Programmable arrays of alkaline earth atoms: qubits, clocks, and the Bose-Hubbard model

by

Aaron William Young

B.A., Wesleyan University, 2017

B.S., California Institute of Technology, 2017

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Department of Physics

2023

Committee Members:

Adam M. Kaufman, Chair

Joshua Combes

Dietrich G. Leibfried

Cindy A. Regal

Ana Maria Rey

Jun Ye

Young, Aaron William (Ph.D., Physics)

Programmable arrays of alkaline earth atoms: qubits, clocks, and the Bose-Hubbard model

Thesis directed by Prof. Adam M. Kaufman

In this thesis, I report on the development of a new experimental platform that features capabilities drawn from optical tweezer arrays, quantum gas microscopes, and optical clocks. We demonstrate that we can trap, cool, image, and manipulate the positions of individual strontium atoms in an optical lattice using a colocated set of optical tweezer arrays. These capabilities allow us to assemble ensembles of hundreds of atoms while maintaining access to controls and observables that have a resolution at the level of a single atom, and a single lattice site.

In the direction of quantum metrology, we demonstrate that the above capabilities can be combined with control of a long-lived optical frequency qubit encoded in the electronic states of strontium to realize a new kind of optical clock. This "tweezer array clock" provides access to measurements not typically available in optical clocks while also providing fractional frequency stabilities that are close to the state-of-the-art for synchronous comparisons. We further engineer interactions between optical frequency qubits encoded in different atoms by exciting the atoms to high-lying Rydberg states. We use the resulting interactions to perform entangling gates, and to prepare large entangled states that enable synchronous optical clock measurements with a precision below the standard quantum limit.

In the direction of quantum simulation, we control the motion of atoms in a tunnel-coupled optical lattice to study single- and multi-particle quantum walks. By taking advantage of the ability to modify the local potential in the lattice using optical tweezers, we show that these quantum walks can be used as a resource in various algorithms, including in the first experimental demonstration of spatial search by quantum walk. By further taking advantage of the ability to prepare, evolve, and detect large ensembles of atoms in the lattice with high fidelity, we significantly advance the state-of-the-art in Fock state boson sampling. Applying the above techniques to interacting atoms allows us to introduce a new approach to studying both ground states and dynamics in the Bose-Hubbard model. We present preliminary results involving the dynamics of hard core bosons, as well as the preparation of a superfluid that is assembled a

single atom at a time using optical tweezers.

The combination of high fidelity control of arrays of atomic qubits with complicated many body dynamics, and state-of-the-art capabilities in frequency metrology, results in a fruitful blurring of the lines between quantum computation, simulation, and metrology. As the capabilities of these systems continue to expand, we are hopeful that ideas from quantum information science can be drawn on to perform ever more precise measurements and, conversely, that metrological tools and techniques can be used as precision probes of complicated and poorly understood quantum many body effects.

Dedication

For mom and dad, who always answered why.

Acknowledgements

I would like to say that the first time I met Adam, we had a deep and stimulating conversation about the future of atomic physics. In actuality, having only communicated with him via email and not knowing his face, I asked him where the library was in the Lyman building, thanked him brusquely for his clear directions, rushed away, and promptly got lost. This interaction served as the blueprint for our work together in the years that followed. Adam is a brilliant, careful physicist who has weaponized his "neuroticism"[1] into an uncanny ability to see to the crux of a problem, and plan well in advance to avoid any pitfalls. It has been an absolute pleasure to work with him these last six years, even on those occasions when, despite careful planning, I get lost.

I feel very lucky to have worked with an incredible group of people that Adam has assembled over the years, including past and present members of the strontium experiment: Nathan Schine, Ben Johnston, Felix Rönchen, Nelson Darkwah-Oppong, Alec Cao, and Theo Lukin-Yelin. I am particularly indebted to two people involved in the early phases of the experiment: Matt Norcia, who taught me, by demonstration, much of what I know about experimental atomic physics, and is the main reason that this thesis isn't being written in 2027; and Will Eckner, who is a very strong physicist both technically, having built all the scariest lasers in the lab, and theoretically. More often than I care to admit, Will has listened patiently to one of my ad hoc explanations, and later discreetly sent me a detailed LaTeX document explaining why I was wrong.

I was similarly lucky with the fantastic people who guided me professionally early on, and led to my eventual decision to pursue a career in experimental physics. I am deeply grateful to all of them. To Greg MacIsaac, who instilled a desire to explore and document the world (through film, rather than atomic physics experiments). To Lutz Hüwel, who taught a film student to document the world with more precision

[1] His words not mine [162].

(and in much closer proximity to high power lasers). And to Oskar Painter, whose impact was twofold: he got a mechanical engineering student so excited about quantum science that they pretended that quantum mechanics courses were relevant to their degree,[2] and instilled the drive to "set lofty goals, and work goddamned hard."

For most of grad school, I have lived with the same group of close friends: Shawn Geller, who has provided years of stimulating conversations about physics that almost make up for the fact that he's actually a mathematician; Will Milner, who is somehow simultaneously the most hardworking and most laid back person I know; and of course Hannah Knaack, who has been an amazing partner, despite her insistence on working with one less electron than God intended.

I will spare you the stories about how I wanted to be a physicist ever since I was a child — that's not the truth. However, I will indulge in sharing a couple of anecdotes, and thanking my earliest supporters: My grandmother Mary Young, who sent each new issue of Popular Science to me from the other side of the world, without fail, well before I knew I had any interest in science. Perhaps she knew all along. My dad Russell Young, who, for years, patiently explained math to what I can only describe as a hostile audience. I wanted things to be self-evident, and connected to specific applications. He would counter by pulling tricks and puzzles out of nowhere that would miraculously fit together to reveal deep insights. This of course still annoys me endlessly. And finally my mom Lynn Hua, who taught me to be generous but effective with my time, and to cherish time with family above all else. "一寸光阴一寸金," which is why in this work we measure 光阴 to the 19^{th} decimal place.

Acknowledgements are always tough, because a complete list would fill this entire thesis and more. The following is an incomplete and unordered list[3] of people who have helped me significantly throughout my research, thank you all: Manny Knill, Scott Brown, Andrew Childs, Scott Glancy, Hans Green, Cindy Regal, Mark Brown, J. R. Raith, Kaizhao Wang, Eric Alvarado, Dhruv Kedar, Hope Whitelock, Michael Fang, Sean Muleady, Ana Maria Rey, Terry Brown, Jie Luo, Alec Jenkins, Alex Downham, Chris Reetz, Felix Vietmeyer, Charlie McLemore, Aruku Senoo, Jenny Wu, Paul Dieterle, Ting-Wei Hsu, Matteo Marinelli,

[2] Thanks to Guillaume Blanquart for letting this fly, and suggesting that scattering problems are "close enough to fluid mechanics and thermal transport" to justify graduation.

[3] Seriously don't read into the ordering, it's been through a random number generator.

Jun Ye, Josh Combes, Peter Zoller, Dietrich Leibfried, Joanna Lis, Lee Liu, Phillip Buelow, Ross Hutson, Jose D'Incao, John Robinson, and William McGrew. Finally, thanks to COVID-19, for making it easier to justify sitting in a windowless basement for most of my twenties.

Contents

Tables

Table

Figures

Figure

Chapter 1

Introduction

"Such awareness as there was/
and there was little
confined itself to spaces smaller
than the head of a pin
where angels once danced\\"

— AI Ummon

1.1 Historical perspective

In the early days of quantum mechanics, many of the properties of the new theory, in particular the intrinsic randomness and uncertainty at its core, were treated as a nuisance that flew in the face of well-behaved, deterministic classical dynamics. In the decades since there has been a shift in understanding: quantum systems, particularly ones made up of many constituents, allow for fundamentally different and often enhanced capabilities in comparison to their classical counterparts. An early hint at these capabilities was provided in 1964 by Bell's theorem [24], which states that quantum states can possess especially strong correlations that cannot be explained by any classical local hidden variable theory. Such correlations can be used in various tasks in communication [121] and sensing [83, 259] to gain explicit advantages over what can be accomplished classically. Another hint came from the general difficulties one encounters when trying to study quantum many body systems, where even the simplest microscopic models can end up being intractable to solve due to the rapid growth of Hilbert space allowed by entanglement. For example, despite intensive research [184], it is still unknown whether Hubbard's eponymous model from 1963 [144] possesses a d-wave superconducting ground state.[1] In contrast to the Fermi-Hubbard model, the ground (and thermal) states

[1] Although it is now suspected that the basic Hubbard model on a square lattice does not host such a ground state [310].

of the Bose-Hubbard model are typically well-understood based on mean field or quantum Monte-Carlo treatments [266, 30].[2] However, dynamics in any Hubbard model are generically believed to be hard to simulate. In recent years, our understanding of the sense in which such problems are hard has sharpened. In general, our best approximate classical methods for finding ground state or thermal expectation values of observables in an interacting, fermionic system are in the complexity class NP-hard [335], and solving for the full dynamics of even a non-interacting, bosonic system is #P-hard [2].[3] Although it is not believed that quantum systems will help us solve arbitrary NP problems or simulate arbitrary quantum behaviors in polynomial time,[4] it is reasonable to believe that they will outpace (or have outpaced) our ability to simulate the world classically [79, 334].[5] In the pursuit of realizing the above advantages, there has been a concentrated effort in recent years to build large, controllable quantum systems using a zoo of different technologies, including trapped atoms and ions, superconducting and photonic circuits, and quantum dots.

We will focus on experiments involving optically trapped and ultracold neutral atoms, which, when we began the work presented in this thesis, had already made impressive advances in three distinct domains: (1) in simulations of various Hubbard models, (2) in programmably controlling individual atoms using optical tweezers, and (3) in optical frequency metrology.

In (1), ultracold atomic gasses confined within optical lattices have been used to realize a variety of different Hubbard models [130]. In particular, the presence of the lattice potential allows one to reduce the kinetic energy of the atoms (by tuning the tunnel-coupling between different lattice sites) to reach regimes where the atoms are strongly interacting. Incorporating a high resolution microscope in such experiments allows one to perform measurements of the resulting strongly correlated states of matter with resolution on

[2] Due to the lack of a sign problem [335].

[3] In fact, even drawing approximate samples from the wavefunction prepared after dynamical evolution is believed to be hard [2].

[4] For example, the general problem of finding the ground state of a k-local Hamiltonian, as applies to generic Hubbard models, is QMA-complete for $k \geq 2$ [168], meaning quantum computers are not expected to solve all such problems in polynomial time. Another seminal result that should give one some pause is the fact that, in general, determining whether the ground state of a quantum system is gapped in the infinite limit is not just a hard problem, but an undecidable one [77, 22]. To see why this is bad, consider the fact that one could run a quantum simulation with 10^{23} particles, and find that the addition of a single particle qualitatively changes the behavior of the system. However, it is not generally expected that such arguments apply to most physical systems of interest, since the properties of real materials are almost always stable in the thermodynamic limit.

[5] In systems where the many body gap closes polynomially in system size, as is expected to be true for many Hubbard models, quantum systems can prepare the appropriate ground and excited states in polynomial time using protocols based on adiabatic evolution [27, 5] or the variational quantum eigensolver [257, 57, 330]. The belief that polynomial time classical algorithms are not possible in these situations stems, in part, from Bell's theorem: because the states of these systems can only be described by taking non-local correlations into account, in general, we don't expect to be able to simulate such systems classically without keeping track of an exponentially large state vector [265].

the scale of a single lattice site, and sensitivity on the scale of a single atom [14, 305]. By further using the microscope to project optical fields onto the atoms, it is possible to isolate a subset of atoms from a larger many-body state to prepare different states of interest [349, 151, 231, 264, 166, 359].

Along similar lines, in (2), tightly focused beams of light called optical tweezers can be used to trap and move individual atoms [19, 99]. This approach is distinct from (1) in the sense that, instead of isolating the desired part of a large (and often challenging to prepare) many body state, one traps and deterministically rearranges initially isolated atoms to rapidly prepare large ensembles of atoms. The motion of the atoms can further be controlled using sophisticated optical cooling techniques [164, 323]. Tweezer-trapped atoms can be made to tunnel between tweezers in similar ways to the experiments in (1) [165, 163], but it has proved challenging to scale such experiments beyond a handful of tweezers due to difficulties associated with disorder in the depths of the tweezers on the scale of the tunneling energy [316, 359].[6] However, a separate toolkit has been developed to engineer strong and coherent interactions between up to hundreds of tweezer-trapped atoms by exciting the atoms to high-lying Rydberg states. Since Rydberg states are strongly polarizable, the energy of a Rydberg state atom can be affected by the presence or absence of a nearby atom in a Rydberg state. The above techniques have been used to engineer entangling gates [190, 191, 127] and run shallow quantum circuits [128, 32], as well as to study a variety of spin models [47, 95, 294, 297].[7]

In (3), one is interested in preparing ensembles of atoms or ions, with each particle in the ensemble serving as an oscillator that keeps track of time [203]. Alkaline earth atoms are particularly suited to such experiments (in contrast to the alkali atoms typically used in the experiments described above), because the presence of the second valence electron allows for ground and excited electronic states that are long-lived, and insensitive to external perturbations (see Sec. 2.1.1.3). These states form the basis of optical frequency qubits (clock qubits) whose evolution can be used to keep track of time. In the setting of frequency metrology, quantum many body effects are often a nuisance. Instead, one wants to control or mitigate interactions between clock qubits to approximate pristine conditions where the atoms sit motionless in a vacuum, with their oscillations unadulterated by external perturbations from other atoms or fields. However, in the pursuit

[6] Note that the stroboscopic technique developed in [359] does hold some promise for scaling to larger tunnel-coupled tweezer arrays.

[7] Specifically, models where the spins are pinned in place, and therefore where particle statistics are irrelevant. This is distinct from the experiments in (1), where both spin and particle statistics can play a central role.

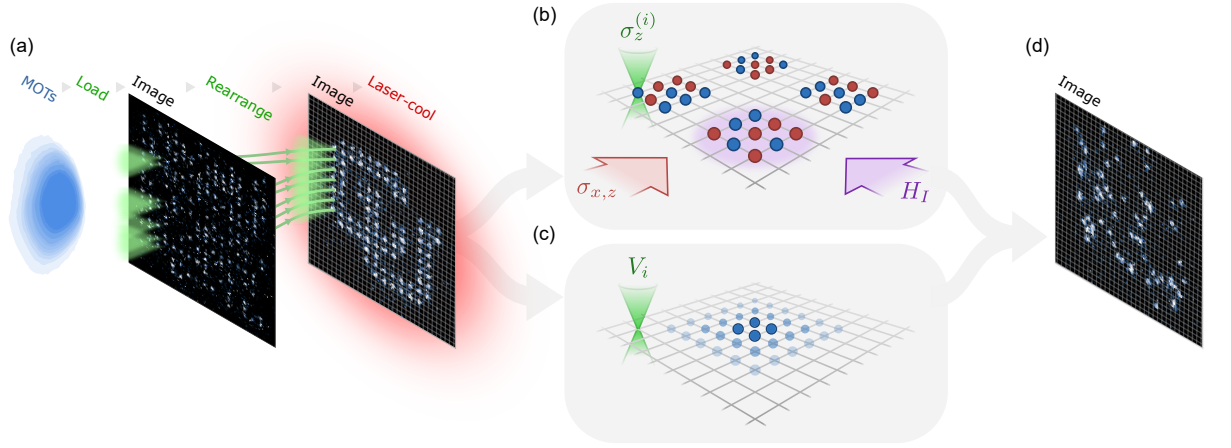

Figure 1.1: Schematic of a typical experiment. (a) A cold thermal cloud of atoms is prepared using several stages of magneto optical traps (MOTs). Atoms are stochastically loaded into a tweezer array (green cones) and transferred into an optical lattice (grey grid). By taking an image of the atoms, we can determine their locations, and deterministically rearrange them into the desired patterns using the optical tweezers (an additional image is typically taken to confirm that the rearrangement was performed successfully). Finally, the atoms are laser cooled to the motional ground state of the lattice sites that they occupy. The result is the on-demand preparation of large ensembles of atoms in states with very low entropy, which serves as the starting point for our experiments. These experiments fall into two broad categories. (b) The atoms (circles) can be pinned in place, and an internal optical frequency qubit degree of freedom (indicated by the blue and red circles) controlled. A universal set of single qubit gates can be performed using global rotations driven by a laser that addresses all the atoms (red arrow), and local phase shifts applied using the optical tweezers (green cone). Interactions and entangling gates can be engineered by driving the atoms to highly excited Rydberg states with a separate laser (purple arrow). (c) Alternatively, the atoms can be allowed to coherently tunnel through the lattice. The resulting quantum walks can be controlled by applying a local chemical potential (with a resolution on the scale of a single lattice site) using the tweezers. (d) In all cases, a final image, or sequence of images, is used to measure the resulting state of the atoms. The images appearing in this figure are single-shot experimental data, where each bright spot indicates the presence of an atom in its electronic ground state.

of these conditions, one often still has to understand and tame complicated quantum dynamics and many body effects. For example, the Fermi-lattice clock at JILA uses properties of the Fermi-Hubbard model to avoid overlaps between the atoms that can lead to uncontrolled shifts in their oscillation frequency [53]. Similarly, with the appropriate engineering, one can use different kinds of interactions in lattice clocks to precisely cancel each other out [35, 3]. Clocks that use ensembles of atoms or ions as their oscillators have achieved amazing feats, including setting bounds on possible candidates for dark matter [169], and performing geodesy [131] and tests of relativity [67] by measuring gravitational red shifts down to the scale of a millimeter [35].

The main objective of the work in this thesis is to apply the flexible and programmable control

provided by optical tweezer arrays (2), which have already been used to study spin models and quantum circuits, to other applications in quantum science, namely quantum simulations of different Hubbard models using optical lattices (1), and optical frequency metrology (3). We will find that by combining features from the three domains described above into a single apparatus, one gains significant advantages in terms of state preparation, readout, isolation from undesirable perturbations, and the kinds of dynamics, and control over those dynamics, that one has access to.

The primary role of the tweezer array in our experiments is to provide programmable control over both state preparation, and the ensuing dynamics. In terms of state preparation, the tweezer array provides fast and flexible preparation of initial states by means of single atom trapping, optical cooling, and, optionally, atom rearrangement (Fig. 1.1). This state preparation minimizes dead time on the experiment and improves the rate at which one can gather statistics. For simulating Hubbard physics, atom rearrangement is particularly useful since one is able to prepare a large class of excited (but still pure) initial states[8] rapidly and with high fidelity, making it easier to study certain kinds of dynamics. In terms of controlling the dynamics of the atoms, the tweezers can be used to implement local single qubit gates, or to modify the chemical potential and tunnel couplings in the lattice (Fig. 1.1). For frequency metrology, access to a (nearly) universal gate set and high-fidelity measurements (in a complete basis) allow one to explore sophisticated measurement schemes, including those involving entanglement, that can improve the performance of a clock. For simulating Hubbard physics, the tweezers allow one to change the lattice geometry, and to implement gate-like operations between sites in the lattice. In both cases, the above controls open the possibility of mixed analog and digital evolution, enabling new schemes for state preparation, and access to new observables. However, tweezer arrays are not a perfect tool, and suffer from limitations on scaling due to limited laser power and disorder in the generated optical potentials. The optical lattice nicely complements these shortcomings, allowing one to achieve similar or improved performance in imaging, cooling, and tunneling in comparison to tweezer arrays, while also allowing one to scale to systems with tens to hundreds of times as many atoms, and sites that the atoms can occupy (Fig. 1.1).

[8] Namely, Fock states corresponding to (nearly) arbitrary patterns with 0 or 1 atoms on each lattice site.

1.2 Summary of main results

With the above big-picture goals and capabilities in mind, the main results in this thesis are as follows:[9] We develop the basic tools to control the alkaline earth atom strontium (Sr) in optical tweezers, including trapping, cooling, and imaging [241] (in parallel with work at Caltech using Sr [73], and Princeton using Yb [286]). Using the above tools, we control an optical qubit degree of freedom in tweezer-trapped Sr atoms, demonstrating similar coherence properties to those achieved in optical lattice clocks [240] (in parallel with work at Caltech [209]). We further show that one can engineer tweezer potentials to preserve atomic coherence while scaling to larger atom arrays, achieving very high quality factors[10] of $Q \sim 10^{17}$ in a tweezer array containing ~ 150 atoms. This high quality factor and relatively high atom number enables synchronous frequency comparisons in the array with a fractional frequency instability of $5.2(3) \times 10^{-17}$ at 1 s of averaging [365].

We show that one can hand atoms between tweezers and an optical lattice with high fidelity, and, using the tightly confining lattice, demonstrate single qubit gates on the optical clock qubit with a π-pulse fidelity of $\gtrsim 99.5\%$, and entangling gates between clock qubits (mediated by Rydberg interactions) with a fidelity of 92.8(2.0)% (with SPAM correction).[11] Using the above gates, we prepare Bell states of optical clock qubits, and show that these states could be useful for optical frequency metrology. Similar Rydberg-mediated interactions are used to prepare squeezed states that allow for synchronous frequency comparisons with a precision beyond what is achievable at the standard quantum limit (SQL), with an enhancement of 3.52(1) dB in variance below the SQL for ensembles of 16 atoms, and 1.94(1) dB in variance for ensembles of 70 atoms [96].

Critically, the lattice is designed not just for providing tight confinement, but also to be compatible with Hubbard physics (namely where there can be a significant tunnel-coupling between adjacent lattice sites). We show that, even in a tightly spaced lattice that can achieve relatively high tunneling energies of $J/h > 100$ Hz,[12] it is possible to rapidly rearrange atoms with high fidelity using optical tweezers.

[9] The figures of merit provided in this section are to give a general feel of the performance of the experiment. Please refer to the relevant chapters and papers for more detailed accounts of errors in different experimental conditions.
[10] Associated with the oscillations of an even superposition of the two optical clock qubit states.
[11] Recent results from the group have improved the two qubit gate fidelities to $\gtrsim 98.5\%$.
[12] Using Sr atoms, which are fairly heavy and thus lead to reduced tunneling energies.

Quantum walks of the rearranged atoms are studied, and we further demonstrate that the tweezers can be used to programmably modify the lattice potential on the scale of the tunneling energy. This capability is used in the first demonstration of spatial search using quantum walks [366]. We benchmark the quality of our state preparation by studying multiparticle quantum walks, finding that the indistinguishability of two atoms is $99.5^{+0.5}_{-1.6}\%$ with postselection on atom survival, and that the on-demand success rate without postselection for perfectly preparing, evolving, and detecting an atom in the lattice is $\sim 92\%$ [367].[13] Given the above performance, we are able to study Fock state boson sampling problems involving up to 180 atoms occupying ~ 1000 sites in the lattice. These demonstrations are beyond the reach of known classical algorithms and current photonics experiments. In addition to the above published results, we will also present some preliminary results studying interacting Bose-Hubbard models in our apparatus.

1.3 Outline

Because the new platform we develop in this work bridges multiple fields, we will endeavor, where possible, to provide pedagogical introductions to the various techniques and applications that our experiments touch on. In Ch. 2, we introduce the basic properties of the alkaline earth atom Sr, and how it is affected by external control fields. In Ch. 3, we describe our experimental apparatus, including design considerations and tradeoffs relating to how the various control fields are applied (especially the various optical potentials, which are fairly unique to this apparatus). We will also cover basic features that are common to all other experiments presented in this work (e.g. initial trapping and cooling of a thermal gas of atoms). In Ch. 4, we describe results relating to preparing arrays of identical Sr atoms — namely preparing and imaging single atoms, optical cooling, and rearrangement using optical tweezers in a lattice. In Ch. 5, we will discuss frequency metrology in tweezer arrays, beginning with a general primer on quantum sensing. We will point out where entanglement can be beneficial in sensing, before discussing the features of tweezer arrays of alkaline earth atoms that are attractive for frequency metrology. The measurements presented in this chapter will focus on using non-entangled states of the atoms to perform comparisons between clock qubits

[13] This number is primarily limited by optimizing for high post-selected indistinguishability. When the experiment is properly tuned up, the imaging fidelity including contributions from atom loss is $\gtrsim 99.8\%$, and the rearrangement success probability is $\sim 99.5\%$ per atom.

and a local oscillator based on an ultrastable laser, as well as on synchronous comparisons between clock qubits. In Ch. 6, we will discuss how the atoms can be excited to Rydberg states, and how these states can be used to engineer an entangling gate between clock qubits. We will benchmark these gates, with an emphasis on how Bell pairs prepared using these gates can be used for frequency metrology. We will also briefly discuss results that use related techniques to generate spin squeezing in ensembles of tens of atoms.

In Ch. 7, we take a step back from systems of atomic qubits pinned in place, and consider the situation where the atoms are free to coherently move through the lattice. We will review the underlying physics of atoms in lattices, before framing our discussion in the language of quantum walks. We will discuss an algorithm for performing spatial search by quantum walk, and show that we can realize this algorithm by programming quantum walks of individual atoms. This demonstration will involve both the adiabatic (and dynamical) preparation of resource states for the algorithm, as well as the execution of the algorithm itself. We will further comment on the degree to which the tools used in our demonstration of spatial search allow for universal control over the quantum walk dynamics. In Ch. 8, we will extend our discussion to multiparticle quantum walks, where different atoms are free to exchange positions. We will find that, unlike in earlier chapters, it is no longer sensible to talk about one qubit being entangled with another because the labels we assign to the different qubits are no longer physically meaningful, resulting in dramatically different physical behavior. We will comment on the mechanisms leading to interactions between the atoms, and present preliminary results studying strongly interacting quantum walks, and a superfluid that is prepared and probed in a programmable fashion using optical tweezers. Much of the discussion in Ch. 8 will be about non-interacting and bosonic quantum walks, which can be framed in terms of the boson sampling problem. We will present techniques for benchmarking both state preparation and the quantum walk dynamics using multiparticle quantum walks, as well as quantum walks with up to 180 atoms occupying ~ 1000 distinct sites. Although the largest-scale experiments we perform are not expected to be efficiently simulable with any known classical algorithm, we will discuss caveats to this claim, including the errors that our system is sensitive to, and how these compare to other experimental demonstrations of boson sampling. Finally, in Ch. 9, we conclude with a general discussion of the strengths and weaknesses of our platform. We will discuss future directions in metrology and quantum simulation using the tools developed in this thesis, as

well as combined directions that blur the lines between quantum simulation and sensing.

1.4 Published content

While pursuing their Ph.D., the author contributed to the following works:

[241] M. A. Norcia, A. W. Young, and A. M. Kaufman. Microscopic control and detection of ultracold strontium in optical-tweezer arrays. **Phys. Rev. X**, 8(4):041054, Dec. 2018

[240] M. A. Norcia, A. W. Young, W. J. Eckner, E. Oelker, J. Ye, and A. M. Kaufman. Seconds-scale coherence on an optical clock transition in a tweezer array. **Science**, 366(6461):93–97, 2019

[365] A. W. Young, W. J. Eckner, W. R. Milner, D. Kedar, M. A. Norcia, E. Oelker, N. Schine, J. Ye, and A. M. Kaufman. Half-minute-scale atomic coherence and high relative stability in a tweezer clock. **Nature**, 588(7838):408–413, Dec. 2020

[364] A. Young, W. G. Graham, T. J. Morgan, and L. Hüwel. Laser-induced breakdown in liquid water: Influence of repeated laser pulses on plasma formation and emission. **Journal of Applied Physics**, 129(18):183303, May 2021

[97] W. J. Eckner, A. W. Young, N. Schine, and A. M. Kaufman. High-power, fiber–laser-based source for magic-wavelength trapping in neutral-atom optical clocks. **Rev. Sci. Instrum.**, 92(9):093001, Sept. 2021

[290] N. Schine, A. W. Young, W. J. Eckner, M. J. Martin, and A. M. Kaufman. Long-lived Bell states in an array of optical clock qubits. **Nat. Phys.**, pages 1–7, Aug. 2022

[366] A. W. Young, W. J. Eckner, N. Schine, A. M. Childs, and A. M. Kaufman. Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice. **Science**, 377(6608):885–889, Aug. 2022

[96] W. J. Eckner, N. Darkwah Oppong, A. Cao, A. W. Young, W. R. Milner, J. M. Robinson, J. Ye, and A. M. Kaufman. Realizing spin squeezing with Rydberg interactions in an optical clock. **Nature**, pages 1–6, Aug. 2023

[367] A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman. An atomic boson sampler. **In review**, 2023

1.5 Units

Unless otherwise noted, all error bars and confidence intervals provided in this work correspond to ± 1 standard deviations.

Chapter 2

Prerequisites

"The answer to what?"

— Deep Thought

In this chapter, we review the basic properties of the alkaline earth atom Sr, and discuss the tools we have for controlling ensembles of these atoms - namely static electric and magnetic fields, and oscillating electromagnetic fields of variable frequency, amplitude, polarization, and geometry.

2.1 Strontium and the alkaline earth-like atoms

Alkaline earth-like atoms (AEAs) have two valence electrons, and otherwise full inner electron shells.[1] The key feature of AEAs is that the spins of the two valence electrons can either cancel each other out, or pair up to produce ground and excited states with net angular momentum $J = 0$, resulting in both a variety of useful optical properties, and a nearly-complete decoupling of the nuclear spin sector from the electronic sector [126]. The latter property is useful for preserving information stored in the state of the nuclear spin while performing other manipulations of the electronic state [155, 206, 18, 194, 146, 207, 237], as well as for simulations of exotic $SU(N)$-symmetric systems [126, 288, 56]. In this work, we primarily take advantage of the former property, using convenient optical transitions to manipulate the electronic and motional states of neutral Sr atoms.

[1] Atoms falling into this classification include the alkaline earth atoms Be, Mg, Ca, Sr, Ba, and Ra, the zinc-family of atoms Zn, Cd, and Hg, as well as the lanthanide Yb and actinide No [319].

2.1.1 Level structure

We will concern ourselves only with the behavior of the two valence electrons of Sr, and define $\boldsymbol{s}_i$ and $\boldsymbol{l}_i$ to be the spin and orbital angular momentum vectors of the i^{th} electron.[2] s_i is the single electron spin quantum number, where $\hbar^2 s_i(s_i+1)$ are the eigenvalues of the $\boldsymbol{s}_i^2$ operator, and similarly for l_i and $\boldsymbol{l}_i^2$. n_i is the principle quantum number of the i^{th} electron. $\boldsymbol{S} = \sum_i \boldsymbol{s}_i$, $\boldsymbol{L} = \sum_i \boldsymbol{l}_i$, and $\boldsymbol{J} = \boldsymbol{S} + \boldsymbol{L}$ are respectively the total spin, total orbital angular momentum, and total angular momentum vectors of the two electrons. Similar to the single electron case, the $\boldsymbol{S}^2$, $\boldsymbol{L}^2$, and $\boldsymbol{J}^2$ operators have associated quantum numbers S, L, and J. We will assume that the atomic Hamiltonian H commutes (at least approximately) with $\boldsymbol{S}^2$, $\boldsymbol{L}^2$, and $\boldsymbol{J}^2$, such that S, L, and J remain good quantum numbers. This is true, for example, when the only spin-orbit coupling terms in H take the form $\boldsymbol{L}\cdot\boldsymbol{S}$ (this situation is known as LS coupling).[3] Note that we will often write the orbital angular momentum quantum numbers $L = 0, 1, 2, 3\ldots$ as $L = \mathrm{S}, \mathrm{P}, \mathrm{D}, \mathrm{F}\ldots$ [318].[4] With the above definitions in mind, we adopt Russell-Saunders notation ($n_1 l_1 n_2 l_2\ {}^{2S+1}L_J$) to label the different electronic states of Sr. These states and their accompanying transitions are summarized in Fig. 2.1.

The electronic states of Sr naturally separate into $S = 0$ (spin singlet), and $S = 1$ (spin triplet) manifolds. For LS coupling, a transition between the singlet and triplet manifolds is dipole forbidden due to the selection rule $\Delta S = 0$.[5] However, as we will see in Sec. 2.1.1.2, the large nuclear charge of Sr results in spin-orbit coupling that can change this situation, making certain singlet-triplet transitions dipole allowed.

The ground state of Sr is the $5s^2\ {}^1\mathrm{S}_0$ spin singlet state.[6] For bosonic isotopes of Sr, $J = 0$ results in a unique ground state that is first order insensitive to external magnetic fields. Note that for all electronic excited states relevant to this work, only a single electron is excited from the ground state, with the other being in the $5s$ state. This is because most states of Sr with two excited electrons have energies that exceed

[2] We will not concern ourselves with hyperfine structure and the nuclear spin $\boldsymbol{I}$, since all results in this work involve isotopes of Sr with $I = 0$.

[3] The mechanism for this coupling can be understood as a relativistic effect: in the frame of the electron, the moving nucleus of the atom generates a magnetic field which interacts with the electron spin. For light elements, terms like $\boldsymbol{s}_i\cdot\boldsymbol{s}_j$ and $\boldsymbol{l}_i\cdot\boldsymbol{l}_j$ dominate over terms like $\boldsymbol{s}_i\cdot\boldsymbol{l}_j$, and so it is reasonable to make the approximation that the spin orbit terms take the form $\boldsymbol{L}\cdot\boldsymbol{S}$. However, we will point out corrections to this approximation as appropriate.

[4] We will also follow the convention that the single electron configurations are provided in lower case, although we will sometimes omit these labels for convenience.

[5] Such a transition is also dipole forbidden from a different perspective: due to the requirement that the overall electronic wavefunction is antisymmetrized, the spin singlet states all have symmetric spatial wavefunctions, and the spin triplet states antisymmetric spatial wavefunctions. For two electrons, the operator for the electric dipole moment is $\boldsymbol{d} \propto \boldsymbol{x}_1 + \boldsymbol{x}_2$, where $\boldsymbol{x}_i$ is the coordinate of the i^{th} electron. Notice that $\boldsymbol{d}$ is unchanged upon exchange of the electron positions, and so the resulting dipole matrix element between states with different exchange symmetries must be zero [26].

[6] Including the configuration of the inner electrons: $1s^2 2s^2 2p^6 3s^2 3p^6 3d^{10} 4s^2 4p^6 5s^2\ {}^1\mathrm{S}_0$.

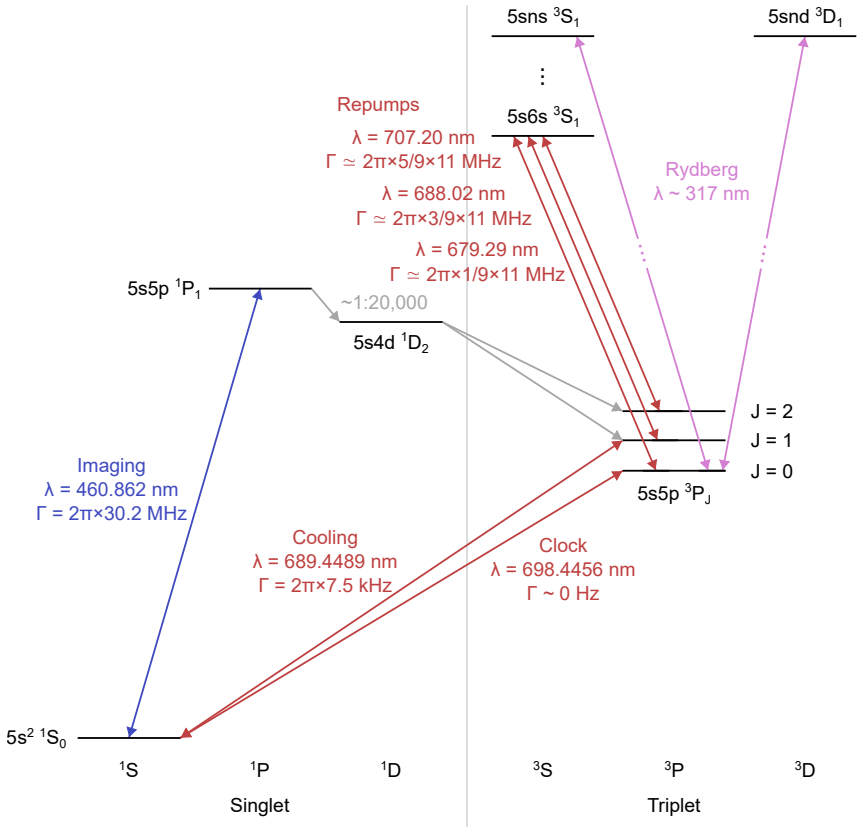

Figure 2.1: Level diagram for ^{88}Sr. We display the relevant low-lying energy levels, and the transitions between these levels that are of interest to our experiments. In particular, note the weakly allowed transitions between the singlet and triplet manifolds (the clock and cooling transitions), and the strongly allowed transitions within a manifold (the imaging and repumping transitions). The indicated Rydberg transitions correspond to a series of Rydberg states with different properties, as described in the text.

the ionization threshold of 5.69485 eV (or 217.7 nm, expressed as a wavelength in vacuum), often resulting in autoionization [279].

2.1.1.1 Dipole-allowed $^1S_0 \leftrightarrow {}^1P_1$ transition

The strongest transition from the ground state is the dipole allowed $^1S_0 \leftrightarrow 5s5p\ {}^1P_1$ "blue" transition, which has an associated vacuum transition wavelength of 460.862 nm. The $^1P_1$ state has a lifetime of 5.263(4) ns [363], and a corresponding linewidth of $\Gamma_b \simeq 2\pi \times 30.2$ MHz. This state decays primarily to $^1S_0$, with a weak branching ratio of $\sim 1:20000$ into the $5s4d\ {}^1D_2$ state [73], which subsequently decays to the states $5s5p\ {}^3P_{1,2}$. As we will see in Secs. 2.1.1.2 and 2.1.1.4, $^3P_1$ decays back to the ground state, and

atoms in $^3\mathrm{P}_2$ can be optically pumped to the ground state. As a result, $^1\mathrm{S}_0 \leftrightarrow {}^1\mathrm{P}_1$ provides a strong cycling transition that can be used to rapidly scatter many photons off of the atoms for initial cooling and trapping of a thermal gas (Sec. 3.1.2), as well as for imaging (Sec. 4.4).

2.1.1.2 Dipole-forbidden $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ transition

Although singlet-triplet transitions are forbidden under LS coupling, the highly charged nucleus of Sr results in spin-orbit coupling terms of the form $\boldsymbol{l}_i \cdot \boldsymbol{s}_i$ which violate the LS coupling assumption.[7] Such terms can mix together states that differ by a spin flip. This applies to the $5s5p\ {}^3\mathrm{P}_1$ state, which gains a small admixture of the $5s5p\ {}^1\mathrm{P}_1$ state [39]. As a result, the $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ "red" transition is (weakly) dipole allowed, with an associated transition wavelength of 689.4489 nm in ^{88}Sr [283]. The $^3\mathrm{P}_1$ state decays to $^1\mathrm{S}_0$ with a lifetime of 21 μs, and a corresponding linewidth of $\Gamma_r = 2\pi \times 7.5(2)$ kHz [283]. As a result, $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ provides a narrow-line cycling transition that can be used for various frequency selective operations (with a resolution on the scale of Γ_r), including high fidelity optical cooling (in the sense that the state of the atom has high overlap with the motional ground state, see Sec. 4.3).

Given that the $^3\mathrm{P}_1$ state has $J = 1$, there are three degenerate states in the $^3\mathrm{P}_1$ manifold with $m_j = -1, 0, +1$, where $\hbar m_j$ are the eigenvalues of the J_z operator. However, these states are split in the presence of a magnetic field $\boldsymbol{B}$, the influence of which is described by the Hamiltonian [318]:

$$H_B \coloneqq \mu_\mathrm{B} g_j \boldsymbol{J} \cdot \boldsymbol{B}, \tag{2.1}$$

where μ_B is the Bohr magneton, and g_j is the Landé g_j factor. For $\boldsymbol{B}$ applied in the z direction, the $^3\mathrm{P}_1$ states experience a linear Zeeman shift of:[8]

$$\Delta E_B = \mu_\mathrm{B} g_j m_j B. \tag{2.2}$$

2.1.1.3 Doubly forbidden $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_{0,2}$ transitions

The $^3\mathrm{P}_0$ state, also known as the "clock" state, can be extremely long lived. This is because the only lower-lying state is the $^1\mathrm{S}_0$ ground state, to which transitions are doubly forbidden in the sense that

[7] The situation where these terms dominate over LS coupling is sometimes referred to as jj coupling.

[8] Note that the treatment in this section is only applicable when ΔE_B is smaller than the fine structure. ΔE_B larger than the fine structure is referred to as the Paschen-Back regime, where J is no longer a good quantum number [318]. This situation will not apply in any of our experiments.

they require both $\Delta S \neq 0$, and that $J = 0$ both before and after the transition. However, as first proposed in [322], one can open up the ${}^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0$ transition, which we will refer to as the clock transition, with the application of a magnetic field.

Specifically, a DC magnetic field vector $\boldsymbol{B}$ can couple the $|{}^3\mathrm{P}_0\rangle$ and $|{}^3\mathrm{P}_1\rangle$ states with:

$$\Omega_B = \langle {}^3\mathrm{P}_0 \,|\, \boldsymbol{\mu} \cdot \boldsymbol{B} \,|\, {}^3\mathrm{P}_1 \rangle / \hbar \tag{2.3}$$

where $\boldsymbol{\mu}$ is the magnetic dipole operator. Assuming approximate values of the Landé g factors $g_L \simeq 1$ and $g_S \simeq 2$, $|\Omega_B| = \sqrt{2/3}\mu_B |\boldsymbol{B}|/\hbar$ for bosonic isotopes of Sr [322]. For $|\Omega_B/\Delta_{01}| \ll 1$, where $\Delta_{01} = 2\pi \times 5.6$ THz [322] is the detuning between the ${}^3\mathrm{P}_0$ and ${}^3\mathrm{P}_1$ states, $|{}^3\mathrm{P}_0\rangle$ acquires a small admixture of $|{}^3\mathrm{P}_1\rangle$:

$$|{}^3\mathrm{P}_0'\rangle = |{}^3\mathrm{P}_0\rangle + \frac{\Omega_B}{\Delta_{01}} |{}^3\mathrm{P}_1\rangle, \tag{2.4}$$

where the prime denotes the new eigenstate in the presence of the perturbing field. As a result, one can directly drive the ${}^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0'$ transition, just with a modified Rabi frequency of $\Omega_c = \Omega_r \Omega_B / \Delta_{01}$, where Ω_r is the Rabi frequency associated with the ${}^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ transition (we will discuss Rabi driving of electronic transitions in more detail in Sec. 2.2.1). By second order perturbation theory, the above coupling also results in an energy shift on the ${}^3\mathrm{P}_0'$ state of:

$$\Delta_B = -\Omega_B^2 / \Delta_{01}, \tag{2.5}$$

which is known as a quadratic, or second order, Zeeman shift. Similarly, the linewidth of the ${}^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0'$ transition is $\Gamma_c = \Gamma_r (\Omega_B/\Delta_{01})^2$. The vacuum transition wavelength for ${}^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0$ is 698.4456 nm [283], but one must be careful when computing the transition wavelength associated with ${}^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0'$ due to both the quadratic Zeeman shift, and any shifts due to the presence of additional control fields (see Ch. 5). For reasonable values of the magnetic field, the lifetime associated with ${}^3\mathrm{P}_0'$ can be in excess of thousands of seconds while still maintaining the possibility of reasonably fast manipulations of the ${}^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0'$ transition (with Rabi frequencies in the kilohertz range for 550 G magnetic fields and ~ 100 mW of optical power). This makes the ${}^1\mathrm{S}_0$ and ${}^3\mathrm{P}_0'$ states attractive as a long-lived and magnetically insensitive qubit degree of freedom, which we will discuss in more detail in Chs. 5 and 6.

Note that the $^1S_0 \leftrightarrow {}^3P_2$ transition is also doubly forbidden, since it requires that $\Delta S \neq 0$ and $\Delta J = 2$, resulting in long lifetimes of ~ 500 s in ^{88}Sr [362]. However, $^1S_0 \leftrightarrow {}^3P_2$ is a magnetic quadrupole-allowed transition, and thus can be driven with the appropriate laser fields [333].

2.1.1.4 Optical pumping transitions

We take advantage of the dipole allowed $^3P_{0,1,2} \leftrightarrow 5s6s\ {}^3S_1$ transitions for optical pumping. The $^3S_1$ state has a lifetime of 13 ns [283], and decays to the ^3P_J manifold with a total rate of $\Gamma \simeq 2\pi \times 11$ MHz, and branching ratios of 1/9, 3/9, and 5/9 for $J = 0, 1, 2$ respectively [320]. The associated vacuum transition wavelengths are 679.29 nm for $^3P_0 \leftrightarrow {}^3S_1$, 688.02 nm for $^3P_1 \leftrightarrow {}^3S_1$, and 707.20 nm for $^3P_2 \leftrightarrow {}^3S_1$ [283].

One can (incoherently) transfer populations from one electronic state to another by turning on the appropriate optical pumping beams. Leaving any given beam on will deplete the population from the corresponding state, with some population decaying to the ground state via $^3P_1$, and some being trapped in the other remaining triplet state. One can return all population in $^3P_{0,2}$ to the ground state by turning on the $^3P_0 \leftrightarrow {}^3S_1$ and $^3P_2 \leftrightarrow {}^3S_1$ optical pumping beams. The reverse process is also possible, where driving the $^1S_0 \leftrightarrow {}^3P_1$ transition while applying the $^3P_1 \leftrightarrow {}^3S_1$ optical pumping beam transfers ground state population to the $^3P_{0,2}$ states. Additionally applying the $^3P_2 \leftrightarrow {}^3S_1$ or $^3P_0 \leftrightarrow {}^3S_1$ pumping beam allows one to pump all the population into either $^3P_0$ or $^3P_2$. All of the above operations can be performed on a timescale of $\lesssim 100\ \mu$s.

2.1.1.5 Rydberg states

For highly excited electronic states with $n_2 \gg n_1$, where $i = 1, 2$ are taken to refer to the second-most and most energetic electrons respectively, the details of the orbits of the inner electrons become comparably unimportant. Instead, the outer electron primarily sees a point charge of $+e$ at the origin, where e is the elementary charge, with all other charges in the nucleus being screened by the inner electrons. The resulting states are hydrogenic, and so we will label these states by the values of n, l, j, and m_j associated with the outer electron, where we have dropped the $i = 2$ index for convenience.[9] The corresponding energy levels

[9] Here, j refers to the quantum number associated with the operator $\boldsymbol{j}^2$, and m_j with j_z, where $\boldsymbol{j} = \boldsymbol{s}_2 + \boldsymbol{l}_2$.

scale approximately like n^{-2}, and small deviations from this scaling can be captured by a quantum defect δ_{nlj}.[10] Given the above definitions, the electronic energy levels are given by the empirical formula [308]:

$$E_{nlj} = E_I - \frac{\mathrm{R}'}{(n-\delta_{nlj})^2}, \tag{2.6}$$

where E_I is the ionization energy threshold, and R' the Rydberg energy corrected by the reduced mass:

$$\mathrm{R}' = \left(\frac{m_e}{m+m_e}\right)\left(\frac{m_e e^4}{8\epsilon_0^2 h^2}\right). \tag{2.7}$$

Here, m is the mass of the atom, m_e is the electron mass, and ϵ_0 is the permittivity of free space. Notice that in the limit of large n the difference in energy between adjacent values of n scales like n^{-3}.

We will discuss these states, and interactions between different atoms in Rydberg states, in more detail in Ch. 6. For now, we point out a few general scaling laws. The average size of the electron orbital scales like $\langle|\boldsymbol{x}|\rangle \propto n^2$, where $\boldsymbol{x}$ is the displacement vector of the electron [308]. As a result, the dipole matrix element $\langle n\,|\,\boldsymbol{d}\,|\,n+1\rangle$ between neighboring Rydberg levels[11] , where $\boldsymbol{d}$ is the electric dipole operator, also scales like $\langle n\,|\,\boldsymbol{d}\,|\,n+1\rangle \propto n^2$. The increase in orbital size also reduces the overlap with the ground state, leading to a dipole matrix element from the ground state $|g\rangle$ of $\langle g\,|\,\boldsymbol{d}\,|\,n\rangle \propto n^{-3/2}$. Considering the overlap with all lower energy states, one finds that the lifetimes of the Rydberg states increase with n. Specifically, states with small orbital angular momentum $l \lesssim 2$ have lifetimes that scale like $\tau \propto n^3$ when considering only spontaneous emission. States with larger l typically have longer lifetimes.[12] At finite temperature, the presence of black body radiation (BBR) can resonantly drive transitions between nearby Rydberg states, leading to a scaling of between $\tau \propto n^3$ and $\tau \propto n^2$ for the lifetimes of the $l \lesssim 2$ states. Typically, one will care about the ratio between the "good" coherent evolution with a timescale scale set by $1/|\langle g\,|\,\boldsymbol{d}\,|\,n\rangle|$ and the "bad" incoherent evolution with a timescale set by τ. A reasonable figure of merit is $|\langle g\,|\,\boldsymbol{d}\,|\,n\rangle|\tau \sim n^{1/2}$ to $n^{3/2}$, which scales favorably with increased n. However, we will see in Sec. 6.1.2 that this naive scaling can sometimes be misleading. In this work, we will be particularly interested in driving transitions between the $^3\mathrm{P}_0$ state and Rydberg states with $n \sim 40$, for which the vacuum transition wavelength is ~ 317 nm.

[10] Note that $\delta_{nlj} \simeq 0$ for $l > 3$ [308], since in these cases the outer electron is concentrated even further away from the nucleus.

[11] We are omitting the other quantum numbers for convenience, but standard selection rules still apply.

[12] For example, the so-called "circular" states with maximal values of $l = n-1$ have lifetimes of up to $\tau \propto n^5$.

	m	Natural abundance	I	Particle statistics	$^1S_0 \leftrightarrow {}^1P_1$	$^1S_0 \leftrightarrow {}^3P_1$	$^1S_0 \leftrightarrow {}^3P_0$
^{88}Sr	87.905612	0.8258(1)	0	Bosonic	-	-	-
^{87}Sr	86.908877	0.0700(1)	9/2	Fermionic	-46.5	-62.187	-62.171
^{86}Sr	85.909260	0.0986(1)	0	Bosonic	-124.8	-163.819	-162.939
^{84}Sr	83.913425	0.0056(1)	0	Bosonic	-270.8	-351.496	-349.656

Table 2.1: Isotopes of strontium. m refers to the atomic mass, and I to the nuclear magnetic moment. The isotope shifts for the indicated transitions are provided as $f(^{8x}\text{Sr}) - f(^{88}\text{Sr})$, where $f(\cdot)$ denotes the frequency of the appropriate transition for the indicated isotope in units of megahertz. The indicated values for ^{87}Sr are for the weighted mean of transitions to different hyperfine states, and are an estimate of the isotope shift in the absence of hyperfine interactions [225]. The values appearing in this table are from [283, 319, 225].

The combined effects of reduced energy splitting between adjacent levels and $\langle n \,|\, \boldsymbol{d} \,|\, n+1 \rangle \propto n^2$ can result in enhanced sensitivity to DC electric fields at large n. Specifically, consider the scalar polarizability of the state $|n, l, j, m_j\rangle$ [318]:

$$\alpha^{(0)}(n,l,j) = 2e^2 \sum_{n'l'j' \neq nlj} \frac{|\langle n,l,j,m_j \,|\, \hat{z} \,|\, n',l',j',m_j \rangle|^2}{E_{n'l'j'} - E_{nlj}}, \tag{2.8}$$

where $\hat{z}$ is the position operator for the outer electron along the axis of the applied field. Note that a static electric field cannot impart angular momentum, and so we have used the selection rule $\Delta m_j = 0$ in the summation. The corresponding DC Stark shift to the energy is $\Delta E_{nlj} = -\alpha^{(0)}(n,l,j)\mathcal{E}^2/2$, where $\mathcal{E}$ is the strength of the applied electric field. Per the arguments above, neighboring energy levels with $\Delta n = 1$ will contribute n^4 to the numerator, and n^{-3} to the denominator, leading to an overall scaling of $\Delta E_{nlj} \propto n^7$. Although this large polarizability can be useful for generating strong interactions between atoms (see Sec. 6.1.2), it also becomes very important to control stray electric fields in order to avoid large uncontrolled DC Stark shifts of the Rydberg energy levels (see Sec. 3.1).

2.1.2 Isotopes of strontium

There are four naturally abundant isotopes of Sr, whose properties are summarized in Tab. 2.1. Critically, the isotope shifts for the relevant transitions described in the previous section are all below a few hundred MHz, meaning that all four isotopes can be controlled with the same laser systems, up to small frequency shifts that are within the bandwidth of acousto-optic modulators. Since there are an even number of electrons and protons in Sr, all isotopes with even mass number (namely ^{84}Sr, ^{86}Sr, and ^{88}Sr) also

have an even number of neutrons, resulting in an even overall number of fermions in each atom, and thus bosonic composite particle statistics. Because the nucleons arrange in spin singlets in the nuclear ground state, the total nuclear spin of these isotopes is $I = 0$ [319]. The isotope with odd mass number, ^{87}Sr, has fermionic particle statistics, and $I = 9/2$. In this work, we will focus on the bosonic isotopes, avoiding details associated with hyperfine structure due to the large nuclear spin in ^{87}Sr. Note that the electronic properties of the bosonic isotopes are almost completely identical, making it easy to switch between different isotopes. The key differences between these isotopes come in the form of their scattering properties, which we discuss in Sec. 8.1.1.

2.2 Electromagnetic control

The basic level structure of Sr described in the previous section forms a Hilbert space that we are interested in manipulating. We briefly described the role of DC electric and magnetic fields in Sec. 2.1.1, but now review a more complete formalism for describing the effect of both DC and AC electric fields on the electronic states of Sr.

Consider an electric field oscillating with angular frequency ω with an initial phase ϕ, and whose amplitude $\mathcal{E}$ varies as a function of position $\boldsymbol{r}$:

$$\boldsymbol{\mathcal{E}}(\boldsymbol{r},t) = \hat{\varepsilon}\mathcal{E}(\boldsymbol{r})\cos(\omega t+\phi) = \boldsymbol{\mathcal{E}}^{+}(\boldsymbol{r},t) + \boldsymbol{\mathcal{E}}^{-}(\boldsymbol{r},t) \tag{2.9}$$

$$\boldsymbol{\mathcal{E}}^{\pm}(\boldsymbol{r},t) \coloneqq \hat{\varepsilon}\frac{\mathcal{E}(\boldsymbol{r})}{2}e^{\mp i(\omega t+\phi)}, \tag{2.10}$$

where $\hat{\varepsilon}$ is a unit vector indicating the polarization of the electric field,[13] and we have defined positive- and negative-rotating components of the field $\boldsymbol{\mathcal{E}}^{\pm}$. For much of the following discussion, we will suppress the position dependence with the understanding that $\mathcal{E}$ varies slowly over space in comparison to the spatial extent of the electronic wavefunctions.

Under the electric dipole approximation, the interaction between a Sr atom and $\boldsymbol{\mathcal{E}}$ is described by the Hamiltonian:

$$H_i = -\boldsymbol{d}\cdot\boldsymbol{\mathcal{E}}, \tag{2.11}$$

[13] In general, $\hat{\varepsilon}$ can be complex to handle circularly polarized components of the field.

where $\boldsymbol{d} = -e(\boldsymbol{x}_1 + \boldsymbol{x}_2)$ is the electric dipole operator for two electrons.

2.2.1 Two level approximation

We first consider the simplified setting where there are only two relevant electronic states $|g\rangle$ and $|e\rangle$. This applies, for example, when ω is close to a resonance between two states, and so the coupling between these two states dominates over the off-resonant coupling to all other states. In this case, the Hamiltonian H_a for the bare atom, and the full Hamiltonian H, in the presence of the drive field are:

$$\begin{aligned} H_a &= \hbar\omega_{ge}|e\rangle\langle e| \\ H &= H_a + H_i, \end{aligned} \tag{2.12}$$

where $\hbar\omega_{ge} = \hbar(\omega_e - \omega_g)$ is the energy difference between the $|e\rangle$ and $|g\rangle$ states, and we have applied an offset by the identity such that the energy of the $|g\rangle$ state in the absence of $\boldsymbol{\mathcal{E}}$ is 0.

Note that in the Schrödinger picture, evolution under H_a will cause the state $|e\rangle$ to rotate with a phase factor of $e^{-i\omega_{ge}t}$. Therefore, for evolution under H, terms in H_i like $\langle e \,|\, \boldsymbol{d} \cdot \boldsymbol{\mathcal{E}}^- \,|\, g\rangle$ will rotate very rapidly in time (with an angular frequency of $\omega_{ge} + \omega$), and thus approximately average to zero.[14] Taking the "rotating wave approximation" where we remove the above terms from H_i,[15] and additionally transforming into the frame rotating at ω (indicated by $|e\rangle \to |e'\rangle$), we have [318]:

$$\begin{aligned} H_a &= -\hbar\Delta|e'\rangle\langle e'| \\ H_i &= \hbar\frac{\Omega}{2}\sigma_x, \end{aligned} \tag{2.13}$$

where σ_x is the Pauli x matrix, and we have defined a detuning $\Delta = \omega - \omega_{ge}$, and Rabi frequency:

$$\Omega = -\frac{\mathcal{E}}{\hbar}\langle g \,|\, \hat{\varepsilon} \cdot \boldsymbol{d} \,|\, e\rangle. \tag{2.14}$$

When $\Delta = 0$, H is proportional to σ_x. When $\Omega = 0$, H is proportional to σ_z up to an offset by the identity. As a result, application of the appropriate laser field for the appropriate duration provides a universal set of unitary operations on the $\{|g\rangle, |e'\rangle\}$ subspace.

[14] Also note that $\langle g \,|\, \hat{\varepsilon} \cdot \boldsymbol{d} \,|\, g\rangle = 0$ due to the parity of $\boldsymbol{d}$ (and similarly for e).

[15] This approximation is valid for $|\omega_{ge} - \omega| \ll \omega_{ge} + \omega$, but notably does not apply to the case of DC fields when $\omega = 0$. We will see how to handle this in the next section.

One can diagonalize H to find its eigenstates and eigenenergies:

$$|+\rangle = \sin\theta|g\rangle + \cos\theta|e'\rangle \tag{2.15}$$

$$|-\rangle = \cos\theta|g\rangle - \sin\theta|e'\rangle \tag{2.16}$$

$$E_\pm = \hbar\frac{-\Delta \pm \Omega'}{2}, \tag{2.17}$$

where we have defined the generalized Rabi frequency $\Omega' \coloneqq \sqrt{\Omega^2 + \Delta^2}$, and the angle $\theta \in [0, \pi/2)$ such that $\tan 2\theta = -\Omega/\Delta$. For $\Delta < 0$ (red detuning) and $-\Delta \gg |\Omega|$, $|-\rangle \simeq |g\rangle - \Omega/(2\Delta)|e'\rangle$ with an associated eigenenergy of $E_- \simeq \hbar\Omega^2/(4\Delta)$. Therefore, the state $|-\rangle$ is approximately equal to $|g\rangle$, and shifted in energy by:

$$V(\boldsymbol{r}) = \hbar\frac{\Omega(\boldsymbol{r})^2}{4\Delta}, \tag{2.18}$$

in comparison to the case where $\Omega = 0$. Note that we have reintroduced a possible dependence on the spatial coordinate $\boldsymbol{r}$. V is known as an AC Stark shift, and is proportional to the intensity $I(\boldsymbol{r}) = c\epsilon_0\mathcal{E}(\boldsymbol{r})^2/2$ of the applied oscillating field.[16] If the motion of an atom is slow enough that the value of V experienced by the atom varies slowly (in time) with respect to Δ, then the atom will simply adiabatically follow the state $|-\rangle$. In this case, one can think of V as a conservative, spatially varying potential which allows one to apply a mechanical force, known as an optical dipole force, to an atom. By appropriately shaping the optical field one can make an optical dipole trap (or optical trap for short), namely a local minimum in V that atoms are unable to escape.[17]

Although the above discussion provides a qualitative understanding of optical dipole forces, when Δ is large it is often no longer reasonable to make the two level approximation, and a more careful accounting of all nearby energy levels, including degeneracies due to angular momentum, is necessary.

2.2.2 Dipole forces

To accurately determine the Stark shift on a given state $|\beta\rangle$, we extend the discussion in the previous section by relaxing the rotating wave approximation (namely by including the fast counterrotating terms),

[16] Note that one can change the sign of V by changing the sign of the detuning.

[17] Note that because in this case V comes from an oscillating electric field, is it possible to make an extremum of the field in free space. This is to be contrasted with the case of DC fields, where no such extremum can exist since $\nabla \cdot \boldsymbol{\mathcal{E}} = \nabla \times \boldsymbol{\mathcal{E}} = 0$.

performing a summation over all other states $|\beta'\rangle$, and more carefully handling the orientation between the atom and the electric field. This leads to the Kramers-Heisenberg polarizability tensor [183, 318]:

$$\alpha_{\mu\nu}(\beta;\,\omega) = \sum_{\beta'} \frac{2\omega_{\beta'\beta}\langle\beta\,|\,d_\mu\,|\,\beta'\rangle\langle\beta'\,|\,d_\nu\,|\,\beta\rangle}{\hbar(\omega_{\beta'\beta}^2 - \omega^2)}, \tag{2.19}$$

where $\hbar\omega_{\beta'\beta} = \hbar(\omega_{\beta'} - \omega_\beta)$ is the energy difference between states $|\beta'\rangle$ and $|\beta\rangle$, and d_μ refers to one Cartesian component of $\boldsymbol{d}$.[18] The associated shift in the energy of the state $|\beta\rangle$ is:

$$\Delta E(\beta;\,\omega) = -\,\mathrm{Re}[\alpha_{\mu\nu}(\beta;\,\omega)]\mathcal{E}_\mu^+\mathcal{E}_\nu^-. \tag{2.20}$$

In both the above expression and throughout the remainder of this section, we are following Einstein notation and summing over repeated indices. Note that $|\beta\rangle$ refers to a state labelled by $n_1 l_1 n_2 l_2\ {}^{2S+1}L_J$ and m_j. Below, when we write J or m_j we are referring to the values of these quantum numbers associated with the state $|\beta\rangle$ (and similarly for $|\beta'\rangle$, J', and m_j').

The summation over β' in Eq. 2.19 can be simplified by observing the appropriate symmetries and angular momentum structure. Specifically, $\boldsymbol{d}$ is a 3-vector, so $\alpha_{\mu\nu}(\beta;\,\omega)$ is a rank-2 Cartesian tensor and is thus reducible to scalar, vector, and tensor irreducible parts. Following the derivation in [183, 318], we arrive at scalar, vector, and tensor polarizabilities of:

$$\alpha^{(0)}(\beta;\,\omega) = \sum_{\beta'} \frac{2}{3}\frac{\omega_{\beta\beta'}|\langle\beta\,||\boldsymbol{d}||\,\beta'\rangle|^2}{\hbar(\omega_{\beta\beta'}^2 - \omega^2)} \tag{2.21}$$

$$\alpha^{(1)}(\beta;\,\omega) = \sum_{\beta'} (-1)^{J+J'+1}\sqrt{\frac{6J(2J+1)}{J+1}}\begin{Bmatrix} 1 & 1 & 1 \\ J & J & J' \end{Bmatrix}\frac{\omega_{\beta\beta'}|\langle\beta\,||\boldsymbol{d}||\,\beta'\rangle|^2}{\hbar(\omega_{\beta\beta'}^2 - \omega^2)} \tag{2.22}$$

$$\alpha^{(2)}(\beta;\,\omega) = \sum_{\beta'} (-1)^{J+J'}\sqrt{\frac{40J(2J+1)(2J-1)}{3(J+1)(2J+3)}}\begin{Bmatrix} 1 & 1 & 2 \\ J & J & J' \end{Bmatrix}\frac{\omega_{\beta\beta'}|\langle\beta\,||\boldsymbol{d}||\,\beta'\rangle|^2}{\hbar(\omega_{\beta\beta'}^2 - \omega^2)}, \tag{2.23}$$

where $\{\vdots\,\vdots\,\vdots\}$ is the Wigner 6-j symbol. These polarizabilities produce a combined energy shift of:

$$\begin{aligned}\Delta E(\beta;\,\omega) = &-\alpha^{(0)}(\beta;\,\omega)|\mathcal{E}^+|^2 \\ &-\alpha^{(1)}(\beta;\,\omega)(i\boldsymbol{\mathcal{E}}^+\times\boldsymbol{\mathcal{E}}^-)_z\frac{m_j}{J} \\ &-\alpha^{(2)}(\beta;\,\omega)\left(\frac{3|\mathcal{E}_z^+|^2 - |\mathcal{E}^+|^2}{2}\right)\left(\frac{3m_j^2 - J(J+1)}{J(2J-1)}\right),\end{aligned} \tag{2.24}$$

[18] We will use similar notation throughout this work, where a bold symbol refers to a vector, the equivalent non-bold symbol with a subscript denotes one component of the vector, and the equivalent non-bold symbol without a subscript denotes the magnitude of the vector.

where the subscript z denotes the component of the indicated vector along z (recalling that m_j is the quantum number associated with J_z).[19] One can think of the first line in Eq. 2.24 as the component of the polarizability that transforms to itself under rotations (namely the spherically symmetric part), the second line is the component that transforms like a vector, and the third line is the irreducible rank-2 component. Note that Eq. 2.24 also applies in the DC limit where $\omega = 0$.[20]

We have reduced the problem of computing the energy shifts imposed on an atom by an electric field of the form in Eq. 2.9 to an expression involving only terms like $\omega_{\beta\beta'}$ and $\langle \beta \,||\boldsymbol{d}||\, \beta' \rangle$. The former terms can be computed directly from the unperturbed energy levels of the atom. The latter terms are known as reduced dipole matrix elements, and are dependent only on the radial part of the electronic wavefunction. The reduced dipole matrix elements can be related to the Einstein A coefficients [318]:

$$|\langle \beta \,||\boldsymbol{d}||\, \beta' \rangle|^2 = \frac{3\pi\epsilon_0\hbar c^3}{\omega_{\beta\beta'}^3}\frac{2J'+1}{2J+1}A_{\beta\beta'}, \tag{2.25}$$

where $A_{\beta\beta'}$ is the decay rate (in population) from the state $|\beta'\rangle$ to $|\beta\rangle$ and c is the speed of light. The values of $A_{\beta\beta'}$ have been extensively characterized and documented in experiments, making it possible to compute $\Delta E(\beta;\, \omega)$ to great accuracy.[21]

2.2.3 Magic wavelengths

Using Eqs. 2.21 to 2.25, and tabulated values of the A coefficients for Sr [378], one can compute the energy shift of different electronic states in Sr as a function of the wavelength of the applied laser field. This is shown for a linearly polarized laser field aligned along the quantization axis z in Fig. 2.2.[22] Naturally, close to specific resonances between electronic states the energy shifts diverge. However, operating away from such resonances is beneficial when trying to make a conservative trapping potential using laser fields. The benefit of operating away from resonance can be seen by returning to the simplified discussion in Sec. 2.2.1:

[19] Note that we typically apply a magnetic field in our experiments, in which case it is natural to define z to be parallel to the applied field.

[20] Also note that the vector contribution vanishes for linear polarization. Since a DC electric field must be linearly polarized, this term vanishes for DC electric fields. Both the vector and tensor terms vanish for states with $J = 0$.

[21] It is worth keeping in mind that the discussion in this section is only second order in perturbation theory. In practice, the approximations made are accurate for regimes we're interested in, where the coupling between different states is far smaller than $\omega_{\beta\beta'}$ and/or ω. For optical dipole traps, it would be hard indeed to enter a regime where the Rabi frequency is comparable to an optical frequency.

[22] We thank Ross B. Hutson for generously providing us with example code in which he performed similar calculations for ^{87}Sr.

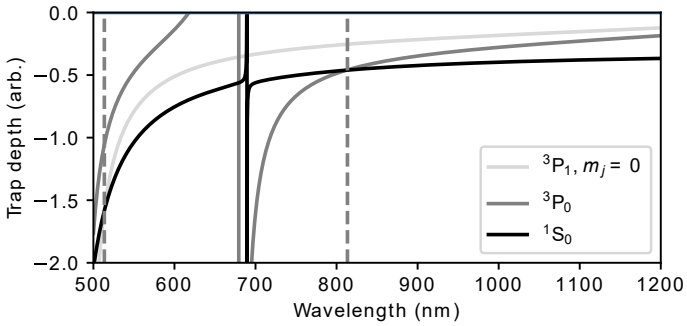

Figure 2.2: Polarizabilities of the $^1S_0$, $^3P_0$, and $^3P_1$ ($m_j = 0$) states. The total energy shift (including contributions from the scalar, vector, and tensor polarizabilities) for the indicated electronic states are shown (in arbitrary units) for an AC electric field which is linearly polarized along the quantization axis (z). The intensity of the field is fixed, and the wavelength varied. We refer to wavelengths at which the energy shift of two levels is equal as "magic" wavelengths. The magic wavelengths near 515 nm for the $^1S_0$ and $^3P_1$ states, and 813 nm for the $^1S_0$ and $^3P_0$ states, are indicated by the dashed lines, and are used extensively in this work.

when coupling to an excited state with a finite decay rate Γ, the overall rate at which photons are scattered off of an atom in the presence of a drive field is $\propto \Gamma\Omega^2/\Delta^2$, whereas the conservative potential due to the light shift is $\propto \Omega^2/\Delta$. As a result, the ratio between the depth of the conservative potential and photon scattering strictly improves with increased Δ.

One notices in Fig. 2.2 various wavelengths at which the energy shift for two electronic states are equal, and so the applied laser field does not perturb the transition energy between the two states. These wavelengths are referred to as "magic" wavelengths. Of particular note to us are the magic wavelengths for the $^1S_0$ and $^3P_0$ states at 813.42757(62) nm [4], and for the $^1S_0$ and $^3P_1$ states near 515 nm [73]. In Sec. 4.1 we will see how it is possible to achieve a stronger magic condition where the $^1S_0$, $^3P_0$, and $^3P_1$ states simultaneously experience the same energy shifts. We will return to the discussion of optical trapping in more detail as appropriate throughout this work, especially with respect to high fidelity ground state cooling of atoms (Ch. 4), and manipulations of the clock transition (Ch. 5).

Chapter 3

Overview of the apparatus

"[It's not that the glass is half full or half empty.] We're engineers. It means the glass has been manufactured to the wrong specifications."

— Host Bernard

Portions of this chapter have appeared in:

M. A. Norcia, A. W. Young, and A. M. Kaufman. Microscopic control and detection of ultracold strontium in optical-tweezer arrays. **Phys. Rev. X**, 8(4):041054, Dec. 2018

A. W. Young, W. J. Eckner, N. Schine, A. M. Childs, and A. M. Kaufman. Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice. **Science**, 377(6608):885–889, Aug. 2022

A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman. An atomic boson sampler. **In review**, 2023

We saw in Ch. 2 that our tools for controlling Sr atoms amount to the targeted delivery of DC and AC electromagnetic fields to the atoms. The design of the optical trapping potentials are particularly important: we are interested in trapping large ensembles of atoms at small length scales, which enhances the energy scales associated with experiments involving high fidelity optical cooling (Ch. 4, and all subsequent experiments), Rydberg interactions (Ch. 6), and tunneling (Chs. 7 and 8). Additional considerations associated with avoiding undesirable perturbations to the atoms lead to the use of several different optical potentials, including optical tweezer arrays and optical lattices at two different wavelengths. Achieving the above properties while also maintaining single-particle-resolved control using the tweezers leads to tradeoffs between the design of the optical lattices, and the high resolution microscope used to generate the tweezers and image the atoms. All the optical potentials must stay referenced to each other at the $\lesssim 50$ nm level, leading to various techniques for maintaining passive stability of the alignment between the different potentials, as well as techniques for active feedback.

The vacuum chamber in which our experiments take place is designed to deliver the aforementioned optical potentials (and other control fields) to the atoms without degradation, while otherwise isolating the atoms from any perturbations. In particular, one is interested in delivering thermal atoms to the control region while maintaining ultra-high vacuum (UHV) conditions (to avoid collisions between stray particles and the atoms), and in keeping surfaces of the vacuum chamber far from the atoms (to avoid unwanted interactions between the atoms and electromagnetic fields emanating from these surfaces).

For various operations, we are also interested in applying strong DC magnetic fields (of up to 550 G) and gradients (of up to 50 G/cm). This necessitates the design of high current magnetic field coils that do not prohibitively restrict optical access to the control region. One could also be interested in controlling the DC electric field at the atoms using control electrodes and Faraday shielding [353]. However, for the experiments described in this work, we find it sufficient to simply keep dielectric surfaces relatively far away from the atoms.

In this chapter, we will summarize our solutions to the above challenges in our apparatus (Fig. 3.1), including the design of the vacuum chamber and atom source in Sec. 3.1, the magnetic field coils in Sec. 3.2, and the optical systems in Sec. 3.3. Finally, in Sec. 3.4, we will briefly discuss the typical sequence of operations that are common to any given experiment in this work — namely how an ensemble of thermal atoms are prepared and measured (this will be discussed in more detail as necessary in subsequent chapters).

3.1 Vacuum chamber and atom source

The core components of our apparatus are shown in Fig. 3.2, including the vacuum chamber, and surrounding magnetic field coils and optical components. Here, we summarize key features relating to the design of the vacuum chamber.

3.1.1 Science chamber

We will refer to the main control region where our experiments take place as the "science chamber" or "cell." The role of the science chamber is to provide UHV conditions,[1] while also allowing various optical

[1] In fact, our vacuum borders on extremely high vacuum (XHV) conditions, with pressures below 10^{-11} Torr.

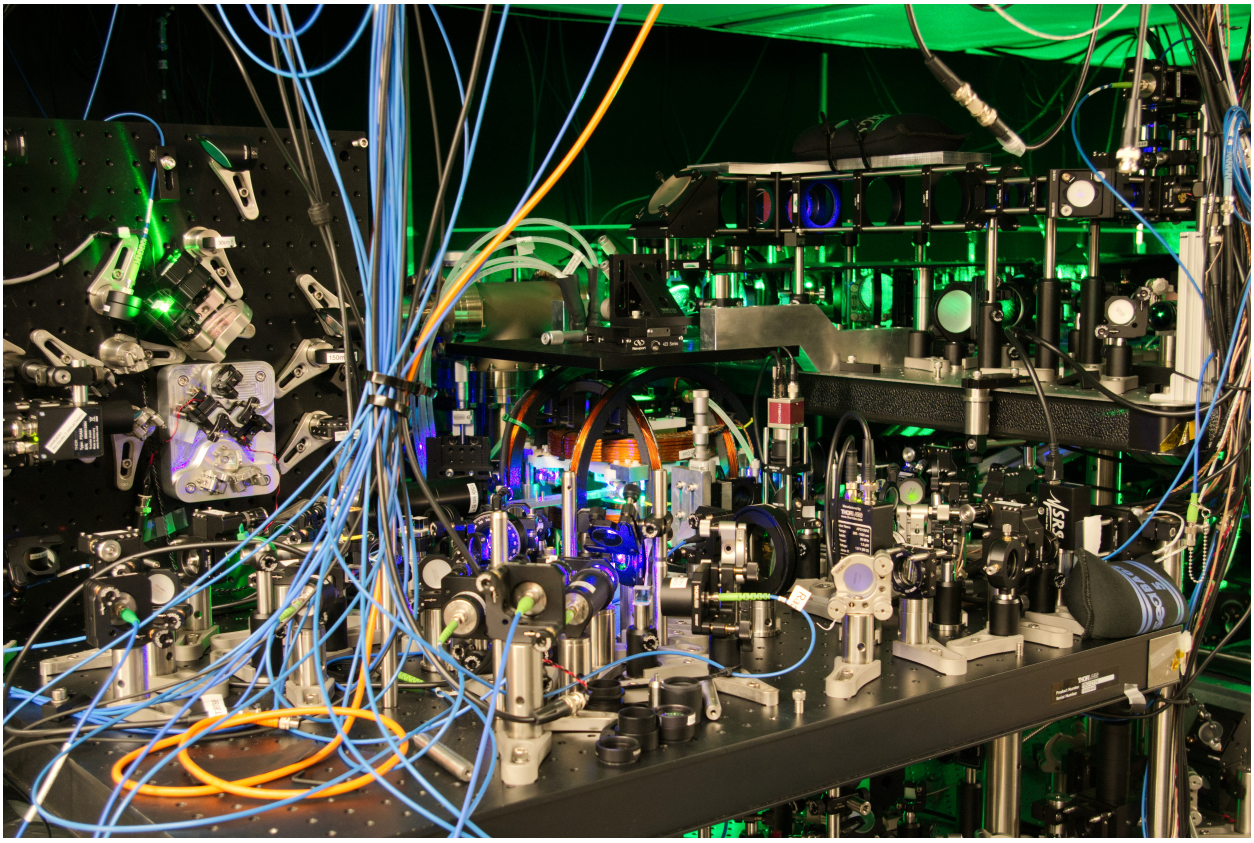

Figure 3.1: Picture of the apparatus circa 2021. To see a timelapse of the first few months of construction, visit doi:10.5281/zenodo.10019752.

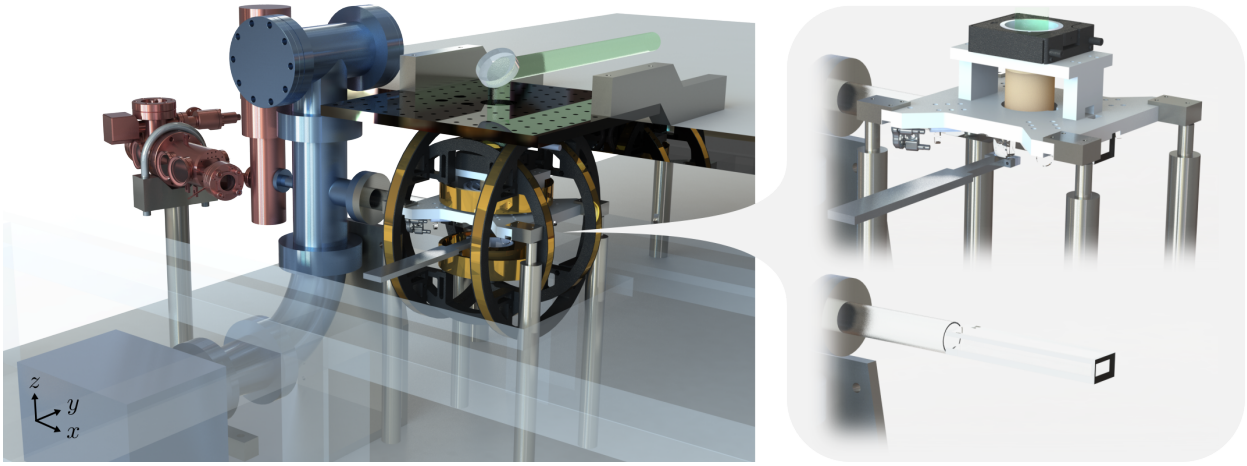

Figure 3.2: Layout and detailed views of the apparatus. The atom source is highlighted in red, and the vacuum pumping chamber in blue. The magnetic field coils are shown in gold. The callout shows detailed views of components that are obscured by the magnetic field coils, including the optical reference plate (and microscope objective and lattice optics, top) and glass science chamber (bottom).

control fields to be delivered to atoms in the chamber without degradation — specifically, without unwanted stray reflections or optical aberrations. To provide UHV conditions, the chamber must be constructed out of materials with low outgassing, and that are compatible with a high temperature bake. To provide low optical aberrations, the chamber must have good optical properties, and must not undergo uncontrolled deformations that change these properties when it is evacuated. For the above reasons, we opt to use a cuvette style vacuum cell, manufactured by Japan Cell, that is constructed entirely of fused silica glass (see Fig. 3.3).

The use of fused silica is helpful due to its optical properties, being highly transparent for the full range of wavelengths that we work with (including wavelengths as low as 317 nm), and its low coefficient of thermal expansion (CTE) of 0.48×10^{-6}/K, which ensures that small fluctuations in the temperature of the experiment do not lead to significant changes in optical performance. A non-magnetic glass-to-metal seal is used to connect the fused silica cell to the metallic components of the vacuum chamber, ensuring that the entire chamber can be baked to high temperatures without cracking, and that magnetically active materials are kept away from the region of the experiment that experiences high magnetic fields. The latter feature helps to minimize uncontrolled fluctuations in the magnetic field that are detrimental to our experiments (see Sec. 3.2).

The cell is anti-reflection (AR) coated on the outside surfaces with a broadband coating that spans from < 450 nm to ~ 950 nm for angles of incidence between 0 and 45 degrees from normal (Fig. 3.4a).[2] The inner surfaces are not AR coated in order to allow the cell walls to be optically contacted together, instead of glued together with epoxy. This all-glass construction allows the cell to be baked at higher temperatures than a glued cell (up to 200 °C), leading to higher quality vacuum.[3]

The lack of AR coatings on the inner surfaces of the cell requires one to ensure that unwanted reflections do not intersect with the laser-trapped atoms. By using a design where most beams enter the cell at a large angle of incidence, we can ensure that reflections up to the first several bounces do not intersect

[2] We purchased a separate cell explicitly for the experiments involving 317 nm light in Ch. 6, since these coatings are not expected to work at 317 nm. However, miraculously, we never had to install the second cell due to an error in the manufacturing of the original cell: during the AR coating run, it was mounted in such a way that the properties of the coating on the end facet are substantially different from on the other surfaces, yielding $\sim 95\%$ transmission of 317 nm light.

[3] Note that another option which has become viable since the original design of our apparatus is to use a nano-textured coating (e.g. from TelAztec) that can survive higher temperatures than typical coatings, and to assemble the cell via a fritting process. Precision Glassblowing has now constructed such cells for a number of groups.

with the atoms.[4] An additional consideration for the geometry of the cell is that it should be large enough that surfaces are relatively far away from the atoms, both to improve atom capture into the magneto-optical traps (MOTs, see Sec. 3.1.4), and to ensure that the atoms are not strongly perturbed by stray charges and dipoles residing on the glass surface (see Sec. 2.1.1.5). Finally, the cell should not be so large that a microscope objective placed outside the vacuum chamber must be prohibitively large to achieve a given numerical aperture (NA).

In order to achieve diffraction-limited resolution of the microscope (Sec. 3.3.2), the windows of the cell must obey certain additional requirements with respect to their flatness, thickness, and parallelism. For our apparatus, the target specifications for the cell are as follows: a flatness of $\leq \lambda/10$ over 20 mm (where $\lambda = 461$ nm is the wavelength of the imaging transition); a window thickness within ~ 100 μm of the specified value;[5] and a parallelism corresponding to a variation of ≤ 10 μm for two sides of the same piece of glass, and ≤ 100 μm for opposing pieces of glass,[6] over the 80 mm length of the cell. Note that the above tolerances must be maintained when the cell is under vacuum. Based on finite element analysis (in COMSOL) we opt to use 5 mm-thick cell windows, for which the deformation of the cell under vacuum is acceptably small (Fig. 3.4b).

Taken collectively, the above considerations lead to the cell geometry shown in Fig. 3.3. Upon receiving the cell, we confirmed that the manufacturing was within tolerances via measurements using a homemade white light interferometer, and some measurements on a Zygo GPI Fizeau Interferometer.

3.1.2 Cold atom source

To deliver atoms to the science chamber, we use a Sr Beam II source from AOSense (see Fig. 3.2). This atom source consists of an oven containing several grams of strontium, a capillary array to collimate the thermal beam of atoms released from the oven (yielding a half-angle of 17 mrad), a Zeeman slower, and two stages of 2D magneto-optical traps (MOTs). The thermal beam released by the oven terminates on a

[4] This is primarily a concern with the lattice beams, where small stray reflections can lead to large amounts of disorder in comparison to the tunneling energy (see Ch. 7). The large (transverse) magneto-optical trapping beams (Sec. 3.1.4) also enter with an angle of incidence of > 30 degrees, and miss the atoms on the first several bounces.

[5] Upon receiving the cell, we measured this to be ~ 20 μm using a white light interferometer.

[6] Note that for a single microscope objective, the specifications for opposing pieces of glass are irrelevant to the performance of the experiment.

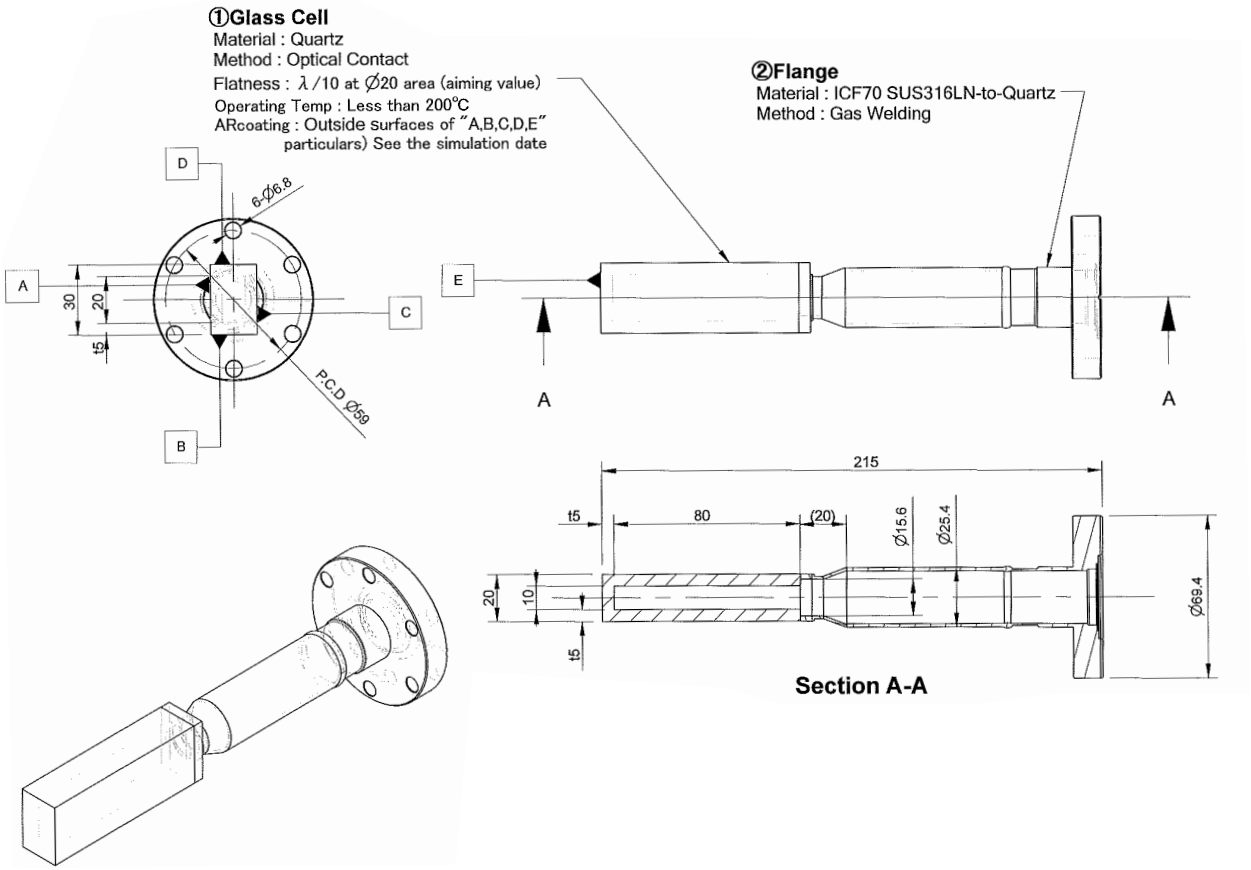

Figure 3.3: Manufacturing drawing for the glass cell.

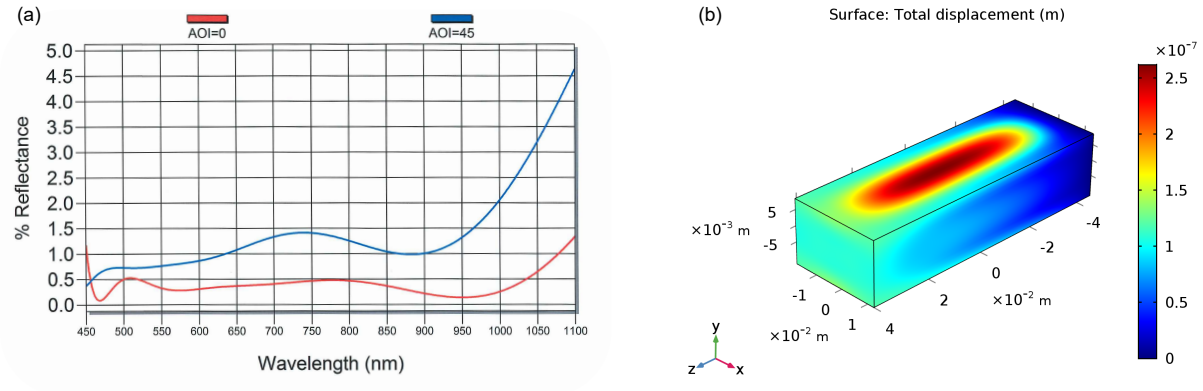

Figure 3.4: Glass cell anti-reflection coating and deformation under vacuum. (a) Measurements of the performance of the anti-reflection coating applied to the outside of the glass cell, including the performance at a 45 degree angle of incidence (AOI). (b) Simulated deformation of the glass cell when evacuated, and under atmospheric pressure (12.1 psi because we're in Boulder), showing maximum deviations of $\lesssim 25$ μm across the entire cell.

heated window (operated at 350 °C to avoid coating the window and reducing its transmission). The two stages of 2D MOTs are tilted relative to the thermal beam, such that only atoms that are appropriately cooled transversely are directed towards the "cold" output aperture of the atom source, and subsequently to the science chamber. When the laser beams forming the 2D MOTs are off very few atoms in the thermal beam make it to the cold output aperture, allowing us to rapidly turn the atom source on and off with no moving parts.

We typically operate the oven at comparably low temperatures of 380 °C, where the atom flux directly out of the oven is 3.6×10^{11} atoms/s. The atoms exiting the cold output of the source have a transverse temperature of < 3 mK and an axial speed ~ 40 m/s,[7] and are thus easily captured directly in a 3D MOT acting on the broad $^1S_0 \leftrightarrow {}^1P_1$ transition in Sr (Sec. 3.1.4). Despite the presence of the oven, the pressure attained in the atom source is fairly low, with a value of 1.6×10^{-9} Torr (as measured directly at the oven).[8] The atom source includes an ion pump and two getter pumps, which help to achieve this low pressure. Additionally, because Sr is an effective getter material, much of the chamber acts as a getter pump once the oven has been activated for a while.

3.1.3 Vacuum system

In order to attain UHV conditions, one must balance outgassing rates (set by the materials and total surface area in the chamber, as well as by the oven in the atom source) with the pumping rate, which is set by the conductances through the chamber as well as the properties of the pump used. Our experiments all take place in the molecular flow regime,[9] where the conductance is only dependent on the geometry of the chamber and the temperature. At room temperature and in the molecular flow regime, the conductance through a pipe is $C \simeq 12.1 d^3/l$, where d is diameter of the pipe and l its length (both in cm), and C is in units of l/s (liters per second).

The atom source is attached to the science chamber via a pumping region to attain even higher quality vacuum in the science chamber (see Fig. 3.2). Between the atom source and the pumping region is

[7] Note that the atoms exiting the cold output make up only a small fraction ($\sim 0.5\%$) of the flux exiting the oven.

[8] We find that low oven temperatures are very important to the lifetime of an optically trapped atom in the vacuum chamber. Operation at 380 °C yields a lifetime of ~ 400 s, whereas operation at 420 °C yields a reduced lifetime of ~ 160 s.

[9] Meaning the mean free path of a gas particle in the vacuum chamber is significantly longer than the dimensions of the chamber.

a differential pumping tube with a diameter of 6.1 mm and a length of 95 mm, leading to $C \sim 0.3$ l/s. The geometry of the pumping region and the science cell is such that $C \gtrsim 5$ l/s through the whole system.[10] The pumping region is constructed out of a large diameter vacuum cross, tee, and elbow. The tee houses a titanium sublimation pump (TSP, Gamma Vacuum 3 filament), which we fired only once after the initial bake of the chamber.[11] At the end of the elbow is a large ion pump (TiTan DI 75S), which has a pumping speed of 75 l/s.

After assembling the vacuum chamber, we baked the metal components at ~ 175 °C, and the science chamber at ~ 125 °C, all for 15 days.[12] We degassed the TSP during the bake, and then fired it after allowing the chamber to slowly return to room temperature over the course of about two days. Although a residual gas analyzer (SRS RGA100/12) was used during the bake, it is not permanently installed on the chamber, and we opt not to include an additional vacuum gauge. Instead, we use the current of the ion pump as a readout, which tells us that the vacuum pressure is $\lesssim 1 \times 10^{-11}$ Torr, and sufficient for our experiments.

3.1.4 Magneto-optical traps

The cold atom beam is trapped directly in the science chamber using a "blue" 3D MOT operating on the broad $^1S_0 \leftrightarrow {}^1P_1$ transition, which cools the atoms to $T \sim 1.5$ mK (slightly higher than the doppler-limited temperature of $T \sim 720$ μK) We then transfer the atoms into a sawtooth wave adiabatic passage (SWAP) MOT [238, 21, 229] operating on the narrow $^1S_0 \leftrightarrow {}^3P_1$ transition to rapidly reduce their temperature to $T \simeq 10$ μK, before performing an additional round of cooling using a conventional "red" MOT also operating on $^1S_0 \leftrightarrow {}^3P_1$. The red MOT compresses the atom cloud to an atom density of $\sim 10^{11}$ cm^{-3}, and cools it further to $T \simeq 5$ μK. Under typical operating conditions, we trap $\sim 10^7$ atoms in the blue MOT in ~ 20 ms, and recapture $\sim 10^5$ atoms from the blue MOT in the red MOT in ~ 100 ms (including ~ 50 ms of operating the SWAP MOT, and ~ 50 ms of operating the red MOT). The general mechanisms and schemes for magneto-optical trapping of Sr are described in detail elsewhere, including in [319] for the red and blue

[10] This is primarily limited by the science chamber. The conductance through the rest of the system is significantly higher, with $C \sim 50$ l/s.

[11] It is not entirely clear that this was necessary for achieving high quality vacuum, but it was included out of an abundance of caution.

[12] Most of the chamber was baked using heater tape and a makeshift oven constructed out of fiberglass insulation and vacuum foil, but we also used a steel canister to protect the science chamber.

	Transverse beams		Vertical beams		Magnetic gradient
	Power (mW)	Waist (mm)	Power (mW)	Waist (mm)	(G/cm)
Blue MOT	10	10	3	2	50
SWAP MOT	0.5	3	0.25	2	4.2
Red MOT	0.1	3	0.05	2	4.2

Table 3.1: Representative conditions for magneto-optical trapping. Note that these are approximate values, and can vary depending on the specific experiment performed.

MOTs, and in [236] for the SWAP MOT. The final temperatures and atom densities attained using a similar cooling scheme to the one outlined above are carefully optimized in [311].

