Scalable Neutral-Atom Arrays and High-Fidelity Universal Quantum Gates on Metastable Qubits

Renhao Tao

SIGILL · UNIVERSITAT · LUDOVIC · MAXIM ·

München 2025

Scalable Neutral-Atom Arrays and High-Fidelity Universal Quantum Gates on Metastable Qubits

Renhao Tao

Dissertation
an der Fakultät für Physik
der Ludwig–Maximilians–Universität
München

vorgelegt von
Renhao Tao
aus Yangzhou, China

München, den 25. August 2025

Erstgutachter: Prof. Dr. Immanuel Bloch
Zweitgutachter: Prof. Dr. Christian Groß
Tag der mündlichen Prüfung: 7. Oktober 2025

Contents

List of Figures

List of Tables

Zusammenfassung

Jahrzehntelange Forschung in der Atom- und Molekülphysik hat Atome von reinen Untersuchungsobjekten zu leistungsstarken Werkzeugen für das Verständnis der Natur transformiert. Es wurden enorme Fortschritte darin erzielt, die quantenmechanische Natur von Atomen nutzbar zu machen, um hochkomplexe und korrelierte Materie zu untersuchen, was zu neuen Erkenntnissen in der fundamentalen Physik, Chemie und Materialwissenschaft geführt hat. Über die Beantwortung grundlegender Forschungsfragen hinaus wird das Gebiet auch durch Innovationen im Bereich des Quanteningenieurwesen vorangetrieben, die eine bislang unerreichte Kontrolle sowohl über die äußeren als auch die inneren Freiheitsgrade einzelner Atome ermöglichen. Diese Entwicklungen sind inzwischen so weit fortgeschritten, dass Forschende erstmals die Realisierung eines universellen Quantenprozessors in Betracht ziehen, der klassische Maschinen übertreffen kann. Zwei herausragende Herausforderungen bestehen weiterhin: Erstens sind Quantenoperationen, die elementaren Bausteine größerer, nutzbarer Quantenschaltkreise, von Natur aus fehleranfällig. Daraus folgt, dass ein Schaltkreis mehr Fehler als nützliche Information akkumulieren kann. Glücklicherweise bietet die Entwicklung der Quantenfehlerkorrektur (Quantum Error Correction) einen Weg, diese gewaltige Herausforderung zu umgehen. Der grundlegende Ansatz besteht darin, Quanteninformation durch Verschränkung über mehrere redundante physikalische Qubits zu delokalisieren. Solche Fehlertoleranzschemata erhöhen jedoch den Ressourcenbedarf erheblich, insbesondere in Bezug auf die Anzahl der auf der Hardware-Ebene benötigten Qubits. In metastabilen Zuständen kodierte Qubits bieten einen möglichen Ausweg, da sie Erasure Conversion unterstützen können. Dies erlaubt es, häufige Fehler zu markieren und korrigieren. Fehlerkorrektur kann dadurch effizienter durchgeführt werden, da solche Fehler an bekannten Orten mit bekannter Struktur auftreten. In dieser Arbeit demonstrieren wir erstmals universelle Quantenoperationen mit hoher Fidelität in einem Erasure-Qubit, das in zwei metastabilen Uhrenzuständen von Strontium kodiert ist. Um die Herausforderung der Skalierung zu adressieren, schlagen wir eine hybride Architektur vor, die eine unkomplizierte Skalierung auf große Systeme in einem optischen Gitter ermöglicht, während gleichzeitig die Einzelatom-Kontrollierbarkeit optischer Pinzetten erhalten bleibt. Wir erwarten, dass diese beiden zentralen Demonstrationen einen wertvollen Beitrag zum Fortschritt fehlertoleranter Quantencomputer leisten werden.

Abstract

Decades of research in AMO physics have transformed atoms from being solely subjects of study themselves into powerful tools for advancing our understanding of nature. Tremendous progress has been made in harnessing the quantum-mechanical nature of atoms to investigate highly complex and correlated matter, yielding insights into fundamental physics, chemistry, and materials science. Beyond addressing fundamental research questions, the field is also propelled by innovations in quantum engineering that enable unprecedented control over both the external and internal degrees of freedom of individual atoms. These developments are now significant enough that, for the first time, researchers have begun to contemplate the realization of a universal quantum processor capable of outperforming classical machines. Two outstanding challenges remain. First, quantum operations, which are the elementary building blocks of larger, utility-scale quantum circuits, are inherently error-prone. As a result, a circuit may accumulate more errors than useful information. Fortunately, the advent of quantum error correction offers a way to circumvent this formidable challenge. The basic approach is to delocalize quantum information across multiple redundant physical qubits via entanglement. However, such error-correction schemes significantly increase the resource overhead in terms of the number of required qubits at the hardware level. Qubits encoded in metastable states offer a potential remedy, as they support erasure conversion, which can transform common errors into heralded errors that are easier to correct. As a result, error correction can be performed more efficiently, since such errors occur at known locations and have a known structure. In this thesis, we demonstrate for the first time high-fidelity universal quantum operations in an erasure qubit encoded in two metastable clock states of strontium. To address the scaling challenge, we propose a hybrid architecture that enables straightforward scaling to large system sizes in an optical lattice, while retaining the single-site controllability of optical tweezers. We expect that these two key demonstrations will prove valuable for advancing fault-tolerant quantum computing.

Chapter 1

Introduction

1.1 Quantum science with neutral atoms

Quantum science is a field devoted to observing and understanding phenomena that are governed by the laws of quantum mechanics [1]. In its early days, the field focused primarily on uncovering the fundamental constituents of nature—first atoms, and later, subatomic particles. Tremendous progress has been made in understanding isolated quantum systems, laying the foundation for much of modern physics. Alongside these advances, there has been a parallel development of experimental techniques that has significantly expanded the capabilities of quantum physicists—so much so that entirely new areas of research have emerged as a direct result [2]. Among the most powerful tools now available are laser cooling and trapping of atoms, which, for the first time, allow physicists to engineer and manipulate quantum systems at the level of individual atoms [3, 4]. Rather than relying solely on naturally occurring systems, we can now design and construct highly artificial quantum systems to investigate fundamental phenomena under controlled conditions [5]. As the fidelity of quantum operations continues to improve, one naturally begins to ask about the practical relevance of these engineered systems beyond purely academic questions [6]. In the following, we outline four research areas that are actively being pursued under the umbrella of quantum science. These areas represent the key scientific frontiers that motivate this thesis work.

Quantum simulation

Quantum simulation involves using a well-controlled and accessible quantum system to emulate the behavior of another, typically more complex or less controllable, quantum system [7]. The concept was originally proposed by Feynman [8] and has since evolved into a distinct and rapidly advancing field [7]. At its core, quantum simulation seeks to address one of the central challenges in many-body quantum physics: the exponential growth of the Hilbert space with system size [9]. This exponential scaling leads to prohibitively high resource requirements, making classical simulation methods infeasible or outright impossible for many relevant problems. While resource-efficient classical techniques, such as

Monte Carlo sampling, have achieved notable success in reducing computational demands, there remain two broad classes of problems that are so far beyond the reach of classical algorithms [10]:

1. Fermionic systems or frustrated spin systems in two or more spatial dimensions, where Monte Carlo methods encounter the notorious sign problem, severely limiting their applicability [11].
2. Quantum dynamics, where time evolution explores a significant portion of the Hilbert space, leading to computational requirements that scale exponentially with time and system size.

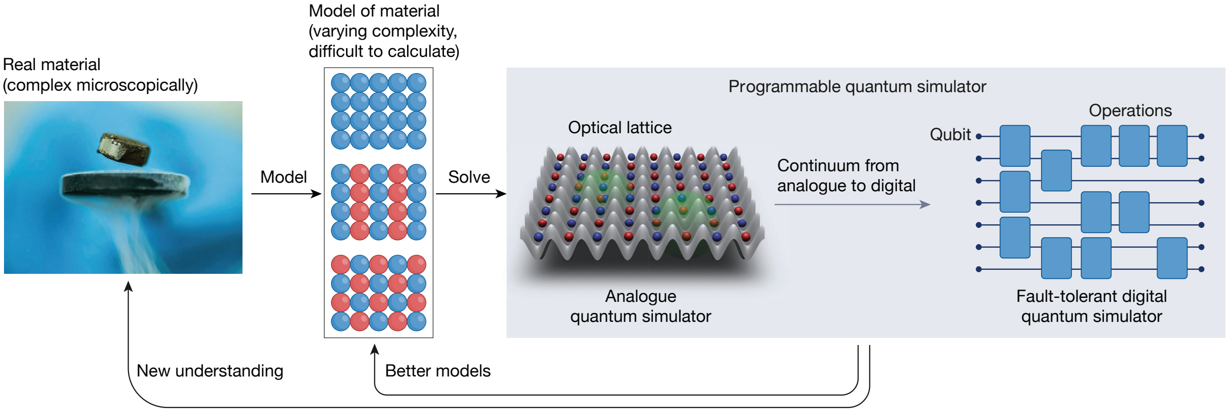

Figure 1.1: Quantum simulation enables the study of real materials using programmable atomic systems. The dynamics of electrons in solids can be captured by simplified theoretical models such as the Hubbard model, which incorporates key assumptions about hopping and interactions. These models can be realized experimentally with cold atoms interacting in optical lattices, allowing direct access to strongly correlated regimes that are difficult to treat theoretically. Insights gained from such analog simulations, in turn, help refine the underlying model assumptions and improve our understanding of quantum materials. Figure adapted from [12] with permission.

Quantum simulation circumvents the resource bottleneck by leveraging inherently quantum-mechanical systems to study complex quantum problems. Physical implementations include neutral atoms, trapped ions, superconducting qubits, and others; however, this thesis focuses exclusively on neutral-atom platforms. A particularly relevant example is the simulation of the fermionic Hubbard model, which is closely linked to the study of high-T_c superconductivity [13]. The general approach involves first distilling a problem of academic or practical relevance into a simplified model Hamiltonian that captures the essential physics. This Hamiltonian is then mapped onto controllable interactions, such as atom-atom collisions and light-matter couplings, that can be naturally realized using ultracold atoms in optical lattices [14]. By varying the parameters of the model Hamiltonian, which is often difficult in real materials but straightforward in the analog atomic system, one can experimentally explore the resulting phase diagram [15, 16]. These insights can then be translated back to the original physical system, offering a deeper understanding of the underlying quantum phenomena [17].

Existing quantum simulators have provided significant insights into a wide range of phenomena in both equilibrium and non-equilibrium quantum physics [7]. Indeed, quantum advantage has already been demonstrated in certain classes of tasks [18]. However, their utility remains fundamentally limited to models that can be implemented natively on the available hardware. In near-term devices, analog quantum simulators may be augmented with a limited set of digital operations, forming hybrid digital-analog platforms, particularly in light of recent advances in digital control [12, 19]. Another drawback of analog quantum simulators, in contrast to error-corrected digital quantum processors, lies in the presence of imperfections arising from calibration errors and uncontrolled noise [12]. Nevertheless, such imperfections are often less critical in practice, as the qualitative features of phase diagrams tend to be robust against local defects—similar to those found in real materials [20]. Furthermore, techniques such as Hamiltonian learning [21] can be employed to calibrate system parameters with high precision, thereby improving the reliability of analog simulations.

Quantum computing

Just as any classical computation can be constructed from a small set of logic gates, any arbitrary quantum operation can be decomposed into a sequence of elementary gates drawn from a universal quantum gate set [22]. In contrast to analog quantum simulation, digital quantum computation offers the potential to simulate a much broader class of problems that do not naturally map onto the native dynamics of available quantum hardware. Prominent examples include quantum chemistry, materials science, and combinatorial optimization [23, 24, 25]. Despite the promising outlook for quantum computing, current hardware remains far from achieving the fidelity levels necessary for solving practically relevant problems [26, 6]. The primary limiting factor is the susceptibility of coherent quantum operations to errors arising from noise and decoherence, which disrupt the delicate unitary evolution required for accurate quantum computation [27, 28].

Fortunately, theoretical breakthroughs in quantum error correction (QEC) have established a realistic pathway toward achieving universal, fault-tolerant quantum computation. QEC is strongly inspired by classical error correction codes, which mitigate errors by redundantly encoding logical information across multiple physical bits. In classical systems, if a bit is flipped due to noise, the original message can often be recovered through majority voting during the decoding process [30]. However, in the quantum domain, the direct application of classical error correction strategies is fundamentally constrained. Quantum information cannot be copied due to the no-cloning theorem [22], and measurement generally collapses the quantum state, making direct inspection for errors infeasible. Instead, QEC introduces redundancy by entangling data qubits with multiple ancilla qubits. By performing stabilizer measurements, one can extract error syndromes without disturbing the underlying encoded state, thereby enabling the detection and correction of errors while preserving coherence.

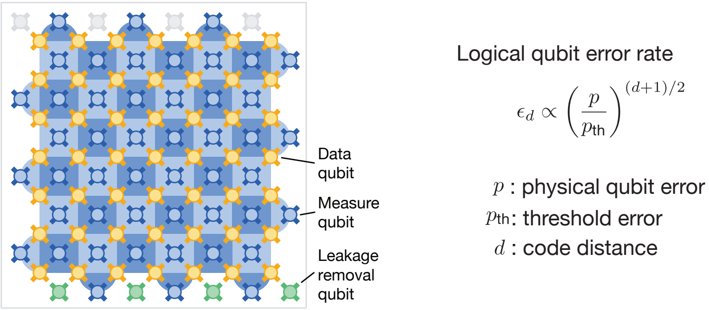

Figure 1.2: Quantum surface code and characteristic threshold behaviour in the logical qubit error. If the physical qubit error rate p is below a threshold p_{th} characteristic of the error correction code, the logical error rate ϵ_d can be exponentially suppressed by spreading quantum information among many qubits. The higher the number of physical qubits (characterized by code distance d) used in constructing a logical qubit, the lower the logical qubit error. Figure adapted with permission from [29].

In QEC, multiple physical qubits are grouped together to form a single logical qubit. If the physical qubit error rate p is below a certain threshold p_{th}, which is a property of the chosen error correction code, the logical error rate can be exponentially suppressed, approximately scaling as $(p/p_{\text{th}})^{(d+1)/2}$, where d is the code distance. Several experiments have already demonstrated below-threshold behavior [31, 32], validating the foundational principles of QEC. Neutral atom systems represent a unique opportunity, as they remain one of the most scalable quantum computing architectures [33]. In addition, the experimental demonstrations of native multi-qubit gates[34, 35], all-to-all connectivity[36], and transversal quantum gates [36] have further fueled optimism regarding the long-term potential of this technology. Substantial efforts are currently underway to both improve native gate fidelities and increase system size. Although a fully digital, universal quantum processor has not yet been realized, the toolkits developed in pursuit of this goal have already been successfully applied to other areas of quantum science, such as metrology and quantum simulation, yielding valuable insights.

Quantum metrology and sensing

Quantum metrology aims to establish a frequency standard by comparing the frequency of a local oscillator to that of a narrow atomic transition [37, 38]. Owing to the intrinsic precision of such frequency comparisons, quantum metrology also forms the basis of many precision sensing applications [39]. Neutral atoms offer several key advantages for metrological applications. First, atoms of a given species are fundamentally identical, allowing for distributed implementations of atomic clocks and their synchronization. A recent proposal to detect gravitational waves using a network of lattice-based optical atomic clocks exemplifies this potential [40]. Second, advances in narrow-linewidth lasers and state-insensitive optical trapping have significantly improved measurement precision, pushing the achievable resolution to the standard quantum limit (SQL).

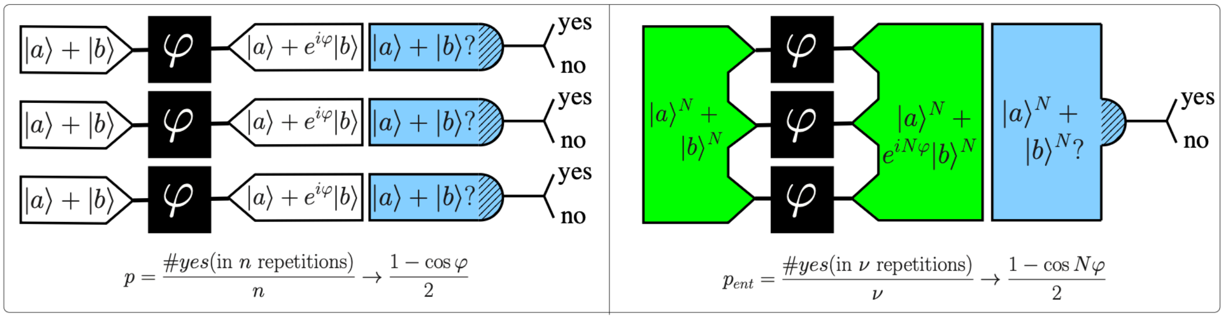

Figure 1.3: Quantum metrology with uncorrelated and correlated ensemble. In estimating an unknown phase φ, using uncorrelated atoms in interferometic measurement yields standard quantum limit scaling $n^{-1/2}$ of phase error to the number of experimental iterations n. In correlated atoms, the sensitivety is enhanced, resulting in an Heisenberg scaling of $(nN)^{-1/2}$ where N is the number of correlated atoms. Figured adapted from [37] with permission.

Recent research has increasingly blurred the boundary that once separated quantum metrology from other areas of quantum science. For example, to suppress various sources of technical noise and increase signal-to-noise ration, precision measurements are typically performed on cold ensembles of quantum-degenerate atoms confined in optical lattices [41]. This experimental configuration closely resembles that used to study the Fermi–Hubbard Hamiltonian. As a result, quantum many-body effects have become increasingly relevant to the atomic clock error budget and must be carefully accounted for. Conversely, the exceptional precision offered by atomic clock measurements has provided quantum many-body physicists with powerful tools for probing system dynamics and higher-order correlations with unprecedented sensitivity [42]. In parallel, a growing synergy has emerged between atomic clocks and universal quantum control [43, 44]. This circuit-based approach promises to surpass the SQL and enables novel ancilla-based readout schemes, opening new possibilities for precision measurement and quantum-enhanced sensing.

Quantum optics

Quantum optics is the study of light–matter interactions at the quantum level. Neutral atoms as a natural light emittor offer an almost ideal platform for such research, thanks to decades of advances in microscopic control of atoms' internal and external degree of freedom [45]. In turn, insights gained from these experiments often inform and enhance future experimental developments. One particularly opportunity in quantum optics is the study of correlated, collective emission processes in ordered atomic arrays [46, 47]. This underlies phenomena such as superradiance and subradiance, in which the emission of a photon by one atom induces a correlated response from the rest of the system [48, 49, 50]. In cold atoms experiment, atoms can be condensed into an ordered array either via collective cooling into a Mott insulator state [51], or atom-by-atom reassembly [52]. As a result, the system exhibits a strongly directional emission and near-perfect optical response. Recently, experiments have shown the possibility of exploting the non-local nature of the Rydberg interaction to switch the transmission and reflection of an atomic array [50], paving way to the engineering of new metasurfaces based on neutral atoms.

In a similar vein, dense atomic ensemble has also been explored in distributed quantum computing and quantum networking [53, 54]. The central idea is a conversion of flying, short-lived qubit into a stationary, atomic qubit for storage, by leaverging atomic nonlinearity. By reversing this process, quantum information is then retrieved and transferred back into a mobile qubit on demand. Second-scale coherent memory has been demonstrated in Rubidium atoms [55].

1.2 Experimental approaches of quantum science

The pursuit in experimental quantum science is driven as much by the profound physics question as by the experimental techniques that enable those discoveries. In the following, we review two well-stablished approaches toward quantum science and a newer trend that aims to combine them.

Top-down approach

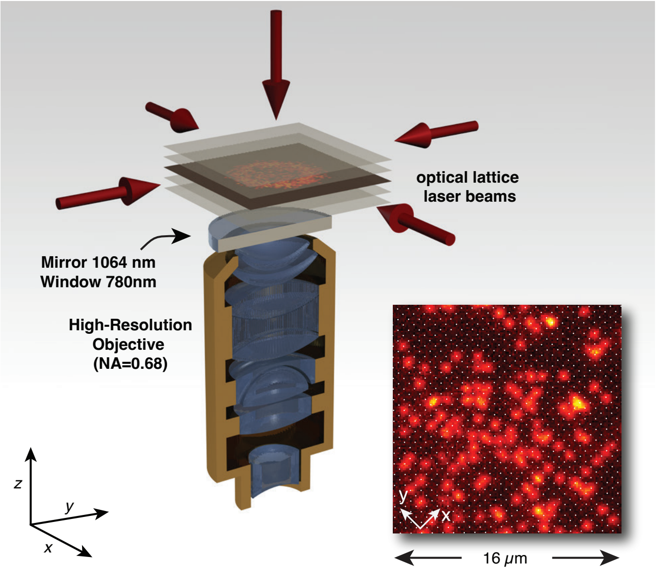

Figure 1.4: A simplified setup of a quantum gas microscope in the top-down engineering of a quantum system [56]. A three-dimensional optical lattice confines quantum gases within the field of view of the objective. Inset is a fluorescence image of atoms loaded in the lattice. Figured adapted with permission.

The top-down approach in quantum science is, to some extent, synonymous with the concept of the quantum gas microscope (QGM) [57, 56, 5]. In this framework, one typically prepares a Bose–Einstein condensate or a degenerate Fermi gas via evaporative cooling, adiabatically loads the gas into an optical lattice, and uses global control to study the many-body physics that emerges in the lattice [58, 16]. A defining feature of a QGM is its ability to detect, with high fidelity, a dense atomic ensemble in an optical lattice with single-site resolution [57, 56, 59]. The use of an optical lattice is essential not only for realizing itinerant many-body physics for which the QSM is tailored, but also for pinning atomic motion during fluorescence imaging, ensuring faithful assignment of fluorescence to its lattice site [60]. Naturally, when atoms densely occupy lattice sites with spacings near the diffraction limit of the imaging system, fluorescence signals from neighboring atoms can substantially overlap, making occupation determination nontrivial. However, by leveraging prior knowledge of the lattice geometry, one can apply deconvolution algorithms to reconstruct the fluorescence signal at each site. As a result, the full lattice occupation can be inferred with high accuracy. Since its inception, the QGM has enabled access to observables that were previously difficult or impossible to measure, such as entanglement entropy, string order, and magnetic polarons [61, 62, 63]. More recently, the same high-NA objective used for high-resolution fluorescence imaging has also been employed to project spatially resolved optical potentials onto the atomic ensemble [64]. In addition, the integration of phase masks and digital micromirror devices (DMDs) has significantly expanded the experimental toolbox, enabling programmable control over local potentials at the single-site level [65, 66, 67].

Bottom-up approach

The bottom-up approach in quantum science treats individual atoms as fundamental building blocks, with the goal of assembling quantum systems atom by atom [68, 69, 70]. The enabling technology is the optical tweezer, which consists of a tightly focused laser beam capable of trapping a single atom [71]. Depending on how the tweezers are generated, they exhibit different characteristics. Tweezers formed using spatial light modulators (SLMs) are highly programmable in their spatial configuration and serve as powerful tools for creating static trapping potentials [72, 73]. The relatively large spacing between tweezer sites ensures that atoms are effectively isolated, which significantly simplifies high-fidelity fluorescence detection. Indeed, site-resolved readout is routinely achieved using straightforward image thresholding techniques. Moreover, the tight confinement provided by optical tweezers enables efficient single-atom preparation and cooling protocols [74]. On the other hand, tweezers generated using acousto-optic deflectors (AODs) are highly dynamic, capable of being reconfigured at acoustic timescales [68, 69, 75]. They are particularly useful for real-time rearrangement of atoms, either within an array created entirely by the AOD itself [69] or in conjunction with a static array formed by an SLM [68]. Together, SLM- and AOD-based systems are complementary in terms of scalability and reconfigurability, offering a flexible platform for bottom-up assembly of quantum many-body systems. Coupling to Rydberg states enables the engineering of strong interatomic interactions, making neutral atoms in trapped optical tweezer arrays an ideal platform for investigating highly correlated spin systems relevant to quantum magnetism [76, 77, 70, 78].

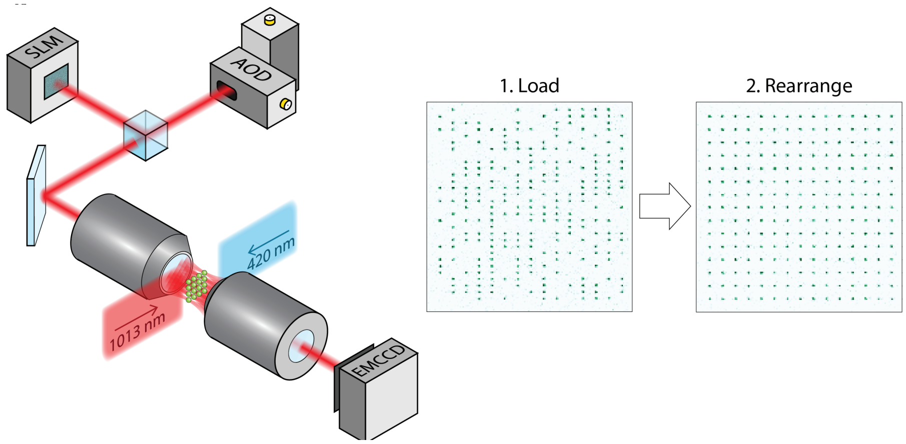

Figure 1.5: The statis and dynamic tweezer array for bottom-up assembly of quantum systems [78] (figure adapted with permission). The AOD tweezers perform real-time rearrangement of atoms stochastically loaded in a SLM tweezer array. Atomic vacancies are removed after the reassembly.

Hybrid approach

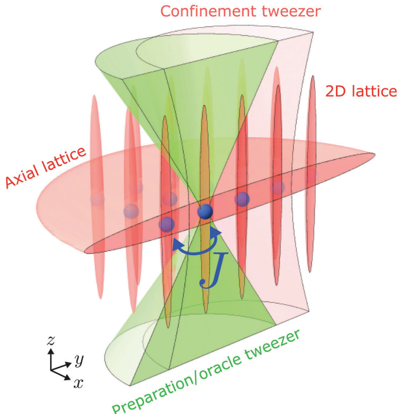

Figure 1.6: Optical tweezers combined with optical lattices in a hybrid platform for studying programmable Hubbard physics [79] (figure adapted with permission). The tweezer (green) acts as a spatially-dependent potential to pin the atoms which can otherwise freely hop in an 2-dimensional optical lattice (red ellipsoids aligned to the tweezer).

The top-down and bottom-up approaches in quantum science are not mutually exclusive. In fact, there is growing interest and utility in hybrid architectures that aim to combine the strengths of both frameworks [79, 80]. For example, the power-efficient pinning lattices used in quantum gas microscopes (QGMs) can be integrated into tweezer-based platforms to improve scalability—particularly during the detection stage. Additionally, the uniformity of lattice potentials enables the study of itinerant quantum physics, which remains challenging in optical tweezer arrays due to site-to-site disorder in trap depths [58]. Conversely, the dynamical control over atom positions afforded by optical tweezers can enhance QGM-based systems by increasing connectivity and enabling more flexible initial state preparation [81]. In principle, optical tweezers could even allow direct assembly of many-body ground states such as Mott insulators within an optical lattice, bypassing the conventional adiabatic loading procedures. If successful, such an approach would yield significantly faster experimental cycle times, which is especially advantageous for measuring statistics-intensive observables like correlation functions.

1.3 Outline of the thesis

This thesis is organized as follows:

Chapter 2 introduces key components of the experimental apparatus, including the ultra-high vacuum chamber, magnetic field coils, and a selected subset of laser systems. Particular attention is given to the ultraviolet (UV) laser systems, which play a critical role in the high-fidelity entangling gates presented in Chapter 4.

Chapter 3 benchmarks high-fidelity fluorescence imaging of atoms in a hybrid experimental platform. A central element of this work is the use of a narrow-line cooling transition in an alkaline-earth species, which enables resolution of the atoms' motional states within the trapping potential. After characterizing the imaging fidelity and survival probability on isolated atom arrays created with optical tweezers, the same cooling technique is applied to a larger trapping region formed by an optical lattice. We demonstrate scalable, site-resolved imaging across more than 10,000 lattice sites. Leveraging the novel lattice geometry, we also achieve loading of over 10,000 parity-projected atoms into a single plane of the lattice.

Chapter 4 presents the implementation of universal quantum gates on a novel qubit encoded in the metastable states of strontium. The advantages of this qubit encoding in the context of quantum information processing are discussed. We detail the state preparation, readout, and single-qubit rotations, along with the implementation of two-qubit entangling gates. In particular, the metastable-state encoding enables a novel state-resolved imaging technique that is sensitive to atom loss, providing a valuable tool for high-fidelity qubit measurement.

Chapter 5 provides an outlook on future experimental directions.

Chapter 2

Experimental apparatus

The experimental apparatus plays a critical role in experimental science: it provides sufficient isolation of the delicate quantum system from environmental noise so that the signal of interest is not drowned out. At the same time, the isolation must not impede the experimentalist's access or prevent precise control over every component of the system. Building on decades of scientific and engineering development, cold-atom experimental platforms have reached a stage of maturity, and a broadly standard design has emerged, even if, externally, individual experiments may appear quite different. Modern tabletop experiments in atomic, molecular, and optical (AMO) physics almost always begin with a high- or ultra-high-vacuum system to prevent collisions between the atoms or molecules and background gas particles. The vacuum chamber typically contains transparent windows to allow optical beams to enter and fluorescence signals from the atoms or molecules to exit. Such experiments often feature externally applied electric or magnetic fields to perturb the atomic or molecular energy levels for quantum engineering. Finally, trapping potentials created from laser beams are commonly used to confine the species under study for extended period of time.

Such a generic design is already sufficiently flexible and powerful to address many open questions in quantum science. Nevertheless, ongoing innovations are still aiming to reduce the experimental footprint, extend operational uptime, and improve overall stability. In this chapter, we present several selected components of our experimental apparatus: the ultra-high-vacuum system, magnetic field coils, laser systems, and optical tweezer generation systems. While not exhaustive, this discussion provides the necessary background for understanding the later sections of the thesis.

2.1 Ultra-high vacuum system

The ultra-high vacuum (UHV) system is functionally divided into three main sections: a cold, high-flux atomic source located at one end of the system; a fused silica science cell positioned at the opposite end; and a central pumping section that connects and maintains vacuum conditions between the two. The entire system is mounted on the translational stage (Lintech, 251242-WC0-2-S022-M00-C000-L00-E00-B00) for convenient retraction of the chamber assembly during for example optical beam pre-alignment near the glass cell.

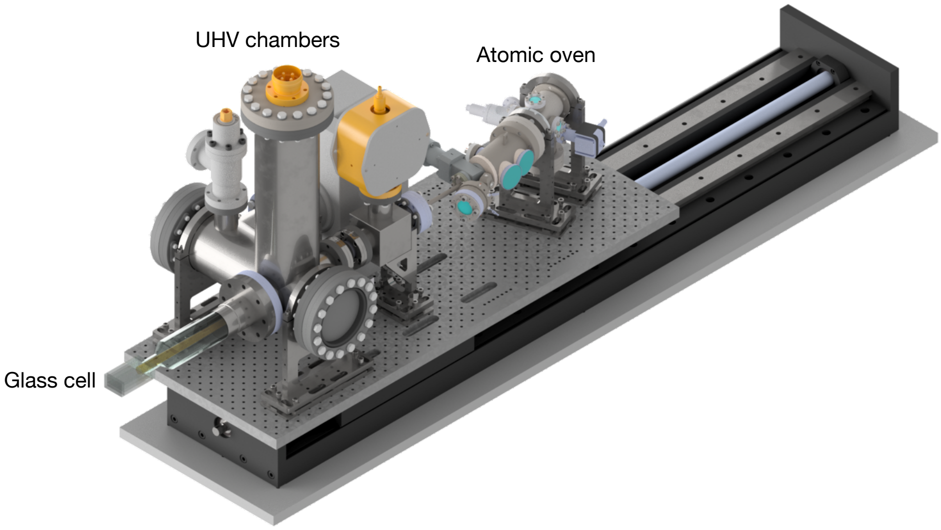

Figure 2.1: A CAD rendering of the ultra-high vacuum assembly used in the experiment.

The atomic source is a commercial unit (AOSense, Beam-RevC-Sr). In addition to the atomic oven, it includes in-vacuum permanent magnets for a Zeeman slower and integrated optics for a two-dimensional magneto-optical trap (2D MOT). This configuration enables the efficient production of a cold, collimated atomic beam with minimal setup, requiring only the application of laser beams for the Zeeman slower and 2D MOT. The 2D MOT deflects the atomic flux at an angle into the output tubing, ensuring that the glass science cell is not in direct line of sight with the oven. This arrangement significantly reduces background collisions from thermal atoms and prevents deposition on the inner surface of the cell. Furthermore, the atomic flux can be effectively shut off by disabling the 2D MOT beams, providing convenient optical control over atomic loading.

In cold atom experiments, the background pressure has a significant impact on the lifetime of atoms in a trap [82]. The primary loss mechanism arises from collisions between trapped atoms and energetic background gas particles, whose thermal energies typically exceed the trap depth by five or more orders of magnitude. Within the vacuum system, the highest pressure is found adjacent to the atomic effusion oven, due to its elevated temperature and resulting outgassing of various material in the vicinity. Despite the continuous operation of an ion getter pump, the pressure near the oven remains at approximately 3×10^{-8} mbar when the oven is held at its typical operating temperature of 400 °C. The purpose of the intermediate pumping stage bridging the atomic source and the glass cell is to create a pressure gradient, reducing the pressure to $\sim 1 \times 10^{-11}$ mbar in the glass cell.

To effectively reduce the pressure in the glass cell, a differential pumping stage is employed. A common approach involves introducing a vacuum component with low conductance to limit gas flow, followed by vacuum pumps to maintain a low downstream pressure. The design of such a system can be guided by exploiting the well-known analogy between vacuum systems and electrical circuits, where pressure corresponds to voltage, gas flow to current, and conductance to electrical conductance. By leveraging the extensive analytical tools developed for circuit analysis, one can quantitatively determine the appropriate parameters for the differential pumping configuration. Finally, special care must be taken to ensure that the conductance-limiting element does not excessively impede the atomic flux, which must still be able to traverse the differential tubing and reach the science chamber.

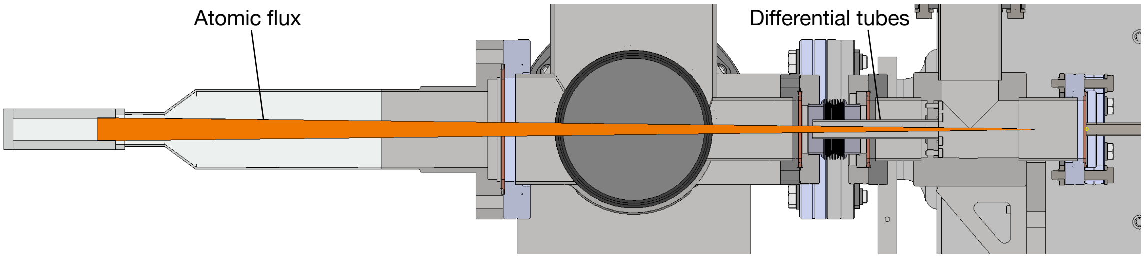

Three different pumps are used for capture the broad range of background gas species. They are a Titanium sublimination (TiSub) pump, a 10 L/s ion pump, and a 50 L/s ion pump. The TiSub pump is essentially a large diameter tube with Titanium sputtered on its inner surface, and is effective against chemically active compounds like O_2 and H_2. The pumping rate can be extremely high due to the large surface area where Titanium is sputtered.

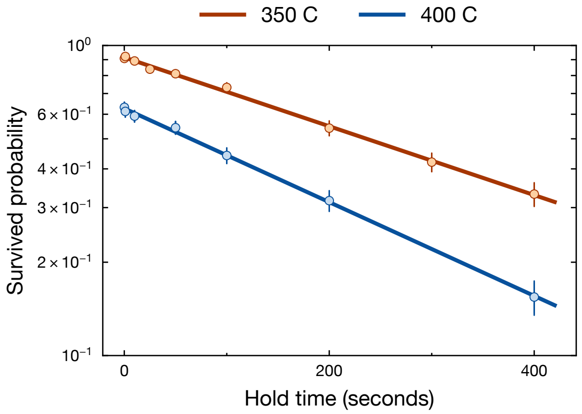

Figure 2.2: The lifetime of trapped atoms under continuous attractive Sisyphus cooling is primarily limited by background gas pressure, as evidenced by its dependence on the oven temperature. Measured $1/e$ lifetimes are 287 s at 400 °C and 391 s at 350 °C, respectively.

