Scalable Analogue Optimisation On A Neutral Atom Array

Elliot Diamond-Hitchcock

A thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

Department of Physics
University of Strathclyde

December 12, 2024

Scalable Analogue Optimisation On A Neutral Atom Array

Elliot Diamond-Hitchcock

Abstract

This thesis describes the development and use of a neutral atom quantum computing system, in particular detailing the methods to create a useful quantum computing platform that can efficiently scale to larger system sizes. The Scalable Qubit Arrays (SQuAre) platform can trap and deterministically arrange over 100 caesium atoms into arbitrary array geometries. These atom arrays can then be used for performing either digital or analogue quantum computation, or for quantum simulation.

Blue-detuned optical dipole traps are demonstrated, using novel holographic techniques to generate arrays of optical potentials capable of coherently trapping arrays of Rydberg atoms. Simulations presented in this thesis are used to confirm the viability of this scheme and find optimal trapping parameters. Experimental results then show blue-detuned optical dipole potentials can be used to trap atoms excited in the Rydberg state for long periods otherwise prohibited by experiments relying on applying gate operations in free-fall.

Graph optimisation problems naturally map onto Rydberg atom quantum computing systems by mapping onto a unit disk graph. A new approach to implement programmable local-light shifts is presented, enabling the first demonstration of Maximum Weighted Independent Set (MWIS) problems on a neutral atom array using annealing. These results provide a route to solving wider classes of classical optimisation problems on arbitrary graphs by encoding onto MWIS, opening the possiblity for near-term application of neutral atom quantum computing to real-world problems.

Contents

List of Figures

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.

Elliot Diamond-Hitchcock
Glasgow, December 12, 2024

Acknowledgements

The list of acknowledgments expands to fill the space available. It's been somewhere between four and twenty-eight years to get here, so I could probably double the word count if I don't start already.

Endless thanks have to go to the SQuAre team who started the project with me and will continue the work after I've left. My contributions have only been what feels like a small part of everything we've done, and a far smaller part of what is to come! The postdocs Nick Spong and Jonny Bass and the former postgradute student (and fellow laser troll) Boyko Nikolov who started the team have moved on to other work, but this thesis feels like it should be at least three-quarters theirs. I can only wish best of luck to Andre de Oliveira, Dan Walker, Maximillian Wells-Pestell and Paul Schroff who are left to keep things moving. If you ever ask "why did they do this?", I'm afraid only Jon's left to say why. Naturally, enormous thanks to Jon for his teaching, supervision, feedback and trust at every step. Finally, a big thanks to everyone at EQOP for helping make such a constructive environment and to all the Strathclyde support staff for keeping this place ticking along.

In what feels like a brag, I think I've too many friends and colleagues to mention here, but I'm sure you all know who you are, in Berkshire, Buckinghamshire, Leeds, Montreal, Glasgow or wherever you've ended up. Special thanks to the folks over at Leeds, some of which passed on their own PhD experience, either as warnings I clearly didn't listen to, or lessons when I sometimes did.

Finally, huge thanks to the parents and the sister for the support and the pet photos every step of the way, even when I thought your confidence was misplaced. Huge thanks to the extended family as well, all the cousins, aunts, uncles, gruncles, the met and unknown grand and great-grandparents who all got me here. A final thank you to Uncle James, who pointed me in this direction so many years ago.

I couldn't have done it without you all.

Chapter 1

Introduction

If quantum physics and information technology each independently revolutionised the 20th Century, quantum information technology promises to revolutionise the 21st. While the cost of classical computation is lower than ever, at rates faster than ever, exponential improvements in performance are finally beginning to slow down [1]. The challenges of fabricating ever smaller classical electronics are beginning to reach fundamental limits. At the same time, the demand for large-scale computing continues to grow in order to tackle problems ranging from material simulations to quantum chemistry.

Quantum computation and simulation provide an alternative to classical methods. Classical simulation of quantum systems scale poorly with system size, where the number of classical bits required for the model increases exponentially with the number of quantum particles. Calculations on this level rapidly become hard to simulate using classical methods, even for modest system sizes. By transitioning from classical to quantum hardware, we can more efficiently and effectively solve a quantum problem. Quantum chemistry expects incredible benefits from the direct simulation of atomic and molecular structure [2], with further applications in materials science [3] and simulating biological processes [4]. Shor's algorithm proved that even some classically hard problems, such as integer factorisation [5] or searches of an unstructured database [6], can expect a quantum speedup.

This thesis focusses on the design and construction of a new platform for computing based on arrays of trapped neutral atoms. Below, an overview of the foundational concepts of quantum computing are introduced alongside the unique advantages of neutral atom quantum computing. Finally, the Scalable Qubit Array (SQuAre) platform is introduced, which is developed over the course of this research.

1.1 Quantum Information

The nature of classical computing is bound by classical physics. One classical bit of data must be either '0' or '1', containing one bit of entropy[1], while N bits contain N bits of entropy, following classical thermodynamic principles [7].

An analogue two level quantum system forms a quantum bit. Quantum bits, or qubits, do not follow these laws. The '0' or '1' states are not exclusive, as the quantum bit can exist in a superposition of both. The state of the qubit is encoded in a wavefunction which can be described by its amplitude and phase. For an N-qubit system, entanglement between qubits leads to further non-classical behaviour, such as negative partial information [8]. By these quantum processes, fewer than one hundred quantum bits can represent states that would require trillions of classical bits to encode. Quantum systems therefore offer efficient storage and manipulation of information, although only a single state can be measured at a time.

To make a perfect simulation of any finite quantum process, the universal computer proposed by Church and Turing must necessarily be itself quantum in nature [9]. As introduced above, this theoretically perfect quantum computer would be able to simulate quantum problems that are classically

---

[1] in base-2

intractable for any realisable classical computer, while also providing new paths to efficiently solve otherwise hard classical problems.

Making such a quantum computer is not a simple task however. The requirements to simulate a quantum system summarised by the DiVincenzo criterea each introduce problems that must be overcome [10]. Firstly, a scalable number of well-characterised qubits is required, which can be initialised to a simple, well-described state. These qubits must be coherent over long times and implement a universal set of gates, and be accurately and individually measured. Further criterea are required for the networking of quantum systems using flying qubits.

The current state of the art is referred to as the 'Noisy Intermediate Scale Quantum' (NISQ) era of quantum computing [11]. Qubits have been generated, entangled, manipulated and measured at greater scales than ever before, but quantum computers must grow far beyond their current scale to outperform classical methods. Noise in state preparation, gates and readout has so far limited systems to 100s of qubits, orders of magnitude below what is required for large-scale computation which provides a conclusive advantage over classical systems. Quantum error correction is necessary for scaling to even larger system sizes, which itself introduces additional overhead and complexity. Plenty of work remains in order to scale quantum computers to this scale, much like the generational efforts to scale classical computing to where it is today.

1.2 Quantum Computation and Simulation

It is important to specify what is meant by a quantum computer and a quantum simulator. An **analogue quantum simulator** replicates the Hamiltonian, or a model of the Hamiltonian, that is to be investigated. These can reproduce dynamics that classical simulators would otherwise struggle with [12]. Just within the field of atomic optics, this has been implemented in Bose-Einstein condensates [13–15], optical lattices [16–19] and ultra-cold Fermi gasses [20, 21], to name a few. **Analogue quantum computing** takes a set of continuous input variables, and uses the evolution of a quantum system to produce a measurable output, the output of a defined input. Analogue quantum computers can be used to implement analogue quantum simulations [22, 23]. Programmable analogue quantum computing is shown in [24] and in Chapter 7 of this work. **Digital quantum computing** uses discrete inputs, where qubits are typically manipulated by discrete quantum gates. Digital quantum algorithms have been shown to scale better than the best known classical algorithms in some cases [5], and can also tackle quantum many-body problems through direct simulation [25]. The 'quantum advantage' over classical digital computing has already been claimed for specific calculations [26–28].

Digital quantum computing has been shown to be equivalent to the analogue method of **adiabatic quantum computing** [29]. Here, the ground state of simple Hamiltonian is prepared. The Hamiltonian is adjusted adiabatically to a final Hamiltonian, such that the ground state of the final Hamiltonian is then prepared. The solution to the problem of interest is encoded in this ground state, and it is using this method that Chapter 7 demonstrates analogue quantum computing.

Without a clear and obvious winner in the field of quantum computation and simulation (yet), many branches of research are tackling the problem of constructing this quantum technology infrastructure concurrently. Superconducting systems have already demonstrated their potential for quantum advantage [26, 28], due to suitable combination of fast operation and the utility of leveraging current nanofabrication techniques used in solid-state engineering. Transmon qubits have demonstrated coherence times > 0.5 μs [30], while readout times of < 100 ns have been achieved with fidelity $F > 0.99$

[31, 32]. Single qubit gates have measured errors $\epsilon < 10^{-4}$, and two-qubit gates have measured errors $\epsilon < 10^{-2}$ [33, 34]. Commercial systems from D-Wave use thousands of interconnected superconducting qubits for quantum annealing, while IBM's roadmap expects commercial chips with hundreds of superconducting logical qubits within the next decade, among a number of commercial competitors [35–37]. These do require cryogenic environments, while the fabrication of each superconducting circuit that makes up each qubit is susceptible to errors and inhomogeneities.

Trapped ions have recorded the best coherence times yet, exceeding one hour for a single qubit[38]. Single qubit and two qubit gate errors have reached $< 10^{-5}, < 10^{-3}$ respectively [39–42], better than any other architecture, but scaling to larger system sizes remains a concern [43]. Ionic interactions depend on $1/r^2$ Coulomb interactions, which remain strong even over long distances. Chains of ions cannot simply get longer to scale to greater qubit numbers and so this fidelity reduces with system size [44]. Photonic systems have also claimed quantum advantage [27] and have applications as quantum simulators [45] and scalable fault-tolerant universal computers [46]. Other architectures such as nitrogen vacancy centres [47, 48], alkali molecules [49, 50] and quantum dots [51] have all been proposed and are under active development for use as quantum computers and/or simulators.

1.3 Neutral Atom Quantum Computation and Simulation

Not yet introduced is the use of neutral atoms as qubits. This work will focus on neutral atom quantum computing, and specifically present results regarding techniques for use in analogue quantum computing on neutral atom arrays.

Neutral atoms provide a excellent framework for quantum information science [52, 53]. Each atom of a given isotope is identical, while atomic interactions with light can be manipulated with exceptional precision. Neutral atoms experience only weak interactions while in the ground state, leading to long ground-state coherence times. Non-destructive readout fidelity has been demonstrated at record levels $F > 0.999$ [54], and large scale experiments have been demonstrated with coherence times > 12 s on arrays of > 6100 atoms at once [55]. This property led to atomic clocks and the redefinition of the second that still holds today, and is a key part of a neutral atom quantum computing system.

Neutral atom quantum computing structures typically involve the construction of an array of tweezer traps, each containing a single atom. This was first shown in a single trap in 2001 [56], in a 1D chain in 2016 [57], in 2D arrays the same year [58] and in 2018 for arbitrary three-dimensional structures [59]. Interactions between neutral atoms are mediated through excitation to Rydberg states [52]. The atoms used for these exploit the hydrogenic properties of alkali metals, although promising developments have been made in the use of alkali-earth elements [60–62]. These Rydberg atoms have exaggerated properties, with strong electric dipole moments and strong sensitivity to these same moments. Excitation to the Rydberg state can therefore be used to switch on and off an atoms sensitivity to its neighbours. The finite lifetime of the Rydberg state remains the limit for the scheme, as accessible Rydberg states have lifetimes on the scale < 1 ms even in cryogenic environments [63].

Current experimental systems are rapidly developing. Global single qubit gates have been experimentally demonstrated with errors $\epsilon < 10^{-4}$ [64] on arrays of neutral atoms, with two qubit gate errors $\epsilon < 10^{-2}$ [65] on up to 60 atoms at once. As gate errors reduce to the level $\epsilon < 10^{-3}$ as required for efficient fault tolerant gate schemes [66], neutral atom quantum computing is a promising and rapidly developing avenue of research. Experimental results have shown the dynamically reconfigured rearrangement of qubits without loss of coherence [67], with logical topological qubits for fault-tolerant encoding [68]. Small-scale algorithms for effective computation on relatively small system sizes have been measured [69], while mid-circuit measurements expand the class of algorithms available [70]. Reservoirs of 108 atoms have been used as a quantum co-processor for enhanced machine-learning [71] and qubits have been rearranged in over 2000 sites [72].

Neutral atoms are capable of more than digital quantum computing operations. Rydberg interactions naturally model many-body physics and can be exploited for quantum simulation in many varieties, as shown in [73]. Ising or Ising-like Hamiltonians map well to Rydberg systems [74, 75] while recent results have used Rydberg atom arrays as simulators of dipolar 2D XY [76, 77] and XXZ [78] Hamiltonians, dipolar Heisenberg models [79]. Topological phase transitions in Rydberg atom arrays has been shown in [80, 81], and topological spin liquids in [82]. Work on topological order is mentioned above in its use to encode topologically protected qubits.

These analogue methods can also be used to solve classical problems. Adiabatic quantum computing has been used to solve integer factorisation problems utilising a mapping to graph problems [83]. Graph optimisation problems, such as finding the maximum independent state (MIS), naturally map onto neutral atom systems, as first shown in large graphs in [84]. Mapping to weighted graphs was first theorised in [85], allowing to solve for the maximum weighted independent set (MWIS). This has since been experimentally demonstrated, as shown in [86] and in this thesis.

The Scalable Qubit Arrays (SQuAre) experiment described in [87] and this thesis is a neutral atom quantum computation and simulation experiment. As shown in [64, 87], this platform is capable of trapping > 100 atoms and implementing high-fidelity global gate operations. This thesis intends to demonstrate scalable techniques implemented on analogue quantum computation, with further uses in quantum simulation and digital computation.

1.4 Publications Arising From This Work

de Oliveira, A. G., *et al.* "Demonstration of weighted graph optimization on a Rydberg atom array using local light-shifts." arXiv preprint arXiv:2404.02658 (2024).

Nikolov, B., *et al.* "Randomized benchmarking using nondestructive readout in a two-dimensional atom array." Physical Review Letters, 131(3), p.030602. (2023).

1.5 Contributions Of Others

The design and plan of much of the experimental hardware was carried out by Prof. J. D. Pritchard, the principal investigator of this experiment, with further contributions from N. Spong (April 2020 - June 2022). Hardware and software implementation of the Artiq system was carried out by both these researchers as well as the infrastructure for experimental logging and analysis.

Initial experimental construction was carried out by these researchers, alongside B. Nikolov (July 2020 - November 2023) and the author. B. Nikolov contributed substantially to the assembly of the MOT lasers and optical systems, as well as characterisation of the laser cooling and vacuum systems. Special care is given in chapter 3 to note the breadth and depth of his work and discussion of experimental hardware and his thesis [87] should be considered a necessary companion to this work. Finally, measurements shown in Sec. 5.3 are covered in both theses.

J. Bass (January 2021 - December 2023) constructed much of the initial Rydberg laser system, including the ultra-low expansion (ULE) cavity locking system for frequency stabilisation. Much of the experimental data taken in Sec. 6.4 and Sec. 6.5.2 represents dedicated work by J. Bass. Writing on this experiment is currently in preparation for publication.

A. G. de Oliveira (September 2022 -) developed experimental controls for the arbitrary wave generator (AWG) hardware, alongside leading the experimental team in data collection for results seen in Chapter. 7.

D. M. Walker developed the Sidewinder algorithm for atom sortin and its implementation to experimental software.

All the above researchers contributed to the installation and maintenance of various parts of the experimental system. Much of the research presented in this work is a result of collaborative work from the entire experimental team, including PhD student M. Wells-Pestell (October 2023 -).

G. Pelegrí collaborated in the theoretical basis of work seen in Chapter 7, alongside writing and providing code for simulating these systems.

Chapter 2

Cooling, Trapping and Coherent Control of Atoms

Encoding and storing binary information is simple enough to be possible with punch cards, while any classical logic can be performed through digital electronics [88]. Control of quantum information, however, requires quantum systems. Quantum information must necessarily be stored on quantum objects, in our case quantum bits.

As introduced earlier, there are a wide range of quantum systems which can act as our quantum bits. Encoding this process on a wavefunction as required for a quantum device is here done on the hyperfine structure of an atom, effectively using utilising the quantised angular momentum of the Caesium-133 (Cs) ground state wavefunction to encode our off-on state. This wavefunction, as described in Sec. 2.1.1, introduces our choice of quantum bit.

Control of the atomic wavefunction requires precise coupling to the system, where the rate and duration of the interaction is crucial, as well as the interactions with the complex phase of the wavefunction. Lasers can generate precise, coherent oscillating electric fields that can address the atomic wavefunctions. Furthermore, the amplitude and phase of the fields can be switched over the nanosecond scale. Microwaves can directly drive the transitions between the hyperfine states, and have long since been used as a stable frequency reference [89]. Addressing these transitions optically requires control of a frequency difference, but are substantially faster and can be focussed to a single atom. Sections 2.1, 2.2 and 2.3 introduce how these oscillating fields interact with the atoms, and how these atoms interact with each other.

As shown through the rest of this chapter, the light used in this experiment can be used to cool atoms to microkelvin temperatures (Sec. 2.4) and trap them individually (Sec. 2.5). The end result is an array of cold atoms, forming the basis preparation of the experiments detailed in later chapters.

2.1 Two-Level Atomic Physics

2.1.1 Hyperfine Cs Structure

The hyperfine states of Cs are used as the quantum information basis. In this basis, an atomic state is given by the quantum numbers $|n, L, J, F\rangle$, referring to the principal orbital, orbital angular momentum, spin-orbit and hyperfine quantum numbers respectively. All these quantum states are energetically distinct, thereby lifting the degeneracy of the system. Furthermore, in a small magnetic field B_z the degeneracy of the measured angular momentum projection is also lifted along the direction of the magnetic field, introducing an energy shift

$$\Delta E_{|F,m_F\rangle} = \mu_B g_F m_F B_z \tag{2.1}$$

where μ_B is the Bohr magneton and g_F is the Landé hyperfine g-factor [90]. This is the anomalous Zeeman effect, which lifts degeneracy of the m_F states in a weak magnetic field.

The key details required for the atomic structure of Caesium-133 (hence-

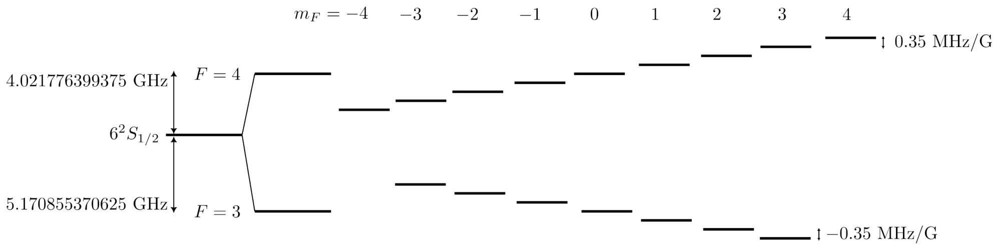

Figure 2.1: Diagram of the $6S_{1/2}$ hyperfine structure of Cs-133. The ground state is split into two states, $F = 3$ and $F = 4$. This energy splitting is exact, due it its use in the definition of the second [91]. This is further split in the presence of a magnetic field into separate m_F levels.

forth referred to simply as Cs) are that it has ground-state atomic quantum number $n = 6$, and nuclear spin $I = 7/2$. The $|n, L, J\rangle$ state is referred to as the $n^{2S+1}(M)_J$, where M is the magnetic quantum number notation for the L state, and S is the electron spin quantum number (1/2). Thus, the $|n = 6,\ L = 0,\ J = 1/2\rangle$ state that is the Cs ground state is written $6^2S_{1/2}$. This state has two hyperfine levels, $F = I \pm S = 3,\ 4$. The transition between $|F = 3\rangle$ and $|F = 4\rangle$ states can be achieved globally using microwave pulses in a single photon transition.

The degeneracy of these levels can be lifted with a weak magnetic field. The $|F = 3, m_F = 0\rangle$ and $|F = 4, m_F = 0\rangle$ states are insensitive to small fluctuations in magnetic field[1], and are thus used as the computational basis states. Note that these states are often referred to as the 'clock states', due to their usage in atomic clocks[2].

For this work, the $|F = 3, m_F = 0\rangle$ and $|F = 4, m_F = 0\rangle$ will be referred to as the $|0\rangle$ and $|1\rangle$ states respectively. These are the basis eigenstates of each qubit.

---

[1] Second order effects are present, but in this case are negligible [92].

[2] The transition frequency between these two states of 9.192631770 GHz is exact, due to its use as the SI definition of the second.

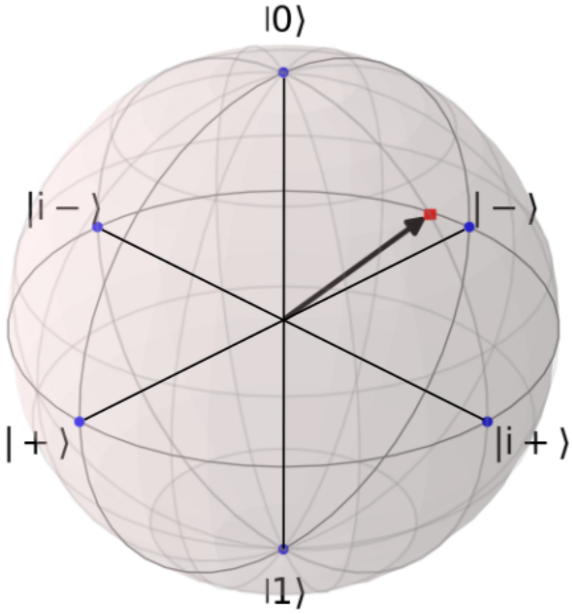

Figure 2.2: A diagram of the Bloch sphere, with blue circles indicating the $x(|\pm\rangle)$, $y(|\mathrm{i}\pm\rangle)$ and $z(|0\rangle, |1\rangle)$ axes (states). The black vector refers to the state $|\psi\rangle = \cos(\pi/8)|0\rangle + \mathrm{i}\sin(\pi/8)|1\rangle$, and is marked on the surface of the sphere with a red square.

2.1.2 Bloch Spheres

Given a system in some quantum state $|\psi\rangle = c_0|0\rangle + c_1|1\rangle$, a useful representation of this system is through the so-called 'Bloch sphere'. Analogous to the Poincaré Sphere [93, Chapter 4], this is a representation of all possible superpositions of a two-level system with complex amplitudes.

As seen in Fig. 2.2, a state vector in the plane given has length 1 when normalised. Along the z-axis, the state vector is in the $|0\rangle$ state for positive z, and $|1\rangle$ for negative. Along the x-axis, the state is in an equal superposition of the two, $|\psi\rangle = \frac{|0\rangle \pm |1\rangle}{\sqrt{2}} \equiv |\pm\rangle$. The y-axis introduces the relative phase between the two basis states, such that along the y axis $|\psi\rangle = \frac{|0\rangle \pm i|1\rangle}{\sqrt{2}} \equiv |i\pm\rangle$.

Note that a series of pulses and techniques will often refer to the Bloch sphere picture. For example, a $\pi/2$-pulse refers to a rotation of angle $\pi/2$ about the y-axis. The axes are indeed chosen to match Pauli matrices, such that a Pauli-X gate matches a rotation of π about the x-axis, and so on. For a single two-level system, all possible states and operations are may be visualised on the Bloch sphere, and it retains its usefulness into more complex systems. This can be understood as a projection of the Hilbert space onto spherical coordinates. For some wavefunction $|\psi\rangle$, the coordinates are given by

$$|\psi\rangle = r\left(\cos\left(\frac{\theta}{2}\right)|0\rangle + \exp\left(i\phi\right)\sin\left(\frac{\theta}{2}\right)|0\rangle\right) \tag{2.2}$$

The normalised wavefunction has a magnitude of the sphere's radius, such that $r = 1$. The altitude θ is proportional to the relative magnitude of each eigenstate, and the azimuth ϕ proportional to the relative phase.

2.1.3 Rabi Oscillations

Our first goal is to demonstrate how light interacts with a two-level atomic system. This will serve as the basis for much of the physics for atom-light interactions. A generic two-level quantum system can be described by the wavefunction

$$|\psi\rangle = c_0|0\rangle + c_1|1\rangle\,, \tag{2.3}$$

where c_i is a complex amplitude and $|i\rangle$ is an electron orbital eigenstate of the atomic Hamiltonian

$$\mathcal{H}_0 = \hbar\left(\omega_0|0\rangle\langle 0| + \omega_1|1\rangle\langle 1|\right). \tag{2.4}$$

In the density matrix formalism, this is given by

$$\rho = |\psi\rangle\langle\psi| = \begin{pmatrix} |c_0|^2 & c_0c_1^* \\ c_1c_0^* & |c_1|^2 \end{pmatrix} \equiv \begin{pmatrix} \rho_{0,0} & \rho_{0,1} \\ \rho_{1,0} & \rho_{1,1} \end{pmatrix}. \tag{2.5}$$

The energy difference between the two states $\omega_1 - \omega_0 = \omega_{0,1}$ is the characteristic transition frequency. We can rescale the Hamiltonian such that the $|0\rangle$

eigenstate is considered zero energy, such that

$$\mathcal{H}_0 = \hbar\omega_{0,1} \left|1\right\rangle \left\langle 1\right| . \tag{2.6}$$

Next, we must consider how this system interacts with an oscillating electric field of frequency ω. The electric field $\mathbf{E}$ is described by

$$\mathbf{E} = E_0 \cos(\omega t + \phi)\hat{\epsilon}_i, \tag{2.7}$$

where E_0 is the electric field amplitude, t is time, ϕ is some arbitrary phase and $\hat{\epsilon}_i$ is the unit vector determining the polarisation of the incident light.

Interactions of light with the atom are dominated by interactions with the atomic dipole moment [94, 95]. This field interacts with the atomic dipole moment $\mathbf{d}$ to form the atom-light interaction Hamiltonian

$$\begin{aligned}\mathcal{H}_{\mathrm{AL}} &= \hat{\mathbf{d}} \cdot \mathbf{E} && (2.8)\\ &= E_0 \cos(\omega t + \phi) \sum_{i,j} \left\langle i\right| e\mathbf{r} \cdot \hat{\epsilon}_i \left|j\right\rangle \left|i\middle\rangle\middle\langle j\right| && (2.9)\\ &= E_0 \cos(\omega t + \phi) \sum_{i,j} X_{i,j} \left|i\middle\rangle\middle\langle j\right| , && (2.10)\end{aligned}$$

where e is the electron charge and $\hat{\mathbf{d}}$ is the dipole moment operator. Here, $X_{i,j}$ is the transition dipole moment along the axis of polarisation, where $X_{i,j} = X^*_{j,i}$. This vanishes when $i = j$.

As also shown in [94, Chapter 2], $\hat{\mathbf{d}}$ is non-zero only on the off-diagonal components of the two-atom Hamiltonian. We can thus condense the total Hamiltonian $\mathcal{H}$ to

$$\begin{aligned}\mathcal{H} &= \mathcal{H}_0 + \mathcal{H}_{\mathrm{AL}} && (2.11)\\ &= \hbar\omega_{0,1} \left|1\middle\rangle\middle\langle 1\right| + E_0 \cos\left(\omega t + \phi\right) \left(X_{0,1} \left|0\middle\rangle\middle\langle 1\right| + X^*_{0,1} \left|1\middle\rangle\middle\langle 0\right|\right) && (2.12)\\ &= \hbar \left(\omega_{0,1} \left|1\middle\rangle\middle\langle 1\right| + \cos\left(\omega t + \phi\right) \left(\Omega_{0\to 1} \left|0\middle\rangle\middle\langle 1\right| + \Omega_{1\to 0} \left|1\middle\rangle\middle\langle 0\right|\right)\right) , && (2.13)\end{aligned}$$

where

$$\Omega_{i\to j} \equiv \frac{E_0}{\hbar} X_{i,j}. \tag{2.14}$$

$\Omega_{i\to j}$ is the Rabi frequency of the transition, and depends on the transition dipole moment and electric field amplitude. Furthermore, $\Omega_{i\to j} = \Omega^*_{j\to i}$

It should be noted that the transition dipole moment is related to the natural linewidth $\Gamma_{i,j}$ of the transition by

$$\Gamma_{i,j} = \frac{\omega_{i,j}^3 |X_{i,j}|^2}{3\pi\epsilon_0\hbar c^3}, \tag{2.15}$$

where ϵ_0 is the permittivity of free space and c is the speed of light [94, Chapter 2].

To find how a wavefunction evolves over time under this Hamiltonian, the time-dependent Schrödinger equation is used. For some wavefunction $|\Psi(t)\rangle$ under some Hamiltonian $\mathcal{H}$, the wavefunction evolves as [90, Chapter 2],

$$i\hbar\frac{\partial}{\partial t} |\Psi(t)\rangle = \mathcal{H} |\Psi(t)\rangle . \tag{2.16}$$

In the density matrix formalism, we use the equivalent Liouville equation,

$$\dot{\rho} = \frac{i}{\hbar} [\rho \mathcal{H}] , \tag{2.17}$$

where the square brackets indicate the commutator operation.

These evolution dynamics still apply under unitary transformations of the Hamiltonian, $\mathcal{H}_T = U\mathcal{H}U^\dagger + i\dot{U}U^\dagger$, where $\mathcal{H}_T$ is the transformed Hamiltonian.

We thus apply the transformation

$$U = |0\rangle\langle 0| + \exp(i\omega_{0,1}t) |1\rangle\langle 1| , \tag{2.18}$$

to $\mathcal{H}_{\mathrm{AL}}$ with an arbitrary shift of the zero energy point to find the transformed

Hamiltonian

$$\begin{aligned}\mathcal{H}_T &= \hbar\omega_{0,1}\left|1\middle\rangle\middle\langle 1\right| \\ &+ \frac{\hbar\Omega_{0\to1}}{2}\left(e^{i((\omega-\omega_{0,1})t+\phi)} + e^{-i((\omega+\omega_{0,1})t+\phi)}\right)\left|0\middle\rangle\middle\langle 1\right| \\ &+ \frac{\hbar\Omega_{1\to0}}{2}\left(e^{-i((\omega-\omega_{0,1})t+\phi)} + e^{i((\omega+\omega_{0,1})t+\phi)}\right)\left|1\middle\rangle\middle\langle 0\right|.\end{aligned} \tag{2.19}$$

When $\omega, \omega_{0,1}$ are of the same scale, such that $|\omega+\omega_{0,1}| \gg |\omega-\omega_{0,1}|$, we can make the approximation that terms of $\exp\left(i\left(\omega+\omega_{0,1}\right)t\right)$ rapidly average to zero. As such, we eliminate these terms from our Hamiltonian. This is the rotating wave approximation, and in the transformed frame is thus

$$\begin{aligned}\mathcal{H}_T^{\mathrm{RWA}} &= \frac{\hbar\Omega_{0\to1}}{2}\left(e^{i(\Delta t+\phi)}\right)\left|0\middle\rangle\middle\langle 1\right| \\ &+ \frac{\hbar\Omega_{1\to0}}{2}\left(e^{-i(\Delta t+\phi)}\right)\left|1\middle\rangle\middle\langle 0\right|,\end{aligned}$$

where $\Delta \equiv \omega - \omega_{0,1}$ is the detuning from the resonant frequency relative to the bare atomic resonance.

Applying the Liouville equation to this Hamiltonian shows

$$\begin{aligned}\dot{\rho}_{0,0} &= \frac{i\rho_{0,1}\Omega_{1\to0}}{2}e^{-i(\Delta t+\phi)} - \frac{i\rho_{1,0}\Omega_{0\to1}}{2}e^{i(\Delta t+\phi)} \\ \dot{\rho}_{1,1} &= \frac{i\rho_{1,0}\Omega_{0\to1}}{2}e^{i(\Delta t+\phi)} - \frac{i\rho_{0,1}\Omega_{1\to0}}{2}e^{-i(\Delta t+\phi)}.\end{aligned} \tag{2.20}$$

We can see here that $\dot{\rho}_{0,0} = -\dot{\rho}_{1,1}$, which is self-consistent. The total probability has remained constant at unity, as we have assumed no loss. Furthermore, if we define the changing coherence in the rotating frame

$$\tilde{\rho}_{0,1} \equiv e^{-i(\Delta t+\phi)}\rho_{0,1}, \tag{2.21}$$

such that

$$\dot{\tilde{\rho}}_{0,1} = e^{-i(\Delta t+\phi)}\left(\dot{\rho}_{0,1} - i\Delta\tilde{\rho}_{0,1}\right), \tag{2.22}$$

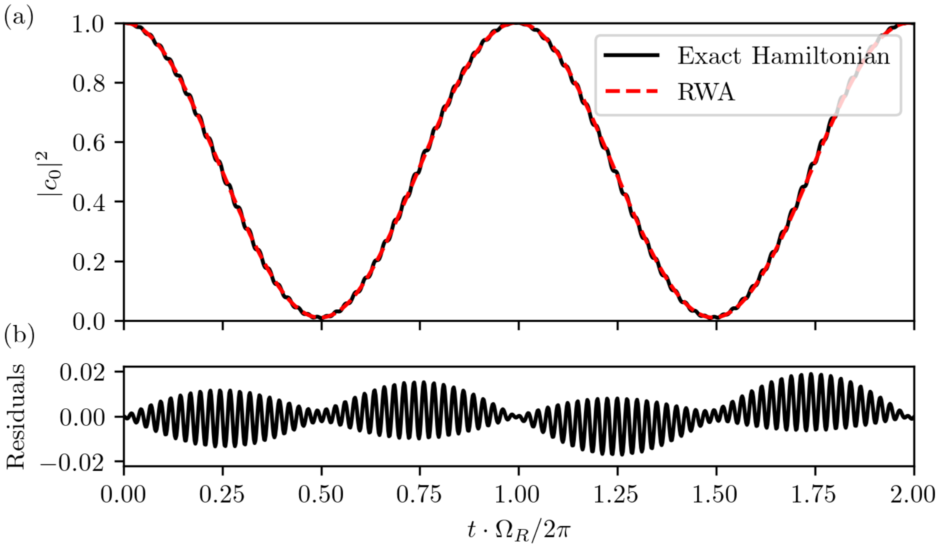

Figure 2.3: Simulated dynamics of the Rabi oscillation Hamiltonian. Plotted is the magnitude of the projection of the wavefunction to the $|0\rangle$ state for both the exact Hamiltonian defined in Eq. 2.11 (in solid black), and the same for the approximate Hamiltonian in Eq. 2.25 (in dashed red). (b) is the residuals, that is, the difference between the two values plotted in (a). This simulation uses $\Omega_R = 0.05i\omega_{0,1}$, $\Delta = +0.1|\Omega_R|$, where $|\psi(t)\rangle = |0\rangle$. Simulations are made using the QuTiP master equation solver [96].

then we can reduce Eq. 2.20 to

$$\begin{aligned}\dot{\rho}_{0,0} &= \frac{i}{2}\left(\Omega_{0,1}\tilde{\rho}_{0,1} - \Omega^*_{0,1}\tilde{\rho}^*_{0,1}\right) = -\sin\left(\arg\left(\Omega_{0,1}\tilde{\rho}_{0,1}\right)\right) \\ \dot{\rho}_{1,1} &= -\dot{\rho}_{0,0}\end{aligned} \tag{2.23}$$

Furthermore,

$$\dot{\tilde{\rho}}_{0,1} = \frac{i\Omega_{0,1}}{2}\left(\rho_{0,0} - \rho_{1,1}\right) - i\Delta\tilde{\rho}_{0,1}. \tag{2.24}$$

Here, we have used $\tilde{\rho}_{0,1} = \tilde{\rho}^*_{1,0}$, and $\Omega_{0,1} = \Omega^*_{1,0}$.

This evolution is well described by the effective Hamiltonian [94]

$$\mathcal{H}^{\mathrm{RWA}}_{T,\mathrm{Eff}} = \frac{\hbar}{2}\begin{pmatrix}\Delta & \Omega' \\ \Omega' & -\Delta\end{pmatrix}, \tag{2.25}$$

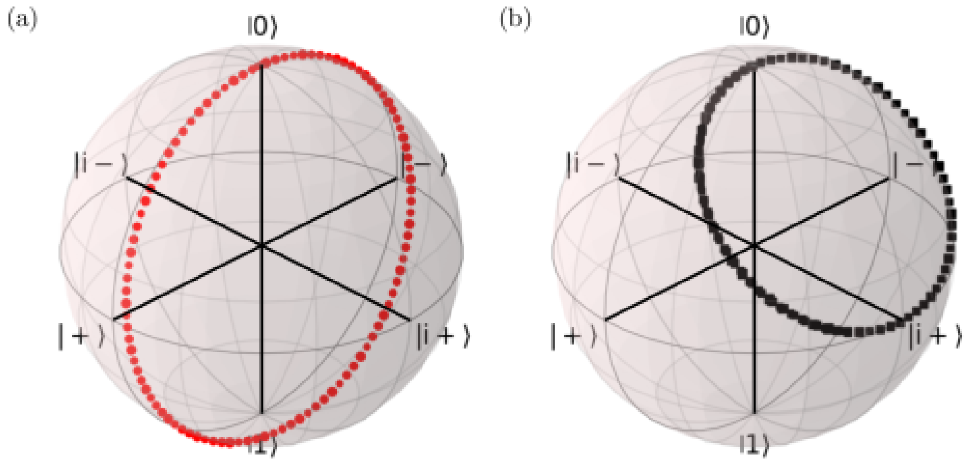

Figure 2.4: Bloch sphere representation of Rabi oscillations. Figure (a) uses the same parameters as Fig. 2.3. Each point (in red circles) plots the state of the wavefunction as it evolves. Figure (b) is, in black squares, the evolution of the state where $\Delta = \Omega_R$. This transition is off resonance, such that $|c_1| \leq |c_0|$ at all times.

where

$$\Omega' = \sqrt{\Delta^2 + |\Omega_{0\to 1}|^2}. \tag{2.26}$$

The dynamics of this Hamiltonian are simulated compared to the bare Hamiltonian in Fig. 2.3, to show how the approximation describes well the system for some small detuning by removing the fast-varying terms.

When $\Delta = 0$ such that the incident driving field is equal in frequency to the transition, and $\phi = 0$, the wavefunction is a transfer of population between the $|0\rangle$ and $|1\rangle$ states at rate Ω. The axis of this evolution on the Bloch sphere is determined by the input phase of light, such that c_0, c_1 gain a complex phase. Finally, as $|\Delta|$ increases above zero, the amplitude of these Rabi oscillations reduces, while the effective Rabi oscillation Ω' increases. If $\Delta \geq \sqrt{\Gamma^2 + |\Omega_{0\to 1}|^2}$, the transition will be substantially off-resonance, and the amplitude of the Rabi oscillations will be suppressed. This is shown diagramatically on the Bloch sphere in Fig. 2.4. The increased Rabi frequency is shown by simulating the exact Hamiltonian at varying deutnings in Fig. 2.5.

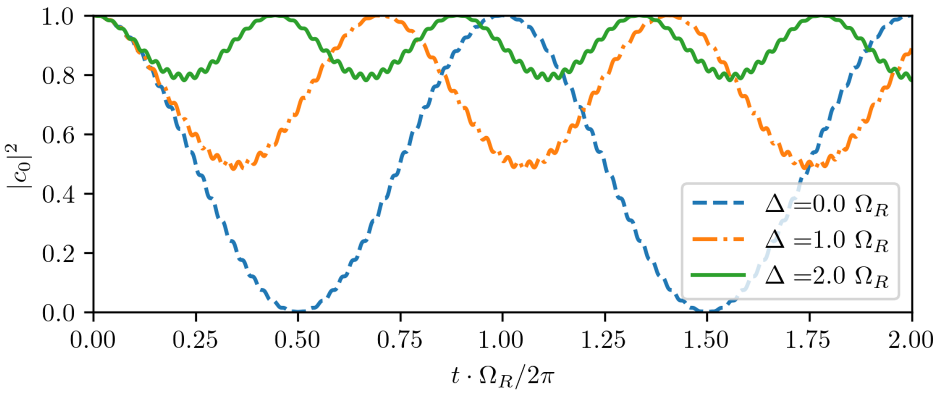

Figure 2.5: Simulated dynamics of the Rabi oscillation Hamiltonian for different detuning values. Plotted is the magnitude of the projection of the wavefunction to the $|0\rangle$ state for the exact Hamiltonian defined in Eq. 2.11, for detunings from $\Delta = 0$ to $\Delta = 2\Omega_R$. Note that as Δ increases, the frequency of oscillations increases, while the amplitude of the oscillations decreases.

2.1.4 AC Stark Shifts

As the Hamiltonian is no longer the bare Hamiltonian $\mathcal{H}_0$, we must characterise the effect of the change in eigenenergies. If we consider Eq. 2.25 to be a perturbation to $\mathcal{H}_0$, we can calculate the eigenenergies of the new Hamiltonian. The eigenenergies of Eq. 2.25 are

$$E_\pm = \pm\frac{\hbar\Delta}{2}\sqrt{1+\frac{\Omega'^2}{\Delta^2}}. \tag{2.27}$$

In the limit where $\Delta \gg \Omega'$, this is approximated by

$$E_\pm \approx \pm\left(\frac{\hbar\Delta}{2}+\frac{\Omega'^2}{4\Delta}\right). \tag{2.28}$$

Notably, as $I \to 0$ and thus $\Omega' \to 0$, a perturbation of $\pm\frac{\hbar\Delta}{2}$ remains. This is due to the treatment of light in the 'dressed' picture, where atoms in the ground state $|0\rangle$ and excited state $|1\rangle$ 'see' an energy difference of $\hbar\Delta$. This is the increased energy from the excitation, minus the energy of the now absorbed photon [97]. An applied field of Rabi frequency Ω' and detuning Δ

thus creates an energy shift $\hbar\Delta_{\mathrm{AC}}$ of magnitude

$$|\Delta_{AC}| \equiv \frac{1}{\hbar}\left(E_{\pm} \mp \frac{\hbar\Delta}{2}\right) = \frac{\Omega'^2}{4\Delta}. \tag{2.29}$$

Here, Δ_{AC} is the AC Stark shift, and is the effective shift of the resonance frequency of a transition as light is incident upon it. This AC Stark shift is to be avoided or exploited as appropriate. If Δ is very large, Ω can be large enough to generate a measurable Stark shift while remaining off resonance, and thus not transfer population.

2.2 Rydberg Physics

2.2.1 Rydberg Principles

For experiments utilising individual atoms, each atom must be distinguishable from the others. The interactions between atoms typically occur on the scale of electron orbitals, that is, of order < 1 nm, such as in molecules or in crystal lattices. Atoms this close are too close to resolve optically and the interactions too strong to reliably control on a single-atom level. We thereby must work with atoms where we can tune the interactions reliably and precisely, that act over distances of microns rather than nanometers.

As introduced in Sec. 1.3, these necessary atom-atom interactions are mediated through Rydberg atom interactions. As shown by Johannes Rydberg experimentally in 1890 [98], the binding energy of an electron W in hydrogen is given by

$$W = -\frac{\mathrm{Ry}}{n^2}, \tag{2.30}$$

where Ry is the Rydberg constant and n is a non-zero integer. Bohr would later go on to calculate this constant exactly for the Hydrogen atom [99], given by

$$\mathrm{Ry} = \frac{\alpha^2 m_e c}{4\pi\hbar}, \tag{2.31}$$

where α is the fine structure constant, m_e is electron mass, c the speed of light and $\hbar$ the reduced Planck constant. Rydberg further demonstrated that such relations exist for other species, most notably alkali metals, such that

$$W = -\frac{\mathrm{Ry}}{(n - \delta_{nlj})^2}, \tag{2.32}$$

where Ry depends on the species and δ_{nlj} is a 'quantum defect' that depends on the n, l, j quantum numbers. With the knowledge of modern quantum mechanics, we know this n value to be the principal quantum number, relating to a quantised radial wavefunction with expected distance from the atom at the centre of $\langle r \rangle \propto a_0 n^2$, where a_0 is the Bohr radius. Differences between alikali atoms and the hydrogen model at low l states are due to the low angular momentum states having large ellipticity, penetrating and thus interacting with the highly charged core [100]. Low-l states thus interact as if the effective radius, and thus n, are smaller than for large n, creating our quantum defect for $l \leq 3$.

Rydberg states have other notable characteristics to be either exploited or managed within Rydberg experiments. Considering only the low-l states, the spacing between atomic levels and the natural linewidths of such levels scales as n^{-3}, The $80S_{1/2}$ state, for example, has natural linewidth $\tau = 567.5$ μs [63].

A major property of Rydberg atoms is the substantially increased sensitivity to electric fields in comparison to ground state atoms. For a small, DC electric field, the second order perturbation creates an effective energy level shift ΔE, such that

$$\Delta E = -\frac{\alpha_0 \mathcal{E}^2}{2}, \tag{2.33}$$

where $\mathcal{E}$ is the electric field strength and α_0 is the scalar DC polarisabilty,

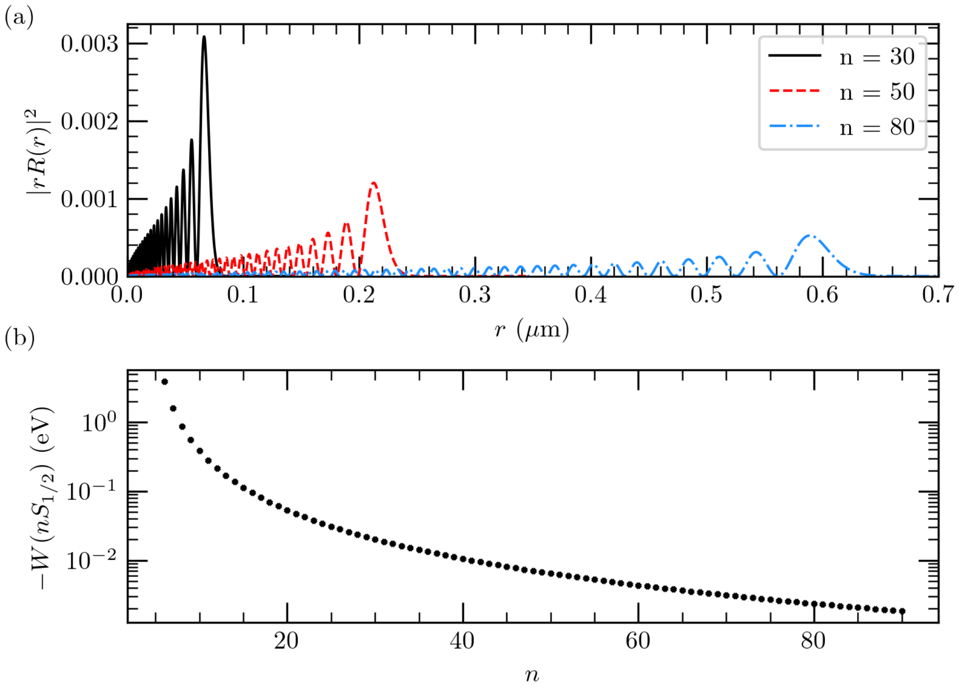

Figure 2.6: Plots of physical properties of Rydberg orbitals. Figure (a) shows the distribution of the radial wavefunction over the distance r from the nucleus. For $n = 80$ states, the most probable radius is 0.589 μm, or $11130a_0$, where a_0 is the Bohr radius. Figure (b) is the binding energy of the Rydberg electron for the Cs $|nS_{1/2}\rangle$ energy levels, as discussed in Eq. 2.32. This is further split in the presence of a magnetic field into separate m_F levels. Properties are calculated using the Atomic Rydberg Calculator library for Python [101].

where $\alpha_0 \propto n^7$ [102]. Thus, Rydberg atoms are highly susceptible to stray electric fields. Static electric field gradients must therefore be minimised across an experimental region. This process is demonstrated experimentally in Sec. 5.1.

Rydberg atoms are also sensitive to electric dipoles. The large orbital radius leads to both the generation of large dipole moments and the increased sensitivity to these dipole moments. For two dipoles, separated by the vector $\mathbf{R}$, the potential energy $V(\mathbf{R})$ of dipole-dipole interactions is given by [102]

$$V(\mathbf{R}) = \frac{\boldsymbol{\mu}_a \cdot \boldsymbol{\mu}_b}{|R|^3} + \frac{3\left(\boldsymbol{\mu}_a \cdot \mathbf{R}\right)\left(\boldsymbol{\mu}_b \cdot \mathbf{R}\right)}{|R|^5}, \tag{2.34}$$

where μ_a and μ_b are the dipole moments of the ground to Rydberg transition and where $|\mu_i| \propto n_i^4$.

If we consider only interactions where dipole moment and separation vectors μ_i and $\mathbf{R}$ are orthogonal, we can reduce the equation to purely scalar terms, such that

$$V(\mathbf{R}) = e^2 \frac{\langle r_a,\ r_b|\, x_a x_b + y_a y_b - 2 z_a z_b\, |g_a,\ g_b\rangle}{|R|^3} \tag{2.35}$$

where $|g_a,\ g_b\rangle\,, |r_a,\ r_b\rangle$ refers to ground and Rydberg states in the transition.

As shown in [90, Chapter 5], the energy shift on the pair state by the dipole-dipole interaction V_{dd} can, at long range, be calculated using second-order perturbation theory, such that

$$V_{\mathrm{dd}} = \frac{e^4}{|R|^6} \sum_{j \neq i} \frac{|\left\langle j^{(0)} \right| x_a x_b + y_a y_b - 2 z_a z_b \left| i^{(0)} \right\rangle|^2}{E_i^{(0)} - E_k^{(0)}}, \tag{2.36}$$

where V_{dd} is the energy shift applied to the pair state $|i\rangle$, and $|j\rangle$ is all other neighbouring Rydberg pair states. The unperturbed energy difference between the neighbouring $|r_{a,N} r_{b,M}\rangle$ and $|r_{a,N\pm1} r_{b,M\mp1}\rangle$ states is the denominator of this equation. This is the energy defect, and is proportional to n^{-3}.

This is simplified to an effective van der Waals interaction of

$$V_{\mathrm{dd}} = -\frac{C_6}{|R|^6} \tag{2.37}$$

where $C_6 \propto n^{11}$. The C_6 van der Waals interaction results in substantial energy shifts to the $|r_a,\ r_b\rangle$ state as Rydberg atoms are close, resulting in the phenomenon of Rydberg blockade.

It should be mentioned that this is true only for when the energy defect $|E_i^{(0)} - E_k^{(0)}| \gg V_{\mathrm{dd}}$, that is when the two atoms are sufficiently separated. At short range, the second-order perturbation in Eq. 2.36 is no longer valid, and in that case

$$V_{\mathrm{dd}} = \pm\frac{C_3}{|R|^3}. \tag{2.38}$$

2.2.2 Rydberg Blockade

The most critical property of Rydberg atoms for use in quantum information technologies is the phenomenon of Rydberg blockade. To summarise, two atoms within close proximity cannot both be excited to the Rydberg state simultaneously using on-resonant interactions. This interaction results in quantum entanglement between the two atoms.

This is due to the energy shift caused by the dipole-dipole interaction (as described in Eq. 2.37) changing the resonant frequency for the interaction. Light that is resonant to the single atom Rydberg transition cannot couple to the state $|r,\ r\rangle$, as it experiences an effective detuning from that state due to the dipole-dipole interaction. If the linewidth of the transition (typically the Rabi frequency) is less than the induced shift, the probability of both atoms being excited to the Rydberg state is substantially reduced, despite each atom being resonant to the field if they were not in close proximity.

To demonstrate this, one must first consider the Hamiltonian for a single atom excited by a resonant field within a Hilbert space consisting of the ground and Rydberg states, or $|g\rangle$ and $|r\rangle$ respectively. Such a Hamiltonian is given by Eq. 2.25. Consider further the two-atom system described by the independent states $|g,\ g\rangle$, $|g,\ r\rangle$, $|r,\ g\rangle$ and $|r,\ r\rangle$. Assuming both atoms are exposed to the same field with equal Rabi frequency, and are otherwise identical aside from position, they experience the same dynamics with individual detuning $\Delta = 0$. However, each atom experiences a dipole-dipole interaction, as detailed in Sec. 2.2.1. This interaction creates an effective detuning for the $|r,\ r\rangle$ state only, of strength $-\frac{C_6}{R^6}$. Thus, if C_6 is substantially large (by

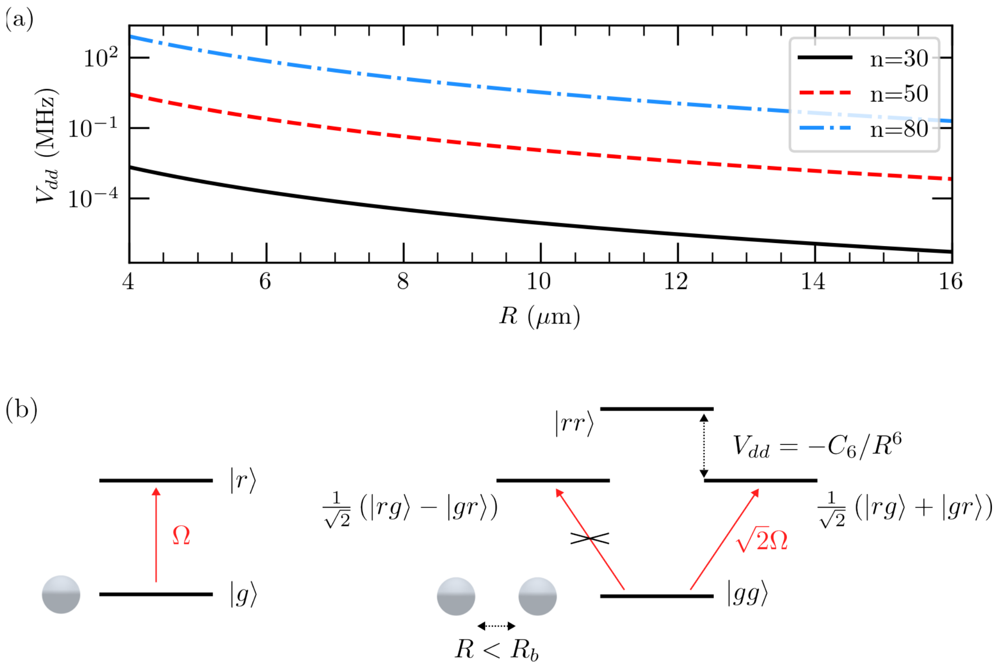

Figure 2.7: Diagram of the Rydberg blockade mechanism. Figure (a) shows the calculated values of the dipole-dipole interaction for various $nS_{1/2}$ states. The magnitude of the shift varies by $V_{dd} \propto n^{11}$ and $V_{dd} \propto R^{-6}$. Properties are calculated using the Atomic Rydberg Calculator library for Python [101]. Figure (b) is a schematic of the dipole-dipole interactions between two Rydberg atoms, separated by a distance R. When $R < R_b$, the $|r,\ r\rangle$ state is shifted off resonance, and coupling to the $\frac{1}{\sqrt{2}}(|r,\ g\rangle + |g,\ r\rangle)$ pair state occurs at the enhanced Rabi frequency $\sqrt{2}\Omega$. The $\frac{1}{\sqrt{2}}(|r,\ g\rangle - |g,\ r\rangle)$ pair state is not coupled.

increasing n) or R is substantially small (by decreasing interatomic spacing), the $|r,\ r\rangle$ state is forbidden. This is what is referred to as the Rydberg blockade. Furthermore, the Rydberg blockade radius R_b is defined such that $|C_6|/R_b^6 = \Omega_R$, such that the $|r,\ r\rangle$ state is off-resonant.

A further useful result is found when considering the evolution of states under this Hamiltonian. This Hamiltonian has eigenstates $\frac{1}{\sqrt{2}}(|r,\ g\rangle \pm |g,\ r\rangle)$, where the rate of evolution of each state is given by

$$\langle gg|\mathcal{H}|\pm\rangle = \frac{\Omega}{2\sqrt{2}}\left(\langle g,\ g|g,\ g\rangle \pm \langle g,\ g|g,\ g\rangle\right). \tag{2.39}$$

Thus, the $\frac{1}{\sqrt{2}}(|r,\ g\rangle + |g,\ r\rangle)$ eigenstate is excited at rate $\sqrt{2}\Omega$, while the

$\frac{1}{\sqrt{2}}\left(|r,\ g\rangle - |g,\ r\rangle\right)$ cannot be coupled. The resultant state is equivalent to the Ψ^+ Bell state, and is thus maximally entangled. This increased Rabi frequency is referred to as the enhanced Rabi frequency. This behaviour is seen in a greater number of atoms as Rydberg atoms are all brought within the collective blockade radius, such that for N atoms within collective blockade,

$$\Omega_N = \sqrt{N}\Omega, \tag{2.40}$$

where Ω_N is the enhanced Rabi frequency [103, 104]. Furthermore, the state generated is the $|W_N\rangle$ state, as defined in [105], which is known to be maximally entangled up to $N = 3$ [106, 107].

2.3 Three-Level Atomic Physics

Restricting ourselves to just two-level systems is not always practical for experimental physics. Exciting Cs to a Rydberg state with a single photon requires UV lasers [108], which come with experimental difficulties, while transferring between hyperfine states requires MW transitions, which are slow relative to Rydberg lifetimes and cannot be focussed to address a single atom. This section will introduce the theory of addressing three-level systems to address both large and small energy differences.

Here, we will be primarily concerned with two states of note, which we shall call $|0, 1\rangle$, and an intermediary excited state $|e\rangle$. The goal of the intended three-level system interactions are to transfer populations between the $|0, 1\rangle$ states, while minimising the population in the intermediate $|e\rangle$ state. If this is done, the problem is reduced to a much simpler two-level system.

Note that while these approximations are both useful and effective at describing the dynamics of the system, the results are only approximate. The exact Hamiltonian can be simulated using numerical methods to compare to

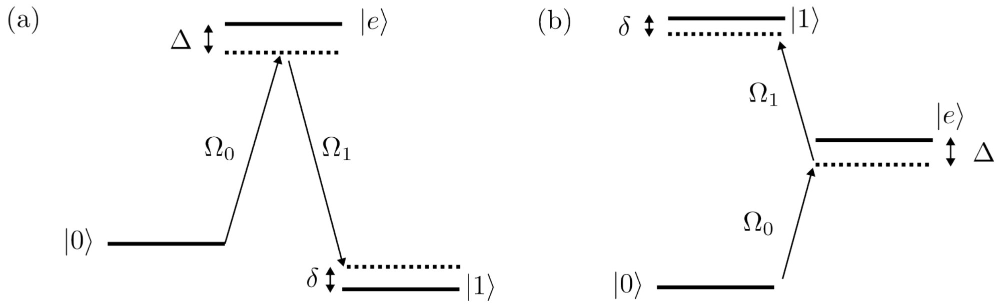

Figure 2.8: Level diagrams demonstrating Λ and ladder systems. These are two example three-level systems, where the population in the excited state is to be minimised. Figure (a) demonstrates a Λ system with average detuning Δ and two-photon detuning δ. As δ is non-zero, this transition is off-resonant. By reducing δ to zero, coherent population transfer occurs between the $|0\rangle$ and $|1\rangle$ states, via the off-resonant coupling to the $|e\rangle$ states. Figure (b) is the same for a ladder system. Here, transfer occurs through a two-photon absorption process, rather than the absorption and emission of virtual photons.

the approximations used, although these simulations can be computationally demanding even for small systems.

