Quantum Gates and Bose-Hubbard Models with Dipolar Systems

Michael Hughes
Keble College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2022

I acknowledge the help of co-authors, advice from group members, especially Jordi Mur-Petit and Martin Kiffner, and the supervision of Dieter Jaksch. Acknowledgements relating to specific chapters are included at the beginning of the relevant chapter. Non-academic acknowledgements include my family, especially Will.

Quantum Gates and Bose-Hubbard models with Dipolar Systems

Michael Hughes

Keble College

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term 2022

Abstract

Ultracold dipolar systems, featuring strong long-range dipole-dipole interactions, have been brought under increasing quantum control over recent years. In this thesis, we study the use of these properties for quantum information processing and simulation of many-body Bose-Hubbard models with off-site interactions. We first investigate robust entangling protocols for polar molecules, finding that shaped microwave pulses provide two-qubit entangling gates based on the dipole-dipole interaction with robustness to experimentally-relevant errors. We then numerically study the ground state properties of hard-core dipolar bosons confined in cylindrical optical lattices, where the curved lattice surface causes the anisotropic dipole-dipole interactions to vary around the ring of the cylinder. This expands the physics of analogous square lattice models due to the cooperation and competition of the interactions in different directions. Finally, we investigate the applicability of the lowest-band dipolar Bose-Hubbard model itself by comparing lattice and continuum methods for small systems of strongly interacting dipolar bosons in optical lattice potentials.

Contents

1

Introduction

Dipolar systems offer opportunities to expand on the quantum applications of short-range interacting atoms. While experimentally more challenging, they interact strongly with external fields and with each other over long distances due to the dipole-dipole interaction. This enables significant interactions between particles in separate traps or separate sites of optical lattices. In this thesis, we focus on the use of these properties for quantum computation and simulation purposes. We first introduce these two aims and then discuss the experimental status of the dipolar systems themselves.

1.1 Quantum Computation

Quantum computation has drawn significant interest since it was first proposed that modestly-sized controlled quantum systems could surpass the information processing capabilities of the most advanced classical computers [1]. While a classical computer's bits are restricted to the values '0' and '1', quantum bits can represent superpositions of the computational basis states and become entangled. For circuit models of quantum computation, information processing takes place through a sequence of quantum gates, which simultaneously perform the desired transformation on all computational states while preserving the phase relationship between them. Because of the extra states available to quantum bits, the number of classical bits required to exactly simulate an arbitrary quantum system grows exponentially with the size of the quantum system, meaning that a universal quantum computer with 50 qubits would be able to perform calculations beyond the realistic capabilities of classical computing. Certain quantum algorithms, such as Grover's algorithm [2] for unstructured search problems or Shor's algorithm [3] for factoring large numbers, are known to perform useful tasks with improved scaling performance compared to the best known classical algorithms.

While the theoretical power of quantum computation is clear, experimental implementation has proved challenging. A practical device for quantum computation must satisfy the DiVincenzo criteria [4] which require the ability to initialise, process, and extract information before it is lost to the environment. This requires extremely precise control and well-isolated systems because the advantages of quantum computing rely on delicate phase coherences and large entangled states. The fact that quantum information cannot be copied or cleanly extracted during a computation also complicates the correction of errors due to imperfect implementation, meaning that a useful near-future device would likely require far more than 50 physical qubits.

Current quantum devices fall into the NISQ (noisy intermediate-scale quantum) regime (reviewed in Ref [5]) where there are enough qubits to surpass classical computation in principle but not enough to implement error correction as well. There have been recent demonstrations of quantum computation using ≈ 50 logical qubits for sampling tasks [6–8]. While these are important milestones, they do not yet offer universal fault-tolerant computing beyond classical capabilities. Multiple platforms (reviewed in Ref [9]) have been investigated for this formidable task with their own advantages and disadvantages, although none has emerged as the single dominant candidate.

Although behind in experimental maturity due to their greater complexity, polar molecules have the advantages of a large controllable electric dipole moment and a large number of long-lived hyperfine states which can be manipulated by microwaves [10, 11] and potentially used to store quantum information [12–14]. This allows for resilience to spontaneous decay and strong interactions despite charge neutrality. These properties have led to multiple proposals to perform entangling gates using polar molecules [15–20].

A necessary condition to harness the potential of polar molecules for quantum information applications is that the entangling protocols operate successfully in the presence of significant experimental imperfections. In chapter 2, we show that the flexibility of choice of qubit states together with shaped microwave pulses allows the dipole-dipole interaction to perform a robust entangling gate for two molecules in separate optical tweezers [21].

1.2 Quantum Simulation of Many-Body Systems

As well as providing a theoretical acceleration for quantum computation compared to classical computation, the difficulty in classically simulating moderate numbers of quantum particles hinders our understanding of the physics of even very simple quantum many-body systems. Strongly correlated quantum materials host a wide range of potentially useful physics (such as high-temperature superconductivity or topological states) whose underlying mechanisms are poorly understood.

Condensed matter systems which demonstrate these properties are often difficult to modify and contain defects, which makes it challenging to explore different physical parameters and explain their behaviour. An alternative approach is to build a well-controlled quantum system to mimic the original system in order to benefit from controllability and the quantum scaling advantage. The system can be designed to simulate the physics of a simple model which may be an abstract model or conjectured to accurately describe the properties of an existing condensed matter system. In both cases, the motivation is to connect the often simple interactions between particles to the diverse emergent behaviour of many particles while, in the latter case, successful quantum simulation may help to verify under which conditions the model describes the condensed matter system [22].

In this thesis we focus on analogue quantum simulation, where the desired Hamiltonian is implemented directly by the system (as opposed to via quantum gates or using variational circuits) [22–24]. This requires a lower degree of control and flexibility than a universal quantum computer, which helps to reach system sizes beyond classical computation, but limits each particular system to simulating a smaller class of Hamiltonians.

Over the past decades, the most widespread platform for this technique has been ultracold alkali atoms in optical lattices. These systems offer excellent control over the Hamiltonian parameters in a pristine and regular optical potential while allowing single-site manipulation and measurement [25–28]. Despite these advantages, alkali atoms in their ground electronic state are limited to scattering interactions whose range is much shorter than the distance between optical lattice sites. This restricts their use for simulating models with significant interactions between distant sites. Dipolar systems overcome this constraint using the dipole-dipole interaction (DDI) which has a much shallower $1/R^3$ decay with distance.

The addition of dipolar interactions is predicted to expand the available ground-state phases [29], including density waves, supersolids [30–32], spin liquids [33, 34], and other topological states [35–41]. While dynamical effects of dipolar interactions have been observed [42–45], such delicate and exotic ground states have not yet been experimentally realised in systems with permanent dipole moments because the required temperatures and filling fractions have not been realised. For applicability to near-future experiments, we focus on systems of $\approx 10 - 1000$ lattice sites where finite-size effects can be significant.

The off-site interactions which create these unusual strongly-correlated states can additionally be made to vary across the surface of the lattice by combining the anisotropic DDI with curved lattices. In chapter 4 we study the dipolar Bose-Hubbard model in a real-space cylindrical lattice, showing that this mechanism leads to a rich phase diagram with enhanced entanglement caused by the competition of these spatially-modulated long-range interactions [46].

The use of ultracold dipolar systems in optical lattices for quantum simulation largely rests upon the simplified description of the physical system using a discrete lattice model. In chapter 5 we quantitatively investigate the validity of this connection by comparing numerical results for the continuous-space model of dipolar interacting bosons in optical lattice potentials with the corresponding lowest-band dipolar Bose-Hubbard model. We find that the nature of the breakdown of the lattice description for very strong DDI depends on the spatial ordering imposed by the interactions.

While we have focussed on Bose-Hubbard models in this thesis, long-range interactions allow non-trivial spin models to be simulated without (or in addition to) particle tunnelling [34, 47–53]. This has recently been demonstrated in other platforms such as $\approx 10^2$ trapped ions [54] and Rydberg atoms in 1D [55] and 2D reconfigurable optical tweezer arrays [56, 57]. These examples show that the ground-state physics of long-range interacting models in moderately-sized systems can now be observed experimentally.

1.3 Dipolar Systems

In this thesis we focus on ultracold systems with a large long-lived dipole moment, specifically magnetic atoms and polar molecules. (We note that Rydberg atoms in highly-excited electronic states have even larger electric dipole transitions and have a range of applications, including quantum computation and simulation, despite faster rates of spontaneous decay). Highly magnetic atoms (such as Chromium [58], Erbium [59], and Dysprosium [60, 61]) have large unpaired electron spins which create a magnetic dipole moment while for heteronuclear polar molecules, an electric dipole moment can be induced along the axis between the two different nuclei. There has also been recent research into using electronically excited states to create short-lived electric dipole moments in magnetic atoms [62] or incorporating magnetic atoms into polar molecules [63, 64], which could lead to doubly-dipolar systems whose electric and magnetic dipoles can be polarised independently. For completeness, we briefly comment on the status of these platforms for quantum information and simulation purposes.

Magnetic atoms (see [65] for a recent review) have been cooled and trapped in optical lattices, which has been used to observe DDI-affected dynamics in the extended Bose-Hubbard model [44] and for simulation of spin models with large spins $S > \frac{1}{2}$ [50–53] including additional tunnelling [66, 67]. More recently, there has been great interest in using gases of magnetic atoms to create supersolid droplets [68–71], where the DDI is important for creating the density modulation.

Polar molecules have seen swift development over recent years. There are two main experimental processes to produce and cool them: The first is association, where two different species of atom are laser-cooled and then a heteronuclear molecule is formed using a Feshbach resonance and subsequent STIRAP pulses or optical transitions alone [72–74]. This approach has produced optical lattices with filling fractions of around 30% [42, 43, 75–77], built individual molecules in tweezers [78–80], and achieved quantum degeneracy for fermionic molecules [81]. The second is to directly cool the molecule [82–84], which is only feasible for molecules whose vibrational structure allows a relatively closed laser cooling cycle. This newer method has so far struggled to achieve the high densities necessary for optical lattice filling but has created molecules in individual traps [85, 86]. For both categories, rapid progress has been made over the past decade in reducing temperatures and extending coherence times.

1.4 Structure and Contribution of Thesis

In this thesis, we first focus on the robust entangling gate for polar molecules in chapter 2. To introduce the quantum simulation section of the thesis, we show the derivation of the lowest-band dipolar Bose-Hubbard model from interacting dipolar bosons in continuous space in chapter 3, which underlies the research on the cylindrical dipolar Bose-Hubbard model in chapter 4 and whose range of validity is investigated through complementary numerical methods in chapter 5. In chapter 4 we present our results for the cylindrical dipolar Bose-Hubbard model and in chapter 5 we show our comparison of continuous-space and discrete-lattice numerical methods for small systems of interacting dipolar bosons in optical lattice potentials.

The primary contributions of this thesis are the numerical results in chapter 2 concerning an entangling gate for polar molecules, made robust to certain experimental imperfections including population of low-lying motional states, the numerical demonstration of the rich ground-state physics of the dipolar Bose-Hubbard model in cylindrical lattices caused by the anisotropy of the DDI in chapter 4, and the quantitative investigation of the breakdown of the lowest-band dipolar Bose-Hubbard model for strong DDI in chapter 5.

2

Robust Entangling Gate for Polar Molecules

This chapter contains and expands on results published in Ref [21]. MH contributed substantially to the manuscript and performed most of the numerical calculations in this chapter. Exceptions to this are clearly marked. Discussions with Carlos Sánchez-Muñoz, Ruth Le-Sueur, and Erik Torrontegui are acknowledged.

Multiple schemes have been proposed with ultracold molecules to implement two-qubit entangling gates [15,16,19] including in the specific case of ${}^2\Sigma$ molecules (molecules with an unpaired electron spin) [17,18,20]. Other proposals also utilise the large number of long-lived hyperfine states offered by polar molecules to implement Hilbert spaces larger than two dimensions (qudits) [11]. We set out a scheme which uses the large choice of levels to reduce the impact of experimental imperfections while using the dipole-dipole interaction to entangle a pair of molecules in separate optical tweezers.

Section 2.1 covers the derivation of the qubit space and DDI interaction Hamiltonian from the molecular energy level spectrum and discusses error sources due to the control and trapping. Section 2.2 illustrates the application of simple entangling protocols based on adiabatic evolution to polar molecules while section 2.3 extends this to a robust entangling quantum gate using optimal control techniques.

2.1 Molecular Hamiltonian

2.1.1 Single Molecular Spectrum

The entangling protocols presented rely on avoided energy level crossings to make the DDI energy significant. In the absence of external fields, diatomic polar molecules feature four broad energy scales, which we list in descending order: Electronic excitations due to the orbital wavefunction of the electrons, vibrational excitations due to relative motion of the two nuclei along their separation vector, rotational energy due to rotation of the nuclei (modelled as a rigid rotor) around an axis perpendicular to the nuclear separation, and hyperfine energy, due to the alignment of the molecular rotation, nuclear spin, and, for molecules with unpaired electrons, electron spin. We consider ultracold molecules cooled to their electronic and vibrational ground state, leaving the rotational and hyperfine levels to serve as the register for the quantum information [10]. This regime corresponds to temperatures below 1mK [87].

Because the spatial density distribution of rotational eigenstates is symmetric with respect to parity, a molecule in a rotational eigenstate has no electric dipole moment so it is necessary to mix rotational manifolds. One method for doing this is to use a DC electric field [88, 89]. This requires an electric field which is at least of comparable magnitude to the rigid rotor rotational splitting divided by the molecular-frame electric dipole moment in order to bridge the rotational energy gap. This gap varies by around a factor of 100 depending on the molecule and can require very large electric fields. For this reason, we use an alternative scheme to bridge the energy gap between rotational manifolds using resonant microwaves [17, 19, 47].

While such a scheme is generally applicable to polar molecules, we focus on the

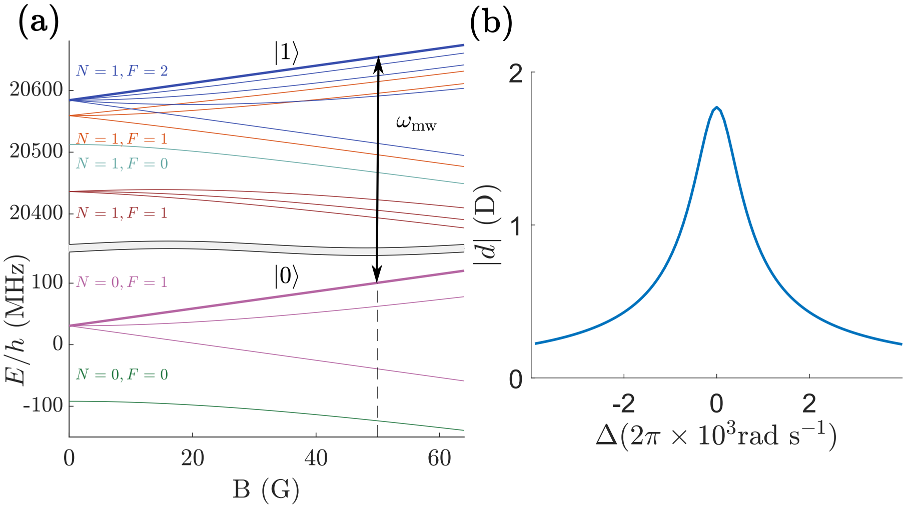

Figure 2.1: (a) Energy levels in the lowest two rotational manifolds of CaF with example qubit states $|0\rangle = |N=0, F=1, m_F=1\rangle$ and $|1\rangle = |N=1, F=2, m_F=2\rangle$ drawn in bold for emphasis and coupled by microwaves. Energy levels provided by M. R. Tarbutt using modelling [10] based on Hamiltonian parameters found experimentally in Refs [90, 91]. (b) Dipole moment $|d|$ of the eigenstates of the single-molecule Hamiltonian as a function of microwave detuning Δ with microwave coupling $\Omega = 2\pi \times 500\text{rads}^{-1}$.

example of CaF, which has recently been trapped and cooled to ultracold temperatures [83, 92–94]. Figure 2.1(a) displays the energy levels of the lowest two rotational manifolds of of the $X^2\Sigma^+(\nu=0)$ state [1] of CaF as a function of magnetic field taken from modelling based on experimental results. CaF is produced from ^{40}Ca, which has two electrons in its outer shell and ^{19}F, which has one hole in its outer shell, meaning that the molecule has one unpaired electron spin $S=\frac{1}{2}$. This isotope of Calcium has no nuclear spin meaning that the total nuclear spin is provided by the Fluorine isotope which has nuclear spin $I=\frac{1}{2}$.

In Figure 2.1(a), the energy levels of the first two rotational manifolds ($N=0$ and $N=1$) are shown, colour coded by the value of the total angular momentum F from the combined nuclear spin, electron spin, and rotational angular momenta.

---

[1] This is molecular spectroscopic notation where X denotes the electronic ground state, $^2\Sigma^+$ means there is one unpaired electron spin, no angular momentum along the internuclear axis, and reflection symmetry about a plane containing the nuclei, and $\nu=0$ denotes the vibrational ground state.

Note that there are two separate branches for $N = 1, F = 1$ where the upper branch corresponds to $J = N + S = \frac{3}{2}$ and the lower branch has $J = \frac{1}{2}$. Each group of a given value of F contains $2F + 1$ energy levels indexed by m_F, the projection of total angular momentum onto the magnetic field quantisation axis, which takes integer values $-F \leq m_F \leq F$. In a finite magnetic field, these eigenstates contain small mixtures of different F (and to a lesser extent N). Nevertheless, we label the states at finite magnetic field by their connection to the zero-field states where these labels are good quantum numbers. At a magnetic field of 50 Gauss, the angular momentum projections are sufficiently split to allow microwaves of the correct polarisation to resonantly couple two chosen levels with negligible probability of exciting others. These two hyperfine states form the 'qubit' space.

As an example, we consider $|0\rangle = |N = 0, F = 1, m_F = 1\rangle$ and $|1\rangle = |N = 1, F = 2, m_F = 2\rangle$. At a magnetic field of 50 Gauss, these states have energy difference $\hbar\omega_{mol} = h \times 20.55$ GHz. These states have a large transition dipole moment and similar magnetic moments. This latter property means that ω_{mol} is relatively stable to unintended Zeeman shifts due to magnetic field instability and has a sensitivity of 0.104(4)kHz/Gauss [95]. Alternative choices of states, such as $|0\rangle = |N = 0, F = 1, m_F = 0\rangle$ and $|1\rangle = |N = 1, J = \frac{1}{2}, F = 0, m_F = 0\rangle$, also offer large dipole moments and good coherence properties [96]. Similar criteria apply to other molecules.

2.1.2 Single Molecule Qubit Space

In this section and section 2.1.3 we show a standard derivation of the Hamiltonian in the molecular qubit space.

The qubit states are linked by an electric dipole transition. In this qubit space, the electric dipole moment operator is given by

$$\hat{\boldsymbol{d}} = \boldsymbol{d}_{10}\sigma_x, \tag{2.1}$$

where $\boldsymbol{d_{10}}$ is the transition dipole moment between states $|0\rangle$ and $|1\rangle$ and σ_x is the Pauli operator $|0\rangle\langle 1| + |1\rangle\langle 0|$. A time-dependent electric field $\boldsymbol{E}(t)$ couples to this electric dipole with energy $-\boldsymbol{d}_{10}.\boldsymbol{E}(t)$. This means that when the electric field consists of microwaves of frequency ω_{mw} and the correct polarisation, the Hamiltonian for a single molecule in the qubit space is

$$H_{mol} = \hbar(\frac{1}{2}\omega_{mol}\sigma_z + \Omega\cos(\omega_{mw}t)\sigma_x), \tag{2.2}$$

where $\sigma_z = |1\rangle\langle 1| - |0\rangle\langle 0|$ and Ω (assumed to be real for simplicity at present) is the maximum microwave coupling strength per cycle. We then transform the Hamiltonian into the interaction picture with respect to the rotating frame $U_{frame}(t) = e^{-i\omega_{mw}t/2}|0\rangle\langle 0| + e^{i\omega_{mw}t/2}|1\rangle\langle 1|$ as standard for driven two-level systems. In order that the time evolution of a general pure state respects this transformation, the Hamiltonian in the rotating frame is related to the Hamiltonian in the original frame by

$$\begin{aligned} H_{mol,rotating} &= i\hbar\frac{\partial U_{frame}(t)}{\partial t}U^{\dagger}_{frame}(t) + U_{frame}(t)H_{mol}U^{\dagger}_{frame}(t) \\ &= \begin{pmatrix} -\frac{\hbar}{2}(\omega_{mol}-\omega_{mw}) & \hbar\frac{\Omega}{2}(1+e^{-i\omega_{mw}t}) \\ \hbar\frac{\Omega}{2}(1+e^{i\omega_{mw}t}) & \frac{\hbar}{2}(\omega_{mol}-\omega_{mw}), \end{pmatrix} \end{aligned} \tag{2.3}$$

where the matrix is written in the $\{|0\rangle,|1\rangle\}$ basis.

In this rotating frame, the transformation introduces ‘counter-rotating’ off-diagonal coupling terms which oscillate with angular frequency ω_{mw}. As we are interested in dynamics on a timescale much longer than $1/\omega_{mol}$ or $1/\omega_{mw}$, we can perform the rotating-wave approximation to remove terms which oscillate at this frequency. (This assumes ω_{mol}, $\omega_{mw} >> \Omega$). We arrive at the single-molecule Hamiltonian

$$H_{mol,RWA} = \hbar(\frac{\Delta}{2}\sigma_z + \frac{\Omega}{2}\sigma_x), \tag{2.4}$$

where Δ is the detuning $\omega_{mol} - \omega_{mw}$ which also obeys $\omega_{mol}, \omega_{mw} >> \Delta$. Due to the microwave coupling, the eigenstates of this Hamiltonian have a non-zero expectation value of $\hat{\boldsymbol{d}}$ which reaches its maximum magnitude of 1.77 Debye when the microwaves are resonant at $\Delta = 0$ as shown in Figure 2.1(b). Around this maximum, the dependence of $\langle \hat{\boldsymbol{d}} \rangle$ with respect to Δ is second-order.

As mentioned, uncertainties in ω_{mol} are an important experimental consideration which translate into errors in Δ. A further possible experimental error is that the implemented magnitude of the microwave intensity may be systematically scaled from its nominal value. When discussing the importance of this error for the entangling protocols, we refer to this scaling factor as f_Ω, where $f_\Omega = 1$ corresponds to implementing the desired microwave intensity.

2.1.3 Two Interacting Molecules

The motivation to create the dipole moment is the dipole-dipole interaction, which becomes important when we have two molecules separated by distances of $\approx 1\mu$m. The far-field electric dipole-dipole interaction between molecules A and B is given by

$$H_{DDI} = \frac{1}{4\pi\epsilon_0 r^3}(\hat{\boldsymbol{d}}_A.\hat{\boldsymbol{d}}_B - 3(\hat{\boldsymbol{d}}_A.\boldsymbol{r}) \otimes (\hat{\boldsymbol{d}}_B.\boldsymbol{r})), \tag{2.5}$$

where ϵ_0 is the permittivity of free space and $\boldsymbol{r}$ is the separation vector between molecules A and B whose magnitude is r. We assume that the two dipoles considered are equally polarised by the microwave field ($\boldsymbol{d}_A = \boldsymbol{d}_B = \boldsymbol{d}$) because their separation is much smaller than the microwave wavelength. Under this condition, the dipole-dipole interaction simplifies to

$$H_{DDI} = \frac{1}{4\pi\epsilon_0 r^3}(1 - 3\cos^2(\theta))\hat{d} \otimes \hat{d}, \tag{2.6}$$

where θ is the angle between the dipole moment of both molecules and their separation and $\hat{d}$ is the dipole operator in the qubit space corresponding to the direction of polarisation. For brevity, we set $V(r) = 2\frac{|\boldsymbol{d}|^2}{4\pi\epsilon_0 r^3}$, corresponding to the strongest interaction at $\theta = 0$, which is $-V\sigma_x \otimes \sigma_x$. We use a representative value of $V = h \times 1847\mathrm{Hz}$, the value for two point dipoles with $|\boldsymbol{d}_{10}| = 1.77$ D at a separation of 0.8μm, although the exact value is dependent on choice of molecule and qubit states. We note that $\theta = 0$ is not possible geometrically if the dipoles are rotating (for example due to circularly-polarised microwaves) or multiple dipoles are arranged in a 2D array. For these cases, the most useful arrangement would orient the dipoles perpendicularly to each other, leading to an interaction of $\frac{V}{2}\sigma_x \otimes \sigma_x$.

The rotating frame we use for the case of two molecules is $U_{frame} \otimes U_{frame}$. In this frame, the Hamiltonian is

$$H_{2mol,rotating} = \begin{pmatrix} -\hbar\Delta & \hbar\frac{\Omega}{2}(1+e^{-i\omega_{mw}t}) & \hbar\frac{\Omega}{2}(1+e^{-i\omega_{mw}t}) & -Ve^{-2i\omega_{mw}t} \\ \hbar\frac{\Omega}{2}(1+e^{i\omega_{mw}t}) & 0 & -V & \hbar\frac{\Omega}{2}(1+e^{-i\omega_{mw}t}) \\ \hbar\frac{\Omega}{2}(1+e^{i\omega_{mw}t}) & -V & 0 & \hbar\frac{\Omega}{2}(1+e^{-i\omega_{mw}t}) \\ -Ve^{2i\omega_{mw}t} & \hbar\frac{\Omega}{2}(1+e^{i\omega_{mw}t}) & \hbar\frac{\Omega}{2}(1+e^{i\omega_{mw}t}) & \hbar\Delta \end{pmatrix} \tag{2.7}$$

in the $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ basis.

Using the rotating-wave approximation for the physically-relevant case of ω_{mol}, $\omega_{mw} >> V$ removes the ‘double flip’ terms of the dipole-dipole interaction ($|00\rangle\langle 11| + |11\rangle\langle 00|$) which are far off-resonant in the static frame and oscillate quickly in the rotating frame, leaving only the ‘flip-flop’ terms ($|01\rangle\langle 10| + |10\rangle\langle 01|$) which transfer a ‘rotational excitation’ from one molecule to the other. In the static frame where they are not time-dependent, the double-flip terms are off-resonant by $\omega_{mol} >> V$, which means the error in energy eigenvalues associated with neglecting them is approximately $\frac{V^2}{\omega_{mol}} \approx 0.5\mathrm{Hz}$, which is much smaller than the existing energy scales of the system.

Overall, the two molecule Hamiltonian in the $\{|00\rangle,|01\rangle,|10\rangle,|11\rangle\}$ basis is

$$H_{2mol} = \begin{pmatrix} -\hbar\Delta & \hbar\frac{\Omega}{2} & \hbar\frac{\Omega}{2} & 0 \\ \hbar\frac{\Omega}{2} & 0 & -V & \hbar\frac{\Omega}{2} \\ \hbar\frac{\Omega}{2} & -V & 0 & \hbar\frac{\Omega}{2} \\ 0 & \hbar\frac{\Omega}{2} & \hbar\frac{\Omega}{2} & \hbar\Delta \end{pmatrix} \tag{2.8}$$

Due to the particle swap symmetry of this Hamiltonian, the Hilbert space is split into a three-dimensional symmetric subspace and an antisymmetric subspace containing only the Bell state $|\Psi_-\rangle = \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle)$. As long as this symmetry is maintained, there can be no transitions between the symmetric and antisymmetric subspaces.

2.1.4 Tweezer Trapped Molecules

In order to investigate the effect of experimental imperfections on the two-molecule Hamiltonian, we consider the case of two molecules trapped in separate optical tweezers. Optical tweezers are focussed beams of laser light which trap particles by using the polarisability of the particle to create a potential minimum. This light is chosen to be off-resonant with the internal transitions of the system to avoid scattering.

In the case of diatomic molecules, the polarisability which allows the potential minimum to be created is higher along the internuclear axis than perpendicular to it [97]. This means the polarisability takes the form

$$\alpha(\theta) = \alpha^{(0)} + \frac{1}{2}\alpha^{(2)}(3\cos^2(\theta) - 1), \tag{2.9}$$

where θ is the angle between the polarisation of the light and the internuclear axis and $\alpha^{(0)}$ and $\alpha^{(2)}$ give the isotropic and anisotropic components of the polarisability respectively. The presence of the anisotropic polarisability $\alpha^{(2)}$ means that different states in excited rotational manifolds have different polarisabilities and therefore experience a different trap strength due to the tweezer light [98]. The polarisability of a given state in the excited rotational manifold depends on the magnetic field, tweezer intensity, and tweezer light polarisation, which can be tuned to reduce the differential polarisability between the two qubit states to much smaller values than their ‘common’ polarisability.

For low temperatures where the molecules are confined close to the trap minima, we model the two optical tweezers as harmonic traps of frequency f_{trap} whose minima are separated by a distance R_e. We consider $R_e = 0.8\mu$m so the traps are sufficiently separated that they do not overlap and lead to collisions between the molecules while also being close enough to enable a large dipole-dipole interaction.

The existence of a finite trapping potential requires the consideration of excited motional states. Comparing the excitation energy hf_{trap} for a large representative value of trap frequency $f_{trap} = 200$kHz with the thermal energy k_BT of CaF molecules cooled to 5.6μK [99], we see that the thermal energy is roughly 60% of the trap excitation energy, which means that the probability of a molecule not being in the ground motional state is significant. For the purpose of this analysis, we assume that there is at most one motional excitation in total for the two molecules. In other words, if $|g\rangle$ and $|e\rangle$ refer to the ground and first excited states of the trap respectively, we limit our consideration to the three two-molecule motional states $\{|gg\rangle, |ge\rangle, |eg\rangle\}$. For now, we also limit our consideration of the extended wavefunction of the positional states to the spatial dimension which separates the tweezers.

The presence of the harmonic traps poses added complications compared to the point dipoles assumed so far. Firstly, the spread of the wavefunctions in a given two-molecule positional eigenstate of the traps increases the expectation value of the DDI compared to point dipoles separated by R_e. Secondly, the DDI creates off-diagonal coupling between the motional eigenstates, again because the positional eigenstates of the trap are extended in space while the positional eigenstates of the DDI are localised in terms of spatial separation.

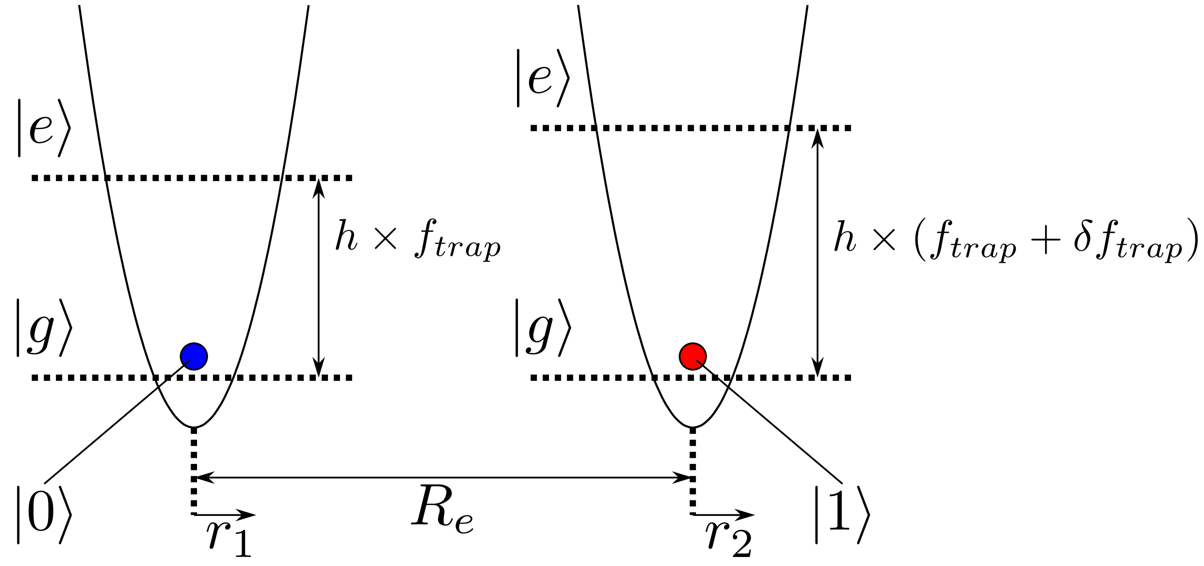

Figure 2.2: Diagram for two molecules in the $|01\rangle$ internal state confined in harmonic traps (distances and energies are not to scale). $|g\rangle$ and $|e\rangle$ denote the ground and first-excited motional states. As defined in the text, $r_{1(2)}$ denotes the position of molecule 1(2) from the centre of its trap while R_e is the distance between traps. f_{trap} is the trap frequency of internal state $|0\rangle$, while the trap frequency of internal state $|1\rangle$ is $f_{trap} + \delta f_{trap}$.

Finally, the differential polarisability between the internal states $|0\rangle$ and $|1\rangle$ causes a difference in trapping frequency i.e. if a molecule in internal state $|0\rangle$ experiences a trapping frequency of f_{trap}, a molecule in internal state $|1\rangle$ experiences a trapping frequency of $f_{trap} + \delta f_{trap}$, where δf_{trap} increases proportionally to the differential polarisability and the trap frequency. This effectively applies an extra error in detuning Δ when the molecule is in motional state $|e\rangle$ rather than $|g\rangle$, which is an additional source of decoherence. For motional states $|ge\rangle$ and $|eg\rangle$, this effect also removes the particle swap symmetry of the internal state Hamiltonian (i.e. $|01\rangle$ is not degenerate with $|10\rangle$) which means the antisymmetric Bell state $|\psi_{-}\rangle$ is no longer decoupled from the symmetric subspace.

2.1.5 Tweezer Trapped Hamiltonian

In this section, we derive how the excited motional states affect the Hamiltonian through these three qualitative mechanisms. To visualise the geometry of the problem, Figure 2.2 shows a schematic of two molecules in the $|01\rangle$ internal state.

We use representative values to calculate the Hamiltonian matrix for the three motional states considered $\{|gg\rangle,|ge\rangle,|eg\rangle\}$ and the four two-qubit internal states $\{|00\rangle,|01\rangle,|10\rangle,|11\rangle\}$, creating a 12×12 matrix. For the effect of the internal state on trapping frequency, this is done by adding δf_{trap} to the relevant elements. To calculate the matrix elements involving the dipole-dipole interaction, we used numerical integration in position space. In appendix A, we discuss an alternative method using ladder operators which improves the physical intuition under extra assumptions. The calculation using position space was the one used to generate the Hamiltonians used for the entangling protocols.

2.1.5.1 Position Basis Derivation

In order to derive how the extended positional wavefunctions affect the dipole-dipole interaction, we express the dipole-dipole interaction in the position basis

$$H_{DDI} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} d_{r_1} d_{r_2} \frac{-2|\boldsymbol{d}|^2}{4\pi\epsilon_0|r_2 - r_1 + R_e|^3} |r_1, 0\rangle |r_2, 1\rangle \langle r_1, 1| \langle r_2, 0| + H.c., \tag{2.10}$$

where $r_{1(2)}$ is the displacement of molecule 1(2) from the minimum of its trap along the separation direction and $H.c.$ means Hermitian conjugate. We use the eigenstates of the quantum harmonic oscillator in the position basis to create wavefunctions for the 12 basis states of the matrix. For this we use the following wavefunctions for the displacement of a molecule from the centre of its trap

$$\langle r|g\rangle = (\pi)^{-\frac{1}{4}} \frac{1}{\sqrt{l_{ho}}} \exp(-\frac{r^2}{2(l_{ho}^2)}) \tag{2.11}$$

$$\langle r|e\rangle = (\pi)^{-\frac{1}{4}}\frac{1}{\sqrt{l_{ho}}}\frac{\sqrt{2}r}{l_{ho}}\exp(-\frac{r^2}{2(l_{ho}^2)}), \tag{2.12}$$

where the harmonic oscillator length $l_{ho} = \sqrt{\frac{\hbar}{2\pi m f_{trap}}}$ for internal state $|0\rangle$ and $l_{ho} = \sqrt{\frac{\hbar}{2\pi m(f_{trap}+\delta f_{trap})}}$ for internal state $|1\rangle$, where m is the mass of a CaF molecule. This lengthscale is around 30nm for CaF molecules if $f_{trap} = 200$kHz, which is much smaller than R_e to prevent tunnelling between the tweezers. We numerically evaluate $\langle\phi_j|H_{DDI}|\phi_{j'}\rangle$ where ϕ_j are the 12 considered states (i.e. all states of the form $\{|gg\rangle,|ge\rangle,|eg\rangle\}\otimes\{|00\rangle,|01\rangle,|10\rangle,|11\rangle\}$ where an individual ϕ_j is made by choosing one positional wavefunction and one internal state from the two lists). Most of these elements are zero due to the form of the dipole-dipole interaction in the internal Hilbert space. We numerically integrate over coordinates r_1 and r_2 between $\pm 3.5 l_{ho}$, which is sufficient to normalise the wavefunctions to 4 significant figures.

2.1.5.2 Effect of Interaction Elements

To describe the outcome of these calculations of the dipole-dipole interaction elements, it is useful to think separately about the matrix elements which leave the motional state unchanged and those which change the motional state of the molecules. In the first category, there is one ‘internal Hamiltonian’ for each of the three motional states. As an example, for the $|ge\rangle$ motional state, the internal Hamiltonian in the $\{|00\rangle,|01\rangle,|10\rangle,|11\rangle\}$ basis is

$$H_{ge} = \begin{pmatrix} -\hbar\Delta & \hbar\frac{\Omega}{2} & \hbar\frac{\Omega}{2} & 0 \\ \hbar\frac{\Omega}{2} & h\delta f_{trap} & -V_{ge} & \hbar\frac{\Omega}{2} \\ \hbar\frac{\Omega}{2} & -V_{ge} & 0 & \hbar\frac{\Omega}{2} \\ 0 & \hbar\frac{\Omega}{2} & \hbar\frac{\Omega}{2} & \hbar\Delta + h\delta f_{trap} \end{pmatrix} \tag{2.13}$$

The different motional states have different effective values of V (i.e. different elements which couple $|01\rangle$ and $|10\rangle$). For two traps of frequency $f_{trap} = 200$kHz separated by $R_e = 0.8\mu$m, these values to the nearest $h \times 1$Hz are $V_{gg} = h \times 1862$Hz and $V_{ge} = h \times 1877$Hz (compared to $V = h \times 1847$Hz for point dipoles separated by R_e).

This interaction is slightly stronger in the more excited motional states because the positional wavefunction is more spread out.

