.83(2)	0.373(11)
Sec 4.5.2	0.2	0.5	0.989(6)	0.979(6)	0.13(5)	-	-
Sec 4.5.3	0.832*	0.6	0.998(2)	0.991(2)	0.139(10)	-	-

Table 4.1: Comparison of different imaging methods, labelled by the section in which the method is detailed. *This sequence includes 0.832 ms of probe-on time in a 1.04 ms combined exposure, taking a total time of 83.2 ms due to fluorescence only being collected on every 80th frame.

Moving forwards from these results, we now look to move from a single 532 nm tweezer to an array of 813 nm (magic wavelength) tweezers. With multiple tweezers we will want to enforce full single atom control, precluding the possibility of multiple atoms loaded into the same tweezer²⁵ using light-assisted collisions [44, 45]. There is also the potential to generate enhanced, near-deterministic single atom loading in tweezer arrays as has recently been demonstrated in ytterbium [50].

At 813 nm the imaging should be greatly improved by the closing of the $5s4d\ ^1D_2$ loss channel, where the state is no longer anti-trapped, and repumping lasers can be used to bring atoms back into the imaging cycle. We will need to work at much lower trap depths due to the reduced ground state polarisability, lack of laser power, larger tweezer waists, and division of power between traps. We therefore will need to return to Sisyphus cooling, though instead using the attractive mechanism (see figure 4.3(b)), and will have to ensure we do not heat out of the much shallower traps. To this end, an adapted version of the interleaved imaging and cooling technique detailed in section 4.5.3 appears to be the most promising imaging technique to bring forwards when we move to the 813 nm tweezer array in chapter 5.

²⁵Having two atoms in the same tweezer is a potentially interesting environment in which to work when using a single tweezer, but when working with a tweezer array having one atom per site at fixed separation is the desired experimental platform.

Chapter 5

Magic wavelength tweezer arrays

In this chapter we implement a holographically generated array of tweezers at the magic wavelength of the clock transition. This begins with the construction and characterisation of the new tweezer system. We then describe the setup required to produce our near-arbitrary tweezer array configurations using a spatial light modulator (SLM). The tweezer array system can have sites separated by just a few microns, so a new custom imaging system is implemented for collecting atomic fluorescence on the SPAD to be able to resolve neighbouring tweezer sites. We then demonstrate some of the tweezer arrays that we can produce, before discussing the imaging protocols that will be best suited to the tweezer array and our applications.

5.1 Magic wavelength tweezers

Implementing the magic wavelength tweezer array required a rebuild of the tweezer system¹ to work at the magic wavelength of the clock transition, and to implement having multiple tweezers using the SLM. We start by considering the benefits and drawbacks of using a magic wavelength array in contrast to the single 532 nm tweezer in chapter 4, before detailing the construction of the system.

¹Aside from the *in-vacuo* lens, which is designed to work for both 532 nm and 813 nm tweezers.

5.1.1 Benefits and drawbacks

Benefits of working at 813 nm

The primary benefit of working with 813 nm tweezers is that 813 nm is the magic wavelength of the clock transition [33, 34] ($5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$), which means that the differential AC Stark shift of the transition is zero. This avoids the clock transition being shifted away from the natural frequency. Furthermore, this means that the frequency of the clock transition does not depend upon the depth of the optical tweezer, so tweezers of different depths will have the same clock transition frequency.²

An additional benefit of moving to 813 nm tweezers is that here the $5s4d\ ^1D_2$ state is trapped, unlike at 532 nm where it is anti-trapped and atoms that decayed on the $5s5p\ ^1P_1 \rightarrow 5s4d\ ^1D_2$ channel were lost. The change in trapping arises because the polarisability of each state is a function of the trapping wavelength, and at 813 nm the sign of the polarisability is opposite for this state. As the $5s4d\ ^1D_2$ state is trapped, it can remain in the tweezer long enough to decay to the $5s5p\ ^3P_{0,1,2}$ manifold, where upon the application of repump lasers (679 and 707 nm) almost all of the atoms that go through the $5s4d\ ^1D_2$ decay channel return to the ground state [131].

Drawbacks of working at 813 nm

There are drawbacks in switching from a single 532 nm tweezer to an 813 nm tweezer array. Some of the reasoning has been touched upon already in section 2.2.3, but it will be considered in more detail here. Two key drawbacks are larger separation required between adjacent tweezers, and the reduction in the possible trap depth. The larger separation between adjacent tweezers is necessary due to the increased waist of tweezers at longer wavelengths. The reduction in trap depth is significant

²It may be a useful tool in addressability of particular sites to sometimes work slightly off-magic such that sites of different trap depths have slightly different clock transition frequencies, but in any case that would involve wavelengths very close to the magic wavelength that are still 813 nm, and many orders of magnitude within our laser tunability.

and arises from four main factors: increased tweezer waist, reduced ground state polarisability, reduced availability of laser power, and splitting the power between multiple tweezers. These four factors combined necessitate working in much shallower tweezers than we did when working in a single 7.5 mK deep tweezer at 532 nm in chapter 4. Each factor will be discussed in more detail below.

When forming a tweezer at 813 nm we fundamentally cannot achieve as tight a tweezer as at 532 nm. The increase in trapping wavelength results in the size of a diffraction limited spot being larger, so both the waist and Rayleigh lengths of the tweezer are necessarily larger. At 532 nm the optimised tweezer waist was measured to be 1.28 μm [72], whereas at 813 nm the best that we can hope to achieve is 1.67 μm .³ There is the potential to create tighter confinement by using superpositions of Laguerre Gauss modes instead of a simple Gaussian beam, which can be done using an SLM [174, 175], though this is left for future work.

The ability to trap atoms at 813 nm is greatly reduced due to the reduced polarisability of the ground state [72, 91]. The overall trend (excluding near other resonances) is that the ground state polarisability decreases the longer we go in wavelength away from the strong $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition (see figure 2.2). The polarisability is only 37 % of that at 532 nm [72], meaning that combined with the best case increase in trap waist, 4.6 times more power would be needed per tweezer to achieve the same trap depth.

The final, and most significant, decrease in possible trap depth is that we seek to make arrays of tweezers rather than just one. In the focal plane where a 2D array of tweezers would form, the laser power must be divided between each of the tweezers, greatly reducing the depths that can be achieved. If one aims to make arrays of tens of tweezers, this effect is significant, though it would of course have been a factor regardless of trapping wavelength. This is combined with the practical constraint that the maximum laser power we can produce at 813 nm (~ 2.7 W) is lower than our available laser power at 532 nm (~ 18 W).

³This value is based upon Zemax OpticStudio simulations.

Combining 532 and 813 nm systems

In moving towards making clock measurements in a tweezer array, changing to the 813 nm magic wavelength is a necessity and benefits far outweigh the drawbacks. That being said, it is worth noting that a combination of the two systems could be a useful tool in the future, though this is left to future work. A 532 nm tweezer could be added into the tweezer system via a dichroic⁴, with a pair of acousto-optic deflectors (AODs) in the beam path to make the 532 nm tweezer movable. This could be a useful tool for shuttling atoms between tweezer sites if we were to do rearrangement or arrays [36, 55, 56]. Movable tweezers can also be implemented with a piezo controlled mirror [176].

Another approach using two tweezer wavelengths has recently been demonstrated, overlapping green⁵ and 813 nm tweezer arrays in the same system [46, 177]. These tweezer arrays are used alternately to make use of the trap depth at short trapping wavelengths where it is needed, and the magic wavelength for clock measurements when it is not [46, 177]. This approach also extends to loading from green tweezer arrays into other trapping geometries such as 813 nm optical lattices [22, 178]. For our purposes at the moment an array at 813 nm is sufficient, though further use of 532 nm tweezers remains a potential avenue of research for the future.

5.1.2 Making a tweezer

To produce the 813 nm light for making the magic wavelength tweezer we use a titanium sapphire laser (Ti:Sa)⁶. The Ti:Sa is etalon locked to provide basic stability to its wavelength (linewidth < 5 MHz), and is stabilised further by locking to a wavemeter⁷. In this future this may be better stabilised by either a scanning transfer cavity [179] or ultrastable cavity.

The output of the laser is single passed through an 80 MHz AOM for switching

⁴Thorlabs DMLP650L longpass dichroic mirror, 650 nm cut-on.

⁵Trapping wavelengths such as 515 nm or 532 nm.

⁶M Squared SolsTiS.

⁷HighFinesse WSU-30 (IR1).

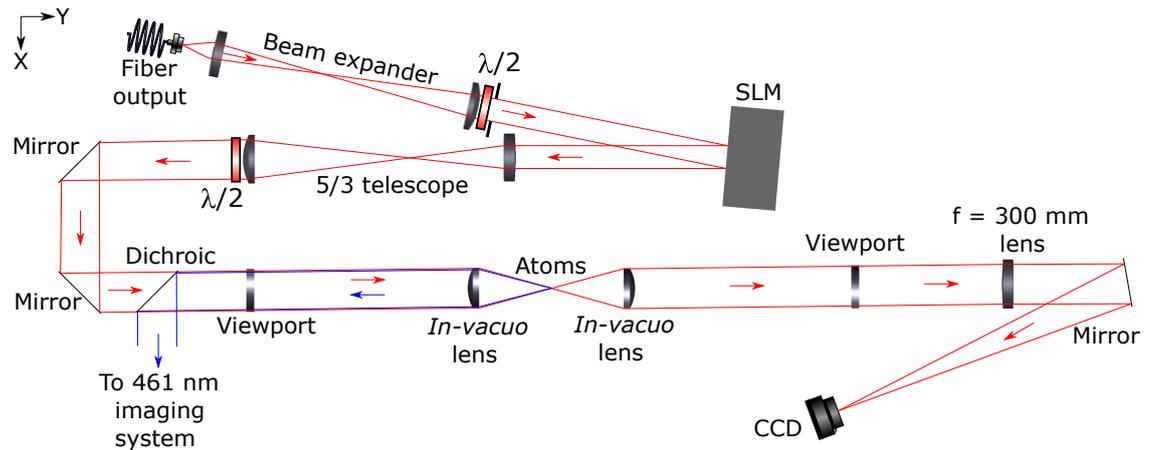


Figure 5.1: Experimental layout diagram for the tweezer array system. Note this diagram is simplified, with some folding mirrors removed for ease of viewing. The 461 nm imaging system will be later shown in figure 5.3.

the light, where the +1 order is coupled into a photonic crystal fiber⁸. The AOM acts as a switch for the tweezer being turned on/off. By varying the RF power sent to the AOM, the power to the tweezer, and thus the trap depth, can be adjusted.

The 813 nm beam path after the fiber is shown in figure 5.1. The fiber outputs into free space into a cage system of two lenses that act as a beam expander to produce a 12 mm diameter collimated beam⁹. A HWP then horizontally polarises the beam as required by the next optic, the SLM¹⁰. The beam is incident upon the window of the SLM at a small angle ($\leq 5^\circ$) [180], where it reflects off the surface. It then passes through a telescope formed of two lenses of focal lengths 75 mm then 125 mm to magnify the beam by a factor of $\frac{5}{3}$ such that the beam diameter is then 20 mm¹¹. At the end of the telescope there is another HWP, allowing us to choose the polarisation of the tweezer. The collimated beam reflects off steering mirrors, a periscope, and then is transmitted through a dichroic¹², before entering

⁸NKT photonics LMA-PM-15.

⁹A 12 mm beam will fill in the SLM window on its narrower axis. There is an iris at the end of the beam expander to apodise the beam at 12 mm diameter.

¹⁰Spatial light modulator, see section 5.2.2.

¹¹Corresponding to the clear aperture of the *in-vacuo* lens.

¹²Thorlabs DMLP490L longpass dichroic mirror, 490 nm cut-on.

the vacuum chamber. The light is then focused with the *in-vacuo* lens to form a tweezer ~ 38 mm from the lens in the centre of the vacuum chamber. The trapping light is then re-collimated using the second *in-vacuo* lens, where it can then be imaged onto a CCD¹³ or Shack–Hartmann wavefront sensor¹⁴ to characterise the trapping light.

The initial alignment of the tweezer system, loading and imaging of atoms in the tweezer, and optimisation of the tweezer waist, are considered in detail in appendix C.1. We are able to prepare tweezers with waist $2.25(1)$ μm loaded with many atoms.

5.2 Generating tweezer arrays

So far we have considered using the *in-vacuo* lens to produce a single tweezer, similar to the previous work done at 532 nm [72]. This has been using the spatial light modulator (SLM) simply as a folding mirror in the setup. We now consider the different ways in which tweezer arrays can be generated, before discussing the specifics of using the SLM based approach which we have chosen to use.

5.2.1 Methods to make arrays

Focusing a Gaussian beam with a high NA lens will generate a single optical tweezer at the focus, as we used in chapter 4. In order to produce a tweezer array, the light field sent to the high NA lens must be changed so that several beams focus after the lens at different transverse (and perhaps longitudinal) positions to form multiple tweezers. We consider four approaches to achieve this: fixed optics, spatial light modulators (SLMs), acousto-optic deflectors (AODs), and digital micromirror devices (DMDs).

The fixed optics approaches to producing arrays are the least flexible, as it does not give the opportunity to vary the spacing between tweezer sites. One method

¹³DMK 33UX273.

¹⁴Thorlabs WFS30-5C/M.

to generate arrays with fixed optics is by using a diffractive beam splitter (DBS), to separate the beam into a grid of traps. A potential benefit to some systems is that this has the option of using calcite displacers to have different polarisations for adjacent tweezers [181], though this is not of interest in our apparatus. Another fixed optic approach is the use of microlens arrays to generate the tweezer arrays [182], though this approach leads to very short working distances and larger tweezer sizes, due to the small size of the microlenses reducing the achievable NA. Microlenses can also be combined with SLMs [183], AODs [184] and DMDs [185] to produce more versatile systems. In this work we want to create changeable tweezer configurations, using long working distance lenses to minimise shifts from stray electric fields that are undesirable for both clock [163] and Rydberg [162, 164] experiments, and with tight confinement of tweezer sites. We therefore do not consider fixed optics any further, instead focusing on comparing SLMs, AODs, and DMDs.

SLMs act as a diffractive optic in reflection, imprinting a phase pattern upon the reflected light, allowing the generation of multiple tweezers [186], as well as control of their two dimensional [39, 40] or three dimensional [35, 36] positions to produce near-arbitrary array configurations [35, 36]. SLMs also give control over the shape of each tweezer spot for the correction of aberration, or the production of more exotic trap shapes such as Laguerre-Gauss beams [174, 175] and bottle beams [62, 85]. As the SLM is a purely diffractive optic, all of the tweezers produced by the SLM have exactly the same trapping wavelength. Furthermore, SLMs can be used to control the individual trap depth of each tweezer for producing homogeneous arrays [39, 40]. One could also produce deliberately inhomogeneous arrays if one wanted have site selective imaging [48], or to work slightly off the magic wavelength to make individual tweezer sites addressable. The key drawback of an SLM is that the refresh rate is too slow for dynamic control of the tweezer positions; in the case of our SLM this rate is only 60 Hz. If movement or rearranging of the tweezers is required then an additional movable tweezer beam must be added. SLMs will be considered in much greater detail when they are required in section 5.2.2.

In an AOD the light is passed through a crystal which is driven by multiple RF tones. Each RF tone causes the formation of a diffracted beam, which is then

used to make a distinct tweezer. A single AOD produces a one dimensional array of tweezers [37, 38, 131], but by placing two AODs consecutively on orthogonal axes, a two dimensional array is produced [44, 45], though the geometry is restricted to square/rectangular arrays. AODs can have trap separations varied while atoms are loaded by changing the frequencies of the RF tones [36, 37, 55, 63, 176]. However, with 2D arrays produced by AODs the traps must be moved as a row or a column at a time,¹⁵ as there are not enough degrees of freedom to give independent control of each trap position. The combination of RF tones is also non-trivial, and requires careful selection of the phases to produce traps of the correct depth. Arrays produced by AODs have slightly different trapping wavelengths for each tweezer [47], due to the absorption or emission of the different frequency phonons from the different the RF tones. The slightly different trapping wavelengths is an issue if one wanted all tweezers to be exactly at the magic wavelength, as is desirable for implementing an atomic clock in the array.

One further alternative approach to consider is the use of a digital micromirror device (DMD). DMDs are deformable mirrors that form a binary phase mask, capable of generating arbitrary 2D tweezer arrays [57, 187]. The 2D arrays offer a more varied geometry than crossed AODs, but have not demonstrated the 3D arrays achievable in SLMs. The DMD acts similarly to an SLM, though instead produces a binary phase mask.¹⁶ The refresh rate of a DMD is 20 kHz, orders of magnitude above SLM refresh rates, allowing DMDs to be used to shuttle atoms [57]. Shuttling atoms using a DMD allows the movement of multiple tweezers at a time, unlike using a crossed AOD as a shuttling beam. However, DMDs have the drawback that they must shuttle atoms using finite steps, which can lead to heating and loss of atoms.

SLMs offer arbitrary trap geometries in three dimensions and control of the tweezer shape and depth, but lack dynamical control during an experimental se-

¹⁵Except when a second 2D AOD setup is overlapped to provide a movable shuttling tweezer [55, 56].

¹⁶As we will see in section C.2.1, our SLM has 215 discrete phase values available, in contrast to the 2 phase values of a DMD.

quence. AODs give control of tweezer positions during an experimental sequence, but do not have individual control of each site, do not have arbitrary geometries or three dimensional control, and all tweezers are not at the exact same wavelength. DMDs offer an interesting approach including 2D arbitrary arrays like the SLM and some rearrangement like the AOD, but lack the 3D arbitrary arrays that we desire. As we desire three dimensional tweezer arrays with arbitrary geometry and separations, with each tweezer at exactly the magic wavelength, we decide to produce our tweezer arrays with an SLM.

Although we have chosen to work with an SLM in this thesis, the other methods outlined here still offer highly valuable platforms for work with tweezer arrays. In fact, SLMs and AODs are readily combined to make the most of both of their advantages. AODs can be used with a SLM generated tweezer array to add a movable tweezer. This movable tweezer can be used to shuttle atoms between neighbouring tweezer sites from rearrangement protocols [36, 55, 56], or to add addressability to tweezer sites by changing the trap depth [48].

5.2.2 Basic principles of the SLM

An SLM is a device that applies a spatially varying modulation to a beam of light. Across the transverse profile of the laser beam it is possible to change one or both of the intensity and the phase. The SLM which we use can only modulate the phase of the wavefront of the reflected light.

The SLM is comprised of an array of pixels containing birefringent liquid crystals, where changing the orientation of the liquid crystals in one of the pixels changes the refractive index of that pixel. Light that passes through the pixel, reflects off the back surface, and re-emerges from the SLM window picks up a phase dependant upon the refractive index. By changing the crystal orientation of each pixel in the array, a phase mask is applied to the wavefront of the light. The phase imprinted on the wavefront will change the intensity distribution of light formed in the experimental chamber by the *in-vacuo* lens, allowing the movement and shaping of tweezers, as well as the generation of tweezer arrays.

The orientation of the liquid crystal pixels in our SLM is controlled via a com-

puter, and thus our SLM is referred to as being electronically addressed (EASLM). Another type of SLM is one that is optically addressed (OASLM), where it is controlled by a light field [40]. In this thesis we only consider the use of our phase modulated EASLM¹⁷.

5.2.3 Phase modulation

The orientation of the liquid crystal pixels imprints a phase mask on the 813 nm light beam as it reflects from the SLM. We now consider how this impacts the intensity distribution of the light when it has been focused by the *in-vacuo* lens by considering the effect the SLM has on a collimated apodised Gaussian beam incident upon its surface. We assume that the incident beam $A_{\text{in}}(x, z)$ ¹⁸ is apodised such that it is smaller than the aperture of the SLM window, so that there is no clipping of the beam¹⁹. The amplitude of the beam is taken to be real, and the phase assumed to be uniform across the wavefront.

The SLM imprints a phase $\theta(x, z)$ upon the reflected beam changing the profile to $A_{\text{SLM}}(x, z)$

$$A_{\text{SLM}}(x, z) = A_{\text{in}}(x, z)e^{i\theta(x, z)}, \quad (5.2.1)$$

where the beam changes from propagating in the $\hat{\mathbf{y}}$ direction to the $-\hat{\mathbf{y}}$ direction. The light then passes through a series of mirrors such that the beam is then travelling in the $\hat{\mathbf{y}}$ direction. It is focused by the *in-vacuo* lens with focal length f , where at the focus of the lens the amplitude of the light A_f is given by the Fourier transform of the beam,

$$\begin{aligned} A_f(\tilde{x}, \tilde{z}) &= \mathcal{F}[A_{\text{SLM}}(x, z)] \\ &= \mathcal{F}[A_{\text{in}}(x, z)e^{i\theta(x, z)}]. \end{aligned} \quad (5.2.2)$$

¹⁷Hamamatsu LCOS-SLM X10468-02. Where LCOS means liquid-crystal on silicon.

¹⁸Here we consider a beam propagating in the $\hat{\mathbf{y}}$ direction such that it is consistent with figure 5.1, though in reality it is propagating in the $\hat{\mathbf{x}}$ direction due to a folding mirror being excluded from the diagram.

¹⁹The beam is already apodised by the iris at the end of the beam expander to 12 mm, so if well aligned there should be no light hitting the SLM that does not reach the LCOS chip.

The intensity distribution of the tweezer light at the focus of the lens I_f is given by

$$I_f = |A_f(\tilde{x}, \tilde{z})|^2. \quad (5.2.3)$$

This intensity distribution I_f ²⁰ depends strongly upon the phase of the incident light, therefore by varying the orientation of the liquid crystals we have significant control of the intensity distribution of the tweezer light in the vacuum chamber.

The next step is to consider for a desired intensity distribution of tweezer light, what SLM mask would be required to generate it. Taking the inverse Fourier transform of A_f can be imagined as reversing the light path from the desired image, to look at which kind of beam is needed after the SLM to produce it. The problem one finds here is that to generate the exact beam needed after the SLM, one would require spatial control of both the phase and amplitude of the light field. As we only have control of the phase $\theta(x, z)$, instead we must therefore find a suitable phase mask that generates an intensity pattern which approximately matches the desired distribution. The task is then to find phase masks that produce approximations to the target distribution at very high quality. Optimising the suitable phase mask will require the use of an adapted version [36, 39, 40] of the Gerchberg-Saxton algorithm [170]. Details of the implementation of the phase masks to the SLM can be found in appendix C.2. The appendix considers: the calibration of the SLM, the different types of phase masks including how to combine them, and the adapted Gerchberg-Saxton algorithm for optimisation of the phase masks to produce desired arrays with equal trap depths.

5.3 High magnification imaging

We now have the means to produce an 813 nm tweezer (see section 5.1) and to generate tweezer arrays using a SLM (see section 5.2 and appendix C.2). Next we consider what kind of imaging system would be best suited to imaging such a tweezer

²⁰The (x, y, z) coordinate system refers to the SLM, the $(\tilde{x}, \tilde{y}, \tilde{z})$ coordinate system refers to in the focal plane of the *in-vacuo* lens.

array. All images in this chapter are taken using the blue MOT beams to cause the atoms to fluoresce.

Throughout the work described so far, we have imaged our atomic fluorescence using the same imaging system described in chapter 4, which is a telescope comprised of the *in-vacuo* lens (see section 3.3.1) and a 400 mm focal length achromat. Together these lenses achieve a magnification of $\times 10.8$. Due to the large size of each pixel of the SPAD array (150 μm), this still corresponds to quite a large size ($\sim 14 \mu\text{m}$) in the focal plane of the *in-vacuo* lens where the atoms are imaged. It is therefore desirable when working with arrays to move to a higher magnification system where we can resolve atoms in tweezers at closer separations.

5.3.1 Motivation of high magnification

Before moving to a higher magnification imaging system, let us first consider what we would be able to see in the original configuration. Applying a two tweezer phase mask to the SLM we can trivially see a pair of tweezers loaded from a NB MOT, the first image we ever took of this is shown in figure 5.2. If we were to compare this to the spatial extent on the NB MOT (see figure 3.6)²¹, one can see that there is a very limited number of pixels in the cloud of the NB MOT for us to be able to image distinct tweezers from. Therefore it becomes necessary for the NB MOT to be on more pixels of the SPAD array, so we must have a higher magnification to work with tweezer arrays of many sites. We also want to remain within the high SLM diffraction efficiency $230 \mu\text{m} \times 230 \mu\text{m}$ region of the focal plane described in appendix C.2.1, so magnifying such that the SPAD array images that region with the best resolution would be beneficial.

An even more pressing reason to increase the magnification of the imaging system is the separation with which we can resolve tweezers. We aim to map each pixel to a region of just a few microns in the focal plane to allow site-resolved imaging of closely packed tweezers, where separations can then be within the blockaded or dressed radii

²¹The two figures 3.6 and 5.2 have the NB MOT on different regions of the SPAD array at the time, but there is no physical significance to this.

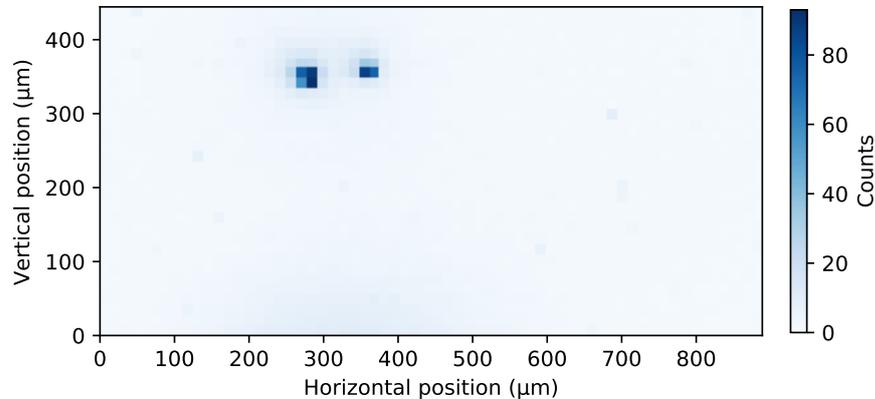


Figure 5.2: Our first image of a pair of tweezers, imaged with the blue MOT beams in the low ($\times 10.8$) magnification system. Here the tweezers are loaded at high power to maximise signal, they are not loaded with the proper protocol of starting in a shallow tweezer then adiabatically ramping the trap depth.

of Rydberg interactions which were discussed in section 2.4. To image the pixels on the SPAD array, we would like a centre to centre separation of at least two pixels, as tweezers being located on neighbouring pixels could complicate resolving where the light has come from. A separation of two pixels corresponds to $300 \mu\text{m}$ on the surface of the SPAD array. Therefore, if we want to have two tweezers separated by $\sim 6 \mu\text{m}$ and be able to clearly image them, then the magnification would have to be increased to $\sim \times 50$.

With a suitable magnification imaging system constructed, there is an ideal 11×11 pixel region on the SPAD array in which we could image a 6×6 tweezer array with two pixel separation, as already detailed in section 3.4.1 and shown in figure 3.20. For the moment we focus on the imaging of a 2D tweezer array, though techniques to extend imaging to 3D exist, such as through the use of electrically-tunable lenses [35, 36].

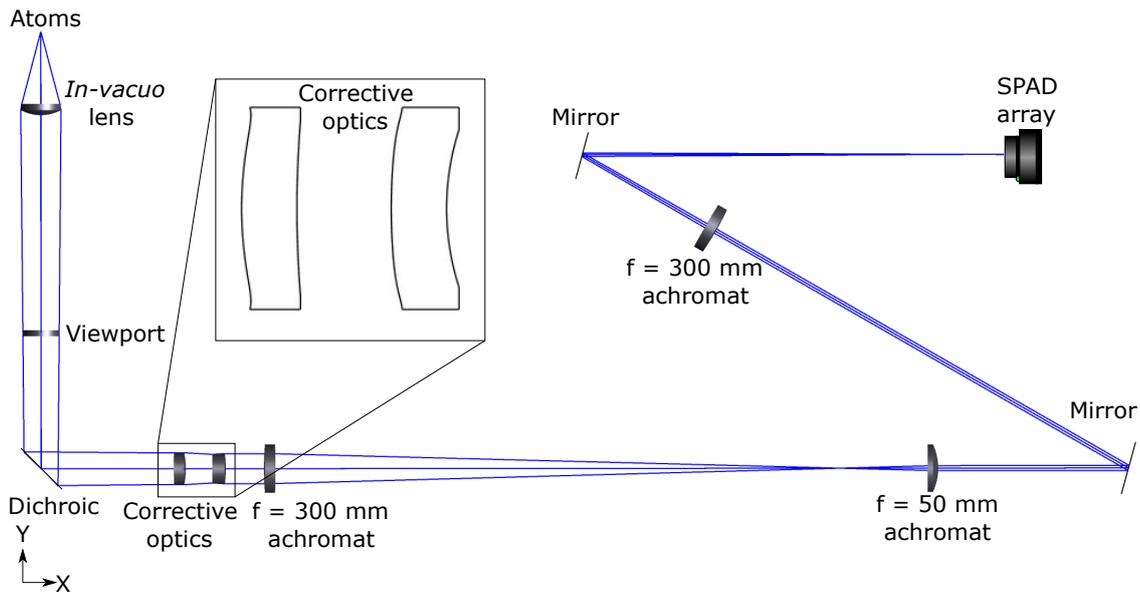


Figure 5.3: The high magnification ($\times 48.6$) double telescope imaging system. Corrective optics are shown in an expanded view with their cross sections, and are used to remove aberration caused by the *in-vacuo* lens. A pick off mirror can be added between the first $f = 300$ mm achromat and the $f = 50$ mm achromat to view the intermediate image on a CCD (not shown in diagram) to provide a lower magnification image.

5.3.2 Double telescope

To achieve a magnification of $\sim \times 50$ using a single telescope including our long working distance *in-vacuo* lens, we would need to use a lens of focal length approaching 2 m, which is an unreasonable size to fit onto the optical table. Instead we use a double telescope. The first telescope is comprised of the *in-vacuo* lens and a 300 mm achromat²², providing a magnification of $\times 8.1$. The second telescope is a 50 mm aspheric lens²³ and again another 300 mm achromat²⁴ and provides an additional magnification of $\times 6$. The double telescope together provides a total magnification of $\times 48.6$. The full imaging system is shown in figure 5.3.

²²Thorlabs AC254-300-A

²³Edmund optics 33-945

²⁴Thorlabs AC254-300-A

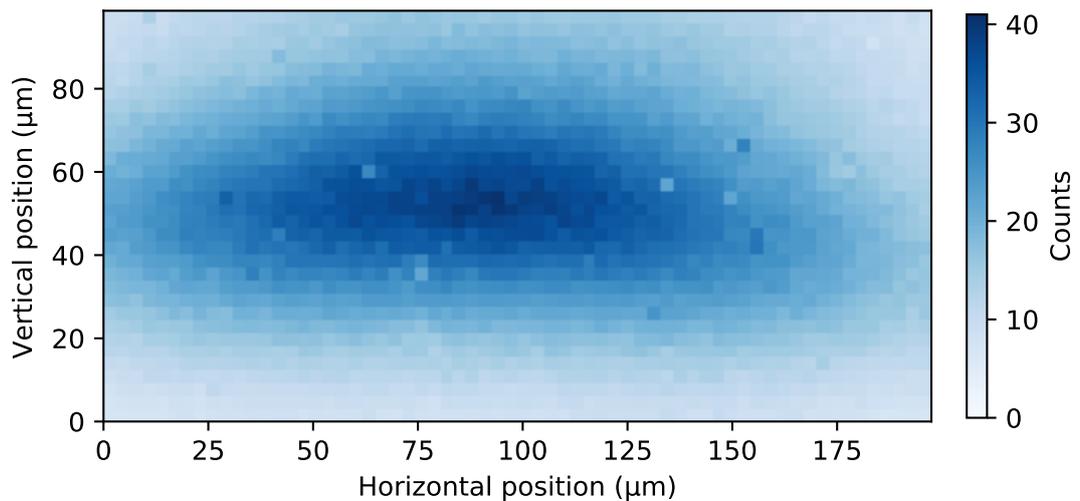


Figure 5.4: A typical averaged fluorescence image of 100 narrowband red MOTs, taken at $\times 48.6$ magnification onto the SPAD array. Each pixel corresponds to approximately $3.1 \mu\text{m}$ at the position of the atoms. The colourmap is given in a false colour of blue to represent the 461 nm light used to image the MOT. Compared to figure 3.6, the image is inverted due to the extra telescope in the imaging system, so when released, the MOT appears to ‘fall’ upwards. The pixels with lower counts within the cloud will need checking to see if they are operating at lower quantum efficiency, or whether there is an obstruction in the imaging system. Several of the pixels with lower counts were also visible when looking at the broadband red MOT in the $\times 10.8$ magnification imaging system (see figure 3.5).

This high magnification imaging system now means that the effective pixel size of our SPAD in the focal plane of the *in-vacuo* lens is now $3.1\ \mu\text{m}$, and means that the full 64×32 pixel sensor area covers roughly a $200 \times 100\ \mu\text{m}$ region of the focal plane. This area is approximately the size of our NB MOT, as shown in figure 5.4, which images the NB MOT under the same $52\ \mu\text{s}$ pulse of the blue MOT beam light used in figure 3.6. The imaged region does not quite show all of the NB MOT, and is certainly not large enough to look at a full BB MOT. Being able to image the full NB MOT is highly useful, not least for being able to easily take ballistic expansion data to measure its temperature. To circumvent this, we can drop in a mirror at 45° into the cage between the two telescopes such that the intermediate image is instead sent to a CCD²⁵.

5.3.3 Aberration and corrective optics

There is an unfortunate design flaw in our *in-vacuo* lenses, they have high order aberrations which cause issues when collecting fluorescence from a point source that near uniformly fills the surface of the lens. High order aspheric terms become significant at the edge of the lens, and the near uniform nature of the fluorescence hitting the lens means that the outer area of the lens has a significant effect on the formation of the image. This is contrary to trapping light passing through the lens, where we send an apodised Gaussian beam, so there is greater intensity towards the centre of the lens. The *in-vacuo* lens is very well optimised for trapping, but inadvertently pays the price for this in imaging.