The main additional challenge in our apparatus is the fact that the high NA microscope objective results in limited optical access on one side of the cell. Although it is possible to operate MOTs with skewed beams that do not intersect at right angles, we instead opt to project the downwards-pointing MOT beams through the objective itself, by focusing these beams on the back focal plane of the objective. Due to space constraints on the apparatus, we are limited to using a $2''$ diameter lens with a 250 mm focal length, yielding a beam waist[13] of ~ 2 mm when collimated by the objective. In the transverse directions, the blue MOT beams have a waist of ~ 10 mm,[14] and the red MOT beams a waist of ~ 3 mm. Although the parameters of the MOTs change throughout our works, representative values of the various beam powers and magnetic gradients appear in Tab. 3.1.

Because the blue MOT beams are at the same wavelength as the imaging light, they cannot be separated from the imaging path exiting the microscope objective using a dichroic mirror. Instead, we use a custom 90:10 wedged non-polarizing beam splitter (Perkins Precision Developments) that results in a negligible loss of photon collection efficiency in the imaging path, at the cost of requiring extra (but readily available) laser power in the downwards facing MOT beams (meaning the beams projected through the objective, see Fig. 3.7). Note that although the downward facing beam is required for a stable blue MOT, this beam is not strictly required for the SWAP MOT and red MOT, since gravity provides a restoring force in the downward direction.[15]

[13] Namely, the radius at which the intensity of the beam drops to $1/e^2$ of its maximal value.
[14] But are clipped in the vertical direction to the 10 mm height of the interior of the science chamber.
[15] This works in the case of the narrow line MOTs, since the temperatures achieved are such that the atoms do not undergo excursions that are large compared to the size of the science chamber [319].

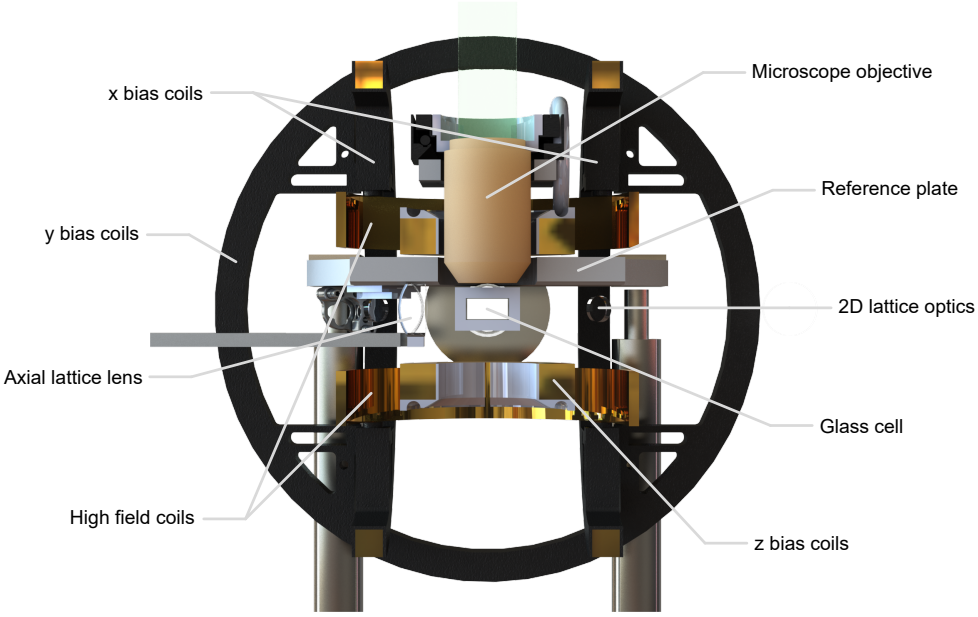

Figure 3.5: Cross-section of the science chamber and surrounding components. Note that the lattice optics (except the axial lattice lens, which is still on a temporary mount visible as a grey bar on the left of this figure) and microscope objective are mounted directly to the reference plate, which is mechanically isolated from the coils.

The position of the red MOT can be controlled on length scales smaller than the MOT beams by simply applying bias magnetic fields which shift the position of the magnetic quadrupole that defines the position of the MOT. We make use of this to align the MOT to our various optical trapping potentials, and to load atoms from the MOT into the optical traps, as described in more detail in Sec. 4.2.

3.2 Magnetic field coils

In order to apply the appropriate magnetic fields and gradients, the chamber is surrounded by a set of magnetic field coils. In particular, we are interested in applying magnetic gradients of up to 50 G/cm, magnetic fields in the vertical (z) direction of up to 550 G, and moderate bias fields of ~ 30 G that can point in an arbitrary direction. The bias fields are applied using pairs of bias coils in the x, y, and z directions, where z is the optical axis of the microscope objective, and x and y are transverse to that axis (see Fig. 3.2). The large fields and gradients are applied using an additional pair of high current coils oriented along the z axis.

In order to ensure homogeneous magnetic fields across the atom array (except when explicitly applying magnetic field gradients), all pairs of coils are arranged close to a Helmholtz configuration, except for the z bias coils, which are separated by 1.7 times their radius due to geometric constraints.[16] Although the z bias coils are fairly far from the Helmholtz condition, we are less sensitive to gradients imposed by these coils, since the resulting gradients are primarily normal to the atom array. The suboptimal design of the z bias coils is because they were originally intended to be used only for applying our MOT gradients. However, out of a concern of heating the microscope objective,[17] we instead use these coils only at low currents. To provide the requisite magnetic gradients, we use an H-bridge circuit to switch the high current coils between a Helmholtz and anti-Helmholtz configuration as needed (see below and Appendix C.4).

The coils and their relative orientation with respect to the vacuum chamber can be seen in Fig. 3.2, and in more detail in Fig. 3.5. The relevant parameters associated with the coils are summarized in Tab. 3.2. The x and y bias coils are designed to be large to avoid unwanted gradients, and use many turns of wire both to allow for the use of low current power supplies, and to avoid heating. The z bias coils are water cooled via an aluminum backplate, and the high current coils are cooled by running water through the core of the hollow wire used in their construction. Each high current coil is composed of two windings with 25 turns each that are connected electrically in series using jumper bars, but in parallel for the water cooling [236]. This approach improves the water flow rate, and is critical for attaining sufficient cooling power. All coils are constructed by wrapping wire around a support structure, and potting them in place with epoxy (Araldite 2011). The support structure is 3D printed nylon plastic for all coils except the inner z coils, which are potted directly to the aluminum cooling plate.

The bias coils are all powered by a JILA designed current controller [52]. The high current coils are powered by a Delta electronica SM30-200 DC power supply, which is capable of outputting 200 A at 30 V. To switch the high current coils between the Helmholtz and anti-Helmholtz configurations, we use a custom H-bridge circuit that switches the polarity of one of the two coils (see Appendix C.4). Note that although the coils are operated at low voltages, there is a significant amount of energy stored in the magnetic field when they are activated. If they are disconnected suddenly, the stored energy can result in dangerous spikes

[16] Namely having to fit in between the high current coils and microscope objective.
[17] A concern that arose only after I melted our first microscope objective.

	x bias	y bias	z bias	High field
Diameter	9″	10.5″	2.6″	6″
Wire	20 AWG	20 AWG	20 AWG	3/16″ hollow square
Turns per coil	200	200	160	50
Resistance (Ω)	4.72	5.88	1.22	0.1
Helmholtz field ($\mathrm{G\,A^{-1}}$)	17.35	15.18	27.07	6.15

Table 3.2: Magnetic field coil parameters. The listed Helmholtz fields are for when the pairs of coils are wired in series (keep in mind that the z bias coils are relatively far from the Helmholtz condition). Other parameters are listed for each individual coil.

in voltage that can destroy electrical and human components [236]. Because of this, the H-bridge circuit includes an interlock that ensures that the polarity of the coil cannot be switched unless the current is below a safe value. A 1 Ω flyback resistor is also included in order to dissipate the accumulated current, allowing the coils to be switched on and off relatively quickly (on a timescale of ~ 25 ms).

The above electrical sources provide fairly stable currents, with fractional root mean squared (RMS) fluctuations at or below the 10^{-4} level, translating to similar stability in the applied magnetic field. However, magnetization of components near the coils can serve as another source of fluctuating magnetic fields. To mitigate this effect, most components near the science chamber are non-magnetic (glass, plastic, ceramic, aluminum, titanium, etc.). Although this limits the effect of magnetization, we still encounter issues with drifting fields when the fields are left at a high value (> 250 G) for a long time (minutes). We suspect that this is due to magnetization of the upper optics breadboard in the experiment (see Fig. 3.2), which is constructed using stainless steel. To monitor the drifting fields, we use a 3 axis magnetometer (AlphaLab VGM) installed near the science chamber. We find our magnetic field stability to be sufficient for most applications, including those involving precision spectroscopy (Ch. 5). However, further reducing the use of magnetic materials near the science chamber, and providing additional magnetic shielding, would be useful for experiments involving high fields and long coherence times.[18]

[18] For example, experiments in which we wish to increase the clock Rabi frequency to perform fast single and two qubit gates, and then use the resulting entangled states for stable frequency comparisons with a laser (see Chs. 5 and 6).

3.3 Optical systems

The optical design of the apparatus is centered around a custom manufactured microscope objective (from Special Optics), which both allows us to efficiently collect photons scattered from the atoms in a spatially well-resolved manner, and to project tightly confining and reconfigurable "optical tweezer" potentials on the atoms. However, arrays of optical tweezers often suffer from disorder in the intensity and shapes of the tweezers due to high sensitivity to optical aberrations in the entire tweezer system. Additionally, optical tweezers are not particularly efficient with respect to optical power, since each tweezer is an independent tightly focused beam of light. This can be contrasted with optical lattices, which are more efficient with respect to laser power,[19] and can have very low disorder.[20]

As a result, we integrate both kinds of optical potentials in our experiment, and realize motional state-preserving, low-loss transfer between these different potentials.[21] We will loosely classify these potentials as "preparation" and "science" potentials. The preparation potential is a tweezer array operating at 515 nm, where a magic trapping condition can be achieved for the $^1S_0 \leftrightarrow {}^3P_1$ cooling transition at 689 nm via tuning of a magnetic field (see Sec. 4.1). 515 nm is a convenient wavelength for generating large arrays of tightly confining tweezers because of the short wavelength (which allows one to generate tweezers with finer spatial features), the availability of large amounts of laser power ($\gtrsim$ 10 W), and the high polarizability of ground state Sr near this wavelength (of $\sim 900 a_0^3$ in atomic units, where a_0 is the Bohr radius).

The science potential is either a tweezer array or a 3D optical lattice operating at 813 nm, which is a magic wavelength for the clock transition in Sr (see Sec. 2.2.3). The 3D optical lattice is composed of a 2D lattice which is in the x-y plane (the same as the tweezer arrays), and an "axial" lattice that is normal to this plane (oriented along the z axis). Note that 813 nm is less appropriate than 515 nm for initial trapping and cooling due to lower available laser power ($\sim 5 - 8$ W, although this may change in the future [97]), longer wavelength, and lower polarizability of the ground state ($282 a_0^3$ in atomic units [39]). As a result, the science tweezers are typically only used during phases of the experiment where shallow traps are preferable,

[19] Since the same light is reused to make many tightly confining optical traps, and because a retroreflected lattice generates features with a resolution that is akin to having a numerical aperture of 1.

[20] Because they are generated using an interference effect, without using complicated high NA optics.

[21] As had been accomplished in tweezer arrays [198], but not between tweezer arrays and lattices prior to our works.

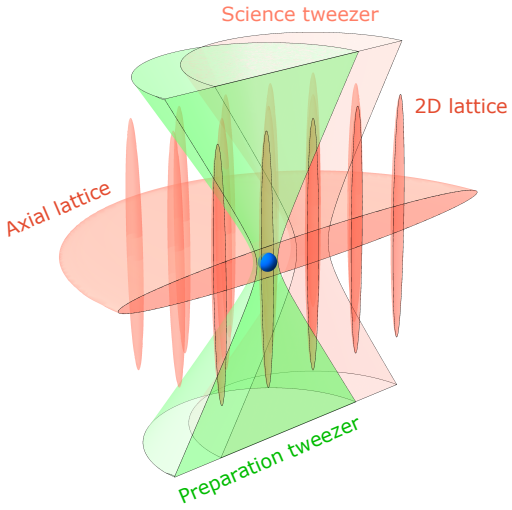

Figure 3.6: Cartoon of the optical potentials applied to the atoms. These potentials include a 3D optical lattice (of which we use only a single 2D layer) composed of axial and 2D lattices, the 515 nm "preparation" tweezers (green) used primarily for initial trapping and atom rearrangement, and the 813 nm "science" tweezers used primarily for clock interrogation. Both sets of tweezers are also used to programmably modify the lattice potential as needed during experiments involving quantum walks in the lattice. For illustration purposes, a single atom (blue sphere) is shown being placed by a preparation tweezer into a minimum of the 3D lattice.

and therefore where these constraints do not impose a limitation on atom number or state preparation. In contrast, the enhanced efficiency of the lattice potential more than makes up for these shortcomings, allowing us to generate thousands of traps that are significantly more tightly confining than the preparation tweezers. Although we use a variety of different combinations of tweezers and lattices in our experiments, a cartoon of the typical combination of optical potentials and how they interface with each other is shown in Fig. 3.6.

3.3.1 Mechanical design

In order to use the various optical potentials described above together, it is imperative that the position of the lattice potential be well-referenced to the positions of the potentials projected through the microscope objective. For example, a variation in the alignment between these potentials at the $a/4$ level, where $a \simeq 813/\sqrt{2}$ nm is the lattice constant of the 2D optical lattice, would cause the experiments appearing in Ch. 7 to fail. The optical clock laser used in the experiments in Chs. 5 and 6 must also be well-referenced to the positions of the atoms at a similar level of precision. Our approach to addressing this challenge is

twofold: we can actively feed back on the position of the tweezers or the phase of the optical clock laser to correct for these drifts, but we additionally want the experiment to be as passively stable as possible to minimize the requirements on this feedback. To improve the stability of the experiment, we provide a single rigid reference surface for all optical beams whose positions (and possibly orientations) at the atoms must be well-controlled. This "reference plate" (see Figs. 3.2 and 3.5) is constructed using Macor glass ceramic due to the rigidity (with a Young's modulus of 66.9 GPa at room temperature), low CTE (9.3×10^{-6} C^{-1} at room temperature), and insensitivity to magnetic field[22] of this material.

To first order, the position of potentials projected through the objective are determined by the position of the objective itself,[23] and so the objective is attached directly to the reference plate. Similarly, the lenses (and folding mirrors) that define the position of the 2D optical lattice are attached to the reference plate.[24] The optical clock laser does not have fine spatial structure, and so we are insensitive to transverse fluctuations in the position of this beam at the location of the atoms. However, it is critical that the spatial phase of this beam at the location of a given atom be consistent in a given experimental trial. To address this, we perform fiber noise cancellation (See Appendix D.2.1), using a mirror that is attached to the reference plate as a proxy for the position of the atoms.

Although the above design addresses most fluctuations in alignment as a result of vibrations of the various mechanical components in the experiment, additional drifts can occur as a result of air currents, which affect the optical path experienced by the various laser beams, and thermal fluctuations, which cause the various mechanical and optical components in the experiment to change their properties. As a result, great care is taken to ensure that the temperature of the experimental apparatus is stable. Due to the large heat load in the lab from various electronic components, the room air temperature is stabilized by pumping cold air from the building into the lab, leading to peak to peak fluctuations of ~ 0.6 K under normal operating conditions.[25] The optical tables are completely enclosed and have less heat load than the rest of the room. Therefore, their temperature can be controlled by electrically reheating the air that is sent

[22] Important because the reference plate sits in the region of the apparatus that experiences high magnetic fields.

[23] Although changes in the angle with which the beams enter the objective can result in differential motion between the objective and the optical potentials, such drifts are strongly suppressed.

[24] Although the initial design intention was to glue the lens defining the position of the axial lattice in place after aligning it with a 5 axis kinematic mount (Newport 9081), we found that the stability of the axial lattice was acceptable without attaching this lens to the reference plate. As a result, we have so far opted to avoid performing this irreversible procedure.

[25] Normal meaning not during one of the many floods and power outages that have occurred in JILA during my PhD.

to the tables via several high efficiency particulate air (HEPA) filters, leading to peak to peak fluctuations in the table temperature of ~ 0.05 K assuming the heat load on the tables doesn't change.[26] To ensure that the stabilized air to the tables doesn't lead to large air currents, the HEPAs are sized for a linear flow rate of up to 0.11 m/s over the tables (with 0.05 m/s being typical for day to day operation). This corresponds to a flow rate of up to 700 cfm for the "science" table which houses the main experimental apparatus, and 350 cfm for the "laser" table which houses many of the lasers involved in the experiment (see Appendix D).

3.3.2 Microscope system

Our microscope objective is custom-made by Special Optics, with a specified numerical aperture (NA) of 0.67, a working distance of 12 mm, an effective focal length of 23.1 mm, an aperture of 27 mm, and a diffraction-limited field of view of 75 μm.[27] The objective is designed to achieve the above specifications when imaging through the 5 mm thick windows of the science chamber (Sec. 3.1.1). Given the diverse range of wavelengths used in the experiment, the objective is also designed to achieve diffraction limited performance between 461 nm and 950 nm. Over this range of wavelengths, the transmission efficiencies of the objective are above 87% (measured to be 91% at 515 nm),[28] and the focal shift due to chromatic aberrations is below 2 μm. Like other components in the high magnetic field region of the experiment, the objective is constructed of non-magnetic materials, namely glass, epoxy, and ULTEM PEI plastic.

When designing high NA objectives for cold atom experiments,[29] it is also important to consider realistic limitations on the manufacturing of the vacuum windows, and the degree to which alignment can be achieved and maintained. Our objective can maintain diffraction limited performance given errors in the window thickness of ± 20 μm by adjusting only the working distance, and errors of ± 100 μm if compensating with both the working distance and the collimation of the input (or output) beam.[30] Additionally, the objective is insensitive to angular misalignment below 0.02 degrees with respect to the cell window. In order

[26] Larger fluctuations can occur when opening the table doors, or installing new components. For the most sensitive experiments we have performed, we find it helpful to avoid opening the doors leading into the lab, since a perturbation to the room temperature and air currents can also lead to measurable changes in the temperature of the optical tables.

[27] The objective acts like an f-tan(θ) lens, and has negligible distortion and field curvature over the entire field of view.

[28] Note that although reflections off of the various surfaces in the objective are not significant in terms of efficiencies, they do make it challenging to project optical lattices through the objective, since even small stray reflections can interfere with the lattice beams and result in unacceptable amounts of disorder for studying Hubbard physics [13].

[29] Specifically, objectives that sit outside the vacuum chamber, although other approaches which put parts of, or all of, the objective inside the chamber come with their own strengths and weaknesses [312].

[30] These considerations lead to the specifications for the science chamber in Sec. 3.1.1

to ease the alignment process, the front surface of the objective is flat, allowing us to use a reference beam projected from below the science cell, and normal to the cell surface, to align the objective normal to the cell with the requisite precision (by measuring the reflection off of the objective).[31] Note that this is the only alignment step required for installing the objective in the experiment, since all other beams and optical potentials will be aligned to the objective itself.

Upon receiving the microscope objective, we performed a series of tests to characterize its performance, including interferometric measurements with a dielectric bead and an optical flat, and direct measurements of the point spread function (PSF) of the objective using a pinhole. We found the latter measurements to be more effective, since it was challenging to ensure that other optical components in the interferometric measurements were of similar or better quality than the objective itself. Specifically, in a test setup that includes both the objective and a sample of the science chamber window material, we align a 350 nm diameter pinhole to the focus of the objective, illuminate the pinhole from the rear, and image the pinhole through the objective. Deconvolving the measured image with the shape of the 350 nm pinhole yields a direct measurement of the PSF of the objective. Scanning the position of the pinhole also allows one to measure the field of view of the objective in a straightforward manner.

Note that the ideal PSF of the microscope system is an Airy pattern, which corresponds to when the aperture of the microscope is uniformly illuminated.[32] When imaging a point source, the ideal PSF yields an intensity distribution of:

$$I(\boldsymbol{\rho}) \propto \left(\frac{R}{\rho} J_1(kR\rho/f)\right)^2 \tag{3.1}$$

where J_1 is the first order Bessel function of the first kind, $\boldsymbol{\rho}$ is the coordinate in the image plane, $k = 2\pi/\lambda$ is the wavenumber of the light (with wavelength λ), R is the radius of the aperture, and f is the focal length of the lens.

We Fourier transform the measured PSF to attain an estimate of the field in a Fourier plane (where the aperture of the system is defined), and fit the resulting distribution with the Zernike polynomials, which are a complete orthonormal set of basis functions on the unit disk. The lowest order Zernike polynomials

[31] The objective is mounted on a Newport HVM-2u positioner, which provides the appropriate controls for alignment.
[32] In other words, it's the Fourier transform of a disk. See [125] for a very nice introduction to Fourier optics.

correspond to well known optical aberrations (defocus, astigmatism, coma, etc.). Understanding these aberrations allowed us to identify and correct a manufacturing error in the original objective, and is generically helpful once the objective is installed in the experiment to correct misalignments in the system. The final two objectives we received were both within specification based on the above tests, with a Strehl ratio of $\gtrsim 0.8$ across the entire field of view.[33] The associated PSF has an effective beam waist of 440(20) nm when imaging the atoms on the $^1S_0 \leftrightarrow {}^1P_1$ transition at 461 nm (Fig. 3.7b).

To align the objective in the experimental apparatus, we first install the tracer beam described above that is normal to the cell. With the objective removed from the apparatus, a tweezer array can be aligned to the tracer beam. The objective can then be installed, and its angle aligned using the tracer beam. An alignment mirror (with markings that denote the size of the entrance pupil) is then placed at the rear of the objective (on the side facing away from the cell, and reflecting the tweezer beam), and shimmed in place such that it retroreflects the tweezer beam. This mirror now serves as a reference for the optical axis of the objective (as well the position of its entrance pupil) that can be removed and reinstalled at will. Both the tweezer beams and imaging system can be coarsely aligned using this reference.[34] Final alignment is performed using the trapped atoms themselves as a probe, by maximizing trap frequencies and minimizing the PSF measured in images of the atoms. The imaging and tweezer paths are separated using a custom wedged[35] dichroic mirror from Laseroptik, which reflects 461 nm and transmits 515 nm, 680-705 nm, and 813 nm, when oriented at a 30 degree angle. The layout of the imaging and tweezer systems are shown schematically in Fig. 3.7a.

Note that a relatively large magnification is required to match the PSF of the imaging system to the size of the pixels on our camera (see Sec. 4.4). Due to challenges associated with aligning a multi-element imaging system without degrading the performance limits imposed by the objective, we instead opt to use a single off-the-shelf achromatic lens with a 1.9 m focal length (Edmund Optics 70-163) to achieve the desired

[33] The Strehl ratio is a useful coarse-grained metric for evaluating the performance of an optical system, and is defined to be the ratio between the peak value of the PSF of the system, and the peak value of the ideal PSF in Eq. 3.1. For small aberrations (namely due to wavefront errors of less than a wavelength), the Strehl ratio is directly related to the wavefront error of the system [153]. A Strehl ratio of > 0.8 is typically accepted as an operational cutoff for being "diffraction limited," with 1 being the maximum possible value.

[34] In the latter case, we use the same trick of once again installing a tracer beam that traverses the optical path in reverse.

[35] The wedge is important to avoid etaloning in the tweezer and imaging systems.

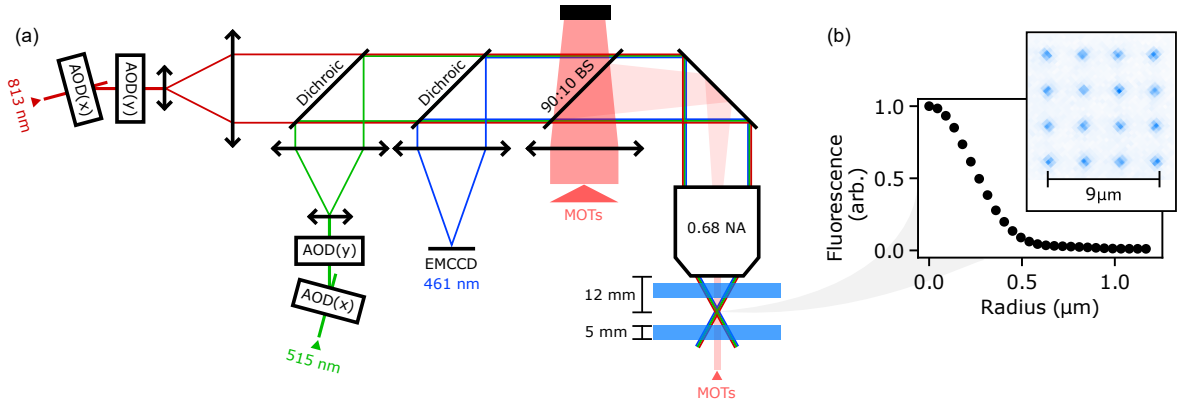

Figure 3.7: Simplified layout of the microscope system and its performance. (a) The 515 nm and 813 nm tweezer systems are generated using acousto-optic deflectors (AODs) (see Fig. 3.8 for more detail), combined on a dichroic mirror, and separated from the 461 nm fluorescence from the atoms using another dichroic mirror. The downwards-facing MOT beams come to a focus at the back focal plane of the objective, and are combined onto the optical path using a $90:10$ beam splitter (BS). The objective focuses (or collimates) the various beams, and has a numerical aperture (NA) of 0.68, a working distance of 12 mm, and is corrected to achieve diffraction-limited performance when focusing through a 5 mm thick piece of glass. (b) An averaged image (at 461 nm) of atoms trapped in a 4×4 array of optical tweezers is displayed, illustrating the PSF of the microscope system. In the background is a radial average of the PSF, yielding a waist of 440(20) nm.

magnification.[36]

3.3.3 Optical tweezers

The optical tweezers are simply tightly focused beams of light projected through the microscope objective. Typically, we do not uniformly illuminate the aperture of the objective, and instead illuminate it with a Gaussian beam, which results in an intensity profile of:

$$I(\boldsymbol{\rho}, z) = I_0 \left(\frac{w_0}{w(z)} \right)^2 e^{-2\rho^2/w(z)^2}, \tag{3.2}$$

$$w(z) \coloneqq \sqrt{1 + (z/z_R)^2}. \tag{3.3}$$

Here, z is the coordinate along the optical axis, and $\boldsymbol{\rho}$ is the coordinate transverse to z. I_0 is the maximum intensity of the beam, $z_R = \pi w_0^2/\lambda$ is the Rayleigh range, $w(z)$ is the beam waist as a function of position along the optical axis, and w_0 is the minimum beam waist. The actual intensity profile of the tweezers deviates only slightly from Eq. 3.2 due to the finite size of the microscope aperture. Per the discussion in

[36] This lens has a (perhaps unnecessarily) large diameter of 128 mm, as it was originally designed for use in refractive telescopes. This makes alignment of the imaging system quite straightforward.

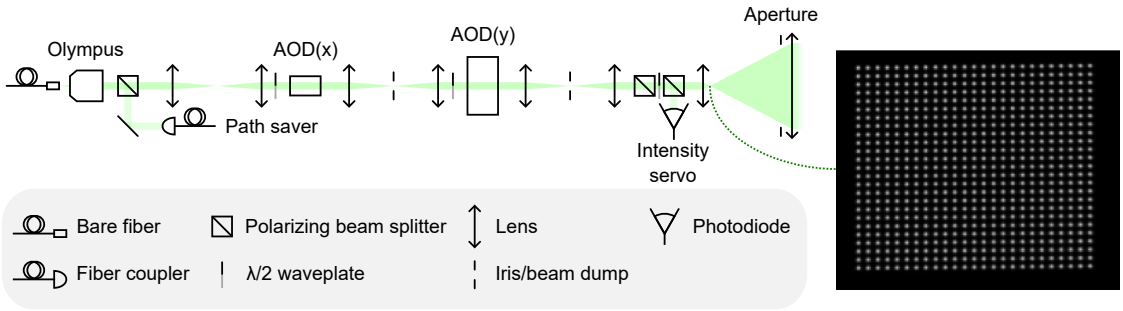

Figure 3.8: Layout of the tweezer systems. Note that similar designs are used for both tweezer systems, and that some optical elements are omitted for simplicity. The first relay from the left is actually a 3 : 4 telescope in order to match the beam size to the optimal size for the AODs, and the waveplates are used to align the beam polarization to the appropriate axis of the AODs. The right-most waveplate is used to set an appropriate pickoff fraction for the intensity servo. The path saving fiber allows the bare input fiber to be replaced or repaired without spoiling the alignment of the rest of the tweezer system (since a newly installed fiber can just be aligned to the path saver). A typical tweezer array used in our experiments, imaged before the objective, is shown on the right.

Sec. 2.2.1, for the appropriate choice of laser frequency the atoms experience a conservative potential that is proportional to Eq. 3.2.[37]

We are interested in generating many such tweezers that can be rapidly reconfigured in a (mostly, see Sec 3.3.3.2) independently controllable fashion, while also maintaining relatively efficient use of laser power. Throughout the entire system used to generate the tweezers, great care must be taken to minimize aberrations that degrade the PSF of the tweezer system (resulting in features that are larger than the diffraction-limited performance of the microscope objective).[38] The system used to generate arrays of tweezers is shown schematically in Fig. 3.8. Laser light originates from a fiber, and is collimated using a high NA lens with low aberrations (Olympus RMS10X for the 515 nm system, and RMS4X-PF for the 813 nm system[39]). In order to increase the optical power available, we opt to use a custom photonic crystal fiber (PCF) with a mode expander (from Tratech Fiberoptics).[40] The collimated beam is passed to an acousto-optic deflector (AOD), then to a second AOD oriented at 90 degrees to the first, and finally to a telescope that expands the beam(s) to fill the aperture of the objective.[41] Each element described above is separated from the next

[37] For atoms that are well-cooled, the resulting potential is effectively harmonic, since the center of a Gaussian function is well approximated by a quadratic function.

[38] This amounts to controlling the wavefront of the laser through an optical system that spans several meters at a level of significantly below a wavelength, which is a testament to the impressive performance of modern optical components.

[39] Note that, due to poor transmission at 813 nm, the RMS4X-PF was later replaced with a Thorlabs C280TMD-B.

[40] We do not find the additional structure in the output mode of the PCF to degrade the performance of the tweezer system in comparison to a normal fiber.

[41] All optical elements that are in the beam path after the beam is expanded (where we are most sensitive to imperfect optics

by a $4f$ relay,[42] which images the optical field at the output of the previous element onto the input (or image plane) of the next element. The integrated intensity of the tweezer array is stabilized using a pickoff[43] located after the second AOD, and actuated with an acousto-optic modulator (AOM) located upstream of the fiber.

AODs (and AOMs) are effectively programmable diffraction gratings, where an acoustic wave launched into a crystal results in a spatially varying strain, and corresponding variation in the index of refraction. The frequency f_a of the modulation sets the wavelength of the acoustic wave $\lambda_a = f_a/v$ (where v is the speed of sound in the crystal), and thus the pitch of the diffraction grating $d = \lambda_a/2$. For diffraction to occur, the incident angle θ of the laser with respect to the grating (in this case the k-vector of the acoustic wave) must be close to the Bragg condition:

$$n\lambda = 2d\sin(\theta), \tag{3.4}$$

where n is the diffraction order (typically 1). Deviations from this phase matching condition cause the diffraction efficiency to drop, which sets the range of modulation frequencies over which an AOD can operate.

The linearity of the above process[44] means that one can drive an AOD with multiple radio frequency (RF) tones to generate multiple superimposed diffraction gratings, each of which generates a separate spot whose deflection angle can be controlled via the frequency of the drive tone. This allows one to turn a single laser beam into many, whose positions can be controlled independently on a line. However, because the grating in an AOD is moving, the diffracted light is shifted in frequency by the frequency of the acoustic wave, resulting in tweezers that are at different frequencies.[45] Although we want to make many tweezers, we also want to ensure that beating between nearby tweezers that overlap slightly does not cause significant parametric heating (see Sec. 5.3.1). For the trap frequencies we typically work with (see Sec. 4.3.3), we want adjacent tweezers to be separated in frequency by more than 1 MHz, and so the number of tweezers we can

introducing aberrations) were tested carefully on a Zygo interferometer (the same one used in Sec. 3.1.1). This yielded several surprises. For example, securing an otherwise flat $2''$ diameter mirror in a standard kinematic mount using a set screw resulted in deviations as large as $\lambda/2$, compromising the tweezers. We found it necessary to instead glue the optics in place using slow curing epoxy.

[42] Two lenses, each of which performs a Fourier transformation, resulting in the optical field at the input of the relay being replicated at its output.

[43] Composed of two polarizing beam splitters separated by a $\lambda/2$ waveplate (see Fig. 3.8).

[44] Although, in real devices the radio frequency components and acousto-optic response of the crystal are not always nicely linear (see Sec. 3.3.3.2).

[45] For a full discussion of the photoelastic effect and coupled mode theory, see [361].

make is primarily limited by the RF bandwidth of the AODs.

To increase the RF bandwidth of an AOD, one can use a phased array of acoustic transducers rather than a single transducer. The relative phase of the transducers controls the angle at which the acoustic wave is launched into the crystal, and so appropriately tailoring this phase as a function of frequency allows one to match the condition in Eq. 3.4 for a large range of different frequencies (without changing the alignment of the laser relative to the AOD). We use custom designed AODs from IntraAction which use this design principle to attain an RF bandwidth of 90 MHz about a center frequency of 180 MHz.[46] To provide reasonable grating pitches for these operating frequencies, the AODs excite longitudinal rather than shear waves in tellurium dioxide (TeO2) crystals.[47] At 515 nm (813 nm), the diffraction efficiency is 85% (70%) at the optimal RF drive power of 1.6 W (2.8 W), the center value of the deflection is 21.8 mrad (34.4 mrad), and the deflection range is 10.9 mrad (17.1 mrad).[48]

The tweezer arrays at 515 nm and 813 nm are combined on a custom dichroic mirror from Laseroptik that reflects 461 nm and 500-600 nm, and transmits 650-850 nm, when oriented at 45 degrees from the incident beams. Based on both direct measurements of the tweezer system, and measurements of the trap frequency experienced by the atoms (see Sec. 4.3.3), we infer that the tweezers generated at 515 nm (813 nm) have a waist of $w_0 = 480(20)$ nm ($w_0 = 740(40)$ nm), which is close to (but slightly larger than) the limit set by the performance of the objective. Note that the number of tweezers we can currently make at 515 nm is limited by thermal lensing in the tweezer system, which precludes optical powers in excess of 2 W out of the fiber. We suspect this is due to lensing in the first AOD, and that using a different material like flint glass would improve the power handling of the system.[49]

[46] The number of resolvable spots given the ~ 4 mm beam diameter used at the AODs ends up being ~ 80 (this is independent of operating wavelength, and applies to both tweezer systems). Ignoring limitations on laser power, this would allow for the generation of several thousand independent tweezers.

[47] The longitudinal waves have a speed of sound of 4260 m/s, in comparison to 617 m/s for shear waves. The higher speed of sound results in longer wavelengths for a given modulation frequency, and thus for larger frequency separations between gratings with a given difference in their pitch.

[48] Note that the higher drive power required by the 813 nm AODs necessitates the use of a heat sink. Without this, thermal changes in the optical properties of the crystal make the alignment of the tweezer system unstable (see next paragraph).

[49] Often, a high photoelastic constant comes hand in hand with sensitivity to thermal lensing, since the latter effect can be related to a change in the index of refraction as a function of strain (although other mechanisms such as the thermo-optic effect sometimes also play a role). In places where we value optical power handling over other properties of the AOD or AOM (for example in the AOM upstream of the fiber used to stabilize the tweezer intensity), we will often use flint glass rather than TeO2.

3.3.3.1 Tweezer electronic control

When generating multiple tweezers by means of driving the AODs with multiple RF tones, one wants the entire system to be as linear as possible. Although the AO response is inevitably non-linear when optimizing for diffraction efficiency, and operating in the regime where the AOD is close to saturation, the high power RF amplifiers used to drive the AODs can also be a significant source of non-linearity. Our amplifiers (Delta RF LA10-1-525-40-A1) are chosen to have a saturation power (10 W 1 dB gain compression point, 20 W saturation power) that will not immediately destroy the AODs, but also to have as linear a response as possible at the desired RF drive powers. In particular, these amplifiers are optimized to have a low third-order intermodulation distortion (IMD3) point of -30 dBc (measured for two tones with a spacing of 10 kHz and a peak envelope power of 10 W), which avoids extra drive tones appearing near the operating range of the AODs when operating at high power and with many tones.

One of the main sources of intermodulation when driving an amplifier with many tones is simply the maximum voltage that the amplifier is capable of outputting. If many tones phase up such that there is a large excursion in the signal in the time domain, this spike in voltage can get cut off. As a result, one can reduce or mitigate this source of intermodulation by judiciously choosing the relative phases of the drive tones such that these large excursions in the time domain are minimized. This is sometimes accomplished through brute force optimization by computing the signal in the time domain, but such approaches are costly with respect to computing resources, and can either slow down the experiment, or require one to choose from a limited selection of pre-computed signals [95]. Instead, we find that applying the following heuristic solution [295][50] provides signals with acceptably low intermodulation:

$$\phi_i = \phi_1 - 2\pi \sum_{j=1}^{i-1} (i-j) p_j, \tag{3.5}$$

where ϕ_i is the initial phase of the i^{th} tone, and p_i the fraction of the total signal power contained in this tone.[51] This simple approach to choosing the phases allows us to synthesize our tweezer control signals in real time without precomputing a time series (for more detail, see Appendix C.3).

[50] Other heuristic approaches with low computational overhead could improve the performance further [340].

[51] The intuition for this formula comes from the observation that the lowest peak factor signal corresponds to the case where the signal is generated using only frequency modulation. By approximating a target signal with a frequency modulated signal that simply jumps (in time) between the different frequency components of the target signal, and enforcing that the frequency modulated signal is continuous, one arrives at Eq. 3.5.

For the experiments in Chs. 7 and 8, we are interested in stabilizing the total power in the tweezer arrays (and lattices) using an additional AOM to modulate the laser power (see Fig. D.1). This stabilization must be done over a large dynamic range of $\sim 10^4$.[52] To achieve this dynamic range in the presence of the non-linear response of the tweezer system, we simply tolerate lower gain for deep traps that are closer to saturation (the electronics used in the experiment are described in brief in Appendix C). However, upgrades are underway to use a digital loop filter that provides greater flexibility than a typical analog loop filter, allowing one to calibrate out the non-linear response of the tweezer system electronically.[53]

3.3.3.2 Balancing large tweezer arrays

When making a square 2D array of tweezers, tweezers along the anti-diagonal of the array will be at the same frequency, since the frequency shifts imposed by the two AODs relative to the mean frequency of the array are equal in magnitude but opposite in sign.[54] This results in static interference patterns leading to uncorrectable modulations of the tweezer depths in large 2D arrays of tweezers. If the array is slightly rectangular, tweezers along the anti-diagonal can be at difference frequencies that are relatively low and can cause significant heating, even if the frequency difference between directly adjacent tweezers is large.[55] As a result, we typically space the two axes of our array differently, with a spacing of 2 lattice sites in one direction, and 3 in the other (where the lattice spacing is $a \simeq 575$ nm).[56]

In order to balance the depths of individual tweezers, we split off a small fraction of the light before the objective, and measure the integrated intensity per tweezer using a CMOS camera. By adjusting the relative power in the different RF tones applied to the AODs, it is possible to balance the total optical power in each spot to within peak to peak fluctuations of 5% for the rectangular pattern described above. The main limitation on this balancing is a lack of fully independent control over each spot. For an $N_x \times N_y$ array, each of the $N_x + N_y$ RF tones has independent amplitude control.[57] As a result, we only have $N_x + N_y$ degrees

[52] Note that when we are interested in applying a shallow tweezer, for example when applying the oracle in Sec. 7.5.4, we can mitigate this difficulty by making additional deep tweezers that are far away from the region of interest, and stabilizing the intensity of the entire tweezer array [48].

[53] Another option is to use an analog circuit that uses a logarithmic amplifier, as in [255].

[54] Assuming the two AODs and surrounding relays are identical.

[55] Avoiding this condition is important in the rearrangement procedure in Sec. 4.5.

[56] Other spacings and patterns are also used extensively throughout this work. See Sec. 4.5 for a discussion of what patterns are compatible with our current rearrangement procedure.

[57] Although the phases can also be controlled, they are constrained by Eq. 3.5.

of freedom to balance a tweezer array with $N_x \times N_y$ spots. In comparison, a 1D array can be balanced to much better precision, with a peak to peak variation of $\sim 1\%$. While it is possible to balance the tweezer depths using the atoms,[58] we find the measurements using a camera to be sufficiently accurate (since the variations in tweezer depth are currently dominated by a lack of independent control).

3.3.4 Optical lattices

The role of the optical lattice is to provide a static potential with many traps that are compatible with high fidelity imaging, cooling, and clock spectroscopy. Additionally, the traps should be sufficiently homogeneous to allow for experiments involving tunneling in large systems (Chs. 7 and 8). To fulfill the above roles, the lattice operates at 813 nm, which is a magic wavelength for the clock transition, and where a magic angle condition (Sec. 4.1) can be met for the narrow line cooling transition. However, to satisfy the magic angle condition, the polarization of the lattice must be uniform and linear. Such a condition is not possible using a retroreflected 3D lattice, since the polarization of a beam must be normal to its propagation direction, and so not all three sets of orthogonal beams can share the same polarization. Instead, we use a 2D bowtie lattice in the plane of the tweezer arrays (the x-y plane), and a 1D crossed-beam lattice along the remaining "axial" or z axis (Fig. 3.9).

3.3.4.1 2D lattice

The requirements on the 2D lattice are twofold: given fixed laser power, we want to generate many lattice sites that provide sufficiently strong confinement for high quality imaging and optical cooling. Additionally, we want our imaging system (with a waist of 440(20) nm) to be able to resolve individual sites in the lattice.[59] To balance the above considerations, we use a bowtie design (Fig. 3.9a) which increases the lattice spacing by a factor of $\sqrt{2}$, and the lattice depth by a factor of 8 in comparison to a retroreflected lattice (using the same total amount of laser power). The resulting potential is given by the interference

[58] E.g. via light shifts (Sec. 7.3) or measurements of the trap frequency (Sec. 4.3.3).

[59] Note that our imaging resolution is potentially good enough to use a retroreflected lattice design at 813 nm, especially if using more sophisticated image analysis techniques like those in [149].

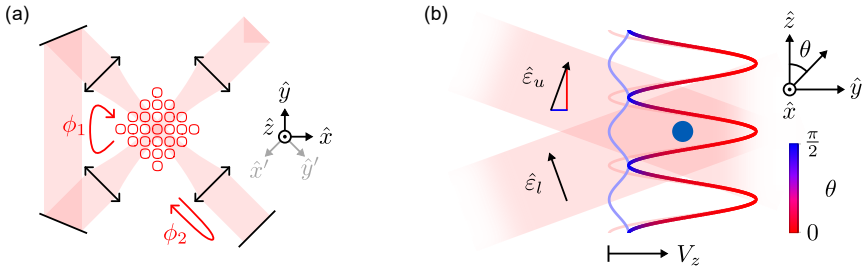

Figure 3.9: Optical lattice geometry and polarization. (a) Schematic of the 2D bowtie lattice, showing folding mirrors (black lines) and $4f$-relays (double arrows). $\phi_{1,2}$ denote the phase accrued when traversing the indicated optical path, and the red contours denote an equipotential surface of the lattice, which is polarized along the z axis. (b) Schematic of the crossed-beam axial lattice, showing the upper and lower lattice beams, with polarizations $\hat{\varepsilon}_u$ and $\hat{\varepsilon}_l$. Coloring denotes the polarization of the resulting standing wave, which lies in the y-z plane at an angle θ from the z-axis. The two polarization components of the lattice beams result in two standing waves with orthogonal polarizations. For the shallow crossing angle of 36 degrees between the two beams used in our experiment, the depth V_z of the z-polarized standing wave dominates. As a result atoms (blue disk) sit at anti-nodes of this lattice, and nodes of the y-polarized lattice.

between four plane waves:

$$V(x',y') = -\frac{V_0}{16}\left|e^{ikx'} + e^{i(ky'+\phi_1)} + e^{i(-ky'+\phi_1+\phi_2)} + e^{i(-ky'+2\phi_1+\phi_2)}\right|^2 \tag{3.6}$$

where V_0 is the maximum depth of the lattice, and $k = 2\pi/813$ nm. As shown in Fig. 3.9a, x' and y' are coordinates that are aligned with the propagation direction of the lattice beams, and rotated with respect to x and y, and $\phi_{1,2}$ are the phase shifts accrued in the folding arms of the lattice.[60] When considering the trapping frequencies that can be achieved in the lattice, the enhancement in depth more than compensates for the reduced lattice constant, allowing for improved optical cooling and imaging (Ch. 4). Note that the shared paths of the four beams generating this lattice ensures that the intensity pattern is stable [296], but this pattern can shift in position if ϕ_1 or ϕ_2 changes. However, the spatial phase of the lattice in the y' direction (normal to the retroreflecting mirror) is stable with respect to the position of the retroreflecting mirror (although the phase in the x' direction is sensitive to the position of all the mirrors).[61]

The bowtie lattice is composed of a pair of $4f$ relays that are folded to cross at 90 degrees at their foci, where the beam waists are 55 μm (see Fig. 3.9a). To ensure that all arms of this lattice fully interfere, the polarizations of these beams are oriented out of the plane of the lattice (in the z direction). With 540 mW of input power, this yields average lattice depths of $416E_\mathrm{R}$ (where $E_\mathrm{R} = h \times 3.43$ kHz is the recoil energy of an 813 nm photon), and corresponding average radial trap frequencies of 98(1) kHz over a $26 \times 26\,\mu$m region (corresponding to 45×45 sites).[62] Representative values of the depth of the 2D lattice during different phases of the experiment are summarized in Tab. 3.3.

3.3.4.2 Axial lattice

The axial lattice is composed of two parallel elliptical beams which are focused by an aspheric lens (Edmund Optics 48-537)[63] such that they cross at their foci, where the beam waists are 25×10 μm, with the 25 μm axis oriented in the plane of the 2D lattice (along the y axis, see Fig. 3.9). These beams cross

[60] Note that, in practice, we must also take the imperfect transmission of the lattice optics and cell into account, which results in unbalanced optical powers in the different arms of the lattice, and slight deviations from Eq. 3.6.

[61] The retroreflecting mirror could serve as a better reference for the optical clock laser than the current reference off of the optical reference plate (Sec. 3.3.1) if the clock laser entered from this direction.

[62] In more recent experiments, we use ~ 1.5 W of optical power to achieve radial trap frequencies of ~ 160 kHz.

[63] The use of an aspheric lens is important in this application if one is interested in dynamically varying the spacing of the axial lattice.

	E_R	E_R^{Ax}	E_R^{2D}	J_0	kHz (E/h)
$1E_R$	-	10.47	2	21.05	3.43
$1E_R^{Ax}$	0.095	-	0.19	2.01	0.328
$1E_R^{2D}$	0.5	5.24	-	10.53	1.72
J_0	-	-	0.095	-	0.163
Axial full	423	4430	846	-	1450
Axial tunneling	2.48	26	4.97	52.3	8.52
2D full	416	-	832	-	1430
2D pre-tunneling	11.9	-	23.7	249	40.7
2D tunneling	2.5	-	5.0	52.6	8.58

Table 3.3: Representative values of the lattice depth during different phases of the experiment. These energies E are provided as a frequency E/h, where h is Planck's constant, as well as in units of the tunneling energy J_0 (see Ch. 7), and the recoil energies of an 813 nm photon, and the axial and 2D lattices (E_R, E_R^{Ax}, and E_R^{2D}, respectively). Note that the provided values are representative values (specifically for the experiments in [366]), but that the lattice depth, and value of the tunneling energy J_0, can differ for different experiments performed in this work.

with an angle of 36 degrees, yielding a lattice constant of 1.32 μm. For these beams to interfere fully their polarization should be oriented in the x direction, normal to the plane in which they cross. However, this would result in a polarization that is orthogonal to that of the 2D lattice. Although it is possible to find an approximate magic-angle magnetic field that is shared by both these lattice polarizations, in this situation the atoms experience very large polarization gradients. For example, when an atom moves from a node to an antinode in either the 2D or axial lattices they experience a 90 degree rotation in the polarization of the light field. This results in substantial broadening of the $^1S_0 \leftrightarrow {}^3P_1$ cooling transition, and thus reduced cooling and imaging performance.

Instead, we orient the polarization of the axial lattice beams along the z axis before the final focusing lens. This results in two components of the lattice with orthogonal polarizations, and spatial phases that differ by π (Fig. 3.9b). Because the crossing angle of our beams is well under 90 degrees, the axial lattice is predominantly z-polarized. As a result, atoms sit at the anti-nodes of the z-polarized lattice, and at the nodes of the y-polarized lattice. As long as the atoms remain cold, their experience of the y-polarized component of the axial lattice is suppressed to first order in their position. In this regime the entire 3D optical lattice, composed of the 2D bowtie lattice and the 1D crossed-beam lattice, is effectively z-polarized, and a robust magic-angle magnetic field condition can be found. This enables high-fidelity resolved sideband cooling and imaging (Sec. 4.3.3), with the primary tradeoff being a slightly reduced effective lattice depth of 90% in comparison to an x-polarized axial lattice. Nonetheless, with 660 mW of total optical power split between the two arms of the axial lattice we achieve average lattice depths of $423E_R$, and average axial trap frequencies of 41(1) kHz over a $26 \times 26\,\mu$m region.[64] The axial trap frequencies have a standard deviation of 7 kHz over this region due to the more tightly focused axial lattice beams in comparison to the 2D lattice beam.

One challenge with the crossed-beam configuration used in the axial lattice is that the relative phase of the two beams must be carefully controlled to preserve the spatial phase of the lattice. To hand atoms from the tweezers to the optical lattice, it is critical to ensure that the anti-node of the axial lattice lines up with the foci of the optical tweezer arrays, as described in more detail in section 3.3.5. To ease this process

[64] More recent experiments operate with ~ 1.6 W of total optical power, and axial trap frequencies in the range of 65 kHz.

we ensure that the optical path length of two arms of the axial lattice are matched to within $\sim 100\,\mu$m of each other, which keeps the spatial phase of the axial lattice passively stable on the several-hour timescale.[65] This is accomplished by sending both beams through the same optical elements, with the only exception being an optical flat whose angle is controlled with a piezoelectric actuator in order to tune the relative phase of the beams, and thus the spatial phase of the axial lattice.[66] To additionally improve the passive stability of this system, optical components that are not shared by both beams are attached to a single monolithic mount (which is visible as a silver rectangle on the left side of Fig. 3.1).

3.3.5 Alignment of optical potentials

Aligning the many optical potentials in the experiment is not a simple task. There are two tweezer arrays, four passes of the 2D lattice beams, and two arms of axial lattice that all need to be overlapped in a volume that is $\sim 20\ \mu$m on a side, and intersect with this volume at the appropriate angles. On top of this, we want to minimize the number of degrees of freedom in the optomechanical components of the apparatus in order to improve the passive stability of optical alignment under normal operation. As a result, introducing the requisite degrees of freedom judiciously, and also ensuring access to fast, robust signals for alignment, is critical for the success of the experiment.

There are a few tricks that we use extensively when aligning almost any sensitive optical system. When we want to save the alignment of a beam, for example when replacing an optical component, a single coupled fiber serves as an excellent reference for both the position and angle of a beam. As described in Sec. 3.3.2, we often use a tracer beam in combination with a retroreflecting mirror to establish a reference for the optical axis of a given system. We also often use machined components and dowel pins to register optical components, particularly components that are installed only temporarily during alignment (e.g. retroreflecting mirrors). In systems where there are beams that do not follow the optical axis, like in the axial lattice, it is often useful to have a permanently installed tracer beam that is on the optical axis to ease alignment to the atoms. For tracer beams that are used to align to the atoms, one typically wants

[65] We suspect that the remaining drift is due to drifts in the position of the final lens in the axial lattice as a result of the temporary mount currently used (see Sec. 3.3.1).

[66] Even in this case, it is important that similar optical flats are installed in both beams in order to match the optical path lengths.

to choose a wavelength that is sufficiently close to the operating wavelength of the system being aligned to avoid issues with chromatic aberrations. However, the chosen tracer wavelength must also have strong and obvious interactions with the atoms to provide a signal for alignment. Good choices for us are 461 nm and 688 nm, which can be used to deplete atoms from the MOTs, or to cause atoms to be lost (or appear to be lost) between two images.[67]

3.3.5.1 Tweezer alignment

As mentioned in Sec. 3.3.2, atoms in the preparation tweezers serve as the reference point to which all other optical potentials in our experiment are aligned. Aligning the 813 nm tweezers to the preparation tweezers is straightforward because they share a common reference, namely the microscope objective used to project them. This shared reference makes the alignment particularly robust, with negligible drifts on the 0.5 μm scale over the course of multiple days. To perform this alignment, one can simply install the retroreflecting mirror described in Sec. 3.3.2, and align the 813 nm tweezers to this reference. Finer adjustments (in particular of the collimation of the 813 nm tweezer system) can be performed by imaging atoms trapped in the 813 nm tweezers.[68] Final alignment of the two tweezer systems is performed by adjusting the RF control tones used to generate the 813 nm tweezers, optimizing for low atom loss and temperature after passing atoms between the two tweezer arrays (see Sec. 3.4).

3.3.5.2 Lattices to tweezers

The lattices are aligned by first aligning a tracer beam to the atoms, and then aligning the relevant lattice beam to the tracer.[69] Final alignment of the lattice beams is performed by maximizing the light shift induced by the lattice beams on the $^1S_0 \leftrightarrow {}^3P_1$ transition in 515 nm tweezer-trapped atoms (in magnetic fields that are not magic for the lattice polarization).[70] With the above optimization complete, one can

[67] 461 nm can be used to blow away atoms, and 688 nm can be used in combination with driving the $^1S_0 \leftrightarrow {}^3P_1$ transition to pump atoms into dark states (see Sec. 2.1.1.4).

[68] This imaging can be destructive and low fidelity, since one is just interested in measuring the shape of the average fluorescence signal.

[69] In the case of the 2D lattice, we first align the first and second pass of the lattice beams, and then align the retroreflection by maximizing the back-coupling into the fiber that supplies light to the lattice. If the alignment is very far off, one can align the back reflection before the first two passes are aligned, and look for a fluorescence signal from atoms trapped in the two resulting 1D lattices.

[70] If working with the 2D lattice it is useful to fully extinguish the tweezers during this measurement, before transferring atoms back into the tweezers for readout, or performing readout directly in the lattice. This approach ensures that the atoms

verify the alignment by maximizing the observed motional trap frequencies provided by the lattices, however these two procedures typically yield similar results.

One important consideration in our experiment is the possibility that the position of the lattice drifts relative to the tweezer array, and to the masks used in image analysis (Sec. 4.4.1), which can occur on a timescale of $\sim$ 30 minutes. To correct this, we use the images taken before and after rearrangement to identify the positions of the lattice-trapped atoms relative to the imaging system, and thus also relative to the optical tweezers. This information can be used to correct any drifts by adjusting the tweezer and mask positions on subsequent runs of the experiment. Because the repetition rate of the experiment is $\sim$ 1 Hz, corrections can be made much faster than drifts can occur. With the above feedforward procedure engaged, the RMS fluctuations in the lattice position relative to the tweezers is estimated to be $\lesssim a/10$. The tweezer positions can drift relative to the imaging system on a timescale of several hours, but this is readily corrected by intermittently taking images of atoms trapped in the tweezers and adjusting their positions to match the lattice.

3.3.5.3 Additional considerations for aligning the axial lattice

Alignment of the axial lattice is particularly challenging, since the foci of two separate lattice beams need to be overlapped with the tweezer array in a $25 \times 32 \times 10\,\mu$m volume. While this can be accomplished via the above procedure, the fringes of the resulting lattice must also be flattened relative to the plane of the tweezers while maintaining this overlap. Performing this flattening without misaligning the beam positions is challenging without the appropriate controls. For this reason, we introduce an additional $4f$ relay to the axial lattice before the final aspheric focusing lens. The relay allows us to install servo-actuated mirrors in both a Fourier plane and an image plane of the tweezer array. The Fourier plane mirror can be used to control a pure displacement of the lattice at the tweezer array, making it possible to compensate for drifts in beam pointing without spoiling the flattening of the lattice. This mirror is back-side polished, such that the transmission can be imaged onto a camera for in-situ monitoring of the axial lattice. The image plane mirror controls a pure tilt (rotation about x) of the lattice fringes, enabling flattening of the lattice in that

sit at the anti-nodes of the lattice, yielding consistent results, instead of being pulled away from the center of sites in the lattice due to imperfect alignment of the tweezers.

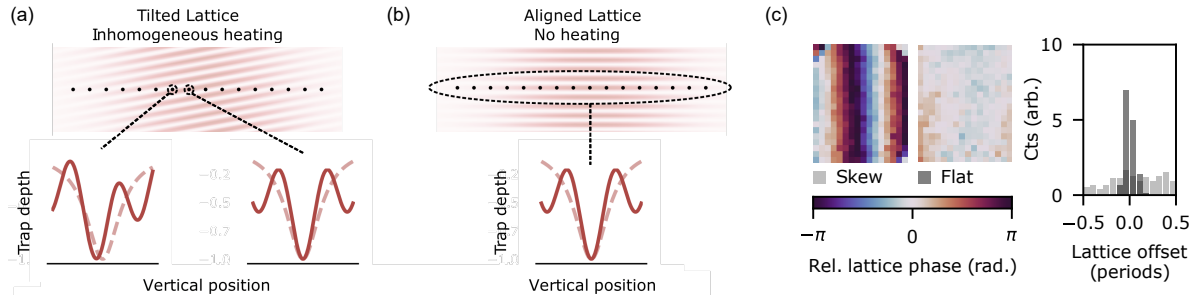

Figure 3.10: Axial lattice alignment via parametric heating. (a) The upper panel shows a side view of a misaligned axial lattice (red fringes) overlapped with the tweezer array (black dots). The lower plots show the tweezer potential (dashed curve) and the combined lattice and tweezer potential (solid red curve) at two locations. When the lattice fringe is tilted relative to the plane of the tweezers, it shifts the centers of the traps (as shown in the bottom left plot) in a position-dependent way. As a result, modulating the depth of the tweezers will cause different parametric heating rates at different sites in the array. (b) When the tweezers and lattice are well aligned, modulating the tweezers at the trap frequency does not cause substantial heating. (c) Measurements of the spatial phase of the axial lattice (via position-dependent heating rates) at each tweezer with an intentional tilt (skew) and properly aligned (flat) show that it is possible flatten the lattice relative to the entire tweezer array to within a tenth of a lattice period. In the left plot, each pixel represents the phase measured at the location of a given tweezer. A histogram of these phases is provided in the right plot.

direction.[71] To control the roll (rotation about y) of the axial lattice, each lattice arm passes through an independently controlled optical flat (the same ones used to control the relative phase of the beams in Sec. 3.3.4.2). By controlling the angle of these optical flats, the lattice beams can be independently displaced in the Fourier plane.

To flatten the lattice fringes relative to the tweezers we parametrically heat atoms trapped in the axial lattice with the tweezers. Specifically, in the case that a tweezer is well-centered on a lattice fringe, modulating the power in the lattice changes the overall trap depth experienced by an atom, but not the position of the trap center (Fig. 3.10b). This results in parametric heating with a resonance at $\omega_m = 2\omega_0$, where ω_0 is the trap frequency in the absence of modulation [287, 154]. In this case, if the modulation frequency is equal to the trap frequency, the atom is not strongly heated. If, on the other hand, the lattice and tweezer are misaligned, modulating the power results in a shaking of the trap center, which can result in heating at $\omega_m = \omega_0$ (Fig. 3.10a) [287, 154]. For an appropriately chosen modulation amplitude and duration at $\omega_m = \omega_0$, we can use the observed probability of loss in a given tweezer as an indication of the relative alignment between that tweezer and the nearest lattice fringe. By scanning the phase of the lattice and fitting the phase of the resultant heating signal at each tweezer, we extract the spatial phase of the lattice at the location of the tweezers (Fig. 3.10c). The above signal allows for alignment of the lattice to the tweezers at the $a_{\mathrm{ax}}/10$ level, where $a_{\mathrm{ax}} = 1.32$ μm is the axial lattice spacing.

Once optimized, the flattening of the axial lattice has not been observed to change over multiple months. However, as mentioned in Sec. 3.3.4.2, the spatial phase of the axial lattice drifts on the several-hour timescale. In practice the loading, cooling, imaging, and tunneling presented in this work are relatively insensitive to drift of the axial lattice phase. As a result, this phase only needs to be adjusted when a node is close to the focal plane of the tweezers, where the transfer of atoms between the tweezers and the lattice is bistable between two layers of the axial lattice, and can result in inconsistent loading and cooling.

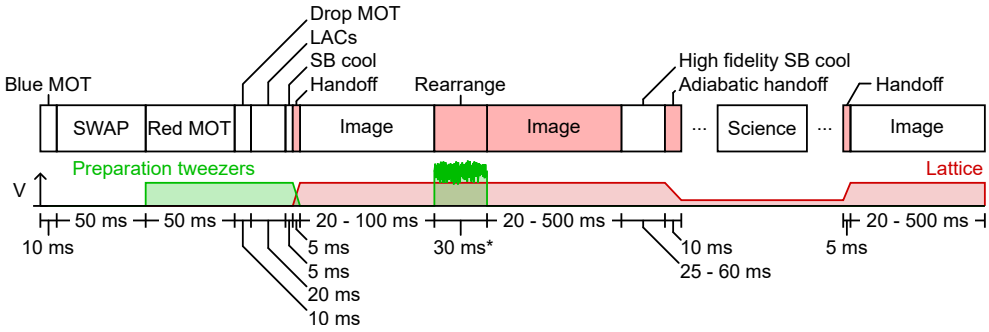

Figure 3.11: Typical experimental sequence. The steps marked in red are optional, depending on the experiment being performed. See the text for a description of the listed operations. Note that certain steps (particularly imaging) can vary dramatically in duration depending on the specific experiment performed. In the lower row, we show the relative depths (V, not to scale) of the lattice and preparation tweezer potentials during a typical experiment involving clock measurements or tunneling in the lattice. However, many other variations exist and are not depicted, including variations involving other tweezer potentials. This diagram is intended only to provide a general feeling for the steps required in a typical experiment, and does not serve as a source of accurate timing information for any given experiment. * The rearrangement step currently includes an additional technical delay of 110 ms, which we intend to address in future upgrades.

3.4 Standard operating procedures

Given the above tools, a typical experimental sequence involves the following steps (see Fig. 3.11): First, atoms are captured from the atom beam (Sec. 3.1.2) into a blue MOT, and then transferred into a SWAP MOT, followed by a narrow line red MOT (Sec. 3.1.4). The red MOT is overlapped with the preparation tweezers such that some of the atoms are trapped in the tweezers, and then the MOT is extinguished.[72] We typically wait for the MOT to dissipate ("dropping" the MOT) before proceeding. Finally, light assisted collisions (LACs) are used to ensure that each preparation tweezer contains exactly 0 or 1 atoms (Sec. 4.2). Depending on the parameters of the light assisted collisions step, between 45% and 75% of the preparation tweezers contain an atom at this stage.

The tweezer-trapped atoms are then sideband cooled. However, this cooling does not yet have to be particularly high fidelity, since the aim is simply to shrink the atomic wavepacket to be substantially smaller than a single site in the lattice.[73] After cooling, the tweezer-trapped atoms are transfered into the lattice via a sequence of ramps in trap depth, each 10 ms long.[74] The axial lattice is ramped up first, followed by the 2D lattice, and finally the preparation tweezers are ramped down. This handoff can be performed with $0.0^{+3}_{-0}\%$ atom loss. Once the tweezers are ramped low we fully shutter the tweezer beam, and then ramp the tweezers back up to full power. This allows us to avoid thermal lensing effects that can cause pointing fluctuations in the tweezers, while also avoiding unwanted light shifts introduced by the tweezers. Once the atoms are in the optical lattice, we switch from a magic magnetic field condition for the 515 nm tweezers [241] to a magic field condition for the 813 nm optical lattices [240] and perform resolved sideband cooling. The above preparation sequence typically takes about 200 ms.

At this stage, we take an initial image of the atoms in the lattice (see Sec. 4.4). In this image, the atoms can only occupy the subset of lattice sites which were overlapped with the preparation tweezers. Given

[71] Note that one must avoid scatterers in an image plane of the atoms, since this could lead to disorder in the potential projected onto the atoms. For this reason, we are very careful to protect the image plane mirror from dust.

[72] The tweezers provide a convenient means of isolating a single 2D layer of the 3D lattice. We have, however, performed proof of principle experiments loading directly into lattice, which enables higher atom numbers. However, some upgrades are required to isolate a single layer of the 3D lattice when loading the lattice directly from the MOTs.

[73] The procedure when working with the 813 nm tweezers as opposed to the lattice is similar, although we will comment on some slight differences below.

[74] The duration of these ramps is likely longer than necessary, especially given how quickly this transfer can happen during rearrangement (Sec. 4.5). However, we opt to be conservative because the handoff step doesn't result in a significant increase in the duration of the experiment, and because the atoms are not necessarily well-cooled at this stage.

that the preparation tweezers are spaced by more than a lattice spacing, good imaging performance can be achieved with a relatively short imaging duration of ~ 75 ms. If rearrangement is being used, we compute the required moves based on this image, and program the RF control system in ~ 5 ms,[75] execute the moves in ~ 30 ms (Sec. 4.5), and take an additional image to verify that the rearrangement has been performed correctly. If the target pattern of atoms is relatively sparse, the same ~ 75 ms image can be used. If the target pattern includes atoms at nearest-neighbor spacing (meaning atoms occupying adjacent lattice sites), then image durations of between $200 - 500$ ms are used. In all experiments, we then perform high fidelity resolved sideband cooling (typically using a pulsed cooling scheme as described in Sec. 4.3.1) to bring the atoms to their 3D motional ground states in the lattice (Sec. 4.3.3). The result of the above sequence is the preparation of a (nearly, see Ch. 8) pure state corresponding to a known pattern of atoms, with each atom being in the electronic ground state, and in the motional ground state of the lattice site that they occupy. This is the starting point of almost all experiments described in this work.

The atoms then undergo evolution of some sort, which can involve manipulations of their electronic state in the space spanned by the ground and clock states (Ch. 5), as well as a Rydberg state (Ch. 6), or coherent manipulations of the positions of the atoms in the lattice (Chs. 7 and 8). After this evolution, the positions and states of the atoms are read out, with an optional additional round of sideband cooling depending on if any heating was expected to occur between the previous round of cooling and readout. Depending on the experiment, readout can either be a single image or a pair of images, and possibly involves additional control pulses such that the measurement is sensitive to the electronic state of the atoms in addition to their position in the lattice (Sec. 5.2). As before, the duration of these final images varies between $75 - 500$ ms depending on the expected atom density. Ignoring the fact that the evolution step can have durations between microseconds and hundreds of seconds, the preparation and readout phases of the experiment typically vary between $250 - 1300$ ms.

For experiments that use only 813 nm tweezer arrays and not the lattice the experimental sequence is similar, with the main difference being that the images and final sideband cooling are instead performed in the 515 nm tweezers, resulting in worse performance (Secs. 4.3.3 and 4.4). However, we find the handoff

[75] There is an additional 30 ms delay when reading the image off of the camera. There is also currently an additional 50 ms delay before the moves are executed due to a computer issue, but we expect this problem to be readily solvable.

performance between tweezers to be similar to the performance between tweezers and the lattice.

Chapter 4

Preparing and detecting single atoms

"I'll get everything else working while you complete this first simple test, which involves deadly lasers, and how test subjects react when locked in a room with deadly lasers."
— GLaDOS

Portions of this chapter have appeared in:

M. A. Norcia, A. W. Young, and A. M. Kaufman. Microscopic control and detection of ultracold strontium in optical-tweezer arrays. **Phys. Rev. X**, 8(4):041054, Dec. 2018

A. W. Young, W. J. Eckner, W. R. Milner, D. Kedar, M. A. Norcia, E. Oelker, N. Schine, J. Ye, and A. M. Kaufman. Half-minute-scale atomic coherence and high relative stability in a tweezer clock. **Nature**, 588(7838):408–413, Dec. 2020

A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman. An atomic boson sampler. **In review**, 2023

This chapter is about state preparation (and by extension, detection). Given the tools described in Chs. 2 and 3, how do we control all degrees of freedom of an atom plucked from a thermal gas? For bosonic Sr atoms, the internal degrees of freedom are well-controlled since the electronic and nuclear ground states are both unique and separated in energy from any excited states by far more than the available thermal energy at room temperature (see Sec. 2.1.1). However, the motional states of the atoms, namely their positions and momenta, are not necessarily controlled. To tame the motion of the atoms, we place them in arrays of tightly confining optical traps (either tweezer arrays or optical lattices, as described in Sec. 3.3). This discretizes the available motional states into states labelled by the index of the specific trap an atom occupies, and the atom's motional state within that trap.[1] In Sec. 4.1, we will describe how our traps are engineered to affect the motion of the atoms without affecting relevant features of their internal degrees of freedom. In Sec. 4.2, we will describe how single atoms are loaded into these traps. In Sec. 4.3, we will describe how the motional

[1] We will discuss more sophisticated manipulations of the atomic motion in Chs. 7 and 8.

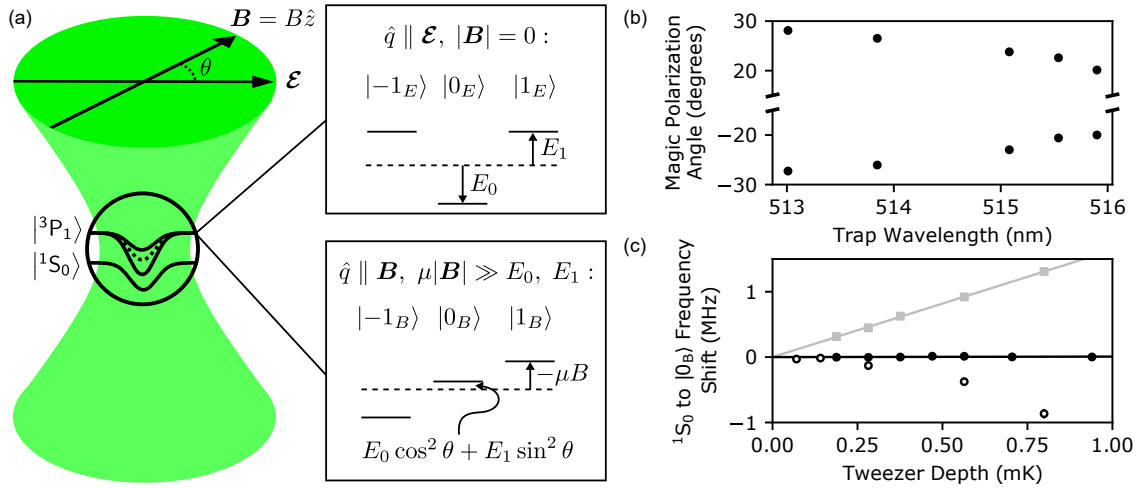

Figure 4.1: Magic angle trapping conditions. (a) By setting the angle θ between an applied magnetic field $\boldsymbol{B} = B\hat{z}$ and the polarization of the trapping light $\boldsymbol{\mathcal{E}}$ (in this case a tweezer), we can achieve equal trap depths between the ground state ($^1\mathrm{S}_0$) and an optically excited state ($^3\mathrm{P}_1$). In the absence of a magnetic field, the tweezer light causes energy shifts to the different spin projections of $^3\mathrm{P}_1$ (where $\hat{q}$ is the quantization axis along which we have defined our states). These shifts are of opposite sign relative to the shift of the ground state (dashed lines denote the shifted ground state offset by the unperturbed transition frequency). At the "magic" polarization angle θ, the shifts cancel for an eigenstate $|0_B\rangle$, leading to state-insensitive trapping. (b) Magic polarization angle as a function of tweezer wavelength, showing that a magic condition exists over a range of wavelengths. (c) Shifts of the $^1\mathrm{S}_0$ to $|0_B\rangle$ transition as a function of tweezer depth for a large (22 G) field oriented at $\theta = 0$ (gray squares), a large (24 G) field oriented at the magic angle (black points), and a small (7 G) field oriented at the magic angle (empty circles). All data in (c) is at a tweezer wavelength of 515.13 nm. Grey and black lines are linear fits.

state of an atom within a given trap can be controlled optically. In Sec. 4.4, we will describe how one can collect spontaneously emitted photons from optically cooled atoms to image the atoms with high fidelity and low loss, and thus identify which trap they occupy. Finally, in Sec. 4.5, we will describe how these images can be combined with rearrangement using mobile optical tweezers to control the specific locations of the prepared atoms. Taken collectively, the above tools allow us to control all degrees of freedom of the atoms. Later, in Sec. 8.3, we will see what this control means in a deeper sense, and conduct experiments that unambiguously demonstrate that, with high probability, the universe is unable to distinguish between different atoms prepared using the prescription in this chapter.

4.1 Magic angle trapping

In Sec. 2.2.3, we discussed how the appropriate choice of laser wavelength can allow one to engineer optical traps that decouple a subspace spanned by two electronic states of an atom from the motion of the atom within the trap. However, one can additionally tune the properties of the trap using the polarization (and shape [353]) of the trap. This additional tunability can be used to achieve magic-like conditions at wavelengths where convenient laser sources exist, or to decouple a larger electronic subspace from the atomic motion.

In our work, we use linearly polarized optical traps, and control the angle of the polarization relative to a quantization axis set by a magnetic field [357, 241]. This approach is particularly robust because one can take advantage of high quality polarizing optics, and, with the appropriate alignment of the polarization relative to the optical system used to project the traps, can often avoid uncontrolled modifications to the laser polarization due to birefringence. Similar tunability can be achieved using elliptically polarized optical traps [73], but this approach requires more careful consideration of how the various optics in the trapping system affect the polarization of the trapping field.

Consider the situation where we are interested in attaining a magic condition where the polarizability of the $^1S_0$ and $^3P_1$ states are equal. We will apply a magnetic field B in the z direction,[2] and a linearly polarized oscillating electric field with field strength $\mathcal{E}$ and with a polarization vector (see Eq. 2.9) of:

$$\hat{\varepsilon} = \cos(\theta)\hat{z} + \sin(\theta)\hat{x}, \tag{4.1}$$

where $\hat{x}$ and $\hat{z}$ are unit vectors along the x and z axes respectively, and θ is the angle of the electric field from the z axis. The resulting Hamiltonian for the atom is:

$$H = H_a + H_i + H_B, \tag{4.2}$$

where H_a is the Hamiltonian for the bare atom (e.g. Eq. 2.12 for a two level atom), H_i describes the effect of the electric field on the atom (Eq. 2.11), and H_B the effect of the magnetic field on the atom (Eq. 2.1). When the coupling terms in H_i and H_B are similar in magnitude, we perform exact diagonalization of Eq. 4.2

[2] With z as shown in Fig. 4.1, and not along the optical axis of the tweezers as in Ch. 3.

numerically [241]. However, in our experiment, B is typically large enough ($\sim$ 20 G, compared to trap depths of $\sim$ 0.1 mK) that the eigenstates in the $^3\mathrm{P}_1$ manifold remain approximate eigenstates of the J_z operator with $m_j = -1, 0, 1$. Therefore, we can proceed by computing the AC Stark shifts on these states.

Note that $\hat{\varepsilon}$ is real, and so the contribution from the vector polarizability vanishes (see Eq. 2.24). The $^1\mathrm{S}_0$ state has angular momentum $J = 0$, and so the tensor polarizability also vanishes, resulting in an AC Stark shift of:

$$\Delta E(^1\mathrm{S}_0) = -\alpha^{(0)}(^1\mathrm{S}_0)\frac{\mathcal{E}^2}{4}, \tag{4.3}$$

where the scalar polarizability $\alpha^{(0)}$ is as defined in Eq. 2.21.

The tensor contribution to the AC Stark shift on the $^3\mathrm{P}_1$ states does not vanish. Plugging Eq. 4.1 into Eq. 2.24 results in a total shift of:

$$\Delta E(^3\mathrm{P}_1, m_j) = -\left(\alpha^{(0)}(^3\mathrm{P}_1) + \alpha^{(2)}(^3\mathrm{P}_1)\left(\frac{3\cos^2(\theta)-1}{2}\right)\left(3m_j^2 - 2\right)\right)\frac{\mathcal{E}^2}{4}, \tag{4.4}$$

where the tensor polarizability $\alpha^{(2)}$ is as defined in Eq. 2.23. We find that the effective polarizabilities of the $^3\mathrm{P}_1$ states can be tuned with the value of θ. For example, for the magnetically insensitive $m_j = 0$ state, one can tune the polarizability between $\alpha^{(0)}(^3\mathrm{P}_1) + \alpha^{(2)}(^3\mathrm{P}_1)$ and $\alpha^{(0)}(^3\mathrm{P}_1) - 2\alpha^{(2)}(^3\mathrm{P}_1)$. As long as $\alpha^{(0)}(^1\mathrm{S}_0)$ falls within the above range for a given trapping wavelength, one can tune to a magic condition were $\Delta E(^3\mathrm{P}_1, m_j = 0) = \Delta E(^1\mathrm{S}_0)$, independent of the value of $\mathcal{E}$.

In Fig. 4.1b, we perform spectroscopy of the $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ transition, and, for different wavelengths of the tweezer potential, identify the angle at which this transition can be driven at its free-space transition frequency. Notice that a magic condition can be reached for a range of wavelengths near 515 nm. In Fig. 4.1c, we show how operation at the magic angle (for a trapping wavelength of 515.13 nm) leads to insensitivity of the transition frequency to trap depth.

We use the above approach to achieve magic trapping conditions for the $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ transition both for a trapping wavelength of 813 nm, and for wavelengths near 515 nm. The former case corresponds to the magic wavelength for the $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_0$ clock transition, resulting in simultaneous magic conditions for the $^1\mathrm{S}_0$, $^3\mathrm{P}_0$, and (one of the) $^3\mathrm{P}_1$ states.

In the 3D optical lattice (see Sec. 3.3), an additional complication is that the polarization of the different lattice potentials must be aligned. We achieve this by setting the fields to the magic angle with respect to the 2D lattice in a similar fashion to the above measurements in tweezers, and then vary the polarization of the axial lattice to minimize the shift of the $^1S_0 \leftrightarrow {}^3P_1$ transition as a function of axial lattice depth.

4.2 Preparing single atoms

To load individual atoms into the optical traps described in the previous section, we overlap the traps with the narrow-line MOT described in Sec. 3.1.4. The traps are typically in the range of $100 - 1000$ μK deep, far exceeding the thermal energy of the atoms in the MOT of ~ 5 μK. Although the MOT conditions do not correspond to the magic conditions in Sec. 4.1, residual Doppler or sideband cooling of the atoms in the presence of the AC Stark shifts imposed by the traps is typically still sufficient to reduce the atom temperature to well below the trap depth. The result is a Poisson distribution of atoms in each trap, with a mean of $N \sim 1 - 2$ atoms per trap for typical experimental parameters.

To prepare exactly one atom in each trap, we take advantage of light assisted collisions [292]. Schematically, the mechanism for these collisions is the presence of a deep molecular potential between electronically excited atoms.[3] When one drives one or more atoms to an electronically excited state with a laser, pairs of atoms roll down this molecular potential, gaining a significant amount of kinetic energy. Eventually, the atoms decay back to the electronic ground state, effectively turning off the molecular potential. The kinetic energy gained by the atoms during this process typically far exceeds the depth of the traps,[4] and so atoms exit the trap in a pairwise fashion [292]. The result is that traps initially filled with an even number of atoms are left with no atoms, whereas traps initially filled with an odd number of atoms are left with a single trapped atom. Given the above behavior, the use of light assisted collisions to map even and odd atom numbers to 0 and 1 atoms is often referred to as "parity projection." For large initial numbers of atoms, 50% of the traps are occupied with exactly one atom after parity projection, and the rest with zero. For the atom

[3] For example, a dipole-dipole interaction between atoms in superpositions of S and P states, with a characteristic scaling of r^{-3}, where r is the interatomic separation.

[4] Often, these collisions are excited by a red-detuned optical beam, as one might use for optical cooling. This excites the atoms to attractive molecular states with a binding energy that is almost always far in excess of reasonable trap depths.

numbers we typically load into each trap, the fraction of sites containing one atom is between 45% and 50%.

Note that by carefully controlling the collision channel one excites, it is possible to arrange for collisions that are comparable in energy to the trap depth, and thus typically only remove a single atom at a time [134, 189, 111, 48].[5] Using these techniques, single atom loading efficiencies as high as 96(1.4)% have been achieved in a single tweezer (and 92.78(8)% in an array of 100 tweezers) [155]. In Sr, we have used similar techniques to achieve loading efficiencies as high as 75% in an array of 384 tweezers [367], but we leave a detailed discussion of these results for future works.

4.3 Controlling atomic motion

We are interested in controlling the motion of a trapped atom with a laser field. So far, we have considered how far off-resonant laser light can generate conservative trapping potentials with different properties, but in Eq. 2.11 we neglected the fact that $\boldsymbol{\mathcal{E}}$ can vary over the spatial extent of the atom. Including the spatially varying phase of the electric field in Eq. 2.13 leads to [356, 185]:[6]

$$H_i = \hbar\frac{\Omega}{2}(\sigma_+ e^{i(\boldsymbol{k}\cdot\boldsymbol{r}-\Delta t)} + \sigma_- e^{-i(\boldsymbol{k}\cdot\boldsymbol{r}-\Delta t)}), \tag{4.5}$$

where $\boldsymbol{r}$ is the position operator for the atom, $\sigma_+ \coloneqq |e\rangle\langle g|$ and $\sigma_- \coloneqq |g\rangle\langle e|$ are the spin raising and lowering operators, t is the time, Δ is the detuning defined below Eq. 2.13, and Ω is the Rabi frequency as defined in Eq. 2.14. One can interpret Eq. 4.5 as describing the process where absorbing or emitting a photon in a given direction results in a corresponding momentum kick in that direction.

Assuming that the trapping potential is harmonic, the Hamiltonian describing the motion of the atom is:[7]

$$H_m = \hbar\omega_m(a^\dagger a + \frac{1}{2}), \tag{4.6}$$

where ω_m is the angular frequency associated with the harmonic oscillator, and $a^\dagger$ and a are the harmonic

[5] The central trick in these "enhanced loading" techniques is to use a blue-detuned beam which excites the atoms to repulsive molecular states. In this case, one can precisely control how far up the potential the atoms are excited to relative to the free space energy, and thus the amount of kinetic energy gained in the collision process.

[6] Keep in mind that this is in the rotating frame of the laser, and after taking the rotating wave approximation.

[7] For now, we will consider the simplified scenario where the atomic motion is constrained to the propagation axis of the control field, x, but this discussion is readily extended to the three-dimensional case [185]. In our experiments, we typically drive three orthogonal axes of a 3D trapping potential with three separate control beams, and so one can just duplicate the equations in this section for each axis.

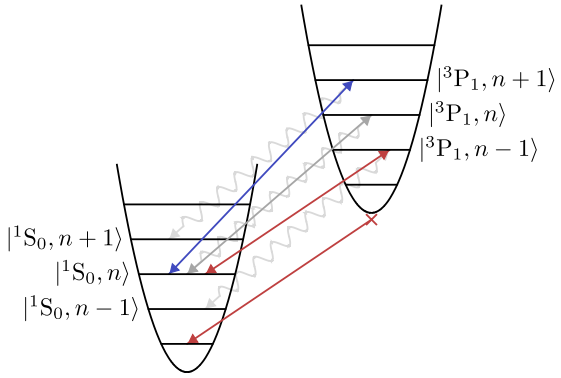

Figure 4.2: Electronic transitions in a harmonic trap. The states of a harmonically trapped atom can be labelled by the internal state of the atom (e.g. the states $^1\mathrm{S}_0$ and $^3\mathrm{P}_1$ in this example), and the quantum number n describing the harmonic motion of the atom. In the Lamb-Dicke regime, spontaneous emission of a photon is unlikely to change the motional state of the atom (grey squiggly lines). However, if motional-state-changing processes (red and blue lines) are well-resolved from processes which do not change the motional state (grey line), then one can resonantly drive these processes with a laser. Driving the red process will lead to an overall reduction of the motional state of the atom, until the atom is in the state $|^1\mathrm{S}_0, n = 0\rangle$. There is no state that differs in energy from $|^1\mathrm{S}_0, n = 0\rangle$ by the energy of a photon in the field of the cooling laser (red X), and so at this point the atom goes dark to the cooling light.

raising and lowering operators. In this case, $x = x_0(a^\dagger + a)$,[8] where $x_0 \coloneqq \sqrt{\hbar/(2m\omega_m)}$ is the harmonic oscillator length (with m being the atomic mass). Given the above definitions, we can rewrite Eq. 4.5 as:

$$H_i = \hbar\frac{\Omega}{2}(\sigma_+ e^{i\eta(a^\dagger+a)-i\Delta t} + \text{H.C.}) \tag{4.7}$$

where we have defined the Lamb-Dicke parameter $\eta \coloneqq kx_0$.

The states of the atom can be written as $|\beta, n\rangle \coloneqq |\beta\rangle \otimes |n\rangle$, where $\beta \in \{e, g\}$ refers to the electronic state of the atom, and n is the quantum number associated with the harmonic motion of the atom.[9] One can compute the coupling induced by H_i between states $|g, n\rangle$ and $|e, n+s\rangle$ exactly [354, 185]:

$$\begin{aligned}\Omega_{n,n+s} &\coloneqq \Omega|\langle g, n \,|\, e^{i\eta(a^\dagger+a)} \,|\, e, n+s\rangle| \\ &= \Omega e^{-\eta^2/2}\eta^{|s|}\sqrt{\frac{n_<!}{n_>!}}L_{n_<}^{(|s|)}(\eta^2),\end{aligned} \tag{4.8}$$

where $L_n^{(\alpha)}(\cdot)$ is the generalized Laguerre polynomial, $n_> = n + s$ and $n_< = n$ for $s \geq 0$, and $n_> = n$ and $n_< = n + s$ for $s < 0$. Note that the states $|g, n\rangle$ and $|e, n+s\rangle$ differ in energy by $\hbar(\omega_{eg} + s\omega_m)$, and so for $\Omega \ll \omega_m$ typically only a single transition is driven at a time for a given detuning of the control field.

When the fluctuations in the atom position are small in comparison to the wavelength of the electromagnetic field, namely when $\eta\sqrt{\langle(a^\dagger + a)^2\rangle} \ll 1$, one is in so-called the Lamb-Dicke regime. In this regime, we can simplify the description of the dynamics by Taylor expanding Eq. 4.7 about $\eta = 0$, resulting in:

$$H_i(t) = \hbar\frac{\Omega}{2}\sigma_+ e^{-i\Delta t}(1 + i\eta a^\dagger e^{i\omega_m t} + i\eta a e^{-i\omega_m t} + \dots) + \text{H.C.} \tag{4.9}$$

Here, Δ is the angular frequency associated with the detuning from resonance with the electronic transition (as in Sec. 2.2.1), and we have kept the explicit time dependence of the spin- and motion-raising and lowering operators in the rotating frame of the atom. It becomes clear that we can address a specific motion-changing transition by picking the appropriate Δ, and perform a second round of the rotating wave approximation to remove all other off-resonant processes. The processes which are resonant at $\Delta = \pm\omega_m$ are known as the first blue and red sidebands respectively, and drive the transitions $|g, n\rangle \leftrightarrow |e, n \pm 1\rangle$. Similarly, one can drive the $|s|^{\text{th}}$-order sidebands, corresponding to higher order terms in the expansion with respect to η, by picking detunings of $\pm|s|\omega_m$.

[8] We have taken $\boldsymbol{r} \to x$, see previous footnote.

[9] These states generically form a complete basis, and are precisely the eigenstates of the atoms for magic trapping conditions.

4.3.1 Sideband cooling

When combined with dissipation due to the finite decay rate γ of the $|e\rangle$ state, driving the red sideband(s) provides a powerful mechanism for cooling. Specifically, in the resolved sideband regime where $\gamma, \Omega \ll \omega_m$, the dynamics associated with relaxation and the electromagnetic drive are slow in comparison to the motional dynamics, and so one can spectroscopically resolve the separate resonances appearing in Eq. 4.9.[10] In this case, starting in the state $|g, n\rangle$, one can coherently drive the atom to $|e, n-1\rangle$ and allow it to spontaneously decay to $|g, n-1\rangle$. This procedure can be repeated until $n = 0$, at which point there are no available states with lower n, and so the atom becomes dark to the drive light (see Fig. 4.2).

The evolution of the atom during sideband cooling is described by the master equation [226]:

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[H_a + H_m + H_i, \rho] + \mathcal{L}(\rho), \tag{4.10}$$

where H_a, H_m, and H_i are as defined in Eqs. 2.13, 4.6, and 4.7 respectively, $[\cdot\,,\cdot]$ is the commutator, ρ is the density matrix describing the state of the atom, and $\mathcal{L}$ is a Liouvillian describing the damping of the atom. To define $\mathcal{L}$, we must consider how spontaneous emission of a photon affects the motion of the atom. When the atom decays, it emits a photon into an emission pattern described by the angular distribution $f(\theta, \phi)$. In our works, we typically use transitions between states with $m_j = 0$, where [318]:

$$f(\theta, \phi) = \frac{3}{8\pi}\sin^2(\theta). \tag{4.11}$$

Emission of a photon in a direction given by the unit vector $\hat{k}_{\theta\phi}$, where θ and ϕ are the polar and azimuthal angles of the vector, is described by the jump operator [185]:

$$L_{\theta\phi} = \sigma_- e^{ik\hat{k}_{\theta\phi}\cdot\boldsymbol{r}}, \tag{4.12}$$

where k is the wavenumber associated with the $|e\rangle \rightarrow |g\rangle$ transition. The Liouvillian $\mathcal{L}$ is given by the combination of all such decay processes, weighted by f [185, 208]:

$$\mathcal{L}(\rho) = \gamma\left(\int d\Omega f(\theta, \phi) L_{\theta\phi}\rho L_{\theta\phi}^\dagger - \frac{1}{2}\{\sigma_+\sigma_-, \rho\}\right) \tag{4.13}$$

[10] Note that there are other cooling mechanisms that take advantage of transitions with small γ to reach very low temperatures. One particularly useful example is a kind of Sisyphus cooling [73, 74] which we use in some of our experiments [240]. We opt not to discuss this cooling mechanism in more detail here, since most of our experiments use sideband cooling to reach lower temperatures at the cost of lower cooling rates. For a nice description of Sisyphus cooling in alkaline earth atoms, see [208].

where $\{\cdot,\cdot\}$ is the anticommutator, and $\int d\Omega$ denotes an integral over all solid angles. Note that for the term involving the anticommutator, we have taken advantage of the fact that the ordering of the operators is such that the angle-dependent terms in the integral cancel. For approximate calculations, it is often enough to reduce the integration over f to a summation over three appropriately weighted decay processes corresponding to photon recoil along the three orthogonal axes of the trap.

With the above evolution in mind, the primary limitation to the temperature reached by sideband cooling is the competition between cooling via driving the red sideband and heating due to off-resonant driving of other processes.[11] Given these mechanisms, one can arrive at an approximate, analytical solution for the ultimate performance of resolved sideband cooling of $\langle n\rangle \simeq (\gamma/\omega_m)^2$ [185, 101]. The performance of this cooling can be modified by a constant factor for different geometries and emission patterns. For example, in [208] they numerically arrive at $\langle n\rangle \simeq 0.17(\gamma/\omega_m)^2$.

Practically, one does not always saturate the above cooling limit due to geometric constraints, and additional sources of technical heating. Noise on the drive laser can result in a spectral feature that is wider than the expectation given γ, leading to an enhanced rate at which the various off-resonant processes are driven. As described in Appendix D, we go to great pains to ensure that the linewidth of the drive laser is comparable to the 7.5 kHz wide $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ transition we use for sideband cooling. Scattering of the trapping light can also result in recoil heating, but typically these rates are very low in comparison to the cooling rates. Additionally, modulations of the trapping potential (either due to intensity or pointing fluctuations) can lead to parametric heating. Again, these processes are normally slow in comparison to the cooling rate, and are only relevant in specific situations. We will discuss these parametric heating mechanisms in more detail in Sec. 5.3 in the context of experiments that are performed on very long timescales (of tens to hundreds of seconds). Parametric heating is also relevant in rearrangement, as described in Sec. 4.5, where tweezer-trapped atoms traversing a lattice potential experience a significant modulation of the trap depth.

In our experiments, we typically utilized a pulsed cooling sequence, where three different axes are cooled sequentially (see Fig. 4.3). This allows us to tune the relative cooling performance along different

[11] Namely, off resonant driving of the carrier followed by decay on the first red sideband, and driving of the first blue sideband followed by decay on the carrier.

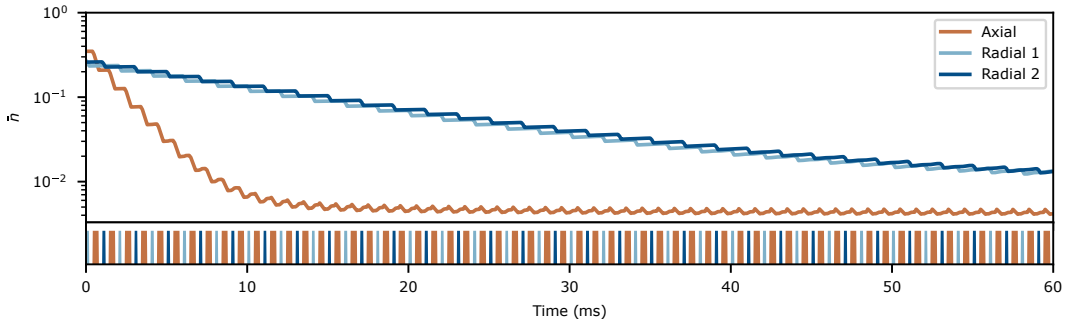

Figure 4.3: Simulated sideband cooling performance. To optically cool lattice-trapped atoms with high fidelity, we typically use a pulsed cooling sequence involving 0.4 ms axial cooling pulses, and 0.2 ms radial cooling pulses (timing diagram pictured in lower panel). We compute the expected average thermal occupation $\bar{n}$ as a function of time in each of three nearly-orthogonal axes of a given site in the lattice via a master equation calculation, yielding reasonable agreement with measured values in the experiment. Note that in the depicted cooling sequence, we optimize for high-fidelity cooling of the axial direction at the cost of slightly worse cooling in the radial directions.

trap axes[12] without constraining the amount of optical power used to cool each axis. The pulsed procedure additionally minimizes the effect of coherent processes where a photon is absorbed from one cooling beam, and undergoes stimulated emission into another beam.