Although the ion getter pump can roughly measure the pressure level, the atom lifetime in the trap remains the most direct and relevant benchmark of the vacuum quality. In the benchmarking sequence, atoms are held in an optical lattice under continuous cooling. Under normal conditions, collisions with energetic background gas particles constitute the dominant loss mechanism. We observe that the atom lifetime in the trap depends on the oven temperature (Fig. 2.2), indicating that the lifetime is limited by the vacuum conditions.

2.2 Arbitrary magnetic field generation

The magnetic field is a crucial control parameter for experimentalists. Practical implementations span from field gradients employed in magneto-optical trapping to spatially-uniform fields utilized for both broadening the effective linewidth of the clock transition and establishing a well-defined quantization axis. Our objective is to design a magnetic field generation system satisfying a diverse (and at times seemingly conflicting) set of requirements.

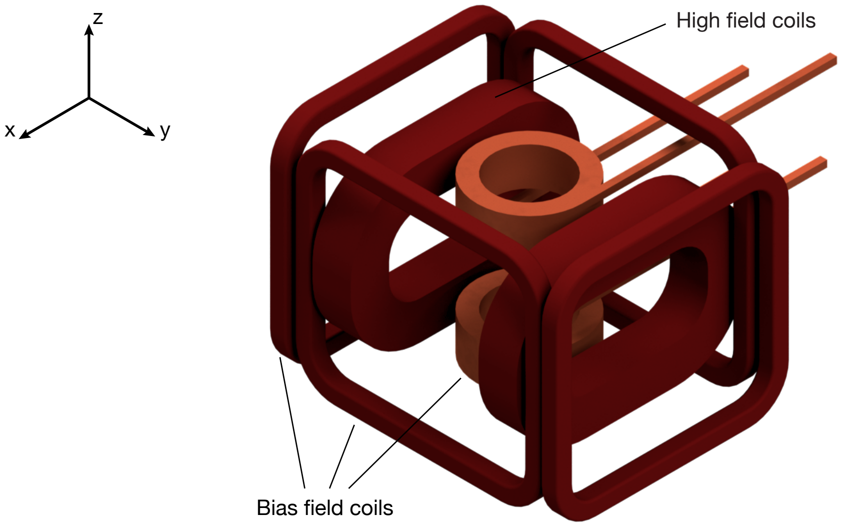

Figure 2.3: A CAD rendering of coil assembly.

The magnetic field is generated by running current through electromagnetic coils. All coils are made of the same type of hollow copper wires, protected on the exterior by an insulating Kapton film. The thermal load dissipated in the wires when the coils are in operation is carried away by water circulating in the wires. The flow and temperature of the water are regulated by a chiller. The coil mounts are made of Tecaform—a material that is insulating enough to prevent the buildup of eddy current in a dynamical magnetic field environment, and conductive enough to dissipate charges in µs. The latter feature which is beneficial to Rydberg atoms sensitive to stray electric field together with the material's superior UV resistance makes it ideal in Rydberg experiments.

2.2.1 High-field coil

Helmholtz configuration

The high-field coils in Helmholtz configurations have currents propagating in the same direction. It produces a nearly uniform magnetic field close to the geometric center of the coils. The high-field coils are designed to provide up to 400 G magnetic field along one axis. Such a strong field is used in enhancing the $^1S_0 \rightarrow {}^3P_0$ coupling strength. Finite-element simulation gives a field conversion factor of 1.625 G/A. Experimentally, the magnetic field strength is calibrated via the probe shift of the $^1\mathrm{S}_0 \to {}^3\mathrm{P}_0$ clock transition as a function of applied current, and the result is 1.45 G/A. The difference in the conversion factor between the numerics and experiments could be attributed to electrical shorts at some coil winding.

Anti-Helmholtz configuration

In the anti-Helmholtz configuration, currents flow in opposite directions in two coils. This gives rise to a gradient of magnetic field close to the geometric center of the coils. So far, magnetic field gradients are only used during the loading of magneto-optical traps. Numerically, the gradients are 0.24 G/cm/A, 0.395 G/cm/A, 0.13 G/cm/A along x, y, and z-axis. Although there has not been a careful calibration of the gradient strength using atomic signals, reduced values along all three axes by the amount by which the field in Helmholtz configuration is attenuated are expected.

In the following table, we summarize the field conversion factor for the strong-field coil in both Helmholtz and anti-Helmholtz configurations:

	Axis	Field conversion factor
Helmholtz	y	1.45 G/A
anti-Helmholtz	x	0.216 G/cm/A
	y	0.352 G/cm/A
	z	0.116 G/cm/A

Table 2.1: Magnetic field conversion factors of strong-field coils in Helmholtz and anti-Helmholtz configurations.

2.2.2 Low-field bias coils

The low-field bias coils consist of three mutually-orthogonal pairs of coils in Helmholtz configuration and can create a magnetic field up to ∼20 G in arbitrary direction. Common uses of these coils include nulling the stray magnetic field and providing the quantization axis. To obtain the conversion factor of bias coils, it is convenient to the 2.1 MHz/G Zeeman splitting between $\sigma^{\pm}$ and π transition on the $^1\mathrm{S}_0$-$^3\mathrm{P}_{1,m_J}$ line. The measured conversion factors are 0.226 G/A, 0.136 G/A, and 0.697 G/A along flux, orthogonal, and vertical axis.

Stray field compensation

One of the earliest applications of low-field bias coils is in zeroing the background magnetic field. The stray field can originate from a variety of sources such as the Earth or the ion pumps in the vicinity. In the regime in which the Zeeman splitting from the background field is well-resolved, one could vary the applied magnetic field along some (say, z) axis, and fit the Zeeman shift using $\Delta f = \mu\sqrt{B_{0x}^2 + B_{0y}^2 + (B_{0z} + \alpha_z I_z)^2}/h$. Here, $\mathbf{B}_0 = B_{0x}\hat{x} +$

$B_{0y}\hat{y} + B_{0z}\hat{z}$ is the stray magnetic field, α_z the conversion factor, I_z the applied current along the x-axis, and h Planck constant. Depending on whether the applied field is anti-parallel or parallel to the stray field along the z-axis, the Zeeman shift either has a dip before rising or rises immediately. In the former case, the compensation field is then $\alpha_z I_z^{\text{dip}}$ where I_z^{dip} is the current of the dip. In the latter case, the direction of the current needs to be flipped to counteract the stray field. This completes the compensation along z-axis. The same technique can be subsequently applied to the other two axes.

Axis	Field conversion factor
x	0.226 G/A
y	0.136 G/A
z	0.697 G/A

Table 2.2: Magnetic field conversion factors of low-field bias coils.

Coil control electronics

To switch between the Helmholtz and the anti-Helmholtz configuration for the coils, we use an H-bridge circuit, see Fig. 2.4a. This circuit consists of 4 pairs of switches (Q_1-Q_4) and catch diodes (D_1-D_4). At any time, only two switches on opposite sides are on to route the current in a definite direction, see Fig. 2.4b, and c. The catch diodes dissipate inductive current during switching to protect the circuit. In our implementation of the H-bridge, the switches are IGBTs primarily for their superior handling of higher voltage in fast switching applications ($U = LdI/dt$) compared to alternative such as MOSFET.

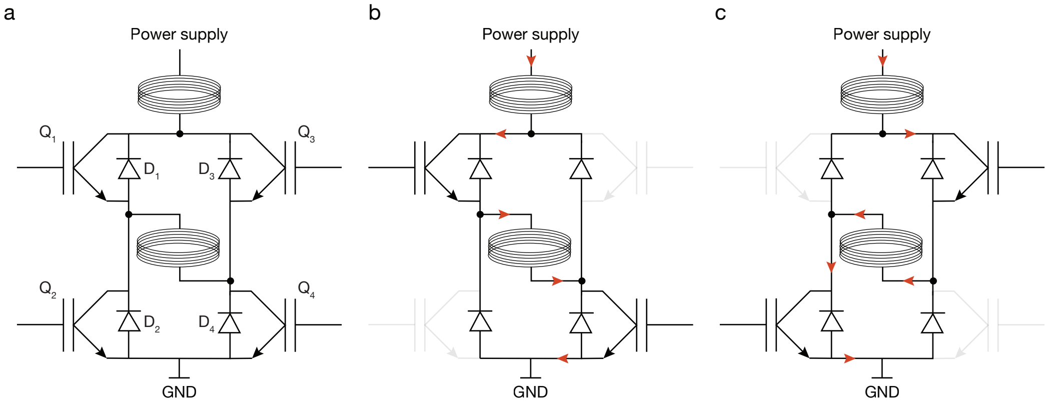

Figure 2.4: **(a)** An H-bridge circuit for switching current polarity. The circuit configured to drive **(b)** Helmholtz coils and **(c)** anti-Helmholtz coils.

To control the magnetic field, one first issues a command, specifying the polarity, magnitude, and switching timing of the current. The control system then parses the command,

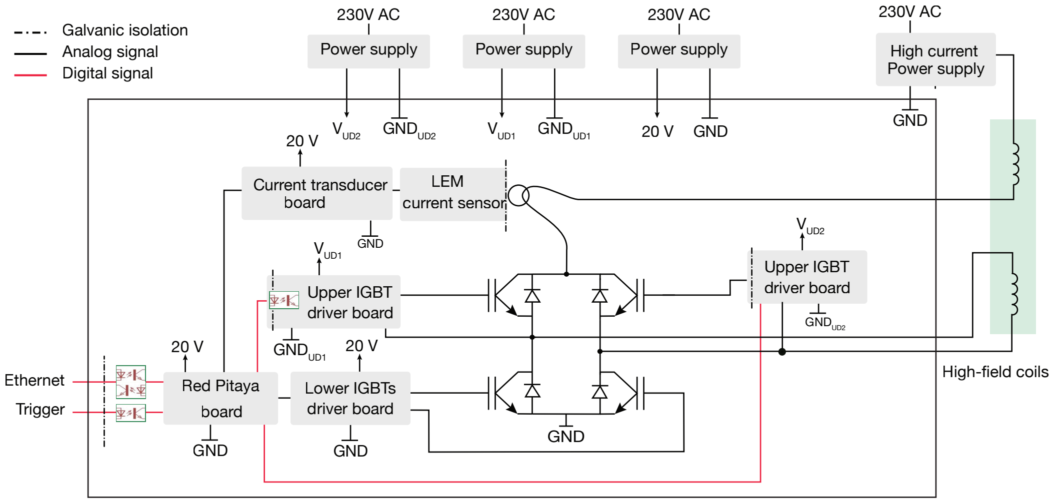

Figure 2.5: Circuit diagram of electronics controlling the coil power supplies. The polarity and magnitude of current are programmed into Red Pitaya via the pulse-width encoding scheme. The FPGA on the board parses the command, configure the H-bridge, and set the current through the coils. A current transducer (LEM it 400-s ultrastab) measures the current circulating in the coils. The rectangle denotes the circuit physically enclosed in a metallic box.

and passes the encoded message via the pulse-width encoding scheme to a microcontroller (Red Pitaya) board, which sets the H-bridge gates and current supplied to the coils. Benefited from the timing determinism in pulse generations of the control system (Adwin), this pulses-based communication protocol fully sidesteps the latency issue frequently encountered in network-based communication. In fact, the measured delay between the switching and the trigger generated from the control system is reproducibly 20 µs. Another feature of the circuit is the use of multiple digital communication layers in the control path, which provide galvanic isolation and improve the noise immunity of digital signaling.

2.3 Laser systems

As the primary tools for atomic manipulation, lasers are among the most essential components in cold atom research labs. In our experiment, we use lasers to engineer trapping potentials and to manipulate the electronic states of atoms. The sheer number of lasers makes a detailed description of each unit unrealistic (see Table 2.3). Therefore, we elaborate on only three representative lasers due to their particular importance in the experiment.

Table 2.3: Lasers used in our lab and their primary applications

Wavelength (nm)	**Laser Type**	**Application**
461	SHG	MOT, imaging, and pushout on $^1S_0 \rightarrow ^1P_1$
689	ECDL	Narrow-line cooling on $^1S_0 \rightarrow ^3P_1$
707	ECDL	$^3P_2$ repump, FS Raman transition
679	ECDL	$^3P_0$ repump, FS Raman transition
688	ECDL	$^3P_1$ repump
716	ECDL	$4d\,^1D_2 \rightarrow 6p\,^1P_1$ repump
496	ECDL	$5p\,^3P_2 \rightarrow 5d\,^3D_2$ repump
698	ECDL	$^1S_0 \rightarrow ^3P_0$ clock transition
671	ECDL	$^1S_0 \rightarrow ^3P_2$ clock transition
520	SHG	$^1S_0 \rightarrow ^3P_1$ magic trapping
813	Fiber lasers + SFG	$^1S_0 \rightarrow ^3P_0$ magic trapping
689.22	TiSa	$^1S_0$ tune-out
1040	Fiber laser	$^1S_0$ trapping in optical lattices
316	SFG + SHG	$5p\,^3P_0 \rightarrow ns\,^3S_1$ Rydberg
322	SFG + SHG	$5p\,^3P_2 \rightarrow ns\,^3S_1$ Rydberg
408	ECDL	Rydberg auto-ionization

2.3.1 The 461 nm laser system

Of all the laser systems ever used in our lab, none has consumed more resources and energy than the 461 nm blue laser system. As the laser responsible for driving the broad $^1S_0 \rightarrow {}^1P_1$ transition, it plays an all-important role in magneto-optical trapping, fluorescence imaging, and resonant pushout. In the early days of our lab, we experimented extensively with the 461 nm ECDL with an injection-locked amplifier techology. While simple and compact, this setup was extremely challenging to maintain. Meanwhile, poor choices in the optical fibers used for power distribution resulted in significant power fluctuations, and exacerbated the experiment's overall instability. After several catastrophic laser failures, we abandoned the old units and switched to a 922/461 second-harmonic generation (SHG) system. The new system substantially outperformed its predecessor. It still requires occasional maintenance, such as replacing the desiccant pack above the SHG cavity for humidity control and optimizing the coupling of seed light into the TA. However, over time, maintenance of the blue laser system became less frequent and considerably less stressful. We also learnt to use end-cap fibers to effectively reduce the optical intensity at the tip, preventing damage to the epoxy adhesive. As a result, optical fibers ended up lasting much longer.

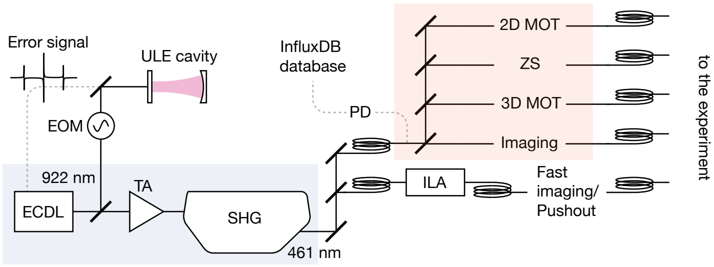

Figure 2.6: A simplified diagram of the 461 nm blue laser system. The majority of the blue power is sourced from a commercial 922/461 second-harmonic generation (SHG) system (shaded blue). The SHG system includes a 922 nm seed ECDL, a tapered amplifier (TA), and a SHG cavity converting 922 nm light to 461 nm light. The blue power is further distributed via optical fibers on a standalone breadboard (shaded orange) to the 2D MOT, Zeeman slower (ZS), 3D MOT, and imaging beam paths. An injection-locked amplifier (ILA) is dedicated to powering the fast imaging and pushout unit. The 922 nm seed light is locked to a high-finesse cavity via an offset lock. To monitor the optical power, a small fraction is sampled after the fiber and measured. The power readings are then logged to InfluxDB for long-term tracking.

The 922/461 SHG system is shown in Fig. 2.6. The laser itself outputs about 1 W 461 nm light for the experiment, and a few mW 922 nm light for the offset locking to an ULE cavity. The blue light is coupled to standalone power distribution breadbords, which house additional downstreams optics, via end-cap optical fibers. In particular, a few hundred mW is shared between 2D MOT, Zeeman slower, 3D MOT, and resonant imaging, while a few tens of mW is used to injection-lock a high-power laser diode for fast imaging/resonant pushout. All breadboards containing mechanical shutters are mounted on Sorbothane feet to prevent the acoustic noise during shuttering from destabilizing nearby lasers. In addition, they are fitted with black hardboard enclosures to block any stray light.

2.3.2 The 698 nm clock laser system

The 698 nm laser addresses the ultra-narrow $^1S_0 \rightarrow {}^3P_0$ transition. The laser system's power unit consists of an ECDL seed and an injection-locked amplifier (ILA), see Fig. 2.7. To reduce the laser linewidth, the seed light is locked to a ULE cavity with a finesse of 200,000. Owing to the aging of the ULE spacer, the cavity resonance which serves as the frequency reference for the clock laser drifts in an almost linear fashion (Fig. 2.7b). To correct this drift, we compare the clock light against a frequency comb synced to a hydrogen maser. The resulting beat frequency is then used to compensate for the drift through feedback on the AOM and keep the clock laser frequency on the atoms absolutely stable. Notably absent is fiber-noise cancellation over the long delivery fiber to the experiment. This limitation is a consequence of placing the pulsing AOM before the fiber, since fiber cancellation can only be engaged during the time when a pulse is present. In practice, pulses we use in experiments are almost always too short for the noise cancellation to respond effectively.

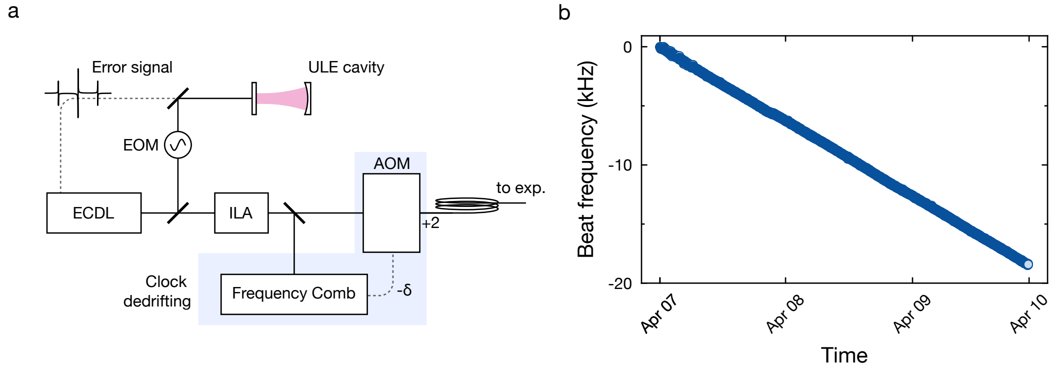

Figure 2.7: **a.** The essential components in the 698 nm clock laser setup. The seed light sourced from an ECDL is locked to a ULE cavity with finesse 200,000. An injection-locked amplifier (ILA) amplifies the optical power to about 100 mW. Some light is picked up and compared against a frequency comb synced to a Hydrogen maser for compensation of slow cavity drift. The frequency difference is then applied to a double-pass AOM to keep the frequency of the light on the experimental side constant. **b.** An example beat frequency drift originating from ULE spacer aging. The drifts is about a few Hertz per minute and largely linear if the laser is unperturbed.

2.3.3 The 316 nm UV system

The 316 nm UV laser system is solely reserved for Rydberg physics. Specifically, it targets the $^3P_0 \rightarrow n = 47\,{}^3S_1$ transition. The laser system sourced from NKT Photonics contains a DFB fiber laser at 1550 nm, optical amplifiers at 1550 nm and 1071 nm, a sum-frequency generation unit at 634 nm, and a second-harmonic generation unit at 316 nm. Although

DFB fiber lasers at 1071 nm with low phase noise are available, we purposefully supplement the UV laser system with our seed light sourced from a Toptica ECDL to obtain a broader tunability in UV wavelength. At the optimal setting, the laser output power can reach 1.4 W at 316 nm. However, it is subject to a slow decline at a rate of 0.8 mW/h under continuous-wave operation [83].

To stabilize the UV frequency, we lock the 634 nm light to a high-finesse ($\mathcal{F} = 45000$) cavity through offset sideband locking [84]. The error signal is derived from comparing one of the tunable sidebands of 634 nm to the nearest cavity resonance and applied to the current modulation and piezo of the 1071 nm ECDL to close the loop. Such a scheme conveniently ignores the common-mode frequency drift in both seed lights which does not shift the UV frequency and greatly simplifies the setup.

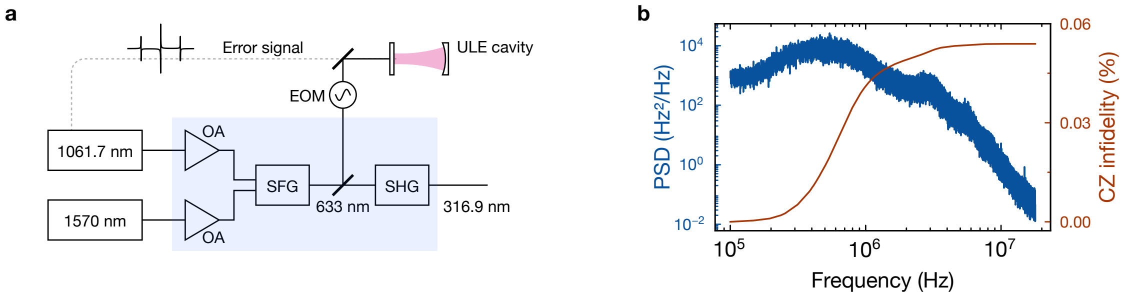

Figure 2.8: **(a)** A simplified diagram of the UV laser system and its frequency stabilization scheme. A small fraction of 634 nm light is picked off from the laser and modulated via a fiber-based electro-optic modulator (EOM) for offset sideband locks to a cavity. The error signal passes through a PID filter before being applied to the 1071 nm ECDL. **(b)** The frequency noise power spectral density measured using the 633 nm light and converted to that in UV. The cumulated CZ infidelity (red) for $2\pi \times 6$ MHz Rabi is estimated from frequency noise using linear repsonse theory [85].

The UV optical system incorporates several features critical for high-fidelity entanglement generation, see Fig. 2.9.

Pulse shaping

For advanced pulse shaping, two acousto-optic modulators (AOMs) are used in tandem. AOM 1 (Gooch & Housego, 1-M110-3C10BB-3-GH2) operates in a single-pass configuration and is used for intensity stabilization only. AOM 2 (Isomer M1212-aQ150-2-M) is in a double-pass configuration dedicated for fast amplitude/phase modulation. It is driven by a arbitrary waveform generator (Spectrum Instrument, M4i6631-x8) with an external RF amplification stage. While the single-pass setup is straightforward, the double-pass configuration is considerably more complex. Conventional double-pass setup isolates the twice-diffracted light from the undiffracted beam via polarizations [86], which requires the light's polarization in two AOM passes to be mutually orthogonal. In the UV range, this

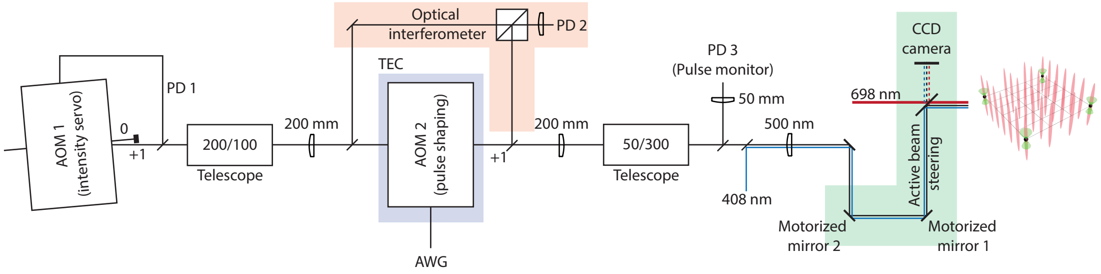

Figure 2.9: A simplified UV beam/pulse shaping system. The acousto-optic modulator (AOM) 1 is in the single-pass configuration for continuous-wave intensity noise suppression along with photo-detector (PD) 1. AOM 2 is in the double-pass configuration for fast UV phase and amplitude modulation. An optical interferometer (shaded red) beats the light after AOM 2 with that before it to measure the optical phase. PD 3 is dedicated to monitoring the pulse intensity and is located at a plane image to the AOM 2. The beam shaping consists of several telescopes with (de)magnification ratio indicated. The optics are designed to image the AOM 2 both to the atomic plane, and to the intensity monitoring PD 3, and a CCD camera monitor. The UV beam pointing is actively regulated in a servo (shaded green) consisting of two motorized mirrors and a CCD camera. The 698 nm and 408 nm light are overlapped onto the UV via a dichroic mirror (FILTROP Dichroic Beam Splitter 430 nm). The thermoelectric cooler (TEC) element shown in shaded blue regulates the temperature on the baseplate of AOM 2 for improved beam pointing and intensity stability.

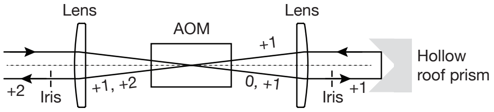

Figure 2.10: A double-pass AOM setup based on a hollow roof prism. The undiffracted beam and twice-diffracted beam are separated spatially while maintaining a fixed polarization state throughout the two passes. The number on the beam denotes the diffraction order with respect to the incoming beam. Two irises block the undiffracted beam on each pass. The beam pointing of the +2 order in principle stays the same independent of the frequency modulation of the AOM.

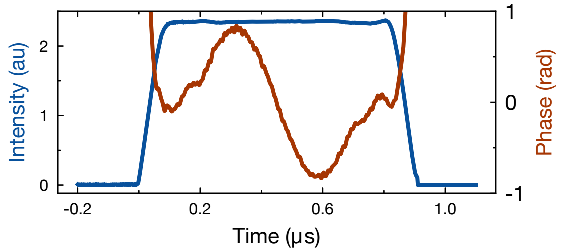

Figure 2.11: A representative feedback-optimized UV control waveform for implementing a time-optimal CZ gate.

approach is less advisable, as the quartz crystal which is almost exclusively used at short wavelengths is a strongly uniaxial birefringent material. As a result, using orthogonal polarization will lead to substantial optical power loss. To maintain a fixed polarization state in the double-pass, one can employ a roof prism (Thorlabs, HR1015-F01) as a reflector and spatially seperate out the twice-diffracted light, see Fig. 2.10. An alternative method, which is not used in this work, is to use two AOMs in single-pass configuration. By imaging one AOM onto another, the beam pointing remains unchanged even as the shared RF drive signal is modulated.

In quantum optimal control, quantum unitaries like a controlled-Z gate are usually implemented via phase modulation [87], see Fig. 2.11. To measure and, if necessary, actively regulate the waveform, it is crucial to have access to both the phase and amplitude of the optical beam. While the optical amplitude can be directly measured on a photodetector (PD 3, Thorlabs, PDA10A2), phase measurement requires an optical interferometer, see Fig. 2.9. To acquire the phase, we perform digital I/Q demodulation on the beat signal recorded by PD 2 (Alphalas, UPD-200-UP). With a high-bandwidth oscilloscope, this digital approach performs considerably better than the analog techniques. In practice, the obtained waveforms always deviate from the programmed ones due to various deleterious effects in the phase modulators. Nevertheless, one can apply feedback to the waveform and smoothen out the irregularities, regardless of their origin [88]. See Fig. 2.11a for an example of a regulated CZ gate waveform.

To implement fast quantum gates, the optical modulator must exhibit a short response time—ideally much shorter than the target gate time. In the case of a single-pass acousto-optic modulator (AOM), a common strategy is to reduce the optical beam diameter within the AOM to minimize the transit time of the acoustic wave across the beam. Although certain trade-offs must be considered in the design, particularly in balancing diffraction efficiency and the beam profile of the diffracted light against response time, pulse shaping with a single-pass AOM is generally straightforward. In contrast, additional complications can arise when using a double-pass AOM configuration. In the following, we describe and formally model a specific type of misalignment observed in our experiment that can significantly degrade the response time of the double-pass AOM [89].

In a double-pass AOM, the second-pass beam can become laterally offset from the first pass. Such misalignment is unlikely to produce a noticeable reduction in total diffraction efficiency, which is the usual metric of concern during double-pass alignment, but it can substantially increase the effective rise time and degrade the modulator's temporal response. Pulse-area linearity measurements on a well-aligned and misaligned system, shown in Fig. 2.12c, clearly illustrates this phenomena. Such a discrepancy can be undersood by considering the action the accoustic wave has on each pass of the AOM. Consider a left-propagating acoustic wave intersecting the optical beam during its first pass through the AOM (Fig. 2.12b). The acoustic wave can be modeled as a window function of width Δx, centered at position x_0. Only the portion of the optical beam that lies within this acoustic window is diffracted during the first pass and subsequently retro-reflected in the second pass. Crucially, the second-pass beam is a mirror image of the first-pass beam. If it does not spatially overlap with the acoustic wave—specifically, as in the case of $|x_0| > \Delta x/2$, it will not experience diffraction during the second pass. In this case, no light is transmitted through the double-pass AOM. Only when $|x_0| \leq \Delta x/2$ can the second pass be diffracted (Fig. 2.12b). The net result is a reduced transmission through the double-pass AOM.

Formally speaking, assuming an optical intensity profile

$$I(x) = \sqrt{\frac{2}{\pi\sigma^2}} \exp\left[-2\frac{(x-d)^2}{\sigma^2}\right],$$

of the first pass where d is the displacement of each pass from the optical axis, and σ the width of the beam, we can express the intensity of the transmitted second pass as

$$I_{\mathrm{sp}} = I(-x) \cdot \mathrm{Rect}\left[\frac{x-x_0}{\Delta x}\right] \cdot \mathrm{Rect}\left[\frac{-x-x_0}{\Delta x}\right] = I(-x) \cdot \mathrm{Rect}\left[\frac{x}{\Delta x - 2|x_0|}\right].$$

Here, the window function $\mathrm{Rect}[x] = 1$ if $|x| < 1/2$ and 0 otherwise. It is straightforward to show that when $\Delta x > 2|x_0|$, the condition of double-pass transmissivity, the total power transmitted through is given by

$$P_t = \int_{-\infty}^{\infty} I_{\mathrm{sp}}\, dx = \frac{1}{2}\left(\mathrm{Erf}\left[\frac{-2d + \Delta x - 2|x_0|}{\sqrt{2}\sigma}\right] + \mathrm{Erf}\left[\frac{2d + \Delta x - 2|x_0|}{\sqrt{2}\sigma}\right]\right).$$

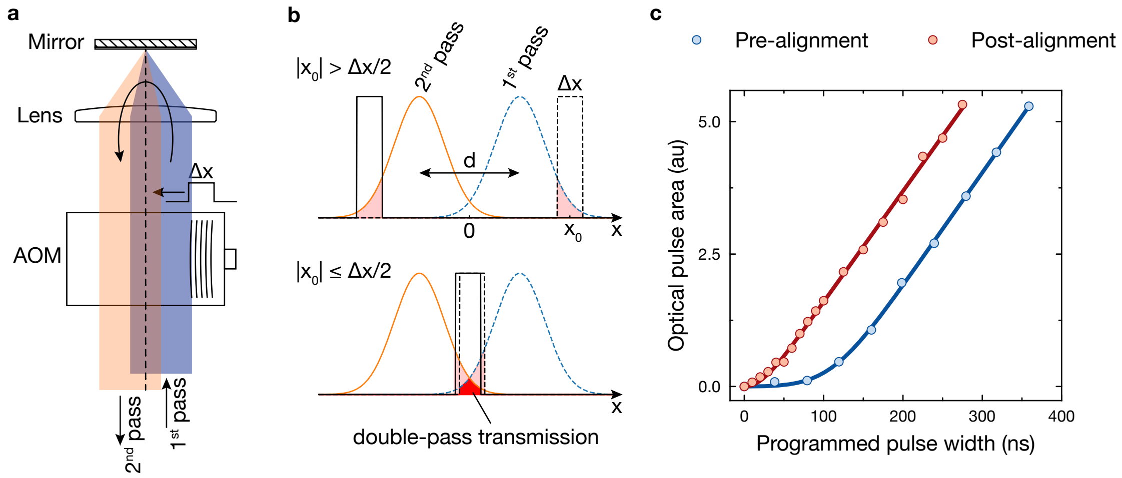

Figure 2.12: **(a)** Schematic illustration of lateral misalignment between the first and second passes in a double-pass AOM. The second passes are retroreflected by a cat-eye reflector. The cat-eye reflector mirrors the two passes. The accoustic pulse enters the AOM from the right. **(b)** Intensity profiles of the first pass (dashed blue) and second pass (solid orange) along the direction orthogonal to the acoustic wave propagation. The dashed box indicates the spatial extent of the acoustic pulse, which defines the region over which light can be diffracted during the first pass (red shaded area). Light transmitted during the first pass is reflected and reenters the AOM from the left in the second pass; however, due to the misalignment, it does not overlap with the acoustic wave and is therefore not diffracted. As a result, no net transmission occurs through the double-pass AOM. Double-pass transmission is only possible when the accoustic wave reaches the optical axis ($|x_0| \leq \Delta x/2$). The accoustic pulse width is defined as $\Delta x = v_s t$ where v_s is the speed of sound and t the programmed pulse width. **(c)** Measured optical pulse area as a function of the programmed acoustic pulse width. In the presence of lateral misalignment, strong nonlinear behavior at short pulse width is observed, consistent with the mechanism described in **(b)**. By using the optical pulse area at a fixed programmed pulse width as a feedback signal in an optimization loop, the linearity can be significantly improved. The fit function is defined in Eqn. (2.1).

The pulse area is then

$$\begin{aligned}\int_{-\Delta x/2}^{\Delta x/2} P_t\, dx_0 = &-\sqrt{\frac{2}{\pi}}\sigma \exp\left[-\frac{2d^2}{\sigma^2}\right] + \frac{\sigma}{\sqrt{2\pi}} \exp\left[-\frac{(\Delta x - 2d)^2}{2\sigma^2}\right] + \frac{\sigma}{\sqrt{2\pi}} \exp\left[-\frac{(\Delta x + 2d)^2}{2\sigma^2}\right] \\ &- 2d\,\mathrm{Erf}\left(\frac{\sqrt{2}d}{\sigma}\right) + \frac{1}{2}(\Delta x - 2d)\,\mathrm{Erf}\left(\frac{\Delta x - 2d}{\sqrt{2}\sigma}\right) \\ &+ d\,\mathrm{Erf}\left(\frac{\Delta x + 2d}{\sqrt{2}\sigma}\right) + \frac{1}{2}\Delta x\,\mathrm{Erf}\left(\frac{\Delta x + 2d}{\sqrt{2}\sigma}\right). \end{aligned} \tag{2.1}$$

The last expression is used to fit the linearity measurement in Fig. 2.12c, along with $\Delta x = v_s t$, where t is the programmed pulse width and v_s the speed of sound. After the alignment, the fitted displacement d approaches zero, which is consistent with the model under ideal alignment.

Beam shaping and pointing

The beam-shaping optics are designed with several constraints. First, the UV beam at the atoms must be an image of AOM 2 to ensure a faithful transfer of the dynamic phase from the AOM to the atoms. Second, the beam at AOM 2 must be appropriately sized to achieve a short risetime while maintaining high diffraction efficiency. Finally, the beam at the atoms must be large enough to homogeneously illuminate the atom arrays, but no larger, to maximize Rydberg coupling strength. In practice, we slightly focus the UV beam into AOM 2 to achieve a single-pass risetime 40 µs while retaining 80% diffraction efficiency. From an accoustic velocity of 5.6 mm/µs, one can deduce a beam waist of approximately 115 µm. By design, the beam waist on the atoms is approximately 40 µm. Even though spherical (non-cylindrical) optics are used for convenience, the final laser beam recorded on the CCD camera still exhibits non-negligible ellipticity, due to aberrations in the optical system.