While not used explicitly in the calculations of the entangling protocols, it is important to consider the effect of the second category of dipole-dipole interaction matrix elements which couple different motional states. Due to the form of the dipole-dipole interaction, these always involve the internal state operators $|01\rangle\langle 10| + H.c.$. The term coupling $|gg\rangle$ and the singly-excited motional states $|ge\rangle$ and $|eg\rangle$ has a magnitude of $h \times 145$Hz. While this is a significant energy, it couples states which are separated by the trap excitation energy hf_{trap} and so is far off-resonant. The coupling between the $|eg\rangle$ and $|ge\rangle$ motional states, which effectively swaps a motional excitation and internal excitation between the particles, is smaller ($h \times 15$Hz) but acts between states which can in principle be degenerate and therefore could cause population transfer. A 2% difference between the trap frequencies of the two traps is enough to render this coupling off-resonant as well, so for the rest of the chapter we consider $f_{trap} = 204$kHz for the second molecule [2]. While these 'motional coupling' terms were not included in the design of the entangling protocols for simplicity, they were included in the numerical simulation of the evolution.

2.2 Adiabatic Entangling Protocol

At a more basic level than full quantum logic gates, quantum information applications rely on the generation of entangled states. An entangling protocol in the symmetric 'triplet' subspace in the ground motional state $|gg\rangle$ would evolve a product state (e.g. $|00\rangle$) to an entangled state such as $|\psi_+\rangle = \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle)$. Ideally, such a protocol would be simple and robust to the most likely experimental imperfections. A well-established

---

[2] The difference in tweezer intensity may introduce small violations of the particle-swap symmetry of the internal states. We assume this effect is negligible as it will be smaller than the existing δf_{trap} terms.

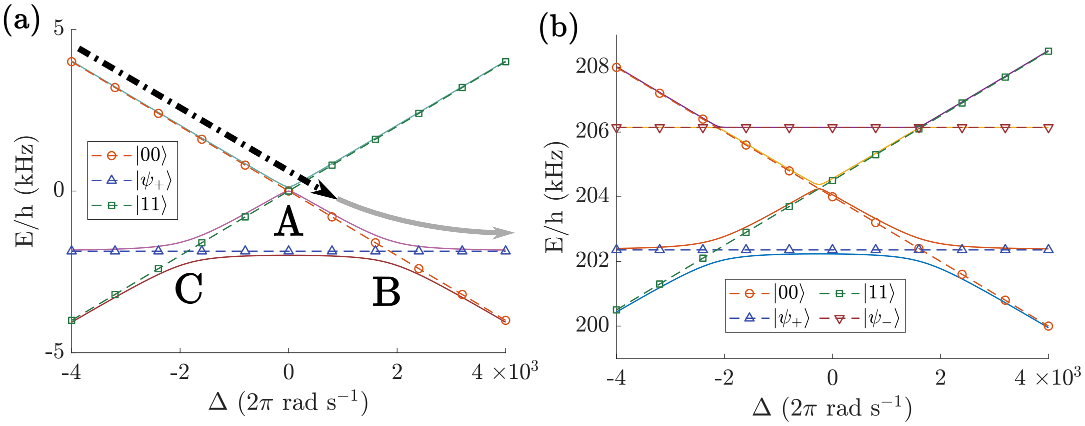

Figure 2.3: (a)Energy levels of $H_{triplet}$ in the ground motional state as a function of detuning Δ with microwave coupling strength $\Omega = 2\pi \times 500\text{rads}^{-1}$. The dashed lines show $\langle \hat{H}_{triplet} \rangle$ for the labelled basis states, while the solid lines show the eigenvalues of the Hamiltonian. The three avoided crossings are labelled A, B, and C. An entangling protocol is shown with the dotted black arrow denoting diabatic evolution and the solid grey marking adiabatic evolution. (b) Energy levels for the motional state $|ge\rangle$ where the antisymmetric state $|\psi_-\rangle$ is no longer decoupled from the symmetric subspace, using the additional trap frequency for the $|1\rangle$ internal state set to $\delta f_{trap} = 500\text{Hz}$.

such protocol is to drive the system adiabatically through an avoided crossing in the two-molecule energy spectrum. Adiabatic protocols change the parameters of the Hamiltonian slowly such that the state of the system remains almost exactly in an eigenstate of the instantaneous Hamiltonian [100]. The avoidance of relying on resonant couplings makes such protocols robust as long as the eigenstates of the initial and final Hamiltonians are not significantly altered by the imperfections in the Hamiltonian.

Within the triplet subspace with the basis states $\{|00\rangle, |\psi_+\rangle, |11\rangle\}$, the two-molecule Hamiltonian is

$$H_{triplet} = \begin{pmatrix} -\hbar\Delta & \hbar\frac{\Omega}{\sqrt{2}} & 0 \\ \hbar\frac{\Omega}{\sqrt{2}} & -V & \hbar\frac{\Omega}{\sqrt{2}} \\ 0 & \hbar\frac{\Omega}{\sqrt{2}} & \hbar\Delta \end{pmatrix} \tag{2.14}$$

The energy levels of this Hamiltonian in the ground motional state are shown in Figure 2.3(a). The dipole-dipole interaction reduces the energy of $|\psi_+\rangle$ by V, which separates a single three-level crossing into three two-level crossings labelled A, B, and

C at $\Delta_A = 0$, $\Delta_B = V/\hbar$, and $\Delta_C = -V/\hbar$ respectively. Crossings B and C, between $|00\rangle$ and $|\psi_+\rangle$ or $|11\rangle$ and $|\psi_+\rangle$ respectively, are avoided due to the microwave coupling Ω between the relevant basis states. Crossing A (between $|00\rangle$ and $|11\rangle$) is avoided due to weaker second-order coupling through $|\psi_+\rangle$ because $|00\rangle$ and $|11\rangle$ are not directly coupled in the Hamiltonian. In the excited motional state $|ge\rangle$, the detuning is effectively shifted by $\frac{\delta f_{trap}}{2\times 2\pi}$ and the antisymmetric state $|\psi_-\rangle$ couples to the symmetric subspace, producing two further second-order avoided crossings with $|00\rangle$ and $|11\rangle$ respectively.

A simple entangling protocol which was previously proposed in Ref [101] is shown using the arrowed path in Figure 2.3. The two molecules are initially prepared in the product state $|00\rangle$ at negative detuning Δ. Δ is increased linearly, passing through the weakly-coupled crossing A and then the directly-coupled crossing B. To ensure that crossing B is traversed adiabatically (with state evolution following the instantaneous eigenstate of the Hamiltonian) the microwave coupling Ω is increased in a Gaussian pulse as Δ approaches the crossing following the time dependence $\Omega(t) = \Omega_0 \exp(\frac{-(t-t_B)^2}{T^2})$, where Ω_0 and T are chosen constants and t_B is the time at which crossing B is traversed. By staying in the instantaneous eigenstate of the Hamiltonian, the entangled state $|\psi_+\rangle$ is prepared as the detuning Δ is increased beyond the crossing region. This protocol relies on the separate addressability of crossings A and B (C) due to the dipole-dipole interaction energy V.

An example of this protocol applied to the two-molecule system is shown in Figure 2.4. This figure shows the population transferring from $|00\rangle$ to $|\psi_+\rangle$ without coherent oscillations. For the ground motional state, the entangled state $|\psi(t)\rangle$ is produced with high probability over a wide range of detuning errors and is largely unaffected by a systematic scaling of the microwave intensity due to the adiabatic nature of the protocol. For the excited motional state $|ge\rangle$, the transfer from $|00\rangle$ to $|\psi_+\rangle$ is shifted earlier in the protocol due to δf_{trap}, while the overall transfer probability is reduced by around 0.002 compared to the ground motional state due to leakage into the antisymmetric state $|\psi_-\rangle$ at crossing B.

While this protocol benefits from the lack of 'fine-tuning' needed, there are weaknesses: The first is that it is slow due to the requirement of adiabaticity at crossing B. Ω_0 must be small compared to V to keep the crossings A and B separate, which limits the rate of change of Δ and therefore increases the minimum entangling time. The adiabatic protocol entangling time is much longer than the timescale $\frac{h}{V}$ for a full phase rotation of $|\psi_+\rangle$ due to the DDI energy, which is around 0.54ms in this example. Secondly the requirement that crossings A and B must be separately addressed means that errors in Δ (and therefore the δf_{trap} energy of excited motional states) are more problematic than for a purely adiabatic protocol because errors in Δ decrease the energy gap which resolves the crossings. Thirdly, for excited motional states, the additional avoided crossings causing population transfer to $|\psi_-\rangle$ are difficult to address separately from the desired adiabatic crossing. Therefore, we investigate more recent techniques to improve the performance of adiabatic protocols in the next subsection.

2.2.1 Shortcut to Adiabaticity

Since the initial proposal of transitionless quantum driving [102], 'shortcuts to adiabaticity' (STA) have sought to reproduce the intuitive nature and robustness of adiabatic evolution in a shorter time [103]. While many approaches exist (review [104]), one particular method for developing shortcuts to adiabaticity which allows significant flexibility to optimise the Hamiltonian controls for simplicity or robustness is to inverse engineer the Hamiltonian based on the evolution of a Lewis-Riesenfeld Invariant (LRI) [105].

The goal is to keep the system in an eigenstate of an 'invariant' $I(t)$ at all times rather than an eigenstate of the Hamiltonian itself and to derive the necessary form of

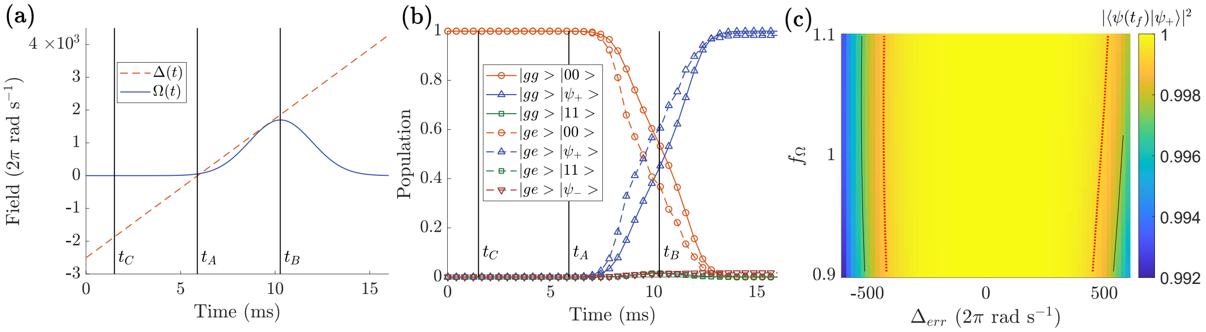

Figure 2.4: Adiabatic entangling protocol. (a) Microwave detuning and coupling intensity during protocol for microwave coupling strength parameter $\Omega_0 = 2\pi \times 1700$rad s^{-1} and Gaussian pulse width $T = 2.4$ms. (b) Evolution of population of basis states during protocol, starting from either $|gg\rangle|00\rangle$ (solid lines) or $|ge\rangle|00\rangle$ (dashed lines). (c) Numerically calculated probability to create state $|gg\rangle|\psi_+\rangle$ from $|gg\rangle|00\rangle$ at end of protocol as a function of constant error in detuning and systematic microwave power scaling. The red dotted contour and black solid contour bound the regions of transfer probability > 0.999 and > 0.997 respectively. The times at which the detuning Δ traverses the three avoided crossings in the ground $|gg\rangle$ motional state are labelled in subfigures (a) and (b).

$H(t)$ from this. This is done by ensuring that the full operator matrix of the invariant $I(t)$ does not evolve in the Heisenberg picture. The invariant is designed to commute with the Hamiltonian at the start and endpoints of the evolution to ensure that the eigenstates coincide at these times so that initial and final states of the system are stationary. By choosing the form of $I(t)$, one can create the correct evolution while exerting a degree of choice regarding the form of the driving fields.

We followed a modified version of the LRI protocol from Ref [106]. The details of this derivation are included in appendix B. This shortcut is only derived for the adiabatic part of the evolution (i.e. crossing B) and does not explicitly involve the state $|11\rangle$. We include the full Hilbert space in our numerical simulation of the evolution however to investigate the conditions under which it creates an entangled two-qubit state.

Given how the STA was derived, its success requires that $|11\rangle$ is not highly populated. Whether this assumption is true depends to a large extent on the microwave coupling $\Omega(t)$ remaining sufficiently small for crossings A and B to be separate two-level crossings.

This limits the speed at which the STA can be performed, as a fast STA will require larger values of $\Omega(t)$. This significantly degrades the transfer probability for protocols quicker than $t_f = 0.8$ms as $|11\rangle$ becomes populated.

Figure 2.5(b) shows the evolution of the populations in the ground motional state during the adiabatic crossing for this example protocol in three different bases: the triplet basis $\{|00\rangle,|\psi_+\rangle,|11\rangle\}$, the eigenbasis of the instantaneous internal Hamiltonian $H(t)$ labelled in ascending order (so that $|H1\rangle$ is the lowest energy state), and the initially-populated eigenstate of the invariant $I(t)$. The population evolution in the triplet basis is similar to the true adiabatic protocol, with a smooth transfer from $|00\rangle$ to $|\psi_+\rangle$ and a small probability to excite $|11\rangle$. Unlike the adiabatic protocol, the STA protocol evolution does not follow the instantaneous eigenstates of $H(t)$, as shown by the fact that the lowest two eigenstates of the $H(t)$ are almost equally occupied in the middle of the protocol. The evolution follows the eigenstates of $I(t)$ closely as expected.

In Figure 2.5(c) we show the transfer probability and Hamiltonian controls for the LRI protocol with $t_f = 1$ms. This demonstrates the ability for the LRI method to significantly shorten the 'adiabatic' crossing. However, the robustness of this protocol is significantly reduced, especially with respect to microwave power scaling (as shown in Figure 2.5(c) for the ground motional state) and it performs poorly for excited motional states because of poor robustness to δf_{trap}. We note that LRI protocols which are optimally-robust to systematic errors have been derived for two-level systems where the full Hilbert space is included in the LRI ansatz [107–109] but we find these protocols cause significant leakage into $|11\rangle$.

This protocol favours a small constant positive offset in $\Delta(t)$, an effect which increases for shorter protocol times t_f. Analysis of the populations of the three basis states during the protocol suggests this is so that the higher values of $\Omega(t)$ are reached when $\Delta(t)$ has a larger value (i.e. is further from the 'diabatic' crossing A). The

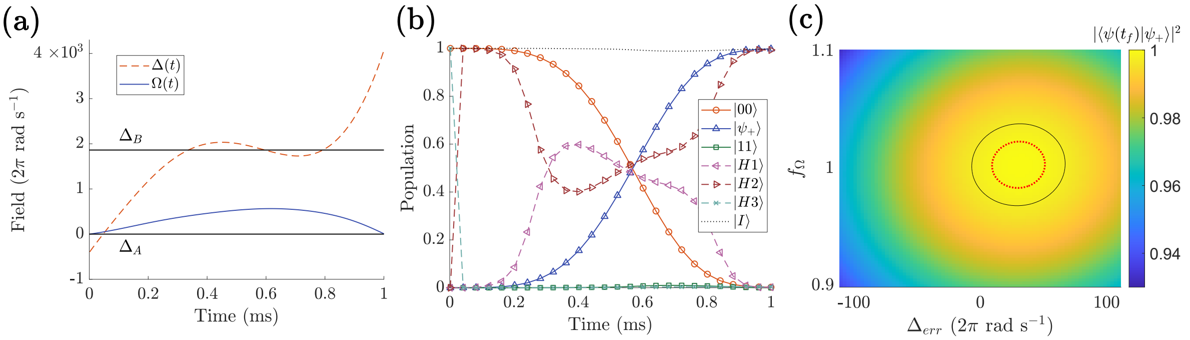

Figure 2.5: LRI entangling protocol. (a) Example control field parameters detuning $\Delta(t)$ and microwave coupling strength $\Omega(t)$. Values of Δ corresponding to crossings A and B are marked with black solid lines. (b) Evolution of population starting from the $|00\rangle$ internal state in the ground motional state. Basis states $|00\rangle$, $|\psi_+\rangle$, and $|11\rangle$ are shown in solid lines, symmetric eigenstates of the internal Hamiltonian for the ground motional state are shown in ascending order of eigenvalue in dashed lines, while the initially-populated eigenstate of the invariant is shown in a dotted line. (c) Final transfer probability $|\langle\psi(t_f)|\psi_+\rangle|^2$ in the ground motional state for systematic offsets in detuning and microwave power scaling. The red dotted contour and black solid contour bound the regions of transfer probability > 0.999 and > 0.997 respectively.

observation that the third state $|11\rangle$ limits the speed of the protocol and favours detunings which are larger than those of the 'ideal' two-state STA was also seen when treating a similar problem with transitionless quantum driving and optimal control protocols [110].

2.3 Entangling Gate

While these adiabatic protocols offer simple state transfer, they perform less well when implementing quantum logic gates, where the relative phases as well as the populations of states are important. For example, rather than demanding only that the protocol transfers a chosen product state to a chosen maximally-entangled state, a full quantum logic gate would specify the entire unitary evolution operator up to a global phase, for example (in the $\{|00\rangle,\ |\psi_+\rangle,\ |11\rangle,\ |\psi_-\rangle\}$ basis)

$$U_{aim} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \exp(i\beta) \end{pmatrix} \tag{2.15}$$

where β is a chosen phase. Importantly when aiming for a gate, errors in Δ would lead to an unexpected phase between the $|0\rangle$ and $|1\rangle$ states being accumulated at a constant rate during the protocol. To solve this problem, we employ optimal control methods to design pulses which drive a more complicated state evolution, allowing accumulated errors to cancel out and producing greater robustness. For this purpose, the duration of the gate must be longer than the minimum time for the DDI to create a maximally-entangled state.

2.3.1 Quantum Optimal Control Methods

Quantum optimal control refers to the process of optimising driving fields in order to achieve the desired evolution (e.g. state-to-state transfer or quantum logic gate). One particular method which has been widely used for its simplicity and versatility since its introduction in the field of NMR quantum computing is GRadient Ascent Pulse Engineering (GRAPE) [111]. The idea of this method is to build the pulse from a large number of short timesteps δ_t during which the control fields are constant, and optimise the values of the fields at each timestep iteratively.

Mathematically, the Hamiltonian is expressed as the sum of an uncontrollable ‘drift’ Hamiltonian H_0 and a number of control fields. In this example, we have a fixed dipole-dipole interaction and parametrise the microwave control according to its intensity along the X and Y quadratures with fixed detuning $\Delta_0 = 0$. The Hamiltonian at the j^{th} timestep in the symmetric subspace in the $\{|00\rangle,|\psi_+\rangle,|11\rangle\}$ basis is therefore

$$H_j = \begin{pmatrix} -\hbar\Delta_0 & 0 & 0 \\ 0 & -V & 0 \\ 0 & 0 & \hbar\Delta_0 \end{pmatrix} + \frac{\Omega_{X,j}}{\sqrt{2}}\hbar \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} + \frac{\Omega_{Y,j}}{\sqrt{2}}\hbar \begin{pmatrix} 0 & -i & 0 \\ i & 0 & -i \\ 0 & i & 0 \end{pmatrix}$$
$$= H_0 + \Omega_{X,j} h_{\Omega_X} + \Omega_{Y,j} h_{\Omega_Y} \quad (2.16)$$

where $\Omega_{X(Y),j}$ refers to the microwave coupling in the X(Y) quadrature at timestep j and $h_{\Omega_{X(Y)}}$ refers to the matrix which appears multiplied by the relevant microwave coupling $\Omega_{X(Y),j}$ in the total Hamiltonian. This is also commonly considered using the magnitude Ω_j an phase ϵ_j of the microwave coupling, where $\Omega_{X,j}$ and $\Omega_{Y,j}$ are the real and imaginary parts of $\Omega_j e^{i\epsilon_j}$ respectively. Control of the X and Y quadratures (or equivalently the phase ϵ) is achieved by varying the oscillation phase of the driving microwaves with respect to the phase which is chosen to represent the X quadrature ($\epsilon = 0$) [3], where Y is one quarter of a full oscillation out of phase with X. Note that a constant rate of change of phase is an alternative parametrisation for a constant effective detuning Δ.

Ignoring errors at present, the total unitary operator implemented by the protocol is

$$U_{GRAPE} = \prod_{j=1}^{N_{steps}} \exp(-iH_j\delta_t/\hbar) = \prod_{j=1}^{N_{steps}} U_j \quad (2.17)$$

meaning the task is to optimise the fidelity $f = \frac{1}{D}^2|\operatorname{Tr}(U_{aim}^\dagger U_{GRAPE})|^2$ with respect to the values of $\Omega_{X,j}$ and $\Omega_{Y,j}$, where D is the number of dimensions in the Hilbert space and N_{steps} is the number of timesteps. In appendix C, we qualitatively outline the approach developed in Ref [111] to achieve this efficiently, with the computational cost of each iteration scaling proportionally to N_{steps} and D^3 [112].

We utilise an extension to GRAPE which targets pulses that are robust to systematic errors (in this case in Δ and f_Ω) by performing the same optimisation procedure over

---

[3] This choice of phase between the $|0\rangle$ and $|1\rangle$ states must be consistent with that used when performing an operation to rotate the state of the qubit to map to the computational basis, in order to measure the σ_x operator for example.

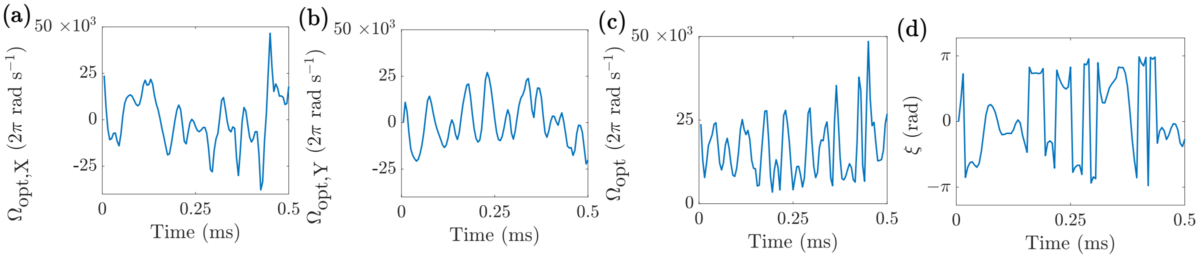

Figure 2.6: Optimised GRAPE microwave pulse displayed using (a) X and (b) Y quadratures or (c) magnitude and (d) phase of the optimised microwave pulse $\Omega_{opt}(t)$.

multiple 'offset' versions of the drift and control Hamiltonians and aiming to optimise the average fidelity amongst these [113]. This produces pulses which have high average fidelity over the chosen systematic errors, which often (but not always) corresponds to high fidelity over the corresponding region of systematic error parameter space rather than at the chosen errors only (this can be checked easily after optimisation). We also use this functionality to optimise for the considered motional states. The computational cost of each iteration scales proportionally to the number of offset Hamiltonians used.

Each iteration of the GRAPE algorithm causes the pulse to approach a local maximum of the fidelity. For pulses with a large number of timesteps and multiple offset Hamiltonians, the optimisation landscape generally has many different local fidelity maxima that can be reached depending on the initial pulse. We repeated the algorithm until an acceptable pulse was found, although more sophisticated approaches to perturb previously-found local maxima can be used [113].

2.3.2 Results

2.3.3 Motion along Separation Direction

The GRAPE pulse in this subsection was produced by Jonathan A. Jones and Gaurav Bhole in the symmetric subspace using Hamiltonians chosen with MH.

The GRAPE pulse of 0.5ms was optimised in 100 steps of 5μs. This pulse is shown

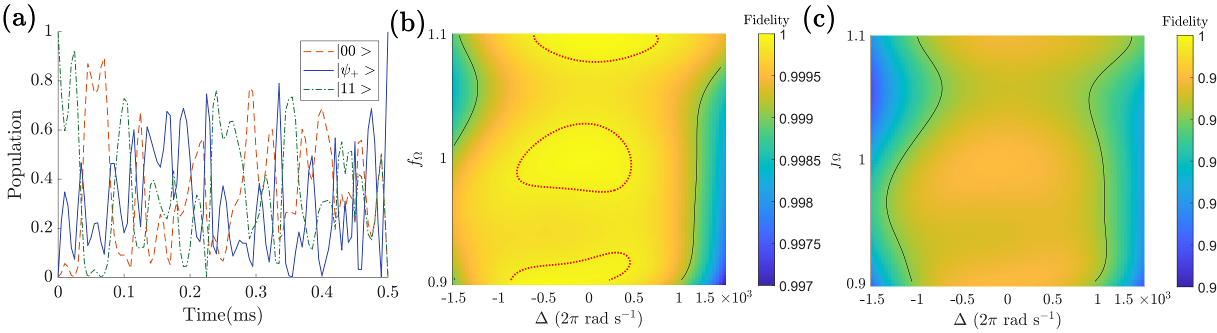

Figure 2.7: (a) Evolution of initial internal state $|11\rangle$ in the ground motional state. Heatmaps show internal state fidelity of the GRAPE optimised pulse subject to errors in detuning Δ and microwave intensity scaling f_Ω for (b) the ground motional state $|gg\rangle$ (c) an equal mixture of all three considered motional states. The red dotted contour bounds the region of fidelity > 0.9999 while the black solid contour bounds the region of fidelity > 0.999. The parameters used for these simulations were trap frequency $f_{trap} = 200$kHz for one molecule and 204kHz for the other with the additional trap frequency of the $|1\rangle$ internal state $\delta f_{trap} = 500$Hz for both molecules (the performance of this pulse does not require fine-tuning these trap parameters, as shown in appendix D).

in Figure 2.6. Compared to the intuitive adiabatic and STA protocols, it uses much more intense and sharply varying microwave fields, as shown by the peak microwave coupling of just under 50kHz in subfigure (c). This pulse drives the system along a complicated path through the internal Hilbert space, which allows the specified systematic errors to cancel out. This is shown in Figure 2.7(a), where the evolution of the initial internal state $|11\rangle$ for the ground motional state is plotted, which involves many more population transfers to the other symmetric internal states than for either the adiabatic or STA protocols. In the ground motional state, the particle-swap symmetry of the internal Hamiltonian is preserved so population within the symmetric subspace cannot be transferred to $|\psi_-\rangle$.

This pulse is designed to implement U_{aim} as specified in equation 2.15 in the symmetric subspace, with the singlet phase β calculated from the relative phase of the decoupled singlet and triplet manifolds [4] and used to evaluate the fidelity. $|\psi_-\rangle$

---

[4] In general, the relative phase of the singlet can be important in determining whether a gate is entangling, although in this case the gate is entangling regardless of the value of β because it transfers

is only coupled to the rest of the internal Hilbert space through imperfections which break the particle swap symmetry, so this phase is largely set by the gate time and dipole-dipole interaction strength rather than the pulse.

We initially assume that the molecules remain in the ground motional state $|gg\rangle$ to numerically evolve the pulse using the internal Hamiltonian H_{gg} and calculate the gate fidelity. These results are shown in Figure 2.7(b). The fidelity is > 0.999 for errors in $\Delta < 1$kHz and for up to 10% errors in the microwave intensity, which means this region is above the threshold for quantum error correction [114]. The three regions of fidelity > 0.9999 are centred around $f_\Omega = 0.9, 1, 1.1$ because these were the values of f_Ω used in the GRAPE optimisation to target robustness to this type of error.

We then numerically evolve the same pulse using the full 12×12 Hamiltonian to account for motional excitations, finding the unitary matrix U_{full}. The aim is that the pulse implements the same gate in the internal Hilbert space regardless of the motional state. This condition can be met if the unitary operation in the combined Hilbert space of the internal and motional states obeys

$$U_{full,aim} = U_{aim} \otimes U_{motion,phase}, \tag{2.18}$$

where $U_{motion,phase}$ is a unitary gate in the motional Hilbert space which delivers a phase to the three motional states $|gg\rangle$, $|ge\rangle$, and $|eg\rangle$ without transferring population between them. In this case, any initial motional state which is an incoherent superposition of the three motional states would experience the gate U_{aim} in the internal Hilbert space. Using the calculated unitary operator U_{full}, we extract these motional phases by taking the phase difference between the elements of U_{full} which correspond to the same internal state transition. We use these to construct $U_{motion,phase}$ and by extension the overall target gate $U_{full,aim}$. We then calculate the fidelity using $\frac{1}{12^2}|U^\dagger_{full,aim}U_{full}|^2$.

---

the product state $|11\rangle$ to the entangled state $|\psi_+\rangle$.

We use this particular process (which gives equal importance to the three motional states) in order to justify the claim that the pulse implements the same internal gate on the three motional states to a good fidelity. However, for sufficiently cooled molecules, the population in the ground motional state will be higher than in the singly-excited motional states.

The fidelity when including the three considered motional states is shown in Figure 2.7(c). The fidelity is reduced compared to that of the ground motional state (Figure 2.7(b)) but still maintains a comparable region of parameters in which the fidelity > 0.999. A significant fraction of the reduction in fidelity is due to the phase acquired by the singlet state due to the dipole-dipole interaction energy which varies between motional states. This phase is harder to correct during the swap-symmetric gate protocol because population is only transferred into or from $|\psi_-\rangle$ due to δf_{trap} rather than a control field.

2.3.4 Motion in Three Dimensions

The additional GRAPE pulse in this subsection was produced by MH using the optimal control module of the Spinach library [115, 116] by optimising in the two-qubit space [5].

While the largest effects of the finite particle wavefunctions are along the direction of tweezer separation (x direction), it is important to consider the effects of the positional wavefunctions in the perpendicular directions (y and z). Although in practice, thermal population leads to more motional excitations along the weaker trap direction, we consider one excitation in total between both molecules and

---

[5] We found that population transfer between the symmetric and antisymmetric subspaces could significantly reduce fidelity of pulses which would otherwise implement the desired unitary in the symmetric subspace to high fidelity and therefore included $|\psi_-\rangle$ in our optimisation. This population transfer is reduced by ensuring Ω_j is large enough to reduce the impact of the δf_{trap} terms which break the particle swap symmetry in excited motional states. It was useful to start the optimisation with large Ω_j to ensure that the singlet was largely decoupled throughout the optimisation, because otherwise the algorithm would become drawn to extremely large control fields.

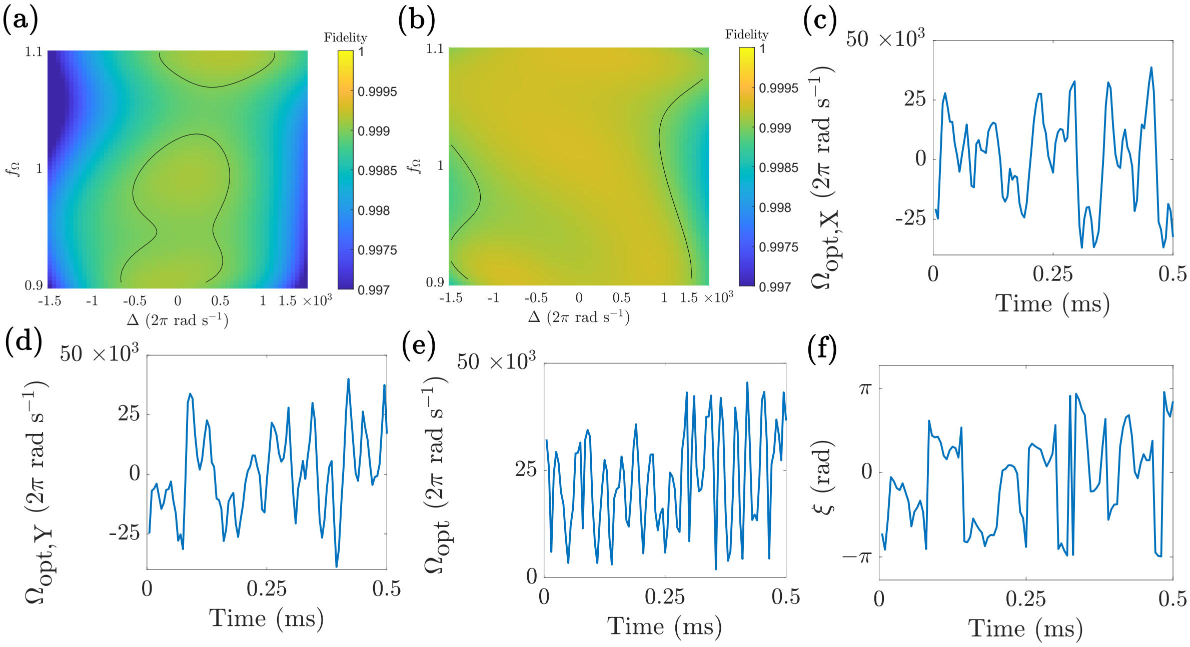

Figure 2.8: (a) Fidelity of GRAPE pulse shown in Figure 2.6 when evolved considering motional states in three dimensions. (b) Fidelity of pulse optimised for motional states in three dimensions (X and Y quadratures of this pulse shown in (c) and (d), while magnitude and phase are shown in (e) and (f)). The evolution of these pulses was simulated using trap frequency $f_{trap} = 200$kHz for the first molecule and 204kHz for the second molecule in the x and y directions, and $f_{trap} = 100$kHz for the first molecule and 102kHz for the second molecule in the z direction. For both molecules, we used additional trap frequency for the $|1\rangle$ internal state $\delta f_{trap} = 500$Hz in the x and y directions and 250Hz in the z direction.

all three directions, which means seven two-molecule motional states denoted by $\{|gg\rangle, |ge_x\rangle, |e_xg\rangle, |ge_y\rangle, |e_yg\rangle, |ge_z\rangle, |e_zg\rangle\}$.

The finite wavefunction spread and excited motional states in the perpendicular directions reduce the overall DDI strength because of increased separation and the anisotropy of the DDI. Numerically, for the first tweezer, we consider equal trap frequencies of $f_{trap} = 200$kHz in the x and y direction with a weaker trap frequency of $f_{trap} = 100$kHz in the z direction which we take as the optical tweezer axis. We again assume the second tweezer to have larger trap frequencies by 2%. These radial and axial trap frequencies are demanding, but could be reduced for cooler samples. The ‘internal’ DDI elements (which do not couple different motional states) are $V_{gg} = h \times 1840$Hz,

$V_{ge_x} = h \times 1854\text{Hz}$, $V_{ge_y} = h \times 1832\text{Hz}$, and $V_{ge_z} = h \times 1825\text{Hz}$ while the elements transferring motional excitations in the x, y and z directions are 14Hz, 7Hz, and 14Hz respectively which as before can be rendered sufficiently off-resonant by the difference in trapping frequencies between the tweezers. The DDI elements which transfer motional excitations between different directions are extremely weak ($< 0.1\text{Hz}$).

Overall, we find that the pulse optimised for motion in only the x direction maintains only a small region of fidelity > 0.999 when evolved including the extra four considered motional states of the perpendicular directions as shown in Figure 2.8(a). This region can be enlarged by specifically optimising a new pulse for the average fidelity of the seven considered motional states, as shown in Figure 2.8(b). The weaker trap frequency in the tweezer axis causes the most significant added reductions in fidelity from considering three dimensions.

2.3.5 Discussion

2.3.5.1 Error Tolerance

Having established the robustness of the gate protocol to certain systematic errors, it is necessary to check that the most significant error mechanisms fall within the high-fidelity region. Robustness to errors in Δ is important due to fluctuations in ω_{mol}, which can be driven by magnetic field and tweezer intensity fluctuations. As mentioned earlier, the susceptibility to magnetic field fluctuations can be reduced by the choice of qubit states. For example, for CaF, choosing $|0\rangle = |N = 0, F = 1, m_F = 1\rangle$ and $|1\rangle = |N = 1, F = 2, m_F = 2\rangle$ reduces the differential magnetic moment, meaning that fluctuations of 1mG cause errors in Δ of around 100Hz.

Fluctuations in tweezer intensity cause errors in detuning because of the differential AC Stark shift between the qubit states. We denote the ‘local’ polarisability of state $|0\rangle(|1\rangle)$ by $\alpha_{|0\rangle(|1\rangle)} = \frac{\partial E_{trap,|0\rangle(|1\rangle)}}{\partial I_{trap}}$ where E_{trap} is the trap energy and I_{trap} is the tweezer intensity, and accordingly define the differential polarisability $\delta\alpha = \alpha_{|1\rangle} - \alpha_{|0\rangle}$ and average polarisability $\bar{\alpha} = \frac{\alpha_{|0\rangle}+\alpha_{|1\rangle}}{2}$. By careful choice of magnetic fields and tweezer light polarisations, the local differential polarisability can be reduced to below 1% of the average polarisability [10, 96]. To estimate the error due to tweezer fluctuations, we compare the differential AC Stark shift δE_{trap} due to tweezer intensity fluctuations δI_{trap} with the intended AC Stark shift trapping potential depth $\bar{E}_{trap} \approx \bar{\alpha} I_{trap}$ [6]. The fluctuation in the differential AC Stark shift between the qubit states is given by $\delta E_{trap} = \delta\alpha\delta I_{trap}$. This means that $\delta E_{trap} \approx (\delta\alpha/\bar{\alpha})(\delta I_{trap}/I_{trap})\bar{E}_{trap}$. Given representative values $E_{trap} = h \times 1\text{MHz}$ and $\delta I_{trap}/I_{trap} = 10^{-3}$, δE_{trap} can be reduced below $\approx 100\text{Hz}$. This uncertainty corresponds to an error in detuning which is within the boundaries of the region with fidelity > 0.999.

2.3.5.2 Assumptions

Our theoretical treatment assumes that the internal qubit Hilbert space is closed and that no transitions to other rotational and hyperfine states are excited by the microwaves. We estimate the probability for such an off-resonant excitation as $(\Omega_{off}/\Delta_{off})^2$, where Δ_{off} is the detuning to the off-resonant level and Ω_{off} is the microwave coupling between the two states involved in that transition. For many choices of qubit states such as our example qubit states at 50G, we see from Figure 2.1(a) that the nearest level outside the qubit space has $\Delta_{off} \approx 10\text{MHz}$, while $\Omega < 55\text{kHz}$ during the pulse if $f_\Omega < 1.1$. As the example qubit states, and most good choices of qubit states in general, have a larger transition dipole moment to each other than to any other states within the $N = 0$ and $N = 1$ rotational manifolds, $\Omega_{off} \leq \Omega$ for a given microwave intensity. This gives an off-resonant excitation probability below 4×10^{-5} which is much smaller than the infidelities already present in the gate.

---

[6] This order-of-magnitude estimate uses the approximation $\bar{\alpha} = \frac{\bar{E}_{trap}}{I_{trap}}$ rather than $\bar{\alpha} = \frac{\partial \bar{E}_{trap}}{\partial I_{trap}}$ [97].

The use of the rotating-wave approximation to remove counter-rotating terms also neglects the Bloch-Siegert energy shift to the resonant frequency. This energy shift has a magnitude $\approx \Omega^2/\omega_{mol}$ [100] which is < 0.2Hz for this protocol and therefore well within the region of high robustness.