2.3.1 Adiabatic Elimination

A typical issue with driving an oscillation between two target states in a three level-system is population remaining in the excited state. Example three-level systems are shown in Fig. 2.8. For example, Rydberg state lifetimes are long, measured in the hundreds of μs scale for $n > 40$ [63, 109], while the natural lifetime of the $7P_{1/2}$ state is $\tau_{1/2} = 165.2$ ns [110]. This means that any process moving from the $6S_{1/2}$ ground state to the $nS_{1/2}$ Rydberg state via the $7P_{1/2}$ state will spontaneously emit photons and scatter far faster from this intermediate state than the Rydberg state. This is the ‘ladder’-type energy level, shown in Fig. 2.8(b). By reducing the population within this intermediate state, we can eliminate this scattering. The process of removing the population of the intermediate state is known as "adiabatic elimination". Conveniently, this will reduce a three-level system to a two-level system as described above.

Using the previously defined initial states $|0\rangle, |1\rangle$, these two states are coupled to the excited state by fields of frequency $\omega_{0,e}, \omega_{1,e}$, defined by Rabi frequency Ω_0, Ω_1 and detuning Δ_0, Δ_1. If we define an average detuning $\Delta = (\Delta_0 + \Delta_1)/2$ and a two-photon detuning $\delta = (\Delta_0 - \Delta_1)$, we can extend the rotating wave approximation as seen in Eq. 2.25 to a three-level system.

$$H = \hbar \begin{pmatrix} -\delta/2 & 0 & \frac{\Omega_0^*}{2} \\ 0 & \delta/2 & \frac{\Omega_1^*}{2} \\ \frac{\Omega_0}{2} & \frac{\Omega_1}{2} & -\Delta \end{pmatrix}. \tag{2.41}$$

Note that here we have assumed no coupling between the $|0\rangle$ and $|1\rangle$ states, such that $\Omega_{0\to1} = \Omega_{1\to0} = 0$.

To adiabatically eliminate population in the excited state, we must first assume the initial population in the excited state is zero. Furthermore, we must assume that the average detuning Δ is large in comparison to the Rabi frequencies Ω_0, Ω_1, and excited state linewidths $\Gamma_{e\to0}, \Gamma_{e\to1}$, such that with only one field there is minimal coupling between the two states. We can then make the assumption that the system reaches a steady-state population, such that $\dot{c}_e = 0$. As shown in [111], solving the time-dependent solution results in an effective Hamiltonian

$$H^{RWA} = -\hbar \begin{pmatrix} \delta/2 + \Delta_{AC}^0 & \frac{\Omega_R^*}{2} \\ \frac{\Omega_R}{2} & -\delta/2 + \Delta_{AC}^1 \end{pmatrix}. \tag{2.42}$$

Here, $\Omega_R = \frac{\Omega_0 \Omega_1^*}{2\Delta}$, which has a complex phase equal to the phase difference between the two fields. The problem is thus reduced to a two-level system, with the effective detuning shifted by the two-photon detuning, and contributions from the differential AC Stark shift from each beam.

This approximation is not entirely rigorous and for a full solution the reader should see [111]. For substantially large Δ however this approximation can be used effectively to understand the evolution of the system. Fig. 2.9 shows

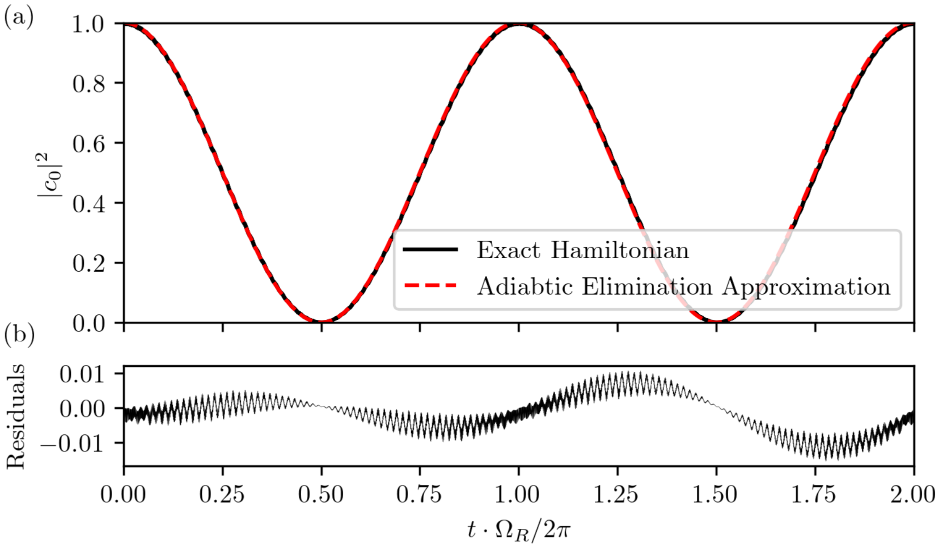

Figure 2.9: Simulated dynamics of the adiabatic elimination Hamiltonian. In (a), plotted is the magnitude of the projection of the wavefunction to the $|0\rangle$ state for both the exact time dependent Hamiltonian (in solid black) compared to the approximate Hamiltonian in Eq. 2.42 (in dashed red). (b) is the residuals, that is, the difference between the two values plotted in (a). This simulation uses $\Omega_{0\to e},\ \Omega_{1\to e} = 0.05\omega_{0,e},\ 0.05\omega_{1,e},\ \Delta_0,\ \Delta_1 = 5\omega_{0,e},\ 5\omega_{1,e}$ and $\delta = 0$, where $|\psi(t)\rangle = |0\rangle$. Simulations are made using the QuTiP master equation solver [96].

a comparison between the approximation the exact Hamiltonian in simulation, where the difference in population is taken as a residual, showing only small deviations. We can thus take it as two fields acting as a single interaction strength term Ω_R, although each field produces an AC Stark shift independently of the others. Also note that this approximation is agnostic as to if the excited state is above, below or between the two initial states.

Even though we assume the population in this intermediate state is very small, we must consider the probability of scattering from this state. Given the well known scattering rate equation [95, 112], we can assume the factors from the large detuning Δ dominate. As such we can approximate the contributions from Ω_0 to be

$$R_{\text{scatt},0} = \frac{\Gamma}{2}\frac{\Omega_0^2/2}{\Delta^2 + \Omega_0^2/2 + \Gamma^2/4} \approx \Gamma\frac{\Omega_0^2}{4\Delta^2}. \tag{2.43}$$

The same contribution applies for contributions from Ω_1. Furthermore, we can estimate the timescale of any transition to be on the order of $t \approx \frac{2\pi}{\Omega_R}$. Thus, the approximate total number of scattered photons from the $|e\rangle$ state is given by

$$R_{\text{scatt},0} t \approx \frac{\Gamma}{2} \frac{\Omega_0^2}{\Delta^2} \frac{2\pi}{\Omega_R} = \frac{\Gamma}{2\Delta} \frac{\Omega_0}{\Omega_1}. \tag{2.44}$$

Following from this, $R_{scatt,1} \approx \frac{\Gamma}{2\Delta}\frac{\Omega_1}{\Omega_0}$. Thus, if Ω_0 and Ω_1 are comparable in size and $\Delta \gg \Gamma$, the number of photons scattering from the excited state is $\ll 1$.

2.3.2 Lambda and Ladder Systems

Adiabatic elimination processes can occur through any two-photon process. As previously established, a two-photon processes can be used to address otherwise unfeasible transitions. In a Λ-system, where the $|0\rangle$, $|1\rangle$ states are close in energy relative to the excited state $|e\rangle$, the dynamics are the same as a two-photon absorption or emission process. A ladder system is one where the excited state is between the $|0\rangle, |1\rangle$ states in energy. Both are described in Fig. 2.8.

We can first look at the scenario where our initial levels are close in frequency, and we wish to drive a transition between the two via an excited state at a much higher frequency. This has the appearance of the stimulated absorption of one virtual photon, and the stimulated emission of another virtual photon. Alternatively, one can consider it as the two fields interfering to create a beat frequency that itself addresses the transition. In either respect, this type of transition is often called a ‘Raman transition’[3].

In this experiment, Raman transitions are used for rapid high-fidelity transfer between hyperfine states. Phase-locked lasers are be capable of state transfer at far greater rates than MW addressing. Such beams can be focussed to high

---

[3] Not to be confused with Raman scattering

intensity while retaining control within experimental timescales of ≈ 1 MHz. These beams can also be focussed to a single atom, thus creating single atom state control, unavailable in a MW source. The phase of these beams must be kept constant, as must the relative frequency difference, as it is this effective detuning δ that determines the resonance condition.

One notable factor of this Raman transition is the impulse on the atoms from this transition. If the two beams have momenta $\mathbf{k}_{a,b}$, the momentum imparted onto the atom is thus $\mathbf{k} = \mathbf{k}_a - \mathbf{k}_b$. This momentum transfer can be used for Raman sideband cooling. This technique involves coupling the atomic transition to the motional states of the atom within a trap (see chapter 2.5.1), such that the impulse of the Raman transition reduces the total energy of the atom. This has been demonstrated to cool neutral atoms [113, 114] and molecules [115, 116] to the motional ground state of dipole traps.

Although care must be taken with the sign of the detuning values, two-photon excitation schemes in Ladder systems use much of the same physics as Raman transitions. For this work, one further consideration must be made regarding the excitation pathway; a transition may couple from the initial to the final state via multiple different intermediate states $|e_i\rangle$, with proportional rates determined by relative dipole matrix elements. The detuning value is itself weighted by these factors as well. In general, these multiple-pathway considerations must be taken into account for Raman transitions as well, but is only relevant for the paths taken for two-photon excitations in this experiment.

In summary, the Ladder system is used for excitation to Rydberg states while the Lambda system is used for Raman transitions between hyperfine levels. These can be approximated to the same form of Hamiltonian, although Raman systems can also be used for Raman sideband cooling.

2.4 Laser Cooling

The first stage of many atomic experiments is typically cooling. A room temperature Cs atom may move at rates exceeding 200 m/s, leading to an extensive Doppler shift while increasing the difficulty of interacting with the same atom for more than a short period. An atomic beam may have been suitable for the first groundbreaking experiments in atomic spectroscopy [117], but for more precise quantum technologies, the same atom must be addressed over the period of milliseconds or even seconds. Variations in atomic velocity must also be eliminated, to avoid the effects of Doppler shifts.

Cooling an atom from room temperature to the microkelvin scale takes multiple stages, which are each briefly discussed below.

2.4.1 Doppler Cooling and Magneto-Optical Traps

The magneto-optical trap (MOT) has become a staple of many atomic physics experiments since its introduction. First demonstrated in 1987 at Bell Labs [118], this technique uses radiation pressure to cool and trap neutral atoms into a cloud of < 1K. Radiation pressure here refers to how atoms, when absorbing a photon, increase their momentum by the momentum of the photon absorbed. When this photon is spontaneously emitted, it is scattered in a random direction. Thus the absorption process leads to a mean change in momentum over time (and thus force) while the scattering does not.

Firstly, two beams red-detuned from resonance are incident on an atom. An atom at rest is equally off-resonant from either beam. If the atom's velocity is counter to the direction of the light, the Doppler effect also moves the atom into resonance, while it moves further off-resonance for the beam along its velocity. This results in a net counter-force proportional to its velocity, This effect is known as 'Doppler cooling', and is a frequently used laser cooling technique.

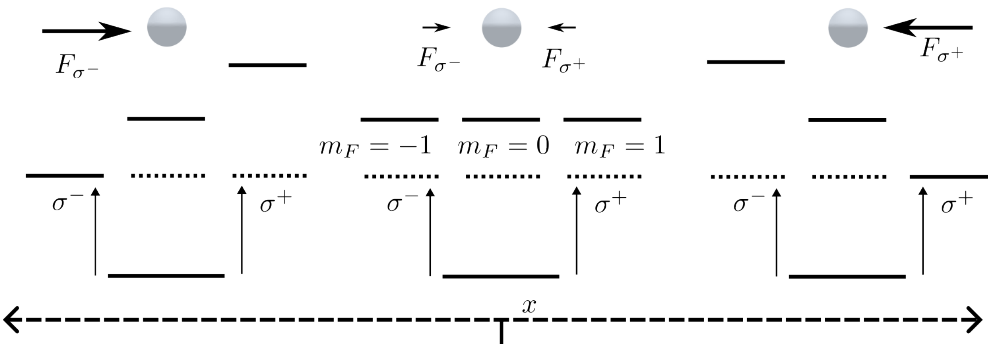

Figure 2.10: Energy level diagram of a MOT along a single axis. As the atom moves away from the zero position, the Zeeman shifts of the atomic levels bring energy levels into and out of resonance. Beams in resonance propagate in the direction opposite to that of the atom, acting as a restoring force.

While Doppler cooling slows atoms, the addition of a magnetic field gradient can create trapping. The operating principle of the MOT requires two counter-propagating beams of orthogonally polarised light, red detuned from resonance by a few linewidths, creating a Doppler cooling effect. A magnetic field gradient is also applied across this axis, with zero field at the centre. As an atom moves in the positive direction, the magnetic field becomes stronger and the Zeeman shift brings it into resonance with the beam opposing its motion. A restoring force is thus applied to accelerate it towards the origin. In the other direction, as the magnetic field becomes weaker, the second beam is in resonance, and so the force restores the atom again. A magnetic field gradient and counter-propagating beams can thus cool and trap atoms in each axis. The result is the containment of a cloud of atoms, capable of trapping $> 10^8$ atoms [119–121]. As atoms are always in receipt of a force counter to their displacement and velocity, this traps the atoms in the centre with a much reduced mean velocity compared to room temperature. The effect of this is a dense, cold cloud of atoms.

This technique has limits, however. Even when off resonance, the atoms still absorb photons at a rate determined by the scattering equation, Eq. 2.43. Thus, an atom at zero velocity in zero field will still absorb photons, and stochastically emit them. This forms a random walk process, such that at a

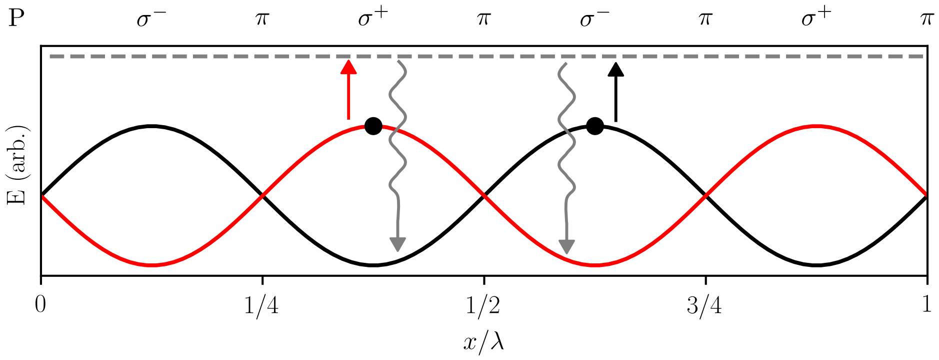

Figure 2.11: Diagram of the PGC process. Red and black lines are the energy levels for for alternating F levels while the dashed line is the excited state. On the top row, "P" refers to the polarisation of incident light. As atoms move in x, they reach the top of an energy potential, where they are excited and spontaneously decay, losing energy in the process.

steady-state the atoms will reach a limit of cooling. As proven in [112], this limit T_D is known as the Doppler-cooling limit.

$$k_B T_D = \hbar\Gamma/2. \tag{2.45}$$

For Cs cooled along the $D2$ transition (as done in this experiment), this limit is reached at 125 μK, or a velocity of 0.15 m/s. To cool atoms below this limit, a further cooling stage is required.

2.4.2 Polarisation Gradient Cooling

Cooling atoms below the Doppler-cooling limit was first achieved before it was known to be possible [122]. A standing wave pattern formed by the MOT beams forms a spatial variation of the field polarisation. This forms the basis for the highly effective polarisation gradient cooling (PGC) technique, also known as Sisyphus cooling.

Firstly, when counterpropagating beams of wavelength λ and orthogonally but linearly polarised light interfere along an axis, the resultant standing wave has spatially varying polarisation. It will smoothly vary from linear, to circular, to orthongonally linear, to orthogonally circular, and back to linear again. This effect repeats in a standing wave pattern over spatial period $\lambda/2$, as shown in Fig. 2.11. As an atom moves from linear to circular polarisation, it begins to encounter a negative AC Stark shift at the $|F, mF\rangle$ energy level, with the largest shift when $mF = \pm F$. As such, atoms that are highly shifted no longer see optical pumping effects. An atom moving out of this shifted region reduces its potential energy, but is also subject to further optical pumping. Said atoms undergo a repump stage and are most likely to eventually end in a position of again being under AC shift (as atoms in the unshifted state simply repeat the process). What the atom 'sees' in this case is starting in a region of low potential energy but high velocity. As it moves into a region of high potential energy and low velocity, this potential energy is removed by the optical pump. The atom continues moving to a region of low potential energy, and is repumped back to the target state.

As was first shown in [123], this cooling effect only occurs in limited circumstances. The optical pumping effect must be slow, such that an atom has time to 'climb' to the top of a potential hill and reduce its velocity before being optically pumped. Also shown is how the effect on the velocity is only strongly seen when atoms are already slow, such that it acts best as a second stage of cooling after a MOT.

2.5 Single Atom Traps

The first step, trapping atoms in a MOT, produces a cloud of cold atoms. However, our experiment requires isolated, individually distinguishable atoms. Knowing that our Rydberg interaction is strongly dependent on separation, we require fine control of the atomic positions. Given that we want readout from a single atom at a time, they must also be distinguishable and identifiable. In Sec. 2.5.1 I will discuss the theory of using optical dipole traps (also called atomic tweezers) to hold atoms in place, while Sec. 2.5.2 will discuss using a light-assisted collisions stage to reduce occupancy of these traps to a single atom.

2.5.1 Optical Dipole Traps

Aside from the force of absorbing and emitting photons, light can interact with atoms through a dipole force. In brief, the force applied to a dipole in an oscillating field is known to be [95]

$$F_{dipole}\hat{\epsilon}_i = \frac{-\hbar\Delta}{2}\frac{\Omega}{\Delta^2+\Omega^2/2+\Gamma^2/4}\frac{\partial\Omega}{\partial x_i}\hat{\epsilon}_i, \tag{2.46}$$

where $\hat{\epsilon}_i$ is the unit vector along the x_i axis. As before with two-photon transitions (see Sec. 2.1), the probability of interaction with a given transition is given by its relative dipole moment, and so the actual force is a weighted average over all possible transitions.

On resonance, this interaction power is zero. If we take the same assumptions for AC Stark shift calculations in Sec. 2.1.4, $\Delta \gg \Gamma, \omega$ we can reduce this to

$$F_{dipole}\hat{\epsilon}_i = \frac{\partial}{\partial x_i}\frac{-\hbar\Omega^2}{4\Delta}\hat{\epsilon}_i = \frac{\partial}{\partial x_i}\left(-\hbar\Delta_{AC}\right)\hat{\epsilon}_i \tag{2.47}$$

In brief, the dipole force is applied to the atom across the gradient of the AC Stark shift. To extend this to three dimensions,

$$\mathbf{F}_{dipole} = -\hbar\nabla\Delta_{AC}(\mathbf{r}) = -\nabla U_{dipole}(\mathbf{r}). \tag{2.48}$$

Here we have reduced the dipole force to the gradient of a scalar potential energy field. For a far-off resonance transition, it can be useful to understand this in terms of the electric field intensity I [95, 124].

$$U_{dipole}(\mathbf{r}) = \frac{3\pi c^2 \Gamma}{2\omega_{a,e}^3 \Delta} I(\mathbf{r}) = U_0 I(\mathbf{r}), \tag{2.49}$$

where $\omega_{g,e}$ is the resonant frequency of the interacting transition between the ground and excited states. For sufficiently detuned beams, this beam will couple to multiple transitions simultaneously. By generating an energy well, atoms can be trapped using this dipole force. This is an optical dipole trap (ODT).

Using this formalism, and conversion of energy to temperature through $U = k_B T$ where k_B is Boltzmann's constant, dipole traps are typically described in terms of temperature. A theoretical atom with no energy at the bottom of a potential well requires additional energy equal to the 'trap depth' to escape. A trap of depth U_{TD} sets the upper limit of the system temperature.

One property of Eq. 2.49 is its dependence on the sign of the detuning. If the incident light is of a longer wavelength / lower frequency than the resonance frequency, $\Delta < 0$, or 'red' detuned. The opposite is considered 'blue' detuned.

Red-Detuned ODTs

In the red detuned case, the potential is minimised where $I(\mathbf{R})$ is maximised, such that atoms are trapped in regions of high intensity. This somewhat counterintuitive result means that a highly focussed spot of light acts to trap atoms.

A typical ODT is generated by focussing a Gaussian beam to a small point of $1/e^2$ waist w_0. Along the radius at the focus,

$$I(r) = I_0 \exp\left(\frac{-2r^2}{w_0^2}\right) \tag{2.50}$$

which, for $r \ll w_0$,

$$I(r) \approx I_0 \left(1 - \frac{2}{w_0^2} r^2\right), \tag{2.51}$$

and so, using Eq. 2.49,

$$U(r) \approx U_0 I_0 + \frac{kr^2}{2}, \tag{2.52}$$

where

$$k = -U_0 I_0 \frac{4}{w_0^2}. \tag{2.53}$$

As long as $U_0 < 0$, this has the form of a harmonic oscillator potential, where

$$U(r) = \frac{1}{2} kr^2 + C, \tag{2.54}$$

where C is a constant energy offset and k is the spring constant. This corresponds to a characteristic radial trap frequency

$$\omega_{\text{radial}} = \sqrt{\frac{4U_0 I_0}{m w_0^2}}. \tag{2.55}$$

Following the same steps along the axis orthogonal to the trap radius, we can see that the atoms are trapped axially as well, such that

$$\omega_{\text{axial}} = \sqrt{\frac{2U_0 I_0}{m z_R^2}}, \tag{2.56}$$

where $z_R \equiv \frac{\pi w_0^2}{\lambda}$ is the Rayleigh range.

Atoms that are much colder than the depth of the trap closely match the conditions for this potential. We can therefore assume, for cold atoms, that a red-detuned optical dipole trap acts as a quantum harmonic oscillator. The distribution of occupied modes for a thermal atom matches a thermal Maxwell-Boltzmann distribution (conveniently providing a method of describing temperature for a single atom), although this can be further modified. Raman transitions can couple to these harmonic modes, leading to Raman thermometry for directly measuring atom temperature. Raman sideband cooling (RSBC) can then cool atoms by reducing occupation of harmonic modes to the ground state of the trap, as discussed in Sec. 2.3.2.

The focussed Gaussian beam configuration also traps along the axis of propagation of the beam, although with substantially weaker forces. Crossed-beam trap configurations can therefore improve trapping uniformity and confinement along this final axis.

Blue-Detuned ODTs

The alternative to red-detuned ODTs is the blue-detuned ODT. Here, we can consider the bottom of the trap to be the zero of potential energy, and is surrounded by potential walls. Atoms are trapped in regions of low intensity, surrounded by high intensity.

Such fields can be generated through a variety of methods, although all require a degree of optical ingenuity beyond simply focussing a single Gaussian beam. They are typically used in situations where strong electric fields would be destructive to the trapped object. Most relevant to us, this includes Rydberg atoms.

A wide variety of optical designs have been historically used to trap atoms in blue-detuned light, for example ‘light sheets’ [125] or using higher order Gaussian modes [126, 127]. We will give special attention to the method of generating ‘bottle beams’ [128], trap schemes where a Gaussian beam is ‘hollowed out’, as it is described in [124]. One such method of generating such beams involves using two Gaussian beams of the same wavelength, different waists and in antiphase, such that at the focus there is a region of destructive interference in the centre [128, 129]. This can be reduced to a single beam interfering with itself, as first shown in [130] and expanded to diffractive optical elements in [131]. In [132], this method uses a spatial light modulator (SLM) to generate a blue-detuned ODT to trap Rydberg atoms. This thesis will demonstrate the same method to generate arrays of red-detuned and blue-detuned ODTs to trap ground state and Rydberg arrays in Sec. 3.3.3 and Chapter 6 respectively.

2.5.2 Light Assisted Collisions

As we can now trap a large number of atoms in a MOT, then small numbers in a dense ODT array, we only need one final step to reduce trap occupancy to a single atom. Light assisted collisions are the mechanism for which this is achieved. Using the same beams used for PGC (see Sec. 2.4) but closer to resonance, dense collections of atoms will be attracted and collide [133]. Atoms undergoing collisions will exchange spin and each increase kinetic energy. This ejects atoms from the trap in a pair-wise fashion. Traps that initially loaded with an odd number of atoms will thus be left with a single atom that, absent further collisions with background atoms, will remain.

This is a stochastic process and so a trap can be expected be filled by a single atom with a probability of $\approx 50\%$. For further discussion on this process, see Sec. 4.1.3.

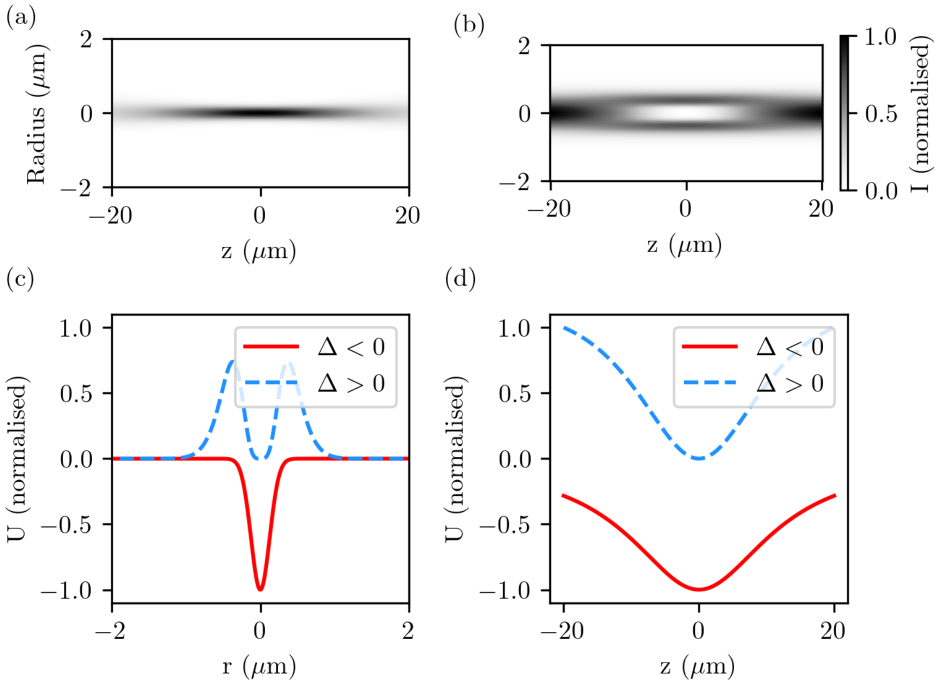

Figure 2.12: Comparison between simulated red and blue detuned ODTs. Figure (a) is a color plot of a typical red detuned ODT, generated by a focussed beam of $1/e^2$ waist 1 μm. The figure is a cross-section of the beam profile, as it is rotationally symmetric about the axis of propagation (z). Figure (b) is a color plot of a bluebottle blue-detuned ODT, generated by a focussed beam of $1/e^2$ waist 1 μm, destructively interfering with a beam of waist 2.5 μm. Both intensity profiles are normalised to their respective maximum intensity. Figures (c) and (d) demonstrate energy profiles along the radial and axial directions, normalised to the same degree as above. Red solid lines are for the red-detuned ODT as shown in (a), while blue dashed lines are for the blue-detuned ODT as in (b).

Chapter 3

Experimental Setup

This Chapter details the realisation of the SQuAre atom array experimental system. Included here is the vacuum chamber containing the atoms, the optics to address and control the atoms, and the various hardware systems required to control the experiment as a whole.

This Chapter will be a brief overview of the system to discuss its capabilities such that they can be assessed, its limitations understood, and the potential for expansion highlighted. Here, special consideration must go to the thesis of B. Nikolov [87], who also provided substantial contributions to the construction of the experimental system and has described its workings in detail.

3.1 Vacuum Chamber System

The core of the experiment is the vacuum chamber. This consists of a glass cell to contain a 2D magneto-optical trap (MOT, see Sec. 2.4), which feeds into a main chamber through a differential pumping tube. This main chamber is an octagonal vacuum system, with windows on the cartesian axes for input light. Further smaller windows at a ±25° offset give access for the 2D MOT

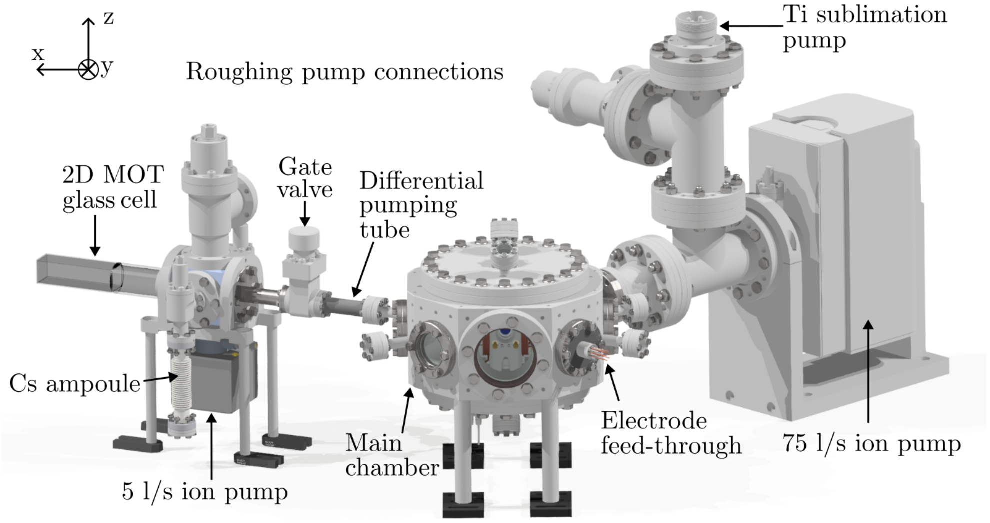

Figure 3.1: Annotated diagram of the experiment vacuum system, with laboratory coordinate reference frame. Reproduced from [87] with permission.

connection. As mentioned in Sec. 2.3.2, a second offset axis for Raman beams gives access for future Raman sideband cooling.

3.1.1 2D MOT Cell

The source of atoms for the experiment is an ampoule of metallic Cs, such that it is fed Cs at vapour pressure. This is at a pressure $P_V \approx 7 \times 10^{-7}$ Torr at the system temperature of 20° C [91]. This is the highest pressure of the system. The 2D MOT cell is connected to the ampoule through a valve. It is opened for 30 minutes, once a week, such that atom numbers in the 2D MOT cell remain low but roughly constant. The 2D MOT cell is also connected to a 5 l/s ion pump to maintain low pressure. By measuring the ion pump current, we can infer an initial pressure of $P \approx 8 \times 10^{-9}$ Torr after baking. This ion pump was on during the initial vacuum preparation process, but is not in regular use. This is due to accumulation of Cs in the pump leading to inaccurate readout.

3.1.2 Main Experiment Chamber

The main experiment chamber is held at a lower pressure than the 2D MOT cell. The main chamber is connected to the 2D MOT cell through a differential pumping tube, which limits the movement of atoms from the high pressure 2D MOT cell to the main chamber. Use of a titanium sublimation pump at initialisation reduces the chamber pressure. A further 75l/s ion pump is kept in continuous operation, connected to the main chamber. As such, the pressure is maintained at $P \approx 1 \times 10^{-11}$ Torr. This is measured by the pump rate of the ion pump, which reached a minimum measurable pressure of $1.0 \cdot 10^{-11}$ Torr. For details on measuring vacuum performance, see [87]. This pressure slowly increases over time, due to outgassing of the materials and slow leaks. The pressure increase rate is estimated to be at a rate of $2 \cdot 10^{-11}$ Torr per year, and has been reduced down to the minimum measurable pressure once through further activation of the titanium sublimation pump. This UHV pressure results in trap lifetimes of 17(1) s, as will be demonstrated in Sec. 3.3.3.

3.1.3 Chamber Optics Assembly

The primary chamber imaging optics are based on work in [134–136]. Here, the lens used is a single high-NA lens fixed in-vacuum with a short working distance on each side of the chamber centre. Only a single optical element with high efficiency of light collection is used, compared to compound microscope objectives. Two identical lenses are used on either side of the atoms, such that the primary optics are equal on each side. This experiment uses Geltech 355561 lenses (as used in [135]). These lenses have focal length f = 10 mm, numerical aperture NA = 0.45 at the imaging wavelength of 852 nm, and are placed at a working distance of 7 mm. Given the close distances of the lenses to the atoms, the insulating properties of glass could lead to charge build up and thus residual electric fields. These lenses were thus

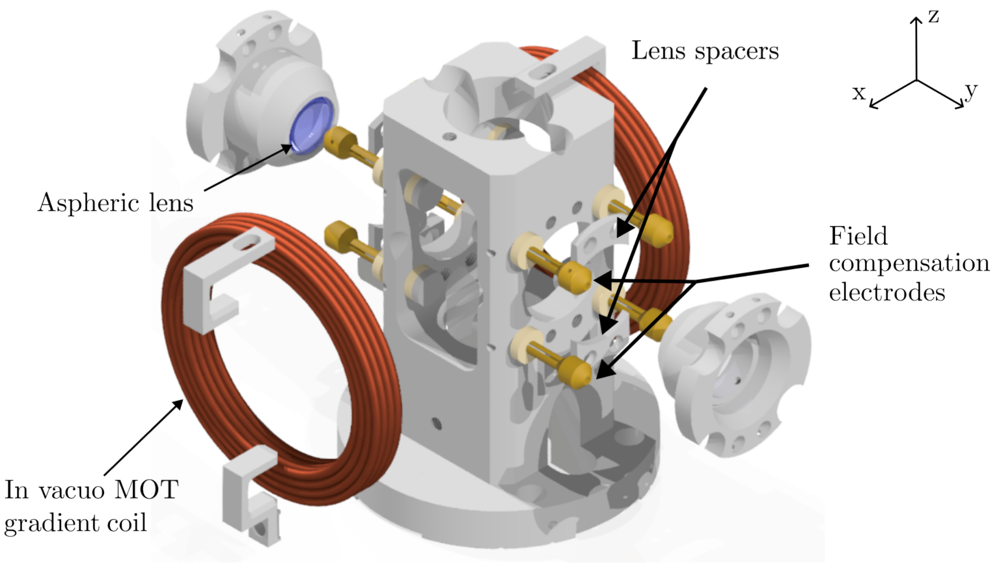

Figure 3.2: Annotated exploded-view diagram of the in vacuo optics system, with laboratory coordinate reference frame. Included in this system are the MOT gradient coils to define the experiment quantisation axis, field compensation electrodes to null electric fields, and the in vacuo lenses. Reproduced from [87] with permission.

coated in a 200 nm layer of indium tin oxide (ITO), a transparent conductor. This layer reduces transmission to approximately $\approx 90\%$ across the range of wavelengths used, with sheet resistance 250 $\Omega/\square$.

Other parts of the lens assembly include the electric and magnetic field components. To compensate for residual electric fields, a series of eight electrodes are fixed in place around the atoms. By controlling the potential on each pin, the field can be corrected. Also included are the in vacuo MOT gradient coil. This coil is made of 1.8 mm diameter Kapton coated wire, for low outgassing in the chamber, and provides the largest magnetic field in the experiment, acting as the field quantisation axis. The in vacuo MOT coil operates with a current of 7 A during the MOT loading sequences, generating a magnetic field gradient of 10 G/cm.

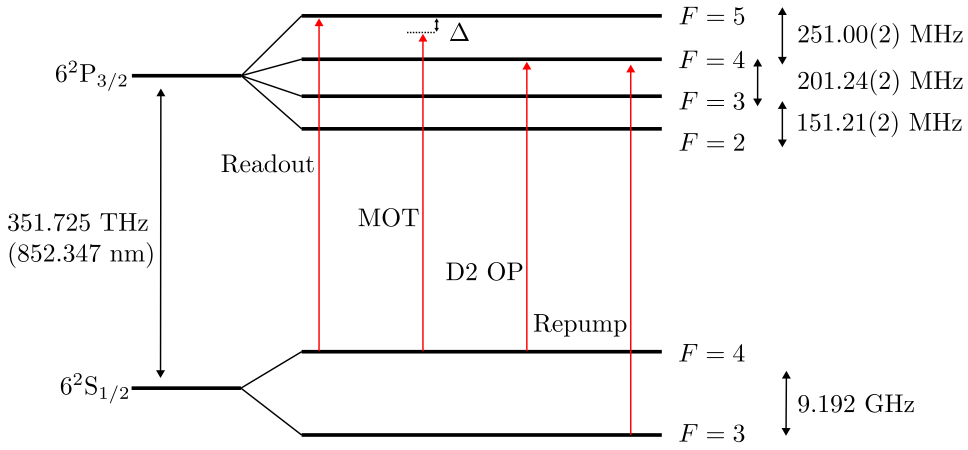

Figure 3.3: Cs Hyperfine structure for the D2 transition. Annotated in red are the transitions used for cooling, imaging, D2 optical pumping and repump. Data from [91].

3.2 Magneto-Optical Traps

3.2.1 Optical Design

For cooling, trapping and state preparation, three Toptica DL Pro diode lasers are used, each controlled by a Toptica DLC Pro. These are kept at a stable frequency through the use of saturation-absorption sub-Doppler spectroscopy. Two diode lasers are used for cooling and trapping atoms in the MOT along the D2 line at 852 nm. One laser is locked to the frequency used for the $|F = 4\rangle \rightarrow |F'\rangle$ transition, and is amplified from 18 mW to 1.5 W through the use of a tapered amplifier. This laser is used for the optical pump, MOT and readout (see Sec. 3.4) beams. The other addresses the $|F = 3\rangle \rightarrow |F'\rangle$ repump transition. The energy levels of these two beams are shown in Fig. 3.3. The third laser is resonant with the D1 line at 895 nm, and is discussed further in Sec. 3.5.3.

Fine frequency control is managed through the use of acousto-optic modulators (AOMs). With the use of a series of AOMs, each laser can be split into a series of paths, where each path can be controlled in both frequency and intensity by the ARTIQ control system [137]. Voltage-controlled amplifiers

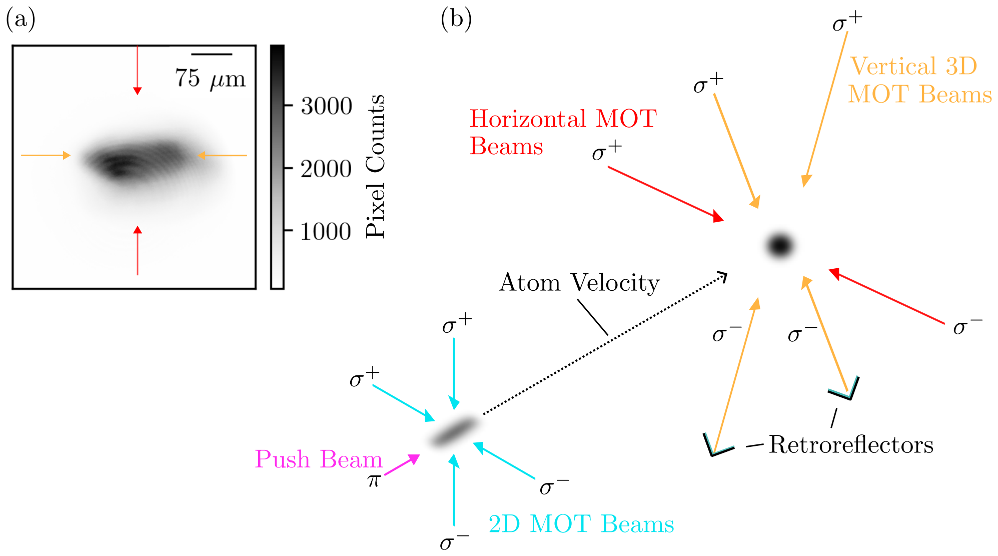

Figure 3.4: Diagram of the 2D and 3D MOT. Figure (a) shows the 3D MOT at a steady state after 5 seconds of loading, with 2 ms of exposure through the primary imaging system. Labelled are directions of the MOT beams as shown in (b). The vertical MOT beams (in yellow) are angled 25° into and out of the image plane. Figure (b) is a schematic of the beam paths. The vacuum assembly is omitted from this diagram for clarity.

are used to dynamically adjust beam intensity and transistor-transistor logic (TTL) signals are used to switch the amplifier on or off, with a rise time of < 20 ns. See Sec. 3.7.3 for further details on TTL hardware.

All of these lasers make use of physical shutters for the total elimination of any leakage light through any axis during the experiment. The shutter design used is based on those in [138], with characterisation in [87].

Atom cooling starts in the 2D MOT in a dedicated cell, as discussed in sec. 3.1.3. A near-resonant, linearly polarised push beam applies an impulse to atoms within the 2D MOT to send low temperature atoms through to the 3D MOT in the main experiment chamber. This is orthogonally to the 2D MOT trapping beams and through the differential pumping tube. A flux of low temperature atoms allows for rapid loading of the 3D MOT for fast experimental rates. See Fig. 3.1 for the spatial arrangement of the vacuum system.

The 2D MOT is operated with a $1/e^2$ beam waist of 10 mm, at a power of 200 mW and detuning $\Delta = -3\Gamma$, where $\Gamma/2\pi = 5.22$ MHz, the natural linewidth of the transition [91]. A resonant repump intensity of 7 mW is incident along the same path. Permanent bar magnets are used to provide a 16(1) G/cm magnetic field gradient within the cell. Trapping only occurs along the transverse plane, such that a large number of atoms can be continuously cooled, with forces orthogonal to the impulse applied by the push beam.

In the 3D MOT, the horizontal MOT beams typically operate with 7 mW of power. These are collimated to a $1/e^2$ waist of $w_0 = 5$ mm, but are clipped to a diameter of 8 mm at the fibre optic outputs by a metal ring to minimise scatter within the chamber by stray light. The two vertical retroreflected MOT beams operate at 5 mW at a $1/e^2$ waist of $w_0 = 2.5$ mm and are clipped to a beam diameter of 4 mm. All paths operate at a detuning of $\Delta = -4\Gamma$, where $\Gamma/2\pi = 5.22$ MHz, the natural linewidth of the D2 transition [91]. Repump light is incident through the vertical paths at a total power of 0.6 mW and through the horizontal σ^+ path at a power of 0.35 mW. The magnetic field gradient is produced through the in-vacuum MOT coils at strength 12 G/cm.

Fine tuning of magnetic field gradients is managed through the use of external shim coils. These coils provide a bias to the zero-field position of the MOT, such that the atom cloud is aligned with the focus of the high-NA imaging lenses. Two coil pairs are active around the 2D MOT, and three pairs for the 3D MOT, for auxiliary field control along each axis.

3.2.2 MOT Performance

MOT loading performance follows an exponential saturation, such that the atom number as a function of time $N(t)$ is given by [139, 140]

$$N(t) = N_{\max}\left(1 - \exp(-t/\tau)\right), \tag{3.1}$$

where $N_{\max}$ is the steady-state atom number and τ is a loading timescale. Imaging the 3D MOT over a period of time demonstrates a maximum atom number of $N_{\max} \approx 10^5$ atoms, and is loaded over a timescale $\tau = 1.9(1)$ s [87].

Atoms within the MOT cloud are at a temperature of 250(10) μK, measured through time-of-flight measurements [122]. Two PGC stages are then applied. The first stage operates at $\Delta = -5.8\Gamma$ for 3 ms, with horizontal MOT beams at a power of 0.8 mW. The vertical MOT beams are reduced in power by the same fraction to 0.5 mW. The second PGC stage operates at $\Delta = -8.0\Gamma$ for 2 ms and at half the beam intensity of the first stage. This reduces the atom temperature to 37(4) μK, also measured through time-of-flight measurements. This is substantially colder than the Cs Doppler temperature of $T_D = 125$ μK. For experimental details on these measurements, see [87].

This sequence is capable of loading a 225 atom ODT array after a 0.25 s MOT loading period (see Sec. 3.3.3), and is thus is sufficient for further experiments on single atoms within the traps.

3.3 Optical Dipole Traps

The generation of the ODTs used in this experiment is achieved through light fields shaped by holographic techniques. Lasers and optical design is discussed below, but it is first necessary to introduce the spatial light modulator (SLM), a tool for the generation of arbitrary holograms.

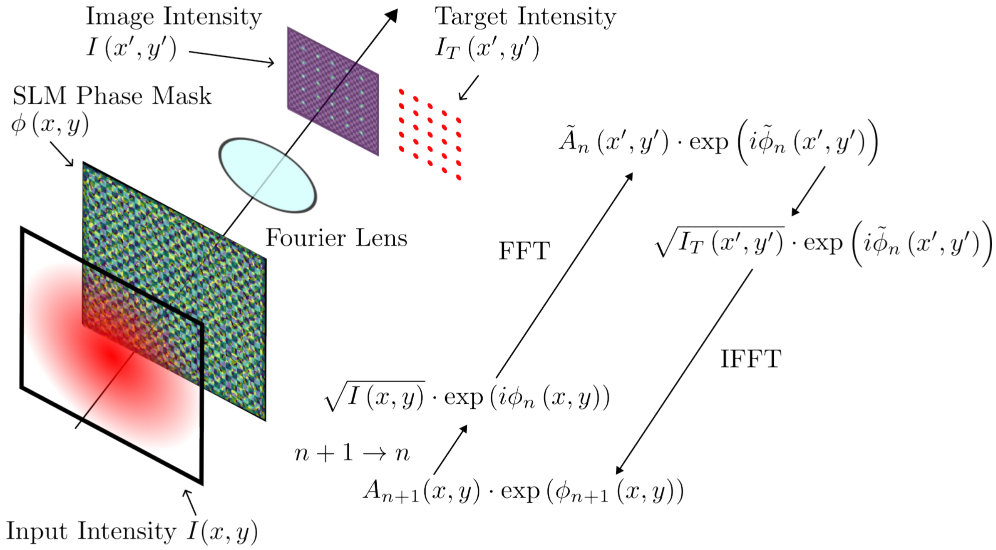

Figure 3.5: Data demonstrating shaping of light potentials using an SLM. Incident light of known intensity and phase (here a collimated Gaussian beam) is incident on an SLM. This modulates the phase, but not the intensity. A lens acts as a Fourier transform between focal planes for the intensity and image fields [141]. Also shown is the process of the Gerchberg-Saxton algorithm for phase retrieval, which modifies the SLM phase $\phi(x, y)$ to match a target intensity $I_T(x', y')$.

3.3.1 Spatial Light Modulators

The principle of the spatial light modulator (SLM) is the same that of a typical display screen. An SLM instead adjusts the phase of light for each pixel deterministically, much like how a display screen adjusts light intensity, using liquid-crystal on silicon (LCOS) cells. This principle extends beyond analogy for the LCOS-SLM that is used in this experiment. Liquid crystal cells on a silicon control panel can be individually controlled as per-pixel retarders. This is effectively a liquid crystal display screen, operating in reflection mode and without polarisation filtering to adjust intensity. The first-order diffraction pattern created by the holographic pattern is used, while the zeroth-order reflection is blocked.

With careful phase control, an input light field can be reshaped into a required output field [142]. A lens acts as a Fourier transformer for this process[141]. This is the fundamental principle of holography [143], although instead of the retrieval of phase information from an image, we are interested in the generation of an image by controlling said phase. Established phase retrieval algorithms, such as the Gerchberg-Saxton algorithm [142], calculate the phase required for an arbitrary intensity. The Gerchberg-Saxton algorithm uses a fast Fourier transform (FFT) to compare a measured amplitude $\tilde{A}_n(x', y')$ and phase $\tilde{\phi}_n(x', y')$ with a target intensity $I_T(x', y')$. The measured amplitude is replaced with the target intensity, and an inverse FFT (IFFT) is applied. The transformed intensity is again replaced with the known intensity, and the new phase is used to update the SLM. See Fig. 3.5 for a diagram of this process.

This experiment makes use of the adaptive Weighted Gerchberg-Saxton (aWGS) algorithm [144], a modification of this scheme for faster convergence. Thus, we can calculate and generate a hologram, display this onto our SLM, illuminate it with our laser, and thus generate an arbitrary intensity distribution. For hologram generation, the phase profile is initialised as a random phase distribution. Repeated FFT and IFFT processes carry out the phase retrieval process without requiring experimental feedback. This generates the intended pattern, with some small error. These errors in the process can be corrected by using the measured intensity on a camera rather than theoretical intensity, thus corrected for systematic differences between the optics and the calculation.

This process has some limitations. Each pixel also has a degree of cross talk, although this can be measured and corrected for [145]. The Gerchberg-Saxton algorithm and its derivatives do not specify phase in the Fourier plane, which can lead to unexpected interference in the final image plane. Real-time feedback can also correct these issues to some extent [144]. Finally, it is limited by the size and number of the SLM pixels. To maximise the control over the field, the largest number of SLM pixels available should be used.

For the experiments described here, two SLMs were used, one for each wave-

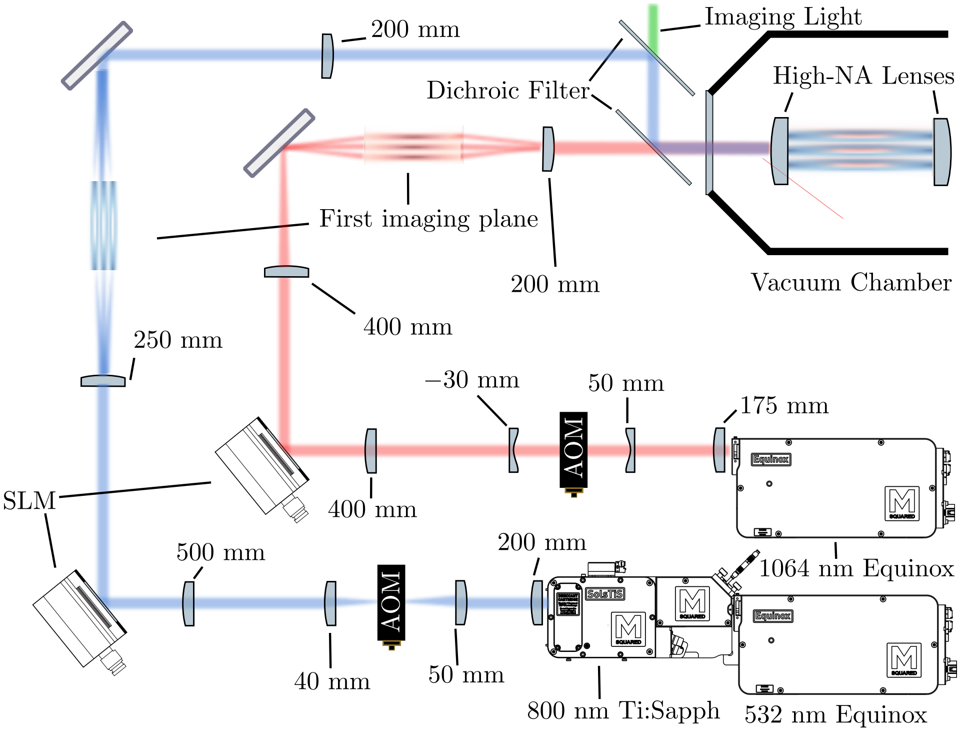

Figure 3.6: Layout of optics in creating red and blue detuned ODTs. The two paths combine on dichroic filters, which are shared with the imaging path. The optical dipole traps overlap inside the chamber. Each path uses an AOM for control of power and fast switching. Measurements refer to the focal length of the lens. The schematic shown is not to scale, and omits mirrors used for beam pathing. Angle of incidence on each SLM is 10° in the experiment to match the manufacturer's specification.

length used in generating ODTs. Each SLM is a Hamamatsu X13138-02WR SLM. Each SLM was used to generate an array of ODTs that overlap in the vacuum chamber. For the red-detuned ODT arrays at 1064 nm, this is an established technique that splits a single high power trapping beam into an arbitrary array of traps, first demonstrated in [146].

In Chapter 6 we will discuss novel techniques that for the first time generate stable arrays of blue detuned optical dipole traps containing Rydberg atoms.

3.3.2 Optical Design

Two lasers are used for the generation of two different types of optical dipole trap.

The first laser, at a wavelength of 1064 nm, is used for the generation of red-detuned ODTs. This laser is a custom M Squared Equinox system modified to provide high power, single mode and narrow linewidth light at 1064 nm.

A second laser, at a wavelength of 800 nm, is used for the generation of blue-detuned ODTs. This laser is an M Squared Equinox diode laser used to pump an M Squared SolsTiS Ti:Sapph laser. The Ti:Sapph laser system allows for a wide range of narrow linewidth output frequencies, but has limited efficiency. As such, with a maximum pump power of 19.5 W, the 800 nm system is capable of outputting 5.1 W. For the 800 nm beam, the shorter wavelength means that the effective numerical aperture is 0.44. For further details on the 800 nm ODT performance, see chapter 6.

Each laser, after beam shaping, passes through a single AOM. This AOM is used for fast switching and intensity control. These beams are expanded through telescope systems and are incident on their respective SLM. The telescopes used for beam expansion are designed such that the maximum area of the SLM is exposed to the incident beam without clipping, in this case a $1/e^2$ beam radius of 6.7 mm is used. This gives the best control over the resultant field, i.e. the greatest number of SLM pixels are used in beam shaping.

Each beam combines on a dichroic along the main imaging path. For the 1064 nm beam, this creates a diffraction limited spot size of 1.04 μm. The calculated spot size, measured by profiling the beam outside the chamber and simulating the optics, creates a focus with $1/e^2$ radius of 1.1 μm, matching this limit.

When further near-resonant light is on the atoms, the resultant AC Stark shift can reduce the effectiveness of the trap, or even cause the atom to be ejected. As such, a ‘chopping’ sequence is used. Here, the ODT light is switched on and off rapidly, with a period substantially less than the fundamental trap frequency [147]. Incident light is also chopped like this but out of phase, such that it is incident only when the ODT is off. As such, interactions that would otherwise be destructive can be extended into the millisecond scale while atoms remain trapped. This includes cooling, imaging, optical pumping and repumping, all of which excite Cs into anti-trapped states.

To generate our chopping signal, we use an FPGA[1] to generate a 50 % duty cycle 1 MHz square wave. This square wave acts as a transistor-transistor logic (TTL) signal to the RF amplifiers that drive the switch AOMs[2] for each path. Other channels from the FPGA can be adjusted in phase, such that the incident beams are in phase or π out of phase as required.

3.3.3 Loading Atoms In Red Detuned ODTs

A typical experimental sequence initialises loading at a trap depth of 3 mK. Light-assisted collisions (LAC) reduce trap occupancy to a single atom per trap, at a rate of 57% (see section 2.5.2). Once loaded, a further PGC stage cools atoms within the traps.

Atoms remain trapped and are reimaged after time t with probability $P(t)$, given by

$$P(t) = e^{-t/\tau}, \tag{3.2}$$

where τ is the trap lifetime. At a background pressure of $\approx 1.0 \times 10^{-11}$ Torr, we measured $\tau = 17(1)$ s, as seen in Fig. 3.7(a).

If an atom is slow and near the centre of a red detuned ODT, the potential acts like a harmonic oscillator (see section 2.5.1). Thermal atoms occupy the

---

[1] Altera Cyclone II EP2C5T144 FPGA

[2] Gooch & Housego 3080-122 at 80 MHz for the 800 nm path, and 3110-197 at 110 MHz for the 1064 nm path.

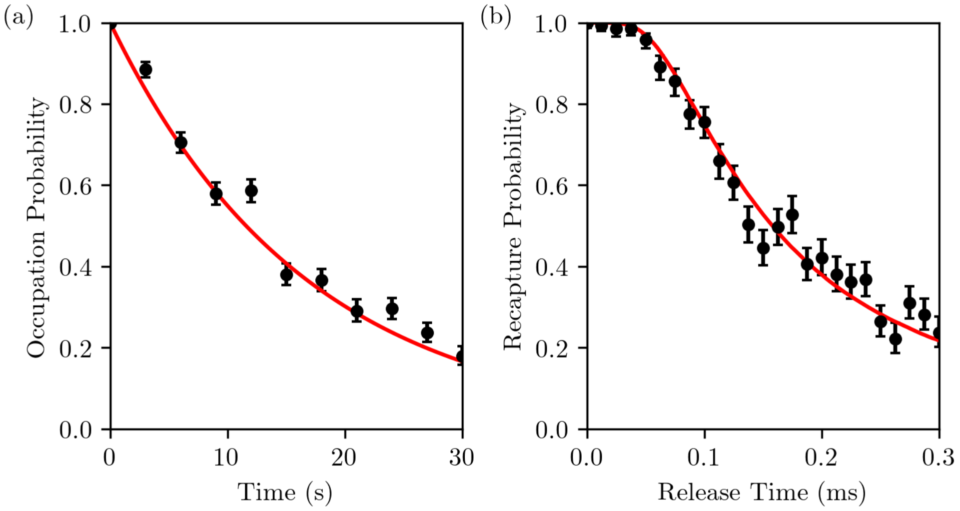

Figure 3.7: Data demonstrating red detuned ODT performance. Figure (a) shows the probability of an atom remaining in a trap over time, with a decay constant of $\tau = 17(1)$ s. The red line is a fit to the data of Eq. 3.2. Figure (b) uses the release and recapture method to find the temperature of atoms within the red detuned traps. The red line is a fit of the recapture probability calculated using the Monte-Carlo method discussed in [148]. This gives a fitted temperature of 1.5(1) μK for a trap depth of 30 μK.

quantised motional states $|n\rangle$ of the harmonic oscillator. The wavefunction $|\psi\rangle$ is, at thermal equilibrium, given by the distribution

$$|\psi(T,k)\rangle = \sum_{N=0} c_N(T,k)\,|n\rangle \tag{3.3}$$

where k is the spring constant of the quantum harmonic oscillator, and $c_n(T,k)$ follows a Boltzmann distribution of loading temperature T of states where $\mathcal{H}\,|n\rangle = \hbar\sqrt{k/m}\,(n+1/2)\,|n\rangle$ and where m is the atomic mass [148]. Thus, the temperature of a single atom is defined by its expected velocity within the trap. This temperature can be measured using a release and recapture technique [148]. Turning off the trap potential allows atoms to travel in free-space. After a period of time, the trap power is ramped back up and the atoms are recaptured. Colder atoms will be recaptured more often given their lower average velocities when dropped from the traps.

As the trapping light is slowly reduced in power[3], the wavefunction evolves adiabatically. In this case, k reduces, but the probability amplitude of states c_N remains constant. Adiabatic ramps of the ODT power thus acts as a further cooling stage for the atoms, although at the cost of delocalisation of the atoms within the traps.

For experiments, the traps are ramped down to a minimum power of 0.5% of the maximum intensity, with 400 μW per trap, or a trap depth 30 μK. Atom temperatures at this stage are measured to be 1.5(1) μK using the release and recapture method, as shown in Fig. 3.7(b). This uses a Monte-Carlo simulation of recapture probability, assuming a thermal distribution of atom velocities within the traps.

This adiabatic ramping sequence is also used during imaging when traps are ramped up to a trap depth of 1.2 mK. This is such that atoms remain trapped during the large number of scattering recoils that occur during imaging.

3.4 Imaging Single Atoms

Atom readout is accomplished through imaging atomic fluorescence. A single atom in an optical dipole trap can scatter sufficient light to be imaged by a camera without ejecting the atom from the trap. We can thus image individual trapping sites to determine if an atom is trapped in that site.

Atoms are trapped in the focal plane of the high-NA in vacuo lenses. These lenses both focus the input trapping light and act as the front lens for imaging. This provides diffraction limited imaging performance[4] up to an object offset of 40 μm, such that a field of view of 80 μm is achieved.