While beneficial for achiving strong Rydberg coupling strength, tightly focusing the UV beam onto the atoms can significantly amplify the adverse effects of beam pointing drift. We adddress this problem by actively locking the beam pointing using a CCD camera (Allied Vision, E0022017, Manta G-145B mono PoE RCG) and a pair of moterized mirrors (Newport, Piezo Mirror Mount 8885), see Fig. 2.9. The alignment of the UV beam is initially optimized by maximizing the atomic Rydberg Rabi frequency on the atom array and possibly further fine-tuned by minimizing spatial inhomogeneity in a tweezer-resolved measurement. Once an optimal alignment is found, the beam's position on the CCD camera is recorded, and motorized mirror 1 steers the beam back to this reference position at the start of each experimental cycle. It is important that the CCD camera is placed at the focal plane of the last lens, just as the atoms are. So far, we rely on only one motorized mirror. However, it is conceivable that the second motorized mirror may also be required as angular misalignment in addition to positional drift can develope over time.

2.4 Optical tweezers arrays

Optical tweezers are tightly focused laser beams formed using high numerical aperture (NA) objectives. With beam waists typically below one micron, they provide strong atom-trapping potentials and have become a fundamental tool across many areas of quantum science. The properties of optical tweezers vary significantly depending on the method used to generate them. For instance, tweezers produced using a spatial light modulator (SLM) offer a high degree of programmability, enabling the creation of arbitrary tweezer arrays within a fixed optical setup. In contrast, tweezers generated using an acousto-optic deflector (AOD) are highly dynamic, allowing for real-time repositioning of trapped atoms. In the following chapter, we present a detailed discussion of SLM-based optical tweezers, including the tweezer projection system, the principles of holographic tweezer generation, and methods for wavefront correction. The basics of how a SLM works can be found elsewhere [90].

2.4.1 Optical alignment

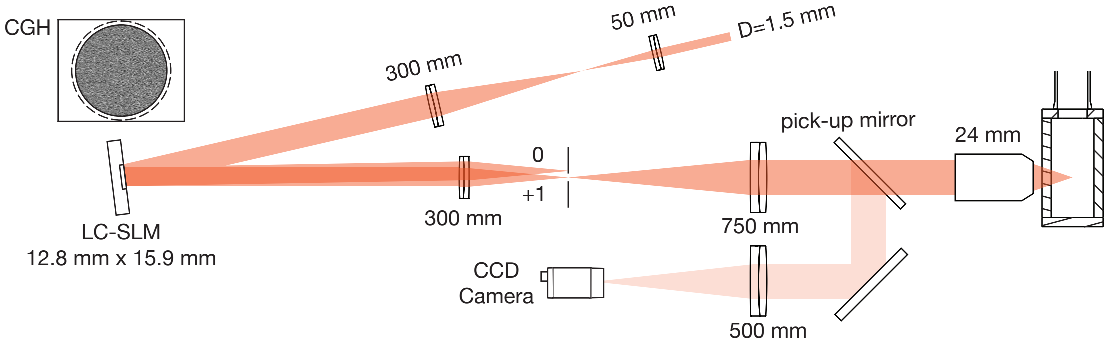

Figure 2.13: Optical setup of a phase-only SLM and the projection of light into the glass cell by a high-NA objective.

Several factors must be considered to generate high-quality optical tweezers. First, the input beam is expanded to a waist of $w_0 = 4.5\,\mathrm{mm}$ at the SLM, which has a half-width of 6.4 mm. If the beam is too small, it illuminates too few SLM pixels, resulting in poor diffraction efficiency. Conversely, if the beam is too large, it will be clipped at the edges of the SLM, leading to power loss. To enforce rotational symmetry and removing spurius diffraction from the circuit wiring at the edge of a SLM chip, a circular hard aperture with a radius of 5.3 mm is digitally applied on the SLM. Various computer-generated holograms (CGHs) are superimposed onto the SLM. Among them, a blazed grating is applied to produce a diffraction pattern in the image plane. Only the first diffraction order is retained, as it corresponds to the portion of light under full spatial control. All higher-order and zero-order beams are blocked. Subsequently, the first-order beam is magnified to fill approximately 90% of the objective's input aperture. This maximizes the trap depth by balancing optical power throughput with the waist size of the focused spot. For wavefront correction and tweezer monitoring, a high-quality glass plate with a nanotextured anti-reflection coating is placed immediately before the objective. This pickoff element reflects less than 0.5% of the light to a CCD camera for diagnostic purposes. To reduce aberration from strain, mirrors are glued into the mirror mount.

2.4.2 Tweezer hologram generation

For a phase-only SLM with no amplitude modulation capabilities, there is currently no analytical method to determine a hologram that transforms an incident Gaussian laser beam into a desired optical tweezer array. Consequently, one must rely on iterative numerical algorithms to approximate such a hologram. In the standard approach, the algorithm begins with a randomly initialized phase pattern on the SLM, which defines the complex electric field $E(x, y)_{\mathrm{SLM}}$. This field is numerically propagated to the focal plane to obtain $E(x, y)_{\mathrm{focus}}$. The amplitude at the focal plane is then replaced with the target intensity distribution $\sqrt{I(x, y)}$, while retaining the computed phase $\phi_{\mathrm{focus}} = \mathrm{Arg}[E(x, y)_{\mathrm{focus}}]$, resulting in a modified field $\sqrt{I}e^{i\phi_{\mathrm{focus}}}$. This modified field is then propagated backward to the SLM plane. At this point, a new SLM field is constructed by combining the known Gaussian amplitude profile of the input beam with the phase extracted from the back-propagated field. The process is repeated iteratively until convergence is achieved, typically defined by minimal deviation between the simulated and target intensity profiles in the focal plane.

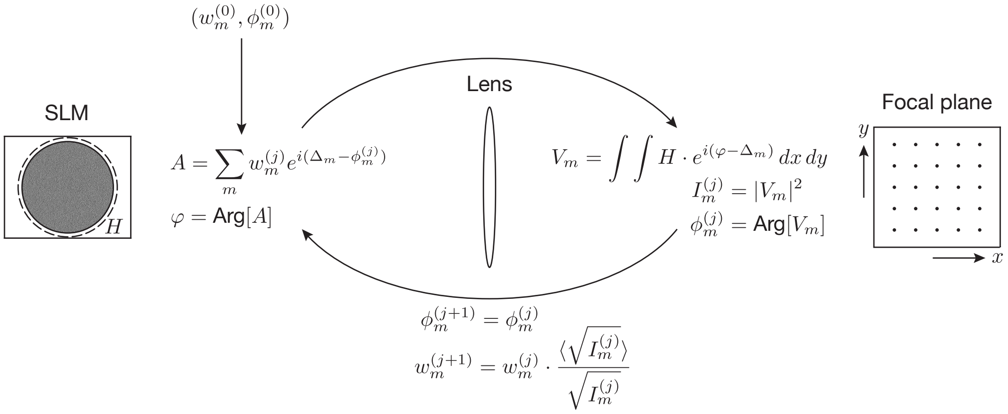

Figure 2.14: Phase-retrieval algorithm based on continuous-variable propagators

The phase-retrieval algorithm used in our experiment is adapted from [91, 92]. The algorithm can be understood from the following arguments. First, each tweezer at coordinate (x_m, y_m) is associated to a diffraction pattern

$$\Delta_m(x, y) = \frac{2\pi}{\lambda f}(x x_m + y y_m),$$

where λ is the laser wavelength and f is the focal length of the lens. The complex electric field on the SLM corresponding to the full tweezer array is then constructed as

$$A = \sum_m w_m^{(j)} \exp\left[i(\Delta_m - \phi_m^{(j)})\right],$$

summing over all tweezer m. Here, the phase offset $\phi_m^{(j)}$ serves to suppress amplitude modulation in $|A|$, ensuring compatibility with the phase-only constraint of the SLM. $w_m^{(j)}$ is the weight meant to equalize the tweezer amplitude across the array. The superscript (j) denotes the iteration index in the feedback loop, and the parameters $(w_m^{(j)}, \phi_m^{(j)})$ are iteratively updated during the optimization process. The phase pattern displayed on the SLM is then given by

$$\varphi(x, y) = \mathrm{Arg}[A].$$

To evaluate the focal-plane intensity of each tweezer, we apply the continuous-variable diffraction equation

$$V_m = \int\int H(x, y) \exp\left[i(\varphi - \Delta_m)\right]\, dx\, dy$$

where H is the aperture function defined on the SLM. The amplitude and phase of each tweezer on the focal plane is then used as input to the next iteration of feedback. The resulting complex amplitude V_m provides both the intensity and phase of the m-th tweezer in the focal plane, which are then used to update the parameters for the next feedback iteration. At certain feedback round, the phase $\phi_m^{(j)}$ is held to constant values to achieve faster convergence. Ultimately, as more rounds of feedback are applied, the uniformity of the tweezer array improves significantly.

A significant advantage of this elementary construction of complex holograms is the precise control it offers over each tweezer's position, focus, and shape at the individual tweezer level [92]. Moreover, by employing a continuous-variable propagator instead of a conventional discrete Fourier transform (e.g., FFT), the resulting tweezer positions on the focal plane are free from the pixelation-induced uncertainty caused by discrete sampling. For Rydberg physics in the non-blockaded regime, such a advantage is crucial.

2.4.3 Wavefront correction through self-interferometry

In an optical setup, wavefront errors can significantly degrade imaging quality. For passive optics, the best practice typically involves careful selection of optical components with minimal wavefront distortion, along with precise alignment. In some cases, wavefront aberrations can be partially mitigated using dedicated correction elements, provided the aberrations are known. However, the degree of control achievable with passive elements is inherently limited. In contrast, a spatial light modulator (SLM), as an active device, offers distinct advantages for wavefront correction by enabling programmable, spatially

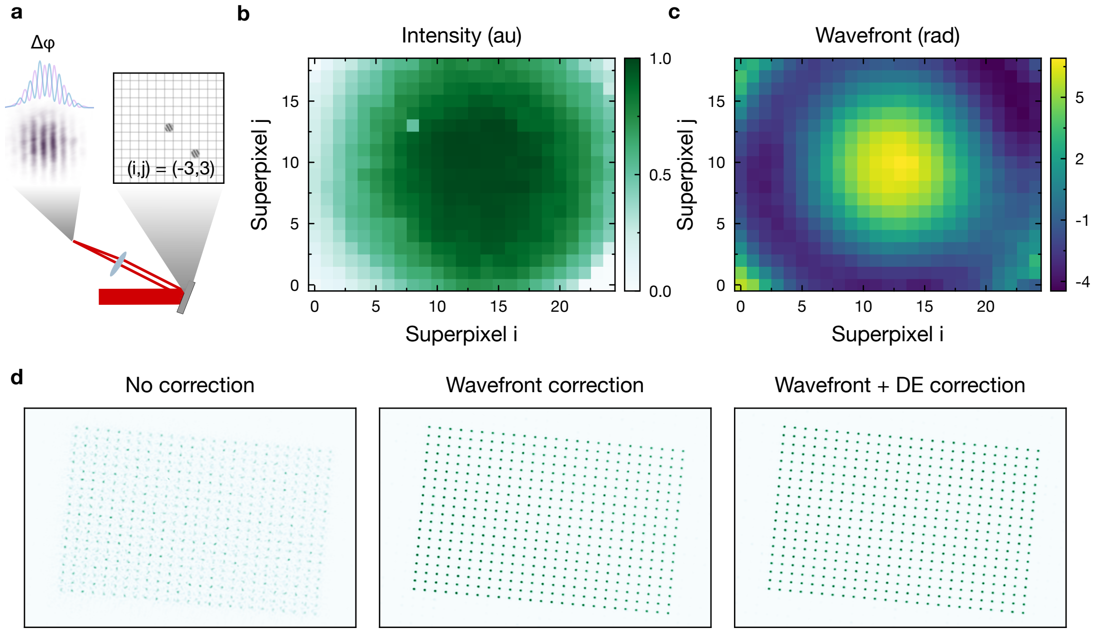

Figure 2.15: Wavefront correction based on phase-shifting interferometry. **(a)** Self-interferometry setup [93]. Light reflected from two superpixels (inset) interferes spatially, revealing the relative phase $\Delta\varphi$ they accumulate as they propagat from the SLM pixels to the point of interference. By keeping one superpixel fixed and scanning the other across the entire SLM pixel grid, the full wavefront of the laser bounded by the SLM sensor can be reconstructed, as shown in **(b)** and **(c)**. **(d)** With wavefront correction applied, the optical tweezers exhibit significantly higher peak intensity due to improved focusing. Residual intensity gradients can be largely eliminated using independently measured diffraction efficiency corrections. Note that all three tweezer images share an identical color scale and range, highlighting the effectiveness of each correction method with striking clarity.

resolved phase adjustments. Here, we describe a simple technique extensively used in our experiment which can reconstruct the laser wavefront with an SLM and a CCD camera.

The self-interferometry method is based on phase-shifting interferometry [66], although active phase shifting is no longer required in its modern implementation. The core idea is to interfere light emanating from one region of the SLM with light from another, thereby extracting the relative phase difference. To enhance the signal-to-noise ratio, several neighboring pixels are typically grouped into a single superpixel Fig. 2.15a. By keeping one superpixel fixed as a reference and scanning the other across the entire SLM surface, the phase profile of the laser beam over the SLM can be reconstructed. This method is particularly effective when the wavefront error varies slowly over spatial scales larger than the size of a superpixel, allowing for reliable wavefront reconstruction Fig. 2.15b,c. Once the wavefront measurement is obtained, it is added to the tweezer hologram due to the additive feature of the phase modulation. See Fig. 2.15d for a comparison between various corrections.

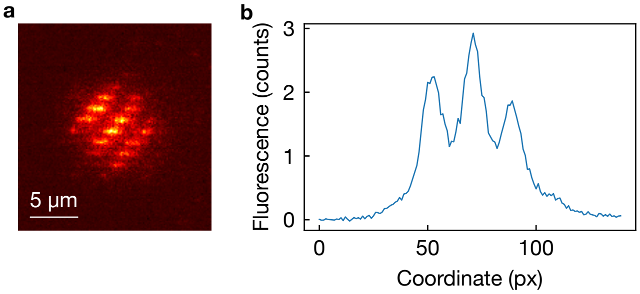

Figure 2.16: *In situ* characterization of wavefront correction using the atoms themselves. **(a)** Self-interference fringes revealed via atom loading in a hybrid tweezer-lattice apparatus. **(b)** The fringe is clearly visible in the column-averaged fluorescence signal along one of the principal axes.

The effectiveness of wavefront correction is fundamentally limited by the accuracy of the detector used to measure the wavefront. If the detector is positioned along an optically aberrated beam path, it may record unreliable phase information. Furthermore, between the detector (e.g., a camera) and the atomic plane, there is typically a glass cell and a high-numerical-aperture objective, both of which introduce additional aberrations that cannot be corrected through standard calibration Fig. 2.13. Ultimately, the most accurate wavefront correction can only be achieved by treating the atoms themselves as the phase-sensitive detectors [66]. By leveraging the fine spatial sampling of an optical lattice over the tweezer array, we can infer the wavefront phase in a manner analogous to camera-based detection. In a proof-of-principle experiment, we employ a loading sequence in which atoms are transferred from a dense lattice into the interference fringes created by the SLM-generated trapping light. Subsequently, atoms are returned to the lattice, allowing the fringe pattern to be identified through fluorescence imaging in the lattice Fig. 2.16a,b. Although the lattice spacing imposes a fundamental limit on the smallest resolvable spatial features in fluorescence imaging, the fringe spacing can be kept fixed. By translating the superpixel pair in a common-mode fashion and tracking the incremental changes in relative phase, the full wavefront phase profile can be reconstructed. While not shown in this thesis work, such an *in situ* characterization method represents an unique advantage of our hybrid tweezer-lattice apparatus.

2.4.4 Tweezer equalization

For equalizing tweezer trap depths, we leverage an observed correlation between the optimal Sisyphus cooling condition on the $^1S_0$ to $^3P_1$ transition and tweezer trap depth. This probe remains robust even in the early stages of optimization, where the tweezer array exhibits large variations in trap depth that renders imaging on the global scale challenging. We begin with a tweezer array that is sufficiently uniform in trap depth to enable atom loading in every site. If this condition is not met, an initial optimization using a separate low-NA imaging system with a camera is typically necessary. We then scan the Sisyphus cooling frequency during imaging and identify the frequency that maximizes fluorescence for each tweezer. This site-resolved optimal cooling frequency serves as an experimental observable that can be used for equalization, see Fig. 2.17a,b.

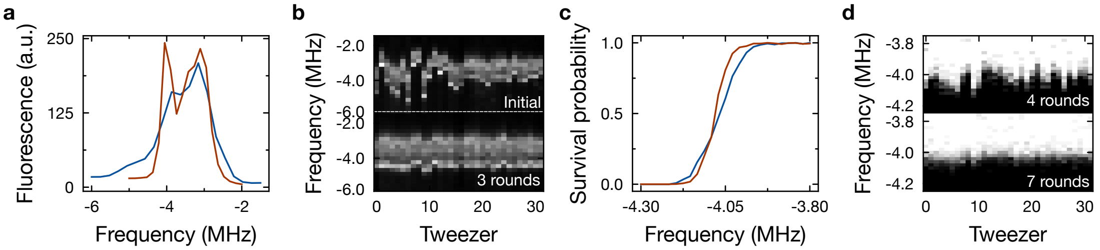

Figure 2.17: Tweezer trap depth equalization baesd on Sisyphus cooling on $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_1$ transition. **(a)** Tweezer-averaged fluorescence shows pronounced spectral narrowing from uncorrected tweezers (blue) in tweezers corrected with 3 rounds of amplitude feedback (red). **(b)** When tweezer-resolved features are well separated initially, the mean cooling frequency that maximizes the fluorescence is used to balance the amplitude difference between tweezers. The numbers label the index of the iteration round. **(c)** After a few iterations of fluorescence-based feedback, tweezers become sufficiently homogeneous and reasonably high imaging fidelity can be reached. For further optimization, we focus on the blue-detuned edge of the cooling feature in an imaging block owing to its better correlation to trap depth in the attractive cooling regime. Arrays with better amplitude homogeneity (red) exhibit sharper transition between low and high survival in imaging than the ones with lower homogeneity (blue). **(d)** For feedback purpose, the midpoint of the slope between low and high survival probability is used. The numbers label the index of the iteration round, continuing from the labeling in the fluorescence-based optimization.

To equalize trap depth below the percent level, we try to overlap the steep edge of the Sisyphus cooling feature at blue detunings, and sharpen the transition in survival probability from low to high, see Fig. 2.17c,d. In holographic arrays created using a spatial light modulator (as in our case), this is achieved by tuning the weight of each tweezer in the phase-retrieval algorithm. Our protocol shows a strong convergence and is expected to further reduce inhomogeneity beyond $\sim 0.3\%$ achieved here with more iterations, see Fig. 2.18a. To confirm a trap depth homogeneity across the entire array after the optimization, we show we can reach more than 209 s lifetime under continuous illumination, see Fig. 2.18b. As a comparison, our vacuum lifetime is estimated to be 273 s [94]. The imaging lifetime corresponds to an imaging survival probability > 0.99928 for 150 ms-long imaging.

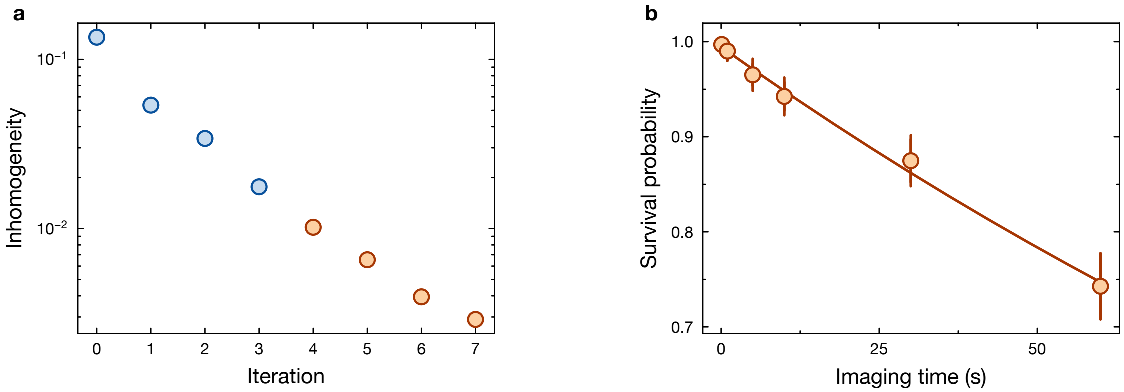

Figure 2.18: Convergence plot for tweezer amplitude equalization and imaging characterization. **(a)** The atomic signal used is initially the central cooling frequency, which maximizes fluorescence during iterations 0–3, and later the blue-detuned edge of the cooling frequency, which optimizes survival probability during iterations 4–7. **(b)** With a tweezer array optimized by this homogenization strategy, the continuous imaging lifetime exceeds 200 seconds, corresponding to a survival probability of > 0.99928 for a 150 ms-long imaging.

Chapter 3

High-fidelity imaging of large-scale arrays

Neutral atoms have emerged as a powerful platform in quantum science. Among their numerous advantages, scalability stands out as one of the most distinguishing features compared to alternative architectures. In quantum metrology, large atomic ensembles enhance measurement sensitivity by improving signal-to-noise ratios and enabling entangled state protocols [37]. In quantum simulation, a substantial number of particles is often necessary to faithfully capture the emergent phenomena of complex many-body systems [95]. For fault-tolerant quantum computing, scalability is a critical requirement for reducing error rate of logical qubits and implementing deep quantum circuits.

In neutral atom platforms, it is typically possible to create atomic ensembles containing $\sim 10^8$ atoms, generated from an effusive atomic oven and captured using magneto-optical traps (MOTs). For individual atom control, the ensembles are usually loaded into discrete optical traps with each trap ideally containing at most one atom. Two mainstream approaches exist for creating such optical trapping arrays. The first employs optical tweezers, generated using high-numerical-aperture (high-NA) objectives in combination with spatial light modulators (SLMs). System sizes exceeding 6,100 individually trapped atoms have already been demonstrated [96]. However, further scaling of tweezer arrays faces several challenges. Notably, the diffraction efficiency of SLMs decreases as the array size increases, meaning that significantly more optical power is required to generate a trap far from the optical axis than one near the axis. This limitation can be partially mitigated by employing metasurfaces with finer pitch sizes, which offer improved efficiency for large arrays [97, 98, 99, 100]. Another constraint arises from the finite field of view of the objective, which leads to significant degradation of the trapping potential away from the optical axis due to optical aberrations. Optical lattices, by contrast, offer a promising alternative that circumvents many of the limitations of tweezer arrays. Formed by optical interference, lattices can generate tens of thousands of well-defined trapping sites with high power efficiency. This enables substantially higher atom densities per unit area and supports the realization of large-scale, highly ordered atomic arrays suitable for quantum simulation and computation.

In this chapter, we present our work on high-fidelity imaging of large-scale atom arrays in an optical lattice [94]. By combining optical lattice trapping with cooling techniques tailored to alkaline-earth atoms, we demonstrate non-destructive imaging of more than 10,000 lattice sites. Simultaneously, we achieve direct loading of over 10,000 atoms into the lattice, with site-resolved detection via parity projection. This system represents the state of the art in scaling up neutral atom arrays. In Section 3.1, we introduce two distinct cooling mechanisms implemented in strontium. Section 3.2 introduces the tweezer-lattice hybrid platform and benchmarks the repulsive Sisyphus cooling using isolated atoms in a lattice loaded via optical tweezers. Section 3.3 applies the scalable imaging methods to a lattice directly loaded from a MOT.

3.1 Cooling on the intercombination line

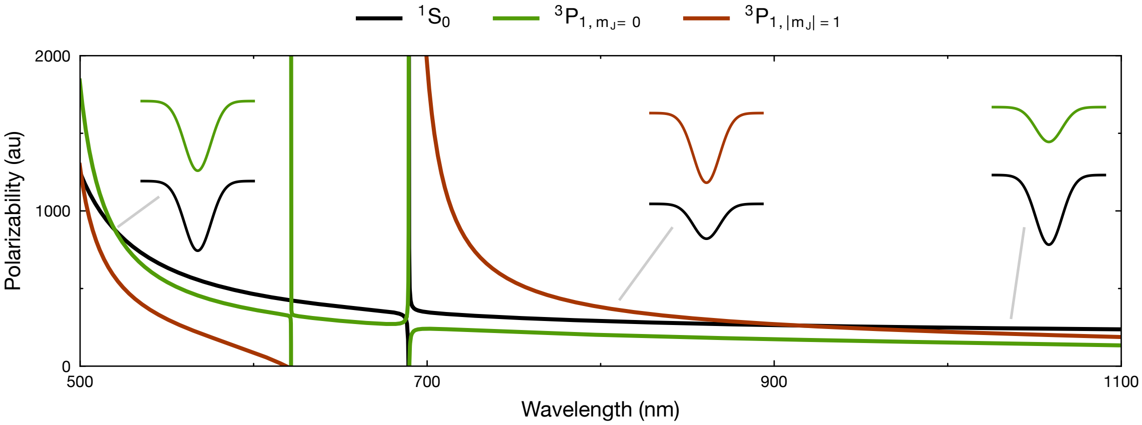

Figure 3.1: The atomic polarizability of $^1S_0$, $^3P_{1,m_j=0}$, and $^3P_{1,|m_j|=1}$ at trapping wavelengths between 500 nm and 1100 nm [101]. The inset illustrates cooling regimes at three representative wavelengths: resolved sideband cooling at 520 nm, attractive Sisyphus cooling at 813 nm, and repulsive Sisyphus cooling at 1040 nm. The colors of the potential landscapes in the inset correspond to the specific states used in the cooling.

Efficient cooling of atoms is an critical requirement in quantum science. First, it suppresses deleterious motional dephasing, which can degrade the performance of quantum operations. Second, it enables a highly controlled dissipation of excess energy during resonant imaging and ensures a high-fidelity, non-destructive fluorescence read-out. The narrow intercombination line between $^1S_0$ and $^3P_1$ in alkaline-earth species such as Sr and Yb offers unique opportunities for laser cooling, owing to the ability to resolve motional states on this transition. Specifically, one can engineer *repulsive* or *attractive Sisyphus cooling*, depending on whether the excited state $^3P_1$ on the cooling transition is less or more tightly confined than the ground state $^1S_0$ (Fig. 3.1). In strontium where the intercombination line has a narrow linewidth of 7.5 kHz, one can access the *resolved sideband cooling* regime if ground and excited states experience similar trapping potential, and reach high fraction of motional ground state of the trap.

3.1.1 Resolved sideband cooling

In resolved sideband cooling (RSC), motional quanta of a trapped atom are removed through repeated cycles of coherent excitation and dissipation [102]. By tuning the coherent drive to the first red motional sideband, the atom can be eventually transferred to its motional ground state—the only dark state of the cooling process. In contrast to RSC in alkaline species where two lasers are required for coherent and dissipative steps [103], the radiative decay from the excited state in strontium naturally provides the dissipative leg, eliminating the need for a second laser and simplifying the overall scheme [74, 104]. The conditions for effective RSC are as follows. First, the cooling transition should be *magic* in the trapping potentials, meaning that the differential polarizability vanishes across the entire potential. At certain trapping wavelengths, this condition can be naturally realized. Examples include $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_{1,m_j=0}$ at 520 nm and $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_{1,|m_j|=1}$ at 914 nm, see Fig. 3.1.

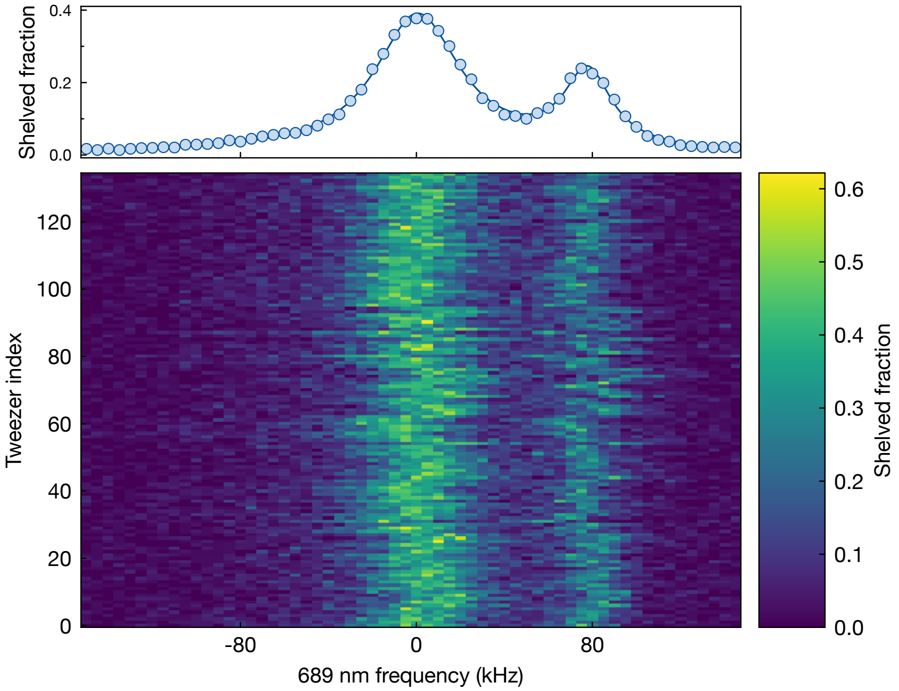

Figure 3.2: Tweezer-averaged (top) and tweezer-resolved (bottom) sideband spectra cooled and probed on the intercombination line in the resolved sideband regime. The average motional occupation is below 0.1.

However, through the technique of polarizability mixing, magic trapping conditions can be extended to a much broader wavelength range [104]. For example, in 813 nm optical tweezers, this technique is used to engineer magic trapping on the intercombination line. Second, the power-broadened transition linewidth should be much smaller than the trapping frequency. Third, the Lamb-Dicke parameter which characterizes how tightly an atom is confined in the trap should be much less than 1. All of these conditions can be reasonably satisfied in strontium at suitable trapping conditions. In Fig. 3.2, the sideband spectra after sideband cooling in 813 nm tweezers are shown.

3.1.2 Sisyphus cooling

At trapping wavelengths where a differential polarizability exists between the ground and excited states of the cooling transition, a Sisyphus cooling mechanism can be realized. Depending on whether the excited or ground state experiences stronger confinement, the cooling is referred to as either attractive or repulsive. The core principle behind this cooling mechanism involves the selective excitation of atoms such that they are more likely to absorb photons while ascending a steeper potential and decay while descending a shallower one (see Fig. 3.3). In trapping conditions where the differential polarizability is large, the resulting cooling rate can be high.

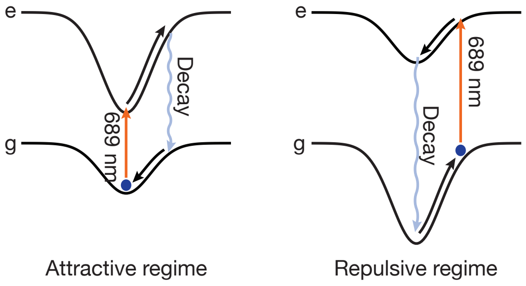

Figure 3.3: Potential of cooling states in attractive and repulsive Sisyphus cooling regime

In the attractive regime, the excited state is more tightly confined than the ground state. The cooling light at 689 nm is tuned to be resonant with the cooling transition at the center of the optical trap. Due to the narrow linewidth of the transition, only atoms near the trap center are excited. Because the radiative lifetime of the excited state is comparable to the trap oscillation period, the atom has time to climb the steeper excited-state potential before spontaneously decaying. After returning to the ground state, it descends a shallower potential back to the trap center, losing net energy in the process. In the repulsive regime, the 689 nm light is red-detuned from the resonance at the trap center, and the roles of the ground and excited states are reversed. A similar cooling effect can still be achieved, although with additional complexities. Cooling begins to occur at sufficiently large detunings, where the atoms preferentially absorb light near the edges of the trap. As the detuning decreases, the equilibrium temperature also drops. However, if the detuning becomes too small, the cooling mechanism can reverse, causing atoms to gain energy instead, leading to runaway heating and eventual atom loss. Therefore, careful optimization of the cooling frequency is required to balance the cooling rate with atom survival probability. Note that repulsive Sisyphus cooling is much less sensitive to trap depth inhomogeneity than attractive Sisyphus cooling, making it ideal in large-scale imaging.

3.2 High-fidelity detection in tweezer-loaded lattice

3.2.1 Infared pinning lattice

In quantum science with neutral atoms, trapping wavelengths are almost never one-size-fits-all. In strontium, for example, 520 nm is ideal for ground-state cooling and scaling to larger systems, but it is not compatible with high-fidelity fluorescence imaging [74, 104]. 813 nm is well-suited for interrogation on the ultra-narrow clock transition, but scaling up is difficult due to limited laser power [105, 106]. A powerful technique to address this issue is to introduce multiple trapping wavelengths into the experimental cycle, each optimized for its specific purpose [107]. A particularly useful choice for imaging is a wavelength in the infrared where high-power lasers are readily available. In addition, infrared wavelengths are less likely to be near-resonant with atomic transitions than visible wavelengths, reducing their negative impact on imaging. In this section, we introduce a imaging lattice at 1040 nm for a science tweezer array at 520 nm. Owing to the availability of high laser power at 1040 nm, the lattice can be sufficiently deep to support high-fidelity imaging over $> 10^4$ lattice sites. We demonstrate high-fidelity atom transfer between the lattice and tweezer array, showing repeated imaging in the lattice and motional state reset in tweezers. Such a tweezer-lattice hybrid system is expected to facilitate the scaling of atom arrays. All experimental results in this and next section are adapted from [94].

Infared lattice

Our optical lattice is formed from an single laser beam at 1040 nm folded multiple times, see Fig. 3.4. With a 4f alignment of relaying lens, the beam is focused to a waist of 100 µm (20 µm) in (out of) plane at each pass of the crossing. To maximize the interference contrast, the polarization is linearly out of plane. At a full crossing angle of 53°, the lattice spacing of the resulting interference pattern is designed to be 581 nm along the x axis and 1165 nm along the y-axis. This choice of 2:1 lattice spacing ratio is primarily motivated by the potential use of higher loading along the short spacing axis as a atom resevoir in assembling a unity-filled square array with longer spacing *in the lattice*. Atom assembly with such a local resevoir is supposed to scale more favorably to larger arrays due to the availability of shorter transports which have lower loss probability. To the ends of quantum simulation of itenerant model, a full crossing angle of 90° would have been more advantangeous for exploring the full 2-dimensional connectivity in the lattice. However, the optical acess required by this configuration is unfortunately blocked by the coils.

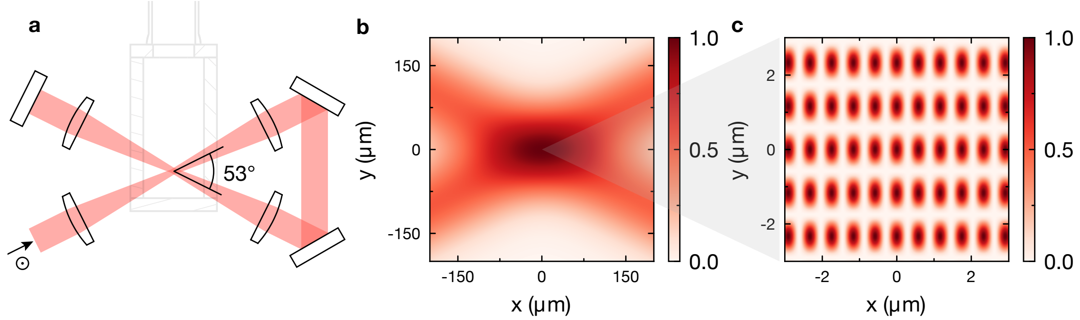

Figure 3.4: **(a)** Folded optical lattice geometry formed by a single 1040 nm laser beam. The beam is folded and retroreflected to create an interference pattern at the center of the glass cell at an angle of 53° between the two arms. The polarization of the beam is linear and out of plane. The four lens in the beam path focus the beam to a waist of 100 µm on the plane and 20 µm out of the plane. **(b)** Simulated optical intensity of the folded lattice near the crossing point, using realistic parameters. **(c)** A close-up view of lattice, showing asymmetric confinement in the plane. The designed lattice spacing is 581 nm along the x axis and 1165 nm along the y-axis.