A simulation using Eq. 4.10 of a typical pulsed cooling sequence, specifically the one used for cooling in a 3D optical lattice (with typical trap parameters listed in Tab. 3.3), as is used in the experiments in Ch. 8, is shown in Fig. 4.3. Here, the cooling is composed of 120 pulses which alternate between cooling pulses with a duration of 200 μs on two nearly-orthogonal "radial" axes in the plane of the 2D lattice separated by 400 μs cooling pulses on the "axial" out of plane axis. Each cooling pulse is separated by a delay of 200 μs in order to allow the atoms to completely relax to the electronic ground state between pulses, leading to an overall cooling sequence that is 60 ms in duration. The described cooling sequence is expected to result in an axial motional ground state occupation of 99.58 %, and a combined ground state occupation of 97.37 % in the radial directions, leading to a 3D motional ground state occupation of 96.97 %. This sequence is biased towards high-fidelity cooling of the less strongly-confining axial direction at the cost of slightly worse radial cooling (by controlling the durations of the cooling pulses along different axes).[13] The ultimate observed cooling performance in our experiment is consistent with this simulation, as we will discuss in more detail in Secs. 4.3.3 and 8.3.

4.3.2 Sideband thermometry

There are many ways to characterize the final temperature of the atoms after cooling, including release and recapture techniques [323, 73], and the interference techniques described in Ch. 8. However, one particularly simple and effective technique is known as sideband thermometry, where one performs spectroscopy that can resolve the motional sidebands, and compares the heights of the first red and blue sidebands.

Consider the evolution of a thermal state under a drive field on the red or blue sidebands. The initial

[12] This capability is important, for example, when optimizing for high quality cooling only along the direction of a given control beam as in Sec. 5.2, or when some thermal excitations can be filtered out using postselection as in Sec. 8.3.

[13] In this specific case, we choose to make this trade-off because the axial cooling performance is more important than the radial for the experiments described in Ch. 8.

state of the atom is described by the density matrix [185]:[14]

$$\rho(t=0) = \sum_{n=0}^{\infty} \frac{\langle n\rangle^n}{(\langle n\rangle+1)^{n+1}} |g,n\rangle\langle g,n|. \tag{4.14}$$

Under the influence of a drive on the red sideband with duration t, the probability of observing the atom in the $|e\rangle$ state is:

$$P_e^{(r)}(t) \coloneqq \sum_{n=0}^{\infty} \langle e,n \,|\, \rho(t) \,|\, e,n\rangle = \sum_{n=1}^{\infty} \frac{\langle n\rangle^n}{(\langle n\rangle+1)^{n+1}} \sin^2(\Omega_{n-1,n}t), \tag{4.15}$$

and similarly for the blue sideband:

$$P_e^{(b)}(t) \coloneqq \sum_{n=0}^{\infty} \langle e,n \,|\, \rho(t) \,|\, e,n\rangle = \sum_{n=0}^{\infty} \frac{\langle n\rangle^n}{(\langle n\rangle+1)^{n+1}} \sin^2(\Omega_{n+1,n}t). \tag{4.16}$$

Shifting the summation index, and keeping in mind that $\Omega_{n,n+s} = \Omega_{n+s,n}$, we find that:

$$\frac{P_e^{(r)}}{P_e^{(b)}} = \frac{\langle n\rangle}{\langle n\rangle+1}, \tag{4.17}$$

and so one can measure the ratio $P_e^{(r)}/P_e^{(b)}$ to extract the quantity $\langle n\rangle$. This measurement of $\langle n\rangle$ is particularly robust, since Eq. 4.17 is independent of all other relevant parameters in the dynamics, including η, t, Ω, and ω_m. However, when characterizing cooling along multiple axes, one must be careful that the dynamics that occur during sideband thermometry do not perturb the temperature of the atoms.[15] In practice, we simply simulate the thermometry sequence via Eq. 4.10, and confirm that it does not significantly perturb the atomic motion [241].

In our experiments, the procedure for sideband thermometry involves preparing the cooled atoms, driving the $^1\mathrm{S}_0 \leftrightarrow {}^3\mathrm{P}_1$ transition with variable detuning Δ and $\Omega \sim 2\pi \times 10$ kHz for a duration of ~ 100 μs. This duration is chosen because it is long in comparison to the decay time of $^3\mathrm{P}_1$, allowing the populations to stabilize,[16] but short compared to the cooling duration, preventing the perturbations to the atomic temperature described above. A $\simeq 3$ μs blow-away pulse of light resonant with the $^1\mathrm{S}_0 \leftrightarrow {}^1\mathrm{P}_1$ transition is used to rapidly heat and eject the atoms in $^1\mathrm{S}_0$ from the traps before the $^3\mathrm{P}_1$ atoms decay back to $^1\mathrm{S}_0$ (see

[14] Note that this is often a somewhat crude approximation, since the distribution prepared by sideband cooling is not necessarily thermal.

[15] During the probe pulse, one can heat the probed axis on the blue sideband, and cool it on the red sideband, leading to a greater asymmetry in the resulting distribution, and thus artificially good inferred performance. This is particularly problematic when there are tradeoffs associated with cooling multiple axes.

[16] Although Eq. 4.17 is not dependent on extraneous experimental parameters, it can be a problem if these parameters fluctuate over multiple experimental trails such that the conditions under which the red and blue sidebands are measured differ.

Sec. 2.1.1.2). Typically, the traps are extinguished for the final 1 μs of the blow-away pulse before being turned back on, which aids in removing the heated $^1S_0$ atoms, but does not result in substantial loss of the $^3P_1$ atoms. The remaining $^3P_1$ atoms are allowed to decay back to $^1S_0$, before being imaged (see Sec. 4.4). By taking images before and after performing the above sequence, the number of surviving atoms can be converted into a probability of the atoms being in $^3P_1$. To minimize sensitivity to imperfect calibration of the magic angle, and fluctuations in trap frequency, we typically perform the above procedure at many values of the detuning (not just at the location of the first red and blue sidebands), and perform maximum likelihood estimation to fit the resulting line shape with three Lorentzians — one each for the carrier, and the red and blue sidebands. $\langle n \rangle$ is then inferred via Eq. 4.17 from the fitted heights of the sidebands. Note that one must be careful when inferring ω_m from the separation of the red and blue sidebands from the carrier. When Ω approaches ω_m, probing near a given sideband also results in a probe-induced AC Stark shift (see Sec. 2.2.1) of the entire spectroscopic feature including the carrier and sidebands. The result is that, when scanning across the carrier, the red and blue sidebands appear to shift towards the carrier, resulting in artificially low inferred values of ω_m.

4.3.3 Cooling results

We perform resolved sideband cooling and characterize the resulting cooling performance using sideband thermometry in a few different conditions, which are summarized as appropriate in the works listed in Sec. 1.4. Here, we will focus on two conditions that are particularly relevant to our later works:

(1) Where we cool the atoms in tightly confining 515 nm tweezers and a 1D 515 nm lattice oriented along the weakly confined axis of the tweezers (collectively referred to as the preparation potential), and then hand the atoms into weakly confining 813 nm tweezers (science potential)

(2) Where we cool the atoms directly in a tightly confining 3D optical lattice at 813 nm.

In both cases, the trapping potentials are tuned to a magic condition for the $^1S_0 \leftrightarrow {}^3P_1$ transition as described in Sec. 4.1.

In condition (1), which applies to some of the experiments described in Ch. 5, we are interested in

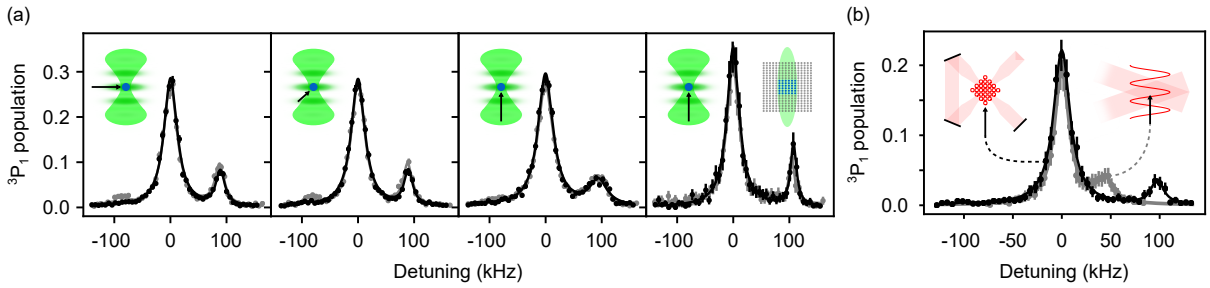

Figure 4.4: Sideband thermometry in tweezers and optical lattices. (a) Performing sideband spectroscopy before (black points) and after (grey points) adiabatically transferring the atoms to and back from the science tweezers (after performing sideband cooling in a preparation potential composed of tweezers and a 1D axial lattice), we measure an average phonon occupation of $\bar{n} = 0.07^{+0.14}_{-0.07}$, $0.06^{+0.08}_{-0.06}$, and 0.07 ± 0.06 ($\bar{n} = 0.25 \pm 0.12$, 0.31 ± 0.13, and 0.27 ± 0.10) before (after) the handoff in the axial, first, and second radial directions respectively. Cartoons in the top left of each frame indicate the orientation of the probe beam relative to the traps, showing probes in two orthogonal radial directions (left two subplots) and in the axial direction (right two subplots). The right-most spectra show that, in a reduced 6×6 region at the center of the array (denoted by the right-most cartoon), the axial cooling performance is vastly improved, with an average phonon occupation of $\bar{n} = 0.00^{+0.06}_{-0.00}$ ($0.06^{+0.10}_{-0.06}$) before (after) the handoff. This is due to the comparable extent of the lattice beams to the tweezer array (light-green contour in right-most cartoon shows the region over which the lattice intensity stays within 90% of its maximal value). (b) Measurements in the 3D lattice show similar performance over a larger region spanning more than 2000 lattice sites, with a thermal occupation of $\bar{n} = 0.000^{+33}_{-0}$ and 0.000^{+29}_{-0} in the axial (grey) and radial (black) directions respectively (the relevant orientations of the probe beams are once again indicated by the inset cartoons).

confirming that the atoms remain cold when handing them from the preparation potential to the science potential (as described in Sec. 3.4). We perform sideband thermometry in the preparation potential both immediately after sideband cooling, and after adiabatically passing the atoms to the science potential, holding for 25 ms, and passing them back. As shown in Fig. 4.4a, before the handoff we observe an average phonon occupation of $\bar{n} = 0.07^{+0.14}_{-0.07}$, $0.06^{+0.08}_{-0.06}$, and 0.07 ± 0.06 in the axial, first, and second radial directions respectively. After the handoff we observe an average phonon occupation of $\bar{n} = 0.25 \pm 0.12$, 0.31 ± 0.13, and 0.27 ± 0.10 (again in the axial, first, and second radial directions). Since we expect that heating occurs during both steps of the handoff, the mean of these two measurements serves as an estimate of the temperature of the atoms in the science potential.

Although the tweezers, and thus the radial trap frequencies, can be balanced across the entire array, there is substantial inhomogeneous broadening of the axial trap frequencies. This is due to the relatively small 25 μm waists of the lattice beams, which are comparable to the extent of the tweezer array (Fig. 4.4a).[17] In a smaller 6×6 region at the center of the array the axial cooling and handoff performance is vastly improved, with an average phonon occupation of $\bar{n} = 0.00^{+0.06}_{-0.00}$ ($\bar{n} = 0.06^{+0.10}_{-0.06}$) before (after) the handoff.

In condition (2), the tight confinement provided by the optical lattice allows for significantly higher quality cooling, as well as larger numbers of traps. Sideband spectroscopy after resolved sideband cooling in these traps indicates average phonon occupations of $\bar{n} = 0.000^{+33}_{-0}$ and 0.000^{+29}_{-0} in the axial and radial directions respectively, and thus a 3D ground state fraction of $100^{+0}_{-7}\%$ over a 26×26 μm region corresponding to more than 2000 lattice sites (Fig. 4.4b). This performance is consistent with the expectation of 96.97 % 3D motional ground state occupation from Sec. 4.3, albeit with somewhat large statistical uncertainty. At these very low temperatures, it can often be hard to achieve precise measurements, since as the red sideband approaches zero, it disappears into the spectral feature associated with the carrier transition. In Sec. 8.3, we will discuss a complementary measurement that provides tighter bounds on the cooling performance, and is also in agreement with the expectation based on simulations.

[17] This condition was chosen as a compromise between atom number and high quality state preparation and detection, given the optical power we had available at the time.

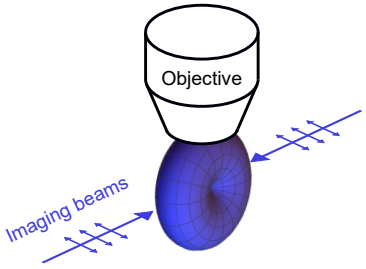

Figure 4.5: Layout of the imaging beams. For improved imaging performance, it is important that the polarization of the imaging beams (small blue arrows) is oriented such that the dipole emission pattern from the atoms (blue donut) is oriented appropriately relative to the microscope objective. Additionally, it is useful to align the propagation axis of the imaging beams (large blue arrows) to be along the most tightly confining axis of the optical traps, and to use two counterpropagating imaging beams with balanced intensities.

4.4 Imaging single atoms

We are interested in imaging atoms in the $^1S_0$ ground state by collecting photons that are scattered off of the atoms. We would like to be able to do this with high fidelity (in the sense that the presence or absence of an atom is correctly identified with high probability) and low loss (in the sense that atoms are not lost during the measurement). Additionally, we would like to identify the positions of the atoms with good spatial resolution, and we would like to acquire the images as quickly as possible.

To balance the above considerations, our imaging scheme consists of exciting the $^1S_0 \leftrightarrow {}^1P_1$ transition, and collecting the spontaneously emitted photons. These photons have a convenient wavelength of 461 nm, which is short enough to provide good spatial resolution (see Sec. 3.3), and energetic enough to provide good detection efficiencies using commercially available camera sensors.[18] However, as photons are scattered off of this transition, the atoms experience recoil heating. To prevent atom loss due to this heating, we simultaneously perform sideband cooling using the $^1S_0 \leftrightarrow {}^3P_1$ transition as described in Sec. 4.3. Due to the multiple decay paths from $^1P_1$ (see Sec. 2.1.1.1), and Raman scattering of the trapping light when occupying the $^3P_1$ state, the metastable $^3P_{0,2}$ states can be populated during this process, and so it is important to repump these states back to the ground state by additionally applying 707 nm and 679 nm repumping

[18] It is also helpful that these photons are more energetic that the photons in the trapping beams. If the opposite is true, uncontrolled fluorescence in the microscope objective as a result of the high power trapping beams can result in emitted photons that are close in energy to the photons emitted from the atoms, and thus contribute to a background signal that is challenging to filter out.

light (see Sec. 2.1.1.4). Although we have experimented with pulsing these various control fields during imaging [241], we have not found this to make a significant difference in imaging performance, and therefore typically apply all the control fields continuously and simultaneously.

The orientation of the 461 nm drive beam(s), which we will refer to as the imaging beam, is important for optimizing imaging performance (see Fig. 4.5). For scattering rates that are fast compared to the trap frequency, photons absorbed from the imaging beam result in momentum kicks that add coherently, resulting in faster heating of the atoms. As a result, it is helpful to orient the imaging beam along the most tightly confining axis of the trap, and to use two counter-propagating imaging beams so that these momentum kicks average to zero, even in the absence of the trap. To avoid interference between these beams leading to spatially varying scattering rates, and thus varying imaging performance, we used two counter-propagating imaging beams that are balanced in intensity, and detuned from each other by 20 MHz (resulting in a travelling interference pattern that varies on timescales comparable to the linewidth of the imaging transition). The two beams are additionally blue-detuned (on average) by 690 MHz from the free space resonance. This detuning can help with imaging performance in tightly confining traps [217]. Whereas resonantly exciting the atoms to the excited state can result in significant heating due to fluctuations in the applied dipole force (since the excited state is often anti-trapped, or at least experiences a very different trapping potential), weakly dressing the atoms with the imaging light results in a dressed state that remains trapped while still intermittently scattering photons.

To efficiently collect the spontaneously emitted 461 nm photons, one must consider the angular distribution of the emitted photons (see Eq. 4.11), and orient the polarization of the 461 nm drive field and the microscope objective appropriately. The optimal orientation in most of our experiments is where the polarization of the 461 nm beam is orthogonal to the axis of the microscope objective. For our microscope and imaging setup (see Sec. 3.3), the ultimate collection efficiency of a spontaneously emitted photon from an atom onto the camera is $\sim 9\%$, including imperfect transmission through the various imaging optics, and the quantum efficiency of the camera.

The collected photons are imaged onto an electron-multiplying charge-coupled device (EMCCD) camera (Andor iXon Ultra 897). The tradeoffs for different modes of operation of this camera are described

in detail in [9]. Here, we briefly summarize a few particularly important tradeoffs between signal-to-noise and image resolution that are relevant in our experiments. The main source of noise on an EMCCD camera is due to clock induced charges (CICs).[19] These extra charges are generated during the charge transfer process that moves the collected signal to the readout register, and are thus indistinguishable from real photoelectrons generated by the atoms. As this source of noise is a per-pixel effect, it can be advantageous to concentrate the collected signal on a single pixel to maximize signal-to-noise. Additionally, CICs can be minimized by increasing the vertical shift speeds used to transfer the collected signal to the readout register on the sensor. However, both these choices result in reduced spatial resolution. The former because one is limited by the pixel size, rather than the optical properties of the imaging system, and the latter because increased vertical shift speeds result in reduced charge transfer efficiency, and thus a smearing of the collected signal over multiple pixels. As a result, we opt to design our imaging system such that the waist of the PSF (see Sec. 3.3) corresponds to ~ 2.5 pixels, spreading the collected atomic fluorescence over multiple pixels to preserve the spatial resolution of the imaging system. This choice enables the super-resolution techniques described in Sec. 4.4.2. We also do not attempt to perform very fast readout of the images, using a moderate (but still fairly fast) vertical shift speed of 1.7 μs.[20]

4.4.1 Image analysis

For each image, we are interested in identifying which sites in the array contain an atom in the $^1S_0$ state (occupied sites), and which do not (empty sites). To minimize the required signal and thus imaging time, we take advantage of a few extra pieces of information. First, the array of optical traps results in a discrete set of sites that the atoms can occupy, and the location of these sites can be characterized through independent measurements (Fig. 4.6a). For each site, we multiply the raw images by a mask corresponding to the measured point spread function (PSF) of the imaging system at that location (Figs. 4.6b and 4.7b) and

[19] In a conventional charge coupled device (CCD) camera, one typically also suffers from read noise associated with digitizing the collected photoelectric signal, and thermally excited extra charges. The latter effect can be suppressed by cooling the sensor, and the former effect can effectively be eliminated using electron-multiplying gain, at the cost of a constant reduction in signal-to-noise of $\sqrt{2}$ from the optimal case. Below signals of a few thousand photons per pixel, read noise typically dominates, and so it is advantageous to use EM gain [9].

[20] Since our exposure times are typically tens to hundreds of milliseconds, this does not pose a significant limitation to the experiment.

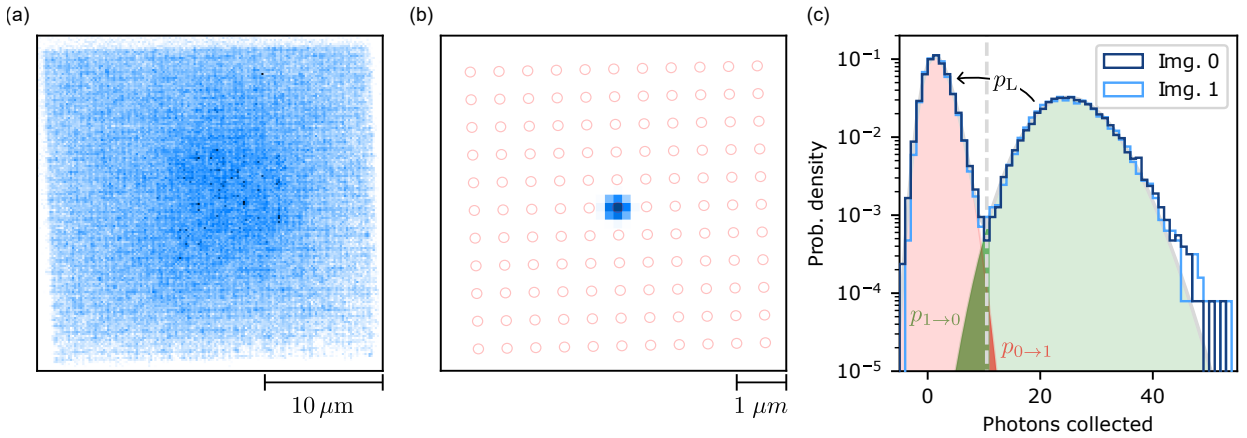

Figure 4.6: Image analysis for sparse arrays. (a) An averaged image of atoms loaded into ~ 2500 sites in the lattice is used to calibrate the positions of the lattice sites for subsequent analysis. Note that only a small subset of these sites are loaded on any individual run of the experiment. (b) An averaged image of a single implanted atom, surrounded by empty lattice sites denoted by the red circles. This serves as a direct measurement of the PSF of our imaging system. (c) A histogram of the binned counts for two consecutive images (img. 0 and 1) in typical imaging and loading conditions in the lattice. Specifically, in this data a 16×24 array of preparation tweezers is loaded with $\sim 50\%$ filling, and the atoms are implanted in a sublattice that is spaced by 3 lattice sites in one direction (the direction with 16 tweezers), and two lattice sites in the other (the direction with 24 tweezers). Finally, the atoms are imaged for 100 ms. The resulting distribution of photon counts on each site (considering only sites where a preparation tweezer was present) is clearly bimodal, with empty sites corresponding to the red peak, and occupied sites corresponding to the green peak. We estimate whether an atom is present or absent in a given experimental trial by evaluating if the measured counts fall above or below an appropriately chosen threshold (grey dashed line). The loss probability p_L can be estimated by looking for an excess of empty sites in image 1 relative to image 0, and takes a value of 0.015(15)% in the displayed data. The false negative probability $p_{1\to0}$ corresponds to the portion of the atom peak that falls below threshold (dark green region), and the false positive probability $p_{1\to0}$ to the portion of the empty peak that falls above threshold (dark red region). The shaded regions in this plot correspond to a fit to a skewed Gaussian distribution for the empty peak, and a Poisson distribution for the atom peak. Note that these fits are only to guide the eye, and that our procedure for estimating the various error probabilities is independent of any assumptions about the functional form of the measured distribution of photon counts.

sum the result to obtain the number of counts collected for a given site in a given image.[21] For experiments where the atoms are pinned in place (namely those appearing in Chs. 5 and 6), we ensure that the sites are far enough apart that the overlap between the PSF from different sites is negligible.

The histograms of the counts are bimodal, with one peak near zero counts associated with the empty sites, and one peak with higher counts associated with the occupied sites (see Fig. 4.6c). The former peak is approximately Gaussian, with scatter that is set by thermal noise and CICs on the EMCCD, as well as leakage light. The latter peak is approximately Poissonian, due to the discrete and probabilistic nature of the collected signal. We set a threshold between these peaks to binarize the signal where, for a given image, sites with counts that exceed the threshold are identified as being occupied (a 1 event), and sites that are below the threshold are identified as being empty (a 0 event).

Three types of errors can occur during imaging:

(1) A false positive where an empty site is incorrectly identified as being occupied, which occurs with probability $p_{0\to1}$.

(2) A false negative where an occupied site is incorrectly identified as being empty, which occurs with probability $p_{1\to0}$.

(3) A false negative where an atom is lost during the course of the image, which occurs with probability p_{L}.

We will sometimes collectively refer to the first two processes as infidelity. Optimizing the imaging conditions — including the parameters of the imaging and cooling beams, the duration of the image, and the selection of the masks and threshold — amounts to minimizing $p_{0\to1}$, $p_{1\to0}$, and p_{L}, with the secondary goal of minimizing the duration of the image. The degree to which these three errors are problematic depends on the specific conditions of the experiment. For example, if a larger fraction of the sites are occupied, $p_{0\to1}$ errors are less problematic than $p_{1\to0}$ errors (and vise versa when a small fraction of sites are occupied). To avoid concerns of biasing, we typically opt to make a best effort attempt at minimizing the sum of all three

[21] This procedure can be thought of as a version of image deconvolution which takes advantage of the additional fact that bright points in the underlying image should come from only a discrete set of locations [305].

errors for conditions where half the sites are occupied, and do not subsequently adjust the image analysis or parameters as a function of experimental conditions.

To characterize the above errors, we take two images back to back and compare the histograms of the collected counts. One approach is to take the assumption that the empty distribution is Gaussian and the occupied distribution is Poissonian. In this case, one can perform a maximum likelihood fit of the measured histograms of the counts to extract the characteristic parameters of the different distributions. $p_{0\to1}$ corresponds to the part of the empty distribution that exceeds the threshold, and $p_{1\to0}$ to the part of the occupied distribution that is under the threshold. $1 - p_{\mathrm{L}}$ corresponds to the ratio of the area contained under the occupied peak in the second and first images.

Although the occupied distribution is Poissonian with high confidence, it is not the case that the empty distribution is always Gaussian. In particular, leakage light, imperfect charge transfer on the EMCCD, and high-order aberrations in the imaging optics can lead to a slightly skewed distribution. Additionally, atoms that are lost during imaging can result in a long tail that connects otherwise well-separated empty and occupied distributions. However, we typically do not care to distinguish $p_{1\to0}$ from p_{L}, since both false negative errors affect the measured signal in the same way. When only considering the overall false positive and false negative rates, we can measure the relevant errors in a model-independent way. Specifically, for calibration data involving pairs of images as described previously, we can directly measure the probability p_{01} of events x_{01} in which an empty site in the first image was occupied in the second, and probability p_{10} of events x_{10}, where an occupied site in the first image was empty in the second. Because $p_{0\to1}$, $p_{1\to0}$, and p_{L} are small, we neglect any effects that are second order in these errors. For well-isolated traps with no tunneling or thermal hopping, there is no physical process wherein an atom can be generated in the second image, and so $p_{0\to1} = p_{01}$. Similarly, any process that results in a missing atom leads to a false negative measurement error, and so $p_{10} = p_{1\to0} + p_{\mathrm{L}}$. Characterizing these errors allows us to separate the effect of measurement errors from other aspects of the experiment, as we will see in more detail in Chs. 6 and 8.

Although the imaging performance can fluctuate depending on the specific experiment at hand, using the above procedure we achieve a false positive probability of $p_{01} \simeq 0.1\%$ and a false negative probability of $p_{10} \simeq 0.1\%$ when performing 100 ms images in the 813 nm optical lattice. In deeper 813 nm traps, even better

imaging performance is possible [74] at the cost of reduced atom number given a fixed amount of available optical power.[22] The imaging performance in the 515 nm tweezers is a fair bit worse, with $p_{01} \simeq 1\%$ and $p_{10} \simeq 5\%$ for 75 ms images. It is not clear exactly what the mechanism leading to increased imaging loss in 515 nm traps is. Based on the branching ratio from the $^1P_1$ state to the $^1D_2$ state, we do not believe the observed loss is fully explained by anti-trapping of the $^1D_2$ state at 515 nm. We observe that the loss appears to be related to the number of 689 nm photons scattered during cooling, and not directly to the number of 461 nm photons scattered (except to the extent that more scattered 461 nm photons necessitates additional cooling). Note that although we do not observe any loss during the optical cooling step (without imaging), such loss has been observed in other experiments [73]. One hypothesis is that 515 nm (and 532 nm) photons have an energy above the two-photon ionization threshold from the $^3P_1$ state, leading to loss, whereas 813 nm is below this threshold and thus does not allow for this particular loss channel [207].

4.4.2 Image analysis in dense arrays

Image analysis is more challenging when the spacing between sites is comparable to the PSF of the imaging system, since an atom in a neighboring site can lead to increased background counts on an adjacent site. In our apparatus, the imaging system is designed to be able to resolve individual lattice sites, and so this is not a significant concern. However, to improve the fidelity of the procedure described in the previous section for denser arrangements of sites, we optimized the applied masks by training a single-layer neural network on simulated data corresponding to 30 % random filling of the lattice (see Fig. 4.7). The resulting masks are similar to the ones based on the measured PSF of our imaging system, but with a slight negative bias on adjacent lattice sites that reduces errors due to leakage light from one lattice site to another (see Fig. 4.7d). Although a deep neural network can result in better performance [149], the resulting performance gains are marginal in our setup due to the already high imaging resolution in comparison to the

[22] Images using the Sisyphus cooling mechanism described in [73, 208] can be a fair bit faster than those using sideband cooling, since each scattered 689 nm photon can remove more thermal energy. In 813 nm tweezers, we achieve $p_{01} \simeq 1\%$ and $p_{10} \simeq 0.1\%$ for 30 ms images using the Sisyphus cooling technique [240]. Note that although the same procedure is expected to work in 813 nm optical lattices, we have not had great success with this at the depths we work at (although we are optimistic that increases of $\sim 2-3\times$ in lattice depth could result in good performance). One additional consideration in a lattice is that tunneling can become significant in higher bands. As a result, in the relatively shallow lattices we work with, we must ensure that the atoms stay primarily in the ground band for the full duration of the image, eliminating the benefit of removing multiple motional quanta with a single scattered 689 nm photon.

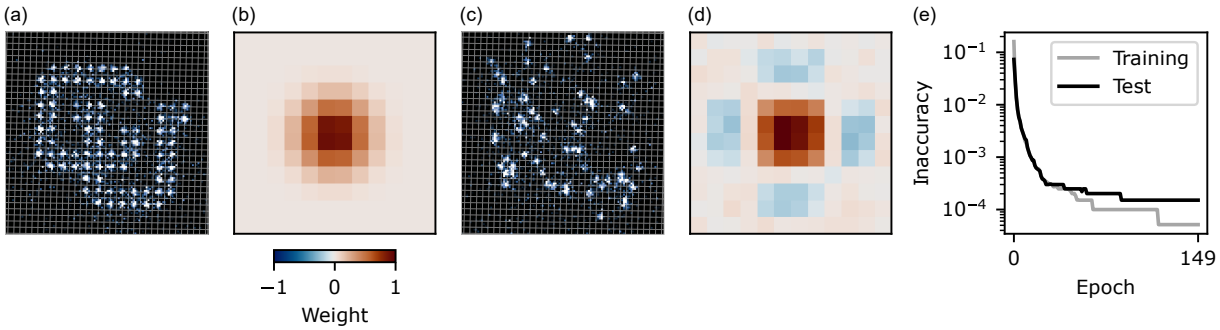

Figure 4.7: (a) Atoms separated by next-nearest-neighbor spacing (or more) in the lattice (lattice sites are indicated by where the grey lines cross) are well-resolved from each other, as is clear from the depicted single shot image of rearranged atoms. (b) When the atoms are well-resolved, it is sufficient to use a completely positive mask (as depicted) to identify the presence or absence of an atom. (c) Atoms at nearest-neighbor spacing in the lattice have slightly overlapping point spread functions, as seen in this single shot image taken after the atoms are allowed to tunnel (resulting in some atoms being in neighboring sites). (d) We train a single layer neural network to generate the appropriate mask for identifying the presence or absence of an atom when other atoms can occupy neighboring sites in the lattice. The resulting mask is similar to before, but with slightly negative weights on the neighboring lattice sites to suppress spurious detection events due to leakage light (from atoms on neighboring sites). Color scale is shared with (b). Note that the masks in (b, d) are representative examples, but that the masks applied in our analysis differ from lattice site to lattice site depending on the alignment of a given lattice site to the pixels on the camera. (e) Performance of the optimized masks as a function of training epoch. The inaccuracy is defined to be the proportion of events where the presence or absence of an atom on a given lattice site is incorrectly identified. Notice that whereas the performance continues to improve for the training data set, additional training does not improve the performance with respect to a test dataset that is independent of the training data. The ultimate performance for 500 ms images is similar to the performance for 100 ms images in sparse arrays.

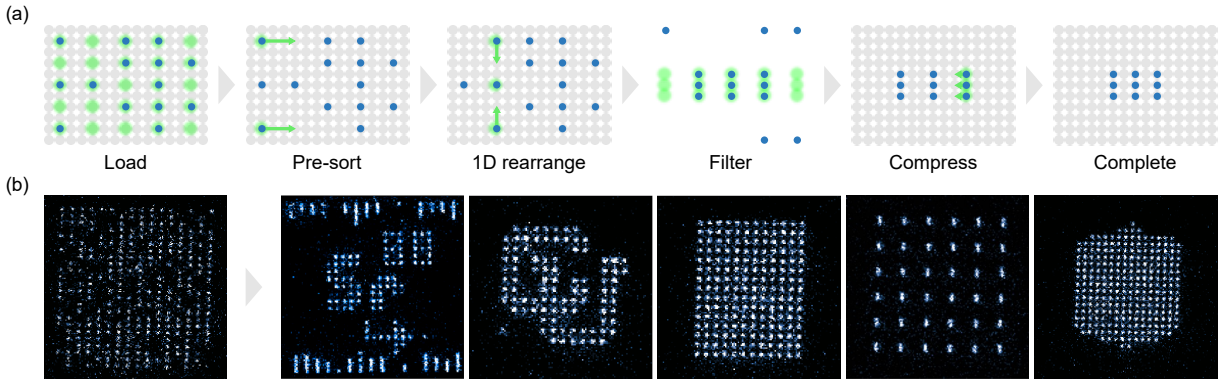

Figure 4.8: Procedure for atom rearrangement. (a) Rearrangement of atoms (blue circles) in an optical lattice (grey circles denote sites in the lattice) using optical tweezers (green) must balance several conflicting requirements, leading to the multistep algorithm described in the text. (b) A selection of single-shot images from the experiment, where each bright spot indicates the presence of an atom. Starting from a stochastically loaded array (left), one can prepare nearly arbitrary patterns of atoms in the lattice (right). Of particular note is the second-from-the-right pattern, which shows pairs of atoms positioned in neighboring sites in the lattice, and the far-right pattern, where the target pattern is adapted in real time to make use of every stochastically loaded atom in a given experimental trial.

lattice spacing. This being the case, we opt to use a single-layer neural network to ease the characterization of errors and image fidelity.

The resulting analysis yields a false negative probability of $p_{10} = 0.2(1)\%$ for a calibration data set using 500 ms-long exposure times, however, p_{10} can fluctuate day to day by $\sim 0.2\%$ for the experiments appearing in Chs. 7 and 8. The false positive probability is much lower, with a value of $p_{01} \sim 0.001\%$. However, for experiments with very low density (for example when one atom occupies an analysis region containing ~ 1000 sites) the effect of false positives can become comparable to or larger than the effect of false negatives when trying to correctly identify the presence and position of the atom.

4.5 Atom rearrangement

Once the atom positions are identified, we can use a mobile set of optical tweezers to deterministically rearrange the atoms. Rearrangement techniques have been used to great success in tweezer arrays [99, 19], and in programmable optical lattices [179], but to our knowledge, our experiments are the first to perform rearrangement in an optical lattice that allows for appreciable tunnel coupling (as we will discuss in detail in Chs. 7 and 8). Rearrangement in such a lattice must balance several conflicting requirements. In particular,

our tweezers have a beam waist that is comparable to the lattice spacing (see Sec. 3.3). This means that tweezers directed at a specific lattice site still have substantial overlap with adjacent sites, and so atoms trapped in those adjacent sites can experience higher loss rates, especially when the tweezers are moving. Additionally, overlapping tweezers will beat against each other due to our use of crossed acousto-optic deflectors to project the tweezer array [241]. Tweezers on adjacent lattice sites are separated in frequency by 1.95 MHz, far higher than the relevant trap frequencies of 50-200 kHz, and so 1D tweezer arrays with equal spacing to the lattice don't cause significant heating (see Sec. 3.3.3). However, because the lattice is slightly rectangular, placing a tweezer on every lattice site in 2D can result in lower beat frequencies along the diagonals of the array that cause significant heating. Our algorithm for atom rearrangement seeks to balance these concerns against the desire to minimize the total distance travelled by the atoms and the total duration of the rearrangement procedure.

The basic premise is to perform most operations in parallel with 1D arrays of tweezers [95, 324], and to perform as many operations as possible on a sublattice that does not lead to undesirable beat frequencies between overlapping tweezers. We choose to work with either nearest- or next-nearest-neighbor spacing along one axis, and third-nearest-neighbor spacing along the other axis. Atoms are stochastically loaded into tweezers on this sublattice with $50-75$ % filling (depending on whether enhanced loading is used, see Sec. 4.2) and implanted into the lattice [366]. Rearrangement proceeds in four stages (Fig. 4.8a):

(1) **Pre-sorting:** If a column is loaded with more atoms than are required in the target pattern, any excess atoms are pushed to the neighboring column in a single step using a 1D tweezer array. Similarly, if the column has too few atoms the missing atoms are pulled from the adjacent column. In the rare case that this step fails, the experiment is terminated, and we simply load a new ensemble of atoms.

(2) **1D rearrangement:** Each column is then rearranged in 1D, with excess atoms pushed to the edges of the array.

(3) **Filtering:** A subset of atoms is transferred back into a 2D tweezer array that only addresses the correct sublattice, and the lattice potential is extinguished. Laser light that addresses the free-space

$^1S_0 \leftrightarrow {}^3P_1$ transition is applied to resonantly blow away the excess atoms while optically cooling the tweezer-trapped atoms (since the AC Stark shifts imposed by the tweezers shifts the $^1S_0 \leftrightarrow {}^3P_1$ transition in the tweezer-trapped atoms to higher energies). The remaining atoms are subsequently transferred back into the lattice.

(4) **Compression:** Optionally, the columns of rearranged atoms can be translated closer together one column at a time to prepare denser target patterns.

The resulting algorithm allows for the preparation of nearly arbitrary patterns with 0 or 1 atoms on a given site. However, we are currently unable to translate columns of atoms to less than next-nearest-neighbor spacing without incurring significant loss in step (4). This means that although the prepared patterns can have atoms at nearest-neighbor spacing in one axis, the other axis must be spaced at next-nearest-neighbor spacing or greater.[23] The number of steps in the algorithm scales as $O(\sqrt{n})$, where n is the number of atoms in the target pattern. For target patterns with similar density to the loaded sublattice, the mean distance travelled by each atom is $O(1)$, and the runtime of the algorithm scales like $O(\sqrt{n})$. For dense target patterns the worst-case scaling of the mean distance travelled is $O(\sqrt{n})$ leading to a runtime of $O(n)$.

Throughout rearrangement the lattices are held at a constant depth of $U_{\mathrm{ax}}/\hbar = 2\pi \times 2.4$ MHz for the axial lattice, and $U_{\mathrm{2D}}/\hbar = 2\pi \times 1.7$ MHz for the 2D lattice. Resolved sideband cooling is continuously applied to the lattice-trapped atoms. To move a given set of atoms, tweezers are ramped on and off over 60 μs to a depth of $U_{\mathrm{tw}}/\hbar \simeq 2\pi \times 30$ MHz and moved with a constant speed of 26 μm/ms. Rearrangement occurs under magic conditions in the lattice for the $^1S_0 \leftrightarrow {}^3P_1$ transition (see Sec. 4.1), and so the tweezer-trapped atoms are shifted out of resonance from the cooling light while they are being translated (since the magic conditions for the lattice are different than for the tweezers). Because the appropriate moves in our algorithm can be computed efficiently, and due to the architecture of our tweezer control system [365], analyzing the atom images and programming the control system takes only $\lesssim 5$ ms for target patterns containing up to 270 atoms. However, there is currently a technical delay of 110 ms between taking the image of the stochastically loaded atoms and initiating rearrangement due to the time it takes to extract

[23] We expect that this limitation can be circumvented by improving the lattice alignment such that it is closer to being square, but this has yet to be confirmed.

the image data from the camera. This could be addressed in the future by performing image processing on the same FPGA hardware used for atom rearrangement [345]. For patterns containing up to 270 atoms, the rearrangement itself is performed in $\lesssim 30$ ms. The per-atom success rate for this rearrangement can be as high as 99.5 %, and is primarily limited by imaging fidelity and loss. However, due to drifts on the experiment, the typical per-atom success rate is 98 % for the experiments discussed in this work. Once the success of the rearrangement procedure is verified via a second image, the atoms are optically cooled to near their 3D motional ground state as described in Sec. 4.3. The resulting prepared states have very low entropy in all degrees of freedom of the atoms, and serve as the starting point for all subsequent experiments described in this work.[24]

[24] In some cases where the specific pattern of atoms is unimportant, we opt not to use rearrangement for convenience.

Chapter 5

Optical frequency metrology with tweezer arrays

"Typically, the subject being copied is terminated."
— The Terminator

Portions of this chapter have appeared in:

M. A. Norcia, A. W. Young, W. J. Eckner, E. Oelker, J. Ye, and A. M. Kaufman. Seconds-scale coherence on an optical clock transition in a tweezer array. **Science**, 366(6461):93–97, 2019

A. W. Young, W. J. Eckner, W. R. Milner, D. Kedar, M. A. Norcia, E. Oelker, N. Schine, J. Ye, and A. M. Kaufman. Half-minute-scale atomic coherence and high relative stability in a tweezer clock. **Nature**, 588(7838):408–413, Dec. 2020

In Ch. 4, we showed that we can prepare ensembles of atoms in well-determined initial states, and detect each atom individually. In this chapter, we will see how we can control a convenient two level (qubit) degree of freedom encoded in the electronic states of the atoms for frequency metrology. We will begin in Sec. 5.1 with a brief primer on quantum sensing, and identify the relevant features that make for a good sensor for frequency metrology (a clock). In Sec. 5.2 we will argue that the above qubits are particularly suited to frequency metrology, and mention some unique benefits and challenges associated with our approach. We will also discuss how we prepare, control, and read out the qubits, and the limitations on the fidelity of these operations. In Sec. 5.3, we will discuss the limitations on the coherence of the qubits, and how the carefully designed optical potentials in our experiment allow us to tailor the environment of the qubits to enhance their coherence. Although the record-high quality factor of these qubits is beneficial for clock measurements, we will discuss systematic perturbations to the qubit transition frequency that can affect the accuracy of such measurements. In Sec. 5.4, we will describe some specific clock measurements we have performed, including frequency comparisons between the atoms and a laser, as well as synchronous comparisons between the

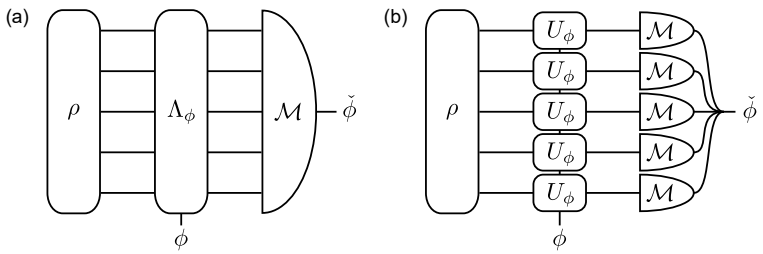

Figure 5.1: Schematic for quantum sensing. (a) A quantum sensing problem can generally be formulated as the task of estimating some parameter (or set of parameters) ϕ by preparing a state described by a density matrix ρ, evolving the state under a CPTP map Λ_ϕ which encodes the parameter of interest, performing a measurement described by a POVM $\mathcal{M}$, and using the results in an estimation procedure $\check{\phi}$. (b) Often, the quantum sensor is composed of multiple similar subsystems. Phase estimation is a restricted measurement problem where the signal Λ_ϕ acts identically but independently (for example, with a local unitary U_ϕ) on each subsystem. Although one could perform a joint measurement on the entire system, the same performance can be achieved using only local measurements $\mathcal{M}$.

atoms. We will conclude in Sec. 5.5 with a brief discussion of how the so-called "tweezer clock" developed in our group, and in parallel at Caltech [209], fits into the broader field of frequency metrology.

5.1 A primer on quantum sensing

In almost any experiment, we are interested in extracting some parameter ϕ relating to the physical world from experimental data.[1] This problem can, in general, be framed in terms of three steps (see Fig. 5.1): preparation of some probe state described by a density matrix ρ, evolution of that state under a completely positive trace preserving (CPTP) map Λ_ϕ that encodes the parameter ϕ of interest (and can capture both unitary and non-unitary evolution), and measurement of the resultant state described by a positive operator-valued measure (POVM) $\mathcal{M}$ (which is a generalization of a typical projective measurement) [306]. This sequence is then repeated M times to estimate ϕ to greater precision.[2] The question is how to achieve the lowest uncertainty in the estimate of ϕ given restricted physical resources (typically some combination of limitations on the total duration of the measurement, and physical constraints on the available probe states).

For starters, let us fix ρ, Λ_ϕ, and $\mathcal{M}$. In this case each measurement gives us an outcome x sampled from the random variable X which follows the probability density distribution $p(X|\phi)$. The result is a dataset

[1] In principle, ϕ can be a set of multiple parameters, but we omit that possibility in this simplified discussion.

[2] One can employ Bayesian techniques to adaptively adjust ρ, Λ_ϕ, and $\mathcal{M}$ [227], but these techniques are outside the scope of this work.

$\boldsymbol{x} = \{x_1, x_2 \ldots x_M\}$ of independent and identically distributed (IID) samples. Our goal is to construct an estimator $\check{\phi}(\boldsymbol{x})$ that takes the data $\boldsymbol{x}$ and outputs an estimate of ϕ.[3] $\check{\phi}$ should be "unbiased," in the sense that its expectation value is equal to ϕ:[4]

$$\langle \check{\phi} \rangle \coloneqq \int d\boldsymbol{x}\, p(\boldsymbol{x}|\phi)\check{\phi}(\boldsymbol{x}) = \phi, \tag{5.1}$$

and "consistent," in the sense that it converges to ϕ as M increases:[5]

$$\lim_{M\to\infty} \check{\phi} = \phi. \tag{5.2}$$

With this in mind, we are interested in minimizing the mean squared error[6] of $\check{\phi}$ with respect to $\boldsymbol{x}$:

$$\Delta^2\check{\phi} \coloneqq \langle(\check{\phi} - \phi)^2\rangle = \int d\boldsymbol{x}\, p(\boldsymbol{x}|\phi)(\check{\phi}(\boldsymbol{x}) - \phi)^2. \tag{5.3}$$

If $\check{\phi}$ is unbiased, it obeys the Cramér-Rao bound [120]:[7]

$$\Delta^2\check{\phi} \geq \frac{1}{M\mathcal{I}\left(p(X|\phi)\right)}, \tag{5.4}$$

where $\mathcal{I}$ is the classical Fisher information (see Appendix A.1). An estimator is known as "efficient" if it saturates this bound. Notice that Eq. 5.4 includes the effect of the central limit theorem: M different IID random variables can be combined into an estimate that averages down with a variance of $1/M$.

In the previous case, the estimation was dictated by the classical probability distribution $p(X|\phi)$. However, since we are dealing with a quantum state, we are free to choose our measurement $\mathcal{M}$, resulting in a probability distribution $p(X|\phi, \mathcal{M})$ and associated Fisher information $\mathcal{I}(\phi, \mathcal{M}) \coloneqq \mathcal{I}(p(X|\phi, \mathcal{M}))$. Maximizing the Fisher information over all $\mathcal{M}$ leads to the quantum Fisher information $\mathcal{Q}(\phi) \coloneqq \max_{\mathcal{M}} \mathcal{I}(\phi, \mathcal{M})$ [40] and an associated quantum Cramér-Rao bound (see Appendix A.2):

$$\Delta^2\check{\phi} \geq \frac{1}{M\mathcal{Q}\left(\phi\right)}. \tag{5.5}$$

[3] As in [309], we use $\vee$ to denote an estimator rather than $\wedge$ to avoid confusing an estimator with a quantum operator.

[4] For discrete random variables, one can replace the integrals in this section with sums.

[5] Note that there is a subtle difference between being consistent and being unbiased. An unbiased estimator can never converge and therefore be inconsistent, and a consistent estimator can be biased for a finite sample size.

[6] This is equivalent to the variance in the case that $\check{\phi}$ is unbiased [306].

[7] In fact, $\check{\phi}$ needs only be asymptotically locally unbiased to obey this bound [76], meaning that it is unbiased in the large M limit, and about some specific value of ϕ.

The above discussion provides a general framework to evaluate different kinds of measurements, where optimization of a given measurement amounts to optimizing over ρ, Λ_ϕ, and $\mathcal{M}$. Next, we look at a specific example that is relevant to our work.

5.1.1 Phase estimation

Consider a measurement apparatus composed of N qubits.[8] The single qubit Hilbert space $\mathcal{H}_1$ is spanned by the states $\{|0\rangle, |1\rangle\}$, and the combined Hilbert space is $\mathcal{H} = \mathcal{H}_1^{\otimes N}$. We are interested in the situation where $\Lambda_\phi = U_\phi^{\otimes N}$ is a unitary transformation that acts independently but identically on each qubit.[9] Specifically, each qubit undergoes the unitary transformation:

$$U_\phi = e^{-i\phi H}, \tag{5.6}$$

where the generator H is some known Hermitian operator acting on $\mathcal{H}_1$. The generator acting on the whole system is $H_\mathcal{H} = \sum_{i=1}^N H_i$, where H_i refers to H acting on the i^{th} qubit (and zeroes elsewhere). Additionally, let $\rho = |\psi\rangle\langle\psi|$ be a pure state, and $\mathcal{M}$ a projection-valued measurement (PVM) in the $\{|0\rangle, |1\rangle\}$ basis of each qubit. In this case, the quantum Fisher information is [40, 41] (see Sec. A.2):

$$\mathcal{Q} = 4(\langle\psi \,|\, H_\mathcal{H}^2 \,|\, \psi\rangle - |\langle\psi \,|\, H_\mathcal{H} \,|\, \psi\rangle|^2) = 4\Delta^2 H_\mathcal{H}, \tag{5.7}$$

where we have defined $\Delta^2 H_\mathcal{H}$ to be the variance in the expectation of $H_\mathcal{H}$ given the state $|\psi\rangle$.

Let the lowest and highest eigenvalues of H be λ_- and λ_+ respectively, with corresponding eigenstates $|\lambda_-\rangle$ and $|\lambda_+\rangle$. If we restrict $|\psi\rangle$ to product states of the different qubits then $\Delta^2 H_\mathcal{H} = \sum_{i=1}^N \Delta^2 H_i$, where $\Delta^2 H_i$ is the variance of the measurement outcome on the i^{th} qubit. As a result, the largest $\Delta^2 H_\mathcal{H}$ we can achieve is with the state [119]:

$$|\psi\rangle = \left(\frac{|\lambda_-\rangle + |\lambda_+\rangle}{\sqrt{2}}\right)^{\otimes N}. \tag{5.8}$$

In this case $\Delta^2 H_\mathcal{H} = N(\delta\lambda)^2/4$, where $\delta\lambda = \lambda_+ - \lambda_-$, leading to a bound of:

$$\Delta^2\check{\phi} \geq \frac{1}{MN(\delta\lambda)^2}. \tag{5.9}$$

[8] The analysis in this section applies just as well to qudits (quantum systems with more than two levels).

[9] This applies whenever the signal of interest acts globally on all qubits, but the qubits are otherwise decoupled from each other. For example, when a uniform electric or magnetic field acts on isolated atoms which each encode a qubit degree of freedom.

Eq. 5.9 is known as the standard quantum limit (SQL)[10] , which can be interpreted as the lowest-variance estimation of ϕ that one can achieve when using multiple unentangled quantum sensors. Notice that the SQL simply corresponds to combining MN different IID measurements, where M characterizes the number of separate measurements, N the size of the sensor (number of qubits), and $\delta\lambda$ the energy cost.

Removing the restriction that $|\psi\rangle$ be a product state, $\Delta^2 H_{\mathcal{H}}$ is maximized by an even superposition of the states with maximum and minimum eigenvalues of $H_{\mathcal{H}}$, namely:[11]

$$|\psi\rangle = \frac{|\lambda_-\rangle^{\otimes N} + |\lambda_+\rangle^{\otimes N}}{\sqrt{2}}, \tag{5.10}$$

which is known as a Greenberger-Horne-Zeilinger (GHZ) state. For the above GHZ state, $\Delta^2 H_{\mathcal{H}} = (N\delta\lambda)^2/4$, leading to a bound of:

$$\Delta^2\check{\phi} \geq \frac{1}{MN^2(\delta\lambda)^2}. \tag{5.11}$$

Eq. 5.11 is known as the Heisenberg limit (HL), which can be interpreted as the lowest variance estimation of ϕ that one can achieve with a quantum sensor composed of many identical constituents. Expressed as a standard deviation ($\Delta\check{\phi} \coloneqq \sqrt{\Delta^2\check{\phi}}$), the SQL scales like $\Delta\phi \propto 1/\sqrt{N}$, whereas the HL scales like $\Delta\phi \propto 1/N$, leading to the often quoted quadratic enhancement with respect to N provided by entangled sensors. The scaling of the SQL and HL with respect to all other parameters are equal. Additionally, note that the HL can be saturated with only local measurements, meaning that entanglement is only necessarily useful in the preparation of ρ, and not required in $\mathcal{M}$ [119]. The $1/\sqrt{N}$ scaling of uncertainty in the SQL is often referred to as the limit imposed by quantum projection noise (QPN) or "shot noise," since it is noise associated with the discrete nature of the measured signal.

The HL and SQL, at their core, simply describe a tradeoff between the coherent and incoherent combinations of different signals (namely, adding amplitudes or probabilities), with the former performing up to quadratically better than the latter. Note that, depending on the sensing problem, entanglement is not always necessary to achieve the best possible performance. For example, if one takes a single qubit and evolves it for N times longer than an N qubit measurement, the effect is an N times enhancement of

[10] This terminology is somewhat unfortunate. Although I don't know where it originated, people often lament that the SQL "is neither standard, quantum, nor a limit."

[11] Note that there are other entangled states that can achieve similar performance, like the spin squeezed states that we discuss in Ch. 6.

$\delta\lambda$, and thus equal performance to the HL for the shorter-duration measurement with N qubits. Instead, entanglement should be thought of as a way to trade off between resources in time and space, as is often true in quantum algorithms [268, 43, 120]. As a result, when discussing enhancements in metrology relating to entanglement, one must be very careful to specify the specific sensing problem at hand, and what the relevant restricted resources are [145, 10, 120].[12]

5.1.2 On good clocks, qubits, and decoherence

Turning to the case of a typical clock measurement, the measured quantity is the elapsed time per measurement run $\phi \to t$, and U_ϕ simply corresponds to time evolution under the Hamiltonian $H = \delta\lambda\sigma_z/2$, where σ_z is the Pauli-z matrix and $\delta\lambda$ is now an energy. We are interested in attaining the best possible measurement of the fractional uncertainty $\Delta t/t$ for a given total measurement duration $T = Mt$ (assuming no dead time, see Sec. 5.1.3.2). If one saturates the SQL, the resulting performance is:[13]

$$\frac{\Delta t}{t} = \frac{1}{2\pi f\sqrt{NTt}}, \tag{5.12}$$

where $f \coloneqq \delta\lambda/h$ is the frequency associated with the precession of the qubits. Note that in the above expression we have fixed f to a known value, and associate fluctuations in the measurement with a fluctuation in t. This situation applies, for example, when using a frequency reference to distribute a time standard. Alternatively, one can fix the value of t, and take $\phi \to f$ and $H \to t\sigma_z/2$, arriving at a similar expression:

$$\frac{\Delta f}{f} = \frac{1}{2\pi f\sqrt{NTt}}. \tag{5.13}$$

This situation applies, for example, when measuring a frequency difference between two clocks. Given the scalings of the above measurement precisions, in this chapter, we focus on maximizing f, N, and t. [14] We will describe the preparation of states which improve the scaling with respect to N, including those of the form in Eq. 5.10, in Ch. 6.

Maximizing f has led to atomic clocks operating at increasingly high frequencies, first in the microwave domain [342], and more recently in the optical and ultraviolet domains [203]. Maximizing t by improving

[12] This is generically true when discussing quantum enhancements of any kind, like in Sec. 7.7.
[13] Note that the 2π comes from associating H with a Hamiltonian, rather than a generator of the form in Eq. 5.6.
[14] We leave the maximization of 2π for later works.

the lifetime of the qubits is important in the SQL, but also has significant implications for any entanglement that leads to performance beyond the SQL. For example, consider the case where the qubits decay from $|\lambda_+\rangle$ to $|\lambda_-\rangle$ with a time constant τ. For the GHZ state in Eq. 5.10, decay of a single qubit completely decoheres the state, and so the effective decay time is τ/N. Similar arguments apply when τ characterizes a more general decoherence model, for example involving both decay and dephasing [145, 299, 85]. If t is limited by τ/N, then any enhancement due to GHZ-like entanglement is cancelled by a reduction in t. In this scenario, the only benefit coming from entanglement is a tradeoff between time and space: one can perform N parallel measurements of duration $t \sim \tau$ with unentangled atoms, or N sequential measurements of duration $t \sim \tau/N$ with maximally entangled atoms. The latter case can be beneficial in the sense that the measurement has increased bandwidth (one can be sensitive to variations in the detected signal at higher frequencies), but does not change the ultimate sensitivity of the measurement to a static signal. Taking the expected decoherence into account and preparing the appropriate partially entangled state can lead to a constant enhancement over the SQL of $1/e$ in variance [145], but not to a modification of the asymptotic scaling of the procedure with N. However, an enhanced scaling is possible when t is limited by other effects [10], as we will see in more detail in Ch. 6. Although outside the scope of this work, more sophisticated protocols involving error correction can restore the HL in the presence of certain kinds of errors [171, 86, 84, 377].[15]

With the above discussion in mind, a useful figure of merit is the unitless product of τ and the angular frequency $\omega \coloneqq 2\pi f$, leading to the definition of the quality factor $Q = \omega\tau/2$. If τ is exclusively due to decay, Q can be interpreted as the total angular displacement (in radians) accrued by an oscillator before its energy decays to $1/e$ of its original value.

Notice that the desirable features for quantum sensing are very similar to those for a good qubit or quantum memory: we want the ability to prepare specific (possibly entangled) pure states with high fidelity, and to isolate these states from uncontrolled influences of the environment. The key additional requirement for quantum sensing is that we want our qubit to be as sensitive to a specific signal of interest as possible, which, in the case of frequency metrology, means that we want our qubit states to be as far apart in energy as possible. By contrast, for typical computing applications there is no requirement on the energy

[15] The necessary and sufficient condition for error correction restoring the HL in metrology ends up being quite intuitive: the generator H of the signal must not be contained in the linear span of all errors [377].

separation between qubit states, and reducing this energy separation can be useful for reducing sensitivity to other sources of noise [155]. Another figure of merit is the number of gates that can be applied within the coherence time of a qubit. To a certain extent this is characterized by Q, in the sense that the value of Q tells you how many σ_z gates one can apply, but one must also consider the other gates making up a universal gate set. Even when considering these other gates, we will see that the qubit we choose to work with in the context of frequency metrology also ends up being a reasonably attractive qubit for quantum information applications.

5.1.3 Additional considerations in optical clocks

In an atomic clock, the mechanism for timekeeping is a periodic oscillation in the electronic states of the atoms, or, in the language of the previous section, the energy difference between the qubit levels. However, in optical clocks these oscillations are much too fast to detect directly.[16] Instead, we need access to a reference frame that rotates at or near the difference frequency of the qubit. This is typically done with a laser,[17] which additionally allows one to apply single qubit gates as described in Sec. 2.2.1. With the above requirement in mind, the typical operation of an optical atomic clock involves a local oscillator (LO) in the form of an ultrastable laser whose frequency is periodically compared to the atomic reference frequency, and adjusted accordingly [203]. The ultimate limit on the frequency instability (defined in Sec. 5.1.3.1) of the LO is set by the ability to measure the difference frequency between the LO and the atoms, and thus by Eq. 5.13 for unentangled atoms.

To understand the joint performance of the atom-LO system, one must consider effects relating to noise and uncertainty on both the LO and the atoms, as well as effects relating to how the LO frequency is referenced to that of the atoms. Besides QPN,[18] and decoherence leading to finite interrogation times, one must also consider various sources of technical noise that result in uncertainty or fluctuations in H. We will discuss the relevant perturbations to H as they apply to our experiments in more detail in Sec. 5.3. For now, we introduce a standard figure of merit for optical clocks, point out one additional source of error that

[16] The fastest photodiodes can attain impressive bandwidths of up to hundreds of gigahertz [195], but this is still a far cry from the hundreds of terahertz at which optical clocks operate.

[17] We will describe measurements where other atoms serve as the LO instead of a laser in Sec. 5.4.2.

[18] Typically, it is the QPN of the atoms that dominates, since measurements of the LO benefit from large numbers of available photons.

often prevents lattice clocks from reaching the performance limits set by QPN, and discuss some key sources of perturbations to H in optical lattice clocks.

5.1.3.1 Allan deviation

As introduced in the previous section, we are interested in characterizing the fractional frequency difference $y \coloneqq (f - f_{\mathrm{LO}})/f$ between the atomic reference frequency f and the LO frequency f_{LO}.[19] Although the variance of y comes to mind as a reasonable figure of merit, a typical variance can diverge under commonly occurring sources of noise in a clock. Instead, we define the Allan variance [273, 275]:[20]

$$\sigma_y^2(T) \coloneqq \frac{1}{2M-1} \sum_{i=1}^{M-1} \left(\langle y(T) \rangle_{i+1} - \langle y(T) \rangle_i \right)^2 , \tag{5.14}$$

where $\langle y(T) \rangle_i$ denotes the i^{th} independent measurement of the average value of y over a measurement duration T, and M total measurements are performed. The Allan variance differs from a standard variance in the sense that one is computing deviations between adjacent measurements, not deviations from the mean. As a result, the Allan variance converges for most realistic sources of noise in a clock [275]. In practice, one often uses the overlapping Allan variance, which behaves similarly to the Allan variance, but provides tighter error bars on the value of the variance inferred from a given data set by removing the restriction that the different samples $\langle y(T) \rangle_i$ are fully independent [273, 275]. One often quotes the square root of the Allan variance, or the Allan deviation $\sigma_y(T)$, which is also known as the "instability".[21]

Based on the discussion so far, $\sigma_y(T)$ scales like $1/\sqrt{T}$ due to the SQL, but different sources of noise can lead to different behaviors. For example, drifts or random walks in frequency can lead to a levelling off or even increase of $\sigma_y(T)$ with T [275]. It is common to quote the instability at one second of averaging, with the understanding that a $1/\sqrt{T}$ scaling applies only at short averaging times when the contribution from sources of noise that break this scaling is negligible.

[19] Recall from Eqs. 5.12 and 5.13 that measuring frequency and elapsed time are similar.

[20] Note that we use the symbol Δ^2 to denote standard variances, and σ^2 to denote Allan variances. However, the contribution to the Allan variance from QPN is the same as in Eq. 5.13 when replacing Δ with σ.

[21] Note that people sometimes interchangeably refer to this as the stability. However, in these cases, "high stability" refers to low values of $\sigma_y(T)$.

5.1.3.2 Dick effect noise

Often, optical atomic clocks operate in a cyclic fashion where, on each cycle of duration t_c, a new ensemble of atoms is prepared, evolved, and then measured destructively.[22] If the interrogation time is t, the duty cycle of a clock is defined as $\zeta \coloneqq t/t_c$. For neutral atom-based optical clocks initial trapping and cooling can take ~ 100 ms (see Sec. 3.1.2) in comparison to evolution times of $t \sim 100 - 1000$ ms, and so ζ can be fairly low.

Cyclic operation of the atomic reference in combination with $\zeta < 1$ results in increased sensitivity to frequency noise in the LO through what is known as the Dick effect [90]. Specifically, frequency noise at harmonics of $1/t_c$ will be downconverted to near DC, resulting in errors in the inferred center frequency of the LO and thus higher Allan variance. The contribution to the Allan variance from the Dick effect, $\sigma^2_{y,\mathrm{D}}$, can be calculated by computing the Fourier components of the sensitivity function associated with a given measurement [90]. For Ramsey measurements (see Sec. 5.2.4), where the sensitivity function is a square wave that is high during t, and low for the remainder of t_c,[23] $\sigma^2_{y,\mathrm{D}}$ takes the form [285]:

$$\sigma^2_{y,\mathrm{D}}(T) = \frac{1}{Tf\zeta^2} \sum_{m=1}^{\infty} S_f\left(\frac{m}{t_c}\right) \left(\frac{\sin(\pi m \zeta)}{\pi m}\right)^2, \tag{5.15}$$

where $S_f(\cdot)$ is the one-sided power spectral density of the difference frequency $f - f_{\mathrm{LO}}$.

For clocks which operate continuously, like thermal beam clocks or masers, $\zeta = 1$, and so the Dick effect vanishes [203]. For ion clocks, QPN typically dominates over Dick effect noise due to both high duty cycles,[24] and small ion numbers (leading to high QPN). However, in neutral atom-based optical clocks Dick effect noise can dominate over QPN due to low duty cycles and large atom numbers. In Sec. 5.4, we attempt to mitigate the Dick effect by increasing ζ. Alternatively, one can circumvent the Dick effect by interleaving measurements with two or more clocks such that the sensitivity function never drops to zero [291, 242].

[22] Terminated.
[23] Assuming instantaneous control pulses.
[24] In ion clocks the same ion is typically used for multiple interrogation cycles (not terminated), with only a few milliseconds of cooling and optical pumping being required to reinitialize the ions [203].

5.1.4 Fluctuations and systematic shifts

Fluctuations in H lead to decoherence and thus impose limits on t, and on the resulting stability of a clock, whereas uncertainty in the characterization of H limits the accuracy of a clock. The distinction between stability and accuracy is key when considering the performance of a clock as a frequency standard. Whereas stability relates to how quickly a measurement converges, accuracy relates to the degree to which the measurement provides an absolute estimate of the frequency of a given unperturbed atomic transition. Although good stability and accuracy is always preferable, certain applications can be fulfilled with only good stability and poor accuracy, or vise-versa.[25]

Variations in H can arise from any of the perturbations discussed in Ch. 2, including AC and DC electromagnetic fields. The pristine sensitivity of optical clocks means that these perturbations can sometimes come from surprising places. For example, assuming that the control and trapping fields are appropriately designed, the largest systematic shifts in optical lattice clocks often come from Stark shifts imposed by black body radiation (BBR).[26] For Sr, room temperature BBR results in a systematic shift of $\sim 5 \times 10^{-15}$ in fractional frequency, and is the dominant source of uncertainty for state-of-the-art lattice clocks with fractional frequency uncertainties at the low 10^{-18} level [234, 220, 34]. Because these BBR shifts scale with the radiant energy density of the BBR field, they follow the T^4 scaling of the Stefan-Boltzmann law (where T here refers to temperature). As a result, experiments are underway that try to minimize BBR-related effects by operating optical clocks in a cryogenic environment [336].

At the 10^{-18} level in uncertainty, gravitational redshifts also start to play a significant role. Specifically, a fractional frequency uncertainty of 10^{-18} corresponds to an uncertainty in altitude at the Earth's surface of ~ 1 cm [67, 220], which is comparable to or better than the best measurements of altitude using traditional survey techniques [253]. This means that, in the pursuit of lower uncertainty, optical clocks could also soon become our best tools for performing geodesy [203, 220]. In fact, for differential measurements that do not

[25] For example, accuracy is paramount if one wants to compare measurements to ab-initio calculations, or to compare measurements performed in different apparatuses or at very different times. If one is instead interested in a time-varying or differential signal between two or more clocks, stability with low accuracy is sometimes sufficient.

[26] At room temperature, the peak in the black body spectrum is at ~ 9.7 μm. This is a fair bit longer than relevant transition wavelengths in Sr, and so the dominant effect of BBR is a DC-Stark shift (see Sec. 2.1.1.5) that is proportional to $\langle \mathcal{E}^2 \rangle$, where $\mathcal{E}$ is the electric field that is present due to thermal fluctuations. However, for clocks operating with state-of-the-art accuracy, corrections from the AC polarizability can be relevant too.

require a full accuracy evaluation, measurements of differences in the gravitational redshift at the Earth's surface on the millimeter scale have already been achieved using optical clocks [35].

5.2 Optical clock qubits

To summarize the last section, for the purposes of frequency metrology we are interested in working with qubits that balance several features. For starters, we want qubits with a large energy separation, long coherence times, and well-characterized evolution. These ingredients are met by using neutral Sr atoms trapped in magic wavelength optical potentials (see Sec. 2.2.3), which have been extensively characterized and used to great effect in state-of-the-art optical lattice clocks [203]. Additionally, we want to balance the desire for large ensembles of qubits with fast, high fidelity operations (including both manipulations of the qubits and readout). Optical lattice clocks tend to operate with large ensembles of thousands to tens of thousands of atoms, but do not provide access to controls and observables at the single atom level.[27] By contrast, trapped ion clocks typically work with a single ion, but achieve fast and high fidelity control and readout of that ion [42]. In our work, we attempt to balance the above considerations by scaling to ensembles of intermediate size ($\sim$ 100 atoms), while maintaining high fidelity, single-atom-resolved control and readout.

Throughout this chapter, we will report on results in three slightly different experimental conditions, differentiated by the specific science potential[28] used:

(1) An array of 10 tightly confining clock-magic wavelength tweezers [240].

(2) An array of 320 shallow clock-magic wavelength tweezers [365].

(3) A clock-magic wavelength 3D optical lattice [290, 96].

In (2) and (3), atoms are prepared via implantation from a separate set of tightly confining preparation tweezers (see Ch. 3). In all cases, the optical traps are initialized with atoms that are laser-cooled to near their 3D motional ground states. The cooling performance in (1) and (2) are similar, and worse than that

[27] Although some progress has been made in this direction [213].
[28] As defined in Sec. 3.3.

achieved in (3) (see Sec. 4.3.3). With the above considerations in mind, we will use a qubit degree of freedom that is encoded in the $|^1S_0\rangle$ ground and $|^3P_0\rangle$ clock electronic states (see Sec. 2.1.1) of an optically trapped Sr atom. That being the case, we will sometimes refer to the atoms as "optical clock qubits" or just "clock qubits", with states $|0\rangle$ and $|1\rangle$ referring to the ground and clock states respectively.

The three different modes of operation are the result of balancing different conflicting requirements. To minimize the effect of QPN, we are interested in preparing and interrogating large ensembles of clock qubits in parallel. However, given finite optical power in a tweezer array, an increase in sample size comes at the expense of trap depth and atomic confinement, with implications for detection fidelity, cooling performance, qubit coherence, and atomic loading.[29] Using different optical potentials optimized for different stages of the experiment allows us to circumvent some of these challenges. By using an optical lattice potential in condition (3) [290, 365], we are able to make many traps that are at the clock-magic wavelength (Sec. 2.2.3) and tightly confining, enabling high fidelity imaging, ground state cooling (Ch. 4), and single qubit gates (Sec. 5.2.2). The combination of relatively long-lived atomic coherence with high fidelity single qubit gates will be critical in Ch. 6. However, as we will see in Sec. 5.3, a conventional optical lattice is not ideal when optimizing for atomic coherence as opposed to gate fidelity, motivating the experiments in condition (2). Working in condition (2) results in relatively high atom loss incurred when imaging in 515 nm potentials [73, 241] compared to the performance possible in (1) or (3) [74, 240, 290, 367] (see Sec. 4.4). However, this loss is not a significant issue in the specific case where one is optimizing for high precision measurements with non-entangled clock qubits, since the constant reduction in precision due to preparation and measurement errors is more than offset by the enhancement in coherence time.

5.2.1 Measurement

In all experiments, we take an initial image to identify the positions of the atoms (see Sec. 3.4). After running a given experimental sequence, we can either measure the states of the qubits non-destructively or destructively. In the former case, we simply image atoms in the $|0\rangle$ state as described in Sec. 4.4. Optionally, after this image one can perform a π-pulse to swap the $|0\rangle$ and $|1\rangle$ populations, and take an additional image

[29] These considerations are what limits the system size to 10 tweezers in (1) [240].

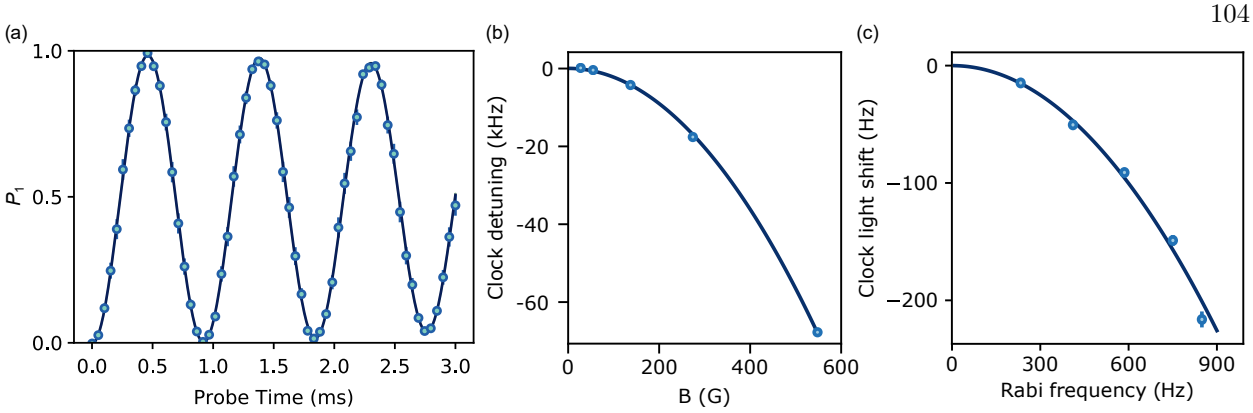

Figure 5.2: Clock Rabi oscillations. (a) P_1 refers to the (measurement error corrected) probability of observing an atom in the $^3P_0$ state, and is measured after initializing the atoms in the $^1S_0$ state and applying a laser probe resonant with the $^1S_0 \leftrightarrow {}^3P_0$ clock transition for variable duration. (b) Due to the quadratic Zeeman shift, the resonant frequency of the clock transition varies with the applied magnetic field. We measure the strength of this shift to be $\beta = -226(5)$ mHz/G^2. (c) The probe laser can also apply a differential light shift to the clock transition, leading to an increasing energy shift for higher clock Rabi frequencies. For normal operating conditions in our experiment, we observe a probe shift of $-280(10)$ mHz/Hz^2.

in order to distinguish between atoms that are in the clock state and atoms that are lost. The above scheme has the downside that $|1\rangle$ can decay to $|0\rangle$ during detection, leading to incorrect identification of the state of a qubit. In condition (1), for 3900 to 6000 E_R deep tweezers, where E_R is the recoil energy of an 813 nm photon, this leads to a 5 to 9% probability that $|1\rangle$ decays to $|0\rangle$ during a measurement [240]. In the case of destructive measurements, one can mitigate errors due to decay by applying a "blow-away" pulse of light resonant with the $^1S_0 \leftrightarrow {}^1P_1$ transition to remove atoms that are in the $|0\rangle$ state. The blow-away operation is very fast in comparison to the lifetime of $|1\rangle$ (< 100 μs in comparison to seconds), and so mapping $|0\rangle$ to the vacuum state with no atom can be performed with negligible infidelity. The remaining atoms in $|1\rangle$ can then be repumped to the ground state (see Sec. 2.1.1.4), which can also be performed with effectively no errors, and imaged. As a result, the destructive measurement scheme is limited only by imaging fidelity and loss, resulting in readout fidelities as high as 99.9% (see Sec. 4.4). These high detection fidelities mean that, unlike typical lattice clocks [203], tweezer clocks have effectively no uncertainty due to fluctuations in atom number or detection errors.

5.2.2 Single qubit gates

After loading ground-state-cooled atoms into the science potential of choice, we can interrogate the clock transition. We apply a magnetic field between 1.4 G and 550 G to mix the $^3P_1$ state into the $^3P_0$ state (see Sec. 2.1.1.3). By applying "probe" laser light that is referenced to an ultra-stable crystalline cavity [242] (see Appendix D) and resonant with this transition, we can drive Rabi oscillations between $|0\rangle$ and $|1\rangle$ (see Sec. 2.2.1). These Rabi oscillations correspond to global rotations about the x axis of the Bloch sphere, and can be performed with Rabi frequencies as high as $\Omega \simeq 2\pi \times 3$ kHz at $B = 550$ G and with ~ 100 mW of optical power. A measurement of these Rabi oscillations under typical operating conditions is shown in Fig. 5.2a. Global rotations about the z axis can be performed by simply detuning the probe laser from the atoms, changing the phase of the rotating frame.[30] Local rotations about the z axis can be performed by applying a local differential light shift between the ground and clock states for a given atom using a non-magic optical tweezer. Similar local operations are used to apply the oracle in Ch. 7 [366], and are demonstrated in the context of optical clock measurements in [96].[31] The combination of local z rotations and global x rotations constitutes a universal set of single qubit gates. We will discuss entangling gates and the possibility of a complete universal gate set in Ch. 6.

The above scheme for performing single qubit gates comes with a few additional considerations. Although both the ground and clock state have total angular momentum $J = 0$, and thus a vanishing linear Zeeman shift, one does have to consider the quadratic Zeeman shift (Fig. 5.2b). Rewriting Eq. 2.5, we have:

$$\Delta_B = \beta B^2, \tag{5.16}$$

where B is the strength of the applied magnetic field, and $\beta = -233$ mHz/G^2 [247]. In combination with magnetic field noise, the quadratic Zeeman shift leads to a tradeoff between the speed of the σ_x gates, which prefers larger B, and the sensitivity of the qubits to magnetic field noise, which prefers smaller B. For example, a B-field stability in the range of 1 mG is typical in most experimental labs. For $B = 550$ G, a 1 mG fluctuation leads to a shift of 0.26 Hz, which is enough to dramatically reduce the coherence times of

[30] These z rotations can be performed with essentially perfect fidelity and on the scale of 100 ns, since they amount to changing the phase of a digitally synthesized RF signal that controls the AODs used to switch the probe beam on and off.

[31] We leave a more detailed discussion of local z gates to later works from the group.

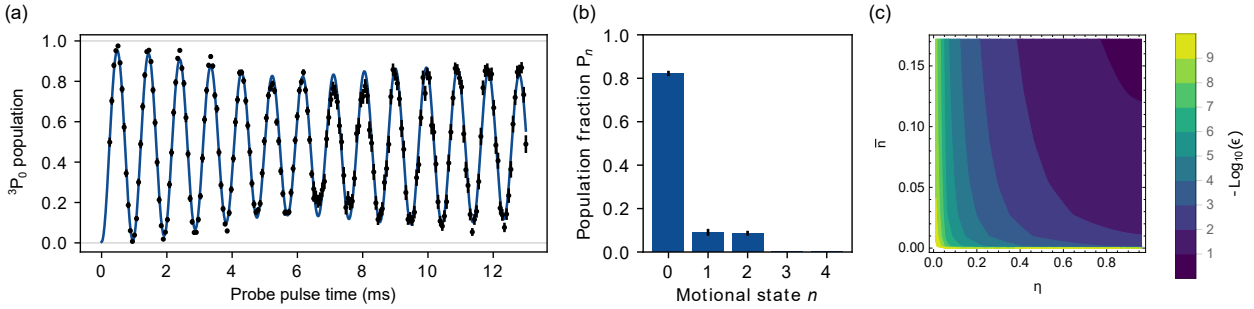

Figure 5.3: Effect of motion on Rabi oscillations. (a) Exaggerated example (with more heating than is present in a typical experiment) of how thermal motional excitations affect Rabi oscillations between the ground and clock electronic states. (b) One can fit the measured Rabi oscillations to extract the occupied motional states along the axis of the probe laser, in this case yielding a ground state fraction of ~ 0.8. (c) Predicted π-pulse error ϵ as a function of Lamb-Dicke parameter η and the average motional state occupation $\bar{n}$ (along the axis of the probe laser, and assuming a thermal distribution). Error bars in this figure correspond to 95% confidence intervals.

the atoms in comparison to the measured values (see Sec. 5.3). For $B = 20$ G, the equivalent shift is 9 mHz, and negligible compared to the measured coherence times.

It is also important to note that the presence of the probe light results in a differential AC Stark shift between the ground and clock states (see Sec. 2.2.2), leading to a shift of the qubit energy which must be accounted for when applying control pulses (Fig. 5.2c). This probe shift is only present when the probe beam is on, and so it has a negligible impact on clock measurements with a long dark time during which the atoms are allowed to evolve freely. However, it is important to control the homogeneity of the probe beam so that the gate applied to different atoms in a given ensemble is identical.

The final consideration for the fidelity of the applied single qubit gates is the motion of the atoms.[32] To calculate the expected clock π-pulse fidelity achievable given our known cooling performance and confinement, consider the motion-dependent Rabi frequency associated with the $s = 0$ or carrier transition in Eq. 4.8:

$$\Omega(n) = \Omega e^{-\eta^2/2} L_n^{(0)}(\eta^2). \tag{5.17}$$

We can average the evolution resulting from Eq. 5.17 over a thermal ensemble (of different values of n) given by the measured atom temperatures in our experiment to predict the optimal π-pulse fidelities achievable in our apparatus. These fidelities are shown for different combinations of cooling performance and confinement

[32] Note that, based on characterizations of the probe laser system (see Appendix D), we do not expect intensity or frequency noise to contribute significantly to errors in the applied single qubit gates.

in Fig. 5.3c, along with an example where significant occupation of motionally excited states results in dynamics predicted by Eq. 5.17.

In condition (2) [365], we use weakly confining magic wavelength tweezers with a typical depth of $15E_{\mathrm{R}}$, leading to a radial trap frequency of $f_{\mathrm{trap}} = 6.8$ kHz and $\eta = 0.83$. For measurements in these conditions, the cooling and handoff procedure led to an inferred average phonon occupation of $\langle n \rangle \simeq 0.14$ along the direction of the clock probe (see Sec. 4.3). Additionally, imaging and imperfect calibration of the handoff procedure led to $\sim 9\%$ additional losses in these measurements. Given these losses, the expected π-pulse fidelity is 0.81, which is consistent with the measured value of 0.82(2) without correcting for state preparation and measurement (SPAM) errors.

In condition (3) [290], we optimize for the fidelity of single qubit operations[33] by using more tightly confining potentials, and improving the quality of state preparation and readout. Specifically, we use a $420E_{\mathrm{R}}$ deep 2D lattice,[34] resulting in $\eta \simeq 0.22$ and $\bar{n} = 0.04^{+0.11}_{-0.04}$ in the direction of the clock probe. In this case, we expect to achieve a π-pulse fidelity of 0.9998^{+2}_{-9}. However, by fitting the observed clock Rabi oscillations with Eq. 5.17, we can extract the distribution of populated motional states during the single qubit gates (see Fig. 5.3b), and find that some additional heating occurs during the experimental sequence which degrades the initial state preparation (see Ch. 6 for details). This heating leads to an expected clock π-pulse fidelity of 0.9979, which is to be compared with the measured value of 0.9941(57) (with SPAM correction). In more recent (preliminary) measurements in the lattice with reduced heating and $\eta \simeq 0.18$, we have measured π-pulse fidelities as high as 0.999^{+1}_{-2} with SPAM correction, and 0.995(2) without.

5.2.3 Rabi spectroscopy

One way to compare the LO frequency to the atomic frequency is to perform Rabi spectroscopy, where one calibrates a resonant π-pulse as in the previous section, and then applies the same pulse to $|0\rangle$ atoms as a function of LO frequency before projectively measuring the atoms in the z basis. For square pulses in the

[33] At the cost of other performance metrics, as described in Sec. 5.3.

[34] And a $400E_{\mathrm{R}}$ deep axial lattice, but this is less relevant for the performance of the clock rotations, since the probe beam lies in the plane of the 2D lattice, and orthogonal to the axial lattice.

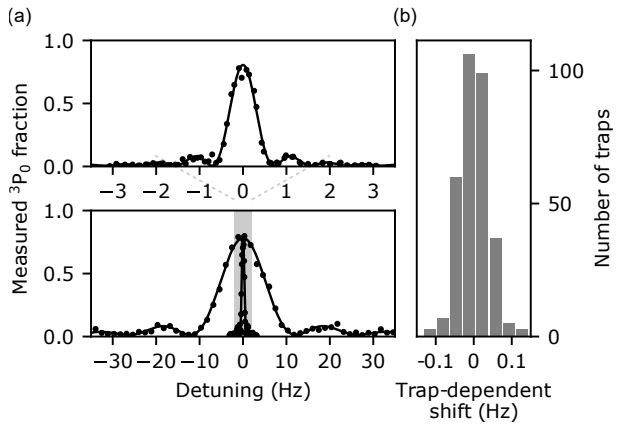

Figure 5.4: Clock Rabi spectroscopy. (a) Array-averaged Rabi spectroscopy of the clock transition in an array of 320 tweezers, with Fourier-limited linewidths of 10.1(2) Hz and 0.62(1) Hz (full width at half maximum). Solid lines correspond to the expected spectral feature given the known probe durations used in each case. Callout (top) shows the 0.6 Hz feature in detail. Note that the amplitude of the 0.6 Hz feature is not reduced in comparison to the 10 Hz feature, suggesting that this feature is not yet significantly affected by inhomogeneous broadening or other effects resulting in decoherence. (b) A histogram of the fitted transition frequency for different tweezers yields a standard deviation of 0.039(2) Hz, providing an estimate of the magnitude of inhomogeneous, tweezer-dependent shifts of the clock transition.

intensity of the probe beam of duration t_p, the resulting spectral feature is:

$$p_1 = \frac{\Omega^2}{\Omega^2 + \delta^2} \sin^2\left(\frac{t_p}{2}\sqrt{\Omega^2 + \delta^2}\right), \tag{5.18}$$

where p_1 refers to the probability of measuring the atoms in $|1\rangle$, and $\delta \coloneqq f - f_{\mathrm{LO}}$ is the detuning between the atomic reference and the LO. For lower Rabi frequencies and thus longer probe times, the spectral feature in Eq. 5.18 narrows, giving a better discriminator for fluctuations in δ. At very low Rabi frequencies, differential frequency noise between the LO and the atoms leads to a lineshape that is independent of Rabi frequency, and typically Gaussian. In condition (1), and for $B = 1.4$ G and $\Omega = 2\pi \times 125$ mHz, we reach a linewidth as low as 350(50) mHz, suggesting several seconds of coherence between the atoms and the LO (also see Sec. 5.2.4). Similar performance is obtained in condition (2), as shown in Fig. 5.4.

In a typical clock operating using Rabi spectroscopy, one probes the Rabi feature at its steepest point, where the excitation probability is $p_1 \simeq 0.54$. By additionally probing either side of the feature sequentially, one is able to decouple variations in $\langle z \rangle$ due to fluctuations in δ from variations due to other effects, like fluctuations in atom number or laser power. Such a scheme is often employed to simplify the control sequence in a clock measurement and to minimize probe shifts [242], but does not saturate the SQL (although it does follow the same scaling with respect to M and N).

5.2.4 Ramsey spectroscopy

An alternative approach to inferring δ from the atomic population is to perform a Ramsey measurement. In this case, one performs a $\pi/2$ rotation about the y axis of the Bloch sphere to prepare the state $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, then allows the state to evolve under the Hamiltonian $H = h\delta\sigma_z/2$ (in the rotating frame) for interrogation time t, and finally performs another $\pi/2$ rotation about the x axis before protectively measuring the atoms in the z basis (see Fig. 5.5a). This leads to a signal of $\langle \sigma_z \rangle = C \sin(2\pi\delta t)$, where C is the contrast of the Ramsey fringe (with $C = 1$ in the ideal case). Notice that $|+\rangle$ is precisely the state in Eq. 5.8 which we used to derive the SQL. In fact, for $\delta \sim 0$ and $C = 1$, where the slope of the signal is maximized, the estimator $2\pi\delta = \langle \sigma_z \rangle / t$ is locally unbiased, consistent, and efficient in the sense described in Sec. 5.1. As a result, perfect Ramsey measurements saturate the SQL [120].

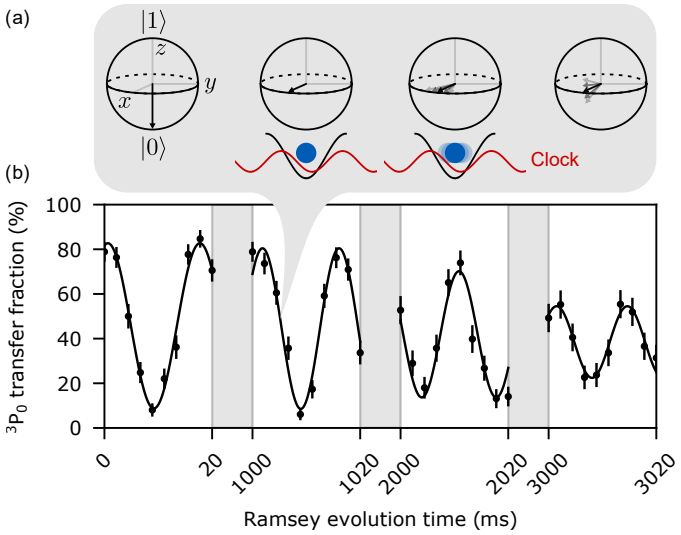

Figure 5.5: Clock Ramsey spectroscopy. (a) Under ideal conditions, the Ramsey sequence described in the text results in the atoms being on the equator of the Bloch sphere, and therefore maximum sensitivity to deviations in the phase of the clock laser. However, uncontrolled fluctuations unrelated to the laser itself, for example residual motion leading to random displacements of the atoms along the propagation axis of the laser, can also result in phase shifts. Averaging over these phase shifts results in a shortening of the Bloch vector, and therefore a reduction in the contrast of the Ramsey fringe. (b) Measurements of Ramsey oscillations as a function of free evolution time between the first and second Ramsey pulses. Note that in this case the laser is intentionally detuned slightly from the atomic resonance, leading to oscillations in time. At late times, the contrast of these oscillations decays due to the loss of coherence between the atoms and the laser.

In our experiments, Ramsey measurements performed as a function t show that the contrast of the Ramsey oscillations typically decay with a Gaussian time constant of 3.6(2) s [240, 365] (Fig. 5.5b).[35] This can be thought of as the amount of time over which the rotating frame defined by the LO remains coherent with the rotating frame of the atoms. Because our single qubit gates depend on the presence of the LO, this is also the coherence time against which one should compare the duration of the single qubit gates.[36]

In the following sections, we will see that the coherence time of the atoms relative to each other is significantly higher than the atom-LO coherence time, and so most of this dephasing can be associated with the LO (although at high magnetic fields, magnetic field noise and correlated fluctuations of the atomic reference frequency can also play a role). Our clock probe light comes from a laser that is injection locked with light stabilized to a cryogenic silicon reference cavity [242]. In principle, this laser system can reach remarkable linewidths of a few millihertz [218, 167]. However, the performance of this laser is degraded in the reference frame of the atoms, primarily due to mechanical vibrations that can lead to phase shifts in the applied laser light (see Sec. D.2.1). The clock laser system is described in more detail in Appendix D. For now, we will just point out one key challenge associated with tweezer array clocks: whereas in optical lattice clocks, the cavity mirror which defines the position of the optical lattice serves as a very good phase reference for the position of the atoms [34], a tweezer array lacks a similar reference of the atom positions, and so we instead resort to using the position of the microscope objective as a phase reference (see Sec. 3.3). Although this eliminates the effect of atomic motion due to motion of the objective itself, it does not, for example, eliminate motion due to air currents in the tweezer system.[37] We suspect that the atom-laser coherence is limited by residual motion of the tweezer array (or lattice), and hope to address this shortcoming in future upgrades (for example with a different lattice design, or more sophisticated active feedback of the tweezer positions).

[35] This decay time was measured with $B = 22$ G, but is not significantly affected by higher B.

[36] In Sec. 5.4.2 we will see that useful things can still be done even in the limit where the LO is completely dephased with respect to the atoms.

[37] Due to the bowtie design of our lattice, the final folding mirror could act as a good phase reference for the atom positions, but we do not currently take advantage of this in our experiments.

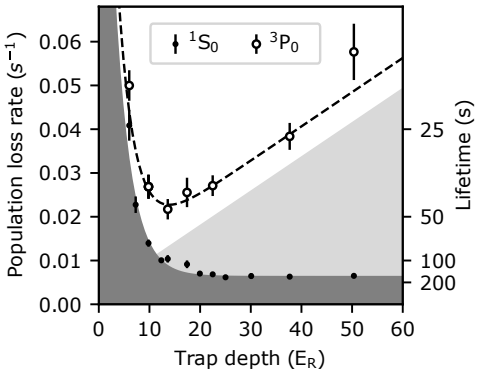

Figure 5.6: Ground and clock state lifetimes. To determine limits on atomic coherence, we measure the lifetime of both the ground ($^1S_0$, black points) and clock ($^3P_0$, open circles) states. For ground-state atoms the lifetime saturates to 160(10) s in deep traps, with additional technical sources of atom loss contributing in shallower traps (exponential fit to $^1S_0$ data, dark grey). For clock state atoms an optimal trap depth arises from a competition between depumping via spontaneous Raman scattering of the trap light (theory prediction with no free parameters, light grey) and atom loss. The combination of the two loss mechanisms (dashed line) is in good agreement with the measured lifetimes, including with the optimal value of 46(5) s at a depth of $14E_R$.

5.3 Preserving coherence and minimizing frequency shifts

So far, we have discussed how to manipulate and read out individual clock qubits. What remains is to engineer an environment that minimizes systematic shifts of the qubit transition frequency, and preserves the coherence of the qubits to enable long interrogation times.

5.3.1 Atom loss

One significant limitation to the atomic coherence time is simply the fact that atoms can be lost from the optical traps. We expect the lifetime of an atom in an optical trap to be limited by collisions with residual background gas. Assuming that the main collision partners are room temperature Σ state H_2 molecules interacting with the atoms via van der Waals forces, and following the procedure in [341], we expect that such collisions are substantially more energetic than the trap depths we have access to. As a result, we expect the vacuum lifetime to effectively be independent of trap depth [16]. For trap depths in excess of $\sim 20E_R$, we see this borne out in measurements of the lifetime of $^1S_0$ ground state atoms (see

Fig. 5.6), where we measure a trap-depth-independent lifetime of $\tau = 160(10)$ s (see Fig. 5.6).[38] More recent measurements where we operate with a lower oven temperature of 380 C instead of 420 C improve the quality of our vacuum, and yield an improved collision-limited lifetime of ~ 400 s.

Similar collision processes apply to atoms that are in the clock state, however atoms in electronically excited states have larger C_6 coefficients, and thus larger scattering cross sections $\sigma \propto C_6^{-2/5}$. The result is a reduced trap lifetime of $\tau \propto 1/\sigma$. With the known C_6 coefficients for collisions between H_2 and ^{88}Sr [224], we calculate that the ratio between the ground and clock-state trap lifetime (τ_g and τ_e respectively) is $\tau_g/\tau_e = 1.10$, which agrees with results from [93].