Light is scattered from atoms through a fluorescence process. Cooling and repump beams are again applied to the atoms to excite atoms across the

---

[3] The rate of change of the oscillation frequency is less than the square of the same frequency, as shown also in [148].

[4] Defined as Strehl ratio > 0.8 [149].

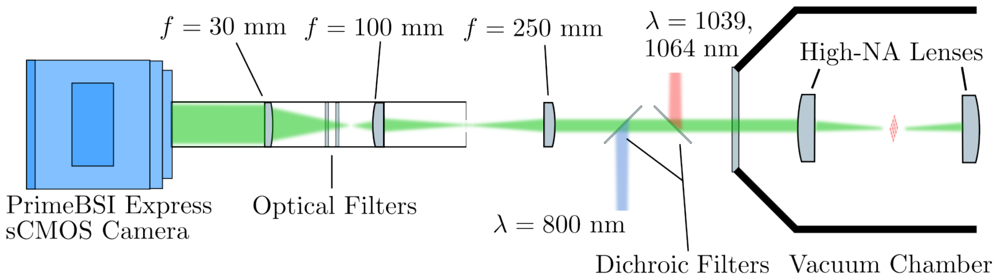

Figure 3.8: Diagram of the optics used for imaging, not to scale. Light at wavelengths of 800 nm and 1039, 1064 nm is combined with the imaging path using dichroic filters. Stray light is blocked from entering the camera through the use of an enclosed path, with a narrow aperture. Optical bandpass filters centered at $\lambda = 850$ nm are used to suppress noise on image acquisition.

D2 transition, with greater detuning and greater power so as to minimise the force on the atoms. This is applied over 40 ms. By adjusting power, detuning and duration of the applied light during the imaging stage, the maximum flourescence signal can be measured without heating the atoms to the degree of removing them from the traps. As this process is across the D2 transition, the atoms emit light at a wavelength of 852 nm.

Light emitted from the flourescence process is isotropically and stochastically scattered. Given a numerical aperture of 0.45, this limits the collection efficiency of the system to 5.4 % of light incident on the first lens. This light is transmitted, magnified and measured across an imaging system seen in Fig. 3.8. Further loss occurs here, due to limited transmission through optics. A further two filters[5] are used to avoid the collection of light outside the range around 850 nm. This suppresses stray light and improves the signal-to-noise for imaging, but further reduces transmission of imaging light. As such, the total transmission of light through the imaging optics is estimated to be 70 % in the best case.

Images are captured with a Teledyne Photometrics Prime BSI sCMOS camera. This camera is capable of rapid imaging and data transfer, with quantum efficiency of 47 % at D2 wavelengths. The camera itself is attached to a trans-

---

[5] Thorlabs FESH0950 shortpass filter and Semrock FF01-850/10-25 bandpass filter.

lation stage, such that it can be precisely positioned for optimal focussing. Considering all factors, the total photon collection efficiency is estimated to be 1.8 %. The camera collects light for 50 ms, longer than the imaging light is applied for, such that the rolling shutter can initialise all pixel rows as required.

For fast calculations and minimising the data storage overhead, experimental images are processed during the experimental sequence. For a given array of traps, calibration images are taken. These calibration images consist of 400 stochastic loads of the trap array, with trap occupied by a single atom. The 10 brightest pixels for each atom across this sample are counted and summed, with the statistical distribution analysed to mark a cutoff count. This cutoff is defined as the minimal infidelity $1 - F$ for this atom, where F is the probability that a present atom produces the counted number of photons or more, as defined by Poissonian statistics. If this sum is above this threshold value, we can assume an atom is present with certainty given by F.. As such, only the pixel counts within a region of interest (ROI) are considered during analysis, where they are compared to these threshold values, rather than the full experimental image. This calibration is measured and stored as preparation for experiments using a given array, and does not occur at the beginning of each experiment.

Atoms are loaded into traps with sub-Poissonian probability p_0, p_1, where p_i is the probability of loading i atoms into the trap [56, 136]. The count distributions are fitted to the distribution [136]

$$P(c) = p_0 G(c_0, \mu_0, \sigma_0) + p_1 G(c_1, \mu_1, \sigma_1), \tag{3.4}$$

where

$$G(c_i, \mu_i, \sigma_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\left(c_i - \mu_i\right)^2 / \left(2\sigma_i^2\right)\right). \tag{3.5}$$

This is a fit to two separate normalised Gaussian curves, one correspond-

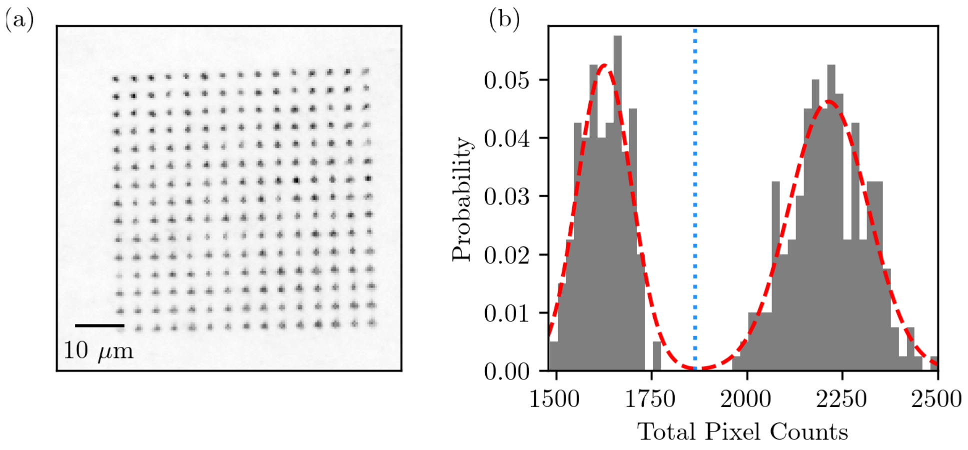

Figure 3.9: Experimental data showing imaging of single atoms within an array. (a) is the average of 400 experimental images of a 15 × 15 atom array, separated by 8 μm. (b) is the histogram of pixel counts across the 10 brightest pixels around the central trap. The red dashed line indicates a fit of two separate Gaussian profiles, with a minimum between the two peaks at the discriminator value (dotted blue line). A summed pixel count above the discriminator value is considered as a present atom. The measured infidelity is here $1 - F < 10^{-4}$

ing to background counts only and other including atom fluorescence. The minimum of this curve between the two Gaussian peaks marks the point of maximum imaging fidelity, and so is used as the discriminator point. An example fit of an array is shown in Fig. 3.9, including the discriminator point at the curve minimum. This method has a predicted fidelity $F > 0.999$ for each atom. Actual imaging fidelity is reduced due to imaging loss, where present atoms are lost during the imaging process. There is a balance to maintaining maximal measurable fidelity while reducing atoms lost during imaging. Imaging loss is measured by the rate of loss of atoms between two subsequent images. Extrapolation of trap lifetime measurements over time t can be extrapolated to $t = 0$ to find this baseline level of loss. This imaging loss is typically of order 1%. Further image processing techniques, such as weighting pixels during counting, can improve measured fidelities, as shown in recent results [55]. Such methods could then reduce the time and intensity of imaging to reduce loss while maintaining acceptable imaging fidelity.

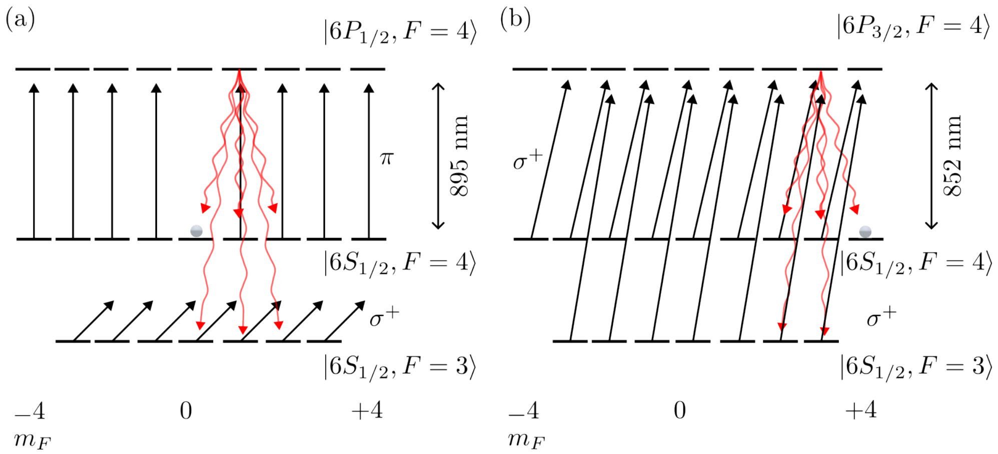

Figure 3.10: Level diagram of optical pumping schema. Figure (a) is clock state preparation, preparing atoms in the $\left|6S_{1/2}, F = 4, m_F = 0\right\rangle$ state using D1 resonant light. π polarised light couples atoms to the $\left|6P_{1/2}, F = 4, m_F \neq 0\right\rangle$ states. Light spontaneously decays from the $\left|6P_{1/2}, F = 4, m_F\right\rangle$ states, shown in red wavy arrows. The σ^+ repump light is along the D2 line, as shown in (b). Atoms accumulate in the $\left|6S_{1/2}, F = 4, m_F = 0\right\rangle$ state. Figure (b) is stretch state preparation, preparing atoms in the $\left|6S_{1/2}, F = 4, m_F = 4\right\rangle$ state using D2 resonant. π polarised light couples atoms to the $\left|6P_{3/2}, F = 4, m_F \neq -4\right\rangle$ states. The σ^+ repump light is also present to prevent atoms being trapped in the $F = 3$ states. Atoms accumulate in the $\left|6S_{1/2}, F = 4, m_F = 4\right\rangle$ state.

3.5 State Preparation And Detection

The execution of many quantum algorithms relies on the deterministic preparation of a target initial quantum state. The light scattering process from the MOT, PGC and LAC stages initialises atoms stochaistically in some state within the hyperfine manifold. We can use two different optical pumping schemes to deterministically prepare all atoms in a single state. To measure state preparation fidelities, microwave pulses are used to coherently couple to hyperfine transitions, and is described in detail in [87].

3.5.1 Clock State Preparation

The $|F = 3, 4, m_F = 0\rangle$ states are referred to as the 'clock' states[6], Given that states with $m_F = 0$ have zero magnetic moment projected along the magnetic field axis, it has a vanishing first-order Zeeman shift applied to it, and is as such least exposed to errors arising from stray magnetic fields[7]. Atomic transition selection rules prohibit the transition $\Delta F = 0, \Delta m_F = 0$, when $m_F = 0$ [95]. As such, all $|F = 4, m_F = i\rangle$ transitions couple to the $|F' = 4, m_F = i\rangle$ state, where $i \neq 0$. Once in the $|F', m'_F\rangle$ state, atoms spontaneously decay to either the $|F = 4, m''_F\rangle$ or $|F = 3, m''_F\rangle$ states. The m''_F quantum number is, due to selection rules, restricted to the set $\{i-1, i, i+1, \}$, such that total angular momentum can only vary a maximum of $\hbar$ per absorbed photon. The same repump laser is used as before to prevent atoms accumulating in the $F = 3$ state, as shown in Fig. 3.10. Atoms continuously scatter through the hyperfine manifold and accumulate in the clock state. For state preparation, a bias field of 5.7 G applies this shift to lift degeneracy between states. If the atom is scattered to the $|F = 3\rangle$ state, it is no longer affected by the optical pumping light.

This optical pumping process is achieved with a single beam, with a second repump beam. The optical pumping beam used for this purpose is aligned orthogonal to the magnetic bias field axis, and is linearly-polarised such that it drives π-transitions in the atom. A Glan-Taylor polariser is used to maximise polarisation purity. These π-transitions are used to maximise scattering fidelity; the D1 OP transition is separated from its nearest states by 1.167 GHz and as such, off-resonant scattering from $F' = 3$ is suppressed. The equivalent transition in a D2 scheme has a detuning of 250 MHz, and so would have worsened optical pumping fidelity. The increased detuning necessitates a third laser as mentioned in sec.3.2.2.

---

[6] Due to its use in atomic clock experiments.

[7] Second-order Zeeman effects do cause small but non-zero energy shifts [92]. For small fields, this is 429 Hz/G^2 by the Breit-Rabi formula [150]

The observed fidelity for clock state preparation reaches a maximum of $\mathcal{F} = 0.974(3)$, as seen by the contrast in microwave spectroscopy [87]. This is believed to be due to polarisation impurities from the D1 light when projected onto the bias field axis. Most likely, this is due to slight misalignments and maladjustments to the magnetic field coil positions, such that the input beam is off-axis to the atomic quantisation axis. This beam has $1/e^2$ waist $w_0 = 800$ μm, with a typical power of 30 μW, and operates on an optical pumping timescale of $\tau_{OP} = 79(1)$ μs [87].

For optimal starting fidelity, a Raman-enhanced optical pumping scheme can be implemented, as first described in [67, 151] and improved in [65]. These schemes utilise stretch state preparation, and then apply composite Raman pulses to transfer to the clock state. Given the large size of the Cs hyperfine manifold, this may prove experimentally challenging for high fidelity operations.

3.5.2 Stretched State Preparation

The $|F = 4, m_F = 4\rangle$ state is known as the 'stretched' state, due to how it receives the largest Zeeman shift of the hyperfine manifold. The same 5.7 G bias field is applied to lift degeneracy.

The $|F = 4, m_F = 4\rangle$ stretched state is generated by optically pumping on the $F = 4 \rightarrow F' = 4$ D2 transition using σ^+ polarised light. This beam is aligned along the magnetic bias field axis. As shown in Fig. 3.10, this beam is resonant along the $\left|6S_{1/2}, F = 4, m_F = i\right\rangle \rightarrow \left|6P_{3/2}, F' = 4, m'_F = i + 1\right\rangle$ transition. As such, atoms are continuously increasing the m_F number to a maximum of $m_F = 4$. Atoms in this state, once excited into the $|F'\rangle$ state, can only decay back to the stretch state, as scattering into the $F = 3$ manifold is forbidden. Once in the stretched state, it is no longer coupled to the σ^+ light.

As with the MOT beams, this beam is clipped by a 4 mm radius aperture, as it shared a fibre with the σ^+ MOT beam. This was measured to prepare atoms into the targeted $|F = 4, m_F = 4\rangle$ state with fidelity $\mathcal{F} = 0.993(7)$. The beam used for this is along the same path as the primary MOT σ^+ beam. This beam has $1/e^2$ waist $w_0 = 5$ mm, with a typical power of 95 μW, operating on an optical pumping timescale of $\tau_{OP} = 9.5(6)$ μs [87].

3.5.3 State-Selective Readout

The simplest and most effective method of state detection for alkali atoms is a destructive, ‘blow-away’ method. A beam near-resonant to the $\left|6S_{1/2}, F = 4\right\rangle \rightarrow \left|6P_{3/2}, F' = 5\right\rangle$ is applied along the MOT- axis. This is applied at high power for long periods (on the atomic timescale) with no repump, such that atoms in the $|F = 4\rangle$ state are rapidly heated out of the trap. When imaging is then applied, the only atoms observed are those in the $|F = 3\rangle$ state. In order to remain trapped, the ODT and blow away light are both chopped out of phase with a 50 % duty cycle.

Beyond imaging fidelity, we must quantify the effectiveness of the state selective imaging. This is the difference between correctly identifying if an atom is present, and ensuring an atom is present when it is in the $|F = 3\rangle$ state or is not present when in the $|F = 4\rangle$ state. This is measured by initially optically pumping in the D1 state, with a further 100 μs repump stage to prevent any remaining population in the $|F = 3\rangle$ state.

A 0.5(2) % false-positive rate is measured, along with a 1.0(1) % false negative rate. The false-positive rate is the probability for an atom to survive the sequence of heating, and is likely due to leakage into the $|F = 3\rangle$ state during blow away. Blow-away occurs through repeated excitation from the $F = 4 \rightarrow F' = 4$ state, and so has a non-zero probability to couple to the $F = 3$ state after repeated excitations. The false-negative rate is due to off-resonant scattering events and loss during imaging. This was measured by, after the blow away sequence, applying the depump light for 20 ms, and measuring the change in survival.

Rydberg atoms are simpler to detect. Given that the extreme electric susceptibility of Rydberg atoms repels them from red detuned ODTs (see section 2.5.1), Rydberg detection is achieved simply by turning on the ODT array used for ground state trapping. Rydberg atoms are thus expelled from the trap and can be distinguished from ground state atoms with some finite infidelity [152].

Non-destructive readout was also implemented, and is described in great detail in [87, Chapter 5]. It is otherwise not utilised in any experiments or procedures described in this work.

3.6 Atomic State Control

Once our atoms are prepared in the desired initial states, we can then manipulate the states that they are in. For atom-atom interactions, we require the use of Rydberg atoms in order to extend the interaction range. For hyperfine control, we can utilise either microwave or Raman transitions. This section will discuss the optics used for generating the Rydberg, Raman and microwave fields.

3.6.1 Rydberg Lasers

As discussed in detail in section 2.2.1, we can use Rydberg states to enable atom-atom interactions. To excite atoms to the Rydberg states, we use two separate lasers to drive a two-photon transition (see section 2.3). This couples the $\left|6S_{1/2}F = 4, m_F = 0\right\rangle$ state to the $\left|nS_{1/2}\right\rangle$ state, and so requires optical pumping to the clock state. This transition acts via the $7P_{1/2}$ state, and is detuned by frequency Δ from the intermediate state centre-of-mass

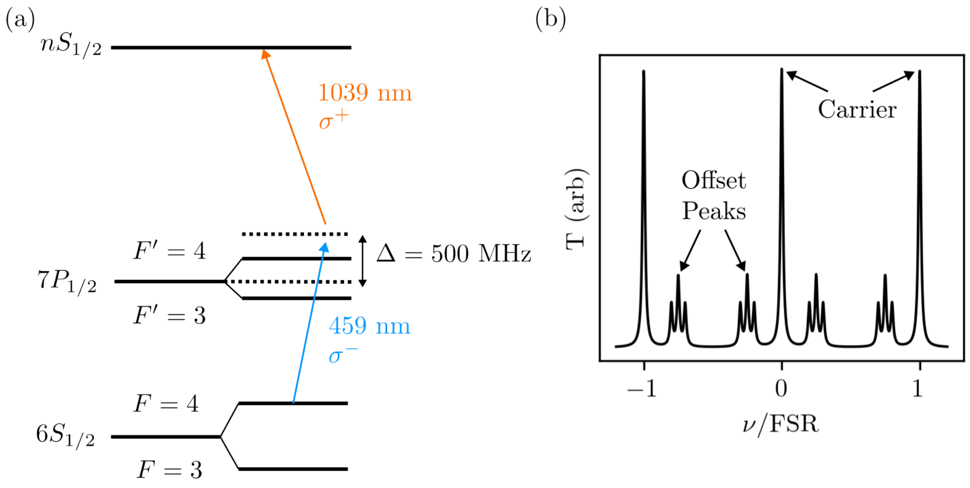

Figure 3.11: Figure (a) is the energy level diagram to the Rydberg transition. There are two optical paths to available, via the $F' = 3$ and $F' = 4$ sublevels. (b) shows the transmission through the ULE cavity using a PDH-offset lock peak. The carrier peaks are when input light is resonant with a cavity mode. This input light is mixed with an offset frequency through an EOM, creating offset peaks with associated sidebands. These offset peaks are separated from the carrier by the offset frequency and are used for locking to arbitrary frequencies.

frequency.

For experiments detailing the effectiveness of the Rydberg system for Rydberg excitations, see Chapter 5.

459 nm Beam

The first beam in this transition, coupling the $|6S_{1/2}\rangle \rightarrow |7P_{1/2}\rangle$ states, is at a wavelength of 459 nm, and is σ^- polarised. Note that this coupling is used, and not the $|7P_{3/2}\rangle$ state, due to the narrower linewidth of the transition. This increases the required power to address this transition, but reduces scattering from the intermediate state by the same fraction. The laser used for this transition is an M Squared SolsTiS Ti:Sapph diode laser, pumped by an M Squared Equinox diode laser at 532 nm. This outputs a beam at a wavelength of 918 nm. Part of this beam is split to couple to the ULE cavity (see below) for locking. The rest is converted to 459 nm using an M Squared

ECD-X second harmonic frequency generation external cavity. At an input pump power of 14.5 W, the ECD-X can output > 1 W of power, although this level of power was not typically available or used.

The 459 nm beam is then split into two paths, ‘global’ and ‘local’. ‘Local’ is yet to be used in an experimental capacity, but remains as a region for development. ‘Global’ utilises a separate double-pass AOM connected to an AWG (see 3.7.3). After fibre coupling, the beam is incident on the atoms in a Gaussian beam with $1/e^2$ waist of $w_{459} = 90$ μm. This beam is typically incident on the atoms with a total power 45 mW.

1039 nm Beam

This beam couples the $\left|7P_{1/2}\right\rangle \rightarrow \left|nS_{1/2}\right\rangle$ transition. It is at a wavelength of $\lambda \approx 1039$ nm, and is σ^+ polarised. Note that the wavelength here depends on the state n that it is coupled to; larger n requires a greater energy, with the highest used in this experiment at $n = 80$. The 1039 nm frequency is adjusted to address the required Rydberg transiton. The laser used for this transition is also an M Squared SolsTiS Ti:Sapph diode laser, pumped by an M Squared Equinox diode laser. At a maximum input pump power of 18.5 W, the SolTiS can output up to 2.1 W of power.

The 1039 nm beam is also split into ‘global’ and ‘local’ paths, however the local 1039 nm is also used for atom array rearrangement. For details on this, see Chapter 4. The global path is, after coupling to a fibre optic cable, reshaped before being incident onto the atoms. In this case, cylindrical optics are used to create a beam of waist $w_{x,y} = 58.9(1), 21.0(1)$ μm. This reshaping is done to maximise intensity on the 2D plane of the atoms while minimising variation across the array. The Rabi frequency across the array is not entirely homogenous, however. The waist can also be found by measuring the Rabi frequency across the array and fitting to a Gaussian profile. This matches measurements found outside the chamber.

During the course of this thesis, a further fibre amplifier was installed[8]. This amplifier can amplify 50 mW of light to 20 W, greatly increasing the accessible power. Only a fraction of this power is currently used due to the power limits of the optics used. Experiments can currently use up to 1 W of power at the atoms without adjusting optics. Future experiment improvements will adjust to higher power optics such that 16 W will be available before beam shaping optics, which will be replaced with an SLM for greater beam homogeneity using a tophat beam [153].

This beam is colinear and counter-propagating to the 459 nm, to provide a degree of Doppler-dephasing cancellation.

ULE Cavity Lock

The fidelity of Rydberg based quantum gates is highly dependent on frequency and phase noise [152, 155], and as such these must be minimised. For this purpose, both lasers are stabilised to an Ultra-Low Expansion (ULE) cavity for narrow-linewidth locking, provided by M Squared Lasers. This cavity was measured to have a finesse of $F = 40.9(6) \times 10^3$ at 918 nm using a ringdown measurement [156] and a finesse of $F = 24.0(1) \times 10^3$ at 1039 nm from the full-width half-maximum transmission signal. These measurements can be seen in Fig. 3.13. Ringdown measurements could not be taken at 1039 nm due to afterpulsing effects causing expectedly large decay times in the photodiodes used [157]. The free spectral range (FSR) for the cavity was measured to be 2.995(1) GHz.

Locking was achieved using a Pound-Drever-Hall offset locking technique [158], using Moglabs Fast Servo Controller systems. While a typical cavity lock would be capable of locking to just the carrier signal of the cavity, an offset lock technique provides the capability of mixing in an arbitrary signal to the cavity mode to modulate the transmission signal, as shown in [159].

---

[8] Azurlight Systems ALS-IR-976/1040 nm High Power Fiber Laser

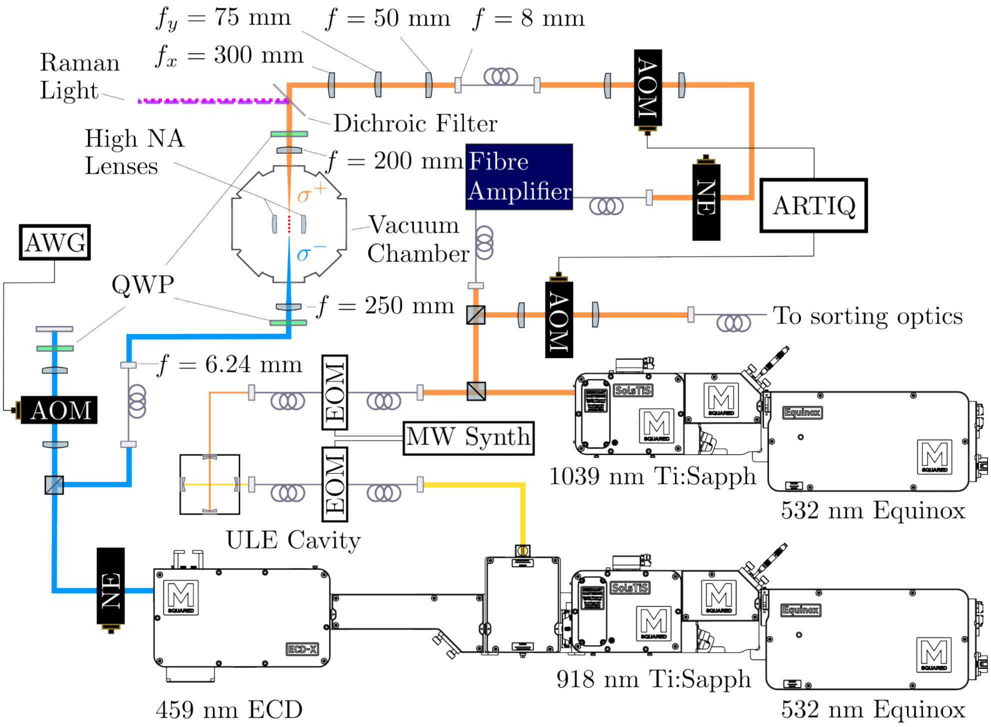

Figure 3.12: Diagram of the Rydberg optical paths. ‘AWG’ refers to the arbitrary wave generator used to modulate the 459 nm light. This path contains a double-pass AOM arrangement, utilising a cat’s-eye retroreflector arrangement [154]. The ‘QWP’ refers to quarter wave plates. The ultra-low expansion (ULE) cavity is used for PDH locking, with feedback to the Ti:Sapph lasers for each path. ‘EOM’ refers to the electro-optic modulators used for the PDH locking. These are connected to microwave synthesizers, controlled by the Artiq system, to provide the PDH locking sideband frequencies. A fibre amplifier is used to amplify the 1039 nm light. $f_{x,y}$ labelled lenses refer to cylindrical lenses, which focus only on one Cartesian axis. These are used to reshape the 1039 nm beam for a narrow focus within the plane of the trapped atoms, with uniform intensity along its width.

With such a method, one can lock to an arbitrary frequency, within range of a cavity mode. See fig.3.11 for an example transmission spectrum.

3.6.2 Microwave Control

Microwaves are used for global state control of qubits within the experiment. By applying a strong oscillating microwave field, the hyperfine transitions of Cs can be directly addressed in single-photon magnetic dipole transitions.

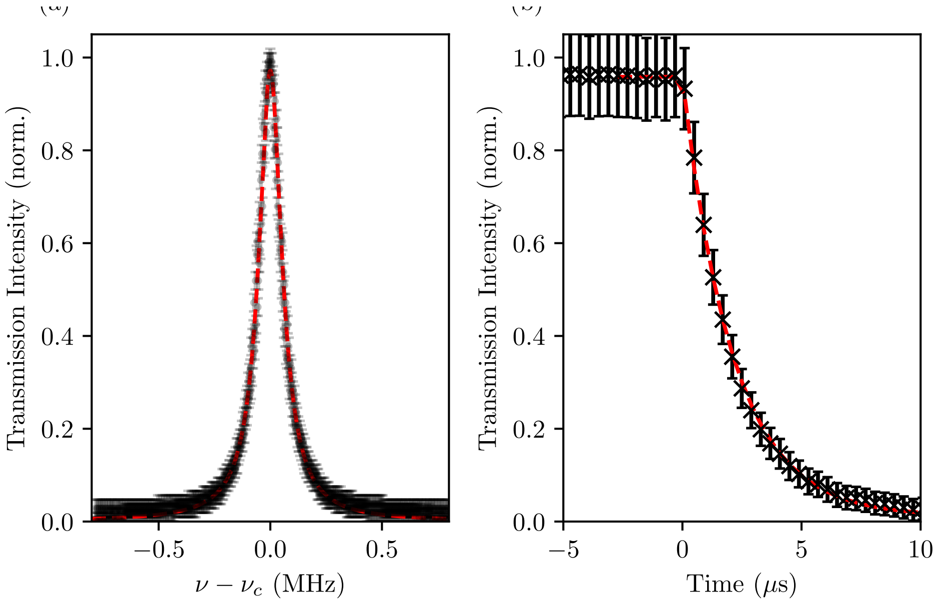

Figure 3.13: Cavity characterisation measurements. Figure (a) is a transmission peak at 1039 nm through the cavity. The red line fits this to a Lorentzian profile with FWHM of 120(1) kHz, indicating a cavity finesse of $F = 24.0(1) \times 10^3$. Fibre (b) is a ringdown measurement at 918 nm, measuring transmission signal as the input light is stopped at $t = 0$. This has an exponential decay of 2.17(3) μs, indicating a cavity finesse of $F = 40.9(6) \times 10^3$.

This can be used for high-fidelity single-qubit gates. These microwave transitions are slow relative to the Rydberg lifetime due to the weak magnetic dipole moment, even for powerful applied MW sources, and so is only suitable for slow but highly coherent global pulses.

To maximise gate fidelity, the phase noise of the MW source must be minimised. For this purpose a Phase Locked Dielectric Resonator Oscillator[9] is used, with the phase noise of −100 dBc/Hz at 1 kHz. This outputs at a frequency of 8.95 GHz, which also acts as the experiment master clock. This is mixed with an input signal from a 14-bit DDS to match the Cs 9.192 GHz transition frequency. As discussed in section 2.1.3, the phase of the qubit in the lab frame oscillates rapidly. As such, it is convenient to use a co-rotating frame with each qubit. Phase offsets in this co-rotating angular velocity correspond to phase delays on the qubits. As such, a phase shift of the MW

---

[9] Polaris Wave Single Loop PLDRO SPLDRO-RE10-8950-P13-2P

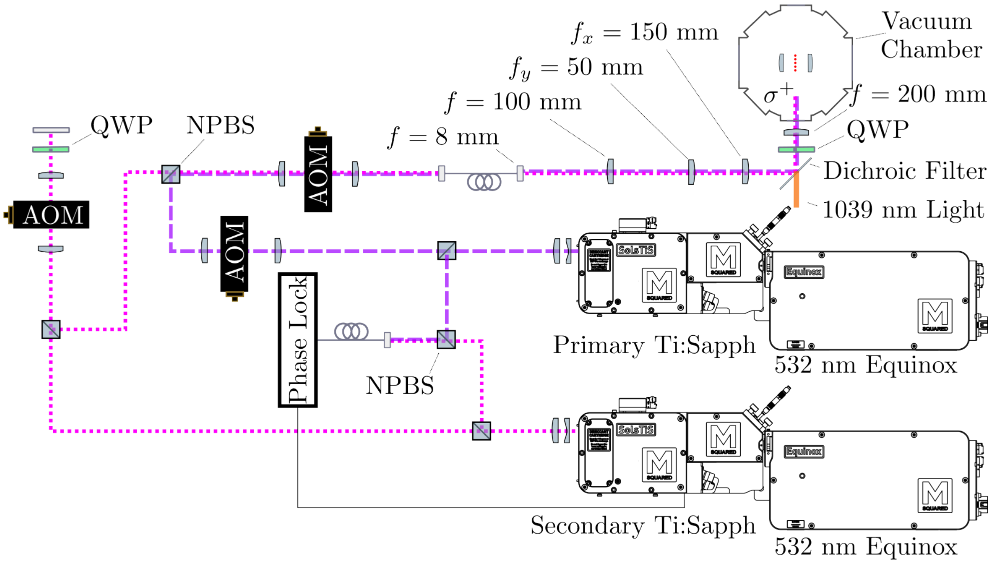

Figure 3.14: Diagram of the Raman optical paths. The two lasers (primary and secondary) are combined on non-polarising beam splitters (NPBS) twice. The first is to generate a beatnote signal to lock the secondary laser to the primary. The second is to combine the beams before coupling into the fibre before the chamber. $f_{x,y}$ labelled lenses refer to cylindrical lenses, which focus only on one Cartesian axis. These are used to reshape the Raman beam for a narrow focus within the plane of the trapped atoms, with uniform intensity along its width. This path is combined with the 1039 nm beam using a dichroic filter.

input acts as a rotation about the y-axis of the Bloch sphere.

Once mixed with the input signal from the DDS, the MW signal passes through a 9.2 GHz bandpass filter to minimise unwanted frequencies, such as the initial 8.95 GHz signal before the mixer. A further three amplifier stages apply low noise amplification to the signal before the microwave horn. The resultant 10 W signal is produced ≈ 15 cm from the atoms. As shown in [64, 87], this system is capable of high-fidelity single qubit gates with errors $7(2) \times 10^{-5}$ as measured using randomised benchmarking.

3.6.3 Raman Lasers

Hyperfine transitions of atoms can also be addressed through Raman transitions. As introduced in Sec. 2.3.2, two lasers with a relative frequency difference δ can address the same hyperfine transition as microwaves through a two-photon process. This requires two separate lasers with stable relative frequency difference.

These lasers are detailed in Fig. 3.14. Each laser is again an M Squared SolsTiS Ti:Sapph diode laser, pumped by an M Squared Equinox diode laser. One laser, the 'primary', addresses the $\left|6S_{1/2}, F = 4, m_F\right\rangle \rightarrow \left|6P_{1/2}\right\rangle$ transition, detuned from this resonance by Δ_{Primary}. This laser is kept at a constant frequency by a slow PID feedback signal from a HighFinesse WS7-60 wavemeter. The second laser, the 'secondary', addresses the $\left|6S_{1/2}, F = 3, m_F\right\rangle \rightarrow \left|6P_{1/2}\right\rangle$ transition, detuned from this resonance by $\Delta_{Secondary}$. The average detuning is defined by $\Delta \equiv (\Delta_{\text{Primary}} + \Delta_{\text{Secondary}})/2$ and two photon detuning $\delta \equiv (\Delta_{\text{Primary}} - \Delta_{\text{Secondary}})$ A small proportion of light is taken from the beams, which is combined light and sent to a fast photodiode to be used as the input for a phase lock. The frequency of the secondary is locked to the primary through this phase-lock system[10]. The 8.95 GHz signal from the MW source is used as the frequency reference to match. These beams are then combined on a non-polarising beam splitter (NPBS) and fibre coupled to the chamber.

This results in a Raman source of average detuning $\Delta = 1.5$ THz and two-photon detuning $\delta = 0$. The average detuning can be adjusted as required to increase the Rabi frequency, at the expense of an increased probability of scattering from the intermediate state.

Each Raman beam is incident on the atoms globally with σ^+ polarisation, up to 400 mW on each beam. Like the 1039 nm Rydberg beam, cylindrical lenses are used to reshape this beam, with a $1/e^2$ waist of $w_{x,y} = 30(1), 101(2),\ \mu$m. As such, the intensity of this beam is spatially inhomogenous over the array due to the Gaussian profile of the beam. As shown in 3.15, the Rabi frequency of atoms on each row changes as the intensity profile of the beam. Given

---

[10] A bespoke system provided by M Squared Lasers

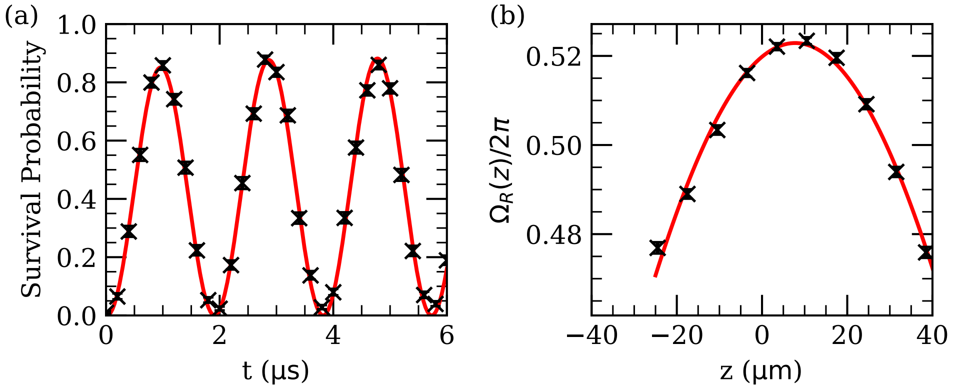

Figure 3.15: Experimental data of Rabi oscillations. Figure (a) shows the Rabi oscillations across the central rows of a 13×13 array at 7 μm spacing. Atoms are prepared in the $|F = 4, m_F = 0\rangle$ state. The Raman light causes a Raman transition to $|F = 3, m_F = 0\rangle$ state with $\Omega_R/2\pi = 0.523(7)$ MHz. Error bars are on scale of point size. These oscillations are fitted to the data in red. Limited contrast of the oscillations is due to imperfect D1 state preparation and off-resonant driving. Figure (b) plots the Rabi frequency of each individual row of the atom array. The red line plots the fit of a Gaussian profile to this data, to show the 101(2) μm $1/e^2$ waist of the input beam in the row axis of the array. One row of atom data is excluded due to poor loading statistics.

that the Rabi frequency $\Omega \propto \sqrt{I}$, and $I \propto \exp\{-2r^2/w_0\}$, the Rabi frequency across the array has a $1/e^2$ waist that matches that of the input beam. For atoms at the centre of the beam compared to z μm away, a π pulse on the centre transferrs > 99 % up to $z = 25$. As such, the number of atoms that can be addressed for high-fidelity qubits is limited by the spatial configuration of the array, and the maximum permitted infidelity with Raman interactions.

3.7 Experimental Sequence Design

The SQuAre experiment uses a large degree of automation to control and execute the experiment. This section will introduce systems that are required for experimental control.

3.7.1 ARTIQ

The SQuAre experimental platform has been designed from the ground up to make use of the 'Advanced Real-Time Infrastructure for Quantum physics' (ARTIQ) control system from M-Labs[137]. ARTIQ is a semi-open source hardware and software control system, specifically designed but not exclusively applicable to, atomic optics experiments. ARTIQ is controlled by a software library, in this experiment the 'Artiq6' package for Python-3[11]. Experiment files to interface with software can thus be written in Python, to be compiled for the experiment hardware.

ARTIQ hardware consists of a collection of interfacing components. The Kasli module acts as a master controller synchronising attached modules with nanosecond resolution. This is based on Artix-7 FPGA components running software precompiled software and so can access fast and precisely timed operations. Such modules from Sinara[12] include transister-transistor logic (TTL) inputs and outputs[13], direct-digital synthesiser (DDS) outputs and microwave synthesisers for signal outputs [14], and digital-to-analogue converters (DAC)[15] for analogue output signals.

3.7.2 Arbitrary Wave Generator

For precise modulation of input frequencies, an arbitrary waveform generator (AWG) is used[16]. This is a high-speed 16-bit digital output device that can generate arbitrary waveforms, with dual-channel output at 1.2 GS/s, with a bandwidth of up to 600 MHz. By connecting the output of the AWG to an

---

[11] Since the installation of the ARTIQ software and hardware into the experiment, Artiq7 is available, and Artiq6 is no longer supported.

[12] https://sinara-hw.github.io/

[13] Sinara DIO BNC and Sinara DIO SMA digital input-output controller

[14] Sinara AD9910 Urukul cards at 1 GS/s, ARTIQ Mirny cards between 50 MHz and 4 GHz

[15] Sinara Zotino 16-bit resolution DAC

[16] Spectrum Instrumentation M4i.6631-x8 Arbitrary Waveform Generator

AOM (after amplification), the input light can be arbitrarily modulated in both amplitude and frequency.

Modulated light can be spatially separated from the carrier signal by modulating a further diffraction order. The first diffracted order of the AOM is what is modulated in these experiments, such that the unmodulated light is blocked.

This AWG is synchronised to the experiment by use of an ARTIQ DDS output acting as a clock signal. Waveform outputs are triggered through ARTIQ TTL connections.

3.7.3 Timeline

A typical experiment takes place over the course of less than one second. See Fig. 3.16 for a schematic demonstrating a typical experimental sequence.

The largest part of an experiment duration involves the loading of the 3D MOT over 250 ms. This is then cooled using PGC stages, trapped in ODTs and reduced to single occupancy in LAC. A MOT fall-away period of 100 ms allows atoms remaining in the MOT cloud to disperse for clear imaging. An imaging stage over 50 ms takes the first image. If required, a sorting stage rearranges atoms. This duration is dependent on the size of the array, but is typically on the order of ≈ 100 ms. The approximate sorting duration is calculated in Sec. 4.6. If sorting has occured, a second image will be taken to measure the initial state before the experiment occurs. This sorting stage can repeat. Optical pumping prepares atoms in the required quantum state over 10 ms. The traps are then ramped adiabatically to a minimum trap depth at experimental power. After the experimental stage concludes, they are ramped adiabatically back up to imaging power and the atoms are imaged for a final time.

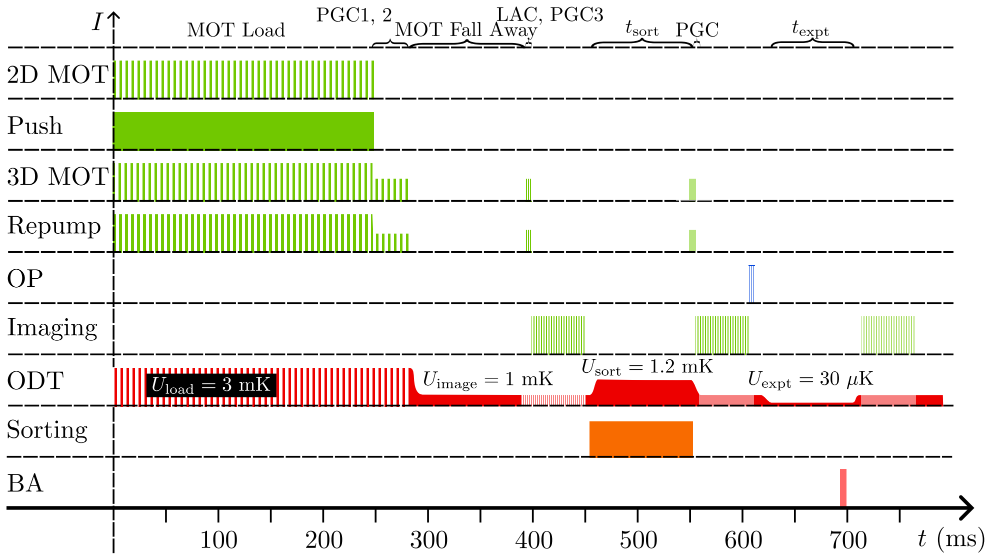

Figure 3.16: Example timeline for experimental sequence. Bar height refers to beam intensity at that time, and is not to scale. Timings are variable, depending on the experimental sequence required. Here, t_{sort} and t_{expt} refer to the time taken for sorting and physics experiment sequences. 'OP' refers to either D1 or D2 optical pumping. 'BA' refers to state-selective blow away. LAC refers to light-assisted collions, and occurs before the third stage of PGC. The final PGC stage repeats PGC3 parameters, after sorting but before the second imaging stage. U_i refers to the trap depth of the trap in that stage, controlled by adjusting the relative power of the ODT array. U_{expt} is the minimum trap depth reached, typically at 30 μK, or 1% of maximum power during the atomic experiment phase, where trap lifetime and atomic coherence must be maximised. The vertical lines on some sections refers to periods where the beam is chopped. This is a 1 MHz square wave intensity, with the ODT π out of phase with other chopped beams.

3.8 Summary

Throughout this chapter, the structure of a typical experiment is detailed. Cs atoms in room temperature vapour are trapped and cooled in a 2D MOT cell before being injected into the main experiment chamber. Further trapping and cooling in three dimensions prepares atoms for trapping in optical dipole traps, where polarisation gradient cooling further reduces atom temperature. Said traps are reduced to single occupancy by light-assisted collisions, such that each trap contains either one or no atoms. This can be seen through sub-Poissonian statistics of photon counts in fluoresence imaging of each trap. Removal of atoms through state-selective blowaway can determine the hyperfine state of these atoms.

Atomic physics experiments from this point involve the preparation and manipulation of the atomic states. Atoms are prepared in the $|F = 4, m_F = 0\rangle$ clock state through D1 optical pumping or the $|F = 4, m_F = 4\rangle$ stretched state through D2 optical pumping. Hyperfine state manipulation is achieved through Raman transitions coupling to the atoms' electric dipole moment, or through microwave light coupling to the atoms' magnetic dipole moment. Atoms interact with each other in Rydberg states, which can be accessed through two-photon transitions. The two lasers used for these transitions are frequency-stabilised by locking to an ultra-low expansion reference cavity. The ARTIQ experimental system is used for synchronisation and control of the sequence.

Later chapters will use these processes as a baseline for experimental procedure.

Chapter 4

Atom Array Assembly

4.1 Introduction

This chapter will discuss the procedure for the deterministic preparation of arrays atoms in arbitrary geometries. The transport and arrangement of neutral atoms within an array is a recent but established technique. Any neutral atom system that requires the level of discrete and deterministic control of a quantum computer or quantum simulator necessarily requires use of these techniques. This chapter will introduce, describe and characterise the systems used for this process in the SQuAre experiment, and is the first comprehensive look at atomic transport of Cs atoms at scale.

Experimental preparation of any neutral atom array system requires array assembly. As established in Sec. 3.3.3, holographic tweezer traps can be made to trap hundreds of atoms in our system. Each trap site contains identical qubits, distinguishable by position. Qubit position is then used for programming. Atom-atom interactions (primarily the Rydberg dipole-dipole interaction) are dependent on the relative positions of other atoms. Furthermore, the atom-light interactions can change depending on the position of the atoms, e.g. a Gaussian beam being less intense on the edges of a beam as shown in Fig. 3.15. For deterministic and repeatable experiments, arrays must be initialised in the same configuration each time, or post selected such that only the correct configuration is selected for.

This particular property of atomic arrays raises a problem in that single atom loading into traps is stochastic [56]. As detailed in section 2.5.2, the initial loading of atoms within an ODT is random, while LAC removes atoms by pairs; if an odd number of atoms is initially loaded, then the trap will be loaded with a single atom. A loading rate of approximately 50% for each trap can then be expected, where each trap's loading probability is statistically independent of any other trap. For an experiment requiring an array of N atoms, the ODT array must have $\geq 2N$ traps. After LAC has reduced each site to single occupancy, these atoms must be moved to the traps required to match the target array.

The probability of a successful sort depends on the success rate of each arrangement step, or 'move'. For m moves to target N sites, the probability of a defect free array P_{DF} is given by

$$P_{\mathrm{DF}} = (P_{\mathrm{Move}})^m \cdot (1 - P_{\mathrm{Loss}}(t))^N , \tag{4.1}$$

where a 'defect' refers to a target ODT that is not loaded. Only defect free arrays can be post-selected as a valid preparation of the experiment. The first term is best understood as the per-move statistics. P_{Move} is the probability that the atom is transferred into the mobile tweezer as the power is ramped up at the start point, transferred to a different site, and then recaptured as the power is ramped down at the destination, all without loss. P_{Loss} is the probability that any atom in any site is lost through the process. This loss is independent of any other atoms and is caused by vacuum limited lifetime and imaging loss, and as such is approximated by

$$P_{\mathrm{Loss}} \approx 1 - \exp\left(\frac{-t_{\mathrm{sort}}}{\tau}\right) + P_{\mathrm{IL}}, \tag{4.2}$$

where τ is the trap lifetime and P_{IL} is the probability of imaging loss. Each imaging step requires each atom to scatter many photons (see Sec. 3.4). This process can eject atoms from traps, and so for each image there is a non-zero probability that atoms otherwise present are lost.

Here we can begin to see the dominant effects. Even if $1 - P_{\mathrm{Move}}$ or P_{Loss} is small, large m or N can result in low experimental yield. A trap site that is not loaded can be considered a 'defect', such that an array that is fully loaded without errors is 'defect-free'. If just 1% of atoms in a 1000 atom array are expected to be lost, $P_{\mathrm{DF}} = 4 \times 10^{-5}$, an unacceptably low yield. Therefore, future neutral atom system must minimise probabilities of loss and maximise any sorting efficiencies in order to remain scalable. Furthermore, since the probability of lifetime limited loss or errors in sorting will never be reduced to zero due to finite sorting times, scaling to multiple sorting rounds will be necessary to correct for loss to correct for errors.

The SQuAre experiment is capable of trapping 289 atoms in a 17×17 grid, limited by available trapping laser power and field-of-view of the imaging lens. This Chapter will detail the process of using a single movable tweezer to sort these atoms into a 100 atom 10×10 grid with an experimental yield of 8%, as well as the deterministic arrangement of atoms into arbitrary geometries. Also detailed is the requirements to maximise efficiency and establish effective scaling, as required for a neutral atom quantum computer.

4.1.1 Sorting Benchmarks

The most often found method of generating neutral atom arrays is through the combination of a static array of red detuned ODTs and one or more moving arrays for atom rearrangement, as shown in Fig. 4.1. Static arrays that then undergo rearrangement can be generated by acousto-optic deflectors (AODs) [160–162], spatial light modulators [163, 164], optical lattices [69, 165] or by microlens arrays [166, 167]. While these experiments differ in

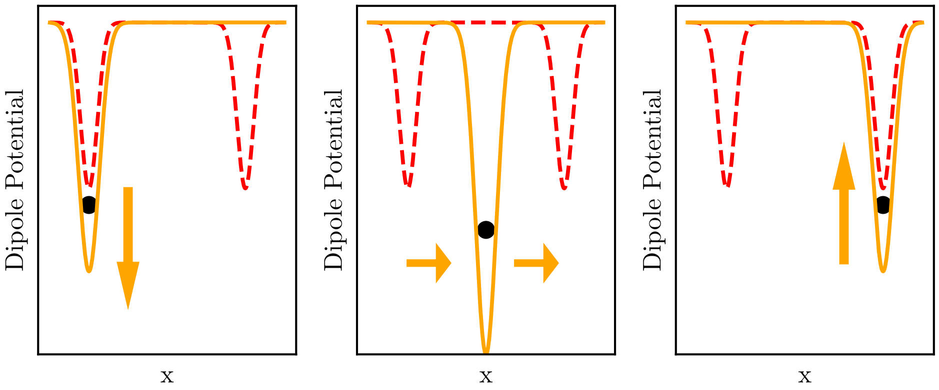

Figure 4.1: Diagram of the trap scheme with static and mobile tweezers. A series of relatively shallow static traps (in red dashed lines) hold atoms (black circles) in place. A deep mobile potential picks up atoms and moves them to other static potentials.

how generate a static array of ODTs, all use one or more secondary beams that can address individual sites. These beams are controlled using a 2D acousto-optic deflector (AOD) to adjust beam angle, and thus position of the trap in the focal plane.

Figures of merit for each experiment depend on a number of factors. Firstly, the target size N must be considered, where N is the total number of traps that must be filled deterministically. Secondly, the average fill fraction F must also be considered. This is the expected fraction of the N target sites that are filled with each experiment. Thirdly, the defect free probability P_{DF} is also a critical measurement.

Table 4.1 is a non-exhaustive overview of sorting results in the literature. Given that each experiment is unique and has separate design goals, it should be taken as a rough overview of what is capable with atomic rearrangement techniques.

Publication	Species	F	N	$P_{DF}(\%)$
This work	Cs-133	0.959(1)	100	6
Paris Saclay, 2016 [58]	Rb-87	>0.96**	25	40
Paris Saclay, 2022 [164]	Rb-87	0.972(4)*	196	2.5
Harvard, 2021, [163]	Rb-87	"∼0.985"	256	Unstated
Singapore, 2023, [161]	Rb-87	0.97766(4)*	225	1.5(3)
Wuhan, 2021, [168]	Rb-87	0.945(7)	30	Unstated
Darmstadt, 2019 [166]	Rb-85	0.88(7)	100	3
Darmstadt, 2024 [167]	Rb-85	0.97(1)	256	14
MPQ Garching, 2024 [169]	Sr-88	>0.98	>1247	0
Atom Computing, Inc., 2024 [170]	Yb-171	>0.99	1225	0

Table 4.1: Comparison of sorting performance in static arrays. Data refers to filling a target of N atoms with fill fraction F, filling perfectly at a rate P_{DF}. "*" indicates data quoted after correcting for loss, such as the vacuum limited lifetime. "**" indicates measurements taken after a single rearrangement stages. Otherwise, the data indicates multiple rounds of sorting. Note that the final two experiments both intend as large an atom number as possible on a large array over dozens of cycles, and do not claim to ever generate a defect free array.

4.1.2 Acousto-Optic Deflector Arrays

In some neutral atom array experiments, the light used for the ODT array is not necessarily static. An ODT array can be generated by using acousto-optic deflectors with a mix of RF input signals. As the deflection angle (and thus trap position) is dependent on the input frequency, the array traps can be moved in situ by changing the input frequency spectrum, rather than using a separate sorting beam.

This experimental design has some benefits over the static/moving experiment paradigm in regard to atomic rearrangement. As atoms do not need to be transferred between tweezers, this eliminates the time required for increasing and decreasing the power of sorting tweezers, and potential loss in doing so. Furthermore, it reduces laser power requirement, as an auxilliary beam is not used for sorting. This is demonstrated in [57, 162] along a 1D chain of atoms, but may be limited in its use to such cases (see [162, Chapter 9] for the use of a static SLM generated array in order to make a 2D array). This limitation to a single dimension is due to the arrangement of the optical elements. A single beam is incident onto the first deflector, and each frequency component ν_i^x of the input beam corresponds to a position in the deflector's axis of deflection, focussing to position set $\{x_i\}$. A second deflector with frequency components ν_j^y, for deflected positions $\{y_j\}$, cannot address a single deflected input beam in this case. The rearrangement of one row will thus cause the rearrangement of all other rows, which require different rearrangments.

This technique is not available to trap generation through SLMs with any efficiency. Moving the traps generated by an SLM requires a different hologram for each movement stage. The refresh rate of an SLM (typically 60 Hz) is too slow for efficient rearrangement, and would require substantial computing resources to precalculate the required holograms before playback for large N. As such, using the SLM for generating mobile traps is not feasible for large arrays. Nevertheless, this has been demonstrated for up to 9 atoms at once [171].

A technique using static tweezers combined with mobile AOD arrays for 'qubit shuttling' has also been demonstrated [67]. Using this technique, many qubits can be rearranged with high efficiency and speed, while maintaining coherence through transport, and forms a promising method of controlling qubit interactions.

4.1.3 Enhanced Loading

As established in chapter 2.5.2, the atom arrays are initially loaded with a single atom through a light-assisted collisions process. While this would appear to initially create an occupancy of $\approx 50\%$, the energy released by the collisions can be low enough such that only one atom is ejected. The remaining atom is trapped, undergoes no further collisions and is subject to the same cooling as other atoms in the system. For this reason, the SQuAre experiment regularly demonstrated a maximal loading rate of $\approx 57\%$.

This same effect can be exploited to further improve load rates in neutral atom array experiments. As demonstrated in a number of Rubidium experiments [160, 172–174], careful preparation of the frequency and duration of the incident light during the light assisted collisions stage can increase the single-atom load rate to a yield of $> 80\%$. This process is referred to as 'enhanced loading'.

This increased yield can be useful in improving the yield of experiments and efficiency of sorting systems. It cannot by itself create defect free arrays with any acceptable yield and scalability, however. Even if the probability of a single atom-load in a trap P_{load} could be made substantially higher, the probability of a defect free array (P_{DFA}) is given by

$$P_{\mathrm{DFA}} = P_{\mathrm{load}}^{N}, \tag{4.3}$$

where N is the number of atoms. As N becomes large, even an exceptionally high load rate would quickly vanish. Thus, while enhanced loading is a useful tool in scaling atom array preparation, an experiment at scale will necessarily require rearrangement of atoms.

4.2 Sorting Algorithms

Given an initial random distribution of atoms in an array of trap sites, the first task in the experiment is to plan the arrangement of moves. As discussed in section 4.1, any rearrangement must be done with a minimal number of moves, in a minimal time, with maximal success per move. In particular, Eq. 4.1 shows how the number of moves m must contend with the sort duration t_{sort}.

In an ideal case for N targets, of which $\approx N/2$ are initially empty, then

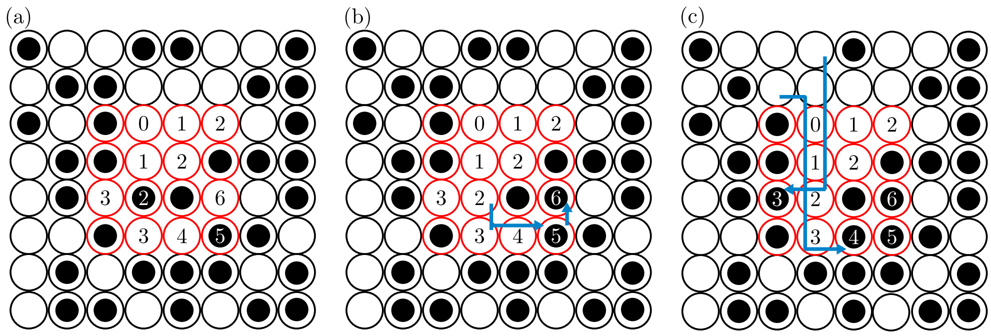

Figure 4.2: Example sorting process for sorting a 4 × 4 target (in red circles) within an 8 × 8 grid (in black). Filled circles represent atom positons, randomly loaded. Figure (a) shows labelling of sites by depth relative to the minimum spanning tree of cluster sites. Unlabelled atoms within the target do not require sorting. Figure (b) shows the first moves performed, where atoms within the graph are moved to the deepest points to remove collisions, from depth 5 to depth 6 and then from depth 2 to depth 5. Finally, figure (c) shows atoms from outside the target moved into target positions along unobstructed paths.

$\langle m \rangle = N/2$, where $\langle m \rangle$ is the average number of moves. This is a trivial lower bound, but is very difficult to optimise for. Firstly, if two atoms collide in transit (i.e. if an occupied mobile tweezer moves over a filled site) then this collision will cause both atoms to be lost. Secondly, exceptionally long moves may take a longer time with a lower success rate than a number of shorter moves. Finally, calculations for these moves must be done rapidly. As shown in Eq. 4.2, any improvements in efficiency must be balanced with a possible increase in calculation duration. As such, optimisation for calculation time and experiment time generate non-linear scaling relations of order $\mathcal{O}\left(N^{\alpha}\right)$ [161, 175] in calculation time and $\langle m \rangle > 0.5N$.

4.2.1 Sidewinder Algorithm

The algorithm used, named ‘Sidewinder’, was developed for the SQuAre experiment by D. Walker. This is developed from two established algorithms. The first is the Heuristic Cluster Algorithm (HCA) [168], developed specifically for atom array sorting with minimal moves. This algorithm connects clusters of unfilled target sites with the shortest path, such that efficient pathfinding can then take place. This algorithm is adapted using the A* pathfinding algorithm [176], where atoms from filled target sites are moved into empty target sites through the shortest path, in order to connect clusters with paths of no obstructions. Once the empty target sites are connected into a single cluster, pathfinding can take place to move atoms from the reservoir (that is, non-target sites that are loaded) to the target sites without collisions.