We saw the effects of this aberration during the work in chapter 4, where despite the effective $14\ \mu\text{m}$ pixel size, only $\sim 90\%$ of the light from the tweezer was collected on a single pixel, even after 3D optimisation of the SPAD position. This caused minimal issues when working with the single $532\ \text{nm}$ tweezer, as the large effective size of the SPAD pixel allowed us to collect most of the light on one pixel. The move to higher magnification will only exacerbate the single pixel collection problem, as the effective pixel size is much smaller.

²⁵PCO Pixelfly QE.

The preferable response upon realising this issue would be to replace the *in-vacuo* lens with a superior design with less aberration. However, due to the circumstances arising from working in a pandemic, a vacuum break was not a desirable option. We instead chose to design two air-side corrective optics that reduce the aberration in the imaging system, without the risk of a vacuum break. The design of the corrective optics was carried out by Frédéric Leroux, based on a combination of use of Zemax OpticStudio software and advanced geometric optics methods for aspheric surfaces. The solution of using air-side corrective optics is not perfect, and in the next iteration of our experiment the *in-vacuo* lenses ought to be upgraded. Given the circumstances however, we press on with an air-side approach and demonstrate how this can lead to a good imaging performance.

Appendix C.3 contains a detailed description of the issues with the aberration in the *in-vacuo* lens and the design of the pair of corrective optics. The theoretical performance of the imaging system is also considered, including reductions in performance arising from the tolerance of the highly precise corrective optic system construction, and the spread of atoms within the tweezers.

Measured performance

Figure 5.5 shows an image of a single tweezer, loaded from the NB MOT, after optimisation of the positions of the two corrective optics (described in appendix C.3). We see that 31% of the detected fluorescence is collected on the single target pixel. 68% of the light is found in the 3×3 pixel region around the target pixel, the rest of the light is more diffusely spread. Losing over two thirds of the fluorescence to being on the wrong pixel is a significant loss, though this can hopefully be improved with further optimisation of the imaging system. If additional fluorescence is required, the tweezer image could be centred on a 2×2 pixel region while still having a tweezer to tweezer separation of two SPAD pixels ($6.2 \mu\text{m}$ in the vacuum chamber), though this would be at the cost of quadrupling the background rate.

Persisting with the planned approach of only using one SPAD pixel per tweezer, we must characterise the cross talk between tweezer sites for a square array spaced by two SPAD pixels. Seven of the eight nearest two pixel spaced sites on the SPAD

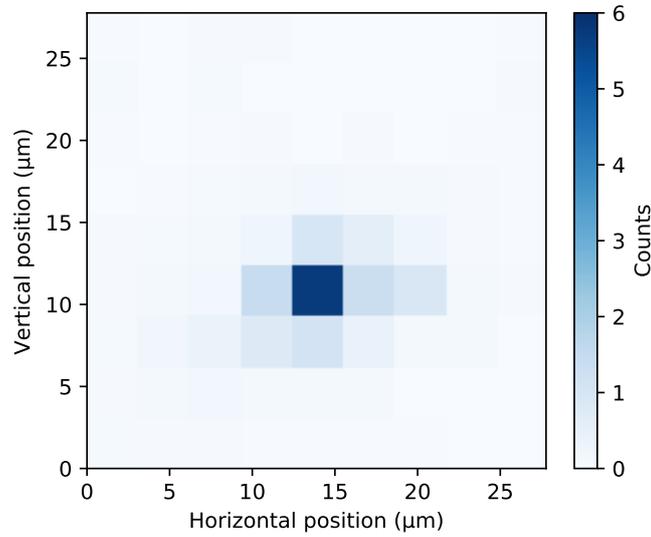


Figure 5.5: Imaging a single tweezer in the high ($\times 48.6$) magnification system. Here the tweezers are properly loaded from shallow tweezers before adiabatically ramping the trap depth. 31% of the fluorescence is collected the single target pixel.

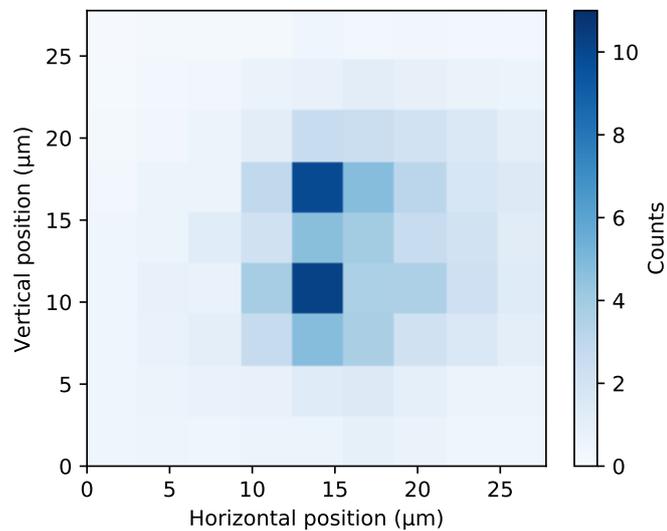


Figure 5.6: Imaging a pair of tweezers in the high ($\times 48.6$) magnification system. Here the tweezers are properly loaded from shallow tweezers before adiabatically ramping the trap depth. The tweezers are centred two pixels apart in the vertical direction, a separation of $6.2\ \mu\text{m}$. The two tweezers are clearly resolvable from one another, though there is spilling of the fluorescence over multiple pixels.

array have detected fluorescence of ≤ 0.17 counts in the data used in figure 5.5. This would mean that the light from an adjacent tweezer would add to the fluorescence by $\leq 3\%$ each, which ought to be sufficiently small cross-talk to clearly resolve whether a tweezer is occupied or not when we push towards the single atom regime. There is however one of the adjacent sites that has a high cross-talk. Two pixels to the right of the tweezer, pixel $(X + 2, Y)$ has a notable amount of fluorescence on it, 0.9 counts compared to the 5.7 counts on the brightest pixel. This is a cross talk of 16%, which is highly significant. When looking for single atoms this cross-talk could be a challenge, so it would be best to reduce it as much as possible. This suggests we need to further tweak the imaging alignment to reduce the asymmetry.

Residual cross-talk between pixels could potentially be managed through post-analysis of images, taking into account the estimated occupation of neighbouring tweezers to help identify atom numbers when the fluorescence is close to the identification threshold. We do note however that if the horizontal spacing was instead three pixels instead of two, and we have a rectangular $9.3\mu\text{m} \times 6.2\mu\text{m}$ array, then the cross-talk in all directions would be less than $\leq 3\%$ from each of the adjacent sites.

Figure 5.6 shows that we can resolve tweezers easily with two pixel separation in the vertical direction, showing that the corrective optics have worked sufficiently well to resolve on the desired length scale. There is potential to further optimise the placement of the corrective optics (see appendix C.3), the absolute tweezer alignment to produce the smallest tweezer waist, and in ensuring that we are imaging through the best part of the *in-vacuo* lens. We are however now at the point where we have a tweezer array with site-resolved imaging.

5.4 Loaded tweezer arrays

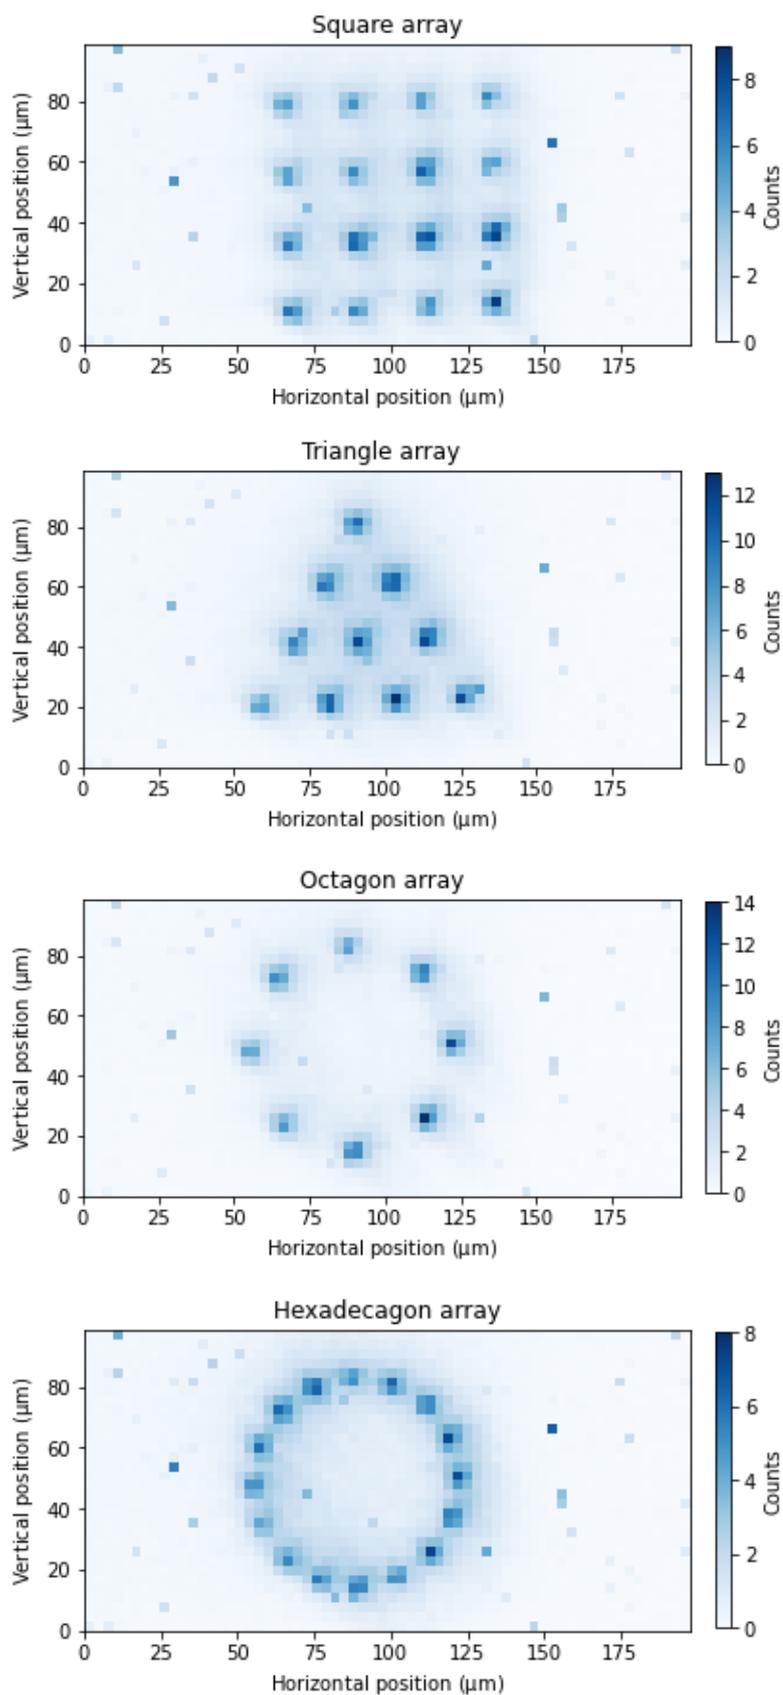
With the implementation of the magic wavelength tweezer, the spatial light modulator, and a high magnification imaging system, we can now create and resolve tweezer arrays. In this section we demonstrate what kind of 2D tweezer arrays we can create, including images of the fluorescence from atoms in these arrays.

5.4.1 Loading from a BB MOT

We find that for a 2D tweezer array (all of the tweezer focii in a single plane) the shape of the NB MOT as an ellipsoidal shell means that loading between vertically separated tweezer sites can vary quite strongly when the tweezer array spans many tens of microns. There are potential strategies to improve this loading from the NB MOT to widen the variety of arrays that we can load. In the first instance we only wish to demonstrate the tweezer arrays themselves, and for that we can easily load a larger tweezer array by loading directly from the broadband (BB) MOT.

The BB MOT has a much larger volume than the NB MOT, and fills the ellipsoidal shape rather than just a thin shell. This lends itself to the atomic distribution covering all of the tweezer sites even for large arrays. Here we can forego our typical loading of a shallow tweezer before adiabatically ramping the trap depth detailed in section 3.1.4, instead just directly loading into a deeper tweezer of constant depth. Due to the atomic temperature in a BB MOT being a factor of ~ 20 higher than the NB MOT (see section 3.1.3), we can no longer load into very shallow tweezers. Here we load to trap depths of $\sim 200\mu\text{K}$, where occupation of any secondary maxima in the tweezer potential should not be an issue²⁶.

That we do not ramp the trap depth to deeper tweezers means that during imaging there will be a greater increase in the atomic spread when imaging, due to heating from recoil of the scattered 461 nm photons. Images will also look worse as we no longer seek to overlap the tweezers with individual pixels here, as we are making arbitrary shapes that do not necessarily fit into a square array. In addition, the images shown in the following section were taken before the optimisation of the imaging system to the performance shown in figure 5.5 was completed, so are not representative of the best we ought to be able to achieve.



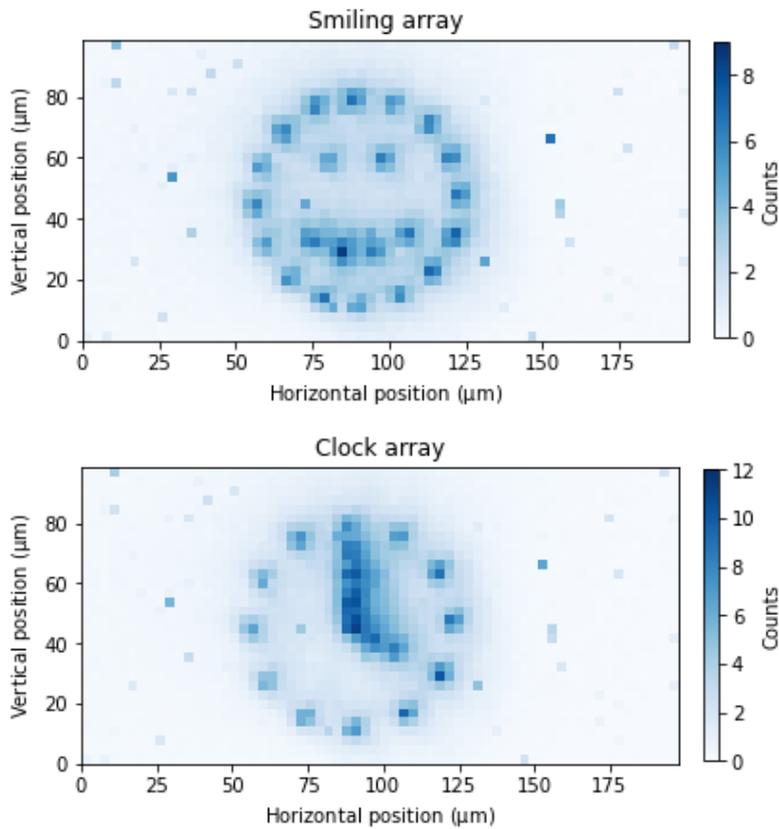


Figure 5.7: Gallery of tweezer array configurations. Images are fluorescence collected from atoms loaded into each tweezer array. The atoms are loaded into the tweezer array directly from the BB MOT, and the images shown here are the array produced when we do not use the trap intensity equalisation scheme. A slight rotation of the pattern compared to the camera axes is seen here, which we attribute to the steering mirrors between the SLM and the *in-vacuo* lens causing a rotation of the pattern. The associated phase masks corresponding to the tweezer arrays are shown in figure C.6.

5.4.2 Gallery of tweezer arrays

A variety of example arrays loaded from BB MOTs are shown in figure 5.7, where each site contains many atoms. We show both square and triangular arrays, demonstrating the flexibility of changing between geometries when working with a SLM to produce tweezer arrays. We also show points on a circle in the form of an octagon and hexadecagon, as well as arbitrary patterns inside a hexadecagon and a dodecagon to make a smile and an analogue clock at 4 o'clock. Each of the tweezers are clearly resolvable from one another, with the potential exception of the clock where we were slightly overzealous on the tweezer density to form the hands. Each array is $< 100 \mu\text{m}$ across²⁷, and would fit on the cross-section of a human hair.

One issue that is common to all of the arrays in figure 5.7 is that they appear to be rotated slightly compared to the axes of the SPAD. This is something we expected from what we found earlier in the construction of the array, where use of the steering mirrors to align the trapping light onto the *in-vacuo* lens appears to cause a rotation of the tweezer pattern. We see this in the trapping light both before and after the chamber, so it is not induced by our *in-vacuo* lenses, and is a real rotation of the atoms in the chamber from the lab axes. We can however simply correct for this by characterising the rotation, and modifying all target array patterns to include a fixed rotation to offset this.

5.4.3 Loading schemes

Directly loading from a BB MOT is an effective way to fill large arrays of tweezers, particularly for arrays of large size in the vertical (\hat{z}) direction. This however comes with the drawback of atoms being loaded at higher temperatures into the tweezers, and the need to work with deeper traps. On the other hand, the spatial extent of the NB MOT limits the size of array which we can load but offers much lower temperatures. We consider three different ways to achieve the loading of large arrays

²⁶Atoms at $\sim 20 \mu\text{K}$ should be too hot to be trapped in the secondary maxima of the tweezer potential.

²⁷Each pixel of figure 5.7 corresponds to $3.1 \mu\text{m}$ in the plane of the atoms.

for scaling our system in the future.

One approach would be to continue with our approach of loading from the BB MOT. Initial investigations suggest that this has been an effective way to load the tweezers, the drawback being the atomic temperature is higher than we would like. Once the atoms are loaded into the tweezer, Sisyphus cooling could be employed on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1\ |m_j| = 1$ transition, where cooling would be on the attractive cooling mechanism for an 813 nm tweezer [131] as shown in figure 4.3(b).

The two other approaches instead consider how can we load from a NB MOT into a shallow tweezer array in our normal way, but over a larger region. The solution here is to physically translate the NB MOT across the target region to sweep across a larger volume, loading tweezers as it overlaps with them. The NB MOT can either be translated by changing its detuning, or changing the current through the \hat{z} compensation coil. For sensible values and a slow scan, both approaches should load across a larger region than the standard NB MOT. Shifting the compensation coil would seem the preferable approach, as the dimensions and temperature of the NB MOT should not be changed, whereas over large changes of detuning these properties would begin to vary. For very large arrays beyond the size of the field of view of our SPAD array, the magnetic field approach would also allow the NB MOT to be scanned in three dimensions using each of the three \hat{x} , \hat{y} and \hat{z} compensation coils.

5.5 Future imaging protocol

The trap depths used in a magic wavelength tweezer array will need to be a lot lower than the 7.5 mK we used in the single 532 nm tweezer, as we discussed in section 5.1.1. For a tweezer of our target (actual) waist of 1.67 (2.25) μm , we would require a power of 70 (127) mW at the focal plane per 1 mK tweezer of depth. For an example trap depth of 400 μK with 36 tweezers, this would require a power of 1.01 (1.83) W to reach the centre of the vacuum chamber in the diffracted tweezer array out of the original 2.7 W from our Ti:Sa.

Improving the waist further from 2.25 μm would be beneficial for increasing the

number of possible tweezer sites, and in tighter confinement for both controlling atomic separation, imaging atoms in a smaller spatial extent, and tighter confinement in terms of the Lamb-Dicke regime for clock measurements.

Continuing with the example of 400 μK deep tweezers, we will then need to image the trapped atoms. As in section 4.1.1, we plan on returning to using a 461 nm probe beam to image the atoms.²⁸ This gives us better control of the imaging light parameters, and means that the dipole radiation pattern is emitting maximally into the *in-vacuo* lens. Given there are now multiple tweezer sites to image, it may be desirable to use a probe beam with a larger waist so that there is a minimal spatial dependence on the fluorescence across the array. The differential AC Stark shift on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ imaging transition is much smaller due to the reduced trap depths, and shifts in the opposite direction due to the differing sign of the polarisability [72]. Here we only expect differential AC Stark shifts of -20 and -15 MHz on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ $|m_j| = 0$ and 1 states respectively, which will not be resolvable within the linewidth of the transition, and can easily be tuned around with the beat lock and AOM frequencies in the 461 nm laser system (see section 3.2.1).

In section 4.5 we demonstrated three imaging schemes for atom(s) in deep 532 nm tweezers, two with Sisyphus cooling and one without. Here in much shallower 813 nm tweezers we will almost certainly need to employ cooling techniques, due to the reduced trap depth. A single 461 nm photon recoil can heat by $\frac{E_r}{k_B} = 0.514\ \mu\text{K}$, meaning that in a 400 μK deep trap we can scatter ~ 800 photons before there is a risk of heating the atom out of the tweezer. Given we only expect to detect $\sim 0.5\%$ of the emitted photons [72], and of that only 31% of the light is on the target pixel, we are looking at only one or two counts for a single atom on a single pixel before heating out of the trap, so cooling is necessary going forwards.

We plan to once again employ Sisyphus cooling²⁹ during imaging, though this time it is also to prevent heating out of the tweezer. As discussed in section 4.2.2,

²⁸The fluorescence images in this chapter all used the blue MOT beams to image the atoms.

²⁹For atoms in 532 nm tweezers (see section 4.2) we cooled with the repulsive Sisyphus mechanism as shown in figure 4.3(a).

at 813 nm the differential polarisabilities on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition are not the same as at 532 nm, where this time the $|m_j| = 0$ and 1 transitions have the differential AC Stark shift in opposite directions [91]. The $|m_j| = 0$ can still operate on the repulsive Sisyphus mechanism as in chapter 4, whereas the differential AC Stark shift on the $|m_j| = 1$ transition has the opposite sign and instead operates on the attractive Sisyphus mechanism [131] which is shown in figure 4.3(b). The latter is a more desirable cooling mechanism, being attracted towards a set temperature cap rather than be repelled away from one. Using the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ $|m_j| = 1$ transition for Sisyphus cooling is likely the best approach to take to avoid heating out of the trap so that we can collect enough photons for imaging.

An additional benefit of working in an 813 nm tweezer is that the $5s4d\ ^1D_2$ state is trapped, and can decay into the $5s5p\ ^3P_{0,1,2}$ manifold. We can then apply light on the 679 nm and 707 nm repump transitions to transfer population to the $5s5p\ ^3P_1$ state, which will quickly decay back to the ground state and into the imaging cycle again. This should remove the $5s5p\ ^1P_1 \rightarrow 5s4d\ ^1D_2$ loss mechanism that caused issues in chapter 4, and allow not just sufficient imaging time to view single atoms, but to be able to repeatedly image them too [131]. However, in the future when doing a normalised electron shelving detection scheme on the clock transition [94, 131], the repumping will not be a viable approach to imaging, as we do not want to change the population in the clock ($5s5p\ ^3P_0$) state while imaging the population in the ground ($5s^2\ ^1S_0$) state. We can instead blow atoms out of the trap with a strong 461 nm probe beam, or image while Sisyphus cooling. As we will want to collect as much fluorescence as possible, improving the collection efficiency through a higher NA lens and a less aberrant imaging system would be beneficial.

To maximise the utility of the SPAD array, combined with the necessity of Sisyphus cooling, we ought to implement an interleaved imaging and cooling approach similar to that of section 4.5.3. We can use a series of strong imaging pulses at a low duty cycle, separated by large blocks of Sisyphus cooling. This will make the most of the low dark count rate of the detector, while also not heating atoms out of the trap. Heating should be easier to control with the attractive Sisyphus cooling mechanism, where we should not need to worry about heating above a repulsive

Sisyphus cap.

Given the variation in the density of the NB MOT over the spatial extent of the tweezer array, particularly in the \hat{z} direction, stochastic loading of the atom number as done in chapter 4 will not produce similar mean atom numbers across the array. Furthermore, now that we have a tweezer array, it is more desirable to trap a single atom per tweezer to look at interactions between atoms in neighbouring tweezers, rather than having multiple atoms loaded in the same tweezer. To load single atoms we will use light-assisted collisions to parity project each tweezer to contain a single atom or be empty [45]. This can be achieved by applying 689 nm light to induce photoassociation to a molecular resonance [188, 189], inducing pairwise loss of atoms from the tweezer [45].

5.6 Summary

In this chapter we have implemented a tweezer array at 813 nm, the magic wavelength of the clock transition. The tweezer array is produced via a spatial light modulator, allowing the generation of arbitrary 3D arrays of tweezers. We detailed the construction of the tweezer array system, as well as the methods with which optimised tweezer arrays can be produced.

A new high magnification imaging system was constructed to allow a 2D tweezer array to have site-resolved low background imaging for tweezer arrays with a spacing of 6.2 μm . The imaging system included two corrective elements to reduce the effects of an aberrant *in-vacuo* lens, without the need for a vacuum break. The imaging system was then used to image various tweezer arrays in figure 5.7, where we can resolve atoms in tweezer sites shaped into different array geometry and shapes.

There is a clear path to high fidelity single atom imaging in these tweezer arrays, through use of attractive Sisyphus cooling, and repumping within the $5s5p\ ^3P_{0,1,2}$ manifold. From here we then have a platform of regularly spaced resolvable trapping sites for implementing a tweezer array clock.

Chapter 6

Rydberg electrometry

The use of Rydberg atoms and Rydberg-dressed atoms is a powerful toolbox for a variety of applications. In this chapter, we focus on the use of Rydberg atoms from tweezers as precision sensors of the electric field. Rydberg atoms are great sensors of electric fields, as they have very large electric field sensitivities scaling as $\alpha \propto n^7$ (see section 2.4.3). The Rydberg atoms localised using optical tweezers measure the electric field at the position of the optical tweezers inside the vacuum chamber, giving a precise measurement of the electric field at the location of the atoms over a small region.

This chapter features the measurement and cancellation of the stray electric field in our vacuum chamber using spectroscopy of Rydberg states and our twelve *in-vacuo* electrodes. The characterisation and cancellation of electric fields is important for both controlling systematic shifts in atomic clocks [163], as well as for producing an environment compatible with the precision measurement of Rydberg states [162, 164].

For clock measurements, the DC Stark shift from the electric field is one of the terms in the error budget which must be characterised and preferably also minimised. The DC Stark shift of the clock state causes a fractional shift of $7.17 \times 10^{-17} \left(\frac{|E|}{\text{Vcm}^{-1}} \right)^2$ [190]. To provide context, this has a shift of 1 Hz in the clock transition for an electric field of 5.70 Vcm^{-1} [134], a shift of a second in the age of the universe (2.3×10^{-18}) for an electric field of 180 mVcm^{-1} , and a shift of the best measured fractional

frequency precision in an atomic clock (7.6×10^{-21}) [8] for an electric field of 10.3 mVcm^{-1} .

As for Rydberg measurements, the cancellation of electric fields is critical for measuring the true frequencies of the Rydberg series. Precision measurement of Rydberg states [65, 66] is important for refining models of Rydberg state energy levels, and even for testing predictions of physics beyond the standard model [69, 70, 191]. In particular, precision measurements within the triplet Rydberg series are important for testing and improving multi-channel quantum defect models [67, 68].

In the work undertaken in this thesis we excite to Rydberg states using the scheme $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1 \rightarrow 5snl\ ^3L_J$ on the 689 nm and 319 nm (UV) transitions (see figure 2.1). We begin by discussing how we choose a Rydberg state, and perform spectroscopy on the feature. The 3D electric field control is introduced, before varying the electric field and looking at how the energy of the Rydberg state changes. We use this data to generate Stark maps in 3D, measuring and cancelling the stray electric field in the vacuum chamber. We then consider the residual electric field remaining in the vacuum chamber after cancellation, and make comparisons between our data and predictions made by ARC 3.0 [1].

6.1 Measuring the DC Stark shift

To characterise the electric field, we need to take spectra of the Rydberg states for different applied electric fields, where the states will be shifted through the DC Stark shift for the total electric field in the vacuum chamber. By applying a series of different electric fields and measuring these shifts, we can determine and cancel the stray electric field in the vacuum chamber.

6.1.1 Choosing a Rydberg state

For the measurement of the electric field at the atoms, we choose to use the Rydberg state $5s79d\ ^3D_1$. Details of finding and identifying the appropriate Rydberg states can be found in appendix D.1. The choice of a 3D_1 state gives us access to two different $|m_j|$ levels, that will each experience a different DC Stark shift, providing

two sets of data from one experiment. The choice of principal quantum number n is coarsely set by wanting to have shifts of tens of MHz in the Rydberg state on the scale of the electric field range achievable with our precision electrode control (see section 3.3.3). The exact value of n is unimportant, so was chosen to match a previously measured state [155] for convenience, and to be similar to prior work in the singlet Rydberg series [76, 141].

The predictions of the expected shifts were made using ARC 3.0¹, an openly available python library which includes the ability to simulate the Rydberg states of divalent atoms [1]. The state we chose is also interesting in that ARC predicts a non-parabolic response from the $|m_j| = 1$ level, unlike the parabolic shape expected from equation 2.4.3. As we will see later, this arises from mixing of the states², where ARC can calculate which state will be mixed to cause the non-parabolic shape.

6.1.2 Experimental sequence

To measure the energies of the Rydberg states, we look for depletion in the atom number than can be imaged in a cloud. Here we use atoms that were loaded from a NB MOT into a 532 nm tweezer, that are then released before Rydberg excitation. As we do not work directly in the BB MOT like in appendix D.1, we do not have the broad 689 nm light that is on during the BB MOT phase used in the excitation path, and the atoms can be colder than in the BB MOT stage. By holding atoms in the tweezer for a short while, we can allow the quadrupole magnetic field to decay, or could even flip the coils to Helmholtz configuration to produce a uniform magnetic field. In this case however, we simply ensure that the quadrupole field has decayed.

Atoms are then released from the tweezer³ to remove the differential AC Stark shift on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ and $5s5p\ ^3P_1 \rightarrow 5s79d\ ^3D_1$ transitions which would have shifted the required frequencies of the 689 nm and UV lasers. Excitation to the Rydberg state is done by applying both the 689 nm and UV lasers together for

¹Alkali(ne) Rydberg Calculator.

²Such that the polarisability becomes a function of the applied electric field.

³This chapter uses many atoms trapped in tweezers, our measurements do not yet use single atoms.

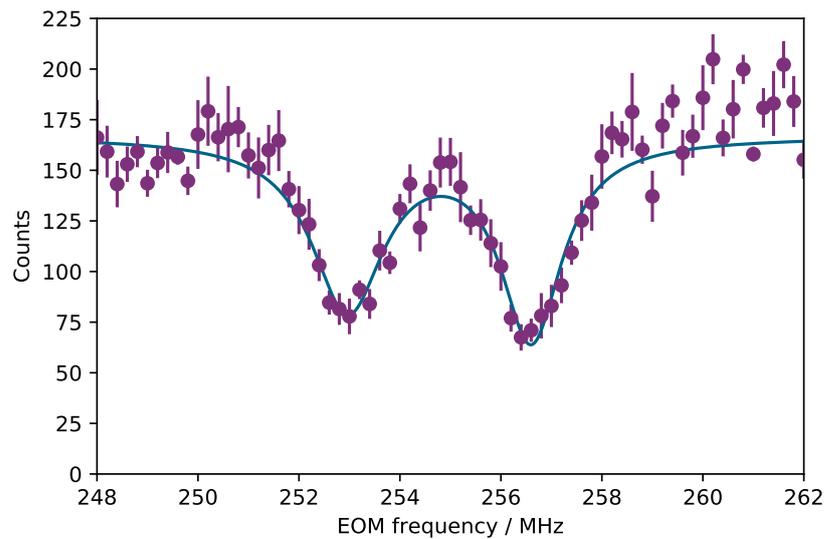


Figure 6.1: Depletion in fluorescence from atoms released from an optical tweezer, as a function of the 638 nm laser frequency. The left and right minima correspond to the $|m_j| = 0$ and 1 levels respectively. This data was taken for no applied electric field (all electrodes grounded), so shows the splitting due to the stray electric field in the chamber at the position of the atoms. The two minima are split by 7.26 ± 0.18 MHz in the UV, which is double the splitting seen in terms of the 638 nm laser, due to the SHG conversion from 638 nm to UV. The linewidth here is significantly more than the natural linewidth, likely arising from power broadening, where in future less intense UV light should be used to improve resolution of the line centres.

1 ms. Atoms which are excited to the Rydberg state can no longer be imaged on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^1P_1$ transition, and are seen as depletion of the ground state atom number when 461 nm light is applied. By scanning the UV frequency we generate a graph of depletion against laser frequency, and can fit to the minima in 461 nm fluorescence detected (see figure 6.1). From here we find the UV frequencies at which depletion occurs, which gives us the splitting of the $|m_j| = 0$ and 1 levels. Next we need to apply a variety of electric fields and study the shift of the two features, as a function of the applied electric field in all three dimensions.

6.2 Electric field control

To control the electric fields, we have twelve *in-vacuo* electrodes around the atoms in a hexagonal split-ring geometry (see section 3.3.3). This allows us to produce electric fields along any axis required [72], and cancel a three dimensional stray electric field. We have high-precision independent control of each electrode voltage to apply the cancellation field, generated by our high precision low noise voltage source⁴, which was detailed in section 3.3.3. Each electrode voltage can be controlled individually via LabVIEW software.

The electrodes are mounted at 15° to our lab axes, such that a set of four of the electrodes are on the axis of the MCP, which sits between two viewports. To cancel the electric field in three dimensions, we can sequentially minimise the DC Stark shifts along three orthogonal axes. To that end we define the coordinates $\hat{\mathbf{x}}''$ and $\hat{\mathbf{y}}''$ coordinates⁵ in the horizontal plane that are not on our lab axes, but are orthogonal directions using simple electrode configurations. To cancel the stray electric field (\vec{E}_{stray}) at the atoms, we need to apply a field (\vec{E}_{applied}) with the same magnitude

⁴Stahl Electronics, BS 1-12-0.3.