We have also ignored the momentum kick from the photons of the driving field, which can populate excited motional states in other platforms such as Rydberg atoms under optical frequency driving [117], because the photon momentum of the microwave driving is smaller by a factor of around 10^5.

A further implicit assumption in our calculations is the ability to implement the shaped pulses made from 100 individual 5μs rectangular pulses (allowing for the considered errors in Δ and f_Ω). Microwave pulses of this duration are currently used to manipulate the quantum state of polar molecules [118]. Smoother variations between time steps can be targeted by penalty terms within GRAPE if required, although this may come at the expense of pulse performance in other respects.

We considered zero or one motional excitation between the two molecules. While this could include a large fraction of the population of the coldest near-future experiments, small populations in higher motional states could lead to significant infidelities. Higher motional states could be included in the optimisation (which could be weighted by their estimated thermal population). This would increase the demands on the pulse and the numerical optimisation cost due to the greater number of iterations and algorithm repeats needed for this more difficult optimisation problem as well as the increased cost of each iteration, although the latter aspect of the computational cost may be mitigated by increasing the number of considered motional states during optimisation [113].

2.3.5.3 Context

As mentioned earlier, there are multiple examples of quantum gates proposed for polar molecules. Our scheme differs from that of Refs [16, 19] in that we do not demand a third molecular level to be resonantly coupled and from the former in that we do not demand a static electric field. Our two-level proposal requires significantly more complicated microwave pulse shapes however. A scheme published since ours [20] considers thermal population of a greater number of motional states, which causes greater decoherence due to the anisotropic polarisability and therefore places greater emphasis on choosing tweezer light intensity and polarisation to minimise this effect.

On the other hand, Refs [119] and [120] used optimal control methods to design shaped pulses for gates in ultracold molecules. Unlike these previous works, we used these methods additionally to provide robustness to experimental imperfections, which gives our scheme greater applicability to near-term experiments. Optimised pulses have been used in a similar way for other platforms [121–124].

2.3.5.4 Computation

A necessary criterion for useful quantum computation is that the coherence time of the logical states must be significantly longer than the gate time [4]. The time needed for this gate is on the scale of $\approx 1 - 10$ milliseconds depending on the DDI strength available from the choice of molecule, tweezer separation, and qubit states. Experimentally-recorded coherence times of polar molecules are currently improving, with recent results including a coherence time of 6.4ms between different rotational states of CaF in magnetic traps [95] and 93ms in optical tweezers [96]. Coherence times of several seconds have been demonstrated between different hyperfine states of the same rotational level for other molecules, which is useful for information storage although a significant electric dipole moment cannot be created within these states [13, 14]. It is therefore realistic that near-future experiments will allow coherence and storage times multiple times greater than the gate timescale.

A further criterion for useful quantum computation is qubit addressability. While two-qubit gates are in principle sufficient for universal quantum computation [125], this relies on being able to perform single-qubit gates on a chosen qubit and two-qubit gates on chosen nearby pairs of qubits rather than necessarily all of the qubits at once. This can be done by physically moving the relevant qubits in an array using optical tweezers [19, 126] although this risks causing additional motional excitations. A similar architecture has recently been realised using Rydberg atoms [127], although polar molecules would require reduced tweezer separation to implement gates.

A different method used in optical lattices is to apply tightly focussed laser beams to alter the target transition frequency ω_{mol} of a chosen qubit such that it alone is in resonance with the microwaves. This has been done using a single beam in a 2D optical lattice [25] and in a 3D optical lattice by using two perpendicular beams [128, 129], although the lattice spacings in these cases were $\approx 5\mu$m, much larger than required for strong interactions between polar molecules. For schemes involving optical lattices, it is worth noting that ultracold molecules have not yet achieved close to unit filling, with notable results including the creation of cubic optical lattices with filling fractions of 25% for fermionic ^{40}K^{87}Rb [76] and 30% for bosonic ^{87}Rb^{133}Cs [77]. This significantly reduces qubit connectivity.

2.4 Conclusions and Outlook

We have simulated protocols to entangle two polar molecules using the dipole-dipole interaction. We found that a shortcut to adiabaticity offers a significant acceleration for Bell state preparation compared with an adiabatic protocol, although the STA protocol's robustness, particularly to motional excitation, is relatively poor. While offering simple and intuitive state control, such protocols struggle to control relative phases when affected by experimental uncertainties as necessary for a full quantum logic gate.

To target this, we utilised optimal control methods to design a robust two-qubit entangling gate. We investigated the effects of thermal excitation and positional wavefunction spread on the interaction Hamiltonian and optimised a gate which may offer fidelity > 0.999 across the ground and singly-excited trap states and over a wide range of errors in microwave amplitude and detuning. This suggests ultracold molecules would be able to perform quantum information processing in the presence of low levels of motional excitation under optimised microwave pulses, which is relevant to near-future experiments.

Future work could investigate whether this scheme allows comparable results in the recently proposed state-dependent tweezers [130, 131]. These trap two molecules at different but much closer positions in a single beam, which in principle can increase the dipole-dipole interaction by a factor of 100. This increased interaction strength would allow shorter gate times and therefore reduce the impact of fluctuations and decoherence, but the reduction of the equilibrium separation may increase the effect of wavefunction spread due to motional excitation. This could mean that the motional state has a greater effect on the internal Hamiltonian or that the dipole-dipole interaction itself transfers or creates motional excitations more easily, which may require the consideration of more motional states (even ignoring thermal population) and make it more challenging for a pulse to target the same unitary gate for multiple motional states. A further mechanism to accelerate entangling gates between polar molecules that may be of interest for future work is to mediate the interaction using Rydberg atoms [132, 133], which allows the use of a simpler blockade entangling scheme instead of dipolar exchange.

3

Dipolar Bose-Hubbard Model

This chapter introduces the dipolar Bose-Hubbard model, which underlies the research on the cylindrical dipolar Bose-Hubbard model in chapter 4 and comparison of lattice and continuum methods in chapter 5. We summarise the construction of the short-range interacting Bose-Hubbard model from ultracold atoms in optical lattices and how the dipole-dipole interaction extends this model.

3.1 Bose-Hubbard Model

While the closely related (Fermi-)Hubbard model was originally motivated by the concrete problem of electrons in condensed matter systems [134], the Bose-Hubbard (BH) model has drawn interest due to its simplicity in describing the Mott-insulator to superfluid transition [135] and ability to be simulated by controlled systems, such as alkali atoms in optical lattices [136].

Ignoring interactions between bosons at first, the Hamiltonian for bosons without internal degrees of freedom moving in an external potential $A(\vec{r})$ is the sum of kinetic and potential energy expressed mathematically as

$$H = \int d\vec{r}\ \hat{\Psi}^\dagger(\vec{r})(-\frac{\hbar^2}{2m}\nabla^2 + A(\vec{r}))\hat{\Psi}(\vec{r}), \tag{3.1}$$

where $\hat{\Psi}(\vec{r})$ annihilates a boson at position $\vec{r}$ while obeying bosonic commutation relations and m is the mass of a boson. Taking a square optical lattice as a simple example, the external potential $A(x, y, z)$ due to the AC stark shift from standing waves of counter-propagating laser beams is of the form

$$A(x, y, z) = A_0(\sin^2(\frac{2\pi x}{\lambda}) + \sin^2(\frac{2\pi y}{\lambda})) + A_z(z), \tag{3.2}$$

where A_0 is the scalar magnitude of the potential proportional to the laser intensity, λ is the laser wavelength, and $A_z(z)$ is a confining potential in the z direction with a minimum at $z = 0$. Provided $A_z(z)$ is strong enough to restrict the physics to the $x - y$ plane, it is not relevant to the derivation of the BH model itself (although it can affect the parameter values within the model). To compare the localising energy of the potential to the kinetic energy, A_0 is usually expressed in units of the recoil energy $E_R = \frac{h^2}{2m\lambda^2}$. This potential has minima at $\vec{r}_{i_x,i_y} = (i_x\vec{x} + i_y\vec{y})\frac{\lambda}{2}$, where i_x and i_y are integer indices for the lattice minima.

According to Bloch's theorem, the single-particle eigenstates of this Hamiltonian with a periodic potential in the $x - y$ plane (such as our example $A(x, y, z)$) are 'Bloch functions' of the form $\exp(i\vec{k}.\vec{r})f(\vec{r})$, where $\vec{k}$ is a wavevector in the $x - y$ plane and $f(\vec{r})$ is a function with the same periodicity as the external potential [137]. However, it is inconvenient to describe the physics of many bosons using Bloch functions because these functions extend over all space whereas the interactions between bosons usually do not. For this reason, the state of the system is usually written in the basis of Wannier functions $w_{i_x,i_y}(\vec{r})$, which are localised at their own potential minimum $\vec{r}_{i_x,i_y}$ and are related to each other through translation by integer multiples of the lattice vectors [138]. They are constructed using superpositions of the Bloch functions in the lowest Bloch band [139]. The use of the lowest Bloch band is justified when the lattice potential depth A_0 is large compared to energy scales which might cause excitations to higher bands, such as temperature or particle interactions.

There exists a choice of phases when building the Wannier functions from the Bloch functions which leads to different Wannier functions and lattice models for the equivalent physical potential. This choice can be resolved by using maximally-localised Wannier functions, which are expected to reduce (although not necessarily minimise) the overlaps and interactions between Wannier functions localised at well-separated sites. Information about the processes to compute such maximally-localised Wannier functions can be found in Refs [140, 141].

The Wannier functions provide a useful basis to decompose the boson field operator $\hat{\Psi}(\vec{r})$ using the equation

$$\hat{\Psi}(\vec{r}) = \sum_{i_x,i_y} w_{i_x,i_y}(\vec{r})\hat{b}_{i_x,i_y}, \tag{3.3}$$

where $\hat{b}_{i_x,i_y}$ now annihilates a boson which has a spatial wavefunction given by $w_{i_x,i_y}(\vec{r})$ and obeys bosonic commutation relations. This allows us to express the state of the system and the Hamiltonian in the discrete basis of the Wannier functions at each lattice site. When expanded in the basis of Wannier states, the single-particle continuum Hamiltonian in equation 3.1 creates ‘tunnelling’ terms which cause particles to hop sites. For deep optical lattices, the non-negligible tunnelling terms are those between nearest-neighbour sites. For tunnelling in the $\vec{x}$ direction for example, this element J is calculated using the integral

$$J = -\int d\vec{r} w^*_{i_x\pm1,i_y}(\vec{r})(\frac{-\hbar^2}{2m}\nabla^2 + A(\vec{r}))w_{i_x,i_y}(\vec{r}). \tag{3.4}$$

The tunnelling along both lattice directions produces part of the Bose-Hubbard lattice

Hamiltonian given by

$$H_J = -J \sum_{i_x,i_y} \hat{b}^\dagger_{i_x+1,i_y}\hat{b}_{i_x,i_y} + H.c. + \hat{b}^\dagger_{i_x,i_y+1}\hat{b}_{i_x,i_y} + H.c., \tag{3.5}$$

where $H.c.$ means the Hermitian conjugate of the preceding term.

On top of the single-particle tunnelling, one can add interactions between bosons. The most important interaction between alkali atoms is the short-range density-density repulsion [1]. Because the lengthscale of the interaction is smaller than the lengthscale over which the Wannier functions vary significantly, this is usually modelled as a delta function multiplied by the contact interaction magnitude $g = 4\pi\hbar^2 l_s/m$, where l_s is the range of the interaction. Because this term is also short-ranged compared to the lattice spacing ($l_s << \lambda$), by far the most significant corresponding lattice Hamiltonian element is repulsion between particles on the same site. This element is calculated using the equation

$$U = g \int d\vec{r} |w_{i_x,i_y}(\vec{r})|^4. \tag{3.6}$$

Adding this to the tunnelling Hamiltonian and expressing the 2D lattice position using a single index, we arrive at the standard BH Hamiltonian

$$H = H_J + H_U = -J \sum_{\langle j,k\rangle} \hat{b}^\dagger_j \hat{b}_k + H.c. + \frac{U}{2}\sum_j \hat{n}_j(\hat{n}_j - 1), \tag{3.7}$$

where $\langle j,k\rangle$ are neighbouring lattice minima j and k and $\hat{n}_j = \hat{b}^\dagger_j \hat{b}_j$ is the density operator for the Wannier function localised at site j. This model now contains two non-commuting terms and can therefore display non-trivial quantum behaviour which is difficult to solve through classical methods, which motivates the desire to solve this model through quantum simulation. Unlike in a condensed matter system, the

---

[1] Note that for polar molecules, the short-range interaction behaviour can involve chemical reactions or sticky collisions, which may complicate the on-site BH interactions if multiple site occupation is significant [142].

model parameters can be controlled precisely in ultracold gas experiments using the laser intensity to set the optical potential and using a magnetic field to modify the short-range scattering length and therefore U [143, 144].

It is worth mentioning that this Hamiltonian omits tunnelling terms beyond nearest-neighbouring Wannier functions and contact interactions between any different Wannier functions on the basis that these elements are negligible. These assumptions can be questionable for weak lattice potentials, while for high site occupation, the short-range repulsion between the large number of particles pushes particles further from the lattice potential minimum and therefore decreases the separation between bosons which are nominally on different sites [145]. A useful review including the effects of these terms is Ref [146].

3.1.1 Dipolar Interactions

The physics of the Bose-Hubbard model can be extended significantly by the inclusion of long-range dipole-dipole interactions. The electric dipole-dipole interaction energy between two equally-polarised dipoles $\vec{d}$ is given by

$$V_{DDI} = \frac{|\vec{d}|^2}{4\pi\epsilon_0 R^3}(1 - 3\cos^2(\alpha)), \tag{3.8}$$

where ϵ_0 is the permittivity of free space, R is the separation between the dipoles, and α is the angle between the dipoles' separation and their (identical) polarisation. For magnetic dipoles $\vec{\mu}$, the factor $\frac{|\vec{d}|^2}{\epsilon_0}$ is replaced by $|\vec{\mu}|^2\mu_0$, where μ_0 is the vacuum permeability. The lattice Hamiltonian elements corresponding to this interaction can be calculated by integrals as before, but the long-range nature of the interaction means that a larger number of off-site terms will be relevant to the physics. A general lattice Hamiltonian element is calculated using

$$V_{i,j,k,l} = \int d\vec{r_1} \int d\vec{r_2} w_i^*(\vec{r_1}) w_j^*(\vec{r_2}) V_{DDI}(\vec{r_1} - \vec{r_2}) w_k(\vec{r_2}) w_l(\vec{r_1}), \tag{3.9}$$

giving the coefficient for the operator $\hat{b}_i^\dagger \hat{b}_j^\dagger \hat{b}_k \hat{b}_l$. For $i = j = k = l$, the on-site DDI operator becomes $\hat{n}_i(\hat{n}_i - 1)$ due to bosonic commutation relations. If the lattice potential is deep enough, the overlap between Wannier functions at different sites will be small and the dominant off-site terms will have $i = l \neq j = k$. These terms correspond to the operator $\hat{n}_i \hat{n}_j$, which physically means the DDI causes an energy shift due to the density of bosons at two different sites without changing the location of any bosons. For weak lattice potentials, the term with $i = k \neq j = l$, which enacts the same lattice operator and physically corresponds to switching the positions of two bosons between sites, makes a small quantitative difference. Adding the dominant off-site and on-site dipole-dipole interaction terms to the Bose-Hubbard model results in the dipolar BH model

$$H = H_J + H_U + H_V = -J \sum_{\langle j,k \rangle} \hat{b}_j^\dagger \hat{b}_k + H.c. + \sum_j \left(\frac{U + V_{j,j,j,j}}{2}\right) \hat{n}_j(\hat{n}_j - 1) + \sum_{i \neq j} \frac{V_{i,j,j,i}}{2} \hat{n}_i \hat{n}_j, \tag{3.10}$$

where the on-site DDI has been absorbed into the contact interaction term.

In the literature, the coefficient $V_{i,j,j,i}$ is often approximated by taking the value of V_{DDI} between the lattice minima of i and j (i.e. approximating V_{DDI} as constant across the spatial extent of the two involved Wannier functions), although this is known to cause inaccuracies in the shorter-range off-site DDI Hamiltonian for weak lattice potentials due to the non-zero spread of the Wannier functions [147]. We also use this assumption in chapter 4 to avoid specialising for a particular lattice potential depth and for comparison with literature, whereas in chapter 5 we integrate over the Wannier functions explicitly for comparison with continuum methods.

Our discussion of the density-density DDI terms has neglected, among other even weaker terms, additional density-induced tunnelling ($V_{i,i,i,j}$) and pair-tunnelling ($V_{i,i,j,j}$)

terms introduced by the DDI in soft-core models [148] (also reviewed in Ref [146]). The most significant of these is the density-induced tunnelling, which creates lattice terms of the form $H_T = -T \sum_{\langle i,j \rangle} \hat{b}_i^\dagger (\hat{n}_i + \hat{n}_j) \hat{b}_j$. The magnitude and sign of T compared to the kinetic tunnelling J depend strongly on the lattice potential, including the confinement in the z direction in a square lattice [149]. The interplay between density-induced, and kinetic tunnelling has recently been studied for its effect on the superfluid and supersolid regions of soft-core dipolar BH models [150–153].

3.1.2 Physics

The physics of the BH model has been studied widely using both numerical methods and experimentally as one of the first demonstrations of quantum simulation of a many-body system. The model produces a ground state transition between a superfluid and a Mott insulator as the ratio of U/J is increased [135], which has been observed using alkali atoms in optical lattices [154–156].

The addition of the dipole-dipole interaction to this model creates density-density interactions between different sites. Unlike the standard BH model, this allows for solid states which have a non-integer filling because tunnelling can additionally be suppressed by off-site interactions. Due to the R^{-3} interaction decay, the strongest of these interactions are on neighbouring sites. Models which keep only these terms are referred to as extended BH models and in square lattices they feature density wave states (characterised by density ordering in the bulk of the system which is different from that imposed by the external potential) and soft-core supersolids [157, 158]. In 1D, the soft-core version of this model also features Haldane insulator states with non-local 'string' order [159, 160]. The effect of further-neighbour couplings is weaker but is important for stabilising the hardcore checkerboard supersolid [31, 161–163] and (softcore or hardcore) diagonal (super)solids [164, 165].

We investigate two separate aspects of the dipolar Bose-Hubbard model in this thesis: In chapter 4, we develop and study the dipolar BH model in real-space cylindrical optical lattices, where the anisotropy of the dipole-dipole interaction combines with the curved lattice surface to produce spatially-varying off-site interactions. In chapter 5, we study the validity of the lowest-band dipolar BH model by comparing exact results for small lattice systems with continuous-space methods that solve the same physical problem by directly including the lattice potential.

4

Dipolar Bose-Hubbard Model in Real-Space Cylindrical Lattices

This chapter contains results published in Ref [46]. MH formulated the project, performed the calculations, and wrote the manuscript. Discussions with Paolo Molignini and Joseph Tindall are acknowledged.

The anisotropy of the DDI significantly affects the physics of dipolar Bose-Hubbard models, as has been studied intensely in 2D square lattices [142, 157, 164, 166, 167]. For curved 1D or 2D optical trap geometries, such as a ring, shell, or torus [168, 169], the anisotropy of the DDI leads to spatially-dependent interactions along the curved surface of the trap [170–175]. This effect has remained less explored for discrete curved lattices such as rings [176, 177], where the lattice vectors between neighbouring sites are not translationally-invariant throughout the lattice. This produces frustration in the simplest scenario of a one-dimensional dipolar zigzag chain [178, 179], while previous work on a single octagonal ring [180] showed that this mechanism can favour the occupation of lattice sites on a particular sublattice imposed by the lattice curvature.

Motivated by proposals for real-space cylindrical optical lattices [181], we study the ground state physics of a hard-core dipolar BH model on a finite cylinder using Matrix Product State (MPS) methods, focussing on the case of a half-filled cylinder with octagonal rings. The addition of the axial direction (in which the lattice vectors do not change) means that the spatially-varying interactions in the azimuthal direction are accompanied by translationally-invariant interactions in the axial direction. For the octagonal-ring cylinder, we find that attractive axial interactions mean that the ground state obeys either spontaneous or sublattice ordering depending on the strength of the interaction variation, which is somewhat similar to the single octagonal ring. For repulsive axial interactions, the interplay between the translationally-invariant and spatially-varying interactions becomes richer, leading to greater interaction competition and entanglement, unusual density-wave states, and sublattice-resolved physics. These results show how the characteristic physics of long-range interacting models can be enhanced through the combination of anisotropic dipolar interactions and curved lattice geometries.

This chapter is arranged in the following manner: We introduce the cylindrical dipolar BH Hamiltonian in section 4.1 followed by the specific case of an octagonal ring in section 4.2. The numerical calculations and physical quantities extracted for cylinders of finite axial length are described in sections 4.3 and 4.4 respectively. Numerical results for the ground state properties for octagonal-ring cylinders of small axial length are presented in section 4.5 and compared for varying axial length in section 4.6. These numerical results are discussed in section 4.7. Different azimuthal lattice lengths and considerations for physical implementation are discussed in sections 4.8 and 4.9 respectively.

4.1 Hamiltonian

In this model, the bosons are trapped in a cylindrical optical lattice with its axis along the z direction. We label the axial coordinate as z, while the potential minima in the $x-y$ plane are labelled by the index c. We first describe a general hard-core cylindrical dipolar BH Hamiltonian before focussing on the case of octagonal rings. The polarised bosons can hop to neighbouring sites both around the ring of the cylinder and along the axis and interact with each other on different sites using the DDI as described in chapter 3. We assume that the short-range repulsion is sufficient to preclude double-occupation of sites. This leads to the following Hamiltonian:

$$H = -J \sum_{\langle j,k \rangle} \hat{b}_j^\dagger \hat{b}_k + H.c. + \frac{V}{2} \sum_{j \neq k} \nu_{j,k} \hat{n}_j \hat{n}_k, \tag{4.1}$$

where j and k are sites with defined axial and azimuthal coordinates (z_j, c_j etc.), $H.c.$ means Hermitian conjugate, $\langle j,k \rangle$ are neighbouring sites, $\hat{b}$ is the hard-core boson annihilation operator such that $\hat{b}^2 = 0$, $\hat{n} = \hat{b}^\dagger \hat{b}$ is the on-site density, J is the tunnelling amplitude, V is the overall dipole-dipole interaction strength, and $\nu_{j,k}$ is a factor describing the distance decay and orientation dependence of the interaction between a given pair of sites which is given by $\nu_{j,k} = (1 - 3\cos^2(\alpha_{j,k}))/R_{j,k}^3$, where $R_{j,k}$ is the distance between sites j and k, and $\alpha_{j,k}$ is the angle between the dipole polarisation direction and the separation vector of these sites. We define the interaction Hamiltonian $\hat{H}_V$ to be the term multiplied by V in equation 4.1. We have assumed that the tunnelling amplitude J is equal around the cylinder and along the axis, and we set $J = 1$ throughout this chapter.

The value of $\alpha_{j,k}$ depends not only on the separation of the sites in terms of the cylindrical lattice indices (i.e. $z_k - z_j$ and $c_k - c_j$) but also on the value of c itself. This means for example that the DDI between neighbouring sites around the ring of the cylinder takes different values depending on the azimuthal coordinates of the sites.

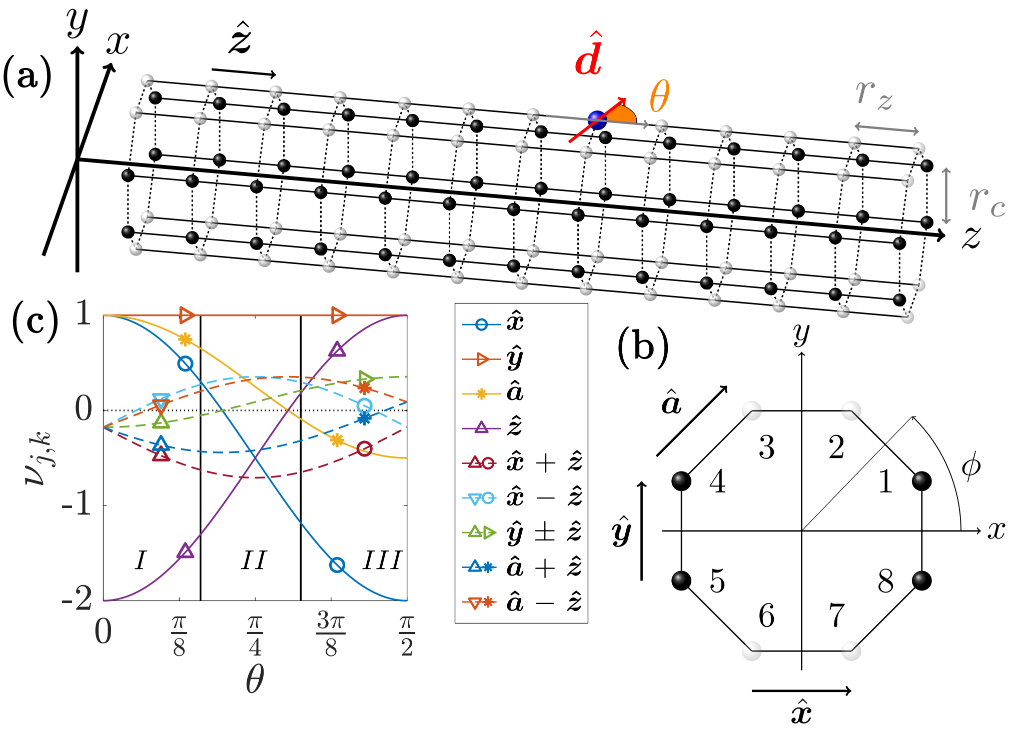

Figure 4.1: (a) Schematic of cylindrical optical lattice with octagonal rings. A single Boson is shown in blue, whose polarisation $\boldsymbol{d}$ is denoted by the red arrow, which is confined to the $x - z$ plane at angle θ to the z axis. The lattice sites are shaded to denote the two sublattices of inequivalent sites, where filled circles show the equatorial sublattice and hollow circles show the polar sublattice. (b) Single ring of the cylindrical lattice with value of azimuthal coordinate c labelled. The nearest-neighbour in-plane lattice vectors are shown. (c) Nearest-neighbour and next-nearest-neighbour dipolar interactions for octagonal ring cylinder. Nearest-neighbour interactions are plotted in solid lines with single markers for identification while next-nearest-neighbour interactions are plotted in dashed lines and have two different markers. A dotted line at zero interaction strength is also added. The domain of θ is split into regions I, II, and III according to the ground-state physics at strong interaction.

4.2 Octagonal Ring Cylinder

We now describe the geometry of the octagonal ring lattice on which the majority of our numerical calculations were based (a brief discussion of other values of L_c is included in section 4.8). Figure 4.1(a) shows a diagram of the octagonal ring optical lattice, which has $L_c = 8$ sites in the azimuthal direction and $L_z = 13$ sites in the axial direction. For comparison with square optical lattices, we have assumed that the separation between neighbouring sites is the same in the axial and azimuthal directions and we set this as our length unit ($r_z = r_c = 1$). We focussed on a fixed particle number N, which unless stated was set at half the number of sites $N = N_H = \frac{L_z L_c}{2}$.

The orientation of the cylinder in the $x - y$ plane is shown in Figure 4.1(b). Due to the curved nature of the octagonal ring, there are three different in-plane lattice vectors. We have focussed on the case where the polarisation direction is in the $x - z$ plane (similar to Ref [180]), where $\theta = 0$ means the polarisation is in the z direction and $\theta = \frac{\pi}{2}$ means the polarisation is in the x direction. This decision leads to two sublattices of inequivalent sites which we call the 'polar' and 'equatorial' sublattices. Sites on the polar sublattice have one neighbour in the $\pm\hat{\boldsymbol{x}}$ direction and one in the $\pm\hat{\boldsymbol{a}}$ direction. Sites on the equatorial sublattice have a neighbour in the $\pm\hat{\boldsymbol{y}}$ direction instead of the $\pm\hat{\boldsymbol{x}}$ direction.

The spatially-dependent interactions can be seen in Figure 4.1(c), where the values of $\nu_{j,k}$ for all different nearest-neighbour and next-nearest-neighbour interactions are plotted as a function of θ for the octagonal-ring cylinder lattice with polarisation in the $x - z$ plane. The domain of θ is split into three regions (labelled I, II, and III) according to the different qualitative mechanisms for stabilisation of the ground state at strong interaction, which are discussed in section 4.5.

In particular, the in-plane nearest neighbour interactions (in the $\hat{\boldsymbol{x}}$, $\hat{\boldsymbol{y}}$, and $\hat{\boldsymbol{a}}$ directions) are equally repulsive at $\theta = 0$ but take maximally-separated values at $\theta = \frac{\pi}{2}$. For $\theta \neq 0$, the interactions in $\hat{\boldsymbol{x}}$ are more attractive than the interactions in $\hat{\boldsymbol{y}}$, which means that the polar sublattice generally provides more attractive interactions than the equatorial sublattice. Meanwhile, the axial interactions in the $\hat{\boldsymbol{z}}$ direction switch from attractive to repulsive at $\theta \approx 57.4°$. The other important point from Figure 4.1(c) is that the next-nearest-neighbour interactions in the $\hat{\boldsymbol{x}} + \hat{\boldsymbol{z}}$ and $\hat{\boldsymbol{a}} + \hat{\boldsymbol{z}}$ directions are always more attractive than their counterparts in the $\hat{\boldsymbol{x}} - \hat{\boldsymbol{z}}$ and $\hat{\boldsymbol{a}} - \hat{\boldsymbol{z}}$ directions respectively because $0 \leq \theta \leq \frac{\pi}{2}$, which becomes relevant when the nearest-neighbour interactions compete against each other.

Other values of the azimuthal polarisation angle ϕ, which we have not studied in depth, lead to significantly altered interactions and sublattice structure. For example, if $\phi = \frac{\pi}{8}$, there are now three sublattices of two, four, and two sites respectively. These sublattices mean the physics of this system is less easily interpretable than for $\phi = 0$ with half of the sites occupied. For less symmetric choices of ϕ and even L_c, each site on the ring is only equivalent to the opposite site halfway around the ring. For odd L_c, all sites on the ring can be made to have inequivalent interactions.

4.3 Finite Lattice and Numerical Calculations

We study the ground states of this system using the Density Matrix Renormalisation Group (DMRG) algorithm to optimise over matrix product states [182–184]. This method is well-suited to one-dimensional or thin two-dimensional systems. It avoids the full numerical cost of the exponentially-growing Hilbert space of many-body quantum systems by truncating the representation of highly entangled states using the maximum bond dimension χ, where a large χ increases both the Hilbert space and numerical cost. A brief description of this method and its application to thin dipolar systems is included in appendix E.

We note that the proposed cylindrical optical lattice [181] exhibits inhomogeneity in the z direction, which limits the number of rings which can be coupled. For this reason, to reduce computational cost, and to ensure the ground states we find are compatible with edge effects, we first investigate the finite-size ground state diagram as a function of V and θ for octagonal-ring lattices with a small finite axial length $10 \leq L_z \leq 16$ in section 4.5. Unless otherwise stated, we report results for $L_z = 13$, which has $N = 52$ Bosons in 104 sites.

In section 4.6, we vary L_z (with $L_c = 8$ fixed) to study the extrapolation of the finite-size phase transitions to the thermodynamic limit. For fixed L_c, the thermodynamic limit for $L_z \to \infty$ is one-dimensional. Further details of the numerical calculations are in appendix F.

4.4 Observables

We use the one-body-density-matrix (OBDM) $\rho_{j,k} = \langle \hat{b}_j^\dagger \hat{b}_k \rangle$ to study the tunnelling coherences of the state, specifically through the bosonic orbital occupations (eigenvalues of the OBDM) and inter-site correlations. While the population fraction in the largest bosonic orbital vanishes in the one-dimensional thermodynamic limit for hard-core bosons, this reduction with cylinder length is slower in a superfluid (SF) state [1] than in a solid state ($\propto \frac{1}{L_z}$), contrasting these states in finite-size systems. To more easily contrast the solid state, we denote the difference between the population fractions of the two largest bosonic orbitals by δ_e, which is zero for a solid and would take the value 1 for a perfect condensate. By studying the spatial distribution of the bosonic orbitals, we can resolve this information according to the sublattices defined by the Hamiltonian. The decay of intersite correlations with axial separation is expected to be polynomial in the SF state, and exponential for solid states.

To study density-wave ordering, we use density-density correlation functions. To mitigate the impact of edge effects on these order parameters, we include pairs of sites in the 'bulk', where the bulk is defined by removing the three rings on each edge of the cylinder. (For most ground states, trimming two rings or even one ring from the edges of the cylinder is sufficient to quantitatively capture the bulk behaviour). To

---

[1] For hard-core bosons in a strictly one-dimensional lattice, the population fraction in the largest orbital scales as $\propto L_z^{-0.5}$ [185]. For our eight-leg cylindrical ladder, this scaling is found to be approximately $\propto L_z^{-0.1}$.

express this as a normalised order parameter and to avoid the effect of inhomogeneous average density between the bulk and the edges, we calculate the order parameter

$$M_{\Delta_z,\Delta_c} = \frac{\sum_{\{z,c\}} \langle \hat{n}_{z,c}\hat{n}_{z+\Delta_z,c+\Delta_c}\rangle - n_b n_0^2}{n_b n_0^2}, \tag{4.2}$$

where $\{z, c\}$ is a site such that both (z, c) and $(z+\Delta_z, c+\Delta_c)$ are in the bulk, n_b is the number of such sites in the bulk, and n_0 is the average density in the bulk. This means that M_{Δ_z,Δ_c} is equal to $1(-1)$ if the presence of a particle on a given site (z_1, c_1) implies the presence(absence) of a particle on site $(z_1 + \Delta_z, c_1 + \Delta_c)$, while it is 0 if there is no correlation beyond the average bulk density. For displacements which only include density-density correlations between sites on the same sublattice, it is also possible to define this order parameter separately for the polar and equatorial sublattices.

The Von-Neumann entanglement entropy S_{ent} is a useful quantity to identify finite-size precursors to phase transitions without needing to know the order parameters of the phases in advance [163, 186, 187]. S_{ent} is defined using the Schmidt decomposition of the MPS approximation of the ground state $|\psi\rangle$ into two subsystems A and B. We can write any state of the full system as

$$|\psi\rangle = \sum_{w}^{W} \lambda_w \, |\mu_w\rangle_A \, |\nu_w\rangle_B \,, \tag{4.3}$$

where $|\mu_w\rangle_A$ are mutually orthogonal states of subsystem A and $|\nu_w\rangle_B$ are defined accordingly for subsystem B. W is the number of terms in the Schmidt decomposition for which the Schmidt value λ is non-zero (the 'Schmidt rank') and is upper-bounded by the dimension of the smaller subsystem and by the maximum bond dimension χ of the MPS. For a product state, $W = 1$. S_{ent} is defined from this decomposition using

$$S_{ent} = -\sum_{w}^{W} |\lambda_w|^2 \ln(|\lambda_w|^2), \tag{4.4}$$

meaning that S_{ent}, and the difficulty of acceptably approximating the state by an

MPS of a given bond dimension χ, increase for highly-entangled states with a large number of different Schmidt values of comparable magnitude.

For continuous phase transitions, the first derivative of the entanglement entropy or other entanglement measures has been used to identify phase transitions instead [188–191]. The second-order Rényi entropy (which provides a lower bound for S_{ent}) can be measured experimentally in optical lattice systems [192, 193], although the cylindrical lattice geometry would complicate the implementation of these schemes relative to 1D chains or square lattices. In this chapter, we use S_{ent} to refer specifically to the entanglement entropy across the centre-most bipartition of the lattice which cuts directly across the cylinder axis, while S_{L_A} refers to the entanglement entropy of a bipartition which again cuts directly across the cylinder axis and for which subsystem A at lower z has length L_A.

4.5 Ground State Diagram for $L_z = 13$

To describe the behaviour of the model as a function of V and θ, we first present a finite-size phase diagram in Figure 4.2. At strong interaction there are different classes of solid ground states in regions I, II, and III, which we label the stripe solid (SS), polar solid (PS), and repulsive solid (RS) respectively. The RS is a collection of states, containing the diagonal stripe (DS) and dual checkerboard (DC), whose boundaries are strongly influenced by the cylinder length. We discuss the three solid regions separately in sections 4.5.1, 4.5.2, and 4.5.3 respectively as the role of the spatially-varying interactions is qualitatively distinct. At weak interaction, the ground state (labelled TU) is dominated by tunnelling but is slightly warped by the interactions. We discuss this state in more detail in section 4.5.4. These states are compared from the perspective of off-diagonal correlation decay in section 4.5.5 and using wavefunction

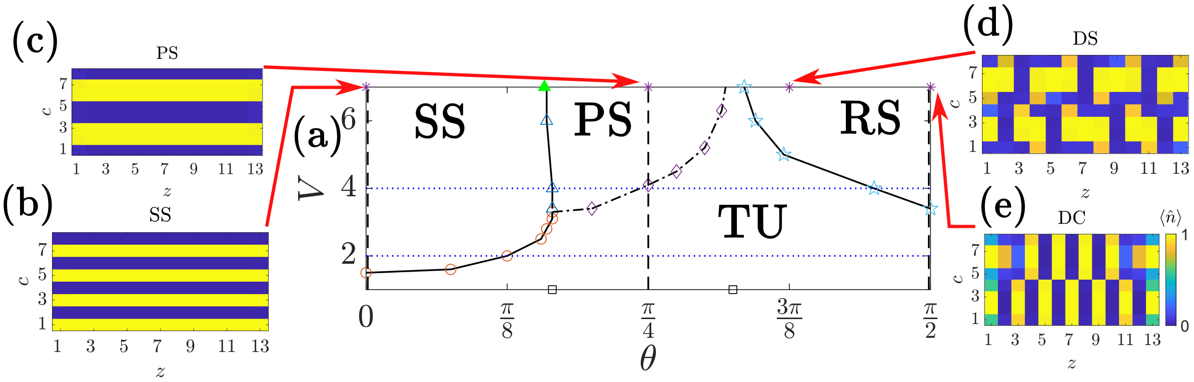

Figure 4.2: (a) Finite-size phase diagram of model using $L_z = 13$ as a function of polarisation angle θ and DDI strength V. Unfilled markers show calculated finite-size transition points, while black lines join neighbouring markers as a guide to the eye (the boundary between the polar solid PS and tunnelling state TU regions is marked in a dash-dotted line to indicate that the transition is not sharp at this cylinder length). The filled triangle shows the value of θ for transition from stripe solid SS to PS as $V \to \infty$. Physical parameters are plotted for the dashed lines at $\theta = \{0, \frac{\pi}{4}, \frac{\pi}{2}\}$ in Figure 4.3 and for the dotted lines at $V = \{2, 4\}$ in Figure 4.4. Black square markers on the θ axis mark the boundaries between regions I, II, and III. (b-e) Expectation value of on-site density for solid ordered states (stripe solid, polar solid, diagonal stripe, and dual checkerboard respectively) using the common colour scale shown next to (e). Arrows point towards the corresponding parameters in (a) marked by asterisks.

overlap in section 4.5.6. The table below summarises the expected values of certain physical parameters for the states which we quantitatively analyse, with the caveat that δ_e would slowly vanish for $L_z \to \infty$ for the TU state.