Green tweezers

We generate our 520 nm tweezer array holographically using a phase-only spatial-light modulator. The algorithm for creating the hologram largely follows earlier work [108]. In the optical setup, we use a collimated beam with a waist of 4.2 mm incident on the SLM chip of dimension 12.8 mm×15.9 mm. A telescope consisting of a 300 mm and 750 mm achromatic lens relays the imprinted phase pattern to the back focal plane of the high-resolution objective (NA 0.65). The objective then transforms the phase pattern into an array of tightly-focused tweezers in the atomic plane. Undiffracted light at the SLM is blocked in an intermediate imaging plane. An added blazing phase grating steers the diffracted beam by about 0.5° away from the undiffracted light and defines the new optical axis, with which we align our optics downstream. To avoid clipping at the input aperture of the objective, we impose a circular aperture of 11.5 mm in diameter on the SLM, such that the light fills less than 90 % of the input pupil of the objective [109]. This aperture is assumed in all phase-retrieval algorithms. To correct wavefront distortion, we sample the SLM light shortly before it enters the objective and focus the beam onto a camera. Using phase-shifting interferometry [66, 110] with super-pixels on the SLM, RMS wavefront error can be suppressed to below $\lambda/14$.

Tweezer-lattice alignment

For adiabatic transfer of population between tweezers and the lattice, alignment of the trapping potential along all three axes is critical. For in-plane alignment, our approach is to first carefully measure the lattice spacings a_x, a_y and angles $\theta_x^{\text{lat}}, \theta_y^{\text{lat}}$ with respect

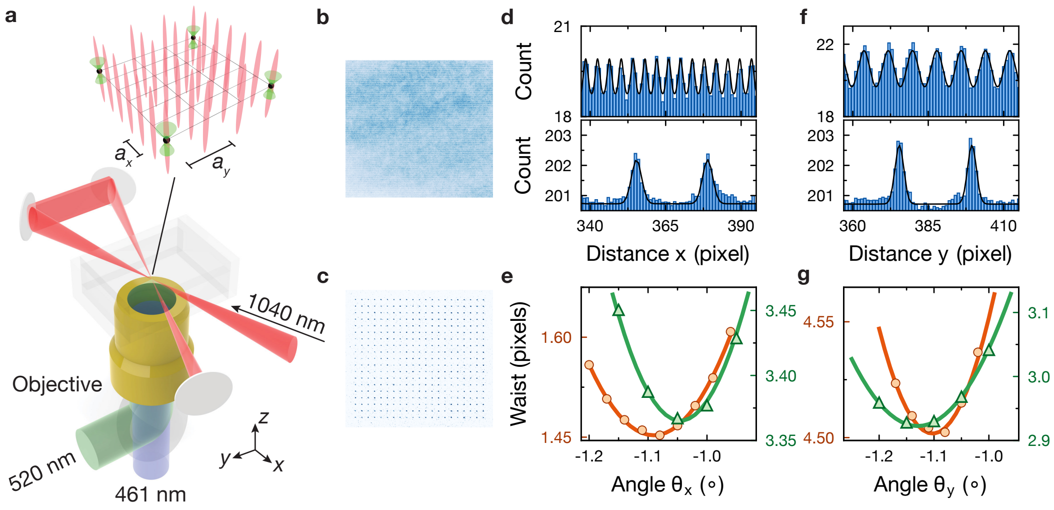

Figure 3.5: **(a)** A sketch of the folded lattice at 1040 nm overlapped with optical tweezers at 520 nm. The atomic fluorescence at 461 nm is split from tweezer light at a dichroic mirror. **(b)** An averaged fluorescence image of a sparsely-loaded lattice used for the analysis. **(c)** An averaged fluorescence image of a tweezer array imaged in tweezers. For images of 2D fluorescence patterns defined by two orthogonal symmetry axes (like **(b)** and **c**), the 1D fluorescence count, which results from projecting the image array along some axis, exhibits smaller peak overlap when the projection axis coincides with either symmetry axis of the array [56, 111]. Such a case for $x(y)$-axis is shown for lattice in the upper panel and the tweezer in the lower panel in **(d)**(**(f)**) Any deviation from such a condition results in an increased peak overlap **(e, g)**. For the analysis, the dependency of the waist on the angle of the projection axis is fitted with $\sigma_0\,[1+\beta\{1-\cos(\gamma(\theta_i-\theta_{i,0}))\}]$ for $i=x,y$ [111]. In both **(e)** and **(g)**, lattice data (circle) and fit (solid line) are red, whereas tweezer data (triangles) and fit (solid line) are green. Once the minimum-waist angle is found, the lattice and tweezer array spacing and waist can be extracted by fitting a grid of Gaussians functions in the rotated frame of reference **(d, f)**. In summary, we extract lattice coordinate angles $\theta_x^{\text{lat}}(\theta_y^{\text{lat}}) = -1.090(0)^\circ(-1.102(1)^\circ)$ and tweezer coordinate angles $\theta_x^{\text{tw}}(\theta_y^{\text{tw}}) = -1.0452(2)(-1.1296(1))^\circ$. For the spacings, we find lattice constants of $a_x(a_y) = 3.92(6)(8.04(9))$ pixels and tweezers spacing of $c_x^{\text{tw}}(c_y^{\text{tw}}) = 23.564(6)(24.047(6))$ pixels.

to the lab reference frame, and then engineer a suitably shaped tweezer array to achieve maximum commensurability with the lattice. Specifically, we define tweezer spacing c_i such that $c_x/a_x = 6$ and $c_y/a_y = 3$, and set $\theta_i^{\text{tw}} = \theta_i^{\text{lat}}$ for $i = x, y$, see Fig. 3.5. Such a conservative tweezer spacing is intended to minimize crosstalk between tweezers, and future experiments with tighter spacing should be feasible. The remaining degree of freedom is a global offset which is usually time-dependent. We remove this offset by periodically checking the individual positions of the lattice and tweezers using atomic fluorescence in their respective potentials and adjusting the tweezer position via an applied blazing phase.

For out-of-plane alignment, we adjust the focus of the tweezers via an applied Fresnel phase on the SLM to make fluoresence imaging in the tweezers as focused as that in the lattice on the camera.

3.2.2 Repulsive Sisyphus cooling in the infared lattice

To establish the infared lattice as a viable imaging lattice, high-fidelity imaging must be demonstrated. At the lattice wavelength of 1040 nm, one can exploit repulsive Sisyphus cooling mechanism on either $^1S_0 \rightarrow {}^3P_{1,m_j=0}$ the π-transition or $^1S_0 \rightarrow {}^3P_{1,|m_j|=1}$ the σ-transition to counteract heating from resonant imaging. To load atoms into the infrared lattice, we use the 520 nm tweezer array as an intermediate step, taking advantage of efficient parity projection at this green wavelength [74]. As a result, we can load arrays of isolated atoms in the lattice and significantly simplify the occupation reconstruction.

The imaging condition involves simultaneous scattering on the $^1S_0 \rightarrow {}^1P_1$ transition and cooling on the $^1S_0 \rightarrow {}^3P_1$ transition. In addition, two repumps are applied concurrently to return the population in $^3P_0$ and $^3P_2$ accumulated through leakages from $^1P_1$ and trap-induced Raman scattering back to the imaging cycles. We begin the experiment with an array of atoms at 399 singly-occupied lattice sites loaded from tweezers. To optimize the cooling performance for imaging, we scan the detuning Δ of the cooling light relative to the free-space resonance, see Fig. 3.6c. We obtain a broad cooling feature at approximately 2.2 MHz, where the atom loss fraction reaches the sub-percent, and a second narrower feature at about 9.8 MHz, where the loss is higher. The two features can be attributed to cooling on $^1S_0 \rightarrow {}^3P_{1,m_J=\pm1}$ and $^1S_0 \rightarrow {}^3P_{1,m_J=0}$ respectively and are consistent with the 689 nm transition split by tensor lightshift in the 1040 nm lattice. To characterize imaging performance under cooling on the $m_J = \pm 1$ transition, we take two consecutive images in the optical lattice, see Fig. 3.6a. The images are binarized based on the tweezer-averaged histogram of the integrated photon count, see Fig. 3.6b. With an optimal threshold, we obtain a classification infidelity of approximately 10^{-4}, demonstrating the feasibility of high-fidelity imaging in our lattice. To benchmark the atom loss probability from imaging, we compare the occupation of the two consecutively acquired images as a function of the exposure time, see Fig. 3.6d. For this measurement, we keep the integrated photon number scattered on the 461 nm transition and hence the classification fidelity constant. We find a robust minimum atom loss probability at an exposure time of approximately 200 ms, where the loss reaches 2×10^{-3}. At shorter exposure times and hence larger imaging beam scattering rate, the atom loss probability increases as a result of recoil heating exceeding the cooling rate from Sisyphus cooling. For longer exposure times, the atom loss probability begins to be dominated by our estimated vacuum lifetime of 273(3) s.

To find an optimal binarization threshold and compute a classification fidelity, we simply fit the background peak and one-atom peak of the histogram with analytical function and optimize the threshold to minimize the misclassification rate. This procedure assumes that the histogram asymptotically approach the underlying emission probability mass function with a sample pool large enough. For better noise suppression, we first apply a deconvolution filter derived from the measured point-spread function to individual atomic fluorescence before summing up the counts in a square roi 3×3 lattice sites large. On account of our use of tweezers in loading the lattice, the position of the atoms in the lattice is susceptible to small variation from shot to shot, as a result of relative time-dependent phase drift which is not regulated in this measurement and residual imperfect time-independent spatial incommensurability between the two potentials. We heuristically find that identifying the location of atoms with a lattice sites-resolved occupation reconstruction aids in centering the roi. Consequently, the histogram exhibits a smooth profile that is amenable to analytical modeling, see Fig. 3.6b. We empirically fit the one-atom emission and background separately using a modified Gaussian function $P(x, x_0) = A \exp\left[-\frac{(x-x_0)^n}{\sigma^2}\right]$ where n is allowed to deviate from 2 and differs depending on whether $x < x_0$ or $x \geq x_0$. An optimal threshold can be found by minimizing the probability of detecting false positive and false negative. We find that the final classification fidelity is insensitive to the histogram bin size, lending credit to the assumption we make for this analysis. The error bar in the classification is estimated from assigning the fitting residual to falsely detected events according to the worse case scenario.

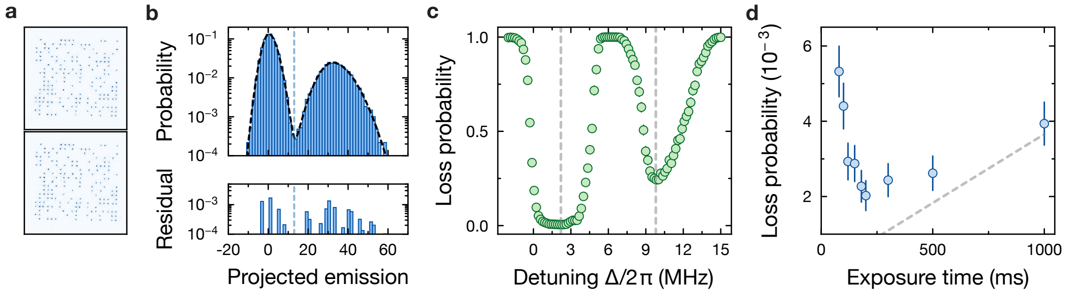

Figure 3.6: **(a)** Exemplary consecutive fluorescence of atoms imaged in the lattice. **(b)** Array-averaged histogram of photon counts taken with exposure time of 300 ms, showing a well-resolved background and one-atom fluorescence peak. The classification fidelity can be as high as 0.99971(1). **(c)**. Array-averaged atom loss probability as a function of 689 nm cooling laser detuning Δ from the ${}^1\mathrm{S}_0 \to {}^3\mathrm{P}_1$ resonance in free space for an imaging exposure time of 300 ms. At $\Delta/2\pi = 2.2$ MHz, imaging loss reduces down to 2×10^{-3}. **(d)**. Imaging loss at constant scattered photon number for classification infidelity of 10^{-4} vs. exposure time. The lower dashed gray line indicates atom loss probability attributed to our estimated vacuum lifetime of 273(3) s, which is reached by our imaging in the limit of long exposure time and low illumination power. The background level remains similarly low all all exposure times.

3.2.3 Post-imaging transfer of atoms to the tweezers

Having demonstrated high-fidelity imaging in the infared lattice loaded using the tweezers, we next show a high-fidelity reloading of atoms back into the tweezers, and resetting its temperature. Such a step is essential in fulfilling the role of the lattice as an imaging potential for tweezers. To benchmark the transfer fidelity, we identify the initial presence of an atom in the lattice and track its survival at the same location after a lattice-tweezer-lattice transfer, see Fig. 3.7a. A challenge in our setup is the weak vertical confinement of our infared lattice, whose waist in the axial direction significantly exceeds the Rayleigh range of 1.5 µm of the tweezers, see Fig. 3.5a. To enable low-loss transfer back to tweezers, we first perform an optimized repulsive Sisyphus cooling in the lattice after imaging (II). Subsequently, we ramp up the tweezers to a depth of 300 µK, before lowering the lattice to an intermediate depth of 150 µK. We perform a second stage of cooling in this combined potential to efficiently transfer the atoms into trapped states in the tweezers. The cooling frequency is chosen to coincide with the lattice-light shifted cooling sideband of the tweezers and the magnetic field is set to the magic cooling transition in tweezers alone. Finally, we ramp the lattices down in 50 ms, completing the transfer to the tweezers. Imaging is then performed once more in the lattice, with an identical tweezer-lattice handover as before the first image. With the images taken before and after lattice-tweezers-lattice transfer, we benchmark the complete round-trip atom loss probability p_n by comparing the reconstructed tweezer occupation between two images taken in the lattice before and after n transfers, see Fig. 3.7e. While the overall atom loss probability increases with the number of round-trips as expected, we find that the probability of atom loss probability per round-trip p_1, extracted under the assumption of a simple power-law scaling of the atom loss probability $1-p_n = (1-p_1)^n$, continuously decreases from 1.3% down to approximately 5×10^{-3} after a few round-trips. We attribute the initially higher atom loss probability predominantly to a systematic spatial inhomogeneity affecting our cooling in the lattice during transfer, which becomes directly apparent in a tweezer-resolved transfer loss map after $n = 80$ round-trips, see Fig. 3.7e inset. Hence, we consider the reported transfer loss as a worst case that can be improved by excluding the traps exhibiting high atom loss or by repositioning the tweezer array relative to the lattice. We find that highly efficient transfers are possible if the tweezer depth in the transfer exceeds approximately 300 µK, see Fig. 3.7c. For the last point beyond 400 µK, the transfer loss probability increases slightly due to non-optimal cooling parameters.

Good transfer efficiency requires a precise alignment of the optical lattice with the tweezer array. For this, we study the dependence of the transfer efficiency on the relative position between lattice sites and tweezers, see Fig. 3.7d,f. When scanning either the vertical or the horizontal position of the tweezer array, we find a pronounced sinusoidal dependence of the transfer loss, which reaches up to 5% in non-optimal conditions. The observed sinusoidal structure is in excellent agreement with the expected dependence for our lattice, and the curves represent a characterization of the underlying lattice structure using a large-scale tweezer array [112]. Finally, we demonstrate the possibility of re-cooling the tweezers close to the radial motional ground state after transfer. As a first step, we characterize the temperature of atoms after the tweezer-lattice-tweezer transfer without cooling. Employing sideband spectroscopy in the tweezers after transfer, we find a significant amount of thermal excitations in the radial motion mode, with average ground-state fraction of 0.37^{+40}_{-17}. After performing sideband cooling for 50 ms in the tweezers, we obtain near ground-state preparation in the radial motional mode with an estimated ground-state fraction of 0.98^{+0}_{-3}, indistinguishable from the initial temperature before the first image in the lattice, see Fig. 3.7g,f. These results are confirmed by release and recapture measurements from the optical tweezers, demonstrating the beneficial effect of sideband cooling after transfer.

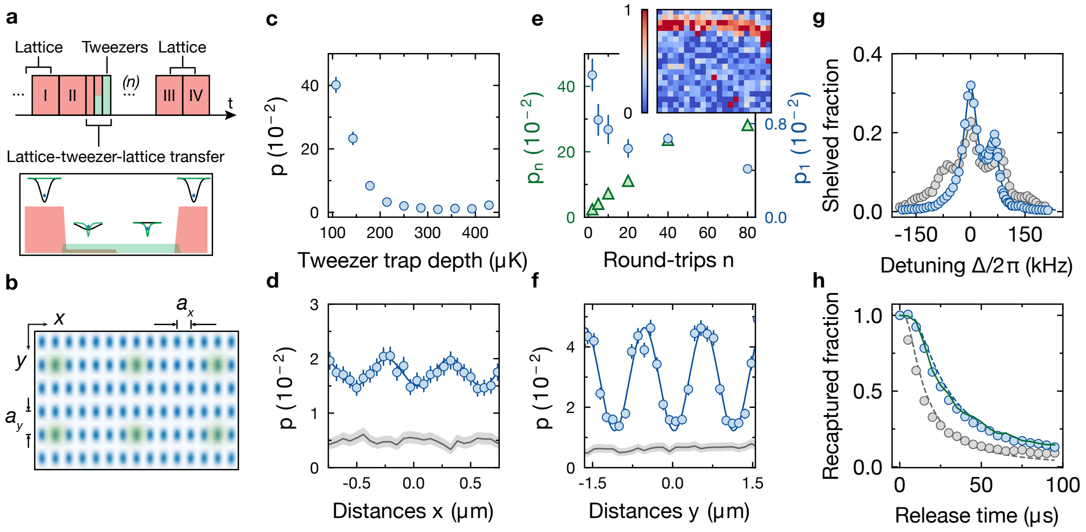

Figure 3.7: **(a)** Schematic of experimental sequence for tweezer-lattice-tweezer transfers. Axial potential landscape of the tweezers (green) and lattice (black) during hand-over is shown in the inset. The sketch is not to scale. **(b)** Sketch of the tweezer traps (green dots) and the lattice potential (blue dots) **(c)** Single-round-trip atom loss probability p as function of tweezer trap depth after transfer. The inset shows a close-up and confirms tweezer averaged single-round-trip losses close to 1%. **(e)** Cumulative loss probability p_n (green triangles) of atoms versus the number of lattice-tweezers-lattice round-trips n. The average single round-trip atom loss probability p_1 (blue round markers) decreases as n increases. The atom loss probability p_n shows a pronounced spatial dependence predominantly at the boundary of the lattice, as apparent from a measurement of the site-resolved atom loss probability after $n = 80$ round-trips (inset). **(d,f)** Single-round-trip atom loss probability p vs. relative position between lattice and tweezer potentials along x-axis and y-axis shown as blue points. The sinusoidal fit reflects the expected lattice potential with a lattice constant $a_x = 579(2)\,\mathrm{nm}$ and $a_y = 1187(18)\,\mathrm{nm}$. The atom loss probability due to imaging alone is indicated by shaded gray lines. **(g)** Radial sideband spectra of atoms loaded back into tweezers without (gray) and with (blue) sideband cooling after transfer. From the asymmetry of the spectrum, we extract ground-state fraction 0.98^{+0}_{-3} with and 0.37^{+40}_{-17} without sideband cooling after the transfer. **(h)** Release-and-recapture measurements in tweezers. We record the recovered fraction of atoms after a variable release time with (blue) and without (gray) sideband cooling after transfer. From a 3D Monte Carlo simulation, we extract temperatures of $5\,\mu\mathrm{K}$ ($17\,\mu\mathrm{K}$) with (without) sideband cooling after transfer, which is consistent with the temperature deduced from mean motional excitation measured by sideband spectroscopy. The fits are shown as blue and gray dashed lines respectively. The result of a release-and-recapture measurement before transferring to the lattice is indicated in green.

3.3 High-fidelity detection of large-scale atom arrays

Optical lattices are highly scalable trapping potentials for quantum science due to the high density of uniformly-packed trapping sites created from optical interference. In addition, because of the low-NA requirement for lattice optics, the lattice generally exhibits low wavefront error, resulting in well-defined trapping potentials free from small-lengthscale defects. To fully take advantage of lattice's scalability, our method of loading the lattice using tweezers is clearly not ideal, due to the significantly limited array size achievable with tweezers. Instead, a more promising approach is to load the lattice directly from a magneto-optical trap which can readily supply millions of cold atoms.

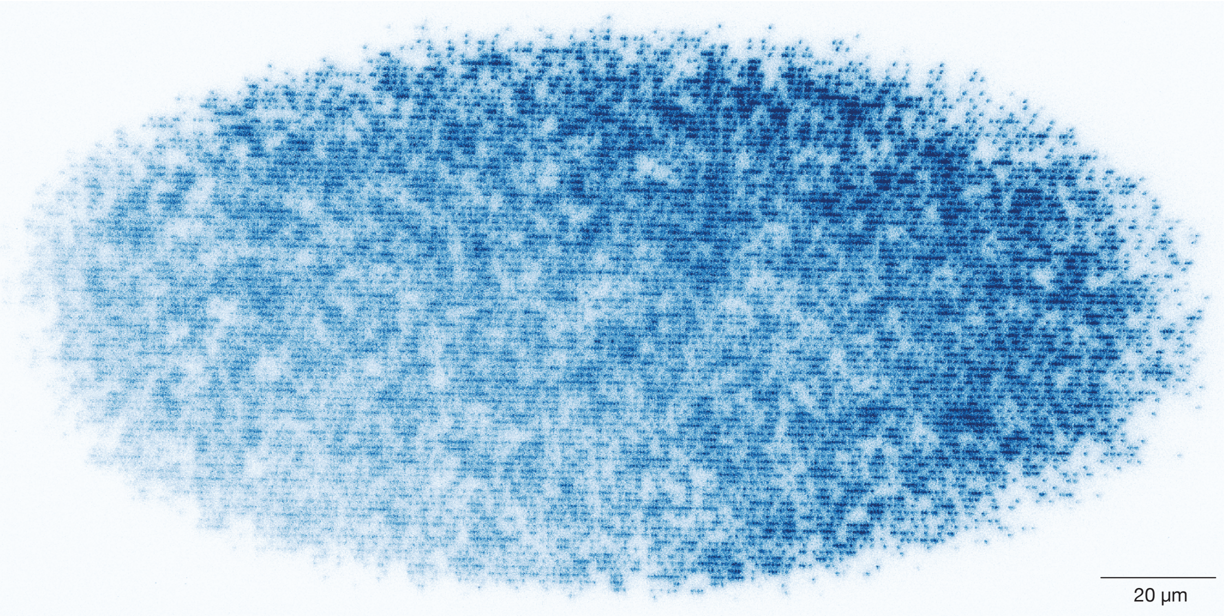

Figure 3.8: A single-shot fluorescence image of > 10000 parity-projected atoms directly loaded in the lattice.

3.3.1 Lattice state reconstruction

To accurately resolve atom occupation at single-site resolution on the lattice grid, it is necessary to apply deconvolution to the image using the known lattice geometry. The deconvolution algorithm we use is a simple, projection-based linear routine that scales favorably to large systems [113, 114]. The algorithm requires as inputs an estimate of the point spread function (PSF) and the lattice geometry (spacing and angle). For the former, we take advantage of our unique ability to load isolated atoms using tweezers and derive a PSF by averaging over individual fluorescence signals. For the latter, see Fig. 3.5. Once the inputs are provided, the algorithm applies a deconvolution kernel, an inverse of the PSF, to the lattice image to identify the fluorescence at each site (Fig. 3.9b,c). Finally, occupation is determined by thresholding the local fluorescence to produce a binary site occupation map (Fig. 3.9d).

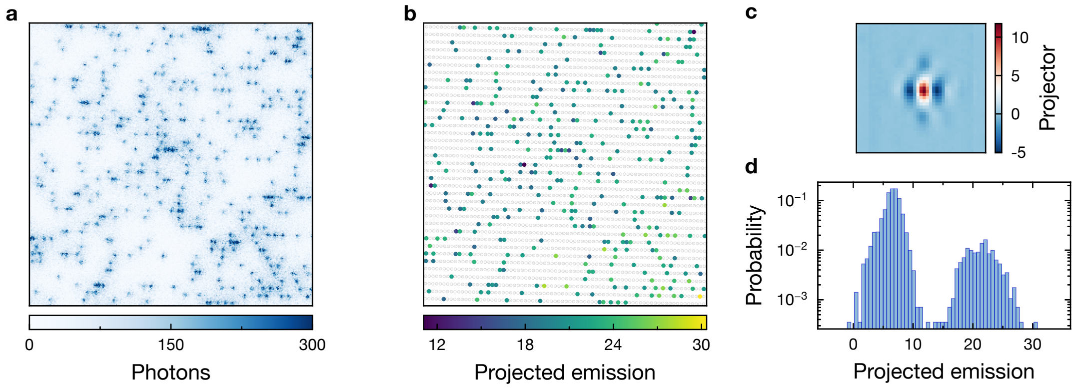

Figure 3.9: **(a)** Example raw fluorescence image of atoms directly loaded from the magneto-optical trap. **(b)** Site-resolved projected emission (trivially scaled down by 10^3) of the lattice image, derived from calculating the scalar product between the projector and raw fluorescence signal. **(c)** Projector used to compute the local fluorescence counts. **(d)** The emission histogram shows clear bimodality, indicating high classification fidelity. The loading fraction is about 13% for this dataset.

3.3.2 Parity projection mechanism in lattice

In loading atoms into a trapping potential from a magneto-optical trap, the number of atoms loaded onsite generally follows Poissionian statistics. The physics of parity projection concerns mechanisms that reduce the atom number in a trap to either zero or one. Common implementations include collisional blockade in tightly focused optical tweezers [115] and light-assisted collisions [116]. For many applications in quantum science, parity projection is an essential ingredient. For example, in quantum information encoding qubits using atoms, ensuring binary (zero or one) atom occupation in traps directly enables single-qubit addressability and high-fidelity multi-qubit entanglement via long-range interactions. Parity projection also plays a crucial role in high-fidelity fluorescence imaging. In systems with binary occupation, the fluorescence signal from a single atom can be well separated from the background signal corresponding to an empty site, for example by increasing the exposure time. As a result, classification fidelity exceeding 99.9% [117, 118, 119] are routinely demonstrated. However, if more atoms are involved, distinguishing N from $N+1$ atoms becomes substaintially more difficult, as the fluorescence histograms increasingly overlap at large N, due to the photon count fluctuations that scale as $\sqrt{N}$. Consequently, the classification fidelity is significantly reduced [120, 121].

In our infrared lattice, we observe a parity projection mechanism under the imaging conditions. For benchmarking, the lattice is always loaded from a SWAP MOT operated on the 689 nm transition [122]. The SWAP MOT is loaded from a blue MOT operated on the

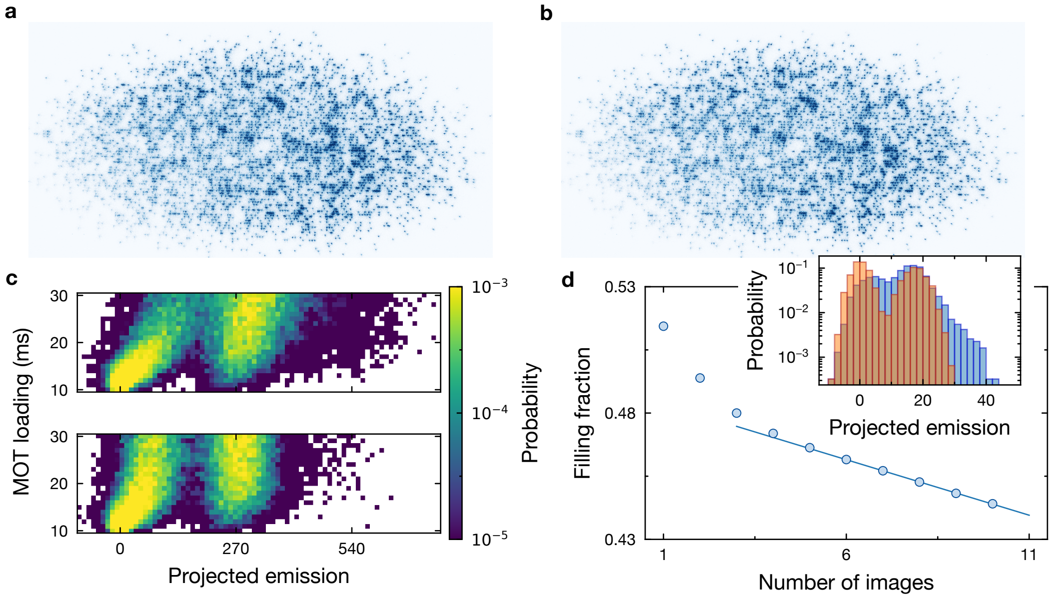

Figure 3.10: **(a, b)** Two consecutive images of the lattice directly loaded from a MOT, used to benchmark parity projection. **(c)** The detected counts versus MOT loading time for two consecutive images reveal the effect of parity projection during an exposure of 1.8 s. This is signaled by the absence of high-count histogram tails in the second image at long loading times. At a MOT loading time of 30 ms, the filling fraction is 0.18. Double occupancy accounts for less than 0.3% of all emission events at this loading time. **(d)** In a more densely loaded lattice, the filling fraction decreases with the number of consecutive images (900 ms each), following two distinct rates. The initial rapid decay (2×10^{-2}/image) is attributed to parity projection, while the slower decay (4×10^{-3}/image) at later times reflects single-particle loss from imaging. In the inset, emission histograms of the first and last image are compared, indicating the removal of events in which more than one atom are loaded.

461 nm transition. We vary the blue MOT loading time, effectively controlling the loading fraction in the lattice, and analyze the emission statistics using two consecutive images (Fig. 3.10a,b) under repulsive Sisyphus cooling conditions. At short loading times, lattice occupation is largely binary, as indicated by two well-resolved fluorescence peaks in the histogram corresponding to background and single-atom events (Fig. 3.10c). As the loading time increases, we observe emission events in the first image with counts higher than those expected for single-atom occupation, signaling the onset of multiple-atom loading. In the second image, however, these high-count events are effectively removed, resulting in well-separated bimodal histograms similar to those seen at short loading times. This observation points to parity projection occurring in the system. While this analysis is performed on a small patch, similar behavior is observed across the entire lattice. At the longest MOT loading time used in the dataset, the filling fraction is approximately 0.18. To quantify the timescale of the parity projection process, we substantially increase the

MOT loading time to 1.8 s and analyze the change in filling fraction as a function of the number of consecutive images taken (Fig. 3.10d). Each image has an exposure time of 900 ms. We observe that during the first three imaging cycles, the filling fraction drops at a significantly higher rate than during later cycles, suggesting that different loss mechanisms are at play in two distinct regimes. A comparison of the emission histograms from the first and last images indeed shows the disappearance of multiple-atom loading events (Fig. 3.10d,inset). Combined with the unique three-dimensional confinement provided by a single lattice beam, this parity projection behavior demonstrates our ability to directly load a large-scale atomic ensemble into a single plane of an optical lattice with singly-occupied lattice sites.

3.3.3 Scalable imaging in the lattice

Having demonstrated high-fidelity imaging in Section 3.2.2 on a small-scale atom array in the lattice, we now focus on characterizing imaging performance at the scale of the entire lattice loaded from the MOT. New challenges arise with upscaling. The first is the fluorescence crosstalk from atoms populating extremely close to each other due to the stochastic loading. In addition, various non-uniformities become more pronounced in a large system, complicating any global imaging optimization routine. Nevertheless, we will show high-fidelity imaging of more than 10^4 lattice sites and comment on the path to further improvement.

The benchmarking sequence begins with MOT loading and parity projection, as described previously. Afterwards, we take two consecutive images, and perform a state reconstruction of the full lattice. To mitigate various inhomogeneously broadening effect on the lengthscale of the lattice, we divide the lattice area into a grid of small patches (10×5 sites each) and calculate all relevant quantities in a patch-resolved manner (Fig. 3.11a). One of the strongest non-uniformity effects manifests in the mean emission count, derived from fitting the one-atom peak in the emission histogram (Fig. 3.11b). A likely origin is intensity modulation of the 461 nm imaging beam, induced by spurious reflections from the uncoated glass cell, as the modulation moves with the imaging beam. From the coarse-grained emission map, we observe an almost 100% relative difference in fluorescence across the lattice. Such variance significantly complicates the global optimization of the scattering rate on the 461 nm transition. Consider patches $(i, j) = (0, 0)$ and $(22, 1)$ as examples, receiving the least and most blue light, respectively. If the scattering rate is set such that patch $(0, 0)$ emits enough photons for high-fidelity classification, patch $(22, 1)$ would scatter nearly twice as many photons as necessary, potentially leading to increased imaging loss.

To quantify classification fidelity, we fit the background and one-atom peak in the emission histogram. We observe a strong correlation between classification fidelity and loading (Fig. 3.11c,d), suggesting fluorescence crosstalk as a potential cause of reduced detection fidelity. Indeed, within the same lattice patch, classification fidelity can be significantly improved by reducing the loading (not shown), thereby excluding lattice-related effects such as the local aberration as the primary explanation. In addition to filling fraction, classification fidelity also correlates with fluorescence inhomogeneity, as evidenced by the skewness both quantities exhibit toward the $(0, 0)$ corner of the lattice. Finally, we benchmark the survival probability between the two images using local thresholding. The coarse-grained imaging loss spatially traces the lattice potential contour, which is consistent with a spatially varying cooling rate that scales with the local lattice trap depth. To improve the imaging survival further, a flatter or deepter lattice potential can be used.

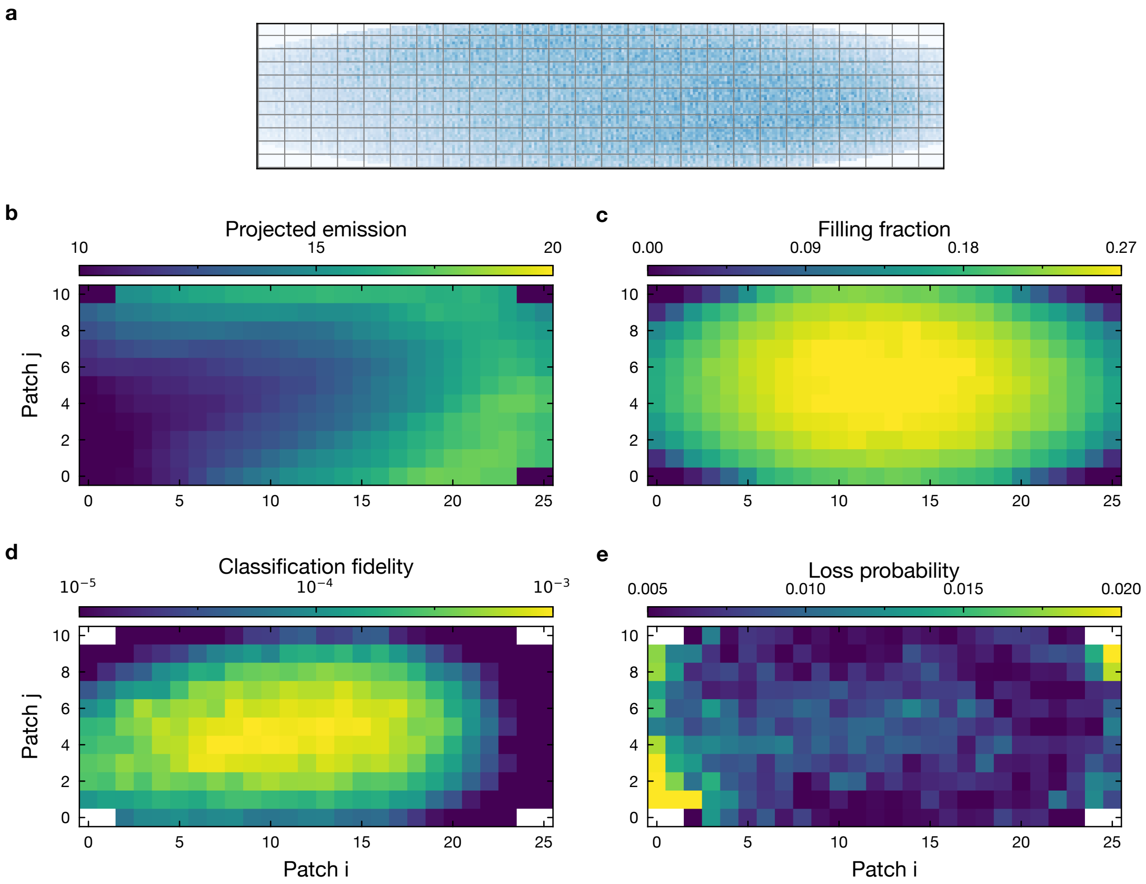

Figure 3.11: **(a)** To mitigate various inhomogeneity effects, imaging analysis is performed on small patches (10×5 sites) into which the full lattice is divided. The exposure time was set to 900 ms for the following four panels below. **(b)** The coarse-grained fluorescence emission is derived by fitting the center of the single-atom distribution in the local histogram. The large peak-to-valley emission difference is the main motivation for local analysis. The lattice filling fraction in **(c)** spatially correlates with the classification fidelity, indicating fluorescence crosstalk in the lattice under dense loading conditions. **(d)** The coarse-grained imaging loss probability reveals the lattice potential landscape, which may originate from lattice potential inhomogeneity.