⁵The double primed coordinates here denote electric fields applied in the xy plane are not on the lab axes, but in fact are on the $\hat{\mathbf{x}}'' = (\cos(15^\circ)\hat{\mathbf{x}} + \sin(15^\circ)\hat{\mathbf{y}})$ and $\hat{\mathbf{y}}'' = (\cos(15^\circ)\hat{\mathbf{y}} - \sin(15^\circ)\hat{\mathbf{x}})$ directions. This is done to keep coordinate systems consistent throughout this thesis.

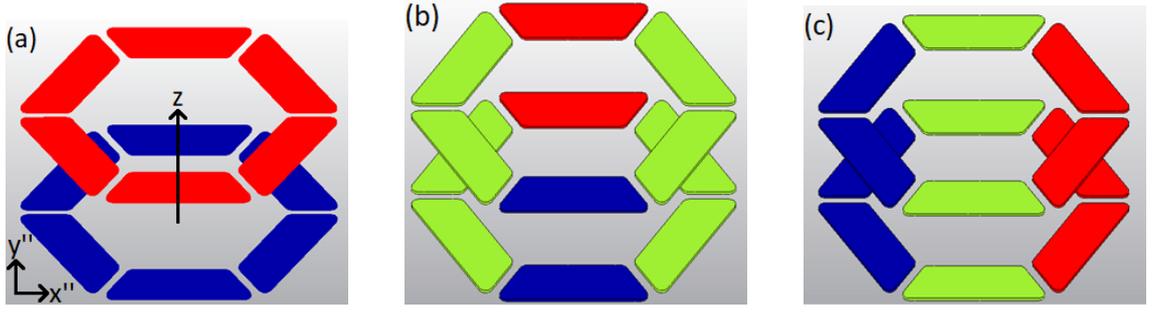


Figure 6.2: The three orthogonal electrode configurations used to cancel the stray electric field. The red/green/blue colour scheme indicates that the electrodes are at $+V_{\text{set}}$, 0 V and $-V_{\text{set}}$ respectively. The two hexagonal rings of electrodes are split in the \hat{z} direction, with the planes of the electrode rings containing the \hat{x}'' and \hat{y}'' directions.

- (a) The voltage configuration to apply electric field in the $-\hat{z}$ direction (cancelling a stray field in the $+\hat{z}$ direction), with a calculated electric field of $0.380 \times V_{\text{set}} \text{cm}^{-1}$.
- (b) The voltage configuration to apply electric field in the $-\hat{y}''$ direction (cancelling a stray field in the $+\hat{y}''$ direction), with a calculated electric field of $0.326 \times V_{\text{set}} \text{cm}^{-1}$.
- (c) The voltage configuration to apply electric field in the $-\hat{x}''$ direction (cancelling a stray field in the $+\hat{x}''$ direction), with a calculated electric field of $0.565 \times V_{\text{set}} \text{cm}^{-1}$.

but opposite direction.

$$\begin{aligned} \vec{E}_{\text{stray}} &= -\vec{E}_{\text{applied}} \\ &= E_{x''} \hat{x}'' + E_{y''} \hat{y}'' + E_z \hat{z}. \end{aligned} \quad (6.2.1)$$

The voltages applied to the electrodes must be converted to the electric field generated using a calibration. The electrode configurations were simulated with voltages applied in three different geometries to calculate the expected electric field at the position of the atoms as a function of applied voltage [72]. These simulations used Autodesk Inventor, simulating just the electrodes, rather than the full CAD model of the vacuum chamber [72]. The orthogonal electrode configurations are shown in figures 6.2(a,b,c), and correspond to the application of an applied electric field in the $-\hat{z}$, $-\hat{y}''$ and $-\hat{x}''$ directions respectively. For applied voltage V_{set} , the configurations should produce fields of $0.380 \times V_{\text{set}} \text{cm}^{-1}$, $0.326 \times V_{\text{set}} \text{cm}^{-1}$ and $0.565 \times V_{\text{set}} \text{cm}^{-1}$ on the respective \hat{z} , \hat{y}'' and \hat{x}'' axes. We expect there to be a

discrepancy in the actual generated field due to inaccuracies in the simulation, or movement of the electrodes during baking out of the vacuum chamber. We will assume that this electrode calibration is accurate for section 6.3.1 where we seek to minimise the shift due to the electric field, before calibrating our electrodes against fitting with ARC [1] in section 6.3.2.

6.3 Stark maps

By varying the applied electric field, then performing the experimental sequence described in section 6.1.2, we scan to see how the energy of the Rydberg states change as a function of the applied electric field. We then fit to find the relative frequencies of the $|m_j| = 0$ and 1 states for each electric field. Plotting the shift of each of these states as a function of the electric field is a Stark map.

6.3.1 Electric field cancellation

We assume the voltage to electric field conversion from section 6.2 to be accurate for the moment, and plot Stark maps against the calculated applied electric field. This is acceptable here as we are seeking to minimise the electric field for the moment, rather than accurately measure it.⁶ We are able to fit to each Stark map with a parabola, as the DC Stark shift is proportional to $|\vec{E}|^2 = E_{x''}^2 + E_{y''}^2 + E_z^2$,⁷ meaning that we expect to see a shift parabolic in the electric field along the axis we are mapping, and an offset arising from the constant fields on the other two axes.

The first Stark map is taken along the $\pm\hat{z}$ axis, and is shown in figure 6.3. The applied voltages used to create this plot were up to $V_{\text{set}} = \pm 250$ mV, which is 83% of the voltage range achievable with our voltage source. The applied field for which we determine the stray field to be minimised is always taken from the shift of the $5s79d\ ^3D_1|m_j| = 0$ state, as that has the largest and best fitted shift⁸.

⁶We consider this calibration more carefully in section 6.3.2.

⁷Except when state mixing becomes significant, we will consider this issue later on.

⁸As we will discuss later, the $5s79d\ ^3D_1|m_j| = 1$ state does not quite have the parabolic shape that equation 2.4.3 would imply.

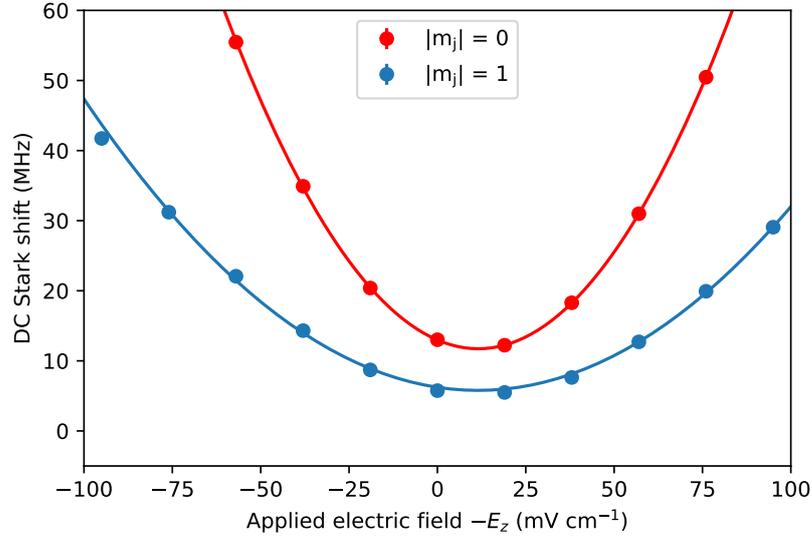


Figure 6.3: Stark map in the $\pm\hat{z}$ direction for the $5s79d^3D_1$ state. The measured DC Stark shift for the $|m_j| = 0$ and 1 states are shown as the red and blue points respectively. Each has a quadratic fitting, where the minima to the $|m_j| = 0$ curve is used to find where the stray field has been cancelled on the \hat{z} axis.

By finding the minimum of the $5s79d^3D_1|m_j| = 0$ fitted parabola in figure 6.3, the stray electric field on the \hat{z} axis is determined. If the electrode calibration were to be exactly correct, the stray electric field on the \hat{z} axis would be $E_z = -11.56 \pm 0.07_{\text{stat}}$ mVcm $^{-1}$. There will be other components of the stray electric field on the other two axes results in the splitting between the $5s5p^3P_1 \rightarrow 5s79d^3D_1$ $|m_j| = 0$ and 1 features in figure 6.3.

We then set the voltages on the electrode to the voltages that would correspond to the fitted minima of the $5s79d^3D_1|m_j| = 0$ state. This becomes the new ‘zero’ state for the measurements on the next axis, where we add the voltages in the configuration shown in figure 6.2(b) to this starting value. We do this as the sum of the three electrode configurations should cancel the stray field. We then add the voltages in the $\pm\hat{y}''$ direction, and take a Stark map along this axis as shown in figure 6.4. The minimum of the $5s79d^3D_1|m_j| = 0$ fitted parabola implies a stray field on the \hat{y}'' axis of $E_{y''} = -27.1 \pm 0.3_{\text{stat}}$ mVcm $^{-1}$, where again we do not consider the electrode calibration error at this point. As before, the electrode voltages are then

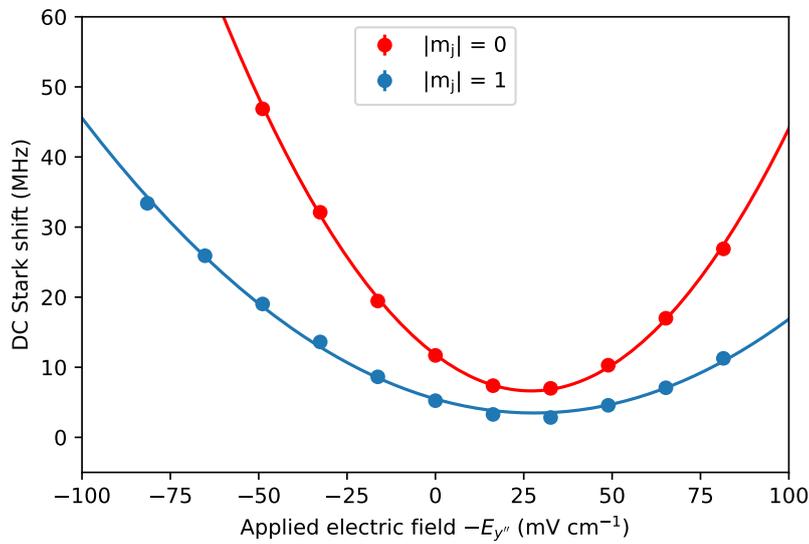


Figure 6.4: Stark map for the $5s79d\ ^3D_1$ state in the $\pm\hat{y}''$ direction, after stray field cancellation in the $\pm\hat{z}$ direction. The measured DC Stark shift for the $|m_j| = 0$ and 1 states are shown as the red and blue points respectively, with a quadratic fitting.

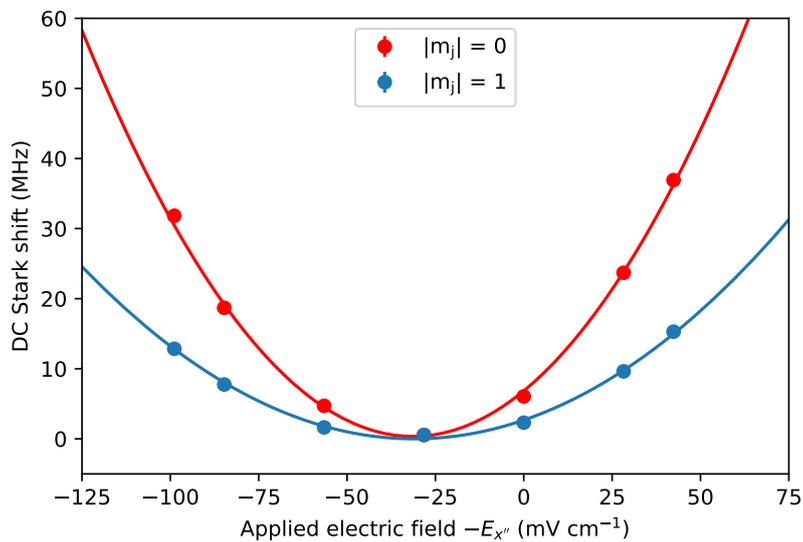


Figure 6.5: Stark map for the $5s79d\ ^3D_1$ state in the $\pm\hat{x}''$ direction, after stray field cancellation in the $\pm\hat{z}$ and $\pm\hat{y}''$ directions. The measured DC Stark shift for the $|m_j| = 0$ and 1 states are shown as the red and blue points respectively, with a quadratic fitting.

set to correspond to the minima of the fitting, and this forms the new ‘zero’ state for the measurements on the final axis. The last stark map is taken on the $\pm\hat{x}''$ axis and is shown in figure 6.5. Taking the fitted minima to the $5s79d^3D_1|m_j| = 0$ state we find that the electric field is best cancelled for $E_{x''} = +31.4 \pm 0.4_{\text{stat}} \text{ mVcm}^{-1}$. This value once again does not yet consider the calibration of the electric field.

To measure just how well the stray field has been cancelled by this method, we can look at the residual splitting between the $|m_j| = 0$ and 1 for the optimal applied field. However, the field is cancelled so well that the peaks for the minima in the depletion features used to make the last Stark map (figure 6.5) are so well overlapped that we cannot resolve the splitting of the features as well as done previously in figure 6.1. To estimate the remaining residual field after cancellation, we look at the residual splitting in the parabolic fits of the $|m_j| = 0$ and 1 peaks. The residual splitting is $0.34 \pm 0.38 \text{ MHz}$, consistent with zero to within errors. A residual splitting of 0.34 MHz is over 20 times smaller than the $7.26 \pm 0.18 \text{ MHz}$ splitting we saw in figure 6.1 for no applied electric field. As the DC Stark shift is proportional to the square of the electric field, the DC Stark shift after stray field cancellation should be reduced by a factor of ~ 400 . This factor of improvement is independent of the calibration of the electric field against the applied electrode voltages, as it arises purely in terms of the residual splitting of the $|m_j| = 0$ and 1 states. We now consider the calibration of the electrodes through comparison with ARC in section 6.3.2, such that we can provide an accurate measure of the stray electric field, and the residual electric field after cancellation.

6.3.2 Comparison with ARC

To better understand the Stark map data we have taken, it is helpful to make comparison with simulations predicting the Stark maps. For that we use the ARC 3.0 python toolbox⁹ [1]. ARC also can be used to calibrate the field produced by our electrodes as a function of the voltage applied, allowing us to use our Stark maps from the previous section to measure the size of the stray field.

⁹Alkali(ne) Rydberg Calculator.

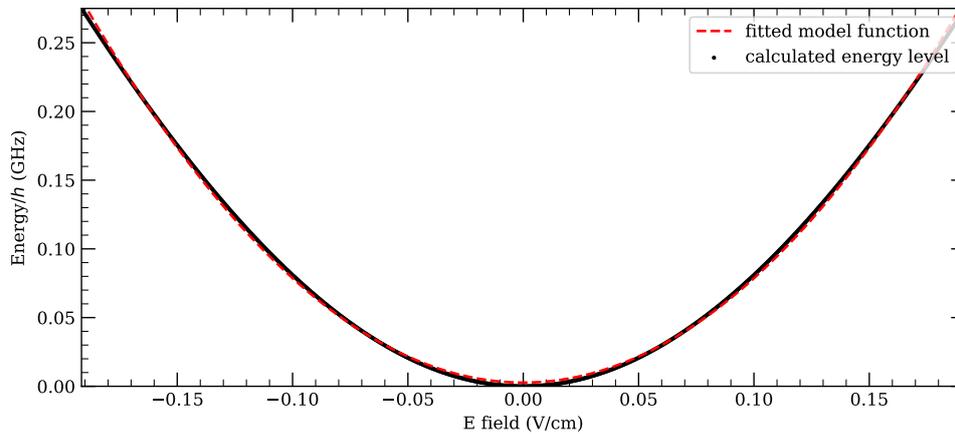


Figure 6.6: Simulated Stark map for $5s79d\ ^3D_1 |m_j| = 0$ over $\pm 200\ \text{mVcm}^{-1}$ range. Where a parabolic fitting of the DC Stark shift works well.

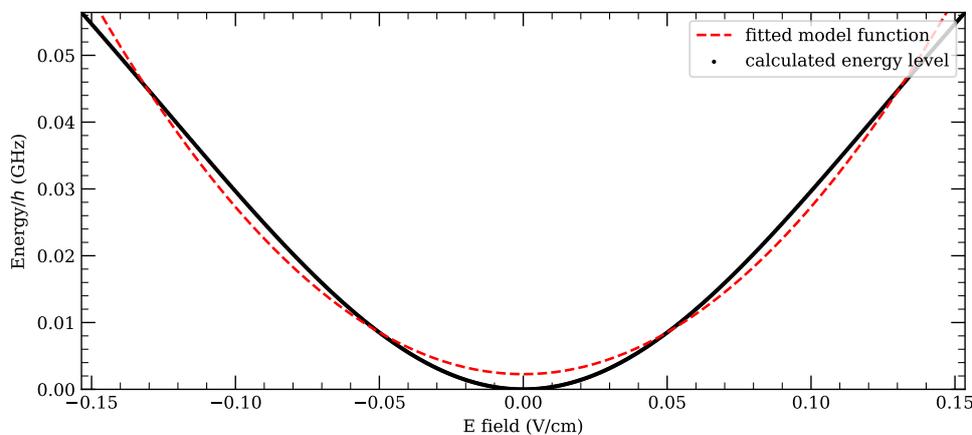


Figure 6.7: Simulated Stark map for $5s79d\ ^3D_1 |m_j| = 1$ over $\pm 200\ \text{mVcm}^{-1}$ range. Clearly a parabolic fitting does not work here, so the DC Stark shift is not quite as we would expect from equation 2.4.3. As we will see in figure 6.8, the polarisability α has become a function of the electric field due to significant mixing of states.

Plotting over a $\pm 200 \text{ mVcm}^{-1}$ range, figures 6.6 and 6.7 show the expected shifts of each of the $5s79d \ ^3D_1 |m_j| = 0$ and 1 states respectively. For the $5s79d \ ^3D_1 |m_j| = 0$ state we see a good parabolic fitting the to the simulation. For the $5s79d \ ^3D_1 |m_j| = 1$ state however, the Stark map is not well fitted by a parabola, suggesting that the DC Stark shift is not well described by the simple form shown in equation 2.4.3 for this Rydberg state when using electric fields this large.

The origin of the non-parabolic shift of the $5s79d \ ^3D_1 |m_j| = 1$ state in figure 6.7 can be found in figure 6.8, which shows the mixing of the neighbouring Rydberg states as a function of the applied electric field. Due to the Stark interaction term in the Hamiltonian of the atom, the atomic state ($|\chi(E)\rangle$) at a particular electric field (E) is a linear superposition of the bare atomic states ($|n, s, l, j\rangle$), with mixing coefficients ($c_{n,l,j}(E)$)¹⁰ which depend upon the electric field [76]:

$$|\chi(E)\rangle = \sum_{n,l,j} c_{n,l,j}(E) |n, s, l, j\rangle. \quad (6.3.2)$$

As the electric field reaches the order of $\pm 100 \text{ mVcm}^{-1}$ (larger than electric fields used in figure 6.5), the Stark interaction term in the Hamiltonian becomes larger and causes more significant mixing between the bare states of the atom. At these large interactions, nearby bare Rydberg states have significant contribution to the eigenstate of the atom, such that the polarisability in equation 2.4.3 has become a function of electric field, and the quadratic scaling no longer holds. The mixing in figure 6.8 primarily comes from the $5s79d \ ^3D_2 |m_j| = 1$ state.

This mixing behaviour is also seen on the $5s79d \ ^3D_1 |m_j| = 0$ state in figure 6.9, primarily due to the $5s79p \ ^3P_0 |m_j| = 0$ state, but is to a lesser extent and does not lead to as significant deviations from a parabola in figure 6.6. Carefully looking at both simulated Stark maps (figures 6.6 and 6.7), we can see that the parabolic fitting it not ideal, overestimating at zero and the edges of the fitted range, and underestimating in between. This is also apparent in our experimental data, particularly for the $5s79d \ ^3D_1 |m_j| = 1$ state in the original \hat{z} Stark map (figure 6.3), where the electric fields were the largest, as none of the stray fields on any axes had been cancelled. This Stark map is re-plotted in figure 6.10, where we can clearly

¹⁰The Stark interaction does not mix states with different s or $|m_j|$ [192].

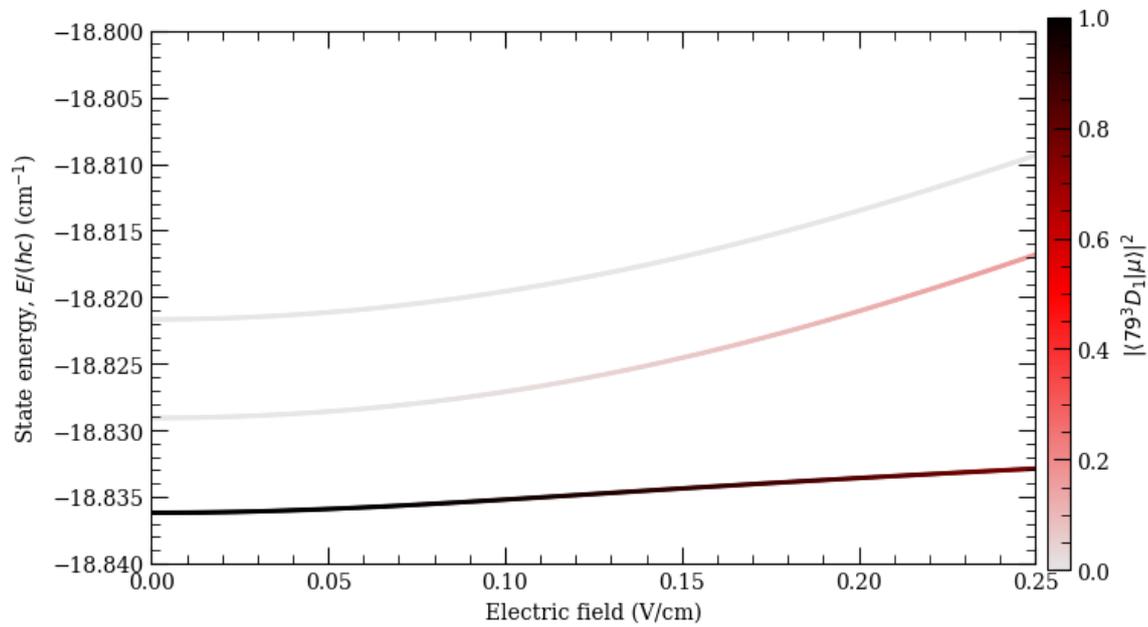


Figure 6.8: Mixing of the Rydberg states as a function of applied electric field for the $5s79d\ ^3D_1|m_j| = 1$ state.

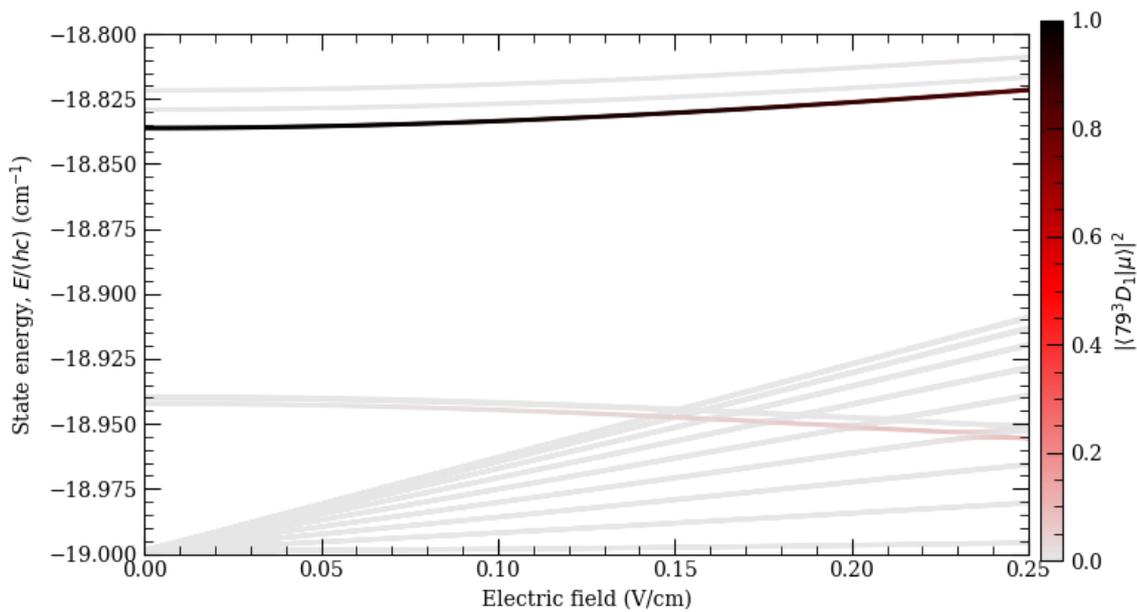


Figure 6.9: Mixing of the Rydberg states as a function of applied electric field for the $5s79d\ ^3D_1|m_j| = 0$ state.

see the deviations in the $5s79d\ ^3D_1|m_j| = 1$ state away from the parabolic fitting.

We must consider what this non-parabolic behaviour means in terms of the analysis of our experimental results. The $5s79d\ ^3D_1|m_j| = 0$ state is well fitted by a parabola over the electric field range of interest, showing reasonable residuals in figure 6.10, so we expect no significant issues as to the cancellation method in section 6.3.1 (where we used the $|m_j| = 0$ fitting to the state to find the cancellation field). The fitting of the $5s79d\ ^3D_1|m_j| = 1$ state is not as good, which raises potential issues. Bad fitting around the two minima of the final Stark map 6.5 would lead to errors in the fitted residual splitting between the two states, and thus the measurement of the residual stray field. Due to the cancellation along the other two axes, the residuals here are small, with root mean square deviation of 0.5 MHz. The deviation is roughly on the scale of the error in the fitting, which should not significantly impact the accuracy of the residual stray field measurement.

The non-parabolic shift of the $5s79d\ ^3D_1|m_j| = 1$ state also challenges an initial assumption of the Stark maps that we can fit parabolas along each direction as the stray field on the other two axes factors out as an offset in the form $|\vec{E}|^2 = E_x^2 + E_{y''}^2 + E_{z''}^2$, as the shift is no longer directly proportional to $|\vec{E}|^2$. This does not cause issues in the electric field cancellations as those used the $5s79d\ ^3D_1|m_j| = 0$ state, but it does mean that we should now use ARC as our model to compare with the experimental data. We now use ARC to extract a calibration of the electric field from the voltages applied to our electrodes, and use this to measure the stray electric field.

As we have Stark maps for all three axes, we begin with the data shown in figure 6.5. Assuming there to be no significant residual stray field after 3D cancellation, we plot the experimental data against the ARC simulations of the Stark maps. This is shown in figure 6.11, where the points show the experimental data, and the curves now show ARC simulations, rather than fitted parabolas. The experimental data is now plotted against the actual electric field at the atoms. We allowed a free parameter as a linear scaling of the field produced by the electrodes away from the Inventor simulation (compared to the conversion factors given in section 6.2), where the parameter has been chosen to fit the data to the ARC prediction. The process

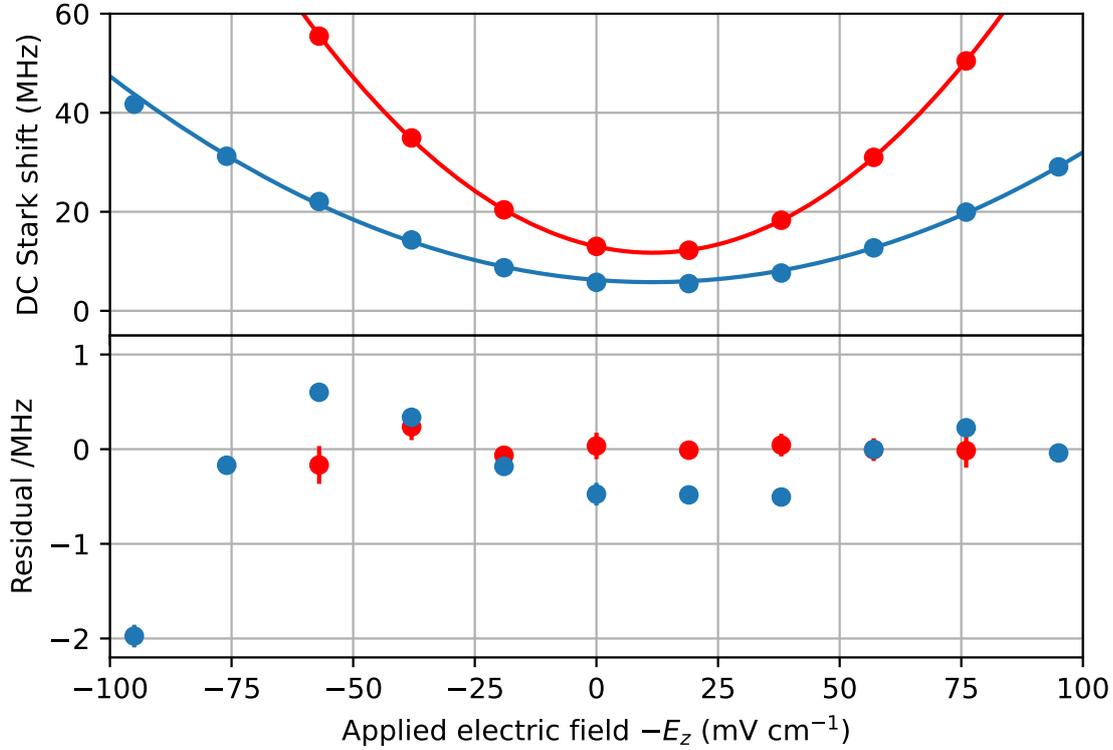


Figure 6.10: Stark map in the $\pm\hat{z}$ direction with residuals, where the points show the experimental data, and the curves are parabolic fits. The $|m_j| = 1$ residuals (shown in blue) behave qualitatively like those seen in the ARC modelled figure 6.7, though any direct comparison would be subject to both the electrode calibration, and accounting for the stray field in the other two directions. The latter point of the stray field in the other two directions can no longer be trivially ignored, as if the shift is not proportional to $|\vec{E}|^2 = E_{x''}^2 + E_{y''}^2 + E_z^2$, then we can no longer factor out the terms in the xy plane to just provide an offset to the Stark map. On the other hand, the $|m_j| = 0$ residuals (shown in red) behave well, suggesting a fitting with a parabola for that state is still appropriate.

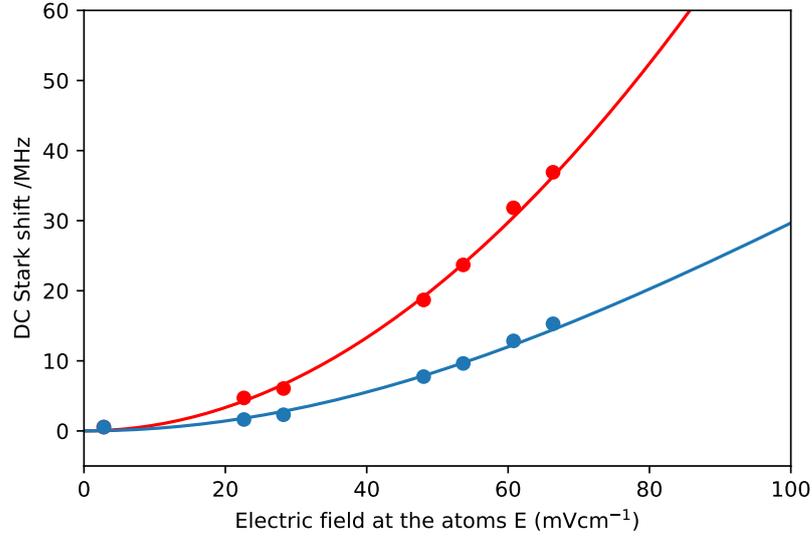


Figure 6.11: Stark map for the $5s79d\ ^3D_1$ state in the $\pm\hat{x}''$ direction as a function of the calibrated electric field at the atoms. The curves show the predicted Stark map produced using ARC. The points show the experimental data after correction with a linear scaling factor in the calibration of the electric fields. The calibration produces the best fit between simulation and experiment for an electric field $10 \pm 2\%$ smaller than predicted in section 6.2.

is then repeated using data from figures 6.4 and 6.3 along the \hat{y}'' and \hat{z} directions, respectively, assuming the calibrated stray field is present on the other axes. The calibration parameters and determined electric fields are summarised in table 6.1

It is reassuring to see that the adjustments in the electric field conversion factors in table 6.1 are to within errors in the xy plane, where we expect the calibration parameters to be similar in the plane of the hexagonal electrode rings. The calibration parameter on the \hat{z} axis being larger than the other two values is not of particular concern, as this depends on the separation of the two hexagonal rings of electrodes, which may be closer together than the model would suggest.

ARC fits the data well in each of the three directions using a single free parameter in the calibration of each electrode. This implies that the dominant difference between ARC and our data was indeed the electric field calibration, which has now been refined using ARC. It is not surprising that the results from the inventor sim-

Axis	Calibration parameter	Stray electric field (mVcm ⁻¹)
\hat{x}''	0.90(2)	$+28.2 \pm 0.4_{\text{stat}} \pm 0.6_{\text{sys}}$
\hat{y}''	0.93(3)	$-25.2 \pm 0.3_{\text{stat}} \pm 0.8_{\text{sys}}$
\hat{z}	1.07(2)	$-12.4 \pm 0.1_{\text{stat}} \pm 0.2_{\text{sys}}$
Total	-	$39.8 \pm 0.5_{\text{stat}} \pm 1.0_{\text{sys}}$

Table 6.1: Calibration of the electric field measurements against ARC, compared to the predicted conversion factors given in section 6.2. The calibrated stray electric fields are also summarised.

ulation in section 6.2 were not perfect, as it did not fully model the entire vacuum system, and components may have shifted during backing out the vacuum chamber. Calibration discrepancies seem reasonable, within 10% of the simulations.