The source of the dramatic reduction of trap lifetime in shallow traps with depths of $\lesssim 20E_R$ (see Fig. 5.6) is not yet fully understood. However, based on the above analysis, we rule out the effect of collisions with background gas. One possible mechanism for the observed loss is parametric heating related to fluctuations in the trapping potential. There are two main mechanisms for such heating: intensity noise in the trapping laser manifests as parametric modulation of the trap frequency which, assuming a flat power spectrum in intensity noise, results in exponential heating (measured in phonon number) with a time constant proportional to f_t^2, where f_t is the trap frequency [287, 115]. Similarly, pointing noise with a flat power spectrum (in terms of fluctuations in the center position of the trap) results in linear heating with a rate proportional to f_t^3 (measured in phonon number) [287, 115]. For comparison, the number of bound states, N, in an optical trap scales roughly like $N \propto \sqrt{V} \propto f_t$, where V is the trap depth. This means that both of these sources of loss should improve with reduced trap depth assuming a flat noise spectrum. The intensity noise of our trapping laser is suppressed below 10 kHz via a servo, and otherwise relatively flat over the frequencies relevant for heating. As a result, we do not expect heating due to intensity noise to be the source of the observed loss in shallow traps. However, there could be increased pointing noise at lower frequencies due to mechanical resonances and acoustic noise, and there is no convenient way of removing such noise with a servo. We therefore hypothesize that pointing noise contributes to our reduced lifetimes at and below trap depths of $15E_R$ (corresponding to 6.8 kHz trap frequencies). Other sources of trap-independent heating, like scattered background light, can also begin to dominate as the traps become very shallow and N becomes

[38] Given this lifetime, we estimate that the fractional frequency shifts due to the background collisions [116, 34] are below the 10^{-19} level, and thus negligible in our experiments.

Measured lifetimes (s)	
$1/\Gamma_g^t$	101(6)
$1/\Gamma_e^t$	92(5)*
$1/\Gamma_e$	43(4)
Expected rates ($\times 10^{-3}$ s^{-1}) [93]	
Γ_0^{BBR} ($^3\mathrm{P}_0 \rightarrow {}^3\mathrm{P}_0$)	3.45(22)
Γ_1^{BBR} ($^3\mathrm{P}_0 \rightarrow {}^3\mathrm{P}_1$)	2.23(14)
Γ_2^{BBR} ($^3\mathrm{P}_0 \rightarrow {}^3\mathrm{P}_2$)	0.105(7)
Γ_0^{R} ($^3\mathrm{P}_0 \rightarrow {}^3\mathrm{P}_0$)	0.557 (V/E$_\mathrm{R}$)
Γ_{12}^{R} ($^3\mathrm{P}_0 \rightarrow {}^3\mathrm{P}_{1,2}$)	0.782 (V/E$_\mathrm{R}$)

Table 5.1: Measured and theoretical values contributing to predicted Ramsey lifetime. All measured values are for a trap depth of 15 E_R, based on interpolating between the nearest points in Fig. 5.6. * indicates that the inferred value of Γ_e^t is dependent on the calculations presented in Sec. 5.3.1. Note that Γ_2^{BBR} is smaller than the error bars on the other processes, and so we neglect this process in our analysis. V is the trap depth and E_R is the recoil energy of an 813 nm photon.

small.

5.3.2 Clock state lifetime and coherence

Besides being lost from the traps, atoms in the $^3\mathrm{P}_0$ clock state can also undergo transitions to other electronic states. Measurements of the $^3\mathrm{P}_0$ state lifetime as a function of trap depth are shown in Fig. 5.6. Note that due to our use of ^{88}Sr, and given the strength of the magnetic fields used, the effects of spontaneous emission from $^3\mathrm{P}_0$ to $^1\mathrm{S}_0$ are negligible. For typical experimental conditions,[39] the dominant process leading to transitions out of $^3\mathrm{P}_0$ is Raman scattering of the trap light, with the main contributions being from $^3\mathrm{P}_0 \rightarrow {}^3\mathrm{P}_1$ and $^3\mathrm{P}_0 \rightarrow {}^3\mathrm{P}_2$. For π-polarized trap light, these two processes occur with rates of 4.98×10^{-4} s$^{-1}E_\mathrm{R}^{-1}$ and 2.84×10^{-4} s$^{-1}E_\mathrm{R}^{-1}$ respectively [93]. Note that while the ratio of these two scattering processes is polarization-dependent, their sum, with a value of $\Gamma_{12}^{\mathrm{R}} = 7.82 \times 10^{-4}$ s$^{-1}E_\mathrm{R}^{-1}$, is conserved. All population driven into $^3\mathrm{P}_1$ can be assumed to immediately decay into the ground state, whereas processes that return population in $^3\mathrm{P}_2$ to the $^3\mathrm{P}_0$ state are negligible. As a result, $V\Gamma_{12}^{\mathrm{R}}$ can be treated as the total rate at which population in $^3\mathrm{P}_0$ is depleted due to Raman scattering of the trap light.

Another mechanism that depopulates the $^3\mathrm{P}_0$ state is scattering of black-body radiation (BBR). The dominant BBR-induced decay at room temperature is due to off-resonant driving from $^3\mathrm{P}_0$ to the $^3\mathrm{D}_1$ state,

[39] Meaning magnetic fields of $B \ll 1$ T, and trap depths of $V \gtrsim 10E_\mathrm{R}$.

which decays to the $^3\mathrm{P_J}$ manifold with branching ratios of $R_J^D = 59.65\%$, 38.52%, and 1.82% for the $J = 0,\ 1$ and 2 states respectively [93]. The dominant mechanism by which BBR contributes to loss from the $^3\mathrm{P}_0$ state is via population that branches from $^3\mathrm{D}_1$ into $^3\mathrm{P}_1$, and subsequently decays into the $^1\mathrm{S}_0$ ground state. This process occurs at a rate of $\Gamma_1^{\mathrm{BBR}} = 2.23 \times 10^{-3}\ \mathrm{s}^{-1}$ [93] at room temperature. The sum of the above effects with the rate at which $^3\mathrm{P}_0$ state atoms are lost from tweezers with depth V, $\Gamma_e^t(V)$, is in good agreement with the measured $^3\mathrm{P}_0$ decay rate, $\Gamma_e(V) = \Gamma_e^t(V) + V\Gamma_{12}^{\mathrm{R}} + \Gamma_1^{\mathrm{BBR}}$ (theory curve in Fig. 5.6).

Measurements of the $^3\mathrm{P}_0$ state lifetime can be confounded by the presence of atoms pumped into the $^3\mathrm{P}_2$ state, which are not distinguished from $^3\mathrm{P}_0$-state atoms during our normal blow-away measurement (see Sec. 5.2.1), and can lead to an artificially long inferred lifetime. To avoid this systematic error, we perform two separate measurements. In one measurement, we add a repumping step that depletes $^3\mathrm{P}_2$ atoms before the blow-away by driving the $^3\mathrm{P}_2 \leftrightarrow {}^3\mathrm{S}_1$ transition at 707 nm (see Sec. 2.1.1.4). In the second measurement, we remove the repumping step to measure the total $^3\mathrm{P}_0 + {}^3\mathrm{P}_2$ population. Based on these two measurements and the branching ratios of $^3\mathrm{S}_1$ we infer the true population in $^3\mathrm{P}_0$, which appears in Fig. 5.6.

Given the decay processes described above, we can compute an expected Ramsey coherence time τ.[40] Due to the magic-wavelength traps, Rayleigh scattering of the trap light does not cause decoherence [337, 214, 93].[41] As a result, trap-induced scattering only contributes to decay of Ramsey contrast through the Raman scattering processes described above that depopulate the $^3\mathrm{P}_0$ state. Unlike Rayleigh scattering of the trap light, BBR-driven processes that transfer population out of and back into the $^3\mathrm{P}_0$ state (predominantly via $^3\mathrm{D}_1$) can serve as an extra source of decoherence that is not directly reflected in the measurement of the $^3\mathrm{P}_0$ state lifetime.[42] Including all the above effects, the inferred Ramsey coherence time is:

$$\tau = 2/(\Gamma_e^t + \Gamma_g^t + V\Gamma_{12}^{\mathrm{R}} + \Gamma_1^{\mathrm{BBR}}(1 + \frac{R_0^D}{R_1^D})), \tag{5.19}$$

where $\Gamma_{g,e}^t$ are the vacuum loss rates of ground and clock state atoms respectively. All the relevant rates in Eq. 5.19 are summarized in Tab. 5.1.

[40] Defined as the interrogation time at which the contrast of the Ramsey fringe decays to $1/e$ of its zero-time value, see Sec. 5.2.4.

[41] One can see this schematically from the fact that it is not possible to identify the electronic state of the atom based on the properties of the scattered photons.

[42] In the case of the weak drive provided by BBR, the dominant process is one in which the atom is transferred to the excited state, and then spontaneously emits a photon. Here, even if the atom returns to the $^3\mathrm{P}_0$ state, it is dephased with respect to an atom that did not scatter a BBR photon. This is to be compared to the Raman process driven by the trapping light, where the strong coherent drive leads to coherent absorption and stimulated emission into the trapping beam.

Given the above analysis and measurements of the trap lifetime in Sec. 5.3.1, the optimal trap depth for performing a Ramsey measurement is approximately $14E_{\mathrm{R}}$. Here, the measured ground state trap lifetime is 96(8) s, leading to a predicted clock state lifetime of 44(6) s. The measured clock state lifetime is 46(5) s, in good agreement with the predicted value. Under these conditions, the Ramsey contrast should decay exponentially with a time constant of $\tau = 55(8)$ s. We will see in Sec. 5.4.2.2 that the measured atomic coherence time is in agreement with the predicted Ramsey coherence time.

5.3.3 Tunneling-induced phase noise

The lifetimes measured in Sec. 5.3.2 are encouraging, but when operating at low trap depths and for long evolution times one must be wary of tunneling or thermal hopping, where atoms are able to move between different sites in the array of optical traps. Because the spatial phase of the LO can differ from site to site, such hopping effectively results in fluctuations in the phase of the LO as experienced by the atoms. A simple way to mitigate this source of phase noise is to eliminate tunneling through appropriate control of the trapping potential [148]. In our system, we can achieve this without adjusting the trap depth by controlling the spacing between tweezers to reduce the tunneling energy. Tunneling is further suppressed by the inherent $\sim 10^{-2}$ level disorder in the depths of the tweezers (see Sec. 3.3.3).[43]

For pairs of images in a given experimental trial, we expect tunneling and thermal hopping to manifest as correlated x_{01} and x_{10} events (see Sec. 4.4.1) on nearest-neighbor sites, where an atom tunnels from one site to an empty adjacent site between images, or pairs of x_{10} events, where an atom tunnels onto an occupied adjacent site after the first image, and both atoms are lost due to light-assisted collisions during the second image (see Sec. 4.2). For tweezers with a spacing of 1.2 μm and depths in excess of $V = 6E_{\mathrm{R}}$, we do not observe an excess of such events beyond what is expected due to loss and imaging infidelity (for hold times of up to hundreds of seconds). The lack of tunneling at low trap depths allows us to operate at the nearly-optimal depth of $V = 15E_{\mathrm{R}}$ to maximize the lifetime of the $^3\mathrm{P}_0$ state without concerns of introducing substantial tunneling and thus LO phase noise.

The ability to operate at very low trap depths for clock operation in a tweezer array is in contrast

[43] See Ch. 7 for a more in depth discussion of tunneling.

to a typical retro-reflected lattice, where the lattice spacing is not independently tunable from the trapping wavelength. In the 3D lattice clock at JILA, tunneling can limit the achievable interrogation times at depths below $\sim 30E_{\mathrm{R}}$ along a single axis (and a combined depth of $\sim 100E_{\mathrm{R}}$ for all three axes) without additional mitigation [148]. However, there are several approaches to eliminating LO phase noise due to tunneling in an optical lattice. For example, one can tune the lattice spacing using other lattice geometries to either eliminate tunneling, or to make the array spacing commensurate with the spatial phase of the LO [148]. Appropriately orienting the lattice with respect to gravity can also eliminate tunneling by means of applying a spatial gradient across the lattice [187, 35, 373, 3]. If the tunneling is suppressed in one direction, for example along the gravitational axis, one can also orient the LO beam such that the spatial phase of the LO is equal for all sites that are tunnel-coupled to each other, resulting in a fixed LO phase even in the presence of tunneling. Finally, in hybrid systems like ours that integrate optical lattices with high NA microscopes, one could use the microscope to project a disorder potential onto a shallow lattice to eliminate tunneling.

5.3.4 Frequency shifts

While a full accuracy evaluation is beyond the scope of this work, the prospects for using the tweezer platform as an absolute frequency reference are fairly good: many of the dominant systematic effects in tweezer clocks are shared with optical lattice clocks [240], which currently have accuracies at or below the 10^{-18} level [31, 234, 34]. However, there are a couple extra considerations that one must be careful of in a tweezer clock.

5.3.4.1 Variations in trapping wavelength

One consideration is the fact that our tweezer array is generated using AODs (see Sec. 3.3.3), leading to different tweezers being detuned from each other, preventing us from attaining magic trapping conditions across the entire array. Although these megahertz-scale detunings are small compared to the hundreds of terahertz-detunings of the trapping beams from resonance, they are still sufficient to result in measurable differential frequency shifts between atoms in different traps.

For deep traps ranging between $V = 450 - 1800E_{\mathrm{R}}$ in condition (1), we perform Ramsey spectroscopy

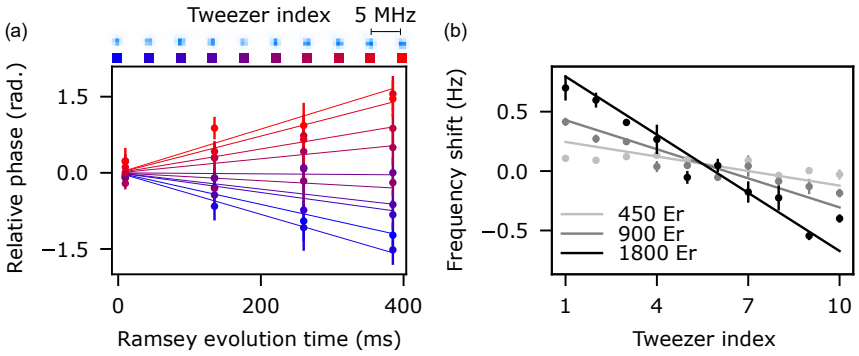

Figure 5.7: Trap-induced clock frequency shifts. (a) Measured relative Ramsey phase shifts for tweezers at different positions (and therefore frequencies), and at a depth of $1800E_{\mathrm{R}}$. (b) Inferred relative frequency shifts as a function of tweezer position at a range of tweezer depths. Solid lines are predictions based on the known sensitivity of the magic wavelength to trap detuning.

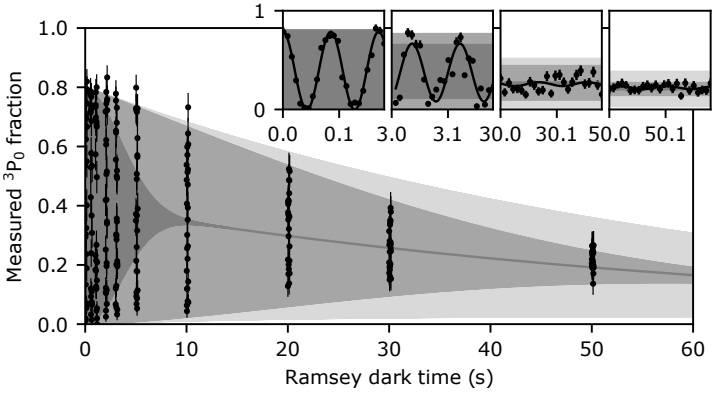

Figure 5.8: Ramsey spectroscopy in large arrays. We perform Ramsey measurements (black points) in an array of 320 tweezers, each with a depth of $15E_{\mathrm{R}}$, which is nearly optimal for preserving atomic coherence. The measured clock-state lifetime (see Fig. 5.6) contributes to a decay of the Ramsey contrast with an expected exponential time constant of 55(8) s (light-grey region). However, tweezer-dependent light shifts further increase site-to-site dephasing, leading to Gaussian decay with an expected time constant of 33(1) s (at $15E_{\mathrm{R}}$, and for the array geometry used in this measurement). The combination of the above two effects is denoted by the medium-grey region. Note that each data point corresponds to a single shot of the experiment. As a result, despite the fact that the atom-laser coherence decays with a Gaussian time constant of 3.6(2) s (dark-grey region, see also Fig. 5.5), the variance of the Ramsey signal decays on a timescale set by atomic coherence. This is clarified by the insets, which share units with the main axes, and show detailed views of the Ramsey fringe at a few different free evolution times. Here, it is possible to see the initial loss of phase coherence with the laser followed, at later times, by total loss of coherence (and thus variance in this signal) in the system.

to measure the frequency shifts imposed by the tweezers (see Fig. 5.7). The results are in line with the expectation of the differential polarizability of the clock states near the 813 nm magic wavelength for Sr (see Sec. 5.3.4.2).

In condition (2), we operate with shallow trap depths of $15E_{\mathrm{R}}$ and with large atom arrays with tweezer frequencies differing by up to ~ 60 MHz (at opposite corners of the array). In our measurements under these conditions, the signal at each Ramsey time is a single-shot measurement such that even though the atom-laser coherence decays over ~ 3 s (see Sec. 5.2.4), we can infer the coherence of the atoms relative to each other from the variance of the Ramsey signal, which remains high on much longer timescales (Fig. 5.8, see also Sec. 5.4.2.2). At short times, the frequency of the Ramsey fringes is set by the differential light shift imposed by the probe beam on the $^1\mathrm{S}_0$ and $^3\mathrm{P}_0$ states. At longer times, the loss of atom-laser coherence manifests as a randomized phase of the second $\pi/2$ pulse in the Ramsey sequence, which obscures the Ramsey oscillations but preserves the probability of large excursions due to the persistence of atomic coherence (see Sec. 5.4.2.1).

The Ramsey contrast inferred from this measurement decays with a $1/e$ time of 19.5(8) s (Fig. 5.8), corresponding to an effective quality factor of $Q = 1.9(1) \times 10^{16}$. The measured decay time is slightly faster than the prediction of 33(1) s based on the measured lifetime, and on the expected dephasing due to the differential AC Stark shifts imposed by the tweezers. Note, however, that dephasing due to differential AC Stark shifts is static (assuming the optical traps are relatively stable and do not fluctuate in intensity or frequency), and thus can easily be eliminated with a spin echo. Additionally, alternative approaches to making optical traps, like tweezer arrays based on spatial light modulators or digital micromirror devices, or accordion lattices, can avoid this issue entirely while still maintaining additional control over the exact shape and geometry of the optical potential.

5.3.4.2 Hyperpolarizability and trap geometry

To accurately determine the energy shift imposed on the clock transition due to the optical traps, one must consider modifications to the low order expressions appearing in Sec. 2.2.2. This is particularly relevant in the case of a tightly confining tweezer array, where the intensity and polarization of the trapping field can vary rapidly in space [164, 323], leading to modified expressions for the light shift. The complications associated with higher order corrections to the trapping potential are discussed in detail in [248, 159, 49] in the context of optical lattices, and applied to tweezer arrays in [209]. The typical approach is to compute the shift in energy of the ground and clock states due to the electric dipole $\alpha^{(\mathrm{E1})}$,[44] magnetic dipole $\alpha^{(\mathrm{M1})}$, and electric quadrupole $\alpha^{(\mathrm{E2})}$ polarizabilities, as well as the hyperpolarizability β. The resulting optical potential is:[45]

$$V(\xi, f, \boldsymbol{r}) = -\left(\alpha^{(\mathrm{E1})}(\xi, f) + \alpha^{(\mathrm{M1})}(\xi, f) + \alpha^{(\mathrm{E2})}(\xi, f)\right) I(\boldsymbol{r}) - \beta(\xi, f) I(\boldsymbol{r})^2, \tag{5.20}$$

where f is the frequency of the laser field, and $\xi \in \{e, g\}$ denotes the electronic state of the atom.[46]

For the purposes of understanding the observed frequency shifts between tweezers, we are interested in isolating the lowest order terms in Eq. 5.20 that contribute to a frequency shift $f_{\mathrm{LS}} = \left(V(e, f, \boldsymbol{r}) - V(g, f, \boldsymbol{r})\right)/h$

[44] The discussion in Sec. 2.2.2 amounts to only handling the $\alpha^{(\mathrm{E1})}$ contribution to the energy shift.

[45] This form assumes a travelling wave, but in general one must be careful about a spatial phase shift between the $\alpha^{(\mathrm{E1})}$, $\alpha^{(\mathrm{M1})}$, and $\alpha^{(\mathrm{E2})}$ contributions both for lattice potentials, and very tightly focused tweezer potentials.

[46] In comparison to Sec. 2.2.2, we have taken $\beta \to \xi$ to avoid confusion with the hyperpolarizability, and are working with frequencies f rather than angular frequencies ω.

of the clock transition, and depend on the detuning $\delta_{\mathrm{m}} = f - f_{\mathrm{m}}$ of the trapping laser from the magic trapping frequency $f_{\mathrm{m}} \simeq 369$ THz. We can solve for f_{LS} in a tweezer potential by Taylor expanding Eq. 5.20 (with I as defined in Eq. 3.2, and with peak intensity I_0) up to fourth order in the longitudinal coordinate z, and in the radial coordinate ρ, and treating the quartic terms as a perturbation to the harmonic (quadratic) solutions. The harmonic solutions can be enumerated by $n_{x,y,z}$, namely the harmonic oscillator state in three orthogonal (Cartesian) directions. Keeping only terms that are up to linear in I, the above approach leads to an approximate expression for the shift f_{LS} in the frequency of the clock transition of [290]:

$$\begin{aligned}\frac{h f_{\mathrm{LS}}}{E_{\mathrm{R}}} &\simeq \left(\sqrt{2}\frac{w_0}{z_{\mathrm{R}}}(n_x + n_y + 1) + \left(\frac{w_0}{z_{\mathrm{R}}}\right)^2 \left(n_z + \frac{1}{2}\right)\right)(\sqrt{u_e} - \sqrt{u_g}) - (u_e - u_g) \\ &\simeq \left[\frac{1}{2}\left(\sqrt{2}\frac{w_0}{z_{\mathrm{R}}}(n_x + n_y + 1) + \left(\frac{w_0}{z_{\mathrm{R}}}\right)^2 (n_z + \frac{1}{2})\right)\sqrt{u_{\mathrm{m}}} - u_{\mathrm{m}}\right] \frac{1}{\alpha^{(\mathrm{E1})}} \frac{\partial \Delta\alpha^{(\mathrm{E1})}}{\partial f}\Big|_{f=f_{\mathrm{m}}} \delta_{\mathrm{m}},\end{aligned} \tag{5.21}$$

where w_0 and z_{R} are the minimum beam waist and Rayleigh range of the tweezer, and $u_\xi = \alpha^{(\mathrm{E1})}(\xi, f) I_0 / E_{\mathrm{R}}$. In the second line, we Taylor expand about the magic wavelength ($\delta_{\mathrm{m}} = 0$) keeping only the first order terms in δ_{m}. Here, we have defined that $u_{\mathrm{m}} \coloneqq u_e = u_g$ and $\alpha^{(\mathrm{E1})} \coloneqq \alpha^{(\mathrm{E1})}(e) = \alpha^{(\mathrm{E1})}(g)$ at $\delta_{\mathrm{m}} = 0$. $\frac{1}{\alpha^{(\mathrm{E1})}} \frac{\partial \Delta\alpha^{(\mathrm{E1})}}{\partial f}\big|_{f=f_{\mathrm{m}}} = -15.5(1.1)$ μHz/MHz [222, 307] is the normalized derivative of the difference in the electric dipole polarizabilities of the e and g states, $\Delta\alpha^{(\mathrm{E1})} \coloneqq \alpha^{(\mathrm{E1})}(e) - \alpha^{(\mathrm{E1})}(g)$, with respect to the trapping frequency f and evaluated at the magic trapping frequency f_{m}.

The above analysis illuminates the fact that the motional state of an atom determines how much of the optical potential an atom samples, and thus how sensitive the clock transition in the atom is to imperfect magic trapping conditions. Minimizing the effect of these shifts in the clock transition frequency is yet another reason it is beneficial to cool the atoms to close to their 3D motional ground state (see Sec. 4.3).[47]

Although so far the tweezer potential seems to be an additional source of inaccuracy, there are situations where it could also be a benefit. For example, the value of β is not particularly well-characterized for Sr [263]. The ability to independently vary the depth of the traps opens up the possibility of performing measurements of β (and the various contributions to the polarizability) via synchronous comparisons that are insensitive to LO noise (see Sec. 5.4.2). Such measurements are, in principle, only limited by QPN (and

[47] Note that, unlike the discussion in Sec. 5.2 where only the motional state along the direction of the probe laser was important, the energy shift in Eq. 5.21 depends on all aspects of the atomic motion in 3D. This is, in part, why we opt to perform high fidelity optical cooling in the axial direction by means of employing an accordion lattice in addition to the tweezers in [365].

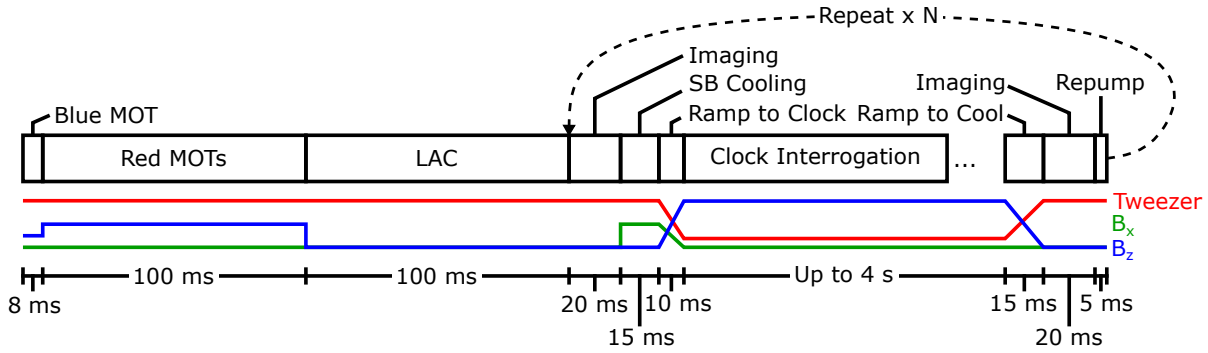

Figure 5.9: Scheme for repeated clock interrogation. After preparing an ensemble of individual atoms in tweezers using several stages of MOTs and light assisted collisions (LACs), the atoms are repeatedly imaged, sideband (SB) cooled, interrogated, and then recycled via repumping back to the ground state. Note that the use of deep $(3900 - 6000 E_{\mathrm{R}})$ tweezers allows for faster images than are typical for the other experiments in this work. The lower traces denote the depth of the tweezers, and the magnitude of the magnetic fields B_x and B_z applied along the x and z axes, during various phases of the experiment. In particular, note that the magnetic field conditions for sideband cooling and for clock interrogation are distinct.

your ability to calibrate relative intensities) and thus can reach arbitrary precision with sufficient averaging time. Based on practical considerations the prospects are good for characterizing these effects at the 10^{-19} level, which would further reduce any inaccuracies imposed by the tweezers.

5.4 Clock measurements

With an understanding of the limitations and capabilities of our system, we are interested in performing precision measurements of the clock transition frequency in tweezer-trapped ^{88}Sr atoms.

5.4.1 Repeated interrogation

In frequency comparisons between a laser and the atoms, we are sensitive to phase noise in the laser due to the Dick effect (see Sec. 5.1.3.2). However, we can minimize Dick effect noise by reusing the same atoms for multiple rounds of interrogation,[48] increasing the duty cycle of the experiment.

For measurements taken in condition (1), our scheme for reusing the atoms takes advantage of the lossless detection described in Sec. 5.2.1, and is shown in Fig. 5.9 [240]. Each interrogation cycle involves 85 ms of dead time during which the atoms are measured and reinitialized, whereas it takes approximately

[48] Similar to ion clocks [203].

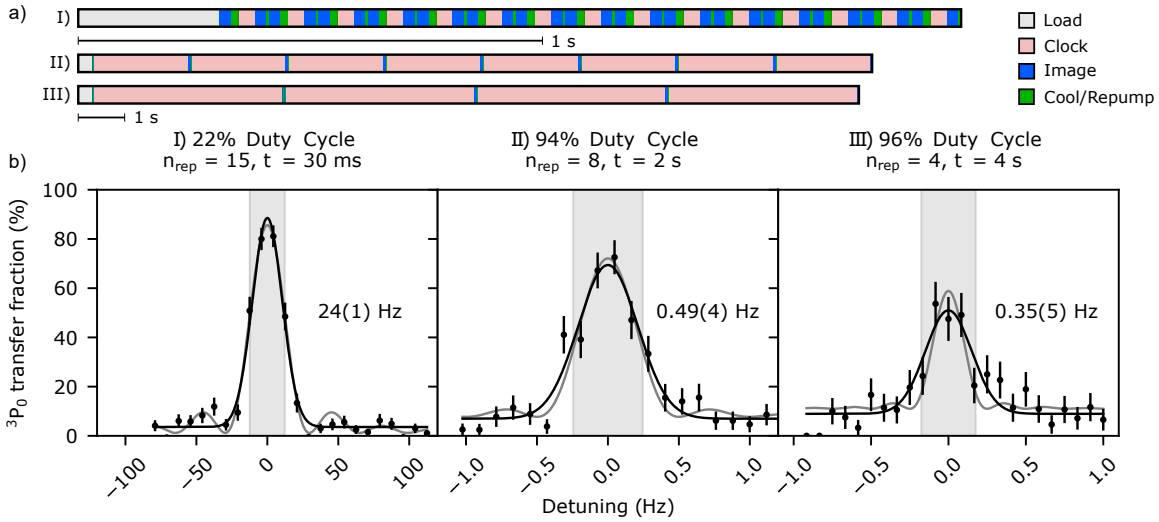

Figure 5.10: Repeated Rabi spectroscopy. (a) Timing diagram for repeated interrogation of the clock transition for various clock interrogation times (t) and number of repeated interrogations (n_{rep}): I) $n_{\mathrm{rep}} = 15$, $t = 30$ ms; II) $n_{\mathrm{rep}} = 8$, $t = 2$ s; and III) $n_{\mathrm{rep}} = 4$, $t = 4$ s. These values of n_{rep} are chosen based on experimental convenience — high repetition numbers have a more significant impact on duty cycle for short interrogation times, but due to stochastic loading require lengthy data scans for long probe times to achieve uniform statistical uncertainty. (b) Corresponding Rabi spectra for sequences I), II), and III). Sinc fits shown in gray, and Gaussian fits in black. Shaded regions and labels correspond to the Gaussian full-width-at-half-maximum.

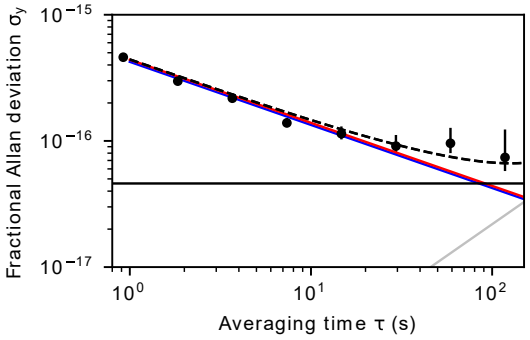

Figure 5.11: Allan deviation of fractional frequency fluctuations between an ensemble of atoms and an ultrastable laser. For Ramsey interrogation with a free evolution time of $t = 501$ ms, we find a short-time frequency instability of $4.7 \times 10^{-16}(T/\mathrm{s})^{-1/2}$, where T is the averaging time (red line). This instability is consistent with the limit imposed by QPN when including the imperfect contrast of the Ramsey fringe (blue line). The long-term stability of this measurement is consistent with the known flicker noise (black) and drift (grey) of the laser (the combined effect of QPN, flicker noise, and drift is indicated by the dashed line).

300 ms to load a new ensemble of atoms into the tweezers. Each cycle of measurement and reinitialization results in a probability of 0.001(1) that a given atom is lost, in addition to losses associated with the vacuum lifetime (see Sec. 5.3.1), which can play a role for very long interrogation times. Practically, these low losses mean that we can reuse the same ensemble of atoms enough times that the duty cycle of the experiment is set predominantly by the dead time during each interrogation cycle, with the dead time due to initial preparation being negligible. Performing repeated Rabi spectroscopy for interrogation times as long as $t = 4$ s allows us to measure spectral features with a linewidth as narrow as 0.35(5) Hz, and with a duty cycle as high as 96% (see Fig. 5.10). Note that in these repeated interrogations, we do not find that the performance of later interrogation cycles is degraded with respect to earlier cycles, suggesting that the only factor limiting the reuse of the atomic ensemble is loss.[49]

We characterize the short-term frequency instability of the atoms by performing Ramsey measurements (see Sec. 5.2.4) with an interrogation time of $t = 501$ ms.[50] At this evolution time, the contrast of the Ramsey fringe is $C = 0.79(3)$, meaning that dephasing between the atoms and the LO does not yet significantly impact the measurement. Measurements of the state of the atoms after the Ramsey sequence corresponds to a measurement of σ_z, with an expectation value of $\langle \sigma_z \rangle = C \sin(2\pi\delta t)$ where δ is the de-

[49] And limited memory in the experimental control system.
[50] Note that in these measurements we did not opt to use repeated interrogation since, even when preparing a new ensemble of atoms on each measurement cycle, Dick effect noise was negligible in comparison to QPN.

tuning of the atomic transition frequency from its mean value on a given experimental trial. Using the single-shot estimates of δ from the experiment,[51] we can compute the Allan deviation (Eq. 5.14) relative to the free-running clock laser, yielding an instability of $4.7 \times 10^{-16}(T/\mathrm{s})^{-1/2}$ at short averaging times T (see Fig. 5.11). This instability is entirely dominated by QPN associated with the average of 4.8 atoms loaded on each experimental trial,[52] with the performance at long averaging times dominated by the known drift and flicker noise floor of the laser [242]. Due to the relatively high duty cycle (and QPN) in these measurements, Dick effect noise contributes to less than 1% of the measured instability and is therefore negligible.

5.4.2 Measurements beyond the atom-laser coherence time

In Sec. 5.1.3 we considered canonical optical clock measurements where a laser serves the role of an LO whose frequency is compared to the atomic reference. In Sec. 5.4.1, we saw how the interrogation time and therefore the instability in such a measurement was limited by the finite coherence time between the atoms and the laser. However, in certain situations one can compare one ensemble of atoms to another in a manner that is insensitive to laser noise, and thus extend the interrogation time to durations that are set by the coherence time between atoms, rather than between the atoms and a laser.[53] A number of related techniques fall under this category, including synchronous clock comparisons [28, 203] and correlation spectroscopy [69, 68, 147, 70, 173].

5.4.2.1 Ellipse fitting

One illustrative example is to extract the relative phase ϕ between two atomic ensembles a and b when the laser phase ϕ_L is unknown, but equal for both ensembles. Consider a simplified model for the probabilities $p^{(a,b)}$ that an atom in a given ensemble is detected in the $|1\rangle$ state after performing a Ramsey

[51] We assume that δ is small and that the conversion between $\langle\sigma_z\rangle$ and δ is linear, and terminate the experiment before this assumption fails due to uncorrected drifts of the LO frequency.

[52] Based on Monte-Carlo simulations, we also include a correction factor of 1.07 to the QPN, which handles excess noise due to fluctuations in the number of atoms loaded on each trial [240].

[53] As we saw in Sec. 5.3.2, we expect the former to be much longer than the latter in our experiments.

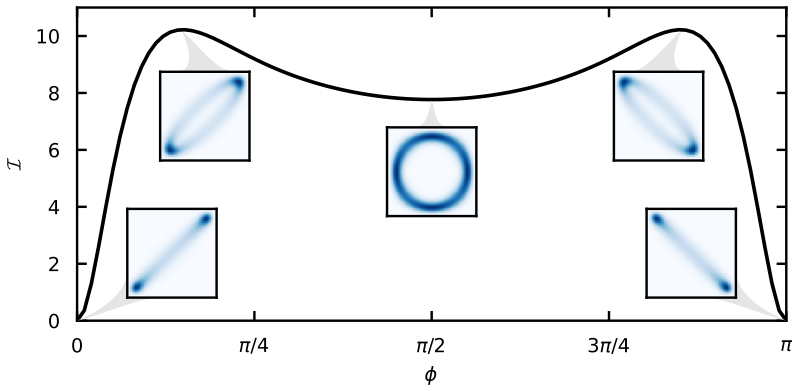

Figure 5.12: Ellipse fitting. Fisher information $\mathcal{I}$ for measuring the relative phase ϕ between two 60 atom ensembles with a Ramsey contrast of 0.8, when marginalizing over the global phase provided by a reference laser. Note that $\mathcal{I}$ vanishes at $\phi = 0$ and π. Insets depict the probability mass function for parametric plots of the excitation fraction of the two ensembles. Color scale is arbitrary, and for illustration purposes only.

measurement:[54]

$$p^{(a)} = \frac{C}{2}(1 + \cos(\phi_L)) \tag{5.22}$$

$$p^{(b)} = \frac{C}{2}(1 + \cos(\phi + \phi_L)), \tag{5.23}$$

where C is the Ramsey contrast. Notice that although one cannot gain any information about ϕ from $p^{(a)}$ or $p^{(b)}$ alone after averaging over all values of ϕ_L, this is not true if one has access to both measurements. For example, the square of $p^{(a)} + p^{(b)}$ averaged over ϕ_L is $\int_0^{2\pi} d\phi_L (p^{(a)} + p^{(b)})^2/(2\pi) = 1 + \cos(\phi)$, which clearly contains information about ϕ.

On each experimental trial, we measure values $x^{(a)}$ and $x^{(b)}$, which correspond to the fraction of atoms in the $|1\rangle$ state in each ensemble, and are samples of the correlated random variables $X^{(a)}$ and $X^{(b)}$. Similar to Sec. 5.1, we end up with a dataset $\{\boldsymbol{x}^{(a)}, \boldsymbol{x}^{(b)}\}$, where the different pairs of values $\{x^{(a)}, x^{(b)}\}$ are IID. There are two similar approaches to efficiently estimate ϕ from $\{\boldsymbol{x}^{(a)}, \boldsymbol{x}^{(b)}\}$, which we find to offer equivalent performance. One approach is simply to perform maximum likelihood estimation on $\{\boldsymbol{x}^{(a)}, \boldsymbol{x}^{(b)}\}$ (see Appendix A.1). The other perhaps more intuitive approach is to plot the points $\{x^{(a)}, x^{(b)}\}$ parametrically. Given Eq. 5.22, when randomizing ϕ_L the points $\{x^{(a)}, x^{(b)}\}$ fall along the edge of an ellipse, where the size of the ellipse is related to C, and the opening angle of the ellipse is related to ϕ (Figs. 5.12 and 5.13a). One can fit an ellipse to the resulting distribution via the singular value decomposition to extract ϕ [213].

We are interested in identifying the optimal experimental conditions (namely the interrogation time t and the mean value of ϕ) for measuring the relative frequencies of the atoms with high stability in condition (2). Note that in order to extract ϕ from $\{\boldsymbol{x}^{(a)}, \boldsymbol{x}^{(b)}\}$, the mean value of ϕ must be non-zero. This can be seen by computing the Fisher information associated with the measurement (using PyTorch), which is zero at $\phi = 0$, and reaches a peak at $\phi \simeq \pi/10$ (see Fig. 5.12 and Appendix A.1).[55] However, operating under conditions where the average clock frequencies $f_{a,b}$ of the two ensembles are equal, and thus $\phi = 0$, is helpful for determining the optimal interrogation time t, since the mean value of ϕ will not change as a function of t (see Fig. 5.13b). Such a measurement does not yield useful information about $\delta = f_a - f_b$, but does yield useful information about other aspects of the measurement, including atom coherence and loss

[54] A more general model including different types of loss and decoherence is provided in [213, 365].
[55] Another way of seeing this is the fact that the measurements near $\phi = 0$ are strongly biased [110, 213, 365].

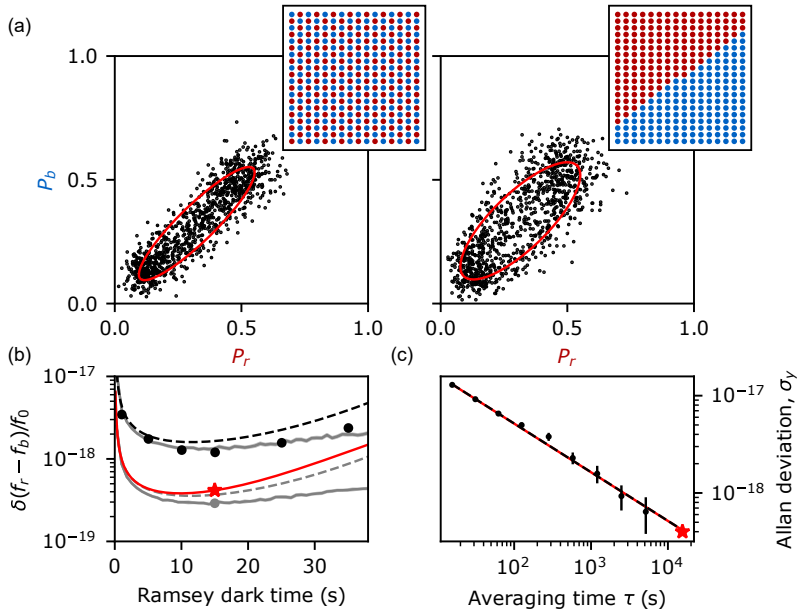

Figure 5.13: Synchronous frequency comparison via ellipse fitting. (a) We perform a synchronous clock comparison by partitioning the atom array into two sub-ensembles (insets, red and blue), and creating a parametric plot of the $^3P_0$ excited-state fraction in the blue ensemble (P_b) versus in the red ensemble (P_r) (in this case at a 15 s interrogation time in $15E_R$-deep tweezers). In the checkerboard (left) partitioning there is no mean frequency shift between the two sub-ensembles, whereas in the diagonal (right) case we expect a 7.0(1.3) mHz shift. The relative frequency between the sub-ensembles can be extracted via ellipse fitting (red lines), which in the diagonal case yields 7.15(18) mHz. Note that such fits are biased near zero phase shift, as is evident in the fit to the checkerboard ensemble, which returns an artificially large phase shift. (b) To identify an optimal dark time, we compute the fractional frequency uncertainty between the sub-ensembles as a function of Ramsey dark time. The black points (grey point) correspond(s) to 13 minutes (4.3 hours) of averaging, and are extracted from the checkerboard partitioning. Note that these values are not representative of a true stability due to biasing. This is made clear by the dashed curves, which correspond to the expected contribution from QPN, and the solid grey curves, which include an additional correction factor calculated via Monte-Carlo simulations to account for the biased fits (shaded regions denote 1-σ confidence interval). At 15 s interrogation times the diagonally separated sub-ensembles have a sufficient phase shift to remove the bias in the fits. This condition (red star) shows the fractional frequency uncertainty of the full 4.3 hour-long measurement, with a value of 4.2×10^{-19}. This value is in good agreement with the expected QPN limit with no bias correction (red curve). (c) We can further compute an Allan deviation associated with this measurement (black points), which averages down with a slope of $5.2(3) \times 10^{-17}$ $(T/\mathrm{s})^{-1/2}$ (black dashed line). This value is in good agreement with the expected value of 5.2×10^{-17} $(T/\mathrm{s})^{-1/2}$ from QPN (red line). Red star is duplicated here as a point of comparison (note that this point is not strictly an Allan deviation, and is extracted via jackknifing [365]).

(namely contributions to C in the simplified model presented in this section).

Because readout occurs in a site-resolved manner in our experiment, the partitioning of a and b can be chosen arbitrarily. Specifically, we choose a "checkerboard" partitioning that yields no net tweezer-induced frequency shift between the two ensembles, and a "diagonal" partitioning that yields a near-maximal

frequency shift (Fig. 5.13a insets). Measurements using the checkerboard partitioning are compared against Monte-Carlo simulations to identify an optimal interrogation time of $t \simeq 15$ s that takes into account both atomic loss and decoherence [365].

Guided by the above measurements, we perform a 4.3 hour-long synchronous comparison using the diagonal partitioning and with an interrogation time of $t = 15$ s. The expected frequency shift between the two ensembles due to the considerations in Sec. 5.3.4 is $\delta = 7.0(1.3)$ mHz, leading to a phase shift of $\phi = 0.219\pi$, and thus nearly optimal performance (see Fig. 5.12). We measure an offset of 7.15(18) mHz over the course of the full 4.3 hour-long measurement, which corresponds to a fractional frequency uncertainty of 4.2×10^{-19}. The Allan deviation associated with the above measurement is determined using a jackknifing procedure [213, 365]. The resulting instability is $5.2(3) \times 10^{-17}$ $(T/\mathrm{s})^{-1/2}$, in good agreement with the expected value of 5.2×10^{-17} $(T/\mathrm{s})^{-1/2}$ from QPN alone (Fig. 5.13c). Moreover, the instability does not deviate from the expected $T^{-1/2}$ scaling for the full duration of the measurement, meaning that, in this specific synchronous comparison, we are not limited by any sources of instability besides QPN down to the 4.2×10^{-19} level.

Note that we do not use repeated interrogation in the measurements described in this section due to the significantly higher losses incurred when imaging in shallow 515 nm tweezers (see Sec. 4.4). However, the long interrogation times used here allow us to match the highest duty cycles achieved in condition (1) of 96%, even without performing repeated interrogation. As a result, although the synchronous measurement in this section is inherently insensitive to LO noise, Dick effect noise is not expected to significantly impact the instability of an asynchronous comparison performed under the same conditions.

5.4.2.2 Correlation spectroscopy

In Sec. 5.4.2.1, we saw that we could extract useful information from the two correlated random variables $X^{(a)}$ and $X^{(b)}$. Although those variables could take on many values, similar arguments apply for binomial random variables, for example when $X^{(i)}$ describes the outcome of a measurement on a single atom i. In fact, $X^{(a)}$ was simply an appropriately normalized sum of the form $\sum_{i \in a} X^{(i)}$ (and similarly for b). This suggests that one can more generally look at correlations between many $X^{(i)}$s to extract useful

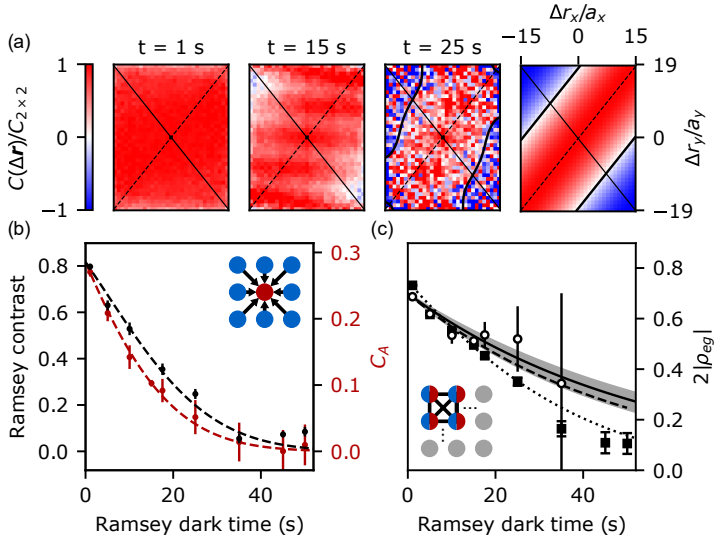

Figure 5.14: Correlation spectroscopy. (a) As a measure of atomic coherence we compute the spatially resolved atom-atom correlation function, $C(\Delta \boldsymbol{r})$, as a function of free evolution time. These plots are normalized by $C_{2\times 2}$ to isolate the effects of dephasing from atom loss and decay (the various correlation functions are defined in more detail in the text). The relative displacements $\Delta r_{x,y}$ are normalized by the array spacing in the relevant direction ($a_x = 1.5$ μm, $a_y = 1.2$ μm) such that the pixel spacing corresponds to the tweezer spacing. Solid (dashed) diagonal lines indicate axes along which the tweezer wavelengths change (remain constant) showing accelerated (reduced) dephasing along the forward (reverse) diagonal of the array due to tweezer-induced frequency shifts (solid contour shows where the correlator passes through zero). The fourth frame is a theoretical prediction at 25 s given our known tweezer frequencies and depths. (b) The coherence of a single atom (red disk in cartoon) can be measured by computing the average correlations C_A between it and an ensemble of reference atoms (blue disks). In this case the reference ensemble is the entire array, and the excess decay of C_A (red points, averaged over a 4×4 block of atoms at the center of the array) compared to the decay of the Ramsey contrast (black points) can be used to quantify the single-atom coherence time. Fits to these quantities (dashed lines) with a Gaussian and exponential component yield overall $1/e$ times of 14.6(7) s and 19.5(8) s respectively. (c) Based on the above measurements, we infer a single-atom coherence time of 48(8) s (dashed line), which is in good agreement with a model based on the measured lifetimes and initial Ramsey contrast (solid line, error in grey). Open circles are C_A with the decay associated with the reference ensemble divided out — this serves as a direct measurement of the single-atom coherence $|\rho_{eg}|$. In the absence of dephasing, an ensemble of atoms would have a Ramsey contrast of $2|\rho_{eg}|$. To extend this measurement from the central 4×4 region to the full array, we consider the average correlation between all atom pairs in a 2×2 block averaged over all such blocks, $C_{2\times2}$. In this case each atom in the block acts as a reference for all other atoms in the block (see cartoon). The square root of $C_{2\times2}$ (black squares) decays with a fitted $1/e$ time of 33(2) s (double-dashed line), and serves as a lower bound on the average atomic coherence across the entire array.

information. For example, correlation spectroscopy techniques with only two ions have been used to extend the interrogation time in a differential measurement to well beyond the ion-laser coherence time [69, 68, 70, 173].

Here, we are interested in using correlations between atoms to characterize tweezer-induced dephasing (see Sec. 5.3.4), and the coherence of the atoms. Specifically, we compute the $g^{(2)}$ correlator between atoms in different tweezers i and j:

$$g_{ij}^{(2)} = \langle \sigma_z^{(i)} \otimes \sigma_z^{(j)} \rangle - \langle \sigma_z^{(i)} \rangle \langle \sigma_z^{(j)} \rangle, \tag{5.24}$$

where $\sigma_z^{(i)}$ is the Pauli z operator acting on the i^{th} atom. After averaging over ϕ_L and otherwise assuming perfect control pulses in the Ramsey measurement, one can show that [365]:

$$g_{ij}^{(2)} = 2|\rho_{eg}^{(i)}||\rho_{eg}^{(j)}| \cos(\phi_i - \phi_j), \tag{5.25}$$

where $\rho^{(i)}$ is the density matrix describing the state of atom i, and $\rho_{eg}^{(i)}$ is an off-diagonal entry of $\rho^{(i)}$.[56] ϕ_i is the phase accrued during the Ramsey interrogation time by atom i. Eq. 5.25 provides information about the relative phase $\phi_i - \phi_j$ between two atoms, and about $|\rho_{eg}^{(i)}|$, which we will refer to as the atomic coherence.

We expect the tweezer-induced frequency shifts to be a function of the relative displacement $\Delta \boldsymbol{r}$ of two atoms, and not of their absolute position. We therefore define $C(\Delta \boldsymbol{r})$ to be the value of $g_{ij}^{(2)}$ averaged over all i and j that are separated by a relative displacement of $\Delta \boldsymbol{r}$. Measurements of $C(\Delta \boldsymbol{r})$ for several interrogation times are shown in Fig. 5.14a. These measurements reveal that along the forward diagonal of the array, where frequency offsets between tweezers — and thus clock frequency offsets — are maximal, the atoms become uncorrelated, and eventually develop negative correlations. Along the anti-diagonal, where there is no frequency offset between tweezers, positive correlations persist over much longer timescales. Although the coarse features of $C(\Delta \boldsymbol{r})$ described above agree with expectations based on independent characterizations of the tweezer potential (and the analysis in Sec. 5.3.4), we further observe the development of fringes in $C(\Delta \boldsymbol{r})$ along the more tightly spaced axis of the array, which we hypothesize are the result of overlaps between tweezers [365].

[56] $\rho^{(i)}$ refers to the state before the final readout pulse in the Ramsey sequence. As a result, $|\rho_{eg}^{(i)}|$ tells you the length of the Bloch vector, or equivalently the Ramsey contrast.

To efficiently extract the atomic coherence, we want to consider groupings of $g_{ij}^{(2)}$ for which $\phi_i - \phi_j \simeq 0$, and thus $g_{ij}^{(2)} \simeq 2|\rho_{eg}|^2$.[57] One approach is to average $g_{ij}^{(2)}$ over 2×2 sub-ensembles of the array, for which we expect tweezer-induced dephasing to be suppressed to a timescale of several hundred seconds. We define $C_{2\times2}$ to be the average over all such 2×2 sub-ensembles, which serves as a lower bound on the quantity $2|\rho_{eg}|^2$.[58] Measurements of $C_{2\times2}$ provide a lower bound on the average $1/e$ atomic coherence time for all atoms in the array of 33(2) s (Fig. 5.14c).

Based on our model for the tweezer-induced clock frequency shift (Sec. 5.3.4) and our balancing of the tweezer intensities (Sec. 3.3.3.2), we expect that the average clock frequency of the entire array is equal to the clock frequency of atoms at the center of the array. Therefore, another approach is to measure the average correlator between the atoms in the central 4×4 sites and the rest of the array, which we define to be C_A (Fig. 5.14b).[59] By looking at the excess decay of C_A relative to the decay of the Ramsey contrast in the entire array, we can isolate the contribution to the decay of C_A from only the central atoms. As a result of the above analysis, we infer a single-atom $1/e$ coherence time of 48(8) s and a resulting atomic oscillator quality factor of $Q = 6.5(1.1) \times 10^{16}$ (Fig. 5.14c) for atoms in the central 4×4 region. This is comparable to the expected value of 55(8) s in Sec. 5.3.2, and corresponds to the useable interrogation time for synchronous frequency comparisons (as in Fig. 5.13) if all tweezers were at the same wavelength, as might be achieved with the use of a spatial light modulator. Note that a significant contribution to the coherence time is simply the vacuum lifetime of the atoms. Often, coherence times are quoted with respect to the surviving atoms in the system. In this case, after removing the contribution from vacuum lifetime (Sec. 5.3.1), our measurements correspond to a coherence time of 96(32) s, and a quality factor of $1.3(4) \times 10^{17}$.[60]

[57] We have assumed here that $|\rho_{eg}^{(i)}| = |\rho_{eg}^{(j)}|$, since the quality of state preparation does not vary significantly site to site, and thus dropped the indices i and j.

[58] This measurement provides a lower bound because any residual frequency differences between atoms can only serve to lower $C_{2\times2}$ due to the $\cos(\phi_i - \phi_j)$ term appearing in Eq. 5.25.

[59] C_A can equivalently be interpreted as the correlator between the central atoms and the total spin projection of the rest of the array.

[60] To put this number in context, if one were to slow the oscillations of the electrons down to more human timescales, like the swinging of the pendulum on a grandfather clock (~ 0.5 Hz), then this quality factor is akin to starting a pendulum at the big bang, and finding that it is still swinging with $\sim 0.5\%$ of its initial amplitude today.

5.5 Tweezer clocks in context

It is worth reflecting on the role of tweezer clocks in the broader context of establishing ever-improving frequency standards. We have discussed a number of benefits to the tweezer system, including lossless and single-atom-resolved readout, very long coherence times, and high fidelity state preparation and control. As a result, tweezer clocks serve as powerful tools for studying new measurement protocols (Sec. 5.4), understanding the limitations in state-of-the-art optical clocks, and exploring new techniques to mitigate those limitations (Sec. 5.3). However, the above benefits come at the cost of limited system size (and thus higher QPN in comparison to lattice clocks), and additional complications associated with the trapping potential (Sec. 5.3.4). It is also important to note that none of the advantages stated above are inherent to the tweezer platform itself. Whereas the tweezer clock is the result of optimizing for flexible control and readout, one could instead optimize for a specific metrological task, incorporating desired controls and measurement techniques as needed.

One can see this tradeoff play out, for example, in recent synchronous comparisons in different optical clocks. The ability to mitigate tunneling at shallow depths in a tweezer array led to an instability of $5.2(3) \times 10^{-17}$ $(T/\mathrm{s})^{-1/2}$ for the synchronous comparison described in Sec. 5.4.2.1 [365]. At the time of publication, this instability was comparable to the state-of-the-art value of 3.1×10^{-17} $(T/\mathrm{s})^{-1/2}$ for a similar comparison in a 3D lattice clock [52], despite using about a factor of 40 fewer atoms. However, these instabilities were soon improved on by an order of magnitude to $4.4 \times 10^{-18}(T/\mathrm{s})^{-1/2}$ in a synchronous comparison in a 1D lattice clock [35].[61] This was accomplished by using carefully engineered tunneling and interactions [3] to extend the achievable interrogation times [148] without sacrificing on atom number.

A similar story arises when considering the accuracy of tweezer clocks, where the tradeoffs made to introduce additional controls have the potential to also introduce additional sources of systematic effects (Sec. 5.3.4). This is in contrast to trapped ion systems, which maintain state-of-the-art accuracy [42] while still offering single-ion-resolved measurements and control that enable advanced interrogation schemes [68, 70, 173], including those involving long range entanglement via quantum networks [233].

[61] Using more atoms than our experiment ($\sim 10^5$ compared to ~ 150), but a shorter interrogation time (~ 3 s compared to 15 s).

With the above considerations in mind, one of the main strengths of tweezer clocks is to serve as flexible testbeds for different measurement protocols and techniques that can eventually be integrated into more specialized sensors. For example, the combination of high fidelity control and relatively large systems sizes makes tweezer clocks ideally suited to study how entanglement[62] can be used to improve the performance of optical clocks, which we will explore in the next chapter.

[62] Particularly in relatively large systems not currently accessible in ion clocks.

Chapter 6

Rydberg interactions between clock qubits

"The light that burns twice as bright burns half as long — and you have burned so very, very brightly."

— Replicant Eldon

Portions of this chapter have appeared in:

N. Schine, A. W. Young, W. J. Eckner, M. J. Martin, and A. M. Kaufman. Long-lived Bell states in an array of optical clock qubits. **Nat. Phys.**, pages 1–7, Aug. 2022

W. J. Eckner, N. Darkwah Oppong, A. Cao, A. W. Young, W. R. Milner, J. M. Robinson, J. Ye, and A. M. Kaufman. Realizing spin squeezing with Rydberg interactions in an optical clock. **Nature**, pages 1–6, Aug. 2023

In Ch. 5 we demonstrated control over large ensembles of optical frequency qubits, and used this control to perform precision measurements of the qubit frequency at the standard quantum limit (SQL). However, in Sec. 5.1 we saw that the SQL, in certain scenarios, can be beaten when one allows for entanglement. As a result, we are interested in engineering entangling gates between clock qubits. There are a variety of mechanisms which one can use to engineer such gates between neutral atoms, including collisional gates based on spin exchange or superexchange [163, 360, 371], and gates based on exchanging photons between atoms and an optical cavity [350, 246]. We opt to excite the atoms to high lying Rydberg states (see Sec. 2.1.1.5), where the polarizabilities of the atoms can be so high that the presence of a nearby atom in a Rydberg state can lead to significant energy shifts. This approach has been used to great success to engineer entangling gates between tweezer-trapped atoms in the past [352, 150, 190, 127, 210, 103], but never in the context of optical frequency qubits and metrology prior to our works. In Sec. 6.1 we will describe the basic properties of Rydberg states, and the mechanisms leading to interactions between Rydberg state atoms. In Sec. 6.2, we will see how the above interactions can be used to generate entanglement between clock qubits, and in

Sec. 6.3 we will benchmark the quality of a two qubit gate using the described mechanism. In Sec. 6.4, we will use the two qubit gate to prepare Bell pairs of clock qubits, and evaluate the utility of these Bell pairs for frequency metrology. Finally, in Sec. 6.5, we will briefly discuss recent results where we use the above tools to prepare larger entangled states that unambiguously achieve performance beyond the SQL.

6.1 Rydberg states

As described in Sec. 2.1.1.5, Rydberg states are high lying electronic states where the outer electron has an orbit that is much larger than usual.[1] These states typically have an energy near the ionization threshold of the atom, and are thus challenging to excite directly with a single photon due to the short wavelengths involved.[2] However, because the $^3\mathrm{P}_0$ excited state in Sr is long-lived, one can drive atoms to the $^3\mathrm{P}_0$ state, and then excite the atoms to high-lying Rydberg states with a reasonable (although still fairly short) wavelength of $\lambda \simeq 317$ nm.[3]

6.1.1 Rydberg spectroscopy

The dipole allowed transitions from the $^3\mathrm{P}_0$ clock state involve states in the $5sns\ ^3\mathrm{S}_1$ and $5snd\ ^3\mathrm{D}_1$ series. As a general rule, the $^3\mathrm{S}_1$ states provide stronger interactions that are also isotropic, whereas the $^3\mathrm{D}_1$ states have weaker and anisotropic interactions, but provide larger dipole matrix elements from the $^3\mathrm{P}_0$ state. As described in Sec. 2.1.1.5, for large n these Rydberg states are approximately hydrogenic, with small deviations in energy captured by a quantum defect parameter δ_n in Eq. 2.6. δ_n can be expressed in terms of a series expansion:

$$\delta_n = \delta^{(0)} + \frac{\delta^{(2)}}{(n-\delta^{(0)})^2} + \frac{\delta^{(4)}}{(n-\delta^{(0)})^4} + \ldots, \tag{6.1}$$

where the constant values $\delta^{(k)}$ are typically determined by fitting to spectroscopic data [338]. For our purposes, it is sufficient to assume that $\delta^{(k)} = 0$ for $k > 4$. The relevant values of $\delta^{(k)}$ for $k \leq 4$ are shown in Tab. 6.1. We find Eq. 6.1 to be accurate for all Rydberg states relevant to our works, but it is important

[1] Typically, on the scale of ~ 0.5 μm in our experiments, in comparison to the size of the ground state orbitals of the valence electrons, which are only a few Bohr-radii in size ($a_0 = 52.9$ pm).

[2] It is worth noting that although laser sources in the deep and extreme ultraviolet exist (with wavelengths of $\lambda \lesssim 280$ nm), they can be very challenging to work with due to absorption in air, photochemical damage, etc.

[3] This is in contrast to approaches that drive a two-photon Raman transition from the ground state, where the intermediate state is typically short-lived and thus can limit the fidelity with which one can drive the ground-Rydberg transition [190, 103].

	$\delta^{(0)}$	$\delta^{(2)}$	$\delta^{(4)}$
$^3S_1$	3.371(2)	0.5(2)	-10(20)
$^3D_1$	2.658(6)	3(2)	-8800(700)

Table 6.1: Quantum defects for the $^3S_1$ and $^3D_1$ Rydberg series in Sr. Values are from [23, 338].

to note that this simple scaling can be broken in rare cases where a Rydberg state from a different series of states is close in energy to the state of interest [113].

Recall from Sec. 2.1.1.5 that the dipole matrix element from the ground state to a Rydberg state with principle quantum number n scales approximately as $\langle g \,|\, \boldsymbol{d} \,|\, n \rangle \propto n^{-3/2}$. Because this argument is based on general scaling laws for the size of the Rydberg orbital, a similar argument applies for the dipole matrix element from the $^3P_0$ state. Also recall that the Rydberg state lifetime τ has a scaling between $\tau \propto n^3$ and $\tau \propto n^2$. In Fig. 6.1, we perform spectroscopy of the $5s40d\ ^3D_1$, $m_j = 0$ state, attaining Rabi frequencies of up to $\Omega = 2\pi \times 18$ MHz using ~ 1 W of power at 317 nm (see Appendix D.3), and measure a lifetime of $\tau = 10.9(4)$ μs.[4]

For the work presented in this thesis, we measure the Rydberg state population by simply relying on the fact that the Rydberg state mostly decays to the ground state, or to states that are lost from the tweezers. Our detection involves blowing away atoms in the ground state, and repumping clock state atoms to the ground state for detection as in Sec. 5.2.1. The above approach results in a detection fidelity of the Rydberg state of 89%, with 11% of the Rydberg state population decaying to the $^3P_0$ clock state.[5] The observed branching ratio is consistent with the expected branching of 1/9 for the $^3S_1$ series of states to decay to the $^3P_0$ clock state in comparison to the $^3P_{J\neq 0}$ states. As we will see below, this low detection fidelity is not important for our work, since the protocols we are interested in leave no population in the Rydberg state. However, in more recent work we have adopted a similar approach to the one demonstrated in [200, 210], where 407 nm light is used to drive a transition of the inner electron,[6] which rapidly autoionizes the outer

[4] Note that this lifetime is measured two ways: by resonantly exciting Rydberg state atoms and observing them decay [290], and by dressing the clock state with the Rydberg state, and observing the dependence of the dressed state lifetime. Both measurements agree to within statistical errors (for lattice depths below $\sim 50 E_R$), and so the value quoted here is based on the dressed state measurement, which has lower statistical error.

[5] Note that these measurements are performed with the trap extinguished. In the case where the trap is on, the strong antitrapping of the Rydberg states causes them to be lost, rather than decaying to low-lying electronic states.

[6] We find this transition to be at 407.886 nm, which corresponds to the unperturbed transition frequency of the Sr^+ ion [282].

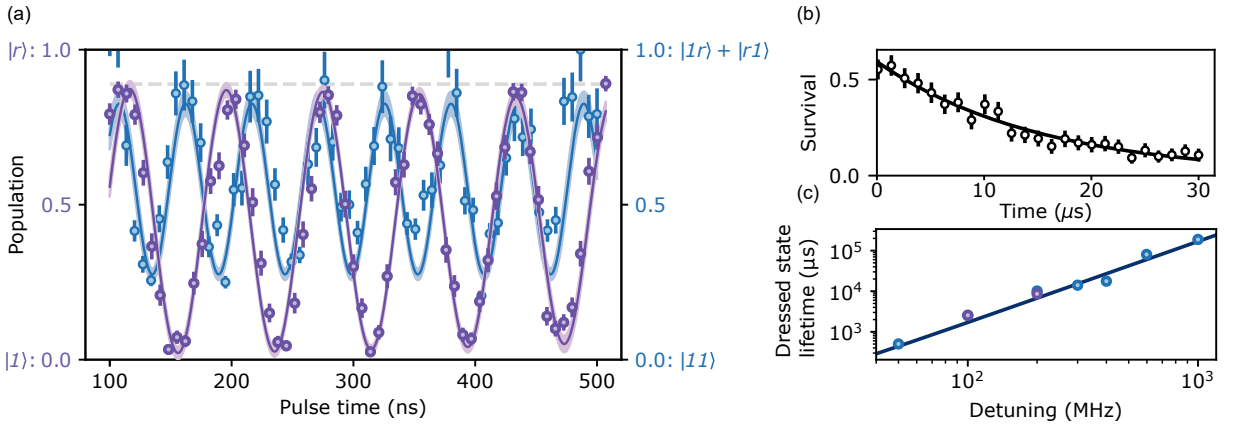

Figure 6.1: Rabi spectroscopy and lifetime of the $5s40d\ ^3\mathrm{D}_1$ Rydberg state. (a) Rabi oscillations of isolated atoms (purple) and atom pairs (blue) when driving the clock to Rydberg transition. Notice that, as expected in the case of Rydberg blockade, the atoms pairs oscillate between a state with no Rydberg excitations ($|11\rangle$) and a state with only a single Rydberg excitation ($|+\rangle \propto |1r\rangle + |r1\rangle$), with an oscillation frequency that is $\sqrt{2}$ higher than that of single atoms. Solid lines are sinusoidal fits to the data, from which we extract a Rabi frequency of $\Omega = 2\pi \times 13$ MHz. The displayed data is representative of the state of the experiment in 2021, where the Rydberg state was detected via loss from the clock state. The dashed line corresponds to the maximum expected value of the observed transfer to the Rydberg state due to the branching ratios associated with decay from the Rydberg state. Recent upgrades have significantly improved both the fidelity of the Rydberg Rabi oscillations, and the fidelity with which we can detect the Rydberg state (see also [210]). We leave a discussion of these upgrades to future works from the group. (b) The lifetime of the Rydberg state can be measured directly by exciting atoms to the Rydberg state, holding for a variable amount of time, and de-exciting the atoms. Atoms that decay from the Rydberg state are lost during this measurement, whereas atoms that are successfully transferred back to the clock state survive. Solid line is an exponential fit to the data, which yields a $1/e$ lifetime of $\tau = 13(3)$ μs. (c) An alternative way of measuring the Rydberg state lifetime is to prepare atoms in the clock state, and to off-resonantly dress the clock state to the Rydberg state (in this case with a Rabi frequency of $\Omega = 2\pi \times 15$ MHz). The detuning is varied (negative values in blue, positive values in purple), resulting in a variation in the measured lifetime of the dressed state. The Rydberg state lifetime extracted from this analysis is $\tau = 10.9(4)$ μs.

electron leading to loss of the atom (now an ion). This loss can occur on timescales much shorter than the Rydberg state lifetime, leading to high detection fidelities of the Rydberg state.

6.1.2 Rydberg interactions

The origin of the strong interactions between Rydberg state atoms is their strong polarizability (see Sec. 2.1.1.5). Schematically, the orbit of the outer electron can be so large in a Rydberg state that the electrons in neighboring Rydberg atoms can push on each other, distorting the electron orbitals and leading to energy shifts that are dependent on the presence or absence of a nearby atom in a Rydberg state.

Under the assumption that only a single outer electron participates in these interactions, and that the atoms are far apart relative to the size of the electron orbitals,[7] the interaction Hamiltonian between two Rydberg state atoms is simply set by Coulomb interactions between the two electrons and nuclei:

$$H_i \propto \frac{1}{|\boldsymbol{r}_1 - \boldsymbol{r}_2 - \boldsymbol{R}|} - \frac{1}{|\boldsymbol{r}_1 - \boldsymbol{R}|} - \frac{1}{|\boldsymbol{r}_2 - \boldsymbol{R}|} + \frac{1}{|\boldsymbol{R}|}, \tag{6.2}$$

where $\boldsymbol{R}$ is the separation between the atomic nuclei, and $\boldsymbol{r}_i$ the coordinate of the outer election in atom i relative to the nucleus of the atom. Having already made the approximation that the atoms are far apart,[8] one can perform a multipole expansion of H_i and neglect all but the lowest order terms. The resulting energy shift for a pair of atoms in a given electronic state β can be expressed as [338]:

$$V(\beta, R) = \sum_N \frac{C_N(\beta)}{R^N}, \tag{6.3}$$

where $C_N(\beta)$ is a constant.

In the absence of a perturbing field, the lowest order non-zero terms in H_i correspond to $N = 5$ and 6.[9] The $N = 5$ contribution is associated with a resonant quadrupole-quadrupole interaction, which vanishes for S states, and is typically relevant only at very long distances. The $N = 6$ contribution is associated with a second order dipole-dipole interaction, which is also known as a van der Waals interaction. In our experiments, $N = 6$ dominates over all other effects, and so the energy shift experienced by two nearby

[7] Such that one can neglect the exchange interaction [338].
[8] Namely, that $\langle|\boldsymbol{R}|\rangle \gg \langle|\boldsymbol{r}_i|\rangle$.
[9] The $N = 1, 2$ interactions correspond to charge-charge and charge-dipole interactions, which vanish for neutral atoms. However, one can induce a permanent dipole moment in the state β with an applied DC or AC electromagnetic field, in which case the leading order term is $N = 3$, corresponding to resonant dipole-dipole interactions. These techniques are outside the scope of this work, but are demonstrated in [8, 269, 20, 82].

atoms in comparison to their free space energy is simply:

$$V(R) = \frac{C_6}{R^6}. \tag{6.4}$$

We will refer to $V(R)$ as the interaction potential. Note that although the above discussion is helpful for isolating the underlying physics, in practice, we often solve for the eigenstates of Eq. 6.2[10] using well-established numerical tools [308].

One can infer the rough scaling of C_6 from the fact that the van der Waals interaction is fourth order in the dipole operator, which scales like n^2 (Sec. 2.1.1.5). Additionally, the energy difference between adjacent Rydberg levels scales like n^{-3} (Sec. 2.1.1.5), leading to an overall scaling of $C_6 \propto n^{11}$. In combination with the increased Rydberg state lifetime with n, this suggests that operating at higher n is strictly beneficial.

However, this is not always the case. Specifically, consider the scenario where one wants to stay in the well-behaved large R limit where the interactions are dominated by the C_6 contribution to the interaction potential.[11] The minimum spacing R_0 where the above condition breaks down is when the dipole coupling between adjacent states is comparable to their energy difference. As in Sec. 2.1.1.5, the dipole matrix element between neighboring states scales like $\Omega_n \coloneqq \langle n \,|\, \boldsymbol{d} \,|\, n+1 \rangle \propto n^2$, and the energy separation like $\Delta E_n \propto n^{-3}$. We therefore have that:

$$\frac{\Omega_n^2}{R_0^3} \sim \Delta E_n \rightarrow R_0 \sim n^{7/3}. \tag{6.5}$$

Given R_0, the maximum value of V is:

$$V(R_0) \sim \frac{\Omega_n^4}{\Delta E_n R_0^6} \sim n^{-3}. \tag{6.6}$$

This is to be compared to the Rydberg state lifetime with a scaling between n^3 (in the absence of BBR) and n^2 (for the worst-case BBR spectrum). As a result, in this setting it is beneficial to operate at lower n, assuming one can arbitrarily reduce the interatomic separation. An additional benefit of operating at lower n is the larger dipole matrix element between the ground (or $^1\mathrm{S}_0$) state and the Rydberg state, meaning

[10] Additionally including the contribution to the full Hamiltonian from the bare atomic Hamiltonians.

[11] One does not necessarily have to make this restriction. For example, with better electric field control one could consider beating many of the scaling arguments made here by going to shorter spacings, or making use of Förster resonances to boost the interaction strength [269]. However, such a restriction can be helpful when one wants to engineer well-behaved interactions between many atoms with variable spacing from each other, as in Sec. 6.5.

higher Rabi frequency for fixed laser power, as well as reduced sensitivity to stray electric fields. Based on the above analysis, we find it useful to design the experiment to minimize the spacing between atoms (see Sec. 3.3). Given fixed atom spacing there is an optimal value of n at which to operate, which for our experiment is $n \sim 40$ (given the interatomic spacing of ~ 0.5 μm in our experiment).[12]

Under the presence of a (global) drive laser on the $|1\rangle \leftrightarrow |r\rangle$ transition with Rabi frequency Ω and detuning Δ,[13] the Hamiltonian for an ensemble of atoms is:

$$H = \sum_i \left(\frac{\Omega}{2} \sigma_x^{(i)} - \Delta |1_i\rangle\langle 1_i| \right) + \sum_{i<j} V(R_{ij}) |r_i r_j\rangle\langle r_i r_j|, \tag{6.7}$$

where $\sigma_x^{(i)}$ is the Pauli x matrix acting on atom i, $|1_i\rangle\langle 1_i|$ is the projector onto the $|1\rangle$ state of atom i, R_{ij} is the separation between atoms i and j, and $|r_i r_j\rangle\langle r_i r_j|$ is the projector onto the subspace where atoms i and j are both in the $|r\rangle$ state. Notice that the first term in Eq. 6.7 simply comes from the familiar single-atom Hamiltonian in Eq. 2.13, whereas the second term shifts the energy of states with two or more Rydberg excitations.

When $V_{ij} \gg \Omega, \Delta$, one is in the so-called "blockade limit," where states with more than a single Rydberg excitation are shifted completely out of resonance, and so the system factorizes into subspaces with one or fewer Rydberg excitations, and more than one Rydberg excitation. One can define a length scale where $V = \Omega$, which is known as the "blockade radius":

$$R_B = \left(\frac{C_6}{\Omega} \right)^{1/6}. \tag{6.8}$$

For systems where all atoms are within R_B of each other, one can solve for the dynamics starting with all atoms in the $|g\rangle$ state, up to errors in the populations on the scale of $(\Omega/V)^2$, by simply eliminating states with more than two Rydberg excitations. For example, the relevant two-atom Hilbert space in the blockade limit is spanned by the states $|11\rangle$, $|+\rangle \coloneqq (|1r\rangle + |r1\rangle)/\sqrt{2}$, and $|-\rangle \coloneqq (|1r\rangle - |r1\rangle)/\sqrt{2}$. Starting in the $|11\rangle$ state, we find that for $\Delta = 0$, $\langle 11 | H | + \rangle = \sqrt{2}\Omega$, and $\langle 11 | H | - \rangle = 0$.[14] Therefore, the dynamics for two blockaded atoms resulting from a drive with Rabi frequency Ω is Rabi oscillations between $|11\rangle$ and $|+\rangle$ with

[12] Note that with fixed laser power, errors due to finite values of Ω/V are not necessarily limiting. In this case it can once again be beneficial to go to higher values of n in order to improve the ratio of the Rydberg state lifetime to the Rabi frequency.

[13] Where $|1\rangle$ refers to the $^3P_0$ state, and $|r\rangle$ to the Rydberg state of choice.

[14] One way to see this intuitively is that for $\langle 11 | H | + \rangle$ there is constructive interference from driving to the $|1r\rangle$ and $|r1\rangle$ components of the state, whereas for $\langle 11 | H | - \rangle$ there is destructive interference. Similar to the discussion in Sec. 2.1.1, the dipole operator is not able to change the parity of a state.

effective Rabi frequency $\sqrt{2}\Omega$. We show an example of the measured time evolution of isolated atoms and pairs of blockaded atoms under the same drive in Fig. 6.1a. Notice that just by driving the above oscillations one can apply an entangling gate that takes a pair of atoms from the product state $|11\rangle$ to the Bell state $|+\rangle$ (which is maximally entangled). This mechanism has been used to prepare Bell states with impressive fidelities of $> 0.991(4)$ (with SPAM correction) [210]. However, in this case the qubits are encoded in the $\{|1\rangle, |r\rangle\}$ electronic states of the atoms, with the lifetime of the $|r\rangle$ state being ~ 80 μs. This short lifetime makes it challenging to realize other operations on the qubits, like single qubit gates.[15]

6.2 Rydberg dressing

One approach to engineering a universal gate set is to use Rydberg interactions to engineer an entangling gate between the long-lived clock qubits discussed in Ch. 5, for which we have already demonstrated high fidelity single qubit gates. Given qubit states $|0\rangle$ and $|1\rangle$, the central idea is to drive closed loops on the $\{|1\rangle, |r\rangle\}$-Bloch sphere that are different depending on the presence or absence of nearby atoms in the $|1\rangle$ state, resulting in populations that remain in the $\{|0\rangle, |1\rangle\}$ subspace at the end of the gate, but with phase shifts that are dependent on the dynamics involving the (interacting) Rydberg states.

Given Eq. 6.7, consider the Hamiltonian H for two atoms in the blockade limit:

$$H = H_0^{(1)} + H_0^{(2)} + H_I \tag{6.9}$$

$$H_0^{(i)} \coloneqq \frac{\Omega}{2}\left(|r_i\rangle\langle 1_i| + |1_i\rangle\langle r_i|\right) - \Delta|r_i\rangle\langle r_i| \tag{6.10}$$

$$H_I \coloneqq \frac{\sqrt{2}\Omega}{2}(|+\rangle\langle 11| + |11\rangle\langle +|) - \Delta(|+\rangle\langle +| + |-\rangle\langle -|), \tag{6.11}$$

where $H_0^{(i)}$ is the single atom Hamiltonian for an isolated atom i, and H_I is the Hamiltonian describing modifications to the single-atom behavior due to Rydberg blockade. Notice that $H_0^{(i)}$ and H_I are quite similar, with the main difference simply being the $\sqrt{2}$ enhancement of Ω.

Note that resonant gate schemes involving evolution under Eq. 6.9 with $\Delta \sim 0$ have been applied to qubits encoded in two hyperfine states of tweezer-trapped alkali atoms [191], including in recent time optimal gate schemes [152] that achieve two qubit gate fidelities as high as 99.52(4)% (with SPAM correction) [103].

[15] For example, the time it takes to translate the atoms by the blockade radius of ~ 1 μm without significant heating is ~ 40 μs (see Sec. 4.5), which is comparable to the lifetime of the qubit.

For the experiments in this work, we focus on the situation where Δ and Ω are varied slowly and continuously. We choose to operate in this regime for two reasons: first, the highest values of Rabi frequency Ω we can reach correspond to timescales that are significantly faster than our ability to switch the drive laser on and off,[16] meaning that gates that take advantage of the full range of Ω that we can access necessarily involve adiabatic ramps. Second, we will see in Sec. 6.5 that staying in the regime where $\Omega/\Delta < 1$ can be beneficial for schemes involving interactions between many atoms at once.

Following the proposed gate scheme and discussion in [223], one can diagonalize Eq. 6.10 to find the dressed eigenstates (Sec. 2.2.1) of a single atom:

$$|\tilde{1}\rangle = \cos(\theta_1)|1\rangle + \sin(\theta_1)|r\rangle, \qquad |\tilde{r}\rangle = \cos(\theta_1)|r\rangle - \sin(\theta_1)|1\rangle, \tag{6.12}$$

where $\tan(2\theta_1) = -\Omega/\Delta$. These states have eigenenergies:

$$E_0^{\pm} = \frac{1}{2}\left(-\Delta \pm \sqrt{\Omega^2 + \Delta^2}\right). \tag{6.13}$$

Similarly, diagonalizing Eq. 6.11, we have:

$$|\tilde{11}\rangle = \cos(\theta_2)|11\rangle + \sin(\theta_2)|+\rangle, \qquad |\tilde{+}\rangle = \cos(\theta_2)|+\rangle - \sin(\theta_2)|11\rangle \tag{6.14}$$

$$\tan(2\theta_2) = -\frac{\sqrt{2}\Omega}{\Delta} \tag{6.15}$$

$$E_2^{\pm} = \frac{1}{2}\left(-\Delta \pm \sqrt{2\Omega^2 + \Delta^2}\right). \tag{6.16}$$

Given the above definitions, we can define an "entangling energy" κ, which corresponds to the difference in energy of two nearby (Rydberg blockaded) $|1\rangle$ state atoms in comparison to two isolated $|1\rangle$ atoms under the influence of a global laser drive with parameters Ω and Δ:

$$\begin{aligned} \kappa^{\pm} &\coloneqq E_2^{\pm} - 2E_0^{\pm} \\ &= \frac{\Delta}{2} \pm \left(\frac{1}{2}\sqrt{2\Omega^2 + \Delta^2} - \sqrt{\Omega^2 + \Delta^2}\right). \end{aligned} \tag{6.17}$$

Note that which sign you take in Eq. 6.17 depends on the sign of Δ. In subsequent discussion, we will omit the "$\pm$" for convenience.

Assuming one applies the laser drive Ω adiabatically, the atoms will follow the dressed eigenstates of the system $\{|00\rangle, |0\tilde{1}\rangle, |\tilde{1}0\rangle, |\tilde{11}\rangle\}$. Specifically, for pulses where Ω starts and ends at zero (see Fig. 6.2), the

[16] Although recent upgrades combined with operation at higher n change this situation.

effective Hamiltonian in the $\{|0\rangle, |1\rangle\}$ subspace is [223]:

$$H = -(E_0 + \frac{\kappa}{2})S_z + \frac{\kappa}{2}S_z^2. \tag{6.18}$$

Here, $S_z \coloneqq \sigma_z \otimes \mathbb{1}/2 + \mathbb{1} \otimes \sigma_z/2$ is the total spin operator, where $\mathbb{1}$ and σ_z are the identity and Pauli z operators acting on a single qubit. Up to a global single qubit phase shift, Eq. 6.18 is simply the one axis twisting Hamiltonian S_z^2 with an energy set by $\kappa/2$. In fact, by performing two identical dressing pulses separated by a spin echo — namely a π-rotation between $|0\rangle$ and $|1\rangle$ — one can remove the contribution from the S_z term in Eq. 6.18, leading to an effective Hamiltonian (see Fig. 6.2):

$$H_{\text{eff}} = \kappa S_z^2. \tag{6.19}$$

Applying H_{eff} for duration $t = \pi/(2\kappa)$ implements the Mølmer-Sørensen (MS) gate [223].[17] Including the effect of the spin echo, one can think of each UV pulse as implementing a controlled phase shift of $\pi/2$ (equivalently, a controlled-$\sqrt{z}$ gate).

Note that, in practice, there are several errors that can affect the fidelity of the applied gate, which we discuss in more detail in Sec. 6.3. One significant source of error is the finite lifetime of the $|r\rangle$ state, which we measure to be $\tau_R = 10.9(4)$ μs for the $5s40d\ ^3\text{D}_1$, $m_j = 0$ state (see Sec. 6.1.1). As a result of this error, optimizing for the fidelity of the above MS gate involves balancing errors associated with decay of the Rydberg state with errors due to non-adiabaticity, where not all the atomic population returns to the $\{|0\rangle, |1\rangle\}$ subspace.

For weak dressing where $|\Delta| \gg \Omega$, $\kappa \simeq -\Omega^4/(8\Delta^3)$. However, in this limit the population in the Rydberg state, namely the $|r\rangle$ contribution to $|\tilde{1}\rangle$, is $\simeq \Omega^2/(4\Delta^2)$. Since the lifetime τ of the $|\tilde{1}\rangle$ state is dominated by the contribution from the short-lived $|r\rangle$ state, it has the same Ω^2/Δ^2 scaling. As a result, $\kappa/\tau \propto 1/\Delta$, suggesting that the quality of the gate should exclusively improve with lower Δ, as long as the functional form of the interaction remains appropriate.

For this reason, we choose to perform trapezoidal ramps of Ω and Δ from a regime where $|\tilde{1}\rangle \simeq |1\rangle$

[17] Note that operations involving the clock require tight confinement to be within the Lamb-Dicke regime, as described in Sec. 5.2. Conversely, operations involving the Rydberg state are improved with shallow traps, due to the strong anti-trapping experienced by the Rydberg state (since the trapping light is blue-detuned from various transitions to nearby Rydberg states). As a result, we perform linear ramps of the lattice depth with a duration of ~ 1 ms between depths of $12E_{\text{R}}$(for Rydberg dressing operations) and $42E_{\text{R}}$(for clock operations) for the gate sequences discussed in this chapter. The axial lattice is held low, at $8E_{\text{R}}$, for all gate operations.

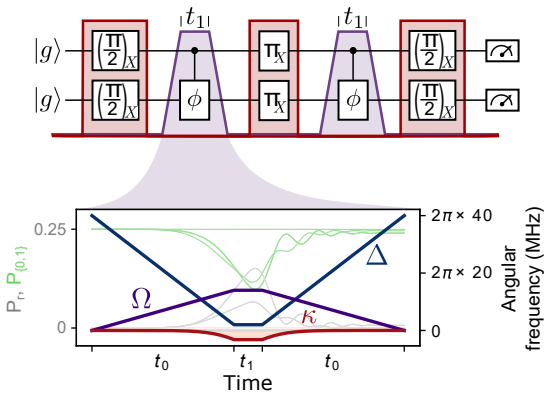

Figure 6.2: Adiabatic gate protocol. The gate is composed of global single qubit rotations (red pulses), and operations involving Rydberg dressing (purple pulses). The Rydberg dressing operations are shown in more detail in the lower panel, where linear ramps of the Rabi frequency Ω and detuning Δ of the laser field used to drive the atoms to the Rydberg state result in a value of the entangling energy κ that varies in time. The atomic populations during these ramps, as predicted via a master equation calculation, are also displayed, with $P_{\{0,1\}}$ denoting the population in the qubit subspace, and P_r the population in any state involving a Rydberg excitation, or a lost atom. Notice that at the end of the ramp, almost all the population returns to the qubit subspace. As a result, up to a single qubit phase shift, the dressing operations can be thought of as applying a controlled-phase gate. The full pulse sequence depicted here implements the Mølmer-Sørensen gate.

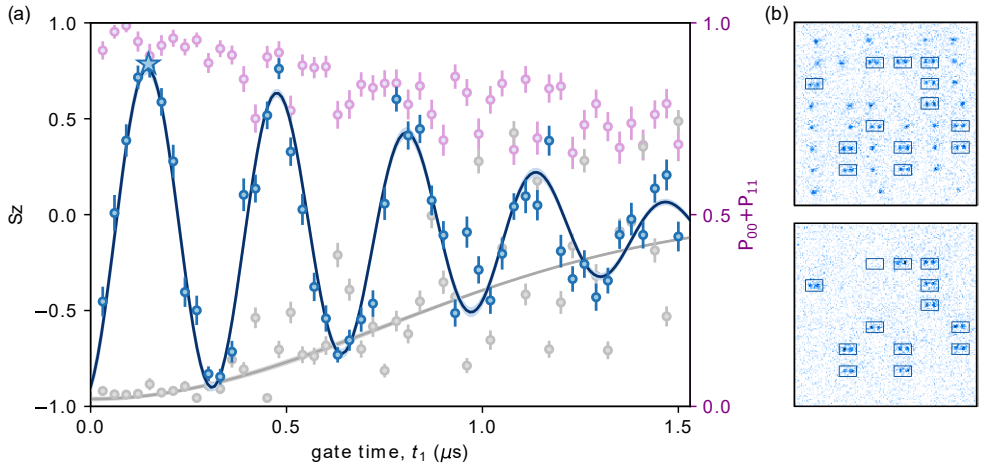

Figure 6.3: Adiabatic gate population dynamics. (a) We apply the gate shown in Fig. 6.2, but varying the time t_1 during which the Rydberg dressing field is held at near-resonant conditions between ramps. As a result, pairs of atoms effectively evolve under the Hamiltonian $H = \kappa S_x^2$ for variable time, resulting in oscillations in the measured spin projection S_z of the atoms along the z axis (blue points). In comparison, the spin projection of isolated atoms (grey points, multiplied by two in vertical direction to match scales) does not evolve significantly. Additionally, note that at early evolution times, pairs of atoms stay primarily in the $\{|00\rangle, |11\rangle\}$ subspace, as evidenced by the high probability $P_{00} + P_{11}$ of observing either the $|00\rangle$ or $|11\rangle$ states. At longer evolution times, the spin echo fails to perfectly cancel the single atom phase shifts, leading to uncontrolled global single qubit rotations and therefore dephasing (decay of the Rydberg state also causes some additional decoherence). (b) An illustrative pair of images corresponding to the evolution time marked by a star. Note that in this measurement we do not use rearrangement, and simply allow an array of doublets to be loaded with 0, 1, or 2 atoms (top image, 2 atom cases circled). In the second image, atoms in the $|1\rangle$ state appear as bright, and we find that most doublets that contained two atoms are in the $|11\rangle$ state after the gate is applied, whereas most single atoms are in the $|0\rangle$ state.

to a regime where $|\tilde{1}\rangle \simeq |r\rangle$ and back. Specifically, Ω and Δ are ramped from initial values of 0 MHz and 40 MHz respectively to values of 18 MHz and 2 MHz using linear ramps with a duration of 350 ns (see Fig. 6.2). The parameters are held at these values for duration t', and then ramped back to their initial values in 350 ns.[18]

In Fig. 6.3, we present results studying the dynamics of pairs of atoms evolving under H_{eff}. Specifically, we arrange our atoms into pairs such that atoms within a pair are fully blockaded, whereas interactions between different atom pairs are negligible. We then perform a global rotation about y such that $z \to x$. This yields $H_{\text{eff}} \to \kappa S_x^2$, which couples the $|00\rangle$ state directly to the $|11\rangle$ state without ever passing through the $|01\rangle$ or $|10\rangle$ states. The resulting dynamics correspond to oscillations in the total spin projection of the atom pair $\langle S_z \rangle$. However, the lack of population in the $|01\rangle$ and $|10\rangle$ states means that when $\langle S_z \rangle = 0$, the

[18] These parameters vary slightly depending on the specific experimental conditions, see [290] for details.

prepared state is of the form $|\psi\rangle = (|00\rangle + e^{-i\phi}|11\rangle)$, namely a Bell state with maximal energy separation between the two components of the state in the computational basis. Notice that this state matches the form of Eq. 5.10, and, in principle, saturates the HL for two qubits.

6.3 Benchmarking entangling gates

6.3.1 Measuring gate fidelity

The ideal solution to measuring the fidelity and general performance of our gates is to perform something like gate set tomography [221] or randomized benchmarking [177]. However, due to the significant experimental overhead posed by such techniques, it is common to perform simpler measurements that can more efficiently (with respect to experimental controls and statistics) extract the quantities of interest (e.g. the Rabi and Ramsey oscillations we use to characterize our single qubit gates) [17, 112].

Specifically, we are interested in applying our entangling gate to a well characterized and separable initial state, and will use the fidelity with which the resulting maximally entangled state is prepared as a stand in for the fidelity of the entangling gate. This approach makes the assumption that the sources of error in the gate are well understood, and that there are no surprises when applying the gate to a different input state. We will apply the MS gate (the pulse sequence in Fig. 6.2) to the $|00\rangle$ state, resulting in the target state $|\psi\rangle = (|00\rangle + e^{i\phi}|11\rangle)\sqrt{2}$ (where ϕ is a well-defined and constant phase that depends on the details of our gate sequence). The fidelity of the prepared state ρ with the target state $|\psi\rangle$ is:

$$\mathcal{F} \coloneqq \langle\psi\,|\,\rho\,|\,\psi\rangle = \frac{1}{2}\left(\rho_{00,00} + \rho_{11,11}\right) + \frac{1}{2}\left(e^{i\phi}\rho_{11,00} + e^{-i\phi}\rho_{00,11}\right). \tag{6.20}$$

We can measure $\rho_{00,00}$ and $\rho_{11,11}$ directly, as in Fig. 6.3, and Fig. 6.4a at the evolution time that corresponds to the MS gate.