Chapter 4

Universal quantum operation on metastable qubits in Strontium-88

4.1 Metastable qubits

The long-lived, metastable atomic states have long played a central role in metrological applications involving alkaline-earth(-like) atomic species such as strontium and ytterbium. The most prominent example is timekeeping based on the ultra-narrow $^1\mathrm{S}_0 \rightarrow {}^3\mathrm{P}_0$ transition in strontium, which serves as the foundation for state-of-the-art optical lattice clocks [123, 124]. More recently, metastable states have started to be explored for quantum information processing, motivated by a set of unique advantages not typically available in other qubit platforms [125]. For instance, metastable states enable fast, single-photon excitation to strongly interacting Rydberg states, based on which high-fidelity entanglement between qubits can be engineered. Additionally, since the absolute ground state lies outside the qubit subspace, mid-circuit measurements can be performed without disturbing the qubit. This property is leveraged in the implementation of erasure conversion, which allows for the identification of leakage errors within a quantum circuit with both temporal and spatial resolution. As a result, this technique has the potential to significantly reduce the resource overhead, particularly the number of physical qubits required to construct logical qubits in fault-tolerant quantum computing [126].

Metastable qubits were first demonstrated in trapped ions [132] and later adapted to neutral atoms [125]. In neutral atoms, high-fidelity universal quantum operations and mid-circuit erasure conversion have been demonstrated in ^{171}Yb, utilizing the nuclear spin degree of freedom in the metastable $^3\mathrm{P}_0$ state. In contrast, similar demonstrations in strontium are still lacking. In ^{87}Sr, the closest fermionic analogue to ^{171}Yb, the nuclear spin $I = 9/2$ presents challenges for isolating a well-defined qubit manifold from the rest of the hyperfine structure [133]. In contrast, bosonic isotopes of strontium have zero nuclear spin, precluding the implementation of nuclear spin qubits. However, qubits can be encoded in the two metastable states $^3\mathrm{P}_0$ and $^3\mathrm{P}_2$, which are separated by 17 THz. Like nuclear spin qubits, it supports erasure conversion and single-photon Rydberg excitation. But

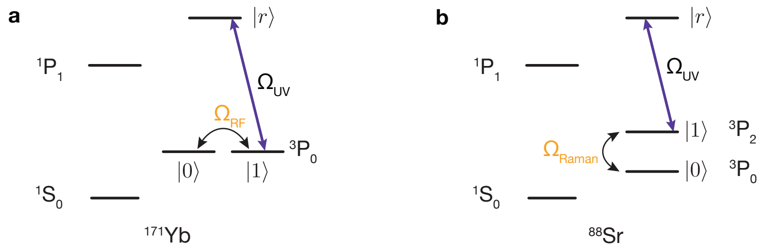

Figure 4.1: Cmmon metastable qubits encoding in alkaline-earth species. **a.** In ^{171}Yb, qubits can be encoded in two hyperfine states of the $^3P_0$ manifold, as a result of the nuclear spin 1/2 [125, 127, 128]. Microwave and two-photon Raman process can be used to drive single-qubit rotations. **b.** In ^{88}Sr, qubits can be encoded in two long-lived clock states $^3P_0$ and $^3P_2$ [129, 130, 131]. Due to the large energy splitting, optical Raman transitions are used to drive single-qubit rotations. In both ^{171}Yb and ^{88}Sr, entanglement between qubits can be induced by coupling of of the qubit states $|1\rangle$ to a highly-energetic Rydberg states $|r\rangle$ via a single-photon transition at ultraviolet wavelength.

additionally, the large energy splitting between the two qubit states enables high-fidelity state preparation and read-out. In the following, we first provide a detailed overview of the relevant background and experimental results concerning state preparation, read-out, and coherent manipulation of individual qubits. We then discuss two-qubit conditional gates mediated by Rydberg interactions and demonstrate the implementation of universal quantum operations.

4.1.1 Qubit preparation

Coherent excitation on the clock transition

To prepare the population in the qubit subspace, we drive transition on the $^1S_0 \rightarrow {}^3P_0$ clock transition. In bosonic species, this transition has extremely small linewidth due to the violation of selection rule in ΔJ and ΔS. One way to drive the clock transition is by perturbatively admixing $^3P_1$ state into $^3P_0$ via a magnetic field $\mathbf{B}$ [134]. As a result, the Rabi frequency acquires a linear dependency on $|\mathbf{B}|$ (see Fig. 4.2a). Due to the off-resonant coupling to $^3P_1$, the clock resonance shifts as $|\mathbf{B}|^2$, which can be conveniently exploited to calibrate the conversion between the current applied in the bias field coil and the actual magnetic field strength experienced by the atoms (see Fig. 4.2b). For maximal coupling strength, the clock beam has a linear polarization aligned to the bias magnetic field.

At the typical magnetic field strengths and optical intensities available in the lab, the Rabi frequency for clock state excitation is on the order of a few kilohertz. Consequently, the shelving process is particularly sensitive to motion of atoms in a trap. To mitigate this thermal dephasing, the atom is cooled on the $^1S_0 \rightarrow {}^3P_1$ intercombination transition in the resolved sideband regime. At the 813 nm magic wavelength for the clock transition, the cooling transition can also be rendered magic through polarizability tuning [105]. Once a high ground-state population is achieved, the magnetic field is ramped up to 420 G in

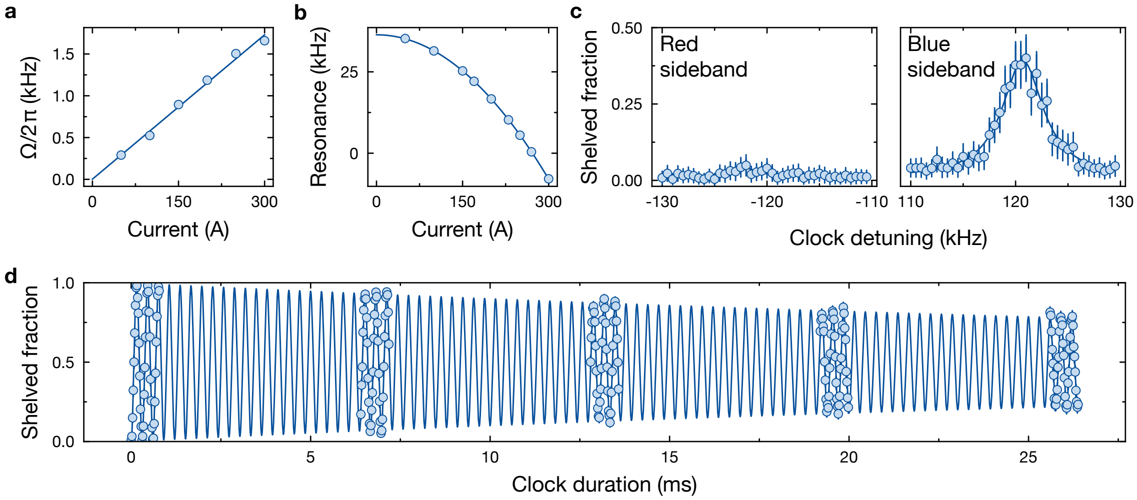

Figure 4.2: **(a)** The Rabi frequency and **(b)** resonance frequency of the $^1S_0 \rightarrow {}^3P_0$ clock transition as a function of the current of the magnetic bias field current. The quadratic dependence of the resonance shift on the bias field strength is used to calibrate the current-to-field conversion ratio. Due to the Doppler sensitivity at kilohertz-scale Rabi frequencies, ground-state cooling is essential for achieving high-fidelity population transfer. With a sideband cooling stage targeting the intercombination line, high ground-state fraction can be reached as shown in **(c)** obtained using the same clock transition. In the sideband spectra, the carrier transition (not shown) is centered at 0 kHz, with the trap frequency at $2\pi \times 122$ kHz. **(d)** Rabi oscillations on the clock transition exhibit approximately 150 coherent cycles at a Rabi frequency of $2\pi \times 3$ kHz.

preparation for clock state excitation. The clock laser itself is used to probe the atomic temperature via sideband thermometry, particularly along the direction of the clock beam (Fig. 4.2c). With a sufficiently narrow-linewidth laser (see 2.3.2), highly coherent Rabi oscillations can be observed (Fig. 4.2d).

4.1.2 Erasure conversion

Encoding qubits in metastable states enables a new class of hardware-level error mitigation known as *erasure conversion*. The key idea is thats common gate errors result in population leakage into a disjoint subspace that can be continuously monitored [135]. As a result, errors can be timely detected and addressed before they accumulate into more complex, uncorrectable faults in fault-tolerant quantum computing. This approach is expected to significantly reduce the overhead in the number of physical qubits required to implement a logical qubit [126]. In alkaline-earth(-like) atoms, the $^1S_0$ ground state serves as the primary reservoir state for erasure-type leakage errors, owing to its strong $^1S_0 \rightarrow {}^1P_1$ cycling transition. Under optimized imaging conditions, the presence of population in $^1S_0$ can be detected within a few tens of µs, without perturbing the remaining qubit states [125, 136].

In this section, we introduce the fast imaging technique relevant to erasure conversion and demonstrate high-fidelity qubit state preparation using this method.

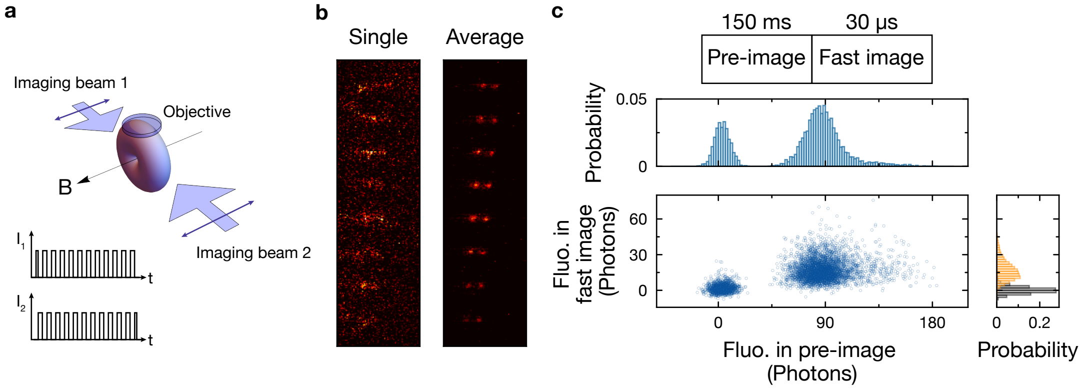

Figure 4.3: **(a)** The layout of the counter-propagating resonant imaging beams, the magnetic field with strength **B**, the entrance pupil of the objective, and the pulsed intensity waveform of the imaging beam during the fast imaging. **(b)** Representative single-shot and averaged fluorescence images obtained in a ∼ 30 μs fast imaging stage. **(c)** A non-destructive, high-fidelity pre-image is employed to benchmark the performance of the fast imaging. Since both images correspond to the same experimental loading event and the pre-image offers significantly higher fidelity, discrepancies in fluorescence-based classification between the two images serve as a measure of fast imaging accuracy. In the histogram on the right, fluorescence counts are color-coded (yellow: atom loaded; black: no atom) according to the atom presence determined from the pre-image.

Conventional fluorescence imaging is poorly suited for erasure detection due to its long acquisition times. Recent work has demonstrated high-fidelity, non-destructive imaging in as little as 5 ms [137]—a substantial improvement, yet still several orders of magnitude longer than state-of-the-art gate times. To scatter a sufficient number of photons before the atom is lost from the trap, imaging must be interleaved with cooling. As a result, the total imaging duration is fundamentally limited by the cooling rate. However, if the radiation pressure from a single imaging beam is balanced by a counter-propagating beam with otherwise identical parameters, atomic motion can be much more effectively constrained [138, 120]. Importantly, the imaging transition can be driven near saturation to maximize the scattering rate, since there is no net force along the imaging beam's **k**-vector in the balanced configuration.

In our setup, a 19 G magnetic field **B** defines the quantization axis (Fig. 4.3a) for the qubit operation. A pair of linearly polarized beams resonantly addresses the π-transition on the $^1S_0 \rightarrow {}^1P_1$ line. An objective positioned at 90° with respect to the magnetic field axis collects photons with an expected efficiency of 0.16. To suppress coherent effects in the strong driving regime [120], the two imaging beams are intensity-chopped out of phase with each other. We collect all the photons in ∼30 μs before atoms are irretrievably lost and image them onto a qCMOS camera (Fig. 4.3b). Despite the short photon emission

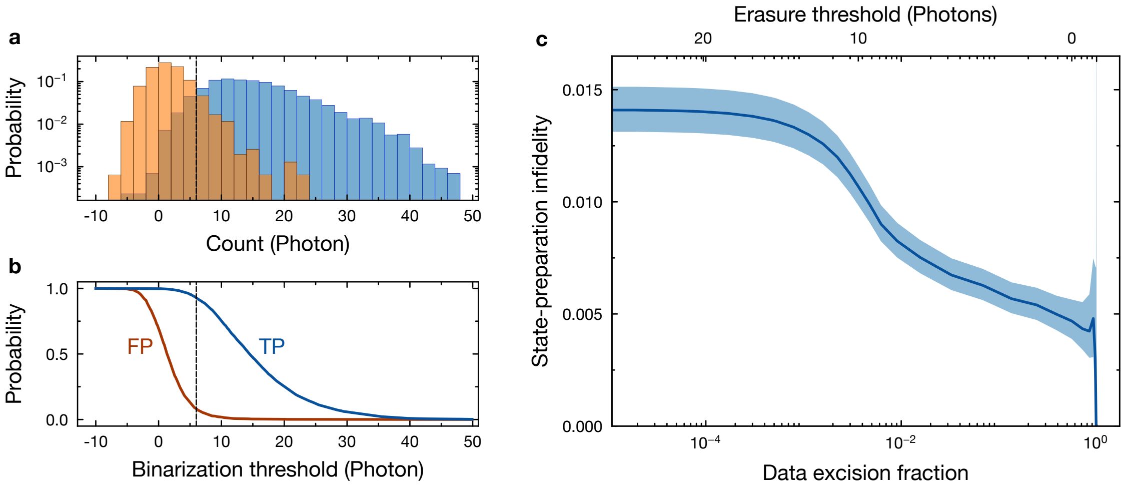

Figure 4.4: **(a)** The histogram of photon counts from an erasure-detection image, binarized using a high-fidelity, low-loss pre-image. The orange bars correspond to counts when no atom was present in the pre-image, and the blue bars correspond to counts when an atom was present. We make a binary classification of whether an atom is detected in an erasure image by comparing the analog fluorescence counts against a threshold (dashed). **(b)** In choosing an appropriate threshold, our primary figure of merit is the true positive (TP, blue) probability, which measures how often we correctly identify events with leakage errors. Lowering the threshold captures a larger fraction of true positives but also increases the false positive (FP, red) probability, which is how often we misclassify events without leakage as if they had leakage. An usually optimal threshold maximizes the classification fidelity (dashed). **(c)** In converting state-preparation error to erasure errors, we see the balancing between lower false positive and higher true positive probability by computing the state-preparation infidelity at various data excision fractions. At a low erasure threshold (i.e. 20 photons), most erasure images appear dark, indicating no leakage error in state-preparation. As a result, few experimental shots are excised and the state-preparation infidelity remains the same as when no erasure conversion is applied. Conversely, at a high erasure threshold (i.e. 5 photons), most leakage events produce bright erasure images. Consequently, removing all shots flagged as erasures lowers the apparent state-preparation infidelity.

window, the actual exposure time of the camera is at least 4 ms, limited by the rolling shutter architecture of the camera (Orca C15550-20UP). To benchmark the classification fidelity of the fast imaging, we acquire a high-fidelity, non-destructive pre-image (Fig. 4.3c). From the fluorescence signal in the pre-image, the presence or absence of an atom can be reliably determined. Given the high survival probability associated with the pre-image, the same atom occupancy should ideally be inferred from the subsequent fast image. Any discrepancy between the two classification results is then used to quantify the fidelity of the fast imaging protocol. At the 50 µK trap depth commonly used for our metastable qubits, the fast imaging classification fidelity is approximately 0.96.

Fast imaging provides a straightforward technique for converting state preparation errors into erasures. During the qubit preparation stage of the experimental sequence, we apply a π-pulse under a 420 G magnetic field, ramp the field down to 19 G for operating the metastable qubit, and take a fast imaging on the $^1S_0$ state. The state preparation error is typically around 1–2%, primarily arising from trap-induced Raman scattering during the long (20 ms) idle period following the magnetic field ramp, as well as imperfections in the π-pulse itself. Crucially, population leakage occurs almost exclusively into the $^1S_0$ ground state, which appears bright in the erasure image. By post-selecting on the data which do not have atom detected in the erasure image, the state preparation fidelity can be increased.

To further enhance the fast imaging fidelity, which has already exceeded 0.993 in other work [120], several potential improvements can be considered. First, the qCMOS camera used in this experiment is limited by its rolling-shutter operation and minimum effective exposure time, leading to unnecessary collection of background stray light. An alternative camera technology, such as an EMCCD, could circumvent this issue by allowing the sensor to be exposed only for the duration of the fast imaging sequence. Second, additional fluorescence from the second objective in a dual-objective setup can be retroreflected into the first objective and overlapped with the primary signal for higher signal-to-noise ratio.

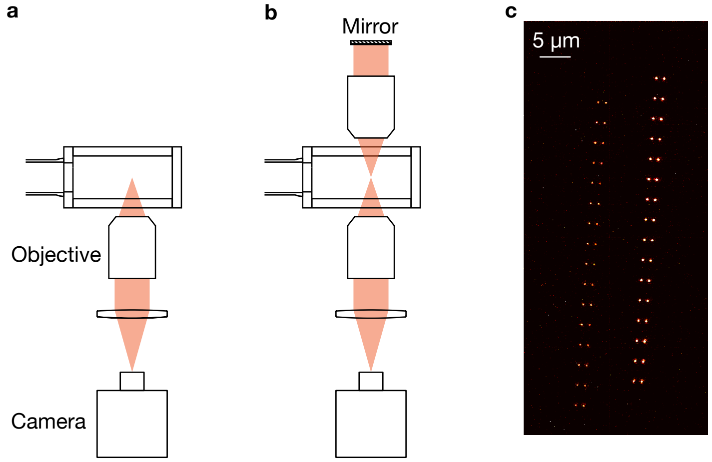

Figure 4.5: **(a)** Conventional single-objective imaging system. **(b)** With an additional objective, twice as much photons can be collected (neglecting loss) by retroreflecting light collected by the second objective into the first one. Note that additional inversion is required to overlap each tweezer to its image. **(c)** Averaged fluorescence image contains primary and secondary image using the setup in **(b)**. Due to optical losses in the additional optical elements, the secondary image has approximately half the intensity of the primary image.

4.1.3 Read-out

State-sensitive read-out

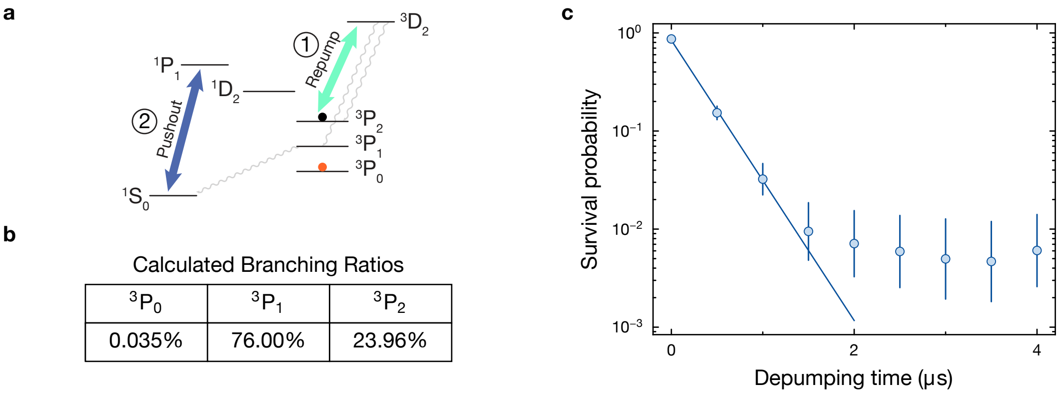

Calculated Branching Ratios

$^3P_0$	$^3P_1$	$^3P_2$
0.035%	76.00%	23.96%

Figure 4.6: **(a)** The level diagram relevant for state-sensitive detection of metastable qubits. Population in $^3P_2$ (black circle) is first incoherently transferred to short-lived $^3P_1$ by strongly driving on the $^3P_3 \rightarrow {}^3D_2$ transition. **(b)** The favorable branching ratio of $^3D_2$ ensures an averaged scattering of 1.315 photon required to exit the repumping cycles [139]. Once atoms decay to the $^1S_0$, a resonant pushout follows to remove atoms from the trap. The remaining population is then that in $^3P_0$ (red circle). **(c)** The $^3P_2 \rightarrow {}^3P_1$ depumping has a $1/e$ time of 303.9(95) ns in the qubit operating condition (19 G bias field). The residual survived fraction can be attributed to imperfect $^3P_2$ state preparation and leakage into $^3P_2$ via $^1P_1 \rightarrow {}^1D_2 \rightarrow {}^3P_2$ pathway during the $^1S_0$ pushout.

To read out the qubit state, we map population in the $^3P_2$ state to atom loss and image the remaining atoms in $^3P_0$ using standard cooling-assisted fluorescence imaging. As the first step to inducing loss of $^3P_2$ atoms, we strongly drive the $^3P_2 \rightarrow {}^3D_2$ repumping transition at 497 nm (Fig. 4.6a). Due to favorable branching ratio (Fig. 4.6b), atoms undergo in average 1.3 scattering events before predominantly decaying from $^3D_2$ to the short-lived $^3P_1$ state, with only a 4.6×10^{-4} probability of branching into the long-lived $^3P_0$ state, which contributes a small but non-negligible bit-flip error. Given the 21 μs lifetime of the $^3P_1$ state, atoms rapidly decay to the $^1S_0$ ground state and are subsequently expelled from the trap via resonant scattering on the $^1S_0 \rightarrow {}^1P_1$ transition at 461 nm. Since the 497 nm transition is a strong dipole transition with a natural linewidth of $2\pi \times 9.8$ MHz, the $^3P_2$ population can be removed from the qubit subspace on a sub-μs timescale (Fig. 4.6c), significantly faster than most error processes (i.e. Raman scattering induced by the trap). For more efficient removal of atoms in $^3P_2$, it is worth noting that repumping on the $^3P_2 \rightarrow {}^3D_3$ transition—detuned by approximately 0.6 nm from $^3D_2$—would be both faster and simpler. The $^3D_3$ state decays back to $^3P_2$ with a branching ratio of 99.93% [139], making this transition nearly closed and thus highly effective for cycling population out of the trap. However, this transition unlike $^3P_2 \rightarrow {}^3D_2$ is not compatable with state-resolved read-out to be introduced next.

State-resolved read-out

While simple and effective, our state-sensitive readout cannot distinguish atoms occupying the $^3P_2$ state from those that were already lost prior to readout. This limitation is particularly relevant given that atom loss is a common and intrinsic feature of neutral-atom quantum computing architectures, arising from unavoidable processes such as background gas collisions in vacuum (Fig. 2.2) or decay from Rydberg states into dark states (Fig. 4.14c). To overcome this limitation, the key is to determine the occupation of both qubit states with high fidelity using spin-resolved detection (SRD). In alkali species, direct imaging of both qubit states has been achieved by spatially separating spin states in strong magnetic field gradients [140, 141] or state-dependent traps [135, 32]. In alkaline-earth(-like) atoms, the SRD was recently demonstrated by imaging on a narrow-line transition and shelving in metastable clock states [137, 142].

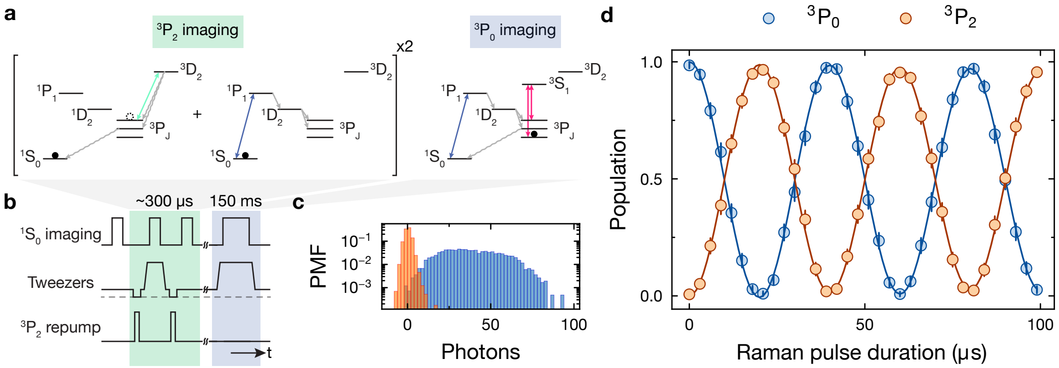

Figure 4.7: **(a)** Internal-state dynamics relevant to the detection of $^3P_2$ (green) and $^3P_0$ (blue). Prior to detecting $^3P_2$, population in this state is incoherently transferred to $^1S_0$ via repumping on the $^3P_2 \rightarrow {}^3D_2$ transition. Then, we perform fast, destructive imaging in deep tweezers to collect sufficient photons for high-fidelity $^3P_2$-state reconstruction. Such a state-transfer plus imaging cycle is repeated to ensure complete removal of $^3P_2$ population, despite small leakage from $^1P_1$ via $^1D_2$ into $^3P_2$. Once $^3P_2$ is destructively imaged, the remaining population is assigned to $^3P_0$ using a subsequent slow image with repumping on the 679 nm and 707 nm transitions. **(b)** Timing diagram (not to scale) for the $^1S_0$ imaging beam, $^3P_2$ repumper, and tweezer trap depth during the state-selective detection. Traps are blinked off (dashed) during the $^3P_2$ repumping pulse to avoid photo-induced losses. **(c)** Probability mass function (PMF) of photons emitted by $^1S_0$ atoms transferred from $^3P_2$ recorded during fast imaging. The resulting histogram, binarized using a high-fidelity pre-image, shows a classification fidelity > 0.993 averaged over all tweezers at the optimal threshold. The orange and blue correspond to zero- and one-atom peak, respectively. **(d)** Rabi oscillation between metastable qubit states with full state resolution. The state preparation errors are removed via additional erasure detection.

Here, we demonstrate an alternative approach to SRD, based on the spin-sensitive detection introduced earlier (Fig. 4.7a,b). First, the incoherent $^3P_2 \rightarrow {}^1S_0$ state transfer is done in free space, to avoid any photo-induced loss in the trap. Second, the $^1S_0$ pushout is replaced by a fast imaging stage in deep traps, to reveal the occupation of $^1S_0$ state. The increased trap depth (1.5 mK) allows the atoms to survive longer and scatter sufficient photons to achieve a classification fidelity of 0.993 (Fig. 4.7c). Note that since $^1P_1$ is not fully closed, a small fraction of population can leak into $^3P_1$ and $^3P_2$ via $^1D_2$. To image the residual leakage, we repeat the incoherent state transfer plus fast imaging cycle. After two complete cycles, the population originally in $^3P_2$ is fully identified and removed from the trap. The remaining population in $^3P_0$ is finally imaged using the standard cooling-assisted technique with repumping on 679 nm and 707 nm. The Rabi oscillation between metastable qubit states with full state resolution is shown in Fig. 4.7d. From the deviation of the total measured population from unity, we estimate an overall SRD fidelity exceeding 0.991, averaged over all possible distributions of population between the two states. One distinct advantage of this loss-based SRD is the speed with which state resolution can be reached which prevents bit-flip error from the imaging itself from happening. In principle, instead of transferring $^3P_2$ population to $^1S_0$ and performing fast imaging there, one could directly image $^3P_2$ population by scattering on the $^3P_2 \rightarrow {}^3D_3$ transition. Such a scheme avoid the overhead in the state transfer and is expected to yield higher imaging fidelity.

4.1.4 Single-qubit rotation

State manpulation by two-photon Raman process

At the highest level of abstraction, neglecting atomic structure and selection rules, the metastable qubit can be approximated by a textbook three-level Λ-system consisting of two long-lived ground states, $|g_1\rangle \equiv |{}^3P_0\rangle$, $|g_2\rangle \equiv |{}^3P_{2,m_j=0}\rangle$, and a short-lived excited state $|e\rangle \equiv |6s\ {}^3S_1\rangle$. Two laser fields couple $|g_1\rangle \rightarrow |e\rangle$ and $|g_2\rangle \rightarrow |e\rangle$ with respective Rabi frequencies Ω_1 and Ω_2, both detuned from resonance by a common detuning Δ. Under appropriate condition ($\Delta \gg \Omega_i$), the excited state $|e\rangle$ can be adiabatically eliminated, yielding an effective Hamiltonian [143]

$$H_{\text{eff}} = -\hbar \begin{pmatrix} \frac{|\Omega_1|^2}{4\Delta} & \frac{\Omega_R^*}{2} \\ \frac{\Omega_R}{2} & \frac{|\Omega_2|^2}{4\Delta} \end{pmatrix}, \qquad \Omega_R \equiv \frac{\Omega_1 \Omega_2^*}{2\Delta} \tag{4.1}$$

that governs coherent dynamics within the $\{|g_1\rangle, |g_2\rangle\}$ subspace. This forms the basis of the coherent two-photon Raman process that enables population transfer within the metastable manifold. However, since the elimination of $|e\rangle$ is only approximate, residual coupling to the excited state introduces an incoherent scattering channel governed by its natural linewidth Γ. While the exact strength of this off-resonant scattering γ_{sc} depends on the population distribution among $|g_1\rangle$ and $|g_2\rangle$, it generally scales as [144]

$$\gamma_{\text{sc}} \sim \frac{\Gamma}{4\Delta^2}(|\Omega_1|^2 + |\Omega_2|^2). \tag{4.2}$$

This simplified treatment provides the minimal theoretical background necessary to understand the coherent manipulation of metastable qubits.

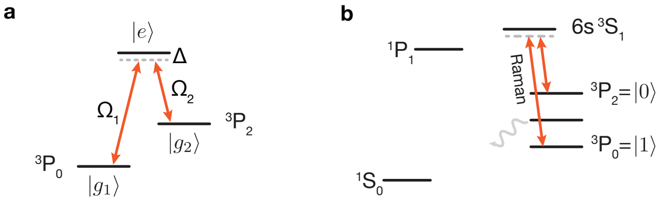

Figure 4.8: **(a)** We consider a three-level Λ-type system comprising two metastable ground states, $|g_1 \equiv {}^3\mathrm{P}_0\rangle$ and $|g_2 \equiv {}^3\mathrm{P}_{2,0}\rangle$, coupled to a common excited state $|e\rangle \equiv |6\mathrm{s}\ {}^3\mathrm{S}_1\rangle$. Each ground state $|g_i\rangle$ is coupled to $|e\rangle$ with a Rabi frequency Ω_i with a common detuning Δ. **(b)** The level diagram of low-lying states in ^{88}Sr relevant to qubit state manipulation is shown. Notably, the qubit states are defined with an inverted mapping: $|0\rangle \equiv |{}^3\mathrm{P}_{2,m_j=0}\rangle$ and $|1\rangle \equiv |{}^3\mathrm{P}_0\rangle$. This choice is motivated by the availability of the Rydberg excitation laser, which can only be tuned to be resonant with the $^3\mathrm{P}_0 \rightarrow |r\rangle$ transition.

In strontium, the relevant atomic parameters are summarized in the table below [101]. Assuming a common beam profile of $(400\,\mu\mathrm{m}, 200\,\mu\mathrm{m})$ and an optical power of 20 mW for each leg of the Raman transition, we estimate a two-photon Rabi frequency of $2\pi \times 1\,\mathrm{MHz}$ and an off-resonant scattering rate of approximately 6 kHz at $\Delta = 2\pi \times 30\,\mathrm{GHz}$. Such a favorable scaling to faster state manipulation represents one strong advantage for metastable qubit in strontium.

Transition	Wavelength (nm)	Transition rate (10^6/s)	I_{sat} (mW/cm^2)
$^3\mathrm{P}_0 \rightarrow 6\mathrm{s}\ {}^3\mathrm{S}_1$	679	8.35	0.55
$^3\mathrm{P}_2 \rightarrow 6\mathrm{s}\ {}^3\mathrm{S}_1$	707	39.03	2.30

Rabi oscillation and coherence

The two-photon Rabi oscillation is shown in Fig. 4.9a. Despite having the optical power to reach higher Rabi frequency, we operate our Raman Rabi frequency at $2\pi \times 17\,\mathrm{kHz}$ due to a much reduced phase noise at lower frequency. To check the coherence property of the metastable qubit, we perform Ramsey-type experiments. We find a $T_2^* = 461\,\mu\mathrm{s}$ averaged over 81 tweezers (Fig. 4.9b). Such a coherence time is understood to be limited by an inhomogeneous broadening of residual lightshifts across all tweezers. One possible effect that can give rise to the non-unifomity is a local variation of trap polarization. As a result, magic condition in general acquires a dependency on tweezers. With XY-4 dynamical decoupling pulse, the T_2 coherence time can be extended to 344 ms in $5\,\mu\mathrm{K}$ traps. Another effect which limits coherence at the single-tweezer level is a position-dependent lightshift induced by polarization mixing near the focal plane of a tightly-focused optical beam [129]. This effect can be mitigated by lowering the atom's temperature.

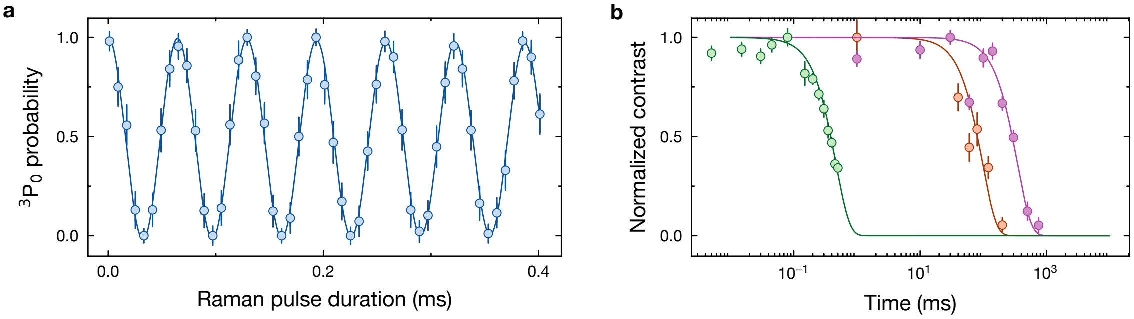

Figure 4.9: **(a)** The Rabi oscillation between metastable qubit states at $2\pi \times 15\,\mathrm{kHz}$ Rabi frequency. **(b)** The coherence property of the mestable qubit [131]. The Ramsey coherent time T_2^* is $463\,\mu\mathrm{s}$ (green). With XY-4 pulses in between two Ramsey $\pi/2$-pulse, the coherence time can be extended to 97 ms at $50\,\mu\mathrm{K}$ (red) and further to 344 ms at $5\,\mu\mathrm{K}$ trap (purple).