	δ_e	S_{ent}	$M_{\Delta_z=0,\Delta_c=2}$	$M_{\Delta_z=0,\Delta_c=4}$
TU	>0	>0	0	0
SS	0	0	$+1$	$+1$
PS	0	0	-1	$+1$
DC	0	0	0	-1

We plot the values of these physical parameters for constant θ in Figure 4.3 and for constant V in Figure 4.4. We describe the qualitative ground states and transitions which lead to this behaviour in the rest of this section.

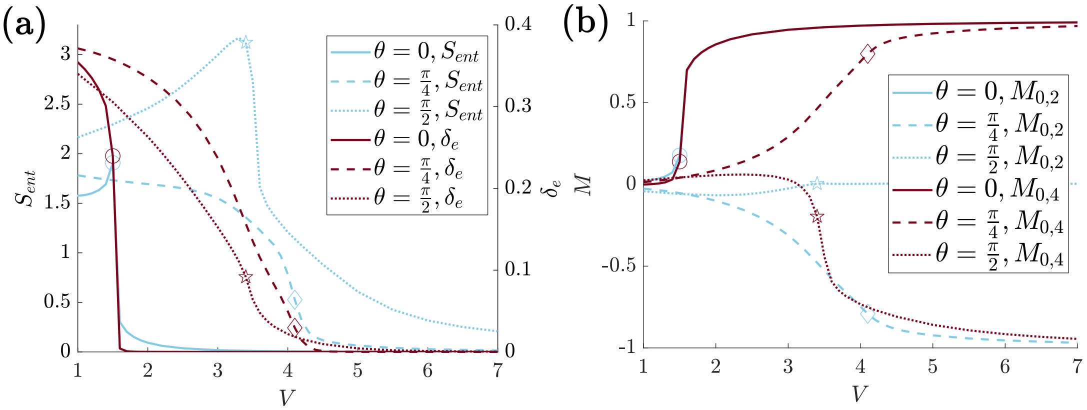

Figure 4.3: Physical parameters for polarisation angles $\theta = 0$, $\frac{\pi}{4}$, and $\frac{\pi}{2}$ as a function of DDI strength V corresponding to the black dashed lines in Figure 4.2(a). (a) Von-Neumann entanglement entropy S_{ent} on the left-hand y-axis and condensation parameter δ_e on the right-hand y-axis. (b) Solid order parameters M_{Δ_z,Δ_c}. Any transition lines from Figure 4.2(a) which are intersected by these graphs are denoted with the corresponding marker.

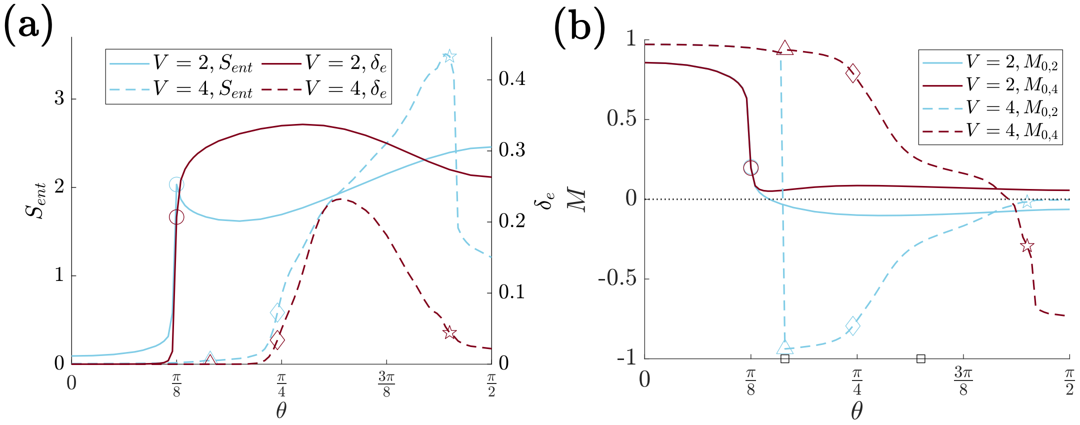

Figure 4.4: Physical parameters for DDI strengths $V = 2$ and $V = 4$ as a function of polarisation angle θ corresponding to the dotted lines in Figure 4.2(a). (a) Von-Neumann entanglement entropy S_{ent} on the left-hand y-axis and condensation parameter δ_e on the right-hand y-axis. (b) Solid order parameters M_{Δ_z,Δ_c}. A black dotted line at $M_{\Delta_z,\Delta_c} = 0$ is added as a guide to the eye. Any transition lines from Figure 4.2(a) which are intersected by these graphs are denoted with the corresponding marker. The boundaries between regions I, II, and III are denoted by black square markers on the θ axis in (b).

4.5.1 Stripe Solid

The simplest solid state we find is the stripe solid (shown in Figure 4.2(b)), in which particles line up along the axis while occupying alternating sites on the ring. This state forms at strong interaction for low θ, where the spatial variation of interactions is small and this model is similar to a flat square lattice Bose-Hubbard model in which the polarisation is along a lattice vector. The SS state spontaneously breaks a discrete symmetry (because there are two degenerate choices of which four values of c are occupied). While any superposition of these two degenerate orthogonal states is also a ground state of the system, DMRG is biased towards the low-entanglement product states and outputs one of the two configurations during optimisation [194], which can be selected by choice of the initial MPS.

Because of the discrete symmetry-breaking property of the SS, the TU-SS transition breaks one discrete symmetry, which means the transition is expected to be second-order. Our numerical results show abrupt changes in order parameters even in this relatively small system, with a sharp increase in the solid order parameters, the vanishing of δ_e, and the peak in S_{ent} all occurring at $V = 1.5$ for $\theta = 0$ as shown in Figure 4.3. We use finite-size scaling to analyse this transition in the thermodynamic limit in section 4.6.1. The same signatures of the transition are observed for $\theta = \frac{\pi}{8}$ for $V = 2$ as shown in Figure 4.4. Increased dipolar interaction strength favours the SS over the TU while increasing θ has the opposite effect because it weakens the attractive interactions in $\hat{\boldsymbol{z}}$ and the repulsive interactions in $\hat{\boldsymbol{x}}$.

4.5.2 Polar Solid

In the polar solid, the bosons line up along the axis occupying the polar sites on each ring, as shown in Figure 4.2(c). While the SS is formed by the same mechanisms that exist in flat lattices, the PS is formed due to the sublattice structure of the spatially-varying azimuthal interactions. The bosons fill the polar sublattice in order to benefit from the interactions in $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{x}} \pm \hat{\boldsymbol{z}}$ which are more attractive than the equivalent interactions in $\hat{\boldsymbol{y}}$ and $\hat{\boldsymbol{y}} \pm \hat{\boldsymbol{z}}$, while still experiencing the attractive interactions in $\hat{\boldsymbol{z}}$. This means that there is a unique sublattice with the lowest energy, an effect which also stabilises ordered states for Bose-Hubbard models in externally imposed alternating potentials [195–197]. As there is no discrete symmetry spontaneously broken by the formation of the PS, the TU-PS transition is much more gradual. It does not feature a peak in S_{ent}.

The TU-PS boundary for $L_z = 13$ in Figure 4.2, which we determined by a maximum in $-\frac{\partial S_{ent}}{\partial V}$ for constant θ, should therefore be interpreted more as a contour at the sharpest changes in physical parameters rather than a sharp phase boundary. For the markers shown in Figure 4.2, this maximum coincided with a local minimum in the overlap between the ground state wavefunctions of neighbouring datapoints at fixed θ and varying V (as shown in Figure 4.5(a)), and also fixed V and varying θ. This means that there is a local maximum in the 'fidelity susceptibility' with respect to V, which has been used to pinpoint phase transitions [187, 198–200]. Due to the smoothness of this transition, we find that the much larger system sizes investigated in section 4.6.2 are necessary to resolve the physics at even a qualitative level, so we defer a full discussion of this transition and the fidelity susceptibility to that section, and discuss here only the difficulties of locating it in a small system.

As a consequence of the facts that the TU-PS transition is smooth and affected by the greater tunnelling at the cylinder edges, the choice of how to locate the transition affects the reported boundary (compared to the TU-SS transition whose location is clear). Our use of the entanglement entropy and wavefunction overlap mean this boundary is quite sensitive to edge tunnelling, particularly at large V. This means that the state immediately on the TU side of the boundary has a large proportion of

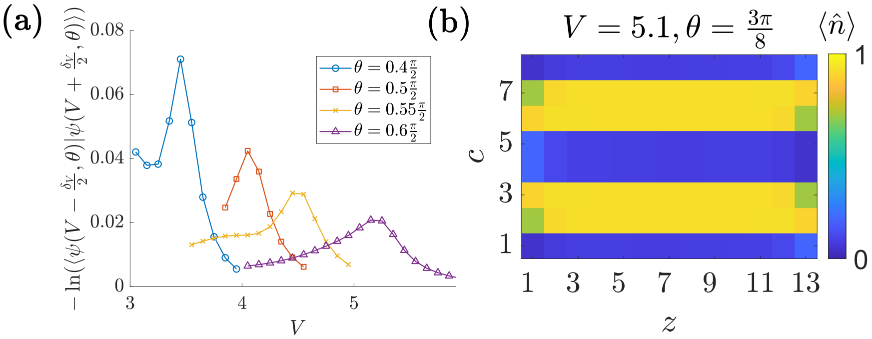

Figure 4.5: (a): Overlaps of neighbouring wavefunctions for PS-TU transition as a function of DDI strength V for different polarisation angles θ, plotted logarithmically to show all on the same scale. $\delta_V = 0.1$ in this plot, which is the minimum spacing between datapoints. The points on this graph are halfway between parameters used to calculate ground states. The chosen transition point is the datapoint which has the lowest and next-lowest overlap on either side as a function of V. (b) Example on-site density for a ground state immediately on the TU side of the boundary.

its tunnelling energy localised at the cylinder edges, as shown in Figure 4.5(b).

States close to the TU-PS boundary experience a subtle edge effect due to the spatially-varying interactions, as shown in Figure 4.5(b). This favours the occupation of sites $c = 3$ and $c = 6$ over $c = 2$ and $c = 7$ on the first ring of the lattice while the opposite is true for the final ring of the lattice, despite all of these sites being in the polar sublattice. This is because the four preferentially-occupied sites experience a more attractive $\hat{\boldsymbol{x}} + \hat{\boldsymbol{z}}$ interaction with the occupied sites on the neighbouring ring inside the lattice but do not experience the less attractive $\hat{\boldsymbol{x}} - \hat{\boldsymbol{z}}$ interaction because the corresponding site is outside the edges of the lattice, while the opposite applies to the other polar sites on the edge rings.

The SS-PS transition meanwhile is sharp. This transition does not cause changes in orbital occupation or S_{ent} because both states are solid, but the $M_{0,2}$ order parameter changes abruptly from $+1$ for the SS to -1 for the PS, as shown in Figure 4.4(b). Both of these solid states experience the attractive interactions in $\hat{\boldsymbol{z}}$ equally, while the PS has additional interactions in $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{x}} \pm \hat{\boldsymbol{z}}$. For $V >> J$, this transition therefore occurs when the sum of the additional interactions in the PS becomes attractive. For interactions restricted to next-nearest-neighbour, this requires a polarisation angle given by

$$1 - 3\sin^2(\theta_{SS-PS}) = \frac{1}{\sqrt{8}}\frac{L_z - 1}{L_z}. \tag{4.5}$$

For an infinitely-long cylinder, $\theta_{SS-PS} \approx 0.307\frac{\pi}{2}$ while for $L_z = 13$ the fact that only the repulsive $\hat{\boldsymbol{x}}$ interaction is present for the first ring pushes the transition to $\theta_{SS-PS} \approx 0.314\frac{\pi}{2}$. This agrees with our numerical results for $V = 6$, where we find an SS ground state for $\theta = 0.31\frac{\pi}{2}$ and a PS ground state for $\theta = 0.32\frac{\pi}{2}$. For $V = 4$ the transition occurs at slightly higher θ because the SS is able to lower its energy more than the PS by incorporating low levels of particle tunnelling at the cylinder edges. We remark that for $\theta = 0.32\frac{\pi}{2}$ we found DMRG would occasionally select the SS or PS when the other solid state was favoured because the optimisation became stuck in a local minimum. To avoid this, we performed calculations using both the SS and PS as starting wavefunctions and compared the final density distributions and energies after DMRG caused them to acquire small amounts of tunnelling.

The PS can be considered as the two-dimensional analogue of the state predicted in the single octagonal ring for polarisation in the $\hat{\boldsymbol{x}}$ direction [180]. It is formed gradually due to the sublattice structure of the azimuthal interactions and is compatible with the axial interactions for $0.32\frac{\pi}{2} \leq \theta \leq 0.65\frac{\pi}{2}$ because the axial interactions are either attractive or very weakly repulsive for these polarisation angles. Unlike for the single octagonal ring, this state becomes disrupted by sufficiently repulsive axial interactions as we describe in the next subsection, even though the spatial variation of interactions which creates the PS becomes stronger as the polarisation tilts away from the axis.

4.5.3 Repulsive Solid

For $0.65\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2}$, repulsive axial interactions compete with the sublattice structure of the azimuthal interactions, which leads to greater frustration and entanglement and more complicated density orderings. At strong interaction, most polarisation angles do not produce a qualitatively ordered state in small systems but rather interpolate between domains of states which do have recognisable order as a function of θ. We therefore begin by describing these recognisable ordered states.

4.5.3.1 Dual Checkerboard

The dual checkerboard (DC) is formed at strong interaction for $\theta \approx \frac{\pi}{2}$ where both the spatial variation of azimuthal interactions and the axial repulsion are strongest. To minimise the repulsion while benefiting from the attractive interactions in $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{a}}$, the bosons occupy sites $c = \{1, 2, 3, 4\}$ and $c = \{5, 6, 7, 8\}$ on alternate rings in the bulk, as shown in Figure 4.2(e). Residual tunnelling at finite V is stronger for the equatorial sublattice ($c = \{1, 4, 5, 8\}$) because the attraction is weaker than for the polar sublattice. While there are clear edge effects, the bulk ordering is equally compatible with all cylinder lengths and is therefore largely unaffected by L_z. This state has a two-fold degeneracy like the SS but unlike the PS.

The TU-DC transition is significantly affected by the spatially-varying interactions, as shown in Figure 4.6. The TU state at high θ features axial density modulation on the polar sublattice (Figure 4.6(b)), which emerges from the cylinder edges due to the axial repulsion. This partial ordering does not penetrate the bulk before the transition to the DC for either even or odd L_z. The solid order and suppression of tunnelling build more quickly on the polar sublattice than the equatorial sublattice due to the spatial variation of interactions. This is shown in Figure 4.6, where we plot sublattice-resolved versions of certain physical parameters. $M_{0,4,P(E)}$ is the $M_{0,4}$ solid

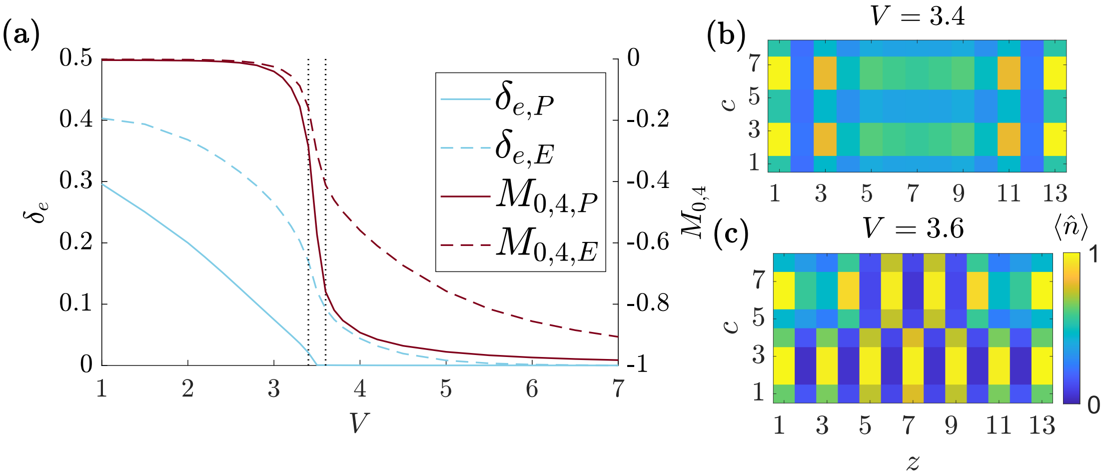

Figure 4.6: (a) Sublattice-resolved physical parameters for the TU-DC transition at polarisation direction $\theta = \frac{\pi}{2}$ as a function of DDI strength V. (b) and (c) On-site densities corresponding to the DDI strengths denoted by vertical black dotted lines in (a), shown on a common colour scale.

order parameter including only pairs of sites which are both on the polar(equatorial) sublattice, while $\delta_{E,P(E)}$ shows the difference in polar(equatorial) sublattice population of the two largest bosonic orbitals. For the polar sublattice, $\delta_{E,P}$ falls sharply to zero across the transition and the DC solid order builds quickly. On the equatorial sublattice, both changes are much more gradual. We have checked that the significant tunnelling on the equatorial sublattice at the cylinder edges for relatively low values of V in the DC region is not responsible for the larger occupation of the largest bosonic orbital, because this orbital has its largest occupation close to the cylinder centre.

4.5.3.2 Diagonal Stripe

The diagonal stripe (DS) family of states is formed from the PS by displacing bosons from two sites of the polar sublattice to the equatorial sublattice per period such that they attach in a diagonal pattern to the remaining blocks of four or six occupied sites on the polar sublattice. The diagonal alignment is set by the fact that the next-nearest-neighbour $\hat{\boldsymbol{x}} + \hat{\boldsymbol{z}}$ interaction is more attractive than the $\hat{\boldsymbol{x}} - \hat{\boldsymbol{z}}$ interaction,

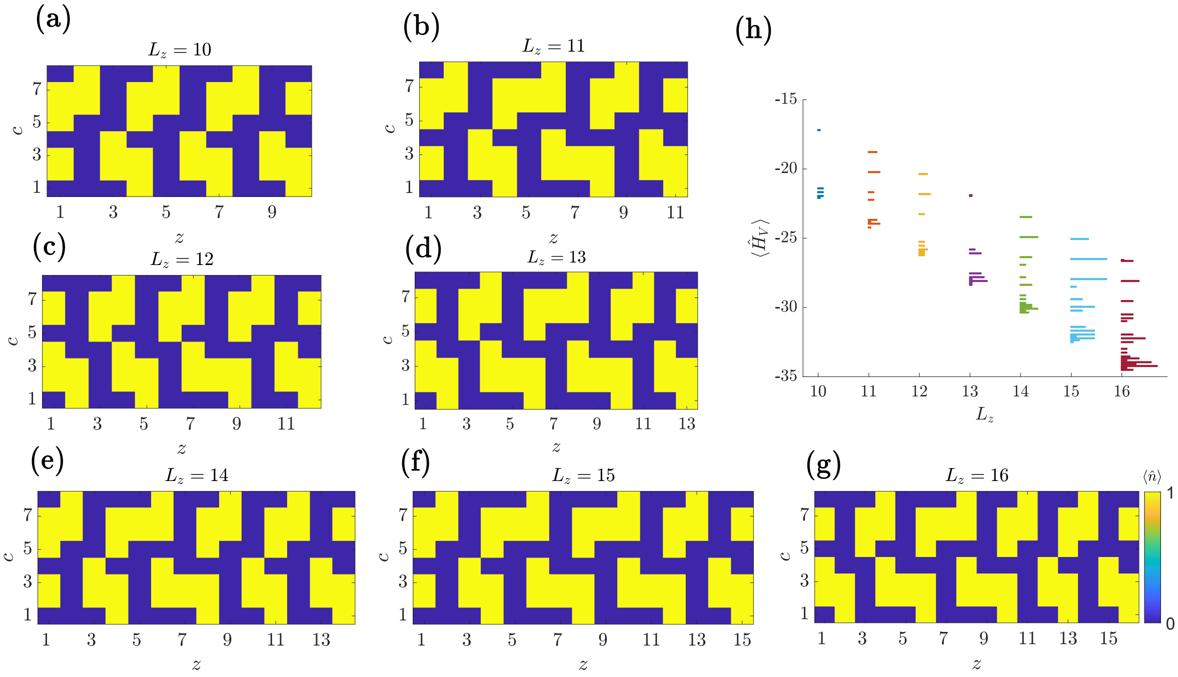

Figure 4.7: (a-g) A DS-type state with the lowest interaction energy for various cylinder lengths at polarisation angle $\theta = \frac{3\pi}{8}$ with common colour scale. (h) Interaction energies $\langle \hat{H}_V \rangle$ of all DS-type states at $\theta = \frac{3\pi}{8}$ for different L_z. Degenerate states are displaced along the x-axis.

showing the importance of the next-nearest-neighbour interactions when the axial and azimuthal interactions compete.

We note that for very strong interaction, a very small range of θ close to the PS, and specific large values of L_z (such as $V = 10$, $\theta = 0.67\frac{\pi}{2}$, and $L_z = 35$), DMRG converges to a state featuring blocks of ten polar sites due to edge effects, even though this block size is not favoured for $V \to \infty$ on an infinitely-long cylinder at any value of θ. For simplicity and due to their reliance on specific values of L_z and extremely strong DDI, we do not include these block sizes in our discussion of the DS family of states.

For interactions truncated to next-nearest-neighbour, the interaction energy densities for infinitely-long cylinders with blocks of four or six polar sites respectively are very similar in region III (crossing from favouring blocks of six to blocks of four at $\theta = 0.76\frac{\pi}{2}$).

Due to this similarity, for a given L_z, the periodicity of this state can be disrupted by edge effects and the interactions between the blocks. This is shown by Figure 4.7 which displays the DS-type state which has the lowest interaction energy for $10 \leq L_z \leq 16$ and $\theta = \frac{3\pi}{8}$ [2]. For all L_z shown except $L_z = 10$, the favoured DS-type state contains blocks of both sizes arranged to minimise their mutual repulsion while ensuring that there are occupied sites at the edges of the lattice due to axial repulsion. It is clear that the length of the cylinder impacts the density distribution deep within the bulk, which means that the cylinder lengths studied here do not produce periodic bulk order. The combination of different block lengths in DS-type states of finite cylinder length leads to multiple degenerate states for interactions truncated at next-nearest-neighbour for $V >> J$, as shown in Figure 4.7(h).

We investigated the existence of DS states for these values of L_z using very strong interaction $V = 7$ and $0.72\frac{\pi}{2} \leq \theta \leq 0.77\frac{\pi}{2}$ by comparing DMRG results when initialised with a random state with small χ or one of the DS-type states with lowest interaction energy (Figure 4.7). We found that all values of L_z converged to qualitative DS-type states for at least one value of θ, but that for other values of θ, the states for $L_z = 12$ and 16 in particular could become significantly distorted during energy minimisation even when initialised with the DS-type state, which is a clear signal that this state is not a local (let alone global) minimum of the energy. This suggests that the edge effects have a significant impact on the stability of the DS state, even at such strong interaction strength.

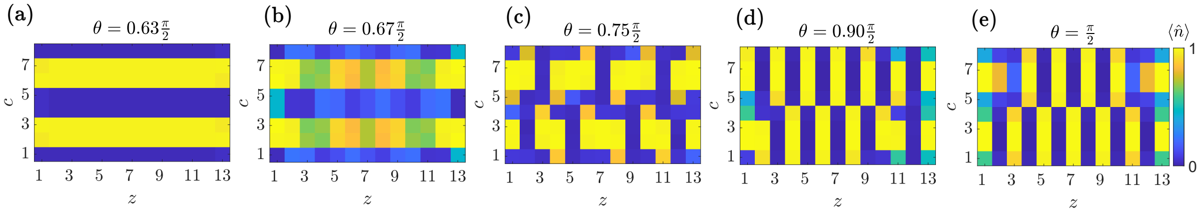

Figure 4.8: On-site density for ground states at strong DDI strength $V = 7$ for varying polarisation angle θ using common colour scale. (a) PS at $\theta = 0.63\frac{\pi}{2}$ (b) Local maximum in Von-Neumann entanglement entropy S_{ent} at $\theta = 0.67\frac{\pi}{2}$ (c) DS-type state at $\theta = 0.75\frac{\pi}{2} = \frac{3\pi}{8}$ (d) Domains of DS and DC at $\theta = 0.9\frac{\pi}{2}$ (e) DC at $\theta = \frac{\pi}{2}$.

4.5.3.3 Changes due to Polarisation Direction

We now describe qualitatively how the ground state changes as a function of θ at strong interaction in Region III. The corresponding on-site densities are plotted in Figure 4.8. Firstly, the PS is destroyed as the axial interactions become sufficiently repulsive to push particles to the cylinder edges where they show slightly increased tunnelling. As θ increases further, the bulk of the PS fractures into smaller blocks on the polar sublattice. For the values of V we studied, we find that this process coincides with sharp decrease in dominant orbital occupation and the local maximum in S_{ent} which we use to pinpoint the boundary of the RS region. If edge effects are satisfied, these blocks then acquire DS order which remains stable over a finite range of θ. As θ increases further, the state interpolates between the DS and the DC through a number of states which do not show any qualitative order in our finite system. The size of the blocks of particles on the polar sublattice decreases as θ increases, eventually arriving at the DC state when all particles on the polar sublattice are arranged in blocks of two. Our numerical results are unable to determine the nature of the transition between the DS and DC states. Similar difficulties have been reported for solid-solid transitions driven by polarisation direction

---

[2] We define a DS-type state on a finite lattice as a product state featuring blocks of four or six occupied polar sites, with the leftover particles on the equatorial sublattice attached to those blocks in the energetically-optimal 'diagonal' fashion. Blocks of two occupied polar sites are allowed at the edges because they are favoured by edge effects.

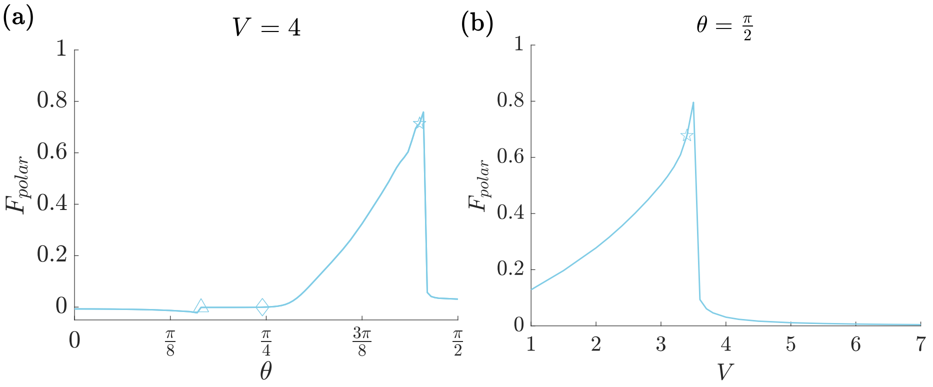

Figure 4.9: Fluctuation parameter F_{polar} in cylinder bulk for (a) DDI strength $V = 4$ and (b) polarisation angle $\theta = \frac{\pi}{2}$.

in flat Bose-Hubbard models, even for relatively large system sizes ($\approx 40 \times 40$) [166].

4.5.4 Tunnelling State

The state at weak interaction is dominated by tunnelling and features a large occupation of a single bosonic orbital in small systems. It has large entanglement between sites because the local particle number fluctuates due to tunnelling while the global particle number is conserved [201]. At $\theta = 0$ it is a superfluid state with algebraically-decaying correlations in the axial direction (see section 4.5.5). At larger θ and moderate V, this relationship becomes warped as the density smoothly acquires an imbalance in favour of the polar sublattice due to the spatial variation of interactions. In region III, the TU also features a weak axial density modulation, especially at the edges of the system, because of the repulsive axial interactions.

In region III, the tunnelling state has very large entanglement entropy $S_{ent} \approx 3.5$ at moderate V as shown in Figures 4.3(a) and 4.4(a). This is due to the competition between the sublattice ordering (favouring the polar sublattice) and the axial repulsion (penalising occupation of the polar sublattice on neighbouring rings). This manifests itself in bulk density fluctuations which favour simultaneous occupation of neighbouring polar sites on the same ring but disfavour simultaneous occupation of neighbouring polar sites along the axis. These fluctuations decrease the interaction energy by increasing the attraction in the $\hat{\boldsymbol{x}}$ direction and reducing the axial repulsion. They are present for both odd and even L_z and are larger at the centre of the system. We show these fluctuations in Figure 4.9 using the quantity

$$F_{polar} = \frac{1}{L_z - 7} \sum_{z=4}^{L_z-4} \sum_{c=2,6} \left(\langle \hat{n}_{z,c} \hat{n}_{z,c+1} \rangle - \langle \hat{n}_{z,c} \rangle \langle \hat{n}_{z,c+1} \rangle \right) - \left(\langle \hat{n}_{z,c} \hat{n}_{z+1,c} \rangle - \langle \hat{n}_{z,c} \rangle \langle \hat{n}_{z+1,c} \rangle \right) \tag{4.6}$$

which has a maximum value of 1 for hard-core bosons. This quantity grows with increasing V and θ in the TU state in region III until just past the TU-RS transition.

It is not clear from our results whether the nominal TU state in region III at moderate V is a distinct phase characterised by large entanglement entropy, or whether it is an extended transition region between the superfluid-type TU state and the RS states. For the system lengths which are numerically feasible despite large S_{ent}, there is no clear transition point between this state and the superfluid-type TU state which exists at weak V and no corresponding local minimum of the overlap of neighbouring wavefunctions. In the absence of a clear signature otherwise, we have used the common TU label for both of these states based on their shared lack of solid order and significant occupation of the dominant bosonic orbital in small systems.

4.5.5 Correlations

In this section, we investigate the decay of off-diagonal correlations in the OBDM as a function of distance along the cylinder axis. This decay is expected to be algebraic for superfluid states, following the relationship

$$\langle \hat{b}^{\dagger}_{z,c} \hat{b}_{z+\Delta_z,c} \rangle \propto \Delta_z^{-p}, \tag{4.7}$$

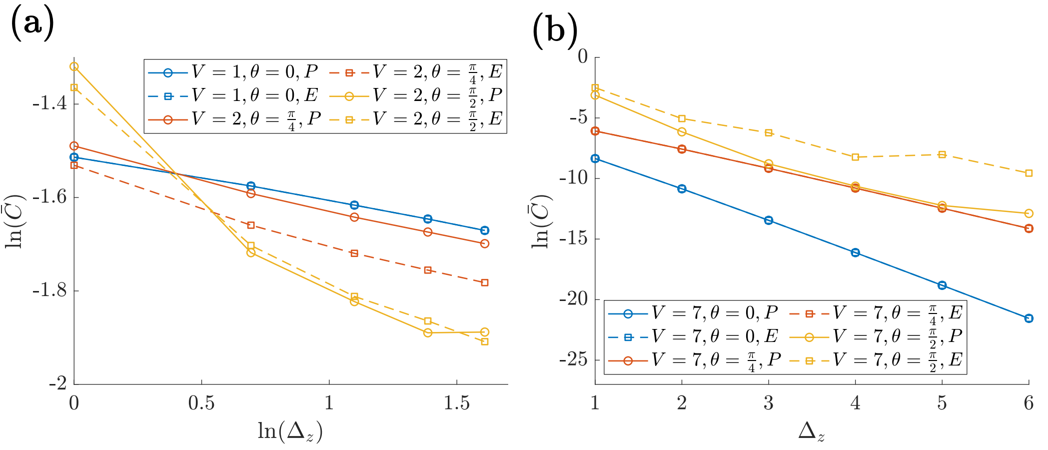

Figure 4.10: Off-diagonal correlations (defined in equation 4.9 as a function of axial separation Δ_z for a polar (P) and equatorial (E) site at $z=4$ for (a) different ground states in the TU region (b) representative SS, PS, and DC states. Hamiltonian parameters for DDI strength V and polarisation angle θ are shown in the legend. Note the very different scales of the y-axes. (a) uses a logarithmic scale for the x axis while (b) uses a linear scale so that the anticipated decay trends (algebraic and exponential respectively) would form straight lines.

where p is a positive power. For solid states, the decay is expected to be exponential, following

$$\langle \hat{b}^{\dagger}_{z,c} \hat{b}_{z+\Delta_z,c} \rangle \propto e^{-\frac{\Delta_z}{\epsilon_z}}, \tag{4.8}$$

where ϵ_z is the correlation length in the axial direction. Specifically, we consider the quantity

$$\bar{C} = \sum_{c}^{L_c} \langle \hat{b}^{\dagger}_{z,c} \hat{b}_{z+\Delta_z,c} \rangle \tag{4.9}$$

which averages correlations over all azimuthal positions c, along with $\bar{C}_{P(E)}$ which are defined analogously for the polar and equatorial sublattices respectively. To reduce edge effects, we again consider the 'bulk' of the system with the three rings at each edge removed and show the correlations for $z=4$ and $1 \leq \Delta_z \leq 6$.

We plot these correlations as a function of axial distance separately for the polar and equatorial sublattices for a range of ground states in Figure 4.10. In subfigure (a), we have used a logarithmic scale for Δ_z so that a polynomial decay of correlations forms a straight trend. While the chosen TU states at $\theta = 0, \frac{\pi}{4}$ broadly follow this pattern (although separately for the polar and equatorial sublattices in the latter case), the correlation decay for the TU ground state at $V = 2, \theta = \frac{\pi}{2}$ deviates from the superfluid trend by displaying faster short-range correlation decay and oscillations at long distances in the polar sublattice.

In subfigure (b) we have used a linear scale for Δ_z, meaning the expected exponential decay of correlations in a solid state would form a straight line. We see that the representative SS and PS states at $\theta = 0, \frac{\pi}{4}$ follow this trend, where the correlations in the PS are identical for the polar and equatorial sublattice, unlike for the TU state at the same polarisation angle. The DC state has a sub-exponential decay of correlations, including oscillations on the equatorial sublattice which are again most notable close to the cylinder edge. These effects reduce but are still noticeable up to $V = 15$.

Overall, we find that the decay of correlations matches well with expectations at low θ but becomes split by sublattice at higher θ. At $\theta = \frac{\pi}{2}$, the ground states of the small system are difficult to categorise by the decay of correlations.

4.5.6 Overlaps with Representative States

As well as the order parameters, a useful way to quantitatively analyse the similarity between ground states is to calculate the overlaps between the ground states and wavefunctions which are clearly representative of one of the ‘idealised’ ground states. For example, we define representative wavefunctions for the TU, SS, PS, and DC states by taking the calculated ground states for the following parameters: $|\psi_{TU}\rangle = |V = 0\rangle$, $|\psi_{SS}\rangle = |V = 7, \theta = 0\rangle$, $|\psi_{PS}\rangle = |V = 7, \theta = \frac{\pi}{4}\rangle$, $|\psi_{DC}\rangle = |V = 7, \theta = \frac{\pi}{2}\rangle$. Using calculated ground states as reference wavefunctions ensures that the wavefunctions incorporate the cylinder edge effects, which would otherwise significantly affect wavefunction overlap.

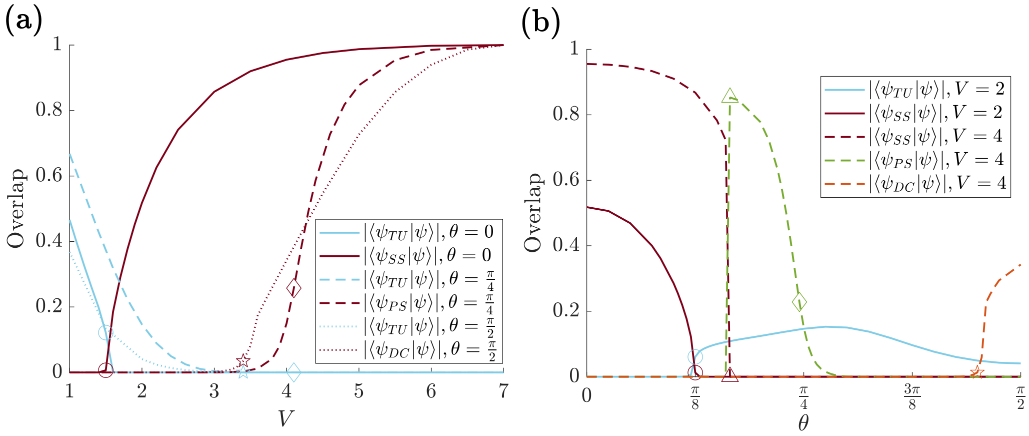

Figure 4.11: Non-negligible overlaps between ground states and reference states (a) as function of DDI strength V (b) as a function of polarisation angle θ. Transition lines from Figure 4.2(a) which are intercepted by these graphs are marked with the corresponding marker.

In Figure 4.11 we plot the overlaps of the calculated ground states $|\psi\rangle$ with these reference states as a function of V and θ. We have only shown the non-negligible overlaps for clarity. To account for the two-fold degeneracy of the SS and DC, we used the initial DMRG state to target both minimally-entangled ground states in separate runs. For example, this means we have two separate $|\psi_{SS}\rangle$ wavefunctions, where the value of $\langle\psi_{SS}|\psi\rangle$ is taken to be the largest value from either choice of $|\psi_{SS}\rangle$.

While the behaviour of these graphs is mostly as expected, they show that the overlap with the reference TU state is quite strongly reduced by any significant value of V, even before solid order develops. As shown in the $V = 4$ line of subfigure (b), there is a large drop in the overlap between the ground state and the reference SS state as the SS-PS transition is approached, which is caused by edge tunnelling.

4.6 Finite-Size Scaling

In this section, we study the ground state physics of the octagonal-ring cylinder for different L_z and perform finite-size scaling analysis. The purpose of finite-size scaling is to use results from finite systems, such as the relatively small systems available to DMRG for quasi-1D lattices, to quantitatively extract the physical behaviour in the thermodynamic limit, which in this case is the one-dimensional thermodynamic limit of $L_c = 8$ and $L_z \to \infty$.

For numerical feasibility, we focus on three values of $\theta = 0, \frac{\pi}{4}, \frac{\pi}{2}$ and analyse the three qualitatively different transitions to solid states. In this section we allow ourselves to use larger system sizes for the numerically easier parts of the phase diagram to improve the reliability of our scaling analysis. We also find that we require smaller increments of interaction strength $\delta_V < 0.1$ for this purpose.