To measure the off-diagonal coherences $\rho_{00,11}$ and $\rho_{11,00}$ we can perform a global $\pi/2$ rotation about a variable axis with angle ϕ_α from the x axis [280], described by a single qubit rotation matrix:

$$R_{\pi/2}(\phi_\alpha) \coloneqq e^{-i\frac{\pi}{4}\sigma_x} e^{-i\frac{\phi_\alpha}{2}\sigma_z}. \tag{6.21}$$

This being a global rotation, the action on the two qubits is $R^{(2)}_{\pi/2}(\phi_\alpha) \coloneqq R_{\pi/2}(\phi_\alpha) \otimes R_{\pi/2}(\phi_\alpha)$. The parity

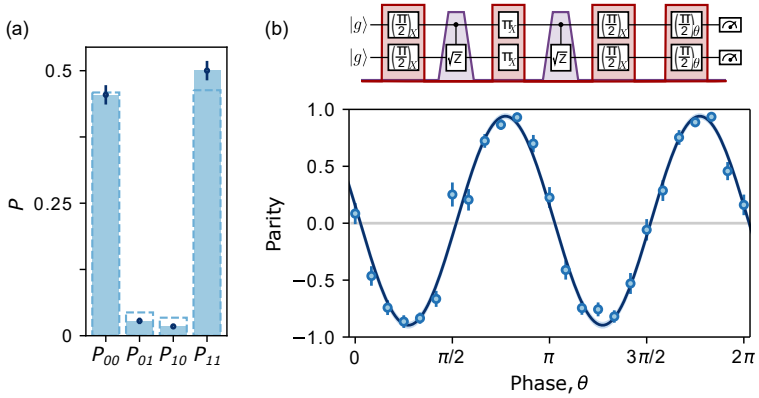

Figure 6.4: Bell state fidelity via population and parity measurements. (a) Measurements of the two atom populations determine the diagonal entries of the two atom density matrix, and yield a value of $P_{00} + P_{11} = 0.955(28)$ (blue) after correcting for measurement errors. Dashed outlines indicate the value of the uncorrected measurement. (b) The off-diagonal coherences in the two atom density matrices can be determined from the contrast of the oscillations in parity when applying an additional $\pi/2$-pulse about a variable axis θ (see circuit diagram). The measured values of the parity are shown in blue (correcting for measurement errors), with the solid line corresponding to a fit to a sin function (shaded region denotes 1σ confidence interval). The fitted value of the oscillation contrast is $C = 0.919(28)$ (including both statistical errors, and systematic errors on the experiment).

$\mathcal{P} \coloneqq \sigma_z \otimes \sigma_z$ of the resulting state is:

$$\begin{aligned}\langle \mathcal{P} \rangle &= \mathrm{Tr}\left(\mathcal{P}\, R^{(2)}_{\pi/2}(\phi_\alpha)\, \rho\, R^{(2)}_{\pi/2}(\phi_\alpha)^\dagger\right) \\ &= \rho_{10,01} + \rho_{01,10} + e^{-2i\phi_\alpha}\rho_{00,11} + e^{2i\phi_\alpha}\rho_{11,00}.\end{aligned} \tag{6.22}$$

Therefore, one can simply measure $\mathcal{P}$ as a function of ϕ_α to isolate the values of $\rho_{00,11}$ and $\rho_{11,00}$ by looking at the contrast of the observed oscillations. We show a measurement of these parity oscillations in Fig. 6.4b.

Note that obtaining an accurate estimate of $\rho_{00,00}$ requires careful handling of the decay of the $|r\rangle$ state, which can lead to a systematic increase in this quantity. For example, the $|11\rangle$ state can decay to the state $|1\varnothing\rangle$, where $|\varnothing\rangle$ refers to the vacuum state where the atom is absent. The state $|\varnothing\rangle$ is outside the $\{0,1\}$ computational basis, but our measurements don't explicitly differentiate between $|0\rangle$ and $|\varnothing\rangle$.[19] As a result, after the final $\pi/2$ pulse in the sequence, $|1\varnothing\rangle$ can incorrectly be measured to be $|00\rangle$.

To take the above effect into account in our estimate of the fidelity, we perform measurements after performing the gate both with and without a blowaway pulse to remove $|0\rangle$ atoms. The resulting measurement provides a direct estimate of the proportion of atoms in the $|\varnothing\rangle$ state, where absent atoms correspond to genuine atom loss. We perform similar measurements with single atoms, or pairs of atoms, and define $p^{(1)}$ and $p^{(2)}$ to be the single and two atom loss probabilities due to Rydberg state decay, if the atoms all start in the $|1\rangle$ state. The above quantities are measured by comparing experiments with and without the UV dressing pulse:

$$\begin{aligned}p^{(1)} &= 2\left(P(a|\text{no UV}) - P(a|\text{UV})\right) \\ p^{(2)} &= 4\left(P(aa|\text{no UV}) - P(aa|\text{UV}) - \frac{p^{(1)}}{2}\right),\end{aligned} \tag{6.23}$$

where $P(a|\cdot)$ and $P(aa|\cdot)$ refer to the conditional probabilities to observe that all the atoms survive in the single and two atom cases respectively.[20] We find that $p^{(1)} = 0.011(6)$ and $p^{(2)} = 0.075(24)$.

When in the gate sequence a decay event occurs affects how likely it is to result in a false positive detection of the $|00\rangle$ state. For example, if the decay occurs early in the sequence, the two atoms are in a product state prior to the decay, and so the state of the remaining atom is not affected. The subsequent

[19] Note that decay to the $\{0,1\}$ subspace is correctly handled, since the loss of coherence in these components of the state will be reflected in the estimate of the off-diagonal elements of the density matrix.

[20] Note that the factors of 2 and 4 are because we perform these measurements in a sequence where the atoms are on the equator of the (collective) Bloch sphere.

evolution simply corresponds to the evolution of one atom. If the decay occurs late in the sequence, the second atom is projected into a mixed state. Averaging over decays that can occur during the pulse sequence, we find that the conditional probability of a false positive detection of $|00\rangle$ given that decay occurs is $P(00|\text{decay}) = 0.75$. We therefore have that the measured population in $|00\rangle$ in the presence of these errors is:

$$P(00) = P(00|\text{no decay})\left(1 - \frac{p^{(1)}}{2} - \frac{p^{(2)}}{4}\right) + 0.75\left(\frac{p^{(1)}}{2} + \frac{p^{(2)}}{4}\right). \tag{6.24}$$

Our measurement of $P(00)$ therefore overestimates the true value of $\rho_{00,00}$ by $0.75(p^{(1)}/2+p^{(2)}/4) = 0.018(5)$. Including an additional uncertainty of 0.006 due to the fact that $P(00|\text{decay})$ is based on simulations, but bounded to be between 0.5 and 1, leads us to an estimate of a systematic error in our measurement of $\rho_{00,00}$ of 1.8(8)%. When including the effect of the above error, the measured Bell state fidelity is 87.1(1.6)%. However, this fidelity includes contributions from a number of state preparation and measurement (SPAM) errors that are unrelated to the applied MS gate.

6.3.2 Correcting SPAM errors

We opt to include the single qubit gates used in the circuit that implements the MS gate in our quoted gate error, but are interested in estimating how much of the gate infidelity is due to measurement errors, as well as due to additional loss during the various ramps in lattice depth. This is particularly relevant since we have made significant improvements to both sources of errors since the above data was taken (see Secs. 3.4 and 4.4). We discussed the relevant errors in Sec. 4.4 for identifying the state of an atom on a given site, including p_{01} false positive errors and p_{10} false negative errors (including atom loss). What remains is to determine how these errors affect the estimate of the Bell state fidelity.

We perform calibration measurements in which the full experimental sequence is executed, except for the clock, UV, and blowaway pulses. These calibrations yield measured values of $p_{01} = 0.004(1)$ and $p_{10} = 0.031(3)$ for the data presented in this chapter (except for in Sec. 6.5). Given these calibrations, we can construct a measurement process matrix M that maps the true underlying results $\boldsymbol{n}$ to the measured

values $\boldsymbol{m}$:

$$\boldsymbol{m} = M \cdot \boldsymbol{n}. \tag{6.25}$$

$\boldsymbol{n} \coloneqq \{n_{11}, n_{10}, n_{01}, n_{00}\}$ is a column vector containing the true number of experimental trials where the atoms were detected[21] in state $|11\rangle$, $|10\rangle$, $|01\rangle$, and $|11\rangle$ (and similarly for $\boldsymbol{m}$ and the observed results). We have that:

$$M = \begin{pmatrix} 1-2p_{10} & p_{01} & p_{01} & 0 \\ p_{10} & 1-p_{10}-p_{01} & 0 & p_{01} \\ p_{10} & 0 & 1-p_{10}-p_{01} & p_{01} \\ 0 & p_{10} & p_{10} & 1-2p_{01} \end{pmatrix}, \tag{6.26}$$

where we have neglected processes that are beyond first order in p_{01} and p_{01}, and assumed that these errors are uniform across the array.[22] We can apply the inverse of M on other measurements $\boldsymbol{m}$ to obtain estimates of the true underlying values $\boldsymbol{n}$, which can in turn be used to estimate the values of different observables (we will refer to these estimates as "SPAM corrected" values). The estimators for loss η and parity $\mathcal{P}$ are:

$$\check{\eta} = \sum_i \frac{n_{10}}{n_{11}+n_{10}} \tag{6.27}$$

$$\check{\mathcal{P}} = \sum_i \frac{n_{11}^{(l)}n_{11}^{(r)} - n_{11}^{(l)}n_{10}^{(r)} - n_{10}^{(l)}n_{11}^{(r)} + n_{10}^{(l)}n_{10}^{(r)}}{n_{11}^{(l)}n_{11}^{(r)} + n_{11}^{(l)}n_{10}^{(r)} + n_{10}^{(l)}n_{11}^{(r)} + n_{10}^{(l)}n_{10}^{(r)}}, \tag{6.28}$$

where the summation is understood to run over all sites in the array for $\check{\eta}$, and all doublets in the array for $\check{\mathcal{P}}$. The superscripts (l) and (r) refer to the left and right site in a given doublet.

The uncertainty in the resulting estimates includes contributions from both statistical uncertainty and uncertainty in the SPAM correction procedure. The statistical uncertainty arises from finite sampling of a multinomial distribution. The uncertainty in the SPAM correction is related to the statistical uncertainty in the elements of M. Although in principle Eq. 6.26 can be inverted analytically, and the errors propagated in the standard fashion, we opt to approximate this inversion numerically. Specifically, given the known statistical errors in the elements of M we generate $50,000$ random samples of M and invert each one to obtain a distribution in M^{-1}. We confirm that the uncertainties in the elements of M^{-1} remain Gaussian,[23] and use the standard deviations of these elements to obtain an estimate of the error in the inversion process.

[21] Keeping in mind that $|0\rangle$ here includes the contribution from $|\varnothing\rangle$.

[22] Both assumptions are reasonable, since higher order errors and the effect of inhomogeneity are below the statistical error in our measurements.

[23] Note that this need not always be the case. Certain process matrices could have inverse matrices containing skew or even multimodal distributions.

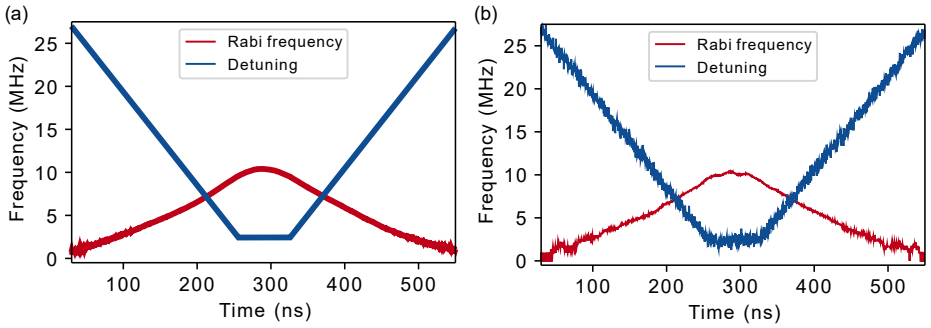

Figure 6.5: Imperfect Rydberg pulses. (a) The average ramp in Rabi frequency applied during the gate (as measured directly in the apparatus using a photodiode), and the ideal ramp in detuning. (b) To generate a more realistic model for the ramps in the experiment, we use a model for the frequency noise on the laser to generate Monte-Carlo samples of the detuning with realistic frequency fluctuations, and use measured single-shot values of the Rabi frequency from the experiment.

Note that this approach neglects covariances between the elements of M^{-1}, which are clearly present given Eq. 6.26. However, because it is straightforward to take calibration measurements with very low statistical error, the characterization of errors in M^{-1} can be seen as a rough confirmation that our measurements errors are dominated by statistical errors in the measurement itself, and not by errors in SPAM correction.

As a result of the above analysis, we infer that with SPAM correction the fidelity of our entangling gate (and of the prepared Bell states) is 92.8(2.0)%, in comparison to 87.1(1.6)% without SPAM correction.

6.3.3 Error budget

As described earlier, the dominant errors expected in our gate are associated with tradeoffs between adiabaticity and the Rydberg state lifetime. We can compare the measured fidelities to a master equation simulation based on the measured Rydberg state lifetime (Sec. 6.1.1), as well as the measured intensity and phase noise of the UV laser (Fig. 6.5). Specifically, we work in the two atom basis, with the state of each atom being spanned by the states $\{|0\rangle, |1\rangle, |r\rangle, |\varnothing\rangle\}$. Time evolution of the two atom density matrix ρ is given by:

$$\frac{\partial \rho}{\partial t} = -\frac{i}{\hbar}[H(t), \rho] + \mathcal{L}\rho \tag{6.29}$$

$$\mathcal{L}\rho := \gamma \left(J \rho J^\dagger - \frac{1}{2}\{J^\dagger J, \rho\} \right), \tag{6.30}$$

Effect	Error
Non-adiabaticity	0.97%
Rydberg state lifetime	0.94^{+13}_{-11}%
Intensity noise	0.73^{+17}_{-17}%
Phase noise	0.00^{+3}_{-4}%
single qubit rotation errors	0.57^{+20}_{-18}%
Rydberg decay systematic	0.90%
Total	4.10^{+40}_{-36}%

Table 6.2: Error budget for the Mølmer-Sørensen gate.

where γ is the rate, and $J \coloneqq |\varnothing\rangle\langle r|$ the jump operator, associated with decay of the Rydberg state.[24] $H(t)$ is given by Eq. 6.7, with the appropriate time dependence of Ω and Δ. Laser intensity noise is incorporated by measuring the intensity of the pulses applied on the experiment directly, and using those traces to determine the values of $\Omega(t)$. Similarly, we measure the spectrum of the frequency noise of the laser using the in loop error signal of the PDH lock (see Appendix D.3), and use this to generate Monte-Carlo-sampled laser frequency trajectories that determine $\Delta(t)$. Errors in the single qubit rotations are characterized independently (see Sec. 5.2), and integrated into the above simulations via Monte-Carlo sampling.

By performing simulations with and without decay, and with and without laser noise, we can isolate the contribution of each effect, and summarize the expected magnitude of the errors originating from these sources in Tab. 6.2. The uncertainties in the estimated contribution of laser intensity and phase noise (including both the clock and UV lasers) are bootstrapped confidence intervals. The uncertainty in the contribution from the Rydberg decay rate is based on the statistical uncertainty in the measurement of γ. The resulting error budget explains much, but not all, of the observed infidelity in our entangling gate. However, recent results from the group that improve the fidelity of the entangling gates to $\gtrsim 98.5\%$ suggest that the remaining infidelity was due to uncharacterized errors in the single qubit rotations.[25]

[24] Note that this model makes the simplifying assumption that decaying Rydberg state atoms are lost. This is not entirely accurate, since the trap depths we operate at during the gate are not so high that anti-trapping of the Rydberg state is the only source of decay (decay to the $|0\rangle$ or $|1\rangle$ states is still possible). However, the approximation is reasonable because multiple rounds of excitation and decay are not expected to play a significant role in our experiment (due to the relatively low error rates).

[25] In particular, the measured atom laser coherence time (see Sec. 5.2.4) indicated that the only contribution to single qubit rotation errors occurred during the rotations themselves, due to residual motion of the atoms and high frequency phase and intensity noise on the laser. However, in more recent measurements we find that there are two timescales associated with the decay of atom laser coherence, leading to larger than expected errors when two single qubit rotations are separated by more than a few milliseconds.

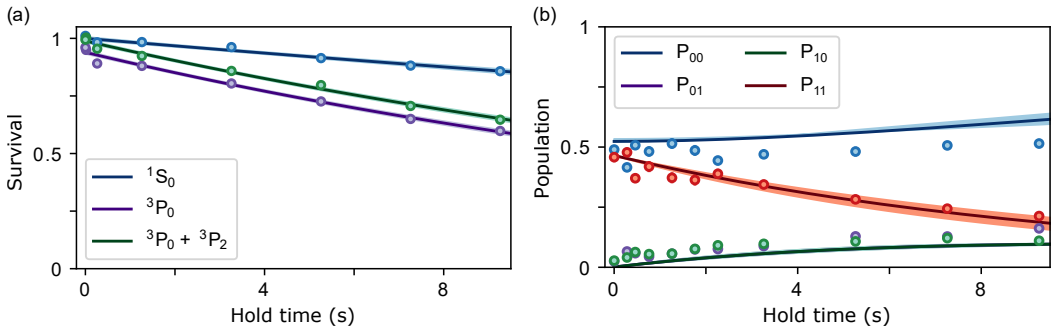

Figure 6.6: Population dynamics of entangled states. (a) Measurements of the $^1S_0$ ground and $^3P_0$ clock state lifetime of single atoms. Note that in order to diagnose certain decay channels, we measure both the population in the clock state, and the population in both the clock state and the $^3P_2$ state. The curves are exponential fits, from which we extract the single atom decay rates. (b) Similar measurements of the population dynamics for Bell pairs, where 0 and 1 refer to the ground and clock states, and the depicted populations are for all two atom states. We primarily observe that P_{11} decays to $\{P_{01}, P_{10}\}$, which can subsequently decay to P_{00}. A master equation model (colored curves), which uses only the measured single atom lifetimes and initial density matrix describing the Bell pairs, is in reasonable agreement with the data. Shaded regions are $\pm 1\sigma$ confidence intervals in the model.

6.4 Characterizing Bell state coherence

With Bell states prepared with relatively high fidelity, we are interested in characterizing their lifetime and coherence, and thus their utility for performing entanglement enhanced measurements in the sense described in Sec. 5.1.1. Similar to Ch. 5, we characterize the coherence of the prepared Bell states by measuring the populations of the atomic states and the contrast of their Ramsey oscillations[26] as a function of time.

Under the conditions of the above experiments, we measure a ground state lifetime of 60.5(5.2) s, and a clock state lifetime[27] of 20.2(8) s (see Fig. 6.6a). Including only these single atom decay processes in a master equation model[28] yields reasonable agreement with the observed two atom populations (Fig. 6.6b), with a relatively long lifetime of 12.3(1.0) s for the $|11\rangle$ state. However, the contrast of the Bell state parity oscillations decays with a Gaussian $1/e$ time constant of $\tau_B = 407(13)$ ms, in comparison to exponential decay of the single atom Ramsey contrast of $\tau = 1.6(2)$ s (Fig. 6.7). The above measurements are performed with an additional spin echo pulse in the middle of the Ramsey dark time to remove effects related to variations in

[26] In the case of Bell states, the parity oscillations used to characterize gate fidelity are analogous to Ramsey oscillations.
[27] Including both decay from the clock state and atom loss.
[28] Similar to Eq. 6.29, except where H corresponds to free evolution of the atoms in the absence of the drive, and with separate jump operators and rates for atom loss, and decay from the clock state to the ground state.

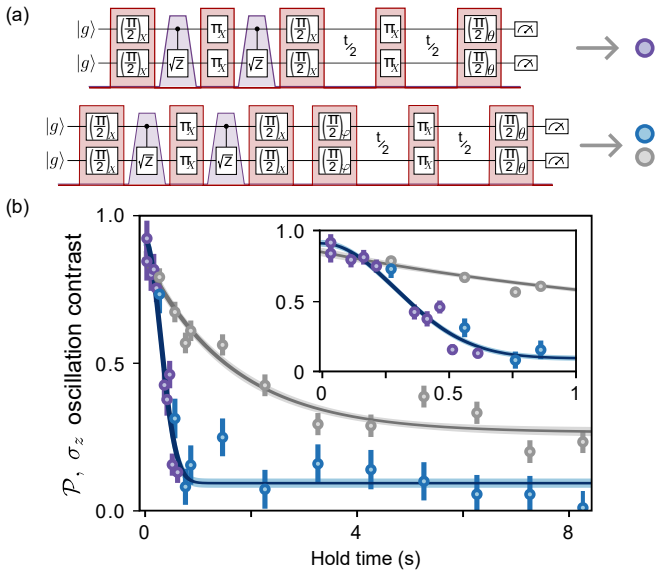

Figure 6.7: Ramsey coherence time of Bell states. (a) We prepare Bell pairs via the procedure described in the text, and inject these Bell pairs into a Ramsey measurement sequence, as shown in the upper panel (this circuit applies to the purple data appearing in this figure). In order to perform parallel measurements with both Bell pairs and isolated atoms, we can introduce an additional rotation about the ϕ axis, which results in a Ramsey measurement for both Bell pairs and isolated atoms, as shown in the lower panel (this circuit applies to the blue and grey data appearing in this figure). (b) As before, the coherence of isolated atoms relative to the laser is characterized by measuring the contrast of the Ramsey oscillations (oscillations in σ_z) as a function of time (grey points). Similarly, the coherence of Bell pairs relative to the laser can be characterized by measuring the contrast of parity ($\mathcal{P}$) oscillations as a function of time (blue and purple points). The solid contours are exponential (Gaussian) fits, allowing for an offset, from which we extract the coherence time of the isolated atoms (Bell pairs), and the shaded regions are $\pm 1\sigma$ confidence intervals.

the lattice intensity in combination with a slight detuning of the lattice from the magic wavelength. Despite removing this systematic, the coherence time between the Bell states and the laser is reduced by more than a factor of two in comparison to between single atoms and the laser.[29]

The above measurements suggest that there are uncharacterized sources of dephasing of the Bell pairs relative to the rotating frame defined by the clock laser, but for certain measurements we are instead interested only in the coherence of the atoms relative to each other. To study these coherences independent of effects relating to the atom-laser coherence, we turn to correlation spectroscopy techniques similar to those in Sec. 5.4.2.2. Specifically, we prepare an ensemble containing both Bell pairs and isolated atoms in the $|0\rangle$ state via the gate sequence described in Sec. 6.2, and perform an additional $\pi/2$ rotation about the x axis, which leaves the parity of the Bell states unperturbed, but rotates the single atoms into the state $(|0\rangle - i|1\rangle)/\sqrt{2}$. This initializes a Ramsey measurement involving both single atoms and Bell pairs in a single experimental trial, allowing us to measure correlation functions between single atoms, between Bell pairs, and between single atoms and Bell pairs. In particular, we will be interested in the single atom $\sigma_z\sigma_z$ (or $g^{(2)}_{ij}$) correlation function as in Eq. 5.24, the parity-parity correlation function between Bell pairs:

$$\mathcal{C}^{(\mathcal{PP})}_{ij} \coloneqq \langle \mathcal{P}^{(i)} \otimes \mathcal{P}^{(j)} \rangle - \langle \mathcal{P}^{(i)} \rangle \langle \mathcal{P}^{(j)} \rangle, \tag{6.31}$$

and the parity-σ_z^2 correlation function between one Bell pair and one single atom:

$$\mathcal{C}^{(\mathcal{P}\sigma_z^2)}_{ij} \coloneqq \langle \mathcal{P}^{(i)} \otimes (\sigma_z^{(j)})^2 \rangle - \langle \mathcal{P}^{(i)} \rangle \langle (\sigma_z^{(j)})^2 \rangle. \tag{6.32}$$

In a Ramsey measurement, the parity-parity correlation function $\mathcal{C}^{(\mathcal{PP})}_{ij}$ plays an analogous role for the Bell pairs as $g^{(2)}_{ij}$ plays for single atoms (see Sec. 5.4.2.2), and by a similar argument one can show that when averaging over the laser phase [290]:[30]

$$\mathcal{C}^{(\mathcal{PP})}_{ij} = 2|\rho^{(i)}_{00,11}||\rho^{(j)}_{00,11}| \cos(\phi^{(i)}_{00,11} - \phi^{(j)}_{00,11}), \tag{6.33}$$

where $\phi^{(i)}_{00,11} \coloneqq \arg(\rho^{(i)}_{00,11})$ is the phase associated with the $\rho_{00,11}$ component of the density matrix for the i^{th} Bell pair. Therefore, by the same arguments as in Sec. 5.4.2.2, one can average $\mathcal{C}^{(\mathcal{PP})}_{ij}$ over all choices of

[29] This is in part explained by the different functional forms used in the fits — although ideally both the single atom and Bell state data would be fit with the same functional form, we found the p-value of a Gaussian fit to the decay of the parity contrast to be lower than that of an exponential fit.

[30] Under the assumption that $\rho_{01,10} \simeq 0$, as confirmed via our characterizations in Sec. 6.3.

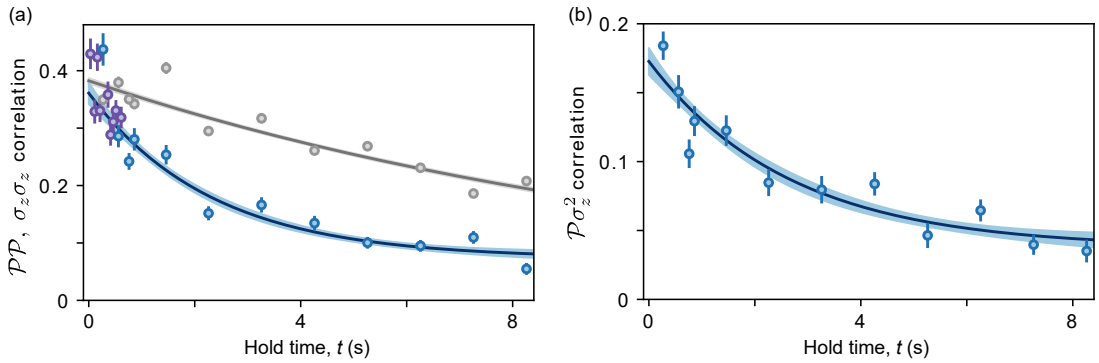

Figure 6.8: Measuring Bell state coherence using correlations. (a) We measure the single atom coherence time by measuring the $g^{(2)}$ (or $\sigma_z\sigma_z$) correlator between isolated atoms (grey), and the Bell state coherence time by measuring the parity-parity ($\mathcal{PP}$) correlator between Bell pairs (blue and purple). Note that, as in Fig. 6.7, the blue data is taken in parallel with the single atom data, and the purple data is taken without the additional rotation about the ϕ axis. Although we do not expect this change in state preparation to affect the coherence time of the resulting state, we separate the data for the two experimental conditions for clarity. (b) To compare Bell pairs to isolated atoms, we compute the $\mathcal{P}\sigma_z^2$ correlator between Bell pairs and atoms. The decay of this correlator is consistent with the expected lifetimes of the Bell pairs and isolated atoms, suggesting that there isn't a significant additional source of decoherence between isolated atoms and Bell pairs. Throughout this figure, the solid lines are exponential fits (allowing for an offset), and the shaded regions are the $\pm 1\sigma$ confidence intervals of the fits.

$\{i, j\}$ to obtain a lower bound for $2|\rho_{00,11}|^2$.[31] We find that the above bound decays with an exponential time constant of 5.0(3) s (Fig. 6.8a), suggesting a Bell state coherence time of $\tau'_B = 10.0(6)$ s (where the $'$ denotes atom-atom coherence, as opposed to atom-laser coherence).[32] This is to be compared to the single atom $g^{(2)}_{ij}$ (also averaged over $\{i, j\}$), which decays with an exponential time constant of 12.2(5) s (Fig. 6.8a), suggesting a Bell state coherence time of the same value, and a single atom coherence time of $\tau' = 24.4(1.0)$ s. We therefore find that the measured Bell state coherence time is close to the expectation based on single atom effects, with possible additional sources of decoherence contributing with a combined time constant of 55(21) s.

The parity-σ_z^2 correlation function $\mathcal{C}^{(\mathcal{P}\sigma_z^2)}_{ij}$ is chosen for a slightly more subtle reason: we are interested in comparing the Bell pairs to a well characterized phase reference (namely single atoms), but under free

[31] Note that the measurements of the various correlation functions are performed with a spin echo, meaning that we expect $\phi^{(i)}_{00,11} - \phi^{(j)}_{00,11} \simeq 0$. Therefore, the decay of this quantity further serves as an estimate of the decay of $2|\rho_{00,11}|^2$

[32] Note that, based on physical grounds, we expect the various correlation functions to decay to zero at long Ramsey interrogation times. However, the hypothesis that the parity-parity correlation data is better fit without an offset is weakly disfavored with a p-value of 0.18. Therefore, in [290], we quote conservative values of a parity-parity decay time of 2.3(4) s, and $\tau'_B = 4.6(7)$ s, based on fit functions that allow for an offset.

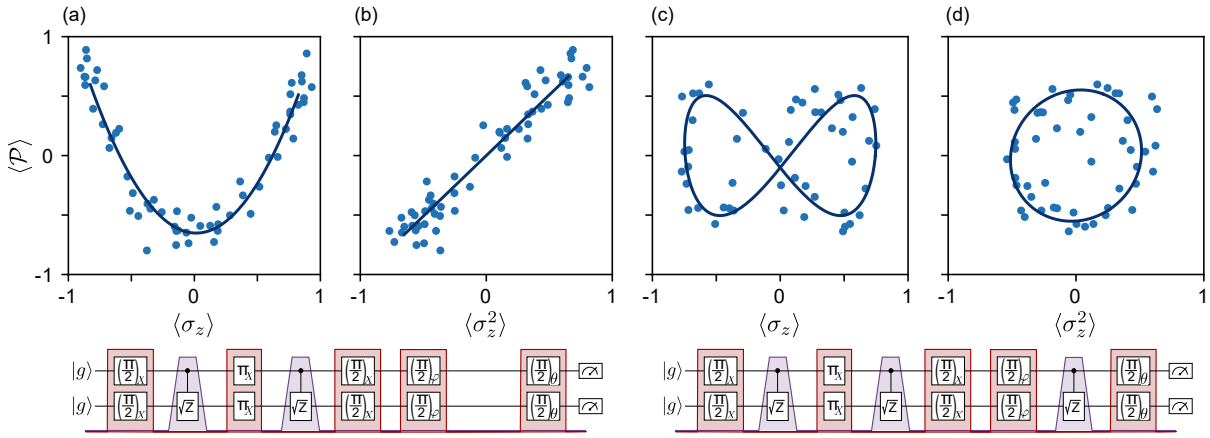

Figure 6.9: Parametric plots of the evolution of Bell pairs and single atoms. (a) Because the phase of the Bell pairs precesses twice as quickly as the phase of a coherent state, a parametric plot of the Bell pair parity ($\langle\mathcal{P}\rangle$) against single atom projection along z ($\langle\sigma_z\rangle$) traces out a $2:1$ Lissajous figure. (b) By squaring the signal from the single atoms (measuring $\langle\sigma_z^2\rangle$), we double the precession frequency, leading to a $1:1$ Lissajous figure. (c, d) Similar measurements with an additional controlled-$\sqrt{z}$ gate (the relevant circuits for state preparation and measurement are shown in the lower panels) applied to the atoms results in a relative phase between the Bell pairs and single atoms, leading to a larger opening angle in the Lissajous patterns. We attempt to apply a phase shift of $\pi/2$, and find the applied phase to be $0.94(7)\pi/2$ via ellipse fitting. Note that the mean has been subtracted from all the above plots, such that their center of mass lies at the origin.

evolution the phase $\phi^{(i)}_{00,11}$ accrues at twice the rate of the phase $\phi^{(i)}_{0,1}$ accrued by a single atom.[33] This effect is illustrated in Fig. 6.9ac, where single-shot measurements of $\mathcal{P}$ for Bell pairs and σ_z for single atoms are plotted parametrically. The resulting parametric plots are $1:2$ Lissajous figures, whose opening angle is controlled by the initial relative phase between the single atoms and the Bell pairs.[34] By squaring the single atom signal, we can generate a signal that precesses at twice the single atom frequency, and is therefore appropriate to compare directly to the Bell state signal. This is reflected in Fig. 6.9bd, where parametric plots of $\mathcal{P}$ and σ_z^2 yield $1:1$ Lissajous figures. Similar to above, averaging over the laser phase yields [290]:

$$\mathcal{C}^{(\mathcal{P}\sigma_z^2)}_{ij} = 2\,|\rho^{(i)}_{0,1}|^2\,|\rho^{(j)}_{00,11}|\,\cos(2\phi^{(i)}_{0,1} - \phi^{(j)}_{00,11}). \tag{6.34}$$

The average of $\mathcal{C}^{(\mathcal{P}\sigma_z^2)}_{ij}$ over $\{i,j\}$ serves as a lower bound on the quantity $2\,|\rho^{(i)}_{0,1}|^2\,|\rho^{(j)}_{00,11}|$, and the decay of this quantity serves as an estimate of the decay rate since $2\phi^{(i)}_{0,1} - \phi^{(j)}_{00,11} \simeq 0$.[35] We measure the above bound to decay with an exponential time constant of 5.3(4) s (Fig. 6.8b). When compared to the $g^{(2)}$ lifetime of 12.2(5) s, we infer that $\tau_B' = 9.5(1.4)$ s from this measurement, which is consistent with the value inferred from the parity-parity measurement.[36]

Note that although our measurements match the expected behavior well at late times, deviations from the expected behavior at early times indicate that two timescales are present in the decay of atomic coherence. This can be seen, for example, when comparing measurements of the various correlations to expectations based on the master equation analysis described above (see Fig. 6.10). In [290], we handle this conservatively by allowing for an offset in the fit function used to extract the decay time, resulting in lower inferred coherence times. However, we suspect that the early time decay is the result of technical noise related to switching the lattice and the magnetic fields to the conditions that are maintained during the Ramsey dark time [290]. This mechanism would result in an initial fixed reduction in coherence, with the subsequent scaling of coherence with respect to Ramsey dark time following the less conservative values listed throughout this section.

[33] $\phi^{(i)}_{0,1} \equiv \phi_i$ in Eq. 5.25. The form here is chosen to make explicit the similarities between the phases in the single- and two-atom states.

[34] We can control this relative phase using the Rydberg gates. For example, we apply an additional controlled-$\sqrt{z}$ pulse, and measure the resulting relative phase to be $0.94(7)\pi/2$, close to the expected value of $\pi/2$.

[35] In the absence of an additional control pulse.

[36] When fitting the parity-σ_z^2 correlator, we do not find the inclusion of an offset to be statistically significant (with a p-value of 0.49). However, we include an offset to provide a conservative estimate in [290] yielding an exponential time constant of 2.7(6) s, and a resulting inferred value of $\tau_B' = 3.5(1.1)$ s.

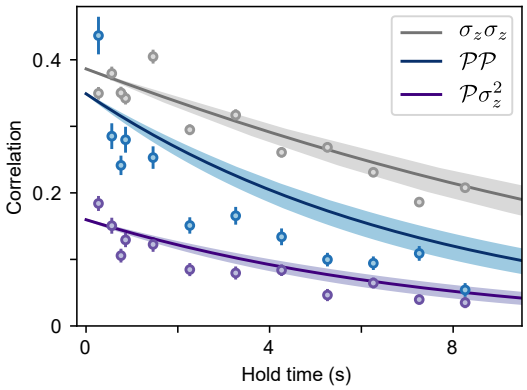

Figure 6.10: Correlation dynamics of entangled states. Similar to the population dynamics in Fig. 6.6b, we can compare the measured correlations to a master equation prediction based only on the measured state prepared after the applied MS gate, and on measured single-atom decay processes. These predictions (solid lines) are compared to the three measured correlation functions in this section. The shaded regions denote $\pm 1\sigma$ confidence intervals in the master equation simulation.

The measurements presented in this section suggest that the coherence of the Bell states (relative to other atoms) is not significantly degraded with respect to the expectation based only on single atom decoherence processes, and therefore that such states could be useful for precision measurements. However, as discussed in Sec. 5.1.2, it is important to note that the benefit associated with using such states for frequency metrology depends on the specific sensing problem at hand. For example, in synchronous measurements where the interrogation time is limited by the coherence of the atoms themselves, such states do not confer a significant benefit. One situation where entangled states can potentially be useful is when one is limited by LO noise, and uses the atomic signal to correct the frequency of the LO [10]. It is important to note that in the experiments in this chapter, the atom-laser coherence is almost certainly limited by common-mode fluctuations of the atomic frequency (due primarily to magnetic field noise), and not by noise in the LO. As a result, in order to realize an advantage from correcting the LO frequency using an entangled state as a reference, more work is required to improve the coherence of the atoms with respect to an ideal phase reference.[37]

6.5 Large entangled states for frequency metrology

Although entangled states involving two particles can already be useful, for example to reject certain kinds of noise in a sensor [233], the central promise of entanglement enhanced metrology is improved scaling of the precision of a sensor with respect to physical resources, as discussed in Sec. 5.1. As a result, we are interested in preparing useful states for frequency metrology involving many entangled particles.

We will comment briefly on experiments we have performed in this direction [96], but defer a detailed discussion to the future thesis of William Eckner. The central idea is to observe that Eq. 6.19 is the well-known one axis twisting Hamiltonian, which generates spin squeezed states [174, 188]. In fact, by staying in the weakly dressed limit where $\Omega \ll \Delta$, one will arrive at Eq. 6.19 from Eq. 6.7 even in the case where many (more than two) atoms are present [37, 117].[38] Therefore, our protocol for applying one axis twisting to large ensembles of atoms (Fig. 6.11) is similar to the gate protocol described in Sec. 6.2 for two atoms.

[37] This is distinct from our demonstrations so far of long-lived coherence between the atoms, since such measurements are, by design, insensitive to common-mode fluctuations of the atomic transition frequency.

[38] When also including the spin echo pulse used in the gate sequence to arrive at Eq. 6.19. Also note that S_z now refers to the projection of the collective spin of all the atoms onto the z axis.

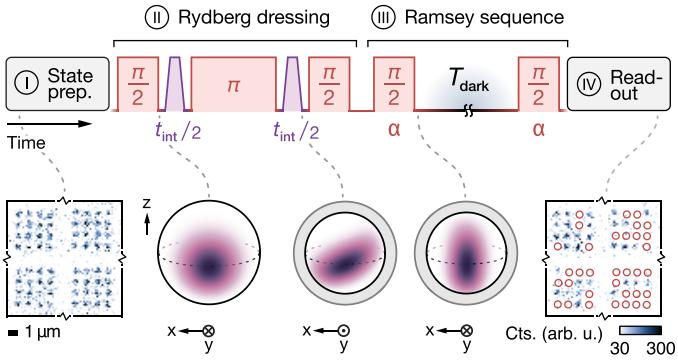

Figure 6.11: One axis twisting in large ensembles of atoms. Coherent states of large ensembles of atoms are prepared on the equator of the Bloch sphere, and then an operation similar to the MS gate applies one axis twisting to the entire ensemble, resulting in a squeezed state with lower variance in its spin projection along a given axis than in the original coherent state. The low variance axis is appropriately oriented, and then injected into a Ramsey sequence as in the measurements in Ch. 5. The resulting measurement can have lower QPN than the equivalent measurement using a coherent state. Note that the pulse diagram and generalized Bloch spheres appearing in this figure are not to scale. The bottom left inset is a single shot image of an atom ensemble after rearrangement, and the bottom right inset corresponds to a projective measurement of this ensemble. The atoms appearing in this image were in the clock state, whereas the missing atoms (highlighted in red circles) were in the ground state.

The states that are generated by evolving under Eq. 6.19 starting from a coherent state on the equator of the Bloch sphere are (ideally) pure states with increased variance in energy in comparison to the original coherent state,[39] and therefore by Eq. 5.7 can achieve enhanced performance in comparison to the SQL. The ultimate limit set by one axis twisting is a reduction in measurement uncertainty by $\propto N^{-1/3}$ in comparison to the SQL, where N is the number of atoms [174]. However, slightly modified protocols involving one axis twisting and global rotations [160, 161] can achieve Heisenberg-limited scaling corresponding to a reduction by $\propto N^{-1/2}$.

The enhancement provided by spin squeezed states is particularly easy to understand, because the mechanism for the reduction in measurement uncertainty is simply a reduction in the variance of the spin projection measured along the axis in which displacements are generated by the signal (without a corresponding reduction in the displacements generated by the signal). This can be characterized by the Wineland squeezing parameter [355]:[40]

$$\xi^2 \coloneqq \frac{N(\Delta J_{\perp,\min})^2}{J^2}, \tag{6.35}$$

where $\boldsymbol{J}$ is the total spin vector (with length J), and $(\Delta J_{\perp,\min})^2$ is the minimum variance achieved when measuring the projection of $\boldsymbol{J}$ onto an axis perpendicular to the direction of $\boldsymbol{J}$. In Fig. 6.12 we show that we do indeed observe such a reduction in variance without significantly reducing J when applying the appropriate evolution, leading to $\xi^2 = -3.6(8)$ dB (in this case for ensembles of 16 atoms). In Fig. 6.13 we use similar states involving ensembles of 16 and 70 atoms in a differential frequency comparison,[41] and achieve stabilities that are respectively 3.52(1) dB and 1.94(1) dB below the SQL.

[39] When properly oriented on the Bloch sphere via a final global rotation

[40] Where any value below 1 indicates an enhancement in comparison to the SQL.

[41] Similar to the one in Sec. 5.4.2.1, but feeding back on the LO to ensure that the atom-laser phase stays near the optimal value for the duration of the measurement, as opposed to averaging the laser phase over all possible values.

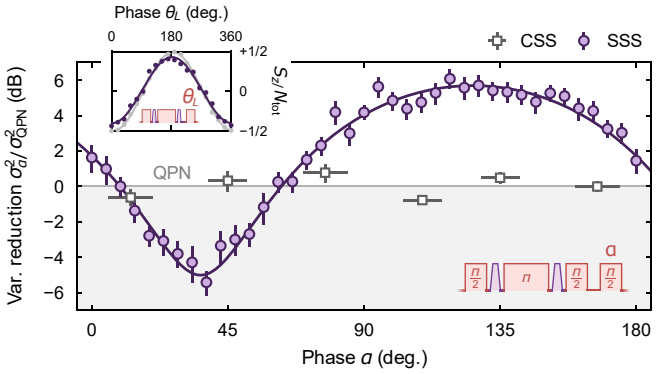

Figure 6.12: Characterizing spin squeezing. Measuring the variance of the spin projection of the atom ensemble σ_a^2, relative to the expectation for a coherent spin state (CSS), as a function of the measurement axis α. Note that there is a measurement axis about which the spin squeezed state (SSS) has significantly lower variance than the limit set by QPN (σ_{QPN}^2). Critically, this reduced variance is not the result of a significantly diminished value of the total spin, which is measured in the inset by measuring the total spin projection S_z/N_{tot} (where N_{tot} is the total atom number) as a function of the laser phase θ_L. Fits to these measurements (solid lines) yield a Ramsey contrast of $C = 0.97(1)$ for the CSS, and $C = 0.83(2)$ for the SSS. The reduction in C in the SSS is not sufficient to counteract the reduction in the variance of the SSS, and so this state can be used in measurements that surpass the SQL.

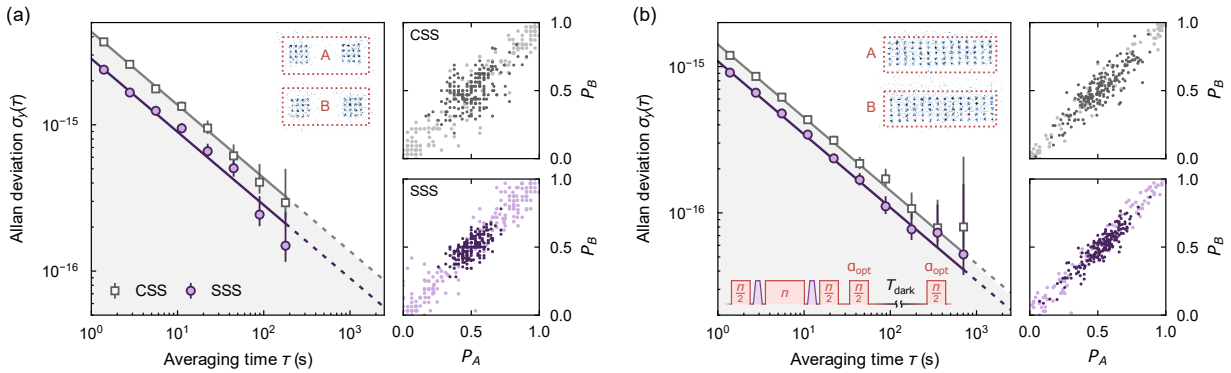

Figure 6.13: Synchronous frequency comparisons using spin squeezed states. We measure the overlapping Allan deviation for differential frequency comparisons between ensembles A and B, using a coherent spin state (CSS) and a spin squeezed state (SSS). In (a), the ensembles each contain 32 atoms, and a Ramsey dark time of $T_{\mathrm{dark}} = 26$ ms is used. The resulting stability of the measurement with the SSS is $2.829(4) \times 10^{-15}(T/s)^{-1/2}$, which is 3.69(2) dB lower than that of the CSS, and 3.52(1) dB lower than the SQL. In (b), the ensembles each contain 70 atoms, and a Ramsey dark time of $T_{\mathrm{dark}} = 54.5$ ms is used. The resulting stability of the measurement with the SSS is $1.087(1) \times 10^{-15}(T/s)^{-1/2}$, which is 2.3(1) dB lower than that of the CSS, and 1.94(1) dB lower than the SQL. The right panels show the excitation probabilities P_A and P_B of atoms in ensembles A and B, with each dark point corresponding to a single experimental trial (only 200 trials are pictured for clarity). The reduction in variance in the SSS is clearly visible. The light points correspond to a similar measurement, in which the laser phase is not locked to the atomic reference signal, resulting in a uniform sampling of the Ramsey fringe. This latter measurement provides an estimate of the contrast of the Ramsey fringe for each measurement setting.

Chapter 7

Single-particle quantum walks

"Nowhere. Everywhere. I'm the sum total of the works, the whole show."
— Neuromancer/Wintermute

Portions of this chapter have appeared in:

A. W. Young, W. J. Eckner, N. Schine, A. M. Childs, and A. M. Kaufman. Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice. **Science**, 377(6608):885–889, Aug. 2022

A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman. An atomic boson sampler. **In review**, 2023

To this point, we have considered the situation where many atomic qubits are pinned in place, allowing us to manipulate their internal degrees of freedom to perform various computational and metrological tasks. In this chapter we will focus on control over the external degrees of freedom of a single atom — namely its position in a tunnel-coupled optical lattice — to enable the study of different kinds of dynamics. We first present a brief review of band theory in Sec. 7.1, and frame our discussion in the language of continuous time quantum walks in Sec. 7.2. In Sec. 7.2, we further demonstrate that such quantum walks can be performed by implanting atoms prepared with optical tweezers into an optical lattice. In Sec. 7.3, we discuss how we calibrate the parameters that determine the behavior of the quantum walks, and in Sec. 7.4 we manipulate these parameters to prepare specific ground states where a single atom is delocalized across the lattice. Next, in Sec. 7.5, we discuss (and demonstrate) how the prepared states and quantum walk dynamics can be used as a resource in various quantum algorithms, with a particular emphasis on spatial search. We conclude in Sec. 7.6 by discussing how the tools developed in this chapter could be adapted to provide universal control over the quantum walk dynamics, and in Sec. 7.7 by commenting on how to interpret the benefits provided by quantum algorithms (based on quantum walks or otherwise).

7.1 Single atoms in a lattice

7.1.1 Band theory

The Hamiltonian for a single atom in a spatially varying potential $V(\boldsymbol{r})$ is:

$$H = -\frac{\hbar^2}{2m}\nabla^2 + V(\boldsymbol{r}), \tag{7.1}$$

where m is the atomic mass. If $V(\boldsymbol{r})$ is a periodic potential, such that $V(\boldsymbol{r}) = V(\boldsymbol{r} + \boldsymbol{R})$ for all Bravais lattice vectors $\boldsymbol{R}$ (and corresponding reciprocal lattice vectors $\boldsymbol{K}$), then the eigenstates ψ and eigenenergies E of Eq. 7.1 can be solved by performing a discrete Fourier transformation:

$$V(\boldsymbol{r}) = \sum_{\boldsymbol{K}} V_{\boldsymbol{K}} e^{i\boldsymbol{K}\cdot\boldsymbol{r}} \tag{7.2}$$

$$\psi(\boldsymbol{r}) = \sum_{\boldsymbol{k}} c_{\boldsymbol{k}} e^{i\boldsymbol{k}\cdot\boldsymbol{r}}, \tag{7.3}$$

where $\boldsymbol{k}$ is a wavevector that runs over all values allowed by the chosen boundary conditions.[1] Plugging the above equations into the time-independent Schrödinger equation ($H\psi = E\psi$), and using the fact that the plane waves form an orthonormal basis, we obtain [11]:

$$\left(\frac{\hbar^2}{2m}|\boldsymbol{k}|^2\right) c_{\boldsymbol{k}} + \sum_{\boldsymbol{K}} V_{\boldsymbol{K}} c_{\boldsymbol{k}-\boldsymbol{K}} = E c_{\boldsymbol{k}}. \tag{7.4}$$

The left-most term in Eq. 7.4 is simply the free space energy, whereas the middle term corresponds to scattering off of the Fourier component of V with wavevector $\boldsymbol{K}$, which couples plane waves with wavevector $\boldsymbol{k}$ to plane waves with wavevector $\boldsymbol{k} - \boldsymbol{K}$. Note that we can restrict $\boldsymbol{k}$ to the first Brillouin zone,[2] since we can always pick some reciprocal lattice vector $\boldsymbol{K'}$ such that $\boldsymbol{k} - \boldsymbol{K'}$ lies in the first Brillouin zone. The key insight in the above derivation is that the periodicity in V causes Eq. 7.1 to factorize into separate eigenvalue problems for each value of $\boldsymbol{k}$ in the first Brillouin zone,[3] since V only couples plane waves which differ by a reciprocal lattice vector, and thus map to the same value of $\boldsymbol{k}$ in the first Brillouin zone. For each such value of $\boldsymbol{k}$ there are an infinite series of solutions to Eq. 7.4, which we label by the band index n. For an infinite lattice, the resulting energy spectrum is given by a set of continuous functions $E_n(\boldsymbol{k})$, which is known as

[1] For example, for a lattice with L unit cells in each direction and cyclic boundary conditions, $\boldsymbol{k} = \sum_i n_i \boldsymbol{b}_i / L$, where $\boldsymbol{b}_i$ is a primitive reciprocal lattice vector, and n_i is an arbitrary integer.

[2] Namely, the region in the reciprocal lattice that is closer to the origin than to any other reciprocal lattice site.

[3] i.e. the Hamiltonian is block diagonal, with each block corresponding to one value of $\boldsymbol{k}$.

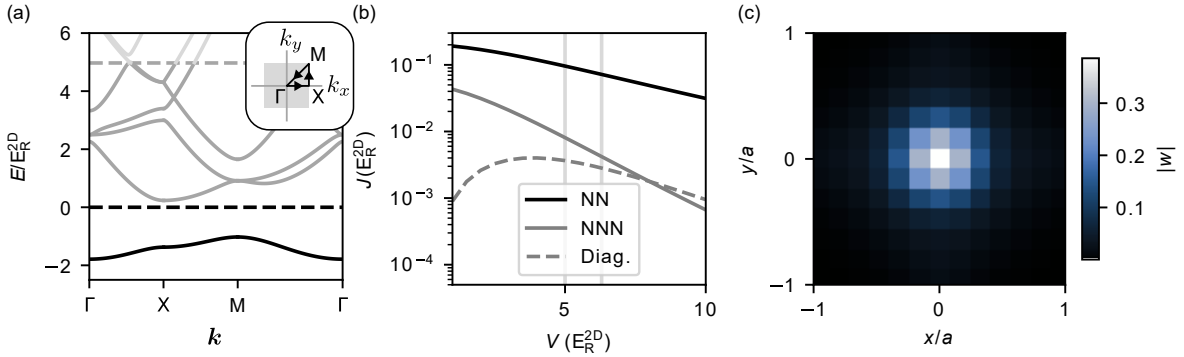

Figure 7.1: Band structure and Wannier functions. (a) Band structure for a bowtie lattice with a depth of $V = 5E_{\mathrm{R}}^{2\mathrm{D}}$, plotted along a triangular path between critical points at the center, edge center, and corner of the first Brillouin zone (Γ, X, and M respectively, see inset), where $\boldsymbol{k}$ is a wavevector and E the energy. $E_{\mathrm{R}}^{2\mathrm{D}}$ is the recoil energy of the 2D lattice. The free space energy is indicated by the dashed black line, and the depth of the axial lattice by the grey dashed line. The inset shows the location of the critical points relative to the first Brillouin zone (shaded region). (b) Under the tight binding approximation, only tunneling processes between nearest-neighbor (NN) sites are significant. However, at the lattice depths we work at (vertical grey lines denote our two typical operating conditions), other processes can become significant. After NN tunneling, the next two dominant processes involve tunneling from one site to its next-nearest neighbors (NNN), or to its neighbors in a diagonal direction (diag.). (c) The amplitude of the Wannier function w in the lattice at a lattice depth of $6.3E_{\mathrm{R}}^{2\mathrm{D}}$. Notice that at this depth, the Wannier function is still well localized relative to the lattice spacing a, and approximately Gaussian.

the band structure of the lattice. The associated eigenstates are the Bloch states $\psi_{n,\boldsymbol{k}}(\boldsymbol{r})$, which, due to the discrete translation symmetry of the problem, are delocalized across the lattice and have the property that $\psi_{n,\boldsymbol{k}}(\boldsymbol{r}+\boldsymbol{R}) = e^{i\boldsymbol{k}\cdot\boldsymbol{R}}\psi_{n,\boldsymbol{k}}(\boldsymbol{r})$ [11].

For most optical potentials, there are only a handful of different relevant values of $\boldsymbol{K}$. For example, for the bowtie lattice potential in Eq. 3.6 the values of $V_{\boldsymbol{K}}$ are:

$$V_{\boldsymbol{K}} = \begin{cases} \frac{V_0}{4}, & \text{if } m_1 = m_2 = 0 \\ \frac{V_0}{8}, & \text{if } |m_1| + |m_2| = 1 \\ \frac{V_0}{16}, & \text{if } |m_1| = 1,\ |m_2| = 1 \\ 0, & \text{otherwise,} \end{cases} \tag{7.5}$$

where V_0 is the lattice depth, $m_1,\ m_2 \in \mathbb{Z}$, and we have introduced the parameterization $\boldsymbol{K} = m_1\boldsymbol{b}_1 + m_2\boldsymbol{b}_2$, where $\boldsymbol{b}_{1,2}$ are the primitive reciprocal vectors of the square 2D lattice. We show the numerically obtained band structure of this lattice for parameters that are relevant to our experiments in Fig. 7.1a.

Note that although our experiments occur in a 3D optical lattice, the axial potential (in the out-of-plane or z direction) is separable from the bowtie lattice. Additionally, the axial lattice is aligned along

gravity and thus the atoms experience a linear gradient of $V(z) = mgz$, where g is the local gravitational field strength. This gradient localizes the atoms to a single site in the axial lattice, and thus a single 2D layer of the 3D lattice, for all parameter regimes of relevance to this thesis.

7.1.2 Wannier functions

Whereas the Bloch states are delocalized across the lattice, the states that we can readily prepare via tweezer implantation involve atoms that are localized on a given site in the lattice. The maximally localized states that can be constructed out of the Bloch states are described by the Wannier functions, which have the form [11]:

$$w_{n,i}(\boldsymbol{r}) = \mathcal{N}^{-1/2} \sum_{\boldsymbol{k}} e^{-i\boldsymbol{k}\cdot\boldsymbol{R}_i} \psi_{n,\boldsymbol{k}}(\boldsymbol{r}), \tag{7.6}$$

where $\mathcal{N}$ is a normalization factor. $w_{n,i}(\boldsymbol{r})$ can be interpreted as the wavefunction of an atom which primarily occupies a specific lattice site i with lattice vector $\boldsymbol{R}_i$ (see Fig. 7.1c). In the limit of a very deep lattice, where each lattice site behaves like an independent harmonic trap, $w_{n,i}(\boldsymbol{r})$ is simply the harmonic oscillator state with motional quantum number n (centered at position $\boldsymbol{R}_i$).

Although the Wannier states form a complete orthonormal basis, they are not eigenstates of H (as defined Eq. 7.1). Considering how H couples the different Wannier functions we define:

$$J_{n,ij} = -\int d^3\boldsymbol{r}\; w_{n,i}(\boldsymbol{r}) H w_{n,j}(\boldsymbol{r}). \tag{7.7}$$

$J_{n,ij}$ is known as a tunneling matrix element, and is associated with the rate at which an atom occupying a given site i in the lattice is tunnel-coupled to another site j.[4]

In practice, our lattice doesn't extend to infinity in all directions. Instead, the finite waists of the lattice beams result in a slowly varying potential (sometimes referred to as a chemical potential for historical reasons) which is approximately harmonic, and confines the atoms to a fixed region in space. Including this chemical potential, and restricting ourselves to only the ground band where $n = 0$, we can rewrite H as:

$$H = -\sum_{\{i,j\}} J_{ij} \left(|i\rangle\langle j| + |j\rangle\langle i|\right) - \sum_i V_i |i\rangle\langle i|, \tag{7.8}$$

[4] Note that, given the definition of the Wannier functions in Eq. 7.6, $J_{n,ij}$ is also directly related to the Fourier transform of the band structure.

where $J_{ij} \equiv J_{0,ij}$, the state $|i\rangle$ refers to the wavefunction $w_{0,i}(\boldsymbol{r})$, V_i is the shift in energy of the state $|i\rangle$ due to the chemical potential, and the first sum runs over all pairs of sites $\{i, j\}$ with appreciable tunnel coupling. Note that although V_i breaks the initially assumed translation symmetry, the harmonic contribution to V_i varies slowly in comparison to the lattice spacing, and so the above analysis remains accurate.[5]

When the Wannier functions are well-localized on a given lattice site, as applies to lattice depths of $V_0 \gtrsim 5E_{\mathrm{R}}$ (where E_{R} is the recoil energy of the lattice, as defined in Sec. 3.3), one typically only needs to consider $\{i, j\}$ that correspond to neighboring sites in the lattice.[6] In this case, and for a region in the lattice where V_i is approximately constant, the Hamiltonian simplifies to:

$$H = -J \sum_{\{i,j\}} \left(|i\rangle\langle j| + |j\rangle\langle i| \right). \tag{7.9}$$

The key result of this analysis is that, when working in the ground band, we can abstract away the exact form of the Wannier functions, and understand the relevant dynamics in terms of a discrete set of locations that a point-like particle can occupy.

7.2 Continuous-time quantum walks

The dynamics of a point-like particle that can hop between a discrete set of locations is conveniently framed in terms of a continuous-time quantum walk (CTQW). A CTQW is defined by a graph $G = \{V, E\}$, which is characterized by a set of vertices $i \in V$ and edges $\{i, j\} \in E$ [272]. G can be represented by its adjacency matrix A, where:

$$A_{ij} = \begin{cases} 1, & \text{if } \{i, j\} \in E \\ 0, & \text{otherwise.} \end{cases} \tag{7.10}$$

We then define that the Hamiltonian for a CTQW is simply:

$$H = -A. \tag{7.11}$$

When considering the square lattice graph where vertices correspond to sites in the lattice, and edges

[5] As a point of reference, under typical operating conditions the difference between V_i on adjacent lattice sites is significantly smaller than the tunnel coupling J between those sites. However, V_i can change by as much as $\sim 10J$ when comparing a site at the center of the lattice to one that is 25 sites away.

[6] Although in our works it is sometimes important to include contributions from tunneling matrix elements that couple diagonally spaced or next-nearest-neighbor sites in the lattice (see Fig. 7.1b).

to the tunnel-coupled, nearest-neighbor sites, Eqs. 7.9 and 7.11 are equivalent up to the proportionality constant J.

To understand the expected behavior of these CTQWs, consider a CTQW on a cycle graph C_N with N total sites (a line with cyclic boundary conditions, see Fig. 7.2). The action of the corresponding Hamiltonian $H_{1\mathrm{D}}$ on a specific site $|i\rangle$, with $i \in [0, N-1]$, is:

$$H_{1\mathrm{D}}|i\rangle = -|(i-1)\%N\rangle - |(i+1)\%N\rangle, \tag{7.12}$$

where % denotes modular division. Based on the same logic as in Sec. 7.1.1, the eigenvectors of $H_{1\mathrm{D}}$ are plane waves $|\psi_k\rangle$ with eigenenergies E_k:

$$|\psi_k\rangle = \frac{1}{\sqrt{N}} \sum_{j=0}^{N-1} e^{ij2\pi k/N} |j\rangle \tag{7.13}$$

$$E_k = -2\cos(\frac{2\pi k}{N}), \tag{7.14}$$

where k is an integer in the interval $[0, N-1]$.

Consider how time evolution under $H_{1\mathrm{D}}$ for time t transfers an atom from site l to site m [272]:

$$\begin{aligned}
\langle m|e^{-iH_{1\mathrm{D}}t}|l\rangle &= \langle m|e^{-iH_{1\mathrm{D}}t} \sum_{k=0}^{N-1} |\psi_k\rangle\langle\psi_k||l\rangle \\
&= \frac{1}{N} \sum_{k=0}^{N-1} \exp\left(i\left((m-l)\,\frac{2\pi k}{N} + 2t\cos\left(\frac{2\pi k}{N}\right)\right)\right) \\
&\simeq i^{m-l} J_{m-l}(2t),
\end{aligned} \tag{7.15}$$

where in the last step we have used the fact that in the limit of large N the summation can be approximated with an integral to yield a result in terms of the Bessel function of the first kind of order $(m-l)$, J_{m-l}.[7] The result is left- and right-moving wavefronts that expand ballistically in time (with a speed of 2 sites/unit time) from the input site l. There is an exponentially suppressed tail outside the wavefronts, and a sinusoidally modulated tail in between the wavefronts (see Fig. 7.2). This is in contrast with the behavior of classical random walks, which exhibit diffusive expansion of a Gaussian probability density distribution.

[7] Note that the large N limit also means that the atom never wraps around the graph. The use of a cycle graph rather than a line graph in this example is just out of convenience, as it yields a simpler expression for the eigenstates of the system.

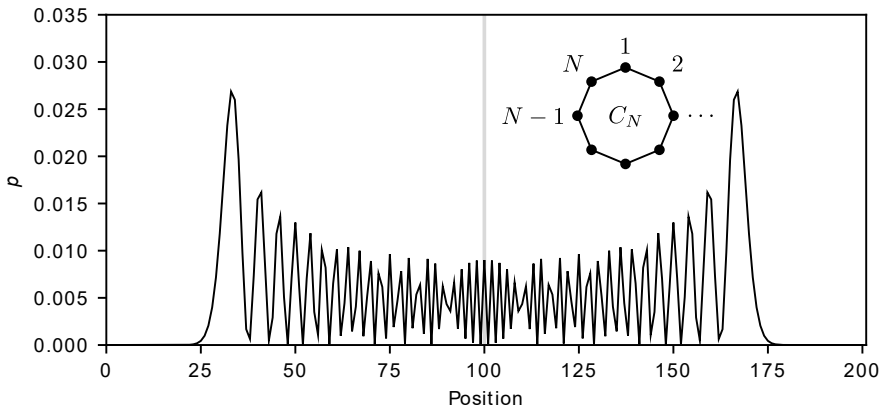

Figure 7.2: A quantum walk on the cycle graph C_N (see inset), with $N = 201$. The probability density p that a quantum walker initialized on a single site (site 100, denoted by the grey vertical line) occupies another site is shown after a unitless evolution time of $t = 35$ (see Eq. 7.15). Note the general features of a sinusoidally modulated inner tail, and an exponentially suppressed outer tail.

7.2.1 Separability of quantum walks in square lattices

When considering quantum walks on a 2D square lattice graph with cyclic boundary conditions, it is useful to work in the basis $|x, y\rangle := |x\rangle \otimes |y\rangle$, where x and y denote the coordinates of the atom in the two orthogonal directions in the square lattice. In this case, the Hamiltonian for the 2D square lattice graph H_{2D} can be written as:

$$H_{\text{2D}} = H_{\text{1D}} \otimes \mathbb{1}_y + \mathbb{1}_x \otimes H_{\text{1D}}, \tag{7.16}$$

where $\mathbb{1}_x$ is the identity operator on the x subspace (and similarly for y). The resulting unitary evolution factorizes into operators that act only on the x and y subspaces:

$$\begin{aligned} U &= e^{-iH_{\text{2D}}t} \\ &= e^{-iH_{\text{1D}}\otimes\mathbb{1}_y t} e^{-i\mathbb{1}_x\otimes H_{\text{1D}}t} \\ &= (U_x \otimes \mathbb{1}_y)(\mathbb{1}_x \otimes U_y), \end{aligned} \tag{7.17}$$

where $U_x = e^{-iH_{\text{1D}}t}$ is a unitary operator that acts only on the x subspace (and similarly for y). Given the separability of the above dynamics, the behavior of a CTQW on a 2D square lattice is precisely that of a CTQW in 1D in each axis of the lattice (with a similar argument applying to higher dimensional square lattices).[8] Along the same lines, if one bins the probability distribution resulting from the 2D CTQW dynamics along one axis, one is simply left with a 1D CTQW along the remaining axis.

We see the above behavior borne out in experiments with a single atom placed in the lattice [366] (Fig. 7.3). For the latest depicted evolution time of $t = 5.0$ ms, the probability density continues to exhibit clear interference fringes, suggesting that the atom has maintained phase coherence while exploring a region spanning approximately 200 lattice sites.

7.3 Calibrating lattice parameters

Our characterizations of the lattice potential primarily rely on either fitting CTQWs of individual atoms, or using the atoms as spectroscopic probes of the intensity of the light field at the position of the

[8] Note that the above arguments do not hold if there are terms that directly couple x to y in the Hamiltonian, for example if there are diagonal tunneling matrix elements, or significant local energy shifts that vary site to site in the lattice and thus depend on a specific set of x, y coordinates.

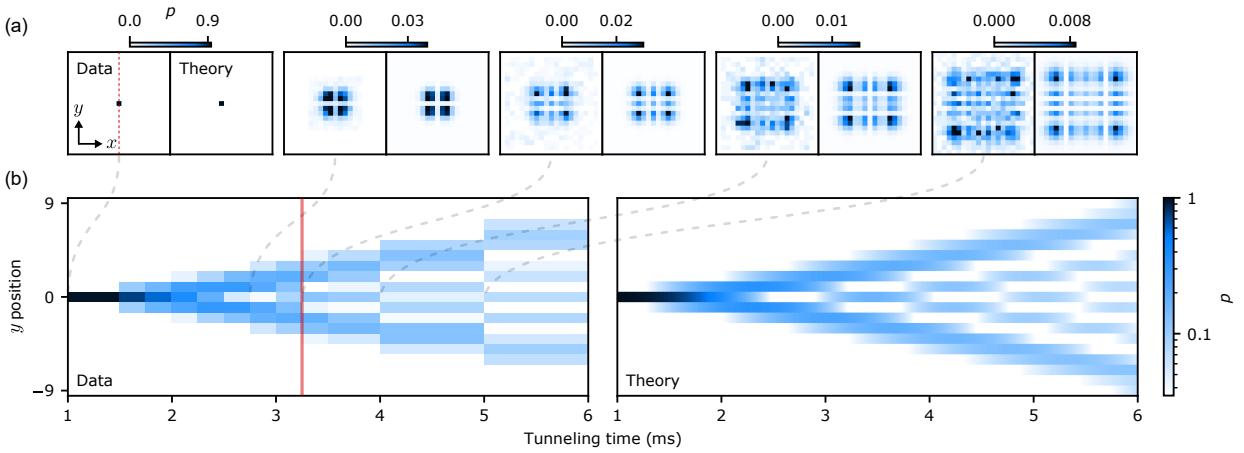

Figure 7.3: Quantum walks in an optical lattice. (a) Atoms implanted in the lattice undergo continuous-time quantum walks in 2D, such that the probability density p corresponding to their measured position (left) exhibits ballistic expansion, and wave-like interference. The behavior of these walks is in good agreement with theory using fitted values of the tunneling energy (right) up to an evolution time of 5 ms, where the atoms have coherently explored a region spanning ~ 200 lattice sites. Each pixel in these plots represents a single lattice site. (b) Tracing out one dimension (in this case the x axis) of a 2D quantum walk on a square lattice, yields a 1D quantum walk along the remaining axis (in this case y, left), which is also in good agreement with theory (right) and more clearly illustrates the ballistic expansion of the wavefunction. Data to the left of the red line is gathered in parallel by preparing atoms in nine different regions in the lattice, since inhomogeneity between regions does not play a significant role for short evolution times. Data to the right of the red line is gathered using a single atom that is placed at the center of the lattice.

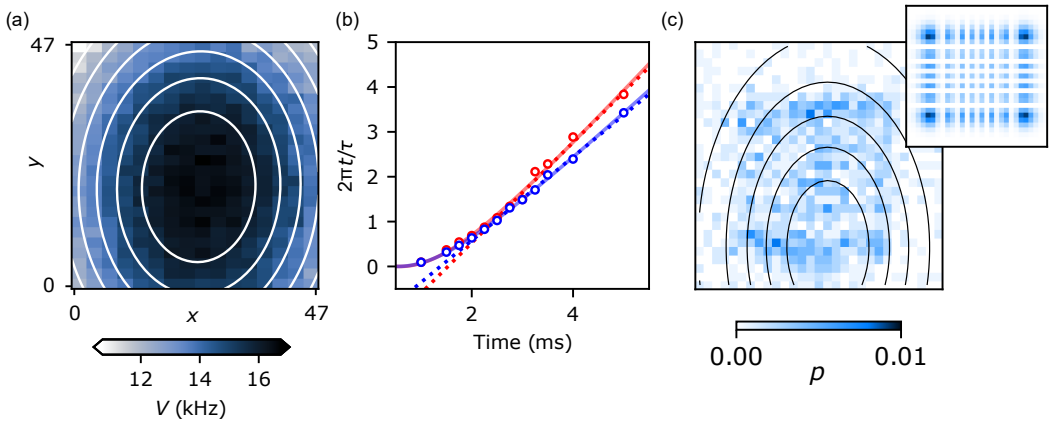

Figure 7.4: Calibration measurements of the lattice potential. (a) Spectroscopic measurements of the lattice-induced AC Stark shift of the $^1S_0 \leftrightarrow {}^3P_1$ transition are used to characterize the lattice depth V as a function of position in the lattice (x and y are in units of the lattice spacing), here pictured at the ratio of 2D and axial lattice depths that is typically used during tunneling. White contours correspond to a Gaussian fit to this potential, which is in good agreement with the expected lattice beam waists. (b) Further calibrations are performed by measuring single particle CTQWs as a function of evolution time (as in Fig. 7.3). At each evolution time, the measured atomic densities are fit to a model that allows for independent values of the unitless tunneling times t_x/τ and t_y/τ in the x and y directions (red and blue circles respectively). At late evolution times, the extracted values grow linearly with evolution time, corresponding to constant values of the tunneling energy of $J_x/\hbar = 2\pi \times 178.1(1.1)$ Hz and $J_y/\hbar = 2\pi \times 148.3(1.1)$ Hz in the two axes (dashed lines are linear fits excluding the first 4 points). At early times the fitted tunneling times deviate from the linear expectation due to the slow settling time of the lattice intensity control. (c) At late evolution times, in this case at 5.2τ, the atoms begin to sample the curvature of the lattice, and the probability density deviates substantially from the prediction in a flat lattice (inset). This behavior is qualitatively consistent with the expectation given the independently characterized lattice potential (contours denote equipotential surfaces of the lattice).

atoms (we describe some more sophisticated characterization techniques in Sec. 8.4)

To spectroscopically measure our lattice potential, namely the chemical potential terms V_i appearing in Eq. 7.8, we change the angle of the magnetic field in the experiment such that the lattice is not magic for the $^1S_0 \leftrightarrow {}^3P_1$ transition [366]. By pinning atoms in the lattice, and performing spectroscopy on the $^1S_0 \leftrightarrow {}^3P_1$ transition, we can measure the local AC Stark shift imposed on the atoms and thus the effective chemical potential on each site in the lattice (see Fig. 7.4a). We repeat the above measurements for the axial and 2D lattices individually, as well as with both lattices on at the ratio of lattice depths that is relevant for tunneling. Note that these measurements are typically performed at $\sim 10-100\times$ higher intensities than those used during the tunneling experiments in order to enhance their signal to noise. We confirm via direct measurements of the power in the optical lattice beams that the ratios between the lattice beam powers stay fixed when operating at different power levels.[9]

The above spectroscopic measurements provide a relative characterization of the chemical potential terms, but it remains to set an overall energy scale for the Hamiltonian. To directly characterize the tunneling rate in the lattice, we perform a maximum likelihood fit of the measured probability density distributions as a function of evolution time t to propagation under Eq. 7.9, assuming a perfectly flat lattice with constant V_i.[10] The free parameters in these fits include independent values of the unitless tunneling times tJ_x/h and tJ_y/h in the x and y axes (Fig. 7.4b), as well as an overall contrast and offset. The offset does not evolve appreciably over the evolution times explored in this chapter, and is consistent with overall error rates due to imperfect imaging and cooling.

The data exhibits constant values of the tunneling energy at intermediate times, reflecting the expected ballistic expansion of the atomic wavefunction. At early times, the data deviates from this linear expansion due to the relatively slow 0.7 ms settling time of the intensity servo used to control the 2D lattice depth, which results in artificially low inferred tunneling rates. At late times ($\gtrsim 4$ ms), the assumption of uniform V_i breaks down, leading to a reduction in the fitted contrast. For the purposes of extracting the tunneling energy, we simply exclude the late time data from our fit. By either fitting the early- to intermediate-time

[9] We make the assumption that the transmission of the glass vacuum cell is independent of optical power. For typical optical intensities at the location of the cell in our experiments, we do not expect effects like thermal lensing to play a role.

[10] Note that including corrections to V_i does not significantly alter the result of these fits.

values of tJ_x/h and tJ_y/h versus t with a hyperbola, or fitting a line to only intermediate-time data, we can extract the tunneling energies. For the work described in this chapter, the lattice has tunneling energies of $J_x/\hbar = 2\pi \times 178.1(1.1)$ Hz and $J_y/\hbar = 2\pi \times 148.3(1.1)$ Hz near the lattice center. This is in good agreement with an ab initio calculation using independent calibration data, specifically: the measured lattice beam waists and powers, the measured transmission of the vacuum cell, the inferred lattice spacing given images of lattice-trapped atoms, and the magnification of the imaging system. Note that the discrepancy between J_x and J_y is because, as a result of our alignment procedure, the lattice is not perfectly square (see Sec. 3.3).

We are able to set an absolute energy scale for V_i (see Fig. 7.4a) by combining our measurements of J_x and J_y with the calculated band structure of our lattice as a function of lattice depth (as described in Sec. 7.1). These calculations also provide a way to convert between the measured chemical potential V_i and the local value of the tunneling energy J_{ij} across the entire lattice, under the assumption that the chemical potential varies smoothly across the lattice. We find this assumption to be accurate given the spectroscopic measurements of V_i. Constructing a full model of our Hamiltonian via the above calibrations yields predictions that are consistent with the late time single atom CTQW data, building confidence in this model. Additionally, we note that inserting calibrated values of V_i into the fit function does not change the inferred values of J_x or J_y in a statistically significant way. For the remainder of this chapter, we define $J_0/\hbar = 2\pi \times 163$ Hz to be the mean of J_x and J_y. This sets a convenient energy scale, and a corresponding characteristic evolution time of $\tau = h/J_0 = 6.12$ ms, for discussing subsequent results.

For the tunneling conditions relevant to this chapter, we operate at a 2D lattice depth of $5.0\mathrm{E}_\mathrm{R}^{2\mathrm{D}}$ (where $\mathrm{E}_\mathrm{R}^{2\mathrm{D}}$ is the lattice recoil energy, see Sec. 3.3).[11] Although this depth is relatively shallow, based on the calculations in Sec. 7.1 we conclude that the effect of diagonal and next-nearest-neighbor tunneling is negligible for subsequent results in this chapter.

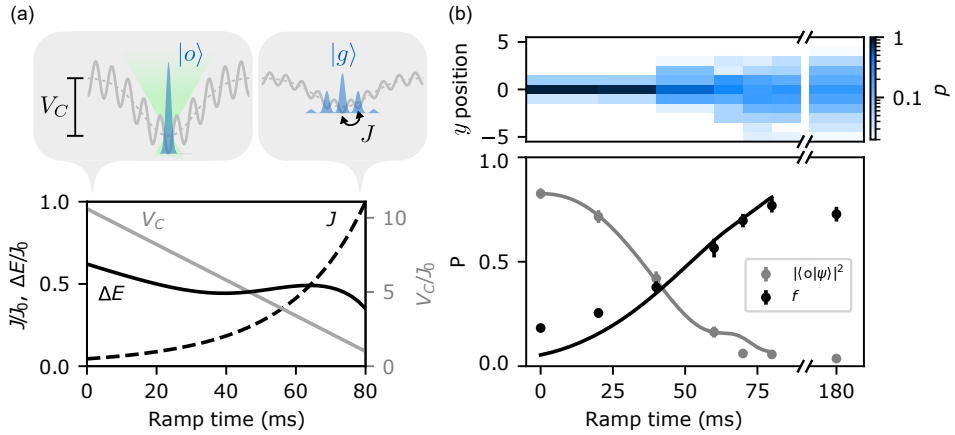

Figure 7.5: Adiabatic preparation of an atom in the ground state of a tunnel-coupled lattice. (a) An atom implanted in the lowest-energy site in a deep lattice with no tunneling (left callout), which we label as being in state $|o\rangle$, can be adiabatically connected to the ground state $|g\rangle$ of a shallow lattice with tunneling (right callout) via an adiabatic ramp of the tunneling energy J and the depth of the confinement tweezer V_C. In the callouts, the grey dashed line denotes the potential provided by the confinement tweezer, and the solid grey line the sum of the tweezer and lattice potentials. The preparation tweezer used to initially implant the atom in $|o\rangle$ is shown schematically in green. By ramping J and V_C together we can maintain a roughly fixed energy gap ΔE between the ground and first excited states of the system throughout the ramp. This relaxes the requirements on adiabaticity, and substantially increases the fidelity with which we are able to prepare the lattice ground state. (b) As an atom evolves under the above ramp, its amplitude spreads over many sites in 2D. Here, the x coordinate has been traced out for illustration purposes, showing the spread in the atom's y position (top). During this ramp, the population on the initial site (grey points) decreases, and the overlap f between the classical probability distributions corresponding to the prepared state $|\psi\rangle$ and the expected lattice ground state $|g\rangle$ (black points) increases, eventually reaching 76.9(3.3)%. This value is in reasonable agreement with theory (solid lines) given the independently measured parameters in our ramp, and an overall scale factor to account for filtering and loss due to imperfect preparation of the atoms in their 3D motional ground states. This serves as an upper bound on the fidelity with which we can prepare the lattice ground state, but does not certify any phase coherence between components occupying different sites in the lattice. Nonetheless, this state is not observed to significantly evolve even after more than 100 tunneling times (right side of the broken axis), further suggesting that this is indeed the lattice ground state.

7.4 Preparing the lattice ground state

Instead of preparing Wannier states by implanting tweezer-trapped atoms in the lattice, one may be interested in preparing the ground state of Eq. 7.8. In order to do this, we take advantage of the fact that we can programmably vary the local chemical potential V_i appearing in Eq. 7.8 by shining optical tweezers on the lattice. Specifically, we use a large 813 nm tweezer (we will refer to this as the "confinement tweezer") with tunable overall depth V_C, and a nearly-Gaussian profile with a beam waist of $5.8a$ (where a is the 2D lattice spacing). At fixed J, the lattice ground state $|g\rangle$ is close to the uniform superposition of the atom occupying all sites in the lattice, except with an approximately Gaussian envelope with a width determined by the value of V_C (see Fig. 7.6).[12]

To prepare the state $|g\rangle$ we implant an atom in the site o, which is the deepest site of the combined potential generated by the lattice and the confinement tweezer ($V_o > V_i \; \forall \; i \neq o$). The prepared state $|o\rangle$ is the ground state of the system when the lattice is deep, and $J \ll V_C$. $|o\rangle$ can be adiabatically connected to $|g\rangle$ in a shallow lattice via a ramp of the tunneling energy (Fig. 7.5a). In practice, we perform a linear ramp of the lattice depth V_L as a function of time t, resulting in a nonlinear ramp in the tunneling energy. We also ramp the depth V_C of the confinement tweezer to maintain an approximately constant value of the energy gap ΔE between the ground and first excited states during the ramp, which substantially relaxes the requirements on adiabaticity and improves the fidelity with which we can prepare $|g\rangle$. The observed evolution during this ramp is in reasonable agreement with theory (Fig. 7.5b), where the atoms start out primarily in $|o\rangle$, and smoothly delocalize over many sites over the course of an 80 ms-long adiabatic ramp of V_C and J.

The prepared state $|\psi\rangle$ can be compared to $|g\rangle$ by computing the overlap between their populations, or the classical fidelity $f = (\sum_i \sqrt{p_{\psi,i} p_{g,i}})^2$, where $p_{\psi,i}$ and $p_{g,i}$ denote the populations on site i in states $|\psi\rangle$ and $|g\rangle$ respectively. This constitutes an upper bound on the fidelity with which we have prepared $|g\rangle$ of 76.9(3.3)%, but does not certify any phase coherence between the amplitudes occupying different sites.

[11] At these shallow depths there is only a single bound band in the 2D lattice, and so atoms occupying higher bands are typically interpreted as lost, either because they actually leave the lattice, or leave the region of the lattice that we analyze in our images.

[12] We will sometimes refer to $|g\rangle$ as the "resource state," for reasons which will become clear in Sec. 7.5.

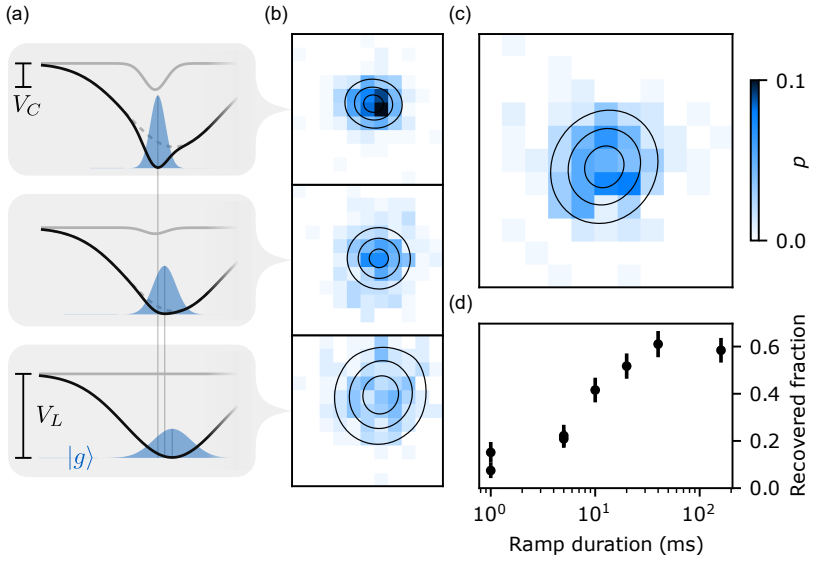

Figure 7.6: Robustness of ground state preparation to misalignment. (a) Cartoons showing how the combination of misalignments between the confinement tweezer (grey curves) and lattice center (dashed grey curves), and changes in their relative depth, can lead to a shift in the minimum (vertical grey lines) of the overall potential experienced by the atoms (black curves), and a corresponding change in the size and position of the atomic ground state $|g\rangle$ (blue shaded regions). These diagrams are for visualization purposes only, and are not to scale. Furthermore, they should be understood as an average envelope spanning many lattice sites, where modulations on the scale of the lattice period have been omitted for clarity. (b) With poor centering between the confinement tweezer and lattice center, measurements of the atomic probability density after performing an adiabatic ramp to the ground state reveal states whose position (and size) vary with the final confinement tweezer depth V_C. From top to bottom, these depths are: 2.2 J_0, the typical depth used for resource state preparation of 1.0 J_0, and 0.0 J_0. (c) Typical atomic probability density in the resource state, with the confinement tweezer well-aligned to the lattice center. Color scale is shared between (b) and (c). (d) As an additional test of the resource state preparation, we perform a forward ramp to the resource state, and then a reverse ramp with variable duration. For longer ramps we recover 57(5)% of the atoms, whereas for shorter ramps the evolution is diabatic, and very little of the atomic probability density returns to the central site, beyond what was already present on that site in the resource state.

However, the prepared state is not observed to substantially evolve over $t \gg \tau$, and the adiabatic ramp can be reversed to recover 57(5)% of the atoms in $|o\rangle$ (excluding loss due to filtering of hot atoms, see Sec. 7.3). This suggests that $|g\rangle$ has been prepared with a fidelity of 76(3)%, in agreement with the bound set by the classical fidelity.

If the center of the confinement tweezer and the lattice are in different locations, the adiabatic ramp also results in a shift in the minimum of the combined potential experienced by the atoms (since the relative depths of the different optical potentials change). For large misalignments, this shift causes the state prepa-

ration to fail. However, for misalignments of $\lesssim 5a$, the center position shifts adiabatically during the ramp, leading to the successful loading of a ground state that is displaced from the initial site that the atom was loaded in (Fig. 7.6ab).

This robustness against misalignment is a useful tool. Because the alignment between the two tweezer arrays is exceptionally stable (see section 3.3.5), the confinement tweezer can be used as an intermediary between the preparation tweezer and the lattice, allowing one to load the lattice ground state even when the tweezers are significantly misaligned from the lattice center (Fig. 7.6b). Specifically, an appropriate ramp in the depths of the confinement tweezer and 2D lattice results in a self-calibrating ramp of the center of the combined optical potential formed by the tweezers and the lattice (Fig. 7.6a). This works even when the confinement tweezer is fully extinguished at the end of the ramp, in which case the center position starts at the location of the implanted atom, and ends at the location of the lattice center.[13] The resulting state (Fig. 7.6b, bottom) behaves similarly to the case when the confinement tweezer is not fully extinguished: it does not significantly evolve over time, and 50(14)% of the atom population can be recovered in the initial site when reversing the ramp. Critically, we measure no atoms in the initial site after a similar set of ramps without the confinement tweezer, meaning that the state preparation fails completely in this scenario.

While we do not do full state tomography to characterize $|\psi\rangle$, there are several additional pieces of evidence indicating that $|\psi\rangle$ is close to $|g\rangle$. First, the size and position of the prepared state is consistent with expectations based on independent characterizations of the various optical potentials in the experiment. For example, the data in Fig. 7.6b upper and middle are taken back to back, with final confinement tweezer depths of $2.2J_0$ and $1.0J_0$ respectively. In this case we expect a ratio in the $1/e^2$ radius of the prepared states of 0.78, and observe 0.83(11). Secondly, we expect forwards and reverse ramps to recover atoms in the central site only when the evolution is adiabatic, and genuinely prepares the lattice ground state. When performing these ramps, the shorter the duration of the reverse ramp, the fewer atoms return to the central site (Fig. 7.6d), as expected. The same is also true when both the forward and reverse ramp times are

[13] A similar shift in the combined center of the axial and 2D lattices can occur. We have attempted to mitigate this by ramping the axial and 2D lattice depths together, maintaining a constant ratio, but this places more stringent requirements on the adiabaticity of these ramps, and results in poorer overall performance. Instead, because these potentials have significantly less curvature than the confinement tweezer, we find it sufficient to align the tweezers to the combined center of the lattices at the final ratio used during tunneling.

varied. Moreover, when the prepared state is perturbed by a quench in V_o (namely a sudden modification of the chemical potential on the central site in the lattice by a few J_0), we expect, and observe, that the reverse ramp also fails despite the fact that immediately after the quench there is more population on the central site than is present after only the forward ramp (see Sec. 7.5.4). This is because the quench results in evolution to a state that is not adiabatically connected to the ground state of the lattice when tunneling is turned off, where the atom sits on the central site.

We have shown so far that we can implement CTQWs of single atoms in an optical lattice, and provide evidence that we can prepare the ground state of a given CTQW Hamiltonian. Next, we discuss how these tools can be combined to perform certain useful computational tasks.

7.5 Search by quantum walk

CTQWs provide a convenient framework for understanding a variety of quantum algorithms. Remarkably, even when restricting oneself to undirected and unweighted graphs, CTQWs are capable of universal quantum computation [59, 64]. This simple but intuitive framework has inspired new quantum algorithms, including those for spatial search [62], graph traversal [60], element distinctness [6], and formula evaluation [104]. In this section, we focus on a spatial search algorithm proposed by Childs and Goldstone [63], and show how this algorithm can be performed with quantum walks of a single atom.

7.5.1 An analog analogue of Grover's search

Spatial search by quantum walk is closely related to Grover's unstructured search algorithm [132]. The problem statement for unstructured search is as follows: Given an unstructured database with items $i \in [0, N-1]$, and access to an oracle function f which has the property

$$f(i) = \begin{cases} 1, & \text{if } i = w \\ 0, & \text{otherwise,} \end{cases} \tag{7.18}$$

how can one find the marked item w with the fewest number of queries to f?

In the classical case, one can do no better than $O(N)$ queries, since on average $N/2$ items must be checked before the marked item is found. However, it is possible to do better if one is allowed to query f in

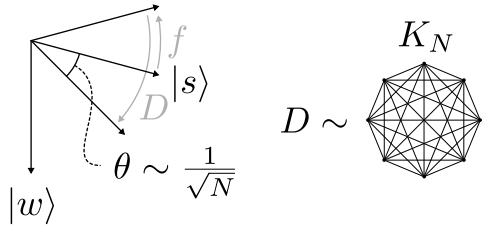

Figure 7.7: Operating principle of Grover's search. Starting from the resource state $|s\rangle$, one iteratively applies the oracle f and the Grover diffusion operator D. Note that D is generated by a CTQW on the complete graph K_N, and so this operation can be implemented using the appropriate CTQW dynamics. The operations f and D leave us in the subspace spanned by $|s\rangle$ and the marked state $|w\rangle$, and their combined action is to rotate the prepared state by an angle $\theta \sim 1/\sqrt{N}$ (where N is the number of states in the search space) closer to the marked state $|w\rangle$. The algorithm terminates after $O(\sqrt{N})$ such rotations, at which point the prepared state is approximately $|w\rangle$, whose value can be read out in a single projective measurement.

superposition. In this case, we redefine f as a unitary operator:

$$f|i\rangle = \begin{cases} -|w\rangle, & \text{if } i = w \\ |i\rangle, & \text{otherwise,} \end{cases} \tag{7.19}$$

where items i are now states $|i\rangle$ in some Hilbert space $\mathcal{H}$. Grover's search (see Fig. 7.7) proceeds by preparing the uniform superposition state over all items:

$$|s\rangle = \frac{1}{\sqrt{N}} \sum_i |i\rangle, \tag{7.20}$$

and sequentially applying f, and then the diffusion operator $D = \mathbb{1} - |s\rangle\langle s|$. The result is a coherent build up of amplitude in the state $|w\rangle$ which is of order unity after $O(\sqrt{N})$ applications of f and D, at which point one can measure the system projectively to identify the state $|w\rangle$. This algorithm yields a quadratic enhancement in comparison to the number of queries to f required in the classical case.

One can translate Grover's search to the setting of continuous, Hamiltonian evolution [105]. In this case, we have an oracle Hamiltonian H_w, a diffusion Hamiltonian H_D, and a combined Hamiltonian H:

$$H_w = -|w\rangle\langle w| \tag{7.21}$$

$$H_D = -N|s\rangle\langle s| \tag{7.22}$$

$$H = H_w + \gamma H_D, \tag{7.23}$$

where we have introduced a parameter γ in order to tune the relative strength of H_w and H_D. The resulting time evolution is reminiscent of standard Grover's search, where H_w results in an inversion of the $|w\rangle$

component of the input state, and H_D an inversion about the uniform superposition state $|s\rangle$. The difference is that in the case of analog Grover's search these operations are applied continuously and simultaneously.

To generalize the query complexity to the continuous case, we want some sense in which the application of H_w is the resource whose use we want to minimize. A natural choice is to minimize $||H_w||t$ [272], where t is the required evolution time of the algorithm, and $||\cdot||$ refers to the spectral norm of a matrix. We will typically enforce that $||H_w|| = 1$, and refer to t as the runtime of the algorithm.

Analog Grover's search proceeds by preparing the state $|s\rangle$, and evolving the system under H for the appropriate choice of γ. Notice that H does not take us out of the subspace spanned by $|s\rangle$ and $|w\rangle$. For $\gamma = 1/N$ the ground and first excited states of H are $|\pm\rangle = (|w\rangle \pm |s\rangle)/\sqrt{2}$ up to corrections of $O(N^{-1/2})$, and separated by an energy gap of $\Delta E = 2/\sqrt{N}$ [63] (see Fig. 7.8b). The resulting dynamics correspond to Rabi oscillations between the state $|s\rangle$ and $|w\rangle$. As a result, one can simply wait for time $\pi\sqrt{N}/2$, then projectively measure the system to identify $|w\rangle$.[14] This algorithm has a runtime of $O(\sqrt{N})$ in comparison to $O(N)$ for a classical continuous time algorithm subject to similar constraints [63], and thus maintains the quadratic speedup associated with Grover's search (see Sec. 7.7). It can be shown that this quadratic runtime is optimal [105].

Notice that $H_D = -A_{K_N}$, where A_{K_N} is the adjacency matrix of a complete graph K_N. This being the case, analog Grover's search can be thought of as a CTQW on K_N in combination with the application of some oracle Hamiltonian H_w.

7.5.2 Spatial search

One might expect that the quadratic speedup exhibited by analog Grover's search relies on a highly connected Hilbert space, where the CTQW on a complete graph K_N allows one to quickly move from one item in the search space to another. In a seminal result, Childs and Goldstone proved this to be false [63]: one can remove almost every edge from K_N while still maintaining a quadratic speedup [58]. Specifically, one can consider the more restrictive scenario (in comparison to unstructured search) of spatial search, where each item in the search space is connected to a constant number of other items, enforcing a sense of

[14] Critically, these dynamics are completely independent of the specific choice of $|w\rangle$, which is what allows one to use them to search for an unknown value of w.

locality [63]. Replacing K_N with $d > 4$ dimensional square lattice graphs (with cyclic boundary conditions) results in a runtime of $O(\sqrt{N})$ and a success probability of $O(1)$ per measurement for the appropriate choice of γ. This saturates the speedup attained by Grover's search despite going from a graph with $O(N^2)$ edges to one with $O(N)$ edges. In $d = 4$ a reduced energy gap leads to a runtime of $O(\sqrt{N \log N})$, and oscillations with imperfect contrast lead to a reduced success probability of $O(1/\log N)$ per measurement. For $0 < d < 4$ the algorithm continues to work qualitatively, but without a significant speedup in comparison to classical search protocols. However, a speedup can be restored in $d \geq 2$ by using a discrete-time quantum walk [7], or by incorporating a single additional spin degree of freedom [62].

7.5.3 Boundary conditions

To understand the asymptotic scaling of CTQW-based search algorithms, it is instructive to solve for the spectrum of the search Hamiltonian as a function of the relative strength of the diffusion and oracle terms. Similar to Eq. 7.23, just with an arbitrary graph G as opposed to K_N:

$$H_w = -|w\rangle\langle w| \tag{7.24}$$

$$H_D = -A \tag{7.25}$$

$$H = H_w + \gamma H_D, \tag{7.26}$$

where A is the adjacency matrix of G.