We have neglected to consider uncertainty arising in the Stark maps calculated by ARC, which will add additional errors to our calibration. The dominant contribution to the uncertainty in the calculated Stark maps is the uncertainty in the zero field energies of the Rydberg energy levels, which are known to absolute accuracy of ~ 10 MHz for the S and D series [66], but less well known for other values of L. Uncertainty in the energy levels leads to uncertainty in the wavefunction and the dipole matrix elements, giving uncertainty in the Stark map [134, 141]. The Rydberg energy series are calculated using optimised parameters of the quantum defects, using the Rydberg Ritz formula (equation 2.4.1). Uncertainties in the quantum defects (equation 2.4.2) characterise the uncertainty in the Rydberg energy levels, and taking extreme values of the quantum defects and propagating that through the calculation can be used to determine the error on the calculated Stark maps [134]. A full treatment of the uncertainty in the predictions from ARC are left for future work. Precision measurements of the Rydberg energy series [65, 66] are of great value in narrowing the uncertainty in the predicted Stark maps, where measurements using a frequency comb (see appendix E.1.1) can be used to assist in precision measurements of these states [161].

Overall, combining the calibrated stray electric fields together gives a total stray electric field at the position of the tweezer of $39.8 \pm 0.5_{\text{stat}} \pm 1.0_{\text{sys}}$ mVcm⁻¹ (neglecting

uncertainty arising from the predictions in ARC). A total stray electric field of this strength corresponds to a shift of $(1.14 \pm 0.03_{\text{stat}} \pm 0.06_{\text{sys}}) \times 10^{-19}$ on the clock transition, below a second in the age of the universe, but above the state-of-the-art fractional frequency precision for optical atomic clocks [8].

The stray electric field is found to be primarily in the horizontal plane, with the vector direction of the stray field being at $18 \pm 1^\circ$ below the horizontal plane. In the horizontal plane, the electric field is at $-33 \pm 2^\circ$ to the \hat{y} axis (where the \hat{y} axis is the axis of the *in-vacuo* lenses as shown in figure 3.1). The origins of the stray electric field in the chamber are unknown, though the d^{-4} dependence on patch potentials [43] means that the source of the field is likely close to the atoms. Possible contributions to the stray electric field include the work function difference between the ITO coated *in-vacuo* lenses and their steel holders, build up of strontium on surfaces, or any of the other *in-vacuo* equipment. Regardless of the origin of the stray electric field, the stray field in our chamber is not particularly large, and we can use our electrodes to cancel it significantly.

Taking the residual splitting after electric field cancellation (shown in section 6.3.1) of 0.34 ± 0.38 MHz, compared to the initial splitting of 7.26 ± 0.18 MHz (shown in figure 6.1), we can now calculate the residual stray field after cancellation using the electric field calibration. The residual stray field is $1.9_{-1.9}^{+2.2}$ mVcm⁻¹, where the shift and uncertainty on the clock transition arising from the DC Stark shift is $2.6_{-2.6}^{+9.5} \times 10^{-22}$, which would not be a limiting factor on any existing optical atomic clock. For a more exhaustive study to determine and cancel the residual stray electric field with greater precision, it would be beneficial to repeat the Stark maps on each axis after electric field cancellation.

6.4 Summary

We have characterised the stray electric field in our vacuum chamber at the position of the tweezers using Rydberg atoms, without the need to use large electric fields or make measurements of the clock transition. We measure the stray field in our vacuum chamber to be $|\vec{E}_{\text{stray}}| = 39.8 \pm 0.5_{\text{stat}} \pm 1.0_{\text{sys}}$ mVcm⁻¹, which would pro-

duce a shift on the clock transition $(1.14 \pm 0.03_{\text{stat}} \pm 0.06_{\text{sys}}) \times 10^{-19}$, small enough that we do not expect to have any limiting impact on performance of future clock experiments.

We can further cancel the electric field using our stable voltage source on the twelve *in-vacuo* electrodes to $1.9_{-1.9}^{+2.2}$ mVcm⁻¹, with a clock shift of $2.6_{-2.6}^{+9.5} \times 10^{-22}$ which would be a negligible error even for state-of-the-art atomic clocks [8]. This electric field cancellation also demonstrates the ability to reduce the electric field environment for precision measurement, where electric field induced shifts of Rydberg states used here can be reduced to the < 1 MHz level. In future work the intensity of the Rydberg excitation laser should be reduced to decrease power broadening and resolve the line centres to better precision and at lower residual fields, where the natural linewidth is roughly two orders of magnitude narrower than we have observed in this preliminary study. Precision measurement of Rydberg states will be of great value in refining predictions made by ARC [1], as well as tests of multi-channel quantum defect models [67].

The choice of the 5s79d ³D₁ state for electric field cancellation offered relatively large DC Stark shifts for the applied field, making it easy to resolve the shifts, but also raising the issue of the non-parabolic behavior and mixing of nearby Rydberg states. If we alternatively had used the 5s79s ³S₁ state, we would see less mixing of nearby Rydberg states and a more parabolic behavior, but this state comes with drawbacks for electric field measurement. The shifts for $|m_j| = 0$ and 1 are predicted by ARC to be only 1.08 and 1.02 $\frac{\text{GHz}}{(\text{Vcm}^{-1})^2}$ respectively, an order of magnitude smaller than those in the 5s79d ³D₁ state. The splitting is also much narrower, so resolving the two lines would be an issue. These means that the 5s79d ³D₁ $|m_j| = 0$ state seems to be better suited for this measurement. The difference in DC Stark shift does however offer the option to cancel the electric field using the 5s79d ³D₁ state, then perform precision measurement of the 5s79s ³S₁ state where the electric field shift is a factor of 14 smaller, such that we would then have the DC Stark shift of the Rydberg state controlled to the < 100 kHz level. Control of the shifts for the 5sns ³S₁ Rydberg series are of interest as those are the states used for Rydberg dressing [21].

Furthermore, if we were to measure the electric field using a Rydberg state with higher n than the Rydberg state to be used in a later experiment, the electric field shift could be even smaller. As an example, by measuring the electric field with a Rydberg state of twice the principal quantum number we want to use, the electric field cancellation ought to be $\sim 2^7 = 128$ times better, leading to the possibilities of controlling the DC Stark shifts of Rydberg states to the < 1 kHz level.

With the electric field controlled such that DC Stark shifts can be at the < 1 kHz level, and with atoms held in optical tweezers, we have a platform which could be used for precision measurements of Rydberg state energies. Optical tweezers allow longer trapping times than searching in MOTs [65]. Single atom/low atom number trapping in tweezers would manage any density dependant shifts, and if desired then tweezers can be turned off during Rydberg excitation to avoid the addition of AC Stark shifts from the trapping light. Further cooling of the atoms should be considered to limit Doppler broadening, and a characterisation of the magnetic and BBR fields should be done, where these shifts should be less than 1 MHz [161]. Along with the depletion detection method used at present, the MCP could also be used to count ions as an alternative detection scheme [155]. A key factor to enable precision measurement in this platform is the use of our optical frequency comb [161] (see appendix E.1.1) which can measure the frequency of both the 698/689 nm and 317/319 nm lasers, where we will not be limited by wavemeter accuracy as is often the case in other work [65]. The challenges are then to manage to systematic shifts of the Rydberg levels, but there is clear scope with which the precision of these measurements can be improved.

Chapter 7

Conclusions and outlook

In this work we have demonstrated progress towards the implementation of a spin-squeezed tweezer array clock. We have demonstrated the ability to image single atoms in optical tweezers with high fidelity, using a low background SPAD array detector. We have implemented 2D magic wavelength optical tweezer arrays using a spatial light modulator (SLM), and have demonstrated site-resolved imaging using a custom high-magnification imaging system. We have used Rydberg atoms to characterise the low electric field environment in our vacuum chamber, demonstrating precision electric field cancellation with applications from precision measurements of Rydberg states to managing clock systematic uncertainties.

Chapter 3 introduced the experimental apparatus required to cool, trap, manipulate and image cold strontium in optical tweezers. In chapter 4 we then demonstrate single atom and number-resolved imaging in a deep 532 nm tweezer [43]. We demonstrate imaging techniques from fast (200 μs) imaging, to high fidelity $F_1 = 0.998(2)$ imaging of single atoms. To do this we combined the negligible readout noise of our SPAD detector through a series of strong imaging pulses, interleaved with long periods of Sisyphus cooling. Moving to magic wavelength tweezers at 813 nm, this imaging fidelity should increase further with the ability to close loss via the $5s5p\ ^1P_1 \rightarrow 5s4d\ ^1D_2$ loss channel through the use of repumping lasers [131].

We implemented magic wavelength tweezer arrays in chapter 5. We produce tweezers of waist 2.25 μm , in arbitrary 3D arrays permitted by our SLM. A high

magnification imaging system was built with corrective elements to allow us to perform site-resolved imaging of 2D tweezer arrays with spacing of $\sim 6 \mu\text{m}$. This imaging performance is despite significant aberration on our *in-vacuo* lens due to a pre-existing design flaw. There is a clear path to high fidelity imaging of single atoms in the magic wavelength tweezer arrays, building upon the work of chapter 4 with the use of attractive Sisyphus cooling and appropriate repumping. The implementation of this tweezer array at the magic wavelength is a key step towards implementation of a tweezer array clock. The tweezer array also serves as a platform for precision measurements of Rydberg states and for implementing Rydberg-dressed clock schemes, as it is a re-configurable array of tweezers with arbitrary separations and geometry.

Future development of the experimental system should include improvements both in and out of the vacuum chamber. When the vacuum chamber is next opened the *in-vacuo* should be replaced to remove aberration and to increase the numerical aperture for enhanced photon collection, and to produce tweezers with smaller waists for better confinement and closer spaced arrays. The MCP connections inside the vacuum chamber should also be improved to reduce ringing on the signal. This should allow us to use the Rydberg laser and electric field control from chapter 6, along with the auto-ionisation laser, to detect single ions from the tweezer array. Outside of the vacuum chamber we could consider a secondary fluorescence detection system using a high quantum efficiency EMCCD, with the option of both an EMCCD and SPAD array offering versatility in detection schemes. The stability of the 813 nm laser for the magic wavelength tweezer array should be improved in order to reduce shifts and instability for future clock measurements. To maximise the number of atoms which can be loaded within a Rydberg dressed radius, we should also consider methods for enhanced loading of single atoms arrays [50, 56], and the extension of the tweezer array and imaging system to 3D [35].

The electric field at location of the atoms inside the vacuum chamber was characterised in chapter 6. The high sensitivity to electric fields of Rydberg atoms was used to both measure and cancel the stray electric field in our chamber. The long working distance of our *in-vacuo* lenses affords us a low electric field environment in

which to do clock and Rydberg experiments. Cancellation of the stray electric field with our *in-vacuo* electrodes improves this further, to the level which would not limit any existing atomic clock [8], and producing an environment compatible with the precision measurement of Rydberg states, which we can do with our frequency comb. Precision measurement of triplet Rydberg states would be of value in refining and testing multi-channel quantum defect models [67], and improving predictions from ARC 3.0 [1]. Potential alternative excitation schemes include two photon excitation to access the less studied P and F series, or one photon excitation using 317 nm light to excite from the $5s5p\ ^3P_0$ state. The use of the Rydberg excitation laser is also essential to studying Rydberg dressing in both MOTs and tweezers via the $5s5p\ ^3P_1$ and $5s5p\ ^3P_0$ states respectively. Rydberg-dressed MOT [149] experiments could be improved through beam profiling optics for the Rydberg excitation light, where the limiting performance is the non-uniformity of the dressing light. Rydberg-dressing in tweezer arrays is key part of our goal of implementing spin squeezing in a tweezer array clock [21], for improvements in optical atomic clocks beyond the QPN limit.

The next steps of our experiment is to use the clock laser light on our strontium atoms. The characterisation of the clock laser as the local oscillator and the frequency comb as a counter are detailed in appendix E.1. The clock laser has a measured linewidth of < 1 Hz at 1 second, with the cavity drift measured. The frequency comb has demonstrated good stability when frequency counting, performing near/at the limit arising from the stability of the GPS reference. We plan to allow the clock transition using MIS as discussed in appendix E.2, considering the different systematic uncertainties in play. We are now ready to search for the clock transition with the atoms, towards implementing an atomic clock in our tweezer array.

Bibliography

- [1] E. J. Robertson, N. Šibalić, R. M. Potvliege, and M. P. A. Jones, *ARC 3.0: An expanded Python toolbox for atomic physics calculations*, *Computer Physics Communications* **261**, 107814 (2021).
- [2] T. P. Wiles, *Dynamics of bright solitary matter-waves*, *Ph.D. thesis*, Durham University (2013).
- [3] T. L. Nicholson, S. L. Campbell, R. B. Hutson, G. E. Marti, B. J. Bloom, R. L. McNally, W. Zhang, M. D. Barrett, M. S. Safronova, G. F. Strouse, W. L. Tew, and J. Ye, *Systematic evaluation of an atomic clock at 2×10^{-18} total uncertainty*, *Nature Communications* **6**, 6896 (2015).
- [4] T. Bothwell, D. Kedar, E. Oelker, J. M. Robinson, S. L. Bromley, W. L. Tew, J. Ye, and C. J. Kennedy, *JILA SrI optical lattice clock with uncertainty of 2.0×10^{-18}* , *Metrologia* **56**, 065004 (2019).
- [5] N. Ohmae, M. Takamoto, Y. Takahashi, M. Kokubun, K. Araki, A. Hinton, I. Ushijima, T. Muramatsu, T. Furumiya, Y. Sakai, N. Moriya, N. Kamiya, K. Fujii, R. Muramatsu, T. Shiimado, and H. Katori, *Transportable Strontium Optical Lattice Clocks Operated Outside Laboratory at the Level of 10^{-18} Uncertainty*, *Advanced Quantum Technologies* **4** (2021).
- [6] W. F. McGrew, X. Zhang, R. J. Fasano, S. A. Schäffer, K. Beloy, D. Nicolodi, R. C. Brown, N. Hinkley, G. Milani, M. Schioppo, T. H. Yoon, and A. D. Ludlow, *Atomic clock performance enabling geodesy below the centimetre level*, *Nature* **564**, 87 (2018).

- [7] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibbrandt, $^{27}\text{Al}^+$ *Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18}* , *Phys. Rev. Lett.* **123**, 033201 (2019).
- [8] T. Bothwell, C. J. Kennedy, A. Aepli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye, *Resolving the gravitational redshift across a millimetre-scale atomic sample*, *Nature* **602**, 420 (2022).
- [9] S. Kolkowitz, I. Pikovski, N. Langellier, M. D. Lukin, R. L. Walsworth, and J. Ye, *Gravitational wave detection with optical lattice atomic clocks*, *Phys. Rev. D* **94**, 124043 (2016).
- [10] R. Hobson, *An Optical Lattice Clock with Neutral Strontium*, *DPhil thesis*, University of Oxford (2016).
- [11] N. Huntemann, B. Lipphardt, C. Tamm, V. Gerginov, S. Weyers, and E. Peik, *Improved Limit on a Temporal Variation of m_p/m_e from Comparisons of Yb^+ and Cs Atomic Clocks*, *Phys. Rev. Lett.* **113**, 210802 (2014).
- [12] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt, *Optical atomic clocks*, *Rev. Mod. Phys.* **87**, 637 (2015).
- [13] M. Safronova, *The Search for Variation of Fundamental Constants with Clocks*, *Annalen der Physik* **531**, 1800364 (2019).
- [14] G. Barontini, L. Blackburn, V. Boyer, F. Butuc-Mayer, X. Calmet, J. R. Crespo López-Urrutia, E. A. Curtis, B. Darquié, J. Dunningham, N. J. Fitch, E. M. Forgan, K. Georgiou, P. Gill, R. M. Godun, J. Goldwin, V. Guarrera, A. C. Harwood, I. R. Hill, R. J. Hendricks, M. Jeong, M. Y. H. Johnson, M. Keller, L. P. Kozhiparambil Sajith, F. Kuipers, H. S. Margolis, C. Mayo, P. Newman, A. O. Parsons, L. Prokhorov, B. I. Robertson, J. Rodewald, M. S. Safronova, B. E. Sauer, M. Schioppo, N. Sherrill, Y. V. Stadnik, K. Szymaniec, M. R. Tarbutt, R. C. Thompson, A. Tofful, J. Tunesi, A. Vecchio, Y. Wang, and S. Worm, *Measuring the stability of fundamental constants with a network of clocks*, *EPJ Quantum Technology* **9**, 12 (2022).

- [15] 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).
- [16] R. Bondarescu, A. Schäfer, A. Lundgren, G. Hetényi, N. Houlié, P. Jetzer, and M. Bondarescu, *Ground-based optical atomic clocks as a tool to monitor vertical surface motion*, *Geophysical Journal International* **202**, 1770 (2015).
- [17] P. Gill, *When should we change the definition of the second?*, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **369**, 4109 (2011).
- [18] P. Gill, *Is the time right for a redefinition of the second by optical atomic clocks?*, *Journal of Physics: Conference Series* **723**, 012053 (2016).
- [19] J. Lodewyck, *On a definition of the SI second with a set of optical clock transitions*, *Metrologia* **56**, 055009 (2019).
- [20] W. M. Itano, J. C. Bergquist, J. J. Bollinger, J. M. Gilligan, D. J. Heinzen, F. L. Moore, M. G. Raizen, and D. J. Wineland, *Quantum projection noise: Population fluctuations in two-level systems*, *Phys. Rev. A* **47**, 3554 (1993).
- [21] L. I. R. Gil, R. Mukherjee, E. M. Bridge, M. P. A. Jones, and T. Pohl, *Spin Squeezing in a Rydberg Lattice Clock*, *Phys. Rev. Lett.* **112**, 103601 (2014).
- [22] N. Schine, A. W. Young, W. J. Eckner, M. J. Martin, and A. M. Kaufman, *Long-lived Bell states in an array of optical clock qubits*, *Nature Physics* **18**, 1067 (2022).
- [23] A. Kuzmich, L. Mandel, and N. P. Bigelow, *Generation of Spin Squeezing via Continuous Quantum Nondemolition Measurement*, *Phys. Rev. Lett.* **85**, 1594 (2000).

- [24] A. Louchet-Chauvet, J. Appel, J. J. Renema, D. Oblak, N. Kjaergaard, and E. S. Polzik, *Entanglement-assisted atomic clock beyond the projection noise limit*, *New Journal of Physics* **12**, 065032 (2010).
- [25] M. H. Schleier-Smith, I. D. Leroux, and V. Vuletić, *States of an Ensemble of Two-Level Atoms with Reduced Quantum Uncertainty*, *Phys. Rev. Lett.* **104**, 073604 (2010).
- [26] O. Hosten, N. J. Engelsen, R. Krishnakumar, and M. A. Kasevich, *Measurement noise 100 times lower than the quantum-projection limit using entangled atoms*, *Nature* **529**, 505 (2016).
- [27] K. C. Cox, G. P. Greve, J. M. Weiner, and J. K. Thompson, *Deterministic Squeezed States with Collective Measurements and Feedback*, *Phys. Rev. Lett.* **116**, 093602 (2016).
- [28] B. Braverman, A. Kawasaki, E. Pedrozo-Peñafiel, S. Colombo, C. Shu, Z. Li, E. Mendez, M. Yamoah, L. Salvi, D. Akamatsu, Y. Xiao, and V. Vuletić, *Near-Unitary Spin Squeezing in ^{171}Yb* , *Phys. Rev. Lett.* **122**, 223203 (2019).
- [29] W. Bowden, A. Vianello, I. R. Hill, M. Schioppo, and R. Hobson, *Improving the Q Factor of an Optical Atomic Clock Using Quantum Nondemolition Measurement*, *Phys. Rev. X* **10**, 041052 (2020).
- [30] M.-Z. Huang, J. A. de la Paz, T. Mazzoni, K. Ott, A. Sinatra, C. L. G. Alzar, and J. Reichel, *Self-amplifying spin measurement in a long-lived spin-squeezed state*, preprint arXiv:2007.01964 (2020).
- [31] E. Pedrozo-Peñafiel, S. Colombo, C. Shu, A. F. Adiyatullin, Z. Li, E. Mendez, B. Braverman, A. Kawasaki, D. Akamatsu, Y. Xiao, and V. Vuletić, *Entanglement on an optical atomic-clock transition*, *Nature* **588**, 414 (2020).
- [32] M. Kitagawa and M. Ueda, *Squeezed spin states*, *Phys. Rev. A* **47**, 5138 (1993).
- [33] M. Takamoto and H. Katori, *Spectroscopy of the $^1S_0-^3P_0$ Clock Transition of ^{87}Sr in an Optical Lattice*, *Phys. Rev. Lett.* **91**, 223001 (2003).

- [34] S. Origlia, M. S. Pramod, S. Schiller, Y. Singh, K. Bongs, R. Schwarz, A. Al-Masoudi, S. Dörscher, S. Herbers, S. Häfner, U. Sterr, and C. Lisdat, *Towards an optical clock for space: Compact, high-performance optical lattice clock based on bosonic atoms*, *Phys. Rev. A* **98**, 053443 (2018).
- [35] D. Barredo, V. Lienhard, S. De Léséleuc, T. Lahaye, and A. Browaeys, *Synthetic three-dimensional atomic structures assembled atom by atom*, *Nature* **561**, 79 (2018).
- [36] S. de Léséleuc, *Quantum simulation of spin models with assembled arrays of Rydberg atoms*, *Ph.D. thesis*, Institut d'Optique Graduate School (2018).
- [37] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, *Atom-by-atom assembly of defect-free one-dimensional cold atom arrays*, *Science* **354**, 1024 (2016).
- [38] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Probing many-body dynamics on a 51-atom quantum simulator*, *Nature* **551**, 579 (2017).
- [39] F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys, *Single-Atom Trapping in Holographic 2D Arrays of Microtraps with Arbitrary Geometries*, *Phys. Rev. X* **4**, 021034 (2014).
- [40] H. Labuhn, *Rydberg excitation dynamics and correlations in arbitrary 2D arrays of single atoms*, *Ph.D. thesis*, Institut d'Optique Graduate School (2016).
- [41] T. Xia, M. Lichtman, K. Maller, A. W. Carr, M. J. Piotrowicz, L. Isenhower, and M. Saffman, *Randomized Benchmarking of Single-Qubit Gates in a 2D Array of Neutral-Atom Qubits*, *Phys. Rev. Lett.* **114**, 100503 (2015).
- [42] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier, *Sub-Poissonian loading of single atoms in a microscopic dipole trap*, *Nature* **411**, 1024 (2001).

- [43] N. C. Jackson, R. K. Hanley, M. Hill, F. Leroux, C. S. Adams, and M. P. A. Jones, *Number-resolved imaging of ^{88}Sr atoms in a long working distance optical tweezer*, *SciPost Phys.* **8**, 38 (2020).
- [44] M. A. Norcia, A. W. Young, and A. M. Kaufman, *Microscopic Control and Detection of Ultracold Strontium in Optical-Tweezer Arrays*, *Phys. Rev. X* **8**, 041054 (2018).
- [45] A. Cooper, J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres, *Alkaline-Earth Atoms in Optical Tweezers*, *Phys. Rev. X* **8**, 041055 (2018).
- [46] M. A. Norcia, A. W. Young, W. J. Eckner, E. Oelker, J. Ye, and A. M. Kaufman, *Seconds-scale coherence on an optical clock transition in a tweezer array*, *Science* **366**, 93 (2019).
- [47] I. S. Madjarov, A. Cooper, A. L. Shaw, J. P. Covey, V. Schkolnik, T. H. Yoon, J. R. Williams, and M. Endres, *An Atomic-Array Optical Clock with Single-Atom Readout*, *Phys. Rev. X* **9**, 041052 (2019).
- [48] A. Urech, I. H. A. Knottnerus, R. J. C. Spreeuw, and F. Schreck, *Narrow-line imaging of single strontium atoms in shallow optical tweezers*, *Phys. Rev. Research* **4**, 023245 (2022).
- [49] S. Saskin, J. T. Wilson, B. Grinkemeyer, and J. D. Thompson, *Narrow-Line Cooling and Imaging of Ytterbium Atoms in an Optical Tweezer Array*, *Phys. Rev. Lett.* **122**, 143002 (2019).
- [50] A. Jenkins, J. W. Lis, A. Senoo, W. F. McGrew, and A. M. Kaufman, *Ytterbium Nuclear-Spin Qubits in an Optical Tweezer Array*, *Phys. Rev. X* **12**, 021027 (2022).
- [51] S. Ma, A. P. Burgers, G. Liu, J. Wilson, B. Zhang, and J. D. Thompson, *Universal Gate Operations on Nuclear Spin Qubits in an Optical Tweezer Array of ^{171}Yb Atoms*, *Phys. Rev. X* **12**, 021028 (2022).

- [52] T. Grünzweig, A. Hilliard, M. McGovern, and M. F. Andersen, *Near-deterministic preparation of a single atom in an optical microtrap*, *Nature Physics* **6**, 951 (2010).
- [53] M. O. Brown, T. Thiele, C. Kiehl, T.-W. Hsu, and C. A. Regal, *Gray-Molasses Optical-Tweezer Loading: Controlling Collisions for Scaling Atom-Array Assembly*, *Phys. Rev. X* **9**, 011057 (2019).
- [54] Y. Miroshnychenko, W. Alt, I. Dotsenko, L. Förster, M. Khudaverdyan, D. Meschede, D. Schrader, and A. Rauschenbeutel, *An atom-sorting machine*, *Nature* **442**, 151 (2006).
- [55] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, *An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays*, *Science* **354**, 1021 (2016).
- [56] K.-N. Schymik, V. Lienhard, D. Barredo, P. Scholl, H. Williams, A. Browaeys, and T. Lahaye, *Enhanced atom-by-atom assembly of arbitrary tweezer arrays*, *Phys. Rev. A* **102**, 063107 (2020).
- [57] D. Stuart and A. Kuhn, *Single-atom trapping and transport in DMD-controlled optical tweezers*, *New Journal of Physics* **20**, 023013 (2018).
- [58] W. Lee, H. Kim, and J. Ahn, *Three-dimensional rearrangement of single atoms using actively controlled optical microtraps*, *Opt. Express* **24**, 9816 (2016).
- [59] A. M. Kaufman and K.-K. Ni, *Quantum science with optical tweezer arrays of ultracold atoms and molecules*, *Nature Physics* **17**, 1324 (2021).
- [60] R. Mukherjee, J. Millen, R. Nath, M. P. A. Jones, and T. Pohl, *Many-body physics with alkaline-earth Rydberg lattices*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **44**, 184010 (2011).
- [61] V. Lienhard, S. de Léséleuc, D. Barredo, T. Lahaye, A. Browaeys, M. Schuler, L.-P. Henry, and A. M. Läuchli, *Observing the Space- and Time-Dependent*

- Growth of Correlations in Dynamically Tuned Synthetic Ising Models with Antiferromagnetic Interactions*, *Phys. Rev. X* **8**, 021070 (2018).
- [62] V. Lienhard, *Experimental quantum many-body physics with arrays of Rydberg atoms: From spin models to topological matter*, *Ph.D. thesis*, Institut d'Optique Graduate School (2019).
- [63] L. R. Liu, J. D. Hood, Y. Yu, J. T. Zhang, N. R. Hutzler, T. Rosenband, and K.-K. Ni, *Building one molecule from a reservoir of two atoms*, *Science* **360**, 900 (2018).
- [64] L. Anderegg, L. W. Cheuk, Y. Bao, S. Burchesky, W. Ketterle, K.-K. Ni, and J. M. Doyle, *An optical tweezer array of ultracold molecules*, *Science* **365**, 1156 (2019).
- [65] R. Ding, J. D. Whalen, S. K. Kanungo, T. C. Killian, F. B. Dunning, S. Yoshida, and J. Burgdörfer, *Spectroscopy of ^{87}Sr triplet Rydberg states*, *Phys. Rev. A* **98**, 042505 (2018).
- [66] L. Couturier, I. Nosske, F. Hu, C. Tan, C. Qiao, Y. H. Jiang, P. Chen, and M. Weidemüller, *Measurement of the strontium triplet Rydberg series by depletion spectroscopy of ultracold atoms*, *Phys. Rev. A* **99**, 022503 (2019).
- [67] C. L. Vaillant, M. P. A. Jones, and R. M. Potvliege, *Multichannel quantum defect theory of strontium bound Rydberg states*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **47**, 155001 (2014).
- [68] C. L. J. Vaillant, *Long-Range Interactions in One- and Two-Electron Rydberg Atoms*, *Ph.D. thesis*, Durham University (2014).
- [69] J.-P. Uzan, *The fundamental constants and their variation: observational and theoretical status*, *Rev. Mod. Phys.* **75**, 403 (2003).
- [70] J. C. Berengut, D. Budker, C. Delaunay, V. V. Flambaum, C. Frugiuele, E. Fuchs, C. Grojean, R. Harnik, R. Ozeri, G. Perez, and Y. Soreq, *Probing New Long-Range Interactions by Isotope Shift Spectroscopy*, *Phys. Rev. Lett.* **120**, 091801 (2018).

- [71] W. Bowden, *Precision Frequency Metrology with Optical Lattice Clocks*, [DPhil thesis](#), University of Oxford (2019).
- [72] R. K. Hanley, *Creation of a strontium microtrap: Towards a spin-squeezed atomic clock*, [Ph.D. thesis](#), Durham University (2018).
- [73] D. Boddy, *First observations of Rydberg blockade in a frozen gas of divalent atoms*, [Ph.D. thesis](#), Durham University (2014).
- [74] Y. N. Martinez, *Studies of the $5s^2 \ 1 \ S_0$ - $5s5p \ 3P_1$ transition in atomic strontium*, [Master's thesis](#), Rice University (2009).
- [75] A. V. Taichenachev, V. I. Yudin, C. W. Oates, C. W. Hoyt, Z. W. Barber, and L. Hollberg, *Magnetic Field-Induced Spectroscopy of Forbidden Optical Transitions with Application to Lattice-Based Optical Atomic Clocks*, [Phys. Rev. Lett.](#) **96**, 083001 (2006).
- [76] J. Millen, *A cold strontium Rydberg gas*, [Ph.D. thesis](#), Durham University (2011).
- [77] G. Lohead, *Excited state spatial distributions in a cold strontium gas*, [Ph.D. thesis](#), Durham University (2012).
- [78] D. P. Sadler, *Many-body interactions in a dissipative frozen strontium Rydberg gas*, [Ph.D. thesis](#), Durham University (2016).
- [79] A. D. Bounds, *A Rydberg-dressed Magneto Optical Trap*, [Ph.D. thesis](#), Durham University (2018).
- [80] N. C. Jackson, *Rydberg spectroscopy and dressing in an ultracold strontium gas*, [Ph.D. thesis](#), Durham University (2018).
- [81] R. Grimm, M. Weidemüller, and Y. B. Ovchinnikov, *Optical Dipole Traps for Neutral Atoms*, [Advances In Atomic, Molecular, and Optical Physics](#), **42**, 95 (2000).
- [82] A. Ashkin, *Optical trapping and manipulation of neutral particles using lasers*, [Proceedings of the National Academy of Sciences](#) **94**, 4853 (1997).

- [83] R. Hobson, W. Bowden, A. Vianello, A. Silva, C. F. A. Baynham, H. S. Margolis, P. E. G. Baird, P. Gill, and I. R. Hill, *A strontium optical lattice clock with 1×10^{-17} uncertainty and measurement of its absolute frequency*, *Metrologia* **57**, 065026 (2020).
- [84] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. D. Opping, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye, *A Fermi-degenerate three-dimensional optical lattice clock*, *Science* **358**, 90 (2017).
- [85] D. Barredo, V. Lienhard, P. Scholl, S. de Léséleuc, T. Boulier, A. Browaeys, and T. Lahaye, *Three-Dimensional Trapping of Individual Rydberg Atoms in Ponderomotive Bottle Beam Traps*, *Phys. Rev. Lett.* **124**, 023201 (2020).
- [86] T. Akatsuka, M. Takamoto, and H. Katori, *Optical lattice clocks with non-interacting bosons and fermions*, *Nature Physics* **4**, 954 (2008).
- [87] T. Akatsuka, M. Takamoto, and H. Katori, *Three-dimensional optical lattice clock with bosonic ^{88}Sr atoms*, *Phys. Rev. A* **81**, 023402 (2010).
- [88] T. Ido, Y. Isoya, and H. Katori, *Optical-dipole trapping of Sr atoms at a high phase-space density*, *Phys. Rev. A* **61**, 061403 (2000).
- [89] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys, *Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms*, *Nature* **595**, 233 (2021).
- [90] J. Beugnon, C. Tuchendler, H. Marion, A. Gaëtan, Y. Miroshnychenko, Y. R. P. Sortais, A. M. Lance, M. P. A. Jones, G. Messin, A. Browaeys, and P. Grangier, *Two-dimensional transport and transfer of a single atomic qubit in optical tweezers*, *Nature Physics* **3**, 696 (2007).
- [91] I. S. Madjarov, *Entangling, controlling, and detecting individual strontium atoms in optical tweezer arrays*, *Ph.D. thesis*, Caltech (2021).

- [92] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres, *High-fidelity entanglement and detection of alkaline-earth Rydberg atoms*, *Nature Physics* **16**, 857 (2020).
- [93] M. Takamoto, H. Katori, S. I. Marmo, V. D. Ovsiannikov, and V. G. Pal'chikov, *Prospects for Optical Clocks with a Blue-Detuned Lattice*, *Phys. Rev. Lett.* **102**, 063002 (2009).
- [94] I. R. Hill, *Development of an Apparatus for a Strontium Optical Lattice Optical Frequency Standard*, *Ph.D. thesis*, Imperial College London (2012).
- [95] M. Barbiero, D. Calonico, F. Levi, and M. G. Tarallo, *Optically Loaded Strontium Lattice Clock With a Single Multi-Wavelength Reference Cavity*, *IEEE Transactions on Instrumentation and Measurement* **71**, 1 (2022).
- [96] L. Essen and J. V. L. Parry, *An Atomic Standard of Frequency and Time Interval: A Cæsium Resonator*, *Nature* **176**, 280 (1955).
- [97] R. Li, K. Gibble, and K. Szymaniec, *Improved accuracy of the NPL-CsF2 primary frequency standard: evaluation of distributed cavity phase and microwave lensing frequency shifts*, *Metrologia* **48**, 283 (2011).
- [98] F. Levi, D. Calonico, C. E. Calosso, A. Godone, S. Micalizio, and G. A. Costanzo, *Accuracy evaluation of ITCsF2: a nitrogen cooled caesium fountain*, *Metrologia* **51**, 270 (2014).
- [99] J. Tennant, *The Jurassic/Cretaceous boundary: a hidden mass extinction in tetrapods?*, *Ph.D. thesis*, Imperial College London (2016).
- [100] G. E. Marti, R. B. Hutson, A. Goban, S. L. Campbell, N. Poli, and J. Ye, *Imaging Optical Frequencies with 100 μ Hz Precision and 1.1 μ m Resolution*, *Phys. Rev. Lett.* **120**, 103201 (2018).
- [101] X. Zheng, J. Dolde, V. Lochab, B. N. Merriman, H. Li, and S. Kolkowitz, *Differential clock comparisons with a multiplexed optical lattice clock*, *Nature* **602**, 425 (2022).