4.6.1 TU-SS

To study the TU-SS transition, we fix $\theta = 0$. To identify the transition point V_{TU-SS}, we use the stripe 'magnetisation' operator $\hat{m}_{SS} = \frac{2}{L_z L_c} \sum_{z,c} (-1)^c (\hat{n}_{z,c} - \frac{1}{2})$ which takes the values ±1 in the SS state depending on which of the two degenerate ground states is present. From this we use the Binder cumulant [202] $U_{SS} = \frac{1}{2}(3 - \frac{\langle \hat{m}_{SS}^4 \rangle}{\langle \hat{m}_{SS}^2 \rangle^2})$, which is zero in the ideal TU state and 1 in the ideal SS state. As shown in Figure 4.12(a) and (c), we find that both U_{SS} and $\langle \hat{m}_{SS}^2 \rangle$ switch from ≈ 0 to ≈ 1 more sharply for larger systems as expected. The crossing points of the Binder cumulants as a function of V converge quickly with increasing system size to $V_{TU-SS} \approx 1.46$ which is close to the corresponding transition point found in previous calculations of analogous models on two-dimensional flat square lattice models [166, 167], although in that case the transition is first-order.

For a second-order phase transition, it is expected that the Binder cumulants for varying L_z will collapse onto each other when plotted against the rescaled interaction strength $L_z^{(\frac{1}{\nu})}(V - V_{TU-SS})$, where ν is the critical exponent for the correlation length. This is shown in Figure 4.12(b) for $\nu = 0.76$. Similar collapse is observed when plotting $L_z^{(\frac{2\beta}{\nu})} \langle m_{SS}^2 \rangle$ against the rescaled interaction strength, where β is the critical

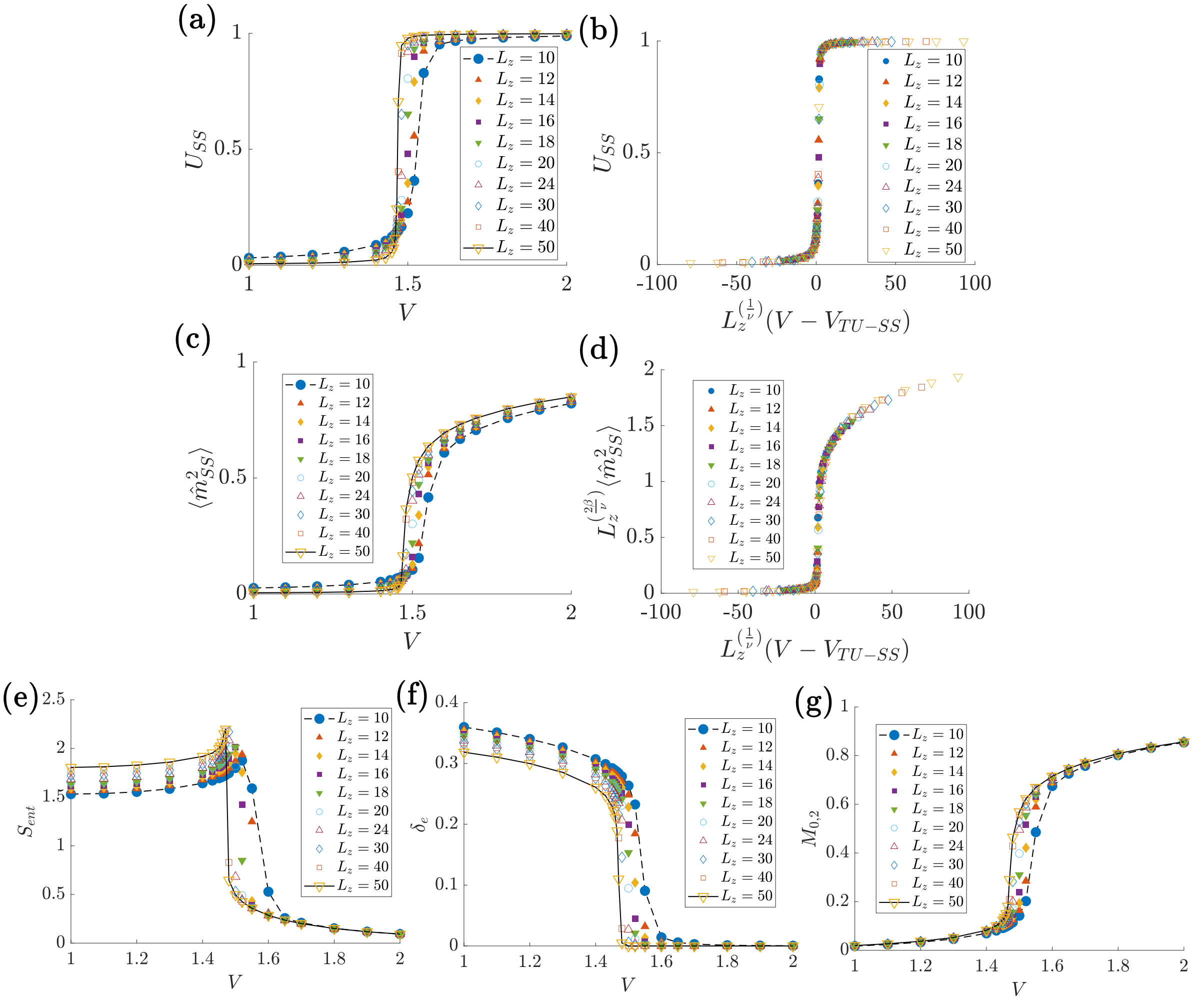

Figure 4.12: Finite-size scaling data for TU-SS transition at polarisation angle $\theta = 0$. (a) Binder cumulant U_{SS} for different axial lengths L_z. As with other panels which do not show data collapse, we link data points for the smallest system $L_z = 10$ in dashed lines and for the largest system $L_z = 50$ in solid lines for clarity. (b) Binder cumulant plotted against scaled interaction strength using critical DDI strength $V_{TU-SS} = 1.46$ and correlation length critical exponent $\nu = 0.76$. (c) Squared magnetisation $\langle m_{SS}^2 \rangle$ against interaction strength. (d) Rescaled squared magnetisation against rescaled interaction strength using magnetisation critical exponent $\beta = 0.08$. (e-g) Physical parameters for the Von-Neumann entanglement entropy S_{ent}, condensation δ_e and stripe solid order $M_{0.2}$ against interaction strength respectively.

exponent for the magnetisation, as shown in Figure 4.12(d) for $\beta = 0.08$, where we used the measure in Ref [203] to quantitatively determine the values of ν and β which showed the smallest deviations from data collapse, with changes of ±0.03 and ±0.005 respectively in these critical exponents leading to a 1% increase in the deviation. The data collapse indicates that the abrupt changes in the behaviour of the small system correspond to a phase transition in the thermodynamic limit. We were not able to convincingly identify this transition with a universality class, with the closest match for the dimensionality (1D quantum), broken symmetry ($\mathbb{Z}2$), and critical exponents being the (classical) 2D Ising transition, for which the expected critical exponents are $\nu = 1$ and $\beta = \frac{1}{8}$ [204]. One possibility for future work to clarify the classification of this transition would be to compare critical exponents with those of the corresponding transition for different fixed even $L_c \neq 8$.

We find that the physical parameters (S_{ent}, δ_e, and $M_{0,2}$) used to identify this transition for the small system show sharper variations in larger systems, which gradually move closer to the critical point V_{TU-SS} for the infinite system. This expected behaviour justifies their use as signatures of the finite-size precursor to the phase transition in the thermodynamic limit.

4.6.2 TU-PS

To study the TU-PS transition in region II, we fix $\theta = \frac{\pi}{4}$. Using larger system lengths allows us to resolve the fact that this transition actually occurs in two distinct parts at a lower value of $V \approx 3.5$ and an upper value of $V \approx 4$. This can be seen clearly in Figure 4.13(a), where S_{ent} decreases smoothly at $V \approx 3.5$ before plateauing and then decreasing sharply at $V \approx 4$ for the largest system sizes. While the most notable change in S_{ent} is at the upper transition, the most significant decrease in δ_e is at the lower transition (Figure 4.13(b)). Meanwhile, the $M_{0,4}$ parameter which denotes the PS order increases

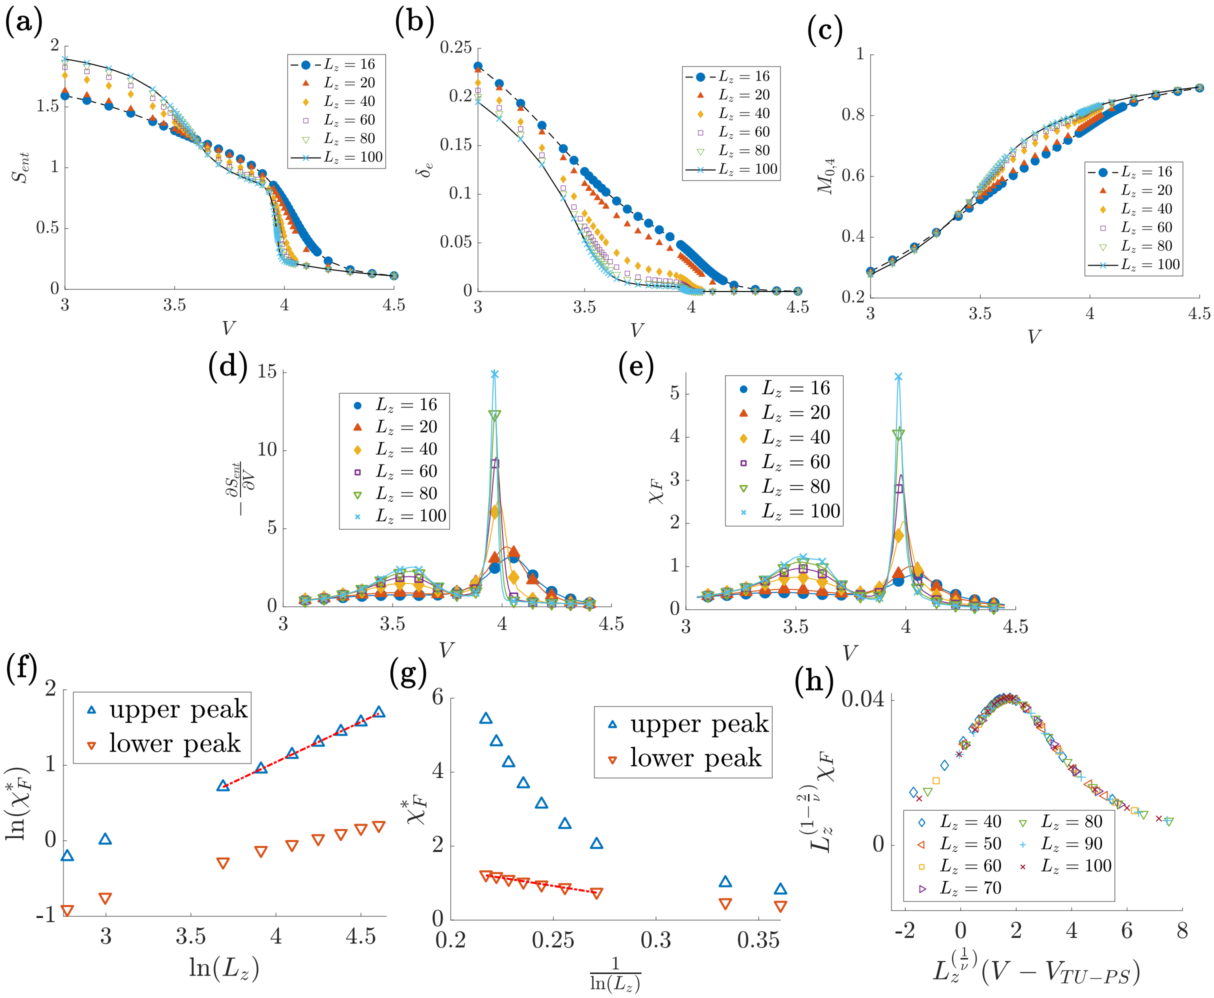

Figure 4.13: Finite-size scaling data for TU-PS transition at polarisation angle $\theta = \frac{\pi}{4}$ as a function of DDI strength V for different axial lengths L_z. (a)-(c) Physical parameters for Von-Neumann entanglement entropy S_{ent}, condensation δ_e, and polar solid order $M_{0.4}$ plotted against interaction strength respectively, with dashed and solid lines shown between datapoints for the smallest and largest system sizes respectively for clarity. (d)-(e) $\frac{\partial S_{ent}}{\partial V}$ and fidelity susceptibility per ring χ_F against interaction strength respectively, where straight lines are plotted between datapoints for all system sizes due to the sharp trends. (f) $\ln(\chi_F^*)$, where χ_F^* is the peak fidelity susceptibility per ring, against $\ln(L_z)$ for both peaks. The dashed red line is a linear fit for the largest seven system sizes of the upper peak, from which $\nu = 0.97$ was extracted. (g) χ_F^* against $\frac{1}{\ln(L_z)}$ for both peaks. The dashed red line is a linear fit for the largest seven system sizes of the lower peak. (h) Data collapse for rescaled fidelity susceptibility $L_z^{(1-\frac{2}{\nu})}\chi_F$ against rescaled interaction strength $L_z^{(\frac{1}{\nu})}(V - V_{TU-PS})$ for the upper peak, using correlation length critical exponent $\nu = 0.97$ and critical DDI strength $V_{TU-PS} = 3.953$.

smoothly and monotonically through both transitions, as shown in Figure 4.13(c).

To gain further insight into these transitions, we consider quantitative measures of sharp changes in the wavefunction. In section 4.5.4, we identified a single nominal transition point by a peak in $-\frac{\partial S_{ent}}{\partial V}$ and a local minimum in the overlap between neighbouring wavefunctions at $V = 4.1$. To quantify the overlap between neighbouring wavefunctions, we use a finite-difference analogue of the fidelity susceptibility per ring [205],

$$\chi_F = \frac{-2\ln(\langle\psi(V - \frac{\delta_V}{2}, \theta)|\psi(V + \frac{\delta_V}{2}, \theta)\rangle)}{L_z \delta_V^2}, \tag{4.10}$$

where δ_V is the difference in V between the Hamiltonians from which the ground-state wavefunctions were calculated. As L_z increases, a second peak in $-\frac{\partial S_{ent}}{\partial V}$ and χ_F emerges at $V \approx 3.5$ and both peaks become sharper as L_z increases for both measures, as shown in Figure 4.13(d) and (e). This supports the conclusion from the behaviour of S_{ent} that there are two separate transitions.

The behaviour of $-\frac{\partial S_{ent}}{\partial V}$ and χ_F across these transitions resemble each other very closely. The fact that they both show signatures of these transitions is expected as both measures are used to identify phase transitions, although the resemblance is not expected to be exact as slight differences between the peak positions of these two measures are visible in Figure 4.13(d) and (e) and have been found in literature [206]. The agreement shows that S_{ent} is an important physical parameter for both transitions, particularly the upper transition where the condensation parameter δ_e and solid order measures are largely unchanged.

To investigate the critical parameters of these transitions, we study the scaling of the peak value of χ_F, which we label χ_F^*. We find that χ_F^* for the upper peak is proportional to L_z^α as derived for second-order transitions with correlation length critical exponent $\nu = \frac{2}{\alpha+1}$ in one dimension [207, 208]. From this fit for $L_z = \{40, 50, 60, 70, 80, 90, 100\}$,

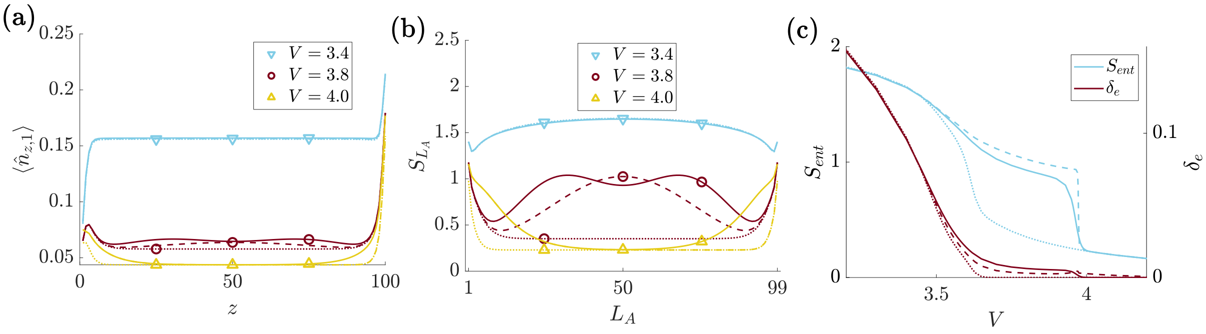

Figure 4.14: TU-PS transition for axial length $L_z = 100$ using different particle numbers $N = N_H$ (solid lines), $N = N_H - 1$ (dashed lines), and $N = N_H - 2$ (dotted lines), where $N_H = \frac{L_z L_c}{2}$ and $L_c = 8$ is the azimuthal length. (a) on-site density for equatorial sites with $c = 1$, where one marker per line is used for identification. (b) Von-Neumann entanglement entropy S_{L_A} for different subsystem lengths $1 \leq L_A \leq 99$. (c) Von-Neumann entanglement entropy across the central bipartition S_{ent} and condensation parameter δ_e as a function of DDI strength V.

we extract $\nu = 0.97$ as shown in Figure 4.13(f). We then find good data collapse for $L_z^{(1-\frac{2}{\nu})}\chi_F$ against the rescaled interaction strength $L_z^{(\frac{1}{\nu})}(V - V_{TU-PS})$ as shown in Figure 4.13(h), where we use $V_{TU-PS} = 3.953$ to obtain the best fit. The quality of the scaling collapse is not itself proof that the upper peak corresponds to a bulk second-order transition, and equivalent data collapse has previously been used to extract ν for different classes of transition, such as topological phase transitions [209]. For the peak at lower V, we find χ_F^* is approximately negatively proportional to $\frac{1}{\ln(L_z)}$ for the largest system sizes as shown in Figure 4.13(g), which is often observed numerically for Berezinkskii-Kosterlitz-Thouless type transitions [206, 210], although significant finite-size corrections mean even longer systems would be required to confirm this relationship and estimate the critical interaction strength.

Having established the presence of these two transitions for $N = N_H$, we discuss the physics of the states below both transitions, between the transitions, and above the transitions. We find it is useful to compare the ground states for three different particle numbers $N = N_H$, $N_H - 1$, and $N_H - 2$ at fixed $L_z = 100$ for this purpose.

We first consider the on-site density on the equatorial sites, showing data for $c = 1$ in Figure 4.14(a) (due to previously-mentioned edge effects for the PS state when affected by finite tunnelling, there is a general trend for higher occupation sites with $c = 1, 8$ at $z = L_z$ than $z = 1$). The equatorial density is uniform in the bulk below both transitions, features oscillations as a function of z between the transitions, and becomes confined to the cylinder edges above the transitions. The exact particle number has little effect on this density below both transitions, while between them extra particles added onto $N = N_H - 2$ add extra peaks in the density distribution, while above both transitions these extra particles become localised almost exclusively at the cylinder edges. (For $N = N_H - 1$, there are two degenerate states with the extra particle swapped between the two cylinder edges.)

Comparison of the entanglement entropy for different bipartitions of the lattice (Figure 4.14(b)) helps to further classify the states. Below both transitions, S_{L_A} grows continuously as the bipartition moves towards the centre of the system, as expected for a superfluid state. More precisely, it fits well to the relation $S_{L_A} = \frac{c_C}{6} \ln(\frac{2L_z}{\pi} \sin(\frac{\pi L_A}{L_z}))$ up to constants, which has been previously found for superfluid states of Bose-Hubbard models in one-dimensional chains and ladders with open boundary conditions with the value $c_C = 1$ [191, 211]. (We numerically find comparable values of c_C between 0.97 and 1.07, although a quantitative extraction of c_C is strongly affected by finite-χ effects (see appendix G)). Between the transitions, S_{L_A} oscillates (except for $N = N_H - 2$), mirroring the equatorial density distribution. Above both transitions (or above the lower transition for $N = N_H - 2$), S_{L_A} saturates in the centre of the system, which is expected for insulating states [212].

The distribution of equatorial density and entanglement entropy along the cylinder therefore suggests that the presence of the upper transition (and by extension the anomalous region between the two transitions) is linked to the cylinder edges. This behaviour is also displayed by the physical parameters S_{ent} and δ_e in Figure 4.14(c), which show that S_{ent} in particular features two distinct reductions for particle numbers $N = N_H$ and $N_H - 1$ (for which the equatorial density becomes localised at the upper transition) but not for $N = N_H - 2$. Based on this analysis, and the lack of a clear symmetry-breaking order parameter or diverging correlation length, the upper transition is not well-described by a bulk second-order transition despite the characteristic scaling of the fidelity susceptibility and is better understood by the interaction of the equatorial sublattice and the cylinder edges. A full characterisation of the nature of this transition would require further research.

Overall, we find that the TU-PS transition splits into two distinct peaks in both $-\frac{\partial S_{ent}}{\partial V}$ and χ_F at large L_z with a narrow intermediate region. The upper peak appears to be strongly influenced by the cylinder edges, so it is unclear whether the upper peak or the intermediate region would exist for periodic boundary conditions in the axial direction. We note that while the low entanglement in this region enables the DMRG calculations to reach larger system lengths $L_z \leq 100$, such system lengths may be more experimentally challenging due to inhomogeneity of the lattice potential in the axial direction [181].

4.6.3 TU-DC

To study the TU-DC transition in region III, we fix $\theta = \frac{\pi}{2}$. As expected, the peak in S_{ent}, decrease in δ_e and increase in the magnitude of $M_{0,4}$ all become sharper as the system length increases, as shown in Figure 4.15(a-c). Unusually, we find that δ_e increases with system length for these small systems despite vanishing in the one-dimensional thermodynamic limit, which appears to be due to the fact the largest bosonic orbital is predominantly supported by the bulk of the system (which forms a larger proportion of larger systems) on both sides of the transition.

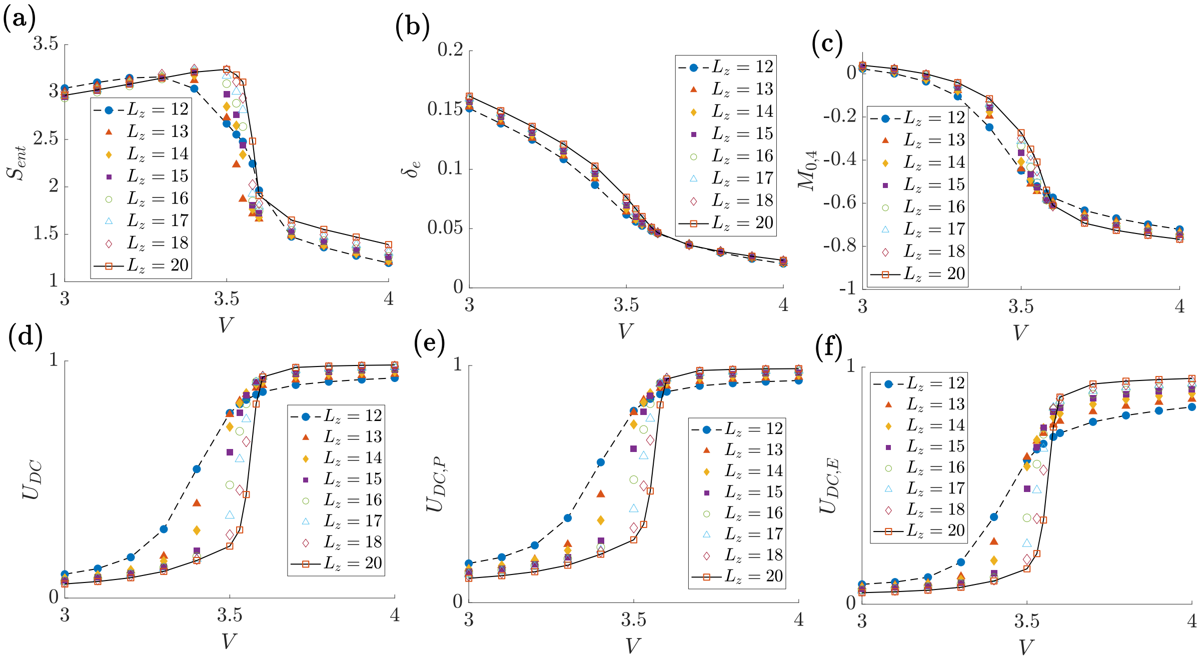

Figure 4.15: Data for different cylinder lengths for TU-DC transition at polarisation angle $\theta = \frac{\pi}{2}$. (a) - (c) Physical parameters for the Von-Neumann entanglement entropy S_{ent}, condensation δ_e, and dual checkerboard solid order $M_{0,4}$. (d) Binder cumulant U_{DC} against interaction strength V. (e)-(f) Sublattice-resolved Binder cumulants. In all panels, neighbouring datapoints for the smallest and largest systems are joined by dashed and solid lines respectively.

As the DC state spontaneously breaks a two-fold symmetry, we define the dual-checkerboard magnetisation analogously to the stripe-solid magnetisation, using the formula $\hat{m}_{DC} = \frac{2}{L_z L_c} \sum_{z,c} (-1)^{z+\lfloor \frac{c-1}{4} \rfloor} (\hat{n}_{z,c} - \frac{1}{2})$. We also define polar and equatorial sublattice versions of the dual-checkerboard magnetisation by including only the terms on the relevant sublattice. We then define the Binder cumulant U_{DC} (and its sublattice-resolved versions $U_{DC,P}$ and $U_{DC,E}$) from these magnetisation operators analogously to the TU-SS transition.

We find that U_{DC} also increases across the transition more sharply for larger L_z as expected, as is shown in Figure 4.15(d), as do its polar and equatorial sublattice counterparts shown in Figure 4.15(e) and (f). This suggests that this transition would become sharp in the thermodynamic limit, however the large entanglement in this region greatly reduces the available system size compared to Regions I and II and means that confirming this would require larger-scale numerical calculations. Additionally, for the available system sizes, large edge effects, particularly on the TU side of the transition, distort the magnetisation. In particular, the values of V for which the U_{DC} curves for successive system lengths cross increase significantly with L_z for even the largest systems we studied, which prevents the identification of a critical point or critical exponents through data collapse.

4.7 Discussion

The octagonal-ring cylindrical lattice leads to more varied ground states than either the flat square lattice or the single octagonal ring because it combines the sublattice structure introduced by the curved lattice surface with translationally-invariant interactions in the axial direction. In region I, the ground state is relatively similar to the flat square lattice model because the spatial variation of interactions is weak. In region II, the attractive axial interactions cooperate with the sublattice ordering which favours the polar sublattice, producing the PS ground state which is a two-dimensional analogue of the equivalent density wave predicted for the single octagonal ring [180]. In region III, the competition of the (repulsive) 'flat' translationally-invariant axial interactions and the 'curved' spatially-varying axial interactions produces qualitatively different physics than either mechanism alone, leading to greater entanglement and importance of next-nearest-neighbour interactions.

The transitions between ground states are also influenced by the combination of these two mechanisms. The TU-SS transition in region I is sharp, even in a small system, due to the discrete symmetry-breaking property of the SS while the TU-PS transition in region II is gradual due to the spatially-varying interactions which favour the polar sublattice alone. It requires relatively large system lengths $L_z \approx 40$ to resolve the two qualitatively separate parts of the transition. In region III, the TU-DC transition is sharper on the polar sublattice than the equatorial sublattice in finite-size systems due to the spatial variation of interactions and the fact that the repulsive axial interactions cause both sublattices to be occupied, although we were not able to verify whether this behaviour persists into the thermodynamic limit. Spatially-resolved transitions in extended Bose-Hubbard models have been predicted due to externally-imposed lattice inhomogeneities, such as additional finite trapping potentials [213] or variation in the number of neighbouring sites for quasicrystal lattices [214], but in this case the inhomogeneity is due to the inter-particle interactions.

Our finite-size scaling analysis suggests that small axial lengths L_z are able in most cases to demonstrate the unusual many-body physics in the bulk of the system and identify finite-size precursors to the phase transitions, which may be important for experimental implementation. The most significant edge effects are in region III as particles preferentially occupy the cylinder edges, although they are also important for part of the TU-PS transition. The DS state in particular depends on L_z due to competition of different block lengths

4.8 Variations of L_c

4.8.1 $L_c \neq 8$

While we have focussed mainly on an octagonal ring cylinder for concreteness, similar models exist for different L_c. We have performed brief calculations for $4 \leq L_c \leq 10$. For even L_c, the SS at low θ is largely unaffected by changing L_c while for odd L_c the density wave periodicity is disrupted by the cylindrical boundary conditions [3]. The

---

[3] The experimental proposal in Ref [181] only includes even L_c.

existence of analogues of the PS and DC states depends on the interplay between the sublattice structure and the particle filling fraction.

4.8.2 $L_c \to \infty$

For low values of L_c (which are the only values for which DMRG is feasible), the exact integer value of L_c greatly influences the physics. The opposite limit of large L_c instead features only small variations in interactions between neighbouring sites around the ring, meaning that it provides a closer connection with square lattice Bose-Hubbard models. To discuss this two-dimensional limit, we consider the case where r_c remains equal to r_z but $L_c \to \infty$ (which means that the radius of the cylinder also $\to \infty$).

In this case, the discrete site index c can be replaced by the continuous azimuthal angle $0 \leq \phi_c < 2\pi$ (where the azimuthal angle is defined as in Figure 4.1(b)). The polarisation angle ϕ now merely translates the Hamiltonian around the cylinder ring and we set $\phi = 0$. Because the curvature of the lattice is very gradual, the interactions between nearby sites are very similar to those of a flat square lattice with lattice vectors $\hat{\boldsymbol{z}}$ and $\hat{\boldsymbol{\phi}}$ (we denote the perpendicular vector, which is in the radial direction, by $\hat{\boldsymbol{R}}$). The fixed-space polarisation direction (defined by θ) can be converted into these local lattice vectors using the following formula:

$$\begin{pmatrix} d_{\hat{\boldsymbol{z}}} \\ d_{\hat{\boldsymbol{\phi}}} \\ d_{\hat{\boldsymbol{R}}} \end{pmatrix} = |d| \begin{pmatrix} \cos(\theta) \\ -\sin(\theta)\sin(\phi_c) \\ \sin(\theta)\cos(\phi_c) \end{pmatrix} \tag{4.11}$$

The effective local polarisation, which maps the nearby Hamiltonian to a flat square lattice Bose-Hubbard model, therefore depends on ϕ_c if the physical polarisation $\hat{\boldsymbol{d}}$ is not along the cylinder axis. In this case, the cylindrical lattice effectively connects a continuous series of flat square lattice Bose-Hubbard models with slowly-varying effective local polarisation.

For example, if the physical polarisation angle $\theta = \frac{\pi}{2}$, then at the azimuthal positions

$\phi_c = 0, \pi$ the polarisation is in the $\pm\hat{\boldsymbol{R}}$ direction (i.e. perpendicular to the lattice plane). This means the DDI is isotropically repulsive, which has been predicted to realise checkerboard and star solids and checkerboard supersolids [31, 162, 163]. The same physical polarisation maps to the $\pm\hat{\boldsymbol{\phi}}$ direction (i.e. along a lattice vector) at $\phi_c = \frac{\pi}{2}, \frac{3\pi}{2}$. The corresponding square lattice model has been predicted to support stripe density waves [166,167] and, in soft-core models, supersolids with the same density ordering [142, 153, 157]. Intermediate values of ϕ_c correspond to flat lattice models with intermediate tilting of the polarisation. For intermediate values of both θ and ϕ_c, the effective local polarisation can point between the lattice vectors $\hat{\boldsymbol{z}}$ and $\hat{\boldsymbol{\phi}}$ with a component in $\hat{\boldsymbol{R}}$ providing isotropic repulsion, which has been recently predicted to stabilise diagonal solids and supersolids which rely on next-nearest-neighbour interactions [164]. In general, the spatial variation of effective polarisation may allow the study of interfaces between these different density wave and supersolid states.

Certain features of the octagonal-ring cylinder at strong interaction would survive for $L_c \to \infty$ while others would be lost. The stripe solid at $\theta = 0$ would remain unchanged but the sublattice structure would be too complicated to allow a single 'polar solid'. Instead, for $0.31\frac{\pi}{2} < \theta < 0.61\frac{\pi}{2}$, the sites closest to the 'poles' ($\phi_c = \frac{\pi}{2}, \frac{3\pi}{2}$) would be favoured due to the spatial variation of interactions, leading to multiple different solid states interpolating between the SS and the $L_c \to \infty$ equivalent of the PS, where all sites for $\frac{\pi}{4} < \phi_c < \frac{3\pi}{4}$ and $\frac{5\pi}{4} < \phi_c < \frac{7\pi}{4}$ would be occupied (further discussion in appendix H).

4.9 Physical Considerations

In this section, we discuss the main considerations regarding the possible physical implementation of this model. The cylindrical optical lattice itself is discussed in Ref [181], so we will focus on the aspects which are specific to this model.

4.9.1 Hard-Core Constraint

We have assumed that the short-range repulsion between bosons is sufficient to prevent more than one particle occupying a single site. In a non-dipolar Bose-Hubbard model, this requires that the on-site repulsion, usually labelled U (see chapter 3), is much larger than the tunnelling. For a dipolar Bose-Hubbard model which includes attractive dipolar interactions, the short-range repulsion must also be strong enough to prevent these attractive dipolar interactions from favouring double occupation.

To estimate how far the attractive interactions could favour double occupation, we consider the stripe solid, which experiences the most attractive dipolar interactions of all states present. Ignoring edge effects and including interactions up to next-nearest-neighbour, the DDI energy per particle is $-2V$. A doubly-occupied stripe solid would instead have DDI energy of $-4V$ per particle but experience on-site repulsion of $\frac{U}{2}$ per particle. This creates the additional requirement that $U \gg 4V$ to prevent double occupation driven by the DDI.

Alternatively for sufficiently weak tunnelling, the possibility of reactive losses for polar molecules on the same site has been observed to suppress coherent tunnelling to occupied sites [42].

4.9.2 Finite Temperature

Our calculations have assumed that there is no thermal excitation and that the system is in a pure quantum state, which will not be exactly satisfied in an optical lattice experiment. As a rough estimate of the robustness of the ground states to temperature, we expect that the superfluid state will be destroyed if the thermal energy k_BT exceeds the tunnelling energy J. In order to observe the SS, PS, and DC as well, we require $V > 2J$, $4J$, and $4J$ respectively which implies $V > 2k_BT$, $4k_BT$, and $4k_BT$ respectively.

The estimate for the temperature requirement for the stripe solid ($V > 2k_BT$) is close to that produced by explicit finite-temperature Quantum Monte Carlo calculations for the analogous state on a flat square lattice [166].

Using typical DDI energies for ultracold dipolar platforms, we can estimate the threshold temperature below which $V > 4k_BT$ is satisfied. As has been experimentally implemented, magnetic ^{168}Er atoms in an optical lattice with 272nm site spacing have $V \approx h\times 34$Hz, which suggests the need for temperatures below 0.4nK while temperatures down to $\approx$ 70nK have been implemented for magnetic atoms [44]. The DDI energy (and therefore temperature threshold) could be increased by a factor of up to 4 using Feshbach molecules made from two magnetic atoms [215]. Polar molecules, interacting using electric dipole moments, can offer much larger interaction energies $V > h \times 1$kHz for lattice spacings of 532nm, which corresponds to a temperature threshold of $\approx$ 10nK. For comparison, temperatures of 60nK have been recorded in bulk gases of fermionic polar molecules [81] and bosonic Feshbach molecules [216]. Overall, the temperature requirements needed to observe density wave states imposed by the DDI are challenging, but rapid progress has recently been made in this regard. The realisation of DS-type states, which rely on the weaker next-nearest-neighbour interactions and therefore need larger DDI strengths to overcome tunnelling, would be more difficult.

4.9.3 Lattice Filling

An additional experimental consideration, which is especially relevant for polar molecules, is that our calculations have assumed one boson per two lattice sites. Recent experiments have recorded filling up to 30% of sites [76, 77], which means moderate improvements to the filling would need to be made.

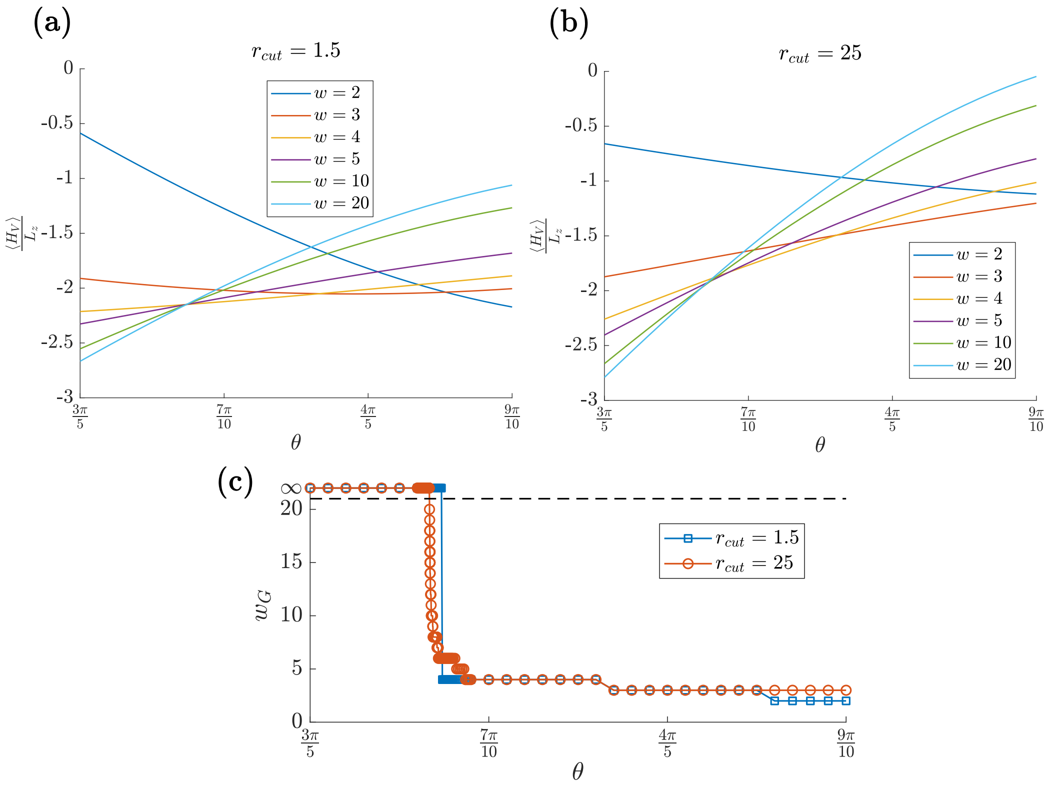

Figure 4.16: (a) DDI energy per cylinder ring for selected periodic DS-type states with periodicity w as a function of polarisation angle θ for interactions truncated at distance $r_{cut} = 1.5$. (b) Equivalent of (a) with $r_{cut} = 25$. (c) Periodicity w_G which has the lowest energy out of DS-type states with $2 \leq w \leq 20$ or the PS. A horizontal dashed line is added to separate DS-type states with a finite w from the PS state with $w = \infty$.