In Fig. 7.8 we study the spectrum of H as a function of γ for several different graphs. In each case we define the resource state, or equivalently the ground state of H_D, to be $|g\rangle$. In the typically studied case of cyclic boundary conditions, the eigenstates of the lattice are plane waves, and so all eigenstates have equal overlap with the oracle, and are affected by varying γ (Fig. 7.8a). In this setting, the search algorithm can be interpreted as the formation of a subspace composed of the ground and first excited states of the system, $|\psi_0\rangle$ and $|\psi_1\rangle$, which is mostly supported by the states $|g\rangle$ and $|w\rangle$.

As in Sec. 7.5.1, for $G = K_N$, and at the critical value of $\gamma = 1/N$ (for large N), $|\psi_0\rangle$ and $|\psi_1\rangle$ have approximately equal contributions from $|g\rangle$ and $|w\rangle$ (Fig. 7.8b). The overlap of $|\psi_0\rangle$ and $|\psi_1\rangle$ with $|g\rangle$ and $|w\rangle$ sets the contrast of the oscillations between $|g\rangle$ and $|w\rangle$ during the search procedure, whereas the

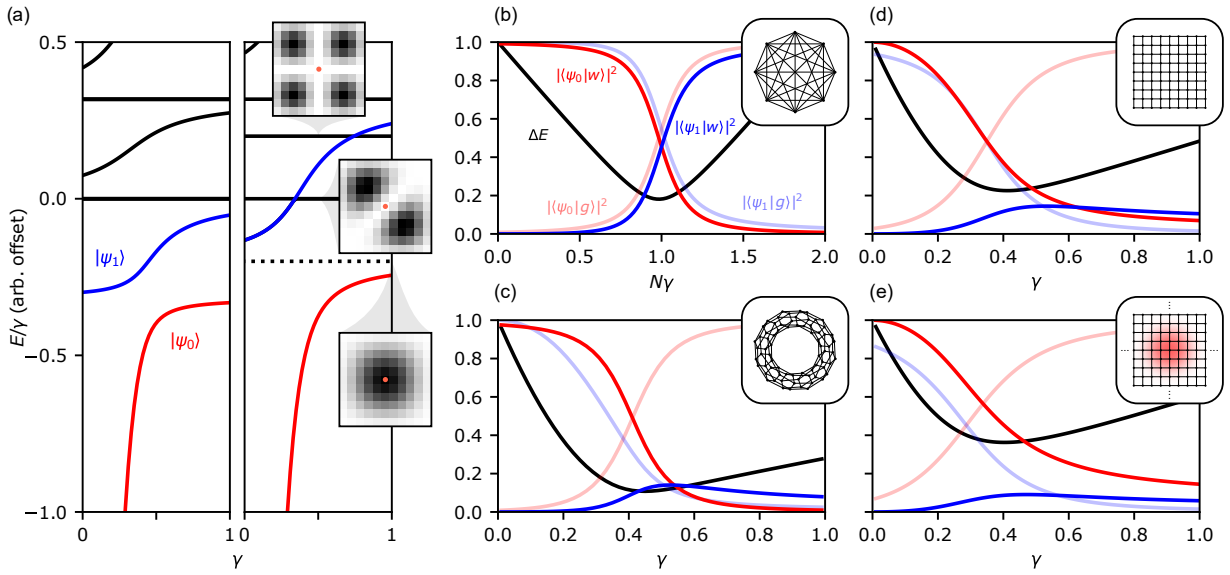

Figure 7.8: Level structure of the spatial search Hamiltonian. (a) Exploring the level diagram of the search Hamiltonian as a function of the relative strength γ of the diffusion and oracle terms, we identify a subspace spanned by states $|\psi_0\rangle$ and $|\psi_1\rangle$. This subspace is primarily supported by the ground state of the diffusion Hamiltonian on the graph, $|g\rangle$, and the marked state, $|w\rangle$. In the case of a lattice with cyclic boundary conditions (left), this subspace corresponds to the ground (red) and first excited (blue) states. In the case of closed boundary conditions (right) and for a Gaussian confining potential (not depicted), the eigenstates of the lattice have spatial structure, resulting in states (top and middle callouts) that do not overlap with a given oracle (positions marked by the red dots). In this case $|\psi_0\rangle$ and $|\psi_1\rangle$ (red and blue) need not be the ground and first excited states, and are instead identified by their overlap with $|g\rangle$ (dotted line and lower callout) and $|w\rangle$. The behavior of $|\psi_0\rangle$ and $|\psi_1\rangle$ vary for different graph geometries, both in terms of their overlap with $|g\rangle$ and $|w\rangle$, and their energy separation ΔE. Insets in (b-e) denote the relevant graph connectivity and boundary conditions, but are for visualization purposes only and are not accurate in terms of graph size. These conditions are: (b) A fully connected graph with $N = 121$ nodes, as applies to the case of continuous-time Grover's search. (c) A 2D square lattice with cyclic boundary conditions and $N = 11 \times 11$. (d) A 2D square lattice with closed boundary conditions and $N = 11 \times 11$. (e) A 2D square lattice with an additional Gaussian confining potential with a waist of 5.8 lattice sites and a depth of $1.0 J_0$, as provided by the confinement tweezer during the search procedure in Sec. 7.5.4. For (b, c), translational symmetry dictates that the behavior of the system is independent of oracle position, whereas in (d, e), the behavior is dependent on oracle position, and we plot the special case of a centered oracle.

energy gap ΔE between $|\psi_0\rangle$ and $|\psi_1\rangle$ sets the frequency of the oscillations. For a 2D square lattice with cyclic boundary conditions (Fig. 7.8c), the ground state $|\psi_0\rangle$ undergoes a similar crossover from $|w\rangle$ to $|g\rangle$ at a critical value of $\gamma = \frac{1}{4\pi} \ln N + 0.0488$ [62], however, at this critical value the overlap between $|\psi_1\rangle$ and $|w\rangle$, as well as the value of ΔE, are substantially reduced relative to the case of a fully connected graph. This results in lower-contrast and lower-frequency oscillations during the search procedure, and thus a less favorable runtime of $\geq O(N/\operatorname{polylog} N)$ [62].

In our experiment we must additionally consider the acyclic boundary conditions in the lattice. In the case of closed boundary conditions, or in the case of an additional confining potential, translational symmetry is broken, and the eigenstates of the lattice have distinct spatial structure. This can result in eigenstates that do not overlap with specific oracles, and thus are not affected by variations in γ (Fig. 7.8a). In this case the search algorithm can still be interpreted in a similar way as in the case of cyclic boundary conditions, with the formation of a subspace composed of $|\psi_0\rangle$ and $|\psi_1\rangle$ supported by $|g\rangle$ and $|w\rangle$, but $|\psi_0\rangle$ and $|\psi_1\rangle$ need not be the ground and first excited states of the system. Furthermore, spatial structure in the ground state can also qualitatively modify the behavior of the search algorithm. For example, in the case of a finite 2D square lattice on a plane (Fig. 7.8d), $|g\rangle$ has a sine-profile along each axis with reduced overlap with vertices near the edges of the graph, since the edge vertices have reduced degree. This results in different behaviors for different oracle positions. For consistency, in this scenario we define a single critical value, γ_c, corresponding to the value of γ that numerically minimizes ΔE for a centered oracle ($w = o$, where o is the central site in the lattice as in Sec. 7.4). For an $N = 11 \times 11$ square lattice with closed boundary conditions, $\gamma_c = 0.412$, which is similar to the critical value of $\gamma = 0.430$ in the same lattice with cyclic boundary conditions. Studying the behavior of the closed system at $\gamma = \gamma_c$ as a function of the distance of the oracle w from the lattice center o (Fig. 7.9a), we find that reduced overlap between $|g\rangle$ and $|w\rangle$ for oracles near the boundary results in reduced (increased) overlap between $|\psi_0\rangle$ and $|g\rangle$ ($|w\rangle$), and reduced ΔE. Interestingly, the contributions from $|g\rangle$ and $|w\rangle$ to $|\psi_1\rangle$ remain relatively unchanged as a function of oracle distance.

Given the above analysis, the main consequence of modified boundary conditions is how those conditions change the ground state $|g\rangle$. As described in Sec. 7.4, the confinement tweezer results in an ap-

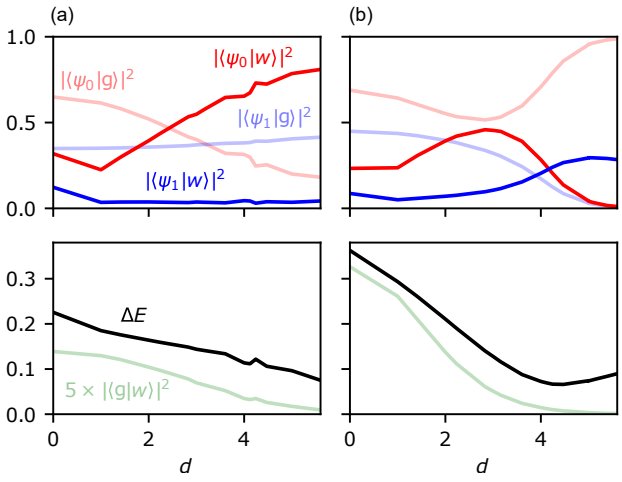

Figure 7.9: Level structure for different oracle positions. When translational symmetry is broken, the spectrum of the search Hamiltonian is dependent on the position of the oracle. Here, we plot the overlap between the states $|\psi_{0,1}\rangle$, and the states $|g\rangle$ and $|w\rangle$ (top), as well as the energy gap ΔE between the states $|\psi_0\rangle$ and $|\psi_1\rangle$ (bottom), all as a function of the magnitude of the oracle's distance from the lattice center, d (in units of the lattice spacing a). We find that the variation of the above quantities with d are primarily related to the spatial structure in $|g\rangle$, and the resulting position-dependent overlap with $|w\rangle$ (bottom, $\times 5$ scale). The above values are shown for closed boundary conditions, as in Fig. 7.8d, in (a), and for a Gaussian confining potential, as in Fig. 7.8e, in (b).

proximately Gaussian ground state, which is similar to the sine-shaped ground state resulting from closed boundary conditions. Consequently, the behavior of the system under these conditions (Figs. 7.8e, 7.9b) is similar to those in the closed-boundary case, leading to a numerically computed critical value of $\gamma_c = 0.404$, and thus an optimal oracle depth of $V_w = J_0/\gamma_c = 2.48J_0$. Note that in this case as we move away from the center of $|g\rangle$, the optimal depth of the applied oracle is reduced, which can be interpreted as the atoms effectively taking more time to reach the marked site. As a result, at intermediate oracle distances, and at the fixed value of $\gamma = \gamma_c$, the application of a deeper-than-optimal oracle leads to atoms being more tightly localized on the marked site in the state $|\psi_0\rangle$. This leads to the peak observed in Fig. 7.9b. At even further distances the oracle has very little overlap with $|g\rangle$, and effectively leaves the system unperturbed such that $|\psi_0\rangle \simeq |g\rangle$.

While the position-dependent behavior of the system with respect to different oracles with fixed parameters seems antithetical to search, in practice, these changes are small enough that the above search procedure still works for fixed system sizes. However, more theoretical analysis is required to understand how different boundary conditions can change the asymptotic scaling of a given search algorithm (see Sec. 7.7).

7.5.4 Experimental demonstration

We perform a proof of principle demonstration of spatial search by CTQW on 2D square lattice. To accomplish this, we prepare the resource state $|g\rangle$ by preparing the ground state of H_D as described in Sec. 7.4, and then quench to H (Eq. 7.26) by turning on the H_w term for some arbitrary choice of w. After performing the quench, we observe coherent oscillations in the population on the marked site $|w\rangle$ for a selection of different values of w (Fig. 7.10b). At the peak of these oscillations, the marked site w is readily identified as the highest amplitude site in the lattice. The amplitudes of these oscillations are in good agreement with theory, and limited in magnitude not by technical noise, but by the expected performance of this search procedure in a 2D square lattice (see Sec. 7.5.3).

H_w is controlled by applying a tweezer at the appropriate location w, and with tunable depth V_w. Although the analysis in Sec. 7.5.3 suggests that $V_w = 2.48J_0$ is optimal, in practice, we operate at a constant, deeper-than-optimal oracle depth of $V_w = 12.55(65)J_0$. At these depths, even when the applied

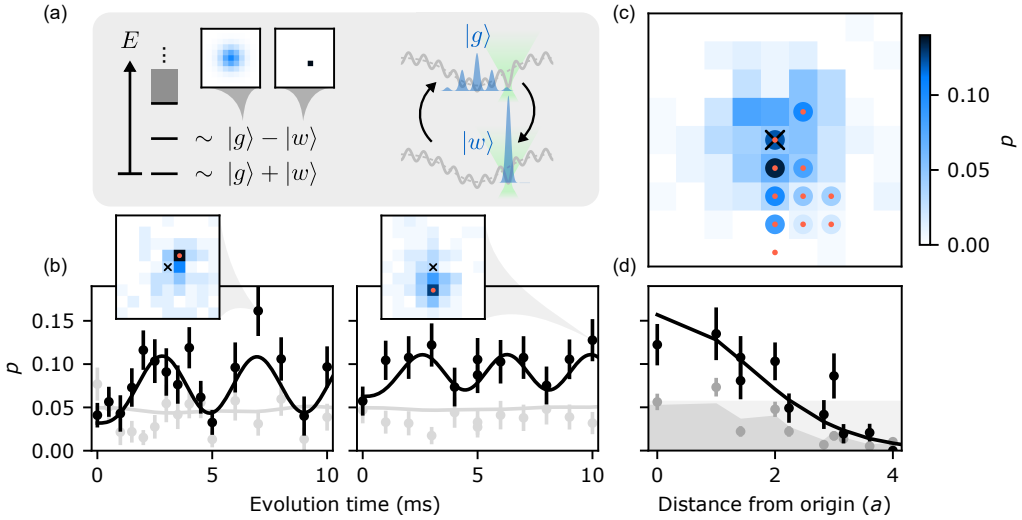

Figure 7.10: Spatial search by quantum walk. (a) We perform the described spatial search algorithm by applying the oracle Hamiltonian H_w using a tweezer (green) that selects for a specific (and arbitrary) marked site $|w\rangle$, in addition to quantum walk dynamics governed by H_D. The spectrum of the system includes ground and first excited states that are approximately the even and odd superpositions of $|w\rangle$ and the lattice ground state, or resource state, $|g\rangle$. As a result, quenching on the contribution from H_w starting from the state $|g\rangle$ results in coherent oscillations between $|g\rangle$ and $|w\rangle$. (b) We observe corresponding oscillations in the population on the marked site (black points) for a selection of oracles. These oscillations are in good agreement with theory, up to an overall offset, which uses independently characterized values of the state preparation fidelity, tunneling rate, and confinement potential (black lines). For comparison, we also plot the population in $|o\rangle$ at the center of the lattice (grey points and theory curves). Insets show the measured populations on different sites at the peak of these oscillations, where the location of the relevant oracle is marked by a red point, and the location of the origin marked by a cross. (c) A given marked site can be found by measuring the state of the system after a half-period of the above oscillations. The population on the marked site after a 2.46 ms quench is plotted for a selection of different oracle positions (each oracle is marked by a red point, with the corresponding population shown in the surrounding circle, and the origin is marked by a cross), showing a noticeable increase in amplitude relative to the amplitudes present in the prepared resource state (background image, with each pixel corresponding to a lattice site. See also Fig. 7.6c). (d) To quantify the range of the search procedure we plot the above amplitudes as a function of the distance of the oracle from the center of the lattice in units of the lattice spacing a (black points). The result is in good agreement with a theory prediction with no free parameters (black curve). At short distances, corresponding to a region spanning ~ 13 lattice sites, the marked site can be found by looking for the highest amplitude site after the quench. At longer distances the amplitude at the origin is not significantly modified by the quench, and remains near its value in the resource state $|g\rangle$ (light grey shading), which can exceed the amplitude on the marked site. Nevertheless, in a region spanning ~ 45 lattice sites there is a several-fold increase in the amplitude on the marked site relative to both the theoretical (dark grey shading) and measured (grey points) amplitude that was present before the quench. Color scale is shared across all parts of this figure.

oracle fluctuates low due to the tweezer-lattice alignment concerns described in section 3.3.5, we expect the resultant dynamics to closely follow a sinusoidal oscillation of fixed amplitude and frequency. This means that a clear signal can be observed even when averaging over different oracle depths. By contrast, if operating closer to the optimal oracle depth, fluctuations to lower depths can lead to more complicated dynamics that are not strictly sinusoidal, and so averaging over these dynamics can lead to misleading signals. We took some data at the closer-to-optimal oracle depth of $V_w = 3.1J_0$, which resulted in higher transfer into the marked site [366]. However, at the time of taking this data, it was challenging to maintain the alignment between the tweezers and the lattice to properly characterize the coherent dynamics contributing to spatial search in this regime.[15]

To calibrate the depth of the oracle, we operate at a fixed evolution time after the quench that is optimal based on calculations of the desired oracle depth, and perform the search procedure as a function of the applied oracle depth at a fixed location. At the depth that maximizes population on this marked site, we characterize the behavior of the quench as a function of both oracle position and time to confirm that the desired coherent behavior is occurring (see Fig. 7.10b). The resultant calibration of the oracle depth is in reasonable agreement with independently characterized values of the tweezer shapes and depths, and time evolution under these conditions matches up well with experimental observations (see Fig. 7.10b). Interestingly, the amplitude of both the calculated and observed oscillations are higher than what is suggested by the overlap between $|\psi_1\rangle$ and $|w\rangle$ alone, indicating that higher-energy eigenstates also play a role in the evolution.

As described in Sec. 7.5.3, in the presence of the confinement tweezer the oscillations between $|g\rangle$ and $|w\rangle$ can be dependent on the specific position of $|w\rangle$. This position-dependent behavior sets the effective size of the search space, where at greater range reduced overlap between $|w\rangle$ and $|g\rangle$ yields oscillations with reduced amplitude. In our experiment, for an oracle a distance of $\sqrt{2}a$ away from the center of the confinement tweezer (site o), the optimal evolution time after the quench is 2.46 ms. Performing this quench and evolution for a variety of different oracles (Fig. 7.10c), we find that the amplitude on the marked site after the evolution decreases as a function of distance from the origin, in agreement with theory (Fig. 7.10d).

[15] Note that for the experiments presented in this chapter, we had not yet implemented the feedforward protocol described in Sec. 3.3.5 for controlling the tweezer-lattice alignment.

Within $2a$, corresponding to a region spanning ~ 13 lattice sites, we can blind ourselves to the position of the oracle tweezer, and identify the marked site by looking for the most probable location of the walker after the quench. At longer range the amplitude that remains at the origin after the quench and subsequent evolution can exceed that of the marked site. However, within $\sqrt{13}a$, corresponding to a region spanning ~ 45 lattice sites, there is still a several-fold increase in the amplitude on the marked site relative to what was present in the resource state. In principle, the effective size of the search space could be increased with constant overhead by measuring the atomic probability density as a function of evolution time after the quench, removing effects relating to variable oscillation frequencies for different oracles, as well as the large amount of amplitude that remains near the origin for distant marked sites.

In the specific context of spatial search, the runtime of the demonstrated algorithm does not exhibit a quadratic speedup in comparison to classical search due to our use of a 2D square lattice [62]. However, as mentioned earlier, a runtime of $O(\sqrt{N}\log N)$ is achievable with a single particle in such a lattice if an additional spin-1/2 degree of freedom is used to implement a Dirac Hamiltonian [63], or a discrete-time quantum walk [7].[16]

7.5.4.1 Lattice ground state preparation by reverse search

One might be concerned that adiabatic preparation of the resource state as described in Sec. 7.4 overwhelms any potential speedups associated with spatial search by CTQW. However, since the timescale of both the search dynamics and the required adiabatic ramp are set by ΔE, the runtime of both the search algorithm and the adiabatic procedure for state preparation have similar asymptotic scaling. An alternative route to preparing $|g\rangle$ is to run the search procedure backwards. By starting with the atom at a known location w, and quenching with an oracle at that location, the atom should oscillate from $|w\rangle$ to $|g\rangle$ (Fig. 7.11). This "reverse search" procedure for state preparation has exactly the same runtime as the search algorithm.

We perform a proof of principle demonstration of reverse search at an oracle depth of $6.2J_0$, leading

[16] One could consider implementing this using a modified optical lattice containing an array of doublets [61], or the optical clock qubit in strontium, especially since it is possible to engineer state-dependent optical potentials for this qubit [138] to realize a broad class of discrete-time quantum walks [158, 175].

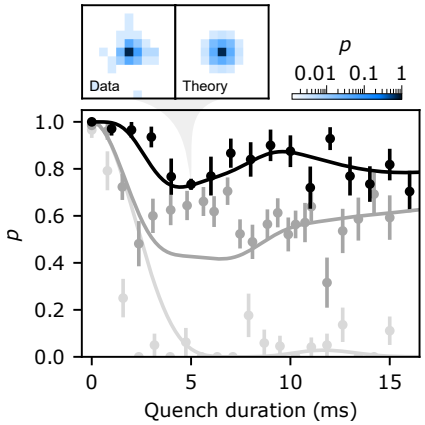

Figure 7.11: Reversed spatial search. By implanting an atom in the central lattice site $|o\rangle$, and quenching with an oracle at that site with a depth of $6.2J_0$, we observe coherent oscillations in the atomic probability density (black points) on this site due to the same dynamics as in the search algorithm. Callout: characterizing the atomic probability density as a function of position at the first minimum in this oscillation (data, left), we observe the expected behavior where much of the atom population remains on the central site, but some has oscillated into a superposition that is qualitatively similar to the desired resource state (theory with no free parameters, right). For comparison, similar quenches are performed with an oracle depth of $4.1J_0$ (dark grey points), and without the oracle (light grey points). In the absence of an oracle, the atom simply leaves the central site, and does not return on the timescales explored in this work. Theory curves with no free parameters match the deeper quench well, and are in reasonable agreement with the shallower quench, whose behavior is more sensitive to fluctuations in the depth of the applied oracle.

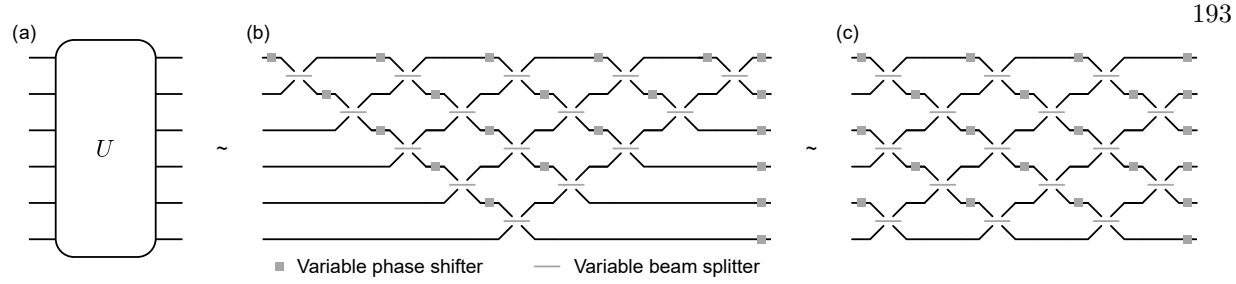

Figure 7.12: Universal control of the quantum walk unitary. An arbitrary unitary transformation between input and output modes of a system (a) can be efficiently decomposed into the appropriate application of variable beam splitters and phase shifters (b) [270]. An arrangement that is more efficient in terms of circuit depth (and still universal) is shown in (c) [71].

to slower and slightly larger amplitude oscillations than those in Sec. 7.5.4. When reducing the oracle depth even further, to $4.1J_0$, we observe less clearly sinusoidal behavior, and the resulting dynamics are more sensitive to fluctuations in the applied oracle depth. The behavior of the reverse searches are in reasonable agreement with theory with no free parameters.[17] However, due to the limitations of spatial search in a 2D square lattice, the state preparation fidelity achieved by reverse search is significantly worse than the adiabatic protocol.[18] Nevertheless, reverse search serves as an additional test of the dynamics contributing to our demonstration of spatial search, and, by extension, that the adiabatic protocol is indeed preparing the appropriate resource state.

7.6 Universal control of the quantum walk unitary

In the previous sections, we have shown that through the appropriate application of optical tweezer potentials to a tunnel-coupled 2D optical lattice, we can program quantum walk dynamics to perform useful tasks. In principle, the demonstrated programmability allows for universal control, in the sense that any unitary transformation between some set of input and output sites in the lattice can be achieved. Note that in this case we are referring to a single particle unitary, namely a unitary transformation of the single particle state space (and not, for example, on the state space of many particles occupying the lattice, which we will discuss in Ch. 8).

[17] Note that the theory predictions for reverse search only use information about the independently characterized optical potentials, since the adiabatic state preparation step is omitted.

[18] However, in a more connected lattice, this protocol could be a viable way of preparing the resource state with identical scaling to the search procedure, removing any additional overhead associated with adiabatic state preparation.

Such control can be realized in a few ways. It has been shown that CTQWs on sparse and unweighted graphs are universal for quantum computation. Specifically, one can construct a graph that prepares a plane wave input from a walker that is initially localized on a single site, and connect that to a graph that realizes an arbitrary scattering process for an incident plane wave [59]. Using only static tweezer potentials, one could implement such a graph.[19] However, this approach suffers from the fact that performing an arbitrary single particle unitary transformation in general requires a polynomial increase in the total number of vertices in the graph in comparison to the number of input and output vertices.

Alternatively, by applying time-varying tweezer potentials, one can apply an arbitrary single particle unitary transformation between all sites in the lattice [230]. Specifically, if one isolates a single site i and applies a tweezer with depth V_i for time t, an atom occupying that site experiences a phase shift $\phi_i = -V_i t/\hbar$ relative to the case where no tweezer is applied. Isolating two adjacent sites i and j and allowing an atom to tunnel between these sites for a variable evolution time implements a tunable beam splitter. For example, evolution for a quarter of a tunneling time, $t = \tau/4$, results in a balanced beamsplitter $U = (|i\rangle\langle i| - i|i\rangle\langle j| - i|j\rangle\langle i| + |j\rangle\langle j|)/\sqrt{2}$. As is well known in the context of linear optics, a network of (local) beam splitters and phase shifters is sufficient to implement any unitary transformation between the inputs and outputs of the network [270, 54, 71] (see Fig. 7.12). Therefore, the appropriate application of time varying tweezer potentials to an optical lattice can similarly implement any single particle unitary transformation between the sites in the lattice.[20]

7.7 A note on resource costs and analog classical simulation

Through the course of working on the results presented in this chapter, I became concerned about what it means to actually have an asymptotic "quantum speedup". For example, a CTQW of a single particle on any graph with N vertices is fully captured by a classical wave equation involving N coupled oscillators, including the dynamics resulting in the $O(\sqrt{N})$ runtime for searching a sufficiently connected graph [133].[21]

[19] In the construction in [59], all required graph elements can be implemented by removing vertices and edges from a 2D square lattice with diagonal tunneling by shining tweezers on and in between vertices.

[20] One does not have to restrict oneself to the described construction of sequential beam splitters and phase shifters. Often, performing these operations in parallel can result in more efficient realizations of arbitrary unitaries [230].

In fact, dynamics in any discrete Hilbert space can be captured in such a manner. That being the case, is this whole endeavor bogus?

The answer, of course, is no. However, when discussing algorithmic speedups, one must be very careful to define precisely what the relevant resource costs are. In the standard oracular model of spatial search [1], an algorithm must alternate between queries that determine whether the current location is marked, and operations that can only perform local interactions on the graph. Schematically, one can think of this as the case where a single local "robot" is allowed to traverse a memory that is spread out in physical space, and report the result of the computation back to the experimenter.[22] In this model, any classical algorithm must make $O(N)$ queries to find the marked item, even if the graph is complete. The classical wave dynamics that reproduce the behavior of a single-particle quantum walk can be classically simulated in this model only with significant overhead. As a result, in this setting there can be an asymptotic quantum speedup for search by quantum walk even when working with a single particle (or robot) traversing a graph whose vertices correspond to physical locations in space.

A distinct formulation of the search problem is to find one of $N = 2^n$ items using a single register of n bits. In the classical case, one can check only one item at a time, leading to a runtime of $O(N)$. The $O(\sqrt{N})$ runtime of Grover's search on N items using a register of n qubits leads to the oft-quoted quadratic speedup. However, unlike in the case of the quantum robot, in this setting the speedup is dependent on the exponential scaling of Hilbert space with additional qubits. In our demonstration each state in Hilbert space corresponds to a physical location in a real lattice, and so one could argue that the physical resource cost scales like $O(N)$ instead of $O(n)$. Additionally, we read out the position of the walker by performing a parallel measurement on all N sites in the lattice. In the classical case, this would correspond to having access to a register with $O(N)$ bits that can be read out in parallel. In this setting, one can simply assign a bit to each item in the search space, flip the bit corresponding to the marked item, and read out the whole register to find the marked item in a runtime of $O(1)$. This illustrates that, in some settings, the speedup

[21] If you are interested in seeing a second machine that I built for performing spatial search which uses wood and string rather than ultra-high vacuum and lasers, visit doi:10.5281/zenodo.10019752.

[22] While this may seem like an absurd restriction at first glance, there are fundamental (albeit somewhat extreme) situations where such a model could apply. For example, if information is distributed in a space that imposes a fundamental speed limit, like the speed of light, and one wants to send out a robot (or collection of robots) to search that space and return with some information [1].

associated with quantum search algorithms requires both reduced runtime, and a state space that scales rapidly with physical resources. We consider such a state space in the next chapter.

Chapter 8

Multi-particle quantum walks

"That's the best thing about being me... there are so many 'me's."
— Agent Smith

Portions of this chapter have appeared in:

A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman. An atomic boson sampler. **In review**, 2023

In this chapter, we will extend the single atom dynamics described in Ch. 7 to the multiparticle case. In Sec. 8.1, we will review the mechanisms leading to interactions between ground state atoms in a lattice, and show how for certain experiments these interactions are negligible. In Sec. 8.2 we will focus on non-interacting multi-particle CTQWs. One might naively expect that in the absence of interactions, the many-particle dynamics arise trivially from that of a single particle. However, we will find that particle statistics play a crucial role. In specific cases, including for what is known as boson sampling [2], non-interacting dynamics can still result in behavior that is not believed to be possible to simulate classically in an efficient manner.

Interestingly, non-interacting bosonic dynamics lie in a convenient middle ground where, on the one hand, their behavior is in general hard to simulate. On the other hand, the lack of interactions introduce additional structure to the dynamics that enable powerful techniques for benchmarking. In Sec. 8.3, we will use the properties of non-interacting dynamics to develop and apply various techniques to benchmark the quality of our state preparation. In Sec. 8.4, we will briefly discuss related measurements that also allow us to characterize the applied unitary evolution directly. In Sec. 8.5, we will extend the above measurements and benchmarking techniques to large systems involving many atoms, including the largest-scale demonstrations

of boson sampling to date, which are beyond our ability to simulate using current classical algorithms. We will comment on the classical simulations we perform in Sec. 8.6. In Sec. 8.7, we will conclude our discussion of non-interacting quantum walks by describing how noise affects the sample complexity of boson sampling, and comparing our results to previous results using photons.

The capabilities demonstrated in this chapter, namely the ability to rapidly prepare nearly arbitrary patterns of identical atoms in a tunnel coupled optical lattice using optical tweezers, provide a new approach to studying Hubbard models using ultracold atoms. We will conclude the chapter in Sec. 8.8 by presenting preliminary results where interactions between atoms are reintroduced in our experiment to study strongly interacting quantum walks, as well as the preparation of a superfluid ground state.

8.1 The Bose Hubbard model

In order to extend the discussion in Sec. 7.1 to multiple non-interacting bosonic atoms, we define bosonic creation operators $a_i^\dagger$ which create an atom with the localized wavefunction $w_{0,i}$ (as defined in Eq. 7.6). Similarly, we can define bosonic annihilation operators a_i and number operators $n_i = a_i^\dagger a_i$. a_i and $a_i^\dagger$ encode the symmetrization required by bosonic particle statistics (namely that the wavefunction remains the same upon exchanging particles), and have the standard bosonic operator algebra:

$$\begin{gathered}
[a_i, a_j] = [a_i^\dagger, a_j^\dagger] = 0, \quad [a_i, a_j^\dagger] = \delta_{ij} \\
a_i^\dagger |n_i\rangle = \sqrt{n_i + 1}|n_i + 1\rangle, \quad a_i|n_i\rangle = \sqrt{n_i}|n_i - 1\rangle, \quad a_i|0_i\rangle = 0 \\
n_i|n_i\rangle = a_i^\dagger a_i |n_i\rangle = n_i|n_i\rangle,
\end{gathered} \tag{8.1}$$

where $|n_i\rangle$ is a Fock state containing n_i atoms on site i (namely atoms with the wavefunction $w_{0,i}$). Given the above definitions, we can write the many-boson equivalent of Eq. 7.8 as:

$$H = -\sum_{\{i,j\}} J_{ij} \left(a_i^\dagger a_j + a_j^\dagger a_i \right) - \sum_i V_i n_i, \tag{8.2}$$

where, as before, the first sum runs over all pairs of sites $\{i, j\}$ with appreciable tunnel coupling. Assuming the V_i's are uniform, the evolution of a Fock state of atoms occupying different sites under this Hamiltonian corresponds to a multi-particle CTQW.

8.1.1 Interactions

It is not always the case that lattice-trapped atoms are non-interacting. When two atoms come close together with relative displacement $\boldsymbol{r}$, they experience an interaction potential $U(\boldsymbol{r})$ (due to van der Waals forces, Coulomb forces, etc.). The details of this potential can, in general, be quite complicated. However, for most atomic gasses $U(\boldsymbol{r})$ can be assumed to be spherically symmetric, and $U(\boldsymbol{r}) \to 0$ for $|\boldsymbol{r}| \gg r_0$, where r_0 is some characteristic length scale of the potential (typically on the scale of angstroms) [65]. Given the above assumptions, we can consider the scattering problem where an atom initially in a plane wave state with momentum k oriented in the z direction scatters off of another atom located at the origin. The steady state atomic wavefunction in the far field ($|\boldsymbol{r}| \gg r_0$) is [281]:

$$\psi(\boldsymbol{r}) = e^{ikz} + f_k(\theta)\frac{e^{ikr}}{r}, \tag{8.3}$$

where θ is the angle $\boldsymbol{r}$ makes with respect to the z-axis, and z is the projection of $\boldsymbol{r}$ onto that axis. f_k is a scattering amplitude that, due to the assumed spherical symmetry of $U(\boldsymbol{r})$, depends only on the angle θ.

One can expand f_k in the basis of spherical harmonics to obtain [281]:

$$f_k(\theta) = \frac{1}{2ik}\sum_{l=0}^{\infty}(2l+1)P_l(\cos\theta)(e^{2i\delta_l} - 1), \tag{8.4}$$

where P_l are the Legendre polynomials, and δ_l is the phase shift acquired by the l^{th} partial wave due to scattering. For $k \ll 1/r_0$, and when there isn't a nearby molecular resonance, $\delta_l \sim k^{2l+1}$ [281]. As a result, for bosonic atoms with very small k (as applies to our experiments), only the $l = 0$ (also known as s-wave) component contributes to scattering. In this low energy case, one can show that [281]:

$$\lim_{k\to 0} k\cot\delta_0 = -\frac{1}{a_s}, \tag{8.5}$$

where a_s is a single constant quantity known as the scattering length which determines the entirety of the low energy scattering behavior. It is challenging to perform an ab initio calculation of a_s for atoms that are significantly more complicated than hydrogen. However, a_s can be measured directly from experiments, the results of which are summarized in Tab. 8.1 for the naturally abundant isotopes of Sr.

Because the low energy scattering is found to be independent of the exact functional form of $U(\boldsymbol{r})$, we can choose a pseudopotential that is easier to work with, but nevertheless results in the same behavior.

For example, we can choose the delta, or contact, potential:

$$U(\boldsymbol{r}) = a_s \frac{4\pi\hbar^2}{m}\delta(\boldsymbol{r}), \tag{8.6}$$

where m is the atomic mass. Using this pseudopotential, we can calculate the energy shift due to interactions between two atoms occupying the same site in the lattice [293]:

$$U_c = \frac{4\pi\hbar^2 a_s}{m}\int d^3\boldsymbol{r}|w_{0,i}(\boldsymbol{r})|^4, \tag{8.7}$$

where $w_{0,i}$ is as defined in Eq. 7.6.

As long as U_c is small compared to the band gap, and the atom density is relatively low (meaning there aren't ever significantly more than two atoms on a given lattice site), we can assume that s-wave interactions between the atoms do not dramatically modify the functional form of $w_{0,i}$. In this case we can write the full, interacting Bose-Hubbard Hamiltonian as:

$$H = -\sum_{i,j} J_{ij}\hat{a}_i^\dagger\hat{a}_j - \sum_i V_i\hat{n}_i + \frac{U_c}{2}\sum_i \hat{n}_i(\hat{n}_i - 1), \tag{8.8}$$

where the last term accounts for all pairwise interactions that can occur on a given site.

Often, one is interested in studying different regimes of Eq. 8.8 by controlling the value of U_c/J. This can be achieved by varying the lattice depth, which alters J much more rapidly than U_c. However, lower J leads to increased sensitivity to noise (e.g. atom loss, or disorder in the lattice potential), resulting in a practical limit to how much one can reduce J. In experiments with alkali atoms, U_c/J can additionally be tuned by adjusting the value of a_s using a magnetic Feshbach resonance [65]. Unfortunately, the $J = 0$ ground state of Sr (where J here refers to the total angular momentum), makes it insensitive to magnetic field tuning.[1] Luckily, the values of a_s for the naturally abundant bosonic isotopes of Sr span a convenient range of values (see Tab. 8.1). Experiments with ^{88}Sr, where $U_c \ll J$, are well captured by the non-interacting Hamiltonian in Eq. 8.2.[2] Experiments with ^{86}Sr, where $U_c \gg J$, are well captured by Eq. 8.8[3] and approximate hard core (infinite strength) interactions. Experiments with ^{84}Sr allow us to access the regime where $U_c \sim J$.

[1] It is possible to achieve similar tuning in alkaline earth atoms using optical Feshbach resonances, where one uses laser light to couple the scattering atoms to a specific molecular state [100, 358]. However, to date, such techniques have also resulted in significant amounts of accompanying inelastic scattering [235].

[2] We will comment on the accuracy of this approximation in Sec. 8.7

[3] Although these large values of U_c begin to approach the limit where $w_{0,i}$ is modified due to the strength of U_c [51].

	^{88}Sr	^{87}Sr	^{86}Sr	^{84}Sr
^{88}Sr	-1.4(6)	55.0(2)	97.4(1)	1790(130)
^{87}Sr	-	92.6(1)	162.5(5)	-56(1)
^{86}Sr	-	-	823(24)	31.9(3)
^{84}Sr	-	-	-	122.7(3)
U_c/h (Hz)	-1.7(7)	109.4(1)	970(30)	144.9(4)

Table 8.1: Scattering properties of the naturally abundant isotopes of strontium. Upper rows list the s-wave scattering length between isotopes, provided in units of the Bohr radius a_0. These values are obtained from [216]. Lower row lists the values of U_c/h for interactions between atoms of the same isotope in units of hertz and for the typical lattice parameters used in this chapter, where $J/h = 119$ Hz.

8.1.1.1 Inelastic interactions

Due to conservation of energy and momentum, the dominant inelastic interactions in our experiment involve three body processes, where two atoms form a molecular bound state and a third atom carries away the binding energy. Such a process typically results in the loss of all three atoms, since the binding energy is much larger than the depth of the lattice potential. A loose upper bound for the three body loss coefficient in ^{88}Sr is $\beta_{ggg} < 10^{-27}$ cm^{6}s^{-1} [314], with an expected value closer to (and likely lower than)[4] the measured value of $\beta_{ggg} = 2.0(2) \times 10^{-30}$ cm^{6}s^{-1} in ^{87}Sr [122]. The resulting rate equation for the total atom number N is:

$$\frac{dN}{dt} = -\beta_{ggg} \int d^3\boldsymbol{r} \rho(\boldsymbol{r})^3, \tag{8.9}$$

where $\rho(\boldsymbol{r}) = \sum_i \langle n_i \rangle |w_{0,i}(\boldsymbol{r})|^2$ is the atom density. At early times, we can define an effective single particle loss rate associated with three body processes:

$$\tau_{ggg}^{-1} = \frac{1}{N} \frac{dN}{dt}\bigg|_{t=0}. \tag{8.10}$$

The highest densities we operate at using ^{88}Sr are $\langle n_i \rangle = 1/4$, which for typical lattice parameters and the above loss coefficients yield $\tau_{ggg} \gtrsim 6 \times 10^2$ s and $\tau_{ggg} \gtrsim 3 \times 10^5$ s respectively. Note that the effect of bosonic interference leads to a worst case correction of 1/3! to the above lifetimes [157, 50]. As a result, we conclude that the effect of inelastic collisions are negligible for all experiments using ^{88}Sr in this work.

[4] Typically, three body collision rates follow the scaling $\beta_{ggg} \propto n^2 a_s^4$ [106], which would indicate a vanishingly small three body loss rate for ^{88}Sr due to the small value of a_s in comparison to the other isotopes of Sr. However, this approximation is only accurate for $|a_s|$ larger than the van der Waals radius of the atom.

We have performed some preliminary experiments involving quantum walks of atoms in both the ground and clock electronic states. In this case, new channels involving electronic-state changing collisions open up, leading to strong two-body loss.[5] Specifically, the two body loss coefficient for a collision between one ground (g) and one clock (e) state ^{88}Sr atom is $\beta_{ge} = 5.3(1.9) \times 10^{-13}$ cm^{3}s^{-1}, whereas the coefficient for a collision between two clock state ^{88}Sr atoms is $\beta_{ee} = 4.0(2.5) \times 10^{-12}$ cm^{3}s^{-1} [196]. The resulting rate equation for the total atom number N is:

$$\frac{dN}{dt} = -\beta_{ge} \int d^3\boldsymbol{r} \rho_g(\boldsymbol{r}) \rho_e(\boldsymbol{r}) - \beta_{ee} \int d^3\boldsymbol{r} \rho_e(\boldsymbol{r})^2, \tag{8.11}$$

where ρ_x denotes the density of atoms in the electronic state x. For typical lattice parameters, and for two clock state atoms occupying the same site, this results in an effective single atom lifetime of $\tau_{ee} \sim 10$ ms. Similarly, for one clock and one ground state atom occupying the same site, this results in an effective single atom lifetime of $\tau_{ge} \sim 200$ ms.

8.2 Non-interacting quantum walks

As we saw in the previous section, ground state ^{88}Sr atoms behave like non-interacting bosons in our experiments, and thus the evolution of these atoms are well-captured by Eq. 8.2. Consider the simplified scenario of a CTQW on a cycle graph, as in Sec. 7.2.[6] As before, the single atom eigenstates are plane waves,[7] and so we can define creation (and annihilation) operators in the eigenbasis of plane waves [230]:

$$b_q^\dagger = \frac{1}{\sqrt{m}} \sum_j e^{ij2\pi q/m} a_j^\dagger, \tag{8.12}$$

where q is the wave number, and m the number of lattice sites (vertices in the graph). The many body time evolution operator is then:

$$U_b(t) = \sum_q e^{it2J\cos(2\pi q/m)} n_q, \tag{8.13}$$

where J is the tunneling energy, and $n_q \coloneqq b_q^\dagger b_q$ is the number operator in the plane wave basis. Although time evolution in this situation is trivial, predicting the outcome of measurements in the position basis is not.

[5] In these cases, the emission of a photon can preserve conservation of energy and momentum.
[6] This corresponds to Eq. 8.2 for a 1D lattice with cyclic boundary conditions, $V_i = 0$, and $J_{ij} = J$.
[7] Eq. 7.13, with energies given by Eq. 7.14.

This is because when making predictions regarding such measurements, one must keep track of the interplay between the phases accumulated during time evolution, and the symmetrization constraints encoded in the bosonic operators. In fact, even in this simple scenario of non-interacting bosonic CTQWs on a cycle graph, it is conjectured that it is not possible to efficiently sample from the output distribution in the position basis classically [230].

More generally, we are interested in the probability of observing a specific pattern of atoms $|\boldsymbol{k}'\rangle$ after arbitrary (non-interacting) time evolution starting from a given initial pattern $|\boldsymbol{j}'\rangle$. The state $|\boldsymbol{j}'\rangle$ (and similarly $|\boldsymbol{k}'\rangle$) is conveniently described in the site occupation basis. Given a site occupation list $\boldsymbol{j}' \coloneqq (j'_1, \ldots, j'_m)$, where j'_i counts the number of atoms occupying a given site $i \in \{1, \ldots, m\}$ (and m is the total number of sites in the lattice), we define that:

$$|\boldsymbol{j}'\rangle \coloneqq \prod_{i=1}^{m} \frac{\left(a_i^\dagger\right)^{j'_i}}{\sqrt{j'_i!}} |0\rangle. \tag{8.14}$$

The effect of the non-interacting Hamiltonian dynamics is to apply a transformation to the bosonic operators [325]:

$$a_j^\dagger \to a_j^\dagger(t) \coloneqq \sum_i U(t)_{ji} a_i^\dagger, \tag{8.15}$$

where $U(t) = e^{-iHt/\hbar}$ is the single atom time evolution operator (t is the evolution time and, for our experiments, H is as defined in Eq. 7.8). Therefore, under time evolution, the many atom state undergoes the transformation:

$$|\boldsymbol{j}'\rangle \to |\boldsymbol{j}', t\rangle \coloneqq \prod_{i=1}^{m} \frac{\left(a_i^\dagger(t)\right)^{j'_i}}{\sqrt{j'_i!}} |0\rangle. \tag{8.16}$$

After time evolution, the probability of observing a specific output state $|\boldsymbol{k}'\rangle$ is:

$$\begin{aligned} P^{\mathrm{B}}(\boldsymbol{k}'|\boldsymbol{j}', U) &\coloneqq |\langle \boldsymbol{k}'|\boldsymbol{j}', t\rangle|^2 \\ &= \left| \langle 0 | \prod_{p=1}^{m} \frac{(a_p)^{k'_p}}{\sqrt{k'_p!}} \prod_{q=1}^{m} \frac{(a_q^\dagger(t))^{j'_q}}{\sqrt{j'_q!}} | 0 \rangle \right|^2 . \end{aligned} \tag{8.17}$$

To simplify Eq. 8.17, we introduce the site-assignment basis $|\boldsymbol{j}\rangle$, which is equivalent to the site occupation basis ($|\boldsymbol{j}\rangle \equiv |\boldsymbol{j}'\rangle$). Here, n atoms occupy the sites $\boldsymbol{j} = (j_1, \ldots, j_n)$, with $j_1 \leq \cdots \leq j_n$, where

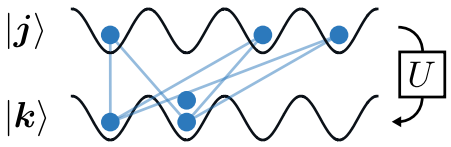

Figure 8.1: Symmetrization constraints on bosonic interference. In this cartoon, we assume that the bosonic particles (blue disks) occupy distinct positions in an optical lattice (black contours), and that the evolution is non-interacting, and thus described by a single particle unitary transformation U. We are interested in computing the probability that one observes an arrangement of particles $|\boldsymbol{k}\rangle$ after preparing the arrangement $|\boldsymbol{j}\rangle$ and applying U. Due to the indistinguishability of the particles, in order to perform this calculation one must coherently sum over all trajectories (blue lines) that can take one from $|\boldsymbol{j}\rangle$ to $|\boldsymbol{k}\rangle$, leading to complicated interference effects.

$j_l \in \{1, \ldots, m\}$. Plugging Eq. 8.15 into Eq. 8.17, and performing some combinatorial manipulations, one can arrive at the expression [289, 2, 325]:

$$P^{\mathrm{B}}(\boldsymbol{k}|\boldsymbol{j}, U) \equiv P^{\mathrm{B}}(\boldsymbol{k}'|\boldsymbol{j}', U) = \frac{1}{\boldsymbol{k}!\boldsymbol{j}!}|\mathrm{Perm}(U_{\boldsymbol{kj}})|^2, \tag{8.18}$$

where $\mathrm{Perm}(\cdot)$ is the matrix permanent, and $U_{\boldsymbol{kj}}$ is the $n \times n$ submatrix of U that contains only the rows corresponding to sites $\boldsymbol{k}$, and columns corresponding to $\boldsymbol{j}$. The normalization constant is $\boldsymbol{k}! \coloneqq \prod_{i=1}^{m} k_i'!$ (and similarly for $\boldsymbol{j}!$), where we have used the representation of the state in the site occupation basis.[8]

A first quantized interpretation of Eq. 8.18 is often adopted in the literature on boson sampling [2], where the interference between multiple particles can be understood as a summation over all permutations of atom labels on the input going to atom labels on the output (see Fig. 8.1). The equivalent calculation for distinguishable atoms yields:

$$P^{\mathrm{D}}(\boldsymbol{k}|\boldsymbol{j}, U) = \frac{1}{\boldsymbol{k}!\boldsymbol{j}!}\mathrm{Perm}(|U_{\boldsymbol{kj}}|^2), \tag{8.19}$$

which can be interpreted as simply combining the probability distributions resulting from the single particle CTQWs of each atom.

8.2.1 Boson sampling

Computing the permanent of an arbitrary matrix is in the complexity class #P-hard [2, 339], and so even sampling from the probability distribution resulting from the evolution described in Sec. 8.2 is believed

[8] Note that these normalization terms are equal to 1 unless a collision occurs (in the sense of multiple atoms occupying the same site). In our experiments, where the lattice is initialized with at most one atom per site, $\boldsymbol{j}! = 1$.

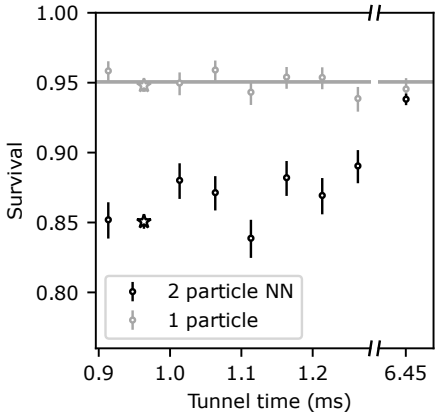

Figure 8.2: Atom loss during quantum walk dynamics. We measure the atom loss in experiments with one atom, and with two atoms initially positioned in nearest-neighbor (NN) sites in the lattice, as a function of evolution time. The pictured evolution times are near the Hong-Ou-Mandel dip (see Sec. 8.3.1), and at the latest evolution time explored in this work of $t = 6.45$ ms (horizontal line corresponds to the mean single particle loss). Data points marked by a star are taken at $t = t_{\mathrm{HOM}} \coloneqq 0.96$ ms. Notice that single particle loss is not observed to vary over the evolution times explored in this work. The excess loss in the two particle data is associated with collisions between atoms during readout. As the atoms spread out and the atom density goes down, the gap between the measured one and two particle loss closes.

to be intractable for more than ~ 50 atoms using classical techniques [72]. By contrast, approximating the permanent of a non-negative real matrix can be performed in polynomial time, and so the corresponding sampling task for distinguishable atoms can efficiently be accomplished classically [326]. Different degrees of atom distinguishability result in behaviors that lie in between these two scenarios [326, 94].

Because most of the discussion surrounding boson sampling is framed in terms of implementations using photonics experiments, it is useful to point out a few differences in terminology: for us, the modes that the bosonic particles (atoms) can occupy are sites in the lattice, and the single particle unitary transformation U is the result of Hamiltonian evolution on the lattice and the associated CTQW dynamics. In photonics, the modes that the bosonic particles (photons) can occupy are optical modes, and U is the transformation applied between the input and output modes of a linear optical network. We will use the above terminology interchangeably as appropriate.

8.2.2 A note on measurements, loss, and postselection

Implementing boson sampling in our apparatus amounts to preparing specific patterns of identical atoms using the prescription in Ch. 4, performing a CTQW by evolving under Eq. 8.2, and measuring the atoms in the position basis. However, due to the parity projection that occurs during imaging (see Sec. 4.2), our measurements are not number resolving, and instead detect atom number parity on each site. For most experiments in this chapter, we will operate in a regime where the probability that more than one atom occupies a given site is low, minimizing the effect of parity projection.[9] Additionally, except for experiments that involve large atom numbers (namely those in Sec. 8.5), we postselect our measurements on perfect rearrangement, and on observing no lost or extra atoms after the quantum walk dynamics.

Postselection is further useful for suppressing the effect of imperfect state preparation. Because atoms occupying higher bands in the lattice experience significantly higher tunneling energies (Sec. 7.1.1), for our experimental parameters these atoms are either lost, or exit the analysis region and are interpreted as lost during the quantum walk dynamics. Therefore, postselection on atom number ensures that all remaining atoms were in the ground band in the plane of the 2D lattice. We observe a loss of 5.0(2) % (Fig. 8.2), which is mostly explained by the expected combined ground state occupation (see Sec. 4.3.1) of 97.37 % in the radial directions (leading to 2.63% loss). The remaining loss is likely due to imperfect adiabaticity of ramps in the lattice depth, as well as imperfect overlap between the Wannier functions of the atoms in the lattice before and after the lattice is quenched to the conditions used for tunneling. Note that we do not see increased single particle loss as a function of evolution time, even at the longest evolution times explored in this chapter, suggesting that the loss is indeed due to filtering of imperfectly cooled atoms. It is important to note that avoiding the effects of loss that scales with evolution time is critical for maintaining the computational complexity of the boson sampling problem when scaling to larger systems, since one must apply a single particle unitary that adequately mixes the different modes in the system [114, 243]. Ultimately, effects like parametric heating and vacuum lifetime will result in single particle loss that is dependent on the evolution time. However, we expect these effects to become relevant only on the timescale of seconds

[9] In fact, although our measurements are not in a complete basis, similar restrictions are also made in the original proposal for boson sampling out of consideration for the use of bucket or click detectors in photonics [2], which distinguish only between no photons, and the presence of one or more photons.

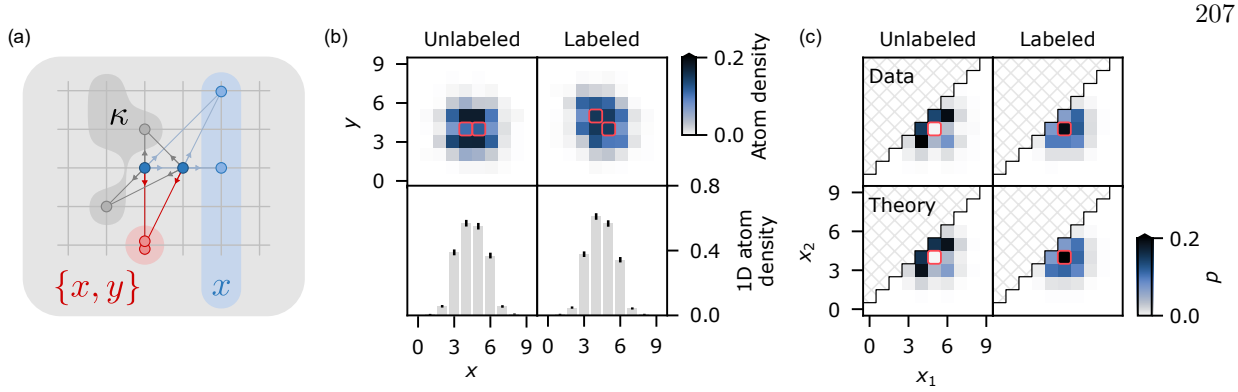

Figure 8.3: Two particle quantum walks on a 2D lattice. (a) The interference between atoms can be studied by measuring correlations in the occupation of specific sites in the lattice (red). Partitioning the lattice into arbitrary subsets of sites κ (grey) can allow one to more efficiently (in terms of the number of experimental trials) study the effects of interference. Due to the separability of quantum walk dynamics on a square lattice, it is particularly useful to partition the lattice into columns indexed by their x coordinates (blue). (b) Specifically, by binning 2D quantum walk data over the y axis, one can study 1D quantum walks along x, where the initial coordinate in the y axis acts as an additional "hidden" degree of freedom that can modify the distinguishability of the atoms in the remaining "visible" x axis. For example, one can prepare two atoms with the same (unlabeled) or different (labeled) y coordinates (sites initialized with an atom are marked in red). The atom number density on each site after the evolution of these atoms is shown here at an evolution time of $t = t_{\mathrm{HOM}} \coloneqq 0.96$ ms. (c) After binning to 1D, each two particle output is uniquely labeled by the x coordinates of the two atoms (x_1, x_2), with $x_2 \leq x_1$. The measured probabilities p of observing a given two particle output are in agreement with theory, where the unlabeled atoms (top left) agree with the expectation for perfectly indistinguishable bosons (bottom left), and labeled atoms (top right) with distinguishable particles (bottom right). The states highlighted in red once again denote the prepared initial states before the atom evolution.

(Sec. 5.3.1) in comparison to the millisecond-scale experiments explored in this chapter.

Finally, the above postselection also results in a quadratic suppression of imaging errors. Given the already low imaging error rates (Sec. 4.4.2), measurements involving such postselection are therefore negligibly impacted by imaging errors. For experiments in which we do not postselect on the number of atoms remaining after the quantum walk dynamics (namely the experiments appearing in Sec. 8.5) imaging errors can play a role. Specifically, for the atom and mode numbers relevant to Sec. 8.5, the combined effects of false positives and negatives lead to overall error rates of ~ 1 % in our ability to correctly identify the presence and location of a given atom. In experiments that are sensitive to imaging errors, we calibrate the effect of the errors and include them in our simulations (see Sec. 8.6).

8.2.3 Effective particle statistics

Although our atoms are fundamentally bosonic composite particles, they may not behave bosonically on the lattice. This is because the single-atom state space includes degrees of freedom (DOFs) other than just the location of the atom in the lattice. For example, the atom can be in different electronic states, or in different motional states in the direction that is normal to the lattice.[10] The single atom state space is thus a tensor product of the Hilbert spaces $\mathcal{H}_V$ ("visible" DOFs) spanned by $|i\rangle$ (sites), and $\mathcal{H}_H$ ("hidden" DOFs) spanned by $|h\rangle$ (labels). In our system, the overall unitary dynamics is non-interacting for the atoms, and the single atom Hamiltonian acts independently on the hidden and the visible DOFs, which means that it is of the form $H_V \otimes \mathbb{1}_H + \mathbb{1}_V \otimes H_H$ (where the subscripts indicate the subspace that each operator acts on, and H is a Hamiltonian acting on the indicated subspace). Due to this independence, the hidden DOFs can affect the visible behavior of the atoms only by changing their effective particle statistics. For example, if we prepare some number of atoms in specific visible sites, but each atom has a different hidden label, then the visible behavior of the atoms is that of perfectly distinguishable particles. Other types of visible particle statistics are also possible, including fermionic statistics if the multi-particle wavefunction is antisymmetric in the hidden DOFs [284, 219]. The tests in this chapter primarily focus on determining limits on the deviation from bosonic behavior due to the hidden DOFs without characterizing the specific particle statistics exhibited in the visible behavior. Note that errors in the visible DOFs are unlikely due to our procedure for state preparation[11] and postselection (see Sec. 8.2.2).[12]

There is flexibility in how to partition between visible and hidden DOFs (Fig. 8.3a), which we take advantage of in tests of indistinguishability involving up to 8 atoms. In a square lattice, H takes the form $H_x \otimes \mathbb{1}_y + \mathbb{1}_x \otimes H_y$ where x and y denote the spatial coordinates of atoms (see Sec. 7.2.1), and so we can consider one of the spatial coordinates as hidden. For example, if we ignore the y coordinate of the atoms the resulting visible behavior is that of a 1D multiparticle quantum walk along x (Fig. 8.3bc). We refer to this procedure as "binning," since it is accomplished by summing the observed numbers of atoms along the y

[10] Recall that in-plane motional excitations lead to loss (Sec. 8.2.2).
[11] Namely, optical cooling and atom rearrangement as described in Ch. 4.
[12] Because both cooling and imaging involve the spontaneous emission of photons from atoms occupying different lattice sites, our state preparation is also expected to be perfectly dephasing. This rules out error models for boson sampling like the mean field sampler that rely on coherences between different input modes [327].

axis in the final image. Binning is convenient for two reasons. First, it allows us to operate in a regime where it is rare that there is more than one atom on a given lattice site at the end of the dynamics. As a result, parity projection with binning is effectively equivalent to number-resolved measurements of the visible sites. Second, because we can control the initial y coordinates of the atoms via rearrangement, we can control one of the hidden DOFs and thus the distinguishability of the atoms on the visible DOF (Fig. 8.3bc). We will refer to this approach to preparing distinguishable atoms as position labeling.[13]

An alternative way of simulating distinguishable particle statistics is to perform a pair of experiments where only one atom is prepared at a time, and then combine the data in subsequent analysis (we will refer to this as time labeling). One can view time labeling as being akin to introducing a time delay between otherwise identical atoms, where the time when a particle is measured serves as a hidden DOF.[14] For data sets in which we use time labeling to simulate n partially or fully distinguishable particles by combining multiple runs with fewer than n particles, we compensate for parity projection during imaging by summing the resulting processed images and taking the result mod 2.

8.3 Benchmarking atom indistinguishability

As stated in Sec. 8.2.1, the degree to which it is computationally hard to sample from non-interacting quantum walk dynamics depends on particle statistics, with distinguishable (and fermionic) particles being easy to simulate classically. We are therefore interested in characterizing the indistinguishability of the atoms in our experiment, which tells us the degree to which they behave like identical bosons.

8.3.1 Multimode Hong-Ou-Mandel interference

The indistinguishability $\mathcal{I}$ of two atoms is defined to be the Hilbert-Schmidt inner product of the density matrices of the two one-atom wavefunctions on the hidden DOFs. It can be inferred from a Hong-Ou-Mandel (HOM) experiment [143] by comparing the coincidence probability of the atoms to the corresponding coincidence probability for perfectly distinguishable atoms. We measure HOM interference of atoms [165,

[13] Referring to the fact that the initial positions of the atoms in y act as labels that cause the atoms to be distinguishable.
[14] This is similar to standard techniques in photonics, where a time delay is used to tune the distinguishability of two or more photons.

201, 276] by studying quantum walks in 1D via the binning procedure described in Sec. 8.2.3 (Figs. 8.3 and 8.4). Schematically, one can associate the evolution time with a variation in the reflectivity of the beam splitter in a standard HOM measurement (Fig. 8.4ab).

However, whereas in a typical HOM experiment there are only two modes, in our experiment there are many modes, and so the probability of coincidence of two atoms in any specific two sites is small. Therefore, it is useful to bundle many sites together. Given partially distinguishable atoms that start on sites k, l with $k \neq l$, the probability that they arrive at sites i, j is:[15]

$$p^{\text{partial}}(ij|kl) = \begin{cases} |U_{ik}|^2|U_{jl}|^2 + |U_{il}|^2|U_{jk}|^2 + 2\mathcal{J}\,\text{Re}\left(U_{ik}U_{jl}U_{jk}^*U_{il}^*\right), & \text{if } i \neq j \\ (1+\mathcal{J})|U_{ik}|^2|U_{il}|^2, & \text{otherwise.} \end{cases} \tag{8.20}$$

One can see the role that interference plays by noting the cross terms that vanish with $\mathcal{J} \to 0$.[16] The probability we observe a coincident detection of partially distinguishable atoms in disjoint sets of sites S_1 and S_2 is:

$$p^{\text{partial}}_{S_1,S_2} = \sum_{i\in S_1, j\in S_2} p^{\text{partial}}(ij|kl). \tag{8.21}$$

The corresponding probability $p^{\text{dist}}_{S_1,S_2}$ for distinguishable particles is obtained when setting $\mathcal{J} = 0$. Analogous to a typical HOM experiment, we can consider the ratio of such coincident detections for partially distinguishable and distinguishable particles:

$$Q^{\text{HOM}}_{S_1,S_2} \coloneqq \frac{p^{\text{partial}}_{S_1,S_2}}{p^{\text{dist}}_{S_1,S_2}} = 1 + \mathcal{J}\underbrace{\frac{2\sum_{\substack{i\in S_1\\ j\in S_2}} \text{Re}(U_{i,k}U_{j,l}U_{j,k}^*U_{i,l}^*)}{\sum_{\substack{i\in S_1\\ j\in S_2}} \left(|U_{i,k}|^2|U_{j,l}|^2 + |U_{i,l}|^2|U_{j,k}|^2\right)}}_{\coloneqq -\tau(S_1,S_2)}. \tag{8.22}$$

Note that the ratio $-\tau(S_1, S_2) \geq -1$, and so we have the inequality:

$$Q^{\text{HOM}}_{S_1,S_2} \geq 1 - \mathcal{J}. \tag{8.23}$$

We therefore have that the distinguishability $(1 - \mathcal{J})$ of two atoms is upper-bounded by the ratio $Q^{\text{HOM}}_{S_1,S_2}$ for any choice of S_1 and S_2.

When equality holds in Eq. 8.23,[17] we say that S_1 and S_2 satisfy the balanced condition, in analogy to the case of a balanced beamsplitter. Such a balanced condition can easily be found in our experiment

[15] Assuming linear optical evolution, and a label degree of freedom that evolves independently of the visible degrees of freedom.

[16] Note that the enhancement by a factor of two for two indistinguishable atoms (with $\mathcal{J} = 1$) to bunch in the same mode i is precisely the two atom case of the more general enhancement of $n!$ for full bunching of n atoms (see Sec. 8.3.2).

[17] So that $\tau(S_1, S_2) = 1$ and $Q^{\text{HOM}}_{S_1,S_2} = 1 - \mathcal{J}$.

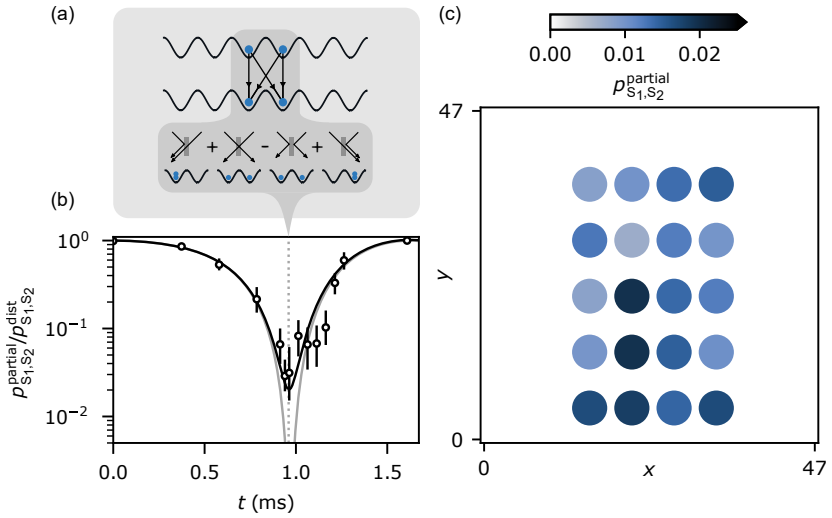

Figure 8.4: Benchmarking two particle quantum walks. (a) Two atoms (blue disks) are prepared on adjacent sites in a 1D chain and undergo CTQW dynamics. At the appropriate evolution time t, the operation applied between the two initially populated sites is precisely a beam splitter, and we can measure the resulting Hong-Ou-Mandel (HOM) interference. (b) p_{S_1,S_2} refers to the probability of coincident detection of one atom on each of the two input columns (sites in 1D after binning). We plot the ratio $p^{\text{partial}}_{S_1,S_2}/p^{\text{dist}}_{S_1,S_2}$, where the superscripts "dist" and "partial" respectively refer to experiments with and without introducing an additional label for distinguishability. At an evolution time of $t = t_{\text{HOM}}$ (grey dotted line) $p^{\text{partial}}_{S_1,S_2}/p^{\text{dist}}_{S_1,S_2}$ should vanish for perfectly indistinguishable bosons, and for an ideal lattice with only nearest-neighbor tunneling (grey theory curve). Higher order tunneling terms in our lattice (diagonal and next-nearest neighbor) result in imperfect visibility even for identical bosons (black theory curve). (c) HOM measurements at different locations in the lattice do not show a strong dependence on location over a region that spans all input sites used in this work (colored circles show the quality of the HOM dip at the corresponding position). Note that the HOM data appearing in (b) is averaged over three regions centered at $x = 27$ and $y = (16.5, 24.5, 32.5)$.

by considering the binning procedure described in Sec. 8.2.3, which converts a 2D quantum walk into a 1D quantum walk (Fig. 8.3). In the 1D quantum walk, an evolution time of $t_{\text{HOM}} \coloneqq 0.96$ ms in our experiment approximates a balanced beam splitter between two adjacent sites. Therefore, at t_{HOM}, $Q^{\text{HOM}}_{S_1,S_2}$ is a tight upper bound for the distinguishability of two atoms when S_1 and S_2 each correspond to all the sites in one of the two columns that initially contained an atom before the quantum walk dynamics.[18]

Performing measurements of $Q^{\text{HOM}}_{S_1,S_2}$, including measurements of distinguishable atoms using both position and time labeling,[19] we infer a lower bound on the atom indistinguishability of $97.1^{+1.0}_{-1.5}$ %. Due to the presence of small diagonal and next-nearest-neighbor tunneling terms (see Sec. 7.1.1), we are not able to perfectly implement the balanced condition for reasonable choices of t, S_1, and S_2. However, via the

[18] One must be careful to include the effect of parity projection in the measurement of $Q^{\text{HOM}}_{S_1,S_2}$, which we discuss in Appendix B, but the effect of this correction is multiplicative, and fairly small.

[19] We find that both approaches agree to within statistical errors.

calibrations discussed in Secs. 7.3 and 8.4, we estimate the magnitude of the resulting error and apply it as a correction. This correction yields an estimated value of $\mathcal{J} = 99.5^{+0.5}_{-1.6}$ %,[20] which is consistent with what we expect based only on errors due to imperfect optical cooling.[21] We do not expect other sources of distinguishability to play a significant role in our experiments.

Note that the above measurements are averaged over three different regions in the lattice. To confirm that the quality of our state preparation does not vary significantly across the lattice, we perform additional HOM measurements as a function of position in the lattice, and do not observe significant variation in performance across a region containing all input sites used in this chapter (Fig. 8.4c).

8.3.2 Bunching and clouding

Similar to HOM interference and the resulting bunching for two atoms, the visible particle statistics of many atoms can be tested by measuring full bunching [315] (Fig. 8.5acd), namely the probability that all n atoms occupy the same output site.[22] In fact, the full-bunching probability is $n!$ larger when the particle statistics are perfectly bosonic as opposed to when they are perfectly distinguishable, independent of the unitary applied [315]. This can be seen by comparing Eq. 8.18 and Eq. 8.19: because full bunching corresponds to a situation where all elements of $\boldsymbol{k}$ are equal, all the rows of $U_{\boldsymbol{kj}}$ are equal. The permanent is then simply a sum of $n!$ identical terms, and so squaring inside or outside the permanent results in a difference of $n!$.[23]

The full-bunching probability is uniquely maximized by bosonic particle statistics in comparison to other particle statistics, and so maximum full bunching for the visible DOFs implies that the visible particle statistics are bosonic (for non-zero full bunching). We note that we normalize the bunching probability of distinguishable particles by the probability of full survival for bosonic particles to account for the effect of parity projection (see Appendix B). We see the expected enhancement in full bunching borne out in

[20] We also performed a separate measurement of the indistinguishability that does not rely on either postselection or the separability of the two dimensions of the quantum walk dynamics, and found that the two measurements were consistent [367].

[21] Based on the master equation simulation in Sec. 4.3.1, we predict an axial motional ground state occupation of 99.58 %, and a corresponding indistinguishability of 99.27 %.

[22] Keeping in mind that in our binned measurements a site refers to a column of the lattice.

[23] This is an example of what is sometimes known as "kinematic" interference, namely interference that only depends on particle statistics, and the structure those statistics impose on the relevant Hilbert space. This is to be contrasted with "dynamical" interference, namely the (likely more familiar) interference that arises due to specific details regarding the evolution of a system. For an excellent discussion of different kinds of interference, see [325].

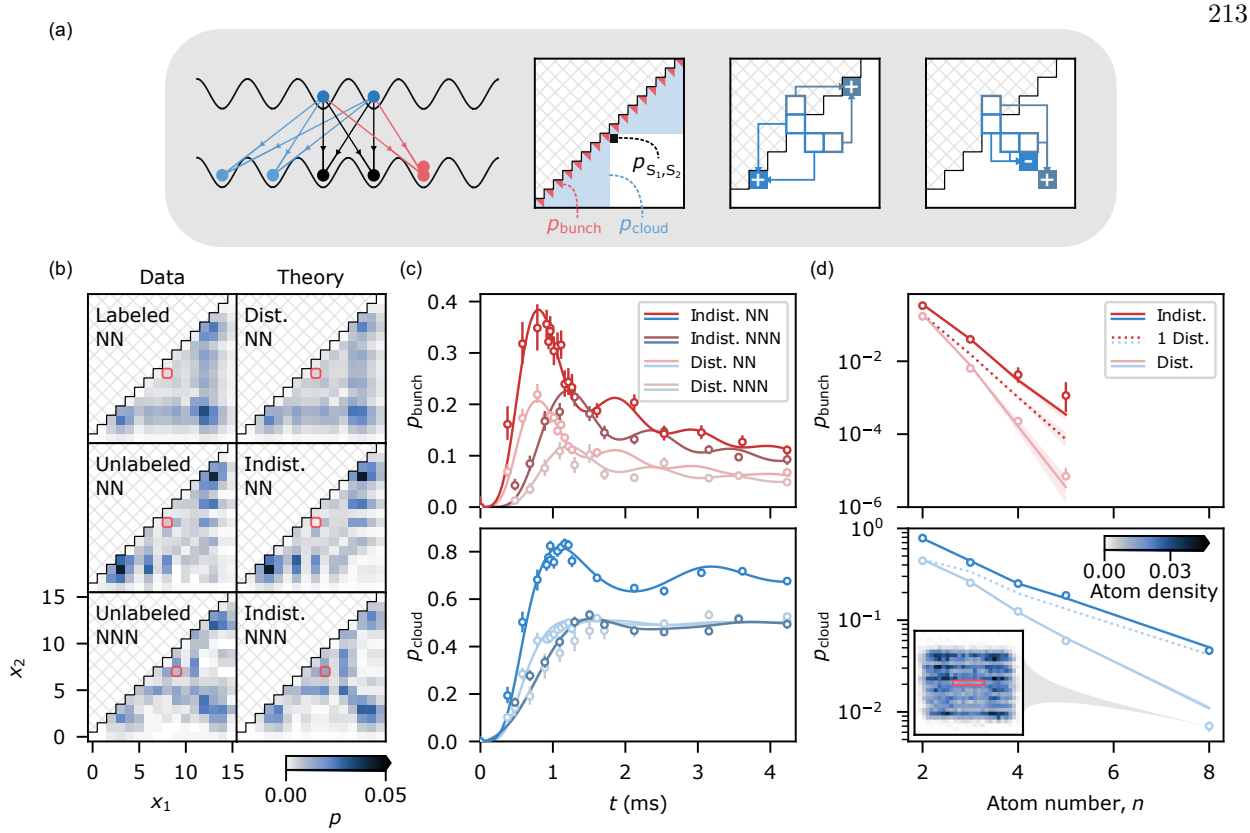

Figure 8.5: Multiparticle quantum walks in 1D. (a) Quantum walks in 1D exhibit dramatically different behaviors for different degrees of distinguishability, and for different input states. We consider three ways of coarse-graining multi-atom distributions (after binning) to characterize the degree to which the atoms are behaving like indistinguishable bosons. p_{S_1,S_2} (black output) refers to the probability of coincident detection of one atom on each of two input sites, p_{bunch} to a coincident detection of all atoms on the same site (summing over red outputs), and p_{cloud} to all atoms appearing on the same half of the array (summing over blue outputs). Example trajectories that fulfill these conditions are shown on the left. The rightmost two panels provide an intuitive, first quantized picture for understanding how constructive (+) and destructive (−) interference of different input states (open squares) lead to the observed probability of occupying a given output state (shaded squares). (b) Two particle quantum walks in 1D at an evolution time of $t = 4.23$ ms. The relevant input state for each subpanel is highlighted in red, with the atoms prepared in either neighboring sites in the lattice (nearest-neighbor, NN), or separated by one site (next-nearest-neighbor, NNN). The theoretical predictions are for perfectly distinguishable particles (Dist.) and indistinguishable bosons (Indist.). The associated measurements either introduce (Labeled) or do not introduce (Unlabeled) a time label to distinguish between the different atoms. For the remainder of this figure we use "Dist." or "Indist." to refer to both data and theory for consistency, with the understanding that the experimental measurements may not correspond to perfectly indistinguishable atoms. (c) Measurements of p_{bunch} and p_{cloud} for two atoms as a function of evolution time, and for both NN and NNN inputs, are in good agreement with theory. (d) Measurements of p_{bunch} and p_{cloud} with up to 5 and 8 atoms respectively are also in good agreement with theory, and indicate that there is no observable degradation of particle indistinguishability when scaling to more atoms. For comparison, we also display a prediction for when each trial of the experiment contains one randomly selected atom that is distinguishable from the rest (1 Dist.). The displayed experiments are performed at an evolution time of $t = (n-1)t_{\text{HOM}}$, where all n input sites are approximately uniformly coupled to each other after the quantum walk dynamics (see inset for an example of the measured atom density in 2D with 8 prepared atoms, red region highlights input sites). The atoms are prepared with NN spacing in 1D, which, unlike NNN spacing, is an arrangement that can generate a significant increase in p_{cloud} for identical bosons in comparison to distinguishable particles. The theory predictions appearing throughout this figure are for error-free preparations of atoms with the indicated particle statistics. The theory curves in (d) are performed using Monte-Carlo methods, and thus the shaded regions denote $\pm 1\sigma$ confidence intervals that include both sampling errors, and systematic errors relating to fluctuations in the applied unitary evolution.

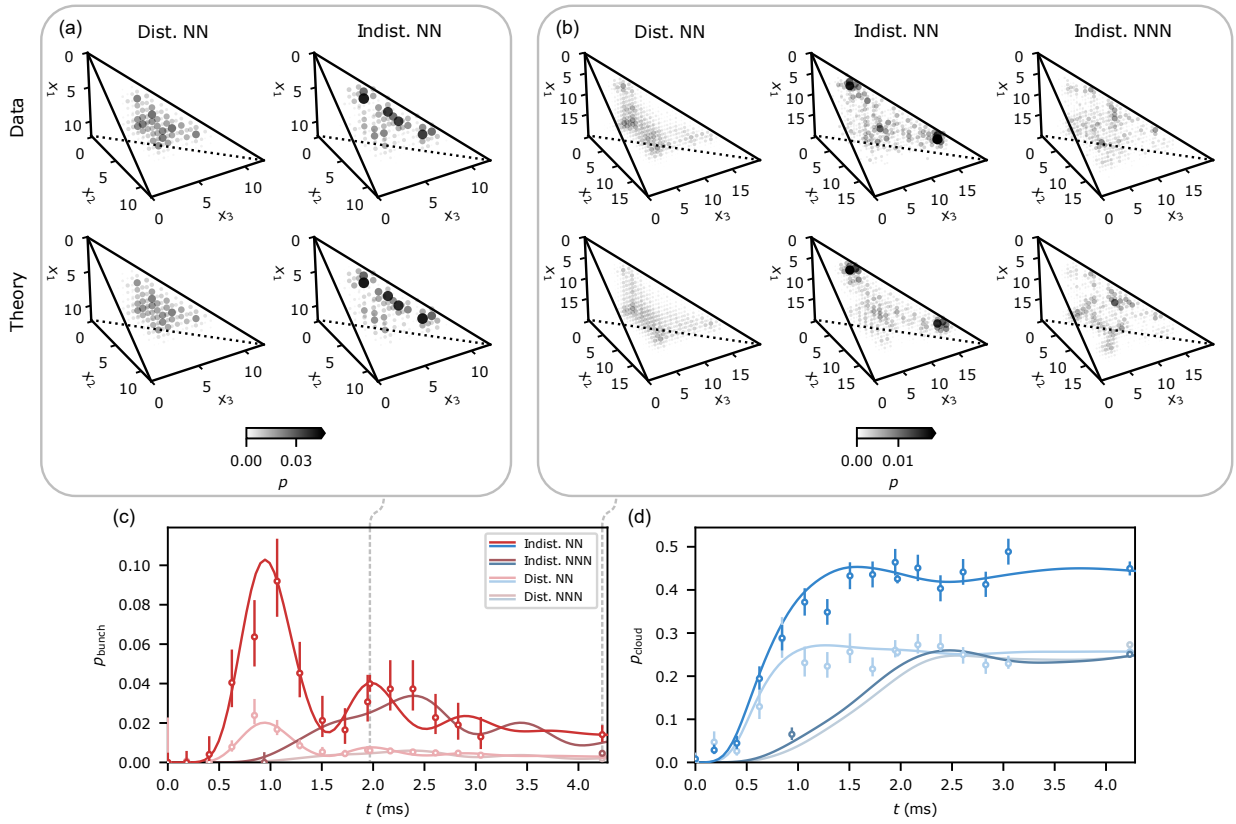

Figure 8.6: Three particle quantum walks in 1D. The output distributions resulting from three particle quantum walks at evolution times of (a), 1.97 ms and (b), 4.23 ms are in good agreement with theory. Similar to the two particle case, each three particle output can be uniquely labeled by the coordinates of the three particles (x_1, x_2, x_3), with $x_3 \leq x_2 \leq x_1$. The probability p of measuring an output state (x_1, x_2, x_3) is denoted by both the size and color of the circle at the corresponding coordinates. Note that, when referring to data, "indistinguishable" denotes the lack of an explicit time label but does not necessarily mean that the atoms behave as perfectly indistinguishable bosons. For input patterns with nearest-neighbor (NN) spacing, indistinguishable bosons (Indist.) exhibit enhanced probability to lie near the leading edge of the distribution along the main diagonal ($x_1 = x_2 = x_3$) in comparison to distinguishable particles (Dist.). This tendency disappears for input patterns with next-nearest-neighbor (NNN) spacing. In both the NN and NNN cases the probability to lie directly on the main diagonal is enhanced in the indistinguishable case, since this corresponds to bunching. (c), (d), Like in the two particle case, we can coarse-grain the three particle distributions by measuring bunching and clouding, and find good agreement with theory as a function of evolution time. All theory predictions in this figure correspond to error-free preparations of atoms with the appropriate particle statistics.

experiments with up to 5 atoms (Fig. 8.5cd). However, the full-bunching probability diminishes exponentially in atom number, becoming infeasible to measure for more than 5 atoms given the data rate currently attained in our setup.

In order to test the indistinguishability of larger numbers of atoms it is beneficial to look for signals that are sensitive to interference, but bundle many correlation functions together into a single quantity that can converge more quickly than any specific correlation function. One can build an intuition for the kinds of interference that occur in multiparticle CTQWs in 1D by noting that a CTQW of two distinguishable particles in 1D is isomorphic to a single particle CTQW in 2D — for example, one can associate the position of particle 1 with the x coordinate and particle 2 with the y coordinate. Similar arguments apply for more particles and higher dimensional square lattices. Therefore, one can apply intuitions about single particle CTQWs in square lattices to the visualizations of the full output distributions for multiparticle CTQWs in 1D appearing in Figs. 8.3 and 8.5a.

In the following discussion, we will refer to the space that these multiparticle output distributions occupy (namely states labeled by the coordinates of multiple particles) as a virtual lattice (by analogy to a real lattice, where the states are labeled by real coordinates). Consider the two particle input state $\boldsymbol{k} = (k_0, k_1)$ in a first quantized picture:

$$|\boldsymbol{k}\rangle = \frac{1}{\sqrt{2}}(|k_0\rangle|k_1\rangle + |k_1\rangle|k_0\rangle), \tag{8.24}$$

where $|i\rangle$ refers to the single atom state where an atom occupies site i in a real 1D lattice (Fig. 8.5a). One can consider $|\boldsymbol{k}\rangle$ to be the state where one is in a uniform superposition of two states across the diagonal of a 2D virtual lattice. Under CTQW dynamics, the two components of $|\boldsymbol{k}\rangle$ will interfere.[24] By virtue of matrix exponentiation, each "hop" from one site to another in the virtual (and real) lattice is associated with a phase shift of $\pi/2$. Note that the minimum path length connecting the two components of $|\boldsymbol{k}\rangle$ to a site lying on the diagonal of the virtual lattice is equal. Therefore, the leading edge of the probability distribution along the diagonal of the virtual lattice always experiences constructive interference, as predicted by the full bunching probability.[25] The interference on the leading edge along the anti-diagonal of the virtual lattice can either

[24] Keep in mind that these two "components" are not physical, since the labels we have assigned the two indistinguishable bosons here are not physical.

[25] Keep in mind that the leading diagonal of the virtual lattice precisely corresponds to states where all the particles have

be constructive or destructive. If, for example, k_0 and k_1 are adjacent sites in the real lattice, then the two components of the wavefunction in Eq. 8.24 are separated by two hops on the virtual lattice,[26] resulting in a phase shift of π and thus destructive interference (Fig. 8.5a, b middle). If k_0 and k_1 are separated by one site in the real lattice, then the two components of the wavefunction are separated by four hops on the virtual lattice, resulting in a phase shift of 2π and thus constructive interference (Fig. 8.5a, b bottom). As one rotates from an axis oriented along the anti-diagonal to the diagonal of the virtual lattice, the relative phase associated with the minimum paths that connect the two components of the state to the leading edge of the distribution along the axis will interpolate between 0 and 2π.[27] Note that for distinguishable particles, one instead starts in a statistical mixture of $|k_0\rangle|k_1\rangle$ and $|k_1\rangle|k_0\rangle$,[28] and therefore all the effects of interference described above disappear (Fig. 8.5b top).

When extending the above picture to more particles it becomes clear that, for particles initially arranged in adjacent sites in 1D, there will be a general tendency for interference to yield an enhanced probability for being close to the main diagonal, and a reduced probability to be close to the antidiagonal(s) of the virtual lattice. This tendency can be quantified by clouding [55], which is defined as the probability that all atoms end up on the same half of the 1D array (Fig. 8.5).[29] We study clouding for atoms prepared in neighboring sites in the lattice as a function of the atom number n, and with an evolution time of $t = (n-1)t_{\mathrm{HOM}}$ such that all input sites are approximately uniformly coupled to each other after the quantum walk dynamics. In clouding measurements with up to 8 atoms (Fig. 8.5cd), the behavior of the atoms is in line with the prediction for ideal bosons and clearly separated from both measurements with time-labeled (and thus distinguishable) atoms, and theoretical predictions for partially distinguishable atoms. Similar to the bunching scenario, we normalize the clouding probability of distinguishable particles by the probability of full survival for bosonic particles to account for the effect of parity projection (see Appendix B). Note that, unlike full bunching, the enhancement of clouding for identical bosons in comparison to distinguishable particles is strongly dependent on both the system evolution, and the specific input state. For example, per

the same coordinate, i.e. bunch on the same site.

[26] Equivalently, it takes two hops to exchange the positions of two distinguishable particles occupying adjacent sites.

[27] As k_0 and k_1 are separated by more and more sites, this phase will interpolate from 0 to larger multiples of π, leading to more lobes of constructive interference, separated from each other by regions of destructive interference.

[28] This state is mixed because we do not keep track of the labels. If keeping track of the labels, one would instead be in a pure state corresponding to a specific site in the virtual lattice.

[29] Equivalently, in a quadrant of the virtual lattice that touches the main diagonal.

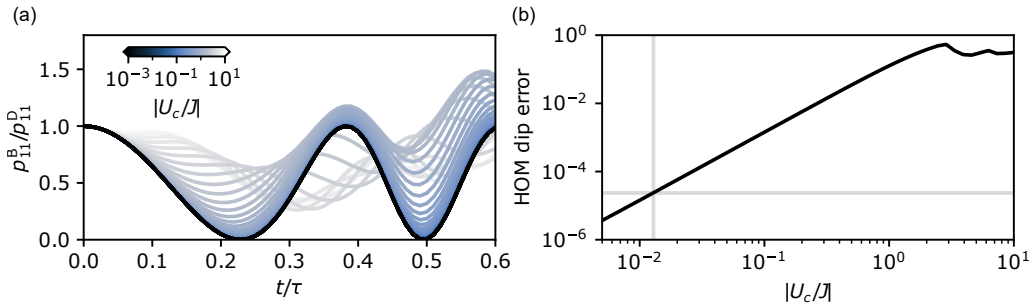

Figure 8.7: The role of interactions in HOM interference. (a) We simulate the evolution of two nearest-neighbor atoms in a uniform 2D lattice under the influence of a variable contact interaction U_c in comparison to the nearest-neighbor tunneling energy J. We define $\tau = 1/J$. (b) The deviation in the visibility of the HOM dip from unity is characterized as a function of interaction strength, with the experimentally relevant values marked by the grey lines.

the arguments above, preparing atoms at next-nearest-neighbor spacing in the lattice can make the difference in clouding for bosonic and distinguishable atoms vanishingly small, as confirmed by measurements with two and three atoms (Fig. 8.5bc, Fig. 8.6). As a result, although our measurements of clouding provide strong evidence of indistinguishable, bosonic particle statistics for up to 8 particles, interpreting these measurements requires precise knowledge of the atom evolution, namely the single-particle unitary U.

8.3.3 The effect of interactions

To confirm that the interactions between ^{88}Sr atoms are indeed negligible for our experiments, we simulate the interacting dynamics exactly for two and three particles (Fig. 8.7) via matrix exponentiation in QuSpin [348], and compare these results to the non-interacting case (see also Sec. 8.6). For all signals involving two and three particles, we find that the resulting errors are at or below the 10^{-4} level. Although we do not explicitly check if these interactions affect any of the signals at higher particle number, we do confirm that the measured signals are not significantly affected by triple occupation or higher of a given site. Since n_i must reach a value of ~ 10 for the interaction term in Eq. 8.8 to be similar to the hopping energy for a single atom, we do not expect these interactions to significantly affect any of the signals at higher particle number [94].

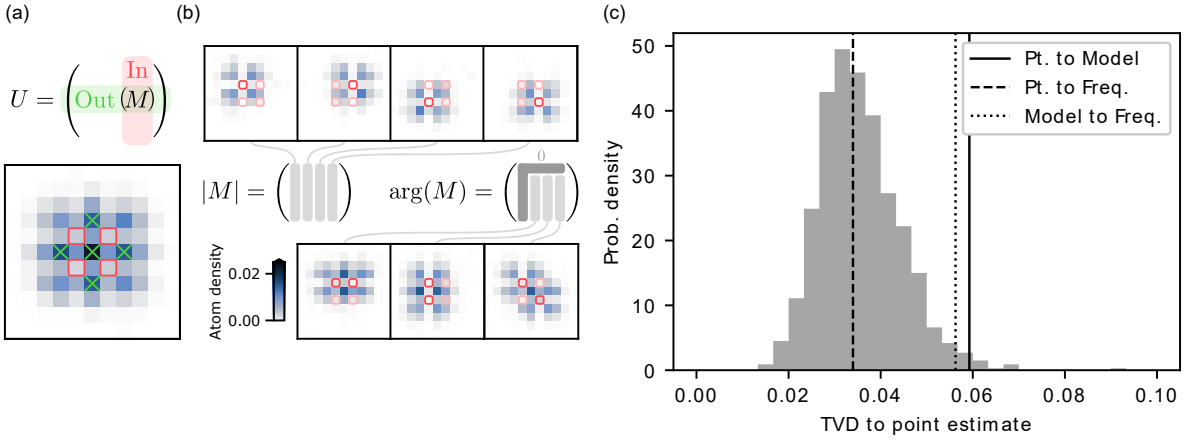

Figure 8.8: Calibrating the single particle unitary using quantum walks. (a) We directly characterize the evolution of pairs of atoms from the sites circled in red to those marked by green crosses, which is determined by a 5×4 submatrix M of the single particle unitary U. (b) Depicted are density plots of the one- and two-particle quantum walks that are used to infer M, with sites that are populated in a given preparation circled in dark red (unpopulated sites in light red). The magnitudes of the entries of M can be inferred using the one-particle quantum walk data, and the phases of the entries of M relative to the first column and row can be inferred using two-particle quantum walks [180]. (c) We perform maximum likelihood (ML) inference using the above data, and compare the point estimate to ML estimates of bootstrap resamples of the data, and to the spectroscopic calibration (see Sec. 7.3). To quantify this comparison, we compute the one- and two-particle distributions generated by the inferred parameters, and compute the total variation distances (TVDs) of these distributions, then take the maximum of the TVDs over the prepared input patterns. We call this quantity the max TVD between two sets of distributions. The depicted histogram is the max TVD between the point estimate and the ML estimates of the bootstrap resampled data. Shown also are the max TVD between the frequencies of the data (Freq.) and the point estimate (Pt.), and that between the point estimate and the spectroscopic model (Model). The bootstrap histogram gives a sense of the size of the statistical fluctuation of the max TVD between the point estimate and the truth. The max TVD between the spectroscopic model and the point estimate is large compared to the bulk of the histogram, which is the expected behavior because statistical fluctuations in the model add to the statistical fluctuations in the point estimate.

8.4 Calibrating the quantum walk unitary using Fock states

The calibrations of U discussed so far involve performing spectroscopic measurements that use the atoms as local probes of the lattice depth, and to combine these measurements with band structure calculations to generate a model of the lattice Hamiltonian (see Sec. 7.3). Although this approach is expected to be accurate for quantum walks in the lattice, direct measurements of U that do not rely on band theory are important for future studies that attempt to program U using local, possibly time-varying potentials imposed by the optical tweezers [366] (see Sec. 7.6).

For comparison, in photonics experiments, direct characterizations of U can be performed efficiently and with very low noise by taking advantage of bright coherent states containing a macroscopic number of photons [376, 211]. A challenge in our experiment is the inability to prepare equivalent states of atoms. Instead, as a proof of principle, we show that U can be inferred directly by measuring the interference resulting from different preparations of Fock states [180] (Fig. 8.8). Specifically, we perform a maximum likelihood (ML) fit of the desired parameters in the unitary using data involving preparations of one and two atoms in the lattice, which results in better precision than previous approaches [180].[30] Using this fitting procedure, we perform characterizations of the terms in U corresponding to four input sites and five output sites at an evolution time of $t = 1.46$ ms using 4440 separate experimental trials (Fig. 8.8). We find that the ML fit is within statistical variation of the spectroscopic characterization of U (see Fig. 8.8c). In principle, these fits allow one to characterize all parameters in U at constant precision with a number of experimental trials that scales polynomially in m. However, the number of parameters in U that we can directly infer using quantum walk data is currently limited by the combination of high constant overhead in the required number of experimental trials, and drifts in the experiment that modify U. As a point of reference, it is reasonable for us to collect $\sim 10^4$ useful experimental trials each day, which is comparable to the number of trials collected in the above calibration. Although U is stable on the scale of a single day (based on the spectroscopic calibrations), it is prone to drifting over the course of a few days. Therefore, it is not currently feasible for us to use the above procedure to perform calibrations of U at significantly later evolution times, and for significantly more input and output sites.