- [102] S. Origlia, M. S. Pramod, S. Schiller, Y. Singh, K. Bongs, R. Schwarz, A. Al-Masoudi, S. Dörscher, S. Herbers, S. Häfner, U. Sterr, and C. Lisdat, *Towards an optical clock for space: Compact, high-performance optical lattice clock based on bosonic atoms*, *Phys. Rev. A* **98**, 053443 (2018).
- [103] R. Schwarz, S. Dörscher, A. Al-Masoudi, E. Benkler, T. Legero, U. Sterr, S. Weyers, J. Rahm, B. Lipphardt, and C. Lisdat, *Long term measurement of the ^{87}Sr clock frequency at the limit of primary Cs clocks*, *Phys. Rev. Research* **2**, 033242 (2020).
- [104] K. Yamanaka, N. Ohmae, I. Ushijima, M. Takamoto, and H. Katori, *Frequency Ratio of ^{199}Hg and ^{87}Sr Optical Lattice Clocks beyond the SI Limit*, *Phys. Rev. Lett.* **114**, 230801 (2015).
- [105] A. A. Madej, P. Dubé, Z. Zhou, J. E. Bernard, and M. Gertszov, *$^{88}\text{Sr}^+$ 445-THz Single-Ion Reference at the 10^{-17} Level via Control and Cancellation of Systematic Uncertainties and Its Measurement against the SI Second*, *Phys. Rev. Lett.* **109**, 203002 (2012).
- [106] P. Dubé, A. A. Madej, M. Tibbo, and J. E. Bernard, *High-Accuracy Measurement of the Differential Scalar Polarizability of a $^{88}\text{Sr}^+$ Clock Using the Time-Dilation Effect*, *Phys. Rev. Lett.* **112**, 173002 (2014).
- [107] R. M. Godun, P. B. R. Nisbet-Jones, J. M. Jones, S. A. King, L. A. M. Johnson, H. S. Margolis, K. Szymaniec, S. N. Lea, K. Bongs, and P. Gill, *Frequency Ratio of Two Optical Clock Transitions in $^{171}\text{Yb}^+$ and Constraints on the Time Variation of Fundamental Constants*, *Phys. Rev. Lett.* **113**, 210801 (2014).
- [108] N. Huntemann, C. Sanner, B. Lipphardt, C. Tamm, and E. Peik, *Single-Ion Atomic Clock with 3×10^{-18} Systematic Uncertainty*, *Phys. Rev. Lett.* **116**, 063001 (2016).
- [109] T. Rosenband, D. B. Hume, P. O. Schmidt, C. W. Chou, A. Brusch, L. Lorini, W. H. Oskay, R. E. Drullinger, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, W. C. Swann, N. R. Newbury, W. M. Itano, D. J. Wineland, and

- J. C. Bergquist, *Frequency Ratio of Al^+ and Hg^+ Single-Ion Optical Clocks; Metrology at the 17th Decimal Place*, *Science* **319**, 1808 (2008).
- [110] Y. Huang, H. Guan, P. Liu, W. Bian, L. Ma, K. Liang, T. Li, and K. Gao, *Frequency Comparison of Two $^{40}Ca^+$ Optical Clocks with an Uncertainty at the 10^{-17} Level*, *Phys. Rev. Lett.* **116**, 013001 (2016).
- [111] N. Ohtsubo, Y. Li, K. Matsubara, T. Ido, and K. Hayasaka, *Frequency measurement of the clock transition of an indium ion sympathetically-cooled in a linear trap*, *Opt. Express* **25**, 11725 (2017).
- [112] S. A. King, L. J. Spieß, P. Micke, A. Wilzewski, T. Leopold, E. Benkler, R. Lange, N. Huntemann, A. Surzhykov, V. A. Yerokhin, J. R. C. López-Urrutia, and P. O. Schmidt, *An Optical Atomic Clock Based on a Highly Charged Ion*, preprint [arXiv:2007.01964](https://arxiv.org/abs/2007.01964) (2022).
- [113] L. von der Wense, B. Seiferle, M. Laatiaoui, J. B. Neumayr, H.-J. Maier, H.-F. Wirth, C. Mokry, J. Runke, K. Eberhardt, C. E. Düllmann, N. G. Trautmann, and P. G. Thirolf, *Direct detection of the ^{229}Th nuclear clock transition*, *Nature* **533**, 47 (2016).
- [114] E. Peik, T. Schumm, M. S. Safronova, A. Pálffy, J. Weitenberg, and P. G. Thirolf, *Nuclear clocks for testing fundamental physics*, *Quantum Science and Technology* **6**, 034002 (2021).
- [115] L. A. Downes, A. R. MacKellar, D. J. Whiting, C. Bourgenot, C. S. Adams, and K. J. Weatherill, *Full-Field Terahertz Imaging at KiloHertz Frame Rates Using Atomic Vapor*, *Phys. Rev. X* **10**, 011027 (2020).
- [116] L. A. Downes, *A High-speed THz Imaging System based on THz-to-optical Conversion in Atomic Vapour*, *Ph.D. thesis*, Durham University (2020).
- [117] J. Ye and S. T. Cundiff, *Femtosecond Optical Frequency Comb: Principle, Operation, and Applications* (Springer, 2005).
- [118] T. Fortier and E. Baumann, *20 years of developments in optical frequency comb technology and applications*, *Communications Physics* **2**, 153 (2019).

- [119] M. A. Jankauski, *Measuring the frequency response of the honeybee thorax*, *Bioinspiration & Biomimetics* **15**, 046002 (2020).
- [120] J.-S. Chen, *Rydberg excitation dynamics and correlations in arbitrary 2D arrays of single atoms*, *Ph.D. thesis*, University of Colorado Boulder (2017).
- [121] J.-S. Chen, S. M. Brewer, C. W. Chou, D. J. Wineland, D. R. Leibbrandt, and D. B. Hume, *Sympathetic Ground State Cooling and Time-Dilation Shifts in an $^{27}\text{Al}^+$ Optical Clock*, *Phys. Rev. Lett.* **118**, 053002 (2017).
- [122] N. Herschbach, K. Pyka, J. Keller, and T. E. Mehlstäubler, *Linear Paul trap design for an optical clock with Coulomb crystals*, *Applied Physics B* **107**, 891 (2012).
- [123] M. Takamoto, F.-L. Hong, R. Higashi, Y. Fujii, M. Imae, and H. Katori, *Improved Frequency Measurement of a One-Dimensional Optical Lattice Clock with a Spin-Polarized Fermionic ^{87}Sr Isotope*, *Journal of the Physical Society of Japan* **75**, 104302 (2006).
- [124] R. Santra, K. V. Christ, and C. H. Greene, *Properties of metastable alkaline-earth-metal atoms calculated using an accurate effective core potential*, *Phys. Rev. A* **69**, 042510 (2004).
- [125] M. Boyd, *High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time*, *Ph.D. thesis*, University of Colorado Boulder (2007).
- [126] J. A. Muniz, D. J. Young, J. R. K. Cline, and J. K. Thompson, *Cavity-QED measurements of the ^{87}Sr millihertz optical clock transition and determination of its natural linewidth*, *Phys. Rev. Research* **3**, 023152 (2021).
- [127] Z. W. Barber, C. W. Hoyt, C. W. Oates, L. Hollberg, A. V. Taichenachev, and V. I. Yudin, *Direct Excitation of the Forbidden Clock Transition in Neutral ^{174}Yb Atoms Confined to an Optical Lattice*, *Phys. Rev. Lett.* **96**, 083002 (2006).

- [128] Z. Barber, *Ytterbium Optical Lattice Clock*, **Ph.D. thesis**, University of Colorado Boulder (2007).
- [129] X. Baillard, M. Fouché, R. L. Targat, P. G. Westergaard, A. Lecallier, Y. L. Coq, G. D. Rovera, S. Bize, and P. Lemonde, *Accuracy evaluation of an optical lattice clock with bosonic atoms*, **Opt. Lett.** **32**, 1812 (2007).
- [130] M. Bober, P. Morzyński, A. Cygan, D. Lisak, P. Masłowski, M. Prymaczek, P. Wcisło, P. Ablewski, M. Piwiński, S. Wójtewicz, K. Bielska, D. Bartoszek-Bober, R. S. Trawiński, M. Zawada, R. Ciuryło, J. Zachorowski, M. Piotrowski, W. Gawlik, F. Ozimek, and C. Radzewicz, *Strontium optical lattice clocks for practical realization of the metre and secondary representation of the second*, **Measurement Science and Technology** **26**, 075201 (2015).
- [131] J. P. Covey, I. S. Madjarov, A. Cooper, and M. Endres, *2000-Times Repeated Imaging of Strontium Atoms in Clock-Magic Tweezer Arrays*, **Phys. Rev. Lett.** **122**, 173201 (2019).
- [132] E. Peik, T. Schneider, and C. Tamm, *Laser frequency stabilization to a single ion*, **Journal of Physics B: Atomic, Molecular and Optical Physics** **39**, 145 (2005).
- [133] L. A. P. Gallagher, *Optical and microwave spectroscopy of Rydberg excitons in Cu_2O* , **Ph.D. thesis**, Durham University (2022).
- [134] W. Bowden, R. Hobson, P. Huillery, P. Gill, M. P. A. Jones, and I. R. Hill, *Rydberg electrometry for optical lattice clocks*, **Phys. Rev. A** **96**, 023419 (2017).
- [135] H. Weimer, M. Müller, I. Lesanovsky, P. Zoller, and H. P. Büchler, *A Rydberg quantum simulator*, **Nature Physics** **6**, 382 (2010).
- [136] M. Saffman, *Quantum computing with atomic qubits and Rydberg interactions: progress and challenges*, **Journal of Physics B: Atomic, Molecular and Optical Physics** **49**, 202001 (2016).

- [137] J. R. Rydberg, *XXXIV. On the structure of the line-spectra of the chemical elements*, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **29**, 331 (1890).
- [138] N. Bohr, *I. On the constitution of atoms and molecules*, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **26**, 1 (1913).
- [139] T. F. Gallagher, *Rydberg Atoms*, Cambridge Monographs on Atomic, Molecular and Chemical Physics (Cambridge University Press, 1994).
- [140] C. L. Vaillant, M. P. A. Jones, and R. M. Potvliege, *Long-range Rydberg–Rydberg interactions in calcium, strontium and ytterbium*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **45**, 135004 (2012).
- [141] J. Millen, G. Lothead, G. R. Corbett, R. M. Potvliege, and M. P. A. Jones, *Spectroscopy of a cold strontium Rydberg gas*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **44**, 184001 (2011).
- [142] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys, *Entanglement of Two Individual Neutral Atoms Using Rydberg Blockade*, *Phys. Rev. Lett.* **104**, 010502 (2010).
- [143] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *High-Fidelity Control and Entanglement of Rydberg-Atom Qubits*, *Phys. Rev. Lett.* **121**, 123603 (2018).
- [144] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Generation and manipulation of Schrödinger cat states in Rydberg atom arrays*, *Science* **365**, 570 (2019).
- [145] M. D. Lukin, M. Fleischhauer, R. Cote, L. M. Duan, D. Jaksch, J. I. Cirac, and

- P. Zoller, *Dipole Blockade and Quantum Information Processing in Mesoscopic Atomic Ensembles*, *Phys. Rev. Lett.* **87**, 037901 (2001).
- [146] D. Tong, S. M. Farooqi, J. Stanojevic, S. Krishnan, Y. P. Zhang, R. Côté, E. E. Eyler, and P. L. Gould, *Local Blockade of Rydberg Excitation in an Ultracold Gas*, *Phys. Rev. Lett.* **93**, 063001 (2004).
- [147] J. E. Johnson and S. L. Rolston, *Interactions between Rydberg-dressed atoms*, *Phys. Rev. A* **82**, 033412 (2010).
- [148] Y.-Y. Jau, A. M. Hankin, T. Keating, I. H. Deutsch, and G. W. Biedermann, *Entangling atomic spins with a Rydberg-dressed spin-flip blockade*, *Nature Physics* **12**, 71 (2016).
- [149] A. D. Bounds, N. C. Jackson, R. K. Hanley, R. Faoro, E. M. Bridge, P. Huillery, and M. P. A. Jones, *Rydberg-Dressed Magneto-optical Trap*, *Phys. Rev. Lett.* **120**, 183401 (2018).
- [150] J. Zeiher, R. van Bijnen, P. Schauß, S. Hild, J.-y. Choi, T. Pohl, I. Bloch, and C. Gross, *Many-body interferometry of a Rydberg-dressed spin lattice*, *Nature Physics* **12**, 1095 (2016).
- [151] R. K. Hanley, P. Huillery, N. C. Keegan, A. D. Bounds, D. Boddy, R. Faoro, and M. P. A. Jones, *Quantitative simulation of a magneto-optical trap operating near the photon recoil limit*, *Journal of Modern Optics* **65**, 667 (2018).
- [152] E. M. Bridge, J. Millen, C. S. Adams, and M. P. A. Jones, *A vapor cell based on dispensers for laser spectroscopy*, *Rev. Sci. Instrum* **80**, 013101 (2009).
- [153] T. Uehara, K. Tsuji, K. Hagiwara, and N. Onodera, *Optical beat-note frequency stabilization between two lasers using a radio frequency interferometer in the gigahertz frequency band*, *Optical Engineering* **53**, 1 (2014).
- [154] R. W. Fox, C. W. Oates, and L. W. Hollberg, *Stabilizing diode lasers to high-finesse cavities*, *Elsevier Science (USA)* **40** (2003).

- [155] E. M. Bridge, N. C. Keegan, A. D. Bounds, D. Boddy, D. P. Sadler, and M. P. A. Jones, *Tunable cw UV laser with <35 kHz absolute frequency instability for precision spectroscopy of Sr Rydberg states*, *Opt. Express* **24**, 2281 (2016).
- [156] J. I. Thorpe, K. Numata, and J. Livas, *Laser frequency stabilization and control through offset sideband locking to optical cavities*, *Optics express* **16**, 15980—15990 (2008).
- [157] L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall, *Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path*, *Opt. Lett.* **19**, 1777 (1994).
- [158] S. M. Foreman, A. D. Ludlow, M. H. G. de Miranda, J. E. Stalnaker, S. A. Diddams, and J. Ye, *Coherent Optical Phase Transfer over a 32-km Fiber with 1 s Instability at 10^{-17}* , *Phys. Rev. Lett.* **99**, 153601 (2007).
- [159] W. Bowden, R. Hobson, I. Hill, A. Vianello, M. Schioppo, A. Silva, H. S. Margolis, P. E. G. Baird, and P. Gill, *A pyramid MOT with integrated optical cavities as a cold atom platform for an optical lattice clock*, *Sci Rep* **9** (2019).
- [160] S. Stellmer, *Degenerate quantum gases of strontium*, *Ph.D. thesis*, University of Innsbruck (2013).
- [161] R. Kliese, N. Hoghooghi, T. Puppe, F. Rohde, A. Sell, A. Zach, P. Leisching, W. Kaenders, N. C. Keegan, A. D. Bounds, E. M. Bridge, J. Leonard, C. S. Adams, S. L. Cornish, and M. P. A. Jones, *Difference-frequency combs in cold atom physics*, *The European Physical Journal Special Topics* **225**, 2775 (2016).
- [162] L. Beguin, *Measurement of the van der Waals interaction between two Rydberg atoms*, *Ph.D. thesis*, Institut d'Optique Graduate School (2014).
- [163] J. Lodewyck, M. Zawada, L. Lorini, M. Gurov, and P. Lemonde, *Observation and cancellation of a perturbing dc stark shift in strontium optical lattice clocks*,

- IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **59**, 411 (2012).
- [164] A. Browaeys, D. Barredo, and T. Lahaye, *Experimental investigations of dipole–dipole interactions between a few Rydberg atoms*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **49**, 152001 (2016).
- [165] J. M. Obrecht, R. J. Wild, and E. A. Cornell, *Measuring electric fields from surface contaminants with neutral atoms*, *Phys. Rev. A* **75**, 062903 (2007).
- [166] M. Meucci, E. Mariotti, P. Bicchi, C. Marinelli, and L. Moi, *Light-Induced Atom Desorption*, *Europhysics Letters (EPL)* **25**, 639 (1994).
- [167] A. M. Hankin, Y.-Y. Jau, L. P. Parazzoli, C. W. Chou, D. J. Armstrong, A. J. Landahl, and G. W. Biedermann, *Two-atom Rydberg blockade using direct $6S$ to nP excitation*, *Phys. Rev. A* **89**, 033416 (2014).
- [168] A. Browaeys, D. Barredo, and T. Lahaye, *Experimental investigations of dipole–dipole interactions between a few Rydberg atoms*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **49**, 152001 (2016).
- [169] D. Bronzi, F. Villa, S. Tisa, A. Tosi, F. Zappa, D. Durini, S. Weyers, and W. Brockherde, *100 000 Frames/s 64×32 Single-Photon Detector Array for 2-D Imaging and 3-D Ranging*, *IEEE Journal of Selected Topics in Quantum Electronics* **20**, 354 (2014).
- [170] R. W. Gerchberg and W. O. Saxton, *A Practical Algorithm for the Determination of Phase from Image and Diffraction Plane Pictures*, *Optik* **35**, 237 (1972).
- [171] T. Ido and H. Katori, *Recoil-Free Spectroscopy of Neutral Sr Atoms in the Lamb-Dicke Regime*, *Phys. Rev. Lett.* **91**, 053001 (2003).
- [172] N. Schlosser, *Étude et réalisation de micro-pièges dipolaires optiques pour atomes neutres*, *Ph.D. thesis*, Université Paris XI (2001).

- [173] D. Bronzi, F. Villa, S. Tisa, A. Tosi, F. Zappa, D. Durini, S. Weyers, and W. Brockherde, *100 000 Frames/s 64×32 Single-Photon Detector Array for 2-D Imaging and 3-D Ranging*, *IEEE Journal of Selected Topics in Quantum Electronics* **20**, 355 (2014).
- [174] M. Walker, *Using Superpositions of Laguerre-Gauss Modes to Generate Tighter Confinement of Ultracold Atoms in Optical Tweezers*, Master's thesis, Durham University (2021).
- [175] J.-B. Béguin, J. Laurat, X. Luan, A. P. Burgers, Z. Qin, and H. J. Kimble, *Reduced volume and reflection for bright optical tweezers with radial Laguerre–Gauss beams*, *Proceedings of the National Academy of Sciences* **117**, 26109 (2020).
- [176] R. V. Brooks, S. Spence, A. Guttridge, A. Alampounti, A. Rakonjac, L. A. McArd, J. M. Hutson, and S. L. Cornish, *Preparation of one 87Rb and one 133Cs atom in a single optical tweezer*, *New Journal of Physics* **23**, 065002 (2021).
- [177] A. W. Young, W. J. Eckner, W. R. Milner, D. Kedar, M. A. Norcia, E. Oelker, N. Schine, J. Ye, and A. M. Kaufman, *Half-minute-scale atomic coherence and high relative stability in a tweezer clock*, *Nature* **588**, 408 (2020).
- [178] A. W. Young, W. J. Eckner, N. Schine, A. M. Childs, and A. M. Kaufman, *Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice*, *Science* **377**, 885 (2022).
- [179] S. Subhankar, A. Restelli, Y. Wang, S. L. Rolston, and J. V. Porto, *Microcontroller based scanning transfer cavity lock for long-term laser frequency stabilization*, *Review of Scientific Instruments* **90**, 043115 (2019).
- [180] Hamamatsu, *LCOS-SLM X10468 Series Operation Manual*, (2018).
- [181] M. J. Piotrowicz, M. Lichtman, K. Maller, G. Li, S. Zhang, L. Isenhower, and M. Saffman, *Two-dimensional lattice of blue-detuned atom traps using a projected Gaussian beam array*, *Phys. Rev. A* **88**, 013420 (2013).

- [182] R. Dumke, M. Volk, T. Mütter, F. B. J. Buchkremer, G. Birkl, and W. Ertmer, *Micro-optical Realization of Arrays of Selectively Addressable Dipole Traps: A Scalable Configuration for Quantum Computation with Atomic Qubits*, *Phys. Rev. Lett.* **89**, 097903 (2002).
- [183] M. Schlosser, S. Tichelmann, J. Kruse, and G. Birkl, *Scalable architecture for quantum information processing with atoms in optical micro-structures*, *Quantum Information Processing* **10**, 907 (2011).
- [184] D. Ohl de Mello, D. Schäffner, J. Werkmann, T. Preuschoff, L. Kohfahl, M. Schlosser, and G. Birkl, *Defect-Free Assembly of 2D Clusters of More Than 100 Single-Atom Quantum Systems*, *Phys. Rev. Lett.* **122**, 203601 (2019).
- [185] D. Schäffner, T. Preuschoff, S. Ristok, L. Brozio, M. Schlosser, H. Giessen, and G. Birkl, *Arrays of individually controllable optical tweezers based on 3D-printed microlens arrays*, *Opt. Express* **28**, 8640 (2020).
- [186] S. Bergamini, B. Darquié, M. Jones, L. Jacubowicz, A. Browaeys, and P. Grangier, *Holographic generation of microtrap arrays for single atoms by use of a programmable phase modulator*, *J. Opt. Soc. Am. B* **21**, 1889 (2004).
- [187] Y. Wang, S. Shevate, T. M. Wintermantel, M. Morgado, G. Lohead, and S. Whitlock, *Preparation of hundreds of microscopic atomic ensembles in optical tweezer arrays*, *npj Quantum Information* **6**, 54 (2020).
- [188] T. Zelevinsky, M. M. Boyd, A. D. Ludlow, T. Ido, J. Ye, R. Ciuryło, P. Naidon, and P. S. Julienne, *Narrow Line Photoassociation in an Optical Lattice*, *Phys. Rev. Lett.* **96**, 203201 (2006).
- [189] G. Reinaudi, C. B. Osborn, M. McDonald, S. Kotochigova, and T. Zelevinsky, *Optical Production of Stable Ultracold $^{88}\text{Sr}_2$ Molecules*, *Phys. Rev. Lett.* **109**, 115303 (2012).
- [190] T. Middelmann, S. Falke, C. Lisdat, and U. Sterr, *High Accuracy Correction of Blackbody Radiation Shift in an Optical Lattice Clock*, *Phys. Rev. Lett.* **109**, 263004 (2012).

- [191] M. P. A. Jones, R. M. Potvliege, and M. Spannowsky, *Probing new physics using Rydberg states of atomic hydrogen*, *Phys. Rev. Research* **2**, 013244 (2020).
- [192] B. H. Bransden and C. J. Joachain, *Physics of Atoms and Molecules* (Longman Group Limited, 1983).
- [193] H. Engler, T. Weber, M. Mudrich, R. Grimm, and M. Weidemüller, *Very long storage times and evaporative cooling of cesium atoms in a quasielectrostatic dipole trap*, *Phys. Rev. A* **62**, 031402 (2000).
- [194] S. Panezai, D. Wang, J. Zhao, Y. Wang, L. Rong, and S. Ma, *Study of oblique incidence characterization of parallel aligned liquid crystal on silicon*, *Optical Engineering* **54**, 1 (2015).
- [195] J. L. M. Fuentes, E. J. Fernández, P. M. Prieto, and P. Artal, *Interferometric method for phase calibration in liquid crystal spatial light modulators using a self-generated diffraction-grating*, *Opt. Express* **24**, 14159 (2016).
- [196] Z. Zhao, Z. Xiao, Y. Zhuang, H. Zhang, and H. Zhao, *An interferometric method for local phase modulation calibration of LC-SLM using self-generated phase grating*, *Review of Scientific Instruments* **89**, 083116 (2018).
- [197] R. M. W. van Bijnen, *Quantum engineering with ultracold atoms*, *Ph.D. thesis*, Technische Universiteit Eindhoven (2013).
- [198] Jackson-Labs, *Fury User Manual*, (2020).
- [199] W. Riley and D. A. Howe, *Handbook of Frequency Stability Analysis* (NIST Special Publication 1065, 2008).
- [200] H. Ludvigsen, M. Tossavainen, and M. Kaivola, *Laser linewidth measurements using self-homodyne detection with short delay*, *Optics Communications* **155**, 180 (1998).
- [201] A. D. Ludlow, X. Huang, M. Notcutt, T. Zanon-Willette, S. M. Foreman, M. M. Boyd, S. Blatt, and J. Ye, *Compact, thermal-noise-limited optical cavity for diode laser stabilization at 1×10^{-15}* , *Opt. Lett.* **32**, 641 (2007).

- [202] S. A. Webster, M. Oxborrow, S. Pugla, J. Millo, and P. Gill, *Thermal-noise-limited optical cavity*, *Phys. Rev. A* **77**, 033847 (2008).
- [203] A. W. Wiederkehr, M. Motsch, S. D. Hogan, M. Andrist, H. Schmutz, B. Lambillotte, J. A. Agner, and F. Merkt, *Multistage Zeeman deceleration of metastable neon*, *The Journal of Chemical Physics* **135**, 214202 (2011).
- [204] L. A. McArd, *A Travelling Wave Zeeman Decelerator For Atoms and Molecules*, *Ph.D. thesis*, Durham University (2018).
- [205] A. D. West, *Interfacing ultracold atoms with nanomagnetic domain walls*, *Ph.D. thesis*, Durham University (2012).
- [206] C. Lisdat, J. S. R. V. Winfred, T. Middelmann, F. Riehle, and U. Sterr, *Collisional Losses, Decoherence, and Frequency Shifts in Optical Lattice Clocks with Bosons*, *Phys. Rev. Lett.* **103**, 090801 (2009).
- [207] N. Poli, M. G. Tarallo, M. Schioppo, C. W. Oates, and G. M. Tino, *A simplified optical lattice clock*, *Applied Physics B* **97**, 27 (2009).
- [208] V. I. Yudin, A. V. Taichenachev, C. W. Oates, Z. W. Barber, N. D. Lemke, A. D. Ludlow, U. Sterr, C. Lisdat, and F. Riehle, *Hyper-Ramsey spectroscopy of optical clock transitions*, *Phys. Rev. A* **82**, 011804 (2010).
- [209] N. Huntemann, B. Lipphardt, M. Okhapkin, C. Tamm, E. Peik, A. V. Taichenachev, and V. I. Yudin, *Generalized Ramsey Excitation Scheme with Suppressed Light Shift*, *Phys. Rev. Lett.* **109**, 213002 (2012).
- [210] R. Hobson, W. Bowden, S. A. King, P. E. G. Baird, I. R. Hill, and P. Gill, *Modified hyper-Ramsey methods for the elimination of probe shifts in optical clocks*, *Phys. Rev. A* **93**, 010501 (2016).
- [211] T. Zanon-Willette, E. de Clercq, and E. Arimondo, *Probe light-shift elimination in generalized hyper-Ramsey quantum clocks*, *Phys. Rev. A* **93**, 042506 (2016).

- [212] C. Shi, J.-L. Robyr, U. Eismann, M. Zawada, L. Lorini, R. Le Targat, and J. Lodewyck, *Polarizabilities of the ^{87}Sr clock transition*, *Phys. Rev. A* **92**, 012516 (2015).
- [213] T. Takano, R. Mizushima, and H. Katori, *Precise determination of the isotope shift of 88Sr – 87Sr optical lattice clock by sharing perturbations*, *Applied Physics Express* **10**, 072801 (2017).
- [214] A. Bruschi, R. Le Targat, X. Baillard, M. Fouché, and P. Lemonde, *Hyperpolarizability Effects in a Sr Optical Lattice Clock*, *Phys. Rev. Lett.* **96**, 103003 (2006).
- [215] P. G. Westergaard, J. Lodewyck, L. Lorini, A. Lecallier, E. A. Burt, M. Zawada, J. Millo, and P. Lemonde, *Lattice-Induced Frequency Shifts in Sr Optical Lattice Clocks at the 10^{-17} Level*, *Phys. Rev. Lett.* **106**, 210801 (2011).
- [216] R. Le Targat, L. Lorini, Y. Le Coq, M. Zawada, J. Guéna, M. Abgrall, M. Gurov, P. Rosenbusch, D. G. Rovera, B. Nagórny, R. Gartman, P. G. Westergaard, M. E. Tobar, M. Lours, G. Santarelli, A. Clairon, S. Bize, P. Laurent, P. Lemonde, and J. Lodewyck, *Experimental realization of an optical second with strontium lattice clocks*, *Nature Communications* **4**, 2109 (2013).

Appendix A

Experimental methods

This appendix provides additional detail on equipment described in chapter 3.

A.1 461 nm beat lock components

A beat lock for two 461 nm lasers is described in section 3.2.1, as shown in figure 3.8. The list of components used are:

- Laser A: Toptica TA SHG laser
- Laser B: MOGLabs laser / Toptica DL Pro laser
- RF Signal: Agilent E4421B
- Photodetector: Thorlabs DET025AL/M
- Amplifier: ZKL-2+ followed by two ZX60-3018-G-S+
- Coupler: ZFDC-20-33-S+
- Mixer 1: ZFM-150-S+
- Low pass filter 1: BLP-30+
- Splitter: ZFRSC-42-S+
- Delay line length: ~ 4.7 m

- Mixer 2: ZLW-1-1
- Low Pass Filter 2: BLP-5+
- Frequency counter: Tektronix FCA3100
- Spectrum Analyser: Agilent CSA Spectrum Analyser N1996A

Appendix B

Number-resolved imaging

This appendix details information supporting the methods used in chapter 4.

B.1 Sisyphus cooling optimisation

We implement Sisyphus cooling on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1\ |m_j| = 0$ transition as described in section 4.2.2 and shown in figure 4.3(a). To test the effectiveness of Sisyphus cooling we load 532 nm tweezers with many atoms,¹ and take images of many SPAD frames with both the 461 nm probe beam and Sisyphus cooling on. From this we yield a lifetime measurement, with two apparent ‘loss’ mechanisms. Firstly, there is the genuine atomic loss due to de-pumping into the $5s4d\ ^1D_2$ state, which is the limiting lifetime. Secondly, there is an apparent loss due to heated atoms scattering less photons as the resonance frequency shifts, reducing the count rate and making it appear like there is loss from the tweezer. Varying the Sisyphus frequency we can then find the appropriate location to set the Sisyphus cap where heating effects are minimised.

We measure the lifetime as a function of Sisyphus frequency for three different trap depths, 1.4, 3.0 and 7.5 mK, to better understand the mechanism. The 461 nm probe beam is red-detuned 20 MHz from the resonance condition in the centre of the

¹Many atoms are loaded by setting the overlap time with the NB MOT $t_{\text{load}} = 50$ ms, as discussed in section 3.1.4.

tweezer in each case. The power of the probe beam is the same for each experiment, set to give a scattering rate of $0.14 \mu\text{s}^{-1}$. The fitted decay rates in atomic fluorescence are shown in figure B.1(a,b,c) as a function of Sisyphus frequency, for each of the three trap depths 1.4, 3.0 and 7.5 mK, respectively. These plots include a measured decay rate in the absence of cooling, and the expected decay rate if purely limited by decay to the $5s4d \ ^1D_2$ loss channel.

For each trap depth we see two features where the Sisyphus light decreases the apparent decay rate compared to without using any cooling. Inspecting figures B.1(a,b,c), it appears that there are two features where the cooling is effective; a narrow feature at low frequency and a broad feature at high frequency. These features are separated by a feature where the 689 nm light actually worsens the lifetime, around 2.5 MHz, 6 MHz and 14 MHz in (a,b,c) respectively. We attribute the low (high) frequency cooling feature to be Sisyphus cooling on the $5s5p \ ^3P_1$ $|m_j| = 0$ (1) states respectively. We attribute the loss mechanism between the two features to be where the Sisyphus cap is set too close to the bottom of the trap for the $5s5p \ ^3P_1$ $|m_j| = 0$ transition, such that the Sisyphus light instead heats the atoms further. This occurs as the Sisyphus mechanism here is repulsive (due to the polarisability of the states at $\lambda_{\text{trap}} = 532 \text{ nm}$), so if the atoms are hotter than the Sisyphus cap due to the cap being set too low, then they will be repelled to higher temperatures.

As figures B.1(a,b,c) use the same 461 nm probe scattering rate, we see that the Sisyphus cooling also appears to be more effective in deeper traps, where the decay from the trap more closely approaches the $5s5p \ ^1P_1 \rightarrow 5s4d \ ^1D_2$ decay rate limit. We expect that this is due to deeper traps having a steeper trapping potential, resulting in more energy being removed from the atom per Sisyphus scattering event. To effectively Sisyphus cool in the shallower tweezers, particularly for the 1.4 mK case, a lower 461 nm probe scattering rate should be used to avoid an excessive heating rate.