4.9.4 Interaction Truncation

While we truncated the DDI at next-nearest-neighbour for numerical reasons, the physical system would feature dipolar interactions between all pairs of sites. Due to the r^{-3} decay of the DDI, the shorter-range interactions we have kept are the dominant contribution but longer-range interactions could have a small quantitative effect. For example, including interactions up to a distance of 2 lattice units increases the magnitude of the DDI energy of the SS and PS when the axial interactions are attractive while it decreases the magnitude of the DDI energy of the DC, due to the largest additional term being the axial interactions between sites separated by $\pm 2\hat{\boldsymbol{z}}$.

The most obvious effect of the smaller long-range interactions is on the DS-type states because there are multiple competing states close to the ground state energy in that parameter region. One qualitative difference which would be caused by increasing the truncation range is that periodic DS-type states with larger block sizes could become favoured. As discussed in Section 4.5.3.2, for interactions truncated at next-nearest-neighbour, the periodic states in the DS family which have the lowest interaction energy feature axial periodicities of 2 (DC), 3, 4, and ∞ (PS). Increasing the range of interactions allows states with a finite periodicity greater than 4 to become the state with the lowest interaction energy of the DS family [4] (although there is no guarantee that a state from this family will be globally optimal when the interaction range is increased).

This is shown in Figure 4.16, where we numerically compare the DDI energy per cylinder ring $\frac{\langle H_V \rangle}{L_z}$ of DS-type states of different periodicities $2 \leq w \leq 20$ with interactions included up to a distance r_{cut}. In subfigures (a) and (b) we show the DDI energy density of certain periodicities for the nearest-neighbour cutoff we used in our DMRG calculations $r_{cut} = 1.5$ [5] and an extremely long truncation range $r_{cut} = 25$. From these subfigures, it can be seen that a lower value of θ and larger truncation range favours DS-type states with larger periodicities. Subfigure (c) shows the value of w, labelled w_G, for which the energy density is lowest out of all DS-type states with $2 \leq w \leq 20$ or the PS ($w = \infty$) at that value of θ. For interactions truncated to next-nearest-neighbour, w_G can only take the values 2, 3, 4, and ∞. For interactions truncated at extremely long range, each integer value of w offers the lowest energy density for some value of θ. This effect is somewhat reminiscent of the 'devil's staircase' predicted in models with interactions kept to very long range [31, 217–219], except in this case the filling fraction

---

[4] As mentioned in Section 4.5.3.2, blocks of more than six occupied polar sites can already reduce the energy in finite lattices for interactions truncated at next-nearest-neighbour provided carefully chosen θ and L_z and extremely strong DDI.

[5] note that any value of $\sqrt{2} < r_{cut} \leq \sqrt{(2+\sqrt{2})}$ also restricts the DDI to next-nearest-neighbour.

remains constant and only the density periodicity would change with θ. However, these states have a very small energy separation, require a very small window of θ, rely on extremely long-range interactions, and are highly susceptible to edge effects, which would reduce their chances of experimental observation. We also emphasise again that we have only checked that these states have the lowest interaction energy of periodic DS-type states, so they are not necessarily the globally optimal ground state.

4.10 Conclusions

Overall, we have shown that the spatially-dependent interactions caused by the combination of the curved surface of a cylindrical optical lattice with octagonal rings and the anisotropy of the DDI lead to varied ground-state behaviour, showing significant entanglement and importance of next-nearest-neighbour interactions. We have investigated the nature of the ordering phase transitions through finite-size analysis.

The most straightforward future extension of this work would be to study this model using parameters other than the ones we chose to focus on, such as $L_c \neq 8$, $\phi \neq 0$, or relaxing the hard-core constraint. A further direct generalisation would be the investigation of dipolar spin models without particle tunnelling in cylindrical or other curved optical lattices, extending existing work on the dipolar zigzag chain [220].

Another avenue for future work is to consider the interplay of the spatially-varying dipolar interactions with an artificial magnetic field, the implementation of which for a cylindrical lattice was discussed in Ref [181]. The combination of complex tunnelling phases and strong interactions between bosons has generated interest in relation to bosonic lattice analogues and extensions of continuum quantum Hall states [221], and has been shown to extend the parameter regimes of interaction-dominated phases in extended Bose-Hubbard models with repulsive off-site interactions [222]. The cylindrical geometry may be advantageous in this context as it contains both the bulk and edge physics [223].

At a broader level, there has been recent interest in simulating lattice Hamiltonians with spatially-varying model parameters [224] motivated by the unusual properties of twisted bilayer graphene. This type of physics would be difficult to capture using a cylindrical lattice dipolar model because the interactions only vary along one dimension (the azimuthal dimension) and the variation lengthscale is exactly commensurate with L_c, which prevents the formation of large Moiré supercells. However, the general concept of embedding an anisotropically-interacting lower-dimensional system into higher-dimensional space to implement spatially-varying Hamiltonians may be flexibly applied to many-body systems built using a 'bottom-up' approach from reconfigurable arrays of individual tweezers [56, 57, 225, 226].

5

Comparison of Lattice and Continuous-Space Methods for Dipolar Bose-Hubbard Models

This chapter contains results from a collaboration between MH, Paolo Molignini, and Axel U. J. Lode, which has been submitted for publication [227]. MH performed the lattice-model calculations. Paolo Molignini performed the continuum calculations. Both contributed to the project direction and comparison between methods. Discussions with Hongmin Gao are acknowledged. We would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work [228]. Computation time on the ETH Zurich Euler cluster and at the High-Performance Computing Center Stuttgart (HLRS) is acknowledged

The single-band BH model and its extension including the DDI, for example the cylindrical lattice model studied in chapter 4, rely on the assumption that the lowest band is sufficient to describe the physics. This assumption has been shown to be questionable for BH models with moderate short-range interactions [229, 230]. In this chapter, we quantitatively compare ground state calculations of bosons interacting via the DDI in lattice potentials using the BH model and direct continuous space methods, thereby testing the validity of the BH model for ground state calculations with strong off-site interactions. We find that strong but feasible DDI can warp the on-site wavefunction, leading to population of higher bands and reduced energy.

In section 5.1 we briefly discuss the two methods before introducing the numerical calculations in section 5.2 and comparing results in section 5.3.

5.1 Methods

5.1.1 Lattice Methods

To find the ground state of the dipolar interacting bosons using lattice methods we derive the corresponding dipolar BH model as described in Chapter 3. While the dipolar BH model is usually derived using Wannier functions which assume an infinitely-large periodic lattice, our physical scenario of small lattices with very large potential walls requires Wannier functions [1] which account for finite-size effects. The process we used to construct these is described in the supplemental material of Ref [231] with a minor extension required for the second band which we describe in appendix I for reproducibility. This results in site-dependent Wannier functions, leading to site-dependent tunnelling, interactions, and chemical potentials which we incorporate into our lattice model. This effect is larger closer to the edges and for weaker lattice potentials. For the examples studied in this chapter, we found that the Hamiltonian parameters were within 5% of those calculated from infinite-lattice Wannier functions using the methods in Ref [141].

We solve the lattice model using exact diagonalisation (ED) [232] implemented

---

[1] We use the term ‘Wannier functions’ for the localised single-particle basis in the finite lattice because their role in our calculations is the same as genuine Wannier functions for infinite lattices.

in the QuSpin library [233]. The resulting wavefunction is denoted by $|\psi_{\text{ED}_1}\rangle$ for calculations using the lowest band and $|\psi_{\text{ED}_2}\rangle$ for calculations using the lowest two bands. Unlike MPS, this method suffers from the exponential growth of Hilbert space but this is not numerically prohibitive for the system sizes that can be accessed with the continuum methods. Quantitatively, the number of states D_H in the Hilbert space for N bosons in N_w Wannier functions [2] is $\binom{N+N_w-1}{N_w-1}$, although this can be reduced if the number of bosons in a Wannier function is limited to $n_{max} < N$. In that case, the number of states in the Hilbert space is given by

$$D_H = \sum_{q=0}^{\lfloor \frac{N}{n_{max}+1} \rfloor} -1^q \binom{N_w}{q} \binom{N_w + N - q(n_{max}+1) - 1}{N_w - q(n_{max}+1)}, \tag{5.1}$$

as implemented by the QuSpin library. Diagonalising the Hamiltonian to find all of its eigenvalues and eigenvectors has a computational cost scaling as $\mathcal{O}(D_H^3)$, but in practice the extremal eigenvalues and eigenvectors (such as the ground and lowest few excited states) of relatively sparse Hamiltonian matrices converge more quickly using the Lanczos algorithm [232].

5.1.2 Continuous-Space Methods

To solve the same physical problem in continuous space, we use the Multi-Configurational Time-Dependent Hartree method for Bosons (MCTDHB) [234, 235], which we will only briefly introduce here, implemented in the MCTDH-X software package [236]. MCTDHB expresses the wavefunction using M optimisable orthonormal single-particle orbitals. These orbitals are the continuum analogue of the natural orbitals of the one-body-density-matrix discussed in Chapter 4. The orbitals are numerically represented on a fine discrete spatial grid.

The wavefunction is built using a superposition of all normalised particle-exchange-

---

[2] $N_w = L$ for the lowest band and $2L$ for the lowest two bands

symmetric combinations $|\vec{n}\rangle$ of different numbers of bosons in each orbital. These combinations are called permanents and for N bosons in M orbitals, there are $D_C = \frac{(N+M-1)!}{N!(M-1)!}$ permanents, which is equal to the dimension of the lattice Hilbert space if there is no restriction on the particle number per site and $M = N_w$. The value of M is chosen to control the size of the considered Hilbert space. For example, $M = 1$ restricts the state to a pure condensate, while larger M is required to represent strongly-correlated states such as Mott insulators or density waves. The orbitals and the coefficients $c_{|\vec{n}\rangle}$ for each permanent are optimised via imaginary time evolution.

The total wavefunction is

$$|\psi_{\mathrm{C}}\rangle = \sum_{|\vec{n}\rangle} c_{|\vec{n}\rangle} |\vec{n}\rangle \,, \tag{5.2}$$

where $|\psi_{\mathrm{C}}\rangle$ stands for 'continuum'. If the chosen number of single-particle orbitals M is equal to the number of Wannier functions in the BH model, this equation has a similar form to the wavefunction of the BH model, where the single-particle orbitals correspond to the Wannier functions and the permanents correspond to states with an integer occupation of each Wannier function [3].

We use this similarity to project the wavefunction produced by MCTDH-X into the Hilbert space of the lowest band (for $M = L$) or lowest two bands (for $M = 2L$) of the BH model, producing the unnormalised wavefunction $|\psi_{\mathrm{C}\to\mathrm{ED}_1}\rangle$ or $|\psi_{\mathrm{C}\to\mathrm{ED}_2}\rangle$ respectively. This projection is calculated by integrating the overlaps of the lattice model's Wannier functions and the continuum model's natural orbitals in real space. The projection magnitude into the lowest band is calculated as $P_1 = \langle\psi_{\mathrm{C}\to\mathrm{ED}_1} \,|\, \psi_{\mathrm{C}\to\mathrm{ED}_1}\rangle$ while the projection to the lowest two bands is calculated as $P_2 = \langle\psi_{\mathrm{C}\to\mathrm{ED}_2} \,|\, \psi_{\mathrm{C}\to\mathrm{ED}_2}\rangle \geq P_1$. The fidelity between the single-band lattice and continuum methods is $f_1 =$

---

[3] This correspondence is only with respect to the form of the equation. Even in cases where the lattice and MCTDH-X wavefunctions are physically almost identical, the single particle orbitals produced by MCTDH-X do not generally resemble the Wannier functions.

$|\langle\psi_{\mathrm{C}\to\mathrm{ED}_1}\,|\,\psi_{\mathrm{ED}_1}\rangle|^2 \leq P_1$ while the fidelity between the two-band lattice and continuum methods is $f_2 = |\langle\psi_{\mathrm{C}\to\mathrm{ED}_2}\,|\,\psi_{\mathrm{ED}_2}\rangle|^2 \leq P_2$.

The computational cost of each imaginary timestep using MCTDHB scales proportionally to $M^4(n_{grid}^2 + D_C)$, where n_{grid} is the number of spatial gridpoints and D_C is the number of many-body permanents, because of the spatial integration needed to calculate the two-body interaction Hamiltonian matrix elements in the basis of the natural orbitals [237]. Note that D_C itself depends on M, such that for large M when the D_C term dominates the computational cost, the cost would scale as $M^4\frac{(N+M-1)!}{N!(M-1)!}$. The expensive computational scaling with respect to the number of orbitals M means that MCTDH-X is limited to small systems like ED if a large number of orbitals is required to represent the physics accurately, although the method would be well-suited for systems where the number of orbitals required is much smaller than the number of bosons [4]. The cost of calculating $\psi_{\mathrm{C}\to\mathrm{ED}_1}$ and $\psi_{\mathrm{C}\to\mathrm{ED}_2}$ to compare the continuum and lattice methods scales proportionally to D_C^2 for $M = N_w$.

5.2 Numerical Calculations

We study the validity of the lowest-band dipolar BH model in a small one-dimensional lattice of $L = 9$ sites with repulsive dipolar interactions. Previous 1D continuum studies of strongly-repulsive dipolar bosons at filling fractions above one particle per site have shown 'splitting' of density within sites to form 'crystals' with smaller particle separation than the lattice spacing [238, 239]. We use filling fractions below one particle per site to focus on the warping of the on-site wavefunction through off-site interactions

---

[4] If the relationship between the number of particles N and the required number of orbitals M is known for a given type of physical system, a simplified form for the scaling $M^4\frac{(N+M-1)!}{N!(M-1)!}$ can be estimated using the approximation $\ln(n!) \approx n\ln(n) - n$ for large $n \in \mathbb{Z}$. For example, in the case of $M = N$, as is needed to represent a Mott insulator with unit filling, the scaling becomes $M^4\frac{2^{2M}}{\sqrt{M}}$, which involves an exponential scaling with system size.

and quantitatively compare this with the corresponding BH model.

We set our unit of length to the distance between neighbouring sites. The continuum calculations use a finite spatial grid of 2048 segments. The lattice potential is defined between $-4.5 \leq x \leq 4.5$ and is bordered by an extremely large potential wall. The spatial domain of the MCTDH-X calculations extends to $-16 \leq x \leq 16$ to make the effect of periodic boundary conditions on the DDI negligible. This means there are 64 spatial segments between neighbouring lattice sites.

We assume a strong transverse harmonic confinement, which regularises the short-range divergence of the DDI [238, 240]. To include this regularisation, we use the DDI potential

$$H_V = \frac{V}{R^3 + \alpha}, \tag{5.3}$$

where $\alpha = 0.05$ in units of the lattice separation cubed.

To reduce the population of multiply-occupied sites within the lowest band to $< 1\%$, we consider a very strong contact interaction represented as a narrow Gaussian in the continuum calculations, although this term becomes irrelevant as the repulsive DDI performs this role more strongly for $V \geq 0.5E_R$. This interaction has the spatial dependence

$$H_G = \frac{V_G}{\sqrt{2\pi\sigma^2}} e^{-\frac{R^2}{2\sigma^2}}, \tag{5.4}$$

where $V_G = 2.5E_R$, R is the particle separation in lattice units and $\sigma = 0.05$ in the same units. The corresponding lattice model parameters are quantitatively similar to those of a true delta function of the same magnitude due to the narrow width.

We consider a 1D optical lattice potential defined by $A(x) = A_0(\sin^2(\pi x))$ with weak to moderate optical lattices of potential depth A_0 from $5E_R$ to $10E_R$. Examples of the corresponding Wannier functions and the most important resulting parameters of the dipolar BH model are shown in Figure 5.1.

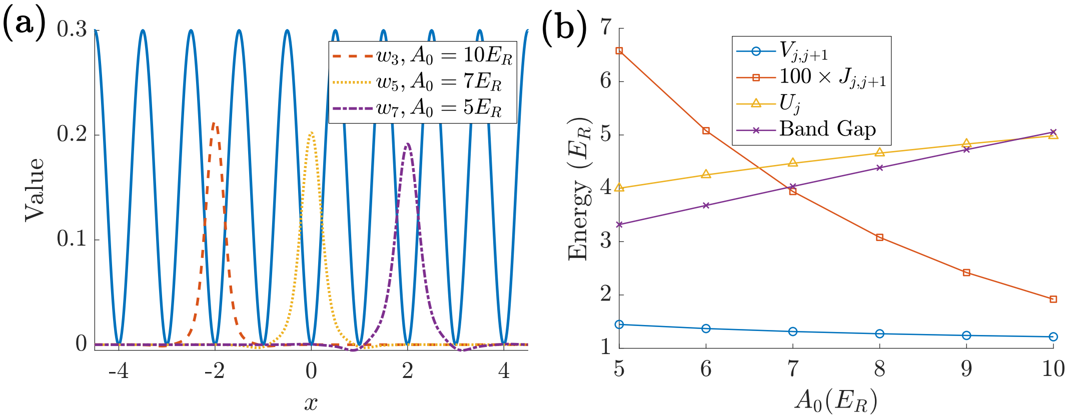

Figure 5.1: (a) Lowest-band Wannier functions for three different optical lattice sites w_i (indexed according to position) at three different potential depths. A schematic of an optical lattice is superimposed for visualisation, although note that the potential depth corresponding to the three shown Wannier functions is different. (b) Bose-Hubbard parameters for centremost sites of the 1D lattice potentials. $V_{j,j+1}$ denotes the nearest-neighbour DDI for DDI strength $V = 1E_R$, while $J_{j,j+1}$ is the nearest-neighbour tunnelling, and U_j is the on-site repulsion due to the contact interaction. The band gap, shown for comparison, is the energy between the lowest and first-excited bands of the corresponding infinite lattice, calculated using the methods in Ref [141].

5.3 Results

We first focus on the simple case of $N = 5$ bosons. The filling fraction here is chosen such that the ground state is a superfluid (SF) at weak DDI and a density wave (DW) [5] $|101010101\rangle$ with a large energy gap at strong DDI. We compare this to the case of $N = 4$, which has a lower filling fraction that is not commensurate with the density wave order. This leads to a much smaller gap between the lowest energy states of the lowest-band dipolar BH model at strong DDI. In both cases $M = 9$ allowed convergence in the MCTDH-X calculations.

While it was sufficient for quantitative agreement for $A_0 = 10E_R$ to include tunnelling to nearest-neighbours and off-site density-density interactions due to the DDI between all pairs of sites while assuming a hard-core constraint, for $A_0 = 5E_R$,

---

[5] We use the term density wave to reference the periodic density modulation, although it does not spontaneously break a symmetry.

the greater spatial overlap between Wannier functions causes more lattice terms to become important, while the reduced on-site repulsion meant it was important to allow up to two bosons per site. In the numerical results for $N = 4$ and $N = 5$, we included tunnelling up to next-nearest-neighbour and density-density DDI terms between all pairs of sites. We also included all DDI and contact interaction terms terms $\hat{b}_i^\dagger \hat{b}_j^\dagger \hat{b}_k \hat{b}_l$ where i, j, k, and l are contained within three consecutive sites. The largest of these extra terms were the finite on-site repulsion due to the contact interaction and DDI, interaction-induced soft-core tunnelling terms of the form $\hat{b}_j^\dagger \hat{n}_j \hat{b}_k$, where j and k are neighbouring sites, and DDI-induced tunnelling terms of the form $\hat{n}_j \hat{b}_k^\dagger \hat{b}_l$, where j, k, and l are neighbouring but different sites.

Figure 5.2(a) compares the energies of the ground states in the two methods. For small DDI, there is good quantitative agreement for $A_0 = 10E_R$, which becomes slightly reduced for weak lattice potentials where the lattice model description is less appropriate. There are larger discrepancies between methods for strong DDI, where the continuum calculations are able to reduce the energy below that of the lattice model, an effect which increases with weaker lattice potential and increased particle density as expected. For $N = 5$, $A = 5E_R$, $V = 10.1E_R$, the difference in energy between methods is $\approx 0.2E_R$, which is $\approx 1\%$ of the total interaction, kinetic, and potential energy.

Figure 5.2(b) shows the projection magnitude of the MCTDH-X ground state onto the lowest band and onto the corresponding ED ground state itself. This shows quantitatively that the reduced energy of the continuum method compared to the lattice method at strong DDI coincides with a small population of higher bands. For the weakest lattice potential $A_0 = 5E_R$ and $N = 5$, there is a small reduction in P_1 and an increase in the energy gap between methods for $V \approx 0.2E_R$ where the larger tunnelling at the weak lattice potential means there is still a small occupation in the nominally-empty sites of the DW. This effect reduces as the DW becomes fully

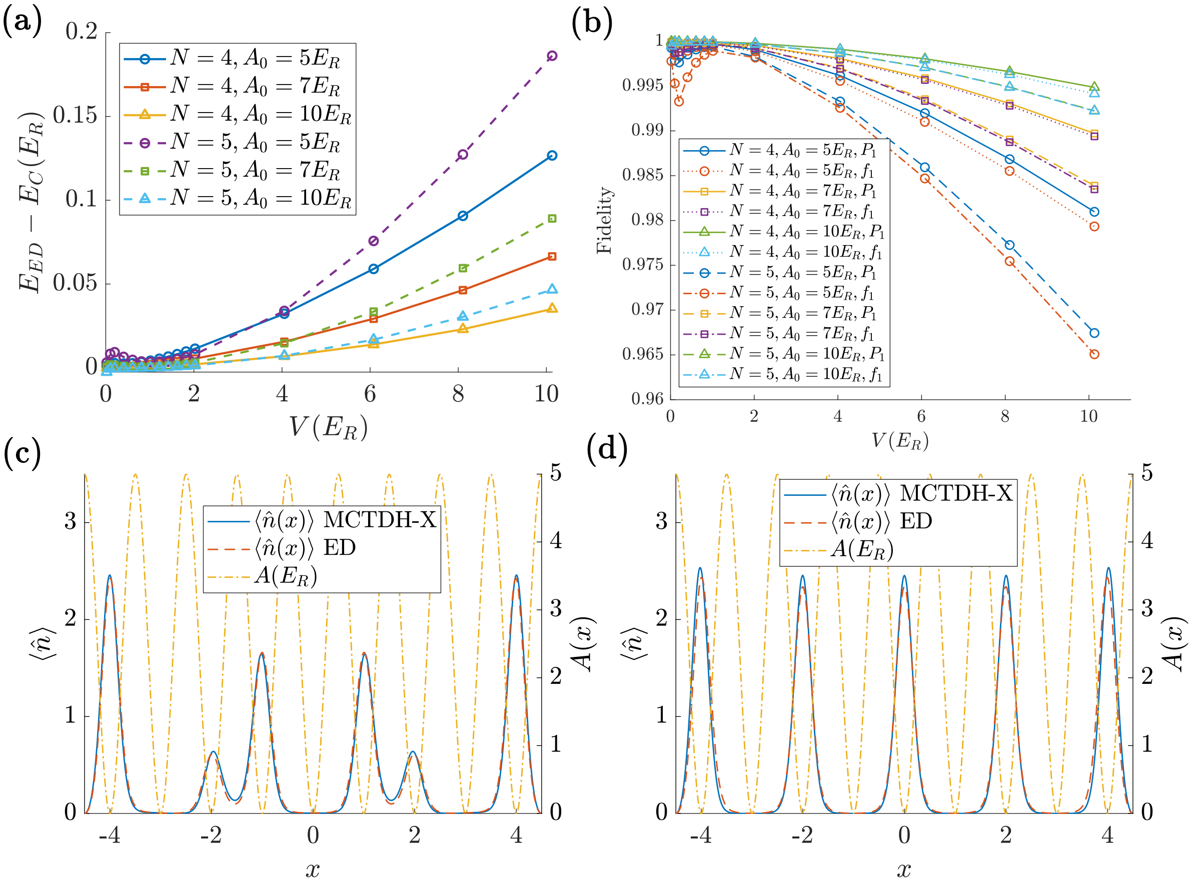

Figure 5.2: Results for particle numbers $N = 4$ and $N = 5$ at varying lattice potential depths A_0 and DDI strengths V. (a) Comparison of ground state energies for ED (E_{ED}) and MCTDH-X (E_C). (b) Projection magnitude P_1 of the MCTDH-X wavefunction $|\psi_C\rangle$ onto lowest band and overlap f_1 with corresponding single-band ED ground state. (c) Comparison of densities between methods for $N = 4$, $A_0 = 5E_R$, $V = 10.1E_R$. (d) Comparison of densities between methods for $N = 5$, $A_0 = 5E_R$, $V = 10.1E_R$.

established and the particles are separated further.

Figures 5.2(c) and (d) show the continuum densities for $A = 5E_R$, $V = 10.1E_R$ for $N = 4$ and $N = 5$ respectively. We extract the continuum densities $\langle \hat{\Psi}^\dagger(x)\hat{\Psi}(x) \rangle$ (on the fine spatial grid) from the lattice wavefunctions using equation 3.3 to express the bosonic annihilation operator $\hat{\Psi}(x)$ in terms of the finite-size Wannier functions. The densities of the two methods are visually similar except for small warping of the MCTDH-X on-site density to minimise dipolar repulsion. For $N = 4$, this slightly displaces the density on the edge sites further towards the edge and slightly displaces the density on sites 3 and 7 towards the centre of the lattice. For $N = 5$, due to the fact that the filling is commensurate with the density wave, there is less benefit in displacing the on-site density except pushing the edge sites further outwards and narrowing the peaks.

For $N = 6$, increased DDI effects cause greater difficulty for the lowest-band dipolar BH model and the continuum numerical methods. Like $N = 4$ the energy gaps to low-lying excited states of the lattice model are small due to incommensurate filling and the lattice ground state is significantly entangled for the range of V we cover. We found that the continuum results for $N = 6$ were not converged with respect to the number of single-particle orbitals for $M = 9$ for $V \geq 0.5E_R$ and instead required $M = 18$. We therefore compared these results with both the lowest-band dipolar BH model and the corresponding two-band model, which features the same number of single-particle orbitals as the continuum calculations.

For the ED calculations for $N = 6$, we also allowed two bosons to occupy each Wannier function and used the same spatial truncation of tunnelling, DDI, and contact interaction terms as for the single-band calculations for $N = 4, 5$, although we note that the second-band Wannier functions have a greater spatial extent. The addition of the second band introduces strong DDI-induced 'tunnelling' between the two bands on the same site, meaning that as well as the density-density interactions $\hat{n}_{j,\sigma_1}\hat{n}_{k,\sigma_2}$, we included terms of the form $\hat{n}_{j,\sigma_1}\hat{b}^\dagger_{k,1}\hat{b}_{k,2}$ and $\hat{b}^\dagger_{j,1}\hat{b}_{j,2}\hat{b}^\dagger_{k,1}\hat{b}_{k,2}$ between all pairs of different sites j and k, where $\hat{n}_{j,\sigma}$ is the density operator for site j and band $\sigma \in \{1, 2\}$. We also included all DDI terms $\hat{b}^\dagger_{i,\sigma_1}\hat{b}^\dagger_{j,\sigma_2}\hat{b}_{k,\sigma_3}\hat{b}_{l,\sigma_4}$ where i, j, k, and l are contained within three consecutive sites, which introduces DDI tunnelling processes between sites and bands. For comparison, we included the same terms in the single-band model for $N = 6$ in the cases where they exist.

Figure 5.3 shows the results for $N = 6$. The energies of the continuum calculations,

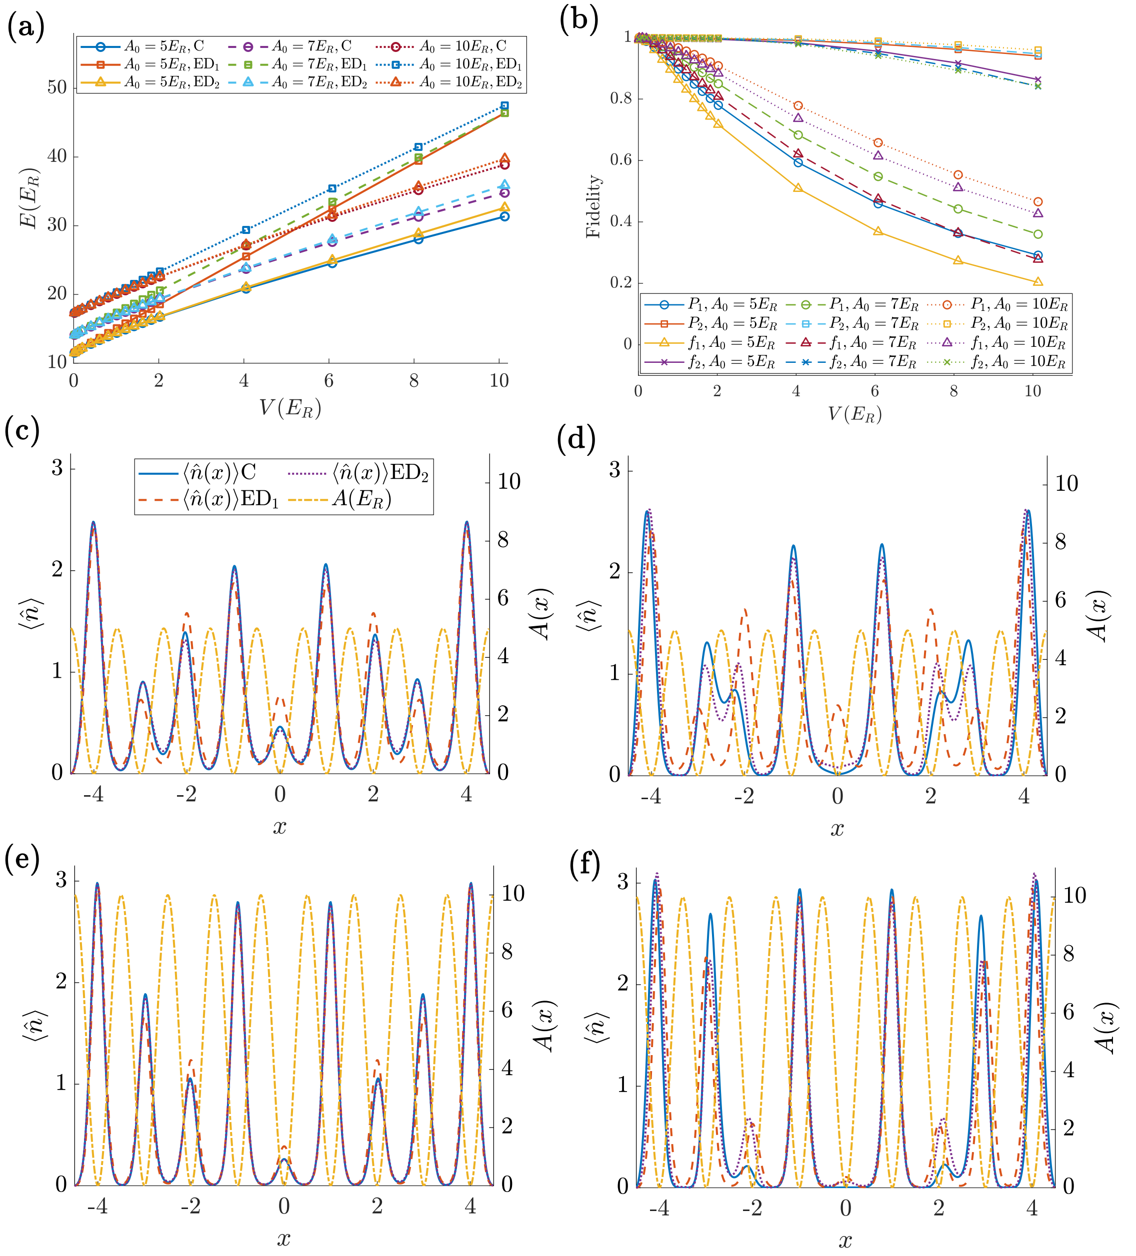

Figure 5.3: Results for particle number $N = 6$ at varying lattice potential depths A_0 and DDI strengths V. (a) Total kinetic, potential, and interaction energies of MCTDH-X (labelled C), ED with one band (ED$_1$), and ED with two bands (ED$_2$). (b) Projection magnitude of MCTDH-X wavefunction $|\psi_{\mathrm{C}}\rangle$ onto the lowest band (labelled P_1) and lowest two bands (P_2) and overlaps with the corresponding ground state for ED using 1 band (f_1) and ED using two bands (f_2). (c-f) Comparison of densities between methods shown with lattice potential A according to the common legend in (c) for physical paramters (c) $A_0 = 5E_R$, $V = 1.01E_R$ (d) $A_0 = 5E_R$, $V = 10.1E_R$, (e) $A_0 = 10E_R$, $V = 1.01E_R$ (f) $A_0 = 10E_R$, $V = 10.1E_R$.

single-band ED, and two-band ED agree well at weak DDI as expected. For strong DDI, the continuum calculations are able to achieve a significantly lower energy than the single-band model and slightly lower than the two-band model, with the energy differences increasing with reduced lattice potential depth A_0 as expected. The projection magnitude of the MCTDH-X wavefunction onto the lowest band is reduced to around 0.5 for the strongest DDI strengths studied, showing a much larger population of excited bands than for $N = 4, 5$. The projection magnitudes onto either the lowest band or lowest two bands decrease monotonically with decreasing lattice potential depth as expected, but this is not always true for the overlap with the ground state found using ED with two bands, where the fidelity between 2-band ED and MCTDH-X is slightly higher at strong DDI for $A_0 = 5E_R$ than $A_0 = 10E_R$. One explanation for this is that the energy gap between the lowest-energy states in the lattice model is small for $N = 6$ because the density wave order is incommensurate, which means small effects from higher bands, which are captured in the continuum model, could cause changes in the ground state within the lowest two bands.

For this filling fraction, where there is greater average density and there must be at least one pair of neighbouring sites with particles, the repositioning of particle density within sites to reduce the DDI repulsion is clearer, as shown in Figure 5.3 (c-f). The two-band model is able to qualitatively capture this effect because superpositions of occupation of the lowest two bands are offset from the lattice minimum, although this model does not minimise the energy as effectively as the continuum calculations with $M = 18$.

5.3.1 Finite-Size Scaling

While the computationally-feasible system sizes are limited and commensurability effects are significant, it is possible to compare the results of three systems which are

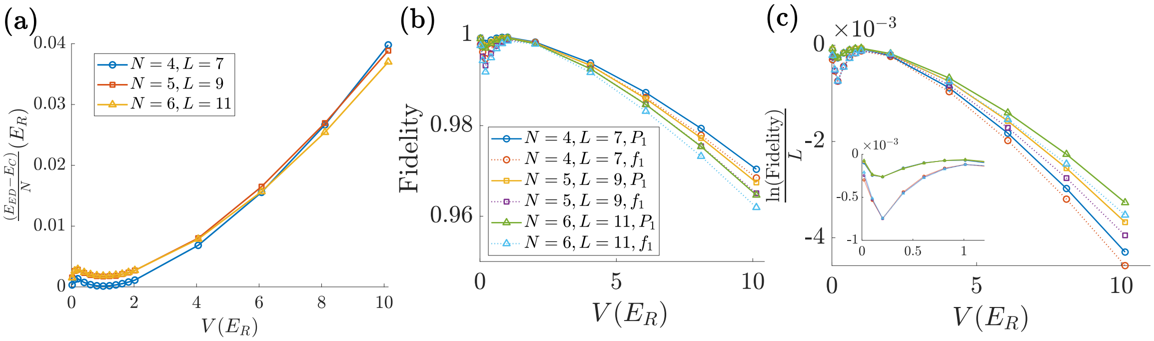

Figure 5.4: Data at fixed lattice potential depth $A_0 = 5E_R$ for three systems with filling fraction commensurate with the density wave order: particle number $N = 4$ and length $L = 7$, $N = 5$ and $L = 9$, and $N = 6$ and $L = 11$. (a) Comparison of energies per particle for continuum and single-band lattice models $\frac{E_{ED}-E_C}{N}$ as a function of DDI strength V. (b) Comparison of projection magnitude of the continuum wavefunction onto the lowest band P_1 and overlap with corresponding single-band ED ground state f_1. (c) Rescaled version of data in (b) showing $\frac{\ln(P_1)}{L}$ and $\frac{\ln(f_1)}{L}$. (inset) Zoomed-in version of (c) focussing on weak DDI region.

compatible with the density wave order: particle number $N = 4$ and length $L = 7$, $N = 5$ and $L = 9$, and $N = 6$ and $L = 11$. The larger length of the latter system is possible because only one MCTDH-X orbital per site (i.e. $M = 11$) is necessary unlike the smaller but more strongly-interacting system with $N = 6$ and $L = 9$. The MCTDH-X calculations use the same fine spatial grid as those for $L = 9$.

Comparisons of the energy per particle and many-body wavefunctions are shown in Figure 5.4. From Figure 5.4 (a), it can be seen that the energy difference between methods per particle is very similar between the two largest systems at relatively weak DDI $V < 1E_R$ but that this difference is larger for smaller systems (which have larger edge effects) at very strong DDI. This suggests that part of the reduction in energy available to the continuum calculations by populating higher bands at very strong DDI is linked to the edges of the system rather than the bulk, as the DDI repulsion can be reduced by displacing the particles in the edge sites slightly outwards.

A similar conclusion can be seen from the comparison of the fidelities in Figure

5.4 (b) and (c). From Figure 5.4 (b), it is clear that the projection magnitude of the continuum wavefunction to the lowest band P_1 and the overlap with the single-band ED wavefunction f_1 vary in a very similar manner with respect to DDI strength, with the larger systems reporting lower values of P_1 and f_1 as expected. Figure 5.4 (c) shows these projection magnitudes and fidelities rescaled according to the ansatz that $\frac{ln(P_1)}{L}$ and $\frac{ln(f_1)}{L}$ should not scale with respect to system size for sufficiently-large equivalent systems. This ansatz is based on a multiplicative reduction in fidelity upon increasing the size of a system, and is equivalent to the ansatz used in Chapter 4 to compare the fidelity susceptibility per cylinder lattice ring χ_F for different system sizes (equation 4.10) [6]. At relatively weak DDI strength $V < 1E_R$, P_1 and f_1 for the three system sizes collapse approximately onto a single curve respectively when rescaled according to this relation (inset of Figure 5.4 (c)). For very strong DDI, the rescaled P_1 and f_1 are lower for smaller systems, suggesting again that the population of higher bands in this parameter regime involves significant edge effects.