[30] We leave a detailed discussion of this fitting procedure to the future thesis of Shawn Geller.

8.5 Measuring interference between many particles

With an understanding of the quality of our state preparation, and the applied CTQW dynamics in the lattice, we perform experiments involving large numbers of particles that are beyond the reach of current classical simulations.

8.5.1 Generalized bunching

A natural way of measuring interference between many identical particles is to measure connected correlation functions of increasingly high order. However, the difference in such correlations for atoms with different particle statistics typically decreases exponentially in atom number. Instead, we are interested in measuring observables that are sensitive to interference, but also converge quickly with respect to the number of experimental trials. One option is to appeal to generalized notions of bunching, where one measures the probability p_κ that all n atoms appear in an arbitrary subset κ of sites. Although it was originally suspected that p_κ was maximized by bosonic particle statistics in comparison to other particle statistics for any κ where the number of sites $|\kappa|$ is $\geq n$ or $|\kappa| = 1$ [304], recent counter-examples have shown that this is false [298, 262].[31] Nevertheless, except in certain fine-tuned cases, p_κ serves as a sensitive signal of interference (meaning that the number of required measurements scales only polynomially with atom number). Specifically, partial distinguishability errors rapidly reduce the value of p_κ in comparison to the value for indistinguishable bosons.

The above generalization of bunching is also useful in that it unifies our earlier measurements, where binning columns of sites and measuring coincidences, full bunching, or clouding simply correspond to specific choices of κ (Fig. 8.3a). For appropriate selections of the size of κ, the difference in generalized bunching probabilities of bosonic behavior in comparison to other behaviors is expected to converge in a number of measurements that is polynomial in n [304, 302]. Specifically, we choose $|\kappa| = \lfloor m - m/n \rceil$ (where $\lfloor \cdot \rceil$ denotes rounding to the nearest integer), and study generalized bunching as a function of atom number for square $\sqrt{n} \times \sqrt{n}$ input patterns with next-nearest neighbor spacing. We perform these measurements at a fixed

[31] These counterexamples are interesting in their own right, in that they break some long-standing mathematical conjectures, including various stronger versions of Lieb's permanental dominance conjecture [303, 346, 262].

evolution time of $t = 6.45$ ms, and with $m = 500$, which is approximately the number of sites that are likely to contain an atom after the quantum walk dynamics. We cannot directly measure p_κ, and instead measure the probability p'_κ that all observed atoms on a given run of the experiment appear in the set of sites κ. p'_κ differs from p_κ because even-numbered occupancy of a site not contained in κ contributes to a "successful" event where all remaining atoms after parity projection appear within κ. We find in numerics that p'_κ, like p_κ, serves as a useful observable: it converges in a reasonable number of measurements, and distinguishes between a family of experimentally relevant models for the hidden DOFs (defined below).

8.5.1.1 Averaging generalized bunching

To avoid concerns of biasing in our choice of κ, we average p'_κ over all choices of a given size $|\kappa|$ to compute the quantity $\overline{p'_\kappa}$. Despite the large number of such selections, this quantity can be computed efficiently via a combinatorial argument. Let $S_m = \{1, \ldots, m\}$ be shorthand for the set of output sites, and let $S \subseteq_k S_m$ denote a subset of sites of size k. Let G be the random variable denoting the site occupation of the output, with g being a single sample of that random variable. Then, the average probability that all n particles arrive in a set of size $k = |\kappa| \geq n$ is

$$\overline{p'_\kappa} = \binom{m}{k}^{-1} \sum_{S \subseteq_k S_m} \Pr(G \subseteq S) \tag{8.25}$$

$$= \binom{m}{k}^{-1} \sum_{g} \Pr(g) \sum_{S \subseteq_k S_m} \mathbb{I}(g \subseteq S) \tag{8.26}$$

$$= \binom{m}{k}^{-1} \sum_{g} \Pr(g) \binom{m - C(g)}{k - C(g)} \tag{8.27}$$

where $\Pr(f)$ refers to the probability of event f occurring, and $\mathbb{I}$ is the indicator function that is 1 when its argument is true, and 0 otherwise. The sum over g ranges over all possible mode occupations of the output, and $C(g)$ is the number of nonzero entries of g. Note that the main contribution to $\overline{p'_\kappa}$ is bunching and the resulting loss as a result of parity projection. As a result, we do not postselect on the number of surviving atoms after the CTQW dynamics in our measurements of $\overline{p'_\kappa}$.

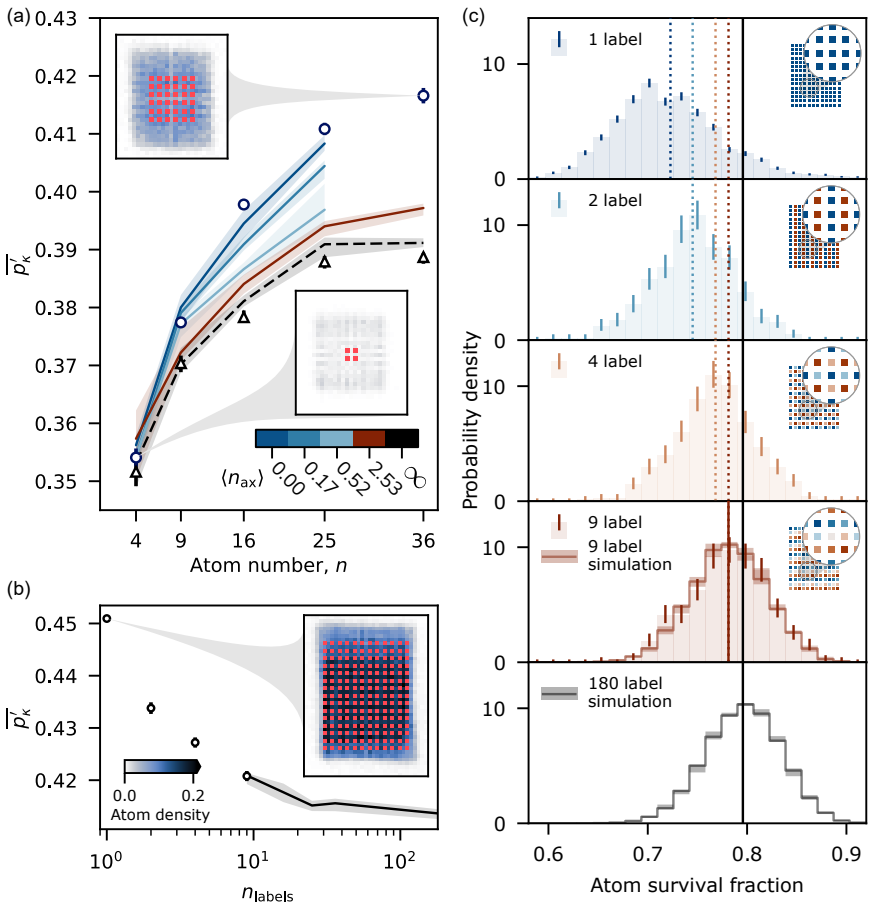

Figure 8.9: Interference of large bosonic Fock states. (a) Measurements of the generalized bunching probability $\overline{p'_\kappa}$ as a function of atom number for square $\sqrt{n} \times \sqrt{n}$ input patterns at next-nearest-neighbor spacing and at a fixed evolution time of $t = 6.45$ ms with (triangles) and without (circles) introducing time labels that make the atoms fully distinguishable. These measurements can be compared to a model for partial distinguishability due to thermal motion normal to the lattice (colored lines), which is described by a harmonic oscillator with quantum number n_{ax} and thermal expectation value $\langle n_{\mathrm{ax}} \rangle$. The fully distinguishable case (dashed line) corresponds to infinite temperature. Notice that simulations for low temperatures and large atom numbers are absent due to computational overhead. Insets in (a) and (b) show the atom density after their evolution, and share a color bar. The input patterns relevant to each inset are highlighted in red. (b) Measurements of $\overline{p'_\kappa}$ with input patterns containing 180 atoms as a function of distinguishability (circles). The atom distinguishability is controlled by partitioning the input state into n_{labels} sub-ensembles, such that only atoms within a sub-ensemble share a time label and thus can interfere. For particles that are sufficiently distinguishable to simulate, we can compare these measurements to theory (solid line). (c) The distribution in the observed fraction of surviving atoms on each shot of the experiment is also sensitive to the effects of interference due to parity projection. For the case of 9 labels (second from bottom panel), simulations capture the measured distribution of atom survival probabilities. Measurements with different numbers of labels (and thus atom distinguishability) are clearly resolved both from each other, and from a simulation of the fully distinguishable case with 180 labels (bottom panel). For ease of comparison, the vertical lines denote the mean of each distribution (dotted lines are measurements, solid lines are theory). Insets in (c) denote the relevant assignment of labels for each data set in both (b) and (c), with each color corresponding to a unique time label for a subset of input sites. Shaded regions about all theory curves in this figure denote $\pm 1\sigma$ confidence intervals, including systematic errors relating to fluctuations in the applied unitary.

8.5.1.2 Measuring generalized bunching

Measurements of $\overline{p'_\kappa}$ show a clear separation between the distinguishable and bosonic visible behaviors (Fig. 8.9a). As expected for p_κ [302], this separation grows with the particle density $\rho = n/m$. We expect that the main source of distinguishability in our experiment is due to thermal motional excitations normal to the lattice (see Sec. 8.3). This motional DOF is well-approximated by a harmonic oscillator with motional quantum numbers n_{ax}, where atoms that possess different values of n_{ax} are distinguishable. We numerically show that this model for distinguishability leads to expected measurements that monotonically interpolate from the distinguishable to the indistinguishable bosonic case as the temperature is reduced. Our measurements of $\overline{p'_\kappa}$ are consistent with a thermal occupation of $\langle n_{\mathrm{ax}} \rangle = 0$ corresponding to the fully indistinguishable case, and inconsistent with significantly higher thermal occupation ($\langle n_{\mathrm{ax}} \rangle \gtrsim 0.167$). It is important to note that, due to the lack of postselection on the number of surviving atoms, measurements of $\overline{p'_\kappa}$ (and atom survival) are significantly more sensitive to the effects of single-particle loss and certain kinds of detection errors (see Sec. 8.2.2) than the measurements in Sec. 8.3. As a result, these effects must be calibrated and included in our simulations (Sec. 8.6). Despite being sensitive to calibration errors, the agreement between our measurements of $\overline{p'_\kappa}$ and low temperature simulations suggest that the indistinguishability measured via few particle calibrations in Sec. 8.3 is not noticeably degraded when scaling our experiments up to more particles. This agreement motivates experiments with large ensembles of atoms whose behavior we are unable to simulate.

The largest input patterns we study in this work contain 180 atoms (Fig. 8.9bc). Although these patterns can be prepared with no defects, in subsequent measurements we no longer enforce perfect rearrangement to avoid incurring significant overhead in the number of required experimental trials. However, because we image the atoms after rearrangement and before their evolution, we can identify the locations of any defects. This procedure results in a version of scattershot boson sampling [204], but with much less variation in input states than is typical [25]. Based on the few-particle characterizations of atom indistinguishability performed across relevant regions in the lattice (see Sec. 8.3.1), we expect the on-demand success rate for preparing a single ground-state atom, evolving it with no loss, and detecting its position

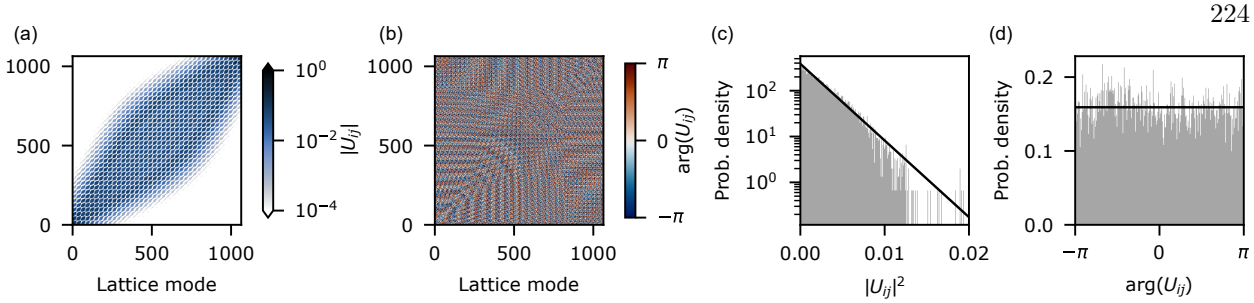

Figure 8.10: Properties of the single particle unitary. The modulus (a) and argument (b) of the single particle unitary U is depicted here for an evolution time of $t = 6.45$ ms, as is relevant to the measurements in Fig. 8.9. The finite waist of the optical lattice beams and resulting harmonic confinement means that U does not significantly couple all sites to each other, and thus is not Haar random. (c), (d) However, U is approximately Haar random when considering only a 15×15-site region near the center of the lattice. In this region, the distribution of the norm-square of the amplitudes in U (c) are well-captured by the Porter-Thomas distribution for 385 outputs (black line). The distribution of the phases in U (d) is well-captured by the uniform distribution (black line).

in an arbitrary site to be ~ 92 %. The dominant source of deviation from perfect boson sampling is from 5.0(2) % atom loss due to imperfect cooling in the in-plane directions, with additional contributions from imaging, rearrangement, and distinguishability. Given the above performance we expect that, on an average run of the experiment, ~ 166 of the 180 input sites are populated with identical bosons that evolve to 1015 output sites in the lattice with no loss or detection errors.

Although U couples all input sites to all output sites, only some of these couplings are significant, primarily due to the finite size of the lattice beams and resulting harmonic confinement (see Fig. 8.10ab and Sec. 7.3).[32] The harmonic confinement is negligible in a 15×15-site region near the center of the lattice, in which the distribution of the norm-square of the elements in U is well-captured by the Porter-Thomas distribution for 385 outputs (Fig. 8.10cd), indicating behavior that is close to Haar-random (in the sense that the distribution is spiky as a result of single-particle interference, and not necessarily that the distribution is truly Haar-random). Over the full 1015-site region considered in this section,[33] U couples each site to an average of 83 of the 180 input sites with an amplitude of $\geq 10^{-3}$. Some sites are coupled by $\geq 10^{-3}$ to as many as 156 input sites. While these numbers suggest that the experiment samples from a large effective state space, it is important to stress that this is not a certification of computational complexity, since there

[32] The effect of this harmonic confinement manifests in the "eye" shape of the amplitudes appearing in Fig. 8.10a, where, at a fixed evolution time, atoms starting further from the center of the lattice spread out less than atoms near the center.

[33] This region includes sites that can contain an atom at the latest evolution times explored, but for which the amplitudes of U are not necessarily well-captured by the Porter-Thomas distribution.

is no known efficient means to verify that the machine is behaving correctly in this regime.

8.5.2 Tuning particle distinguishability

One way to build additional confidence that the many particles maintain coherence during their evolution, leading to high sample complexity, is to tune the visible particle statistics into a regime that is simulable. Specifically, multiple experiments performed with different subsets of atoms in the full input pattern can be combined to simulate partially distinguishable atoms (Fig. 8.9bc), where only atoms that share a time label are able to interfere. We observe the expected qualitative behavior where, as the distinguishability is increased by increasing the number of labels, $\overline{p'_\kappa}$ is reduced. Once the atoms are sufficiently distinguishable it is once again possible to simulate their evolution, yielding good agreement with theory. Due to parity projection, the number of detected atoms after their evolution (Fig. 8.9c) serves as a similar measurement to $\overline{p'_\kappa}$, and is akin to standard measurements performed in the photonics community using click detectors [376, 375]. Measurements of the atom survival do indeed yield a clear separation between experiments with different visible particle statistics. Additionally, a theoretical prediction for the partially distinguishable case with 9 extra labels captures the distribution of atom survival probabilities (Fig. 8.9c). The above measurements serve as indirect evidence that the experiment with no additional time labels is behaving in line with expectations for bosonic particle statistics, and sampling from a very large effective state space.

8.6 Simulating non-interacting many-particle quantum walks

The conclusions drawn in Sec. 8.5 (and some of the conclusions in Sec. 8.3) primarily rely on comparing the performance of the experiment to simulations. For all simulations in Secs. 8.3 and 8.5, we use the spectroscopic calibration of the lattice Hamiltonian (Sec. 7.3) to infer the single particle unitary U. For up to 3 particles, we solve for the full output distribution exactly by evaluating the permanent in Eq. 8.18 (or in Eq. 8.19 for distinguishable particles) for all possible outputs via Glynn's formula. For larger particle numbers, we follow the approach of Clifford and Clifford [72] to sparsely sample from the full output distribution.

The simulations involving thermal occupation of the motional degree of freedom normal to the lattice in Fig. 8.9a assume that the evolution of this hidden degree of freedom is independent of the visible evolution of the atoms, and dephased by our state preparation. As a result, the evolution of atoms with a mixed motional state can be simulated by assigning a specific motional state drawn from the appropriate thermal distribution to each atom in a given simulated sample. For each subset of atoms that share a motional state we draw a sample using the approach of Clifford and Clifford [72], and combine these samples into a single sample of partially distinguishable atoms.

For experiments without postselection, we calibrate the false positive and false negative detection error rates (see Sec. 4.4.1), and take advantage of the fact that the absence of an atom, whether due to loss or detection errors, is expected to be an effect that commutes with the linear optical dynamics [44].[34] Given this observation, our simulations proceed by simulating samples of the experiment without imaging errors, applying single atom loss (including due to detection errors), then parity projection, and finally detection errors that lead to extra sites that are observed to contain atoms.

8.7 A note on noise and the hardness of boson sampling

Boson sampling [2] and most of its variants, including Gaussian boson sampling [137, 88], are strongly believed to be hard to simulate classically in the absence of noise.[35] In this scenario, the hardness of the sampling task is effectively characterized by the number of bosons n and number of modes m involved.[36] However, the hardness of these problems can change in the presence of noise. Such noise can broadly be categorized into non-unitary errors relating to state preparation, detection, and loss, and errors relating to uncontrolled parameters that govern the unitary evolution of the particles. Here, we compare our approach to pioneering experiments that study boson sampling using photons (Tab. 8.2) through the lens of the above errors. In summary, although our experiment achieves significantly lower non-unitary error rates than photonics experiments, an outstanding challenge is to characterize the applied unitary to a level that is comparable to what can be achieved in photonics experiments using large coherent states of many photons.

[34] Due, in part, to the absence of interactions (see Sec. 8.1.1).
[35] The main assumption being that the polynomial hierarchy does not collapse to the third level.
[36] Up to requirements that the unitary transformation applied between modes is sufficiently mixing [114, 243], and that one doesn't use a pathological input state (e.g. all bosons on one mode).

	n	$\mathcal{P}$	$\mathcal{J}$	r	m	Loss	Detection		Input	Evolution
[344]	20	0.975	0.93-0.954(1)	-	60	26 %	60-82 %	Click	Fixed	Fixed
[376]	50*	0.938	-	1.34-1.84	100	$\sim$ 45 %	73-92 %‡	Click	Fixed	Fixed
[211]	216*	-	-	$\sim$ 1.1	216	$\sim$ 67 %	95 %	Counting	Tunable	Tunable
[87]	50*	0.962	-	1.2-1.6	144	57 %**	-	PPNRD	Phase	Fixed
[367] (our work)	180	†	0.995^{+5}_{-16}	-	$\sim$ 1015	5.0(2)%	99.8(1) %††	Parity	Pattern	Hamiltonian

Table 8.2: Comparison of large-scale boson sampling demonstrations. For works involving Fock state boson sampling, n denotes both the particle number and the number of input modes. Works involving Gaussian boson sampling are marked with a *, in which case n corresponds only to the number of input modes. $\mathcal{P} = 1 - g^{(2)}(0)$ is typically referred to as the "purity" in photonics experiments, and is measured via second order correlations in Hanbury-Brown-Twiss-like experiments. †To the extent that these measurements characterize the single-particle nature of the input field [344], in our experiments $\mathcal{P} \simeq 1$ and is lower-bounded by our imaging fidelity of 0.998(1). However, our state purity is primarily limited by thermal motional excitations normal to the lattice, and can be estimated using the measured particle indistinguishability of $\mathcal{J} = 0.995^{+5}_{-16}$, which is an estimate of the purity assuming that the single-particle wavefunctions in the out-of-plane motional DOF are equal. To characterize state preparation, we list $\mathcal{J}$ for Fock state boson sampling results and the squeezing parameter r for Gaussian boson sampling results. m denotes the number of output modes in the linear optical network. "Loss" denotes the fraction of particles lost during evolution, including incoupling from the source to the linear-optical network, loss in the network, and outcoupling to the detectors. Detection is characterized by the detection efficiency, and the type of measurement performed on each output mode. "Click" refers to detecting the presence or absence of particles, "parity" to detecting particle number parity, "PPNRD" to pseudo-photon-number-resolving detection, and "counting" to full particle number-resolved readout. The work marked with ‡ includes fiber coupling loss in the estimate of detection efficiency, and the work marked with ** includes detection efficiency in the quoted value for loss. ††The listed value for our work is a detection fidelity rather than an efficiency, and includes contributions from both particle loss and infidelity. Converting the other listed values to detection fidelities would involve including the effects of leakage light and dark counts, resulting in slightly lower values. "Input" refers to the class of states that can be prepared as inputs to the linear optical network, with "phase" referring to tunability of the phases of the prepared squeezed states, and "pattern" to nearly arbitrary Fock states with occupations of 0 or 1 on each input mode (see Sec. 4.5). "Evolution" refers to the family of linear optical networks that can be applied in a given system, with "Hamiltonian" referring to unitary evolution for variable time under a fixed Hamiltonian. For both "input" and "evolution", "fixed" refers to a single instance, and "tunable" to flexible, but not universal, programmability. The numbers appearing in this table are representative values for approximate comparison only, please refer to the original publications for details.

8.7.1 Non-unitary errors

In photonics experiments, it can be challenging to generate and interfere large Fock states of photons due to transmission losses, and the probabilistic techniques that are often [256, 249, 284, 46, 317, 329, 25, 54, 250] (but not always [202, 344]) used for single photon generation. To circumvent this difficulty, recent experiments have performed modified versions of boson sampling that take advantage of more easily accessible non-classical states of light [204, 25], including demonstrations of Gaussian boson sampling [376, 375, 211, 87].

The above demonstrations have focused on scaling up the number of modes in the linear optical network, as well as the number of photons detected at the output of the network. For example, in [211], they achieve impressive results involving up to 219 photons detected in 216 modes. However, recent theoretical results suggest that the high transmission losses in these experiments ($\sim 67\%$ in [211]) make it possible to efficiently simulate their behavior classically. In particular, because loss occurring either during evolution or detection is expected to commute with the applied linear optical evolution [44],[37] one can treat all realistic non-unitary errors as a state preparation error, resulting in reduced quantum correlations on top of a classical background. State-of-the-art tensor network algorithms take advantage of this simplification to simulate the largest photonics-based boson sampling experiments in about an hour [244]. Therefore, beyond n and m, important figures of merit for the hardness of a given boson sampling demonstration include the state preparation fidelity, loss, and detection errors.

In our experiment, we benefit from low loss that is not observed to scale with evolution time (see Sec. 8.2.2) [114], high state preparation and detection fidelity, and many output modes (Tab. 8.2). Although it is likely that modern tensor network approaches developed for Gaussian boson sampling [244] can be modified to simulate our experiment, it is not expected that these approaches will provide a significant speedup over preexisting simulations [72] (Sec. 8.6) given the loss rates and detection fidelities achieved in our experiment.

[37] Given certain assumptions regarding how the linear optical network is implemented.

8.7.2 Unitary errors

The comparably low non-unitary error rates achieved in our experiment, and the challenges they impose for classical simulation, are encouraging. However, another source of errors comes from the imperfect application or characterization of a given target unitary. In our experiment, these errors primarily come from fluctuations in the intensity of the optical lattice, and thus a Hamiltonian whose energy scale fluctuates:

$$H(s) = sH_0, \tag{8.28}$$

where s is an unknown unitless real number, and H_0 is the target Hamiltonian that generates the desired unitary transformation. Assuming that s follows a Gaussian distribution with mean 1 and standard deviation σ_s, we show in [367] that the fidelity $\mathcal{F}$ of the prepared state to the ideal state with no unitary errors obeys the lower bound:

$$\mathcal{F} \geq \exp\left(-\frac{(n\sigma_s W t)^2}{2}\right). \tag{8.29}$$

Here, $W \simeq 8J$ is the bandwidth of the single particle Hamiltonian (Eq. 7.9) and n is the atom number.[38] In our experiments, $W \simeq 1$ kHz and $t \leq 6.45$ ms. Fast shot-to-shot fluctuations in the lattice intensity result in $\sigma_s \simeq 10^{-3}$,[39] which does not significantly affect our experiments with up to $n = 180$.

However, the uncertainty in our characterization of the applied unitary is much larger, due to the combined effect of drifts and the time-consuming calibration procedures we currently use (see Secs. 7.3 and 8.4). The resulting value of $\sigma_s \gtrsim 10^{-2}$ dramatically affects the fidelity of the prepared state.[40] It is very possible that such unitary errors affect the hardness of the sampling problem realized, although I cannot find much discussion of this limitation in the boson sampling literature.[41]

Additionally, like in previous large-scale demonstrations of boson sampling [376, 375, 211, 87], we currently can only apply a restricted family of unitaries. These unitaries possess additional structure that could, in principle, be taken advantage of in efficient classical simulations. Haar-averaging of the unitary

[38] Note that to obtain an evolution time when all input sites are overlapping (in a 2D lattice), we need that $t \gtrsim \sqrt{n}$, meaning that the exponent in Eq. 8.29 scales like n^3.

[39] These errors are comparable to errors due to shot-to-shot fluctuations in the optical path length achieved in photonics experiments, e.g. [376].

[40] Note that residual interactions between the atoms enter at a similar order (Sec. 8.3.3). However, these interactions do not fluctuate significantly in strength, and are not expected to affect the hardness of the sampling problem [374].

[41] Note that similar errors can occur in photonics experiments given limitations on phase locking large linear optical networks [376, 211].

would remove the possibility of such simulations, and provide access to additional tests of boson sampling that rely on random matrix theory [304, 302, 33]. Based on the results in Ch. 7, universal control over the single particle unitary (see Sec. 7.6) can in principle be implemented using the same optical tweezers we use for atom rearrangement without introducing additional loss. However, an important open problem is to improve the efficiency of protocols for directly calibrating the applied unitary for large systems, as well as improving the stability of our experiment.[42]

8.8 Interacting quantum walks

The tools demonstrated in the previous sections for controlling non-interacting quantum walks are just as easily applied to interacting systems. As discussed in Sec. 8.1.1, one way to reintroduce interactions into our quantum walks is to simply work with a different isotope of Sr. In this section, we discuss preliminary experiments we have performed with ^{86}Sr, which provides strong interactions on the scale of $U_c/J \gtrsim 8$ in our experiments.[43] Note that, despite the fact that non-interacting quantum walk dynamics can be challenging to simulate (as we saw in the previous sections), the lack of interactions still results in a significant simplification in the behavior of the system.[44] However, incorporating almost any interaction into quantum walks of identical particles makes such quantum walks universal for quantum computing [64].[45] Here, we will briefly discuss preliminary experiments we have performed involving both interacting dynamics, and the preparation of a non-trivial and interacting ground state.

[42] In particular, note that the quoted fluctuations in H apply only to the lattice potential. We expect that drifts in the tweezer potential relative to the lattice (see Sec. 3.3.5) will result in significantly larger errors in our current apparatus.

[43] Upgrades are underway to work with ^{84}Sr, which will provide intermediate interaction strengths of $U_c/J \sim 1$.

[44] For example, boson sampling is not believed to be universal for quantum computing, although the addition of mid-circuit measurement and feed forward is enough to recover universality [176].

[45] This is distinct from the universality of single particle quantum walks [59] — although single particle quantum walks can be universal, they operate on a Hilbert space that grows linearly in physical resources (see Sec. 7.7). Interacting multiparticle quantum walks are universal on a Hilbert space that grows exponentially in the particle number, and therefore can in principle fulfill the full promise of efficient and universal quantum computing.

8.8.1 Hard core bosons and fermionization

Our use of ^{86}Sr results in $U_c \gg J$, which is often referred to as the "hard core" limit. In this limit, the bosonic operators in Eq. 8.1 take on modified commutation relations [139]:

$$\begin{aligned} [a_i, a_j] = [a_i^\dagger, a_j^\dagger] = [a_i, a_j^\dagger] = 0 \quad \text{for} \quad i \neq j \\ \{a_i, a_i\} = \{a_i^\dagger, a_i^\dagger\} = 0, \quad \{a_i, a_i^\dagger\} = 1, \end{aligned} \tag{8.30}$$

where the relations on the second line simply encode the fact that double occupancy of a given site is energetically forbidden. Other than the modified operator algebra in Eq. 8.30, the Hamiltonian governing hard core bosons (HCBs) is the same as in Eq. 8.2. Notice the similarity of Eq. 8.30 to the fermionic anticommutation relations:

$$\{f_i, f_j\} = \{f_i^\dagger, f_j^\dagger\} = 0, \quad \{f_i, f_j^\dagger\} = \delta_{ij}, \tag{8.31}$$

where $f_i^\dagger$ is the fermionic creation operator on site i (and f_i the corresponding annihilation operator). One can map the behavior of HCBs to fermions by using the Jordan-Wigner transformation, which restores the full set of anticommutation relations associated with fermions by including the appropriate phase factors in the operators [274]:[46]

$$a_j^\dagger = f_j^\dagger \prod_{l=1}^{j-1} e^{-i\pi f_l^\dagger f_l}, \tag{8.32}$$

and similarly for a_j. For nearest-neighbor tunneling in 1D, namely where:

$$H = -J \sum_i (a_i^\dagger a_{i+1} + a_{i+1}^\dagger a_i) + \sum_i V_i a_i^\dagger a_i, \tag{8.33}$$

the phase factor in Eq. 8.32 cancels[47] leading to a Hamiltonian describing non-interacting (or free) fermions:

$$H = -J \sum_i (f_i^\dagger f_{i+1} + f_{i+1}^\dagger f_i) + \sum_i V_i f_i^\dagger f_i. \tag{8.34}$$

Therefore, for quantum walks on a 1D chain, the only difference between HCBs and free fermions is the phase factor appearing in Eq. 8.32. This phase factor is irrelevant for measurements involving particle positions,

[46] Note that the HCB system can be mapped to a spin model by observing that each lattice site only has two allowed states (0 or 1 bosons), and therefore can be represented with a single spin 1/2 degree of freedom. From this point the Jordan-Wigner transformation can proceed in the standard way.

[47] Note that in the hopping terms, the cancellation of the phase factors additionally relies on the fact that a fermion cannot hop onto an occupied site.

or correlations between particle positions (namely for the sampling problems discussed in this chapter), and so all predictions for such measurements are identical for HCBs and free fermions in a 1D chain. The above equivalence is sometimes referred to as the "fermionization" of HCBs.

Interestingly, the fermionic equivalent of boson sampling (fermion sampling) is trivial to simulate classically,[48] since the anticommutation relations result in a determinant[49] rather than a permanent in Eq. 8.18 [2, 328]:

$$P^{\mathrm{F}}(\boldsymbol{k}|\boldsymbol{j}, U) = |\mathrm{Det}(U_{\boldsymbol{kj}})|^2, \tag{8.35}$$

where P^{F} is the probability of observing a particular output event for fermions. Note that the normalization appearing in Eq. 8.18 is 1 here due to Pauli exclusion and the resulting lack of collisions. In our experiment, we expect Eq. 8.35 to describe the behavior of ^{86}Sr atoms in a 1D lattice (up to small corrections due to imperfect realization of the HCB condition).[50]

Interactions break the separability of quantum walks on square lattices discussed in Secs. 7.2.1 and 8.2.3.[51] Therefore, in order to study the dynamics of HCBs in 1D, we isolate individual rows or columns of the lattice using optical tweezers.[52] This is done either by applying 515 nm tweezers to all sites in the 1D region of interest, or by applying tweezers to sites away from the region of interest to form a pair of barriers.[53] In both cases, we have experienced issues with atoms leaking out of the 1D region of interest into other sites in the lattice on a timescale of $\sim 10 - 100$ ms. Although this is not an issue for short time dynamics,

[48] I have long wondered about whether there is a deep reason that there is an inversion in the computational complexity associated with bosons and fermions with and without interactions. Whereas non-interacting fermionic dynamics are trivial (and non-interacting bosonic dynamics complicated, as we saw in previous sections), the opposite is true when studying interacting ground states. Here, the non-negativity of bosonic ground states makes quantum Monte-Carlo techniques very effective, whereas properties of fermionic ground states remain challenging to compute due to the sign problem [335]. As far as I can tell, the special case of free fermions stands alone because the cancellations that make the determinant efficient to compute are very fine-tuned (i.e. most matrix immanants are hard to compute [78]), but I would be very interested in any insights the reader may have.

[49] Which can be computed efficiently, for example via Gaussian elimination (although more efficient polynomial time algorithms for computing the determinant also exist).

[50] The ability to change the effective particle statistics of the atoms (either by studying HCBs in 1D, or by changing to a fermionic isotope) could be useful in future schemes for benchmarking boson sampling in atomic systems. Specifically, without changing other aspects of the experiment, one can switch between a situation that is easy to simulate and therefore benchmark, and a situation that is hard to simulate. Such a change would be challenging to implement using photons, especially without affecting other aspects of the experiment [284, 219].

[51] Because the presence or absence of interactions depends on the exact location of a given particle in the lattice, rather than just the occupied row or column.

[52] Note that the main utility of the mapping between 2D and 1D in the non-interacting case was to suppress the effect of parity projection (see Sec. 8.2.2). In the case of HCBs, double occupancy and thus parity projection is suppressed due to interactions, and so mapping from 2D to 1D is no longer necessary for achieving measurements in an effectively complete basis.

[53] Note that the 515 nm tweezers are attractive, meaning that the former approach leads to sites in the region of interest being lower in energy than other sites in the lattice, whereas in the latter approach sites in the region of interest are higher in energy.

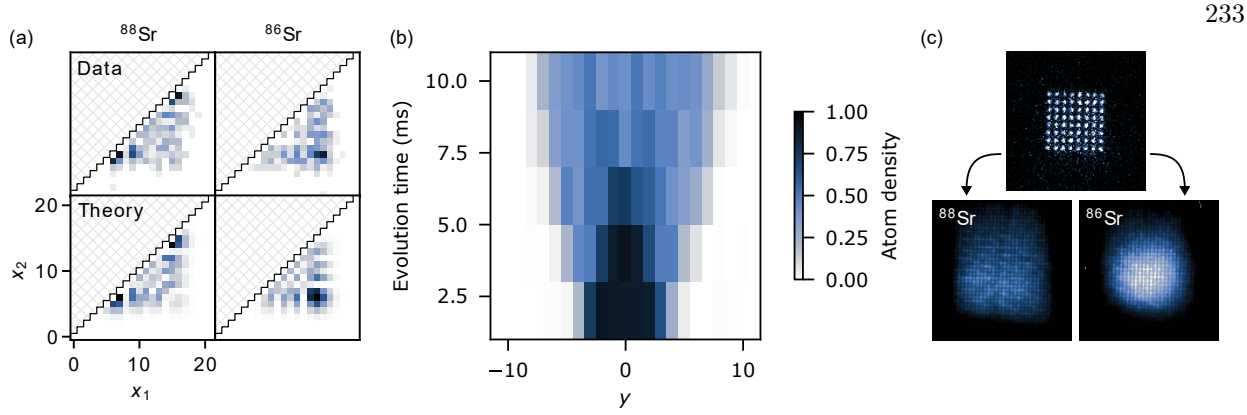

Figure 8.11: Dynamics of hard core bosons (preliminary). (a) Two particle quantum walks in 1D (not appealing to binning 2D quantum walks) exhibit dramatically different behaviors for ^{88}Sr, which behaves like a non-interacting boson, and ^{86}Sr, which behaves like a hard core boson (HCB). In the former case we observe similar interference to that of previous sections, with constructive interference along the diagonal and destructive interference along the anti-diagonal in the space of $\{x_1, x_2\}$ detection events. In the latter case, the fermionization of HCBs leads to signatures of fermionic interference — namely constructive interference along the anti-diagonal, and destructive interference along the diagonal. (b) Similar experiments with quantum walks of many hard core bosons in 1D are underway. Here we display the density distribution for a quantum walk of 7 HCBs in 1D. (c) In 2D, one can qualitatively observe a crossover between ballistic and diffusive transport for non-interacting bosons and HCBs. Here, a regular pattern of atoms is prepared (single shot image, top), and the atoms are allowed to tunnel for 8.5 ms. Averaged images (bottom) exhibit the approximately rectangular distribution characteristic of non-interacting quantum walks for ^{88}Sr (as explored in previous sections), but a diffusive, approximately Gaussian distribution for ^{86}Sr.

it does significantly limit our ability to perform adiabatic ramps of the lattice parameters. One might be concerned that applying tweezers to every lattice site would pose challenges in terms of inhomogeneity, but it is important to note that this approach is distinct from performing tunneling directly in the tweezers. Here, the tweezers are providing a shift in the chemical potential that makes certain tunneling processes off-resonant, requiring depths on the scale of $\sim 10J \simeq 1$ kHz. This is to be compared with depths of $\gtrsim 5E_{\mathrm{R}} \simeq 10$ kHz (where $E_{\mathrm{R}}$is the recoil energy associated with our lattice spacing) for realizing tunneling (in the tight binding limit) directly between the tweezers.

Two particle quantum walks in 1D performed using ^{88}Sr and ^{86}Sr (Fig. 8.11a) nicely illustrate the different dynamics resulting from bosonic and fermionic interference, with the familiar bunching and clouding from Sec. 8.3.2 exhibited in the former case, and anti-bunching and anti-clouding in the latter case.[54] Similar arguments to those in Sec. 8.3.2 apply to the case of multiparticle fermionic quantum walks in 1D (and by extension, HCBs in 1D), leading to constructive interference along the leading edge of the anti-diagonal(s)

[54] Namely, an enhanced probability for the particles to be far apart from each other.

in the virtual lattice.

Note that although the phases in Eq. 8.32 do not affect the HCB sampling problem, they are still physical, and can result in interesting effects. For example, they are responsible for what is known as "dynamical quasicondensation" [274], which has been investigated previously in the quantum gas community [343]. These studies focus on the situation where particles initialized at nearest neighbor spacing are allowed to expand in 1D. Although the density profiles (and correlation functions involving density) of free fermions and HCBs undergoing such expansion in 1D are identical, their description in the momentum basis is different. Free fermions initialized in Fock states are in a superposition of all momentum states, and remain this way because they are non-interacting. In contrast, a significant fraction[55] of the HCBs condense into modes with definite wavenumbers $k = \pm\pi/(2a)$ (where a is the lattice spacing) [274]. This partial (quasi) condensation is precisely due to the non-local phase factors appearing in Eq. 8.32. To investigate dynamical quasicondensation in our system, we have extended the two particle experiments in 1D described above to more particles (Fig. 8.11b). However, it remains to perform direct measurements of the momentum after performing these experiments.

Although the above picture mapping HCBs to free fermions breaks in a 2D lattice, the effect of dynamical quasicondensation is expected to persist [156]. However, simulations of such effects are challenging without making certain mean field approximations [156]. To our knowledge, quasicondensation has not yet been observed experimentally in more than one dimension, and performing such experiments to verify the accuracy of approximate (classical) simulations is an ideal proving ground for quantum simulators. Additionally, quasicondensation has only been explored in the context of quenching a Mott insulating state.[56] The ability to rapidly prepare arbitrary patterns of atoms in our experiment opens up the unique possibility of studying similar dynamics starting from other input patterns, which could dramatically alter the behavior of the atoms in momentum space.[57]

[55] $\sim \sqrt{n}$, where n is the number of particles. Note that the fact that the number of condensed particles does not scale like n is the sense in which this is only "quasi" condensation.

[56] Namely, a state where an equal, integer number of atoms are prepared on each lattice site in a given contiguous region.

[57] Similar to how changing the input pattern can switch clouding on and off in 1D CTQWs of (possibly interacting) bosons.

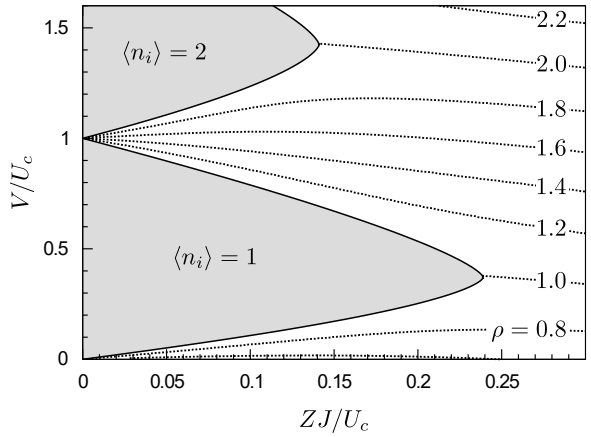

Figure 8.12: Phase diagram of the Bose-Hubbard model as a function of the chemical potential V and the tunneling energy J (normalized to the contact interaction U_c). Z is the coordination number of the lattice (e.g. 2 for a 1D lattice, and 4 for a 2D square lattice). The Mott insulating phase is shown in the shaded regions, where the number of atoms on each lattice site n_i takes on well-defined integer values. The white region corresponds to the superfluid phase. The dotted lines indicate contours of constant particle density ρ. Notice that for any non-integer value of ρ, the system is superfluid. This figure is adapted from [193].

8.8.2 Assembling a superfluid

Beyond the studies of dynamics described in the preceding sections, we are also interested in exploring the utility of optical tweezers for preparing interesting ground states (or other steady states) in different Hubbard systems. Along similar lines to the experiments in Sec. 7.4, we can take advantage of programmable tweezer potentials to adiabatically connect a state that is easy to prepare — namely patterns of ground state cooled atoms in the lattice — to a different ground state of interest.

For bosons with repulsive interactions, the zero temperature ground state phase diagram is shown in Fig. 8.12. There are two possible phases, including a Mott insulating (MI) phase where there is a well-defined, integer number of atoms on each lattice site,[58] and a superfluid phase, where there is long range phase coherence[59] between different sites in the lattice.[60] Typically, the Bose-Hubbard phase diagram is studied in quantum gas experiments by modifying the chemical potential, leading to the observation of isolated MI and superfluid regions depending on the local chemical potential [15]. Our experiment allows for operation in a unique, and perhaps somewhat counter-intuitive regime where the number of atoms in the

[58] The MI phase additionally has zero compressibility, and a finite energy gap for generating particle-hole excitations.
[59] In the sense that the relative phase of the wavefunction on two sites in the lattice is well-defined, in contrast to the MI phase.
[60] The superfluid phase additionally has finite compressibility, and (sometimes) a vanishing energy gap.

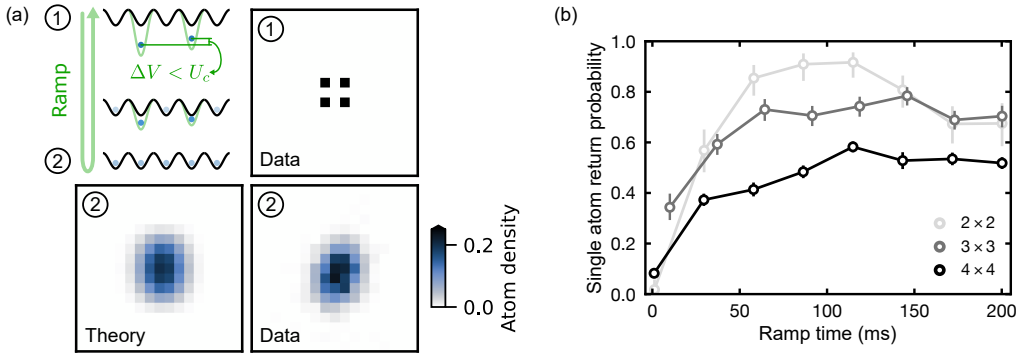

Figure 8.13: Preparing a superfluid ground state using optical tweezers (preliminary). (a) Scheme for preparing a superfluid at fixed atom density (top left) using an adiabatic ramp of optical tweezers (green) interfaced with an optical lattice (black curves). Preparation tweezers are used to implant atoms in a subset of sites in the lattice (top right), and then the tweezers are adiabatically ramped off. When the tweezers are extinguished, the expected ground state of the system is superfluid, with the atoms delocalized across the lattice. After the tweezers are extinguished adiabatically, we measure an atom density (bottom right) that is in line with a theoretical prediction based on independent characterizations of our lattice potential (bottom left). (b) To provide evidence that the prepared state is indeed a ground state (or at least an eigenstate), we reverse the adiabatic ramp, and measure the probability of returning to the original atom distribution as a function of the duration of the ramps. The single atom return probability is the probability that a given site in the initial pattern is populated with an atom after the forward and reverse ramps. For fast ramps that are not sufficiently adiabatic, the return probability is negligible. For slower ramps and for input patterns containing up to 3×3 atoms, we observe a significant probability of recovering the original patterns, even after holding for many tunneling times between the forward and reverse ramps. For 4×4 input patterns, the probability of perfectly recovering the input pattern is negligible, but there is an enhanced probability to recover a pattern that is close to the original (in comparison to a model where the tweezer-addressed sites are randomly populated with atoms, as one might expect for a thermal state).

system is fixed. If the chemical potential is uniform, and the number of atoms in the system is not an integer multiple of the number of lattice sites, the Bose-Hubbard ground state is always superfluid (see Fig. 8.12).[61]

Notice that the states that we can prepare via rearrangement and optical cooling look a lot like a MI state, with exactly one atom on each lattice site (or a subset of sites). With this in mind, our procedure for preparing the superfluid state is as follows (Fig. 8.13a): we prepare a pattern of atoms with the desired density, and place a tweezer on each site that contains an atom. The tweezers are made deep enough that tunneling is off resonant, and so the prepared state corresponds to the MI state, just on the sublattice (or other pattern of sites) that is marked by the tweezers. The tweezers can then be adiabatically ramped

[61] One can think of this in the following way: HCBs obey a kind of Pauli exclusion principle in real space, leading to MI behavior. However, unlike fermions, they do not obey such an exclusion principle in momentum space and so can condense into plane-wave-like states. Even for large atom numbers (> 1 atom per site), one can think of the coherence in the superfluid phase as being the result of a small number of condensed atoms on top of a MI background.

off to prepare the ground state in the regime where there is significant tunneling. Preliminary results performing the above protocol using modest numbers of atoms $(4-16)$ are shown in Fig. 8.13ab. Note that the density profile of the atoms after the forward ramp matches expectations based on simulations using the density matrix renormalization group (DMRG) and independent calibrations of the lattice potential (Sec. 7.3). Additionally, we can reverse the adiabatic ramp and recover the original pattern prepared after rearrangement with a fidelity of $\sim 70\%$ for four atoms, even when holding the atoms for many tunneling times (> 50 ms) between the forward and reverse ramps. This recovery suggests that we have indeed prepared an eigenstate of the Bose-Hubbard Hamiltonian under conditions where we expect the ground state to be a superfluid. However, it remains to perform a direct measurement of the phase coherence expected in the superfluid state, for example via well-established time of flight techniques [129].

The additional control provided by optical tweezers opens up a variety of future directions for both state preparation and measurement. For example, one could isolate subsets of sites in the lattice using the tweezers to efficiently (with respect to computational resources and the number of required experimental trials) measure certain non-standard observables [332], including the phase coherences characteristic of the superfluid state [172].[62] Although it is not yet clear that the tweezers can help prepare arbitrary ground states with higher fidelity than conventional techniques involving evaporation, they are already helpful for preparing specific patterns of identical atoms rapidly and with high fidelity. Critically, the schemes that we have explored so far for preparing the superfluid ground state do not take advantage of the full set of controls that our apparatus provides.[63] It is possible that more sophisticated variational [57, 330], time optimal [91], or circuit inspired [251] techniques could offer significantly improved performance. We have just started to scratch the surface of what can be achieved with Hubbard quantum simulators that possess additional programmable controls, and are excited by what such systems could offer in the future [79].

[62] Many of these ideas come to us from Daniel Mark.

[63] Including control over the initial states of atoms that we can prepare and the chemical potential and tunneling energies in the lattice, and also over internal degrees of freedom as discussed in Ch. 5 and 6.

Chapter 9

Conclusion and outlook

"Thank you for a most enjoyable game."

— HAL 9000

In this work, we have endeavored to apply the flexible control provided by optical tweezer arrays to frequency metrology, namely by working with the alkaline earth atom (AEA) strontium, and to simulations of Hubbard models, by incorporating features of a quantum gas microscope (QGM) in our experiment.

In the direction of frequency metrology, we demonstrate fast assembly of large ensembles of Sr atoms, and engineer the tweezers to preserve the coherence of the atoms. High-fidelity readout in a complete basis provides access to new observables in our experiments, enabling the use of various correlation spectroscopy techniques. We further engineer high fidelity single qubit gates, and entangling gates using Rydberg interactions. The achieved fidelities and coherence times are close to the state-of-the-art for any neutral atomic qubits [103], but our clock qubits have the additional benefit that the qubit states are separated by an optical frequency, and are thus useful for frequency metrology. The above capabilities enable synchronous frequency comparisons with competitive fractional frequency instabilities at the 10^{-17} level at 1 s of averaging, and separate measurements using entangled states that achieve instabilities below the standard quantum limit.

A key outstanding challenge is to preserve the controls provided by tweezer clocks without sacrificing on other desirable features, like low systematic uncertainty and large atom numbers (and thus low instability). Because the controls provided by tweezer arrays are primarily used during state preparation, and possibly readout, and because the various control beams can be fully extinguished during clock interrogation, the prospects for maintaining accuracy in such systems is fairly promising. However, tweezer clocks currently

make significant sacrifices in terms of atom number in comparison to leading lattice clocks, which work with hundreds of thousands of atoms [35]. Depending on the specific sensing problem at hand, universal control over clock qubits is not always necessary [92, 301], and so it could be beneficial to give up some control in favor of larger system sizes, for example by generating useful large-scale entangled states without resorting to the brute force solution of a universal gate set [254, 277, 96]. With the above considerations in mind, it is likely that the main role of tweezer clocks will be to serve as universal testbeds for developing new measurement protocols, especially those that draw on ideas developed in the context of quantum computing [171, 86, 84, 377]. Lessons learned from these experiments can then be applied to more specialized quantum sensors that are optimized for specific measurement tasks.

In the direction of simulating Hubbard models, we demonstrate fast and high fidelity state preparation in a QGM using tweezer rearrangement and optical cooling. We also demonstrate the ability to programmably modify the lattice potential using additional optical potentials projected onto the lattice.[1] Using the above tools, we perform spatial search using quantum walks of single atoms, large-scale boson sampling using quantum walks of many atoms, and proof of principle preparations of a superfluid. It is worth reflecting on the strengths and weaknesses of our approach to state preparation in comparison to standard techniques used in QGMs. One useful capability that is uniquely provided by the rearrangement-based approach to state preparation is the possibility of performing powerful kinds of postselection. For example, postselection based on measurements taken both before and after the dynamics of interest allows us to study many particle quantum walks in the absence of atom loss and collisions. Such postselection would not be possible in conventional QGMs, where the procedure for state preparation precludes taking an image before the dynamics.

For the specific states prepared using rearrangement and optical cooling — namely Fock states of 0 or 1 atom per site — we reach state preparation fidelities that are comparable to or exceed the performance of standard QGMs.[2] These states serve as a strong starting point for studying dynamics, for example problems in transport and thermalization (see Sec. 9.2), or in the demonstrations of boson sampling performed in this

[1] It is worth noting that this capability is not unique to our experiment — similar techniques are already in use in other QGMs (e.g. in [379]).

[2] This is in part due to the additional overhead required to prepare such states using "cookie-cutting" techniques to isolate a subset of atoms from a larger many body state in standard QGMs [349, 151, 231, 264, 166, 359].

work. It is also worth noting that we expect there is significant room to improve the fidelity of our state preparation, since the remaining errors are well understood, and primarily due to the optical cooling step. The performance of the optical cooling is consistent with the ratio between trap frequency and the linewidth of the cooling transition, suggesting that cooling schemes based on the clock transition could result in significantly improved performance [372, 271].

Although the above results are encouraging, the degree to which tweezers help with preparing interesting ground states is not yet clear. For example, we still rely on slow adiabatic ramps to prepare the superfluid ground state, as well as the resource state used for spatial search. The utility of atom rearrangement for studying ground state physics will hinge on the existence of states that are both easy to prepare with the tweezers, and can be adiabatically connected to ground states of interest efficiently.[3] There is some hope that this could be the case, especially when, as in our experiment, one has access to single and two qubit gates in addition to atom rearrangement and cooling (see Sec. 9.2).

Although, in principle, the tweezers provide universal control over quantum walk dynamics in the lattice (Sec. 7.6), and the ability to perform circuit-like operations between sites in the lattice (Sec. 8.8.2), these operations currently are prone to drifts and calibration errors.[4] Even the improved calibration schemes we develop in this work remain costly to run in terms of the required experimental trials, and cannot yet be applied to large systems due to the comparable timescales associated with the calibration procedure, and drifts in the experiment.[5] Improving the fidelity of tweezer-controlled operations in the lattice is a key outstanding challenge that could be addressed by reducing drifts on the experiment, or by developing faster calibration schemes.

9.1 Directions in metrology

There are a number of exciting directions in performing frequency metrology using tweezer arrays. One can show that the states that are most sensitive to a global signal (in the sense of Sec. 5.1.1) are in the

[3] Specifically, that the requirements on adiabaticity are similar to (or more relaxed than) those for adiabatic ramps starting from states that are easily prepared via techniques like evaporation [30].

[4] Primarily uncompensated drifts in the alignment between the tweezers and the lattice.

[5] Although, fortunately, these schemes scale only polynomially in the number of lattice sites, providing some hope for accessing high fidelity operations on large systems in the future.

symmetric subspace, since the quantum Fisher information is upper bounded by the length of the total spin vector [259]. However, often, as in Sec. 6.5, we only have access to interactions that explicitly take us out of the symmetric subspace. One interesting direction is to devise improved schemes that use local interactions to generate metrologically useful states (namely ones that reside in the symmetric subspace). This could be achieved by introducing the appropriate transverse field to the interactions used in Sec. 6.5 [369], or via variationally optimized sequences using the same transverse field and interactions [160, 161, 161]. Given possible challenges associated with collective loss mechanisms in large ensembles of Rydberg atoms [370, 38, 368, 135, 107], another possibility is to perform variational optimization of a simple gate sequence that only involves pairwise interactions (which we know can be high fidelity based on measurements presented in this work).

More broadly, being able to move the atoms without destroying the coherence of the clock qubits would, in combination with the results demonstrated in this thesis, constitute a universal gate set.[6] Such a gate set would open the door to a variety of directions, including schemes to mitigate [170] or correct [182] errors to improve the performance of a quantum sensor. The required moves can be performed with mobile 813 nm tweezers [301], although using a different magic wavelength that is more amenable to making tightly confining traps (like the magic wavelength near 497 nm) could be helpful when attempting to perform rapid moves in a large array.

The additional controls provided by our system could be particularly valuable when considering more general formulations of the sensing problem. For example, one could take a Bayesian approach to sensing and optimize entire sequences of measurements to more rapidly converge on a precise estimate of the parameter of interest given prior knowledge about the parameter [228]. Another interesting possibility is to study situations where the symmetry of the sensing problem is explicitly broken by a signal or noise that does not act uniformly on all qubits [181].[7] In these situations, the enhancements provided by entangled states can be more significant than when one is detecting a global signal.

[6] Similar capabilities have been demonstrated in the alkali atoms [32], however, clock qubits suffer from the additional challenge that the qubit frequency is significantly more sensitive to the exact characteristics of the trap.

[7] For example when the signal varies in space, and acts on an extended array of atoms.

9.2 Directions in simulations of Hubbard models

We have just started to explore interacting Hubbard models in our apparatus, and are excited by the near- and long-term directions enabled by tweezer-controlled QGMs. One particularly exciting near-term direction, mentioned earlier, is to use the fast preparation of excited states in our apparatus to study transport problems. For example, the high repetition rate of the experiment could allow for studies of density- or energy-dependent modifications to the diffusion constant [118, 36], which can be challenging to distinguish from typical diffusion without significant statistics and large system sizes. The ability to repeatedly prepare specific patterns of atoms could also allow one to study how a given input pattern seeds the dynamics leading to universal transport behaviors [252, 36, 347]. The ability to dynamically alter the lattice potential further provides access to new observables. For example, using an appropriately designed quench of the local chemical potential imposed on the lattice using tweezers, one could measure separations between charge and energy transport by directly measuring the local energy density [332].

In the direction of non-interacting quantum walks, fine-tuned programming of the quantum walk unitary could enable stronger certifications of the many particle interference at the core of boson sampling [327]. Additionally incorporating a controllable label degree of freedom, encoded in the axial motion of the atoms, or in their nuclear or electronic states, would enable studies of new and counter-intuitive kinds of interference between partially distinguishable particles [298, 262]. The ability to prepare the appropriate entangled states on the label degree of freedom would allow for more exotic particle statistics, including the ability to simulate fermions [284], and certain kinds of interference problems involving anyons [219].

In some applications, the clock qubit is an attractive candidate for a controllable label. This is especially true in cases where one wants the different qubit states to have dramatically different properties. For example, the large energy separation between the clock qubit states makes it straightforward to engineer spin-dependent optical potentials [142], which has enabled studies of mass-imbalanced Hubbard models [81, 80], and possible future studies of Kondo lattice models [109, 108]. However, the clock qubit introduces additional complications associated with the strong inelastic (and elastic) interactions between pairs of clock state atoms, as well as, to a lesser degree, between clock and ground state atoms (see Sec. 8.1.1).[8] In

[8] We have performed some exploratory measurements with tunneling clock state atoms, and find reasonable performance for

applications where one does not want the label to affect the evolution of the system (other than modifying the particle statistics), one possibility is to map the state of the clock qubit to the axial motion of the atoms (by simply driving a π-pulse on the red sideband of the clock transition).[9] Another possibility is to take advantage of the nuclear spin of fermionic isotopes of Sr (see below).

In this work, we control the strength of the interactions between atoms by simply working with different isotopes of Sr, which provides limited tunability. More flexible control of the interactions could be achieved by taking advantage of the Rydberg dressed interactions we demonstrated in the context of performing entangling gates between clock qubits. Proof of principle demonstrations using itinerant, Rydberg-dressed atoms have already yielded exciting results [135], allowing one to study extended Hubbard models with interactions that are tunable in both strength and range. Although there are challenges associated with minimizing inhomogeneity in the lattice introduced by the dressing light,[10] as well as the collective loss mechanisms described earlier, we suspect that such challenges could be addressed using the appropriate Trotterized control sequences. For example, to address lattice inhomogeneity, one could perform a spin echo sequence similar to the gate sequence in Sec. 6.2 that removes the contribution from the single atom light shift. Trotterized evolution involving interactions generated using the above sequence interspersed with tunneling could introduce long-range interactions without any lattice inhomogeneity. Furthermore, it has been shown that short dressing pulses like those that would appear in the above Trotterized evolution are beneficial because they allow atoms in contaminant Rydberg states to decay before they seed any collective decay processes, leading to a dramatic improvement in coherence time [141]. For the fidelities and lifetimes achieved in our experiments, we expect that one could study dynamics involving up to hundreds of the described Trotter steps.

It is worth noting that the toolset associated with tweezer-controlled QGMs is broadly applicable to other controllable particles, including dipolar atoms [66, 321] or molecules [278, 192] that natively realize long range interactions in their ground states. Working with a fermionic AEA like ^{87}Sr is of particular interest.

small ensembles and short evolution times. However, for denser ensembles and long evolution times (for example for adiabatic ramps of ~ 100 ms in duration) we measure significant loss as a result of inelastic two body collisions.

[9] As we saw in the boson sampling demonstrations in this work, the axial motion serves as a label that is otherwise decoupled from the effectively non-interacting dynamics of ground state ^{88}Sr atoms.

[10] For reasonable parameters in our experiment, controlling the chemical potential introduced by the dressing light at the level of the tunneling energy would require homogenization of the dressing beam at the $\sim 10^{-4}$ level in intensity.

The interactions between AEAs in different nuclear spin states are symmetric in the electronic ground or clock states (due to these states having angular momentum $J = 0$),[11] and therefore provide a direct implementation of highly symmetric $SU(N)$ Hubbard models, with N as large as 10 in ^{87}Sr [126, 56, 212]. Such systems could host a variety of exotic behaviors, including chiral spin liquids that host non-Abelian anyons [140]. The $SU(2)$ case is particularly well-studied and significant, since it corresponds to the standard spin-1/2 Fermi-Hubbard model that is believed to capture the physics of high T_c superconductivity in the cuprates [184]. Therefore, applying the controls developed in this work (namely the ability to perform gates on the spin degree of freedom, and gate-like operations between sites in the lattice) to fermionic atoms could be extremely fruitful [124].[12] For state preparation, variational or circuit-inspired techniques [57, 330, 91, 251] could potentially be used to prepare ground states in these models more effectively than conventional techniques involving evaporation followed by adiabatic ramps [30, 102]. For measurement, gates between spins could allow one to access complicated spin correlation functions, likely with relatively high fidelity. Modifications of the lattice potential and tunneling dynamics could further provide access to nearly arbitrary observables [332], although likely with lower fidelity that the spin gates.[13] Access to non-standard measurements of this type could be extremely valuable. For example, whereas different pairing mechanisms leading to superconductivity can be hard to distinguish using conventional measurements, using the above tools one could directly measure the appropriate order parameter to distinguish s-wave from d-wave pairing in a Hubbard model that hosts high T_c superconductivity [332].

9.3 Combined directions

In addition to the directions outlined above, there are a number of directions that combine aspects of Hubbard physics with precision measurement and control. We discussed in Sec. 9.2 how the clock qubit can be used to alter the distinguishability of different atoms. This serves as a useful control in precision measurements. For example, in [3], the distinguishability of the atoms is used to tune the relative strength of

[11] Unlike collisions between atoms in different electronic states, these collisions also do not suffer from the lossy two body collisions.

[12] Fermionic particle statistics can only be simulated on standard quantum architectures with significant (albeit polynomial) overhead in either runtime or physical resource cost. As a result, implementing analog evolution of a fermionic system in combination with the ability to perform appropriate gates and circuits could be useful for a broad set of simulation problems in quantum science beyond just Hubbard physics, ranging from high energy physics to chemistry [124].

[13] Due to the challenges associated with implementing these operations with low noise, as described above.

p- and s-wave interactions and cancel the interaction shift, enabling clock measurements with large ensembles of atoms and long interrogations times, and thus record-low instability. Distinguishabilty could also be an interesting signal in and of itself. Although landmark experiments have measured the gravitational redshift using optical clocks [67, 35], in these measurements the proper time for each ensemble is well-defined classically (in the sense that the atomic ensembles being compared remain separate). The ability to perform clock measurements with the atoms in a spatial superposition allows for fundamental tests of relativity and quantum mechanics,[14] where one expects to observe interference effects that can only be explained by properly incorporating general relativistic notions of proper time into quantum theory [380, 381].[15] Such measurements also relate to universal sources of decoherence due to time dilation [260, 261], which have not yet been observed.

Beyond measuring fundamental quantities and behaviors, frequency metrology techniques can be applied to probe and control large systems with complicated emergent properties. Already, high resolution spectroscopy has played a key role in studying many body quantum systems, where the clock transition can be used as a spectrally narrow probe which allows one to directly study the energy landscape of a given system with very fine resolution [215]. Such techniques have been used to precisely measure the strength of spin exchange [288] and many-body interactions [122] in systems involving lattice-trapped fermionic AEAs, including in implementations of the Fermi-Hubbard model in 3D [53]. Calibrations based on the above techniques are already relevant in our simulations of the Bose-Hubbard model [367].

One can also use the clock transition to explicitly engineer different kinds of dynamics, for example by driving transitions that couple the electronic (or spin) state of the atoms to their motion. Spin orbit coupling of this kind has been used to engineer artificial gauge fields [199, 178, 45] in the synthetic lattice where one spatial dimension corresponds to an internal degree of freedom of the atoms.[16] Combining the above techniques with state- and site-resolved readout, as well as the preparation of low entropy initial states, could provide a route towards studying exotic many body behaviors in lattices subjected to strong

[14] These tests are distinct from what people normally refer to as quantum gravity, and are not related to detecting gravitons.

[15] Schematically, the idea here is that different components of the atomic wavefunction experience different proper time, leading to entanglement between the trajectories and the internal states of the atoms.

[16] Similar demonstrations have been performed in pioneering experiments using Raman transitions [123]. Implementations using single-photon transitions between long-lived clock states have the benefit of eliminating scattering off of the short-lived intermediate states, potentially improving coherence. However, one must be careful of the additional challenge associated with decoherence due to the inelastic collisions between clock state atoms mentioned above.

fields, like chiral Mott insulators [89, 258] and topological superfluidity [267].

Bibliography

[1] S. Aaronson and A. Ambainis. Quantum search of spatial regions. In 44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings., pages 200–209, Oct. 2003.

[2] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. In Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, STOC '11, pages 333–342, New York, NY, USA, June 2011. Association for Computing Machinery.

[3] A. Aeppli, A. Chu, T. Bothwell, C. J. Kennedy, D. Kedar, P. He, A. M. Rey, and J. Ye. Hamiltonian engineering of spin-orbit–coupled fermions in a Wannier-Stark optical lattice clock. Science Advances, 8(41):eadc9242, Oct. 2022.

[4] T. Akatsuka, M. Takamoto, and H. Katori. Three-dimensional optical lattice clock with bosonic ^{88}Sr atoms. Physical Review A, 81(2), Feb. 2010.

[5] T. Albash and D. A. Lidar. Adiabatic quantum computation. Rev. Mod. Phys., 90(1):015002, Jan. 2018.

[6] A. Ambainis. Quantum walk algorithm for element distinctness. In 45th Annual IEEE Symposium on Foundations of Computer Science, pages 22–31, Oct. 2004.

[7] A. Ambainis, J. Kempe, and A. Rivosh. Coins make quantum walks faster: Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms. pages 1099–1108, 2005.

[8] W. R. Anderson, J. R. Veale, and T. F. Gallagher. Resonant Dipole-Dipole Energy Transfer in a Nearly Frozen Rydberg Gas. Phys. Rev. Lett., 80(2):249–252, Jan. 1998.

[9] Andor. iXon Ultra and Life 897, Sept. 2023.

[10] A. André, A. S. Sørensen, and M. D. Lukin. Stability of Atomic Clocks Based on Entangled Atoms. Phys. Rev. Lett., 92(23):230801, June 2004.

[11] N. W. Ashcroft and N. D. Mermin. Solid State Physics. Holt, Rinehart and Winston, New York, 1976.

[12] X. Baillard, A. Gauguet, S. Bize, P. Lemonde, Ph. Laurent, A. Clairon, and P. Rosenbusch. Interference-filter-stabilized external-cavity diode lasers. Optics Communications, 266(2):609–613, Oct. 2006.

[13] W. S. Bakr. Microscopic Studies of Quantum Phase Transitions in Optical Lattices. PhD thesis, 2011.

[14] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner. A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice. Nature, 462(7269):74–77, Nov. 2009.

[15] W. S. Bakr, A. Peng, M. E. Tai, R. Ma, J. Simon, J. I. Gillen, S. Fölling, L. Pollet, and M. Greiner. Probing the Superfluid–to–Mott Insulator Transition at the Single-Atom Level. Science, 329(5991):547–550, July 2010.

[16] S. Bali, K. M. O'Hara, M. E. Gehm, S. R. Granade, and J. E. Thomas. Quantum-diffractive background gas collisions in atom-trap heating and loss. Phys. Rev. A, 60(1):R29–R32, July 1999.

[17] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits. Phys. Rev. Lett., 117(6):060504, Aug. 2016.

[18] K. Barnes, P. Battaglino, B. J. Bloom, K. Cassella, R. Coxe, N. Crisosto, J. P. King, S. S. Kondov, K. Kotru, S. C. Larsen, J. Lauigan, B. J. Lester, M. McDonald, E. Megidish, S. Narayanaswami, C. Nishiguchi, R. Notermans, L. S. Peng, A. Ryou, T.-Y. Wu, and M. Yarwood. Assembly and coherent control of a register of nuclear spin qubits. Nat Commun, 13(1):2779, May 2022.

[19] D. Barredo, S. de Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys. An atom-by-atom assembler of defect-free arbitrary 2d atomic arrays. Science, 354(6315):1021–1023, Nov. 2016.

[20] D. Barredo, H. Labuhn, S. Ravets, T. Lahaye, A. Browaeys, and C. S. Adams. Coherent excitation transfer in a spin chain of three Rydberg atoms. Phys. Rev. Lett., 114(11):113002, Mar. 2015.

[21] J. P. Bartolotta, M. A. Norcia, J. R. K. Cline, J. K. Thompson, and M. J. Holland. Laser cooling by sawtooth-wave adiabatic passage. Phys. Rev. A, 98(2):023404, Aug. 2018.

[22] J. Bausch, T. S. Cubitt, A. Lucia, and D. Perez-Garcia. Undecidability of the Spectral Gap in One Dimension. Phys. Rev. X, 10(3):031038, Aug. 2020.

[23] R. Beigang, K. Lücke, D. Schmidt, A. Timmermann, and P. J. West. One-Photon Laser Spectroscopy of Rydberg Series from Metastable Levels in Calcium and Strontium. Phys. Scr., 26(3):183, Sept. 1982.

[24] J. S. Bell. On the Einstein Podolsky Rosen paradox. Physics Physique Fizika, 1(3):195–200, Nov. 1964.

[25] M. Bentivegna, N. Spagnolo, C. Vitelli, F. Flamini, N. Viggianiello, L. Latmiral, P. Mataloni, D. J. Brod, E. F. Galvão, A. Crespi, R. Ramponi, R. Osellame, and F. Sciarrino. Experimental scattershot boson sampling. Science Advances, 1(3):e1400255, 2015.

[26] H. A. Bethe and E. E. Salpeter. Quantum Mechanics of One- and Two-Electron Atoms. Springer Berlin Heidelberg, Berlin, Heidelberg, 1957.

[27] J. D. Biamonte, V. Bergholm, J. D. Whitfield, J. Fitzsimons, and A. Aspuru-Guzik. Adiabatic quantum simulators. AIP Advances, 1(2):022126, May 2011.

[28] S. Bize, Y. Sortais, P. Lemonde, S. Zhang, P. Laurent, G. Santarelli, C. Salomon, and A. Clairon. Interrogation oscillator noise rejection in the comparison of atomic fountains. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, 47(5):1253–1255, Sept. 2000.

[29] E. D. Black. An introduction to Pound–Drever–Hall laser frequency stabilization. American Journal of Physics, 69(1):79–87, Jan. 2001.

[30] I. Bloch, J. Dalibard, and W. Zwerger. Many-body physics with ultracold gases. Rev. Mod. Phys., 80(3):885–964, July 2008.

[31] B. J. Bloom, T. L. Nicholson, J. R. Williams, S. L. Campbell, M. Bishof, X. Zhang, W. Zhang, S. L. Bromley, and J. Ye. An optical lattice clock with accuracy and stability at the 10^{-18} level. Nature, 506(7486):71–75, Feb. 2014.

[32] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin. A quantum processor based on coherent transport of entangled atom arrays. Nature, 604(7906):451–456, Apr. 2022.

[33] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven. Characterizing quantum supremacy in near-term devices. Nature Phys, 14(6):595–600, June 2018.

[34] T. Bothwell, D. Kedar, E. Oelker, J. M. Robinson, S. L. Bromley, W. L. Tew, J. Ye, and C. J. Kennedy. JILA SrI optical lattice clock with uncertainty of 2.0×10^{-18}. Metrologia, 56(6):065004, Oct. 2019.

[35] T. Bothwell, C. J. Kennedy, A. Aeppli, D. Kedar, J. M. Robinson, E. Oelker, A. Staron, and J. Ye. Resolving the gravitational redshift across a millimetre-scale atomic sample. Nature, 602(7897):420–424, Feb. 2022.

[36] I. Bouchoule and J. Dubail. Generalized hydrodynamics in the one-dimensional Bose gas: Theory and experiments. J. Stat. Mech., 2022(1):014003, Jan. 2022.

[37] I. Bouchoule and K. Mølmer. Spin squeezing of atoms by the dipole interaction in virtually excited Rydberg states. Physical Review A, 65(4), Apr. 2002.

[38] T. Boulier, E. Magnan, C. Bracamontes, J. Maslek, E. A. Goldschmidt, J. T. Young, A. V. Gorshkov, S. L. Rolston, and J. V. Porto. Spontaneous avalanche dephasing in large Rydberg ensembles. Phys. Rev. A, 96(5):053409, Nov. 2017.

[39] M. M. Boyd. High Precision Spectroscopy of Strontium in an Optical Lattice: Towards a New Standard for Frequency and Time. PhD thesis, 2007.

[40] S. L. Braunstein and C. M. Caves. Statistical distance and the geometry of quantum states. Phys. Rev. Lett., 72(22):3439–3443, May 1994.

[41] S. L. Braunstein, C. M. Caves, and G. J. Milburn. Generalized uncertainty relations: Theory, examples, and Lorentz invariance. Annals of Physics, 247(1):135–173, Apr. 1996.

[42] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibrandt. $^{27}Al^{+}$ quantum-logic clock with a systematic uncertainty below 10^{-18}. Phys. Rev. Lett., 123(3):033201, July 2019.

[43] H. J. Briegel, D. E. Browne, W. Dür, R. Raussendorf, and M. Van den Nest. Measurement-based quantum computation. Nature Physics, 5(1):19–26, Jan. 2009.

[44] D. J. Brod and M. Oszmaniec. Classical simulation of linear optics subject to nonuniform losses. Quantum, 4:267, May 2020.

[45] S. L. Bromley, S. Kolkowitz, T. Bothwell, D. Kedar, A. Safavi-Naini, M. L. Wall, C. Salomon, A. M. Rey, and J. Ye. Dynamics of interacting fermions under spin–orbit coupling in an optical lattice clock. Nature Phys, 14(4):399–404, Apr. 2018.

[46] M. A. Broome, A. Fedrizzi, S. Rahimi-Keshari, J. Dove, S. Aaronson, T. C. Ralph, and A. G. White. Photonic boson sampling in a tunable circuit. Science, 339(6121):794–798, Feb. 2013.

[47] A. Browaeys and T. Lahaye. Many-body physics with individually controlled Rydberg atoms. Nature Physics, 16(2):132–142, Feb. 2020.

[48] M. O. Brown, T. Thiele, C. Kiehl, T.-W. Hsu, and C. A. Regal. Gray-molasses optical-tweezer loading: Controlling collisions for scaling atom-array assembly. Phys. Rev. X, 9(1):011057, Mar. 2019.

[49] R. C. Brown, N. B. Phillips, K. Beloy, W. F. McGrew, M. Schioppo, R. J. Fasano, G. Milani, X. Zhang, N. Hinkley, H. Leopardi, T. H. Yoon, D. Nicolodi, T. M. Fortier, and A. D. Ludlow. Hyperpolarizability and Operational Magic Wavelength in an Optical Lattice Clock. Phys. Rev. Lett., 119(25):253001, Dec. 2017.

[50] E. A. Burt, R. W. Ghrist, C. J. Myatt, M. J. Holland, E. A. Cornell, and C. E. Wieman. Coherence, Correlations, and Collisions: What One Learns about Bose-Einstein Condensates from Their Decay. Phys. Rev. Lett., 79(3):337–340, July 1997.

[51] T. Busch, B.-G. Englert, K. Rzażewski, and M. Wilkens. Two Cold Atoms in a Harmonic Trap. Foundations of Physics, 28(4):549–559, Apr. 1998.

[52] S. L. Campbell. A Fermi-degenerate Three-Dimensional Optical Lattice Clock. PhD thesis, Oct. 2017.

[53] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. Darkwah Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye. A Fermi-degenerate three-dimensional optical lattice clock. Science, 358(6359):90–94, Oct. 2017.

[54] J. Carolan, C. Harrold, C. Sparrow, E. Martín-López, N. J. Russell, J. W. Silverstone, P. J. Shadbolt, N. Matsuda, M. Oguma, M. Itoh, G. D. Marshall, M. G. Thompson, J. C. F. Matthews, T. Hashimoto, J. L. O'Brien, and A. Laing. Universal linear optics. Science, 349(6249):711–716, Aug. 2015.

[55] J. Carolan, J. D. A. Meinecke, P. Shadbolt, N. J. Russell, N. Ismail, K. Wörhoff, T. Rudolph, M. G. Thompson, J. L. O'Brien, J. C. F. Matthews, and A. Laing. On the experimental verification of quantum complexity in linear optics. Nature Photon, 8(8):621–626, Aug. 2014.

[56] M. A. Cazalilla and A. M. Rey. Ultracold Fermi gases with emergent SU(N) symmetry. Reports on Progress in Physics, 77(12):124401, Dec. 2014.

[57] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles. Variational quantum algorithms. Nat Rev Phys, 3(9):625–644, Sept. 2021.

[58] S. Chakraborty, L. Novo, A. Ambainis, and Y. Omar. Spatial Search by Quantum Walk is Optimal for Almost all Graphs. Phys. Rev. Lett., 116(10):100501, Mar. 2016.

[59] A. M. Childs. Universal Computation by Quantum Walk. Phys. Rev. Lett., 102(18):180501, May 2009.

[60] A. M. Childs, R. Cleve, E. Deotto, E. Farhi, S. Gutmann, and D. A. Spielman. Exponential algorithmic speedup by quantum walk. Proceedings of the thirty-fifth ACM symposium on Theory of computing - STOC '03, page 59, 2003.

[61] A. M. Childs and Y. Ge. Spatial search by continuous-time quantum walks on crystal lattices. Phys. Rev. A, 89(5):052337, May 2014.

[62] A. M. Childs and J. Goldstone. Spatial search and the Dirac equation. Phys. Rev. A, 70(4):042312, Oct. 2004.

[63] A. M. Childs and J. Goldstone. Spatial search by quantum walk. Phys. Rev. A, 70(2):022314, Aug. 2004.

[64] A. M. Childs, D. Gosset, and Z. Webb. Universal Computation by Multiparticle Quantum Walk. Science, 339(6121):791–794, Feb. 2013.

[65] C. Chin, R. Grimm, P. Julienne, and E. Tiesinga. Feshbach resonances in ultracold gases. Rev. Mod. Phys., 82(2):1225–1286, Apr. 2010.

[66] L. Chomaz, I. Ferrier-Barbut, F. Ferlaino, B. Laburthe-Tolra, B. L. Lev, and T. Pfau. Dipolar physics: A review of experiments with magnetic quantum gases. Rep. Prog. Phys., 86(2):026401, Dec. 2022.

[67] C. W. Chou, D. B. Hume, T. Rosenband, and D. J. Wineland. Optical Clocks and Relativity. Science, 329(5999):1630–1633, Sept. 2010.

[68] C. W. Chou, D. B. Hume, M. J. Thorpe, D. J. Wineland, and T. Rosenband. Quantum coherence between two atoms beyond $Q = 10^{15}$. Phys. Rev. Lett., 106(16):160801, Apr. 2011.

[69] M. Chwalla, K. Kim, T. Monz, P. Schindler, M. Riebe, C. Roos, and R. Blatt. Precision spectroscopy with two correlated atoms. Appl. Phys. B, 89(4):483–488, Dec. 2007.

[70] E. R. Clements, M. E. Kim, K. Cui, A. M. Hankin, S. M. Brewer, J. Valencia, J.-S. Chen, C.-W. Chou, D. R. Leibrandt, and D. B. Hume. Lifetime-limited interrogation of two independent $^{27}Al^{+}$ clocks using correlation spectroscopy. Phys. Rev. Lett., 125(24):243602, Dec. 2020.

[71] W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walmsley. Optimal design for universal multiport interferometers. Optica, OPTICA, 3(12):1460–1465, Dec. 2016.

[72] P. Clifford and R. Clifford. The classical complexity of boson sampling. In Proceedings of the 2018 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), Proceedings, pages 146–155. Society for Industrial and Applied Mathematics, Jan. 2018.

[73] A. Cooper, J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres. Alkaline-earth atoms in optical tweezers. Phys. Rev. X, 8(4):041055, Dec. 2018.

[74] J. P. Covey, I. S. Madjarov, A. Cooper, and M. Endres. 2000-times repeated imaging of strontium atoms in clock-magic tweezer arrays. Phys. Rev. Lett., 122(17):173201, May 2019.

[75] D. R. Cox and E. J. Snell. A General Definition of Residuals. Journal of the Royal Statistical Society. Series B (Methodological), 30(2):248–275, 1968.

[76] H. Cramér. Mathematical Methods of Statistics. Mathematical Methods of Statistics. Princeton University Press, Princeton, NJ, US, 1946.

[77] T. Cubitt, D. Perez-Garcia, and M. M. Wolf. Undecidability of the Spectral Gap. Nature, 528(7581):207–211, Dec. 2015.

[78] R. Curticapean. A full complexity dichotomy for immanant families. In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2021, pages 1770–1783, New York, NY, USA, June 2021. Association for Computing Machinery.

[79] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller. Practical quantum advantage in quantum simulation. Nature, 607(7920):667–676, July 2022.

[80] N. Darkwah Oppong, G. Pasqualetti, O. Bettermann, P. Zechmann, M. Knap, I. Bloch, and S. Fölling. Probing Transport and Slow Relaxation in the Mass-Imbalanced Fermi-Hubbard Model. Phys. Rev. X, 12(3):031026, Aug. 2022.

[81] N. Darkwah Oppong, L. Riegger, O. Bettermann, M. Höfer, J. Levinsen, M. M. Parish, I. Bloch, and S. Fölling. Observation of Coherent Multiorbital Polarons in a Two-Dimensional Fermi Gas. Phys. Rev. Lett., 122(19):193604, May 2019.

[82] S. de Léséleuc, D. Barredo, V. Lienhard, A. Browaeys, and T. Lahaye. Optical Control of the Resonant Dipole-Dipole Interaction between Rydberg Atoms. Phys. Rev. Lett., 119(5):053202, Aug. 2017.

[83] C. L. Degen, F. Reinhard, and P. Cappellaro. Quantum sensing. Rev. Mod. Phys., 89(3):035002, July 2017.

[84] R. Demkowicz-Dobrzański, J. Czajkowski, and P. Sekatski. Adaptive Quantum Metrology under General Markovian Noise. Phys. Rev. X, 7(4):041009, Oct. 2017.

[85] R. Demkowicz-Dobrzański, J. Kołodyński, and M. Guţă. The elusive Heisenberg limit in quantum-enhanced metrology. Nat Commun, 3(1):1063, Jan. 2012.

[86] R. Demkowicz-Dobrzański and L. Maccone. Using Entanglement Against Noise in Quantum Metrology. Phys. Rev. Lett., 113(25):250801, Dec. 2014.

[87] Y.-H. Deng, Y.-C. Gu, H.-L. Liu, S.-Q. Gong, H. Su, Z.-J. Zhang, H.-Y. Tang, M.-H. Jia, J.-M. Xu, M.-C. Chen, J. Qin, L.-C. Peng, J. Yan, Y. Hu, J. Huang, H. Li, Y. Li, Y. Chen, X. Jiang, L. Gan, G. Yang, L. You, L. Li, H.-S. Zhong, H. Wang, N.-L. Liu, J. J. Renema, C.-Y. Lu, and J.-W. Pan. Gaussian Boson Sampling with Pseudo-Photon-Number-Resolving Detectors and Quantum Computational Advantage. Phys. Rev. Lett., 131(15):150601, Oct. 2023.

[88] A. Deshpande, A. Mehta, T. Vincent, N. Quesada, M. Hinsche, M. Ioannou, L. Madsen, J. Lavoie, H. Qi, J. Eisert, D. Hangleiter, B. Fefferman, and I. Dhand. Quantum computational advantage via high-dimensional Gaussian boson sampling. Science Advances, 8(1):eabi7894, Jan. 2022.

[89] A. Dhar, M. Maji, T. Mishra, R. V. Pai, S. Mukerjee, and A. Paramekanti. Bose-Hubbard model in a strong effective magnetic field: Emergence of a chiral Mott insulator ground state. Phys. Rev. A, 85(4):041602, Apr. 2012.

[90] G. J. Dick. Local Oscillator Induced Instabilities in Trapped Ion Frequency Standards. In Proceedings of the 19th Annual Precise Time and Time Interval Systems and Applications Meeting, pages 133–147, Dec. 1989.

[91] P. Doria, T. Calarco, and S. Montangero. Optimal Control Technique for Many-Body Quantum Dynamics. Phys. Rev. Lett., 106(19):190501, May 2011.

[92] S. Dörscher, A. Al-Masoudi, M. Bober, R. Schwarz, R. Hobson, U. Sterr, and C. Lisdat. Dynamical decoupling of laser phase noise in compound atomic clocks. Commun Phys, 3(1):1–9, Oct. 2020.

[93] S. Dörscher, R. Schwarz, A. Al-Masoudi, S. Falke, U. Sterr, and C. Lisdat. Lattice-induced photon scattering in an optical lattice clock. Phys. Rev. A, 97(6):063419, June 2018.

[94] G. Dufour, T. Brünner, A. Rodríguez, and A. Buchleitner. Many-body interference in bosonic dynamics. New J. Phys., 22(10):103006, Oct. 2020.

[95] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, S. Choi, S. Sachdev, M. Greiner, V. Vuletić, and M. D. Lukin. Quantum phases of matter on a 256-atom programmable quantum simulator. Nature, 595(7866):227–232, July 2021.

[96] W. J. Eckner, N. Darkwah Oppong, A. Cao, A. W. Young, W. R. Milner, J. M. Robinson, J. Ye, and A. M. Kaufman. Realizing spin squeezing with Rydberg interactions in an optical clock. Nature, pages 1–6, Aug. 2023.

[97] W. J. Eckner, A. W. Young, N. Schine, and A. M. Kaufman. High-power, fiber–laser-based source for magic-wavelength trapping in neutral-atom optical clocks. Rev. Sci. Instrum., 92(9):093001, Sept. 2021.

[98] B. Efron and R. J. Tibshirani. An Introduction to the Bootstrap. Chapman and Hall/CRC, New York, May 1994.

[99] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin. Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. Science, 354(6315):1024–1027, Nov. 2016.

[100] K. Enomoto, K. Kasa, M. Kitagawa, and Y. Takahashi. Optical Feshbach Resonance Using the Intercombination Transition. Phys. Rev. Lett., 101(20):203201, Nov. 2008.

[101] J. Eschner, G. Morigi, F. Schmidt-Kaler, and R. Blatt. Laser cooling of trapped ions. J. Opt. Soc. Am. B, JOSAB, 20(5):1003–1015, May 2003.

[102] T. Esslinger. Fermi-Hubbard Physics with Atoms in an Optical Lattice. Annu. Rev. Condens. Matter Phys., 1(1):129–152, Aug. 2010.

[103] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara, H. Levine, G. Semeghini, M. Greiner, V. Vuletić, and M. D. Lukin. High-fidelity parallel entangling gates on a neutral-atom quantum computer. Nature, 622(7982):268–272, Oct. 2023.

[104] E. Farhi, J. Goldstone, and S. Gutmann. A Quantum Algorithm for the Hamiltonian NAND Tree. Theory of Computing, 4(1):169–190, Dec. 2008.

[105] E. Farhi and S. Gutmann. Analog analogue of a digital quantum computation. Phys. Rev. A, 57(4):2403–2406, Apr. 1998.

[106] P. O. Fedichev, M. W. Reynolds, and G. V. Shlyapnikov. Three-Body Recombination of Ultracold Atoms to a Weakly Bound s Level. Phys. Rev. Lett., 77(14):2921–2924, Sept. 1996.

[107] L. Festa, N. Lorenz, L.-M. Steinert, Z. Chen, P. Osterholz, R. Eberhard, and C. Gross. Blackbody-radiation-induced facilitated excitation of Rydberg atoms in optical tweezers. Phys. Rev. A, 105(1):013109, Jan. 2022.

[108] M. Foss-Feig, M. Hermele, V. Gurarie, and A. M. Rey. Heavy fermions in an optical lattice. Phys. Rev. A, 82(5):053624, Nov. 2010.

[109] M. Foss-Feig, M. Hermele, and A. M. Rey. Probing the Kondo lattice model with alkaline-earth-metal atoms. Phys. Rev. A, 81(5):051603, May 2010.

[110] G. T. Foster, J. B. Fixler, J. M. McGuirk, and M. A. Kasevich. Method of phase extraction between coupled atom interferometers using ellipse-specific fitting. Opt. Lett., 27(11):951, June 2002.

[111] Y. H. Fung and M. F. Andersen. Efficient collisional blockade loading of a single atom into a tight microtrap. New J. Phys., 17(7):073011, July 2015.

[112] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland. High-Fidelity Universal Gate Set for $^{9}Be^{+}$ Ion Qubits. Phys. Rev. Lett., 117(6):060505, Aug. 2016.

[113] T. F. Gallagher. Rydberg Atoms. Cambridge University Press, 1 edition, Sept. 1994.

[114] R. García-Patrón, J. J. Renema, and V. Shchesnovich. Simulating boson sampling in lossy architectures. Quantum, 3:169, Aug. 2019.

[115] M. E. Gehm, K. M. O'Hara, T. A. Savard, and J. E. Thomas. Dynamics of noise-induced heating in atom traps. Phys. Rev. A, 58(5):3914–3921, Nov. 1998.

[116] K. Gibble. Scattering of cold-atom coherences by hot atoms: Frequency shifts from background-gas collisions. Phys. Rev. Lett., 110(18):180802, May 2013.

[117] L. I. R. Gil, R. Mukherjee, E. M. Bridge, M. P. A. Jones, and T. Pohl. Spin squeezing in a Rydberg lattice clock. Phys. Rev. Lett., 112(10):103601, Mar. 2014.

[118] S. Gil-Gallegos, R. Klages, J. Solanpää, and E. Räsänen. Energy-dependent diffusion in a soft periodic Lorentz gas. Eur. Phys. J. Spec. Top., 228(1):143–160, May 2019.

[119] V. Giovannetti, S. Lloyd, and L. Maccone. Quantum Metrology. Phys. Rev. Lett., 96(1):010401, Jan. 2006.

[120] V. Giovannetti, S. Lloyd, and L. Maccone. Advances in quantum metrology. Nature Photon, 5(4):222–229, Apr. 2011.

[121] N. Gisin and R. Thew. Quantum communication. Nature Photon, 1(3):165–171, Mar. 2007.

[122] A. Goban, R. B. Hutson, G. E. Marti, S. L. Campbell, M. A. Perlin, P. S. Julienne, J. P. D'Incao, A. M. Rey, and J. Ye. Emergence of multi-body interactions in a fermionic lattice clock. Nature, 563(7731):369–373, Nov. 2018.

[123] N. Goldman, G. Juzeliūnas, P. Öhberg, and I. B. Spielman. Light-induced gauge fields for ultracold atoms. Rep. Prog. Phys., 77(12):126401, Nov. 2014.

[124] D. González-Cuadra, D. Bluvstein, M. Kalinowski, R. Kaubruegger, N. Maskara, P. Naldesi, T. V. Zache, A. M. Kaufman, M. D. Lukin, H. Pichler, B. Vermersch, J. Ye, and P. Zoller. Fermionic quantum processing with programmable neutral atom arrays, Mar. 2023.

[125] J. Goodman. Introduction to Fourier Optics. Roberts and Company Publishers, Englewood, Colo, 3rd edition edition, Dec. 2004.

[126] A. V. Gorshkov, M. Hermele, V. Gurarie, C. Xu, P. S. Julienne, J. Ye, P. Zoller, E. Demler, M. D. Lukin, and A. M. Rey. Two-orbital SU(N) magnetism with ultracold alkaline-earth atoms. Nature Physics, 6(4):289–295, Apr. 2010.

[127] T. M. Graham, M. Kwon, B. Grinkemeyer, Z. Marra, X. Jiang, M. T. Lichtman, Y. Sun, M. Ebert, and M. Saffman. Rydberg-mediated entanglement in a two-dimensional neutral atom qubit array. Phys. Rev. Lett., 123(23):230501, Dec. 2019.

[128] T. M. Graham, Y. Song, J. Scott, C. Poole, L. Phuttitarn, K. Jooya, P. Eichler, X. Jiang, A. Marra, B. Grinkemeyer, M. Kwon, M. Ebert, J. Cherek, M. T. Lichtman, M. Gillette, J. Gilbert, D. Bowman, T. Ballance, C. Campbell, E. D. Dahl, O. Crawford, N. S. Blunt, B. Rogers, T. Noel, and M. Saffman. Multi-qubit entanglement and algorithms on a neutral-atom quantum computer. Nature, 604(7906):457–462, Apr. 2022.

[129] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. Nature, 415(6867):39–44, Jan. 2002.

[130] C. Gross and I. Bloch. Quantum simulations with ultracold atoms in optical lattices. Science, 357(6355):995–1001, Sept. 2017.

[131] J. Grotti, S. Koller, S. Vogt, S. Häfner, U. Sterr, C. Lisdat, H. Denker, C. Voigt, L. Timmen, A. Rolland, F. N. Baynes, H. S. Margolis, M. Zampaolo, P. Thoumany, M. Pizzocaro, B. Rauf, F. Bregolin, A. Tampellini, P. Barbieri, M. Zucco, G. A. Costanzo, C. Clivati, F. Levi, and D. Calonico. Geodesy and metrology with a transportable optical clock. Nature Phys, 14(5):437–441, May 2018.

[132] 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, pages 212–219, New York, NY, USA, July 1996. Association for Computing Machinery.

[133] L. K. Grover and A. M. Sengupta. From coupled pendulums to quantum search. In Mathematics of Quantum Computation. Chapman and Hall/CRC, 2002.

[134] T. Grünzweig, A. Hilliard, M. McGovern, and M. F. Andersen. Near-deterministic preparation of a single atom in an optical microtrap. Nature Phys, 6(12):951–954, Dec. 2010.

[135] E. Guardado-Sanchez, B. M. Spar, P. Schauss, R. Belyansky, J. T. Young, P. Bienias, A. V. Gorshkov, T. Iadecola, and W. S. Bakr. Quench Dynamics of a Fermi Gas with Strong Nonlocal Interactions. Phys. Rev. X, 11(2):021036, May 2021.

[136] J. L. Hall, M. S. Taubman, and J. Ye. Laser stabilization. In Handbook of Optics, volume II of Handbook of Optics. McGraw-Hill Professional, New York, NY, 3 edition, 2009.