The data presented in chapter 4 use a 7.5 mK tweezer and so we will analyse this case further. We have found an optimal detuning for the Sisyphus beams of around $\frac{\Delta}{2\pi} \sim 7.7 \text{ MHz}$, next we must know what range of scattering rates we can use with

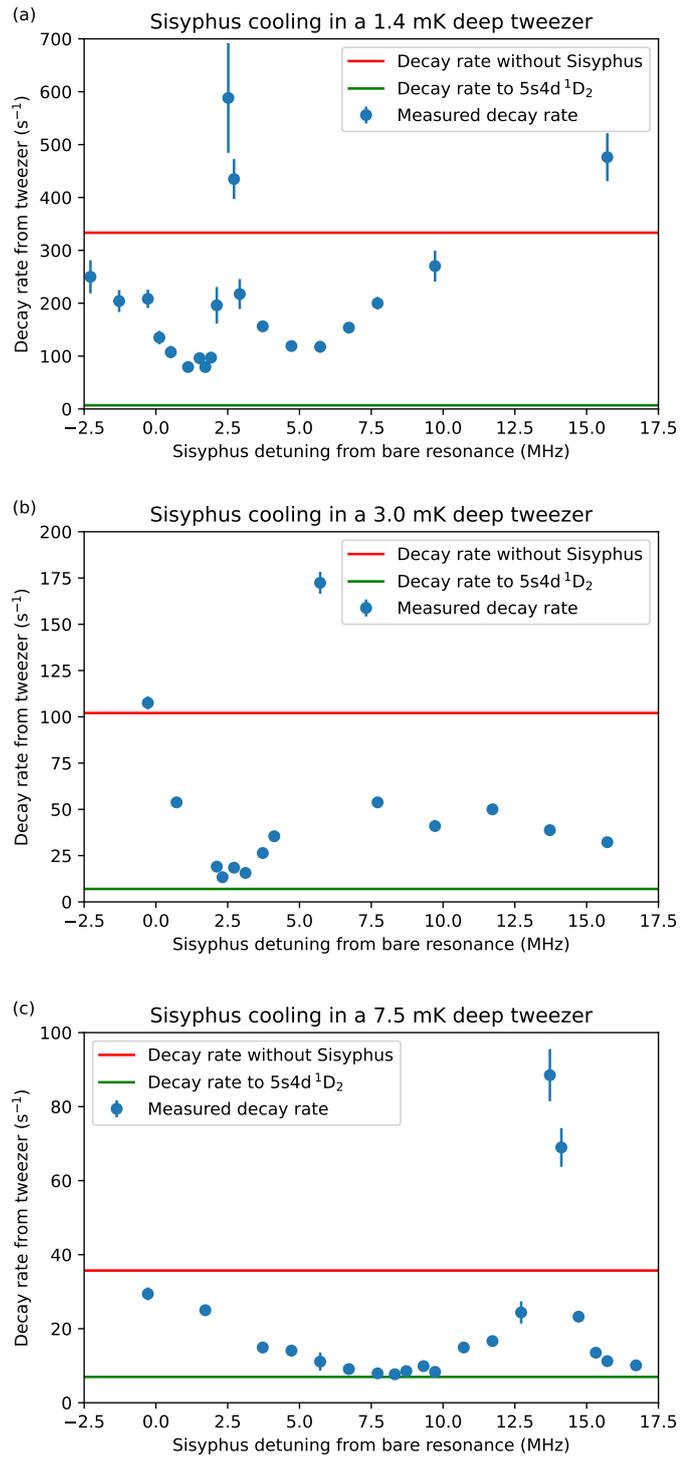


Figure B.1: (a,b,c) Decay rate from the tweezer as a function of Sisyphus detuning from the freespace $5s^2^1S_0 \rightarrow 5s5p^3P_1$ transition for trap depths of 1.4, 3.0 and 7.5 mK respectively. Plots include the measured decay rate in the absence of cooling in red, and the limiting decay rate if loss were to be purely caused by de-pumping to the $5s4d^1D_2$ state at the initial $0.14 \mu s^{-1}$ scattering rate in green.

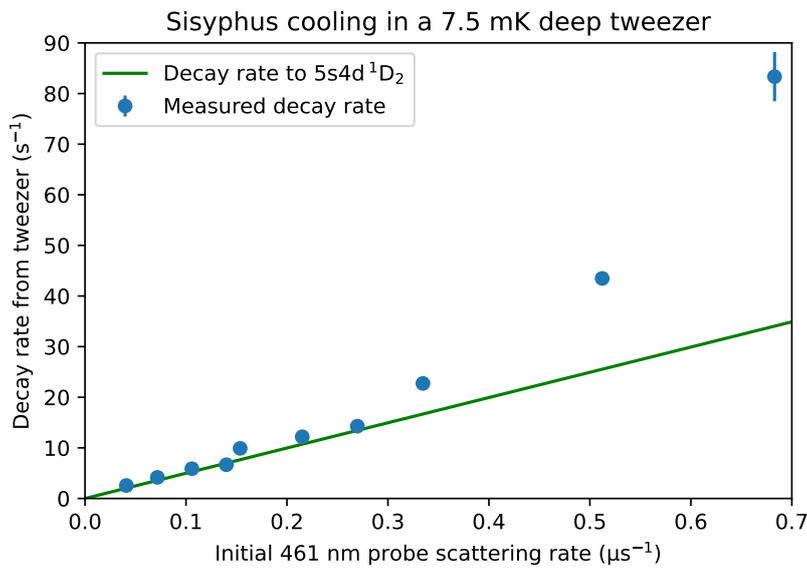


Figure B.2: Decay rate from the tweezer when under constant Sisyphus cooling, as a function of the initial scattering rate of the 461 nm probe beam. The straight line shows the expected decay rate if the Sisyphus cooling were to be totally effective and the only loss mechanism were to be loss via the $5s5p \ ^1P_1 \rightarrow 5s4d \ ^1D_2$ decay channel. We see that Sisyphus cooling seems to be effective for 461 nm probe scattering rates of up to $0.30 \mu\text{s}^{-1}$, beyond which heating effects shorten the lifetime under imaging.

the 461 nm probe beam while the cooling remains effective. If the probe scattering rate is set too high then the heating mechanism will overwhelm the Sisyphus cooling. If the probe scattering rate is set too low then the SPAD exposures will need to be longer to collect fluorescence, resulting in an unnecessarily high background count, n_b . For $5s5p\ ^1P_1 \rightarrow 5s4d\ ^1D_2$ limited decay, we expect the decay rate to be linear with the number of scattered 461 nm photons, deviations above this line imply that heating is taking place. Figure B.2 shows that for trap depths of 7.5 mK, Sisyphus cooling appears to be effective up to a probe scattering rate of $0.30\ \mu\text{s}^{-1}$.

B.2 Histograms including loss

To model the effect of the one-body loss in the histograms described in section 4.3, additional terms must be included. Here we choose to express these in terms of timescales. The two key timescales are the duration of the SPAD exposure t_{exp} , and the exponential decay time of the atom number τ for the set of probe beam parameters chosen during imaging.

For each possible atom number N_{at} loaded into the tweezer, we now must consider the different possible number of atoms that can ‘decay’² during the imaging frame N_{dec} which can range from 0 to N_{at} . This probability of a particular experimental run starting with N_{at} atoms and having $N_{\text{at}} - N_{\text{dec}}$ atoms at the end of the imaging frame will be given by

$$A(N_{\text{at}}, N_{\text{dec}}) = P_{\text{at}}(N_{\text{at}}) \frac{N_{\text{at}}!}{(N_{\text{at}} - N_{\text{dec}})! N_{\text{dec}}!} \left(1 - e^{-\frac{t_{\text{exp}}}{\tau}}\right)^{N_{\text{dec}}} e^{-(N_{\text{at}} - N_{\text{dec}})\frac{t_{\text{exp}}}{\tau}} \quad (\text{B.2.1})$$

which is comprised of the following terms:

- $P_{\text{at}}(N_{\text{at}})$ is the probability of starting with N_{at} atoms as defined in equation 4.3.1
- $\frac{N_{\text{at}}!}{(N_{\text{at}} - N_{\text{dec}})! N_{\text{dec}}!}$ is the number of different permutations of atoms with which N_{dec} atoms can decay given there are N_{at} atoms to begin with.

²‘Decay’ here refers to a reduction in atom number in the tweezer. In reality of course the atom does not decay as such, it is depumped into the anti-trapped $5s4d\ ^1D_2$ state, then the atom is ejected from the tweezer and lost.

- $\left(1 - e^{-\frac{t_{\text{exp}}}{\tau}}\right)^{N_{\text{dec}}}$ is the probability for a specific group of N_{dec} atoms to decay during the exposure
- $e^{-(N_{\text{at}} - N_{\text{dec}})\frac{t_{\text{exp}}}{\tau}}$ is the probability of the rest of the atoms not decaying during the exposure.

In equation B.2.1 there are a set of special cases where $N_{\text{dec}} = 0$ where no reduction in atom number occurs during the imaging frame. If the histogram was to be generated using these cases alone the histogram would be qualitatively like those in section 4.3.1, except with reduced amplitude in the peaks with $N_{\text{at}} \neq 0$. The reduction in the amplitude of each peak is more significant for higher N_{at} , where the chance of a decay occurring is higher. Adding the remaining cases where $N_{\text{dec}} \neq 0$ will result in the addition of a reasonably flat filling between the peaks in the histogram, due to the random decay times³, making it harder to resolve between the peaks. The form of this filling between the peaks shall now be considered in more detail.

We now return to considering that atom(s) may decay during the frame, and determining what number of scattered photons we expect during the imaging frame when decays occur. Note here that we only consider changes to the distribution of counts from scattered photons, as the distribution of background counts is unaffected. The expected value of SPAD counts from the scattered photons depends on the times at which atoms decay during the exposure. If each atom j that decays does so at time t_j , the expected scattered counts collected will be

$$C(N_{\text{at}}, N_{\text{dec}}, \vec{t}) = \alpha \left((N_{\text{at}} - N_{\text{dec}}) + \sum_{j=1}^{N_{\text{dec}}} \frac{t_j}{t_{\text{exp}}} \right),^4 \quad (\text{B.2.2})$$

where \vec{t} is a vector containing the various decay times t_j , i.e. $\vec{t} = (t_1, \dots, t_{N_{\text{dec}}})$.

To simulate this we consider the SPAD exposure to be broken up into M discrete time steps of size Δt at which the atom(s) can decay. These discrete times are used

³Essentially these cases behave like the atom number does not have to take on integer values, so they do not show the peaks that arose from equation 4.3.1.

⁴In the special case $N_{\text{dec}} = 0$ the counts become the original value of the mean counts αN_{at} used in equation 4.3.2.

for ease of computation, and for higher quality models one should use smaller values of $\frac{\Delta t}{t_{\text{exp}}}$. This then allows us to develop from equation B.2.1 which gave us the probability of starting with N_{at} atoms of which N_{dec} decay, to instead have the probability of starting with N_{at} atoms of which N_{dec} decay at specific times \vec{t} . The form of this probability is

$$W(N_{\text{at}}, N_{\text{dec}}, \vec{t}) = \begin{cases} A(N_{\text{at}}, N_{\text{dec}}) \left(\frac{e^{-\frac{\sum_{j=1}^{N_{\text{dec}}} t_j}{\tau}}}{\sum_{n_1, \dots, n_{\text{dec}}=0}^{M-1} e^{-\frac{\Delta t}{\tau} \sum_{l=1}^{N_{\text{dec}}} n_l}} \right) & N_{\text{dec}} > 0 \\ A(N_{\text{at}}, N_{\text{dec}} = 0) & N_{\text{dec}} = 0 \end{cases} \quad (\text{B.2.3})$$

where in the $N_{\text{dec}} > 0$ case the numerator is based upon the atom not decaying before the time t_j and then decaying in that particular time step, and the denominator normalises the weightings of the different decay times. The values n_l are the numbered time steps within which the l^{th} atom decays. Note that the numerator also would have explicitly included the probability of the atom that survived for time t_j then decaying within the next Δt , but that decaying probability is cancelled by the normalisation so is not apparent in the final expression.

There are a few practical considerations that must be made to use these distributions in a simulation. From equation 4.3.3 onward we have been summing over infinite possibilities for the atom number N_{at} , which is obviously not feasible. We therefore limit the possible range of atom numbers to be those which would have a reasonable probability of being loaded for the mean atom number in the tweezer $\overline{N_{\text{at}}}$, and which have a reasonable probability of having enough atoms decaying to contribute counts in the region of the single atom regime. For the data presented later in this chapter it was appropriate to constrain $0 \leq N_{\text{at}} \leq 5$.⁵ As mentioned earlier, the discrete time steps at which we model decay occurring should be small to improve the simulation, but not so small as to slow the simulation drastically. We choose $\frac{\Delta t}{t_{\text{exp}}} = 0.1$ as a reasonable compromise, which is sufficient to provide a smooth filling between the peaks of our histograms. The parameters of the model are well constrained, the mean atom number $\overline{N_{\text{at}}}$ is obtained from the experimental

⁵The highest mean atom number loaded in the results presented in section 4.5 is $\overline{N_{\text{at}}} = 1.2$, where the probability of loading $N_{\text{at}} > 5$ is small.

data, the decay rate τ can be determined from the branching ratio and probe scattering rate, and the background count rate on the SPAD pixel is well characterised. The expected counts per atom α is also predicted by the scattering rate and the efficiency of the imaging system, where the exact value is determined by inspection of the data.

B.3 Selective fidelity

To generalise the concept of a single-atom infidelity error described in section 4.4 to the case where more than one atom may be present, initially consider that an error occurs if a single atom ($N_{\text{at}} = 1$) is identified in the second frame but not the first, with associated probability

$$P(\bar{1}1) = P(01) + P(21) + P(m1) \quad (\text{B.3.4})$$

where $\bar{1}$ means $N_{\text{at}} \neq 1$, and m refers to all events with $N_{\text{at}} > 2$ for which we do not attempt to resolve atom number (as defined in equation 4.4.1).

$1 - P(\bar{1}1)$ is not a suitable definition for selective fidelity. As detailed in section 4.4.2, we seek to define fidelity that excludes contribution from loss of atoms from the tweezer. The terms $P(21)$ and $P(m1)$ will include strong contributions from loss, so would result in an underestimate of the fidelity. The subset of $P(21)$ and $P(m1)$ events that are due to infidelity will occur at a similar rate to the total number of $P(12)$ and $P(1m)$ respectively, so those terms can be used as part of loss-free replacements. Directly swapping these values underestimates the number of infidelity events, as the probability of seeing events with one atom in the first frame and two in the second are suppressed by loss. Therefore $P(12)$ is rescaled by the loss rate to prevent undercounting. Finally we can neglect the $P(1m)$ term, as the measured atom number increasing by two or more between successive frames is extremely unlikely. This leaves us with our definition of selective single atom fidelity

$$F'_1(\bar{N}_{\text{at}}) \approx 1 - P(01) - \frac{P(12)}{1 - L_1}. \quad (4.4.4)$$

This selective fidelity is a function of \bar{N}_{at} , and gets worse as \bar{N}_{at} is increased due to the larger $P(12)$ term.

Lastly we consider the selective fidelity for two-atom detection F'_2 . Loss errors L_2 are calculated directly from the data. However calculation of the fidelity in our data (to be shown in section 4.5) was hampered by the small number of frames with $N_{\text{at}} > 2$. Therefore the same threshold-based method was applied to the model instead, where the infidelity error can be obtained directly from the fraction of occurrences within thresholds B and C that were due to the model starting with two atoms. Comparing this model-based approach yielded values for F'_1 that were almost identical with those obtained empirically.

Appendix C

Magic wavelength tweezers

This appendix details information supporting the methods used in chapter 5.

C.1 813 nm tweezer preparation

The preparation and optimisation of a single 813 nm tweezer (mentioned at the end of section 5.1.2) is described, from initial alignment to the measurement of the 2.25(1) μm waist.

C.1.1 Initial alignment

The alignment procedure of the 813 nm tweezer largely follows the procedure used to align the 532 nm tweezer into the same chamber [72]. The 813 nm beam is first reduced to a ‘pencil beam’¹ by reducing the diameter of the iris at the end of the beam expander. Alternatively this could be done by placing a small aperture² at the end of the cage containing the $\frac{5}{3}$ magnification telescope. Immediately before the pencil beam enters the chamber, it passes through a piece of paper with a small hole at the centre. Using an infrared viewer, the back reflections from the viewport and *in-vacuo* lens are visible on the piece of paper just before the chamber. By walking

¹A narrow collimated beam of diameter ~ 1 mm.

²For example a Thorlabs CPA1 cage alignment plate.

the pencil beam using two of the steering mirrors between the telescope and the vacuum chamber, the back-reflections are optimised such that the beam should be at normal incidence to the centre of the *in-vacuo* lens. This procedure should place us in the region where the best possible tweezer will be formed, subject to some fine tuning to optimise the waist.

C.1.2 Loading and imaging a tweezer

In order to image the tweezer, we must first load atoms into it. To achieve this we must overlap the NB MOT with the tweezer. The position of an optimised tweezer is defined by the location of the *in-vacuo* lens that forms it, whereas the location of the NB MOT is defined by the location of the magnetic field zero, magnetic field gradient, and the detuning of the MOT beams. As the locations of the lenses and MOT coils are fixed, and the detuning of the MOT beams is an optimised parameter for the size and temperature of the NB MOT, we require an additional parameter to overlap the tweezer and the MOT. To do this, we can adjust the current through each of the three pairs of compensation coils (see section 3.3.2) to produce a magnetic field that shifts the position of zero magnetic field, and thus also shifts the MOTs.

As the spatial extent of the NB MOT is quite small, we first set up a probe beam to assist in locating the overlap region. We counter-propagate a 461 nm probe beam with the pencil³ 813 nm beam, overlapping the two together. The 813 nm beam is then turned off, and we use the 461 nm probe as a reference for where the tweezer would be. By releasing atoms from a blue MOT into the chamber, a large volume is filled with atoms, and one can then look for fluorescence to image the probe beam from the side. To image this we used the Pixelfly camera⁴. This was then compared to the location of an imaged red MOT, where the red MOT was then moved to overlap with the probe beam by changing the compensation coils. Fine tuning of this overlap is then done by using the 461 nm probe beam to best destroy the red MOT. This means that the atoms should be overlapped with the

³The 813 nm is still reduced here using the iris at the end of the beam expander.

⁴PCO Pixelfly QE, taken at an angle of 60° to the probe beam axis.

tweezer in both the lab \hat{x} and \hat{z} axes.⁵ There is no guarantee of any overlap on the \hat{y} axis, as the focus of the 461 nm probe and the 813 nm tweezer come from different *in-vacuo* lenses, and the *in-vacuo* lenses have wavelength dependant focal lengths. The 461 nm probe beam was then turned off, but the setup was retained for later.⁶

The iris in the 813 nm beam was now set to 3 mm diameter, with a beam power of 15 mW to the chamber. The reduced diameter of the 813 nm beam results in a longer Rayleigh length tweezer than if we were to have filled the full clear aperture of the *in-vacuo lens*, which makes finding the location of the tweezer with the NB MOT easier. The compensation coil in the \hat{y} direction is moved in steps, and the SPAD array brought back to a focused image of the MOT using a translation stage. At each \hat{y} compensation coil position, the tweezer is overlapped with the NB MOT for 50 ms, before turning off the NB MOT for at least 10 ms to allow the untrapped atoms to fall away. Any atoms that remain in the tweezer are then imaged by shining on 461 nm light from the blue MOT beams and collecting fluorescence on the SPAD array. Once a position is found where the tweezer is loaded, all three compensation coils are optimised to maximise the fluorescence from the tweezer (i.e. higher atom numbers loaded). The iris is then set to the correct size of 12 mm diameter, which corresponds to the narrowest (horizontal) axis of the SLM. This fills the full clear aperture of the *in-vacuo lens* and produces a much tighter tweezer. The SPAD array is translated to find the best image of the atoms trapped in the tweezer.

C.1.3 Optimisation of tweezer waist

Once the atoms have been loaded into the tweezer, we then must optimise the waist of the tweezer to be as small as possible. Decreasing the trap waist increases the intensity of the tweezer, so less laser power is required to achieve the same trap depth, and allows us to generate more tweezers for the same finite total laser power. Smaller waists also allow smaller separations between the tweezers, which is of great

⁵Axes are shown in figure 5.3.

⁶The 461 nm probe beam is used in section 5.3 as an alignment tool when we construct a new imaging system at higher magnification.

use for Rydberg and Rydberg-dressed experiments where we will want as many tweezer sites as possible within the appropriate volume. Tighter packed arrays also allow laser beams that are applied to the array to be more tightly focused, such as the clock and Rydberg lasers, where power should be utilised as efficiently as possible. Having a small trapping waist is particularly important for clocks, where we must minimise the Lamb-Dicke parameter [71].

In order to measure the waist of our tweezer, we measure the trap frequency via release-recapture [193]. The method for using a trap frequency experiment to measure the waist is well documented in prior work using strontium in 532 nm tweezers by Ryan Hanley [43, 72]. Atoms held in a tweezer are released briefly to allow the cloud of atoms to expand, before the tweezer is turned back on and the atoms start oscillating in the tweezer potential. The atoms are held in the tweezer for a variable hold time, followed by a second brief release of the atoms. If at the end of the variable hold time the atoms are spread spatially but have low momentum, then most of the atoms will remain trapped when the tweezer is turned back on. If however the atoms are all near the bottom of the tweezer potential and have maximum kinetic energy, then many atoms will move beyond the region of the tweezer before it is turned back on again, and many atoms are lost. Plotting the recaptured fraction as a function of the variable hold time (see figure C.1) gives oscillations at twice the radial trap frequency ω_r .

The trap frequency is then measured as a function of tweezer power at the atoms⁷, where a linear fit yields the waist of the tweezer (see figure C.2). The fitted gradient (κ) is used to calculate the waist (w_0) from the relation [72]

$$w_0 = \left(\frac{4\alpha}{m\pi\epsilon_0 c \kappa} \right)^{\frac{1}{4}}, \quad (\text{C.1.1})$$

where α is the ground state polarisability, m is the mass of a strontium atom, and all other values are fundamental constants. Fitting to figure C.2 found a tweezer waist of 2.88(1) μm for the initial alignment.

⁷There are losses of tweezer power at several optics, notably including $\sim 5\%$ loss on the ITO coated face of the *in-vacuo* lens.

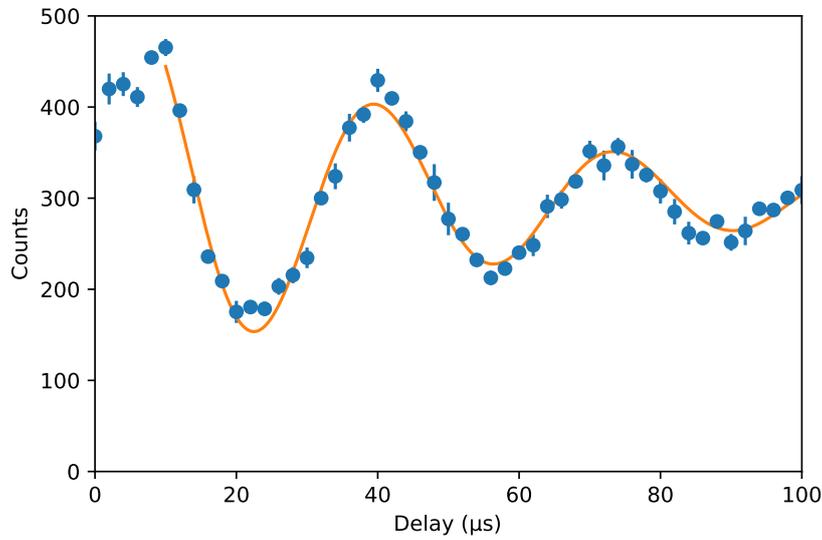


Figure C.1: Example trap frequency measurement for a tweezer power of 36 mW. For this particular trap frequency measurement, the first and second release durations were 30 μs and 70 μs respectively, with a variable hold time in the range of 0 to 100 μs between the two releases. The release durations are empirically chosen to give good contrast in the oscillations. The trap frequency here is $0.093 \pm 0.001 \mu\text{s}^{-1}$.

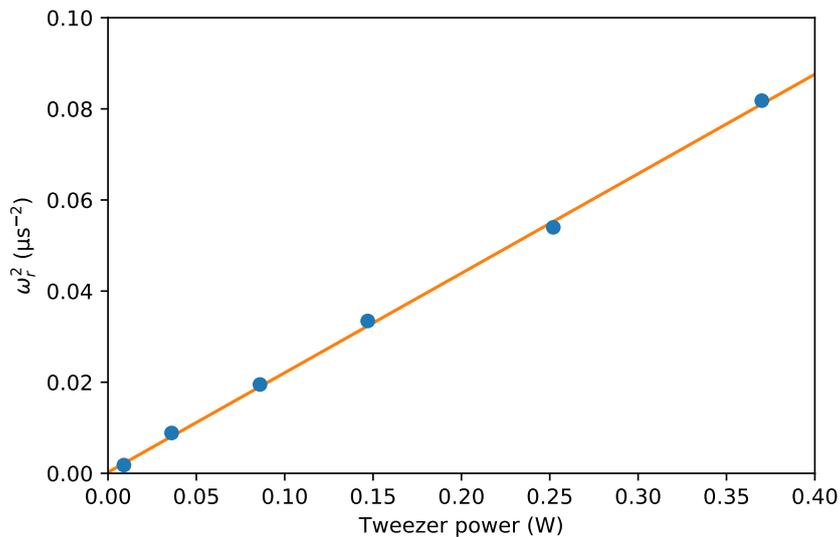


Figure C.2: Measured trap frequency from release-recapture experiments as a function of tweezer power. The linear fitting here finds a tweezer waist of $2.88(1) \mu\text{m}$.

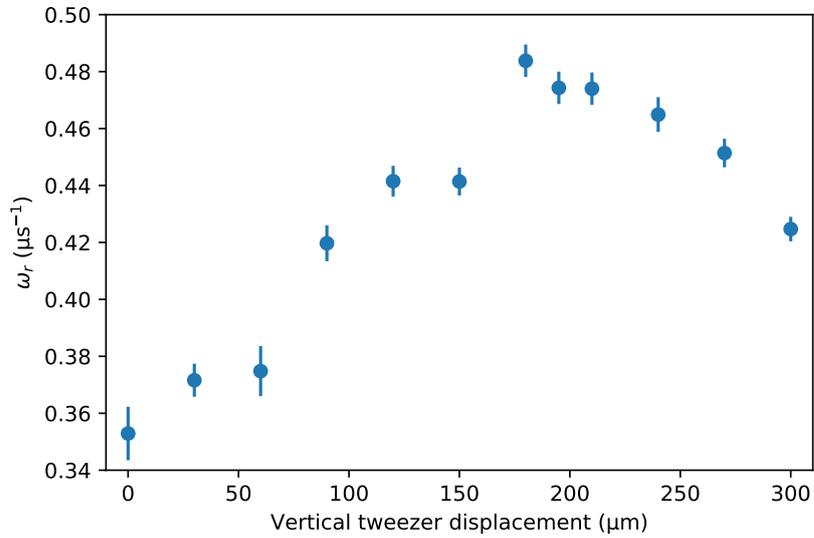


Figure C.3: Vertical alignment of the trapping light into the *in-vacuo* lens to optimise the tweezer waist. The tweezer waist is minimised by maximising the trap frequency, with a notable improvement in this direction.

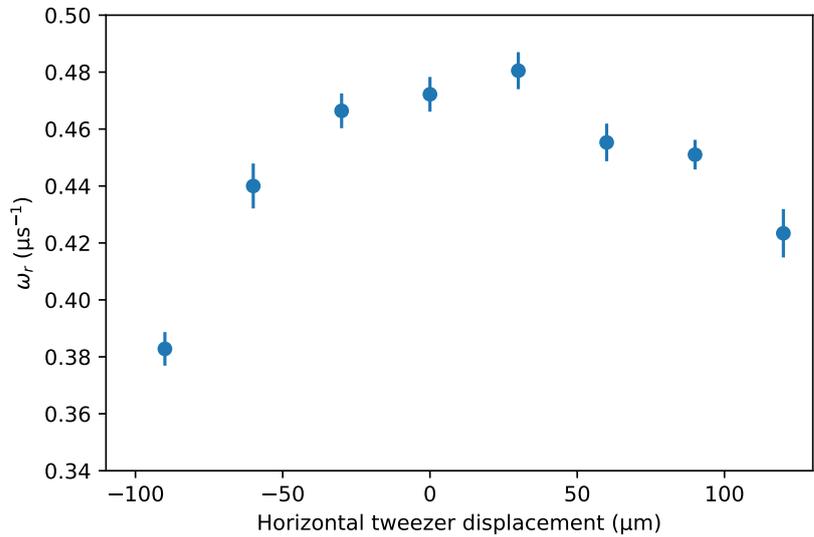


Figure C.4: Horizontal alignment of the trapping light into the *in-vacuo* lens to optimise the tweezer waist following on from the vertical optimisation in figure C.3. The tweezer waist is minimised by maximising the trap frequency.

To optimise the waist, we pick a specific tweezer power, and measure the observed trap frequency as a function of an experimental parameter. Here we adjust a single mirror steering the 813 nm beam before the vacuum chamber to improve the transverse alignment into the *in-vacuo* lens. The optimisation was done in first the vertical and then the horizontal directions, as shown in figures C.3 and C.4. The vertical alignment shows a significant increase in the trap frequency, corresponding to a decrease in the tweezer waist as desired. The trap frequency varies with the alignment comparably with previous work aligning the 532 nm tweezer [72]. The optimised tweezer waist here is 2.25(1) μm . This waist measured from the trap frequency is found to be consistent with both the measured differential AC Stark shift of the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1$ transition compared to the calculated polarisability values, as well as the refocused spot size of the trapping light after the chamber measured on a the CCD.⁸

In between figures C.2 and C.3 the collimation of the trapping light between the $\frac{5}{3}$ magnification telescope and the vacuum chamber was improved, where it had been found to be slightly converging in the initial setup. There are several routes of further optimisation of the trap waist. Finer optimisation of the trapping beam divergence, perhaps using a Fresnel lens pattern on the SLM (see section C.2.2) could lead to a further reduced waist. The transverse alignment done in figures C.3 and C.4 only adjusted a single optic, walking the transverse alignment into the *in-vacuo* lens will likely improve performance even further. In addition, reducing aberration in the tweezer spot using an SLM phase mask (see section C.2.2) may also offer improvement. Improving the tweezer waist further is left for future work, for the rest of this chapter we continue with a tweezer of waist 2.25(1) μm .

C.2 Phase masks for the SLM

This section builds upon the introduction to generating phase masks for the SLM in section 5.2.3. We now consider the calibration of the SLM, the different type of phase masks we can send to the SLM, and the Gerchberg-Saxton algorithm used to

⁸DMK 33UX273.

optimise the phase masks.

We would like to thank our collaborators at the Institut d’Optique Graduate School in the group of Antoine Browaeys for their advice and support in implementing the SLM tweezer array in Durham, in particular for very kindly sharing their code used to generate and optimise phase masks for the SLM. Some details of their work can be found in references [35, 36, 39, 40, 56], where much of the work in this section is based upon the theses of Henning Labuhn [40] and Sylvain de Léséleuc [36].

C.2.1 SLM calibration

To convert the digital phase masks into a pattern on the SLM, there must be a voltage that the computer can send to each pixel that corresponds to the particular phase shift. There is a look up table (LUT) of values for different wavelengths sent to the SLM. In the case of 813 nm, the manufacturers specification would suggest that a value of 215 would correspond to a phase shift of 2π . This allows us to control the phase on each pixel in $\frac{2\pi}{215}$ increments. All phase masks (figures C.5, C.6 & C.7) are shown in grey corresponding to this calibration.

To test the calibration was accurate, we took our own measurement to confirm it. The light to the SLM from the beam expander was passed through two neighbouring pinholes. The light from each pinhole was incident on a different half of the SLM, where each half was configured to have a different phase. One half was kept constant at phase 0, where the other half has its phase varied for different measurements. For each phase on the second half, the interference pattern in the light reflected from the SLM was fitted to, and the shift of the interference pattern as a function of phase yields a calibration value for a 2π phase shift [194]. This method retrieved a calibration value of 212, consistent with the manufacturer specification. For fine refinement to this calibration, a more sophisticated measurement approach [195, 196] could be used.

The SLM is comprised of discrete pixels, in a grid of size $N_x \times N_z = 792 \times 600$, with pixels of size $\Delta_x \times \Delta_z = 20 \mu\text{m} \times 20 \mu\text{m}$. Following the treatment of references [40, 197], $\Delta_x \times \Delta_z$ defines a SLM unit. The maximum physical size of the field at

the SLM is $L_x \times L_z = N_x \Delta_x \times N_z \Delta_z \approx 16 \text{ mm} \times 12 \text{ mm}$. In the focal plane of the of the *in-vacuo* lens, a different set of relevant units is defined, giving a physical size $\tilde{L}_x \times \tilde{L}_z$ and a focal unit of the smallest possible shift of $\tilde{\Delta}_x \times \tilde{\Delta}_z$. These values are related by:

$$\tilde{L}_x \times \tilde{L}_z = \frac{\lambda_{\text{trap}} f_{\text{eff}}}{\Delta_x} \times \frac{\lambda_{\text{trap}} f_{\text{eff}}}{\Delta_z}, \quad (\text{C.2.2})$$

$$\tilde{\Delta}_x \times \tilde{\Delta}_z = \frac{\lambda_{\text{trap}} f_{\text{eff}}}{L_x} \times \frac{\lambda_{\text{trap}} f_{\text{eff}}}{L_z}. \quad (\text{C.2.3})$$

Where f_{eff} is the effective focal length of the *in-vacuo lens* combined with the telescope after the SLM.

$$f_{\text{eff}} = \frac{f}{m}, \quad (\text{C.2.4})$$

where f is the focal length of the *in-vacuo lens* and $m = \frac{5}{3}$ is the magnification of the telescope.

This means that in the focal plane of the *in-vacuo lens* there is a region of size $\tilde{L}_x \times \tilde{L}_z = 930 \text{ } \mu\text{m} \times 930 \text{ } \mu\text{m}$ which can be reached using the SLM. There is however a strong decrease in diffraction efficiency towards the edge of this range as the gratings become pixelated. For diffraction efficiencies $> 75 \%$ [180] we can work in roughly a $230 \text{ } \mu\text{m} \times 230 \text{ } \mu\text{m}$ region, which in the high magnification imaging system in section 5.3 will encompass the full view of the SPAD array.