Randomized benchmarking of single-qubit gates

Having demonstrated coherent control of the metastable qubit states, we now proceed to benchmark the performance of single-qubit gates. In quantum computing, the standard pulse rotation angles are π and $\pi/2$, corresponding to full and half population transfers on the Bloch sphere, respectively. To minimize discretization errors in the single-qubit rotations, we apply a large number of consecutive pulses and optimize the return probability as a function of pulse duration. A typical calibration procedure begins by identifying the pulse duration (π time) required to implement a π-rotation (Fig. 4.10b). Next, the appropriate duration for a $\pi/2$-rotation is determined based on this reference (Fig. 4.10c). For convenience, we parametrize single-qubit rotations using (θ, ϕ), which defines the Bloch vector on a unit sphere. A resonant laser pulse implements the unitary

$$R(\theta,\phi) = \begin{pmatrix} \cos\frac{\theta}{2} & -ie^{i\phi}\sin\frac{\theta}{2} \\ -ie^{-i\phi}\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix}.$$

The single-qubit calibration effectively yields θ_π and $\theta_{\pi/2}$ with which we can construct $X \equiv R(\theta_\pi, 0)$ and $X/2 \equiv R(\theta_{\pi/2}, 0)$.

It is worth highlighting the technique of the *virtual* phase gate, which is commonly used in experimental quantum control [145]. This method relies on the identity:

$$R(\theta,\phi) = Z(-\phi)X(\theta)Z(\phi), \tag{4.3}$$

where $X(\theta) = R(\theta, 0)$ represents a rotation around the x-axis, and $Z(\phi) = \mathrm{Diag}(1, e^{i\phi})$ denotes a phase gate. Using this identity, a phase gate can be effectively implemented by updating the control frame (i.e., virtually), without the need to physically apply an additional pulse. For example, to implement the sequence $X(\theta_1)Z(\phi_1)X(\theta_2)X(\theta_3)$, one

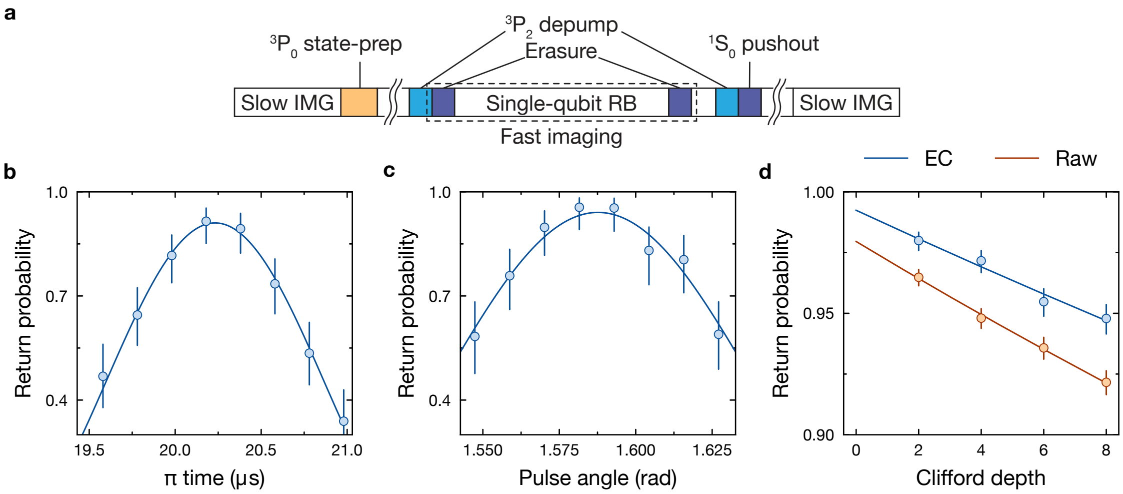

Figure 4.10: **(a)** Experimental sequence. The cooling-assisted, non-destructive imaging (slow IMG) at the beginning of the sequence verifies the loading of atoms at the target tweezer sites, and at the end reads out the $|1\rangle$ population following a $|0\rangle$ depletion stage (using $^3\mathrm{P}_2$ depumping and $^1\mathrm{S}_0$ pushout). Two erasure-conversion pulses placed before and after the single-qubit Clifford randomized benchmarking (RB) sub-sequence remove state-preparation error in $^3\mathrm{P}_0$ and off-resonant scattering errors originating from $^3\mathrm{P}_1$. See main text for details. **(b)** Calibration of the π-pulse duration. A sequence of $N = 16$ consecutive pulses is applied while scanning the π-pulse time. The return probability is fit to the function $A\cdot\cos(\pi N t/2t_\pi)^2$. **(c)** Calibration of the π-half rotation angle. A sequence of $N = 32$ pulses is applied while scanning the pulse angle θ. The return probability is git to $A\cdot\cos(\pi N\theta/4\theta_{\mathrm{half}})^2$. **d.** The average Clifford gate fidelity raw (orange) and after erasure conversion (blue) from randomized benchmarking are 0.992(1) and 0.993(1), respectively.

can instead apply:

$$\begin{aligned}
&X(\theta_1)\cdot Z(\phi_1)X(\theta_2)Z(-\phi_1)\cdot Z(\phi_1)X(\theta_3)Z(-\phi_1)\cdot Z(\phi_1)\\
&= X(\theta_1)\cdot R(\theta_2,-\phi_1)\cdot R(\theta_3,-\phi_1)\cdot Z(\phi_1)
\end{aligned}$$

where the final equality follows from the identity in Eqn. (4.3). In this construction, the frame is effectively rotated after each virtual phase gate, and this frame shift "sticks" for subsequent operations. The final $Z(\phi_1)$ is trivial in the sense that it commutes with measurement in the computational basis and can be ignored in practice. The virtual phase gate has near-unit fidelity, as classical hardware is usually self-calibrated to a stable frequency reference [145]. However, it introduces compilation overhead due to the need for continuous tracking of the reference frame. Moreover, since the phase adjustment is global, it imposes limitations in more sophisticated quantum circuits where local or qubit-specific control is required.

To estimate the average single-qubit gate fidelity, we employ the randomized benchmarking (RB) protocol. Specifically, we focus on the Clifford group, which consists of all gates that map Pauli operators to other Pauli operators under conjugation. Clifford gates play a central role in quantum error correction and stabilizer codes. Moreover, circuits composed entirely of Clifford gates can be efficiently simulated on classical hardware, making them particularly useful for benchmarking gate performance. In the Clifford RB protocol, a sequence of m randomly chosen Clifford gates is applied:

$$C_m \cdot C_{m-1} \cdots C_1,$$

where each C_i is drawn uniformly at random from the Clifford group. Finally, an inversion gate C_{m+1} defined as

$$C_{m+1} = (C_m \cdot C_{m-1} \cdots C_1)^\dagger = C_1^\dagger \cdots C_{m-1}^\dagger \cdot C_m^\dagger$$

is applied at the $m+1$ step to ideally return the system to its initial state in the absence of gate errors. In actual hardware, the deviation of the measured outcomes from the ideal value is used to estimate the gate fidelity, averaged over many repetitions of the RB experiment.

We compile all Clifford gates using the native gate set $\left\{\pm\frac{X}{2}, \pm\frac{Y}{2}, X, Y\right\}$. The benchmarking sequence prepares and reads out the population in state $|1\rangle$ (Fig. 4.10a). Two detection stages before and after the single-qubit RB convert state preparation error and partial leakage error (off-resonant scattering into $^3\mathrm{P}_1$ and $^1\mathrm{S}_0$) into erasures which are subsequently excluded in the post-processing. We extract an average single-qubit Clifford gate fidelity $F_{1\mathrm{q}}$ of 0.992(1) (0.993(1)) without (with) erasure excision, averaged over 16 tweezer sites, at $2\pi \times 17$ kHz Rabi frequency, see Fig. 4.10d. Note that, in principle, the additional off-resonant scattering into $^3\mathrm{P}_{2,m_J\neq 0}$ can be converted to erasure events by two-photon coupling of $^3\mathrm{P}_2$ to the short-lived $^3\mathrm{P}_1$ state.

4.2 Two-qubit gate

4.2.1 Rydberg physics

Rydberg states are highly excited atomic states in which the valence electron occupies an orbital far from the ionic core. These states are characterized by a large principal quantum number, $n \gg 1$, and exhibit exaggerated atomic properties—such as large electric dipole moments, strong long-range interactions, and extended lifetimes—that make them especially valuable for applications in quantum simulation and quantum computation. In the following, we summarize key scaling relationships relevant to the entangling gate discussed later in this work.

The Rydberg spectrum closely resembles that of the hydrogen atom and is given by

$$E_{nlj} = -\frac{\mathrm{Ry}}{(n - \delta_{lj})^2}$$

where δ_{lj} is the quantum defect and Ry the Rydberg constant. The energy spacing between neighboring Rydberg states scales as $\Delta E_n \approx n^{-3}$, as illustrated in Fig. 4.11. The orbital

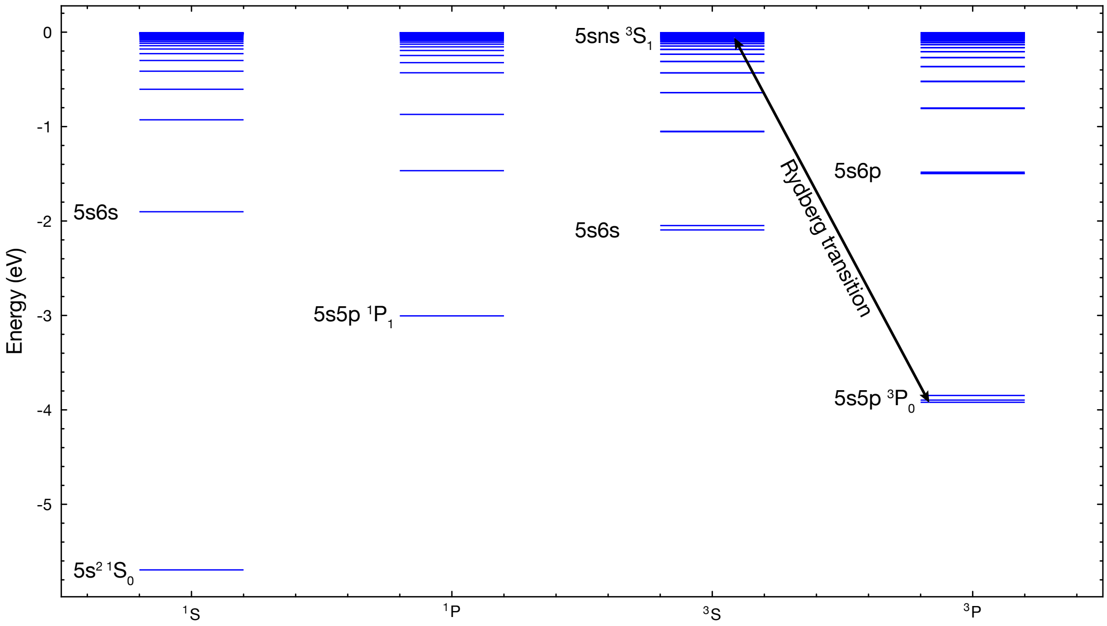

Figure 4.11: Energy level diagram of strontium-88 for angular momentum $L \in [0, 1]$ and principal quantum number $n \in [5, 50]$ calculated from ARC [146]. The Rydberg state targeted in this work is 5s47s ^{3}S$_1$, excited from 5s5p ^{3}P$_0$ by one-photon process at $\sim 317\,$nm.

radius scales as $r \sim a_0 n^2$, where a_0 is Bohr radius. At high principal quantum numbers, Rydberg states have long lifetimes. There are several decay pathways for a Rydberg atom: (1) radiative decay to a nearby Rydberg state with rate $\gamma_{\Delta n=1} \propto n^{-5}$, and (2) decay to low-lying states (including the ground state) with rate $\gamma_{\Delta n \gg 1} \propto n^{-3}$. Both are suppressed at large n. (3) the black-body radiation induced decay $\gamma_{bb} \propto n^{-2}$.

In a system of two Rydberg atoms, the atoms interact via a dipole–dipole potential of the form

$$V = \frac{1}{4\pi\epsilon_0 r^3} \left[\mathbf{d}_1 \cdot \mathbf{d}_2 - 3(\mathbf{d}_1 \cdot \hat{r})(\mathbf{d}_2 \cdot \hat{r}) \right],$$

where $\mathbf{d}_i = e\mathbf{r}_i$ are the dipole operators for atoms $i = 1, 2$, and $\hat{\mathbf{r}}$ is the direction of the interatomic axis. Depending on the states coupled by this interaction, two qualitatively different regimes can emerge: In the exchange interaction regime, two Rydberg states are resonantly coupled. The interaction energy is given by

$$V_{\text{dip}} = \frac{C_3}{r^3}, \qquad C_3 \propto n^4.$$

In the Van der Waals interaction regime where coupling are off-resonant, the interaction takes the form

$$V_{\text{vdW}} = \frac{C_6}{r^6}, \qquad C_6 \propto n^{11}.$$

This van der Waals regime is commonly used for entanglement generation via the Rydberg blockade mechanism.

4.2.2 Rydberg excitations

We demonstrate coherent excitation of atoms to Rydberg states by driving high-contrast Rabi oscillations between the $^3\mathrm{P}_0$ state and the 47s $^3\mathrm{S}_{1,-1}$ Rydberg state with a transition wavelength of $\sim$ 317 nm (Fig. 4.12a). To detect Rydberg atoms, we drive the 5sns $^3\mathrm{S}_1 \rightarrow$ $5\mathrm{p}_{3/2}\mathrm{ns}_{1/2}$ transition at 408 nm, which leads to rapid atom loss via auto-ionization [106]. To avoid anti-trapping effect, the tweezers are switched off by a constant duration. The tweezer array used for single-atom Rabi measurements is identical to that employed in subsequent two-qubit gate benchmarking. Tweezers are laid out in a rectangular pattern with a spacing of 2 µm along the short axis and 6.5 µm along the long axis (Fig. 4.12b). When overlapped with a tightly focused UV beam, this configuration minimizes spatial inhomogeneity in the Rydberg Rabi frequency while maximizing data acquisition rate. In single-atom Rabi experiments, atoms are further rearranged into a zigzag configuration with an interatomic spacing of 13.15 µm to suppress unwanted interactions and further reduce spatial bias. Over a continuous 7-hour measurement period, the relative spread (standard deviation divided by the mean) of the Rabi frequency is found to be 0.27% across different tweezers and 0.08% over time (Fig. 4.12c,d), highlighting the stability of the system.

To characterize the coherence of the Rydberg transition, we perform a Ramsey interferometry sequence (Fig. 4.13a). By varying the dark time between the two Rydberg excitation pulses, we observe a decaying envelope in the $^3\mathrm{P}_0$ population. Independent measurements reveal a slowly drifting Ramsey fringe frequency, which appears to be partially correlated with ambient humidity near the glass cell. Motivated by this observation, we model the decay envelope as $C(t) = \exp\left(-2\pi^2\sigma^2t^2\right)$, corresponding to dephasing due to a Gaussian-distributed detuning error with standard deviation σ. The T_2^* coherence time defined by $C(T_2^*) = e^{-1}$ is fitted to be 4.3(4) µs, consistent in order of magnitude with an independently measured detuning noise of $\sigma \approx 100\,\mathrm{kHz}$. Given the slow timescale of the detuning drift, we implement an atomic servo—in addition to the flashing of UV LED (Nichia NVSU233C-D4) near the glass cell—that regulates the UV laser frequency using an interleaved Ramsey sequence. With the servo engaged, the frequency drift is much suppressed (Fig. 4.13b).

4.2.3 Rydberg lifetime and decay branching

We measure the Rydberg state lifetime using the pair of sequences in Fig. 4.14a. Inspired by recent work [136, 44], we make a distinction between decays into bright states and dark states. Bright states—namely $^1\mathrm{S}_0$ and the $^3\mathrm{P}_J$ manifold—produce a positive signal in our standard fluorescence-based imaging. In contrast, the dark states produce no positive signal and appear dark. For example, population in nearby Rydberg levels can be lost through anti-trapping and thus cannot be detected.

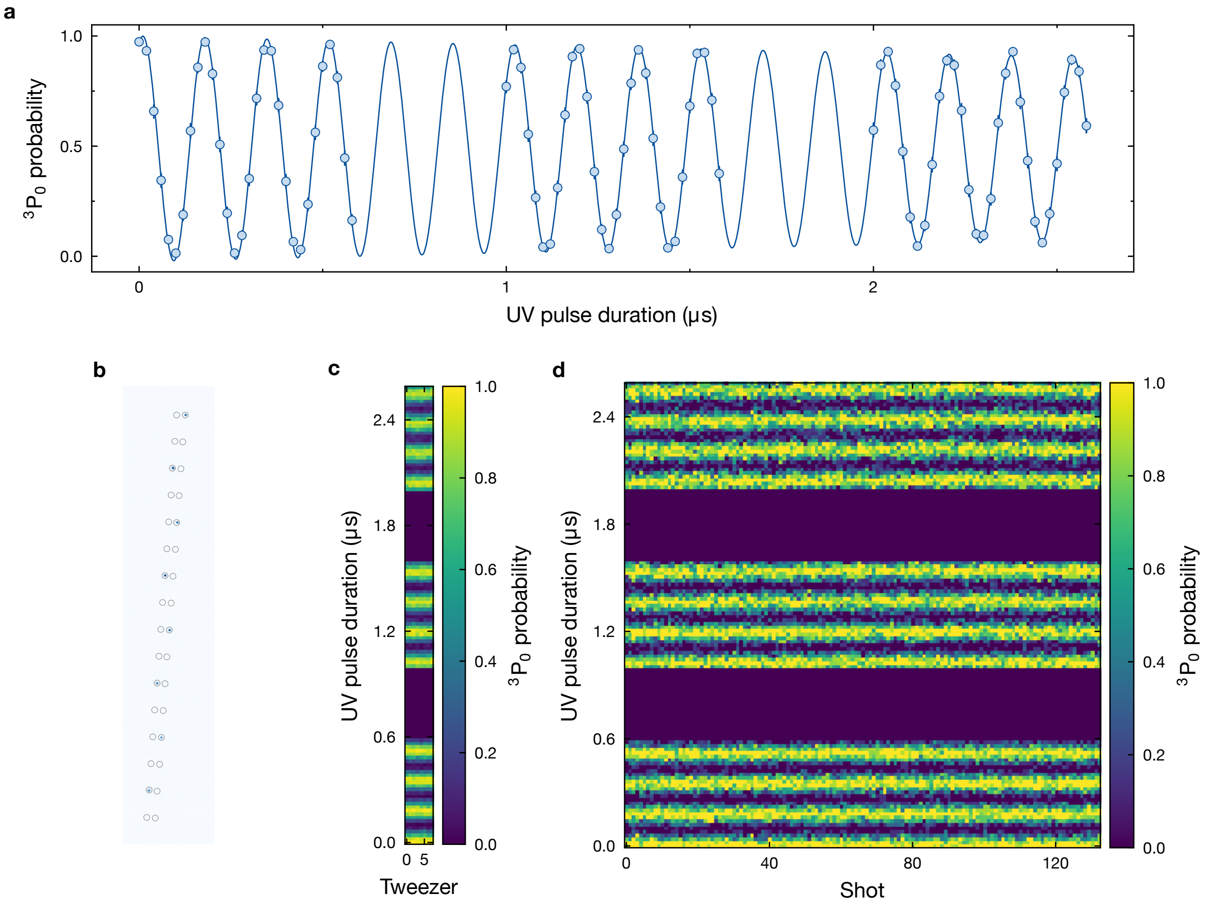

Figure 4.12: **(a)** $^3P_0 \rightarrow 47s\ ^3S_{1,-1}$ Rabi oscillation averaged over all tweezers and time. **(b)** The zigzag atomic array used for Rabi flopping benchmarking. The Rabi flopping measurement analyzed in a tweezer-resolved (**c**) and experimental shot-resolved (**d**) way shows no significant drift. For reference, one shot takes 3-4 seconds.

The population P_1, P_2 measured in Fig. 4.14a do not directly yield the lifetime of interest. To establish the missing link, we employ a simplified description using a three-level system containing Rydberg state $|r\rangle$, dark state $|d\rangle$, and bright state $|b\rangle$. The dynamics during the decay can be described by equations

$$P_r'(t) = -\left(\frac{1}{\tau_b} + \frac{1}{\tau_d}\right) P_r(t), \quad P_d'(t) = \frac{1}{\tau_d} P_r(t), \quad P_b'(t) = \frac{1}{\tau_b} P_r(t). \tag{4.4}$$

To streamline the analysis, we combine the dark-to-bright decay with the Rydberg-to-bright decay. With initial condition $P_r(0) = A, P_d(0) = P_b(0) = 0$, the solution of the rate

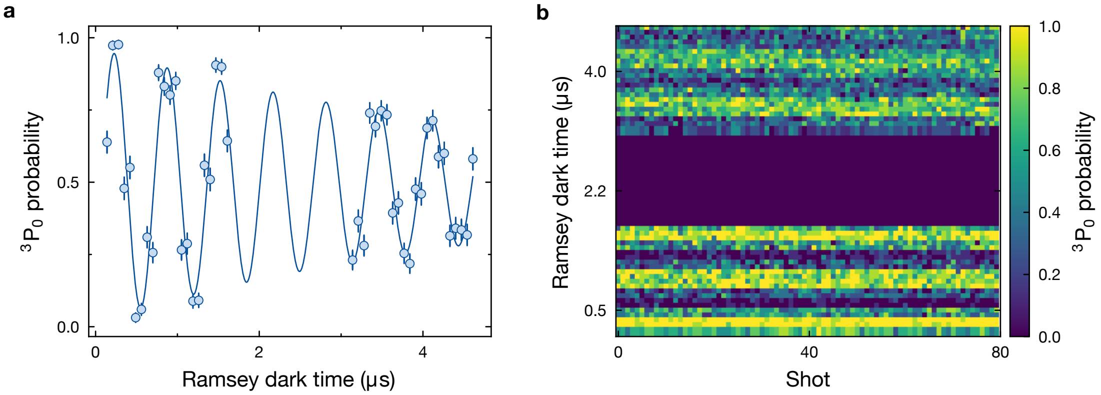

Figure 4.13: **(a)** Rydberg Ramsey measurement with two Rydberg pulses detuned by 1 MHz. The decaying envelop is fitted with $C(t) = \exp\left(-2\pi^2\sigma^2 t^2\right)$, yielding a $1/e$ coherence time of 4.3(4) μs. **(b)** The same Ramsey data analyzed in experimental shots. With atomic servo locking the UV frequency to the resonance, the Ramsey frequency does not show noticeable drift over long timescale.

equations is

$$P_r(t) = A\exp\left[-t\left(\frac{1}{\tau_d}+\frac{1}{\tau_b}\right)\right]$$

$$P_d(t) = A\frac{\tau_b}{\tau_b+\tau_d}\left(1-\exp\left[-t\left(\frac{1}{\tau_d}+\frac{1}{\tau_b}\right)\right]\right)$$

$$P_b(t) = A\frac{\tau_d}{\tau_b+\tau_d}\left(1-\exp\left[-t\left(\frac{1}{\tau_d}+\frac{1}{\tau_b}\right)\right]\right)$$

Here, A is the global normalization factor which is about 0.4 due to loss of atoms in a 70 μs long free-flight independent of T.

It follows that

$$P_1(t) = P_b(t) \tag{4.5}$$

$$P_2(t) = A - P_d(t) - \frac{1}{9}P_b(t). \tag{4.6}$$

The 1/9 accounts for a re-excitation of $^3P_0$ population to $|r\rangle$ at the second Rydberg π-pulse and its subsequent removal by the auto-ionization pulse. Fitting the data simultaneously with these two equations yields the bright-state lifetime $\tau_b = 110(8)$ μs and dark-state lifetime $\tau_d = 37(2)$ μs, see Fig. 4.14b. With a suitable combination of repumpers in the read-out stage, we can further identify the branching into each substate among the bright states, see Fig. 4.14c. Specifically, we measure the population accumulation rate of the bright state γ_b, and the same quantity after (i). 461 nm (ii). 496 nm+461 nm pulses. Based on their ratio, the branching can be calculated. Note that the observed relative branching within the bright states is consistent with the expectation considering only the degeneracy alone.

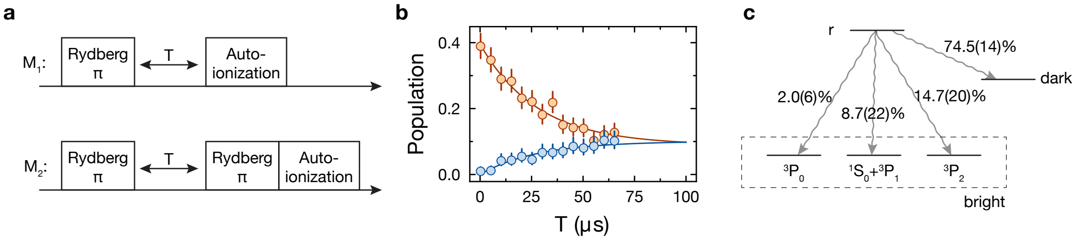

Figure 4.14: **(a)** The two complementary sequences, when analyzed together, determine the lifetimes of Rydberg states decaying into both bright and dark channels. **(b)** The population measured in sequence M_1 (blue) and M_2 (orange). The solid lines are fit functions derived from the decay model, see main text. The dark and bright state lifetimes are 37(2) μs and 110(8) μs, respectively. The experiments are performed in free-space, resulting in the drop in population independent of T. **(c)** Measured branching ratio into dark and low-lying states.

4.2.4 CZ gate and Bell state benchmarking

An ideal controlled-phase (CZ) gate implements an two-qubit unitary U_{CZ} defined as

$$U_{\mathrm{CZ}} = \mathrm{Diag}(U_{00,00}, U_{01,01}, U_{10,10}, U_{11,11}) = \mathrm{Diag}(1, 1, 1, -1). \tag{4.7}$$

That is, the CZ gate applies a π-phase shift to the $|11\rangle$ state while leaving all other computational basis states $|00\rangle, |01\rangle$, and $|10\rangle$ unchanged. It can equivalently be viewed as flipping the phase of the target qubit conditioned on the control qubit being in the $|1\rangle$ state. We forego any historical summary of CZ gate implementations, and instead focus exclusively on the Jandura-Pupillo (JP) gate [87]. Due to its simplicity in the use of a global addressing beam and its time-optimal control design, high-fidelity experimental demonstrations of the JP gate have been widely reported in recent years [35, 125, 43, 85, 128].

Let us first develop a physical understanding of the JP gate. Under continuous UV illumination, the state $|01\rangle$ undergoes Rabi oscillation with $|0r\rangle$ at a Rabi frequency Ω. In contrast, the state $|11\rangle$ couples to the symmetric singly excited Rydberg state $|W\rangle \sim |0r\rangle + |r0\rangle$ and oscillates at an enhanced Rabi frequency $\sqrt{2}\Omega$, due to the Rydberg blockade effect. Intuitively, there should exist a pulse sequence that drives both $|01\rangle$ and $|11\rangle$ through closed trajectories on their respective Bloch spheres, returning them to their original states. However, because the two transitions proceed at different Rabi frequencies, the dynamic phases accumulated by $|01\rangle$ and $|11\rangle$ differ. This relative phase difference is precisely what is exploited by the JP gate to implement the conditional phase. Formally, the JP gate implements a diagonal unitary of the form

$$U'_{\mathrm{CZ}} = \mathrm{Diag}(1, e^{-i\phi_{01}}, e^{-i\phi_{10}}, e^{i(\phi_{11}-\phi_{01}-\phi_{10})}) = \mathrm{Diag}(1, e^{-i\phi_1}, e^{-i\phi_1}, e^{i(\phi_{11}-2\phi_1)}), \tag{4.8}$$

where the second equality assumes exchange symmetry between the two qubits ($\phi_{01} = \phi_{10} = \phi_1$). The conditional phase ϕ_{11}, which arises from the Rydberg interaction, is

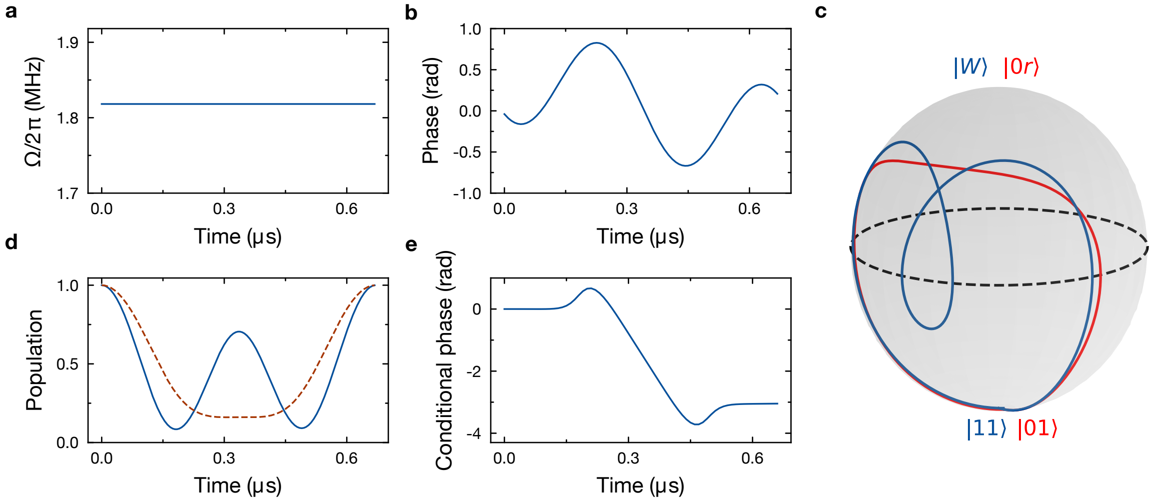

Figure 4.15: The UV waveform amplitude **(a)** and phase **(b)** during a ideal time-optimal gate [87]. For the simulation shown below, the amplitude is assumed without loss of generality to be $\Omega = 2\pi \times 1.81\,\text{MHz}$ and have no rise time. **(c)** Under UV illumination, the state vectors of initial state $|01\rangle$ (red) and $|11\rangle$ (blue, solid) evolve at different rate due to the Rydberg blockade effect. **(d)** At $T = 7.612/\Omega$ which marks the end of the gate pulse, the population of both $|01\rangle$ (red, dashed) and $|11\rangle$ (blue, solid) returns to their initial value. **(e)** At the same time, the conditional phase $\phi_{11} - 2\phi_1$ becomes $-\pi$, as required for a CZ gate.

optimized to $\pm\pi$. To engineer a canonical CZ gate as defined in Eqn. (4.7), the residual single-qubit phase ϕ_1 must be corrected. This can be achieved by appling a global single-qubit phase gate $Z(\phi_1)$ which affects both qubits. As a result, we arrive at $U_{\text{CZ}} = U'_{\text{CZ}} \cdot (Z(\phi_1) \otimes Z(\phi_1))$.

With the CZ gate, two qubits can be readily entangled. To characterize the fidelity of this entangling operation, one typically prepares and verifies the Bell state

$$|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

which not only serves as a convenient benchmark state but also has significant applications in entanglement-enhanced metrology. The circuit used to generate and benchmark the Bell state is shown in Fig. 4.16b [35]. Given an experimentally reconstructed density matrix ρ, the fidelity with respect to the target Bell state is given by

$$F = \langle\Phi^+|\rho|\Phi^+\rangle = \frac{1}{2}(\rho_{00,00} + \rho_{11,11} + \rho_{00,11} + \rho_{11,00}).$$

The first two terms are diagonal elements corresponding to populations in the computational basis and can be measured directly. The last two terms involve coherences between $|00\rangle$ and $|11\rangle$, and require a basis rotation to access. To extract these off-diagonal elements, we employ the analyzer module shown in the dashed box of Fig. 4.16a.

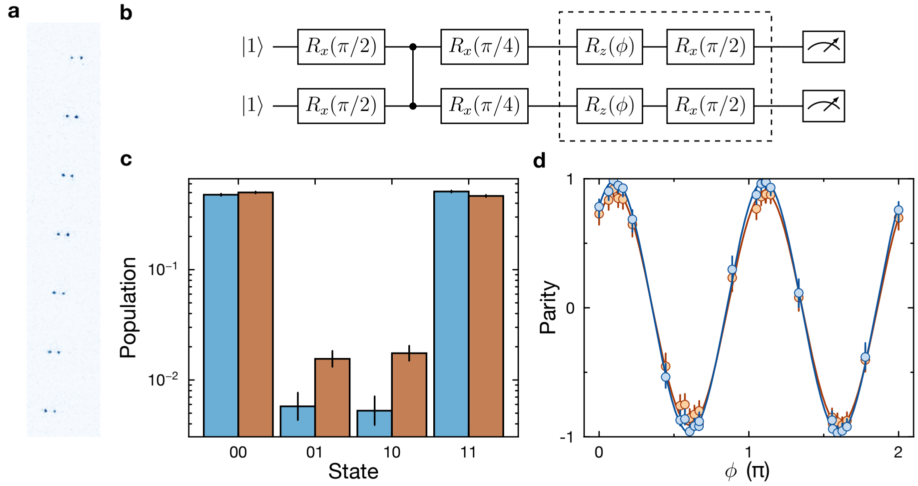

Figure 4.16: **(a)** Atom pairs used for parallel generation of Bell states. **(b)** The circuit for preparing $|\Phi^+\rangle$ and probing the off-diagonal elements by inducing parity oscillation (boxed in dashed line) **(c)** The Bell state tomography and **(d)** parity oscillation measurement, with (orange) and without (blue) sensitivity to atom loss in the state detection. Bell state tomography and parity oscillation together yield a Bell state fidelity of 0.983(8) (blue) and 0.935(9) (red) with and without post-selection on the survival of the atom pairs.

Mathematically, the analyzer applies a local frame rotation described by the unitary $U = \exp(-i\pi X/4)\exp(-i\phi Z/2)$ which rotates the measurement basis. Subsequent measurement of the parity operator Z_1Z_2 in the new basis yields

$$P = \mathrm{Tr}[\rho U^\dagger Z_1 Z_2 U] = -e^{-2i\phi}\rho_{00,11} - e^{2i\phi}\rho_{11,00} + \rho_{01,10} + \rho_{10,01}.$$

As the phase ϕ is scanned, the oscillating component arises from $2\mathrm{Re}[e^{-2i\phi}\rho_{00,11}]$. By fitting the experimetal data to the function $A\cos(2\phi+\phi_0)$, we extract the amplitude $A = 2|\rho_{00,11}|$, which provides an upper bound on the coherence term $\rho_{00,11} + \rho_{11,00} = 2\mathrm{Re}[\rho_{00,11}]$ as $\mathrm{Re}[\rho_{00,11}] \leq |\rho_{00,11}|$. Therefore, the maximum Bell state fidelity can be calculated from the measured populations and the parity oscillation contrast A as

$$F \leq \frac{1}{2}(P_{00} + P_{11} + A). \tag{4.9}$$

where $P_{00} = \rho_{00,00}$ and $P_{11} = \rho_{11,11}$.

The experimental result of the Bell state benchmarking is shown in Fig. 4.16. Atoms are arranged into pairs with a spacing of 2 µm, leading to an interaction shift of $2\pi \times 114$ MHz. The separation between neighboring pairs is 13 µm, ensuring that crosstalk between different pairs remains negligible. The population and parity oscillation measurements as required in (4.9) are shown Fig. 4.16b,c. Each measurement is analyzed both in a loss-insensitive and a loss-sensitive manner. A substantial increase in Bell state fidelity is observed when loss events are excluded from the analysis. This improvement is attributed to partial ionization of atoms in the $^3\mathrm{P}_2$ state due to the UV beam, as well as to Rydberg state decay leading to atom loss. While conceptually simple, Bell state benchmarking is sensitive to single-qubit gate and SPAM errors and tends to underestimate the fidelity of the CZ gate. This limitation, along with the omission of the CZ gate parameter optimization in this section, will be addressed in subsequent sections.

4.2.5 Interleaved randomized benchmarking

To characterize the fidelity of the CZ gate while minimizing the influence of single-qubit and SPAM (state preparation and measurement) errors, we employ randomized benchmarking (RB). Unlike the single-qubit case, where the fidelity of the entire Clifford group is of interest, our focus is on quantifying the fidelity of a specific two-qubit gate—namely, the controlled-Z (CZ) gate. Interleaved randomized benchmarking (IRB) provides a scalable and efficient method for this purpose by inserting the CZ gate between layers of randomized single-qubit Clifford operations. By maintaining a constant number of single-qubit gates across sequences, the decay of the return probability to the initial state follows a depolarizing error model as a function of the number of interleaved CZ gates. The fidelity of the CZ gate can then be extracted by fitting the decay curve to the model function.