To efficiently carry out this process, the Sidewinder algorithm uses a k-d tree method across the entire graph structure. Clusters of empty target sites can be connected such that a path from the atom reservoir can be made to any target without collisions by finding the minimum spanning tree of the graph of average coordinates of the clusters. Sorting is then prioritised for filling the deepest target sites on a spanning tree, where the depth heuristic is the distance from the target site to an available reservoir atom. [1] A demonstration sorting process can be seen in Fig. 4.2, using a modified version of the algorithm where only a single entry point to the target tree is used to calculate depth. This modification is used for clarity; allowing for multiple entry points both improves efficiency and permits the sorting of targets which have disconnected regions.

For performance improvements, this algorithm was written and compiled in the Rust programming language [177], rather than the Python3 software that other experimental software was written in. The Python3 control software instead executes the compiled Rust algorithm. The algorithm is given the positions of all sites in the array, and which sites are to be filled or removed. Upon receiving a list of trap sites that are initially filled, it returns a list of moves to follow. This move list is sent to the microcontroller to execute the rearrangement process (see Sec. 4.3).

---

[1] This depth heuristic requires positional data of each atomic site, and thus is not included in the k-d tree.

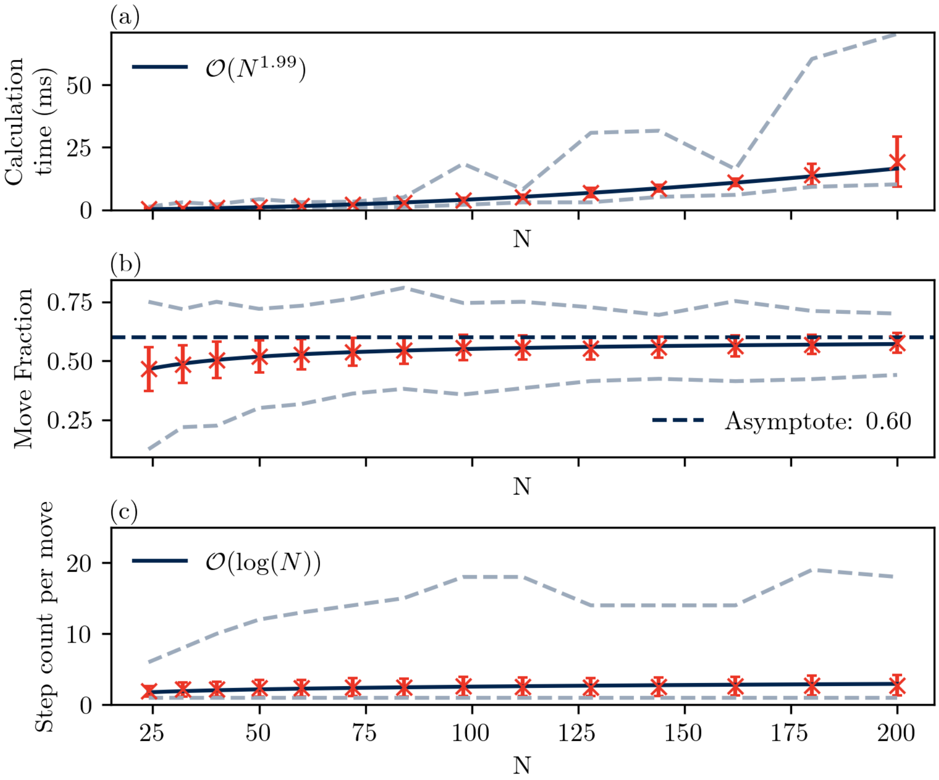

Figure 4.3: Benchmarking results from the Sidewinder algorithm, scaling with number of targets N. Red crosses indicate average simulated performance, while grey dashed lines indicate extremal performance. (a) shows calulcation time scaling, (b) shows the number of moves per N (move fraction) required, while (c) shows how long the average move must be, in steps along the array sites.

4.2.2 Algorithm Performance

Benchmarking was carried out by simulating move lists for a rectangular target of size N, inside a square array of $> 2N$ traps, randomly loaded with $F = 0.57$ to match experiment conditions. As shown in Fig. 4.3, the algorithm is benchmarked on hardware to demonstrate an asymptotic time complexity $O\left(N^2\right)$, with a mean number of moves $\mathcal{O}\left(0.6N\right)$. The average length of each move is of order $\mathcal{O}\left(\log\left(N\right)\right)$, although extremal performance may be substantially greater. Each measurement point is repeated over a random sample of 1000 trials, where red crosses are the measured data points, grey dashed lines are extremal performance, and solid black lines are fits to the data.

4.2.3 Alternative Algorithms

As discussed in [175], the best choice of algorithm depends on the graph that is to be sorted. The SQuAre experiment makes use of a single tweezer, typically across an ordered array, such as a Bravais lattice. Moving to non-grid structures or using multiple tweezers in parallel will necessitate alternative algorithms, such as the Linear Sum Assignment Problem (LSAP) solver algorithm as described in [175] above.

Examples of parallel sorting with multiple tweezers are discussed in [161, 177]. This parallel behaviour is limited to only one row or column at a time, but can still lead to substantial improvements in the total time required. Extensions would require upgrades to the RF chain used such that the AOD can be driven with multiple input tones at once, but the optics would be retained. The process of both installing these upgrades on the experiment and developing a new sorting algorithm to use them is in progress as of writing.

4.3 Experimental Hardware

4.3.1 Optical Components

The optical system used for atom sorting is shown in Fig. 4.4. The static beam is generated using light from an M Squared SolsTiS Ti:Sapph laser, itself pumped by an M Squared Equinox diode laser. This beam is used as one light source for Rydberg interactions, and thus is at wavelength between $1039 - 1043$ nm, depending on the Rydberg state it is coupled to. Some of this power is split into a path dedicated for local Rydberg operations, and so includes an AOM for switching and frequency control.

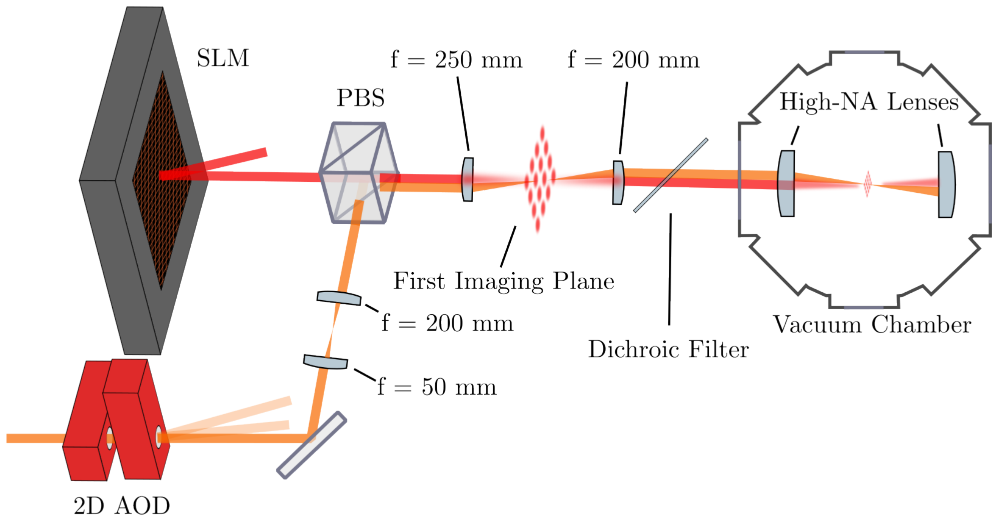

Figure 4.4: Schematic of the optical components used for steering the mobile tweezer (at a wavelength of 1039 nm, in orange), and combining this beam with the red detuned ODT array (at a wavelength of 1064 nm, in red). The high-NA lenses fixed in-vacuum are Geltec 355561 lenses, at a working distance of 7 mm. This diagram is not to scale.

The beam's position in the chamber is controlled through a pair of AODs[2]. These each deflect at central frequency of 77.5 MHz and are controlled through an independent signal. The beam of order $+1, +1$ is aligned to maximise power homogeneity across the array, at an average efficiency of $\approx 85\%$ (see Fig. 4.5).

After collimation is fixed through a relay telescope, the beam is combined with the 1064 nm static array on a non polarising beam splitter (NPBS). The beam is reimaged at a waist size of 1.55(1) μm at the focal plane of the atoms. Despite the use of achromatic lenses and close wavelengths, the final lens used in the optical path is an aspherical lens, and is thus is not chromatically compensated. The two beam wavelengths that would otherwise be identically collimated along this path do not neccessarily focus to the same point. Thus, a reference camera outside the chamber cannot capture both beams in the focal plane. The relay telescope can be adjusted to change the collimation of the 1039 beam, such that the focus plane matches the 1064

---

[2] AA Optoelectronic DTSXY-400-1040.1064 AOD

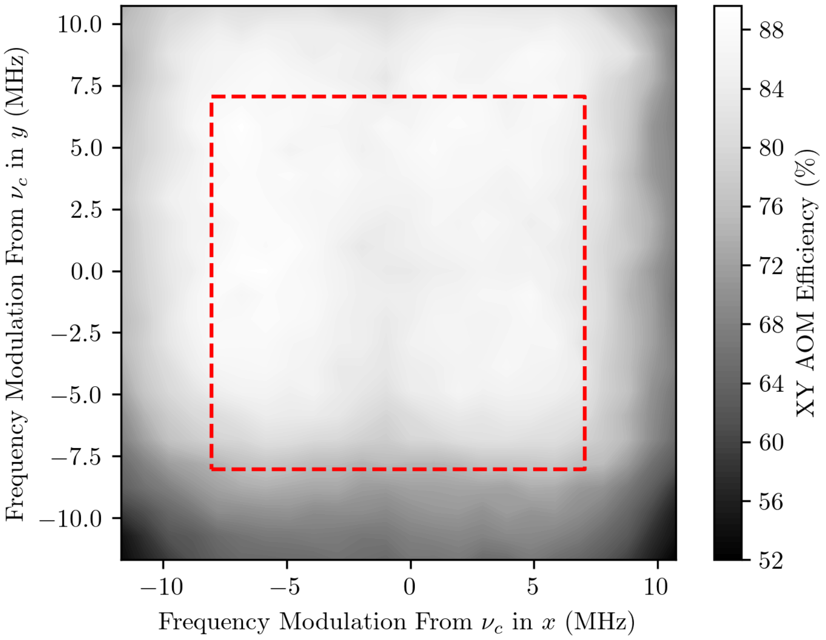

Figure 4.5: Efficiency of the 2D AOD over the range of output frequencies. The dashed red line marks the span of frequencies that address a 112 μm × 112 μm square, which is the footprint of a 15 × 15 atom array at 8 μm spacing. In this region, the average power is 85.4%, with a minimum efficiency of 79.0% and a maximum of 89.4%. The axes refer to the input signal frequency, relative to the central drive frequency $\nu_c = 77.5$ MHz.

nm array inside the chamber.

The resulting ODT has a measured lifetime of $\tau_{ODT} = 1.6(3)$ s, with a transfer efficiency of 98(1)%, measured by extrapolating lifetime to $t = 0$.

4.3.2 Radio-Frequency Components

Each of the two AODs used for controlling the trap position requires an independent frequency source. Channel control is managed through a PJRC Teensy 4.1 microcontroller, which outputs Serial Peripheral Interface (SPI) commands to an AD5686 Eval DAC board, which outputs a DC voltage. Three outputs are used, two for frequency control and one for amplitude control. Only one amplitude source is required, as attenuating a single chan-

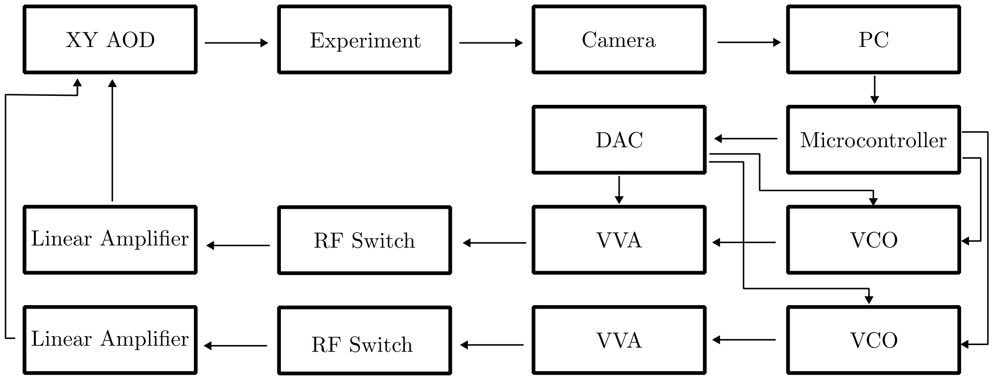

Figure 4.6: Schematic of the RF chain required for the sorting experimental loop. PC refers to the computer where the Sidewinder algorithm instructions are generated. Instructions are executed by the Teensy microcontroller, via the AD5686 Digital Analogue Converter (DAC). 'VCO' and 'VVA' refer to Voltage Controlled Oscillator and Voltage Variable Attenuator respectively. RF Switches can be activated through external TTLs to switch to an auxiliary RF source.

nel input can still control the full dynamic range of output intensities.

Microwave generation is given by a voltage controlled oscillator[3] (VCO) for each channel. This signal is then attenuated by a voltage variable attenuator[4] (VVA) to control beam amplitude. A TTL controlled MW switch[5] allows for fast switching to an auxiliary MW source for higher precision in potential Rydberg control experiments. Finally, this MW signal is amplified using a Moglabs AADPCB linear amplifier[6] and inserted into each AOM. See Fig. 4.6 for a flowchart of the equipment used in this chain.

Given that the DAC has an input resolution of 2^{16} bits, this gives fine control over the potential sent to the MW components. A difference of 283 bits in the DAC input command corresponds to a difference of 139.5(1) kHz in modulation frequency, and $\approx 1.00\ \mu$m at the atoms[7].

---

[3] Minicircuits ZX95-100+ VCO

[4] Minicircuits ZX73-2500-S+

[5] Minicircuits ZYSW-2-50DR+

[6] Moglabs AADPCB linear amplifier

[7] The diffraction angle, and thus the position at the atoms, depends on the frequency of the light.

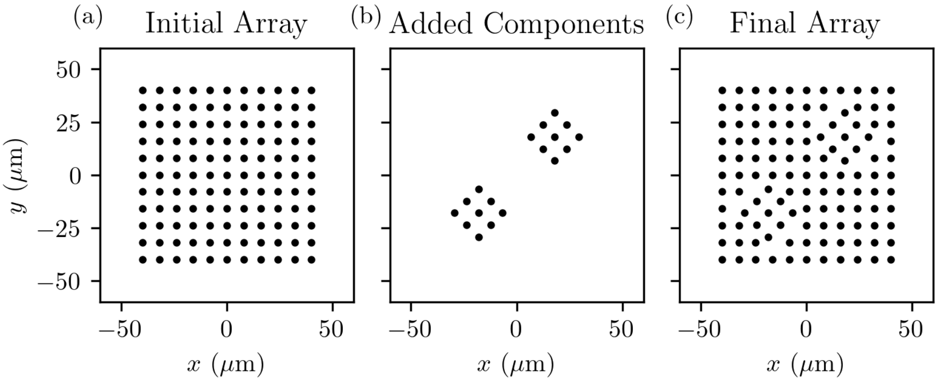

Figure 4.7: Plan of atom placement for arrays. Here, each black point indicates the placement of an optical dipole trap, and thus an atom. Figure (a) is an 11×11 array at 8 μm spacing, which forms the initial reservoir. Figure (b) is two component subarrays, each a 3×3 array at 8 μm spacing, rotated 45° to the initial array. Figure (c) is these arrays added to the underlying array, with atoms closer than 6 μm to the component subarrays removed.

4.4 Array Generation

4.4.1 Arbitrary Arrays

To generate the focussed beams used for atom trapping, an SLM is used, as described in 4.3. This SLM can generate an arbitrary pattern of red detuned ODTs, within the field of view of the high-NA lens. For a typical experiment a regular array is both convenient and practical, such that the conditions for each atom are identical and can be labelled predictably. The experiment is proven to produce high-fidelity single gates on arrays of size 15×15 at a spacing of 8 μm [64].

While the Sidewinder algorithm is capable of sorting an arbitrary array to arbitrary targets, sorting performance is imporved by the use of regular arrays. The k-d tree built for the array (see Sec. 4.2) requires information about which atoms are connected, e.g. atoms are connected to their nearest neighbours only. As such, a typical array requiring sorting that is not on a Bravais lattice is usually overlaid on such a lattice. Traps that then are too close to resolve are removed from the theoretical graph and the array is generated. The original Bravais lattice acts as a reservoir of atoms for sorting into the added components. An example array is shown in Fig. 4.7, where rotated squares of atoms are inserted into a regular grid. This experimental procedure has shown to be effective at the deterministic preparation of atoms in arbitrary geometries.

4.5 Parameter Optimisation

The parameters required for an ideal rearrangement procedure have so far been found empirically, and at that primarily for Rubidium experiments. The experimental results found in Table 4.1 are for species other than Cs-133, and so are not necessarily ideal for a Cs experiment. The following parameters are known to affect the rearrangement efficiency.

Optimisation must be done on multiple parameters which are discussed here. Firstly, the precision and focus of the mobile tweezer requires calibration, such that the tweezer starts and ends at the right place. Next, the power of this mobile tweezer must be carefully controlled. A shallow tweezer will fail to pick up atoms out of the static tweezers, but if too deep then will heat atoms on ramping up the input power or in movement. This ramp rate thus depends on the input power also. Movement of the atoms requires further care. Longer moves with non-ideal movement will induce further parametric heating of atoms. Collisions between atoms during movement will result in atoms being lost as well, and so must be avoided.

Of these parameters, only some can be removed as a concern. By pathing the mobile tweezer such that it only moves through empty sites, collisions can be eliminated. Trap lifetime is limited by vacuum pressure, so improvements must be made in total sorting time (see Chapter 4.2). Furthermore, while these parameters may be optimised for a small subsection of traps, extending this performance to a substantially larger array can result in a reduction in

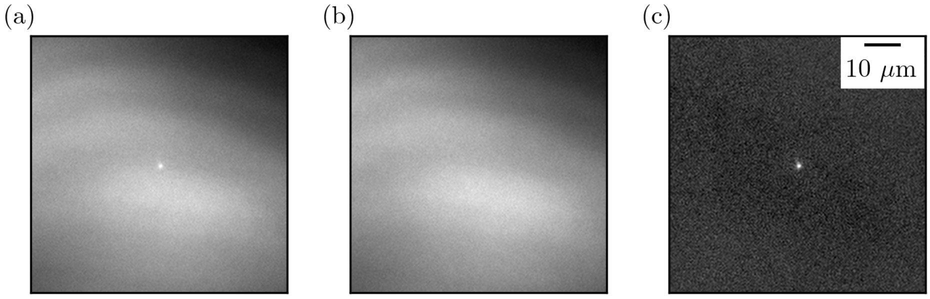

Figure 4.8: Images used for tweezer positional calibration. Figure (a) is an average of 10 images of the mobile tweezer within the MOT. The tweezer light is chopped, much like the ODT light (see section 3.3.3). Figure (b) is the average of images of the MOT, each taken immediately after each of the images used in (a). Figure (c) is the difference between the two, showing a clear image of atom fluorescence within the beam. This background subtracted image is used for calibration of the mobile tweezer position.

efficiency. As such, care must be taken to clarify which results represent ideal behaviour for a limited system, and which results represent the capabilities of a larger system as required for considerations of scalability.

4.5.1 Calibration

Positional Calibration

Calibration is required to match the position of the mobile tweezer to the position of the static array. Given that the light used for each beam is of a different wavelength (see Fig. 4.4), the two beams will not necessarily focus on the same plane, and may be offset. Furthermore, the mobile tweezer's wavelength is variable, and so corrections for this effect are also variable. As such, reimaging both the array and the mobile tweezer onto a single camera or CCD (as seen in [164]) for calibration is not viable.

An alternative calibration scheme was developed by imaging the fluorescence of atoms trapped in the mobile tweezer using the experiment camera. This is compared to images of atoms within the optical dipole trap array. Given that both the mobile tweezer and the static array ideally trap atoms in the same plane, and that both atoms fluoresce at the same wavelength of 852 nm, this comparison can thus be used for in-situ calibration. Example images used for this are seen in Fig. 4.8.

Power Calibration

Further calibration is required to control the power of the mobile tweezer. The proportional power of the mobile tweezer relative to the control signal is trivial to measure and interpolate, as it forms a monotonically increasing function. Keeping this power even over the range of motion of the tweezer is a greater concern. Maximising diffraction efficiency at the central frequency of the two AOMs was found to create high efficiency in the centre of the input range, but poor efficiency at the edges. Instead, the input control signal was randomised in both axes around the full range, and then aligned to maximise the time averaged power. This was found to consistently produce a mean diffraction efficiency of $> 85\%$ over the range covered by the field of view of the system (see Fig. 4.5).

Mapping the 2^{16} bits of DAC output to the measured efficiency of the AOM can thus be used to calibrate the power of the tweezer. Smooth ramps of output power are limited by the cycle rate of the microcontroller, which updates every 5 μs.

4.5.2 Maximising Transport Efficiency

Transport of an atom from one site to the next depends on a number of coupled parameters. The two most crucial are the transfer between mobile and static tweezers, and the movement of the mobile tweezer from the static reservoir tweezer to the static target tweezer. As these parameters are linked, it is required to first discuss movement parameters.

These parameters were maximised on a simple square lattice of fixed spacing,

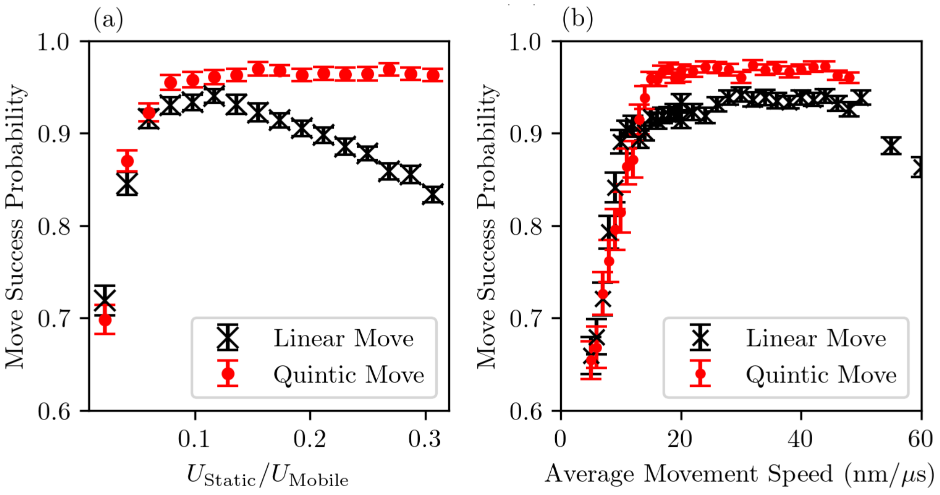

Figure 4.9: This shows a comparative efficiency of a tweezer in a single move over 8 μm. Move success rate is also corrected for background and imaging loss, with the same t_{sort} used for all scans. Figure (a) demonstrates the robustness of the quintic (minimum jerk) movement profile to changes in power. Here, the mobile tweezer power is kept constant at a depth of 9.1 mK and movement success rate is measured as the static trap power is varied during the sorting stage. Figure (b) demonstrates how, at optimal parameters, the quintic profile moves atoms at a greater success rate for the same movement duration.

and moving each atom across by a single site. This eliminates experimental uncertainty regarding pathfinding or movement length.

Movement Profiles

The naïve method to move atoms is a simple constant velocity movement profile. A velocity of 50 nm/μs was found to be the maximum possible velocity without loss. This velocity is comparable to the 100 nm/μs found in rubidium experiments [58, 175].

Constant velocity profiles are potentially limited in scope, however, as this movement profile is known to excite atoms to higher motional states in transit [178]. The rate of change of acceleration, or jerk, is what causes motional heating of the atom. First described in [179], a quintic movement profile can be defined that trajectory minimises the total jerk of the atom, thereby

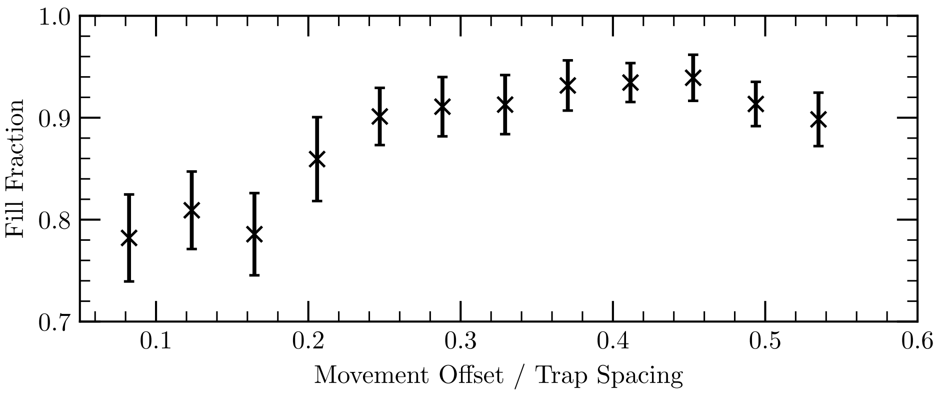

Figure 4.10: Scan demonstrating imporvement in graph sorting with a movement offset. This is an initial move in both x and y axes before further sorting moves to prevent collisions and perturbative potentials.

minimising the heating of an atom within the trap. This movement profile is defined by

$$x(t) = x(0) + (x(T) - x(0)) \left(10 \left(\frac{t}{T} \right)^3 - 15 \left(\frac{t}{T} \right)^4 + 6 \left(\frac{t}{T} \right)^5 \right). \tag{4.4}$$

where x is the position of the mobile tweezer and T is the total duration of movement. This can be compared to a linear profile of the same duration in Fig. 4.11.

The quintic trajectory is utilised in the bottom-up formation of molecules [180, 181] where minimal heating is a necessity. The quintic profile maintains a consistently higher movement performance than linear moves. Furthermore, this profile maintains an optimal efficiency over a larger range of trap powers. This is due to the low acceleration not heating the atom during movement into or out of the static traps.

Movement Offset

Moving the mobile tweezer over static tweezer sites can cause loss even when the static tweezer sites are not occupied by atoms. This is due to sudden

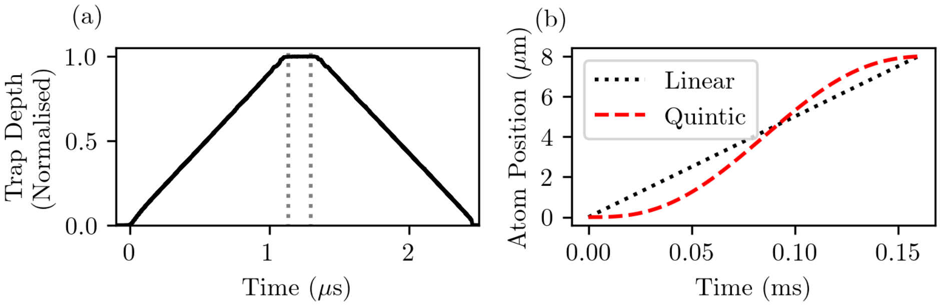

Figure 4.11: Power and position of the mobile tweezer during a single 8 μm move. (a) is the trap depth for mobile tweezer trap depth over time, measured on a Picoscope 4444 oscilloscope. The dashed lines indicate where the tweezer is at constant power through the movement stage of the rearrangement process. (b) is the position of the atom in the mobile tweezer during this movement stage. Both linear (dotted black line) and quintic (dashed red line) profiles are plotted for comparison.

changes in trap depth from the static traps disturbing the atoms from smooth transport, causing heating. To avoid this, and any potential collisions, movements started by offsetting the atoms from the initial point by a small fixed value. This offset is undone at the end of the move to place atoms back in the array.

Optimal transport efficiency was found for an offset of 0.4 times the trap spacing, as seen in Fig. 4.10, although this is relatively insensitive and can change depending on array geometry.

Trap Powers

As discussed above, the relative power of the mobile trap compared to the static trap can substantially change transport efficiency. Furthermore, the total time taken to ramp tweezer power up and down is often substantially longer than the duration of transit. The deeper a mobile tweezer is, the less likely it is to lose an atom in transit, but will take longer to ramp to maximum power.

The highest efficiency seen was at a static trap depth of 1.5 mK[8], with the mobile tweezer at trap depth of 9.1 mK[9]. This requires the static tweezers to be deeper than at the imaging stage, at a trap depth of 1.2 mK.

As seen in Fig. 4.11, the power ramps account for the substantial majority of time spent in each sorting move. Efficient power ramps over 60 μs have been seen in Rb-85 experiments [161], indicating that faster ramps are possible with hardware upgrades.

For linear ramps in power, the mobile tweezer ramps from minimum to maximum output at a constant rate. Optimal performance was found using small increases in power for each update cycle of the microcontroller, such that it is the smoothest possible ramp. This corresponds to a 1.07 ms ramp, or $\approx 0.093\%/\mu$s.

Experiments testing this is seen in 4.12, where step sizes are adjusted with known delay between each step. While success rate and ramp speed are not directly correlated, small (and thus smooth) increases in power are shown to be more effective at transfer. The legend shows the delay between each step, with step sizes chosen to correspond to the same time range, and thus average ramp speed.

Ideal Performance

For a single move over 8 μm, along one Cartesian axis, the highest success rate for movements observed was 98.3(5)%, as seen in 4.12. This is lower than the 99.3% movement success rate quoted in [58], which had lifetime losses as the greatest cause of loss. For an initial fill fraction $F = 0.57$ after initial loading, this comes to an expected upper limit of $F \leq 0.990(3)$ just considering losses from moves. Between images without sorting stages, atoms were found to remain in the traps at a rate 99.03(4)%, indicating loss due to imaging of

---

[8] Averaged over all traps in the array

[9] Calculated from average AOD efficiency over the array

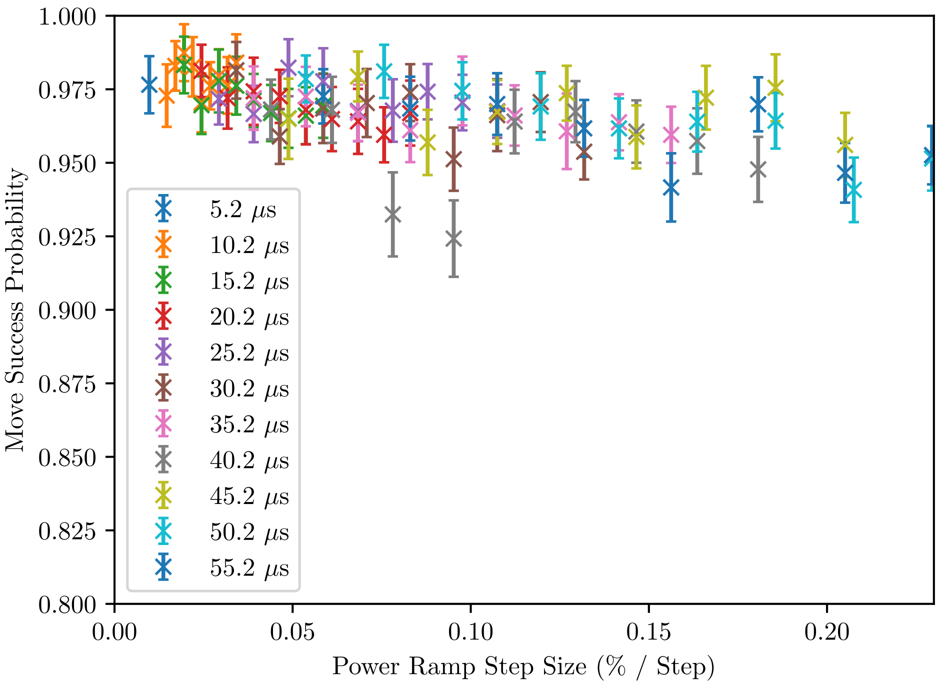

Figure 4.12: Graph of movement success rate as power ramp profiles are changed. Each colour refers to a different delay after each step in increasing the output power. The microcontroller has a measured cycle period of 5.2 μs, and thus is the minimum time in each movement step.

0.97(4)%. Atom loss scales with t_{sort}, which itself is dependent on algorithm calculation time and the number of moves required, which scales linearly with N, such that $t_{\text{sort}} \approx (2.0 \cdot N)$ ms. Atom losses over this timescale were found to be consistent with a trap lifetime of $\tau = 12(1)$ s matching the vacuum-limited lifetime of the experiment during these measurements [87]. Since these measurements, trap lifetime has been improved to 17(1) s by activation of the titanium sublimation pump. This corresponds to an optimal fill fraction of

$$F_{\text{ideal}} \approx 0.981(5) \cdot \exp\left(\frac{-2N}{\tau}\right). \tag{4.5}$$

This is competetive with other experiments seen in Table 4.1, although experimental improvements can be made in reducing loss in imaging and hardware upgrades for faster and smoother power ramps.

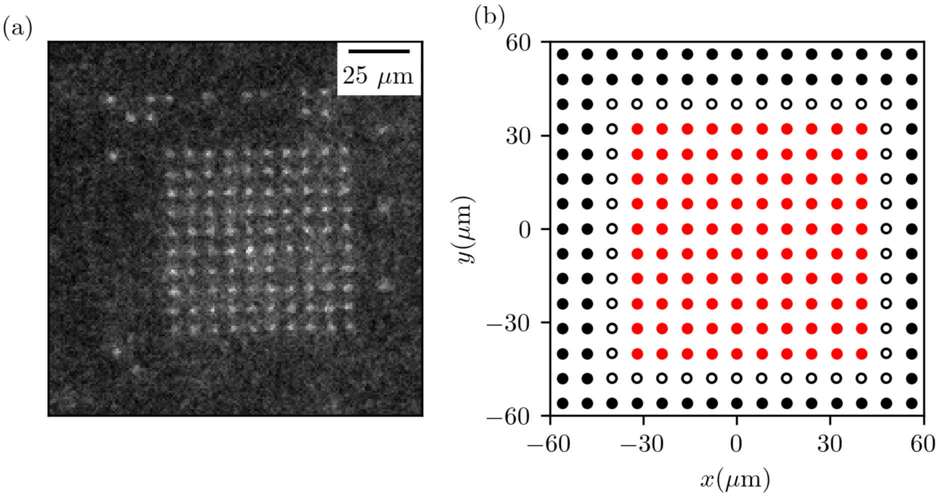

Figure 4.13: Images of sorted arrays. Figures (a) shows an experimental image of a 10 × 10 atom square sorted without defects, within a 15 × 15 trap array spaced by 8 mum. Figure (b) shows the array configurations for figure (a), where red circles are targets, empty black circles are atoms to be removed and filled black circles are the reservoir to be used as appropriate.

4.6 Graph Performance

With power, transit and algorithmic performance characterised, we can discuss the performance of the rearrangement sequence in experimental arrays of different geometries.

4.6.1 Square Targets

To demonstrate the performance of the system in optimal conditions, initial characterisation was made on a smaller square targets inside a much larger array. This is chosen such that there are only predicatable moves with an abundance of reservoir atoms, where $3N \geq T$, where T is the total number of traps. Exemplar data is then taken to fill a target of a 10 × 10 atom square inside a 15 × 15 atom square array, the design specification for this experiment. Fig. 4.13(a) shows a 10 × 10 array filled without defects.

As shown in Fig. 4.14, the fill fraction for small ($N \leq 9$) targets is

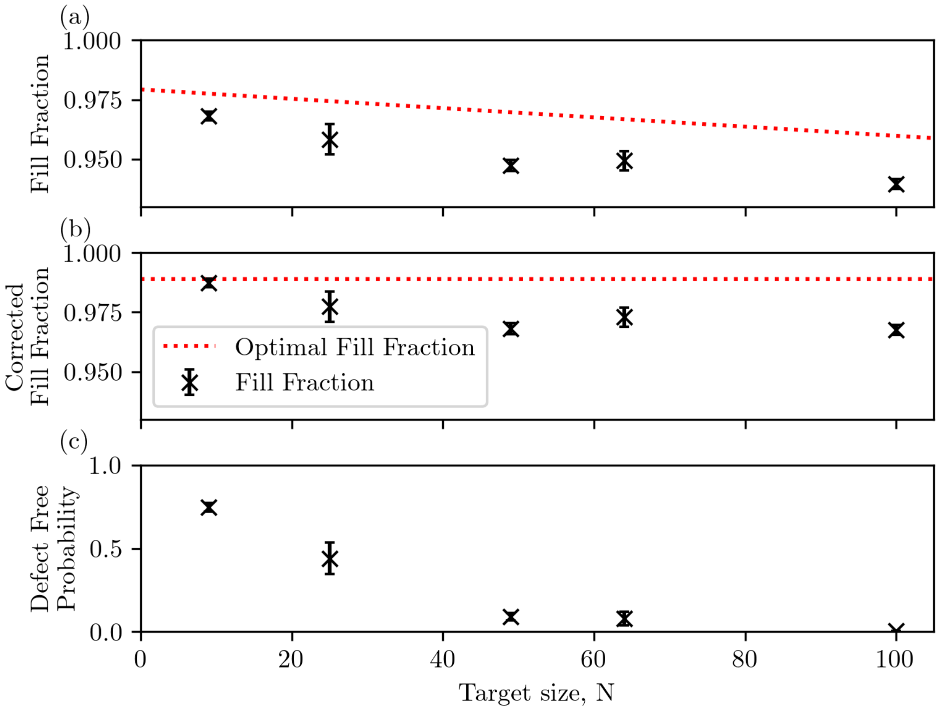

Figure 4.14: Plot of the fill fraction of target size N, as N increases. Measurements are taken from exemplar data, that is with minimum sorting time for each graph. (a) shows how fill fraction reduces for target size, due to the longer durations required for sorting and the increased complexity of moves when filling deeper arrays. (b) is the same data, corrected for imaging loss and background loss over t_{sort}. (c) shows how the probability of making a defect-free array after a single sort drops with N. This is using the same experimental data as in (a). At $N = 100$, a defect free array was prepared once over 300 experimental iterations.

$F = 96.8(2)\%$, or a corrected rate of $F_{\mathrm{Corrected}} = 98.7(3)\%$ when correcting for loss, and is at the optimal limit determined by loss in transit. This drops to $F = 0.940(2)$, $F_{\mathrm{Corrected}} = 96.8(2)\%$ for the largest target, $N = 100$. This is indicative of how with larger targets, sorting efficiency reduces further. As shown in section 4.5, mobile tweezer power changes over the array. Without feed-forward corrections, this results in unoptimal power as the tweezer moves over the array. Using quintic movement profiles reduces sensitivity to power changes when moving into and out of traps, but rapid changes in power during a movement can rapidly heat atoms out of the trap. Other potential loss mechanisms include electrical noise during the tweezer ramp stages.

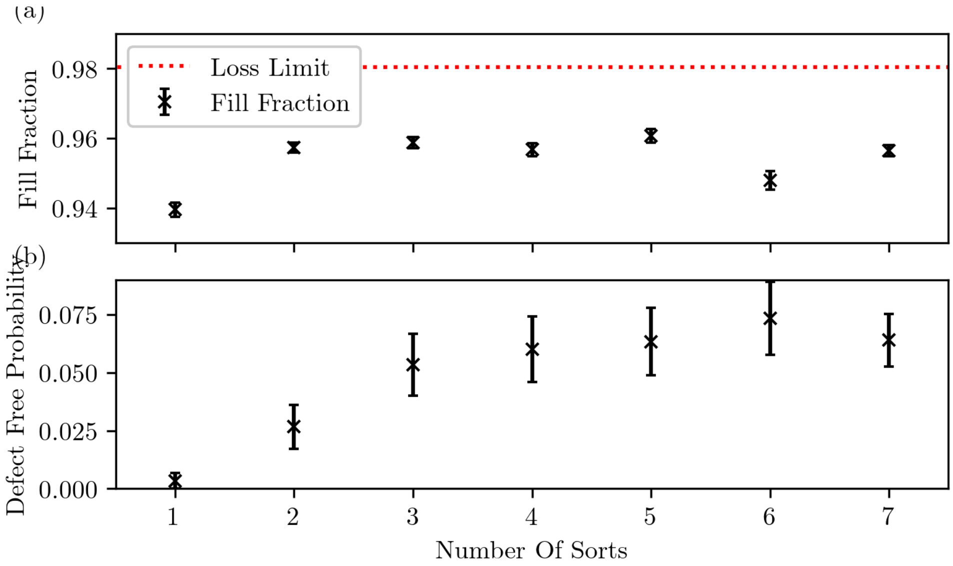

Figure 4.15: Plot of the fill fraction of the same graph as the number of sorting rounds increases. Measurements are taken from exemplar data, that is with minimum sorting time for each graph. (a) shows how fill fraction reduces for target size, due to the longer durations required for sorting and the increased complexity of moves when filling deeper arrays. (b) shows how the probability of making a defect-free array after a single sort drops with N. This is using the same experimental data as in (a).

The probability of longer than average sorting sequences also increases as some loading configurations may require substantially more moves than average. Off-axis performance at the edge of arrays may also be affected, as N increases.

As N is increased, the probability of creating a defect free array drops exponentially. By increasing the number of sorting rounds, this defect free rate can be increased. For $N = 100$, this saturates at a fill fraction of 96%, indicating the rate of loss at each image is at equilibrium with the rate of accurately correcting for this loss. Further experimental upgrades will be able to correctly identify when an array is correctly filled, and stop sorting from this point.

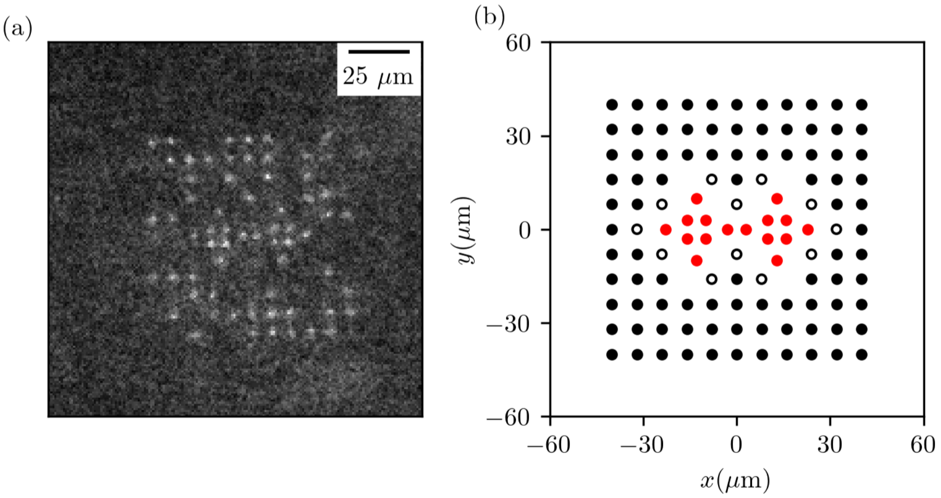

Figure 4.16: Images of sorted arrays. Figures (a) shows an experimental image of a non-standard array, within an 11×11 trap array spaced by 8 mum. Figure (b) shows the array configurations for figure (a), where red circles are targets, empty black circles are atoms to be removed and filled black circles are the reservoir to be used as appropriate.

Off-Grid Targets

Graphs where points do not lay on a regular grid prove harder to fill. Gadgets (discussed further in Sec. 7.4) are subgraphs with standard arrangement that do not confirm to a unit grid. These gadgets have limited paths to fill the target atoms, with and must make diagonal moves off the axis of the Bravais lattice. However, due to the small numbers of targets relative to the large square arrays for experiments so far performed, the probability of a random atom loss is lower, as is the time required for the sorting procedure. An example gadget can be prepared with an average fill fraction of 0.91(1) after a single sorting round, where Fig. 4.16 shows this gadget array filled without defects. After three sorting stages, this grows to a fill fraction of 0.992(3), for a defect free rate of 87(1)%. This demonstrates reliable preparation of atoms in configurations other than square lattices.

4.7 Summary

4.7.1 Atomic Rearrangemnet Limitations and Outlook

As shown above, the experiment can deterministically prepare an arbitrary arrangement of atoms, including up to 100-atom arrays or unconventional off-grid patterns. This forms the initial stage of preparing a neutral atom quantum computer or quantum simulator. Atom sorting rates within this experiment can still be improved for incremental gains in sorting efficiency. Each potential improvement is detailed below, with outlook for the SQuAre experiment included.

Reducing Imaging Loss

Fixed per-image losses can be reduced by optimisation of the imaging sequence to prevent atoms from being lost. Imaging losses have been reduced to $< 10^{-3}$ while retaining imaging fidelity > 0.9999 as seen in other Cs experiments [55, 182]. This would drastically reduce the largest source of sorting losses, that is loss during imaging. Further experiments to reduce sorting infidelity must therefore be considered if larger arrays are required at high yield.

Optimising Sorting Sequence Software

Per-sort loses can be reduced by reducing the duration of sorting and so reducing t_{sort}. Faster, smoother power ramps will do both in this respect, but require upgrades to the experimental hardware. Further software upgrades, such as ending the sorting sequence at the point of finishing the sort instruction and not requesting further sorting stages if the target is defect free after imaging will prevent long sorting sequences from being cut off early and reduce t_{sort}, while also increasing the defect free rate.

Given that fill fraction saturates with the number of sorting rounds, the probability of a successful rearrangement is equal to the probability of no atoms lost in the sequence. With sufficient experiment time and a large enough reservoir, large targets can be filled even with low sorting success rates. For example, in [167], a substantial reservoir of atoms is generated such that targets of 441 atoms can be filled, and a square target of 225 is filled completely with a cumulative probability of $\approx 35\%$ after 50 such rounds of rearrangement. This 225 atom target uses a large number of small squares, such that long, circuitous moves into the centre of the array are not required, cutting down on the sorting time required and power differences as the mobile tweezer moves across the array.

Array Feedback And Mobile Tweezer Feedforward

As shown in [164, Chapter 6], the static ODT array can be optimised using an iterative feedback process, such that the trap depth of all static traps is even. This is demonstrated in the aformentioned thesis to improve the movement success rate. Such a technique is available for this experiment; see [87] for the calibration of ODT power by measuring the AC Stark shift of the D1 transition on individual atoms within the array. The performance of the mobile tweezer may also be improved by utilising a feed-forward process. Given that the diffraction efficiency of the AOM is known across the range used, the RF power sent to the AOM can be adjusted to keep the diffracted power even as it moves across the array. This would require substantial precalculation of the MW signal used to drive the AOM, and would limit the AOM efficiency to the lowest efficiency of its target region, but may improve maximal atomic transport efficiency.

Vacuum Lifetime Improvements

Trap lifetime has been shown to be improved by at least two orders of magnitude using periodic grey-molasses cooling [55], or utilising a cryogenic system [183], although the latter requires an entirely new vacuum system to be constructed. Reducing the total experimental time requires additional experimental hardware, which is being implemented (see below).

Parallel Sorting

A reduction in rearrangement times can also be achieved by sorting atoms in parallel. This requires substantially more laser power dedicated to atom rearrangement, as well as more sophisticated microwave control hardware. Hardware and software for parallel sorting is in the process of implementation on this system. As the majority of the sorting time is spent in ramping power, this must also be optimised. Implementation of the faster, smooth cubic ramps as used in the static ODTs in this experiment requires a different control hardware, due to hardware speed limits in the microcontroller in use here.

In this respect, an arbitrary wave generator (AWG) would be ideal for both power and speed concerns. With total control of the MW signal in both frequency and intensity, multiple mobile tweezers can engage in parallel rearrangement with even power across the entire array. These can be ramped up or down in power smoothly, over much faster timescales. Feedforward procedures can be implemented in this process as well.

4.7.2 Conclusion

In summary, the SQuAre experiment has met its design goals in sorting atoms into 100 atom targets, or into arbitrary arrays. These results are competetive with other experiments in the field of arrays of single species atoms. Experiments requiring larger target arrays or higher yield will require additional hardware and experimental improvements to the imaging sequence, which is to be implemented in the future. With iterative improvements to each part of the sorting process and with additional experimental hardware, even higher yields are possible.

Chapter 5

Global Rydberg Excitation

Interaction between atoms within an experimental array are engineered using the long-range interactions of highly-excited Rydberg atoms.

This Chapter details the methods and results of global excitation of Rydberg atoms. Also discussed are the improvements required for a more precise and effective global excitation. Secs. 2.2.1, 2.2 introduce the theory of Rydberg atoms, while Sec. 3.6.3 introduces the experimental hardware used.

5.1 Coherent Excitation of Rydberg Atoms

The Rydberg atoms used in this experiment are those in the $\left|nS_{1/2}\right\rangle$ state of Cs. As mentioned in Sec. 2.2.1, the linewidth of the Rydberg states scales as n^{-3}, while as shown in Eq. 2.1.3, this further reduces the Rabi frequency of transitions.

As introduced in Sec. 3.6.3, Rydberg atoms are excited from the $\left|6S_{1/2}, F=4\right\rangle$ ground state to the $\left|nS_{1/2}, m_j=+1/2\right\rangle$ state using a two photon transition, via the $\left|7P_{1/2}\right\rangle$ intermediate excited state. The transition used here is detailed in Fig. 3.11. This operates at 459 nm, generated from frequency doubling a 918 nm Ti:Sapph laser, and a second Ti:Sapph laser operating at

1039 nm. Both 918 nm and 1039 nm lasers are are stabilised against a high-finesse ultra-low expansion (ULE) cavity to enable narrowing of the linewidth to a few kHz. This also enables locking of the lasers to frequencies far detuned from the intermediate state to suppress scattering. This intermediate state detuning frequency Δ is typically at 500 MHz. In this section, the initial measurements to document the performance of the system for Rydberg excitation for single Cs atoms are recorded.

5.1.1 Initial Spectroscopy

The first measuremenet required is acccurate determination of the frequency offset of the ULE cavity mode with respect to the $\left|7P_{1/2}\right\rangle$ intermediate state to enable precise definition of the intermediate state detuning Δ[1]. While the WS7 wavemeter used provides relative precision of 10 MHz for input light, the absolute accuracy of measurement is 200 MHz at 918 nm. After frequency-doubling, this absolute accuracy worsens to 400 MHz.

For this measurement, atoms prepared in the $\left|6S_{1/2}, F=4, m_F=0\right\rangle$ state through D1 optical pumping. Atomic resonances to the $\left|7P_{1/2}, F'=3\right\rangle$ and $\left|7P_{1/2}, F'=4\right\rangle$ states are indentified by the frequency at which a short pulse of 459 nm light optically pumps atoms into the $\left|6S_{1/2}, F=3\right\rangle$ state. We can perform spectroscopy to find the exact frequency of the 459 nm beam, and thus the intermediate state detuning. As shown in Fig. 3.11, this intermediate state detuning is the frequency difference of the light relative to the centre of mass of the $7P_{1/2}$ state. The transition targetted for this was the $|F=4\rangle \rightarrow |F'=3,4\rangle$ transitions.

Spectroscopy is performed by scanning the offset frequency of the 918 nm seed light relative to the ULE cavity mode. As the laser is then frequency doubled, the frequency shift at the atoms is twice that of the offset frequenecy. Initially, the closest carrier mode to the relevent transitions is found, as

---

[1] See Fig. 3.11 for a diagram of offset-locking.

measured using the WS7. Optical pumping into the $|F = 3\rangle$ state is measured by state-selective blowaway of atoms in the $|F = 4\rangle$ state, as introduced in Sec. 3.4.

The $|F = 4\rangle \rightarrow |F' = 3\rangle$ transition was found at an offset frequency of 1400.00(2) MHz, while the $|F = 4\rangle \rightarrow |F' = 4\rangle$ transition was found at an offset frequency of $-1405(1)$ MHz, offset to the negative frequency of the carrier of the next highest cavity mode. Given a measured free-spectral range (FSR) of 2.995(1) GHz, the maximum offset frequency possible is $\nu = \text{FSR}/2 = 1498$ MHz. The presence of a very weak degree of unmodulated light from the AOM when measuring the $|F = 4\rangle \rightarrow |F' = 4\rangle$ transition contributed to uncertainty in mesurement.

Combining these measurements of with measurement of the cavity FSR, we estimate a hyperfine splitting of the $7P_{1/2}$ state of 380(3) MHz, in excellent agreement with recent precision measurements of this feature at 377.4 MHz [184].

Following measurement of the two resonant transitions, we proceed to lock at the cavity offset frequency of 1237 MHz, which corresponds to an intermediate state detuning of $\Delta/2\pi = 502(2)$ MHz relative to the centre of mass energy for the $7P_{1/2}$ state. The errors here correspond to measurement error of the absolute frequency and not frequency instability.

5.1.2 Rydberg Spectroscopy

The choice of Rydberg state determines the experimental parameters within the experiment. Higher n increases the blockade radius and state lifetime, but also sensitivity to stray fields. In addition, the dipole matrix element from the intermediate excitate state to the Rydberg state scales as $n^{-3/2}$, resulting in a reduction in Rabi frequency for higher states, as shown in eq. 2.1.3. This Rabi frequency can be increased with the same power by

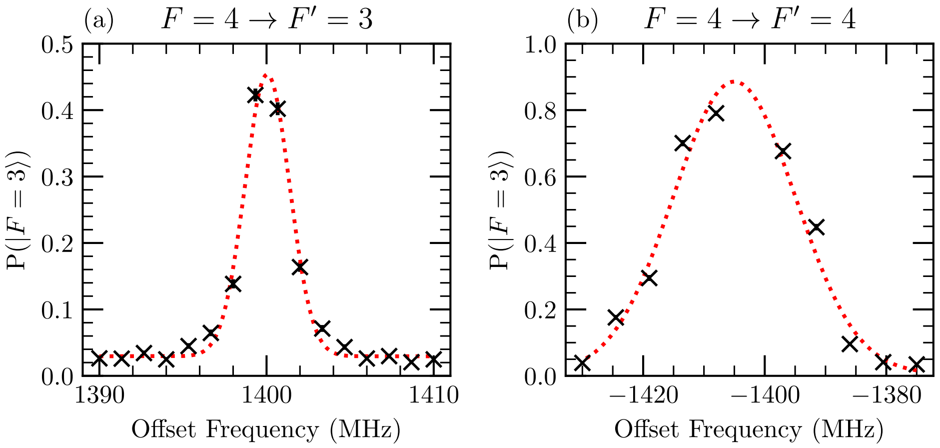

Figure 5.1: Spectroscopy of the $F = 4 \rightarrow F'$ transition. Atoms in the $|F = 4\rangle$ state are removed from trapped by state-selective blowaway, such that the probability of populations in $|F = 3\rangle$ are measured. (a) shows the $F' = 3$ transition, while (b) is the $F' = 4$ transition. One data point in (b) was excluded due to interactions with the 0th order mode of the 459 nm Rydberg AOM. Data points show atomic survival after a blow-away measurement, indicating atoms in the $F = 3$ state. Red dotted lines are fits to a Gaussian profile.

reducing the intermediate state detuning, at the cost of increased AC Stark shifts and intermediate state scattering. This is described through Sec. 2.3.

Polarisation is fixed to σ^+ for 1039 nm light and σ^- for 459 nm light. This enforces selection to only the $\left|nS_{1/2}, m_j = +1/2\right\rangle$ states.

In the experiments below, we perform experiments with two Rydberg states, $\left|50S_{1/2}\right\rangle$ and $\left|80S_{1/2}\right\rangle$. For $n = 50$, the typical Rabi frequency is around $\Omega/2\pi = 5.5$ MHz with $C_6 = -11$ GHz$(\mu\text{m})^6$, meaning $R_b \approx 3.6\ \mu$m. As such, for spacings with $r > 6\ \mu$m, we can neglect dipole-dipole interactions. This an ideal state for electric field spectroscopy and alignment of the Rydberg beams onto the array, as a larger number of closely-spaced atoms can be used through an experiment. Reduced sensitivity to electric fields also allows for greater precision in measurement of the effects of electric field gradients. For experiments involving Rydberg blockade effects, we use $n = 80$, where $\Omega/2\pi \approx 2.5$ MHz and $C_6 = -3599$ GHz$(\mu\text{m})^6$.

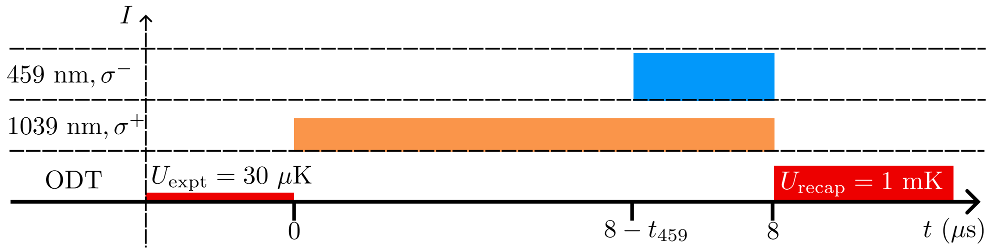

Figure 5.2: Rydberg Excitation Timeline. Bar heights correspond to intensity and are not to scale. Included for red-detuned ODT beams are calculated trap depths U_0 at the given intensity.

5.1.3 Rydberg Excitation Sequence

As Rydberg atoms cannot be trapped by the red detuned ODTs used in this experiment, the experimental sequence must adjust for this. This sequence is shown diagramatically in Fig. 5.2. During the Rydberg excitation procedure, the ODTs are turned off such that the Cs atoms are in free-fall. At the same time, the 1039 nm light is turned on. As this light couples to the $|7P_{1/2}\rangle \rightarrow |nS_{1/2}\rangle$ transition, and no atoms are in either of these states during the state preparation, this does not affect atomic populations until the 459 nm light is turned on. This duration of switching off of the 1064 nm light and constant intensity of 1039 nm light remains fixed through the experiment, at a release time of 8 μs. This is the maximum time that atoms can be released from the ODTs and be recaptured without measurable loss (see fig. 3.7 for release-recapture measurements).

The Rydberg transition only occurs when the 459 nm light is incident on the atoms. This turns on for a variable duration. The ODTs are then turned on at a trap depth of 1 mK, the trap depth used in imaging, to eject any Rydberg atoms, as discussed in sec. 3.4.

5.1.4 n=50

Spectroscopy exciting to the $50S_{1/2}$ state can be seen in Fig. 5.3. Here, the second beam is used, at a wavelength of 1042.8469(1) nm and is locked to

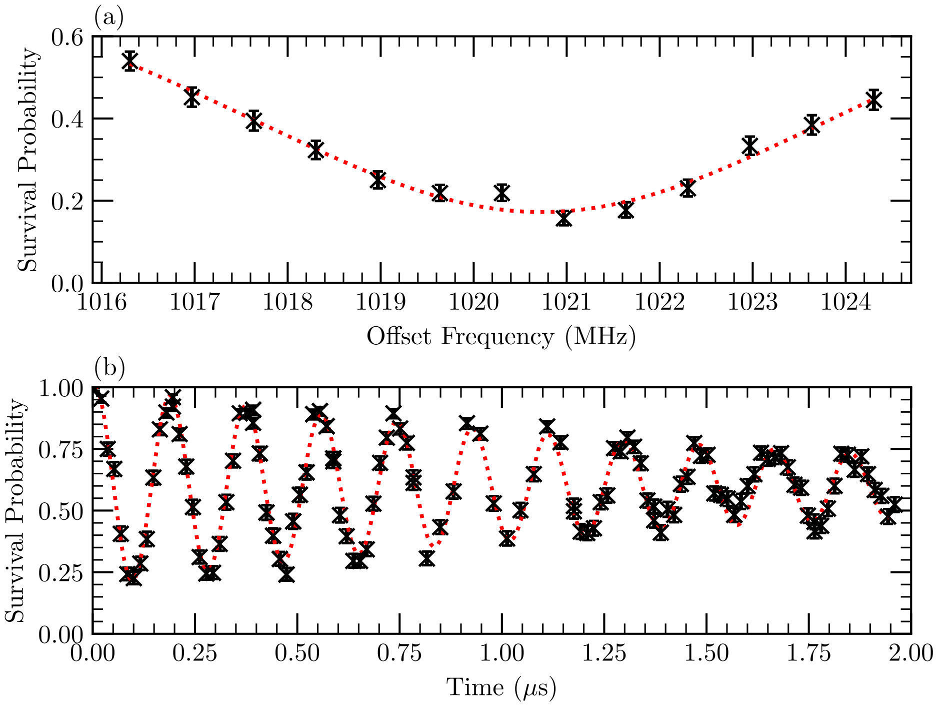

Figure 5.3: Spectroscopy and Rabi oscillations of the $n = 50$ Rydberg transition. Figure (a) shows the linewidth of the transition, relative to the offset frequency of the 1039 nm lock. The red dashed line is a fit of the data to a Lorentzian profile. Figure (b) shows an averaged Rabi oscillation of the central row of the array, with each atom separated by 8μm. The red dashed line is a fit of the data to predicted Rabi oscillations, with frequency $\Omega = 5.42(1)$ MHz and damping time $\tau = 1.7(1)$ μs. Damping is partially due to inhomogeneity; each atom has a different Rabi frequency, and so that averaged oscillations have lower amplitudes at later times.

the ULE cavity using the same offset locking technique as the 459 nm beam. This is incident on an 11×11 array at 8 μm spacing, although only the central row is used for spectroscopy data.