The size of the focal unit in our experiment will not be the maximally achievable value, as we illuminate a 12 mm circle within in the SLM window with light from the beam expander. This sets $L_x = L_z = 12 \text{ mm}$, and leads to the focal unit being $\tilde{\Delta}_x = \tilde{\Delta}_z = 1.55 \text{ } \mu\text{m}$, setting the smallest scale with which we can control our tweezers.

C.2.2 Types of phase masks

Before we detail the exact procedure in how one can arrive at a phase mask that generates a desired tweezer away, we first consider what kind of effects can be achieved using a phase mask. The tweezer position can be translated in three dimensions through blazed diffraction gratings, and Fresnel lenses. The light can be split into multiple tweezers. The tweezers themselves can also have aberration corrections

applied, as well as general shaping of the beam profiles. Before we go into the details of each, we will first consider how these kind of patterns can be combined and displayed to produce the desired intensity distribution.

Combining phase masks

There are two key mathematical properties to consider when preparing the phase mask of a SLM. The first is that a phase shift of 2π is equivalent to no phase shift at all, so all phase patterns can be taken modulo 2π before sending the pattern to the SLM. This means that the SLM need only be able to produce phase shifts of up to 2π . As we saw when discussing the SLM calibration in section C.2.1, this gives us 215 discrete values of phase available for each pixel.

The second key property of phase masks is that we can add different masks together to combine their effects. The *in-vacuo* lens focusing the 813 nm light acts as a Fourier transform of the pattern of the light. The Fourier transform of the sum of two masks is equal to the sum of the Fourier transforms,

$$\mathcal{F}[g(x, z) + h(x, z)] = \mathcal{F}[g(x, z)] + \mathcal{F}[h(x, z)]. \quad (\text{C.2.5})$$

This means that we can simply add the phase masks $\theta(x, z)$ together and it will apply both effects at once. As an example of this we could apply a diffraction grating $\theta_1(x, z)$ to translate the tweezer in the $(\tilde{\mathbf{x}}, \tilde{\mathbf{z}})$ plane and a defocus $\theta_2(x, z)$ to translate it in the $\tilde{\mathbf{y}}$ plane. The phase mask we would then need to send is simply

$$\theta_{\text{total}}(x, z) = \theta_1(x, z) + \theta_2(x, z), \text{ modulo } 2\pi. \quad (\text{C.2.6})$$

This result generally holds, and is not specific to the two masks being translations of a tweezer, as it is a basic property of Fourier transforms.

Translations

Each tweezer has a position controllable in three dimensions. For translations in the focal plane $(\tilde{\mathbf{x}}, \tilde{\mathbf{z}})$, we can simply apply a phase mask that tilts the wavefront on the beam. When this effect is taken modulo 2π , this takes the form of a blazed diffraction grating pattern. To move the focus along the tweezer axis a quadratic

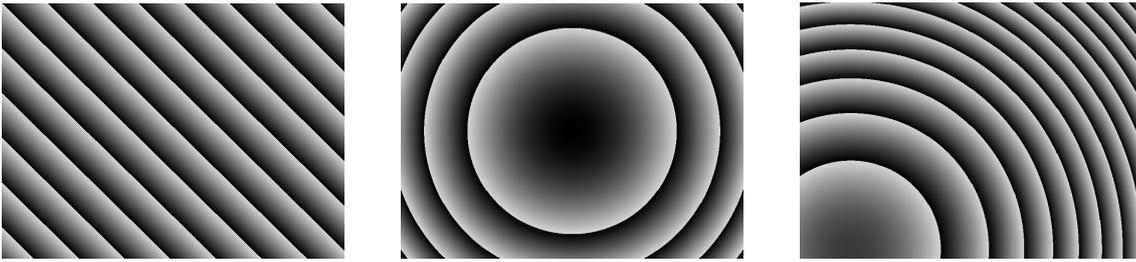


Figure C.5: Phase masks for translating a tweezer. (left) Phase mask for a grating to translate in the focal plane. (middle) Phase mask for Fresnel lens to defocus the tweezer into a different plane. (right) Phase mask combining the previous two effects to a 3D translation.

phase is applied, where the effect is as though the SLM had a lens on it. Taking the quadratic shift of a lens modulo 2π is a Fresnel lens pattern.

We now need to consider the specifics of how to achieve a phase pattern that shifts the tweezer the correct amount. Taking the natural position of the optical tweezer when no phase mask is applied to be at position $(\tilde{x}, \tilde{y}, \tilde{z}) = (0, 0, 0)$, we need to find the correct phase pattern $\theta(x_{\text{SLM}}, z_{\text{SLM}})$.

For a blazed grating in the focal plane, we need a pattern proportional to the distance that we want to move in the focal plane (\tilde{x}, \tilde{z}) . This pattern takes the form

$$\theta_{\text{grating}} = \frac{2\pi}{\lambda_{\text{trap}} f_{\text{eff}}} (\tilde{x}x_{\text{SLM}} + \tilde{z}z_{\text{SLM}}), \text{ modulo } 2\pi. \quad (\text{C.2.7})$$

Where each term corresponds to a grating in the orthogonal lab axes, which summed together make a grating on an arbitrary axis within the plane. The phase mask for this is shown in figure C.5(left).

For a Fresnel lens pattern to shift the focal position of the tweezer (the \tilde{y} axis), we need to apply a quadratic phase shift that is proportional to how far we want to defocus the tweezer \tilde{y} .

$$\theta_{\text{defocus}} = \frac{\pi\tilde{y}}{\lambda_{\text{trap}} f_{\text{eff}}^2} (x_{\text{SLM}}^2 + z_{\text{SLM}}^2), \text{ modulo } 2\pi. \quad (\text{C.2.8})$$

The phase mask for this is shown in figure C.5(middle).

As discussed earlier, SLM phase masks can simply be added to combine their effects, a three dimensional translation of a single tweezer to coordinates $(\tilde{x}, \tilde{y}, \tilde{z})$ is

then produced by the phase mask [36]

$$\begin{aligned}
\theta_{3D} &= \theta_{\text{grating}} + \theta_{\text{defocus}}, \text{ modulo } 2\pi \\
&= \frac{2\pi}{\lambda_{\text{trap}} f_{\text{eff}}} (\tilde{x} x_{\text{SLM}} + \tilde{z} z_{\text{SLM}}) + \frac{\pi \tilde{y}}{\lambda_{\text{trap}} f_{\text{eff}}^2} (x_{\text{SLM}}^2 + z_{\text{SLM}}^2), \text{ modulo } 2\pi \\
&= \frac{\pi}{\lambda_{\text{trap}} f_{\text{eff}}^2} ([2f_{\text{eff}} x_{\text{SLM}}] \tilde{x} + [x_{\text{SLM}}^2 + z_{\text{SLM}}^2] \tilde{y} + [2f_{\text{eff}} z_{\text{SLM}}] \tilde{z}), \text{ modulo } 2\pi.
\end{aligned} \tag{C.2.9}$$

The phase mask for this is shown in figure C.5(right).

Multiple tweezers

Light can be diffracted into multiple tweezers, which is how we arrive at our tweezer array. Each tweezer can be assigned to an independent spatial location in three dimensions, such that arbitrary geometries can be achieved [35, 36]. The ability to form three dimensional arrays is particularly important for maximising the number of sites within a dressed blockade volume [21]. We will see atomic fluorescence from different tweezer arrays that we can produce later in figure 5.7.

Now back to how arbitrary 3D arrays of tweezers can actually be achieved. Simply adding the phases of multiple tweezer positions will not help in this case, as that would just repeatedly adjust the three dimensional position of a single tweezer. Instead we must consider the problem in reverse. If coherent light from N independent point sources at the positions of the desired tweezers,

$$(\tilde{x}_n, \tilde{y}_n, \tilde{z}_n), \quad n \in \mathbb{Z} : n \in [1, N], \tag{C.2.10}$$

were incident upon the SLM, what phase and amplitude would it have? Each tweezer $(\tilde{x}_n, \tilde{y}_n, \tilde{z}_n)$ will have a phase Θ_n and a weight (relative amplitude) w_n in the form $w_n e^{i\Theta_n}$. At the SLM these N sources will interfere together to make a beam of the form [36]

$$A_{\text{SLM}} = \sum_{n=1}^N w_n e^{i[\theta_n + \Theta_n]}, \tag{C.2.11}$$

where θ_n here is defined in the same way as equation C.2.8 for the 3D position of a single tweezer,

$$\theta_n = \frac{\pi}{\lambda_{\text{trap}} f_{\text{eff}}^2} ([2f_{\text{eff}} x_{\text{SLM}}] \tilde{x}_n + [x_{\text{SLM}}^2 + z_{\text{SLM}}^2] \tilde{y}_n + [2f_{\text{eff}} z_{\text{SLM}}] \tilde{z}_n). \tag{C.2.12}$$

If one were able to produce the light field at the SLM described by equations C.2.11 and C.2.12, one ought to then have the desired light distribution of N tweezers at the desired positions $(\tilde{x}_n, \tilde{y}_n, \tilde{z}_n)$. This pattern is not directly achievable using our SLM as it only has access to encoding phases, not amplitudes. However it has been demonstrated that we can start producing an array by taking only the argument of A_{SLM} (as defined in equation C.2.11) as

$$\theta_{\text{estimate}} = \arg [A_{\text{SLM}}], \quad (\text{C.2.13})$$

where all traps are set to the same depth $w_n = 1$, and the phases Θ_n are randomised [36]. This does not produce an idealised array of traps with all the same trap depths, and this will require optimisation using both the Gerchberg-Saxton algorithm [170], and tweezer intensity feedback [36, 39, 40], which we will discuss in section C.2.3. Example phase masks for producing tweezer arrays are shown in figure C.6.

Tweezer shaping

Not only can we make multiple tweezers and move them in three dimensions, we can also change the shape of the tweezers themselves. Aberrations in the tweezer shape arise from imperfections in the optics setup that create the tweezers. These aberrations can be corrected by applying masks corresponding to the Zernike coefficients matching the aberrations. There is even a correction to compensate the imperfect flatness of the SLM itself supplied by the manufacturer⁹, which is shown in figure C.7. Reducing aberration can improve the waist of the tweezers [39, 40], but this can also be taken further to produce tighter confinement than would be given by a Gaussian beam. By using SLM patterns that combine Laguerre-Gauss modes [174, 175], it is possible to produce even smaller trapping volumes. Also where the ground state is anti-trapped it is possible to produce bottle beams [62, 85], where atoms are trapped instead at the point of lowest intensity. Bottle beams are not appropriate for trapping at 813 nm where the ground state is trapped, but Laguerre-Gauss modes to generate tighter confinement are a potential avenue of investigation in the future.

⁹Hamamatsu.

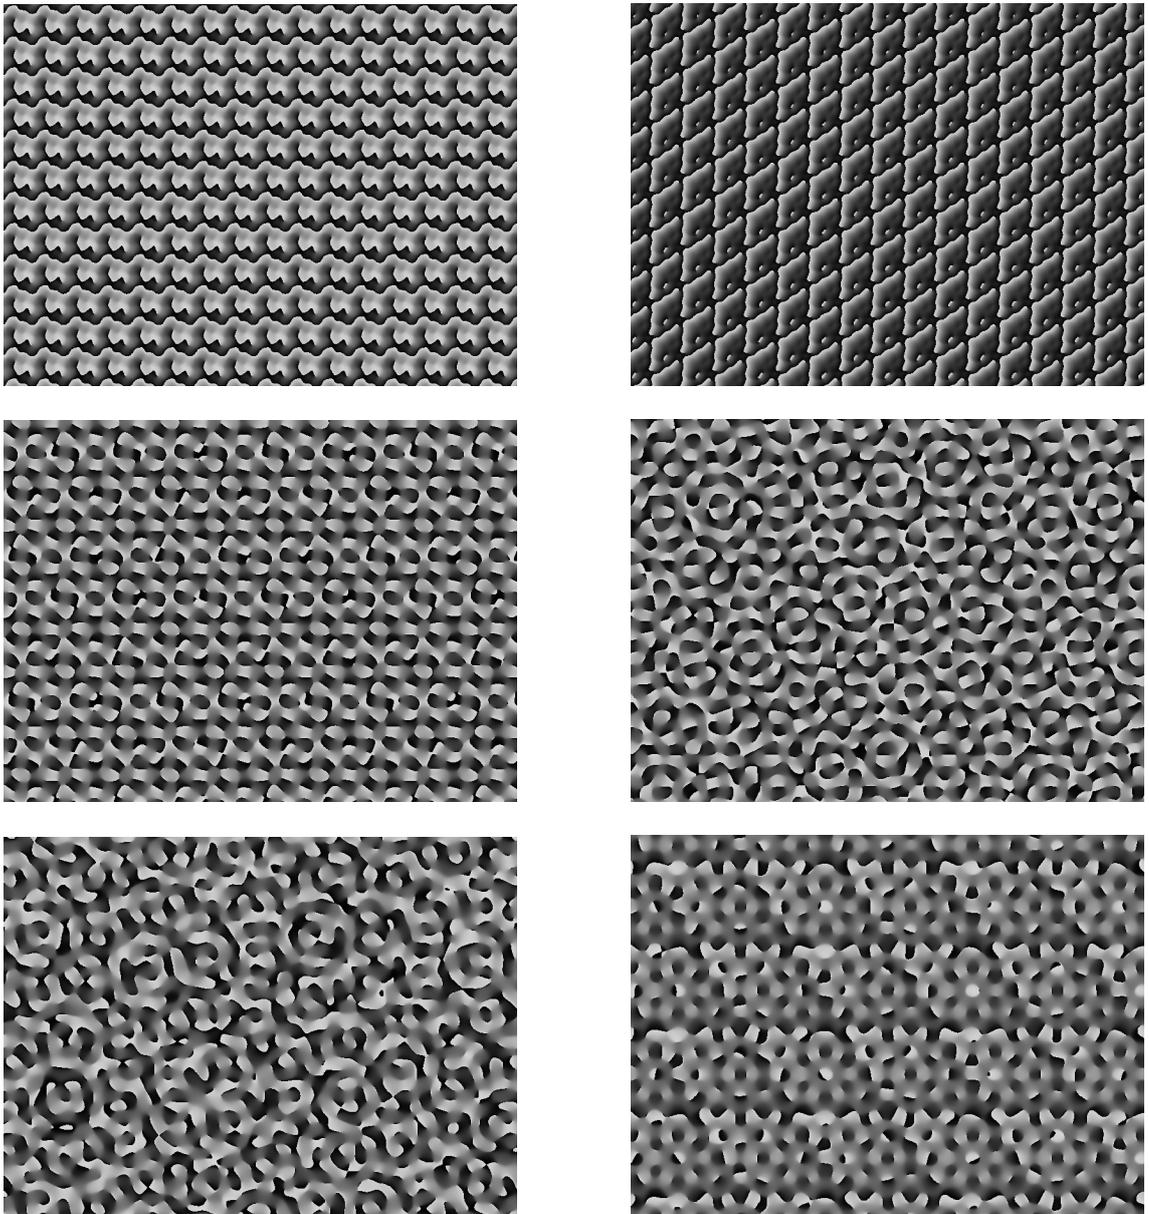


Figure C.6: Gallery of phase masks for producing tweezer arrays. These phase masks correspond to the tweezer arrays shown in figure 5.7, up to a global 3D translation which is not included in the mask here. The phase masks in the top row are for square and triangular tweezer arrays, the middle row are for an octagon and hexadecagon of tweezers, and the bottom row are for non-regular array shapes of a smile and an analogue clock.



Figure C.7: Phase mask provided by Hamamatsu to correct for the deviations in the SLM surface away from being flat. The correct pattern here is for light at 810 nm.

C.2.3 Gerchberg-Saxton algorithm

The toolbox of phase masks available on an SLM to produce 3D tweezer arrays is extremely useful and versatile. However, with the SLM only having access to the phase of the light beam rather than both phase and amplitude, our first guesses at a phase mask to produce a desired tweezer array (see equation C.2.13) are not perfect. To remedy this, we continue to use software kindly shared with us by the group of Antoine Browaeys, including adapted versions of the Gerchberg-Saxton algorithm. For an original description of the algorithm used please see references [36, 40]. Here we describe the basics of the algorithm as it is used in our experiment, where it is shown in figure C.8

We begin with a calculated phase mask (θ_{estimate}) as defined in equation C.2.13, where each tweezer is intended to have equal depths w_n and randomised phases Θ_n . The phase mask is combined with the incident light amplitude, before being virtually propagated to the focal plane of the *in-vacuo* lens, which can be done by FFT¹⁰ [39, 40] or by the diffraction formula [36]. The calculated image in the focal plane has an amplitude for each tweezer, which we then replace by the square root of the calculated target intensity $\sqrt{I_t}$ (the desired amplitude). The new light field is

¹⁰Fast Fourier transform.

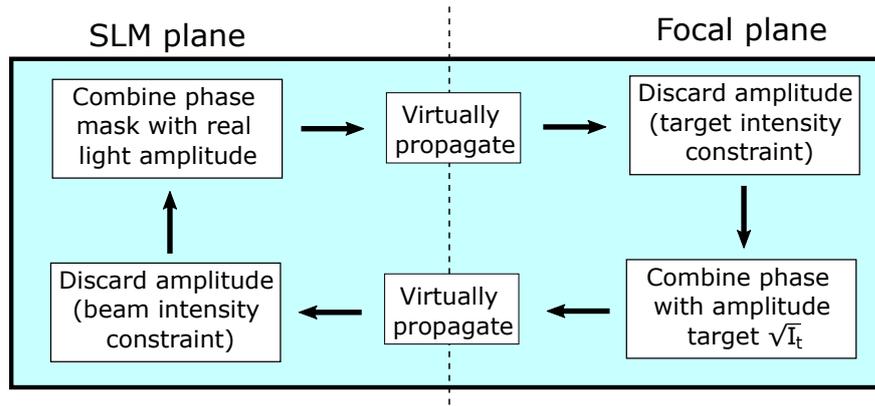


Figure C.8: Basic idea of the Gerchberg-Saxton algorithm. We begin with a calculated phase mask θ_{estimate} , and iteratively propagate between the SLM and focal planes. In each plane the phase is retained, but the amplitude is replaced by either the target intensity or actual light field in the focal and SLM planes, respectively.

then virtually propagated back to the SLM plane to calculate the light field required at the SLM to produce the traps. Once again, we run into the issue that the SLM can only control the phase, so we discard the calculated amplitude information (the amplitude is defined by the actual light beam), and keep the new calculated phase to send to the SLM. This algorithm repeats for several iterations and converges towards a more accurate phase mask, θ_{GS} .

The next step in refining the phase mask is to use a further modified Gerchberg-Saxton algorithm that feeds in the real world intensity distributions [35, 36, 39, 40]. The 813 nm trapping light is imaged by the CCD after the chamber, measuring the intensity of each tweezer directly. The correction algorithm proceeds similarly to the above, where the intensity is now physically measured, and the target intensity I_t for each tweezer is then adjusted based upon the measured intensities. For example, if one tweezer is too shallow, then the target intensity for that tweezer is increased. This is done using a smaller than unity gain factor, and the process is repeated until the uniformity of the trap depths are within tolerance. [39]

Thanks to the application of the Gerchberg-Saxton algorithm including the trap intensity feedback, the arbitrary tweezer array has good trap depth homogeneity. Each tweezer having the same trap depth means that any shifts on the Rydberg

or clock transitions would be the same to across the array¹¹. We can also use this approach to arrive at traps of deliberately different depths, which could be used to provide addressability between tweezer sites, for example working with a deliberately slight off-magic tweezer wavelength.

C.3 Corrective optics for imaging

In section 5.3.3 we introduce the need for a pair of high tolerance airside corrective optics to remove aberration in the *in-vacuo* lens when imaging at 461 nm. This section details the scope of the problem with the *in-vacuo* lens, as well as the design, optimisation and theoretical performance of the corrective imaging system.

C.3.1 Aberration correction

The *in-vacuo* lens surface includes up to 12th order aspheric coefficients, so requires significant air-side correction to undo this. The viewport of the vacuum chamber is 160 mm from the atoms (compared to the ~ 37 mm working distance for the *in-vacuo* lens), and there are spatial constraints on how close any corrective optic can be placed to the chamber. This is a long distance between the aberrant *in-vacuo* lens and the corrective optics, and makes it harder to work with. We wish to both trap and image through the same *in-vacuo* lens, such that the atom-lens distance is better defined, meaning we must include a dichroic¹² on the airside of the viewport before any corrective optic can be placed. This is because we do not want to pass the 813 nm tweezer light through a corrective optic, as that would instead add aberrations to the trapping light, and form a tweezer with a worse waist. We therefore place the corrective optic 140 mm further than the viewport¹³.

In order to sufficiently correct the aberration over a suitable transverse field, we require two surfaces with corrections imprinted upon them, but surfaces on different

¹¹To within the homogeneity of the tweezers, and subject to the atomic distribution within the tweezers.

¹²Thorlabs DMLP490L longpass dichroic mirror, 490 nm cut-on..

¹³300 mm away from the atoms.

optics. Each optic has a convex side where the corrections are machined, and a concave side to control the collimation of the fluorescence. The corrections are on the convex sides of each optic as the high order aspheric terms can be machined to a much higher tolerance on a convex face. The corrective optics require a complicated surface, where a cross-section of each lens can be seen in the expanded view in figure 5.3. The convex face of the first corrective optic requires an upturned edge, which limits what can possibly be manufactured. The corrective elements are designed to work with a 20 mm clear aperture of our *in-vacuo* lens.

The focal length of the *in-vacuo* lens varies with wavelength, being almost 1 mm shorter at 461 nm than at 813 nm [72]. The corrective optics were designed for imaging atoms near the natural focal plane of the 813 nm tweezer array¹⁴. Imaging in a plane further away than the 461 nm focus results in the collected fluorescence from the *in-vacuo* lens being weakly converging, such that a clear aperture of 17.25 mm is sufficient on the first of the corrective optics, and should not result in any loss of fluorescence.

The two corrective optics each have a conic and five non-zero aspheric coefficients, up to 12th order. They were designed by Frédéric Leroux, and manufactured by Kreischer Optics from NBK-7 glass. The corrective optics must be placed precisely relative to the *in-vacuo* lens, and especially relative to each other, in terms of both position and angle. We therefore mounted the two corrective optics in a custom shared V-groove to attempt to get the best possible centration. The separation of the corrective optics was set using a custom machined set of separation blocks at precisely the desired thickness, which is then slid out after the optics are placed and secured on either side of it. The whole V-groove is mounted on a 3D translation stage, and the two adjusters on the dichroic provide control of the two relevant angular degrees of freedom¹⁵. The optimisation of these five degrees of freedom on the position and angle of the lens pair was done iteratively, optimising the collected

¹⁴To image different planes either the final lens or the SPAD array must be translated to compensate, and this should have minimal impact on imaging performance for reasonable changes in the plane being imaged.

¹⁵As the optics are spherically symmetric, only the pitch and yaw matter, the roll is inconsequential.

fluorescence from each tweezer in terms of both the peak intensity and the spot shape.

C.3.2 Theoretical imaging performance

To simulate the performance of the high magnification imaging system with the corrective optics in place, the system is modelled using Zemax OpticStudio. Diffraction limited performance of the system would correspond to a $65\ \mu\text{m}$ Airy radius¹⁶, resulting in 95% of the atomic fluorescence detected by the SPAD array being collected on a single pixel. Zemax predicts for an idealised system that atoms in the imaging plane which are up to $40\ \mu\text{m}$ away from the optical axis should have diffraction limited performance.

In reality, this idealised behaviour is not feasible or expected, we are not going to do this well in collecting fluorescence onto a single pixel. We are not imaging an idealised point source, the atoms have a spatial distribution within the optical tweezer, with some variation in their 3D position. Moreover, the manufacturing and placement of the optics used will not be perfect, where variations in these within tolerances will degrade performance below the diffraction limited 95% of detected fluorescence being on a single pixel. We consider each factor that will limit performance in isolation, as it is important to see how each individual effect could bound our imaging performance, and it is not *a priori* clear how these effects will combine.

We begin by considering a perfectly manufactured and constructed imaging system, what drop in performance could we expect from the finite spatial distribution of the atoms? We now consider the transverse and longitudinal distributions of the atoms in the tweezer separately. If we were to take the transverse spread of the atoms to be on the scale of the tweezer waist ($2.25\ \mu\text{m}$), the expected collection fraction onto a single pixel would fall to 74%. The atoms also spread along the longitudinal direction, where for example if there were an atom on axis but at the

¹⁶Everything within the first minima of the Airy pattern would fall within the $75\ \mu\text{m}$ radius SPAD pixel.

Rayleigh length away from the waist, the collection fraction from that atom to a single pixel would be 69%. Typical atoms will have longitudinal displacements less than the Rayleigh length, so we expect the performance due to the longitudinal spread of atoms to be notably better than this. The finite size of the atomic distribution clearly could reduce the collection onto one pixel, and motivates having a spacing between the pixels used for adjacent tweezers.

We now return to considering imaging of a point source, but instead consider that there are imperfections within the imaging system. The corrective optics were manufactured to high tolerance, where variations in the surface are unlikely to be the leading issue in our limiting performance. The surfaces of the corrective optics were measured with a Zygo interferometer by Kreischer Optics, with RMS (peak to valley) deviations of 14 (70) nm for the first optic, and 11 (40) nm for the second optic. The other key manufacturing parameters are the surface flatness of the viewport of the vacuum chamber and the dichroic, both of which are fortunately in a region of reasonably well collimated light. The key concern in the construction of the imaging system for loss of performance is instead the placement tolerance of the corrective optics. We focus discussion on two of the practical tolerances to give a sense of scale on the sensitivity of the optics placement: the distance between the two corrective optics, and the distance between the corrective optic and the *in-vacuo* lens.

The 15 mm separation of the two corrective optics is critical to their use, where the spacing must be set on the scale of 0.1 mm. A spacing error of 0.1 mm would reduce the diffraction limited region from up to 40 μm off axis to 30 μm , which is still within acceptable performance. If however the separation is wrong by 0.2 mm or 0.5 mm, nothing would be diffraction limited, and the expected collection fraction on a single pixel for a tweezer 15 μm off axis¹⁷ falls to 70% and 15% respectively. This highly sensitive placement is why the optics were separated using a custom precision spacer¹⁸, as this separation is critical. The centration of the two corrective

¹⁷If a 6×6 tweezer array centred around the optical axis with $\sim 6.1 \mu\text{m}$ spacing, we would be imaging up to $\sim 15 \mu\text{m}$ off axis.

¹⁸Made by the Mechanical Engineering Services team in the Durham University physics department.

optics is also key for them to remove aberration from the system, which is why they are placed in a shared V-groove.

The placement of the optics relative to the *in-vacuo* lens is also important, but to a lower sensitivity. There is a 140 mm separation between the *in-vacuo* lens and the first corrective optic, which is difficult to measure as it is between an object in the heart of the vacuum chamber and an object outside of the chamber, with a dichroic folding path. For a 1 mm error in this separation the diffraction limited region on the imaged plane would be reduced to 20 μm off axis. For greater errors of 2 mm or 5 mm the system would not be diffraction limited, with the fraction of light on one pixel for a tweezer 15 μm off the optical axis falling to 92% and 56% respectively. Control of this separation is achieved by mounting the V-groove on a 3D translation stage which has its position optimised.

We can see that the idealised performance of 95% of the collected light being on a single pixel is not practically realisable, where a combination of the tolerances of the imaging system and the finite size of the atomic distribution will decrease our performance. The sensitivity of the construction arises from making an airside correction at such a long distance from the aberrant *in-vacuo* lens, and the need to use two corrective optics. The measured performance of the imaging system is described in section [5.3.3](#).

Appendix D

Rydberg electrometry

This appendix details information supporting the methods used in chapter 6.

D.1 Finding Rydberg states

In order to find and identify Rydberg states chosen in section 6.1, one must first calculate the expected values of the energy levels. This is done using the Rydberg Ritz formula [1, 140] (see equations 2.4.1 and 2.4.2), with values for the quantum defect terms that can be retrieved from software [1] or literature [66, 80]. From here we can then calculate the required laser frequencies for the UV laser system. In the work undertaken in this thesis we choose to excite in the scheme $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_1 \rightarrow 5snl\ ^3L_J$ on the 689 nm and 319 nm (UV) transitions (see figure 2.1).

D.1.1 UV laser optimisation

The choice of excitation via the $5s5p\ ^3P_1$ state means that we must use the 1079 nm laser as one of the seed lasers for the Rydberg laser, and then tune the 1549-1565 nm laser in order to choose the target Rydberg state. We choose to tune to the region $n \sim 80$ to have electric field sensitivities on a sensible scale for section 6.1, where the choice of state will be discussed further. This is achieved by coarsely setting the 1549-1565 nm laser wavelength using the grating angle while measuring the output

on a wavemeter¹, and then adjusting the piezo for finer tuning.

In order to achieve good sum frequency generation (SFG) conversion of optical power in the PPLN frequency summing crystal, the poling period and temperature must be optimised for the two IR seed wavelengths used. The PPLN crystal has five different poling periods² which can be accessed by horizontal translation of the crystal with a translation stage. We operate in the fourth poling period for the measurements taken in this chapter, and find an optimal temperature of 152.5°C, within a few degrees of the predicted value for this poling period. We then optimised the fiber coupling to the main optical table, and optimised the second harmonic generation (SHG) in the bowtie cavity to produce as much UV light as possible.

D.1.2 Depletion in a BB MOT

Once the UV laser has been optimised at the target wavelength, one can then begin locating the Rydberg states. Rydberg states can be localised by continuously scanning the UV frequency on the scale of GHz and measuring the ion signal [155]. However as the calculations are quite accurate, and the Rydberg state that we will choose to perform electrometry in has already been seen in previous work [155], we can focus on precisely locating of the feature where the UV is resonant with the atoms.

Initial searching is done by loading a BB MOT, and scanning the detuning of the UV by changing the sideband frequency sent to the 638 nm fiber EOM. This scans the detuning of the UV light incident upon the BB MOT, and resonance is coarsely found by looking for a depletion in the atom number of the BB MOT as a function of UV detuning. By scanning across several depletion features, we once again find the 5s79d 3D_1 , 5s79d 3D_2 , 5s80s 3S_1 , 5s80d 3D_1 , 5s80d 3D_2 and 5s81s 3S_1 Rydberg states, as previously observed [155]. We do not in this case look for the $^3P_{0,1,2}$ Rydberg states, and we do not deliberately apply an external electric field [80, 155]. These states can be identified in terms of the $^{2S+1}L_J$ levels by the

¹Bristol Instruments 671 Series Laser Wavelength Meter.

²The five poling periods of the PPLN crystal are 11.60, 11.65, 11.70, 11.75 and 11.80 μm .

spacing of the Rydberg states, and in terms of principal quantum number n by the energy scale compared to calculation. Further confirmation of the identification of Rydberg states can be done by measuring their properties, such as how the $|m_j|$ levels split when an electric field is applied.

Now that the Rydberg energy levels have been located, measurements beyond this point are no longer taken in the BB MOT. The BB MOT uses 689 nm light that is artificially broadened across a large frequency range, and there is a large quadrupole magnetic field in the vacuum chamber. Both of these things can lead to broadening of our signals, so we do not want to use this for the subsequent precision measurements. In chapter 6 all measurements will be taken in atoms released from an optical tweezer, and the 689 nm light applied will be at a single frequency. Another benefit of using an optical tweezer is that the atoms from a tweezer are localised much more than in a MOT, allowing the measurement of the electric field in a much smaller region.

Appendix E

Towards clock measurements

The magic wavelength tweezer arrays produced in chapter 5 offer a platform to trap atoms with minimal tweezer induced shifting of the clock transition, and the electric field cancellation in chapter 6 measures and reduces a systematic uncertainty in advance of producing an atomic clock. This appendix details the preparation towards making clock measurements, including three key components needed to make an optical atomic clock:

1. The characterisation of the clock laser as the local oscillator.
2. The combined frequency comb and frequency counter system for counting the ‘ticks’ of the clock laser.
3. The planned implementation of the atomic reference in a magic wavelength array, with a MIS (see section 2.3.3) allowed clock transition.

E.1 Clock stability

A stable, narrow-linewidth local oscillator is critical to atomic clock performance in defining the short term frequency noise. We therefore must characterise the stability of the clock laser (see section 3.2.5) when locked to its high finesse cavity, but without any atomic feedback. To measure the frequency of the clock laser we need to beat the clock laser system with an optical frequency comb so that we can measure the

beat frequency between the two on a frequency counter¹. Recording and analysing the beat signal allows us to characterise the stability of the clock laser, as well as checking the operation of the counter itself. This covers two of the three key parts of an optical atomic clock.

E.1.1 Frequency comb

The clock laser and the frequency comb were both introduced in sections 3.2.5 and 3.2.6 respectively. We will now elaborate on how the frequency comb works. First let us consider the general operation of a frequency comb, before considering the specific case of our difference frequency comb.

The output of a frequency comb is produced by a mode-locked femtosecond laser, which in the frequency domain form a series of sharp frequency peaks² separated by f_{rep} , as shown in figure E.1. f_{rep} is given by the round trip time T_{rep} of the mode-locked laser, such that $f_{\text{rep}} = \frac{1}{T_{\text{rep}}}$. This output of a frequency comb is comprised of a series of optical frequencies f_n ,

$$f_n = nf_{\text{rep}} + f_{\text{CEO}}, \quad (\text{E.1.1})$$

where f_{CEO} is the carrier envelope offset. The carrier envelope offset arises from the difference between the phase and group velocities in the oscillator, such that the carrier frequency and envelope frequency have a different phase in subsequent pulses [161]. In order to know the absolute frequency of each of the comb teeth, one must know n , f_{rep} , and f_{CEO} . When the comb output and the clock laser beat together, the resulting beat frequency will then tell us how far the laser is from one of these comb teeth.