5.4 Discussion

The much greater deviations between the lattice and continuum methods for $N = 6$ compared to $N = 4$ and $N = 5$ for $L = 9$ suggest that the quantitative validity of the lowest-band dipolar BH model at strong DDI is strongly dependent on the filling pattern of the ground state. The occupations of nearby sites significantly affect the warping of the on-site density. This would complicate the application of on-site occupation-dependent Wannier functions and BH parameters, which have been used for strong short-range interactions [241, 241, 242]. While in this example, the small 1D lattice leaves little freedom for particle rearrangement to reduce the energy, this effect may be

---

[6] Note that here, the fidelity is between ground state for the same physical Hamiltonian produced by two methods rather than between ground states for different Hamiltonians.

more important for more complicated particle arrangements in higher dimensions or for attractive interactions where the particles preferentially occupy nearby sites.

While this brief study could not determine general conditions for parameter regimes which cause large populations of higher bands, it is intuitively clear that greater density increases the higher-band population for repulsive DDI because of the reduced particle separation. Quantitatively, the expectation value of the DDI energy per particle $\frac{\langle \hat{H}_V \rangle}{N}$ increases for greater particle density [7] which can be compared to the single-particle band gap (Figure 5.1(b)) for an estimate of whether the interactions can cause significant population of higher bands. Note that this does not completely explain the population of higher bands and in general the many-body physics of the ground state must be taken into account, as shown by the non-monotonic behaviour of the population of the lowest band P_1 as a function of DDI strength in the case where the particle number is compatible with the density wave (e.g. Figure 5.4 (b) and (c)).

In our few examples , incommensurability of the particle numbers and lattice length with the density wave order (for $L = 9$ and $N = 4, 6$) was important for causing displacement of particles within lattice sites in the continuum calculations, which is a clear signal of higher-band population. The incommensurability also causes a reduced many-body energy gap to the lowest-energy excited states of the lowest-band lattice model, although our numerical data suggests that a comparison of the DDI strength and the many-body gap alone is not sufficient to explain the population of higher bands. This can be seen in the example of lowest-band lattice calculations for $L = 9$, lattice potential depth $A_0 = 5E_R$, and strong DDI strength $V = 10.1E_R$, where the gap between the ground state and first excited state of the lowest-band lattice model

---

[7] For length $L = 9$, lattice potential depth $A_0 = 10E_R$, and DDI strength $V = 10.1E_R$, $\frac{\langle \hat{H}_V \rangle}{N} = 0.57E_R$ for $N = 4$ in the single-band lattice calculations, $\frac{\langle \hat{H}_V \rangle}{N} = 1.19E_R$ for $N = 5$ in the single-band lattice calculations, and $\frac{\langle \hat{H}_V \rangle}{N} = 3.34E_R$ for $N = 6$ in the two-band lattice calculations.

is $0.097E_R$, $11.9E_R$, and $0.97E_R$ for $N = 4$, 5, and 6 respectively [8]. The extremely small many-body gap for $N = 4$ in this example does not lead to greater population of excited bands than the larger gap for $N = 6$. Despite the conclusion that a comparison of the DDI strength with the many-body gap does not directly explain the population of higher bands, it is plausible that some systems with a small many-body gap may undergo more dramatic ground-state wavefunction changes due to population of higher bands if favoured by the physics of the inter-particle interactions.

5.5 Experimental Relevance

Bose-Hubbard models showing dynamical effects of the DDI have been implemented using magnetic atoms [44]. In that example, the DDI between the two atoms at a distance of one lattice site was around 1% of the recoil energy, deep within the regime of quantitative agreement between the lowest-band dipolar BH model and the continuum methods for moderate lattice potentials. For the most strongly dipolar polar molecules, this is not necessarily the case. Taking the example of NaCs, which has a large electric dipole moment of up to 4.7D (2.6D of which has been implemented in recent experiments [243]), the repulsive DDI between two molecules polarised perpendicularly to the lattice at a distance of 532nm could be up to $\approx$ 20kHz. This is over 15× the recoil energy E_R in a lattice of wavelength $\lambda = 1064$nm, suggesting that the DDI strengths required to compete with moderate lattice potentials may be realisable in near-future experiments. We considered filling fractions of around one boson per two sites, which is somewhat larger than previously observed filling fractions of 30% for

---

[8] For very strong interaction ($V >> J$) in the lowest band of the lattice of length $L = 9$, $N = 4$ has two low-energy site-occupation basis states, $|101001001\rangle$ and $|100100101\rangle$, which have equal energy, causing the very small many-body gap between the lowest eigenstates of the lattice model (which approach the reflection-symmetric and anti-symmetric superpositions of these basis states). $N = 6$ would have a non-degenerate lowest-energy basis state $|110101011\rangle$, whose energy is split from $|101101011\rangle$ (and the latter state's reflection about the centre of the lattice) by a relatively weak long-range DDI term, causing the slightly larger energy gap.

bosonic polar molecules in optical lattices [77].

5.6 Conclusions

In this short chapter, we have compared lattice and continuum methods for dipolar interacting bosons in finite optical lattice potentials, verifying the validity of the lowest-band dipolar BH model through agreement of energy and wavefunction overlap for moderate lattice potentials and comparatively weak DDI. We extended this to quantitatively study the breakdown of the lowest Hubbard band description in 1D repulsively-interacting systems, finding that the on-site wavefunctions are warped by the presence of particles on other sites. While a general quantification of the validity of the lowest-band dipolar BH model would require significant additional calculations including periodic boundary conditions, our results show that experimentally-feasible parameters can access both the regime in which the lowest-band dipolar BH model is clearly valid and the regime in which the DDI significantly populates higher bands for finite-size lattices.

Future work could extend this brief study to focus on time evolution where significant discrepancies between lowest-band lattice and continuum methods were found for relatively weak contact interactions [229]. The breakdown of the lowest-band BH model may have more significant consequences for the physics at weaker lattice potential or anisotropic DDI, where interactions may encourage particles to occupy neighbouring sites, or in cases where the DDI introduces more intricate competition of interactions within the BH model itself.

6

Conclusions and Outlook

In this thesis, we have considered the application of quantum control over ultracold dipolar systems to the tasks of quantum computation and quantum simulation. In chapter 2, we studied schemes to robustly entangle polar molecules using the DDI, finding that shaped microwave pulses offer useful levels of robustness to experimentally-relevant systematic errors. It is hoped these results will be of relevance to near-future experiments either directly or through incorporation in modified form into later schemes to enhance the DDI.

In chapters 4 and 5, we investigated the role of the DDI between lattice sites in dipolar Bose-Hubbard models. Specifically, in chapter 4, we numerically studied the recently-proposed cylindrical lattice geometry, finding that the combination of the anisotropic DDI and the curved lattice surface expanded the physics of previously-studied square lattice and ring lattice dipolar Bose-Hubbard models by creating competition between spontaneous and sublattice ordering mechanisms. Our numerical results suggest that moderate system sizes of ≈ 100 sites are able to qualitatively identify much of the resultant physics. These results may provide guidance to future experiments and exploration of the role of the DDI in emerging many-body systems of flexible geometry.

In chapter 5 we quantified the discrepancy between the lowest-band dipolar Bose-Hubbard model and direct continuum methods for small systems of strongly-interacting dipolar bosons in optical lattice potentials, finding that the spatial ordering introduced by the DDI influences the qualitative physics of the failure of the lowest-band dipolar Bose-Hubbard model. These results may provide a platform for similar investigation of a wider range of systems, especially those where the DDI causes more delicate competition of ordered states.

Bibliography

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

[2] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 212–219, New York, NY, USA, 1996. Association for Computing Machinery.

[3] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26(5):1484–1509, 1997.

[4] David P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.

[5] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.

[6] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent

Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, Oct 2019.

[7] Han-Sen Zhong, Yu-Hao Deng, Jian Qin, Hui Wang, Ming-Cheng Chen, Li-Chao Peng, Yi-Han Luo, Dian Wu, Si-Qiu Gong, Hao Su, Yi Hu, Peng Hu, Xiao-Yan Yang, Wei-Jun Zhang, Hao Li, Yuxuan Li, Xiao Jiang, Lin Gan, Guangwen Yang, Lixing You, Zhen Wang, Li Li, Nai-Le Liu, Jelmer J. Renema, Chao-Yang Lu, and Jian-Wei Pan. Phase-programmable gaussian boson sampling using stimulated squeezed light. *Phys. Rev. Lett.*, 127:180502, Oct 2021.

[8] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, Ming Gong, Cheng Guo, Chu Guo, Shaojun Guo, Lianchen Han, Linyin Hong, He-Liang Huang, Yong-Heng Huo, Liping Li, Na Li, Shaowei Li, Yuan Li, Futian Liang, Chun Lin, Jin Lin, Haoran Qian, Dan Qiao, Hao Rong, Hong Su, Lihua Sun, Liangyuan Wang, Shiyu Wang, Dachao Wu, Yu Xu, Kai Yan, Weifeng Yang, Yang Yang, Yangsen Ye, Jianghan Yin, Chong Ying, Jiale Yu, Chen Zha, Cha Zhang, Haibin Zhang,

Kaili Zhang, Yiming Zhang, Han Zhao, Youwei Zhao, Liang Zhou, Qingling Zhu, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. Strong quantum computational advantage using a superconducting quantum processor. *Phys. Rev. Lett.*, 127:180501, Oct 2021.

[9] Laszlo Gyongyosi and Sandor Imre. A survey on quantum computing technology. *Computer Science Review*, 31:51–71, 2019.

[10] Jacob A Blackmore, Luke Caldwell, Philip D Gregory, Elizabeth M Bridge, Rahul Sawant, Jesús Aldegunde, Jordi Mur-Petit, Dieter Jaksch, Jeremy M Hutson, B E Sauer, M R Tarbutt, and Simon L Cornish. Ultracold molecules for quantum simulation: rotational coherences in CaF and RbCs. *Quantum Science and Technology*, 4(1):014010, dec 2018.

[11] Rahul Sawant, Jacob A Blackmore, Philip D Gregory, Jordi Mur-Petit, Dieter Jaksch, Jesús Aldegunde, Jeremy M Hutson, M R Tarbutt, and Simon L Cornish. Ultracold polar molecules as qudits. *New Journal of Physics*, 22(1):013027, jan 2020.

[12] Jee Woo Park, Zoe Z. Yan, Huanqian Loh, Sebastian A. Will, and Martin W. Zwierlein. Second-scale nuclear spin coherence time of ultracold ^{23}Na^{40}K molecules. *Science*, 357(6349):372–375, 2017.

[13] Philip D. Gregory, Jacob A. Blackmore, Sarah L. Bromley, Jeremy M. Hutson, and Simon L. Cornish. Robust storage qubits in ultracold polar molecules. *Nature Physics*, 17(10):1149–1153, Oct 2021.

[14] Junyu Lin, Junyu He, Mucan Jin, Guanghua Chen, and Dajun Wang. Seconds-scale coherence on nuclear spin transitions of ultracold polar molecules in 3d optical lattices. *Phys. Rev. Lett.*, 128:223201, Jun 2022.

[15] D. DeMille. Quantum computation with trapped polar molecules. *Phys. Rev. Lett.*, 88:067901, Jan 2002.

[16] S. F. Yelin, K. Kirby, and Robin Côté. Schemes for robust quantum computation with polar molecules. *Phys. Rev. A*, 74:050301, Nov 2006.

[17] Felipe Herrera, Yudong Cao, Sabre Kais, and K Birgitta Whaley. Infrared-dressed entanglement of cold open-shell polar molecules for universal matchgate quantum computing. *New Journal of Physics*, 16(7):075001, jul 2014.

[18] Mallikarjun Karra, Ketan Sharma, Bretislav Friedrich, Sabre Kais, and Dudley Herschbach. Prospects for quantum computing with an array of ultracold polar paramagnetic molecules. *The Journal of Chemical Physics*, 144(9):094301, 2016.

[19] Kang-Kuen Ni, Till Rosenband, and David D. Grimes. Dipolar exchange quantum logic gate with polar molecules. *Chem. Sci.*, 9:6830–6838, 2018.

[20] Eunmi Chae. Entanglement via rotational blockade of MgF molecules in a magic potential. *Phys. Chem. Chem. Phys.*, 23:1215–1220, 2021.

[21] Michael Hughes, Matthew D. Frye, Rahul Sawant, Gaurav Bhole, Jonathan A. Jones, Simon L. Cornish, M. R. Tarbutt, Jeremy M. Hutson, Dieter Jaksch, and Jordi Mur-Petit. Robust entangling gate for polar molecules using magnetic and microwave fields. *Phys. Rev. A*, 101:062308, Jun 2020.

[22] Tomi H. Johnson, Stephen R. Clark, and Dieter Jaksch. What is a quantum simulator? *EPJ Quantum Technology*, 1(1):10, Jul 2014.

[23] Philipp Hauke, Fernando M Cucchietti, Luca Tagliacozzo, Ivan Deutsch, and Maciej Lewenstein. Can one trust quantum simulators? *Reports on Progress in Physics*, 75(8):082401, jul 2012.

[24] I. M. Georgescu, S. Ashhab, and Franco Nori. Quantum simulation. *Rev. Mod. Phys.*, 86:153–185, Mar 2014.

[25] 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, Mar 2011.

[26] 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, Nov 2009.

[27] Florian Schäfer, Takeshi Fukuhara, Seiji Sugawa, Yosuke Takasu, and Yoshiro Takahashi. Tools for quantum simulation with ultracold atoms in optical lattices. *Nature Reviews Physics*, 2(8):411–425, Aug 2020.

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

[29] M. A. Baranov, M. Dalmonte, G. Pupillo, and P. Zoller. Condensed matter theory of dipolar quantum gases. *Chemical Reviews*, 112(9):5012–5061, Sep 2012.

[30] S. Yi, T. Li, and C. P. Sun. Novel quantum phases of dipolar bose gases in optical lattices. *Phys. Rev. Lett.*, 98:260405, Jun 2007.

[31] B. Capogrosso-Sansone, C. Trefzger, M. Lewenstein, P. Zoller, and G. Pupillo. Quantum phases of cold polar molecules in 2D optical lattices. *Phys. Rev. Lett.*, 104:125301, Mar 2010.

[32] L. Pollet, J. D. Picon, H. P. Büchler, and M. Troyer. Supersolid phase with cold polar molecules on a triangular lattice. *Phys. Rev. Lett.*, 104:125302, Mar 2010.

[33] N. Y. Yao, M. P. Zaletel, D. M. Stamper-Kurn, and A. Vishwanath. A quantum dipolar spin liquid. *Nature Physics*, 14(4):405–410, Apr 2018.

[34] Haiyuan Zou, Erhai Zhao, and W. Vincent Liu. Frustrated magnetism of dipolar molecules on a square optical lattice: Prediction of a quantum paramagnetic ground state. *Phys. Rev. Lett.*, 119:050401, Jul 2017.

[35] A. Micheli, G. K. Brennen, and P. Zoller. A toolbox for lattice-spin models with polar molecules. *Nature Physics*, 2(5):341–347, May 2006.

[36] N. Y. Yao, C. R. Laumann, A. V. Gorshkov, S. D. Bennett, E. Demler, P. Zoller, and M. D. Lukin. Topological flat bands from dipolar spin systems. *Phys. Rev. Lett.*, 109:266804, Dec 2012.

[37] N. Y. Yao, A. V. Gorshkov, C. R. Laumann, A. M. Läuchli, J. Ye, and M. D. Lukin. Realizing fractional Chern insulators in dipolar spin systems. *Phys. Rev. Lett.*, 110:185302, Apr 2013.

[38] J. Levinsen, N. R. Cooper, and G. V. Shlyapnikov. Topological $p_x + ip_y$ superfluid phase of fermionic polar molecules. *Phys. Rev. A*, 84:013603, Jul 2011.

[39] Alexey V. Gorshkov, Kaden R.A. Hazzard, and Ana Maria Rey. Kitaev honeycomb and other exotic spin models with polar molecules. *Molecular Physics*, 111(12-13):1908–1916, 2013.

[40] Joana Fraxanet, Daniel González-Cuadra, Tilman Pfau, Maciej Lewenstein, Tim Langen, and Luca Barbiero. Topological quantum critical points in the extended Bose-Hubbard model. *Phys. Rev. Lett.*, 128:043402, Jan 2022.

[41] Sergi Julià-Farré, Daniel González-Cuadra, Alexander Patscheider, Manfred J. Mark, Francesca Ferlaino, Maciej Lewenstein, Luca Barbiero, and Alexandre

Dauphin. Revealing the topological nature of the bond order wave in a strongly correlated quantum system. *Phys. Rev. Research*, 4:L032005, Jul 2022.

[42] Bo Yan, Steven A. Moses, Bryce Gadway, Jacob P. Covey, Kaden R. A. Hazzard, Ana Maria Rey, Deborah S. Jin, and Jun Ye. Observation of dipolar spin-exchange interactions with lattice-confined polar molecules. *Nature*, 501(7468):521–525, Sep 2013.

[43] Kaden R. A. Hazzard, Bryce Gadway, Michael Foss-Feig, Bo Yan, Steven A. Moses, Jacob P. Covey, Norman Y. Yao, Mikhail D. Lukin, Jun Ye, Deborah S. Jin, and Ana Maria Rey. Many-body dynamics of dipolar molecules in an optical lattice. *Phys. Rev. Lett.*, 113:195302, Nov 2014.

[44] S. Baier, M. J. Mark, D. Petter, K. Aikawa, L. Chomaz, Z. Cai, M. Baranov, P. Zoller, and F. Ferlaino. Extended Bose-Hubbard models with ultracold magnetic atoms. *Science*, 352(6282):201–205, 2016.

[45] Lysander Christakis, Jason S. Rosenberg, Ravin Raj, Sungjae Chi, Alan Morningstar, David A. Huse, Zoe Z. Yan, and Waseem S. Bakr. Probing site-resolved correlations in a spin system of ultracold molecules, 2022.

[46] Michael Hughes and Dieter Jaksch. Dipolar Bose-Hubbard model in finite-size real-space cylindrical lattices. *Phys. Rev. A*, 105:053301, May 2022.

[47] Alexey V. Gorshkov, Salvatore R. Manmana, Gang Chen, Jun Ye, Eugene Demler, Mikhail D. Lukin, and Ana Maria Rey. Tunable superfluidity and quantum magnetism with ultracold polar molecules. *Phys. Rev. Lett.*, 107:115301, Sep 2011.

[48] Alexey V. Gorshkov, Salvatore R. Manmana, Gang Chen, Eugene Demler, Mikhail D. Lukin, and Ana Maria Rey. Quantum magnetism with polar alkali-metal dimers. *Phys. Rev. A*, 84:033619, Sep 2011.

[49] Salvatore R. Manmana, Marcel Möller, Riccardo Gezzi, and Kaden R. A. Hazzard. Correlations and enlarged superconducting phase of $t - J_{\perp}$ chains of ultracold molecules on optical lattices. *Phys. Rev. A*, 96:043618, Oct 2017.

[50] A. de Paz, A. Sharma, A. Chotia, E. Maréchal, J. H. Huckans, P. Pedri, L. Santos, O. Gorceix, L. Vernac, and B. Laburthe-Tolra. Nonequilibrium quantum magnetism in a dipolar lattice gas. *Phys. Rev. Lett.*, 111:185305, Oct 2013.

[51] S. Lepoutre, J. Schachenmayer, L. Gabardos, B. Zhu, B. Naylor, E. Maréchal, O. Gorceix, A. M. Rey, L. Vernac, and B. Laburthe-Tolra. Out-of-equilibrium quantum magnetism and thermalization in a spin-3 many-body dipolar lattice system. *Nature Communications*, 10(1):1714, Apr 2019.

[52] A. Patscheider, B. Zhu, L. Chomaz, D. Petter, S. Baier, A.-M. Rey, F. Ferlaino, and M. J. Mark. Controlling dipolar exchange interactions in a dense three-dimensional array of large-spin fermions. *Phys. Rev. Research*, 2:023050, Apr 2020.

[53] Lucas Gabardos, Bihui Zhu, Steven Lepoutre, Ana Maria Rey, Bruno Laburthe-Tolra, and Laurent Vernac. Relaxation of the collective magnetization of a dense 3d array of interacting dipolar $S = 3$ atoms. *Phys. Rev. Lett.*, 125:143401, Sep 2020.

[54] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe. Observation of a many-body dynamical

phase transition with a 53-qubit quantum simulator. *Nature*, 551(7682):601–604, Nov 2017.

[55] 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, Nov 2017.

[56] Pascal Scholl, Michael Schuler, Hannah J. Williams, Alexander A. Eberharter, Daniel Barredo, Kai-Niklas Schymik, Vincent Lienhard, Louis-Paul Henry, Thomas C. Lang, Thierry Lahaye, Andreas M. Läuchli, and Antoine Browaeys. Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms. *Nature*, 595(7866):233–238, Jul 2021.

[57] 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, Jul 2021.

[58] Axel Griesmaier, Jörg Werner, Sven Hensler, Jürgen Stuhler, and Tilman Pfau. Bose-Einstein condensation of Chromium. *Phys. Rev. Lett.*, 94:160401, Apr 2005.

[59] K. Aikawa, A. Frisch, M. Mark, S. Baier, A. Rietzler, R. Grimm, and F. Ferlaino. Bose-Einstein condensation of Erbium. *Phys. Rev. Lett.*, 108:210401, May 2012.

[60] Mingwu Lu, Nathaniel Q. Burdick, Seo Ho Youn, and Benjamin L. Lev. Strongly dipolar Bose-Einstein condensate of Dysprosium. *Phys. Rev. Lett.*, 107:190401, Oct 2011.

[61] Mingwu Lu, Nathaniel Q. Burdick, and Benjamin L. Lev. Quantum degenerate dipolar Fermi gas. *Phys. Rev. Lett.*, 108:215301, May 2012.

[62] Maxence Lepers, Hui Li, Jean-François Wyart, Goulven Quéméner, and Olivier Dulieu. Ultracold rare-earth magnetic atoms with an electric dipole moment. *Phys. Rev. Lett.*, 121:063201, Aug 2018.

[63] Hui Li, Goulven Quéméner, Jean-François Wyart, Olivier Dulieu, and Maxence Lepers. Purely long-range polar molecules composed of identical lanthanide atoms. *Phys. Rev. A*, 100:042711, Oct 2019.

[64] Matthew D. Frye, Simon L. Cornish, and Jeremy M. Hutson. Prospects of forming high-spin polar molecules from ultracold atoms. *Phys. Rev. X*, 10:041005, Oct 2020.

[65] Lauriane Chomaz, Igor Ferrier-Barbut, Francesca Ferlaino, Bruno Laburthe-Tolra, Benjamin L. Lev, and Tilman Pfau. Dipolar physics: A review of experiments with magnetic quantum gases, 2022.

[66] A. de Paz, P. Pedri, A. Sharma, M. Efremov, B. Naylor, O. Gorceix, E. Maréchal, L. Vernac, and B. Laburthe-Tolra. Probing spin dynamics from the Mott insulating to the superfluid regime in a dipolar lattice gas. *Phys. Rev. A*, 93:021603, Feb 2016.

[67] Petra Fersterer, Arghavan Safavi-Naini, Bihui Zhu, Lucas Gabardos, Steven Lepoutre, L. Vernac, B. Laburthe-Tolra, P. Blair Blakie, and Ana Maria Rey. Dynamics of an itinerant spin-3 atomic dipolar gas in an optical lattice. *Phys. Rev. A*, 100:033609, Sep 2019.

[68] Fabian Böttcher, Jan-Niklas Schmidt, Matthias Wenzel, Jens Hertkorn, Mingyang Guo, Tim Langen, and Tilman Pfau. Transient supersolid properties in an array of dipolar quantum droplets. *Phys. Rev. X*, 9:011051, Mar 2019.

[69] L. Tanzi, E. Lucioni, F. Famà, J. Catani, A. Fioretti, C. Gabbanini, R. N. Bisset, L. Santos, and G. Modugno. Observation of a dipolar quantum gas with metastable supersolid properties. *Phys. Rev. Lett.*, 122:130405, Apr 2019.

[70] L. Chomaz, D. Petter, P. Ilzhöfer, G. Natale, A. Trautmann, C. Politi, G. Durastante, R. M. W. van Bijnen, A. Patscheider, M. Sohmen, M. J. Mark, and F. Ferlaino. Long-lived and transient supersolid behaviors in dipolar quantum gases. *Phys. Rev. X*, 9:021012, Apr 2019.

[71] Matthew A. Norcia, Claudia Politi, Lauritz Klaus, Elena Poli, Maximilian Sohmen, Manfred J. Mark, Russell N. Bisset, Luis Santos, and Francesca Ferlaino. Two-dimensional supersolidity in a dipolar quantum gas. *Nature*, 596(7872):357–361, Aug 2021.

[72] J. Deiglmayr, A. Grochola, M. Repp, K. Mörtlbauer, C. Glück, J. Lange, O. Dulieu, R. Wester, and M. Weidemüller. Formation of ultracold polar molecules in the rovibrational ground state. *Phys. Rev. Lett.*, 101:133004, Sep 2008.

[73] K.-K. Ni, S. Ospelkaus, M. H. G. de Miranda, A. Pe'er, B. Neyenhuis, J. J. Zirbel, S. Kotochigova, P. S. Julienne, D. S. Jin, and J. Ye. A high phase-space-density gas of polar molecules. *Science*, 322(5899):231–235, 2008.

[74] Yichao Yu, Kenneth Wang, Jonathan D. Hood, Lewis R. B. Picard, Jessie T. Zhang, William B. Cairncross, Jeremy M. Hutson, Rosario Gonzalez-Ferez, Till Rosenband, and Kang-Kuen Ni. Coherent optical creation of a single molecule. *Phys. Rev. X*, 11:031061, Sep 2021.

[75] Amodsen Chotia, Brian Neyenhuis, Steven A. Moses, Bo Yan, Jacob P. Covey, Michael Foss-Feig, Ana Maria Rey, Deborah S. Jin, and Jun Ye. Long-lived dipolar molecules and Feshbach molecules in a 3D optical lattice. *Phys. Rev. Lett.*, 108:080405, Feb 2012.

[76] Steven A. Moses, Jacob P. Covey, Matthew T. Miecnikowski, Bo Yan, Bryce Gadway, Jun Ye, and Deborah S. Jin. Creation of a low-entropy quantum gas of polar molecules in an optical lattice. *Science*, 350(6261):659–662, 2015.

[77] Lukas Reichsöllner, Andreas Schindewolf, Tetsu Takekoshi, Rudolf Grimm, and Hanns-Christoph Nägerl. Quantum engineering of a low-entropy gas of heteronuclear bosonic molecules in an optical lattice. *Phys. Rev. Lett.*, 118:073201, Feb 2017.

[78] L. R. Liu, J. D. Hood, Y. Yu, J. T. Zhang, N. R. Hutzler, T. Rosenband, and K.-K. Ni. Building one molecule from a reservoir of two atoms. *Science*, 360(6391):900–903, 2018.

[79] L. R. Liu, J. D. Hood, Y. Yu, J. T. Zhang, K. Wang, Y.-W. Lin, T. Rosenband, and K.-K. Ni. Molecular assembly of ground-state cooled single atoms. *Phys. Rev. X*, 9:021039, May 2019.

[80] William B. Cairncross, Jessie T. Zhang, Lewis R. B. Picard, Yichao Yu, Kenneth Wang, and Kang-Kuen Ni. Assembly of a rovibrational ground state molecule in an optical tweezer. *Phys. Rev. Lett.*, 126:123402, Mar 2021.

[81] Luigi De Marco, Giacomo Valtolina, Kyle Matsuda, William G. Tobias, Jacob P. Covey, and Jun Ye. A degenerate Fermi gas of polar molecules. *Science*, 363(6429):853–856, 2019.

[82] E. S. Shuman, J. F. Barry, and D. DeMille. Laser cooling of a diatomic molecule. *Nature*, 467(7317):820–823, Oct 2010.

[83] S. Truppe, H. J. Williams, M. Hambach, L. Caldwell, N. J. Fitch, E. A. Hinds, B. E. Sauer, and M. R. Tarbutt. Molecules cooled below the doppler limit. *Nature Physics*, 13(12):1173–1176, Dec 2017.

[84] Shiqian Ding, Yewei Wu, Ian A. Finneran, Justin J. Burau, and Jun Ye. Sub-doppler cooling and compressed trapping of YO molecules at μK temperatures. *Phys. Rev. X*, 10:021049, Jun 2020.

[85] Loïc Anderegg, Lawrence W. Cheuk, Yicheng Bao, Sean Burchesky, Wolfgang Ketterle, Kang-Kuen Ni, and John M. Doyle. An optical tweezer array of ultracold molecules. *Science*, 365(6458):1156–1158, 2019.

[86] Lawrence W. Cheuk, Loïc Anderegg, Yicheng Bao, Sean Burchesky, Scarlett S. Yu, Wolfgang Ketterle, Kang-Kuen Ni, and John M. Doyle. Observation of collisions between two ultracold ground-state CaF molecules. *Phys. Rev. Lett.*, 125:043401, Jul 2020.

[87] Lincoln D Carr, David DeMille, Roman V Krems, and Jun Ye. Cold and ultracold molecules: science, technology and applications. *New Journal of Physics*, 11(5):055049, may 2009.

[88] B. Friedrich and D. R. Herschbach. On the possibility of orienting rotationally cooled polar molecules in an electric field. *Zeitschrift für Physik D Atoms, Molecules and Clusters*, 18(2):153–161, Jun 1991.

[89] Bretislav Friedrich and Dudley R. Herschbach. Spatial orientation of molecules in strong electric fields and evidence for pendular states. *Nature*, 353(6343):412–414, Oct 1991.

[90] M. A. Anderson, M. D. Allen, and L. M. Ziurys. The Millimeter-Wave Spectrum of the CaF Radical ($X^2\Sigma^+$). *Astrophysical Journal*, 424:503, March 1994.

[91] W.J. Childs, G.L. Goodman, and L.S. Goodman. Precise determination of the v and N dependence of the spin-rotation and hyperfine interactions in the CaF $X^2\Sigma_{\frac{1}{2}}$ ground state. *Journal of Molecular Spectroscopy*, 86(2):365–392, 1981.

[92] H. J. Williams, L. Caldwell, N. J. Fitch, S. Truppe, J. Rodewald, E. A. Hinds, B. E. Sauer, and M. R. Tarbutt. Magnetic trapping and coherent control of laser-cooled molecules. *Phys. Rev. Lett.*, 120:163201, Apr 2018.

[93] Lawrence W. Cheuk, Loïc Anderegg, Benjamin L. Augenbraun, Yicheng Bao, Sean Burchesky, Wolfgang Ketterle, and John M. Doyle. Λ-enhanced imaging of molecules in an optical trap. *Phys. Rev. Lett.*, 121:083201, Aug 2018.

[94] Loïc Anderegg, Benjamin L. Augenbraun, Yicheng Bao, Sean Burchesky, Lawrence W. Cheuk, Wolfgang Ketterle, and John M. Doyle. Laser cooling of optically trapped molecules. *Nature Physics*, 14(9):890–893, Sep 2018.

[95] L. Caldwell, H. J. Williams, N. J. Fitch, J. Aldegunde, Jeremy M. Hutson, B. E. Sauer, and M. R. Tarbutt. Long rotational coherence times of molecules in a magnetic trap. *Phys. Rev. Lett.*, 124:063001, Feb 2020.

[96] Sean Burchesky, Loïc Anderegg, Yicheng Bao, Scarlett S. Yu, Eunmi Chae, Wolfgang Ketterle, Kang-Kuen Ni, and John M. Doyle. Rotational coherence

times of polar molecules in optical tweezers. *Phys. Rev. Lett.*, 127:123202, Sep 2021.

[97] B. Neyenhuis, B. Yan, S. A. Moses, J. P. Covey, A. Chotia, A. Petrov, S. Kotochigova, J. Ye, and D. S. Jin. Anisotropic polarizability of ultracold polar ^{40}K^{87}Rb molecules. *Phys. Rev. Lett.*, 109:230403, Dec 2012.

[98] Philip D. Gregory, Jacob A. Blackmore, Jesus Aldegunde, Jeremy M. Hutson, and Simon L. Cornish. AC stark effect in ultracold polar ^{87}Rb^{133}Cs molecules. *Phys. Rev. A*, 96:021402, Aug 2017.

[99] L. Caldwell, J. A. Devlin, H. J. Williams, N. J. Fitch, E. A. Hinds, B. E. Sauer, and M. R. Tarbutt. Deep laser cooling and efficient magnetic compression of molecules. *Phys. Rev. Lett.*, 123:033202, Jul 2019.

[100] J.H. Eberly L. Allen. *Optical Resonance and Two-Level Atoms.* Wiley, 1974.

[101] R. G. Unanyan, N. V. Vitanov, and K. Bergmann. Preparation of entangled states by adiabatic passage. *Phys. Rev. Lett.*, 87:137902, Sep 2001.

[102] M V Berry. Transitionless quantum driving. *Journal of Physics A: Mathematical and Theoretical*, 42(36):365303, aug 2009.

[103] Xi Chen, A. Ruschhaupt, S. Schmidt, A. del Campo, D. Guéry-Odelin, and J. G. Muga. Fast optimal frictionless atom cooling in harmonic traps: Shortcut to adiabaticity. *Phys. Rev. Lett.*, 104:063002, Feb 2010.

[104] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga. Shortcuts to adiabaticity: Concepts, methods, and applications. *Rev. Mod. Phys.*, 91:045001, Oct 2019.

[105] H. R. Lewis and W. B. Riesenfeld. An exact quantum theory of the time-dependent harmonic oscillator and of a charged particle in a time-dependent electromagnetic field. *Journal of Mathematical Physics*, 10(8):1458–1473, 1969.

[106] Koushik Paul and Amarendra K. Sarma. High-fidelity entangled Bell states via shortcuts to adiabaticity. *Phys. Rev. A*, 94:052303, Nov 2016.

[107] A Ruschhaupt, Xi Chen, D Alonso, and J G Muga. Optimally robust shortcuts to population inversion in two-level quantum systems. *New Journal of Physics*, 14(9):093040, sep 2012.

[108] Xiao-Jing Lu, Xi Chen, A. Ruschhaupt, D. Alonso, S. Guérin, and J. G. Muga. Fast and robust population transfer in two-level quantum systems with dephasing noise and/or systematic frequency errors. *Phys. Rev. A*, 88:033406, Sep 2013.

[109] Xiao-Tong Yu, Qi Zhang, Yue Ban, and Xi Chen. Fast and robust control of two interacting spins. *Phys. Rev. A*, 97:062317, Jun 2018.

[110] Dionisis Stefanatos and Emmanuel Paspalakis. Efficient generation of the triplet Bell state between coupled spins using transitionless quantum driving and optimal control. *Phys. Rev. A*, 99:022327, Feb 2019.

[111] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser. Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms. *Journal of Magnetic Resonance*, 172(2):296–305, 2005.

[112] Gaurav Bhole and Jonathan A. Jones. Practical pulse engineering: Gradient ascent without matrix exponentiation. *Frontiers of Physics*, 13(3):130312, May 2018.

[113] Benjamin Rowland and Jonathan A. Jones. Implementing quantum logic gates with gradient ascent pulse engineering: principles and practicalities. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 370(1976):4636–4650, 2012.

[114] E. Knill. Quantum computing with realistically noisy devices. *Nature*, 434(7029):39–44, Mar 2005.

[115] H.J. Hogben, M. Krzystyniak, G.T.P. Charnock, P.J. Hore, and Ilya Kuprov. Spinach - a software library for simulation of spin dynamics in large spin systems. *Journal of Magnetic Resonance*, 208(2):179–194, 2011.

[116] P. de Fouquieres, S.G. Schirmer, S.J. Glaser, and Ilya Kuprov. Second order gradient ascent pulse engineering. *Journal of Magnetic Resonance*, 212(2):412–417, 2011.

[117] F. Robicheaux, T. M. Graham, and M. Saffman. Photon-recoil and laser-focusing limits to Rydberg gate fidelity. *Phys. Rev. A*, 103:022424, Feb 2021.

[118] Philip D. Gregory, Jesus Aldegunde, Jeremy M. Hutson, and Simon L. Cornish. Controlling the rotational and hyperfine state of ultracold ^{87}Rb^{133}Cs molecules. *Phys. Rev. A*, 94:041403, Oct 2016.

[119] K. Mishima and K. Yamashita. Quantum computing using rotational modes of two polar molecules. *Chemical Physics*, 361(1):106–117, 2009.

[120] P. Pellegrini and M. Desouter-Lecomte. Quantum gates driven by microwave pulses in hyperfine levels of ultracold heteronuclear dimers. *The European Physical Journal D*, 64(1):163, Jul 2011.

[121] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas. Fast quantum logic gates with trapped-ion qubits. *Nature*, 555(7694):75–78, Mar 2018.

[122] Shaowei Li, Anthony D. Castellano, Shiyu Wang, Yulin Wu, Ming Gong, Zhiguang Yan, Hao Rong, Hui Deng, Chen Zha, Cheng Guo, Lihua Sun, Chengzhi Peng, Xiaobo Zhu, and Jian-Wei Pan. Realisation of high-fidelity nonadiabatic CZ gates with superconducting qubits. *npj Quantum Information*, 5(1):84, Oct 2019.

[123] C. H. Yang, K. W. Chan, R. Harper, W. Huang, T. Evans, J. C. C. Hwang, B. Hensen, A. Laucht, T. Tanttu, F. E. Hudson, S. T. Flammia, K. M. Itoh, A. Morello, S. D. Bartlett, and A. S. Dzurak. Silicon qubit fidelities approaching incoherent noise limits via pulse engineering. *Nature Electronics*, 2(4):151–158, Apr 2019.

[124] Florian Dolde, Ville Bergholm, Ya Wang, Ingmar Jakobi, Boris Naydenov, Sébastien Pezzagna, Jan Meijer, Fedor Jelezko, Philipp Neumann, Thomas Schulte-Herbrüggen, Jacob Biamonte, and Jörg Wrachtrup. High-fidelity spin entanglement using optimal control. *Nature Communications*, 5(1):3371, Feb 2014.

[125] Michael J. Bremner, Christopher M. Dawson, Jennifer L. Dodd, Alexei Gilchrist, Aram W. Harrow, Duncan Mortimer, Michael A. Nielsen, and Tobias J. Osborne. Practical scheme for quantum computation with any two-qubit entangling gate. *Phys. Rev. Lett.*, 89:247902, Nov 2002.

[126] D. Kielpinski, C. Monroe, and D. J. Wineland. Architecture for a large-scale ion-trap quantum computer. *Nature*, 417(6890):709–711, Jun 2002.

[127] 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, Apr 2022.

[128] Yang Wang, Xianli Zhang, Theodore A. Corcovilos, Aishwarya Kumar, and David S. Weiss. Coherent addressing of individual neutral atoms in a 3D optical lattice. *Phys. Rev. Lett.*, 115:043003, Jul 2015.