The first IRB sequence, commonly referred to as the "echo sequence," is shown in Fig. 4.17b. It includes a characteristic CZ–X–CZ block designed to echo out single-qubit phase, which otherwise require separate calibration. That is, $U'_{\mathrm{CZ}} \cdot (X \otimes X) \cdot U'_{\mathrm{CZ}} = e^{2i\phi_1} V$, where a global phase ϕ_1 is factored out and V depends only on ϕ_{11}. The remainder of the sequence consists of a fixed number of single-qubit rotations, denoted R_{rand} and X, which scramble coherence between the CZ gates while maintaining a constant overall sequence length. Finally, R_f maps the population back to its initial state in the absence of gate errors. A typical measurement with varying number of CZ gates is shown in Fig. 4.17a.

Experimentally, IRB can be used to optimize the CZ gate parameters. As discussed earlier, a time-optimal waveform consists of a constant UV amplitude with a phase modulation, which is numerically derived for a given Rydberg Rabi frequency. In practice, however, the engineered UV waveform is often subject to distortions due to experimental imperfections (Fig. 2.11), resulting in a suboptimal CZ gate operation. Although real-time measurement of the UV waveform combined with active feedback control has been explored to mitigate such distortions, we adopt a more pragmatic method by experimentally scanning relevant gate parameters to maximize the return probability in an IRB sequence. To this end, we parametrize the time-optimal gate waveform driving the optical modulator using five parameters, $(T, \alpha, \omega, \phi_0, \delta)$, as follows

$$\Omega(t) = \begin{cases} \Omega, & 0 < t < T, \\ 0, & \text{otherwise} \end{cases}$$

$$\phi(t) = \alpha \cos(\omega t - \phi_0) + \delta t.$$

The optimization routine essentially involves searching for a global maximum of P_{11} at a fixed N_{CZ} within a five-dimensional parameter space. Inspired by recent work [147], we begin with an initial estimate of the experimental parameters and numerically compute the Hessian matrix, which identifies the eigen-directions that yield the steepest ascent in the multi-dimensional gradient landscape. In the experiment, we then scan along the parameter basis defined by the Hessian, rather than over individual bare parameters, to achieve faster convergence. The Hessian, along with the results of the experimental optimization, is presented in Fig. 4.17c,d. Convergence can usually be reached within one round of optimization. In the final benchmarking of gate fidelity, we often observe strong oscillation in P_{11} for small N_{CZ} when using the echo sequence (Fig. 4.17a), along with slow convergence in data fitting. Furthermore, the CZ-gate fidelity extracted from the echo sequence exhibits greater sensitivity to single-qubit gate errors compared to alternative protocols [85]. As a result, we do not use the echo sequence to measure the gate fidelity, but rather employ it as a tool for optimizing gate parameters.

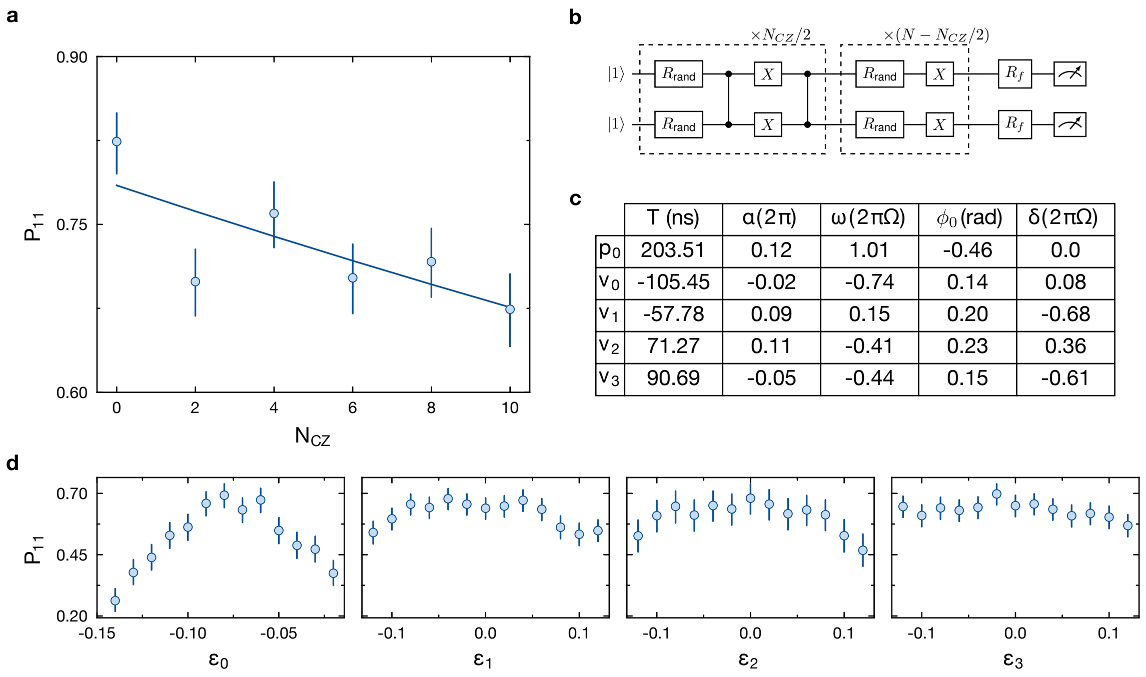

	T (ns)	α(2π)	ω(2πΩ)	ϕ_0(rad)	δ(2πΩ)
p_0	203.51	0.12	1.01	-0.46	0.0
v_0	-105.45	-0.02	-0.74	0.14	0.08
v_1	-57.78	0.09	0.15	0.20	-0.68
v_2	71.27	0.11	-0.41	0.23	0.36
v_3	90.69	-0.05	-0.44	0.15	-0.61

Figure 4.17: **(a)** A typical result obtained from executing the echo benchmarking sequence shown in **(b)**. See main text for a detailed discussion of the sequence. **(c)** Numerically estimated initial gate parameters and corresponding eigen-directions used for accelerated convergence. The final parameter vector is given by $\mathbf{p}_0 + \sum_i \epsilon_i \mathbf{v}_i$. **(d)** Measured P_{11} as a function of scan weight along different eigen-directions in a sequential optimization.

4.2.6 Symmetric stabilizer benchmarking

Complementary to the echo sequence is the symmetric stabilizer benchmarking (SSB) sequence [43, 85]. SSB is less sensitive to single-qubit gate errors, which is particularly relevant in our system, where such errors are non-negligible. Moreover, the sequence is sensitive to single-qubit phase ϕ_1, which is the final undetermined parameter in constructing a canonical CZ gate. The SSB sequence is illustrated in Fig. 4.18a. An initialization unitary U_{init} consisting of five $\pi/2$ pulses and a single CZ gate prepare the two-qubit system in one of twelves symmetric states. The main body of the sequence interleaves random single-qubit rotations, R_c, with CZ gates. Finally, a recovery unitary, U_{rec}, is applied. This unitary is precomputed to return the system to the $|11\rangle$ state in the absence of gate errors. The SSB sequence is sensitive to the single-qubit phase accumulated during the JP gate. This phase is compensated by applying a phase gate, $Z(\phi)$, via frame tracking (Sec. 4.1.4) immediately after the entangling operation (Fig. 4.18a). To determine the compensation phase, we scan ϕ and maximize the return probability of the SSB sequence for a fixed number of CZ gates, N_{CZ} (Fig. 4.18b).

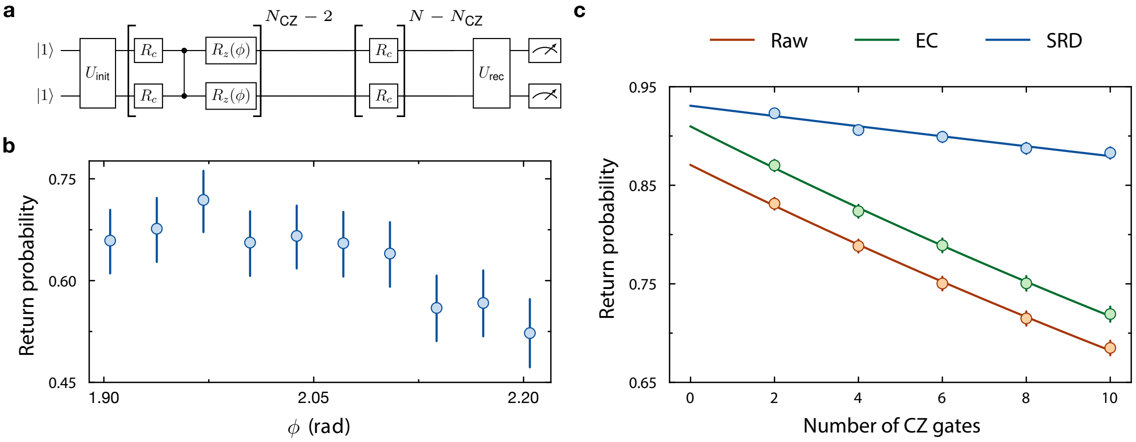

Figure 4.18: **(a)** The circuit of symmetric stabilizer benchmarking (SSB) sequence. **(b)** The return probability exhibits a maximum as the single-qubit phase ϕ is scanned, at a constant N_{CZ}. **(c)** The SSB result shows 0.9945(6) CZ gate fidelity, after removing events in which atoms loss is detected (SRD). The erasure-excised (EC) gate fidelity of 0.9764(8) is higher than the raw fidelity of 0.9759(5) due to the expected conversion of 8.7% decay leakage into $^1S_0$ and $^3P_1$ (Fig. 4.14c) into erasure errors.

Once the single-qubit phase of the CZ gate is identified, the gate becomes fully defined. To benchmark its performance, we perform a scan over the number of CZ gates in a SSB sequence and fit the resulting data to the model function $aF^{N_{\text{CZ}}}$, where a accounts for state preparation and measurement (SPAM) errors, and F represents the average fidelity of the CZ gate (Fig. 4.18b). The benchmarking protocol includes two erasure conversion stages: one preceding the initialization unitary U_{init}, and one following the recovery unitary U_{rec}. The sequence concludes with state-resolved detection (SRD) of $|0\rangle$ and $|1\rangle$ capable of identifying atom loss. If only erasure conversion is applied, without the full state resolution in the read-out, the measured CZ gate fidelity is 0.9764(8), which is only marginally higher than the raw fidelity of 0.9759(5) (Fig. 4.18c). This observation suggests that leakage into

^{3}P$_1$ and ^{1}S$_0$ during the CZ gate plays only a minor role in the overall gate infidelity. However, when full state resolution is utilized, the system becomes sensitive to atom loss, enabling post-selection of runs in which both atoms survive the entire benchmarking sequence. Under such post-selection, the extracted CZ gate fidelity improves significantly to 0.9945(6), compared to the raw fidelity of 0.9759(5) without post-selection (Fig. 4.18c). This discrepancy can be largely attributed to dominant error mechanisms that manifest as atom loss. For example, during UV illumination used to couple ^{3}P$_0$ to the Rydberg state, the same laser can ionize population in ^{3}P$_2$, leading to loss events that are not directly tied to the CZ gate's coherent dynamics. As a result, this source of infidelity is relatively insensitive to the precise gate parameters. In addition, infidelity arising from spontaneous decay of the Rydberg state into long-lived dark states can be substantially mitigated through post-selection, further enhancing the observed fidelity.

Error budget

We estimate the contributions to the infidelity of our gate using the effective super-operator associated with each error source. The super-operator is obtained, for each error source, by solving the associated master equation, taking into account the coherent error sources in the Hamiltonian and the collapse operators for spontaneous decay, pure dephasing on the Rydberg transition, and ^{3}P$_2$ ionization. The infidelity is calculated in two ways, through the process fidelity between the simulated super-operator with the ideal one, and using a simulated SSB sequence with the obtained super-operator. We observe no notable difference in results between the two methods. For gate infidelity originating from laser phase noise, we compared our result to the one derived from the recently proposed linear response calculation [85] and find good agreement. Numerically, aside from the qubit and Rydberg states, we define an additional "bucket" state for each atom to which it is assumed to be transferred in the case of either Rydberg decay to outside the qubit manifold or ^{3}P$_2$ ionization loss, see Fig. 4.19a. To estimate infidelities taking the loss-correcting SRD into account, we use the simulated SSB sequence once more, but only consider the valid computational states. Population decayed from Rydberg states into $m_J \neq 0$ states of ^{3}P$_2$ which are not qubit states but nevertheless manifest as if they are, are also taken into account. We identify Rydberg state decay, ^{3}P$_2$ ionization (see above), and gate parameter estimation errors as the most prominent, known, error sources in our two-qubit gate. The effect of other known noise sources are presented in Fig. 4.19b. The modeled noise sources give rise to a total raw infidelity of 1.84% and loss-corrected infidelity of 0.16%. We attribute the higher infidelities measured in the experiment to experimental drifts such as in Rabi frequency and UV pulse area, phase, and shape. Note that the ionization loss will be eliminated with a UV laser system coupling ^{3}P$_2$ to the Rydberg state, and the infidelity from Rydberg decay can be further reduced by increasing the laser power.

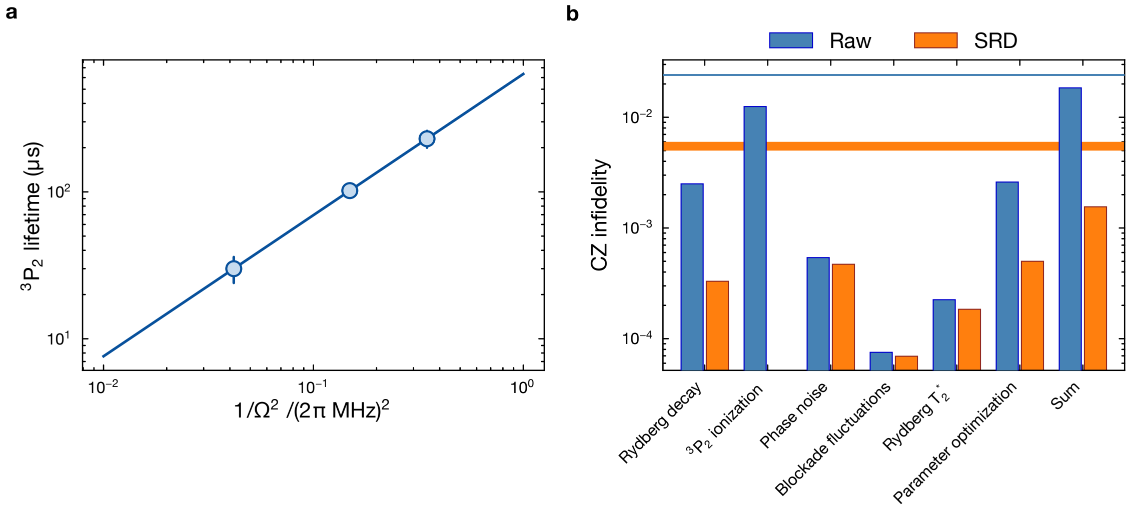

Figure 4.19: **(a)** $^3\mathrm{P}_2$ lifetime τ scales with UV rabi frequency Ω as $\tau = A/\Omega^2$ with a proportionality constant $A = 610/\mu\mathrm{s}$. **(b)** Leading contributions to the infidelity of the CZ gate with (orange) and without (blue) loss correction at a Rabi frequency of $2\pi \times 6\,\mathrm{MHz}$. The experimental infidelity are indicated by the corresponding horizontal lines.

Chapter 5

Further directions

High-fidelity analogue quantum simulation

One class of applications that becomes immediately accessible using the experimental tools developed in this thesis is quantum simulation based on Rydberg interactions. In the following, we outline a specific proposal aimed at studying lattice gauge theories using Rydberg interactions in the facilitation regime. Rydberg facilitation refers to a phenomenon in which atoms are more likely to be excited simultaneously to Rydberg states, even in the presence of strong interaction-induced energy shifts. This behavior is illustrated schematically in Fig. 5.1a for a system of two atoms. We consider a Rydberg Hamiltonian of the form:

$$H = \frac{1}{2}\sum_{i}\left(\Omega\sigma_i^x - \Delta\sigma_i^z\right) + \frac{V}{2}\sum_{i>j}\left|r_i r_j\right\rangle\left\langle r_i r_j\right|,$$

where interactions are assumed to act only between nearest neighbors, and set $\hbar = 1$. In this regime, the two-atom ground state $|gg\rangle$ can be resonantly coupled to the doubly excited state $|rr\rangle$ when the detuning satisfies $\Delta = -V/4$. The resulting effective two-photon Rabi frequency then scales as $\Omega_{2p} \propto \Omega^2/V$. A simulated Rabi oscillation is shown in Fig. 5.1b for parameters $V = 2\pi \times 80\,\mathrm{MHz}$ and $\Omega = 2\pi \times 1\,\mathrm{MHz}$. Crucially, at this specific detuning, the presence of a atom within the blockade radius can facilitate the excitation of a neighboring atom, enabling a form of correlated dynamics well-suited for simulating creation-annihilation process in lattice gauge theory [148].

Our platform offers several advantages for quantum simulation in the Rydberg facilitation regime. First, Rydberg facilitation is highly sensitive to the interaction strength, $V = C_6/R_6$ where R is the interatomic distance. Consequently, any positional disorder leads to inhomogeneous broadening of the interaction strength. Our ultraprecise tweezer generation technique significantly mitigates technical positioning errors at the algorithmic level. It is believed that the residual uncertainty is comparable to the spatial extent of the atomic ground-state wavefunction. Moreover, if some residual disorder remains, one can leverage the periodicity of an optical lattice as a spatial reference to calibrate and further regulate tweezer positions, thereby improving positional uniformity.

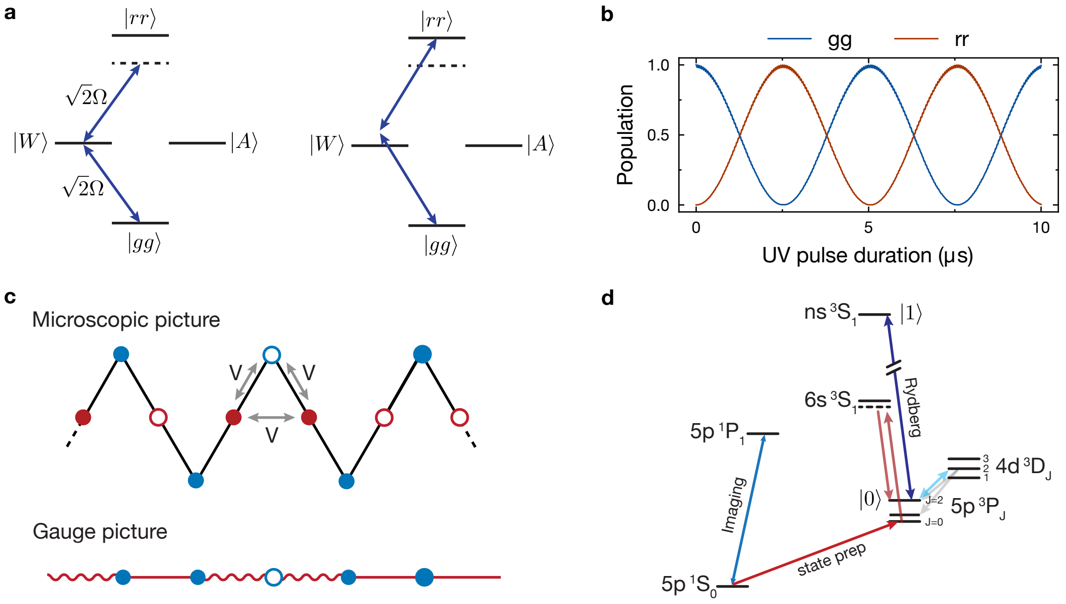

Figure 5.1: **(a)** Rydberg facilitation. **(b)** Two-photon Rabi oscillation between $|gg\rangle$ and $|rr\rangle$ in the Rydberg facilitation regime. The detuning satisfies $\Delta = -V/4$, where V is the nearest-neighbour Rydberg interaction strength. **(c)** Illustration of quantum simulation of lattice guage theory (LGT) [148]. In the microscopic pictures, atoms occupy trapping sites arranged in a zig-zag geometry. Sites located on the links (red) are distinguished from vertex sites (blue) via local light shifts. The angle between adjacent links is 60°, ensuring that all neighboring atoms are equidistant. Open circles denote atoms in the ground state $|0\rangle$, while filled circles indicate atoms in the Rydberg state $|r\rangle$. To map the system to a gauge-theoretic model, the occupation of link sites corresponds to electric field degrees of freedom pointing left or right, while the occupation of vertex sites maps to the presence or absence of matter. Under appropriate driving conditions, the dynamics are constrained to a gauge-invariant subspace where Gauss's law is satisfied. **(d)** Atomic level structure relevant for simulating the 1D lattice gauge theory.

Second, analogue Rydberg simulations are prone to errors, such as Rydberg state decay and leakage errors. This imperfection is particularly troublesome in facilitation physics where the dynamics happens at a much longer timescale. As demonstrated in chapter 4.1.2, we have developed an erasure conversion technique that identify and remove errors arising from imperfect state preparation and approximately 30% of off-resonant scattering. With further system upgrades of the fast $^3P_2 \longrightarrow {}^3P_1$ Raman lasers, we anticipate converting up to 77.8% of off-resonant scattering errors into erasures. More importantly, we can exploit a dissipative imaging protocol to map population in highly excited Rydberg states to lower-lying levels, enabling state-resolved readout. As discussed in chapter 4.1.3, readout with state resolution is sensitive to atom loss, which constitutes a major infidelity source due to the finite lifetime of the Rydberg state. To perform the mapping, one first applies a fast two-photon π-pulse to transfer $|0\rangle$ population into $^3\mathrm{P}_0$ state. After the state transfer, the remaining population in the non-qubit substates can be optionally push out. Subsequently, both the $^3\mathrm{P}_2 \longrightarrow |r\rangle$ Rydberg excitation light and the $^3\mathrm{P}_2 \rightarrow {}^3\mathrm{D}_2$ repumping light are applied simultaneously. Owing to the strong dissipation, Rydberg population is efficiently transferred down to the absolute ground state, where fast, high-fidelity fluorescence imaging can be employed for final state readout. Finally, the ground-state population is removed, allowing selective readout of the remaining atoms in $^3\mathrm{P}_0$.

Collective light-matter interaction

Another promising research direction that leverages the strengths of our platform is the study of collective light–matter interactions in optical lattices. Recent proposal of the 2.9 µm transition has only accelerated the effort [47]. While both superradiant and subradiant effects have been observed in other systems, combining these phenomena with the scalability, programmability, and site-resolved addressability of our hybrid setup is expected to offer new insights into the underlying physics.

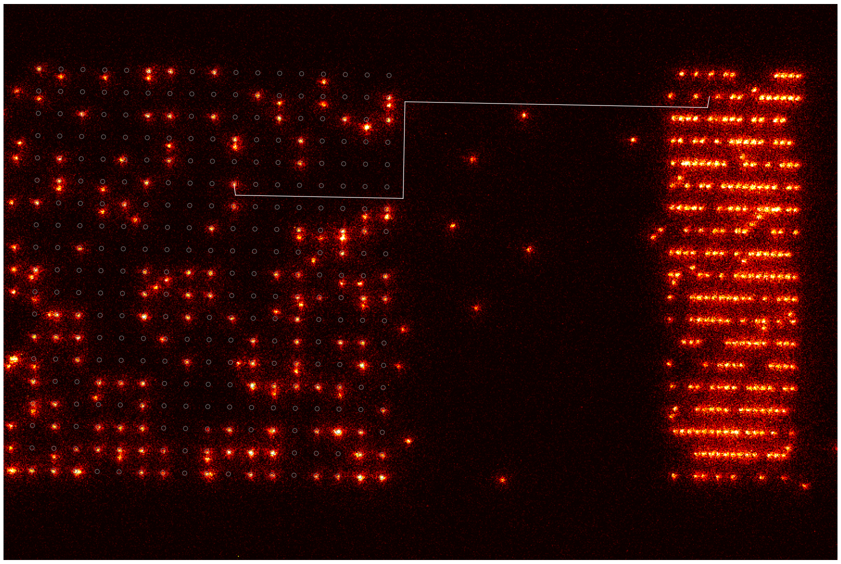

Figure 5.2: Atoms can be rearranged in the lattice, offering unique opportunity to study collective light-matter interaction effect [149].

Our platform presents several opportunities for investigating collective light–matter interactions. First, the hybrid architecture enables near-complete control over atomic positions within the optical lattice. This controllability allows for flexible atom rearrangement, facilitating dense packing of atoms at sub-wavelength separations—a key requirement for observing collective optical phenomena, which rely on emitter spacing comparable to or smaller than the transition wavelength. Such collective behavior arises from quantum interference between multiple emitters and is highly sensitive to spatial disorder. Therefore, minimizing positional inhomogeneity is essential for realizing strong collective effects. Optical lattices, with their intrinsically regular site spacing and minimal positional uncertainty, offer an ideal environment for studying coherent many-body optical responses. In contrast, optical tweezer arrays often exhibit larger site-to-site variations, primarily due to residual aberrations in the optical path.

Second, our Raman coupling scheme enables rapid shelving of the ${}^3\mathrm{P}_2$ population into the non-interacting ${}^3\mathrm{P}_0$ state. This capability provides a means to probe population inversion on short timescales. Additionally, the same laser system may be leveraged for lattice-based cooling, helping to suppress thermal broadening and enhance coherence.

Assembled Hubbard systems in the lattice

The Hubbard model describes systems governed by the Hamiltonian

$$\frac{H_\mathrm{B}}{\hbar} = -J \sum_{\langle i,j \rangle} a_i^\dagger a_j + \sum_i \epsilon_i \, n_i + \frac{U}{2} \sum_i n_i (n_i - 1), \tag{5.1}$$

for bosons, with $n_i = a_i^\dagger a_i$, and

$$\frac{H_\mathrm{F}}{\hbar} = -J \sum_{\langle i,j \rangle, \sigma} c_{i\sigma}^\dagger c_{j\sigma} + \sum_{i,\sigma} \epsilon_i \, n_{i\sigma} + U \sum_i n_{i\uparrow} n_{i\downarrow}, \tag{5.2}$$

for fermions, with $n_{i\sigma} = c_{i\sigma}^\dagger c_{i\sigma}$. Here J is the nearest-neighbour tunnelling rate, ϵ_i the on-site energy offset (take $\epsilon_i = \epsilon_0$ for a homogeneous lattice), and U the on-site interaction strength. For fermions the interaction is between opposite spins; the $(U/2)$ factor is absent because Pauli exclusion prevents same-spin double occupancy.

In neutral atoms, both the Bose–Hubbard [57, 56] and Fermi-Hubbard [150, 151, 16] models have been realized. The Bose–Hubbard model was instrumental in early demonstrations of the superfluid–Mott-insulator transition and, subsequently, studies of synthetic gauge fields [152, 153] and out-of-equilibrium dynamics [154, 61]. The Fermi–Hubbard model is widely believed to underpin strongly correlated spin systems in real materials [155]. Conventionally, Hubbard systems are prepared in a top-down approach by adiabatically loading a degenerate quantum gas into an optical lattice [15]. In our hybrid tweezer–lattice platform, one can instead view the preparation as a bottom-up assembly of a Hubbard system directly in the lattice. Our ability to load $\sim 10{,}000$ atoms into the lattice offers an appealing prospect for scalability. There are, however, a couple of prerequisites for realizing the Hubbard model. First, atoms should ideally occupy only the lowest Bloch band of the lattice, which requires efficient in-lattice cooling and strong axial confinement. Second, the on-site interaction U must be nonzero; for fixed lattice geometry $U \propto a/m$ (more precisely, $U = 4\pi\hbar^2 a/m \int |w(\mathbf{r})|^4, \mathrm{d}\mathbf{r}$), where a is the s-wave scattering length and m the atomic mass. Among strontium isotopes, ^{88}Sr has a small negative scattering length ($a_{88} \approx -a_0$), leading to weak interactions, whereas fermionic ^{87}Sr has $a_{87} \approx 97a_0$, comparable to ^{87}Rb. In addition, bosonic ^{84}Sr has $a_{84} \approx 124a_0$, which is also reasonably large. In our apparatus, all three isotopes are naturally occurring, making the system well suited to exploring Hubbard physics across different isotopes.

Bibliography

[1] Richard P. Feynman, Robert B. Leighton, and Matthew Sands. *The Feynman Lectures on Physics, Volume III: Quantum Mechanics.* Addison-Wesley, Reading, Massachusetts, 1965.

[2] David J. Wineland. Nobel lecture: Superposition, entanglement, and raising schrödinger's cat. *Reviews of Modern Physics*, 85(3):1103–1114, 2013.

[3] Claude N. Cohen-Tannoudji. Nobel lecture: Manipulating atoms with photons. *Reviews of Modern Physics*, 70(3):707–719, 1998.

[4] William D. Phillips. Nobel lecture: Laser cooling and trapping of neutral atoms. *Reviews of Modern Physics*, 70(3):721–741, 1998.

[5] Christian Gross and Immanuel Bloch. Quantum simulations with ultracold atoms in optical lattices. *Science*, 357(6355):995–1001, 2017.

[6] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[7] I. M Georgescu, S. Ashhab, and Franco Nori. Quantum simulation. *Reviews of Modern Physics*, 86(1):153–185, 2014.

[8] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6):467–488, 1982.

[9] Seth Lloyd. Universal quantum simulators. *Science*, 273:1073–1078, 1996.

[10] J. Ignacio Cirac and Peter Zoller. Goals and opportunities in quantum simulation. *Nature Physics*, 8(4):264–266, 2012.

[11] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Physical Review Letters*, 94(17):170201, 2005.

[12] Andrew J. Daley, Immanuel Bloch, Christian Kokail, Stuart Flannigan, Natalie Pearson, Matthias Troyer, and Peter Zoller. Practical quantum advantage in quantum simulation. *Nature*, 607(7920):667–676, 2022.

[13] Patrick A. Lee, Naoto Nagaosa, and Xiao-Gang Wen. Doping a mott insulator: Physics of high-temperature superconductivity. *Reviews of Modern Physics*, 78(1):17–85, 2006.

[14] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller. Cold bosonic atoms in optical lattices. *Physical Review Letters*, 81(15):3108–3111, 1998.

[15] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W. Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms. *Nature*, 415(6867):39–44, 2002.

[16] Anton Mazurenko, Christie S. Chiu, Geoffrey Ji, Maxwell F. Parsons, Márton Kanász-Nagy, Richard Schmidt, Fabian Grusdt, Eugene Demler, Daniel Greif, and Markus Greiner. A cold-atom fermi–hubbard antiferromagnet. *Nature*, 545(7655):462–466, 2017.

[17] Maciej Lewenstein, Anna Sanpera, Veronica Ahufinger, Bogdan Damski, Aditi Sen, and Ujjwal Sen. Ultracold atomic gases in optical lattices: mimicking condensed matter physics and beyond. *Advances in Physics*, 56(2):243–379, 2007.

[18] S. Flannigan, N. Pearson, G. H. Low, A. Buyskikh, I. Bloch, P. Zoller, M. Troyer, and A. J. Daley. Propagation of errors and quantitative quantum simulation with quantum advantage. *Quantum Science and Technology*, 7(4):045025, 2022.

[19] T. I. Andersen, N. Astrakhantsev, A. H. Karamlou, J. Berndtsson, J. Motruk, A. Szasz, J. A. Gross, A. Schuckert, T. Westerhout, Y. Zhang, E. Forati, D. Rossi, B. Kobrin, A. Di Paolo, A. R. Klots, I. Drozdov, V. Kurilovich, A. Petukhov, L. B. Ioffe, A. Elben, A. Rath, V. Vitale, B. Vermersch, R. Acharya, L. A. Beni, K. Anderson, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, B. Ballard, J. C. Bardin, A. Bengtsson, A. Bilmes, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, D. A. Browne, B. Buchea, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, A. Cabrera, J. Campero, H. S. Chang, Z. Chen, B. Chiaro, J. Claes, A. Y. Cleland, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, S. Das, D. M. Debroy, L. De Lorenzo, A. Del Toro Barba, S. Demura, P. Donohoe, A. Dunsworth, C. Earle, A. Eickbusch, A. M. Elbag, M. Elzouka, C. Erickson, L. Faoro, R. Fatemi, V. S. Ferreira, L. Flores Burgos, A. G. Fowler, B. Foxen, S. Ganjam, R. Gasca, W. Giang, C. Gidney, D. Gilboa, M. Giustina, R. Gosula, A. Grajales Dau, D. Graumann, A. Greene, S. Habegger, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, S. Heslin, P. Heu, G. Hill, M. R. Hoffmann, H. Y. Huang, T. Huang, A. Huff, W. J. Huggins, et al. Thermalization and criticality on an analogue–digital quantum simulator. *Nature*, 638(8049):79–85, 2025.

[20] Rahul Trivedi, Adrian Franco Rubio, and J. Ignacio Cirac. Quantum advantage and stability to errors in analogue quantum simulators. *Nature Communications*, 15(1):6507, 2024.

[21] Dominik Hangleiter, Ingo Roth, Jonáš Fuksa, Jens Eisert, and Pedram Roushan. Robustly learning the hamiltonian dynamics of a superconducting quantum processor. *Nature Communications*, 15(1):9595, 2024.

[22] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[23] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum approximate optimization algorithm. *arXiv:1411.4028*, 2014.

[24] Yudong Cao, Jonathan Romero, Jonathan P. Olson, Matthias Degroote, Peter D. Johnson, Mária Kieferová, Ian D. Kivlichan, Tim Menke, Borja Peropadre, Nicolas P. D. Sawaya, Sukin Sim, Libor Veis, and Alán Aspuru-Guzik. Quantum chemistry in the age of quantum computing. *Chemical Reviews*, 119(19):10856–10915, 2019.

[25] F. Segatta, L. Cupellini, M. Garavelli, and B. Mennucci. Quantum chemical modeling of the photoinduced activity of multichromophoric biosystems: Focus review. *Chem. Rev.*, 119:9361, 2019.

[26] Austin G. Fowler et al. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86:032324, 2012.

[27] Maximilian Schlosshauer. Decoherence, the measurement problem, and interpretations of quantum mechanics. *Reviews of Modern Physics*, 76(4):1267–1305, 2005.

[28] Nathalie P. de Leon, Kohei M. Itoh, Dohun Kim, Karan K. Mehta, Tracy E. Northup, Hanhee Paik, B. S. Palmer, N. Samarth, Sorawis Sangtawesin, and D. W. Steuerman. Materials challenges and opportunities for quantum computing hardware. *Science*, 372(6539):eabb2823, 2021.

[29] Rajeev Acharya, Dmitry A. Abanin, Laleh Aghababaie-Beni, Igor Aleiner, Trond I. Andersen, Markus Ansmann, Frank Arute, Kunal Arya, Abraham Asfaw, Nikita Astrakhantsev, Juan Atalaya, Ryan Babbush, Dave Bacon, Brian Ballard, Joseph C. Bardin, Johannes Bausch, Andreas Bengtsson, Alexander Bilmes, Sam Blackwell, Sergio Boixo, Gina Bortoli, Alexandre Bourassa, Jenna Bovaird, Leon Brill, Michael Broughton, David A. Browne, Brett Buchea, Bob B. Buckley, David A. Buell, Tim Burger, Brian Burkett, Nicholas Bushnell, Anthony Cabrera, Juan Campero, Hung-Shen Chang, Yu Chen, Zijun Chen, Ben Chiaro, Desmond Chik, Charina Chou, Jahan Claes, Agnetta Y. Cleland, Josh Cogan, Roberto Collins, Paul Conner, William Courtney, Alexander L. Crook, Ben Curtin, Sayan Das, Alex Davies, Laura De Lorenzo, Dripto M. Debroy, Sean Demura, Michel Devoret, Agustin Di Paolo, Paul Donohoe, Ilya Drozdov, Andrew Dunsworth, Clint Earle, Thomas Edlich, Alec Eickbusch, Aviv Moshe Elbag, Mahmoud Elzouka, Catherine Erickson, Lara Faoro, Edward Farhi, Vinicius S. Ferreira, Leslie Flores Burgos, Ebrahim Forati, Austin G. Fowler, Brooks Foxen, Suhas Ganjam, Gonzalo Garcia, Robert Gasca,

Élie Genois, William Giang, Craig Gidney, Dar Gilboa, Raja Gosula, Alejandro Grajales Dau, Dietrich Graumann, Alex Greene, Jonathan A. Gross, Steve Habegger, John Hall, Michael C. Hamilton, Monica Hansen, Matthew P. Harrigan, Sean D. Harrington, Francisco J. H. Heras, Stephen Heslin, Paula Heu, Oscar Higgott, Gordon Hill, Jeremy Hilton, George Holland, Sabrina Hong, Hsin-Yuan Huang, Ashley Huff, William J. Huggins, et al. Quantum error correction below the surface code threshold. *Nature*, 638(8052):920–926, 2025.