This spectroscopy was at typical operating powers[2] of 47.8(1) mW of light at 459 nm and 792(1) mW at 1043 nm. At the central focus of the 1043 nm light, the Rabi frequency was measured as $\Omega/2\pi = 5.42(1)$ MHz, with a $1/e$ damping time of $\tau = 1.7(1)$ μs. Decay is caused by a combination of noise, off-resonant driving, scattering from the intermediate state and spontaneous decay. This decay is suppressed at higher n due to the longer lifetime.

---

[2] This is power at the atoms, including measured loss through dichroics and other optics.

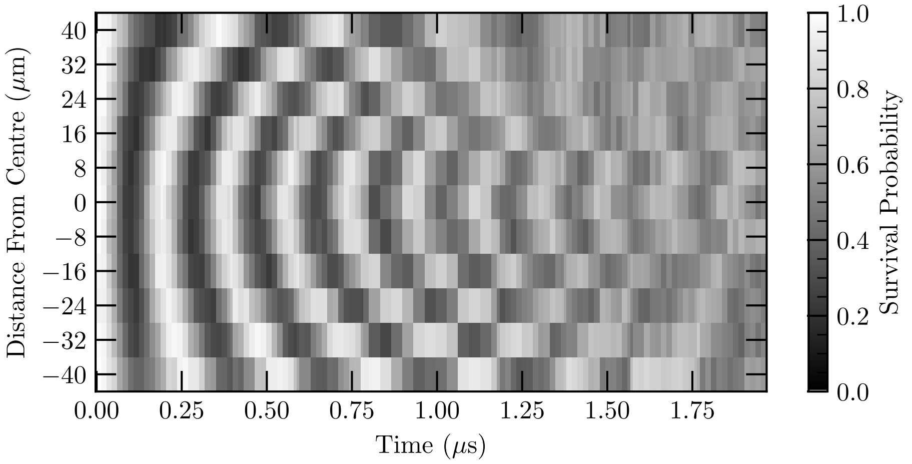

Figure 5.4: Colour plot of Rabi oscillations across the atomic array. The reduced intensity at the edge of the array reduces the frequency of the observed Rabi oscillations. This reduced intensity also provides less AC Stark shift, meaning these atoms have a different resonant driving frequency, and so experience off-resonant driving.

The amplitude of the $n = 50$ oscillations is lower than initially expected. This is due to the detection method, where Rydberg atoms are expelled from the ODTs. At $n = 50$, not all Rydberg atoms are ejected from the deeper traps, and instead decay to trapped states. Imaging from these trapped ground states thus gives a false positive reading of the atoms. For a full analysis of the imperfections of measuring Rydberg state population, see [152].

One further part to note of this data is the effects of the finite beam waists used. As seen in Fig. 5.4, the Rabi frequency is not constant across the array. Furthermore, the AC Stark Shift created by the 1043 nm beam also causes off-resonant driving on the outer rows of the array. This data can be used to estimate the waist of the elliptical 1043 nm beam of 60(1) μm,

The same measurements were used for the Raman beam in Sec. 3.6.3, although the Raman beam is substantially further off-resonance and generates only a negligible AC Stark shift.

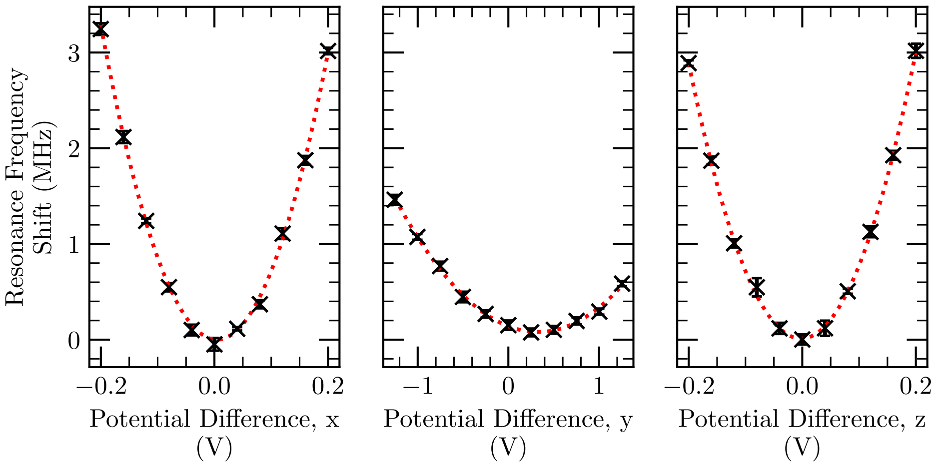

Figure 5.5: Observed electric field shifts on Rydberg transitions. The effect of the shift on the y-axis is less severe. This is due to the alignment of the electrodes along the axis of the ITO-coated lenses, where the electrodes have a greater spacing than along the x and z axes.

5.1.5 Electric Field Nulling

Given the sensitivity of Rydberg atoms to external electric fields, the absolute electric field and electric field gradient must be minimised across the array before working at higher n. As shown in Eq. 2.2.1, electric fields create an energy shift proportional to the square of the electric field strength. This corresponds to a shift in the resonant frequency, where the minimum observed shift corresponds to the minimum electric field strength. The data shown in the section above are demonstrations of the system after this electric field nulling has taken place.

The electric field is controlled through a series of electrodes within the chamber, each of which is connected to a high-voltage supply. This is controlled by an external control voltage connected to a high-voltage supply system. With eight individual electrodes, the electric field at the atoms can be adjusted in each of the three orthogonal axes. This requires a series of spectroscopy measurements at varying control voltages, to find the point of minimum shift. Data from these experiments are demonstrated in fig. 5.5.

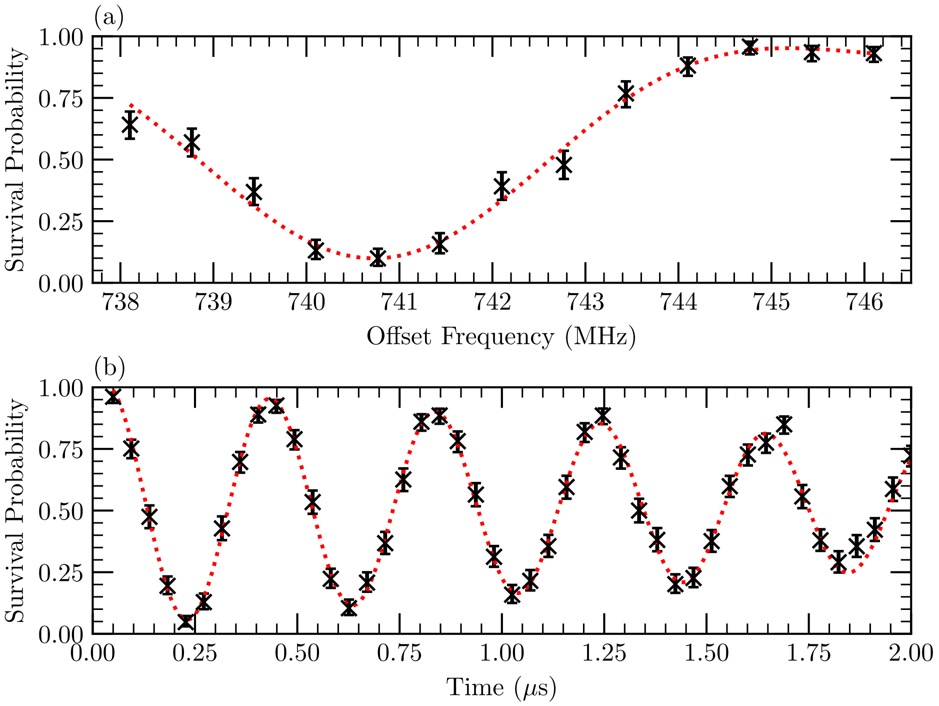

Figure 5.6: Spectroscopy and Rabi oscillations of the $n = 80$ Rydberg transition. Figure (a) shows the linewidth of the transition, relative to the offset frequency of the 1039 nm lock. The red dashed line is a fit of the data to a sinc profile. The linewidth of this transition is narrower than compared to $n = 50$ due to the reduced natural linewidth and reduced Rabi frequency. This is measured at $t\Omega_R = \pi$. Figure (b) shows an averaged Rabi oscillation of the central row of the array, with each atom separated by 20 μm, driven at the resonant frequency measured at $t\Omega_R = 3\pi$. The red dashed line is a fit of the data to predicted Rabi oscillations, including loss due to spontaneous emission. This has measured Rabi frequency of $\Omega_R/2\pi = 2.477(8)$ MHz and $1/e$ damping time 3.0(2) μs. Decay effects also occur due to off-resonant driving as the power of the Rydberg beam changes with time.

The sensitivity of the atoms to electric fields was less than predicted along the y-axis, which is along the axis of the ODTs and high-NA aspheric lenses, which are coated in indium-tin oxide (ITO). As seen in Fig. 3.2, this is the axis also along the plastic lens spacers. This may cause a reduction in the actual electric field gradient caused by the electrodes. This same effect was seen in [185, Chapter 3], which used the same lens/electrode configuration.

5.1.6 n=80

The $n = 80$ transition requires an adjustment to the 1039 nm beam, to a wavelength of 1039.2759(1) nm. Given the increased blockade radius, the prepared array was of size 5×5 at 20 μm spacing. The transition was found with a Rabi frequency of $\Omega_R/2\pi = 2.477(8)$ MHz, as shown in fig. 5.6. The $1/e$ damping time was measured to be 3.0(2) μs. Optimal contrast was much improved compared to $n = 50$, reaching a maximum Rydberg population of $|\langle 80S_{1/2}|\psi\rangle|^2 = 0.943(5)$ after $\Omega_R t = \pi$ (or a 'π-pulse'), as shown in fig. 5.6. Microwave spectroscopy shows D1 optical pumping is limited to a fidelity of 95.7(9)% [87], which is thus the primary limit of Rydberg oscillation amplitude.

5.2 Few-Atom Rydberg Blockade

As introduced in sec. 2.2.1, the dipole-dipole interaction between Rydberg atoms blocks the excitation of more than a single atom to the Rydberg state. Instead, the $|W_N\rangle$ state is prepared, at the enhanced Rabi frequency. This section will show observations of the Rydberg blockade interaction and measurements of the collectively enhanced Rabi frequency.

Firstly, the red-detuned ODT array is prepared such that N atoms are all within blockade radius of each other, but each atom group is separated from any other atoms by a distance much greater than the blockade radius. This configuration is seen in Fig. 5.7. The sorting procedure discussed in Chapter 4 is used to increase experimental yield, with the targets specified such that atoms outside the blockade radius are used as a reservoir, while atoms within are removed unless to be measuremd in the experiment. Experimental data are taken from images post-selected where the sorting conditions are met, such that all atoms are sorted into the target sites for the given N-atom group. Atoms are loaded, sorted and prepared in the $|F = 4, m_F = 0\rangle$ clock

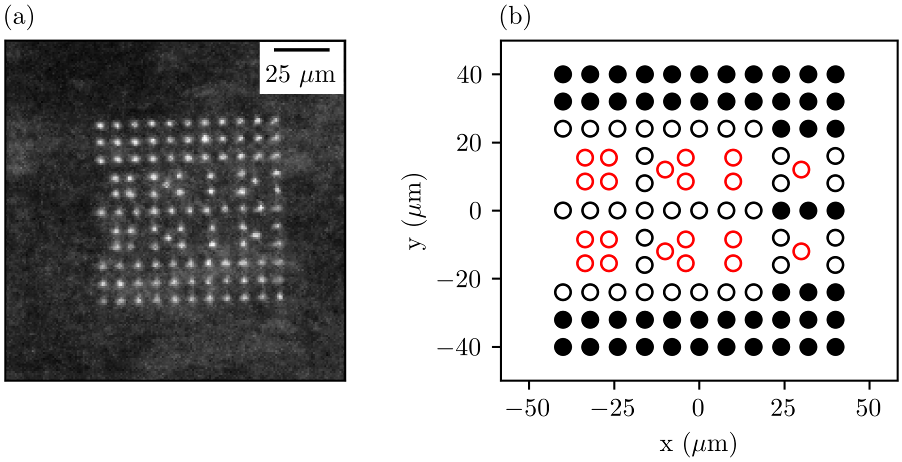

Figure 5.7: Diagram of atom arrays used in measurements of small blockades. Figure (a) is the maximum of pixel counts taken from 400 images of the array used for calibrating array positions. Figure (b) is the requested atom array. Red hollow circles are targets, black hollow circles are atoms to be removed, and black filled circles represent the reservoir.

state using D1 optical pumping. The groups of atoms are prepared in a region of uniform intensity of 1039 nm light, such that the Rabi frequency is equal for all atoms.

Atoms are prepared alone as a control measurement, in pairs separated by 7 μm, in an equilateral triangle of sides 7 μm, and in a square of side 7 μm. For atoms in the square pattern, nearest-neighbour distance is 7 μm, while next-nearest-neighbour is 9.9 μm. This is within the blockade radius, but will affect a reduced shift in comparison to nearest-neighbours.

A lost atom is expected to be a Rydberg atom, while a retained atom is expected to be a ground state atom. Rydberg blockade effects are thus seen as only a single atom lost, where the probability of losing a single atom follows a Rabi oscillation at rate $\Omega_N = \sqrt{N}\Omega$, for N atoms within blockade.

To show that a single Rydberg atom is prepared at the collectively enhanced Rabi frequency, the probability of k excitations $P(k)$ is measured for each of the N atom configurations. As shown in Fig. 5.8, the probability of a single excitation $P(1)$ was seen at greater rates than single excitations. Single atom

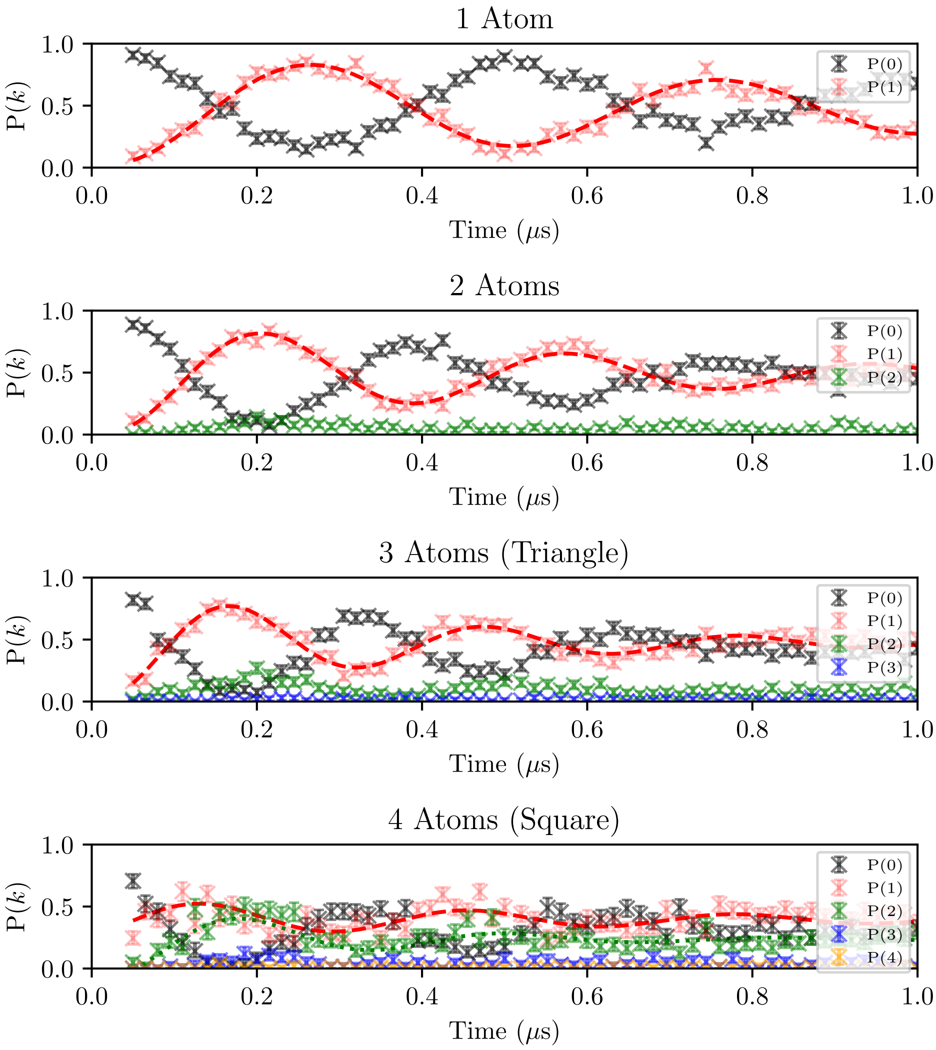

Figure 5.8: Measurements from blockaded Rydberg excitations. This shows the probability of observing k excitations $P(k)$, where $k = 0$ is in black, $k = 1$ in red, $k = 2$ in green, $k = 3$ in blue and $k = 4$ in yellow. Data points with error bars are measured data points, while lines are fits to the single excitation Rabi frequency. Excitations of more than one atom at once is greatly suppressed in all plots, but not entirely eliminated. $N = 4$ plots show a much increased probability of multiple atoms being excited, indicating reduced blockade effects.

Rydberg excitations followed a Rabi frequency at rate $\sqrt{N}\Omega_R$, with the trend shown in Fig. 5.9, with the exception of the $N = 4$ square patterns, where the measured Rabi frequency is reduced. Measured Rabi frequencies are tabulated in Table 5.1. In strongly blockaded patterns ($N = 2, N = 3$),

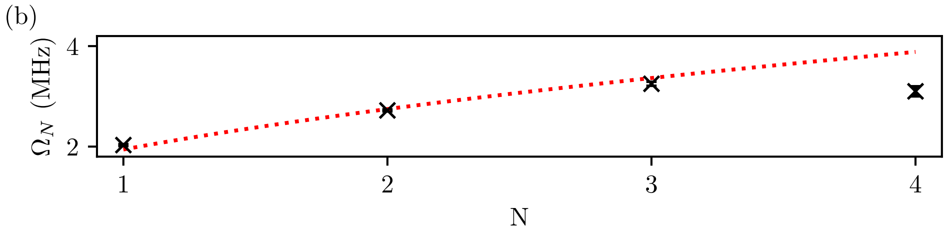

Figure 5.9: Figure (b) shows the measured collective Rabi frequency Ω_N for each configuration, with a $\Omega_N = \Omega\sqrt{N}$ fit used, where Ω is the only free parameter, and ignoring $N = 4$.

N	$\Omega'/2\pi$ (MHz)	$\frac{1}{\sqrt{N}}\Omega'/2\pi$(MHz)
1	2.03(2)	2.03(2)
2	2.72(2)	1.92(2)
3	3.25(4)	1.88(3)
4	3.1(1)	1.54(6)

Table 5.1: Table comparing oscillation frequencies within N-atom Rydberg blockade.

loss of more than one atom was still seen. This is due to a combination of atom loss in recapture, atom loss in imaging such that not all atoms are present despite post-selection, and incomplete state preparation in the $|F = 4, m_F = 0\rangle$ clock state.

The square pattern is less strongly blockaded. The ratio of the resonance shift along the diagonals of the square compared to the edges is given by $\left(\sqrt{2}\right)^{-6} = 1/8$. Calculations using the PairInteraction software expect a shift of 24.6 MHz at 7 μm, and 3.1 MHz across the diagonals. This weakly blockaded effect is seen in the $N = 4$ data, where the probability of losing 2 atoms is far less suppressed than other patterns. The diagonals of the squares are not blockaded without using smaller spacing. For a summary of results, see Table. 5.1.

The measured oscillation frequencies for a single atom are close to the predicted $\sqrt{N}$ scaling, but are slightly lower than expected. This is expected to be due to unaccounted for decoherence effects, imperfect state preparation

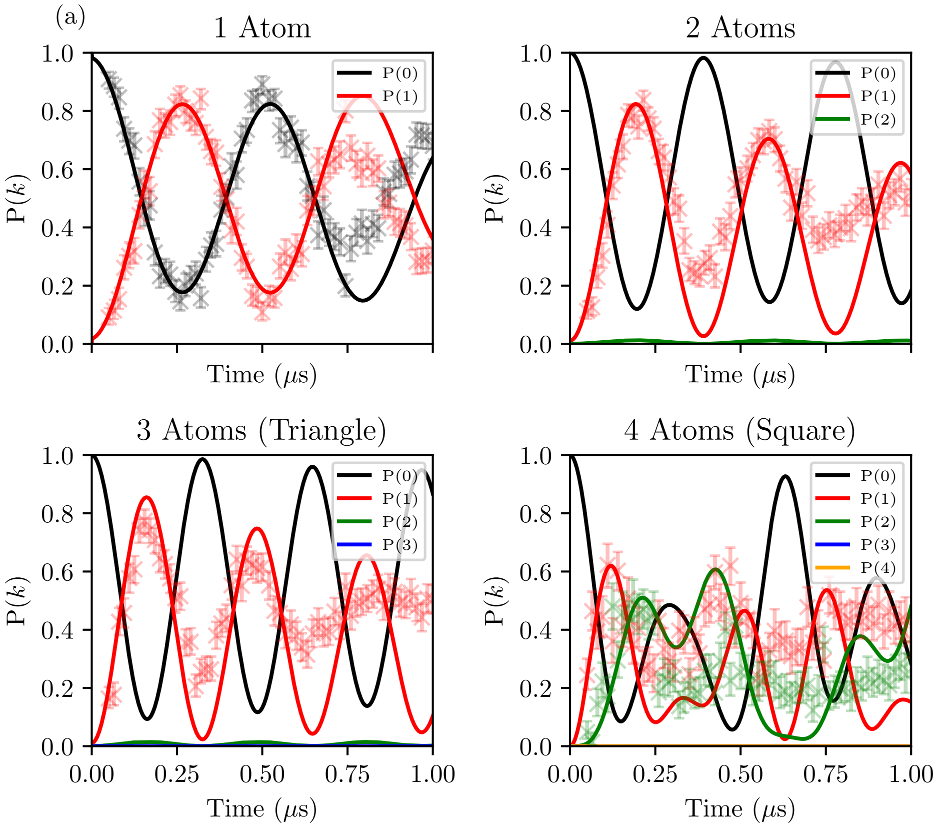

Figure 5.10: Simulation of the system described above. Initial parameters fit to single atom Rabi oscillations as measured in the experiment. The same parameters are then used to plot other parameters.

and atoms lost imaging. Ground state atoms were measured to be lost (as a false postive) at rate 0.9892(5). Given an imperfect D1 state preparation rate of 0.019, the contrast of the Rydberg oscillations implies Rydberg atoms were not ejected at rate 0.063(5). As such, there are experiments where an atom is lost or improperly prepared and does not contribute to the blockade effect, which may correspond to the discrepency in the $\sqrt{N}$ scaling. The single atom Rabi frequency is lower than expected, while dephasing is more rapid than expected. Fits to predicted Rabi frequency where detuning is a free parameter indicates a degree of off-resonant driving and lower-than-expected Rabi frequency. This can be attributed to a reduction in laser power during experiment preparation. This both reduces the Rabi frequency and the AC

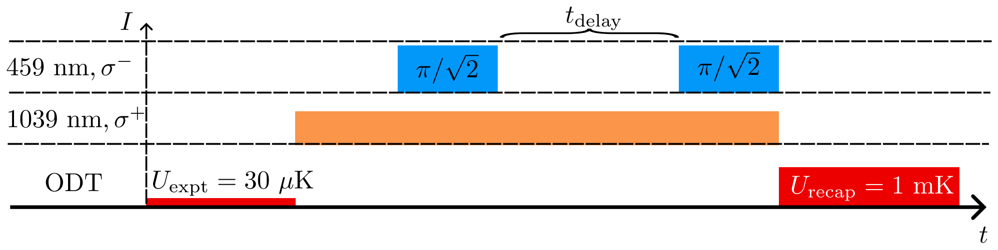

Figure 5.11: Timeline of the V_{dd} measurement.

Stark shift, thus causing the previously resonant offset frequency to be inaccurate. A degree of loss not due to Rydberg excitation is also misattributed to Rydberg excitation.

This system can be modelled mathematically. Using the QuTiP package [96], the N-atom results were fitted to a model of the N-atom Hamiltonian. This model accounts for imperfect D1 state preparation, and that all atoms not in the Rydberg state are detected without loss. As seen in Fig. 5.10, the model matches the Rabi oscillations measured, although as there are no additional loss factors, a degree of decoherence is not accounted for. Nevertheless, the weakly-blockaded behaviour is shown to be responsible for the absence of the $\sqrt{N}$ speedup expected in the square arrangement.

5.3 Dipole-dipole Interaction Measurements

A useful measurement for later work (see Sec. 7.3) is the direct measurement of the Rydberg dipole-dipole interaction. This energy, V_{dd}, is dependent on the Rydberg state and separation between atoms, as detailed in Eq. 2.37.

This uses the method demonstrated in [162], where atoms are prepared in red-detuned ODTs in the $|F = 4, m_F = 0\rangle$ clock state. A $\pi/\sqrt{2}$ pulse[3] at the $|80S_{1/2}\rangle$ transition resonant frequency prepares atoms within blockade in the $|W\rangle$ state. A second $\pi/\sqrt{2}$ pulse detuned by some amount δ then excites to the $|r, r\rangle$ state. A delay period t_{delay} between the two Rydberg pulses

---

[3] Equivalent to a π pulse at the blockade-enhanced Rabi frequency

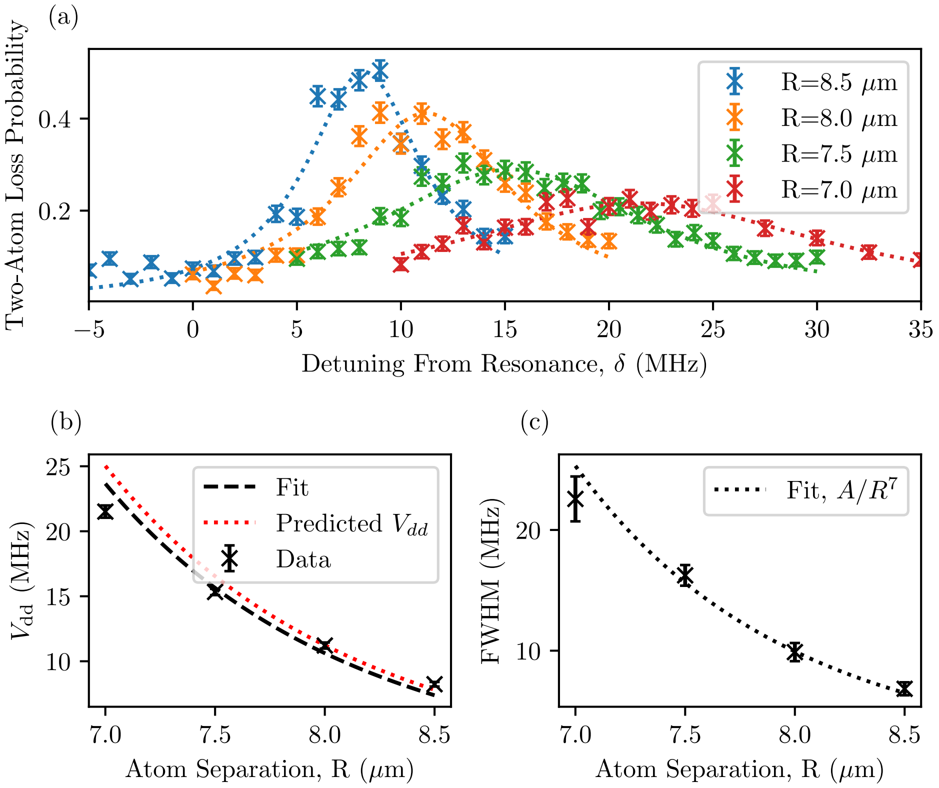

Figure 5.12: Spectroscopy of the Rydberg dipole-dipole interaction. In (a), δ is measured relative to the resonance frequency for a single atom without blockade effects. Each colour refers to a different atomic separation R. The scaling of the centre of these peaks is measured in (b), and the full-with half-maximum (FWHM) in (c). Both are fitted with (A/R^{-6}) curves (black dotted line), where in (b), $A = C_6$ and is thus fixed. In (b), a simulated curve is also used.

increases contrast. Recapture in deep traps ejects these Rydberg atoms, and so is seen as two-atom loss. A timeline for this experiment is seen in Fig. 5.11.

The hologram used for the ODT array uses pairs of atoms at a variety of separations, from $R = 8.5$ μm to 7 μm. Atoms are sorted into pairs, and other atoms within blockade radius are removed. The rate of 2-atom loss is then measured as the input frequency is scanned over a range of values. Contrast was maximised at a delay period of $t_{\text{delay}} = 0.2$ μs. Measurements of the probability of measuring two-atom loss can be seen in Fig. 5.12(a).

The frequency shift measured does demonstrate the R^{-6} value measured, as seen in Fig. 5.12(b). This is compared to simulations using the PairInteraction software package [186]. Fits to the function

$$V_{\mathrm{dd}} = \frac{-C_6}{(sR)^6} \tag{5.1}$$

where s is an arbitrary scaling factor and the only free parameter are also shown. This finds $s = 1.006(7)$. As such, any mismatch between the trap spacing requested by the SLM and the actual atomic spacing is within measurable error.

The maximal probability of two-atom loss reduces as R decreases, while the spectroscopic feature is broadened. Due to the R^{-6} scaling of V_{dd}, the effect of V_{dd} grows rapidly over a small range of measurements. Furthermore, since the atoms are not static within the traps due to the harmonic nature, the separation is only well-defined as to the mean position of the traps. For small perturbations dR, the broadening effect is proportional to the perturbations in energy dV_{dd}, such that

$$dV_{dd} = \frac{6C_6}{R^7} dR. \tag{5.2}$$

The mean distribution of atoms within the traps is constant and dependent on trap frequency and atom temperature, forming a constant Maxwell-Boltzman distribution of relative spacing in R, that is δR is independent of R. The wavefunction of a Cs atom in the ground state of a $2\pi \cdot 8$kHz trap has spatial distribution 68.9 nm. These are the parameters for the radial trapping potential of atoms in the trapping array, although in the experiment the atoms can be expected to be above the ground state. For two atoms, the standard deviation in separation dR would then be 97.5 nm. At 6.8 μm, $dV_{dd} = 2\pi \cdot 2.55$ MHz $> \Omega$, and thus variations in potential are greater than the linewidth of the transition and are entirely off-resonance. As such, transitions are suppressed using this experimental method.

Fits of the full-width half-maximum (FWHM) to relative to R show an approximate $1/R^7$ dependence, as in Fig. 5.12(c), where a $1/R^7$ curve is given as a guide to the eye. Calculating $d\left(\ln(\text{FWHM})\right)/d\left(\ln(R)\right)$ with a linear fit shows a gradient of $-6.4(5)$. Deviations from the expected value of 7 as per Eq. 5.2, where FWHM $\propto dV_{\text{dd}}$, are due to the large shifts in V_{dd}, such that the assumptions of small dR are no longer valid.

5.4 Entangled State Preparation

The Rydberg blockade effect entangles particles, as required for a quantum computing system. It does not, however, entangle particles in the hyperfine basis. The qubit states $|0\rangle, |1\rangle$ the SQuAre experiment intends to use for digital quantum gate operations are the Cs clock states $|F=3, m_F=0\rangle$ and $|F=4, m_F=0\rangle$. The state generated by a π pulse at the enhanced Rabi frequency is, for two particles, $\frac{1}{\sqrt{2}}\left(|1,r\rangle + |r,1\rangle\right)$. To generate an entangled state in the hyperfine basis, we must also implement a rotation about the Bloch sphere to the $|1\rangle$ state. The Raman system is used in this case.

This section details the methods used to generate the $|\Psi^+\rangle$ maximally entangled state within the hyperfine basis, with calculations of the fidelity of the state generated.

5.4.1 State Preparation

A pulse of resonant Raman light at t_{Raman} acts as a $R_X(\theta)$ gate in the hyperfine basis, where

$$\begin{aligned} R_X(\theta)\,|0\rangle &= \cos(\theta/2)\,|0\rangle - i\sin(\theta/2)\,|1\rangle \\ R_X(\theta)\,|1\rangle &= \cos(\theta/2)\,|1\rangle - i\sin(\theta/2)\,|0\rangle\,. \end{aligned} \tag{5.3}$$

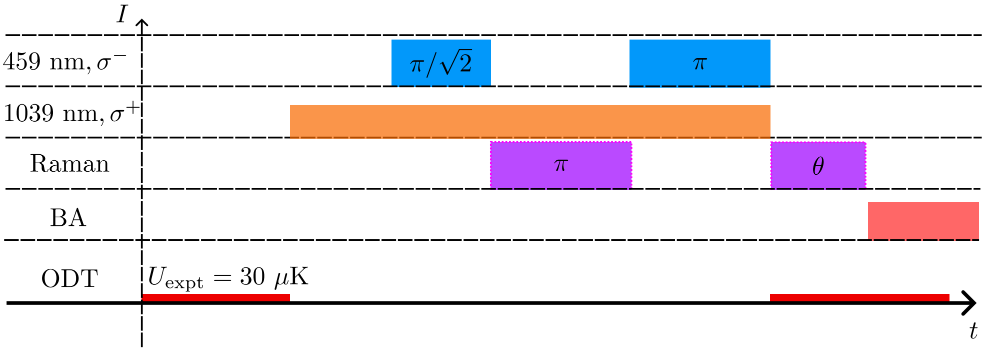

Figure 5.13: Timeline of entangled state fidelity measurement experiment. Two Rydberg pulses are applied with area $\Omega_{\text{Rydberg}}t = \pi/\sqrt{2}$ and π. Two Raman pulses are also applied with area $\Omega_{\text{Raman}}t = \pi$ and θ, where θ acts as the independent variable of the experiment. 'BA' refers to state-selective blow-away, to remove atoms in the $|0\rangle$ state.

Here, $\theta = \Omega_{\text{Raman}}t_{\text{Raman}}$. By applying the pulse with $\theta = \pi$, the Raman light can act as an X gate, where

$$\begin{aligned} X\,|0\rangle &= |1\rangle \\ X\,|1\rangle &= |0\rangle\,, \end{aligned} \tag{5.4}$$

where arbitrary global phase terms have been ignored. Once atoms are in the $\frac{1}{\sqrt{2}}\left(|1, r\rangle + |r, 1\rangle\right)$ state, this X gate is applied globally to prepare the $\frac{1}{\sqrt{2}}\left(|0, r\rangle + |r, 0\rangle\right)$ state.

Atoms within the $|0\rangle$ state are off-resonant to the Rydberg light and so are not coupled. As such, another Rydberg π-pulse transfers population from the Rydberg state to the $|1\rangle$ state only. This prepares the state

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|1, 0\rangle + |0, 1\rangle\right) \equiv \left|\Psi^+\right\rangle, \tag{5.5}$$

with some arbitrary global phase. $|\Psi^+\rangle$ is the one of the four Bell states, which are four orthogonal and maximally entangled two-qubit states [187].

Proof that the state generated is the Bell state requires further experimentation. Following the method laid out in [188], we can apply the Raman light for a further period of time t_{Raman}, using this as an analysis pulse to extract information about the prepared state. Atoms in the $|1\rangle$ state are removed with state-selective blowaway. A timeline of this experiment is seen in Fig. 5.13. The fidelity of the prepared state F is given by

$$F = \left\langle \Psi^{+} \middle| \rho \middle| \Psi^{+} \right\rangle = (P_{1,0} + P_{0,1}) + \mathrm{Re}\left(\rho_{1,0}\rho_{0,1}\right). \tag{5.6}$$

The $\mathrm{Re}\left(\rho_{1,0}\right)$ term is the real part of the off-diagonal coherences. This is referred to as the coherence term. As the experiment only measures the probability after projection of the wavefunction onto either the $|0\rangle$ or $|1\rangle$ states, measurement of the initial probabilities is not enough to measure the coherence term. This is achieved by comparing measured probabilities as t_{Raman} is changed.

State fidelity can be measured through observation of parity oscillations, which give information on coherences. The parity Π is defined by

$$\Pi \equiv (P_{1,1} + P_{0,0}) - (P_{0,1} + P_{1,0}) \tag{5.7}$$

which oscillates with angular frequency $2\Omega_{\mathrm{Raman}}$ under the Raman analysis pulse if perfectly prepared in the $|\Psi^{+}\rangle$ state [188]. Evidence of imperfect state preparation occurs as a further oscillation of angular frequency Ω_{Raman}. As such, the fitted function is

$$\Pi\left(\theta\right) = \Pi_0 + A\cos(\theta) + B\cos(2\theta) \tag{5.8}$$

where $\theta = \Omega_{\mathrm{Raman}} t_{\mathrm{Raman}}$.

As the measurement taken is only the probability that an atom is observed or not, this limits the direct measurements that can be taken. Given that each atom has an independent probability of being lost before measurement, loss during recapture of atoms, state-selective blow-away or during imaging can result in a false-positive reading of the Rydberg state. Imperfect preparation of atoms in the $|1\rangle$ state will also result in detection of atoms in the $|0\rangle$ state, as a false negative. These losses must be considered for calculating the prepared state fidelity. Furthermore, it can be that atoms fail to be prepared in the Rydberg state due to imperfect optical pumping or variations in optical power. There is a distinction between two atoms being present but not in the $|\Psi^+\rangle$ state and only one atom being loaded, but having the appearance of not being in the $|\Psi^+\rangle$ state. Loss corrections must correct for only the second source of error here, and not the first, when determining F.

5.4.2 Experimental Results

This was measured over 16 pairs of atoms, where atoms are separated by 8 μm within the pair and 16 μm between pairs. Exemplar measurements of the experimental probabilities can be seen in Fig. 5.14. The Rydberg Rabi frequency was measured as $\Omega_{\text{Rydberg}}/2\pi = 1.01(2)$ MHz, and the Raman Rabi frequency as $\Omega_{\text{Raman}}/2\pi = 0.940(4)$ MHz. This experiment was performed before the addition of cylindrical optics to reduce field inhomogeneity and increase field strength, leading to reduced Rabi frequencies compared to earlier in this work.

Using the same methods as in [188], the coherence term can be extracted from the data from parity oscillations, such that

$$\Pi(\pi/2) = 2\,\mathrm{Re}\,(\rho_{1,0}\rho_{0,1}) + L_a L_b, \tag{5.9}$$

where L_i is the probability of failing to prepare atoms in the Rydberg state. Measurements of the Rydberg Rabi oscillations for individual atoms are used to determine L_i, where the contrast of the Rabi oscillations is determined by failure to prepare atoms in the Rydberg state. This is measured to be $L_a = 0.055(1), L_b = 0.063(1)$. Microwave spectroscopy can be used to determine

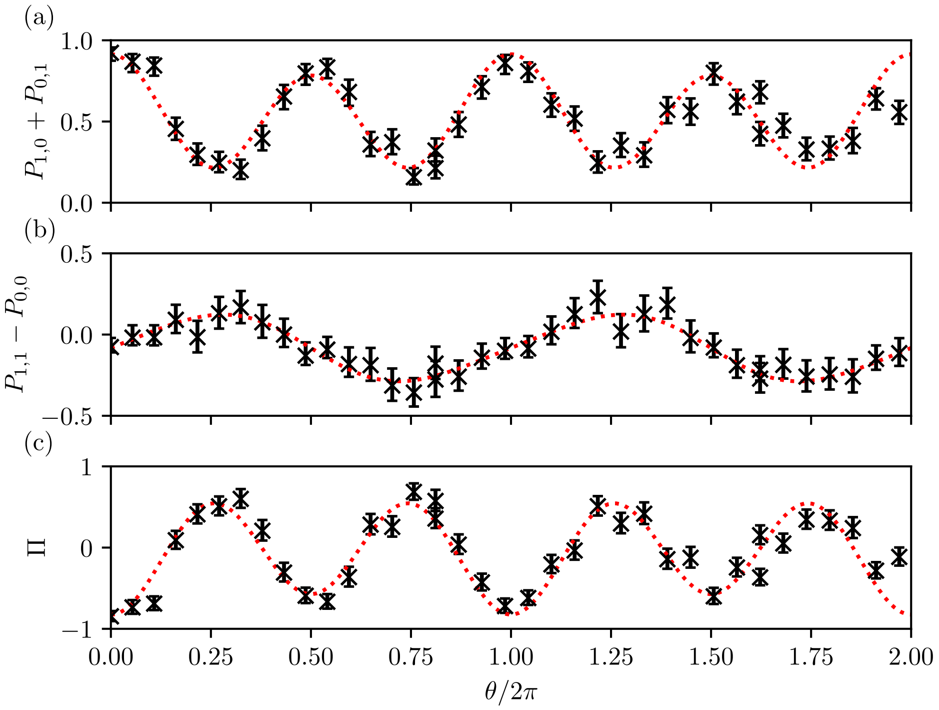

Figure 5.14: Examplar coherence measurements of the $|\Psi^+\rangle$ entangled state. Figure (a) shows the probability of measuring either the $|1,0\rangle$ or $|0,1\rangle$ odd-parity state, with large oscillations over a period $\theta/2$ and small oscillations over period θ. Figure (b) shows the difference in probability of measuring the $|1,1\rangle$ and $|0,0\rangle$ states. This evolves over period θ, and would have amplitude 0 if the $|Psi^+\rangle$ is perfectly prepared. Figure (c) is the oscillation in parity Π, as defined in Eq. 5.4.1. (a) and (c) use a fit to Eq. 5.4.1 with a red dashed line, while (b) fits to the same with a phase shift of $\pi/2$.

the proportion of atoms that are not prepared in the $|F=4, m_F=0\rangle$ state. The probability of failing to prepare atoms in the intended ground state is measured to be 4.3(9)% as mentioned in Sec. 5.1, and as such the remaining factors for L_a and L_b are due to power instability and beam inhomogeneity across the array.

For the best prepared pair, fits to Eq. 5.4.1 find $\mathrm{A}=-0.13(4), \mathrm{B}=-0.62(4)$, such that the measured fidelity is $F=0.73(6)$ [188]. When corrected for the probability of neither atom being lost during the experiment $P_{\text{recap}}=0.95(1)^2$, this increases to the corrected fidelity $F_C=0.81(6)$. For the eight pairs on resonance with both the Rydberg and Raman beams, the average fidelity measured was $F = 0.69(3), F_C = 0.77(5)$.

These experiments were performed early in the experiment cycle when the first evidence of Rydberg excitation was obtained. Since then, further improvements have since been made to the experiment including cylindrical optics for improved Rabi frequencies and homogeneity. Also introduced are noise-eating circuits to reduce per-shot noise from $> 2\%$ to $< 0.5\%$ for the Rydberg lasers, and increased power available for the 1039 nm beam.

5.5 Summary

The Rydberg excitation lasers used in the SQuAre project have been demonstrated to implement global excitation of atoms to Rydberg states. This has been demonstrated with both single atom spectroscopy and the measurement of atom-atom interactions, such that the atomic arrays form a basis for further quantum computing experiments. Lower-n Rydberg states (such as $n = 50$) can be used to inform the response of atoms to Rydberg light within an array that at higher-n states is affected by Rydberg blockade effects. The entanglement of these atoms is measured and the fidelity of this coherence is reported, acting as a 2-qubit quantum gate.

Rydberg excitation fidelity is the greatest limiting factor for Rydberg gates. Improving this fidelity would require greater optical pumping fidelity, which would require Raman-assisted state preparation as introduced in Sec. 3.5.1. Greater homogeneity of the Rydberg intensity across the array is also a concern for global gates over large numbers of atoms. For this purpose, an SLM system is planned for introduction on the 1039 nm beam path.

Currently, only 1 W of power is coupled to the chamber due to the power limits of the fibre optic cables used. The replacement of these fibres will thus allow for an order of magnitude more power from the Azurlight fibre amplifier, which can output up to 20 W power. At these higher powers, greater intermediate state detuning Δ becomes more viable. As the Rabi frequency scales with Δ^{-1} but the scattering rate scales as Δ^{-2}, a greater power means the same Rabi frequency can be used with a reduced scattering rate, thus improving Rydberg excitation fidelity.

Rydberg blockade effects will also be exploited in Chapter 7 for uses in analogue quantum computing, in particular to provide conditions for weighted graph problems. These effects also provide a route to further experiments with quantum simulation in Rydberg arrays. Such work has already been demonstrated in investigating quantum quench dynamics [189, 190] or quantum many-body scars [67, 75, 191].

Chapter 6

Blue Detuned ODTs

6.1 Introduction

6.1.1 Motivation

The experiments involving neutral atom trapping so far presented in this work all share the use of red-detuned optical dipole traps (ODTs). Whether at 1064 nm for the static ODT array or at 1039 nm for the sorting ODT, these traps have negative detuning relative to the D1 and D2 resonances that they couple to at $895, 852$ nm respectively. As shown in Sec. 3.3.3, these red-detuned traps are used to trap atoms in regions of high intensity.

Complications occur with trapping Rydberg atoms. Rydberg atoms have the polarisability of a free electron, which are repelled from regions of high intensity, and can be used for detection of Rydberg states, as introduced in Sec. 3.5.3. This does, however, require the red-detuned ODTs to be turned off periods of excitation to the Rydberg states. This limits interactions with Rydberg atoms to a period where atoms can be released from the traps and recaptured without minimal loss. In Sec. 5.1, this procedure was used, limiting interactions with Rydberg atoms to 8 μs at most.

Experiments performed in free-fall prohibit long-term interactions with Rydberg atoms. For a quantum computing experiment, a confined Rydberg atom that can be contained for long periods is ideal. This increases the duration of gate sequences that can be applied, while increasing fidelity as atoms are static through the operation. Furthermore, trapping ground-state atoms in blue-detuned ODTs minimises the AC Stark shift applied to the atoms while they are trapped. Random movements within the trap cause decoherence in qubits that is irreversible through spin-echo techniques. As such, the use of blue-detuned ODTs for the trapping of Rydberg atoms is an attractive prospect.

Experiments involving arrays of blue-detuned ODTs have involved overlapping Gaussian beams to create an array of high fidelity qubits [192–194], or a 'grid of blue detuned lines' generated by diffractive optical elements [195]. This chapter will focus on the use of holographic techniques to create arrays of blue-detuned ODTs, with the intention of finding scalable methods for the coherent trapping of both ground state and Rydberg state atoms in arbitrary geometries. This is a development of work shown in [132], using a phase mask on an SLM to generate a single blue detuned ODT.

6.1.2 Optical Interactions

The optical system used for these blue-detuned ODTs is detailed in Sec. 3.3.3, where a 800 nm Ti:Sapph laser is used as a coherent light source. This beam is incident on a Spatial Light Modulator (SLM) to generate the holographic trapping potential. The wavelength used has a negative scalar polarisability for both the $\left|6S_{1/2}\right\rangle$ ground state and the $80S_{1/2}$ Rydberg state. As such, it is a suitable wavelength for trapping.

To determine dipole potential of a neutral atom in the 800 nm light requires a different approach for ground and Rydberg state atoms. The dipole potential

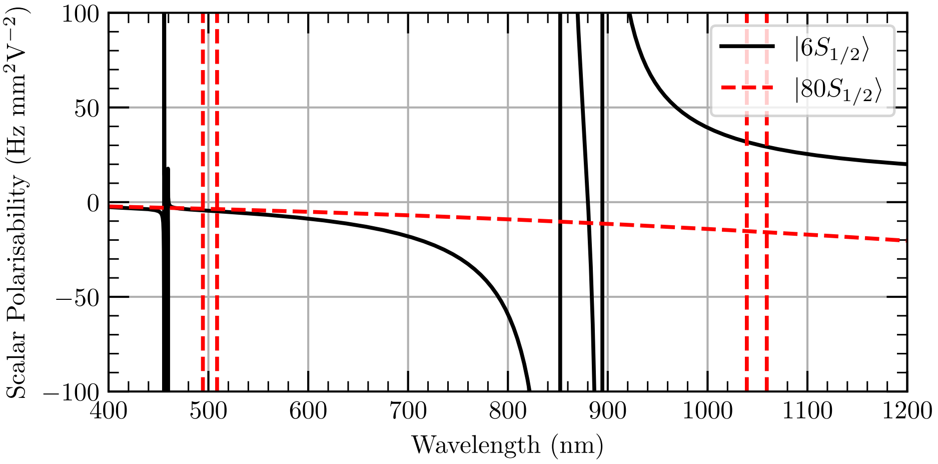

Figure 6.1: Plot of the scalar polarisability for ground and Rydberg states. Ground state polarisability is plotted in black solid lines, while red dashed lines are for the $\left|80S_{1/2}\right\rangle$ state, assuming a constant intensity over the integrated region. Aside from at atomic transitions, the Rydberg state has negative scalar polarisability, unlike the ground state. Data are calculated using ARC[101].

at the ground state is given by Eq. 2.49, where

$$U_0 = -\frac{3\pi c^2}{2}\left(\frac{1}{3}\frac{\Gamma_{D1}}{\omega_{D1}^3\Delta_{D1}} + \frac{2}{3}\frac{\Gamma_{D2}}{\omega_{D2}^3\Delta_{D2}}\right). \tag{6.1}$$

The relative factors of $\frac{1}{3}$, $\frac{2}{3}$ are due to the relative branching ratios of the D1 and D2 paths respectively.

As the electron is sufficiently delocalised from the nucleus, it can no longer be assumed that the electron is in the presence of a homogenous oscillating field. Instead, the Born-Oppenheimer approximation is applied, where the nucleus is taken as a static point in the field and the potential of the electron in the field must be integrated over the spatial variation of the wavefunction [196]. As such, the potential of an atom at point $\mathbf{R}$ is given by [132]

$$U_{n,l,j,m_j}(\mathbf{R}) = \frac{e^2}{2m_e\epsilon_0 c\omega^2}\int d^3r I(\mathbf{r}+\mathbf{R})|\psi_{n,l,j,m_j}(\mathbf{r})|^2 \tag{6.2}$$

where m_e is the mass of an electron and $\psi_{n,l,j,m_j}(\mathbf{r})$ is the spatial wavefunction of the Rydberg electron. This wavefunction for $n = 80$ requires the integration over a sphere of diameter > 1 μm.

6.1.3 Generation of Bottle-Beam Arrays

As introduced above and in [132], the generation of a single blue-detuned ODT can be achieved with a single hologram. The first step of generating a blue-detuned ODT is the same as generating a red-detuned ODT. The target intensity used for the adaptive weighted Gerchberg-Saxton (awGS) algorithm [144] is a Dirac δ-function. Since the incident beam is a collimated Gaussian beam, and with no hologram the Fourier transform of this beam is simply a tight Gaussian beam, the hologram pattern serves to simply adjust the position of the tightly focussed Gaussian beam.

Blue detuned ODTs can be generated by modifying this phase profile. The phase at the centre of the input beam is modulated by π, over a radius r_π. This adjusts the generated field such that at the focus of the Fourier plane, there is now destructive interference. As the phase modulation does not adjust the total power of the beam, this lost power instead forms a peak intensity around the focal point as a sort of bubble of intensity. With well-chosen r_π, this destructive interference is total, such that at the focus, there is instead zero field. For an input electric field $E(r)$, this is given by the condition

$$\int_0^{r_\pi} rE(r)dr = \int_{r_\pi}^{\infty} rE(r)dr. \tag{6.3}$$

When $E(r) = \exp\left(-r^2/w_0^2\right)$, the electric field strength of a Gaussian beam with $1/e^2$ waist w_0, this is met by the condition

$$r_\pi = w_0\sqrt{\ln(2)}. \tag{6.4}$$

Fig. 6.2 shows simulations of how the addition of a phase mask can generate the desired profile in two dimensions.

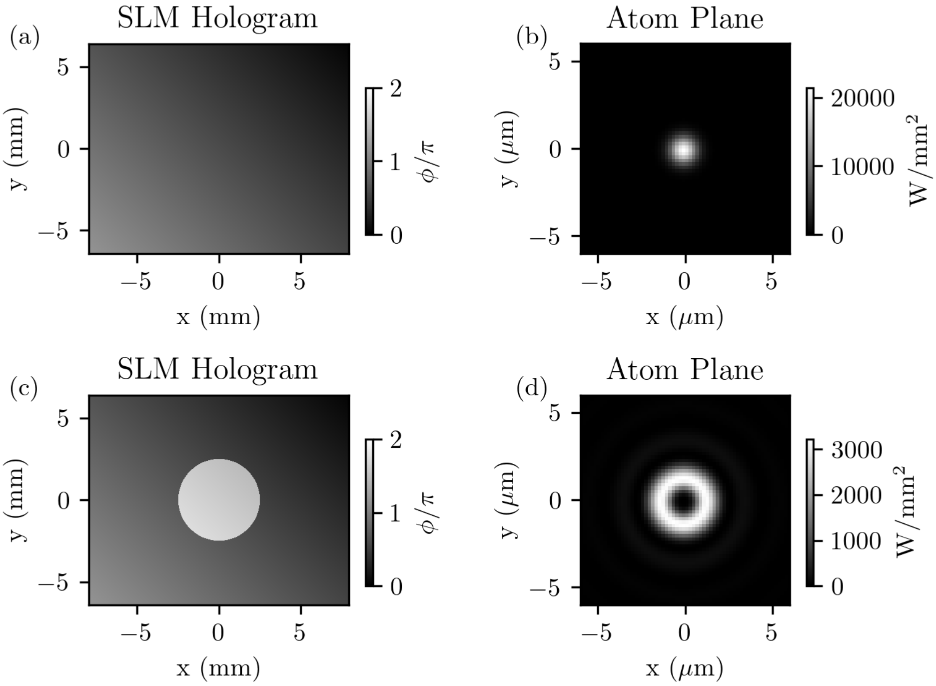

Figure 6.2: Simulations of a single optical dipole trap, with and without the phase mask applied, using an input beam of $1/e^2$ waist 3 mm. Figure (a) shows the hologram generated by the awGS algorithm for creating a focussed spot at the origin of the image plane, where ϕ is the complex phase applied by the SLM to the input beam. Dimensions and resolution matches the SLM used in the experiment. Figure (b) shows the simulated intensity generated by this phase profile. Figure (c) shows the same hologram as in (a), with the π-phase mask applied. Figure (d) is the simulated intensity, with zero intensity in the centre of the trap, and a ring of trapping light around it.

This principle scales up to any input array. A δ-function point in the focal plane is dependent on all spatial parts of the input field. As such, the π-phase mask required is the same for any spot generated. Generating an array of focussed points and then adding the π-phase mask would thus apply the same effect to all points simultaneously.

6.2 Fourier Optics

To assess the requirements and limitations of an optical system, a model of the system was developed. This section will introduce the techniques used to calculate the trapping fields generated by the holographic techniques.

6.2.1 Introduction to Fourier Optics

The principle of holography relies on how a lens acts as a Fourier transformer for an input field [141]. The Fourier transform acts as a transform between position x, y and a set of input k vectors k_x, k_y, such that the input light to be transformed can be decomposed into a set of plane waves, each with vector $\mathbf{k} = k_x\hat{\mathbf{x}} + k_y\hat{\mathbf{y}}$. For a monochromatic input field, $|k| = \omega/c$. The transform is thus defined as

$$\begin{aligned} \tilde{E}(k_x, k_y) &= \frac{1}{2\pi}\int_{\infty}^{\infty}\int_{\infty}^{\infty} dxdyE(x,y)\exp\left(i\left(k_x x + k_y y\right)\right), \\ &= \mathcal{F}_{x,y}\left[E(x,y)\right] \end{aligned} \tag{6.5}$$

This model assumes the paraxial approximation, and that the lens is infinitely large. A further assumption is that the beam is linearly polarised and no ellipticitly is applied by any phase mask.

The computation input here uses the fast Fourier transform (FFT) algorithm to calculate the electric field at the output focal plane, given that the field at the input plane is known. The input field is modelled as a grid of $N_x \times N_y$ pixels, each with width δ_x, δ_y. After the Fourier transform, the size of these pixels is given by $\tilde{\delta_x} = \frac{\lambda f}{N_x \delta_x}$, which defines the resolution for this model[1]. For the model to match actual physics, the input field must contain effectively all the total field of the beam, such that edge effects are not included. Furthermore, effects on the scale below this resolution are not considered.

---

[1] While the amplitude in the Fourier plane represents the relative amplitude of corresponding plane wave in the input beam, when considering Fourier optics, these pixels map to a real image with size $\tilde{\delta_x}$.

6.2.2 Angular Spectrum Analysis

The Fourier transform technique only provides information for the field at the focal planes of the Fourier lens, so far labelled x, y. Propagation along the z axis is therefore required to find the full intensity profile, to assess if the field generated can trap atoms.

As shown in [197], the spatial variation of the electric field is given by the Helmholtz equation,

$$\left(\nabla^2 + k^2\right) E(x, y, z) = 0. \tag{6.6}$$

Given that

$$k^2 = k_x^2 + k_y^2 + k_z^2, \tag{6.7}$$

a further transform operation can be applied to find how a given field evolves into the z plane. This propagator is the Helmholtz propagator, defined by

$$\hat{H}(k_x, k_y, z) \equiv \exp\left(iz\alpha\right), \tag{6.8}$$

where

$$\begin{aligned} \alpha &= +\sqrt{k^2 - k_x^2 - k_y^2},\ k^2 > k_x^2 + k_y^2, \\ &= i\sqrt{k_x^2 + k_y^2 - k^2},\ k^2 < k_x^2 + k_y^2. \end{aligned} \tag{6.9}$$

This follows the definition used in [197], where the sign of the square root is chosen to remove evanescent effects. Within the Fourier plane, this acts as a further Fourier transform operation, such that

$$\begin{aligned} E(x, y, z) &= \mathcal{F}_z^{-1}\left[\hat{H}(k_x, k_y, z)\mathcal{F}_{x,y}\left[E(x, y, 0)\right]\right], \\ E(k_x, k_y, k_z) &= \hat{H}(k_x, k_y, z)\tilde{E}(k_x, k_y, 0). \end{aligned} \tag{6.10}$$

Choice of z is arbitrary, but as z grows, errors arising from the approximations used become more pronounced. The computational domain of z is not necessarily restricted to the same computational domain in x, y, but in the Fourier plane, $\tilde{\delta}_z$ is limited to the same range and resolution as $\tilde{\delta}_x$.