[137] C. S. Hamilton, R. Kruse, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex. Gaussian boson sampling. Phys. Rev. Lett., 119(17):170501, Oct. 2017.

[138] A. Heinz, A. J. Park, N. Šantić, J. Trautmann, S. G. Porsev, M. S. Safronova, I. Bloch, and S. Blatt. State-Dependent Optical Lattices for the Strontium Optical Qubit. Phys. Rev. Lett., 124(20):203201, May 2020.

[139] I. Hen and M. Rigol. Superfluid to Mott insulator transition of hardcore bosons in a superlattice. Phys. Rev. B, 80(13):134508, Oct. 2009.

[140] M. Hermele, V. Gurarie, and A. M. Rey. Mott Insulators of Ultracold Fermionic Alkaline Earth Atoms: Underconstrained Magnetism and Chiral Spin Liquid. Physical Review Letters, 103(13), Sept. 2009.

[141] J. A. Hines, S. V. Rajagopal, G. L. Moreau, M. D. Wahrman, N. A. Lewis, O. Marković, and M. Schleier-Smith. Spin Squeezing by Rydberg Dressing in an Array of Atomic Ensembles. Phys. Rev. Lett., 131(6):063401, Aug. 2023.

[142] T. O. Höhn, E. Staub, G. Brochier, N. D. Oppong, and M. Aidelsburger. State-dependent potentials for the $^1S_0$ and $^3P_0$ clock states of neutral ytterbium atoms, May 2023.

[143] C. K. Hong, Z. Y. Ou, and L. Mandel. Measurement of subpicosecond time intervals between two photons by interference. Phys. Rev. Lett., 59(18):2044–2046, Nov. 1987.

[144] J. Hubbard. Electron Correlations in Narrow Energy Bands. Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 276(1365):238–257, 1963.

[145] S. F. Huelga, C. Macchiavello, T. Pellizzari, A. K. Ekert, M. B. Plenio, and J. I. Cirac. Improvement of Frequency Standards with Quantum Entanglement. Phys. Rev. Lett., 79(20):3865–3868, Nov. 1997.

[146] W. Huie, L. Li, N. Chen, X. Hu, Z. Jia, W. K. C. Sun, and J. P. Covey. Repetitive Readout and Real-Time Control of Nuclear Spin Qubits in ^{171}Yb Atoms. PRX Quantum, 4(3):030337, Sept. 2023.

[147] D. B. Hume and D. R. Leibrandt. Probing beyond the laser coherence time in optical clock comparisons. Phys. Rev. A, 93(3):032138, Mar. 2016.

[148] R. B. Hutson, A. Goban, G. E. Marti, and J. Ye. Engineering quantum states of matter for atomic clocks in shallow optical lattices. Phys. Rev. Lett., 123(12):123401, Sept. 2019.

[149] A. Impertro, J. F. Wienand, S. Häfele, H. von Raven, S. Hubele, T. Klostermann, C. R. Cabrera, I. Bloch, and M. Aidelsburger. An unsupervised deep learning algorithm for single-site reconstruction in quantum gas microscopes. Commun Phys, 6(1):1–8, July 2023.

[150] L. Isenhower, E. Urban, X. L. Zhang, A. T. Gill, T. Henage, T. A. Johnson, T. G. Walker, and M. Saffman. Demonstration of a Neutral Atom Controlled-NOT Quantum Gate. Phys. Rev. Lett., 104(1):010503, Jan. 2010.

[151] R. Islam, R. Ma, P. M. Preiss, M. Eric Tai, A. Lukin, M. Rispoli, and M. Greiner. Measuring entanglement entropy in a quantum many-body system. Nature, 528(7580):77–83, Dec. 2015.

[152] S. Jandura and G. Pupillo. Time-Optimal Two- and Three-Qubit Gates for Rydberg Atoms. Quantum, 6:712, May 2022.

[153] A. Janssen, S. Van Haver, P. Dirksen, and J. Braat. Zernike representation and Strehl ratio of optical systems with variable numerical aperture. Journal of Modern Optics, 55(7):1127–1157, Apr. 2008.

[154] R. Jáuregui. Nonperturbative and perturbative treatments of parametric heating in atom traps. Phys. Rev. A, 64(5):053408, Oct. 2001.

[155] A. Jenkins, J. W. Lis, A. Senoo, W. F. McGrew, and A. M. Kaufman. Ytterbium nuclear-spin qubits in an optical tweezer array. Phys. Rev. X, 12(2):021027, May 2022.

[156] M. Jreissaty, J. Carrasquilla, F. A. Wolf, and M. Rigol. Expansion of Bose-Hubbard Mott insulators in optical lattices. Phys. Rev. A, 84(4):043610, Oct. 2011.

[157] Yu. Kagan, B. V. Svistunov, and G. V. Shlyapnikov. Effect of Bose condensation on inelastic processes in gases. Soviet Journal of Experimental and Theoretical Physics Letters, 42:209, Aug. 1985.

[158] M. Karski, L. Forster, J.-M. Choi, A. Steffen, W. Alt, D. Meschede, and A. Widera. Quantum Walk in Position Space with Single Optically Trapped Atoms. Science, 325(5937):174–177, July 2009.

[159] H. Katori, V. D. Ovsiannikov, S. I. Marmo, and V. G. Palchikov. Strategies for reducing the light shift in atomic clocks. Phys. Rev. A, 91(5):052503, May 2015.

[160] R. Kaubruegger, P. Silvi, C. Kokail, R. van Bijnen, A. M. Rey, J. Ye, A. M. Kaufman, and P. Zoller. Variational spin-squeezing algorithms on programmable quantum sensors. Phys. Rev. Lett., 123(26):260505, Aug. 2019.

[161] R. Kaubruegger, D. V. Vasilyev, M. Schulte, K. Hammerer, and P. Zoller. Quantum Variational Optimization of Ramsey Interferometry and Atomic Clocks. Phys. Rev. X, 11(4):041045, Dec. 2021.

[162] A. Kaufman. Laser Cooling Atoms to Indistinguishability: Atomic Hong-Ou-Mandel Interference and Entanglement through Spin Exchange. PhD thesis, 2015.

[163] A. M. Kaufman, B. J. Lester, M. Foss-Feig, M. L. Wall, A. M. Rey, and C. A. Regal. Entangling two transportable neutral atoms via local spin exchange. Nature, 527(7577):208–211, Nov. 2015.

[164] A. M. Kaufman, B. J. Lester, and C. A. Regal. Cooling a single atom in an optical tweezer to its quantum ground state. Phys. Rev. X, 2(4):041014, Nov. 2012.

[165] A. M. Kaufman, B. J. Lester, C. M. Reynolds, M. L. Wall, M. Foss-Feig, K. R. A. Hazzard, A. M. Rey, and C. A. Regal. Two-particle quantum interference in tunnel-coupled optical tweezers. Science, 345(6194):306–309, July 2014.

[166] A. M. Kaufman, M. E. Tai, A. Lukin, M. Rispoli, R. Schittko, P. M. Preiss, and M. Greiner. Quantum thermalization through entanglement in an isolated many-body system. Science, 353(6301):794–800, Aug. 2016.

[167] D. Kedar. A Fully Crystalline Cryogenic Reference Cavity. PhD thesis, 2023.

[168] J. Kempe, A. Kitaev, and O. Regev. The Complexity of the Local Hamiltonian Problem. In K. Lodaya and M. Mahajan, editors, FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Computer Science, pages 372–383, Berlin, Heidelberg, 2005. Springer.

[169] C. J. Kennedy, E. Oelker, J. M. Robinson, T. Bothwell, D. Kedar, W. R. Milner, G. E. Marti, A. Derevianko, and J. Ye. Precision Metrology Meets Cosmology: Improved Constraints on Ultralight Dark Matter from Atom-Cavity Frequency Comparisons. Phys. Rev. Lett., 125(20):201302, Nov. 2020.

[170] E. M. Kessler, P. Kómár, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin. Heisenberg-limited atom clocks based on entangled qubits. Phys. Rev. Lett., 112(19):190403, May 2014.

[171] E. M. Kessler, I. Lovchinsky, A. O. Sushkov, and M. D. Lukin. Quantum Error Correction for Metrology. Phys. Rev. Lett., 112(15):150802, Apr. 2014.

[172] S. Keßler and F. Marquardt. Single-site-resolved measurement of the current statistics in optical lattices. Phys. Rev. A, 89(6):061601, June 2014.

[173] M. E. Kim, W. F. McGrew, N. V. Nardelli, E. R. Clements, Y. S. Hassan, X. Zhang, J. L. Valencia, H. Leopardi, D. B. Hume, T. M. Fortier, A. D. Ludlow, and D. R. Leibrandt. Improved interspecies optical clock comparisons through differential spectroscopy. Nat. Phys., 19(1):25–29, Jan. 2023.

[174] M. Kitagawa and M. Ueda. Squeezed spin states. Phys. Rev. A, 47(6):5138–5143, June 1993.

[175] T. Kitagawa, M. S. Rudner, E. Berg, and E. Demler. Exploring topological phases with quantum walks. Phys. Rev. A, 82(3):033429, Sept. 2010.

[176] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. Nature, 409(6816):46–52, Jan. 2001.

[177] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. Phys. Rev. A, 77(1):012307, Jan. 2008.

[178] S. Kolkowitz, S. L. Bromley, T. Bothwell, M. L. Wall, G. E. Marti, A. P. Koller, X. Zhang, A. M. Rey, and J. Ye. Spin-orbit-coupled fermions in an optical lattice clock. Nature, 542(7639):66–70, Feb. 2017.

[179] A. Kumar, T.-Y. Wu, F. Giraldo, and D. S. Weiss. Sorting ultracold atoms in a three-dimensional optical lattice in a realization of Maxwell's demon. Nature, 561(7721):83–87, Sept. 2018.

[180] A. Laing and J. L. O'Brien. Super-stable tomography of any linear optical device, Aug. 2012.

[181] D. Layden and P. Cappellaro. Spatial noise filtering through error correction for quantum sensing. npj Quantum Inf, 4(1):1–6, July 2018.

[182] D. Layden, S. Zhou, P. Cappellaro, and L. Jiang. Ancilla-free quantum error correction codes for quantum metrology. Phys. Rev. Lett., 122(4):040502, Jan. 2019.

[183] F. Le Kien, P. Schneeweiss, and A. Rauschenbeutel. Dynamical polarizability of atoms in arbitrary light fields: General theory and application to cesium. Eur. Phys. J. D, 67(5):92, May 2013.

[184] P. A. Lee, N. Nagaosa, and X.-G. Wen. Doping a Mott insulator: Physics of high-temperature superconductivity. Rev. Mod. Phys., 78(1):17–85, Jan. 2006.

[185] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. Reviews of Modern Physics, 75(1):281–324, Mar. 2003.

[186] D. R. Leibrandt and J. Heidecker. An open source digital servo for AMO physics experiments. arXiv:1508.06319 [physics], Nov. 2015.

[187] P. Lemonde and P. Wolf. Optical lattice clock with atoms confined in a shallow trap. Physical Review A, 72(3), Sept. 2005.

[188] I. D. Leroux, M. H. Schleier-Smith, and V. Vuletić. Implementation of Cavity Squeezing of a Collective Atomic Spin. Phys. Rev. Lett., 104(7):073602, Feb. 2010.

[189] B. J. Lester, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal. Rapid production of uniformly filled arrays of neutral atoms. Phys. Rev. Lett., 115(7):073003, Aug. 2015.

[190] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin. High-fidelity control and entanglement of Rydberg-atom qubits. Phys. Rev. Lett., 121(12):123603, Sept. 2018.

[191] H. Levine, A. Keesling, G. Semeghini, A. Omran, T. T. Wang, S. Ebadi, H. Bernien, M. Greiner, V. Vuletić, H. Pichler, and M. D. Lukin. Parallel implementation of high-fidelity multiqubit gates with neutral atoms. Phys. Rev. Lett., 123(17):170503, Oct. 2019.

[192] J.-R. Li, K. Matsuda, C. Miller, A. N. Carroll, W. G. Tobias, J. S. Higgins, and J. Ye. Tunable itinerant spin dynamics with polar molecules. Nature, 614(7946):70–74, Feb. 2023.

[193] F. Lin, E. S. Sørensen, and D. M. Ceperley. The superfluid-insulator transition in the disordered two-dimensional Bose-Hubbard model. Phys. Rev. B, 84(9):094507, Sept. 2011.

[194] J. W. Lis, A. Senoo, W. F. McGrew, F. Rönchen, A. Jenkins, and A. M. Kaufman. Mid-circuit operations using the omg-architecture in neutral atom arrays, May 2023.

[195] S. Lischke, A. Peczek, J. S. Morgan, K. Sun, D. Steckler, Y. Yamamoto, F. Korndörfer, C. Mai, S. Marschmeyer, M. Fraschke, A. Krüger, A. Beling, and L. Zimmermann. Ultra-fast germanium photodiode with 3-dB bandwidth of 265 GHz. Nat. Photon., 15(12):925–931, Dec. 2021.

[196] Ch. Lisdat, J. S. R. V. Winfred, T. Middelmann, F. Riehle, and U. Sterr. Collisional Losses, Decoherence, and Frequency Shifts in Optical Lattice Clocks with Bosons. Phys. Rev. Lett., 103(9):090801, Aug. 2009.

[197] J. Liu, H. Yuan, X.-M. Lu, and X. Wang. Quantum Fisher information matrix and multiparameter estimation. J. Phys. A: Math. Theor., 53(2):023001, Jan. 2020.

[198] L. R. Liu, J. D. Hood, Y. Yu, J. T. Zhang, K. Wang, Y.-W. Lin, T. Rosenband, and K.-K. Ni. Molecular assembly of ground-state cooled single atoms. Phys. Rev. X, 9(2):021039, May 2019.

[199] L. F. Livi, G. Cappellini, M. Diem, L. Franchi, C. Clivati, M. Frittelli, F. Levi, D. Calonico, J. Catani, M. Inguscio, and L. Fallani. Synthetic Dimensions and Spin-Orbit Coupling with an Optical Clock Transition. Phys. Rev. Lett., 117(22):220401, Nov. 2016.

[200] G. Lochead, D. Boddy, D. P. Sadler, C. S. Adams, and M. P. A. Jones. Number-resolved imaging of excited-state atoms using a scanning autoionization microscope. Phys. Rev. A, 87(5):053409, May 2013.

[201] R. Lopes, A. Imanaliev, A. Aspect, M. Cheneau, D. Boiron, and C. I. Westbrook. Atomic Hong–Ou–Mandel experiment. Nature, 520(7545):66–68, Apr. 2015.

[202] J. C. Loredo, M. A. Broome, P. Hilaire, O. Gazzano, I. Sagnes, A. Lemaitre, M. P. Almeida, P. Senellart, and A. G. White. Boson sampling with single-photon Fock states from a bright solid-state source. Phys. Rev. Lett., 118(13):130503, Mar. 2017.

[203] A. D. Ludlow, M. M. Boyd, J. Ye, E. Peik, and P. O. Schmidt. Optical atomic clocks. Rev. Mod. Phys., 87(2):637–701, June 2015.

[204] A. P. Lund, A. Laing, S. Rahimi-Keshari, T. Rudolph, J. L. O'Brien, and T. C. Ralph. Boson sampling from a Gaussian state. Phys. Rev. Lett., 113(10):100502, Sept. 2014.

[205] L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall. Delivering the same optical frequency at two places: Accurate cancellation of phase noise introduced by an optical fiber or other time-varying path. Opt. Lett., 19(21):1777, Nov. 1994.

[206] S. Ma, A. P. Burgers, G. Liu, J. Wilson, B. Zhang, and J. D. Thompson. Universal Gate Operations on Nuclear Spin Qubits in an Optical Tweezer Array of ^{171}Yb Atoms. Phys. Rev. X, 12(2):021028, May 2022.

[207] S. Ma, G. Liu, P. Peng, B. Zhang, S. Jandura, J. Claes, A. P. Burgers, G. Pupillo, S. Puri, and J. D. Thompson. High-fidelity gates with mid-circuit erasure conversion in a metastable neutral atom qubit, May 2023.

[208] I. S. Madjarov. Entangling, Controlling, and Detecting Individual Strontium Atoms in Optical Tweezer Arrays. PhD thesis, 2021.

[209] I. S. Madjarov, A. Cooper, A. L. Shaw, J. P. Covey, V. Schkolnik, T. H. Yoon, J. R. Williams, and M. Endres. An atomic-array optical clock with single-atom readout. Phys. Rev. X, 9(4):041052, Dec. 2019.

[210] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres. High-fidelity entanglement and detection of alkaline-earth Rydberg atoms. Nature Physics, pages 1–5, May 2020.

[211] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins, A. E. Lita, T. Gerrits, S. W. Nam, V. D. Vaidya, M. Menotti, I. Dhand, Z. Vernon, N. Quesada, and J. Lavoie. Quantum computational advantage with a programmable photonic processor. Nature, 606(7912):75–81, June 2022.

[212] M. Mamaev, T. Bilitewski, B. Sundar, and A. M. Rey. Resonant Dynamics of Strongly Interacting SU(n) Fermionic Atoms in a Synthetic Flux Ladder. PRX Quantum, 3(3):030328, Aug. 2022.

[213] G. E. Marti, R. B. Hutson, A. Goban, S. L. Campbell, N. Poli, and J. Ye. Imaging optical frequencies with 100μHz precision and 1.1μm resolution. Phys. Rev. Lett., 120(10):103201, Mar. 2018.

[214] M. J. Martin. Quantum Metrology and Many-Body Physics: Pushing the Frontier of the Optical Lattice Clock. PhD thesis, 2013.

[215] M. J. Martin, M. Bishof, M. D. Swallows, X. Zhang, C. Benko, J. von-Stecher, A. V. Gorshkov, A. M. Rey, and J. Ye. A quantum many-body spin system in an optical lattice clock. Science, 341(6146):632–636, Aug. 2013.

[216] Y. N. Martinez de Escobar, P. G. Mickelson, P. Pellegrini, S. B. Nagel, A. Traverso, M. Yan, R. Côté, and T. C. Killian. Two-photon photoassociative spectroscopy of ultracold ^{88}Sr. Phys. Rev. A, 78(6):062708, Dec. 2008.

[217] M. Martinez-Dorantes, W. Alt, J. Gallego, S. Ghosh, L. Ratschbacher, and D. Meschede. State-dependent fluorescence of neutral atoms in optical potentials. Phys. Rev. A, 97(2):023410, Feb. 2018.

[218] D. G. Matei, T. Legero, S. Häfner, C. Grebing, R. Weyrich, W. Zhang, L. Sonderhouse, J. M. Robinson, J. Ye, F. Riehle, and U. Sterr. 1.5 μm lasers with sub 10 mHz linewidth. Physical Review Letters, 118(26), June 2017.

[219] J. C. F. Matthews, K. Poulios, J. D. A. Meinecke, A. Politi, A. Peruzzo, N. Ismail, K. Wörhoff, M. G. Thompson, and J. L. O'Brien. Observing fermionic statistics with photons in arbitrary processes. Sci Rep, 3(1):1539, Dec. 2013.

[220] W. F. McGrew, X. Zhang, R. J. Fasano, S. A. Schäffer, K. Beloy, D. Nicolodi, R. C. Brown, N. Hinkley, G. Milani, M. Schioppo, T. H. Yoon, and A. D. Ludlow. Atomic clock performance enabling geodesy below the centimetre level. Nature, 564(7734):87–90, Dec. 2018.

[221] S. T. Merkel, J. M. Gambetta, J. A. Smolin, S. Poletto, A. D. Córcoles, B. R. Johnson, C. A. Ryan, and M. Steffen. Self-consistent quantum process tomography. Phys. Rev. A, 87(6):062119, June 2013.

[222] T. Middelmann, S. Falke, C. Lisdat, and U. Sterr. High accuracy correction of blackbody radiation shift in an optical lattice clock. Phys. Rev. Lett., 109(26):263004, Dec. 2012.

[223] A. Mitra, M. J. Martin, G. W. Biedermann, A. M. Marino, P. M. Poggi, and I. H. Deutsch. Robust Mølmer-Sørensen gate for neutral atoms using rapid adiabatic Rydberg dressing. Phys. Rev. A, 101(3):030301, Mar. 2020.

[224] J. Mitroy and J. Y. Zhang. Dispersion and polarization interactions of the strontium atom. Molecular Physics, 108(15):1999–2006, Aug. 2010.

[225] H. Miyake, N. C. Pisenti, P. K. Elgee, A. Sitaram, and G. K. Campbell. Isotope-shift spectroscopy of the S 0 1 $\rightarrow$ P 1 3 and S 0 1 $\rightarrow$ P 0 3 transitions in strontium. Phys. Rev. Research, 1(3):033113, Nov. 2019.

[226] K. Mølmer, Y. Castin, and J. Dalibard. Monte Carlo wave-function method in quantum optics. J. Opt. Soc. Am. B, JOSAB, 10(3):524–538, Mar. 1993.

[227] M. Mullan. Quantum Algorithms for Atomic Clocks. PhD thesis, 2013.

[228] M. Mullan and E. Knill. Optimizing passive quantum clocks. Phys. Rev. A, 90(4):042310, Oct. 2014.

[229] J. A. Muniz, M. A. Norcia, J. R. K. Cline, and J. K. Thompson. A Robust Narrow-Line Magneto-Optical Trap using Adiabatic Transfer. arXiv:1806.00838 [physics, physics:quant-ph], June 2018.

[230] G. Muraleedharan, A. Miyake, and I. H. Deutsch. Quantum computational supremacy in the sampling of bosonic random walkers on a one-dimensional lattice. New J. Phys., 21(5):055003, May 2019.

[231] S. Murmann, A. Bergschneider, V. M. Klinkhamer, G. Zürn, T. Lompe, and S. Jochim. Two fermions in a double well: Exploring a fundamental building block of the Hubbard model. Physical Review Letters, 114(8):080402, Feb. 2015.

[232] L. Neuhaus, R. Metzdorff, S. Chua, T. Jacqmin, T. Briant, A. Heidmann, P.-F. Cohadon, and S. Deléglise. PyRPL (Python Red Pitaya Lockbox) — An open-source software package for FPGA-controlled quantum optics experiments. In 2017 Conference on Lasers and Electro-Optics Europe European Quantum Electronics Conference (CLEO/Europe-EQEC), pages 1–1, June 2017.

[233] B. C. Nichol, R. Srinivas, D. P. Nadlinger, P. Drmota, D. Main, G. Araneda, C. J. Ballance, and D. M. Lucas. A quantum network of entangled optical atomic clocks. Nature, 609(7928):689–694, Sept. 2022.

[234] T. Nicholson, S. Campbell, R. Hutson, G. Marti, B. Bloom, R. McNally, W. Zhang, M. Barrett, M. Safronova, G. Strouse, W. Tew, and J. Ye. Systematic evaluation of an atomic clock at 2×10^{-18} total uncertainty. Nat Commun, 6(1):6896, Nov. 2015.

[235] T. L. Nicholson, S. Blatt, B. J. Bloom, J. R. Williams, J. W. Thomsen, J. Ye, and P. S. Julienne. Optical Feshbach resonances: Field-dressed theory and comparison with experiments. Phys. Rev. A, 92(2):022709, Aug. 2015.

[236] M. A. Norcia. New Tools for Precision Measurement and Quantum Science with Narrow Linewidth Optical Transitions. PhD thesis, 2017.

[237] M. A. Norcia, W. B. Cairncross, K. Barnes, P. Battaglino, A. Brown, M. O. Brown, K. Cassella, C.-A. Chen, R. Coxe, D. Crow, J. Epstein, C. Griger, A. M. W. Jones, H. Kim, J. M. Kindem, J. King, S. S. Kondov, K. Kotru, J. Lauigan, M. Li, M. Lu, E. Megidish, J. Marjanovic, M. McDonald, T. Mittiga, J. A. Muniz, S. Narayanaswami, C. Nishiguchi, R. Notermans, T. Paule, K. Pawlak, L. Peng, A. Ryou, A. Smull, D. Stack, M. Stone, A. Sucich, M. Urbanek, R. van de Veerdonk, Z. Vendeiro, T. Wilkason, T.-Y. Wu, X. Xie, X. Zhang, and B. J. Bloom. Mid-circuit qubit measurement and rearrangement in a ^{171}Yb atomic array, June 2023.

[238] M. A. Norcia, J. R. K. Cline, J. P. Bartolotta, M. J. Holland, and J. K. Thompson. Narrow-line Laser Cooling by Adiabatic Transfer. New Journal of Physics, 20(2):023021, Feb. 2018.

[239] M. A. Norcia and J. K. Thompson. Simple Laser Stabilization to the Strontium ^{88}Sr Transition at 707 nm. Review of Scientific Instruments, 87(2):023110, Feb. 2016.

[240] M. A. Norcia, A. W. Young, W. J. Eckner, E. Oelker, J. Ye, and A. M. Kaufman. Seconds-scale coherence on an optical clock transition in a tweezer array. Science, 366(6461):93–97, 2019.

[241] M. A. Norcia, A. W. Young, and A. M. Kaufman. Microscopic control and detection of ultracold strontium in optical-tweezer arrays. Phys. Rev. X, 8(4):041054, Dec. 2018.

[242] E. Oelker, R. B. Hutson, C. J. Kennedy, L. Sonderhouse, T. Bothwell, A. Goban, D. Kedar, C. Sanner, J. M. Robinson, G. E. Marti, D. G. Matei, T. Legero, M. Giunta, R. Holzwarth, F. Riehle, U. Sterr, and J. Ye. Demonstration of 4.8×10^{-17} stability at 1 s for two independent optical clocks. Nature Photonics, 13(10):714–719, Oct. 2019.

[243] C. Oh, Y. Lim, B. Fefferman, and L. Jiang. Classical simulation of boson sampling based on graph structure. Phys. Rev. Lett., 128(19):190501, May 2022.

[244] C. Oh, M. Liu, Y. Alexeev, B. Fefferman, and L. Jiang. Tensor network algorithm for simulating experimental Gaussian boson sampling, June 2023.

[245] N. D. Oppong. Towards a Degenerate Fermi Gas of Strontium in an Optical Lattice. PhD thesis, 2015.

[246] T. Đorđević, P. Samutpraphoot, P. L. Ocola, H. Bernien, B. Grinkemeyer, I. Dimitrova, V. Vuletić, and M. D. Lukin. Entanglement transport and a nanophotonic interface for atoms in optical tweezers. Science, 373(6562):1511–1514, Sept. 2021.

[247] S. Origlia, M. S. Pramod, S. Schiller, Y. Singh, K. Bongs, R. Schwarz, A. Al-Masoudi, S. Dörscher, S. Herbers, S. Häfner, U. Sterr, and Ch. Lisdat. Towards an optical clock for space: Compact, high-performance optical lattice clock based on bosonic atoms. Phys. Rev. A, 98(5):053443, Nov. 2018.

[248] V. D. Ovsiannikov, V. G. Pal'chikov, A. V. Taichenachev, V. I. Yudin, and H. Katori. Multipole, nonlinear, and anharmonic uncertainties of clocks of Sr atoms in an optical lattice. Phys. Rev. A, 88(1):013405, July 2013.

[249] J. O. Owens, M. A. Broome, D. N. Biggerstaff, M. E. Goggin, A. Fedrizzi, T. Linjordet, M. Ams, G. D. Marshall, J. Twamley, M. J. Withford, and A. G. White. Two-photon quantum walks in an elliptical direct-write waveguide array. New J. Phys., 13(7):075003, July 2011.

[250] S. Paesani, Y. Ding, R. Santagati, L. Chakhmakhchyan, C. Vigliar, K. Rottwitt, L. K. Oxenløwe, J. Wang, M. G. Thompson, and A. Laing. Generation and sampling of quantum states of light in a silicon chip. Nat. Phys., 15(9):925–929, Sept. 2019.

[251] F. A. Palm, J. Kwan, B. Bakkali-Hassani, M. Greiner, U. Schollwöck, N. Goldman, and F. Grusdt. Growing Extended Laughlin States in a Quantum Gas Microscope: A Patchwork Construction, Sept. 2023.

[252] P. B. Patel, Z. Yan, B. Mukherjee, R. J. Fletcher, J. Struck, and M. W. Zwierlein. Universal sound diffusion in a strongly interacting Fermi gas. Science, 370(6521):1222–1226, Dec. 2020.

[253] N. K. Pavlis and M. A. Weiss. A re-evaluation of the relativistic redshift on frequency standards at NIST, Boulder, Colorado, USA. Metrologia, 54(4):535–548, Aug. 2017.

[254] E. Pedrozo-Peñafiel, S. Colombo, C. Shu, A. F. Adiyatullin, Z. Li, E. Mendez, B. Braverman, A. Kawasaki, D. Akamatsu, Y. Xiao, and V. Vuletić. Entanglement on an optical atomic-clock transition. Nature, 588(7838):414–418, Dec. 2020.

[255] A. W.-C. Peng. Quantum Gas Microscope With Optical Lattice. PhD thesis, May 2010.

[256] A. Peruzzo, M. Lobino, J. C. F. Matthews, N. Matsuda, A. Politi, K. Poulios, X.-Q. Zhou, Y. Lahini, N. Ismail, K. Wörhoff, Y. Bromberg, Y. Silberberg, M. G. Thompson, and J. L. OBrien. Quantum walks of correlated photons. Science, 329(5998):1500–1503, Sept. 2010.

[257] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien. A variational eigenvalue solver on a photonic quantum processor. Nat Commun, 5(1):4213, July 2014.

[258] A. Petrescu and K. L. Hur. Chiral Mott Insulators, Meissner Effect, and Laughlin States in Quantum Ladders. Phys. Rev. B, 91(5):054520, Feb. 2015.

[259] L. Pezzè, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein. Quantum metrology with nonclassical states of atomic ensembles. Rev. Mod. Phys., 90(3):035005, Sept. 2018.

[260] I. Pikovski, M. Zych, F. Costa, and Č. Brukner. Universal decoherence due to gravitational time dilation. Nature Physics, 11(8):668–672, Aug. 2015.

[261] I. Pikovski, M. Zych, F. Costa, and Č. Brukner. Time dilation in quantum systems and decoherence. New J. Phys., 19(2):025011, Feb. 2017.

[262] L. Pioge, B. Seron, L. Novo, and N. J. Cerf. Enhanced bunching of nearly indistinguishable bosons, Aug. 2023.

[263] S. G. Porsev, M. S. Safronova, U. I. Safronova, and M. G. Kozlov. Multipolar Polarizabilities and Hyperpolarizabilities in the Sr Optical Lattice Clock. Phys. Rev. Lett., 120(6):063204, Feb. 2018.

[264] P. M. Preiss, R. Ma, M. E. Tai, A. Lukin, M. Rispoli, P. Zupancic, Y. Lahini, R. Islam, and M. Greiner. Strongly correlated quantum walks in optical lattices. Science, 347(6227):1229–1233, Mar. 2015.

[265] J. Preskill. Quantum Information Notes, 2015.

[266] W. Purwanto and S. Zhang. Quantum Monte Carlo method for the ground state of many-boson systems. Phys. Rev. E, 70(5):056702, Nov. 2004.

[267] X.-L. Qi and S.-C. Zhang. Topological insulators and superconductors. Rev. Mod. Phys., 83(4):1057–1110, Oct. 2011.

[268] R. Raussendorf, D. E. Browne, and H. J. Briegel. Measurement-based quantum computation with cluster states. Physical Review A, 68(2), Aug. 2003.

[269] S. Ravets, H. Labuhn, D. Barredo, L. Béguin, T. Lahaye, and A. Browaeys. Coherent dipole–dipole coupling between two single Rydberg atoms at an electrically-tuned Förster resonance. Nature Phys, 10(12):914–917, Dec. 2014.

[270] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani. Experimental realization of any discrete unitary operator. Phys. Rev. Lett., 73(1):58–61, July 1994.

[271] I. Reichenbach and I. H. Deutsch. Sideband cooling while preserving coherences in the nuclear spin state in group-II-like atoms. Phys. Rev. Lett., 99(12):123001, Sept. 2007.

[272] D. Reitzner, D. Nagaj, and V. Bužek. Quantum Walks. Acta Phys. Slovaca., 61(6), Dec. 2011.

[273] F. Riehle. Frequency Standards: Basics and Applications. Wiley, 1 edition, Sept. 2003.

[274] M. Rigol and A. Muramatsu. Emergence of Quasicondensates of Hard-Core Bosons at Finite Momentum. Phys. Rev. Lett., 93(23):230404, Dec. 2004.

[275] W. J. Riley. Handbook of Frequency Stability Analysis. 2008.

[276] C. Robens, I. Arrazola, W. Alt, D. Meschede, L. Lamata, E. Solano, and A. Alberti. Boson sampling with ultracold atoms, Aug. 2022.

[277] J. M. Robinson, M. Miklos, Y. M. Tso, C. J. Kennedy, D. Kedar, J. K. Thompson, and J. Ye. Direct comparison of two spin squeezed optical clocks below the quantum projection noise limit, Nov. 2022.

[278] J. S. Rosenberg, L. Christakis, E. Guardado-Sanchez, Z. Z. Yan, and W. S. Bakr. Observation of the Hanbury Brown–Twiss effect with ultracold molecules. Nat. Phys., 18(9):1062–1066, Sept. 2022.

[279] J. R. Rubbmark and S. A. Borgström. Rydberg Series in Strontium Found in Absorption by Selectively Laser-Excited Atoms. Phys. Scr., 18(4):196, Oct. 1978.

[280] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe. Experimental entanglement of four particles. Nature, 404(6775):256–259, Mar. 2000.

[281] J. J. Sakurai and J. J. Napolitano. Modern Quantum Mechanics. Pearson Education, 2nd edition, 2011.

[282] J. E. Sansonetti. Wavelengths, Transition Probabilities, and Energy Levels for the Spectra of Strontium Ions (Sr II through Sr XXXVIII). Journal of Physical and Chemical Reference Data, 41(1):013102–013102–119, Mar. 2012.

[283] J. E. Sansonetti and G. Nave. Wavelengths, Transition Probabilities, and Energy Levels for the Spectrum of Neutral Strontium (SrI). Journal of Physical and Chemical Reference Data, 39(3):033103, Sept. 2010.

[284] L. Sansoni, F. Sciarrino, G. Vallone, P. Mataloni, A. Crespi, R. Ramponi, and R. Osellame. Two-particle bosonic-fermionic quantum walk via integrated photonics. Phys. Rev. Lett., 108(1):010502, Jan. 2012.

[285] G. Santarelli, C. Audoin, A. Makdissi, P. Laurent, G. J. Dick, and A. Clairon. Frequency stability degradation of an oscillator slaved to a periodically interrogated atomic resonator. IEEE Trans Ultrason Ferroelectr Freq Control, 45(4):887–894, 1998.

[286] S. Saskin, J. T. Wilson, B. Grinkemeyer, and J. D. Thompson. Narrow-line cooling and imaging of ytterbium atoms in an optical tweezer array. Phys. Rev. Lett., 122(14):143002, Apr. 2019.

[287] T. A. Savard, K. M. O'Hara, and J. E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. Phys. Rev. A, 56(2):R1095–R1098, Aug. 1997.

[288] F. Scazza, C. Hofrichter, M. Höfer, P. C. De Groot, I. Bloch, and S. Fölling. Observation of two-orbital spin-exchange interactions with ultracold SU(N)-symmetric fermions. Nature Phys, 10(10):779–784, Oct. 2014.

[289] S. Scheel. Permanents in linear optical networks. arXiv:quant-ph/0406127, June 2004.

[290] N. Schine, A. W. Young, W. J. Eckner, M. J. Martin, and A. M. Kaufman. Long-lived Bell states in an array of optical clock qubits. Nat. Phys., pages 1–7, Aug. 2022.

[291] M. Schioppo, R. C. Brown, W. F. McGrew, N. Hinkley, R. J. Fasano, K. Beloy, T. H. Yoon, G. Milani, D. Nicolodi, J. A. Sherman, N. B. Phillips, C. W. Oates, and A. D. Ludlow. Ultrastable optical clock with two cold-atom ensembles. Nature Photonics, 11(1):48–52, Jan. 2017.

[292] N. Schlosser, G. Reymond, I. Protsenko, and P. Grangier. Sub-poissonian loading of single atoms in a microscopic dipole trap. Nature, 411(6841):1024–1027, June 2001.

[293] P.-I. Schneider, S. Grishkevich, and A. Saenz. Ab initio determination of Bose-Hubbard parameters for two ultracold atoms in an optical lattice using a three-well potential. Phys. Rev. A, 80(1):013404, July 2009.

[294] P. Scholl, M. Schuler, H. J. Williams, A. A. Eberharter, D. Barredo, K.-N. Schymik, V. Lienhard, L.-P. Henry, T. C. Lang, T. Lahaye, A. M. Läuchli, and A. Browaeys. Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms. Nature, 595(7866):233–238, July 2021.

[295] M. Schroeder. Synthesis of low-peak-factor signals and binary sequences with low autocorrelation (Corresp.). IEEE Transactions on Information Theory, 16(1):85–89, Jan. 1970.

[296] J. Sebby-Strabley, M. Anderlini, P. S. Jessen, and J. V. Porto. Lattice of double wells for manipulating pairs of cold atoms. Phys. Rev. A, 73(3):033605, Mar. 2006.

[297] G. Semeghini, H. Levine, A. Keesling, S. Ebadi, T. T. Wang, D. Bluvstein, R. Verresen, H. Pichler, M. Kalinowski, R. Samajdar, A. Omran, S. Sachdev, A. Vishwanath, M. Greiner, V. Vuletić, and M. D. Lukin. Probing topological spin liquids on a programmable quantum simulator. Science, 374(6572):1242–1247, Dec. 2021.

[298] B. Seron, L. Novo, and N. J. Cerf. Boson bunching is not maximized by indistinguishable particles. Nat. Photon., 17(8):702–709, Aug. 2023.

[299] A. Shaji and C. M. Caves. Qubit metrology and decoherence. Phys. Rev. A, 76(3):032111, Sept. 2007.

[300] J. Shao. Mathematical Statistics. Springer Texts in Statistics. Springer, New York, 2nd ed edition, 2003.

[301] A. L. Shaw, R. Finkelstein, R. B.-S. Tsai, P. Scholl, T. H. Yoon, J. Choi, and M. Endres. Multi-ensemble metrology by programming local rotations with atom movements, Mar. 2023.

[302] V. Shchesnovich. Distinguishing noisy boson sampling from classical simulations. Quantum, 5:423, Mar. 2021.

[303] V. S. Shchesnovich. The permanent-on-top conjecture is false. Linear Algebra and its Applications, 490:196–201, Feb. 2016.

[304] V. S. Shchesnovich. Universality of generalized bunching and efficient assessment of boson sampling. Phys. Rev. Lett., 116(12):123601, Mar. 2016.

[305] J. F. Sherson, C. Weitenberg, M. Endres, M. Cheneau, I. Bloch, and S. Kuhr. Single-atom-resolved fluorescence imaging of an atomic Mott insulator. Nature, 467(7311):68–72, Sept. 2010.

[306] N. Shettell. Quantum Information Techniques for Quantum Metrology. PhD thesis, Jan. 2022.

[307] C. Shi, J.-L. Robyr, U. Eismann, M. Zawada, L. Lorini, R. Le Targat, and J. Lodewyck. Polarizabilities of the ^{87}Sr clock transition. Phys. Rev. A, 92(1):012516, July 2015.

[308] N. Šibalić, J. D. Pritchard, C. S. Adams, and K. J. Weatherill. ARC: An open-source library for calculating properties of alkali Rydberg atoms. Computer Physics Communications, 220:319–331, Nov. 2017.

[309] J. S. Sidhu and P. Kok. Geometric perspective on quantum parameter estimation. AVS Quantum Science, 2(1):014701, Feb. 2020.

[310] Simons Collaboration on the Many-Electron Problem, M. Qin, C.-M. Chung, H. Shi, E. Vitali, C. Hubig, U. Schollwöck, S. R. White, and S. Zhang. Absence of Superconductivity in the Pure Two-Dimensional Hubbard Model. Phys. Rev. X, 10(3):031016, July 2020.

[311] S. Snigirev, A. J. Park, A. Heinz, I. Bloch, and S. Blatt. Fast and dense magneto-optical traps for Strontium. Phys. Rev. A, 99(6):063421, June 2019.

[312] M. Sohmen, M. J. Mark, M. Greiner, and F. Ferlaino. A ship-in-a-bottle quantum gas microscope setup for magnetic mixtures. SciPost Physics, 15(5):182, Nov. 2023.

[313] L. I. Sonderhouse. Quantum Gas Engineering for Atomic Clocks. PhD thesis, 2020.

[314] F. Sorrentino, G. Ferrari, N. Poli, R. Drullinger, and G. M. Tino. Laser cooling and trapping of atomic strontium for ultracold atom physics, high-precision spectroscopy and quantum sensors. Mod. Phys. Lett. B, 20(21):1287–1320, Sept. 2006.

[315] N. Spagnolo, C. Vitelli, L. Sansoni, E. Maiorino, P. Mataloni, F. Sciarrino, D. J. Brod, E. F. Galvão, A. Crespi, R. Ramponi, and R. Osellame. General rules for bosonic bunching in multimode interferometers. Phys. Rev. Lett., 111(13):130503, Sept. 2013.

[316] B. M. Spar, E. Guardado-Sanchez, S. Chi, Z. Z. Yan, and W. S. Bakr. Realization of a Fermi-Hubbard Optical Tweezer Array. Phys. Rev. Lett., 128(22):223202, June 2022.

[317] J. B. Spring, B. J. Metcalf, P. C. Humphreys, W. S. Kolthammer, X.-M. Jin, M. Barbieri, A. Datta, N. Thomas-Peter, N. K. Langford, D. Kundys, J. C. Gates, B. J. Smith, P. G. R. Smith, and I. A. Walmsley. Boson sampling on a photonic chip. Science, 339(6121):798–801, Feb. 2013.

[318] D. A. Steck. Quantum and Atom Optics. 2017.

[319] S. Stellmer. Degenerate Quantum Gases of Strontium. PhD thesis, May 2014.

[320] S. Stellmer and F. Schreck. Reservoir spectroscopy of $5s5p\ ^3P_2$–$5snd\ ^3D_{1,2,3}$ transitions in strontium. Phys. Rev. A, 90(2):022512, Aug. 2014.

[321] L. Su, A. Douglas, M. Szurek, R. Groth, S. F. Ozturk, A. Krahn, A. H. Hébert, G. A. Phelps, S. Ebadi, S. Dickerson, F. Ferlaino, O. Marković, and M. Greiner. Dipolar quantum solids emerging in a Hubbard quantum simulator. Nature, 622(7984):724–729, Oct. 2023.

[322] A. V. Taichenachev, V. I. Yudin, C. W. Oates, C. W. Hoyt, Z. W. Barber, and L. Hollberg. Magnetic field-induced spectroscopy of forbidden optical transitions with application to lattice-based optical atomic clocks. Phys. Rev. Lett., 96(8):083001, Mar. 2006.

[323] J. D. Thompson, T. G. Tiecke, A. S. Zibrov, V. Vuletić, and M. D. Lukin. Coherence and Raman sideband cooling of a single atom in an optical tweezer. Phys. Rev. Lett., 110(13), Mar. 2013.

[324] W. Tian, W. J. Wee, A. Qu, B. J. M. Lim, P. R. Datla, V. P. W. Koh, and H. Loh. Parallel assembly of arbitrary defect-free atom arrays with a multitweezer algorithm. Phys. Rev. Appl., 19(3):034048, Mar. 2023.

[325] M. C. Tichy. Entanglement and Interference of Identical Particles. PhD thesis, Uni. Freiburg, 2011.

[326] M. C. Tichy. Sampling of partially distinguishable bosons and the relation to the multidimensional permanent. Phys. Rev. A, 91(2):022316, Feb. 2015.

[327] M. C. Tichy, K. Mayer, A. Buchleitner, and K. Mølmer. Stringent and efficient assessment of Boson-Sampling devices. Phys. Rev. Lett., 113(2):020502, July 2014.

[328] M. C. Tichy, M. Tiersch, F. Mintert, and A. Buchleitner. Many-particle interference beyond many-boson and many-fermion statistics. New J. Phys., 14(9):093015, Sept. 2012.

[329] M. Tillmann, B. Dakić, R. Heilmann, S. Nolte, A. Szameit, and P. Walther. Experimental boson sampling. Nature Photon, 7(7):540–544, July 2013.

[330] J. Tilly, H. Chen, S. Cao, D. Picozzi, K. Setia, Y. Li, E. Grant, L. Wossnig, I. Rungger, G. H. Booth, and J. Tennyson. The Variational Quantum Eigensolver: A review of methods and best practices. Physics Reports, 986:1–128, Nov. 2022.

[331] G. Toth and I. Apellaniz. Quantum metrology from a quantum information science perspective. J. Phys. A: Math. Theor., 47(42):424006, Oct. 2014.

[332] M. C. Tran, D. K. Mark, W. W. Ho, and S. Choi. Measuring Arbitrary Physical Properties in Analog Quantum Simulation, Dec. 2022.

[333] J. Trautmann, D. Yankelev, V. Klüsener, A. J. Park, I. Bloch, and S. Blatt. The $^1S_0$-$^3P_2$ magnetic quadrupole transition in neutral strontium, Nov. 2022.

[334] R. Trivedi, A. F. Rubio, and J. I. Cirac. Quantum advantage and stability to errors in analogue quantum simulators, Dec. 2022.

[335] M. Troyer and U.-J. Wiese. Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations. Phys. Rev. Lett., 94(17):170201, May 2005.

[336] I. Ushijima, M. Takamoto, M. Das, T. Ohkubo, and H. Katori. Cryogenic optical lattice clocks. Nature Photonics, 9(3):185–189, Mar. 2015.

[337] H. Uys, M. J. Biercuk, A. P. VanDevender, C. Ospelkaus, D. Meiser, R. Ozeri, and J. J. Bollinger. Decoherence due to Elastic Rayleigh Scattering. Phys. Rev. Lett., 105(20):200401, Nov. 2010.

[338] C. L. Vaillant, M. P. A. Jones, and R. M. Potvliege. Long-range Rydberg–Rydberg interactions in calcium, strontium and ytterbium. J. Phys. B: At. Mol. Opt. Phys., 45(13):135004, July 2012.

[339] L. Valiant. The complexity of computing the permanent. Theoretical Computer Science, 8(2):189–201, 1979.

[340] A. Van Den Bos. A new method for synthesis of low-peak-factor signals. IEEE Transactions on Acoustics, Speech, and Signal Processing, 35(1):120–122, Jan. 1987.

[341] J. Van Dongen, C. Zhu, D. Clement, G. Dufour, J. L. Booth, and K. W. Madison. Trap-depth determination from residual gas collisions. Phys. Rev. A, 84(2):022708, Aug. 2011.

[342] R. F. C. Vessot. A review of atomic clock technology, the performance capability of present spaceborne and terrestrial atomic clocks, and a look toward the future. In Relativistic Gravitational Experiments in Space, Aug. 1989.

[343] L. Vidmar, J. P. Ronzheimer, M. Schreiber, S. Braun, S. S. Hodgman, S. Langer, F. Heidrich-Meisner, I. Bloch, and U. Schneider. Dynamical Quasicondensation of Hard-Core Bosons at Finite Momenta. Phys. Rev. Lett., 115(17):175301, Oct. 2015.

[344] H. Wang, J. Qin, X. Ding, M.-C. Chen, S. Chen, X. You, Y.-M. He, X. Jiang, L. You, Z. Wang, C. Schneider, J. J. Renema, S. Höfling, C.-Y. Lu, and J.-W. Pan. Boson sampling with 20 input photons and a 60-mode interferometer in a 10^{14}-dimensional Hilbert space. Phys. Rev. Lett., 123(25):250503, Dec. 2019.

[345] S. Wang, W. Zhang, T. Zhang, S. Mei, Y. Wang, J. Hu, and W. Chen. Accelerating the assembly of defect-free atomic arrays with maximum parallelisms. Phys. Rev. Appl., 19(5):054032, May 2023.

[346] I. M. Wanless. Lieb's permanental dominance conjecture. In The Physics and Mathematics of Elliott Lieb, volume 2, pages 501–506. EMS Press, Berlin, June 2022.

[347] D. Wei, A. Rubio-Abadal, B. Ye, F. Machado, J. Kemp, K. Srakaew, S. Hollerith, J. Rui, S. Gopalakrishnan, N. Y. Yao, I. Bloch, and J. Zeiher. Quantum gas microscopy of Kardar-Parisi-Zhang superdiffusion. Science, 376(6594):716–720, May 2022.

[348] P. Weinberg and M. Bukov. QuSpin: A Python Package for Dynamics and Exact Diagonalisation of Quantum Many Body Systems. Part II: Bosons, fermions and higher spins. SciPost Phys., 7(2):020, Aug. 2019.

[349] C. Weitenberg, M. Endres, J. F. Sherson, M. Cheneau, P. Schauß, T. Fukuhara, I. Bloch, and S. Kuhr. Single-spin addressing in an atomic Mott insulator. Nature, 471(7338):319–324, Mar. 2011.

[350] S. Welte, B. Hacker, S. Daiss, S. Ritter, and G. Rempe. Photon-Mediated Quantum Gate between Two Neutral Atoms in an Optical Cavity. Phys. Rev. X, 8(1):011018, Feb. 2018.

[351] C. E. Wieman and L. Hollberg. Using diode lasers for atomic physics. Rev. Sci. Instrum., 62(1):20, 1991.

[352] T. Wilk, A. Gaëtan, C. Evellin, J. Wolters, Y. Miroshnychenko, P. Grangier, and A. Browaeys. Entanglement of two individual neutral atoms using Rydberg blockade. Phys. Rev. Lett., 104(1):010502, Jan. 2010.

[353] J. T. Wilson, S. Saskin, Y. Meng, S. Ma, R. Dilip, A. P. Burgers, and J. D. Thompson. Trapping Alkaline Earth Rydberg Atoms Optical Tweezer Arrays. Phys. Rev. Lett., 128(3):033201, Jan. 2022.

[354] D. Wineland, C. Monroe, W. Itano, D. Leibfried, B. King, and D. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. Journal of Research of the National Institute of Standards and Technology, 103(3):259, May 1998.

[355] D. J. Wineland, J. J. Bollinger, W. M. Itano, and D. J. Heinzen. Squeezed atomic states and projection noise in spectroscopy. Phys. Rev. A, 50(1):67–88, July 1994.

[356] D. J. Wineland and W. M. Itano. Laser cooling of atoms. Physical Review A, 20(4):1521–1540, Oct. 1979.

[357] R. Yamamoto, J. Kobayashi, T. Kuno, K. Kato, and Y. Takahashi. An ytterbium quantum gas microscope with narrow-line laser cooling. New Journal of Physics, 18(2):023016, Feb. 2016.

[358] M. Yan, B. J. DeSalvo, B. Ramachandhran, H. Pu, and T. C. Killian. Controlling Condensate Collapse and Expansion with an Optical Feshbach Resonance. Phys. Rev. Lett., 110(12):123201, Mar. 2013.

[359] Z. Z. Yan, B. M. Spar, M. L. Prichard, S. Chi, H.-T. Wei, E. Ibarra-García-Padilla, K. R. A. Hazzard, and W. S. Bakr. Two-dimensional programmable tweezer arrays of fermions. Phys. Rev. Lett., 129(12):123201, Sept. 2022.

[360] B. Yang, H. Sun, C.-J. Huang, H.-Y. Wang, Y. Deng, H.-N. Dai, Z.-S. Yuan, and J.-W. Pan. Cooling and entangling ultracold atoms in optical lattices. Science, June 2020.

[361] A. Yariv. Quantum Electronics. John Wiley & Sons, 3 edition, 1989.

[362] M. Yasuda and H. Katori. Lifetime measurement of the $^3P_2$ metastable state of strontium atoms. Phys. Rev. Lett., 92(15):153004, Apr. 2004.

[363] M. Yasuda, T. Kishimoto, M. Takamoto, and H. Katori. Photoassociation spectroscopy of ^{88}Sr: Reconstruction of the wave function near the last node. Phys. Rev. A, 73(1):011403, Jan. 2006.

[364] A. Young, W. G. Graham, T. J. Morgan, and L. Hüwel. Laser-induced breakdown in liquid water: Influence of repeated laser pulses on plasma formation and emission. Journal of Applied Physics, 129(18):183303, May 2021.

[365] A. W. Young, W. J. Eckner, W. R. Milner, D. Kedar, M. A. Norcia, E. Oelker, N. Schine, J. Ye, and A. M. Kaufman. Half-minute-scale atomic coherence and high relative stability in a tweezer clock. Nature, 588(7838):408–413, Dec. 2020.

[366] A. W. Young, W. J. Eckner, N. Schine, A. M. Childs, and A. M. Kaufman. Tweezer-programmable 2D quantum walks in a Hubbard-regime lattice. Science, 377(6608):885–889, Aug. 2022.

[367] A. W. Young, S. Geller, W. J. Eckner, N. Schine, S. Glancy, E. Knill, and A. M. Kaufman. An atomic boson sampler. In review, 2023.

[368] J. T. Young, T. Boulier, E. Magnan, E. A. Goldschmidt, R. M. Wilson, S. L. Rolston, J. V. Porto, and A. V. Gorshkov. Dissipation-induced dipole blockade and antiblockade in driven Rydberg systems. Phys. Rev. A, 97(2):023424, Feb. 2018.

[369] J. T. Young, S. R. Muleady, M. A. Perlin, A. M. Kaufman, and A. M. Rey. Enhancing spin squeezing using soft-core interactions. Phys. Rev. Res., 5(1):L012033, Mar. 2023.

[370] J. Zeiher, R. van Bijnen, P. Schauß, S. Hild, J.-y. Choi, T. Pohl, I. Bloch, and C. Gross. Many-body interferometry of a Rydberg-dressed spin lattice. Nature Physics, 12(12):1095–1099, Dec. 2016.

[371] W.-Y. Zhang, M.-G. He, H. Sun, Y.-G. Zheng, Y. Liu, A. Luo, H.-Y. Wang, Z.-H. Zhu, P.-Y. Qiu, Y.-C. Shen, X.-K. Wang, W. Lin, S.-T. Yu, B.-C. Li, B. Xiao, M.-D. Li, Y.-M. Yang, X. Jiang, H.-N. Dai, Y. Zhou, X. Ma, Z.-S. Yuan, and J.-W. Pan. Scalable Multipartite Entanglement Created by Spin Exchange in an Optical Lattice. Phys. Rev. Lett., 131(7):073401, Aug. 2023.

[372] X. Zhang, K. Beloy, Y. S. Hassan, W. F. McGrew, C.-C. Chen, J. L. Siegel, T. Grogan, and A. D. Ludlow. Subrecoil Clock-Transition Laser Cooling Enabling Shallow Optical Lattice Clocks. Phys. Rev. Lett., 129(11):113202, Sept. 2022.

[373] X. Zheng, J. Dolde, V. Lochab, B. N. Merriman, H. Li, and S. Kolkowitz. Differential clock comparisons with a multiplexed optical lattice clock. Nature, 602(7897):425–430, Feb. 2022.

[374] Y.-G. Zheng, W.-Y. Zhang, Y.-C. Shen, A. Luo, Y. Liu, M.-G. He, H.-R. Zhang, W. Lin, H.-Y. Wang, Z.-H. Zhu, M.-C. Chen, C.-Y. Lu, S. Thanasilp, D. G. Angelakis, Z.-S. Yuan, and J.-W. Pan. Efficiently Extracting Multi-Point Correlations of a Floquet Thermalized System. arXiv:2210.08556, Oct. 2022.

[375] H.-S. Zhong, Y.-H. Deng, J. Qin, H. Wang, M.-C. Chen, L.-C. Peng, Y.-H. Luo, D. Wu, S.-Q. Gong, H. Su, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, J. J. Renema, C.-Y. Lu, and J.-W. Pan. Phase-programmable Gaussian boson sampling using stimulated squeezed light. Phys. Rev. Lett., 127(18):180502, Oct. 2021.

[376] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, P. Hu, X.-Y. Yang, W.-J. Zhang, H. Li, Y. Li, X. Jiang, L. Gan, G. Yang, L. You, Z. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan. Quantum computational advantage using photons. Science, 370(6523):1460–1463, Dec. 2020.

[377] S. Zhou, M. Zhang, J. Preskill, and L. Jiang. Achieving the Heisenberg limit in quantum metrology using quantum error correction. Nat Commun, 9(1):78, Dec. 2018.

[378] X. Zhou, X. Xu, X. Chen, and J. Chen. Magic wavelengths for terahertz clock transitions. Physical Review A, 81(1), Jan. 2010.

[379] P. Zupancic, P. M. Preiss, R. Ma, A. Lukin, M. E. Tai, M. Rispoli, R. Islam, and M. Greiner. Ultra-precise holographic beam shaping for microscopic quantum control. Opt. Express, OE, 24(13):13881–13893, June 2016.

[380] M. Zych, F. Costa, I. Pikovski, and Č. Brukner. Quantum interferometric visibility as a witness of general relativistic proper time. Nat Commun, 2(1):505, Oct. 2011.

[381] M. Zych, I. Pikovski, F. Costa, and Č. Brukner. General relativistic effects in quantum interference of "clocks". J. Phys.: Conf. Ser., 723:012044, June 2016.

Appendix A

Notes on estimation theory

A.1 Classical estimation theory

To determine how much information a probability distribution $p(\boldsymbol{x}|\phi)$ contains about some parameter ϕ, we are interested in something like the variance of $\partial p(\boldsymbol{x}|\phi)/\partial\phi$. It is useful to include a logarithm in the above quantity to ensure that the resulting measure of information is additive when considering joint probability distributions of independent random variables. The above considerations lead to the definition of the "score" s:

$$s(\phi) \coloneqq \frac{\partial \ln p(\boldsymbol{x}|\phi)}{\partial\phi}. \tag{A.1}$$

Taking the variance of the score leads to the definition of the Fisher information $\mathcal{I}$:[1]

$$\mathcal{I}\left(p(\boldsymbol{x}|\phi)\right) \coloneqq \int d\boldsymbol{x} s(\phi)^2. \tag{A.2}$$

$\mathcal{I}$ is non-negative and additive, meaning that when $\boldsymbol{x}$ corresponds to N independent measurements of the random variable X, it has the property:

$$\mathcal{I}\left(p(\boldsymbol{x}|\phi)\right) = \mathcal{I}\left(Np(X|\phi)\right). \tag{A.3}$$

By taking the derivative of Eq. 5.1 with respect to ϕ and applying the Cauchy-Schwarz inequality, one arrives at the Cramér-Rao bound [76]:

$$\Delta^2\check{\phi} \geq \frac{1}{\mathcal{I}\left(p(\boldsymbol{x}|\phi)\right)}. \tag{A.4}$$

[1] Note that it is straightforward to generalize the Fisher information to multiple parameters [197], but we omit these details in this simplified discussion.

As discussed in Sec. 5.1, a protocol for estimating ϕ is efficient if it saturates this bound.

One general-purpose scheme for parameter estimation is maximum likelihood estimation (MLE). The likelihood function is defined as $L(\phi|\boldsymbol{x}) = p(\boldsymbol{x}|\phi)$, and can be interpreted as a measure of how well the observed data $\boldsymbol{x}$ is explained by the parameter ϕ. Maximizing $L(\phi|\boldsymbol{x})$ over ϕ provides an estimator $\check{\phi}$ that is consistent and efficient in the infinite-data limit. Although MLE can yield results that are biased for a finite number of samples, a number of techniques can be used to correct this bias [75, 300], typically resulting in leading-order contributions to the bias of only $O(1/N^2)$ (where N is the number of independent measurements).

A.2 Quantum estimation theory

Extending the discussion in the previous section to the quantum case is, at least conceptually, relatively straightforward. Once one has picked a measurement basis, each measurement of a quantum system provides a sample of an IID random variable, and all the statements in the previous section apply. As described in Sec. 5.1, maximizing the Fisher information $\mathcal{I}$ over all possible measurements leads to the definition of the quantum Fisher information $\mathcal{Q}$.

More explicitly, we can generalize the score $s(\phi)$ (Eq. A.1) to the quantum case by considering how quickly a density matrix $\rho(\phi)$ varies with respect to a parameter ϕ. Note that $s(\phi)$ has the property:

$$\frac{\partial p(\boldsymbol{x}|\phi)}{\partial \phi} = p(\boldsymbol{x}|\phi)s(\phi). \tag{A.5}$$

We can analogously define the symmetric logarithmic derivative L [331]:[2]

$$\frac{\partial \rho(\phi)}{\partial \phi} = \frac{1}{2}\{L, \rho(\phi)\}, \tag{A.6}$$

where the anticommutator avoids any ambiguity about operator ordering. Recall from Sec. 5.1 that the generator H of the signal transforms $\rho(\phi)$ according to the von Neumann equation:

$$\frac{\partial \rho(\phi)}{\partial \phi} = -i[H, \rho(\phi)]. \tag{A.7}$$

[2] See [309] for a nice geometric interpretation of the score and symmetric logarithmic derivative.

The equality of Eqs. A.6 and A.7 yields [331]:

$$L = 2i \sum_{kl} \frac{\lambda_k - \lambda_l}{\lambda_k + \lambda_l} |k\rangle\langle l| \langle k \,|\, H \,|\, l\rangle, \tag{A.8}$$

where $|k\rangle$ is an eigenstate of $\rho(\phi)$ with eigenvalue λ_k. Given the above definitions, the quantum Fisher information $\mathcal{Q}$ is simply the expectation of the variance of L:

$$\mathcal{Q} \coloneqq \mathrm{Tr}(\rho(\phi) L^2). \tag{A.9}$$

Note that substituting a pure state into the above expressions leads to Eq. 5.7.

Appendix B

Estimating indistinguishability, bunching, and clouding in the presence of parity projection

Continuing the discussion in Sec. 8.3.1, we are interested in estimating the quantity $Q^{\mathrm{HOM}}_{S_1,S_2}$.[1] For the discussion below, we will need to keep track of the input sites defining S_1 and S_2, and therefore make the substitution $S_1 \to C_k$ and $S_2 \to C_l$. C_k is the column of lattice sites that includes site k, and similarly for C_l and l.

We will construct the estimate of $Q^{\mathrm{HOM}}_{S_1,S_2}$ ($Q^{\mathrm{HOM}}_{C_k,C_l}$) from separate estimates of $p^{\mathrm{dist}}_{C_k,C_l}$ and $p^{\mathrm{partial}}_{C_k,C_l}$, which are respectively the joint detection probabilities of one atom in C_k and one atom in C_l for distinguishable atoms, and nominally indistinguishable atoms.[2] $p^{\mathrm{dist}}_{C_k,C_l}$ can be estimated from data where a single atom is prepared in either k or l on a given experimental trial:

$$p^{\mathrm{dist}}_{C_k,C_l} = p(C_k|k,\neg\lambda)p(C_l|l,\neg\lambda) + p(C_k|l,\neg\lambda)p(C_l|k,\neg\lambda), \tag{B.1}$$

where $p(C_i|j,\neg\lambda)$ is the probability that an atom is observed in C_i, given that a loss event λ did not occur ($\neg\lambda$), and that the atom was initially prepared in state j (with $i,j \in \{k,l\}$). Eq. B.1 is a multilinear polynomial of the single atom probability distributions, and so the plug-in estimator[3] is unbiased.

$p^{\mathrm{partial}}_{C_k,C_l}$ is estimated directly from measurements with two atoms. However, the observed coincidence probability in the experiment $p^{\mathrm{obs}}_{C_k,C_l}$ is distinct from $p^{\mathrm{partial}}_{C_k,C_l}$ due to parity projection:

$$p^{\mathrm{obs}}_{C_k,C_l} = p(\nu_2|kl)p^{\mathrm{partial}}_{C_k,C_l}, \tag{B.2}$$

[1] The analysis in this section is primarily the work of Shawn Geller.
[2] With postselection on no atom loss.
[3] Namely, where one estimates $p^{\mathrm{dist}}_{C_k,C_l}$ by plugging the estimates of the single atom probabilities into Eq. B.1.

where ν_2 refers to an event where neither of the two atoms was lost due to single-atom loss processes (i.e. loss not due to parity projection). Assuming that loss acts independently and identically on each site, we can estimate $p(\nu_2|kl)$ by estimating the single atom loss probability p_λ. Specifically, we measure the probability of an event β, where exactly one of the two atoms is lost:

$$p(\beta|kl) = 2(1-p_\lambda)p_\lambda. \tag{B.3}$$

We therefore have:

$$p_\lambda = \frac{1-\sqrt{1-2p(\beta|kl)}}{2}, \tag{B.4}$$

and:

$$p(\nu_2|kl) = (1-p_\lambda)^2. \tag{B.5}$$

Using Eqs. B.2 to B.5, we can construct a plug-in estimate of $Q^{\mathrm{HOM}}_{C_k,C_l}$. Note that the plug-in estimator is not unbiased in this situation, and so we use the delta method [300] to obtain an estimate that is first order unbiased. We perform 1000 bootstrap estimates of $Q^{\mathrm{HOM}}_{C_k,C_l}$ and apply the bias-corrected percentile method [98] to obtain a confidence interval.

Note that small diagonal and next-nearest neighbor tunneling terms in the lattice break the balanced condition for C_k and C_l (see Sec. 8.3). Based on our calibration of the lattice Hamiltonian (see Secs. 7.3 and 8.4), we calculate a correction factor $\tau(C_k, C_l)$, leading to a corrected indistinguishability estimate of:

$$\mathcal{J} = \frac{1}{\tau(C_k,C_l)}\left(1-Q^{\mathrm{HOM}}_{C_k,C_l}\right). \tag{B.6}$$

Similar to before, the confidence interval in our estimate of $\mathcal{J}$ is computed using 1000 bootstrap estimates, and the bias-corrected percentile method. The combination of statistical fluctuations and $\tau < 1$ can lead to bootstrap samples that go above 1, and so we clip the upper end of our distribution to 1 in both our confidence interval and our point estimate. We do not expect the uncertainty in τ (due to imperfect calibration of the lattice potential) to be significant in comparison to the statistical fluctuations in our estimate of $\mathcal{J}$, and so we ignore this source of uncertainty in our confidence interval.

The point estimates and confidence intervals for full bunching, clouding, and generalized bunching appearing in Ch. 8 use a similar treatment to what is described above, where we account for the differing

effects of parity projection on multi- and single-atom measurements, and then construct point estimates and bootstrapped confidence intervals for the desired quantities. For a full accounting of the above estimation procedures, see [367].

Appendix C

Electronics

A significant number of different custom and off-the-shelf electronics are used in our experiment. Summarizing the designs of all of these systems is beyond the scope of this work, but we will endeavor in this section to give a few guiding principles of electronics design, and provide some useful references.

C.1 Clock distribution

In order to synchronize various parts of the experiment and generate the appropriate control signals, a radio frequency (RF) clock signal must be distributed to various electrical systems in the lab. We use a 10 MHz reference, either from a rubidium frequency standard (Stanford Research Systems FS725), or a 10 MHz maser standard distributed from NIST Boulder. The 10 MHz reference is distributed around the lab as necessary using a distribution amplifier (Stanford Research Systems FS730).

To generate digital clock signals, we use a high performance jitter attenuator (EVAL-HMC7044). This device is locked to the 10 MHz reference signal, and is able to output up to 14 independent digital clock signals with different frequencies, and with different clock standards (including LVDS, LVPECL, and CML). Although analog signals with fixed frequencies are generated in a variety of different ways in the lab, signals that must be well-referenced to the 10 MHz clock are typically generated by converting a digital clock from the HMC7044 board to an analog signal using direct digital synthesis (e.g. using an EVAL-AD9959 board).

C.2 General control electronics

The experimental control system involves multiple devices that are programmed with a set of instructions, and are then synchronized and triggered using TTL signals generated using a Spartan 6 field-programmable gate array (FPGA). Analog control voltages are generated using a National Instruments control system (PXI-1042Q, PXI-6733, BNC-2110). To communicate between the control computer and the various control electronics, we use either USB connections,[1] or, in cases where data transfer rate is critical (for example when programming tweezer moves as in Sec. C.3), UDP over ethernet.[2]

There are a number of electronic servos in the lab, most of which have been described in detail elsewhere. To control the currents in our magnetic field coils, we use a JILA designed current controller similar to the one in [52]. To control laser diode currents and temperatures, we typically use off the shelf controllers from Thorlabs (LDC2xx and TED200C), or JILA controllers of a similar design.

We use a number of different high speed loop filters, both analog and digital. Besides standard JILA designs [136], we also use the Newport LB1005-S, Toptica FALC pro, and Vescent D2-125 controllers. For applications that do not require very high bandwidth ($\lesssim 1$ MHz), we often opt to use digital loop filters due to the increased convenience they offer.[3] We have used digital loop filters based on Red Pitaya boards [232], as well as other open source designs [186]. For stabilizing laser powers using AOMs, we use both the digital loop filters described above, and specialized analog loop filters that take the significant phase delay imposed by AOMs into account [245].

C.3 RF synthesis with FPGAs

There are a number of cases where we have to synthesize RF signals that are modulated in frequency and/or amplitude. In general, we do this via direct digital synthesis (DDS) controlled with an FPGA, however the precise manner in which this is done depends on the specific application.[4]

For the synthesis of general purpose signals with carrier frequencies of $\lesssim 150$ MHz and relatively simple

[1] For long USB cable runs, we opt to use USB to fiber optic converters.

[2] For local connections with short cable runs the chances of an error are extremely low, and so UDP ends up being both faster and easier to deal with than TCP. With UDP, it is relatively straightforward to achieve data rates in excess of 1 Gbps.

[3] For example, the ability to easily include non-linear behaviors, auto re-locking features, etc.

[4] Most of our FPGA-based control electronics were developed by Felix Vietmeyer in the JILA electronics shop.

modulation (jumps and/or slow ramps of frequency and amplitude), we control several EVAL-AD9959 boards using a Spartan 6 FPGA.[5]

The tweezer control system has more stringent requirements, where up to four separate outputs (one for each AOD in the two tweezer systems) each needs to be driven with multiple RF tones (where each tone defines the position of one row or column of tweezers and therefore must be independently tunable). To generate the above control tones, we use a custom frequency synthesizer composed of an FPGA (Kintex UltraScale KCU105) driving a high speed (2.5 Gs/s) quad-channel 16-bit digital to analog converter (DAC, Texas Instruments DAC38J84EVM).[6] Specifically, the FPGA runs 512 DDS cores,[7] which are interleaved to generate 256 outputs with independently tunable frequency, phase, and amplitude. The 256 outputs are split among the four DAC channels, which each drive one of the four AODs used in our tweezer systems. The outputs are clocked at 750 MHz (but can be clocked in the gigahertz range if desired), corresponding to a maximum usable frequency of $\sim$ 300 MHz (for our AODs, we operate in the range of 125-235 MHz). Each analog output is amplified using two stages of linear RF amplifiers, with the final stage being a high power (10 W) amplifier that delivers $\sim$ 2 W ($\sim$ 5 W) of total RF power to each of the 515 nm (813 nm) AODs (see Sec. 3.3.3.1). The final result is that each AOD can be controlled with up to 64 independent RF drive tones, where each tone has 36 bits of frequency resolution, 12 bits of phase resolution, and 10 bits of amplitude resolution.

The tweezer control system can be programmed with multiple settings, and can linearly interpolate between these settings. Therefore, to execute a move between two patterns of tweezers, one only needs to program the appropriate frequencies, amplitudes, and phases of each RF drive tone at the beginning and end of the move, as well as the rates at which the frequencies and amplitudes should be varied during the move. This simple scheme allows the tweezer control system to be programmed with hundreds of moves involving hundreds of distinct tweezers in under 1 ms (using a gigabit Ethernet connection).[8]

The system used to perform arbitrary waveform generation (AWG)[9] is similar to the tweezer control

[5] We refer to this system as the "Minimoog," in honor of the Moog synthesizer that revolutionized electronic music.
[6] We refer to this system as the "Gigamoog," for obvious reasons.
[7] Implemented on the fabric of the FPGA itself, and not in a standalone DDS chip as in the Minimoog.
[8] To specify a similar control sequence as a time series, as is typically required when using off the shelf arbitrary waveform generators, many orders of magnitude more data would have to be transferred from the control computer to the signal generator, possibly resulting in significant experimental delays.
[9] This is used, for example, for modulation of the UV laser to perform the entangling gates described in Ch. 6.

system in terms of hardware, with the only difference being the use of a different FPGA (Zynq UltraScale+ ZCU106).[10] Zynq FPGAs integrate an FPGA and an ARM CPU on the same die, with the CPU providing a number of convenient features for steps that are not sensitive to timing (like communication over Ethernet). The main difference between the tweezer control system and the AWG system is how the desired signals are programmed into the device and subsequently synthesized. Note that specifying the output of an AWG as a time series can be very costly in terms of how much data must be transferred and stored in memory. However, the signals we are interested in generating typically involve comparably slow modulation of a fast carrier signal. Therefore, instead of specifying a time series defining the desired output voltage on each clock cycle, we program the AWG with a series of "snapshots." Each snapshot includes a timestamp, and the desired phase,[11] amplitude, and frequency of the output signal at that time. The AWG will linearly interpolate from the frequency and amplitude settings of one snapshot to the next, allowing us to specify the desired modulated signals without sending a significant amount of data to the AWG.[12] 16384 snapshots can be stored in memory, and the maximum update rate to go from one snapshot to another is 153.6 MHz.

C.4 High current coil control

Our high current coils (see Sec. 3.2) are switched between a Helmholtz and anti-Helmholtz configuration using a custom H-bridge circuit (see Fig. C.1). The circuit is designed with an interlock which ensures that the H-bridge does not switch while there is a significant amount of current in the coils. We are also careful to optically isolate sensitive components of the circuit (and other electronics in the experiment) from the parts of the circuit that can experience large voltage spikes when the coils are switched.

[10] We refer to this system as the "Subharmonicon," for less obvious reasons.

[11] Each snapshot also includes a setting that allows one to ignore this phase, and simply roll over continuously based on the phase of the signal at the end of the previous snapshot.

[12] For example, a linear ramp between two settings of the amplitude and frequency involves only two snapshots, and the transfer of a few bytes of data (even if the duration of the ramp is very long in comparison to the period of the carrier).

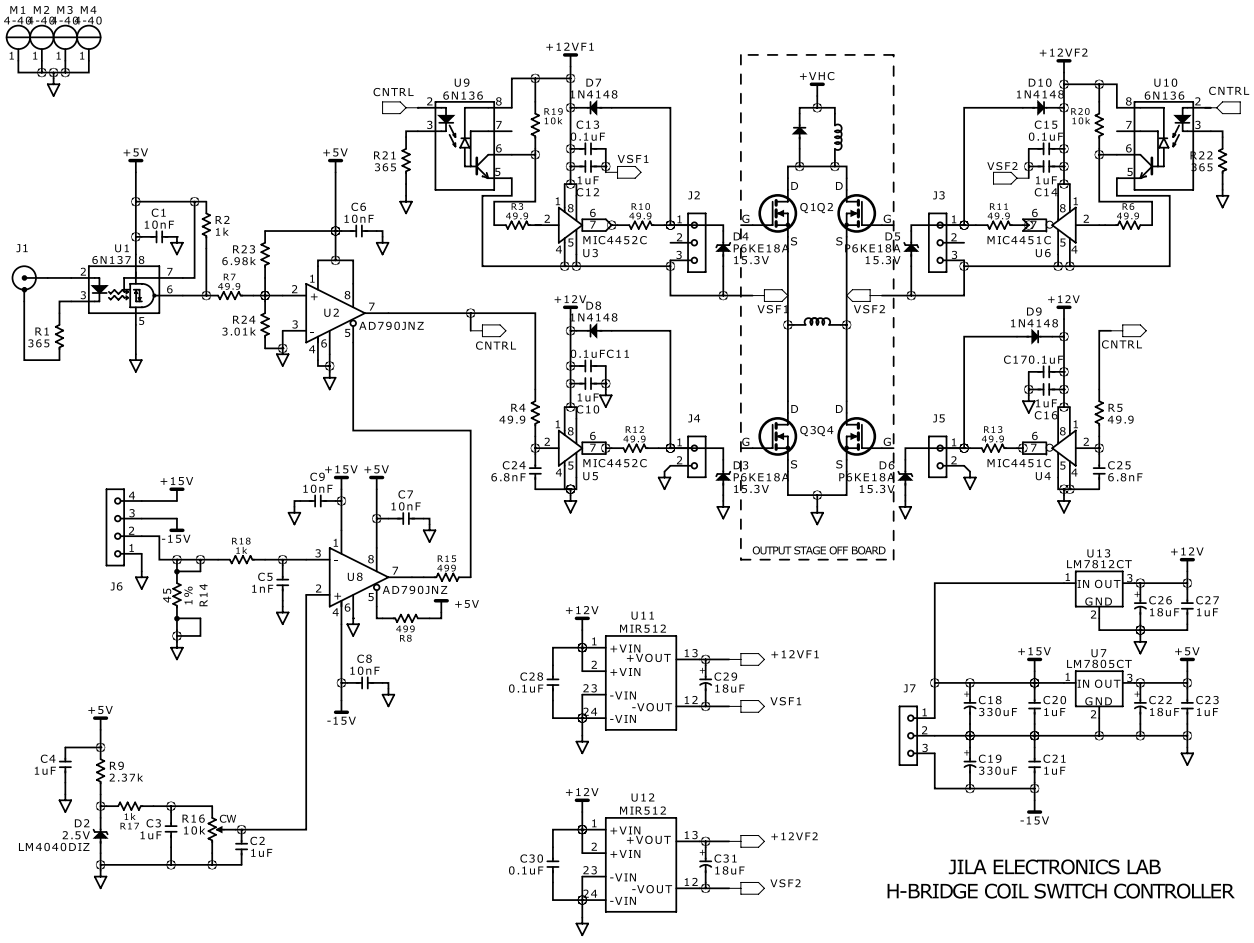

Figure C.1: Schematic for H-bridge coil switch. This circuit was designed by Carl Sauer in the JILA electronics shop.

Appendix D

Laser systems

As described in Sec. 2.1.1, we are interested in addressing a number of different electronic transitions in Sr using the appropriate lasers. Here, we briefly summarize the key laser systems in the lab.

D.1 Diode lasers

For applications involving low power ($\lesssim 100$ mW), we typically use homemade external cavity diode lasers (ECDLs). An ECDL uses a laser diode as the gain medium, and provides feedback using an external cavity. In order to isolate a single mode of the cavity, it is helpful to have an additional frequency selective element that is broader than the cavity linewidth, but narrower than the bandwidth of the gain medium. An optical grating is commonly used both as the intermediate frequency selective element, and to form the optical cavity [351]. We instead opt to use an interference filter[1] as the intermediate frequency selective element. The corresponding laser design is shown schematically in Fig. D.1a, and is sometimes referred to as an interference filter diode laser (IFDL). The main benefit of the IFDL design is that tuning the interference filter (via its angle) does not significantly change the optical length of the cavity, improving both the stability and tunability of the laser, as described in detail in [12].

Note that ECDLs are sensitive to optical feedback from reflections off of optical elements downstream of the laser.[2] To minimize the effect of these reflections, we typically use two stages of optical isolators directly after the laser output.

[1] The specific parts differ for different lasers, but we have had good luck with filters from Iridian Spectral Technologies and Semrock.

[2] In particular, coupling into fibers can often lead to reflections that are well-aligned to the laser diode, and form an additional cavity mode that can compete with the external cavity in the laser.

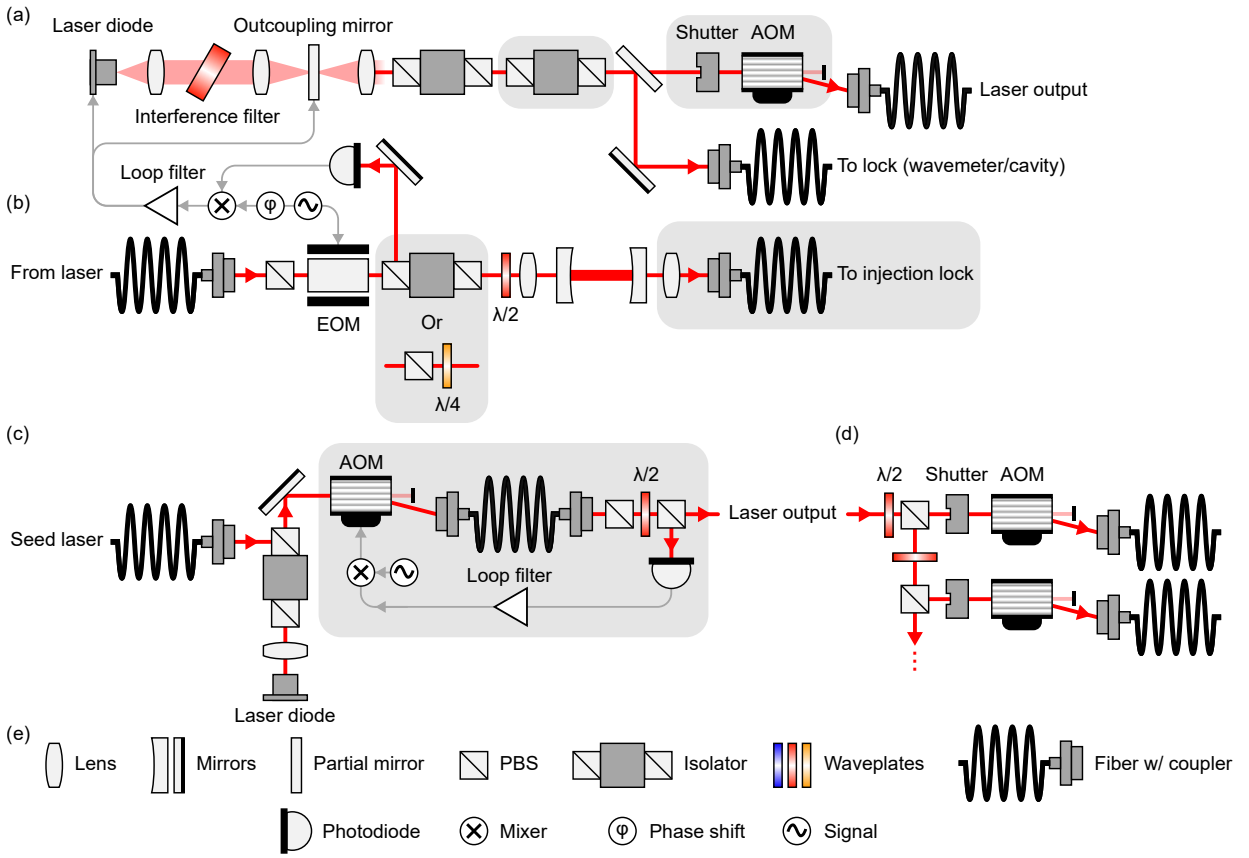

Figure D.1: Diode laser layout, stabilization, and distribution. Note that the depicted diagrams do not necessarily correspond to any specific diode laser system, and simply illustrate design elements that are commonly used in the various laser systems in the lab. Additionally, many optical elements (including those for mode shaping and alignment) are omitted for clarity. Throughout this figure, the grey shaded regions correspond to optical elements that are sometimes omitted, or replaced with a different design, depending on the specific laser system. (a) IFDL design, corresponding to a laser diode that is stabilized using an external cavity in a "cat's eye" configuration, with an intracavity interference filter used for coarse tuning of the dominant (highest quality) cavity mode. One or two optical isolators are used to prevent optical feedback from stray reflections downstream of the laser. Part of the laser output is used for monitoring or frequency stabilization, whereas the output to the atoms is switched using an AOM and optical shutter. (b) Some lasers are stabilized to a high finesse optical cavity using a conventional PDH locking scheme. In the case of the 689 nm laser system, the cavity is used to filter out unwanted laser noise at frequencies higher than the cavity linewidth. (c) When more laser power is required, we injection lock an additional laser diode using $\sim 0.5 - 1$ mW of seed power. In some systems, the output laser power is stabilized using the depicted scheme (in the grey region). (d) In cases where the laser light must be distributed to multiple places in the experiment, we use an optical switchyard similar to the depicted design. Each beam can be switched using an independent AOM, and fully extinguished using an optical shutter. (e) Legend for the optical and electronic components appearing in this figure. PBS refers to a polarizing beam splitter. Note that the figures in this section use graphics that are adapted from the ComponentLibrary package by Alexander Franzen.

D.1.1 Repump systems

Our IFDLs typically have a free-running instantaneous linewidth of $\sim$ 100 kHz (full width at half maximum, FWHM).[3] For repumping on the $\sim$ 1 MHz wide transitions at 679, 688, and 707 nm (see Sec. 2.1.1.4), no additional narrowing of the IFDLs are required. However, the center frequency of the lasers can drift, necessitating additional stabilization on long timescales. For our repump lasers, we simply monitor the output frequency of a given laser on a wavemeter (HighFinesse WS6-600), and adjust the laser diode current, and the cavity length (using a piezo-actuated mirror), a few times a second to maintain the appropriate output frequency.

The repumping beams are delivered to the experimental apparatus using optical fibers, and switched using acousto-optic modulators (AOMs) located before the fibers. Especially for experiments where the atoms are held in place for a long time, it is important to prevent even small amounts of leakage light from reaching the atoms. For this reason, we also incorporate optical shutters that fully extinguish the laser beams when they are not in use (see Fig. D.1a).

D.1.2 Cavity locking

For applications where the laser linewidth must be below $\sim$ 100 kHz, we perform additional narrowing of the laser output by locking it to a high finesse optical cavity (Stable Laser Systems SLS-6010-1-4bore). We use a conventional Pound-Drever-Hall (PDH) locking scheme [29], where the output of the laser is frequency modulated using an electro-optic modulator (EOM), and its reflection off of the optical cavity monitored (see Fig. D.1b).

D.1.3 689 nm system

Although PDH locking can substantially narrow the central spectral feature in the output of the laser (in our case to a few kilohertz), it can also result in added noise when the phase delay of the locking system approaches π (the resulting noise peaks are often referred to as "servo bumps"). In our laser system at 689 nm, which is used for optical cooling, these servo bumps are offset by $\sim$ 500 kHz from the carrier.

[3] Whereas we typically use cavity length of $\sim$ 5 cm, free running linewidths as low as 10 kHz have been achieved using a similar design with a 15 cm cavity in James Thompson's lab.

The servo bumps therefore add a significant amount of noise at frequencies close to multiples of the trap frequency (see Sec. 4.3.3), which could lead to reduced cooling performance by resonantly driving processes that can heat the atoms. To avoid such reductions in cooling performance, we opt to filter out the servo bumps by using light that is transmitted through the PDH locking cavity (Fig. D.1b).[4]

The cavity transmission is significantly attenuated relative to the full output power of the IFDL,[5] and must be amplified to achieve the optical powers required by the experiment. We perform this amplification by injection-locking an additional laser diode, as depicted in Fig. D.1c.[6] The optical power provided by the injection lock (as measured at the output of an optical fiber) is stabilized by actuating the laser with an AOM. The frequency- and power-stabilized laser light is then distributed to various parts of the apparatus using the scheme shown in Fig. D.1d.

D.1.4 461 nm system

Our 461 nm laser (for driving the $^1S_0 \leftrightarrow {}^1P_1$ transition) is a commercial laser system (Toptica DLC TA-SHG pro), which is composed of an ECDL operating at 922 nm that is amplified using a tapered amplifier, and doubled to 461 nm in a resonant doubling cavity. We use the scheme described in [239] to lock this laser to a see-through hollow cathode lamp (Spectrolamps HC054ST), as depicted in Fig. D.2. The scheme for power stabilization and distribution of this light is similar to what is shown in Fig. D.1cd.

D.2 Clock laser

Our clock (698 nm) laser system has undergone a number of revisions during the course of this work, depending on the specific experiment being performed. Typically, this laser is referenced to laser light stabilized to a cryogenic silicon reference cavity in the Ye lab at JILA [242]. This referencing has been done either by beatnote locking an independent laser to the reference laser, or by injection locking a laser diode using the stabilized laser light.[7] Here, we will describe a simplified setup used in many of our experiments

[4] The cavity bore used for stabilizing the 689 nm laser has a linewidth of ~ 70 kHz, which is significantly smaller than the offset of the servo bumps from the carrier.

[5] ~ 1 mW transmitted power, in comparison to ~ 15 mW directly out of the IFDL.

[6] In some cases, we use multiple stages of injection locks of a similar design.

[7] It is worth noting that injection locks are almost unreasonably effective in preserving the spectral properties of the seed light, and can maintain linewidths on the scale of tens of millihertz [218, 242].

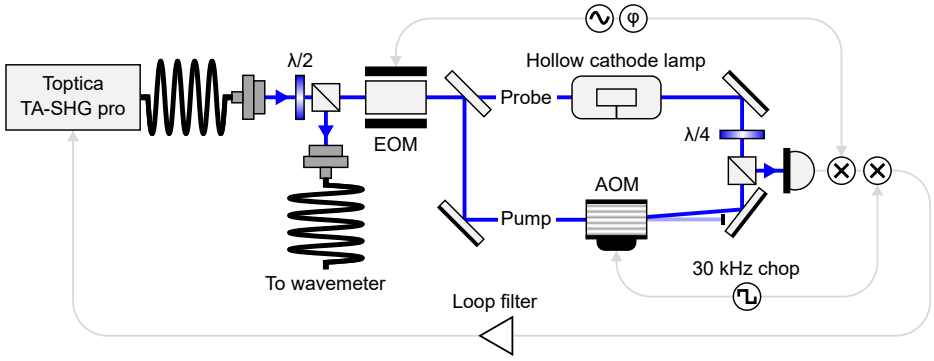

Figure D.2: Laser stabilization to a hollow cathode lamp. Doppler free spectroscopy of atoms in a hollow cathode lamp is used to stabilize the frequency of our 461 nm laser via the scheme described in [239]. Note that the output from the left-most fiber is not the full output of the TA-SHG pro system, and contains only ~ 5 mW of power. "30 kHz chop" refers a square wave at 30 kHz. The actual signal delivered to the AOM is a sinusoidal signal at 220 MHz that is modulated on and off with this square wave. This figure shares the legend presented in Fig. D.1e.

where we forgo any additional amplification, and simply deliver stable laser light from the Ye lab to the atoms in our lab (via a 50 m optical fiber) with minimal degradation of the spectral properties of the light.

D.2.1 Fiber noise cancellation

A dominant source of noise when delivering stabilized laser light to a distant location is fluctuations in the optical path length, leading to a modulation of the laser phase (and thus frequency). To counter this, one can measure the path length fluctuations interferometrically, and apply the appropriate correction using an AOM [205].[8] Three additional considerations inform the design of our path length stabilization: first, we wish to minimize the number of optical and RF components present in the Ye lab; second, we wish to avoid laser beat notes at uncontrolled frequencies in the Ye lab; and third, we must apply the appropriate frequency shift to the clock laser light to compensate for the different isotope shifts of the clock transition in ^{87}Sr (used in the Ye lab) and ^{88}Sr (used in our lab). To balance the above considerations, we operate in a feed-forward configuration as shown in Fig. D.3. One AOM applies a fixed frequency shift to a beam that is reflected off of a surface that is referenced to the atom positions (Sec. 3.3.1), generating the beat note used to measure fluctuations in optical path length. A second AOM is used to undo the measured fluctuations, to apply the requisite frequency offset to address ^{88}Sr atoms, and to stabilize the laser power delivered to the atoms. Note that we endeavor to keep the two laser beams (the reference beam and the beam going to the atoms) as close together as possible to avoid differential fluctuations in the optical path length.[9]

D.3 Rydberg laser

For driving transitions between the ^{3}P$_0$ clock state and different Rydberg states, we are interested in preparing high power and spectrally pure ultraviolet (UV) laser light near 317 nm. For a detailed description of this laser system, see [290] and the future thesis of William Eckner. Schematically, the UV laser system is composed of a pair of high power ($\sim$ 15 W and $\sim$ 10 W respectively) fiber lasers operating at 1066 nm and 1565 nm. These two lasers undergo frequency summing in a MgO:PPLN crystal, resulting in up to $\sim$ 9.5 W

[8] See [313] for a nice description of the path length stabilization up stream of our experiment in the Ye lab.

[9] Based on measurements with different lengths of uncanceled optical paths, we do not expect that differential fluctuations between the reference beam and the atom beam pose a significant limitation to our experiment [240].

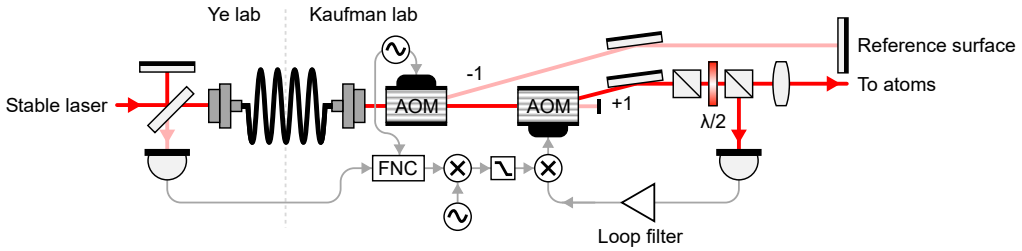

Figure D.3: Fiber noise cancellation. In order to deliver stabilized laser light to our lab (from the Ye lab) without degradation, we measure fluctuations in the optical path length experienced by the laser along a reference path, and apply the appropriate correction to the laser light delivered to the atoms. This setup additionally stabilizes the laser power delivered to the atoms, and allows us to scan the laser frequency. "FNC" refers to a JILA-designed electronic circuit for fiber noise cancellation, which locks a voltage-controlled oscillator (VCO) to the measured beat note, and divides the output frequency by 2. ± 1 denotes the diffraction order of the AOM. This figure shares the legend presented in Fig. D.1e.

($\sim$ 4.5 W typical) of optical power at 634 nm. 634 nm is a particularly convenient wavelength, since it can be locked to the same optical cavity described in Sec. D.1.2. Finally, the 634 nm light is doubled in a resonant doubling cavity using a BBO crystal, resulting in $\sim$ 1 W of optical power at 317 nm, and delivered to the atoms using similar intensity stabilization techniques to those described above.

D.4 Trapping lasers

In addition to resonantly (or near-resonantly) addressing various electronic transitions, we are also interested in optically trapping the atoms using far-detuned laser light at 515 nm and 813 nm (see Sec. 2.2.2). Our 515 nm laser source is a commercial fiber laser system (Azurlight Systems ALS-GR-514.85-10-A-SF), which delivers up to 10 W of optical power with a linewidth below 200 kHz. We find the output frequency of the 515 nm system to be sufficiently stable for our experiments without any additional stabilization. The output power is stabilized using the scheme shown in Fig. D.1c (see also Fig. 3.8). Our 813 nm laser system is a Ti:Sapphire laser (M Squared SolsTiS or Spectra Physics Matisse), which is stabilized to a wavemeter. We have also explored fiber-based laser sources at 813 nm [97], which could offer a higher power and more stable (but less tunable) solution.