[30] Raymond Hill. *A First Course in Coding Theory.* Oxford Applied Mathematics and Computing Science Series. Oxford University Press, 1986.

[31] Benjamin W. Reichardt, David Aasen, Rui Chao, Ivan Basov, Juan M. Bello-Rivas, Parsa Bonderson, ..., and Benjamin J. Bloom. Fault-tolerant quantum computation with a neutral atom processor. *arXiv preprint*, 2024. Demonstration of logical-qubit encoding and fault-tolerant operations across 256 neutral-atom qubits, including Bernstein–Vazirani algorithm and error detection via erasure conversion.

[32] Dolev Bluvstein, Alexandra A. Geim, Sophie H. Li, Simon J. Evered, J. Pablo Bonilla Ataides, Gefen Baranes, Andi Gu, Tom Manovitz, Muqing Xu, Marcin Kalinowski, Shayan Majidy, Christian Kokail, Nishad Maskara, Elias C. Trapp, Luke M. Stewart, Simon Hollerith, Hengyun Zhou, Michael J. Gullans, Susanne F. Yelin, Markus Greiner, Vladan Vuletić, Madelyn Cain, and Mikhail D. Lukin. Architectural mechanisms of a universal fault-tolerant quantum computer. *arXiv preprint*, 2025.

[33] Loïc Henriet, Lucas Béguin, Adrien Signoles, Thierry Lahaye, and Antoine Browaeys. Quantum computing with neutral atoms. *Quantum*, 4:327, 2020.

[34] Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Tout T. Wang, Sepehr Ebadi, Hannes Bernien, Markus Greiner, Vladan Vuletić, Hannes Pichler, and Mikhail D. Lukin. Parallel implementation of high-fidelity multiqubit gates with neutral atoms. *Physical Review Letters*, 123(17):170503, 2019.

[35] Simon J. Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H. Li, Alexandra A. Geim, Tout T. Wang, Nishad Maskara, Harry Levine, Giulia Semeghini, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. High-fidelity parallel entangling gates on a neutral-atom quantum computer. *Nature*, 622(7982):268–272, 2023.

[36] Dolev Bluvstein, Harry Levine, Giulia Semeghini, Tout T. Wang, Sepehr Ebadi, Marcin Kalinowski, Alexander Keesling, Nishad Maskara, Hannes Pichler, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. A quantum processor based on coherent transport of entangled atom arrays. *Nature*, 604(7906):451–456, 2022.

[37] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Advances in quantum metrology. *Nature Photonics*, 5(4):222–229, 2011.

[38] Andrew D. Ludlow, Martin M. Boyd, Jun Ye, E. Peik, and P. O Schmidt. Optical atomic clocks. *Reviews of Modern Physics*, 87(2):637–701, 2015.

[39] Jun Ye and Peter Zoller. Essay: Quantum sensing with atomic, molecular, and optical platforms for fundamental physics. *Physical Review Letters*, 132(19):190001, 2024.

[40] Shimon Kolkowitz, Igor Pikovski, Nicholas Langellier, Mikhail D. Lukin, Ronald L. Walsworth, and Jun Ye. Gravitational wave detection with optical lattice atomic clocks. *Physical Review D*, 94:124043, December 2016.

[41] 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, 2017.

[42] William R. Milner, Stefan Lannig, Mikhail Mamaev, Lingfeng Yan, Anjun Chu, Ben Lewis, Max N. Frankel, Ross B. Hutson, Ana Maria Rey, and Jun Ye. Coherent evolution of superexchange interaction in seconds-long optical clock spectroscopy. *Science*, 388(6746):503–508, 2025.

[43] Ran Finkelstein, Richard Bing-Shiun Tsai, Xiangkai Sun, Pascal Scholl, Su Direkci, Tuvia Gefen, Joonhee Choi, Adam L. Shaw, and Manuel Endres. Universal quantum operations and ancilla-based read-out for tweezer clocks. *Nature*, 634(8033):321–327, 2024.

[44] Alec Cao, William J. Eckner, Theodor Lukin Yelin, Aaron W. Young, Sven Jandura, Lingfeng Yan, Kyungtae Kim, Guido Pupillo, Jun Ye, Nelson Darkwah Oppong, and Adam M. Kaufman. Multi-qubit gates and schrödinger cat states in an optical clock. *Nature*, 634(8033):315–320, 2024.

[45] Vitaly Fedoseev, Hanzhen Lin, Yu-Kun Lu, Yoo Kyung Lee, Jiahao Lyu, and Wolfgang Ketterle. Coherent and incoherent light scattering by single-atom wave packets. *Physical Review Letters*, 135(4):043601, 2025.

[46] Stuart J. Masson and Ana Asenjo-Garcia. Universality of dicke superradiance in arrays of quantum emitters. *Nature Communications*, 13(1):2285, 2022.

[47] Stuart J. Masson, Jacob P. Covey, Sebastian Will, and Ana Asenjo-Garcia. Dicke superradiance in ordered arrays of multilevel atoms. *PRX Quantum*, 5(1):010344, 2024.

[48] Jun Rui, David Wei, Antonio Rubio-Abadal, Simon Hollerith, Johannes Zeiher, Dan M. Stamper-Kurn, Christian Gross, and Immanuel Bloch. A subradiant optical mirror formed by a single structured atomic layer. *Nature*, 583(7816):369–374, 2020.

[49] Zhenjie Yan, Jacquelyn Ho, Yue-Hui Lu, Stuart J. Masson, Ana Asenjo-Garcia, and Dan M. Stamper-Kurn. Superradiant and subradiant cavity scattering by atom arrays. *Physical Review Letters*, 131(25):253603, 2023.

[50] Kritsana Srakaew, Pascal Weckesser, Simon Hollerith, David Wei, Daniel Adler, Immanuel Bloch, and Johannes Zeiher. A subwavelength atomic array switched by a single rydberg atom. *Nature Physics*, 19(5):714–719, 2023.

[51] 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, 2010.

[52] Manuel Endres, Hannes Bernien, Alexander Keesling, Harry Levine, Eric R. Anschuetz, Alexandre Krajenbrink, Crystal Senko, Vladan Vuletic, Markus Greiner, and Mikhail D. Lukin. Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science*, 354(6315):1024–1027, 2016.

[53] Andreas Reiserer and Gerhard Rempe. Cavity-based quantum networks with single atoms and optical photons. *Reviews of Modern Physics*, 87(4):1379–1418, 2015.

[54] Lijun Ma, Oliver Slattery, and Xiao Tang. Optical quantum memory based on electromagnetically induced transparency. *J. Opt.*, 19(4):043001, 2017. PMCID: PMC5562294.

[55] Rui Zhang, Sean R. Garner, and Lene Vestergaard Hau. Creation of long-term coherent optical memory via controlled nonlinear interactions in bose-einstein condensates. *Physical Review Letters*, 103(23):233602, 2009.

[56] Jacob F. Sherson, Christof Weitenberg, Manuel Endres, Marc Cheneau, Immanuel Bloch, and Stefan Kuhr. Single-atom-resolved fluorescence imaging of an atomic Mott insulator. *Nature*, 467(7311):68–72, 2010.

[57] Waseem S. Bakr, Jonathon I. Gillen, Amy Peng, Simon Fölling, and Markus Greiner. A quantum gas microscope for detecting single atoms in a hubbard-regime optical lattice. *Nature*, 462(7269):74–77, 2009.

[58] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. Many-body physics with ultracold gases. *Reviews of Modern Physics*, 80(3):885–964, 2008.

[59] Lawrence W. Cheuk, Matthew A. Nichols, Melih Okan, Thomas Gersdorf, Vinay V. Ramasesh, Waseem S. Bakr, Thomas Lompe, and Martin W. Zwierlein. Quantum-gas microscope for fermionic atoms. *Physical Review Letters*, 114(19):193001, 2015.

[60] Elmar Haller, James Hudson, Andrew Kelly, Dylan A. Cotta, Bruno Peaudecerf, Graham D. Bruce, and Stefan Kuhr. Single-atom imaging of fermions in a quantum-gas microscope. *Nature Physics*, 11(9):738–742, 2015.

[61] Marc Cheneau, Peter Barmettler, Dario Poletti, Manuel Endres, Peter Schauß, Takeshi Fukuhara, Christian Gross, Immanuel Bloch, Corinna Kollath, and Stefan Kuhr. Light-cone-like spreading of correlations in a quantum many-body system. *Nature*, 481(7382):484–487, 2012.

[62] Rajibul Islam, Ruichao Ma, Philipp M. Preiss, M. Eric Tai, Alexander Lukin, Matthew Rispoli, and Markus Greiner. Measuring entanglement entropy in a quantum many-body system. *Nature*, 528(7580):77–83, 2015.

[63] Philipp M. Preiss, Ruichao Ma, M. Eric Tai, Alexander Lukin, Matthew Rispoli, Philip Zupancic, Yoav Lahini, Rajibul Islam, and Markus Greiner. Strongly correlated quantum walks in optical lattices. *Science*, 347(6227):1229–1233, 2015.

[64] Christof Weitenberg, Manuel Endres, Jacob F. Sherson, Marc Cheneau, Peter Schauß, Takeshi Fukuhara, Immanuel Bloch, and Stefan Kuhr. Single-spin addressing in an atomic mott insulator. *Nature*, 471(7338):319–324, 2011.

[65] K. Henderson, C. Ryu, C. MacCormick, and M. G. Boshier. Experimental demonstration of painting arbitrary and dynamic potentials for bose–einstein condensates. *New Journal of Physics*, 11(4):043030, 2009.

[66] Philip Zupancic, Philipp M. Preiss, Ruichao Ma, Alexander Lukin, M. Eric Tai, Matthew Rispoli, Rajibul Islam, and Markus Greiner. Ultra-precise holographic beam shaping for microscopic quantum control. *Opt. Express*, 24(13):13881–13893, 2016.

[67] G. Gauthier, I. Lenton, N. McKay Parry, M. Baker, M. J. Davis, H. Rubinsztein-Dunlop, and T. W. Neely. Direct imaging of a digital-micromirror device for configurable microscopic optical potentials. *Optica*, 3(10):1136–1143, 2016.

[68] Daniel Barredo, Sylvain de Léséleuc, Vincent Lienhard, Thierry Lahaye, and Antoine Browaeys. An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science*, 354(6315):1021–1023, 2016.

[69] Manuel Endres, Hannes Bernien, Alexander Keesling, Harry Levine, Eric R. Anschuetz, Alexandre Krajenbrink, Crystal Senko, Vladan Vuletic, Markus Greiner, and Mikhail D. Lukin. Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science*, 354(6315):1024–1027, 2016.

[70] Antoine Browaeys and Thierry Lahaye. Many-body physics with individually controlled rydberg atoms. *Nature Physics*, 16(2):132–142, 2020.

[71] N. Schlosser, G. Reymond, and P. Grangier. Collisional blockade in microscopic optical dipole traps. *Physical Review Letters*, 89(2):023005, 2002.

[72] F. Nogrette, H. Labuhn, S. Ravets, D. Barredo, L. Béguin, A. Vernier, T. Lahaye, and A. Browaeys. Single-atom trapping in holographic 2d arrays of microtraps with arbitrary geometries. *Physical Review X*, 4(2):021034, 2014.

[73] Daniel Barredo, Vincent Lienhard, Sylvain de Léséleuc, Thierry Lahaye, and Antoine Browaeys. Synthetic three-dimensional atomic structures assembled atom by atom. *Nature*, 561(7721):79–82, 2018.

[74] Alexandre Cooper, Jacob P. Covey, Ivaylo S. Madjarov, Sergey G. Porsev, Marianna S. Safronova, and Manuel Endres. Alkaline-earth atoms in optical tweezers. *Physical Review X*, 8(4):041055, 2018.

[75] Kai-Niklas Schymik, Vincent Lienhard, Daniel Barredo, Pascal Scholl, Hannah Williams, Antoine Browaeys, and Thierry Lahaye. Enhanced atom-by-atom assembly of arbitrary tweezer arrays. *Physical Review A*, 102(6):063107, 2020.

[76] Hannes Bernien, Sylvain Schwartz, Alexander Keesling, Harry Levine, Ahmed Omran, Hannes Pichler, Soonwon Choi, Alexander S. Zibrov, Manuel Endres, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Probing many-body dynamics on a 51-atom quantum simulator. *Nature*, 551(7682):579–584, 2017.

[77] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin. Generation and manipulation of schrödinger cat states in rydberg atom arrays. *Science*, 365(6453):570–574, 2019.

[78] Sepehr Ebadi, Tout T. Wang, Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Dolev Bluvstein, Rhine Samajdar, Hannes Pichler, Wen Wei Ho, Soonwon Choi, Subir Sachdev, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature*, 595(7866):227–232, 2021.

[79] Aaron W. Young, William J. Eckner, Nathan Schine, Andrew M. Childs, and Adam M. Kaufman. Tweezer-programmable 2d quantum walks in a hubbard-regime lattice. *Science*, 377(6608):885–889, 2022.

[80] Aaron W. Young, Shawn Geller, William J. Eckner, Nathan Schine, Scott Glancy, Emanuel Knill, and Adam M. Kaufman. An atomic boson sampler. *Nature*, 629(8011):311–316, 2024.

[81] Benjamin M. Spar, Elmer Guardado-Sanchez, Sungjae Chi, Zoe Z. Yan, and Waseem S. Bakr. Realization of a fermi-hubbard optical tweezer array. *Physical Review Letters*, 128(22):223202, 2022.

[82] T. Arpornthip, C. A. Sackett, and K. J. Hughes. Vacuum-pressure measurement using a magneto-optical trap. *Physical Review A*, 85(3):033420, 2012.

[83] Conversion on the laser lifetime with gleb Maslennikov from NKT photonics, 2025. Personal communication, March 2025.

[84] J Thorpe, K Numata, and J Livas. Laser frequency stabilization and control through offset sideband locking to optical cavities. *Optics Express*, 16(20):15980, 9 2008.

[85] Richard Bing-Shiun Tsai, Xiangkai Sun, Adam L. Shaw, Ran Finkelstein, and Manuel Endres. Benchmarking and Fidelity Response Theory of High-Fidelity Rydberg Entangling Gates. *Phys. Rev. X Quantum*, 6(1):010331, 2025.

[86] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts. Double-pass acousto-optic modulator system. *Review of Scientific Instruments*, 76(6), 2005.

[87] Sven Jandura and Guido Pupillo. Time-Optimal Two- and Three-Qubit Gates for Rydberg Atoms. *Quantum*, 6:712–712, 2022.

[88] Shuzhe Yang, Guido Masella, Vase Moeini, Amar Bellahsene, Chang Li, Tom Bienaimé, and Shannon Whitlock. Compact arbitrary optical waveform modulator with digital feedback. *Physical Review Applied*, 23(5):054009, 2025.

[89] Ohad Lib. A note on the pulse phase and amplitude modulation scheme, 2024. Unpublished note.

[90] Yiqian Yang, Andrew Forbes, and Liangcai Cao. A review of liquid crystal spatial light modulators: devices and applications. *Opto-Electronic Science*, 2(8):230026–1–230026–29, 2023.

[91] Sylvain de Léséleuc. *Quantum Simulation of Spin Models with Assembled Arrays of Rydberg Atoms*. Ph.d. thesis, Université Paris-Saclay, 2019. Available at `https://tel.archives-ouvertes.fr/tel-02414602`.

[92] Y. T. Chew, M. Poitrinal, T. Tomita, S. Kitade, J. Mauricio, K. Ohmori, and S. de Léséleuc. Ultraprecise holographic optical tweezer array. *Physical Review A*, 110(5):053518, 2024.

[93] Maximilian Ammenwerth, Hendrik Timme, Veronica Giardini, Renhao Tao, Flavien Gyger, Ohad Lib, Dirk Berndt, Dimitrios Kourkoulos, Tim Rom, Immanuel Bloch, and Johannes Zeiher. Dynamical spatial light modulation in the ultraviolet spectral range. *arXiv preprint arXiv:2504.09529*, April 2025. `https://arxiv.org/abs/2504.09529`.

[94] Renhao Tao, Maximilian Ammenwerth, Flavien Gyger, Immanuel Bloch, and Johannes Zeiher. High-fidelity detection of large-scale atom arrays in an optical lattice. *Physical Review Letters*, 133(1):013401, 2024.

[95] Adam L. Shaw, Zhuo Chen, Joonhee Choi, Daniel K. Mark, Pascal Scholl, Ran Finkelstein, Andreas Elben, Soonwon Choi, and Manuel Endres. Benchmarking highly entangled states on a 60-atom analogue quantum simulator. *Nature*, 628(8006):71–77, 2024.

[96] Hannah J. Manetsch, Gyohei Nomura, Elie Bataille, Kon H. Leung, Xudong Lv, and Manuel Endres. A tweezer array with 6100 highly coherent atomic qubits. arXiv:2403.12021.

[97] T. W. Hsu, W. Zhu, T. Thiele, M. O. Brown, S. B. Papp, A. Agrawal, and C. A. Regal. Single-atom trapping in a metasurface-lens optical tweezer. *PRX Quantum*, 3(3):030316, 2022.

[98] Zhe Shen and Xinyu Huang. A review of optical tweezers with metasurfaces. *Photonics*, 10(6):623, 2023.

[99] Donghao Li, Qiming Liao, Beining Xu, Thomas Zentgraf, Emmanuel Narvaez Castaneda, Yaoting Zhou, Keyu Qin, Zhongxiao Xu, Heng Shen, and Lingling Huang. In vacuum metasurface for optical microtrap array. *arXiv preprint arXiv:2407.05918*, July 2024. Submitted July 8, 2024; revised May 23, 2025 (v2).

[100] Aaron Holman, Yuan Xu, Ximo Sun, Jiahao Wu, Mingxuan Wang, Bojeong Seo, Nanfang Yu, and Sebastian Will. Trapping of single atoms in metasurface optical tweezer arrays. *arXiv preprint arXiv:2411.05321*, November 2024. Submitted November 8, 2024; revised November 12, 2024.

[101] Parinaz Barakhshan, Adam Marrs, Akshay Bhosale, Bindiya Arora, Rudolf Eigenmann, and Marianna S. Safronova. Portal for high-precision atomic data and computation (version 2.0). *University of Delaware*, Newark, DE, USA. `https://www.udel.edu/atom`, 2022. Accessed February 2022.

[102] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75(1):281–324, 2003.

[103] A. M. Kaufman, B. J. Lester, and C. A. Regal. Cooling a single atom in an optical tweezer to its quantum ground state. *Physical Review X*, 2(4):041014, 2012.

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

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

[106] Ivaylo S. Madjarov, Jacob P. Covey, Adam L. Shaw, Joonhee Choi, Anant Kale, Alexandre Cooper, Hannes Pichler, Vladimir Schkolnik, Jason R. Williams, and

Manuel Endres. High-fidelity entanglement and detection of alkaline-earth rydberg atoms. *Nature Physics*, 16(8):857–861, 2020.

[107] Aaron W. Young, William J. Eckner, William R. Milner, Dhruv Kedar, Matthew A. Norcia, Eric Oelker, Nathan Schine, Jun Ye, and Adam M. Kaufman. Half-minute-scale atomic coherence and high relative stability in a tweezer clock. *Nature*, 588(7838):408–413, 2020.

[108] Donggyu Kim, Alexander Keesling, Ahmed Omran, Harry Levine, Hannes Bernien, Markus Greiner, Mikhail D. Lukin, and Dirk R. Englund. Large-scale uniform optical focus array generation with a phase spatial light modulator. *Opt. Lett.*, 44(12):3178–3181, 2019.

[109] P. S. Salter and M. J. Booth. *Designing and aligning optical systems incorporating Liquid crystal spatial light modulators (SLMs)*, February 2020.

[110] Dimitrios Tsevas. Precise shaping of coherent laser beams with a liquid-crystal-on-silicon spatial light modulator, unpublished.

[111] Ryuta Yamamoto, Jun Kobayashi, Takuma Kuno, Kohei Kato, and Yoshiro Takahashi. An ytterbium quantum gas microscope with narrow-line laser cooling. *New. J. Phys.*, 18(2):023016, 2016.

[112] Emma Deist, Justin A. Gerber, Yue-Hui Lu, Johannes Zeiher, and Dan M. Stamper-Kurn. Superresolution Microscopy of Optical Fields Using Tweezer-Trapped Single Atoms. *Phys. Rev. Lett.*, 128(8):083201, 2022.

[113] David Wei. *Microscopy of spin hydrodynamics and cooperative light scattering in atomic Hubbard systems*. Dissertation, LMU Munich, 2023.

[114] David Wei. state_reconstruction: Fluorescence emission state reconstruction from microscopic fluorescence images of atoms on a lattice. `https://github.com/david-wei/state_reconstruction`, 2023. Accessed: 2025-07-15.

[115] Nicolas Schlosser, Georges Reymond, Igor Protsenko, and Philippe Grangier. Sub-poissonian loading of single atoms in a microscopic dipole trap. *Nature*, 411(6841):1024–1027, 2001.

[116] Yin Hsien Fung, Pimonpan Sompet, and Mikkel F. Andersen. Single atoms preparation using light-assisted collisions. *Technologies*, 4(1):4, 2016.

[117] Jacob P. Covey, Ivaylo S. Madjarov, Alexandre Cooper, and Manuel Endres. 2000-times repeated imaging of strontium atoms in clock-magic tweezer arrays. *Physical Review Letters*, 122(17):173201, 2019.

[118] M. A. Norcia, H. Kim, W. B. Cairncross, M. Stone, A. Ryou, M. Jaffe, M. O. Brown, K. Barnes, P. Battaglino, T. C. Bohdanowicz, A. Brown, K. Cassella, C. A. Chen, R. Coxe, D. Crow, J. Epstein, C. Griger, E. Halperin, F. Hummel, A. M. W. Jones, J. M. Kindem, J. King, K. Kotru, J. Lauigan, M. Li, M. Lu, E. Megidish, J. Marjanovic, M. McDonald, T. Mittiga, J. A. Muniz, S. Narayanaswami, C. Nishiguchi, T. Paule, K. A. Pawlak, L. S. Peng, K. L. Pudenz, D. Rodríguez Pérez, A. Smull, D. Stack, M. Urbanek, R. J. M. van de Veerdonk, Z. Vendeiro, L. Wadleigh, T. Wilkason, T. Y. Wu, X. Xie, E. Zalys-Geller, X. Zhang, and B. J. Bloom. Iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices. *PRX Quantum*, 5(3):030316, 2024.

[119] J. Scott, M. Lim, H. U. Singla, Q. Meece, T. Choy, J. S. Kolkowitz, M. Graham, T. and M. Saffman. Laser cooling and qubit measurements on a forbidden transition in neutral cs atoms. *arXiv preprint*, quant-ph/2507.01720, 2025. Submitted 2 July 2025.

[120] Lin Su, Alexander Douglas, Michal Szurek, Anne H. Hébert, Aaron Krahn, Robin Groth, Gregory A. Phelps, Ognjen Marković, and Markus Greiner. Fast single atom imaging for optical lattice arrays. *Nature Communications*, 16(1):1017, 2025.

[121] Mareike Hetzel, Luca Pezzè, Cebrail Pür, Martin Quensen, Andreas Hüper, Jiao Geng, Jens Kruse, Luis Santos, Wolfgang Ertmer, Augusto Smerzi, and Carsten Klempt. Tomography of a number-resolving detector by reconstruction of an atomic many-body quantum state. *Physical Review Letters*, 131(26):260601, 2023.

[122] S. Snigirev, A. J. Park, A. Heinz, I. Bloch, and S. Blatt. Fast and dense magneto-optical traps for strontium. *Physical Review A*, 99(6):063421, 2019.

[123] 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, 2019.

[124] Tobias Bothwell, Dhruv Kedar, Eric Oelker, John M. Robinson, Sarah L. Bromley, Weston L. Tew, Jun Ye, and Colin J. Kennedy. Jila sri optical lattice clock with uncertainty of 2.0×10^{-18}. *Metrologia*, 56(6):065004, 2019.

[125] Shuo Ma, Genyue Liu, Pai Peng, Bichen Zhang, Sven Jandura, Jahan Claes, Alex P. Burgers, Guido Pupillo, Shruti Puri, and Jeff D. Thompson. High-fidelity gates and mid-circuit erasure conversion in an atomic qubit. *Nature*, 622(7982):279–284, 2023.

[126] Yue Wu, Shimon Kolkowitz, Shruti Puri, and Jeff D. Thompson. Erasure conversion for fault-tolerant quantum computing in alkaline earth rydberg atom arrays. *Nature Communications*, 13(1):4657, 2022.

[127] Yuma Nakamura, Toshi Kusano, Rei Yokoyama, Keito Saito, Koichiro Higashi, Naoya Ozawa, Tetsushi Takano, Yosuke Takasu, and Yoshiro Takahashi. Hybrid atom tweezer array of nuclear spin and optical clock qubits. *Physical Review X*, 14(4):041062, 2024.

[128] Bichen Zhang, Genyue Liu, Guillaume Bornet, Sebastian P. Horvath, Pai Peng, Shuo Ma, Shilin Huang, Shruti Puri, and Jeff D. Thompson. Leveraging erasure errors in logical qubits with metastable ^{171}Yb atoms. *arXiv preprint arXiv:2506.13724*, June 2025.

[129] G. Unnikrishnan, P. Ilzhöfer, A. Scholz, C. Hölzl, A. Götzelmann, R. K Gupta, J. Zhao, J. Krauter, S. Weber, N. Makki, H. P Büchler, T. Pfau, and F. Meinert. Coherent Control of the Fine-Structure Qubit in a Single Alkaline-Earth Atom. *Phys. Rev. Lett.*, 132(15):150606, 2024.

[130] S. Pucher, V. Klüsener, F. Spriestersbach, J. Geiger, A. Schindewolf, I. Bloch, and S. Blatt. Fine-structure qubit encoded in metastable strontium trapped in an optical lattice. *Physical Review Letters*, 132(15):150605, 2024.

[131] Maximilian Ammenwerth, Hendrik Timme, Flavien Gyger, Renhao Tao, Immanuel Bloch, and Johannes Zeiher. Realization of a fast triple-magic all-optical qutrit in strontium-88. arXiv:2411.02869.

[132] H. X. Yang, J. Y. Ma, Y. K. Wu, Y. Wang, M. M. Cao, W. X. Guo, Y. Y. Huang, L. Feng, Z. C. Zhou, and L. M. Duan. Realizing coherently convertible dual-type qubits with the same ion species. *Nature Physics*, 18(9):1058–1061, 2022.

[133] Katrina Barnes, Peter Battaglino, Benjamin J. Bloom, Kayleigh Cassella, Robin Coxe, Nicole Crisosto, Jonathan P. King, Stanimir S. Kondov, Krish Kotru, Stuart C. Larsen, Joseph Lauigan, Brian J. Lester, Mickey McDonald, Eli Megidish, Sandeep Narayanaswami, Ciro Nishiguchi, Remy Notermans, Lucas S. Peng, Albert Ryou, Tsung-Yao Wu, and Michael Yarwood. Assembly and coherent control of a register of nuclear spin qubits. *Nature Communications*, 13(1):2779, 2022.

[134] 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. *Physical Review Letters*, 96(8):083001, 2006.

[135] Tsung-Yao Wu, Aishwarya Kumar, Felipe Giraldo, and David S. Weiss. Stern–Gerlach detection of neutral-atom qubits in a state-dependent optical lattice. *Nat. Phys.*, 15(6):538–542, 2019.

[136] Pascal Scholl, Adam L. Shaw, Richard Bing-Shiun Tsai, Ran Finkelstein, Joonhee Choi, and Manuel Endres. Erasure conversion in a high-fidelity Rydberg quantum simulator. *Nature*, 622(7982):273–278, 2023.

[137] M. A Norcia, W. B Cairncross, K. Barnes, et al. Midcircuit Qubit Measurement and Rearrangement in a ^{171}Yb Atomic Array. *Phys. Rev. X*, 13(4):041034, 2023.

[138] Andrea Bergschneider, Vincent M. Klinkhamer, Jan Hendrik Becher, Ralf Klemt, Gerhard Zürn, Philipp M. Preiss, and Selim Jochim. Spin-resolved single-atom imaging of ^{6}Li in free space. *Phys. Rev. A*, 97(6):063613, 2018.

[139] Simon Stellmer and Florian Schreck. Reservoir spectroscopy of $5s5p\ {}^3p_2$–$5snd\ {}^3D_{1,2,3}$ transitions in strontium. *Physical Review A*, 90(2):022512, 2014.

[140] Martin Boll, Timon A. Hilker, Guillaume Salomon, Ahmed Omran, Jacopo Nespolo, Lode Pollet, Immanuel Bloch, and Christian Gross. Spin- and density-resolved microscopy of antiferromagnetic correlations in Fermi-Hubbard chains. *Science*, 353(6305):1257–1260, 2016.

[141] Joannis Koepsell, Sarah Hirthe, Dominik Bourgund, Pimonpan Sompet, Jayadev Vijayan, Guillaume Salomon, Christian Gross, and Immanuel Bloch. Robust Bilayer Charge Pumping for Spin- and Density-Resolved Quantum Gas Microscopy. *Phys. Rev. Lett.*, 125:010403, Jul 2020.

[142] Joanna W. Lis, Aruku Senoo, William F. McGrew, Felix Rönchen, Alec Jenkins, and Adam M. Kaufman. Midcircuit Operations Using the *omg* Architecture in Neutral Atom Arrays. *Phys. Rev. X*, 13(4):041035, 2023.

[143] E. Brion, L. H. Pedersen, and K. Mølmer. Adiabatic elimination in a lambda system. *Journal of Physics A: Mathematical and Theoretical*, 40(5):1033, 2007.

[144] Daniel A. Steck. Quantum and atom optics. `http://steck.us/teaching`, 2021. Revision 0.13.5, available at `http://steck.us/teaching`.

[145] David C. McKay, Christopher J. Wood, Sarah Sheldon, Jerry M. Chow, and Jay M. Gambetta. Efficient z gates for quantum computing. *Physical Review A*, 96(2):022330, 2017.

[146] Nikola Šibalić, Jonathan D. Pritchard, Charles S. Adams, and Kevin J. Weatherill. ARC: An open-source library for calculating properties of alkali rydberg atoms. *Computer Physics Communications*, 220:319–331, 2017.

[147] J. A. Muniz, M. Stone, D. T. Stack, M. Jaffe, J. M. Kindem, L. Wadleigh, E. Zalys-Geller, X. Zhang, C. A. Chen, M. A. Norcia, J. Epstein, E. Halperin, F. Hummel, T. Wilkason, M. Li, K. Barnes, P. Battaglino, T. C. Bohdanowicz, G. Booth, A. Brown, M. O. Brown, W. B. Cairncross, K. Cassella, R. Coxe, D. Crow, M. Feldkamp, C. Griger, A. Heinz, A. M. W. Jones, H. Kim, J. King, K. Kotru, J. Lauigan, J. Marjanovic, E. Megidish, M. Meredith, M. McDonald, R. Morshead, S. Narayanaswami, C. Nishiguchi, T. Paule, K. A. Pawlak, K. L. Pudenz, D. Rodríguez Pérez, A. Ryou, J. Simon, A. Smull, M. Urbanek, R. J. M. van de

Veerdonk, Z. Vendeiro, T. Y. Wu, X. Xie, and B. J. Bloom. High-fidelity universal gates in the ^{171}Yb ground-state nuclear-spin qubit. *PRX Quantum*, 6(2):020334, 2025.

[148] Lukas Homeier, Annabelle Bohrdt, Simon Linsel, Eugene Demler, Jad C. Halimeh, and Fabian Grusdt. Realistic scheme for quantum simulation of Z_2 lattice gauge theories with dynamical matter in (2 + 1)d. *Communications Physics*, 6(1):127, 2023.

[149] Flavien Gyger, Maximilian Ammenwerth, Renhao Tao, Hendrik Timme, Stepan Snigirev, Immanuel Bloch, and Johannes Zeiher. Continuous operation of large-scale atom arrays in optical lattices. *Physical Review Research*, 6(3):033104, 2024.

[150] Robert Jördens, Niels Strohmaier, Kenneth Günter, Henning Moritz, and Tilman Esslinger. A mott insulator of fermionic atoms in an optical lattice. *Nature*, 455(7210):204–207, 2008.

[151] U. Schneider, L. Hackermüller, S. Will, Th. Best, I. Bloch, T. A. Costi, R. W. Helmes, D. Rasch, and A. Rosch. Metallic and insulating phases of repulsively interacting fermions in a 3d optical lattice. *Science*, 322(5907):1520–1525, 2008.

[152] M. Aidelsburger, M. Lohse, C. Schweizer, M. Atala, J. T Barreiro, S. Nascimbène, N. R Cooper, I. Bloch, and N. Goldman. Measuring the chern number of hofstadter bands with ultracold bosonic atoms. *Nature Physics*, 11(2):162–166, 2015.

[153] M. Lohse, C. Schweizer, O. Zilberberg, M. Aidelsburger, and I. Bloch. A thouless quantum pump with ultracold bosonic atoms in an optical superlattice. *Nature Physics*, 12(4):350–354, 2016.

[154] David Chen, Matthew White, Cecilia Borries, and Brian DeMarco. Quantum quench of an atomic mott insulator. *Physical Review Letters*, 106(23):235304, 2011.

[155] Daniel P. Arovas, Erez Berg, Steven A. Kivelson, and Srinivas Raghu. The hubbard model. *Annual Review of Condensed Matter Physics*, 13(Volume 13, 2022):239–274, 2022.

Acknowledgement

I would like to thank Immanuel Bloch for introducing me to the field of AMO physics through his skillful delivery of Quantum Optics lectures at LMU, and for connecting me to my very first project in the polar molecule lab at MPQ. It was absolutely fascinating to watch abstract, and sometimes seemingly academic, concepts take on a tangible, physical form in the research lab. I learned a great deal about cavity tuning and laser trapping from Roman Bause, Marcel Duda, Xingyan Chen, Andreas Schindewolf, the regulars in the lab, and of course, Xinyu Luo—the mild-tempered group leader with a deep understanding of what makes a molecule tick.

I would like to thank Johannes Zeiher for taking me on as the very first graduate student in the brand-new Strontium Rydberg experiment, where the work for this thesis was carried out. It has been both rewarding and fun to build an experiment from the ground up under his tireless guidance. To say that his approach has shaped my perspective on how to manage a research lab and foster collaboration would be a serious understatement. I am also grateful to Flavien Gyger and Maximilian Ammenwerth, with whom I spent most of my time in the lab. It is thanks to our teamwork that the experiment has developed into what it is today. Working with self-motivated colleagues who are unafraid to challenge each other when opinions differ has been both stimulating and valuable. I am especially indebted to Flavien, whose kindness and willingness to help never wavered. On more than one occasion, I took advantage of the proximity of his apartment to the lab by asking him to come in at off-hours to fix experimental issues while I run measurements remotely in the comfort of my couch. Last, I want to thank Hendrik Timme and Ohad Lib, who brought fresh energy into our experiments. Each of them contributed valuable expertise and enthusiasm to the project. I have no doubt that, in their hands, the experiment has a world of exciting results ahead.

Our lab has greatly benefited from a steady influx of talented junior students, whether during their internships or while completing their master's theses with us. They are Aron, Alexandra, Efe, Hannah, Pranav, Simon, Niklas, Esha, Maurice, Ben, Taylor, and Veronica. Their substantial contributions have expanded the experiment's capabilities and significantly accelerated our progress. As someone working in an experimental lab, I also want to highlight the invaluable work of our engineers—Anton Meyer, Karsten Förster, and Olivia Mödl. And as an international student in Germany, I would be remiss not to acknowledge the tremendous support from our secretaries—Kristina Schuldt, Doreen Seidl, Veronika Seidl, and Ildiko Kecskesi.

Finally, I want to thank my parents, who have always let me pursue my interests without interfering in my decisions. This is something that, I must admit, is quite uncharacteristic of Asian parents. I am especially grateful to my wife, Jingdi, whose companionship has meant so much to me and who took on most of the day-to-day responsibilities so I could focus entirely on my research.