6.3 Blue-Detuned ODT Simulations

The optics used in the modelled system use $f = 12.5$ mm, $\lambda = 800$ nm, Crucially, δ_x, δ_y is limited by the pixel size of the SLM, as this is the spatial resolution of the experiment. The Hamamatsu X13138-02WR SLM uses square pixels of size 12.5 μm, on a 1272×1024 pixel grid. To simulate resolutions below this pixel limit, the field must be supersampled, such that the phase across a region of pixels is considered to be constant, although control at this subsampled resolution is not possible [145]. The total simulated power used was 5 mW per trap site in the trapping plane; this assumes ideal efficiency along all paths, and so will require a scaling factor between input power and trap power.

Initial hologram generation uses the adaptive weighted Gerchbeg-Saxton (awGS) algorithm, as defined in [144]. For fast computation, a GPU[2] is used for awGS calculations using the CUDA package [198]. The maximumum resolutions N_x, N_y here is up to 2^{12}, such that the memory required for calculation was < 5 GB, the dedicated memory of the GPU. For the highest possible resolution as presented in this chapter, $N_x, N_y = 2^{13}$. These simulations took place on CPU using 20 GB of internal memory, trading speed for precision. The simulated field covers the entire SLM screen without loss and the simulated spatial resolution at the trapping plane is $\delta_i = 97$ nm.

6.3.1 Single Traps

To first reproduce the results of [132] and demonstrate that the trapping scheme can theoretically trap both ground state and Rydberg atoms, the

---

[2] NVIDIA Quadro P2000

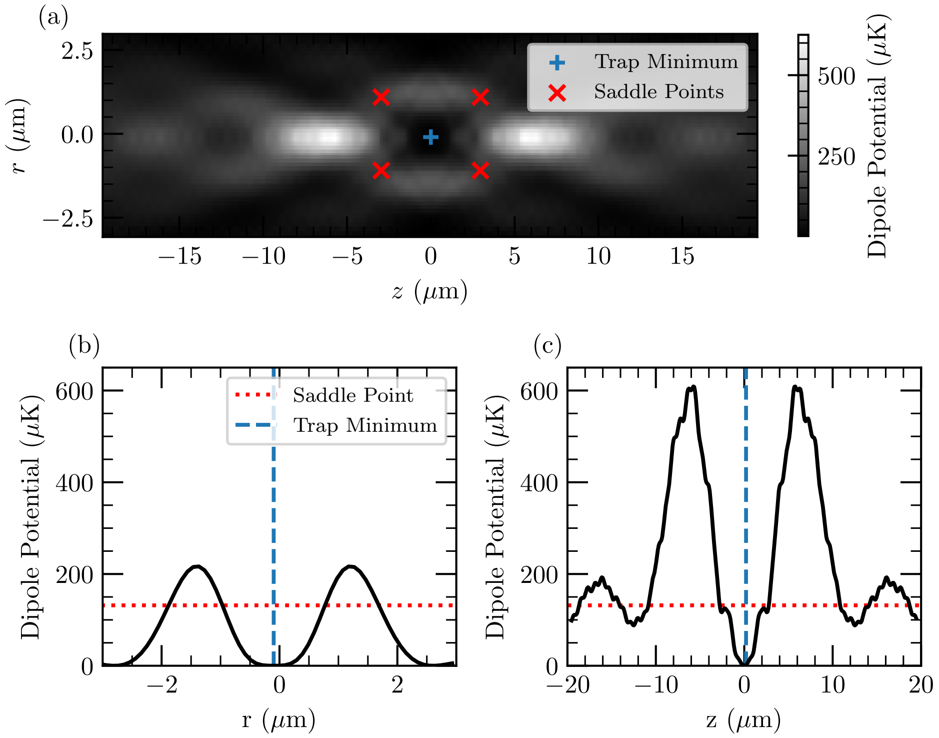

Figure 6.3: Energy profiles of the dipole potential of a ground state atom in a blue-detuned ODT. Total power is 5 mW in the trapping plane. Figure (a) shows a colour plot of the trap in two dimensions, with the trap minimum and saddle points marked by a blue + and red crosses respectively. Figures (b) and (c) show the energy spectrum across the radial and axial planes, with the trap minimum in that axis marked by a blue dashed line, and saddle point energy by a red dotted line.

generation of a single trap is simulated. The results of the angular spectrum method on an input beam of $1/e^2$ waist size 3 mm are shown in Fig. 6.3.

The trap depth of the blue-detuned ODT is determined by the minimum energy required to escape the trap. This is determined by the saddle-points of the trap in three dimensions. This forms a locus of points in a circle around the trap. By taking a cut in the $x - z$ or $y - z$ planes, a sample of these points can be measured to find the trap depth. Simulations here were calculated using 5 mW total power in the trapping plane. This shows a trap depth of 131.7 μK, or 26.34 μK/mW. Trap frequencies can also be measured, by taking the second derivative of the trap potential at the bottom of the

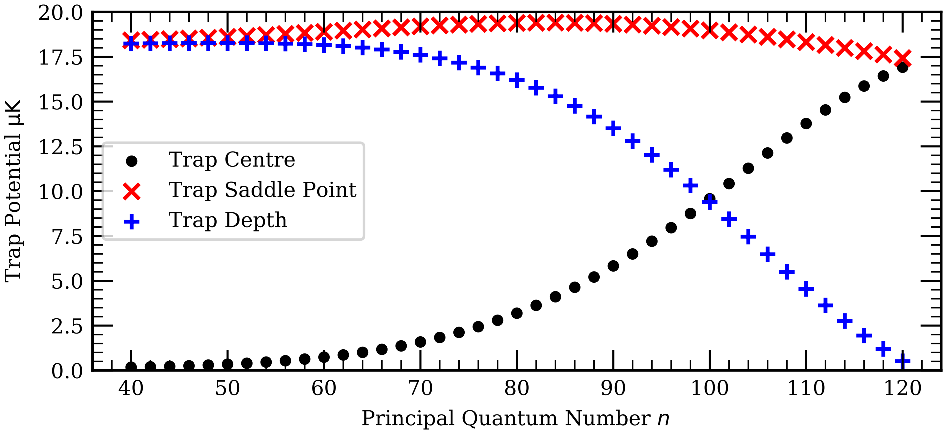

Figure 6.4: Calculated Rydberg potential energy within a blue-detuned ODT. Potential is measured for a range of n values at the trap minimum and saddle points, to find the effective trap depth. Calculated using a total trap power of 5 mW in the trapping plane.

trap. Fits to the energy profiles are used to find $\omega_r/2\pi = 23.7$ kHz, $\omega_z/2\pi = 10.7$ kHz in the radial and axial directions respectively.

Analysis of the pondermotive Rydberg potential at the trap minima and saddle points demonstrates trap effectiveness in trapping Rydberg states. This is shown in Fig. 6.4. Cubic splines are used to interpolate simulated potentials, which are assumed to be smooth. Numerical integration is calculated using RK4 methods [199]. As n grows, the trap minimum becomes much greater than zero as the wavefunction intersects with a region of non-zero intensity. Both the trap depth is reduced and the force of the field on the Rydberg atom is increased, such that Rydberg atoms will be both heated and less effectively trapped. At $n = 80$ and 5 mW in the trapping plane, the traps simulated above have a minimum potential of 3.2 μK and trap depth of 16.2 μK. For atoms cooled to 1.5 μK, as shown in Sec. 3.3.3, the energy added as the blue detuned ODT is turned on is not enough to eject the atom from the trap. As such, the blue detuned ODTs can be expected to trap Rydberg atoms at $n = 80$, if there is no heating as atoms are loaded into the traps. This is not true for $n = 100$ atoms, where the calculated potential at

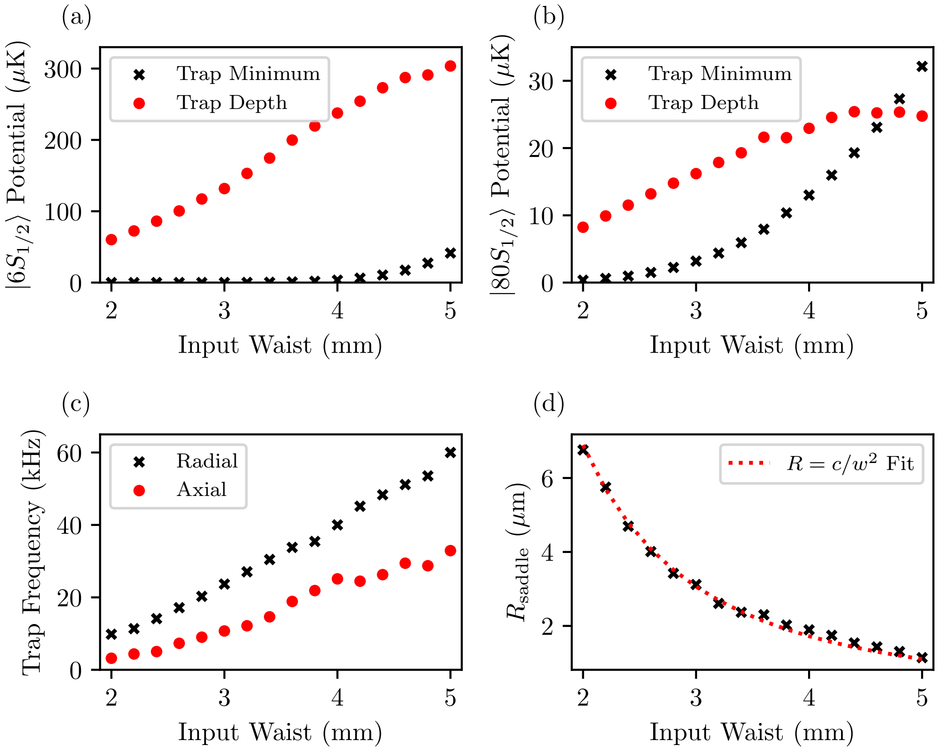

Figure 6.5: Plots showing how trap characteristics change with input waist. Figures (a) and (b) show trap energies at ground state dipole potential and Rydberg ponderomotive potential respectively. Figure (c) plots trap frequencies, calculated from fitting a quadratic function to the bottom of the trap. Figure (d) plots the distance from the trap minimum to the saddle points R as a function of input waist. A $R = 1/w^2$ line is included as a guide for the eye. A total power of 5 mW in the trapping plane is used for all calculations.

the minimum is close to the trap depth, while at $n = 120$, the trap depth approaches zero.

The measured potential, and thus the trap minimum and depth, scales linearly with intensity, and thus input power. Trap frequency is based on the second derivative of power at the minimum and so scales with the square of input power. As such, these values, calculated at 5 mW of input power, can be scaled appropriately.

Scaling With Input Waist

Keeping all other parameters identical, the traps can be scaled by adjusting the input waist w. A smaller input waist scales to a larger trap in the Fourier plane. These larger traps have lower peak intensity and looser trapping, leading to a reduced trap depth and frequency, in exchange for trapping higher-n Rydberg atoms. This is seen in 6.5(a) and (b). The trap minimum at the ground state rises above 1 μK for $w = 3.6$ mm. Given that the SLM is 12.8 mm wide on its longest axis, the distance from the centre of the beam is within two beam waists of the edge. As such, a small but signficant proportion of the input field is not diffracted from the hologram and the destructive interference condition is no longer met, and so the intensity at the centre of the trap increases. This can be corrected for by scaling r_π.

This remaining intensity at the minimum is not the primary cause of an increase of the minimum Rydberg potential. For an $n = 80$ Rydberg atom, a 5 μm input waist leads to a trap minimum comparable to the trap depth, indicating that a Rydberg atom not at the centre of the trap would experience strong heating relative to free-fall and fail to be trapped. An increase in the trap minimum occurs both due to incomplete destructive interference in the centre of the trap and the trap being small enough that the Rydberg atom begins to be affected strongly by the sides of the potential.

Trap frequencies for ground state atoms are used as a measure of the tightness of the trap, which scale roughly linearly with input waist in both radial and axial directions, as shown in 6.5(c). The size of the trap can be measured as the distance from the trap minimum to the saddle points, R_{saddle}. This scales as $1/w^2$, as shown in 6.5(d).

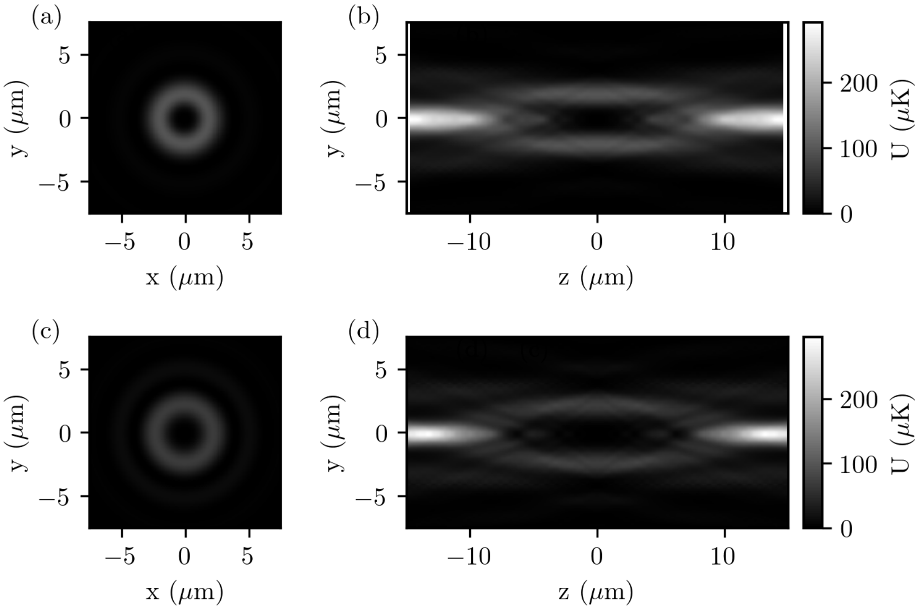

Figure 6.6: Plots showing how the trap dipole potential changes using an apertured waist, compared to an equivalent input beam waist. Figures (a) and (b) show trap potential in the $x - y$ and $x - z$ planes respectively, for an unapertured beam with input waist w_0 = 2.013 mm. Figures (c) and (d) show the same, for an apertured beam with input waist w_0 = 3 mm, w_{eff} = 2.013 mm.

Scaling With Input Beam Aperture

Changing w exactly is not always possible in an experimental configuration. As the initial conditions require a collimated beam incident on the SLM, a collimated beam can be reshaped using a combiantion of lenses in a beam scaling telescope. The beam incident on the SLM is thus limited by the combination of lenses available and the input beam waist.

To effectively reduce the input beam waist without requiring any further optics, an aperture mask is used on the SLM. By setting the phase of the SLM to zero in a radius around the input waist and scaling the π phase mask appropriately, the effective waist of the beam is smaller. The light reflecting off the zero-phase section is reflected without contibuting to the trap, thus increasing trap size at the cost of efficiency. In the experiment, the trap

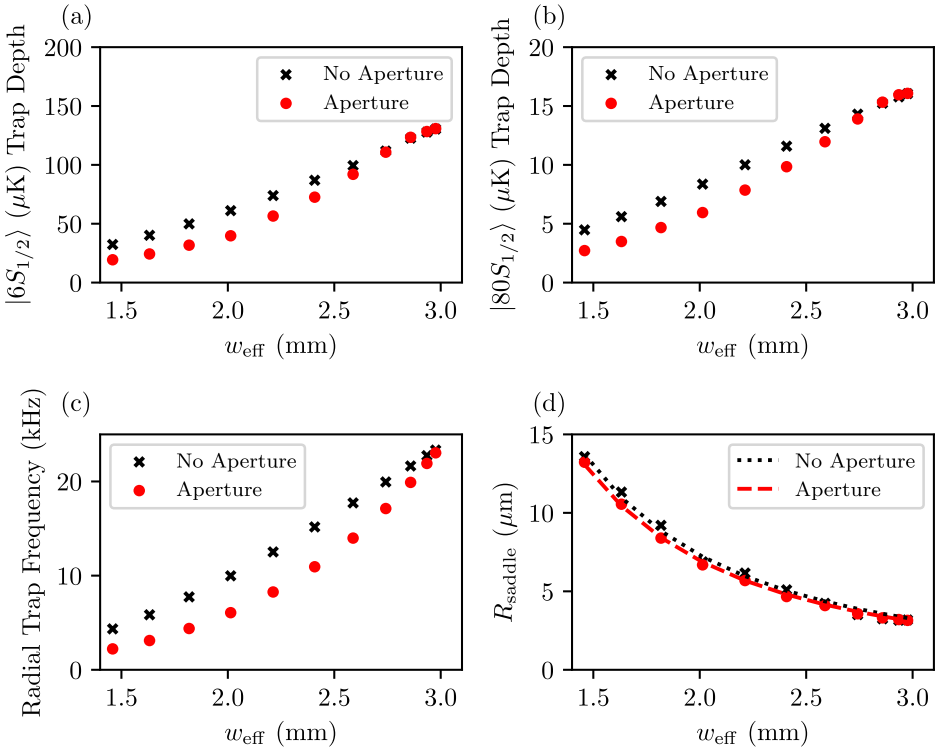

Figure 6.7: Plots showing how trap characteristics change with clipped waists. Figures (a) and (b) show trap depth at ground state dipole potential and Rydberg ponderomotive potential respectively, for apertured and non-apertured beams. Figure (c) plots radial trap frequency, while (d) plots R_{saddle}. Trap size is close for both plots, but trap depths and frequencies are lower, indicating looser potential gradients and reduced extrema. All simulations use a total power of 5 mW in the trapping plane, although mask apertures would be less efficient in utilising input light.

light uses a diffraction grating pattern to shift the light in the imaging plane away from the origin, while in simulation this light can be removed directly by setting amplitude to zero. The mask aperture adds further diffraction effects, as shown in Fig. 6.6.

For an aperture of radius r_a, the required π-phase radius r_π for maximal destructive intereference scales as

$$\begin{aligned} r_\pi &= w_0 \sqrt{\ln\left(\frac{2}{1+\exp\left(-\frac{r_a^2}{w^2}\right)}\right)} \\ &= w_{\text{eff}}\sqrt{\ln 2}, \end{aligned} \tag{6.11}$$

where

$$w_{\text{eff}} = w_0 \sqrt{1 - \frac{\ln\left(1+\exp\left(-\left(\frac{r_a}{w}\right)^2\right)\right)}{\ln(2)}}, \tag{6.12}$$

in comparison to Eq. 6.4.

Comparitive plots for apertured and unapertured waists are seen in Fig. 6.7. These plot trap parmeters using an input waist $w_0 = 3$ mm with an aperture to provide w_{eff}, compared to simply using $w_{\text{eff}} = w_0$ with no aperture. Power is rescaled such that the power at the atoms is both 5 mW, although the aperture technique does reduce the total efficiency of the system. This shows that trap depths are consistently lower for apertured beams, despite the size of the traps being consistent. Radial trap frequencies are also greatly reduced, indicating a looser trap potential. As such, adding an aperture induces a loss in efficiency twice, both by reducing trap depth and reducing the power at the atoms relative to input power.

6.3.2 Blue Detuned ODT Arrays

The dynamics of a single trap differ from those generated in an array. As the awGS algorithm targets only the intensity in the Fourier plane, the focussed Gaussian spots in the target plane will be of different phases. When these spots are expanded into the traps, the traps can interfere inhomogenously. As such, the traps become asymmetric, as seen in Fig. 6.8.

Figure 6.9 is the intensity profile traps in a 5×5 array of blue detuned ODTs, spaced by 8 μm, with $1/e^2$ input waist of 3 mm. In this case, trap depths

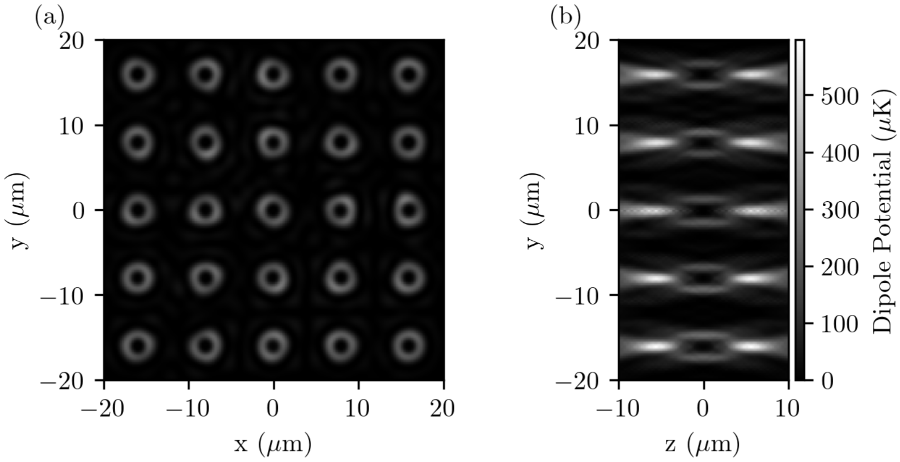

Figure 6.8: Plots showing the energy potential of an array of blue-detuned ODTs. Figure (a) is in the $x - y$ plane and (b) in the $x - z$ plane, taking a cut through the central column of traps.

average to (129.6 ± 24.7) μK across all traps for ground state atoms, and (13.8 ± 2.1) μK for Rydberg atoms. The minimum potential for Rydberg atoms varies from a minimum of 1.9 μK to a maximum of 4.8 μK. The trap size parameter R_{saddle} averages to (3.13 $\pm\, 0.23$) μm. Quoted errors here are the standard deviation in values over 25 traps. The intensity profile shown in 6.9 shows how along z axis there exists interference effects, causing rapid oscillation in trap potential of period $\lambda/2$. As such, the traps less resemble a harmonic potential.

6.4 Ground State Trapping Experiments

This section will detail the experiments performed to realise the trapping of an array of atoms into blue detuned ODTs at the ground state. As the total power of the 800 nm laser is less than that of the 1064 nm laser, and the lower trap depths generated by blue-detuned ODTs relative to red-detuned ODTs, initial loading is performed in the red-detuned ODTs and atoms are transferred to the blue-detuned ODTs.

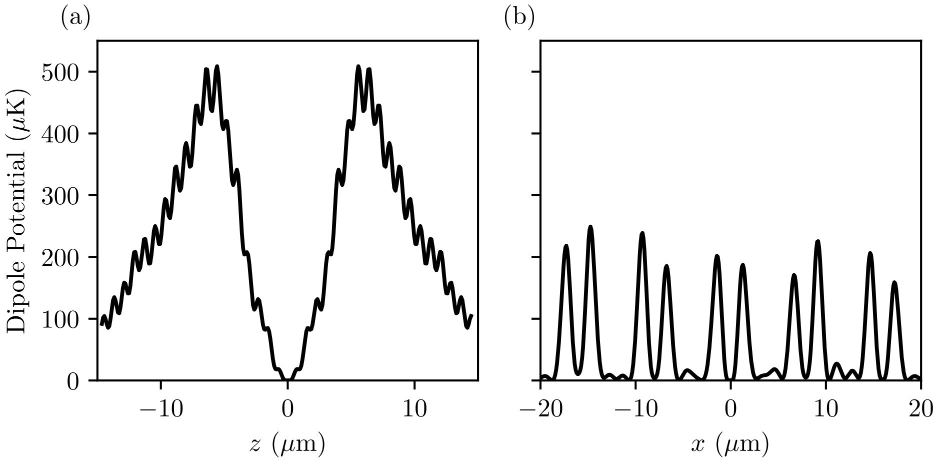

Figure 6.9: Plots showing ground state dipole potential in an array of traps. Figure (a) plots the axial profile, with figure (b) along the x-axis. In this case, multiple traps can be seen.

For these experiments, a 7×7 array of red-detuned ODTs was used at 20 μm spacing, such that Rydberg experiments at $n = 80$ will have no atoms within the blockade radius, as seen in Fig. 6.10(a). Traps are indexed from 0 to 48 based on their position in the red-detuned ODT array. The spacing of the blue-detuned ODTs can be seen in the first imaging plane, before magnification, in Fig. 6.10(b). These traps are substantially far apart, such that the traps do not interfere, as well as avoiding Rydberg blockade effects. Transfer between traps occurs over a period t_{transfer}, typically 50 μs. The full experiment timeline is seen in Fig. 6.10(c).

6.4.1 SLM Alignment Methods

The first stage of experimentation for generating effective traps is to have the generated light fields be at the atoms. Given the scale of magnification in this experiment and the small scale of the structured light, initial alignment with mirrors into the chamber can only give limited results. For fine alignment, more reliable and measurable small adjustments are needed. Also required is information regarding the centre of the input beam and the size and shape

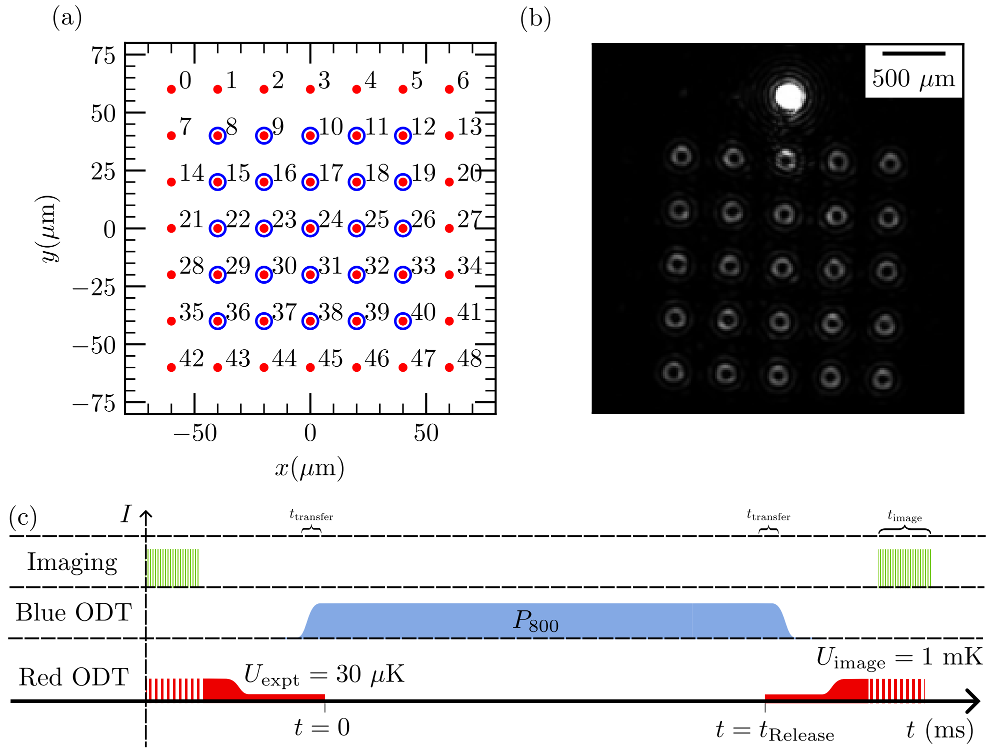

Figure 6.10: Arrangement of ODTs used in this section. Red detuned arrays form a 7 × 7 grid, with a 5 × 5 grid of blue-detuned traps overlaid. Fig (a) is the expected trap positions with indices, while (b) is an image of the blue-detuned ODTs in the first imaging plane, before 19.5× demagnification into the chamber. The large bright spot is the zeroth order of the hologram. Fig. (c) is the experimental timeline.

of such. This subsection will detail how to use the SLM itself to correct for beam alignment.

Zernike Polynomial Corrections

For the purpose, phase modifications based on Zernike polynomials were used to modify the input hologram. The Zernike polynomials are a set of continuous and orthogonal polynomials over a unit circle that well describe aberrations [200]. The use of these polynomials in correcting for beam aberrations is a demonstrated technique in holography[201–203], although in this case it is to correct for alignment without information about the aberrations.

Zernike polynomials are defined by $Z_n^m(r, \theta)$ over the unit circle $0 \leq r \leq 1$.

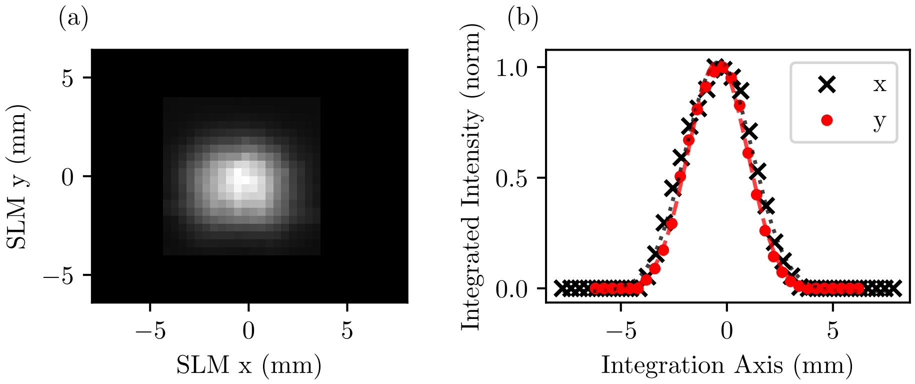

Figure 6.11: Input beam profile using the SLM as an analysis tool. Figure (a) plots the relative proportion of diffracted light from the SLM, where each region plotted in the colour plot is the size of the generated diffraction grating. Figure (b) plots the sum of these intensities across an axis, with fits to a Gaussian waist included, where $w_x = 3.18(6)$ mm, $w_y = 2.81(4)$ mm

This unit circle is recentered and rescaled over the input beam. Of the polynomials, the Z_0^0 order polynomial applies only a global phase shift to the beam, and so is discarded. The two first order polynomials $Z_1^{\pm 1}$ apply a shift along an orthogonal Cartesian axis in the focal plane, while the Z_2^0 polynomial acts as a shift in the depth of the focal plane. These are the most critical polynomials for initial alignment, as they correct for small misalignments through mirrors or lens position. The remaining polynomials correspond to a set of aberrations to be applied or corrected for, in order to maximise trap effectiveness. The experimental was capable of alignment up to the full range of order $n = 4$ polynomials. Polynomials only up to the $n = 3, m$ order are used however, due to insenstivity to higher orders. The form of the polynomials and alignment data can be seen in Chapter A.

SLM Beam Profiling

The Zernike polynomials must be applied to the SLM centered on the input beam. This also applies to the π phase mask and the aperature radius used. For this purpose, an SLM beam profiling process was used.

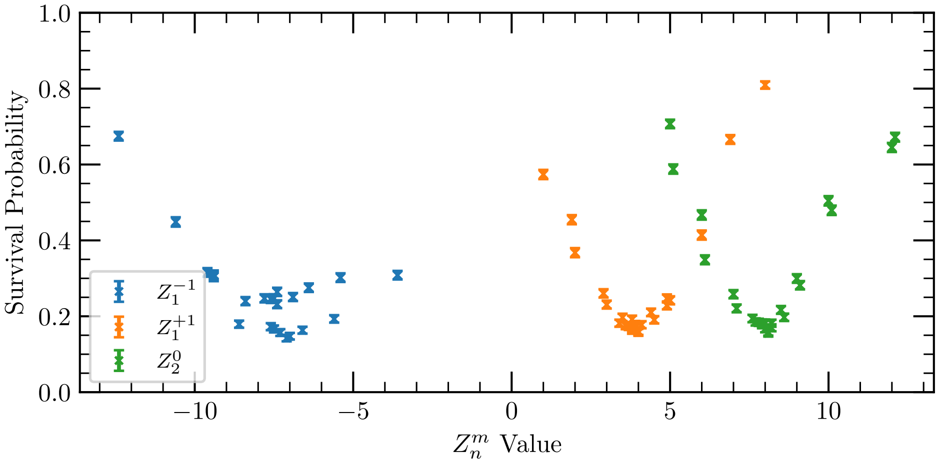

Figure 6.12: Initial Zernike polynomial alignment process, where trapped atoms are blown away by applied 800 nm light. Each data set corresponds to the Zernike polynomial coefficient.

Using the method first described in [204], the SLM is set to a constant phase across the screen, except for small regions that display a diffraction-grating style pattern. When the resulting field after the lens is reimaged, the power diffracted from grating region is proportional to the power incident on that region of the SLM. Undiffracted light is blocked, and the power of the diffracted regions is measured by imaging on a camera[3].

As shown in Fig. 6.11, this was used to measure an input beam of $1/e^2$ waist $w_x = 3.18(6)$ mm, $w_y = 2.81(4)$ mm, with a beam centre at $x = -0.39(2)$ mm, $y = -0.44(2)$ mm relative to the centre of the SLM. The input beam waist is therefore not symmetric as is ideal, but as $\sqrt{w_x w_y} = 2.99(3)$ mm, this closely approximates the previously introduced simulations. This adjusted centre is used for the Zernike corrections, the π-phase mask and the pseudo-aperture.

6.4.2 Blow-Away Alignment

Initial alignment is performed through using the SLM to remove atoms from the traps with minimum power. By not applying a π phase mask and thus generating tightly focussed Gaussian beams, the 800 nm light potential removes atoms from the 1064 nm red-detuned ODTs. Maximising the destructive potential at minimum power is achieved with optimal alignment, when the focal plane of the 800 nm light matches that of the red-detuned ODTs across the entire array, and each point on the 800 nm array generated matches a corresponding trap.

As the most critical to initial alignment, only the $n = 1$ and $n = 2, m = 0$ Zernike polynomials are used for this initial alignment. Results are seen in Fig. 6.12, for a single round of optimisation.

6.4.3 Trapping In Blue-Detuned ODTs

Once initial alignment is completed, the π-phase mask is applied. The experimental sequence used must maximise transfer from the red-detuned ODTs to the blue-detuned ODTs and hold them. This hold period should be long enough that any atoms remaining must have been trapped in the blue-detuned ODTs. The atoms must then be transferred back to the red-detuned ODTs for imaging. A period of $t_{\mathrm{Release}} = 100$ ms is used for this (see Fig. 6.10 for experiment timeline).

As shown in Fig. 6.13, initial trapping populations of 55(23)% was seen, which improved to 89(6)% after optimisation with Zernike polynomials up to and including $n = 3$. The range of iterative scan data is shown in A. Measurements of trapping rate do not correct for loss in imaging or through limited trap lifetime, and does not discriminate between loss transferring into the blue trap and loss transferring out of the blue trap.

---

[3] Thorlabs Zelux CS165MU1

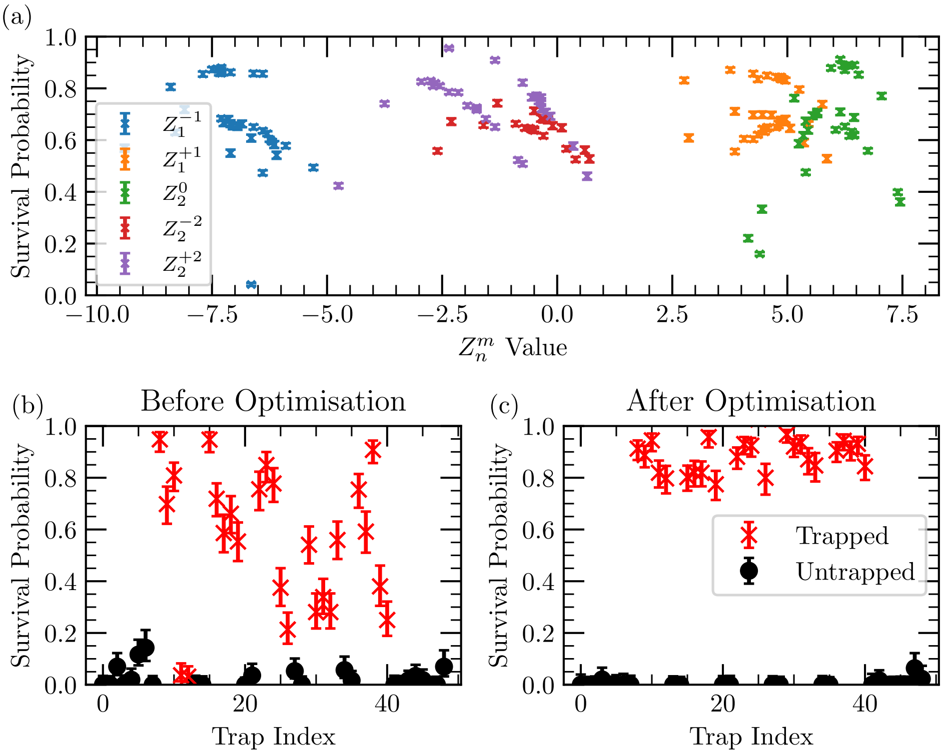

Figure 6.13: Demonstration of optimising blue trapping performance with Zernike polynomial scans. Fig. (a) shows results of scanning selected Zernike polynomials of $n = 1$ and $n = 2$ to find maximal trapping probability. Further scans were performed up to $n = 3$, but are not shown for clarity. These can be seen explicitly in the Appendix. Fig. (b) is trap performance after maximising trap destruction but before Zernike polynomial scans, and (c) is after 7 rounds of optimisation, averaging 89(6)% survival. The trap index refers to labelling of traps based on the underlying array (see Fig. 6.10).

Trapping performance was maximised at a pseudo-aperture radius of 2.0 mm, for an effective waist of $w_{\text{eff}} = 1.60(2)$ mm. With this radius in place, the maximum power used was 9.7 mW per trap, although trapping performance was maintained at reduced powers down to 6.0 mW, as shown in Fig. 6.14.

For this maximum power, the minimum trap potential is expected to be $0.004, 0.047\ \mu$K for ground state and $n = 80$ Rydberg atoms respectively. The trap depths for the same are predicted to be $47.2, 6.77\ \mu$K, while ground state trap frequencies are $\omega_r/2\pi = 4.33$ kHz$, \omega_z/2\pi = 1.00$ kHz.

Also shown is trapping performance as r_π varies, which is maximised at the

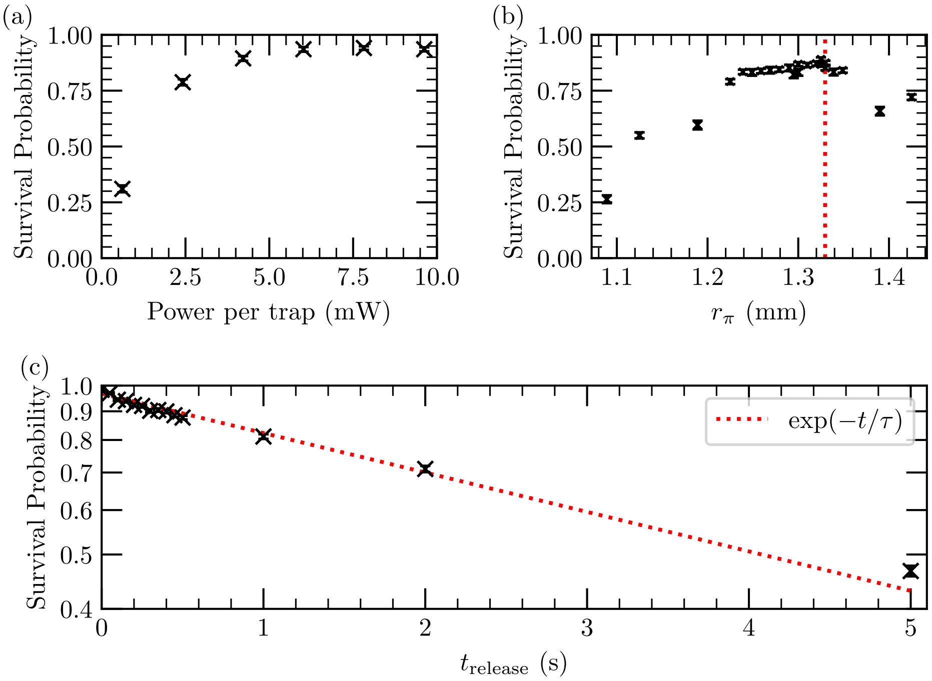

Figure 6.14: Trapping performance as power per trap, r_π and $t_{release}$ are adjusted, in (a), (b) and (c) respectively. The red dotted line in (b) is the predicted ideal value for r_π for $w_0 = 3.0$ mm, $r_a = 2.0$ mm. The line in (c) is a fit to an exponential decay, demonstrating a lifetime of 6.5(4) s. The y-axis is set to a log scale.

theoretical optimum for the input waist and aperture radius used.

Atoms held in the trap are subject to the same losses due to collisions with background atoms and trap heating. Data shown in Fig. 6.14(c) shows an average lifetime of $\tau_{800} = 6.2(4)$ s across the array. From this, the lifetime at zero hold time can be taken as the total transfer efficiency, plus losses in imaging, for a transfer efficiency of 96.8(5)%.

This measured lifetime is lower than the $\tau_{1064} = 12(1)$ s lifetime of the red-detuned ODTs during this experiment, but as the atoms spend substantially less time in the blue-detuned ODTs, this has a reduced impact to experimental efficiency.

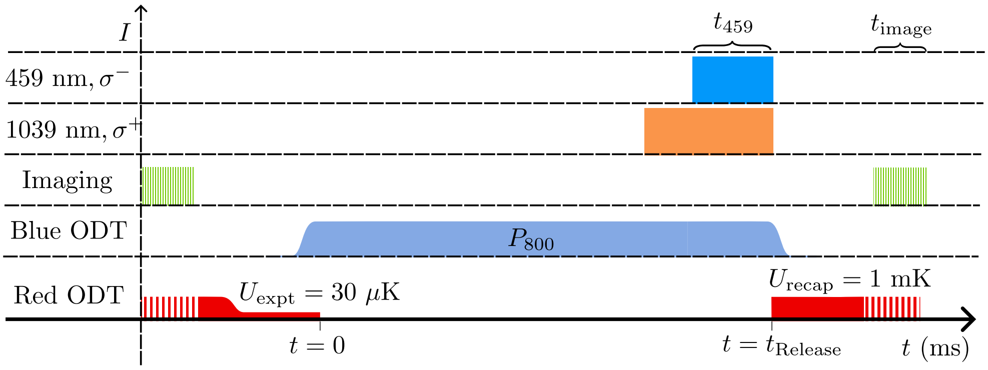

Figure 6.15: Timeline for Rydberg excitation in blue-detuned ODT experiments. Timeline is not to scale. . The 1039 nm light is on for 8 μs consistently to reduce warm-up effects and to maintain parity with experiments in free-fall (see Chapter 5).

6.5 Trapping Rydberg Arrays

With arrays of atoms now trapped in blue-detuned ODTs, experimentation can take place to trap Rydberg atoms. The experiment timeline is similar to that used for typical Rydberg experiments seen in Chapter 5, although with the addition of the 800 nm light to trap atoms. To smoothly transfer from blue- to red-detuned ODTs, a reduced recapture power for red-detuned traps is used. This reduced power is less effective at ejecting Rybderg atoms, and so contrast is reduced. The timeline for these experiments is seen in Fig. 6.15.

6.5.1 Rydberg Excitation In Blue Traps

Experiments involving the Rydberg sequence were performed in identical conditions, where atoms were trapped in blue traps for 300 μs. For a control experiment, the blue-detuned ODTs were released as the atoms undergo Rydberg excitation in free fall. This free-fall period is for 8 μs (as used through Chapter 5), and are recaptured in red-detuned ODTs afterwards, and held for the remaining period to match the blue-detuned trap experiment duration. This is compared to the same traps being held in blue-detuned ODTs

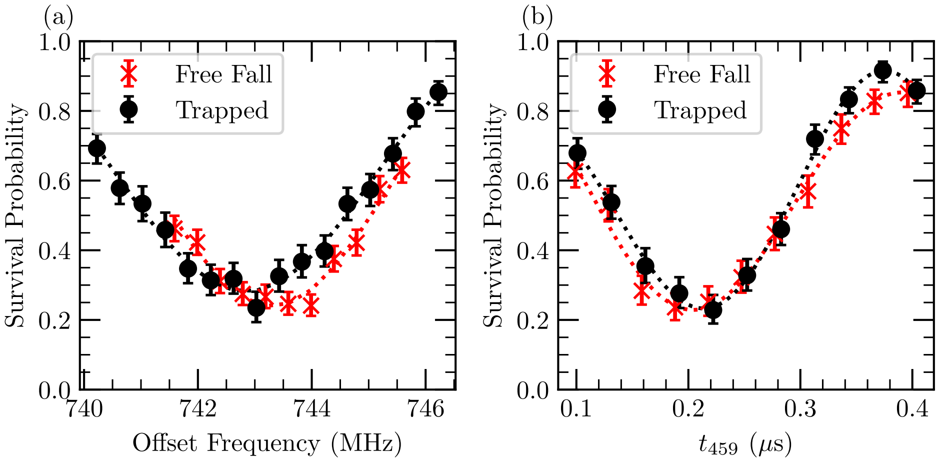

Figure 6.16: Rydberg excitation in blue-detuned ODTs. Fig. (a) is the data for spectroscopy exciting to the $|80S_{1/2}\rangle$ state, including fits to a Gaussian profile. The $x-$axis refers to the offset of the 1039 nm frequency from the ULE cavity carrier. Trapped atoms experience a resonance frequency shift of 0.45(7) MHz. Fig. (b) is a Rabi oscillation at the same offset frequency, including fits to a predicted Rabi oscillation. Fits to trapped atoms show a Rabi frequency of $\Omega_R/2\pi = 3.09(8)$ MHz, and 2.7(2) MHz for atoms in free-fall, where both oscillations are at the resonance frequency for atoms in free-fall.

and recaptured in red-detuned ODTs at the end of the sequence. As such, any difference in atom survival is due to the presence of the blue traps during Rydberg excitation.

The data can be compared in Fig 6.16, which shows the effects of Rabi oscillations in Rydberg light on atoms in the ground state in both blue-detuned ODTs and in free-fall. Atoms within blue-detuned traps have a resonance frequency shift of 0.45(7) MHz, caused by the differential AC Stark shift on the Rydberg transition by the blue-detuned ODTs.

Rabi oscillations at the free-fall resonant frequency demonstrate a Rabi frequency of 2.7(2) MHz in free-fall and 3.09(8) MHz in blue-traps. The faster driving is caused by off-resonant driving, further confirming the presence of an AC shift caused by the blue-detuned ODTs.

6.5.2 Rydberg Decay In Blue Traps

To determine if the blue traps could trap Rydberg atoms beyond a few periods of Rabi oscillations, the decay rate of the Rydberg state was measured. For this experiment, atoms are excited to the $\left|75S_{1/2}\right\rangle$ state with Rabi frequency $\Omega/2\pi = 2.69(1)$ MHz. The decay period for this oscillation was measured at $\tau = 2.5(4)$ μs. After a π pulse, atoms are held in blue traps for t_{hold}, and recaptured in red-detuned ODTs. Remaining Rydberg atoms are ejected, while atoms that have decayed to the $\left|6S_{1/2}\right\rangle$ ground state are recaptured.

Results for this experiment are seen in Fig. 6.17. The minimum hold time used was 10 μs, which is two orders of magnitude greater than the π time, but the recapture population remains the same, within error. Recapture probability shows the expected $1/e$ decay curve, with decay period $\tau = 0.33(3)$ms. Recapture probability is limited to a maximum of 0.46, as a proportion of atoms decay to anti-trapped states in blue traps.

This decay period is double the expected 0.164 ms lifetime for $n = 75$ at room temperature, using data from [63]. This is believed to be due to how the Rydberg state decays; some decay paths are into other long-lived Rydberg states or states that are ejected from the 1064 nm states. As such, recapture probability does not necessarily map directly to population in the prepared Rydberg state. This data does indeed show that Rydberg atoms can be held in blue-detuned ODTs for long periods of time, comparable to the Rydberg state lifetime.

6.6 Summary

Through this chapter, the process to generate arrays of blue-detuned ODTs using an SLM has been shown. At the atoms, 9.4 mW of light is used per trap, and can hold atoms with average lifetime $\tau_{800} = 6.2(4)$ s. Optimisation of atomic lifetime was achieved by implementation of Zernike polynomial

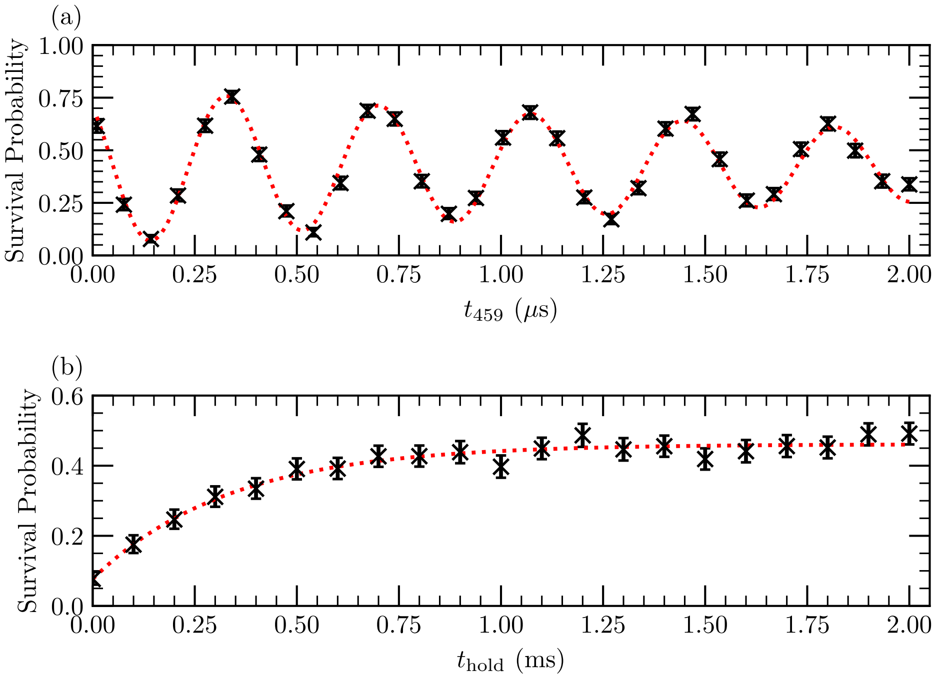

Figure 6.17: Rydberg decay experiments in blue-detuned ODTs. Figure (a) shows coherent Rabi oscillations in blue-detuned ODTs, with minimum recapture probability at maximum population in the Rydberg state. The red-dashed line is a fit to the expected Rabi oscillation, with $\Omega/2\pi = 2.69(1)$ MHz and decay period $\tau = 2.5(4)$ μs. Figure (b) is an experiment where atoms are prepared in this maximum population and held in the blue-detuned ODTs for a period t_{hold}. Fits are a $1 - \exp(t/\tau)$ curve, constrained to match the initial minimum probability after a π pulse. This decay period is measured to be $\tau = 0.33(3)$ms.

phase shifts on the hologram for small adjustments, which would otherwise be unfeasible for manual alignment. These same traps are also capable of trapping Rydberg atoms, including for coherent Rabi oscillations and holding in place for long lifetimes. Some 800 nm intensity remains on the atoms, and produces an AC Stark shift on the Rydberg transition. The small AC Stark shift seen on the Rydberg transition is used in Chapter 7 to implement analogue Rydberg quantum computing.

To enlarge the traps to better suit trapping of Rybderg atoms and increase survival, a zero-phase pseudo-aperture was used. This is a phase pattern on the SLM that was simulated to increase the size of the features in the Fourier plane, at the expense of efficiency. Simulations also show that diffraction effects from limited beam aperture also occur. Another cause of inhomogeneity is interference between traps, which causes asymmetry between and within individual traps.

Further experimental upgrades are planned to adjust the relay telescope for greater magnification into the chamber. The current system produces traps too small for effective trapping without reducing the input waist (and thus reducing the SLM effectiveness) and also adding the pseudo-aperture. Adjusting the relay telescope would allow for larger traps with improved efficiency and reduced diffractive effects. These diffractive effects reduced trap efficiency substantially when within the blockade radius of 10.5 μm. With these improvements, dipole-dipole interactions between Rydberg atoms while remaining trapped can be shown.

These experiments led to a greater understanding of both the experimental procedure required for alignment of the blue traps, and the iterative processes needed for beam shaping based on feedback from interactions with the atoms.

Chapter 7

Local Rydberg Control For Graph Optimisation

7.1 Introduction

The SQuAre experiment is designed as a scalable quantum computing platform. The previous chapters have detailed the construction of various elements used for such a platform, but have not yet shown quantum computation, rather elements to be used in quantum computation.

This chapter will show how the elements introduced in previous chapters can be used or modified to be used in solving graph optimisation problems. The particular problems solved here are first discussed in Sec. 1.3, which are the Maximum Independent Set (MIS) and Maximum Weighted Independent Set (MWIS) problems. Experimental details of solving both the MIS and MWIS problems on the SQuAre platform are shown in Sec. 7.2 and Sec. 7.3. This can potentially scale to larger graphs, as shown in Sec. 7.4.

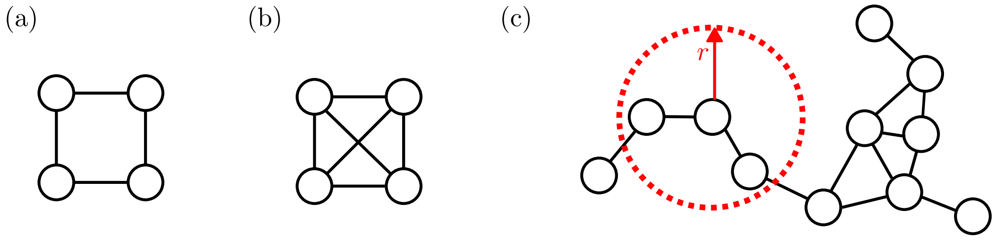

Figure 7.1: Illustration of some basic graphs. Figures (a) and (b) both contain four vertices, with different connecting edges. (b) is the maximally connected four-atom graph. (c) is an example unit disc graph. Vertices separated by $\leq r$ are connected by an edge, as shown in a red dashed circle around one vertex.

7.1.1 Defining Graphs

Graphs are defined by a set of vertices and a set of edges. Each edge starts at a vertex and ends at a vertex, defining a subset of the vertex set. Vertices are connected if an edge set exists that contains both. Following the definition laid out in [205], this thesis is only concerned with graphs where edges are undirected, duplicate edges are ignored, and an edge cannot contain the same vertex twice. An edge defines the connection between two vertices.

In the abstract sense, vertices and edges are simply defined by two sets containing elements, but each element of the set can contain properties as well. By granting each vertex or edge a weight, subsets that would otherwise look identical can change. Vertices can also be defined by position.

Of particular use in this chapter will be the unit-disk graph (UDG). This is the graph defined by the set of vertices where all vertices within a unit distance of each other are connected by an edge. Vertices outside this unit distance do not have a connection. Graphs of this class are useful for use due to the simple mapping to Rydberg atom systems; atoms are the vertices, while the unit distance is the Rydberg blockade radius. Example graphs are shown in Fig. 7.1.

7.1.2 Graph Problems

Graph problems typically take the form of finding a subset of edges or vertices that meet a set of conditions. An entire field of mathematics is based on the study of these graph problems and so this chapter will not go into detail on the specifics of the field. Of more immediate interest is methods of computing solutions to given problems.

For example, a well-known and well-studied graph problem is the Travelling Salesman problem. Given a set of towns, where each town is connected by a known distance, in which order should one visit each town such that the total distance is minimised, ending at the starting point? Reducing the statement to a graph problem, towns are vertices, and distances are weighted edges. Which selection of edges connects every vertex while minimising the total weight? This problem is well-studied in part because, firstly, it maps closely to many decision problems common today, from stock-picking in warehouses to organising bus routes, and secondly, finding exact solutions is known to be NP-complete [206].

Since the development of Shor's algorithm in 1994, some problems have been shown to be scale better in quantum systems in comparison to the best known classical algorithms. Shor's algorithm provides a quantum speedup in integer factorisation that could dramatically change modern cryptography and cryptanalysis [207], while Grover's quantum search algorithm is known to be optimal [208]. Solving difficult graph problems is an attractive proposal for quantum computers, given that they are known to be classically hard while also directly impactful in applications outside the laboratory.

7.1.3 Independent Sets

Of immediate interest to classical quantum computing and this thesis are problems involving finding independent sets. These are the subsets of the

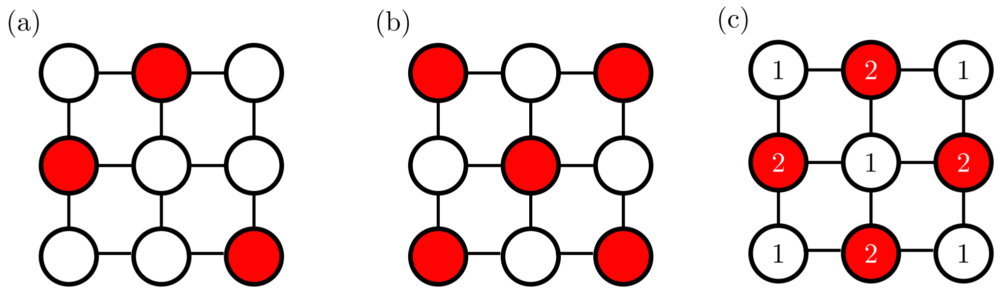

Figure 7.2: Examples of maximal independent sets for the same graph. In red are selected vertices that make up a maximal independent set. Figure (a) is maximal, but not the maximum. Figure (b) is the MIS for this graph. For each vertex in (c), a weight is specified. The selected vertices constitute the MWIS.

vertices of a graph where no selected vertices are connected [209]. These graphs are defined as maximal independent sets when no further vertices can be added to the subset without violating the adjacency condition. For example maximal independent sets, see Fig. 7.2.

The largest possible maximal set for a given graph is the maximum independent set (MIS)[1]. Finding the maximum independent set of a graph is known to be an NP-hard optimisation problem [209], although approximate solutions on UDGs can be found in polynomial time [210]. Experimental data of Rydberg atom optimisation on a king's graph of up to 289 atoms has been investigated, with some speedup for the hardest graphs compared to the experimentors' tensor network code [84]. However, these graphs have been shown to be soluble for other classical annealing algorithms [211].

Should each vertex be given a weight, maximal independent sets can be scored on the sum of the weights of the constituent vertices. The maximal weighted independent set (MWIS) is the maximal independent set with this highest score for a given graph and weighting. This problem is NP-hard not just for finding exact solutions but also for finding approximate solutions for some classes of graphs [212, 213]. For UDGs, this approximation is harder

---

[1] The MIS is not necessarily unique for a graph; degeneracy introduces further complexity to the problem.

than for the MIS case, but remains polynomial [214].

Solving for the MIS or MWIS of arbitrary graphs is a useful problem that is difficult to solve classically for exact solutions. Neutral atom quantum computing is naturally suited to these problems, and are a promising avenue for demonstration of a real-world quantum speedup. Sections 7.2 and 7.3 will show how the SQuAre platform can be used to solve both.

7.2 Finding The Maximum Independent Set

7.2.1 Quantum Annealing

Cost Functions

For some UDG with a set of vertices V containing points $\{r_i\}$ on a graph with edges E defined by $\{r_i, r_j\}\ \forall\ |r_i - r_j| < r_d$, where r_d is the unit-disc radius, a selection of vertices can be made. For each point r_i on the graph, $n_i = 1$ if selected, and $n_i = 0$ if not, making a set $N \equiv \{n_i\}$. There exists also some adjacency matrix $A_{i,j}$ such that $A_{i,j} = 1$ if $\{r_i, r_j\} \in E$, $A_{i,j} = 0$ if $\{r_i, r_j\} \notin E$.

This selection can be scored as H_{MIS}, which is defined by

$$H_{\text{MIS}} = -\sum_i n_i + \sum_{i,j,i<j} U A_{i,j} n_i n_j. \tag{7.1}$$

This is the maximum independent set cost function, where U is an arbitrarily large cost penalty for selecting adjacent vertices. If $U \gg 1$, this score is minimised for the selection N that defines the maximum independent set (MIS).