To remove the need to know f_{CEO} , a difference frequency comb (DFC) can be used. The specifics of how this works can be found in reference [161], where essentially a highly non-linear fiber broadens the light inside the comb to a super continuum spanning 850 nm to 1880 nm. These two extreme parts of the super

¹The frequency of the clock laser is too high to be directly counted, so a frequency comb is used, as discussed in section 2.3.2.

²The sharp frequency peaks will be referred to as ‘comb teeth’ throughout this section.

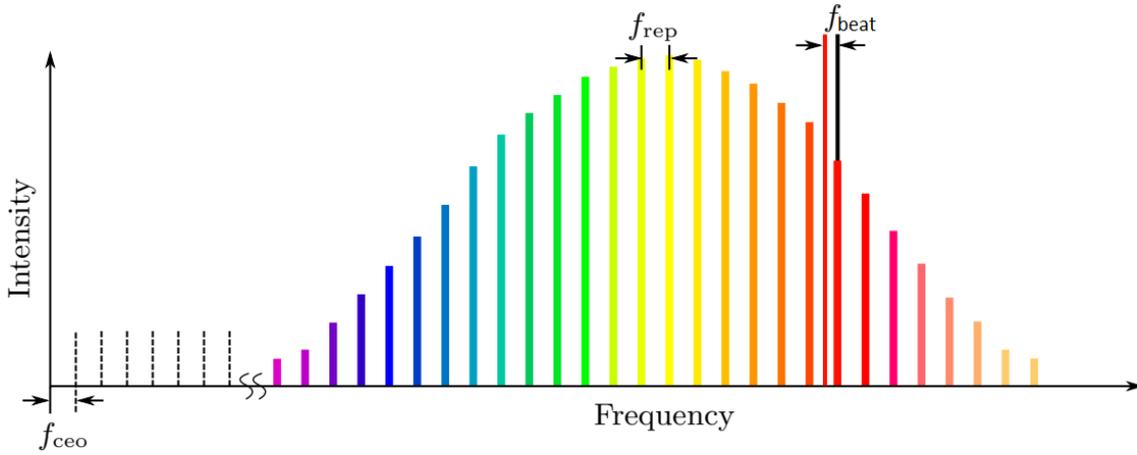


Figure E.1: The output of a frequency comb in terms of frequency. There a regular series of ‘comb teeth’ spaced by $f_{\text{rep}} = 80$ MHz. The clock laser is also shown on here as a single frequency beating with a nearby comb tooth. This figure is modified from a figure in the thesis of Niamh Jackson [80].

continuum are then subtracted using difference frequency generation (DFG). The CEO from each part of the super continuum then cancels, producing light without a carrier envelope offset, simplifying the expression for the comb output to

$$\begin{aligned}
 f_n &= f_j - f_k \\
 &= (j f_{\text{rep}} + f_{\text{CEO}}) - (k f_{\text{rep}} + f_{\text{CEO}}) \\
 &= n f_{\text{rep}},
 \end{aligned}
 \tag{E.1.2}$$

where j , k and $n = j - k$ are integers. This expression means that only f_{rep} and the number of the comb tooth n must be known to find out the absolute frequency of a comb tooth.

The frequency comb used within our research group has previously been used in characterisation of our 689 nm³ and UV⁴ lasers [80, 161]. Our frequency comb has now been extended to also produce an output at 698 nm, allowing us to take a beat frequency between the frequency comb and the clock laser.

³The 689 nm laser is described in section 3.2.2.

⁴The UV (319 nm) laser is described in section 3.2.4, where the frequency comb could beat with the 638 nm light before it undergoes SHG to the UV.

E.1.2 GPS reference

As we saw in the last section, using a difference frequency comb we can compare our lasers against the absolute frequency of the frequency comb output, where to know the frequency of a specific comb tooth we only need to find out n and f_{rep} . Finding n is simple and shall be discussed in the next section, but first we must consider f_{rep} in more detail. The fractional uncertainty in f_{rep} is the same as the fractional uncertainty in the frequency of our comb output. As we are trying to make precision measurements of optical transitions (100s of THz), we cannot afford a large fractional error in the frequency of the comb output. The round trip time of the mode-locked laser (and therefore f_{rep}) must be carefully stabilised against a reliable reference for the comb to be useful.

To stabilise $f_{\text{rep}} = 80$ MHz, it is locked to an oven controlled crystal oscillator (OXCO) at the 10th harmonic at 800 MHz. The oven controlled oscillator provides good frequency stability at short times. The OXCO is then also locked to a GPS reference [198]⁵, which provides long term stability for f_{rep} . Further details of the system layout can be found in previous work [161]. The performance of the OXCO and GPS on different timescales will be discussed in section E.1.4.

E.1.3 Beat measurement

We now have a well controlled set of comb teeth that we know the absolute frequency of. In order to use this to determine the frequency of one of the lasers in our lab, we must beat our laser with one of the comb teeth (as shown in figure E.1), and measure the beat frequency. The clock laser and the comb light are first overlapped using a combo unit⁶. The overlapped light is then fiber coupled to a monochromator unit⁷, which uses a diffraction grating to select light around the wavelength of interest, which is then fiber coupled to a photodiode. Selecting the particular wavelength with the monochromator removes much of the comb light that is not in the nearest

⁵Using a Jackson Labs Fury GPS disciplined oscillator.

⁶Optica DFC BC.

⁷Optica DFC MD.

comb teeth, reducing the excess light to the photodiode that is not contributing to the beat signal.

The photodiode output when viewed on a spectrum analyser shows the beat of the comb with each of the nearest comb teeth, as well as a sharp spike at 80 MHz where adjacent teeth of the comb beat with each other. There will always be two beats between the comb and the laser in the 0-80 MHz range, so we use a 30 MHz low pass filter to remove the higher beat frequency. The single remaining beat frequency is then amplified, and has its signal amplitude and signal to noise ratio (SNR) optimised to ensure that the frequency can be accurately recorded [80]. Here we measure +2 dBm of signal on a -26 dBm background (SNR of 28 dBm), where the bandwidth of the spectrum analyser is set to 510 kHz.

In order to find the absolute frequency of the laser from the beat frequency, one must also send light from the laser in question to a calibrated high accuracy wavemeter⁸. The measured frequency on the wavemeter can then be compared with the beat frequency to work out the integer comb tooth number n , given that

$$f_{\text{laser}} = n f_{\text{rep}} \pm f_{\text{beat}}. \quad (\text{E.1.3})$$

This works as f_{rep} is known to be 80 MHz, f_{laser} is known to high accuracy from the wavemeter, and f_{beat} has been measured on the frequency counter. The corresponding integer value of n and the sign of the beat are then chosen to most closely match the high accuracy wavemeter measurement. f_{laser} is then taken to be defined in terms of equation E.1.3 (instead of using the wavemeter measurement) and is known with very high accuracy. For the measurement of our laser we find ourselves on comb tooth $n = 5365350$ and the \pm sign in equation E.1.3 is taken to be a minus sign.

The beat frequency from the photodiode is sent to a zero dead time frequency counter⁹, which measures the frequency in back-to-back 100 ms measurements. These are then logged via LabVIEW, and record the beat frequency as a function of time. The beat frequency was recorded for several days, where the measurement

⁸Here we use a HighFinesse WSU-30 (IR1). The wavemeter is calibrated against the resonance condition of the 689 nm being set to just prevent a NB MOT forming.

⁹Tektronix FCA3100 Timer/Counter/Analyzer.

was stopped due to sufficient data being collected, rather than any technical limitation¹⁰. The recorded beat frequency is shown in figure E.2, where the linear drift arises from the drift of the clock laser's ULE cavity. This long-term recording of the beat frequency is fitted to with a straight line, which measures the cavity drift rate to be 8.033 ± 0.001 kHz per day¹¹. This cavity drift rate is as expected for a cavity of this type, which would typically be ~ 10 kHz per day, and is very similar to the drift rate of our 689/638 nm ULE cavity discussed in section 3.2.2 and shown in figure 3.14. The change in beat frequency is not perfectly linear, with the residuals from the linear fit to the beat frequency shown in figure E.3.

E.1.4 Overlapping Allan deviation (OADEV)

The long term measurement of the beat frequency between the comb and the clock laser contains significantly more information than may be apparent at first glance. It contains a wealth of information on the noise acting over a variety of timescales, but with over five million data points contributing to figure E.2, all that can be seen is a rough sense of scale of the noise amplitude, and the cavity drift.

To break down the information into a more useful form we use an Allan deviation (ADEV), or more specifically, we will use an overlapping Allan deviation (OADEV). Before considering the different kinds of ADEV, first one may wonder why we do not simply take the standard deviation, s .

$$s^2 = \frac{1}{N-1} \sum_{i=1}^N (y_i - \bar{y})^2, \quad (\text{E.1.4})$$

where N is the number of datapoints in the sample, $y_i = (f_i - f_{\text{beat}})/f_{\text{laser}}$ are the fractional frequency values, f_i is the i^{th} measurement of beat frequency, $f_{\text{beat}} = \frac{1}{N} \sum_{i=1}^N f_i$ is the mean beat frequency, and f_{laser} is defined as in equation E.1.3. $\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i$ is the mean fractional frequency, where convention defines y_i such that $\bar{y} = 0$.

¹⁰The clock laser remained locked to the cavity for over two weeks surrounding the measurement window, and there were no issues with the DFC or frequency counter.

¹¹On timescales relevant for spectroscopy, this is a drift of < 100 mHz per second.

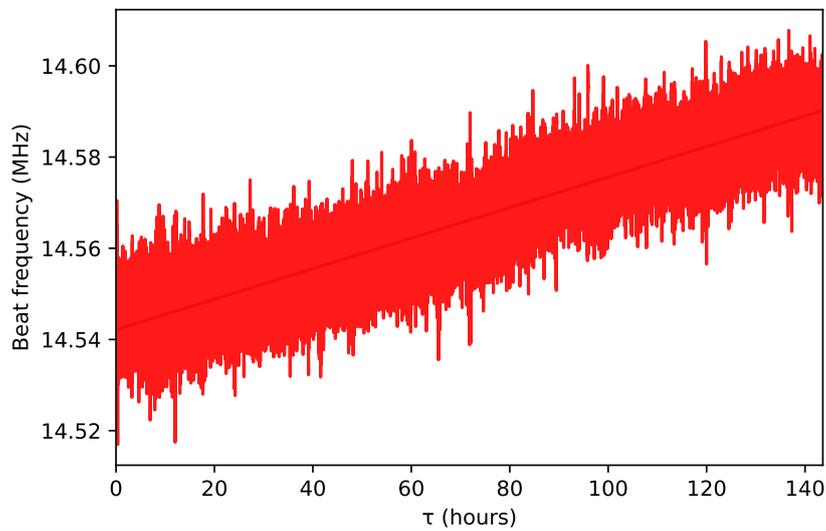


Figure E.2: The recorded beat frequency between the clock laser and the frequency comb over several days. There is an approximately linear drift arising from the drift of the ULE cavity that the clock laser is locked to.

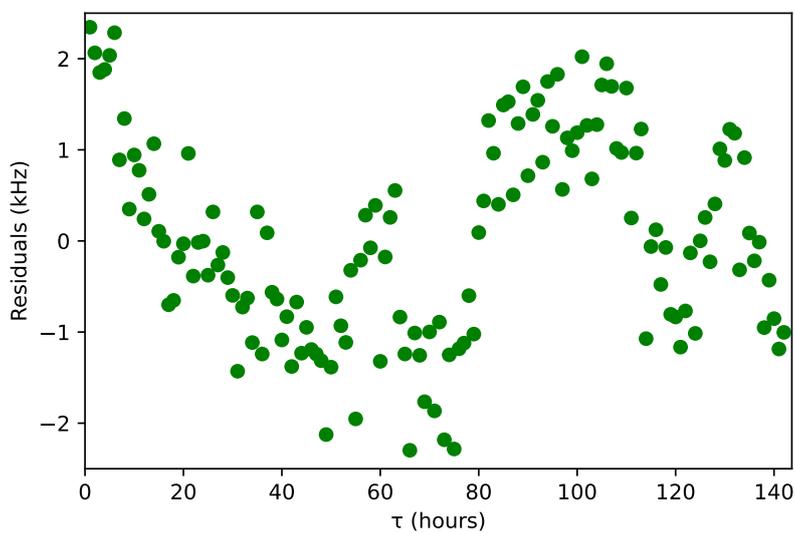


Figure E.3: Residuals for figure E.2 compared to the fitted linear drift. Residuals are taken with bins of 1 hour (compared to the 100 ms bins in figure E.2) to resolve the variations around the long term linear drift of the beat frequency.

The problem with the standard deviation in this context is that the statistical properties of our frequency measurements are dependant on the averaging time used, and the different types of noise that are present. For some types of noise that arise with this kind of measurements, particularly with clocks, the standard deviation is not convergent. An easy way to see this is the cavity drift in figure E.2. As the measurement time increases and the drift continues, the standard deviation is just going to keep on increasing!

To improve upon this we have several kinds of ADEV available [199], which characterise the frequency noise over many different timescales within the measurement. Instead of comparing to the mean value \bar{y} , which depends upon the length of the measurement, we instead compare measurements taken at different times. The basic Allan deviation (ADEV), $\sigma_y(\tau)$, is then expressed as

$$\sigma_y^2(\tau) = \frac{1}{2(M-1)} \sum_{i=1}^N (y_{i+1} - y_i)^2, \quad (\text{E.1.5})$$

where M is the number of fractional frequency values when using a sampling interval τ .

The basic ADEV has since been superseded by alternate versions that have better statistical confidence and can better distinguish between different types of noise. Therefore we use the OADEV, where the different kinds of ADEV are contrasted in reference [199]. The OADEV takes overlapping samples to make the maximum number of available comparisons from the dataset, offering a significant improvement to the statistical confidence. The OADEV, $\sigma_y(\tau)$, shall now be defined for the rest of this thesis as

$$\sigma_y^2(\tau) = \frac{1}{2m^2(M-2m+1)} \sum_{j=1}^{M-2m+1} \left(\sum_{i=j}^{j+m-1} (y_{i+m} - y_i) \right)^2, \quad (\text{E.1.6})$$

where m is defined by $\tau = m\tau_0$, where $\tau_0 = 100$ ms is the basic measuring interval that the data was collected with.

Using an OADEV we can view the frequency stability over different timescales; the OADEV of the beat data is shown in figure E.4. The red trace shows the noise associated with the beat frequency (data in figure E.2) on different timescales. The black trace shows the OXCO/GPS manufacturers specification. The frequency

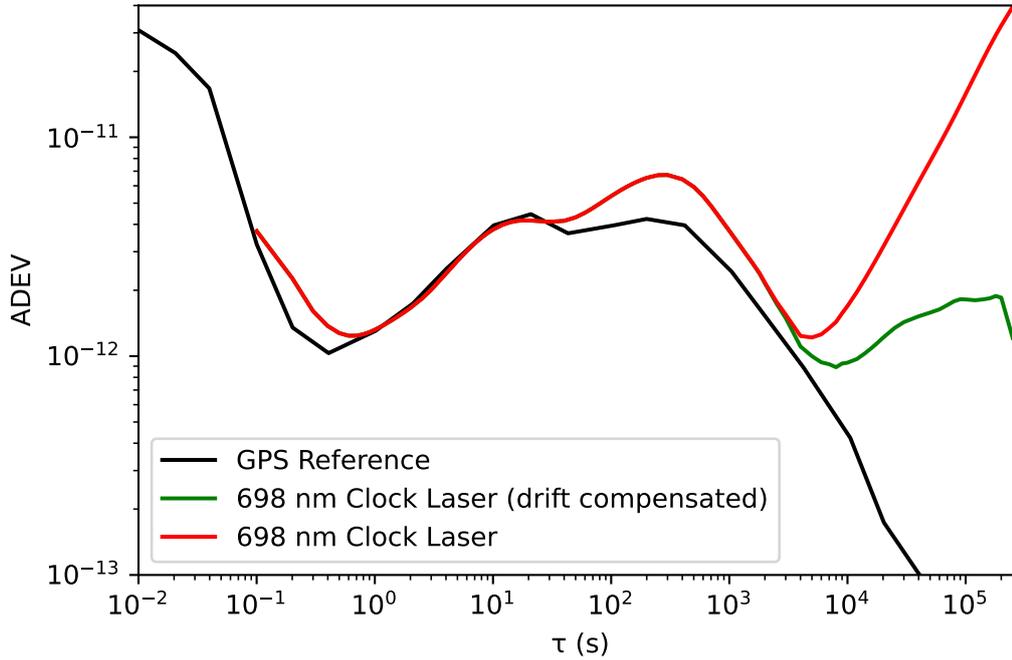


Figure E.4: The OADEV of the clock laser frequency data is shown in red. This is compared to the black trace, which shows the manufacturers specification of the OXCO/GPS [198] that acts as a reference to the frequency comb. The green trace shows how the OADEV changes if one were to post-remove the 8.033 ± 0.001 kHz per day fitted cavity drift from the data.

measurement is incapable of resolving noise on the laser below the black trace, as the noise in f_{rep} dominates, therefore where the red and black traces meet the red trace is the upper bound of how large the noise can be on that timescale. The green trace shows the noise on the beat frequency after cancelling the linear drift of the cavity, where the cavity drift dominates out to long times so the green trace dips below the red trace.

We suspect that the reason the clock laser’s OADEV rises above the GPS reference in the 10^2 s region is due to the GPS reference performance not being as good as the manufacturer specifications, which was not measured on our particular device. OADEVs being limited by the GPS reference above the manufacturer specification have been observed before in our experiment, where the deviation in performance found in characterising our UV laser [155] was later found to be due

to a faulty GPS reference¹². In the future the hypothesis that we are limited by the GPS reference could be checked by taking an OADEV of either the 638 nm or 689 nm lasers that use a different ULE cavity. If they were to show the same trace, it would be reasonable to conclude that the GPS reference would be the common factor resulting in the higher ADEV. It would also be of value to have OADEVs of both lasers for the characterisation of the lasers themselves. The linewidth of the UV laser could be checked to see if it is below the < 35 kHz characterisation [80, 155] done before the most recent maintenance of the frequency comb system. The 689 nm laser was implemented during this work (see section 3.2.2), and has not yet been characterised against the comb.

At timescales of 10^4 s and above in figure E.4, we see a linear¹³ increase in the OADEV with time. This corresponds to the drift of the ULE cavity dominating, as we also saw in figure E.2. To demonstrate that this trend is due to the cavity drift, we redo the analysis on a dataset that is corrected by the 8.033 ± 0.001 kHz linear drift, and plot it as the green trace on figure E.4. As can be seen in the plot, this only changes the trace noticeably on timescales of an hour or greater, and that the noise without the linear part of the cavity drift instead remains down at the 10^{-12} level.

E.1.5 Laser linewidth

To put the fractional frequency noise given by the OADEV into context, it is helpful to consider what that means compared to f_{laser} , as that is the laser frequency we originally sought to understand the properties of. On timescales below 10^4 s (or indeed at all timescales when the cavity drift is removed), the noise does not exceed 7×10^{-12} . On the scale of the clock laser frequency, this corresponds to 3 kHz. This means that we can apply an upper bound to the clock laser linewidth when locked to the cavity of 3 kHz. As mentioned earlier, if the GPS is not performing exactly to specification as we suspect, then the true linewidth will be less than that, and

¹²The GPS reference was replaced before the OADEV in figure E.4 was taken.

¹³Linear trend on this log-log scale.

our measurement is limited by the stability of the frequency comb.

On a timescale of 1 second, our OADEV indicates a fractional instability of 1.3×10^{-12} , corresponding to a linewidth of 0.6 kHz. Here the OADEV corresponds to the expected limit in what we can measure with our GPS reference. This means that the OADEV is at short times is a measure of the performance of our frequency comb, not of the clock laser itself. To bound the performance of the clock laser, the feedback loop can be analysed following the treatment in the thesis of Danielle Pizzey [73], which leads to an upper bound in the laser linewidth of 0.4 kHz.

To further analyse the frequency stability of the laser, a beat measurement should be made with the laser. The two options are to compare the laser with itself after a lengthy (several km) delay line [200], or against another laser at the same wavelength. As we have neither facility available in our laboratory, the latter approach was undertaken by the manufacturer (Menlo Systems) before delivery of the laser system. An optical reference system at 1542 nm had its stability transferred to 698 nm via an ultra-low noise optical frequency comb, where the photodetection beat between our clock laser and the comb was measured. The beat signal was mixed down to 10 MHz, bandpass filtered, and counted on a FXE frequency counter with 10 ms sampling time. After removing a linear drift of 1.03 Hz per second, the OADEV for this showed the fractional frequency instability on a timescale of 1 s to be 1.1×10^{-15} . This indicates a laser linewidth of 0.5 Hz on a 1 s timescale, three orders of magnitude better than we can measure. The limiting factor in this performance is likely to be Brownian motion in the cavity, where there will be thermal fluctuations in the cavity mirror substrates and coatings [201, 202].

We have provided upper bounds of the linewidth of the clock laser, measured the upper bound to be 3 kHz (though 0.6 kHz on a timescale of 1 s) through beating with our frequency comb, and at 0.4 kHz through investigating the feedback loop. Comparisons with another optical reference system by the manufacturer measure the linewidth to be just 0.5 Hz at 1 s. Linewidths this narrow mean that we can expect coherence times of ~ 1 s, so a Rabi frequencies larger than ~ 1 Hz should be sufficient, and the clock laser appears to be a suitable local oscillator. We consider the implications in the choice of Rabi frequency further in appendix E.2. While the

beat measurement with the frequency comb was not the limiting measurement of the laser linewidth (as expected due to the 10 MHz reference), it has shown that the comb is working well as an optical frequency counter, offering one of the other key parts of an atomic clock.

E.2 Clock excitation

We have now detailed that our clock laser is performing adequately as the local oscillator for our clock, and the frequency comb and frequency counter together function well as a counter. We now look ahead to our plans for connecting these to the atomic reference.

E.2.1 Magnetically induced spectroscopy (MIS)

As discussed in section 2.3.3, we use ^{88}Sr which has nuclear spin $I = 0$, so we must apply a magnetic field to allow the clock transition through MIS. The achievable Rabi frequency on the clock transition was given in equation 2.3.2, and is reproduced here as a reminder:

‘The applied magnetic field allows a Rabi frequency on the clock transition, Ω_{clock} of

$$\frac{\Omega_{\text{clock}}}{2\pi} = \alpha\sqrt{IB} \cos\theta. \quad (2.3.2)$$

α is a constant depending on the element being used, here $\alpha = 6.26 \frac{\text{Hz}}{\text{mT}\sqrt{\text{Wcm}^{-2}}}$ [10, 75]. I is the intensity of the clock laser light, B is the uniform external magnetic field used to allow the clock transition, and θ is the angle between the clock laser polarisation and the magnetic field direction.’

Using the MOT coil relay to flip the MOT coils into Helmholtz configuration, we can apply a uniform magnetic field to the atoms. In Helmholtz configuration, a magnetic field per unit current of $12 \text{ G/A} = 1.2 \text{ mT/A}$ is produced. The MOT coils are made from Kapton-dipped single-strand copper wire¹⁴ [72], which are rated

¹⁴Allectra 311-KAP-180-5M.

to ~ 7 A when coiled for continuous currents. If we were to apply 4 A using the maximum power output of the power supply for the MOT coils, described in section 3.3.2, we would expect a magnetic field of $47 \text{ G} = 4.7 \text{ mT}$. To apply currents > 4 A but $\lesssim 7$ A, an alternative current source could be used.

If we were to require coil currents $\gtrsim 7$ A, we could potentially use a pulsed generating circuit [203–205] to apply larger currents but for short periods of time. This approach would unfortunately lower the repetition rate achievable overall, as the time for an excitation pulse scales with $t \propto \frac{1}{\Omega_{\text{clock}}} \propto \frac{1}{B}$, whereas the power into the coils scales as the square of the current through the coils as $P \propto B^2$. This makes the heating per pulse $E = Pt \propto B$, so larger magnetic fields cause more heating per pulse in the coils. As we will see in section E.2.2, the stability of the magnetic field is also a key consideration, where a constant field ought to be more stable.

The clock light should be sent through one of the viewports in the xy plane, polarised in the $\pm \hat{z}$ direction (parallel to the Helmholtz magnetic field) such that $\cos \theta = 1$. Ideally the propagation axis of the clock light will be along a tightly confined axis of the tweezer to minimise the Lamb-Dicke parameter and the effects of Doppler broadening [10, 71, 94]. The axis with tightest confinement that the clock light can be sent along is the $\pm \hat{x}$ direction along the waist of the tweezer.

The clock laser outputs powers of ~ 10 mW. Focusing the clock light to increase the intensity I will increase Ω_{clock} , though at the expense of being less uniform across the array. Taking an example waist size of ~ 0.1 mm with our $47 \text{ G} = 4.7 \text{ mT}$ magnetic field, we would expect to be able to drive a π -pulse on the clock transition in $t_\pi = 2.1$ ms, much shorter than the expected coherence time. For shorter excitation times, the clock light could be focused, decreasing the time inversely proportional with the waist for focusing on both axes, or inversely proportional with the square root of the waist for focusing on a single axis with a cylindrical lens.

E.2.2 Systematic shifts

There are a variety of systematic shifts that will change the clock transition away from its natural frequency, which can be tabulated in an error budget, with as-

sociated uncertainties. There are environmental shifts from both magnetic¹⁵ and electric¹⁶ fields, as well as the black-body radiation (BBR) environment [83]. Applied laser light also induces shifts, namely the trapping (tweezer/lattice) light and clock laser light. There also can be terms arising from collisions, atomic density [206], Doppler shifts, and from the clock servo itself. We now consider the effects of some of these terms.

Chapter 6 characterises that stray electric fields will not be in any way a limiting issue for constructing an atomic clock in our experiment. If we do not cancel the stray electric field the shift due to the electric field is $49 \pm 4 \mu\text{Hz}$ ($1.14_{-0.09}^{+0.09} \times 10^{-19}$), and if we cancel it then it is reduced to $0.11_{-0.11}^{+0.41} \mu\text{Hz}$ ($2.6_{-2.6}^{+9.5} \times 10^{-22}$).

Considering the clock excitation scheme with MIS in optical tweezers, what systematic shifts can we expect to arise from the scheme? We leave several of the shifts for future work, and only consider a selection of the shifts that we will encounter. In particular we now consider the shifts that arise due to the specifics of our excitation scheme in tweezers. These shifts are due to the applied magnetic field, the application of the clock light, and the tweezer light used to trap the atoms.

Magnetic field and probe shifts

The application of the Helmholtz magnetic field from the MOT coils to allow the clock transition through MIS will give rise to a quadratic Zeeman shift,

$$\Delta\nu_B = \beta B^2, \quad (\text{E.2.1})$$

where $\beta = -23.3 \frac{\text{MHz}}{\text{T}^2}$ [75]. For an applied field of $47 \text{ G} = 4.7 \text{ mT}$, this would correspond to a shift of 0.51 kHz (1.2×10^{-12}). This systematic shift must be accounted for, and the variation in this value minimised. For example, fluctuations of the power supply of $< 0.1\%$ would result in an uncertainty arising from variations in the Zeeman shift of $\sim 1 \text{ Hz}$ ($\sim 2 \times 10^{-15}$).

The application of the clock light itself also causes an AC Stark shift of the clock

¹⁵Second order Zeeman shift.

¹⁶DC Stark shift.

transition [128],

$$\Delta\nu_P = \kappa I, \quad (\text{E.2.2})$$

where $\kappa = -18 \frac{\text{mHz}}{(\text{mW cm}^{-2})}$ [75]. For 10 mW of light in a 0.1 mm waist Gaussian beam, this would lead to a shift of 1.1 kHz (3×10^{-12}). It will therefore be important to not have significant fluctuations in the power or beam pointing of the clock interrogation light. For example, fluctuations of the beam intensity of $< 0.1\%$ would result in an uncertainty arising from variations in the probe shift of ~ 1 Hz ($\sim 2 \times 10^{-15}$). Spatial variations in the power over 6.2 μm , which is one tweezer site away from the maximum intensity of the clock light, would then be 0.8% ¹⁷ for a 0.1 mm waist beam. This would correspond to a shift of 9 Hz (2×10^{-14}) between neighbouring tweezer sites, where the shift would be significantly smaller if the clock laser were less tightly focused. For example, if the beam waist were 1 mm then the larger shift would be 11 Hz, with a 1 mHz shift between neighbouring sites.

There is a trade off between the shifts induced by the applied magnetic field and clock light. To reach a desired Rabi frequency in equation 2.3.2, B and I should be chosen to minimise the total uncertainty from the two shifts. As one of the next experimental steps beyond the work in this thesis is to locate the clock transition, having large values of B and I is helpful in offering a large Rabi frequency, and making the transition easier to find [207]. For future clock operation, the Rabi frequency used can be smaller, considering the expected coherence time of our laser, resulting in smaller shifts.

To further reduce the AC Stark shift from the applied clock light, several interrogation schemes are available. Interleaved measurements at high and low probe power can be used to extrapolate to find the frequency of the unperturbed resonance [107], though this relies on how well the ratio of the powers of the two clock probes are known. Beyond this there are hyper-Ramsey protocols [208] which have been experimentally realised [209] to suppress the AC Stark shift, with a discriminant at a frequency unperturbed by the AC Stark shift. Methods of modified/generalised

¹⁷Assuming the tweezer displacement is orthogonal to the direction of the propagation of the clock light, which is the direction with the largest spatial variation.

hyper-Ramsey schemes [10, 210, 211] have been developed to suppress the shift more efficiently, meaning that the probe induced shift from the clock light can be well controlled when using bosonic clocks via MIS.

While the shift from the clock light can be controlled with hyper-Ramsey schemes, there is still a significant quadratic Zeeman shift from the applied magnetic field. To avoid the large Zeeman shifts, an excitation scheme other than MIS could be used. Rather than exciting directly on the $5s^2\ ^1S_0 \rightarrow 5s5p\ ^3P_0$ transition, a multi-photon excitation could be used [71], this scheme avoids the need for a large magnetic field, and resulting shifts from the applied excitation light could also be managed with hyper-Ramsey protocols. We do not consider multi-photon schemes further, as exciting with 698 nm light using MIS is a sensible starting point for constructing a bosonic strontium clock.

Tweezer shifts

The wavelength of the tweezer light, λ_{trap} , is another important source of systematic shift. If we were to perform the clock excitation in a 400 μK deep tweezer, we would expect a shift of -38 ± 3 mHz of the clock transition, per 1 MHz λ_{trap} is away from the magic wavelength [46, 212, 213] (9×10^{-17} /MHz).

When etalon locked, our 813 nm Ti:Sa should have linewidth < 5 MHz, so if correctly set to the magic wavelength we would anticipate variations in shifts on the 0.2 Hz (4×10^{-16}) level. The drift of the tweezer wavelength will also significantly increase the expected shifts, which would cause issues. To reduce the shift due to the wavelength of the tweezer light, the Ti:Sa is wavemeter¹⁸ locked, with absolute accuracy 30 MHz. For a 400 μK deep tweezer this would correspond to shifts of ~ 1 Hz ($\sim 3 \times 10^{-15}$), greater than the expected linewidth of the clock laser.

Steps will need to be taken to reduce the shift of the clock transition due to the tweezer wavelength deviating from the magic wavelength. The shifts could be reduced by improving the locking of the laser, through methods such as: regular calibration of the wavemeter (improving accuracy to < 3 MHz), locking the Ti:Sa

¹⁸HighFinesse WSU-30 (IR1).

to a scanning transfer cavity [179], or locking the Ti:Sa to an ultrastable cavity to improve the stability and linewidth of the trapping light. Another method to reduce the shift due to the wavelength of the trapping light is to adiabatically ramp down the trap depth for clock interrogation [47]. A combination of improved laser stabilisation and reduced trap depth during clock interrogation would offer the best control of the systematic shifts arising from the wavelength of the trapping light.

In addition to deviations away from the magic wavelength, we must also consider the magic wavelength condition itself in more detail. The situation is more complex than is described in section 2.2.3. The polarisability has vector and tensor components, which thankfully for ^{88}Sr are zero due to the lack of hyperfine structure, as discussed in section 2.3.3. There are also higher order shifts that must be considered, in particular the hyperpolarisability.

Stark shifts arise due to off-resonant excitation of two-photon transitions out of the $5s5p\ ^3P_0$ state with the 813 nm trapping light [214], leading to a hyperpolarisability coefficient [215, 216]

$$\Delta\nu_{\text{hyper}} = \left(\frac{U_0}{E_r}\right)^2 \times (0.45 \pm 0.10)\mu\text{Hz}, \quad (\text{E.2.3})$$

where E_r is the recoil energy of an 813 nm photon. For a 400 μK deep tweezer, this shift corresponds to 2.7(0.6) Hz ($6.3(1.4) \times 10^{-15}$). The hyperpolarisability as a function of trap depth can be managed somewhat by detuning the trapping laser [47], and doing clock excitation in shallower traps would also be beneficial in managing the hyperpolarisability shift. Beyond the hyperpolarisability there are also other terms such as the electric quadrupole and magnetic dipole polarisabilities [215] which are minimised for shallower traps and atoms in lower motional states.

E.3 Summary

We have characterised the clock laser to be stable to the level resolvable using our GPS referenced frequency comb. The linewidth of the clock laser is confirmed to be ≤ 3 kHz on all timescales, after accounting for the measured 8.033 ± 0.001 kHz per day cavity drift. The clock laser should act as a suitable local oscillator, with < 1

Hz demonstrated by the manufacturer during testing. We have also detailed the operation of the frequency comb and frequency counter, which will form a second key part of future work towards an atomic clock.

We have then considered plans for the clock excitation scheme with the atomic reference, noting the effects of different systematic shifts. Shifts will unavoidably arise from the Helmholtz magnetic field used to allow the clock transition through magnetically induced spectroscopy (MIS), and will need to be well characterised. The stability of the intensity of the applied clock laser, and the wavelength of the tweezers, will also be of key concern. The former can be mitigated through different interrogation schemes, and the latter suggests that we should use lower trap depths and consider improving the laser locking of our trapping light. We are now in a position to work towards finding the clock transition in our system, and to continue towards implementing a tweezer array based atomic clock.