For a quantum computing experiment, we must find a Hamiltonian that maps well onto this cost function. If this mapping holds, the ground state of the Hamiltonian is also the minimum cost for the cost function. Finding the MIS is then a process of preparing this ground state.

Rydberg Hamiltonian

Compare now the Rydberg Hamiltonian H_{Ryd}. Atoms are prepared in an array, defined by positions $\{r_i\}$, where the unit radius is defined by the Rydberg dipole-dipole potential V_{dd}, matching the UDG defined above. The Hamiltonian is given by [75]

$$H_{\mathrm{Ryd}}/\hbar = -\sum_i \left(\frac{\Omega_{g,r}}{2}\sigma_i^x - \delta_i \hat{n}_i \right) + \sum_{i,j,i<j} V_{\mathrm{dd}}\left(|\mathbf{r}_i - \mathbf{r}_j|\right) \hat{n}_i \hat{n}_j \tag{7.2}$$

where $\Omega_{g,r}$ is the Rabi frequency addressing the Rydberg transition $\left|6S_{1/2}, F = 4, m_F = 0\right\rangle \to \left|80S_{1/2}\right\rangle$, σ_i^x is the Pauli x-operator, $\hat{n}_i = |r_i\rangle\langle r_i|$, δ_i is the local detuning on the ith atom, and $V_{\mathrm{dd}}(R)$ is given by Eq. 2.37. This is the generic Rydberg Hamiltonian, which we can map onto the classical cost function. In this case, the detuning term $\delta_i \hat{n}_i$ in Eq. 7.2 maps to the selection term n_i in Eq. 7.1, while V_{dd} maps to the adjacency energy penalty U. For this mapping to hold, $V_{\mathrm{dd}} \gg \delta_i \forall i$, in the same way $U \gg 1$. For experiments in finding the MIS, we measure $\hat{n}_i$ by finding the probability of occupying the ground state. As this experiment takes place in free-fall, the red-detuned ODTs are used to eject Rydberg atoms at the end of the process. Missing atoms are thus considered to be Rydberg atoms, and remaining atoms are considered to be ground state atoms. As such, $\hat{n}_i |g\rangle = 0$, $\hat{n}_i |r\rangle = |r\rangle$.

If atoms are prepared in the $|g\rangle$ state, when $\delta_i \ll 0$, oscillations are far off resonance and the system is in or near to the minimum energy eigenstate of H_{Ryd}. If $V_{\mathrm{dd}} > \delta_i > 0$, such that δ_i is again off-resonance, the minimum energy eigenstate of H_{Ryd} is the selection of atoms that represent the MIS.

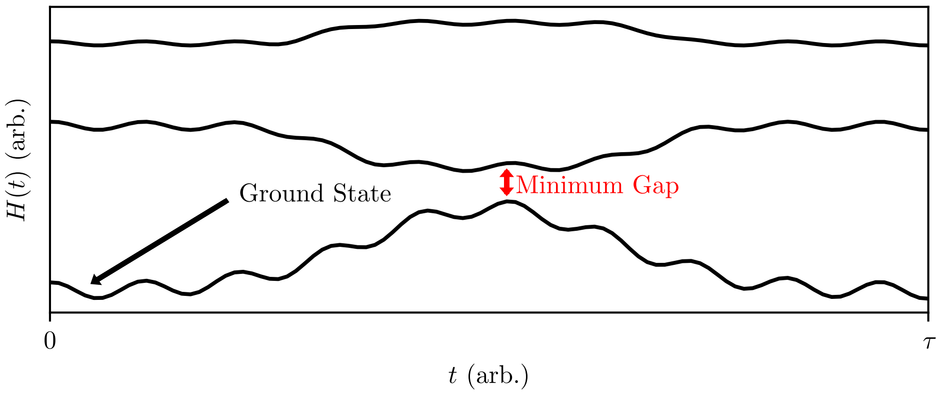

Figure 7.3: Example energy level diagram for a quantum annealing process. $H(t)$ is the energy of some eigenstate (shown in black lines) over time. This can be the classical cost function as the input state is varied, or the eigenenergies of a varying Hamiltonian. By slowly varying H(t), a system prepared in the ground state remains in the ground state, where its probability of moving to the next excited state depends on the energy gap and rate of variation.

Quantum Annealing

By slowly varying the detuning value δ_i from negative to positive values over time, the state can be annealed into the MIS state. A diagram of this process is shown in Fig. 7.3, where the Hamiltonian is expressed as a function of time. This can be compared to the classical cost function picture. For the MIS case, the cost starts at 0, and is slowly reduced to the minimum possible for that graph. The computation process then becomes a recurring iterative procedure.

This holds true for slow variations of any Hamiltonian by the adiabatic theorem, and has been long considered as a quantum computing method with potential to solve classically NP-hard problems with a quantum speedup [215, 216]. This is known as a quantum annealing algorithm (QAA). In general, after state preparation, the Rydberg Hamiltonian is applied, where the detuning is varied from far-negative to far-positive detuned. This variation is defined by some parameter set. For each parameter set, the classical cost function is measured, and the parameter set varied. A closed-loop optimisation process can measure how the cost function varies with the parameter set to find the optimal parameter set. The parameter set that produces the lowest classical cost is then used to determine the lowest cost result. This is the ground state of the Rydberg Hamiltonian and the lowest cost state measured, and as such the solution to our input problem.

7.2.2 Optimising The Annealing Profile

Quantum annealing is limited by the timescale of the system. If $\dot{\delta}_i$ is too large, the state is excited beyond the ground state and the MIS state is not prepared. If the total time is too large, then ramp times exceed the coherence time of the excitation. As such, $\delta(t)$ is crucial to define for this experiment.

To define the detuning ramps, an AWG is used to modulate the 459 nm global beam. For operation on resonance, this is connected to an AOM in a double-pass configuration at central frequency 200 MHz. Detuning at the atoms is thus defined by a frequency chirp about this central frequency, such that the AOM is driven at a frequency $(200+\delta(t)/2)$ MHz. When $\delta = 0$, the beam is on resonance.

The definitions of the detuning ramps $\delta(t)$ are should be smooth, monotonically increasing and well defined for all t. The detuning ramps follow a cubic profile of the form

$$\delta(t) = at^3 + bt + c. \tag{7.3}$$

The parameters a, b, c are defined by

$$\begin{aligned} a &= 8sc/\tau^3, \qquad (7.4)\\ b &= 2c/\tau - a\tau^2/4, \\ c &= \delta_{\text{max}} - \delta_{\text{min}}, \end{aligned}$$

where s is a shape parameter, $0 \leq s \leq 1$, and τ is the total ramp duration.

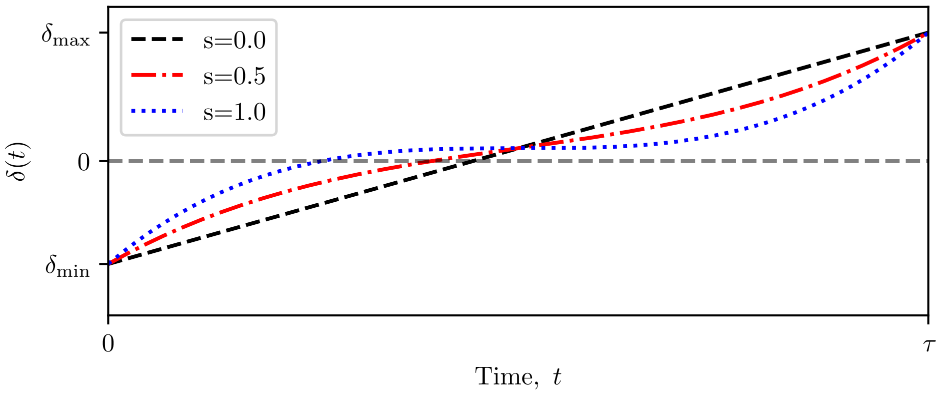

Figure 7.4: Diagram of a generic cubic detuning ramp, where $|\delta_{\max}| > |\delta_{\min}|$. Three values of s are plotted between 0 and 1. A value of 0 is a linear profile ($d\delta/dt =$ const.), while a value of 1 is a strictly monotonic cubic profile where $d\delta/dt = 0$ at some point.

This defines a cubic ramp starting at $\delta_{\min}$ and ending at $\delta_{\max}$. A diagram of this ramp is seen in Fig. 7.4, with varying shape parameter. When $s = 0$, the rate of change of detuning $d\delta/dt$ remains constant, while when $s = 1$, $d\delta/dt = 0$ at the midpoint of the ramp. This defines the bounds of the detuning function such that the detuning profile is always monotonic, where $d\delta/dt$ is at a minimum in the centre of the ramp.

Optimising $\delta(t)$ is achieved through a hybrid quantum-classical approach. The final state is evaluated through finding the average cost $\langle H_{\text{MIS}} \rangle$, as defined in 7.1. A classical optimisation algorithm is then used to vary the parameters used to define $\delta(t)$, in order to find the minimum cost.

This approach belongs to a class of adiabatic quantum optimisation algorithms. Other similar experiments using a cubic detuning ramp have demonstrated success in generating many-body ground states in other neutral atom systems [82, 84, 163, 217].

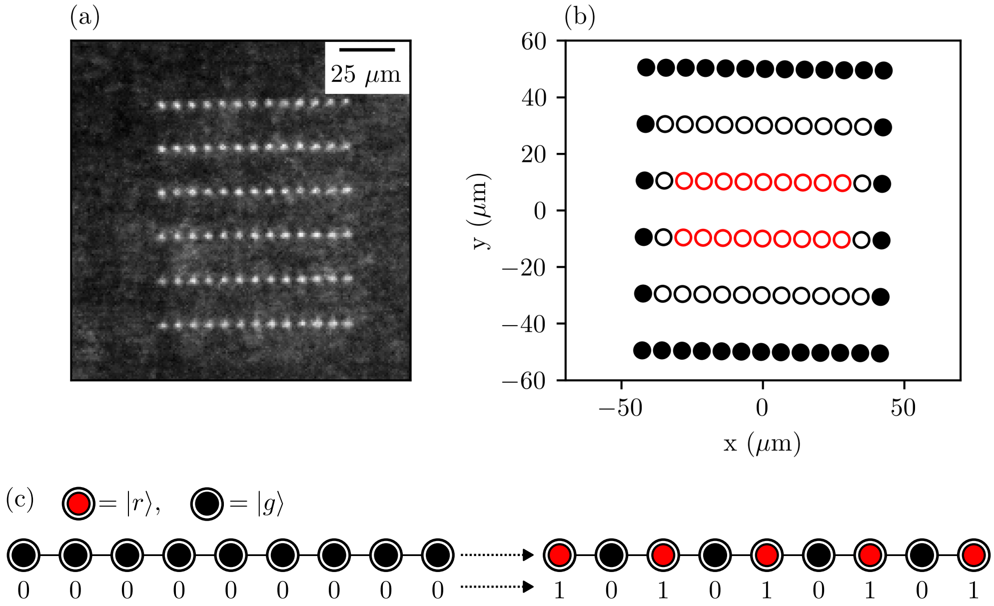

Figure 7.5: Array arrangement for 1D chain MIS state preparation. (a) is the per-pixel maximum of 400 experimental images used for calibration to show the initial trap arrangement, while (b) shows the targets to be filled. Red empty circles define the arrays where atoms to be moved to in sorting, while white circles are where atoms are to be removed. Filled black circles are the reservoir. (c) shows the initialisation of an array of ground state atoms, prepared in the MIS state.

7.2.3 Graph Construction And Labelling

To demonstrate MIS state preparation, a line of atoms is prepared. Lines are defined along the axis of constant Rabi frequency such that Ω_i is identical for all atoms. Each atom is separated by 7 μm. Atom sorting prepares each row, and data is post-selected for correct initial array preparation. The measured Rabi frequency is $\Omega/2\pi = 2.70(2)$ MHz, such that the blockade radius $r_{\rm B} = 10.4$ μm. MIS preparation is thus indicated by excitations of every other atom, forming a Z_2 ordered phase. This replicates the results seen in [75], where states up to Z_4 order were observed in 1D chains.

The array used for this experiment is shown in Fig. 7.5. Rydberg detection is determined as before in previous chapters by destructive ejection from red-detuned ODTs. Following the definitions used in Eq. 7.2, a missing atom is considered to be in the Rydberg state and thus labelled ‘1’, while a present atom is labelled ‘0’. As each atom’s position is distinguishable, each atomic position can be assigned to a position in a bit string.

7.2.4 Experimental Results

The MLOOP closed-loop optimisation package [218] was used to optimise $\delta(t)$ to maximise the probability of generating the Z_2 state. Parameters to optimise are δ_{min}, δ_{max}, τ and s. Each parameter set is evaluated by applying the ramp over 200 experimental runs and optimised to minimise $\langle H_{\text{MIS}} \rangle$, where $U = +1$. The Rabi frequency was measured to be $\Omega/2\pi = 2.75(2)$ MHz.

The lowest scoring parameter sets over a series of experiments are then verified with 1000 experimental runs. Results are seen in Fig. 7.6 for the lowest scoring run, where $\langle H_{\text{MIS}} \rangle = -3.96(3)$, where errors are statistical. This prepared the Z_2 state with probability 0.32(2) for the top row and 0.24(2) for the bottom row. This was at values $\delta_{\text{min}} = -7.6213$ MHz, $\delta_{\text{max}} = +1.00$ MHz, $\tau = 2.415$ μs and $s = 0.0689$, which is shown diagramatically in Fig. 7.6(c). These results show that the MIS state can be robustly prepared using the experimental system.

7.3 Finding The Maximum Weighted Independent Set

7.3.1 Experimental Scheme

The optimisation function of the MWIS is given by

$$H_{\text{MWIS}} = -\sum_i w_i n_i + \sum_{i,j,i<j} U_{i,j} A_{i,j} n_i n_j, \tag{7.5}$$

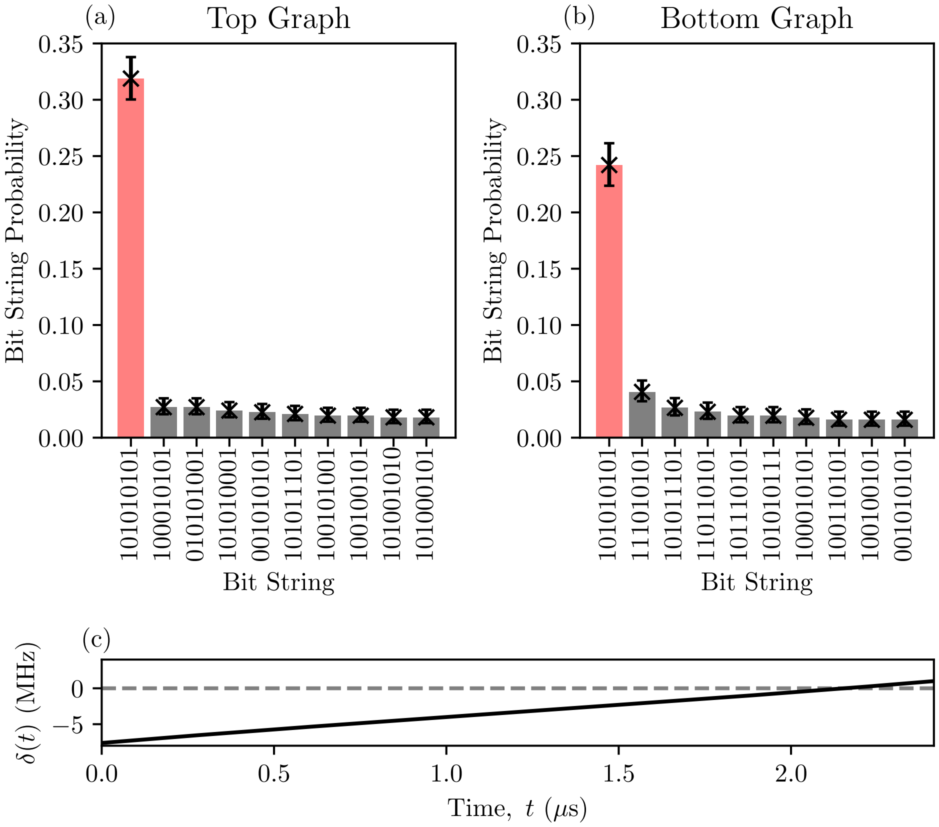

Figure 7.6: Figs (a) and (b) show the measured probabilities after MIS experiment optimisation for the top and bottom rows respectively. Only the 10 most probable strings are plotted for clarity. Highlighted with red bars is the Z_2 phase. (c) shows a diagram of the optimal ramp found for the 1 dimensional Rydberg chain state preparation

where $U_{i,j} \geq w_i w_j, w_i \geq 1 \; \forall \; i$.

Extending the scheme scene in Sec. 7.2 to solve MWIS problems requires further development of the experimental scheme. The detuning applied in this previous section $\delta(t)$ is a global detuning, controlled by applying a detuning to the global 459 nm Rydberg beam. As per Eq. 7.2, δ_i is equal for all atoms i in a global scheme. Weightings must be site-specific and therefore cannot be applied by global beams alone.

As seen in Sec. 6.5.2, a differential AC Stark shift can be applied to the Rydberg transition without applying such force to the Rydberg atoms that they are removed from the experiment outright. In this case, the AC Stark

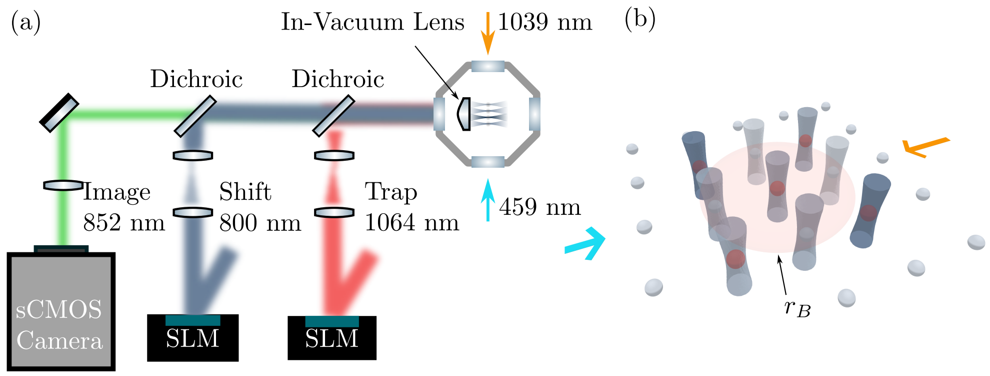

Figure 7.7: Schematics for representing the graph weighting process. Figure (a) shows the overlap of 1064 nm and 800 nm light onto atoms within the chamber. Figure (b) shows a 3D render of atoms in a square lattice. 800 nm light is applied to atoms within a square array, with intensity (shown by different colours) varying depending on weight. Atoms excited to the Rydberg state are shown in red, with the Rydberg blockade radius shown as a red circle around the central atom..

shift is applied by blue-detuned ODTs and so the AC shift is smaller than the linewidth of the transition. This subsection will detail how these blue-detuned holographic light potentials can be used to apply a differential AC Stark shift to the Rydberg transition. Furthermore, this shift can be adjusted to apply site-specific light shifts, leading to local Rydberg control. This is shown in Fig. 7.7.

One caveat of this scheme is that the differential AC Stark shift is only capable of positive shifts. To corectly anneal into the quantum ground state from a negative to positive detuning, detuning ramps using only the AC shift would thus cause sites with heavier weights to pass through resonance first. This would blockade other less weighted sites and prevent maximal entanglement. This property, of creating a highly entropic or 'magic' state, as defined in [219, 220], is necessary for a quantum speedup effect [221, 222].

In this case, the global detuning scheme is retained for negative detunings, where input frequency is chirped from $\delta = \delta_{\text{min}}$ to $\delta = 0$. At resonance, the global detuning is held constant and only then is the local light shift applied by the holographic light potential applied. This scales linearly with

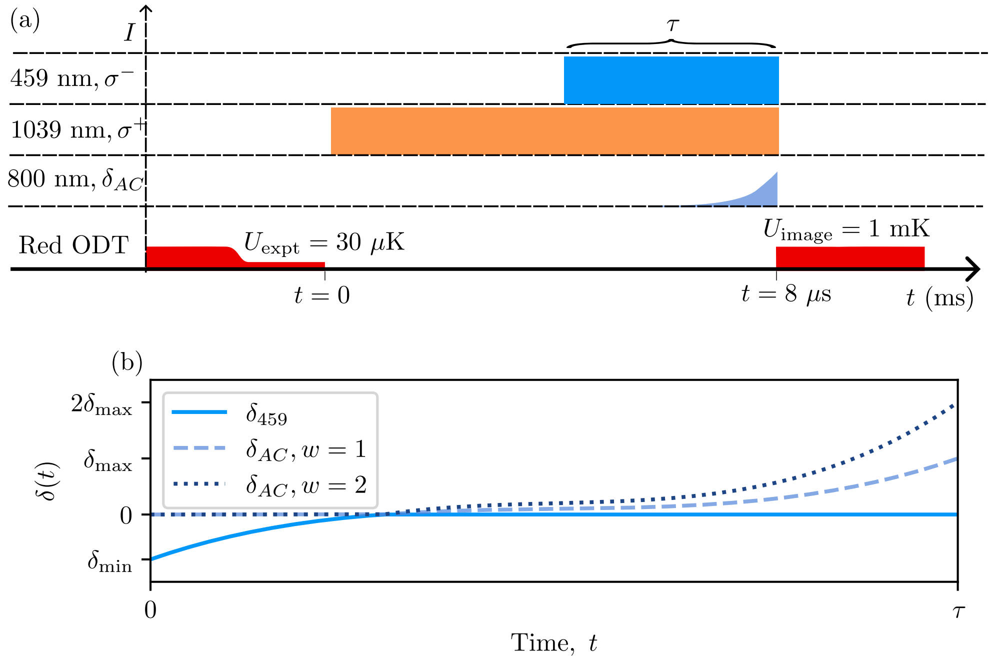

Figure 7.8: Experimental timeline for the MWIS experiment. Figure (a) is a schematic of the powers of the various input beams over the experiment period, and is not to scale. Figure (b) is the detuning of the atoms relative to resonance over the Rydberg interaction period. The detuning caused by 459 nm light through the frequency chirp, δ_{459}, is shown in solid blue. The detuning caused by 800 nm light through holographic dipole potentials, δ_{AC}, is shown dashed and dotted blue lines for weight 1 and weight 2 light respectively.. Each of the two lines represents a different weighting, where $\delta_{\max}$ is the shift on atoms with weight $w = 1$.

the input power of the light potential, from $\delta = 0$ to $\delta_i = \delta_{i,\max}$. At the end of this ramp the input Rydberg interaction and the holographic light shift is switched off, the atoms recaptured in red-detuned ODTs and the survival measured. The timeline for this experiment is shown in Fig. 7.8.

The power of the 800 nm light, and thus the δ_{AC} applied to the atoms, is controlled also by the AWG. A second channel is used, where the power of the output RF waveform is calibrated to the diffraction power of the AOM. Frequency is kept constant in this ramp. Through this method, δ_{AC} is kept synchronised with δ_{459}, and is at the intended power.

This alignment is, at first, to the pattern seen in Fig. 7.9. Subfigs. (a) and

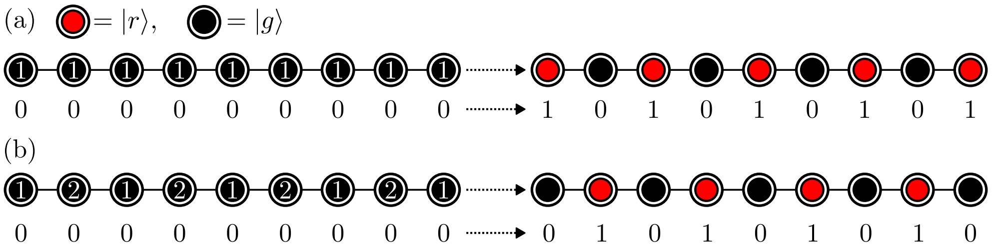

Figure 7.9: Arrangement of atoms in a weighted 1D line. The weighting is given by the number within each atom, represented by a filled circle. Figure (a) shows a pattern of equal weights, such that the MWIS matches that of the MIS. Figure (b) shows an alternating pattern of weights 1 and 2, such that the MWIS has opposite parity in each readout bit relative to the MIS.

(b) correspond to the top and bottom rows respectively. Here, two rows are weighted using the holographic potentials, with the same array geometry used in Fig. 7.5. In this case, the top row is weighted all at $w = 1$, while the bottom row alternates between $w = 1$ and $w = 2$. Highlighted in red are the MWIS solutions for each row. These MWIS solutions are the states that minimise the cost for the classical cost function as defined in Eq. 7.5 and are the ground states of the Hamiltonian defined in Eq. 7.5.

7.3.2 Calibration Of Site Specific Weightings

Initial alignment of the holographic light potential follows the method of Sec. 6.4, where focussed spots of 800 nm light generated through SLM potentials eject atoms from red-detuned ODTs. As the dipole potential of a ground-state atom $U \propto I(r)$ as per Eq. 2.49, the force on the atom $F \propto -\nabla I(\mathbf{r})$ can be minimised by reducing spatial intensity gradient at the atoms.

Once aligned, the SLM potential that generates the focussed spots is modified with a pseduo-aperture, as introduced in Sec. 6.3. This acts as a spatial broadening of the intensity, such that $\nabla I(\mathbf{r})$ is reduced at constant central intensity. This pseduo-aperture reduces efficiency of the SLM system.

The AC Shift seen on each site is, at first, different for each atom, and not in proportion to the required weights for this atom, as seen in Fig. 7.10. Direct

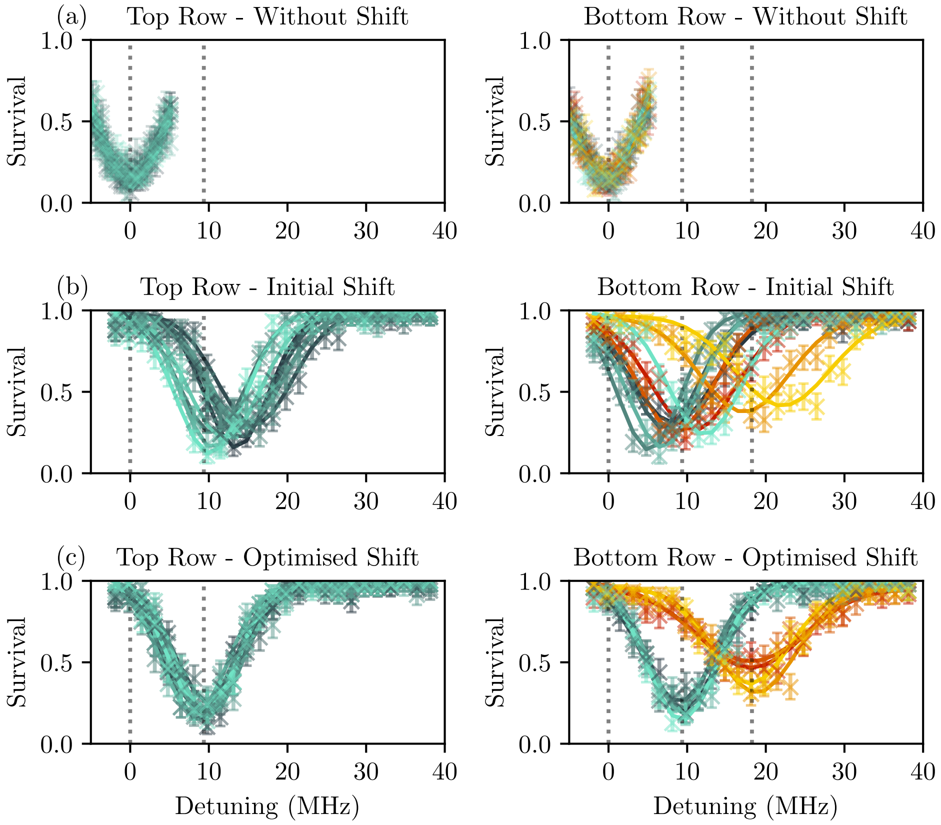

Figure 7.10: Spectroscopy for the top and bottom rows of the array, on the left and right columns of the plot respectively. Each colour refers to a single atom. Data in black or blue corresponds to $w = 1$, while red and yellow corresponds to $w = 2$. Figure (a) shows the two rows without weightings applied, where all spectroscopic features are centered at the same point. The x-axis is rescaled to the detuning from the resonance frequency, as found in (a). Figure (b) is the same transition with the initial holographic potential. Figure (c) is after multiple rounds of hologram iteration, such that the shift converges to the required weights.

Rydberg spectroscopy measurements as the AC shift is applied through the sequence measures this shift applied. This uses the process shown in Sec. 5.1, addressing the $6S_{1/2} \rightarrow 50S_{1/2}$ transition. This can be used for feedback to further update the hologram used for the 800 nm SLM through the awGS algorithm.

Direct measurement of this spectroscopy is seen in Fig. 7.10. The initial hologram used does not correspond well to the required weights, and it is only

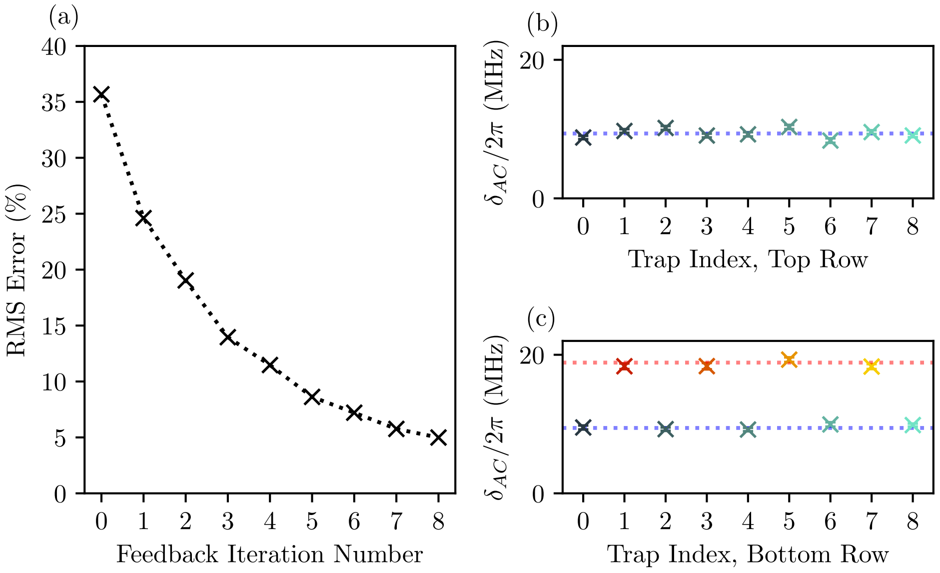

Figure 7.11: Convergence process of the holographic potential for weighting on 1D chains. Figure (a) is the RMS error for each iteration. Iteration 0 is measured before any feedback is applied, and corresponds to data seen in 7.10(b). Figures (b) and (c) are the measured maximum shifts applied by the holographic potential after 8 iterations of feedback, for the top and bottom rows respectively. The dotted lines correspond to the detuning that is rescaled to $w = 1$ in blue and $w = 2$ in red respectively.

after a procedure of iteration that δ_{AC} converges. In this iteration procedure, the average shift is calculated relative to each atom's intended weight to find the average weighted shift. Where the atoms are shifted below the average require an increase in power, and the hologram is adjusted accordingly. The opposite applies for where atoms are above the average shift, where power is reduced.

This is the same procedure as reweighting the AC Stark shifts seen in [87], where the cause of the shift was the red-detuned ODTs, with the AC Stark shift measured in a MW Ramsey sequence. Weighting can be controlled by adjusting the relative weights δ_i. By rescaling the intended power, the site-specific detuning can be rescaled relative to each beam. With multiple iterations, this converges to a site-specific shift in the intended ratios. This light is applied with minimal loss to the atoms when off-resonance. The iter-

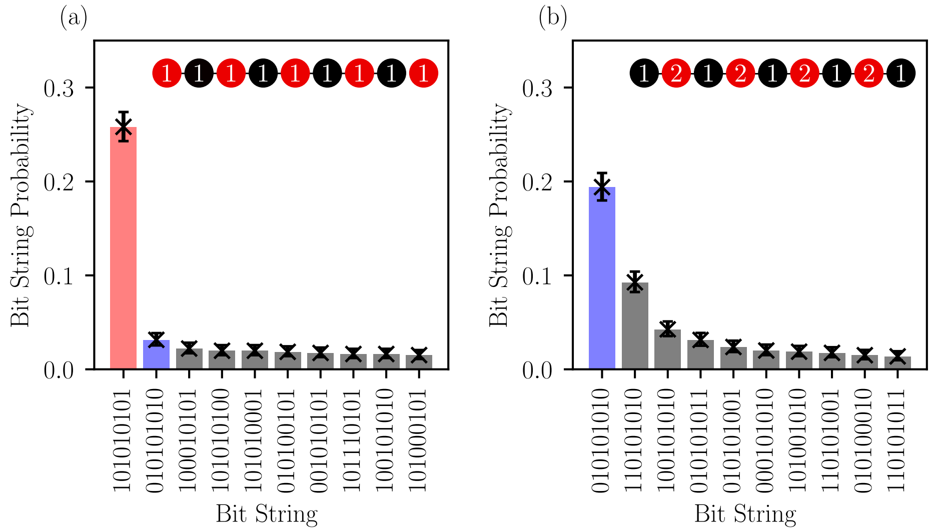

Figure 7.12: Results for both top and bottom graphs after closed loop optimisation, based on weights as shown in Fig. 7.9, and also shown in inset with the MWIS highlighted. Figure (a) shows the results for all weights being equal on the top row, while (b) shows the results for alternating weights, on the bottom row. The MWIS solutions for equal and alternating weights are highlighted in red and blue respectively. Only the 10 most common strings are plotted for clarity.

ative feedback process is shown in Fig. 7.11, where the figure of merit used is the RMS error from the average weight. The top row at $w = 1$ converges for $\delta_{\mathrm{max}}/2\pi = 9.4(6)$ MHz. The bottom row has two weights, converging where for $w = 1$, $\delta_{\mathrm{max}}/2\pi = 9.4(3)$ MHz, and for $w = 2$, $\delta_{\mathrm{max}}/2\pi = 18.9(6)$ MHz. The Rabi frequency was at the same as for the unweighted case, where $\Omega/2\pi = 2.75(2)$ MHz.

Direct comparison of the AC Stark shift at $n = 50$ and $n = 80$ shows a small increase in magnitude, such that $\delta_{\mathrm{AC}}^{\mathrm{n=80}} = 1.10(3)\delta_{\mathrm{AC}}^{\mathrm{n=50}}$.

7.3.3 MWIS in 1D Chains

The closed-loop optimisation for this process used the bottom row for feedback, using Eq. 7.5 with $U_{i,j} = w_i w_j$. This experiment is set up such that the ramp is applied to both rows of atoms simultaneously, allowing for direct comparison between the two weighting schemes.

The best result was then repeated with 1000 experiments to verify the results, which is shown in Fig. 7.12. This prepared the target state with 26(2)% probability for the top row and 19(2)% for the bottom row. As the same ramp produced the MWIS solution simultaneously, a single ramp scheme produces an approximately adiabatic Hamiltonian for a range of weights. This was found with parameters $\delta_{\text{min}}/2\pi = -7.74$ MHz, $\delta_{\text{max}}/2\pi = +7.65$ MHz, $\tau = 2.97\ \mu$s and $s = 1.00$.

With this ramp prepared, this experimental sequence is stopped at points along the ramp and the experimental state is read. This shows how the initial state is annealed into a highly entangled state, before the desired MIS/MWIS[2] state is generated. Evolution of this state is seen in Fig. 7.13.

7.3.4 MWIS in 2D

MWIS Preparation

The theoretical construction of arbitrary graphs for quantum optimisation in [85] introduces a 5-vertex weighted graph, which cannot be constructed as a UDG. It can, however, be mapped to a larger 9-vertex weighted UDG. With careful choice of weights for the 4 ancilla vertices, the ground state of the 9-atom UDG maps to the 5-vertex graph.

Three weighting schemes are demonstrated, labelled graphs **A**, **B** and **C**, which are shown in Fig. 7.14. For this graph, atoms are considered connected if within 8 μm, an increase from the 7 μm shown in Sec. 7.3.3. This is due to the restriction that distance between next-nearest neigbours (or diagonals on the square grid) at this lower spacing is $\approx 9.90\ \mu$m, and as such is still within the blockade radius of the interaction. This wider spacing restricts the dipole-dipole interaction energy V_{dd} to 11.2(2) MHz, as measured in Sec. 5.3.

---

[2] In this case of the top row, the MIS and MWIS solutions are identical.

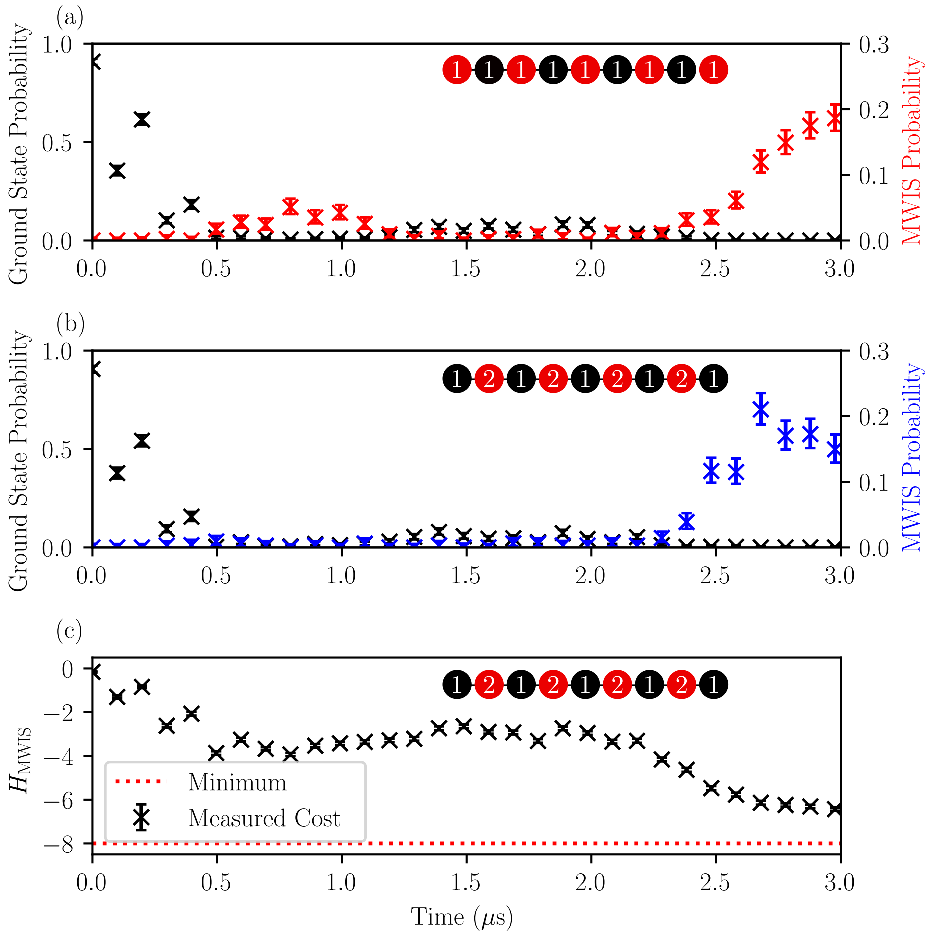

Figure 7.13: Plots show the evolution of the measured state in the two graphs. Each plot includes inset of graph with weights, with the MWIS highlighted. Figure (a) is the top row, with the initially prepared state probability in black and the MWIS solution probability in red. Figure (b) is the bottom row, with the solution here in blue. Figure (c) is the evolution of the cost function H_{MWIS}, evaluated using data from the bottom row.

The ancilla weights w_α and w_β are given by

$$\begin{aligned} w_\alpha &= (w_1 + w_3)/2, \\ w_\beta &= (w_2 + w_4 + w_5)/2. \end{aligned} \tag{7.6}$$

These ancilla weights are not unique, but will define the required mapping to find the MWIS ground state as long as $w_i \geq 1$. These graphs have weights that correspond to different MWIS solutions, where for each graph the closed-

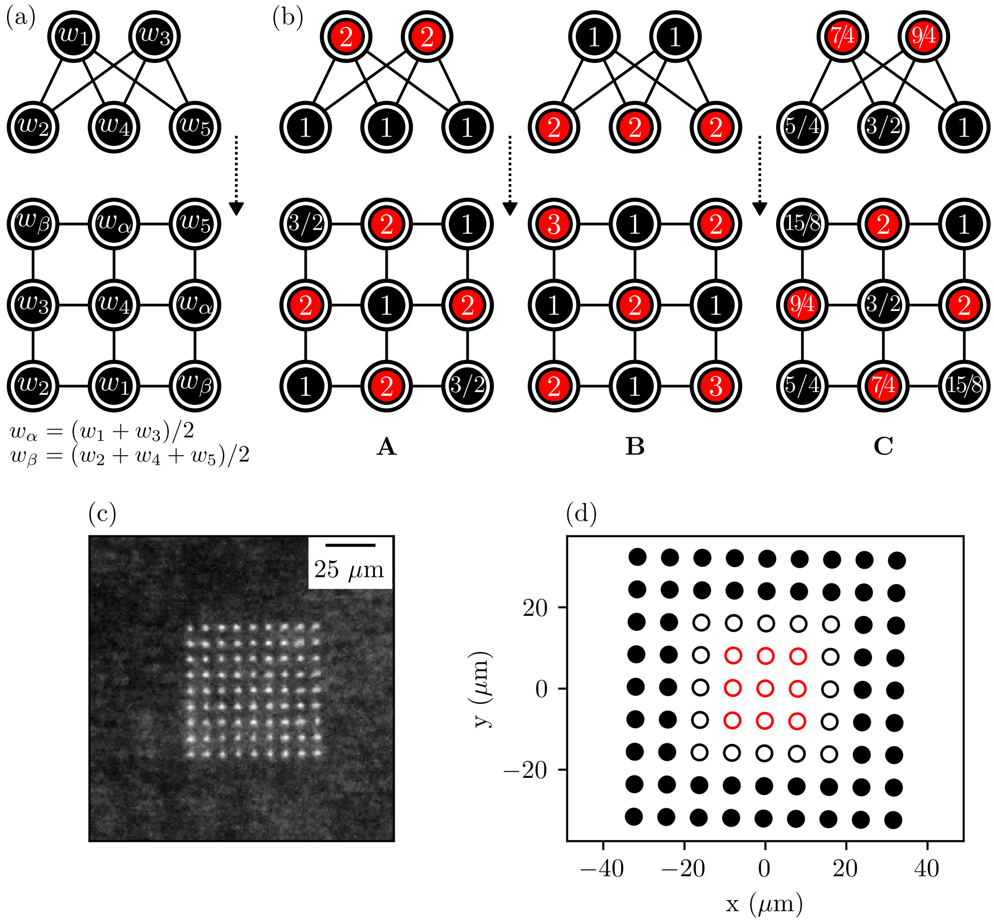

Figure 7.14: Graphs used for 2D MWIS experiments. Figure (a) is the mapping from a 5-vertex graph to a 9-atom UDG. Figure (b) shows the three weightings tested, labelled **A**, **B** and **C**. Figure (c) is the per-pixel maximum of 400 experimental images used for calibration showing the trap site arrangement, and (d) is the array configuration, where hollow red circles are target circle, black hollow circles are to be removed and filled black circles act as a reservoir.

loop optimisation procedure is carried out. Measurements to calibrate for the three graphs are shown in Fig. 7.15. The Rabi frequency was measured for each graph as required for reweighting, such that $\Omega/2\pi = 2.75(6)$ MHz for graph **A**, 2.83(3) MHz for graph **B** and 2.84(6) MHz for graph **C**.

The cost function used for this experiment is the cost H_{MWIS} of the 9-vertex graph. The detuning ramp $\delta(t)$ was optimised for minimising the cost of graph **A**, and the same ramp used without adjustment on each of the two

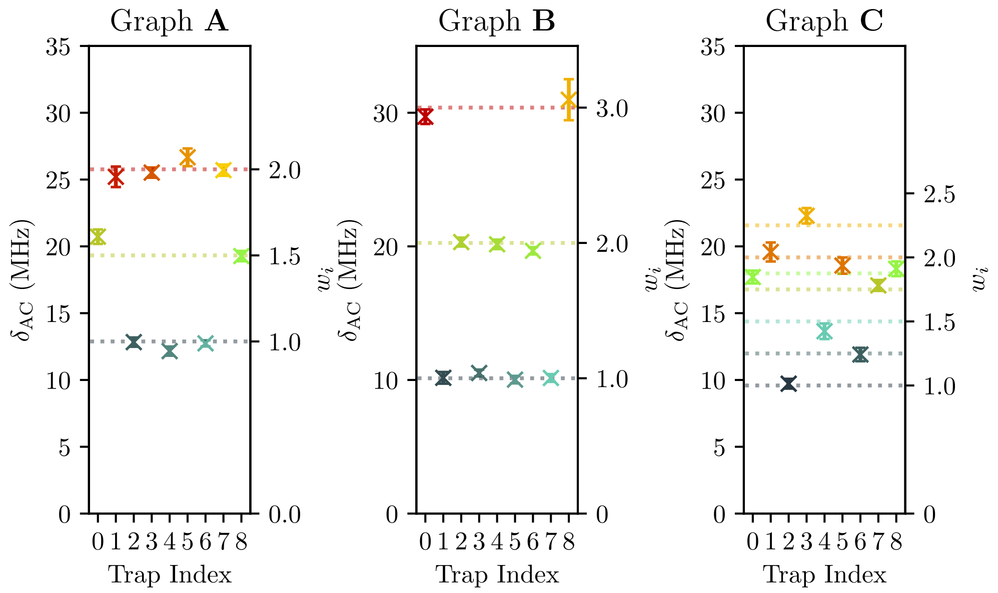

Figure 7.15: Measured weightings for graphs **A**, **B** and **C**. The left of each graph shows the shift in MHz, while the right is the relative weighting this corresponds to, rescaled to the minimum weight as $w = 1$.

other graphs **B** and **C**. This was found to be with parameters $\delta_{\text{min}}/2\pi = -3.416$ MHz, $\delta_{\text{max}}/2\pi = +4.326$ MHz, $\tau = 1.673$ μs and $s = 0.686$. Given that the data shown in Fig. 7.15 is the maximum potential power from the holographic potential, and that in all cases $w_i w_j < U_{i,j} A_{i,j}$ for the algorithm to be valid, δ_{max} was limited to a value such that this condition was not broken. For the highest weight seen ($w = 3$), this corresponds to $3 \cdot \delta_{max} = 13.58 < V_{dd}/2\pi$, at the limits of the acceptable range.

The probability distribution for each graph is shown in Fig. 7.16, with measured data shown in Table. 7.1. For each case, the most-probable state

Graph	Min(H_{MWIS})	$\langle H_{\text{MWIS}} \rangle$	P(MWIS, 9-Vertex)	P(MWIS, 5-Vertex)
A	-8	-3.6(1)	0.11(1)	0.22(2)
B	-12	-6.9(2)	0.17(+0.03, -0.02)	0.31(3)
C	-8	-3.5(1)	0.10(+0.02, -0.01)	0.22(2)

Table 7.1: Table comparing measured probabilities for 2D MWIS experiment. Min(H_{MWIS}) is the value of H_{MWIS} in the ground state, and its minimum value. $\langle H_{\text{MWIS}} \rangle$ is the average value measured. P(MWIS, 9-Vertex) and P(MWIS, 5-Vertex) are the probabilities found of measuring the ground state for the 9-vertex graph and the 5-vertex subgraph respectively.

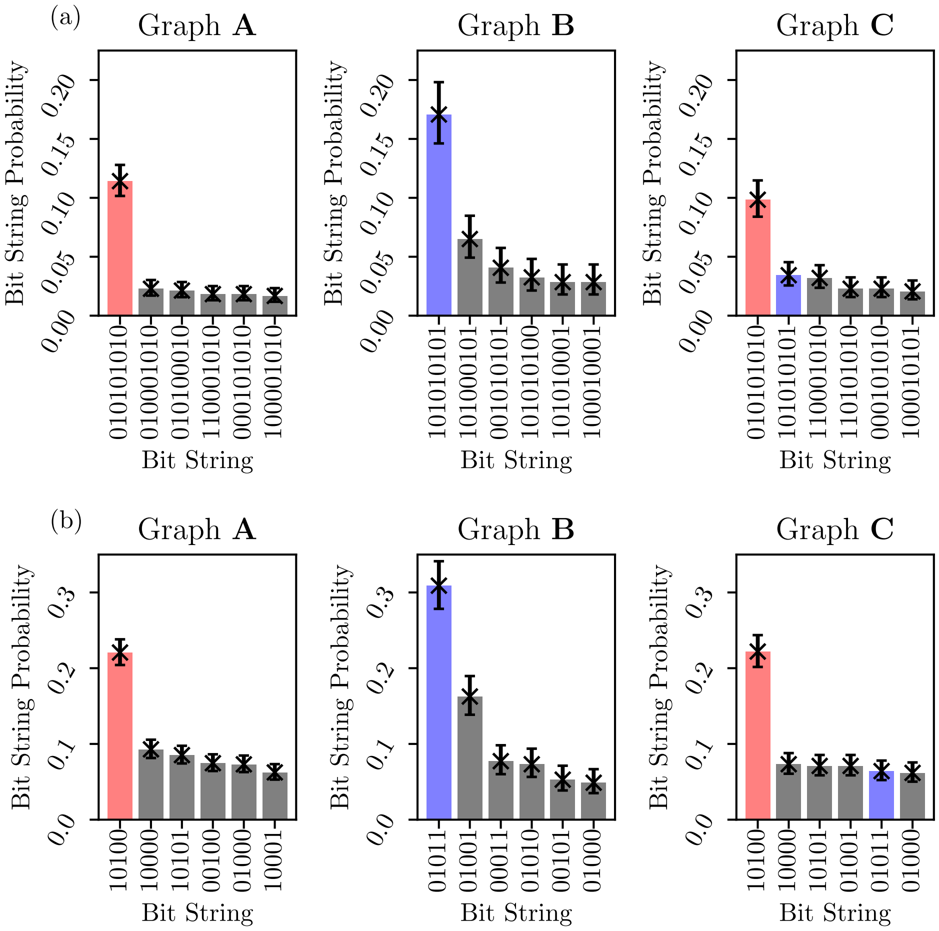

Figure 7.16: Measured probabilities for graphs **A**, **B** and **C**, showing only the 6 most common strings for clarity. Highlighted in red is the MWIS solution for graphs **A** and **C**, while **B** is highlighted in blue. Ramp used for this experiment is optimised for **A** and repeated for **B** and **C**. Figure (a) is results for the 9 vertex graph, with the correct MWIS solution for that graph highlighted in red. Figure (b) is results for the 5 vertex graph found by disregarding the ancilla qubits, with the same highlighting. For all graphs, the MWIS solution is generated with statistical signficance.

generated in the experiment was the expected ground state of the 9-vertex graph. This ground state is the solution for the 9-vertex UDG and the initial 5-vertex graph. The probabilities measured just by reading the data qubits and ignoring the readout qubits also showed a preference for the 5-vertex MWIS solution.

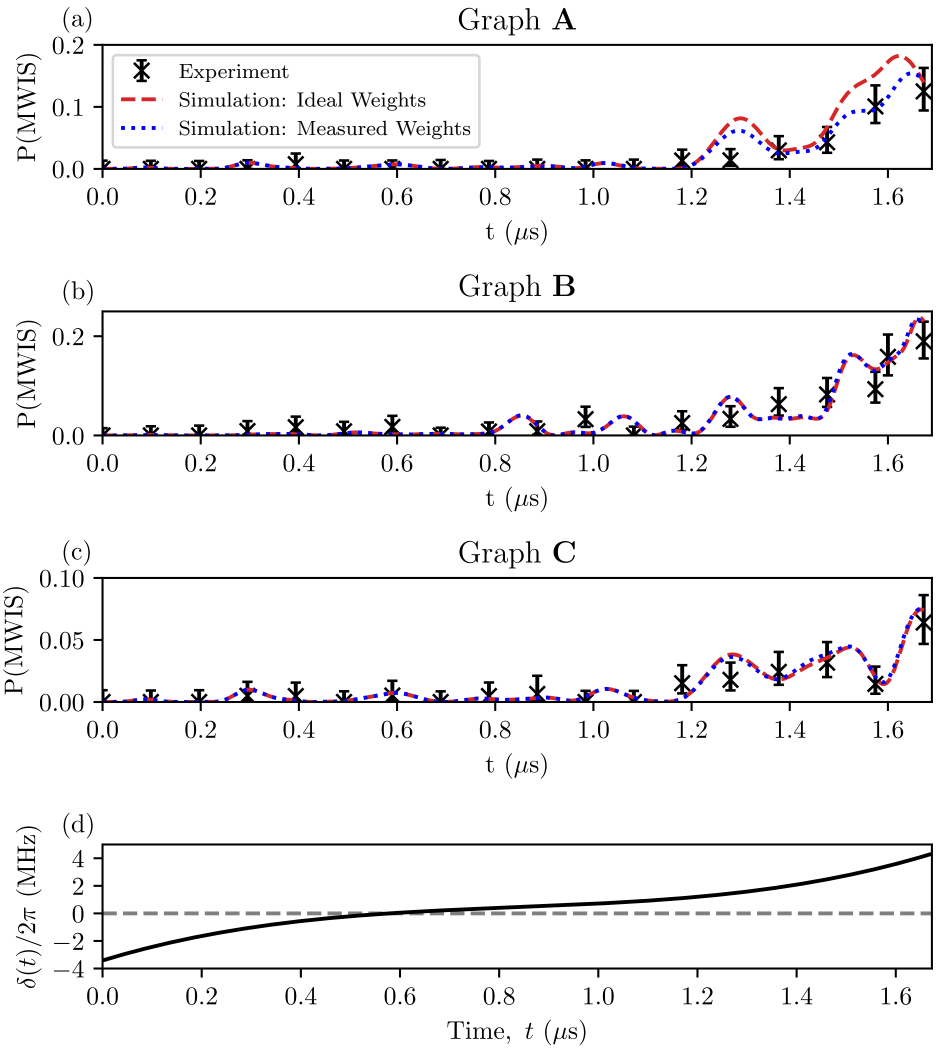

Figure 7.17: Measured MWIS probability over time for graphs **A**, **B**, and **C**, shown in figures (a), (b) and (c) respectively. These are compared with simulations using the ideal detunings (red, dashed line) and with measured detunings (blue, dotted line). Figure (d) shows the detuning profile for this ramp for comparison. This is the detuning for $w = 1$.

State Preparation Dynamics

Using simulation code by G. Pelegri, the dynamics for the 9-vertex graphs were simulted directly, with parameters matching those used in the experiment. This includes long-range dipole-dipole interactions beyond the blockade radius, spontaneous decay from the Rydberg state and scattering from the intermediate state.

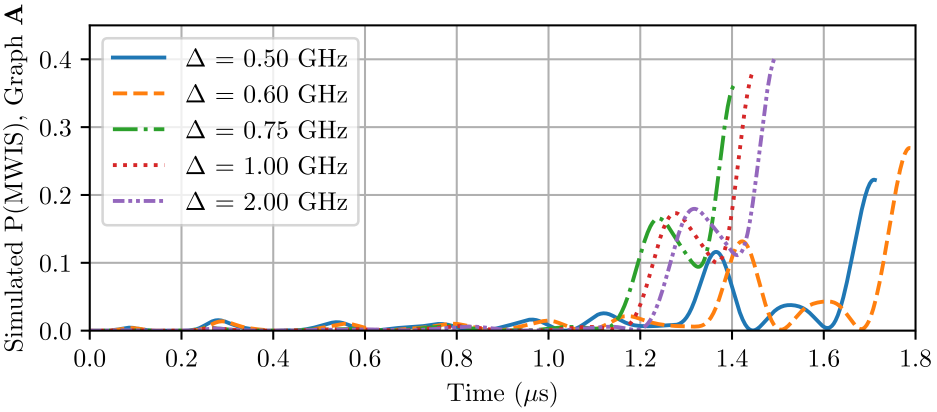

Figure 7.18: Simulated optimisation results for Graph **A**. Experiment parameters are kept identical, with adjusted simulated power to maintain Rabi frequency while adjusting intermediate state detuning Δ. Maximum probability of preparing the MWIS solution state increases with Δ.

As shown in Fig. 7.17, these simulations show comparable final state populations and dynamics, although it should be considered that the simulation does not consider imperfect state preparation or imaging loss. For these simulations, $\chi^2_\nu = 0.59$ for Graph **A**, $\chi^2_\nu = 1.10$ for Graph **B**, and $\chi^2_\nu = 0.40$ for Graph **C** with no free parameters. This indicates a good simulation of the experimental system.

Using these simulations, the maximum MWIS state probability that can be generated is shown to be dependent on the intermediate state detuning. Using a Nelder-Mead optimisation algorithm, the cubic ramp parameters were optimised for the simulated MWIS state probability, as seen in Fig. 7.18. This optimisation process uses the same experimental parameters, changing only intermediate state detuning Δ and changing the simulated power of the 1039 nm beam to match the previous Rabi frequency. The optimisation parameters used for the cubic ramps are the same as used in the closed-loop experimental feedback experiment. This shows higher probabilities of MWIS state preparation can be achieved, and thus indicates that one limiting factor in this experimental scheme is scattering from the intermediate state.

Given that this annealing process follows the Kibble-Zureck mechanism [223,

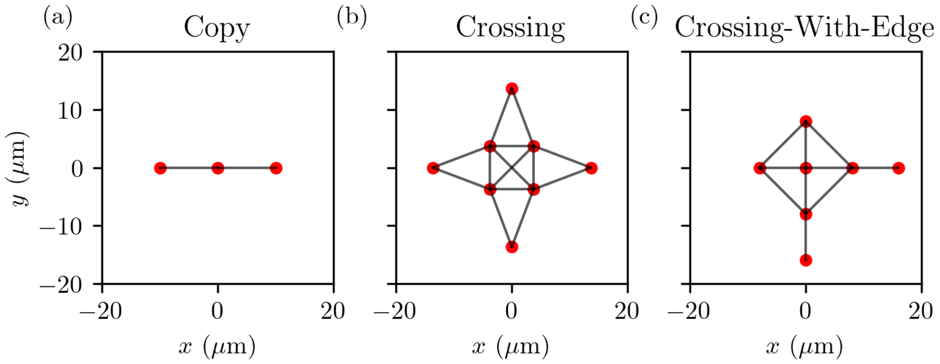

Figure 7.19: Plan of atom placement for gadgets. Here, each red point indicates the placement of an optical dipole trap, and thus an atom. Black lines are guides for the eye to indicate atoms within a Rydberg blockade of 10 μm. For further details on the weightings for each subgraph, see [85].

224], a reduced Rabi frequency and greater detuning ramp parameter range is predicted to further improve annealing performance. This must, however, be fast relative to the lifetime of the Rydberg state, the coherence time of the Rydbeg driving system and within the limits posed by experiments in free-fall.

7.4 Gadgets For Graph Problems

7.4.1 Gadgets For UDG Mapping

In order for a Rydberg simulator to find solutions to an arbitrary weighted graph, this weighted graph must be encoded in a unit-disk graph. This mapping process is shown in Sec. 7.3, but a process for efficient mapping of a generic weighted graph to a weighted UDG is required that can be implemented on neutral atom systems. Solutions creating specific connectivity between nodes can be managed by using anisotropic Rydberg states [225] or by utilising 'quantum wires' [226], further strings of atoms to connect otherwise distant nodes.

Special consideration is required for a selection of the sub-graphs used, named ‘gadgets’, as shown in Fig. 7.19. First described in [85], this set of subgraphs can be utilised to encode any weighted graph problem in a neutral atom framework, with at worst an $\mathcal{O}(N^2)$ overhead in trap sites, where N is the number of traps used. This corresponds to a copy gadget, similar in process to the aformentioned quantum wires, a crossing gadget and crossing-with-edge gadget.

For each gadget, there are conditions that must be met to be a valid encoding. Firstly, each graph is not necessarily on a Bravais lattice; any gadget encoded within a greater graph will likely include atoms offset from said graph. Secondly, all gadgets require weightings to remove degeneracy, such that the ground state of each gadget produces the ground state of the encoded graph problem, even if the graph that is to be encoded has equal weightings on all nodes. Finally, the spacing of each gadget is not uniquely defined; connectivity of the graph is the crucial factor, and this connectivity is defined by strong dipole-dipole interactions between atoms. This adds complexity to the problem, given that $V_{dd} \propto |r_i - r_j|^6$ and so does not have a hard edge. The so-called ‘tails’ of this interaction can lead to a small but not insignificant energy penalty when not intended, where Rydberg transition is not blockaded but still to a degree suppressed. These gadgets have use beyond MWIS optimisation, as shown in [227] for a theoretical use in modelling quantum dimers.

7.4.2 Crossing Gadget Configuration

The copy gadget and crossing-with-gadgets are comparable in geometry and weighting to the 1D and 2D geometries shown in Sec. 7.3 respectively. The ‘crossing’ gadget requires a heavily weighted and dense central square, for a high degree of connectivity. The array geometry used for the experiment is shown in Fig. 7.20, where the central square is spaced by 6 μm on each edge.

Simulations be G. Pelegri optimised the weighting scheme and geometry of

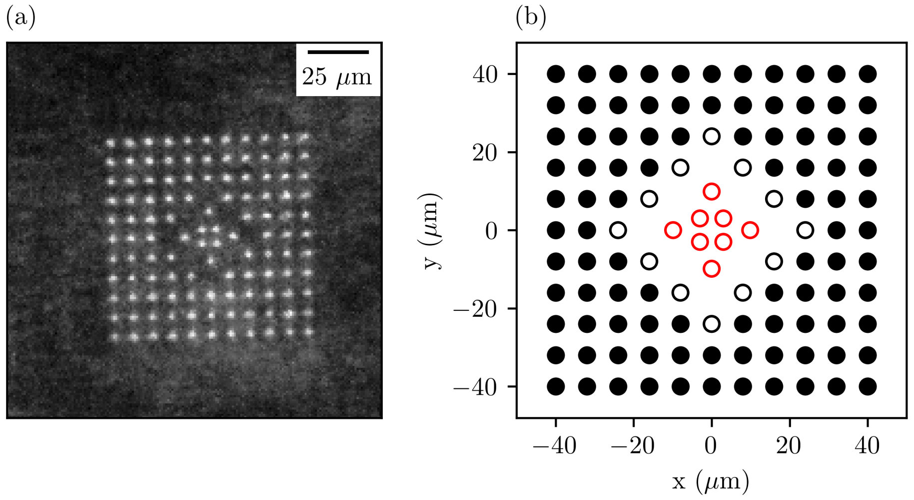

Figure 7.20: Trap array geometry for a crossing gadget experiment. Figure (a) shows the standard deviation of experimental images of the array. Figure (b) shows the array sites to be targetted. Red empty circles define the arrays where atoms to be moved to in sorting, while white circles are where atoms are to be removed. Filled black circles are the reservoir.

this scheme to optimise performance on the SQuAre platform. This is measured by maximising fidelity of the ground state of the crossing gadget with the MWIS solution after a simulated adiabatic ramp. This was found to use the minimum spacing of 6 μm for the internal square of atoms, with weight 2.86, while the outer atoms are spaced 7.50 μm from their nearest neighbours. Given the weak blockade regime in the central set, where $V_{dd}/2\pi = 7.8(2)$ MHz along the diagonals, this corresponds to a relatively small maximum shift on atoms with $w = 1$ relative to the Rydberg transition linewidth, which is dominated by the Rabi frequency.

Input 1039 nm frequency was adjusted, such that the intermediate state detuning is instead $\Delta/2\pi = 1000(2)$ MHz. Rabi frequency was measured to be $\Omega/2\pi = 1.24(6)$ MHz at this detuning, using the same input powers as before. The weighting scheme, with experimental results, is seen in Fig. 7.21(a) and (b).

The crossing gadget has four unique MWIS solutions, corresponding to the

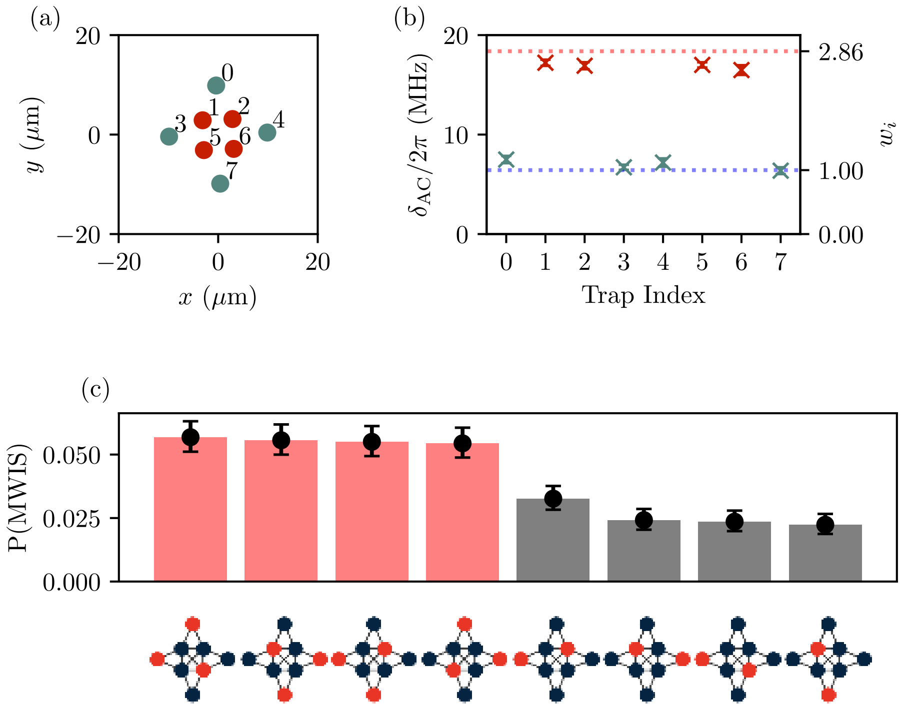

Figure 7.21: Crossing gadget experimental results. Figure (a) shows the arrangement and indexing of atoms within the crossing gadget, where inner red/yellow points indicate higher weightings compared to the outer, blue atoms. Figure (b) shows measurements of the maximum applied AC Stark shift, used for calibrating relative weightings. Figure (c) shows resultant probabilities for these graphs, with MWIS solutions highlighted in red. Corresponding states are shown pictorally.

four-fold rotational symmetry of the graph arrangement. The total probability of finding any of the four MWIS solutions was found to be 22(1)%, after closed-loop optimisation. Each of the four solutions had measured probabilities of $5.6(6), 5.6(6), 5.5(6), 5.4(6)\%$, as shown in Fig. 7.21(c).

This was found with parameters $\delta_{\rm min}/2\pi = -2.00$ MHz, $\delta_{\rm max}/2\pi = +2.00$ MHz, $\tau = 3.00$ μs and $s = 1.00$, at the limits of the parameter range set. Even with these limitations, these experiments show promise for the use of gadgets as a scalable method of analogue quantum computing in neutral atom systems.

7.5 Summary

The results shown in this chapter demonstrate analogue quantum computing on a neutral atom system, which can be reconfigured to arbitrary graphs. These results are also shown in [86]. While the largest graph demonstrated in this work is 9 atoms, distributed in a 2D plane, the SQuAre platform has been shown to be capable of arranging up to 100 atoms (see Chapter 4). Graph optimisation problems themselves are not limited to MIS/MWIS problems. Related problems including minimum (connected) dominating sets and minimum vertex cover are also proposed problems for a neutral atom quantum computer to tackle [228].

Limitations on intermediate state detuning and Rydberg excitation power are due to be addressed with the introduction of high power optics, while field homogeneity is due to be addressed with the use of a further SLM for beam shaping on the 1039 nm optical path. Reconfiguration of the 800 nm holographic potential optics are also underway, as discussed in Sec. 6.6. In summary, the SQuAre platform is capable of further investigation into the optimisation of graph problems as an analogue quantum computer.

Recent work from [229] has used a similar process, where an SLM is used for a bichromatic detuning ramp on a Rydberg system. This is a promising development in using this scheme beyond graph optimisation, demonstrating use as a quantum simulator in studying thermalisation dynamics of highly entangled systems. These techniques and methods may prove of further use in analogue quantum computation and simulation with neutral atom arrays.

Chapter 8

Conclusion and Outlook

8.1 Summary of Results

This thesis presents the research towards building a scalable neutral atom quantum computing platform. Through this work, the experimental design for constructing the system were introduced, with complementing details on chamber construction, non-destructive readout and randomised gate benchmarking shown in [87].

Firstly, the experiment is initialised by starting from an ultra-high environment with a Cs vapour cell and resulting in an array of atoms, each solely occupying an optical dipole trap in an arbitrary and reconfigurable array. These atoms can then be addressed through Rydberg, Raman or microwave interactions. This hardware and the experimental processes are described through Chapter 3.

Site occupancy in these arrays is stochastic, and so a sorting process is required for the construction of arbitrary arrays without defects. This is shown through Chapter 4. Defect free array construction has been shown to be possible up to 100 atoms in this experiment, with experimental improvements underway to improve experimental yield, such that this technique can scale to larger systems. As such, a 100 qubit array is possible within this system.

The dipole-dipole interaction between Rydberg atoms is used to mediate qubit interactions within this experiment. Spectroscopic measurements to the $n = 50$ and $n = 80$ Rydberg states are shown in Chapter 5, alongside evidence of Rydberg blockade and direct measurements of the Rydberg dipole-dipole van der Waals differential shift. Finally, the $|\Psi^+\rangle$ is prepared and its fidelity measured, forming the initial basis of a digital quantum computing scheme.

Chapters 6 and 7 each introduce novel techniques for using holographic light potentials in neutral atom quantum computing experiments. The first demonstrates using these holographic potentials to trap arrays of Rydberg atoms in blue-detuned optical dipole traps. These traps can hold Rydberg atoms in place, rather than experiments occurring in free-fall, and as such may be of use in high-fidelity quantum gates or in increasing Rydberg interaction periods.

The second of these two chapters uses this holographic light shift to instead demonstrate the platform's capability of solving graph optimisation problems. The unweighted graph optimisation problem case is first demonstrated, followed by weighted graph optimisation problems, specifically solving for the maximum weighted independent set. Work on gadget graph structures extends this scheme to arbitrary graphs, with at-worst $\mathcal{O}(N^2)$ overhead in the numbe of traps required, where N is the number of graph vertices.

8.2 Outlook

8.2.1 Outlook For The SQuAre Experiment

The SQuAre experiment is capable of further optimisation and further experimental schemes. This includes further experiments on weighted graph optimisation as it scales to larger graph sizes. Also possible is the exploration of other quantum analogue computing and simulation experiments, such as developing on previous work involving probing non-equilibrium many-body dynamics using quantum adiabatic algorithms [75, 229], processes that seek shortcuts to adiabaticity [190, 230, 231] or other quantum phenomena such as quantum many-body scars [232, 233].

The same platform is also capable of digital quantum computing. Single qubit fidelities on the SQuAre platform have already shown state-of-the-art single-qubit gate fidelities [64]. Implementation of Rydberg gates through adiabatic rapid passage [234] or other schemes [65, 151] can help pave the way towards large and error-tolerant neutral atom digital quantum computing.

Assisting in these experiments is a range of hardware upgrades that are planned or in current implementation. Mentioned before in this thesis is the increase in accessible power of the Rydberg system, allowing for greater transition rates to the Rydberg state, or greater intermediate state detunings at the same transition rate. The greater power is produced by a fibre amplifier, with the limits of power currently set by the optical fibre used. Also due to be implemented is an SLM for beam shaping to improve Rydberg interaction homogeneity across the array.

Further upgrades in implementation are local addressing for both 459 nm and Raman light. This uses the same hardware as the 1039 nm local addressing beam (also used for array sorting), but would allow for site specific Rydberg or Raman addressing. Hardware improvements to the sorting system promise parallel sorting for substantial gains in sorting performance. Replacement of the Teensy microcontroller with an arbitrary wave generator (AWG) would allow for more precise and rapid control of the sorting tweezer, along with parallel sorting operations. Calculation times for parallel sorting waveforms can be calculated rapidly using hardware acceleration from a GPU [235]. This same hardware may be implemented into the local addressing systems for parallel single-qubit interactions with Rydberg or Raman beams.

Reconfiguration and optics adjustments are also underway to improve the efficiency of the 800 nm hologram system, such that more atoms can be addressed at once. Also possible is the introduction of further amplification of the 1064 nm beam, such that an even greater number of atoms can be trapped in an array at once.

8.2.2 Outlook For Neutral Atom Quantum Computing

Since the start of construction of the SQuAre experiment, neutral atom quantum computing systems have spread rapidly. Recent results have demonstrated 48 logical qubits in a reconfigurable array with two- and three-qubit gates [68]. Private ventures such as Pasqal, QuEra and Atom Computing all have or expect to have $>$1000 qubit systems [236], with goals of $>$100 logical qubits encoded in these systems. Former limits on array sizes of > 1000 atoms at once have been met using using cryogenic systems for trap lifetimes of > 6000 s [183]. The largest atomic array experiment has demonstrated the trapping of > 6100 atoms concurrently with lifetimes > 1000 s, with long lifetimes at room temperature through developments in titanium sublimation pumps and repeated PGC stages [55].

Further developments in neutral atom system design is also progressing. Local addressing at high speeds is undergoing development to address up to 10^4 atoms at once with uniform frequency and intensity, using combinations of acousto-optic deflectors, spatial light modulators and digital micromirror devices [237, 238]. Recent proposals using quantum Low Density Parity Check (qLDPC) codes show dense data encoding using large arrays of identical atoms, while retaining error-correcting capability [66, 239, 240].

The architecture of these quantum computing platforms is not static. A dual-species approach promises improved multi-qubit gate fidelities and capabilities in error correction [241, 242], while alkaline-earth species experiments promise erasure conversion, correcting errors and providing exceptional 2-qubit gate fidelities [60]. These are expected to be new tool in a variety of research fields as these quantum information technologies mature. Already, a quantum speed-up is predicted for generative machine learning tasks on neutral atom systems [71], while quantum simulators have shown dynamics of interest to materials science that cannot be simulated classically [243, 244].

As the field of neutral atom quantum computing continues to develop, the SQuAre platform can be expected to be part of that development.

Appendices

Appendix A

Zernike Polynomial Alignment

The Zernike polynomials are defined over a radius $0 \leq r \leq 1$ and angle $0 \leq \theta \leq 2\pi$ for some integers n and m, where $n \geq 0$ and m is in the range $-n, -n+2, ..., n-2, n$. The full definition of the Zernike polynomials is given in [200]. For the polynomials used here, the unnormalised form of the polynomials are shown in Table A.1.

This appendix includes the full range of calibration data for optimising blue-detuned optical dipole traps as partially shown in Fig. 6.13. This covers the range of Zernike polynomial prefactors up to $Z_3^{\pm 3}$, where the scan is adjusting the scaling prefactor for each polynomial. The magnitude of this polynomial

n	m	$Z_m^n(r, \theta)$
0	0	1
1	-1	$\sin(\theta)$
1	1	$\cos(\theta)$
2	-2	$r^2 \sin(2\theta)$
2	0	$2r^2 - 1$
2	2	$r^2 \cos(2\theta)$
3	-3	$r^3 \sin(3\theta)$
3	-1	$(3r^3 - 2r)\sin(\theta)$
3	1	$(3r^3 - 2r)\cos(\theta)$
3	3	$r^3 \sin(3\theta)$

Table A.1: Table of unnormalised Zernike polynomials up to $n = 3$.

at some coordinate refers to the degree of retardance applied by the SLM at that coordinate. Should the prefactor be zero, this is equivalent to a constant phase adjustment across the unit disc, and so no beam adjustment is applied. The scans of the range of prefactors is shown in Fig. A.1.

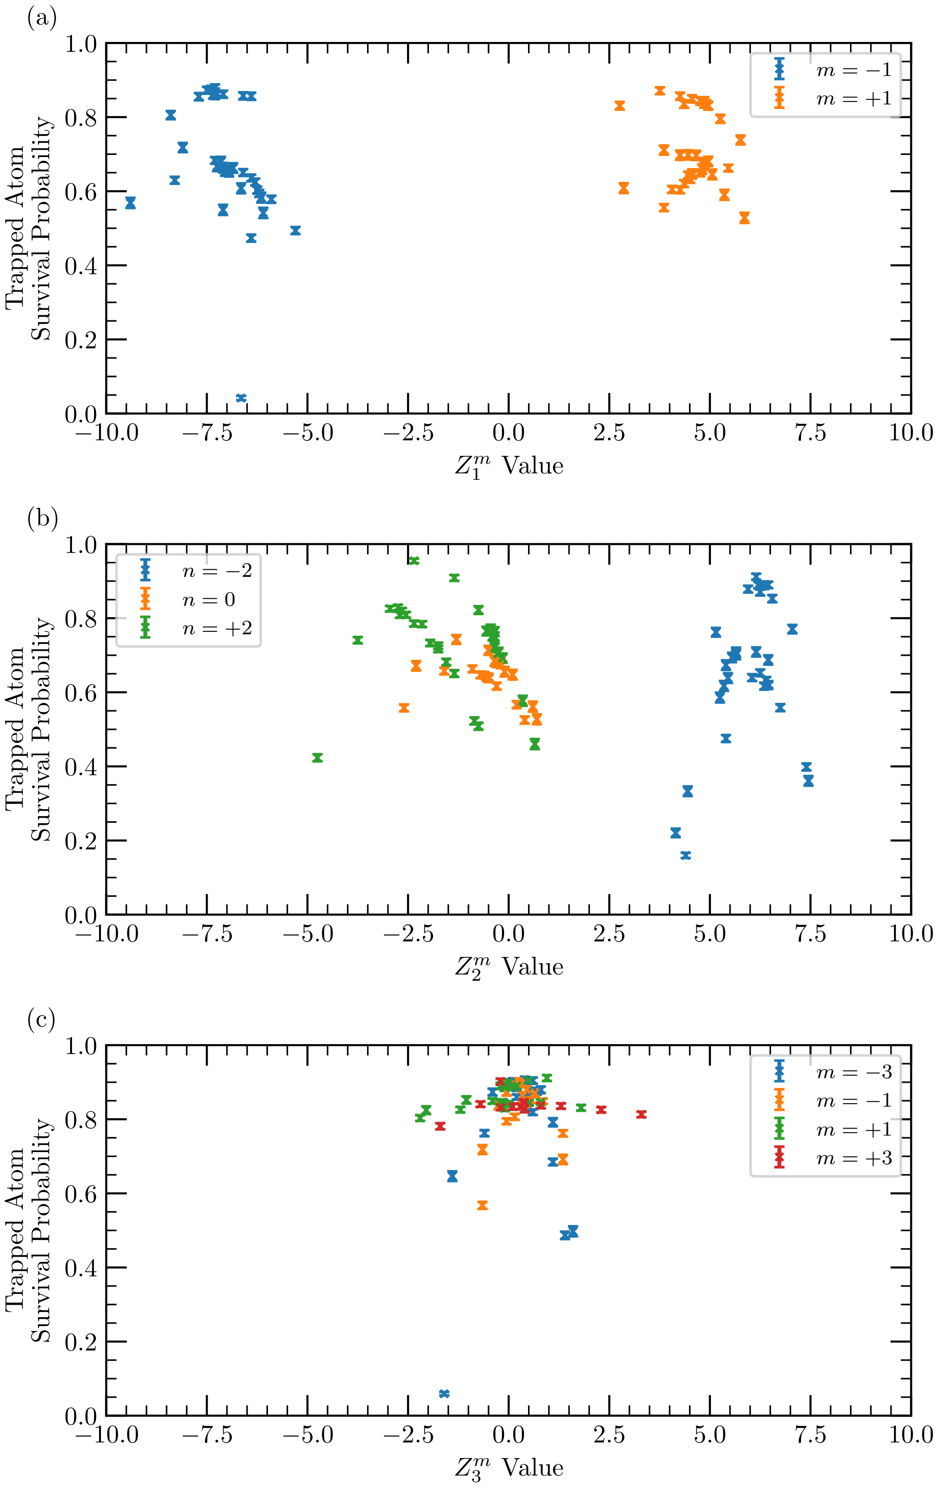

Figure A.1: Combined scans for Zernike polynomial alignment. Figures (a), (b) and (c) show scans for Z_1^m, Z_2^m and Z_3^m respectively. Figure (c) shows that by these higher orders, alignment is either insensitive over a wide range or must remain at zero. Higher order Zernike polynomials are thus neglected.

Bibliography

[1] D. Coyle and L. Hampton, "21st century progress in computing," *Telecommunications Policy*, vol. 48, p. 102649, 2 2024.

[2] S. McArdle *et al.*, "Quantum computational chemistry," *Reviews of Modern Physics*, vol. 92, p. 015003, 3 2020.

[3] H. Liu *et al.*, "Prospects of quantum computing for molecular sciences," *Materials Theory*, vol. 6, p. 11, 3 2022.

[4] P. S. Emani *et al.*, "Quantum computing at the frontiers of biological sciences," *Nature Methods*, vol. 18, pp. 701–709, 7 2021.

[5] P. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proceedings 35th Annual Symposium on Foundations of Computer Science.* IEEE Comput. Soc. Press, 1994, pp. 124–134.

[6] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing - STOC '96.* ACM Press, 1996, pp. 212–219.

[7] R. Landauer, "Irreversibility and heat generation in the computing process," *IBM Journal of Research and Development*, vol. 5, pp. 183–191, 7 1961.

[8] M. Horodecki *et al.*, "Partial quantum information," *Nature*, vol. 436, pp. 673–676, 8 2005.

[9] D. Deutsch, “Quantum theory, the church–turing principle and the universal quantum computer,” *Proc. R. Soc. Lond. A*, vol. 400, pp. 97–117, 7 1985.

[10] D. P. DiVincenzo, “The physical implementation of quantum computation,” *Fortschritte der Physik*, vol. 48, pp. 771–783, 9 2000.

[11] J. Preskill, “Quantum computing in the nisq era and beyond,” *Quantum*, vol. 2, p. 79, 8 2018.

[12] I. Buluta and F. Nori, “Quantum simulators,” *Science*, vol. 326, pp. 108–111, 10 2009.

[13] T. Bravo *et al.*, “Analog quantum simulation of gravitational waves in a bose-einstein condensate,” *EPJ Quantum Technology*, vol. 2, p. 3, 12 2015.

[14] J. R. M. de Nova *et al.*, “Observation of thermal hawking radiation and its temperature in an analogue black hole,” *Nature*, vol. 569, pp. 688–691, 5 2019.

[15] V. I. Kolobov *et al.*, “Observation of stationary spontaneous hawking radiation and the time evolution of an analogue black hole,” *Nature Physics*, vol. 17, pp. 362–367, 3 2021.

[16] E. Zohar *et al.*, “Quantum simulations of lattice gauge theories using ultracold atoms in optical lattices,” *Reports on Progress in Physics*, vol. 79, p. 014401, 1 2016.

[17] C. Gross and I. Bloch, “Quantum simulations with ultracold atoms in optical lattices,” *Science*, vol. 357, pp. 995–1001, 9 2017.

[18] J. Koepsell *et al.*, “Imaging magnetic polarons in the doped fermi–hubbard model,” *Nature*, vol. 572, pp. 358–362, 8 2019.

[19] F. Schäfer *et al.*, "Tools for quantum simulation with ultracold atoms in optical lattices," *Nature Reviews Physics*, vol. 2, pp. 411–425, 7 2020.

[20] I. Bloch *et al.*, "Quantum simulations with ultracold quantum gases," *Nature Physics*, vol. 8, pp. 267–276, 4 2012.

[21] X. Li *et al.*, "Observation and quantification of the pseudogap in unitary fermi gases," *Nature*, vol. 626, pp. 288–293, 2 2024.

[22] V. M. Kendon *et al.*, "Quantum analogue computing," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, pp. 3609–3620, 8 2010.

[23] A. J. Daley *et al.*, "Practical quantum advantage in quantum simulation," *Nature*, vol. 607, pp. 667–676, 7 2022.

[24] A. L. Shaw *et al.*, "Benchmarking highly entangled states on a 60-atom analogue quantum simulator," *Nature*, vol. 628, pp. 71–77, 4 2024.

[25] B. Fauseweh, "Quantum many-body simulations on digital quantum computers: State-of-the-art and future challenges," *Nature Communications*, vol. 15, p. 2123, 3 2024.

[26] F. Arute *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, pp. 505–510, 10 2019.

[27] H.-S. Zhong *et al.*, "Quantum computational advantage using photons," *Science*, vol. 370, pp. 1460–1463, 12 2020.

[28] Y. Wu *et al.*, "Strong quantum computational advantage using a superconducting quantum processor," *Physical Review Letters*, vol. 127, p. 180501, 10 2021.

[29] D. Aharonov *et al.*, "Adiabatic quantum computation is equivalent to standard quantum computation," *SIAM Review*, vol. 50, pp. 755–787, 1 2008.

[30] C. Wang *et al.*, “Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds,” *npj Quantum Information*, vol. 8, p. 3, 1 2022.

[31] Y. Sunada *et al.*, “Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic purcell filter,” *Physical Review Applied*, vol. 17, p. 044016, 4 2022.

[32] L. Chen *et al.*, “Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier,” *npj Quantum Information*, vol. 9, p. 26, 3 2023.

[33] Y. Sung *et al.*, “Realization of high-fidelity cz and zz-free iswap gates with a tunable coupler,” *Physical Review X*, vol. 11, p. 021058, 6 2021.

[34] P. Jurcevic *et al.*, “Demonstration of quantum volume 64 on a superconducting quantum computing system,” *Quantum Science and Technology*, vol. 6, p. 025020, 4 2021.

[35] A. D. King *et al.*, “Scaling advantage over path-integral monte carlo in quantum simulation of geometrically frustrated magnets,” *Nature Communications*, vol. 12, p. 1113, 2 2021.

[36] A. K. Fedorov *et al.*, “Quantum computing at the quantum advantage threshold: a down-to-business review,” 3 2022.

[37] J. Gambetta, “Ibm quantum roadmap,” 2024. [Online]. Available: QuantumRoadmap

[38] P. Wang *et al.*, “Single ion qubit with estimated coherence time exceeding one hour,” *Nature Communications*, vol. 12, p. 233, 1 2021.

[39] T. Harty *et al.*, “High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit,” *Physical Review Letters*, vol. 113, p. 220501, 11 2014.

[40] C. R. Clark *et al.*, “High-fidelity bell-state preparation with ca-40 + optical qubits,” *Physical Review Letters*, vol. 127, p. 130505, 9 2021.

[41] R. Srinivas *et al.*, “High-fidelity laser-free universal control of trapped ion qubits,” *Nature*, vol. 597, pp. 209–213, 9 2021.

[42] C. M. Löschnauer *et al.*, “Scalable, high-fidelity all-electronic control of trapped-ion qubits,” 7 2024.

[43] M. Malinowski *et al.*, “How to wire a 1000-qubit trapped-ion quantum computer,” *PRX Quantum*, vol. 4, p. 040313, 10 2023.

[44] S. IonQ, “Ionq forte: The first software-configurable quantum computer,” 3 2023. [Online]. Available: https://ionq.com/resources/ionq-forte-first-configurable-quantum-computer

[45] A. Aspuru-Guzik and P. Walther, “Photonic quantum simulators,” *Nature Physics*, vol. 8, pp. 285–291, 4 2012.

[46] J. E. Bourassa *et al.*, “Blueprint for a scalable photonic fault-tolerant quantum computer,” *Quantum*, vol. 5, p. 392, 2 2021.

[47] J. Zhang *et al.*, “Efficient implementation of a quantum algorithm in a single nitrogen-vacancy center of diamond,” *Physical Review Letters*, vol. 125, p. 030501, 7 2020.

[48] P. Rembold *et al.*, “Introduction to quantum optimal control for quantum sensing with nitrogen-vacancy centers in diamond,” *AVS Quantum Science*, vol. 2, 6 2020.

[49] D. DeMille, “Quantum computation with trapped polar molecules,” *Physical Review Letters*, vol. 88, p. 067901, 1 2002.

[50] S. L. Cornish *et al.*, “Quantum computation and quantum simulation with ultracold molecules,” *Nature Physics*, vol. 20, pp. 730–740, 5 2024.

[51] G. Burkard *et al.*, "Semiconductor spin qubits," *Reviews of Modern Physics*, vol. 95, p. 025003, 6 2023.

[52] M. Saffman *et al.*, "Quantum information with rydberg atoms," *Reviews of Modern Physics*, vol. 82, pp. 2313–2363, 8 2010.

[53] C. S. Adams *et al.*, "Rydberg atom quantum technologies," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 53, 2020.

[54] M. N. H. Chow *et al.*, "High-fidelity low-loss state detection of alkali-metal atoms in optical tweezer traps," *Physical Review A*, vol. 108, p. 032407, 9 2023.

[55] H. J. Manetsch *et al.*, "A tweezer array with 6100 highly coherent atomic qubits," 3 2024. [Online]. Available: http://arxiv.org/abs/2403.12021

[56] N. Schlosser *et al.*, "Sub-poissonian loading of single atoms in a microscopic dipole trap," *Letters To Nature*, vol. 411, pp. 1024–1027, 2001.

[57] M. Endres *et al.*, "Atom-by-atom assembly of defect-free one-dimensional cold atom arrays," 2016. [Online]. Available: https://www.science.org

[58] D. Barredo *et al.*, "An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays," *Science*, vol. 354, pp. 1021–1023, 11 2016, doi: 10.1126/science.aah3778. [Online]. Available: https://doi.org/10.1126/science.aah3778

[59] ——, "Synthetic three-dimensional atomic structures assembled atom by atom," *Nature*, vol. 561, pp. 79–82, 9 2018.

[60] P. Scholl *et al.*, "Erasure conversion in a high-fidelity rydberg quantum simulator," *Nature*, vol. 622, pp. 273–278, 10 2023.

[61] W. Huie *et al.*, “Repetitive readout and real-time control of nuclear spin qubits in yb-171 atoms,” *PRX Quantum*, vol. 4, p. 030337, 9 2023.

[62] M. Mohan *et al.*, “Robust control and optimal rydberg states for neutral atom two-qubit gates,” *Physical Review Research*, vol. 5, p. 033052, 7 2023.

[63] R. Song *et al.*, “Lifetime measurement of cesium atoms using a cold rydberg gas,” *Applied Sciences (Switzerland)*, vol. 12, 3 2022.

[64] B. Nikolov *et al.*, “Randomized benchmarking using nondestructive readout in a two-dimensional atom array,” *Physical Review Letters*, vol. 131, 7 2023.

[65] S. J. Evered *et al.*, “High-fidelity parallel entangling gates on a neutral-atom quantum computer,” *Nature*, vol. 622, pp. 268–272, 10 2023.

[66] Q. Xu *et al.*, “Constant-overhead fault-tolerant quantum computation with reconfigurable atom arrays,” *Nature Physics*, vol. 20, pp. 1084–1090, 7 2024.

[67] D. Bluvstein *et al.*, “A quantum processor based on coherent transport of entangled atom arrays,” *Nature*, vol. 604, pp. 451–456, 2022. [Online]. Available: https://doi.org/10.1038/s41586-022-04592-6

[68] ——, “Logical quantum processor based on reconfigurable atom arrays,” *Nature*, 12 2023.

[69] T. M. Graham *et al.*, “Multi-qubit entanglement and algorithms on a neutral-atom quantum computer,” *Nature*, vol. 604, pp. 457–462, 4 2022.

[70] T. Graham *et al.*, “Midcircuit measurements on a single-species neutral alkali atom quantum processor,” *Physical Review X*, vol. 13, p. 041051, 12 2023.

[71] M. Kornjača *et al.*, “Large-scale quantum reservoir learning with an analog quantum computer,” 7 2024.

[72] G. Pichard *et al.*, “Rearrangement of single atoms in a 2000-site optical tweezers array at cryogenic temperatures,” 5 2024.

[73] A. Browaeys and T. Lahaye, “Many-body physics with individually controlled rydberg atoms,” *Nature Physics*, vol. 16, pp. 132–142, 2020. [Online]. Available: https://doi.org/10.1038/s41567-019-0733-z

[74] H. Labuhn *et al.*, “Tunable two-dimensional arrays of single rydberg atoms for realizing quantum ising models,” *Nature*, vol. 534, pp. 667–670, 6 2016.

[75] H. Bernien *et al.*, “Probing many-body dynamics on a 51-atom quantum simulator,” *Nature*, vol. 551, pp. 579–584, 11 2017.

[76] C. Chen *et al.*, “Continuous symmetry breaking in a two-dimensional rydberg array,” *Nature*, vol. 616, pp. 691–695, 4 2023.

[77] E. C. Domanti *et al.*, “Floquet-rydberg quantum simulator for confinement in z2 gauge theories,” *Physical Review Research*, vol. 6, p. L022059, 6 2024.

[78] P. Scholl *et al.*, “Microwave engineering of programmable xxz hamiltonians in arrays of rydberg atoms,” *PRX Quantum*, vol. 3, p. 020303, 4 2022.

[79] G. Bornet *et al.*, “Scalable spin squeezing in a dipolar rydberg atom array,” *Nature*, vol. 621, pp. 728–733, 9 2023.

[80] S. Weber *et al.*, “Topologically protected edge states in small rydberg systems,” *Quantum Science and Technology*, vol. 3, p. 044001, 10 2018.

[81] S. de Léséleuc *et al.*, “Observation of a symmetry-protected topological phase of interacting bosons with rydberg atoms,” *Science*, vol. 365, pp. 775–780, 8 2019.

[82] G. Semeghini *et al.*, “Probing topological spin liquids on a programmable quantum simulator,” *Science*, vol. 374, pp. 1242–1247, 12 2021, doi: 10.1126/science.abi8794. [Online]. Available: https://doi.org/10.1126/science.abi8794

[83] J. Park *et al.*, “Rydberg-atom experiment for the integer factorization problem,” *Physical Review Research*, vol. 6, p. 023241, 6 2024.

[84] S. Ebadi *et al.*, “Quantum optimization of maximum independent set using rydberg atom arrays,” *Science*, vol. 376, pp. 1209–1215, 6 2022, doi: 10.1126/science.abo6587. [Online]. Available: https://doi.org/10.1126/science.abo6587

[85] M. T. Nguyen *et al.*, “Quantum optimization with arbitrary connectivity using rydberg atom arrays,” *PRX Quantum*, vol. 4, 1 2023.

[86] A. G. de Oliveira *et al.*, “Demonstration of weighted graph optimization on a rydberg atom array using local light-shifts,” 4 2024.

[87] B. Nikolov, “Square: Scalable qubit arrays for quantum computation and optimisation,” 12 2024.

[88] W. D. Worthy, “Fundamentals of an electronic digital computer. an explanatory introduction for engineers of the principles of operation of a computer,” *Aircraft Engineering and Aerospace Technology*, 1955.

[89] L. Essen and J. V. L. Parry, “An atomic standard of frequency and time interval: A cæsium resonator,” *Nature*, vol. 176, pp. 280–282, 8 1955.

[90] J. J. Sakurai and J. Napolitano, *Modern Quantum Mechanics*, 2nd ed. Addison-Wesley, 2011.

[91] D. A. Steck, "Cesium d line data," 1998. [Online]. Available: http://steck.us/alkalidata,

[92] P. O'Mahony and K. Taylor, "The quadratic zeeman effect in caesium: departures from hydrogenic behaviour," *Journal of Physics B: Atomic and Molecular Physics*, 1986.

[93] D. S. Kliger *et al.*, *Polarized Light In Optics And Spectroscopy*, 1st ed. Academic Press Limited, 1990.

[94] R. Loudon, *The quantum theory of light*, 3rd ed. Oxford University Press, 2000.

[95] C. Foot, *Atomic Physics*, 1st ed. Oxford University Press, 2005.

[96] J. R. Johansson *et al.*, "Qutip 2: A python framework for the dynamics of open quantum systems," *Computer Physics Communications*, vol. 184, pp. 1234–1240, 4 2013.

[97] C. N. Cohen-Tannoudji, "Manipulating atoms with photons," *Rev. Mod. Phys.*, vol. 70, 1998.

[98] 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*, vol. 29, pp. 331–337, 4 1890.

[99] N. Bohr, "Xxxvii. on the constitution of atoms and molecules," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 26, pp. 476–502, 9 1913.

[100] J. Pritchard and David, "Durham e-theses cooperative optical nonlinearity in a blockaded rydberg ensemble," 2011. [Online]. Available: http://etheses.dur.ac.uk/782/

[101] N. Šibalić *et al.*, "Arc: An open-source library for calculating properties of alkali rydberg atoms," *Computer Physics Communications*, vol. 220,

pp. 319–331, 2017. [Online]. Available: http://matplotlib.org/.http://creativecommons.org/licenses/by/4.0/

[102] T. F. Gallagher, *Rydberg Atoms*, 1st ed. Cambridge University Press, 1994.

[103] A. Gaëtan *et al.*, “Observation of collective excitation of two individual atoms in the rydberg blockade regime,” *Nature Physics*, vol. 5, pp. 115–118, 2009.

[104] M. Schlosser *et al.*, “Assembled arrays of rydberg-interacting atoms,” *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 53, 7 2020.

[105] W. Dür, “Multipartite entanglement that is robust against disposal of particles,” *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 63, pp. 1–4, 2001.

[106] ——, “Entanglement molecules,” 6 2000. [Online]. Available: https://arxiv.org/abs/quant-ph/0006105

[107] M. Koashi *et al.*, “Entangled webs: Tight bound for symmetric sharing of entanglement,” *Phys Rev A*, vol. 62, 2000.

[108] J. Wang *et al.*, “Single-photon cesium rydberg excitation spectroscopy using 3186-nm uv laser and room-temperature vapor cell,” *Optics Express*, vol. 25, p. 22510, 9 2017.

[109] I. I. Beterov *et al.*, “Quasiclassical calculations of blackbody-radiation-induced depopulation rates and effective lifetimes of rydberg ns, np, and nd alkali-metal atoms with n80,” *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 79, 5 2009.

[110] G. Toh *et al.*, “Measurement of the lifetimes of the 7p 3/2 and 7p 1/2 states of atomic cesium,” *Physical Review A*, vol. 100, 11 2019.

[111] E. Brion *et al.*, “Adiabatic elimination in a lambda system,” *Journal of Physics A: Mathematical and Theoretical*, vol. 40, pp. 1033–1043, 2 2007.

[112] P. D. Lett *et al.*, “Optical molasses,” 1989.

[113] A. M. Kaufman *et al.*, “Cooling a single atom in an optical tweezer to its quantum ground state,” *Physical Review X*, vol. 2, p. 041014, 11 2012.

[114] Y. Yu *et al.*, “Motional-ground-state cooling outside the lamb-dicke regime,” *Physical Review A*, vol. 97, p. 063423, 6 2018.

[115] Y. Bao *et al.*, “Raman sideband cooling of molecules in an optical tweezer array to the 3d motional ground state,” *Physical Review X*, vol. 14, p. 031002, 7 2024.

[116] Y. Lu *et al.*, “Raman sideband cooling of molecules in an optical tweezer array,” *Nature Physics*, vol. 20, pp. 389–394, 3 2024.

[117] I. Estermann, “History of molecular beam research: Personal reminiscences of the important evolutionary period 1919–1933,” *American Journal of Physics*, vol. 43, pp. 661–671, 8 1975.

[118] E. L. Raab *et al.*, “Trapping of neutral sodium atoms with radiation pressure,” 1987.

[119] A. B. Deb *et al.*, “Versatile laser system for experiments with cold atomic gases,” *J. Opt. Soc. Am. B*, vol. 29, pp. 3109–3113, 11 2012.

[120] A. Camara *et al.*, “Scaling behavior of a very large magneto-optical trap,” *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 90, 12 2014.

[121] J. Mahnke *et al.*, "A continuously pumped reservoir of ultracold atoms," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 48, 8 2015.

[122] P. D. Lett *et al.*, "Observation of atoms laser cooled below the doppler limit," 1988.

[123] J. Dalibard and C. Cohen-Tannoudji, "Laser cooling below the doppler limit by polarization gradients: simple theoretical models," 1989.

[124] R. Grimm *et al.*, "Optical dipole traps for neutral atoms," 1999.

[125] N. Davidson *et al.*, "Long atomic coherence times in an optical dipole trap," 1995.

[126] T. Kuga *et al.*, "Novel optical trap of atoms with a doughnut beam," 1997.

[127] S. Tempone-Wiltshire *et al.*, "High efficiency, low cost holographic optical elements for ultracold atom trapping," *Optics Express*, vol. 25, p. 296, 1 2017.

[128] J. Arlt and M. J. Padgett, "Generation of a beam with a dark focus surrounded by regions of higher intensity: the optical bottle beam," 2000.

[129] L. Isenhower *et al.*, "Atom trapping in an interferometrically generated bottle beam trap," p. 1159, 2009.

[130] R. Ozeri *et al.*, "Long spin relaxation times in a single-beam blue-detuned optical trap," 1998.

[131] V. V. Ivanov *et al.*, "Atom trapping in a bottle beam created by a diffractive optical element," 5 2013. [Online]. Available: http://arxiv.org/abs/1305.5309

[132] D. Barredo *et al.*, "Three-dimensional trapping of individual rydberg atoms in ponderomotive bottle beam traps," *Physical Review Letters*, vol. 124, 1 2020.

[133] J. Weiner, *Cold and ultracold collisions in quantum microscopic and mesoscopic systems*, 1st ed. Cambridge University Press, 2003.

[134] Y. R. P. Sortais *et al.*, "Diffraction-limited optics for single-atom manipulation," *Physical Review A*, vol. 75, p. 13406, 1 2007. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.75.013406

[135] L. Béguin *et al.*, "Direct measurement of the van der waals interaction between two rydberg atoms," *Physical Review Letters*, vol. 110, 6 2013.

[136] C. J. Picken *et al.*, "Single atom imaging with an scmos camera," *Applied Physics Letters*, vol. 111, 10 2017.

[137] S. Bourdeauducq *et al.*, "Artiq 6.0," 2021. [Online]. Available: https://zenodo.org/records/6619071

[138] G. H. Zhang *et al.*, "Note: Fast compact laser shutter using a direct current motor and three-dimensional printing," *Review of Scientific Instruments*, vol. 86, 12 2015.

[139] A. Arnold, "Preparation and manipulation of an 87rb bose-einstein condensate," 11 1999.

[140] K. Ladouceur *et al.*, "Compact laser cooling apparatus for simultaneous cooling of lithium and rubidium," *Journal of the Optical Society of America B*, vol. 26, p. 210, 2 2009.

[141] E. Hecht, *Optics*, 4th ed. Pearson Education, 8 2013, vol. 1.

[142] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," pp. 237–246, 1969.

[143] D. Gabor, “A new microscopic principle,” *Nature*, pp. 777–778, 1948.

[144] D. Kim *et al.*, “Large-scale uniform optical focus array generation with a phase spatial light modulator,” *Optics Letters*, vol. 44, p. 3178, 6 2019.

[145] P. Schroff *et al.*, “Accurate holographic light potentials using pixel crosstalk modelling,” *Scientific Reports*, vol. 13, 12 2023.

[146] F. Nogrette *et al.*, “Single-atom trapping in holographic 2d arrays of microtraps with arbitrary geometries,” *Physical Review X*, vol. 4, p. 21034, 5 2014. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevX.4.021034

[147] S. Chu *et al.*, “Experimental observation of optically trapped atoms,” *Physical Review Letters*, vol. 57, 1986.

[148] C. Tuchendler *et al.*, “Energy distribution and cooling of a single atom in an optical tweezer,” *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 78, 9 2008.

[149] N. Bobroff and A. E. Rosenbluth, “Evaluation of highly corrected optics by measurement of the strehl ratio,” *Applied Optics*, vol. 31, pp. 1523–1536, 1992.

[150] Y. Bai *et al.*, “The application of low-frequency transition in the assessment of the second-order zeeman frequency shift,” *Sensors*, vol. 21, p. 8333, 12 2021.

[151] H. Levine *et al.*, “Parallel implementation of high-fidelity multiqubit gates with neutral atoms,” *Physical Review Letters*, vol. 123, 10 2019.

[152] S. D. Léséleuc *et al.*, “Analysis of imperfections in the coherent optical excitation of single atoms to rydberg states,” *Physical Review A*, vol. 97, 5 2018.

[153] G. Y. Belay *et al.*, "Dynamic optical beam shaping system to generate gaussian and top-hat laser beams of various sizes with circular and square footprint for additive manufacturing applications," *Procedia CIRP*, vol. 111, pp. 75–80, 2022.

[154] J. J. Snyder, "Paraxial ray analysis of a cat's-eye retroreflector," *Applied Optics*, vol. 14, p. 1825, 8 1975.

[155] H. Levine *et al.*, "High-fidelity control and entanglement of rydberg-atom qubits," *Physical Review Letters*, vol. 121, 9 2018.

[156] G. Sridhar *et al.*, "Cavity ring-down technique for measurement of reflectivity of high reflectivity mirrors with high accuracy," *Pramana*, vol. 75, pp. 1233–1239, 12 2010.

[157] M. Zhao *et al.*, "Afterpulsing and its correction in fluorescence correlation spectroscopy experiments," *Applied Optics*, vol. 42, p. 4031, 7 2003.

[158] E. D. Black, "An introduction to pound–drever–hall laser frequency stabilization," *American Journal of Physics*, vol. 69, 2001.

[159] P. D. Gregory *et al.*, "A simple, versatile laser system for the creation of ultracold ground state molecules," *New Journal of Physics*, vol. 17, p. 055006, 5 2015.

[160] M. O. Brown *et al.*, "Gray-molasses optical-tweezer loading: Controlling collisions for scaling atom-array assembly," *Physical Review X*, vol. 9, 3 2019.

[161] W. Tian *et al.*, "Parallel assembly of arbitrary defect-free atom arrays with a multitweezer algorithm," *Physical Review Applied*, vol. 19, 3 2023.

[162] H. J. Levine, "Quantum information processing and quantum simulation with programmable rydberg atom arrays," 4 2021.

[163] S. Ebadi *et al.*, "Quantum phases of matter on a 256-atom programmable quantum simulator," *Nature*, vol. 595, 2021.

[164] K.-N. Schymik, "Scaling-up the tweezer platform-trapping arrays of single atoms in a cryogenic environment," 3 2022.

[165] A. Kumar *et al.*, "Sorting ultracold atoms in a three-dimensional optical lattice in a realization of maxwell's demon," *Nature*, vol. 561, pp. 83–87, 9 2018.

[166] D. O. D. Mello *et al.*, "Defect-free assembly of 2d clusters of more than 100 single-atom quantum systems," *Physical Review Letters*, vol. 122, 5 2019.

[167] L. Pause *et al.*, "Supercharged two-dimensional tweezer array with more than 1000 atomic qubits," *Optica*, vol. 11, p. 222, 2 2024.

[168] C. Sheng *et al.*, "Efficient preparation of two-dimensional defect-free atom arrays with near-fewest sorting-atom moves," *Physical Review Research*, vol. 3, 4 2021.

[169] F. Gyger *et al.*, "Continuous operation of large-scale atom arrays in optical lattices," *Physical Review Research*, vol. 6, p. 033104, 7 2024.

[170] M. A. Norcia *et al.*, "Iterative assembly of yb-171 atom arrays with cavity-enhanced optical lattices," *PRX Quantum*, vol. 5, p. 030316, 7 2024.

[171] H. Kim *et al.*, "In situ single-atom array synthesis using dynamic holographic optical tweezers," *Nature Communications*, vol. 7, p. 13317, 10 2016.

[172] T. Grünzweig *et al.*, "Near-deterministic preparation of a single atom in an optical microtrap," *Nature Physics*, vol. 6, pp. 951–954, 2010.

[173] M. M. Aliyu *et al.*, “D1 magic wavelength tweezers for scaling atom arrays,” *Physical Review Research*, vol. 3, 12 2021.

[174] K. N. Schymik *et al.*, “In situ equalization of single-atom loading in large-scale optical tweezer arrays,” *Physical Review A*, vol. 106, 8 2022.

[175] ——, “Enhanced atom-by-atom assembly of arbitrary tweezer arrays,” *Physical Review A*, vol. 102, 12 2020.

[176] P. Hart *et al.*, “A formal basis for the heuristic determination of minimum cost paths,” *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, pp. 100–107, 1968.

[177] N. D. Matsakis and F. S. Klock, “The rust language,” *ACM SIGAda Ada Letters*, vol. 34, pp. 103–104, 11 2014.

[178] C. Weitenberg *et al.*, “Quantum computation architecture using optical tweezers,” *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 84, 9 2011.

[179] R. Shadmehr and S. Wise, *The computational neurobiology of reaching and pointing: a foundation for motor learning*, 1st ed. The MIT PRess, 2005.

[180] L. R. Liu *et al.*, “Molecular assembly of ground-state cooled single atoms,” *Physical Review X*, vol. 9, 5 2019.

[181] S. Spence *et al.*, “Preparation of 87rb and 133cs in the motional ground state of a single optical tweezer,” *New Journal of Physics*, vol. 24, 10 2022.

[182] Z. Zhang *et al.*, “A cryogenic neutral atom optical tweezer array,” in *55th Annual Meeting of the APS Division of Atomic, Molecular and Optical Physics*. Bulletin Of The American Physical Society, 6 2024.

[183] K. N. Schymik *et al.*, "Single atoms with 6000-second trapping lifetimes in optical-tweezer arrays at cryogenic temperatures," *Physical Review Applied*, vol. 16, 9 2021.

[184] W. D. Williams *et al.*, "Spectroscopic study of the 7p 1/2 and 7p 3/2 states in cesium-133," *Laser Physics Letters*, vol. 15, p. 095702, 9 2018.

[185] R. Legaie, "Coherent control of rydberg atoms using sub-khz linewidth excitation lasers," 2020.

[186] S. Weber *et al.*, "Calculation of rydberg interaction potentials," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 50, p. 133001, 7 2017.

[187] S. L. Braunstein *et al.*, "Maximal violation of bell inequalities for mixed states," *Physical Review Letters*, vol. 68, pp. 3259–3261, 1992.

[188] T. Wilk *et al.*, "Entanglement of two individual neutral atoms using rydberg blockade," *Physical Review Letters*, vol. 104, p. 010502, 1 2010.

[189] D. Bluvstein *et al.*, "Controlling quantum many-body dynamics in driven rydberg atom arrays," *Science*, vol. 371, pp. 1355–1359, 3 2021.

[190] A. Lukin *et al.*, "Quantum quench dynamics as a shortcut to adiabaticity," 5 2024.

[191] C. J. Turner *et al.*, "Quantum scarred eigenstates in a rydberg atom chain: Entanglement, breakdown of thermalization, and stability to perturbations," *Physical Review B*, vol. 98, p. 155134, 10 2018.

[192] M. J. Piotrowicz *et al.*, "Two-dimensional lattice of blue-detuned atom traps using a projected gaussian beam array," *Physical Review A - Atomic, Molecular, and Optical Physics*, vol. 88, 7 2013.

[193] T. Xia *et al.*, "Randomized benchmarking of single-qubit gates in a 2d array of neutral-atom qubits," *Physical Review Letters*, vol. 114, p. 100503, 3 2015.

[194] Y. Wang *et al.*, "Coherent addressing of individual neutral atoms in a 3d optical lattice," *Physical Review Letters*, vol. 115, p. 043003, 7 2015.

[195] T. M. Graham *et al.*, "Rydberg-mediated entanglement in a two-dimensional neutral atom qubit array," *Physical Review Letters*, vol. 123, 12 2019.

[196] S. K. Dutta *et al.*, "Ponderomotive optical lattice for rydberg atoms," 2000.

[197] K. Khare, *Fourier Optics and Computational Imaging.* Wiley, 9 2015.

[198] P. Vingelmann and F. H. P. Fitzek, "Cuda, release: 10.2.89," 2020.

[199] W. Press *et al.*, *Numerical Recipes*, 1st ed. Cambridge University Press, 1987.

[200] J. Schwiegerling, "Scaling zernike expansion coefficients to different pupil sizes," *Journal of the Optical Society of America A*, vol. 19, no. 10, p. 1937, 10 2002.

[201] A. Arias and R. Castaneda, "Estimation and compensation of aberrations in spatial light modulators," *Journal of Physics: Conference Series*, vol. 274, p. 012060, 1 2011.

[202] M. Baum, "Towards dynamic holographic laser beam shaping," *Journal of Laser Micro/Nanoengineering*, vol. 10, pp. 216–221, 5 2015.

[203] T. Häfner *et al.*, "Tailored laser beam shaping for efficient and accurate microstructuring," *Applied Physics A*, vol. 124, p. 111, 2 2018.

[204] T. W. Clark *et al.*, "Comparison of beam generation techniques using a phase only spatial light modulator," *Optics Express*, vol. 24, p. 6249, 3 2016.

[205] F. Harary, *Graph Theory*, 1st ed. Addison-Wesley, 1969.

[206] D. Davendra, *Traveling Salesman Problem, Theory and Applications*, D. Davendra, Ed. InTech, 12 2010.

[207] V. Mavroeidis *et al.*, "The impact of quantum computing on present cryptography," *International Journal of Advanced Computer Science and Applications*, vol. 9, 2018.

[208] C. Zalka, "Grover's quantum searching algorithm is optimal," *Physical Review A*, vol. 60, pp. 2746–2751, 10 1999.

[209] K. Erciyes, *Distributed Graph Algorithms for Computer Networks*, 1st ed., A. J. Sammes, Ed. Springer-Verlag, 2013, vol. 1.

[210] E. J. van Leeuwen, *Graph-Theoretic Concepts in Computer Science*, D. Kratsch, Ed. Springer Berlin Heidelberg, 2005, vol. 3787.

[211] R. S. Andrist *et al.*, "Hardness of the maximum-independent-set problem on unit-disk graphs and prospects for quantum speedups," *Physical Review Research*, vol. 5, p. 043277, 12 2023.

[212] S. Arora *et al.*, "Proof verification and the hardness of approximation problems," *Journal of the ACM*, vol. 45, pp. 501–555, 5 1998.

[213] A. Kasperski and P. Zieliński, "Complexity of the robust weighted independent set problems on interval graphs," *Optimization Letters*, vol. 9, pp. 427–436, 3 2015.

[214] G. K. Das *et al.*, "Efficient independent set approximation in unit disk graphs," *Discrete Applied Mathematics*, vol. 280, pp. 63–70, 6 2020.

[215] E. Farhi *et al.*, “Quantum computation by adiabatic evolution,” 2000. [Online]. Available: https://arxiv.org/abs/quant-ph/0001106

[216] ——, “A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem,” *Science*, vol. 292, pp. 472–475, 4 2001.

[217] K. Kim *et al.*, “Quantum computing dataset of maximum independent set problem on king lattice of over hundred rydberg atoms,” *Scientific Data*, vol. 11, p. 111, 1 2024.

[218] P. B. Wigley *et al.*, “Fast machine-learning online optimization of ultracold-atom experiments,” *Scientific Reports*, vol. 6, p. 25890, 5 2016.

[219] L. Leone *et al.*, “Stabilizer rényi entropy,” *Physical Review Letters*, vol. 128, p. 050402, 2 2022.

[220] K. Goto *et al.*, “Probing chaos by magic monotones,” *Physical Review D*, vol. 106, p. 126009, 12 2022.

[221] S. F. E. Oliviero *et al.*, “Measuring magic on a quantum processor,” *npj Quantum Information*, vol. 8, p. 148, 12 2022.

[222] P. S. Tarabunga, “Critical behaviors of non-stabilizerness in quantum spin chains,” *Quantum*, vol. 8, p. 1413, 7 2024.

[223] W. H. Zurek, “Cosmological experiments in superfluid helium?” *Nature*, vol. 317, pp. 505–508, 10 1985.

[224] G. Biroli *et al.*, “Kibble-zurek mechanism and infinitely slow annealing through critical points,” *Physical Review E*, vol. 81, p. 050101, 5 2010.

[225] A. W. Glaetzle *et al.*, “Quantum spin-ice and dimer models with rydberg atoms,” *Physical Review X*, vol. 4, 2014.

[226] M. Kim *et al.*, “Rydberg quantum wires for maximum independent set problems,” *Nature Physics*, vol. 18, pp. 755–759, 7 2022.

[227] Z. Zeng *et al.*, "Quantum dimer models with rydberg gadgets," 2 2024. [Online]. Available: http://arxiv.org/abs/2402.10651

[228] J. Wurtz *et al.*, "Industry applications of neutral-atom quantum computing solving independent set problems," 5 2022.

[229] T. Manovitz *et al.*, "Quantum coarsening and collective dynamics on a programmable quantum simulator," 7 2024.

[230] D. Guéry-Odelin *et al.*, "Shortcuts to adiabaticity: Concepts, methods, and applications," *Reviews of Modern Physics*, vol. 91, 10 2019.

[231] C. Whitty *et al.*, "Quantum control via enhanced shortcuts to adiabaticity," *Physical Review Research*, vol. 2, p. 023360, 6 2020.

[232] C.-J. Lin *et al.*, "Quantum many-body scar states in two-dimensional rydberg atom arrays," *Physical Review B*, vol. 101, p. 220304, 6 2020.

[233] B. Andrade *et al.*, "Observing quantum many-body scars in random quantum circuits," *Physical Review A*, vol. 109, p. 052602, 5 2024.

[234] G. Pelegrí *et al.*, "High-fidelity multiqubit rydberg gates via two-photon adiabatic rapid passage," *Quantum Science and Technology*, 2022.

[235] J. Tu and S. Subhankar, "Fast real-time arbitrary waveform generation using graphic processing units," 3 2024.

[236] F. Tennie *et al.*, "Quantum computing for nonlinear differential equations and turbulence," 6 2024.

[237] T. M. Graham *et al.*, "Multiscale architecture for fast optical addressing and control of large-scale qubit arrays," *Applied Optics*, vol. 62, p. 3242, 4 2023.

[238] B. Zhang *et al.*, "Scaled local gate controller for optically addressed qubits," *Optica*, vol. 11, p. 227, 2 2024.

[239] Q. Xu *et al.*, "Fast and parallelizable logical computation with homological product codes," 7 2024.

[240] C. Poole *et al.*, "Architecture for fast implementation of qldpc codes with optimized rydberg gates," 4 2024.

[241] J. M. Auger *et al.*, "Blueprint for fault-tolerant quantum computation with rydberg atoms," *Physical Review A*, vol. 96, p. 052320, 11 2017.

[242] P. M. Ireland *et al.*, "Interspecies förster resonances for rb-cs rydberg d-states for enhanced multi-qubit gate fidelities," *Physical Review Research*, vol. 6, p. 013293, 3 2024.

[243] C. Cao *et al.*, "Ab initio quantum simulation of strongly correlated materials with quantum embedding," *npj Computational Materials*, vol. 9, p. 78, 5 2023.

[244] A. Michel, "Quantum simulation for strongly interacting fermions with neutral atoms array: towards the simulation of materials of interest," 6 2024.