[129] Yang Wang, Aishwarya Kumar, Tsung-Yao Wu, and David S. Weiss. Single-qubit gates based on targeted phase shifts in a 3d neutral atom array. *Science*, 352(6293):1562–1565, 2016.

[130] L. Caldwell and M. R. Tarbutt. Enhancing dipolar interactions between molecules using state-dependent optical tweezer traps. *Phys. Rev. Lett.*, 125:243201, Dec 2020.

[131] L. Caldwell and M. R. Tarbutt. General approach to state-dependent optical-tweezer traps for polar molecules. *Phys. Rev. Research*, 3:013291, Mar 2021.

[132] Kenneth Wang, Conner P. Williams, Lewis R. B. Picard, Norman Y. Yao, and Kang-Kuen Ni. Enriching the quantum toolbox of ultracold molecules with Rydberg atoms, 2022.

[133] Chi Zhang and M. R. Tarbutt. Quantum computation in a hybrid array of molecules and Rydberg atoms, 2022.

[134] J. Hubbard and Brian Hilton Flowers. Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 276(1365):238–257, 1963.

[135] Matthew P. A. Fisher, Peter B. Weichman, G. Grinstein, and Daniel S. Fisher. Boson localization and the superfluid-insulator transition. *Phys. Rev. B*, 40:546–570, Jul 1989.

[136] D. Jaksch, C. Bruder, J. I. Cirac, C. W. Gardiner, and P. Zoller. Cold bosonic atoms in optical lattices. *Phys. Rev. Lett.*, 81:3108–3111, Oct 1998.

[137] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. Holt-Saunders, 1976.

[138] Gregory H. Wannier. The structure of electronic excitation levels in insulating crystals. *Phys. Rev.*, 52:191–197, Aug 1937.

[139] W. Kohn. Analytic properties of Bloch waves and Wannier functions. *Phys. Rev.*, 115:809–821, Aug 1959.

[140] Nicola Marzari and David Vanderbilt. Maximally localized generalized Wannier functions for composite energy bands. *Phys. Rev. B*, 56:12847–12865, Nov 1997.

[141] R. Walters, G. Cotugno, T. H. Johnson, S. R. Clark, and D. Jaksch. Ab initio derivation of Hubbard models for cold atoms in optical lattices. *Phys. Rev. A*, 87:043613, Apr 2013.

[142] Soumik Bandyopadhyay, Rukmani Bai, Sukla Pal, K. Suthar, Rejish Nath, and D. Angom. Quantum phases of canted dipolar bosons in a two-dimensional square optical lattice. *Phys. Rev. A*, 100:053623, Nov 2019.

[143] S. Inouye, M. R. Andrews, J. Stenger, H.-J. Miesner, D. M. Stamper-Kurn, and W. Ketterle. Observation of Feshbach resonances in a Bose–Einstein condensate. *Nature*, 392(6672):151–154, Mar 1998.

[144] Michael W. Jack and Makoto Yamashita. Bose-Hubbard model with attractive interactions. *Phys. Rev. A*, 71:023610, Feb 2005.

[145] Gretchen K. Campbell, Jongchul Mun, Micah Boyd, Patrick Medley, Aaron E. Leanhardt, Luis G. Marcassa, David E. Pritchard, and Wolfgang Ketterle. Imaging the Mott insulator shells by using atomic clock shifts. *Science*, 313(5787):649–652, 2006.

[146] Omjyoti Dutta, Mariusz Gajda, Philipp Hauke, Maciej Lewenstein, Dirk-Sören Lühmann, Boris A Malomed, Tomasz Sowiński, and Jakub Zakrzewski. Non-standard Hubbard models in optical lattices: a review. *Reports on Progress in Physics*, 78(6):066001, may 2015.

[147] M L Wall and L D Carr. Dipole–dipole interactions in optical lattices do not follow an inverse cube power law. *New Journal of Physics*, 15(12):123005, dec 2013.

[148] Tomasz Sowiński, Omjyoti Dutta, Philipp Hauke, Luca Tagliacozzo, and Maciej Lewenstein. Dipolar molecules in optical lattices. *Phys. Rev. Lett.*, 108:115301, Mar 2012.

[149] Michał Maik, Philipp Hauke, Omjyoti Dutta, Maciej Lewenstein, and Jakub Zakrzewski. Density-dependent tunneling in the extended Bose–Hubbard model. *New Journal of Physics*, 15(11):113041, nov 2013.

[150] Rebecca Kraus, Krzysztof Biedroń, Jakub Zakrzewski, and Giovanna Morigi. Superfluid phases induced by dipolar interactions. *Phys. Rev. B*, 101:174505, May 2020.

[151] Kuldeep Suthar, Rebecca Kraus, Hrushikesh Sable, Dilip Angom, Giovanna Morigi, and Jakub Zakrzewski. Staggered superfluid phases of dipolar bosons in two-dimensional square lattices. *Phys. Rev. B*, 102:214503, Dec 2020.

[152] Dean Johnstone, Niclas Westerberg, Callum W. Duncan, and Patrik Öhberg. Staggered ground states in an optical lattice. *Phys. Rev. A*, 100:043614, Oct 2019.

[153] Chao Zhang, Jin Zhang, Jin Yang, and Barbara Capogrosso-Sansone. Ground states of two-dimensional tilted dipolar bosons with density-induced hopping. *Phys. Rev. A*, 103:043333, Apr 2021.

[154] 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, Jan 2002.

[155] I. B. Spielman, W. D. Phillips, and J. V. Porto. Mott-insulator transition in a two-dimensional atomic Bose gas. *Phys. Rev. Lett.*, 98:080404, Feb 2007.

[156] 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.

[157] Ippei Danshita and Carlos A. R. Sá de Melo. Stability of superfluid and supersolid phases of dipolar bosons in optical lattices. *Phys. Rev. Lett.*, 103:225301, Nov 2009.

[158] Pinaki Sengupta, Leonid P. Pryadko, Fabien Alet, Matthias Troyer, and Guido Schmid. Supersolids versus phase separation in two-dimensional lattice bosons. *Phys. Rev. Lett.*, 94:207202, May 2005.

[159] Emanuele G. Dalla Torre, Erez Berg, and Ehud Altman. Hidden order in 1d Bose insulators. *Phys. Rev. Lett.*, 97:260401, Dec 2006.

[160] Xiaolong Deng and Luis Santos. Entanglement spectrum of one-dimensional extended Bose-Hubbard models. *Phys. Rev. B*, 84:085138, Aug 2011.

[161] Daisuke Yamamoto, Akiko Masaki, and Ippei Danshita. Quantum phases of hardcore bosons with long-range interactions on a square lattice. *Phys. Rev. B*, 86:054516, Aug 2012.

[162] Takahiro Ohgoe, Takafumi Suzuki, and Naoki Kawashima. Novel mechanism of supersolid of ultracold polar molecules in optical lattices. *Journal of the Physical Society of Japan*, 80(11):113001, 2011.

[163] Paolo Rosson, Martin Kiffner, Jordi Mur-Petit, and Dieter Jaksch. Characterizing the phase diagram of finite-size dipolar Bose-Hubbard systems. *Phys. Rev. A*, 101:013616, Jan 2020.

[164] Huan-Kuang Wu and Wei-Lin Tu. Competing quantum phases of hard-core bosons with tilted dipole-dipole interaction. *Phys. Rev. A*, 102:053306, Nov 2020.

[165] Jin Zhang, Chao Zhang, Jin Yang, and Barbara Capogrosso-Sansone. Supersolid phases of lattice dipoles tilted in three dimensions. *Phys. Rev. A*, 105:063302, Jun 2022.

[166] C Zhang, A Safavi-Naini, Ana Maria Rey, and B Capogrosso-Sansone. Equilibrium phases of tilted dipolar lattice bosons. *New Journal of Physics*, 17(12):123014, dec 2015.

[167] Takahiro Ohgoe, Takafumi Suzuki, and Naoki Kawashima. Quantum phases of hard-core bosons on two-dimensional lattices with anisotropic dipole-dipole interaction. *Phys. Rev. A*, 86:063635, Dec 2012.

[168] 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, apr 2009.

[169] K. C. Wright, R. B. Blakestad, C. J. Lobb, W. D. Phillips, and G. K. Campbell. Driving phase slips in a superfluid atom circuit with a rotating weak link. *Phys. Rev. Lett.*, 110:025302, Jan 2013.

[170] O. Dutta, M. Jääskeläinen, and P. Meystre. Thomas-Fermi ground state of dipolar fermions in a circular storage ring. *Phys. Rev. A*, 73:043610, Apr 2006.

[171] Sascha Zöllner, G. M. Bruun, C. J. Pethick, and S. M. Reimann. Bosonic and fermionic dipoles on a ring. *Phys. Rev. Lett.*, 107:035301, Jul 2011.

[172] M. Abad, M. Guilleumas, R. Mayol, M. Pi, and D. M. Jezek. Dipolar condensates confined in a toroidal trap: Ground state and vortices. *Phys. Rev. A*, 81:043619, Apr 2010.

[173] S. K. Adhikari. Dipolar Bose-Einstein condensate in a ring or in a shell. *Phys. Rev. A*, 85:053631, May 2012.

[174] Pedro C. Diniz, Eduardo A. B. Oliveira, Aristeu R. P. Lima, and Emanuel A. L. Henn. Ground state and collective excitations of a dipolar Bose-Einstein condensate in a bubble trap. *Scientific Reports*, 10(1):4831, Mar 2020.

[175] Maria Arazo, Ricardo Mayol, and Montserrat Guilleumas. Shell-shaped condensates with gravitational sag: contact and dipolar interactions. *New Journal of Physics*, 23(11):113040, nov 2021.

[176] Luigi Amico, Andreas Osterloh, and Francesco Cataliotti. Quantum many particle systems in ring-shaped optical lattices. *Phys. Rev. Lett.*, 95:063201, Aug 2005.

[177] S. Franke-Arnold, J. Leach, M. J. Padgett, V. E. Lembessis, D. Ellinas, A. J. Wright, J. M. Girkin, P. Öhberg, and A. S. Arnold. Optical ferris wheel for ultracold atoms. *Opt. Express*, 15(14):8619–8625, Jul 2007.

[178] Qingyang Wang, Johannes Otterbach, and Susanne F. Yelin. Interacting in-plane molecular dipoles in a zigzag chain. *Phys. Rev. A*, 96:043615, Oct 2017.

[179] Niraj R. Ghimire and Susanne F. Yelin. Frustrated plane-polarized dipoles in one dimension. *Phys. Rev. B*, 100:115110, Sep 2019.

[180] Michał Maik, Pierfrancesco Buonsante, Alessandro Vezzani, and Jakub Zakrzewski. Dipolar bosons on an optical lattice ring. *Phys. Rev. A*, 84:053615, Nov 2011.

[181] Mateusz Łącki, Hannes Pichler, Antoine Sterdyniak, Andreas Lyras, Vassilis E. Lembessis, Omar Al-Dossary, Jan Carl Budich, and Peter Zoller. Quantum Hall physics with cold atoms in cylindrical optical lattices. *Phys. Rev. A*, 93:013604, Jan 2016.

[182] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69:2863–2866, Nov 1992.

[183] Steven R. White. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B*, 48:10345–10356, Oct 1993.

[184] Stellan Östlund and Stefan Rommer. Thermodynamic limit of density matrix renormalization. *Phys. Rev. Lett.*, 75:3537–3540, Nov 1995.

[185] Marcos Rigol and Alejandro Muramatsu. Universal properties of hard-core bosons confined on one-dimensional lattices. *Phys. Rev. A*, 70:031603, Sep 2004.

[186] Rhine Samajdar, Wen Wei Ho, Hannes Pichler, Mikhail D. Lukin, and Subir Sachdev. Complex density wave orders and quantum phase transitions in a model of square-lattice Rydberg atom arrays. *Phys. Rev. Lett.*, 124:103601, Mar 2020.

[187] Rhine Samajdar, Wen Wei Ho, Hannes Pichler, Mikhail D. Lukin, and Subir Sachdev. Quantum phases of Rydberg atoms on a Kagomé lattice. *Proceedings of the National Academy of Sciences*, 118(4), 2021.

[188] A. Osterloh, Luigi Amico, G. Falci, and Rosario Fazio. Scaling of entanglement close to a quantum phase transition. *Nature*, 416(6881):608–610, Apr 2002.

[189] M. Kargarian, R. Jafari, and A. Langari. Renormalization of entanglement in the anisotropic Heisenberg (xxz) model. *Phys. Rev. A*, 77:032346, Mar 2008.

[190] Luigi Amico, Rosario Fazio, Andreas Osterloh, and Vlatko Vedral. Entanglement in many-body systems. *Rev. Mod. Phys.*, 80:517–576, May 2008.

[191] Christian Romen and Andreas M. Läuchli. Chiral Mott insulators in frustrated Bose-Hubbard models on ladders and two-dimensional lattices: A combined

perturbative and density matrix renormalization group study. *Phys. Rev. B*, 98:054519, Aug 2018.

[192] 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, Dec 2015.

[193] Adam M. Kaufman, M. Eric Tai, Alexander Lukin, Matthew Rispoli, Robert Schittko, Philipp M. Preiss, and Markus Greiner. Quantum thermalization through entanglement in an isolated many-body system. *Science*, 353(6301):794–800, 2016.

[194] E.M. Stoudenmire and Steven R. White. Studying two-dimensional systems with the density matrix renormalization group. *Annual Review of Condensed Matter Physics*, 3(1):111–128, 2012.

[195] Huaiming Guo, Yuchuan Wen, and Shiping Feng. Cold atoms on a two-dimensional square optical lattice with an alternating potential. *Phys. Rev. A*, 79:035401, Mar 2009.

[196] Xiao-Juan Li and Yu-Chuan Wen. Hard-core bosons on a two-dimensional square optical superlattice. *Communications in Theoretical Physics*, 63(3):308–316, mar 2015.

[197] Itay Hen, M. Iskin, and Marcos Rigol. Phase diagram of the hard-core Bose-Hubbard model on a checkerboard superlattice. *Phys. Rev. B*, 81:064503, Feb 2010.

[198] Marco Cozzini, Radu Ionicioiu, and Paolo Zanardi. Quantum fidelity and quantum phase transitions in matrix product states. *Phys. Rev. B*, 76:104420, Sep 2007.

[199] Huan-Qiang Zhou and John Paul Barjaktarevič. Fidelity and quantum phase transitions. *J. Phys. A: Math. Theor*, 41(41):412001, sep 2008.

[200] Shi-Jian Gu. Fidelity approach to quantum phase transitions. *International Journal of Modern Physics B*, 24(23):4371–4458, 2010.

[201] N. Schuch, F. Verstraete, and J. I. Cirac. Nonlocal resources in the presence of superselection rules. *Phys. Rev. Lett.*, 92:087904, Feb 2004.

[202] K. Binder. Finite size scaling analysis of Ising model block distribution functions. *Zeitschrift für Physik B Condensed Matter*, 43(2):119–140, Jun 1981.

[203] Somendra M Bhattacharjee and Flavio Seno. A measure of data collapse for scaling. *Journal of Physics A: Mathematical and General*, 34(33):6375–6380, aug 2001.

[204] Géza Ódor. Universality classes in nonequilibrium lattice systems. *Rev. Mod. Phys.*, 76:663–724, Aug 2004.

[205] Wen-Long You, Ying-Wai Li, and Shi-Jian Gu. Fidelity, dynamic structure factor, and susceptibility in critical phenomena. *Phys. Rev. E*, 76:022101, Aug 2007.

[206] Jin Zhang. Fidelity and entanglement entropy for infinite-order phase transitions. *Phys. Rev. B*, 104:205112, Nov 2021.

[207] David Schwandt, Fabien Alet, and Sylvain Capponi. Quantum Monte Carlo simulations of fidelity at magnetic quantum phase transitions. *Phys. Rev. Lett.*, 103:170501, Oct 2009.

[208] A. Fabricio Albuquerque, Fabien Alet, Clément Sire, and Sylvain Capponi. Quantum critical scaling of fidelity susceptibility. *Phys. Rev. B*, 81:064418, Feb 2010.

[209] Silvano Garnerone, Damian Abasto, Stephan Haas, and Paolo Zanardi. Fidelity in topological quantum phases of matter. *Phys. Rev. A*, 79:032302, Mar 2009.

[210] G. Sun, A. K. Kolezhuk, and T. Vekua. Fidelity at Berezinskii-Kosterlitz-Thouless quantum phase transitions. *Phys. Rev. B*, 91:014418, Jan 2015.

[211] Andreas M Läuchli and Corinna Kollath. Spreading of correlations and entanglement after a quench in the one-dimensional Bose–Hubbard model. *Journal of Statistical Mechanics: Theory and Experiment*, 2008(05):P05018, may 2008.

[212] G. Vidal, J. I. Latorre, E. Rico, and A. Kitaev. Entanglement in quantum critical phenomena. *Phys. Rev. Lett.*, 90:227902, Jun 2003.

[213] J.M. Kurdestany, R.V. Pai, and R. Pandit. The inhomogeneous extended Bose-Hubbard model: A mean-field theory. *Annalen der Physik*, 524(3-4):234–244, 2012.

[214] Dean Johnstone, Patrik Öhberg, and Callum W. Duncan. Mean-field phases of an ultracold gas in a quasicrystalline potential. *Phys. Rev. A*, 100:053609, Nov 2019.

[215] A. Frisch, M. Mark, K. Aikawa, S. Baier, R. Grimm, A. Petrov, S. Kotochigova, G. Quéméner, M. Lepers, O. Dulieu, and F. Ferlaino. Ultracold dipolar molecules composed of strongly magnetic atoms. *Phys. Rev. Lett.*, 115:203201, Nov 2015.

[216] Aden Z. Lam, Niccolò Bigagli, Claire Warner, Weijun Yuan, Siwei Zhang, Eberhard Tiemann, Ian Stevenson, and Sebastian Will. High phase-space density gas of NaCs Feshbach molecules. *Phys. Rev. Research*, 4:L022019, Apr 2022.

[217] Per Bak and R. Bruinsma. One-dimensional Ising model and the complete devil's staircase. *Phys. Rev. Lett.*, 49:249–251, Jul 1982.

[218] F. J. Burnell, Meera M. Parish, N. R. Cooper, and S. L. Sondhi. Devil's staircases and supersolids in a one-dimensional dipolar Bose gas. *Phys. Rev. B*, 80:174519, Nov 2009.

[219] Philipp Hauke, Fernando M Cucchietti, Alexander Müller-Hermes, Mari-Carmen Bañuls, J Ignacio Cirac, and Maciej Lewenstein. Complete devil's staircase and crystal–superfluid transitions in a dipolar XXZ spin chain: a trapped ion quantum simulation. *New Journal of Physics*, 12(11):113037, nov 2010.

[220] Haiyuan Zou, Erhai Zhao, Xi-Wen Guan, and W. Vincent Liu. Exactly solvable points and symmetry protected topological phases of quantum spins on a zig-zag lattice. *Phys. Rev. Lett.*, 122:180401, May 2019.

[221] G. Möller and N. R. Cooper. Composite fermion theory for bosonic quantum Hall states on lattices. *Phys. Rev. Lett.*, 103:105303, Sep 2009.

[222] K. Suthar, Hrushikesh Sable, Rukmani Bai, Soumik Bandyopadhyay, Sukla Pal, and D. Angom. Supersolid phase of the extended Bose-Hubbard model with an artificial gauge field. *Phys. Rev. A*, 102:013320, Jul 2020.

[223] Paolo Rosson, Michael Lubasch, Martin Kiffner, and Dieter Jaksch. Bosonic fractional quantum Hall states on a finite cylinder. *Phys. Rev. A*, 99:033603, Mar 2019.

[224] Tymoteusz Salamon, Alessio Celi, Ravindra W. Chhajlany, Irénée Frérot, Maciej Lewenstein, Leticia Tarruell, and Debraj Rakshit. Simulating twistronics without a twist. *Phys. Rev. Lett.*, 125:030504, Jul 2020.

[225] 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, Sep 2018.

[226] Benjamin M. Spar, Elmer Guardado-Sanchez, Sungjae Chi, Zoe Z. Yan, and Waseem S. Bakr. Realization of a Fermi-Hubbard optical tweezer array. *Phys. Rev. Lett.*, 128:223202, Jun 2022.

[227] Michael Hughes, Axel U. J. Lode, Dieter Jaksch, and Paolo Molignini. Accuracy of quantum simulators with ultracold dipolar molecules: a quantitative comparison between continuum and lattice descriptions, 2022.

[228] Andrew Richards. *University of Oxford Advanced Research Computing*, August 2015. doi:10.5281/zenodo.22558.

[229] Kaspar Sakmann, Alexej I. Streltsov, Ofir E. Alon, and Lorenz S. Cederbaum. Exact quantum dynamics of a Bosonic Josephson junction. *Phys. Rev. Lett.*, 103:220601, Nov 2009.

[230] Jan Major, Mateusz Łącki, and Jakub Zakrzewski. Reexamination of the variational Bose-Hubbard model. *Phys. Rev. A*, 89:043626, Apr 2014.

[231] Shovan Dutta, Anton Buyskikh, Andrew J. Daley, and Erich J. Mueller. Density matrix renormalization group for continuous quantum systems. *Phys. Rev. Lett.*, 128:230401, Jun 2022.

[232] J M Zhang and R X Dong. Exact diagonalization: the Bose–Hubbard model as an example. *European Journal of Physics*, 31(3):591–602, apr 2010.

[233] Phillip Weinberg and Marin Bukov. QuSpin: a Python Package for Dynamics and Exact Diagonalisation of Quantum Many Body Systems. Part II: bosons, fermions and higher spins. *SciPost Phys.*, 7:20, 2019.

[234] Alexej I. Streltsov, Ofir E. Alon, and Lorenz S. Cederbaum. General variational many-body theory with complete self-consistency for trapped bosonic systems. *Phys. Rev. A*, 73:063626, Jun 2006.

[235] Ofir E. Alon, Alexej I. Streltsov, and Lorenz S. Cederbaum. Multiconfigurational time-dependent Hartree method for bosons: Many-body dynamics of bosonic systems. *Phys. Rev. A*, 77:033613, Mar 2008.

[236] Rui Lin, Paolo Molignini, Luca Papariello, Marios C Tsatsos, Camille Lévêque, Storm E Weiner, Elke Fasshauer, R Chitra, and Axel U J Lode. MCTDH-x: The multiconfigurational time-dependent Hartree method for indistinguishable particles software. *Quantum Science and Technology*, 5(2):024004, mar 2020.

[237] Camille Lévêque and Lars Bojer Madsen. Time-dependent restricted-active-space self-consistent-field theory for bosonic many-body systems. *New Journal of Physics*, 19(4):043007, apr 2017.

[238] Budhaditya Chatterjee, Marios C Tsatsos, and Axel U J Lode. Correlations of strongly interacting one-dimensional ultracold dipolar few-boson systems in optical lattices. *New Journal of Physics*, 21(3):033030, mar 2019.

[239] Budhaditya Chatterjee, Camille Lévêque, Jörg Schmiedmayer, and Axel U. J. Lode. Detecting one-dimensional dipolar bosonic crystal orders via full distribution functions. *Phys. Rev. Lett.*, 125:093602, Aug 2020.

[240] S. Sinha and L. Santos. Cold dipolar gases in quasi-one-dimensional geometries. *Phys. Rev. Lett.*, 99:140406, Oct 2007.

[241] Ulf Bissbort, Frank Deuretzbacher, and Walter Hofstetter. Effective multibody-induced tunneling and interactions in the Bose-Hubbard model of the lowest dressed band of an optical lattice. *Phys. Rev. A*, 86:023617, Aug 2012.

[242] Mark Kremer, Rashi Sachdeva, Albert Benseny, and Thomas Busch. Interaction-induced effects on Bose-Hubbard parameters. *Phys. Rev. A*, 96:063611, Dec 2017.

[243] Ian Stevenson, Aden Z. Lam, Niccolò Bigagli, Claire Warner, Weijun Yuan, Siwei Zhang, and Sebastian Will. Ultracold gas of dipolar NaCs ground state molecules, 2022.

[244] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen. Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework. *Phys. Rev. A*, 84:022305, Aug 2011.

[245] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, 2011. January 2011 Special Issue.

[246] S Al-Assam, S R Clark, and D Jaksch. The tensor network theory library. *Journal of Statistical Mechanics: Theory and Experiment*, 2017(9):093102, sep 2017.

[247] M B Hastings. An area law for one-dimensional quantum systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2007(08):P08024–P08024, aug 2007.

[248] J. Eisert, M. Cramer, and M. B. Plenio. Colloquium: Area laws for the entanglement entropy. *Rev. Mod. Phys.*, 82:277–306, Feb 2010.

[249] Gregory M. Crosswhite and Dave Bacon. Finite automata for caching in matrix product algorithms. *Phys. Rev. A*, 78:012356, Jul 2008.

[250] Jonathan Coulthard. *Engineering quantum states of fermionic many-body systems*. PhD thesis, University of Oxford, 2018.

[251] Steven R. White. Density matrix renormalization group algorithms with a single center site. *Phys. Rev. B*, 72:180403, Nov 2005.

[252] C. Hubig, I. P. McCulloch, U. Schollwöck, and F. A. Wolf. Strictly single-site DMRG algorithm with subspace expansion. *Phys. Rev. B*, 91:155115, Apr 2015.

[253] Shoudan Liang and Hanbin Pang. Approximate diagonalization using the density matrix renormalization-group method: A two-dimensional-systems perspective. *Phys. Rev. B*, 49:9214–9217, Apr 1994.

A

Ladder Operator Derivation

While the positional-basis numerical process was used to calculate the Hamiltonian elements of the dipole-dipole interaction that were used for numerical evolution, an alternative process to derive these matrix elements when $\delta f_{trap} << f_{trap}$ offers more physical insight for one dimension. The idea is to express the position operator in the dipole-dipole interaction using the raising and lowering operators of the two traps. The operator for the distance between the two molecules is

$$\hat{R} = R_e(\hat{1} + \frac{\hat{r}_2 - \hat{r}_1}{R_e}), \tag{A.1}$$

where $\hat{1}$ is the unit operator. As the harmonic oscillator length is much smaller than the equilibrium separation of the tweezers, the dimensionless operator $\hat{s} = \frac{\hat{r}_2 - \hat{r}_1}{R_e}$ has eigenvalues much smaller than 1. This means that $\hat{R}^{-3}$ can be found using the standard binomial expansion for negative integer powers

$$\hat{R}^{-3} = \frac{1}{R_e^3} \sum_{k=0}^{\infty} \frac{(-1)^k (2+k)!}{k!2} \hat{s}^k, \tag{A.2}$$

where due to the small eigenvalues of $\hat{s}$, the terms with higher values of the index k can be neglected. This operator can then be evaluated explicitly using the formula

$$\hat{s} = \frac{1}{\sqrt{2}} \frac{l_{ho,2}(\hat{a}_2 + \hat{a}_2^\dagger) - l_{ho,1}(\hat{a}_1 + \hat{a}_1^\dagger)}{R_e} \tag{A.3}$$

As only small powers of $\hat{s}$ are necessary for convergence, the action of the operator $\hat{R}^{-3}$ on the three considered motional states can be calculated using a finite number of motional states. Note that this derivation has treated the harmonic oscillator length l_{ho} as separate from the internal state which is not quite accurate due to the difference in trap frequency for $|0\rangle$ and $|1\rangle$, although comparison with the position space derivation confirms that the effect is very small for our parameters.

Having evaluated the 'motional' part of the dipole-dipole interaction in this manner, there is improved physical intuition behind the matrix elements of the full Hamiltonian calculated using the positional wavefunctions. The $k = 0$ term, which is independent of l_{ho}, gives the dipole-dipole interaction between point particles at R_e for all motional states ($h \times 1847$Hz). The $k = 1$ term gives elements which create or destroy one motional excitation (i.e. transfer $|gg\rangle \leftrightarrow |ge\rangle, |eg\rangle$) and has magnitude of roughly $h \times 145$Hz. The $k = 2$ term causes two distinct physical effects: The first is to increase the dipole-dipole interaction within a given motional state due to the variance of the position wavefunction. The second is processes which transfer a motional excitation between the traps ($|ge\rangle \leftrightarrow |eg\rangle$). Both of these contributions from $k = 2$ have magnitude of around $h \times 15$Hz. Higher order terms make small adjustments to these elements, with energy decreasing by around a factor of 10 for each successive value of the index k. This gives an approximate guide to how the dipole-dipole interaction elements would be altered by changing the relative values of the harmonic oscillator lengths and the equilibrium tweezer separation, as the magnitude of $\hat{s}$ is determined by the ratio $\frac{l_{ho}}{R_e}$.

B

Shortcut to Adiabaticity Derivation

The LRI inverse-engineering method is used to accelerate the adiabatic crossing (crossing B), which involves the basis states $|00\rangle$ and $|\psi_+\rangle$. Within this subspace, the Hamiltonian is

$$H_B(t) = -\hbar\Delta(t)\,|00\rangle\,\langle 00| + \hbar\frac{\Omega(t)}{\sqrt{2}}(|00\rangle\,\langle\psi_+| + |\psi_+\rangle\,\langle 00|) - V\,|\psi_+\rangle\,\langle\psi_+|\,. \tag{B.1}$$

The effect of V is to shift crossing B from $\Delta = 0$ to $\Delta = V/\hbar$. We can therefore take this energy shift into account by defining $\bar{\Delta} = \Delta - V/\hbar$. The form of the Hamiltonian in this subspace guarantees that we can write a traceless and Hermitian 2×2 invariant, which has the following form [107]:

$$\begin{aligned} I(t) \propto{} & \sin(\gamma(t))\cos(\beta(t))(|00\rangle\,\langle\psi_+| + |\psi_+\rangle\,\langle 00|) - \sin(\gamma(t))\sin(\beta(t))(-i\,|00\rangle\,\langle\psi_+| + i\,|\psi_+\rangle\,\langle 00|) \\ & + \cos(\gamma(t))(|00\rangle\,\langle 00| - |\psi_+\rangle\,\langle\psi_+|), \end{aligned} \tag{B.2}$$

where β and γ are real angles and the constant of proportionality does not change in time. To keep the state of the system in an eigenstate of the invariant, we require that the invariant does not evolve in the Heisenberg picture, i.e.

$$\frac{dI(t)}{dt} = \frac{\partial I(t)}{\partial t} - i[I(t), H_B(t)] = 0. \tag{B.3}$$

This ensures that the expectation value of $I(t)$ will be conserved for any initial state. As $I(t)$ has two constant, non-degenerate eigenvalues, this means that if the initial state is an eigenstate of $I(t)$, the final state must be the eigenstate of $I(t)$ corresponding to the same eigenvalue. Substituting the definitions of $I(t)$ and $H_B(t)$ into this equation gives the following differential equations that link the invariant parameters $\beta(t)$ and $\gamma(t)$ with the Hamiltonian controls $\bar{\Delta}(t)$ and $\Omega(t)$:

$$\begin{aligned} \dot{\gamma}(t) &= \sqrt{2}\Omega(t)\sin(\beta(t)) \\ \bar{\Delta}(t) + \dot{\beta}(t) &= \sqrt{2}\Omega(t)\cot(\gamma(t))\cos(\beta(t)) \end{aligned} \tag{B.4}$$

Having guaranteed that the populations of the eigenstates of $I(t)$ will remain constant during the evolution, we now need to ensure that $|00\rangle$ is one of the eigenstates of the invariant at the start of the protocol $I(t_0)$ and that this same eigenstate becomes $|\psi_+\rangle$ by the end of the protocol (t_f). This creates the following boundary conditions:

$$\begin{aligned} \gamma(t_i) &= 0 \\ \gamma(t_f) &= \pi \\ \dot{\gamma}(t_i) &= 0 \\ \dot{\gamma}(t_f) &= 0 \end{aligned} \tag{B.5}$$

(To avoid singularities in equation B.4, we slightly displaced $\gamma(t_i)$ and $\gamma(t_f)$ in our numerics). The form of $\beta(t)$ is not defined by these conditions but it is desirable to reduce the maximum value of $\Omega(t)$ to avoid populating the third state $|11\rangle$ that we ignored in H_B. For this reason, we used the following boundary conditions for $\beta(t)$ [1]

$$\begin{aligned} \beta(t_i) &= \frac{\pi}{2} \\ \beta(t_f) &= \frac{\pi}{2} \\ \dot{\beta}(t_i) &= \frac{-3\pi}{2t_f} \\ \dot{\beta}(t_f) &= \frac{3\pi}{2t_f} \end{aligned} \tag{B.6}$$

---

[1] We multiplied the values of $\dot{\beta}(t_i)$ and $\dot{\beta}(t_f)$ by a factor of $\frac{3}{2}$ compared to Ref [106] to ensure that the protocol also traversed the 'diabatic' crossing A.

Using the ansatz that $\gamma(t)$ and $\beta(t)$ are fourth-order and fifth-order polynomials in time respectively [106], the constraints in equations B.5 and B.6 can be solved using linear algebra, leaving one free parameter for $\gamma(t)$ and two free parameters for $\beta(t)$. We set these parameters by numerically minimising the quantity

$$q_L = \int_{t_i}^{t_f} |\frac{\Omega(t)^2}{\bar{\Delta}(t) + \frac{V}{\hbar}}|dt, \tag{B.7}$$

which is a figure of merit for the population leakage into $|11\rangle$ at the diabatic crossing. This optimisation delivers a small improvement in fidelity compared to the protocol in Ref. [106] which was derived without considering leakage into $|11\rangle$ and pushes the maximum value of $\Omega(t)$ to coincide with larger values of $\Delta(t)$. $H(t)$ is then found numerically from $\gamma(t)$ and $\beta(t)$ using equation B.4.

C
GRAPE Method

The GRAPE algorithm [111] starts with a chosen or, as we used for the entangling gate, random set of control parameters. These are iteratively optimised using gradient methods to find a local maximum of the fidelity. To do this, one must calculate the gradient of the overall gate fidelity with respect to the control parameters at each timestep. The algorithm does this efficiently by storing at each timestep j the unitary matrices corresponding to both forward propagation of the time evolution operator up to timestep j and reverse propagation of the target operator U_{aim} back through the subsequent timesteps to timestep j, which reduces the number of computationally-costly calculations of time evolution operators . Mathematically, this means storing the following two matrices for each timestep j:

$$F_j = U_j ... U_2 U_1 \tag{C.1}$$

$$R_j = U_{j+1}^{\dagger} ... U_{N_{steps}}^{\dagger} U_{aim} \tag{C.2}$$

where U_j is the evolution operator for timestep j alone.

Using a first order approximation for the evolution during timestep j (which requires that the eigenvalues of $H_j\delta_t << 1$), the gradient of the overall fidelity with respect to the control field at timestep j is calculated

$$\frac{\partial F}{\partial \Omega_{x(y),j}} = \frac{-2}{D^2}\,\mathrm{Re}(\mathrm{Tr}(R_j^\dagger(i\delta_t h_{\Omega_{X(Y)}})F_j)Tr(F_j^\dagger R_j)) \tag{C.3}$$

These gradients are then used to update the control parameters at each iteration using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) numerical method, which additionally uses the gradient history to aid the optimisation (known as BFGS-GRAPE), which is particularly beneficial when approaching a local maximum of the fidelity, where the gradient of the fidelity becomes very small [116, 244]. By iterating this procedure, a local maximum of the fidelity is found.

D

Effect of Trap Frequencies in One Dimension

While we optimised and simulated the entangling gate using specific estimated values for the trap frequency f_{trap} and the additional trap frequency δf_{trap} for internal state $|1\rangle$, fine-tuning these values is not required for the success of the gate as long as the exchange of motional excitations between the traps is made off-resonant. Qualitatively, this is because the DDI is not strongly-dependent on the trap frequency and the effects of δf_{trap} (increased effective detuning and small leakage into the $|\psi_-\rangle$ internal state due to the off-diagonal coupling element $\frac{\delta f_{trap}}{2}$ between $|\psi_+\rangle$ and $|\psi_-\rangle$) are specifically corrected by the optimisation and suppressed by the microwave driving respectively.

Quantitatively, Figure D.1 shows the fidelity for the pulse shown in Figure 2.7 considering motional states $|gg\rangle$, $|ge\rangle$, and $|eg\rangle$ in the x direction for varying parameters of the second trap. The fidelity is largely unaffected by these parameters provided that the frequency of the second trap is not either much smaller than assumed or extremely close to that of the first. When considering motion in three dimensions, changes in the trap frequency in the tweezer axis direction have the largest effect on the fidelity because this trap frequency is smaller.

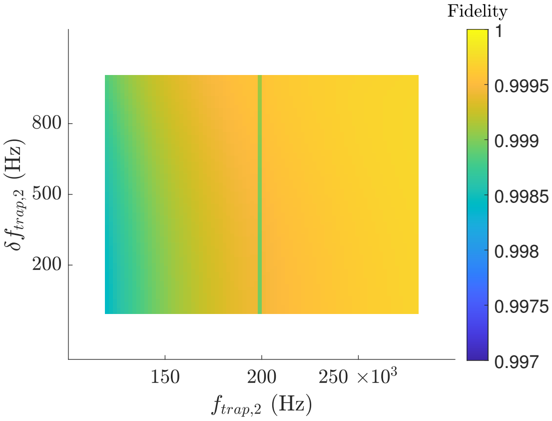

Figure D.1: Fidelity of pulse shown in Figure 2.7 for detuning $\Delta = 0$, systematic microwave power scaling factor $f_\Omega = 1$, trap frequency for internal state $|0\rangle$ of the first molecule $f_{trap,1} = 200$kHz, and extra trap frequency for the internal state $|1\rangle$ of the first molecule $\delta f_{trap,1} = 500$Hz considering up to one motional excitation in the x direction.

E

DMRG for Dipolar Systems

This appendix introduces the important considerations of using the DMRG numerical methods used in chapter 4 to study the cylindrical dipolar Bose-Hubbard model rather than the numerical implementation, which is presented in Refs [245, 246] for all cases which we cover. We firstly introduce matrix product states, matrix product operators, and DMRG themselves before discussing their use in two dimensions.

E.1 Matrix Product States

One of the main reasons for the interest in controlled quantum many-body systems (which applies for both quantum computation and quantum simulation) is the exponential growth of the Hilbert space with system size. This growth creates difficulty for exact simulations of quantum many-body systems because even moderately-sized systems of ≈ 50 particles or lattice sites would require impossible numerical resources to specify the state of the system classically. There are many quantum many-body methods to reduce the computational cost by describing the physics approximately instead, of which matrix product states (MPS) have been one of the most successful, especially in 1D.

E.1.1 Ansatz

To define the MPS form and describe its approximations, we use the example of a system of L sites (indexed by x) which have a local Hilbert space of dimension d (indexed by j), where we assume d is the same for all sites (as is satisfied for all results in this thesis). The MPS form places a tensor A at each site which has one 'physical' index j corresponding to the local Hilbert space and two 'virtual' indices (the tensors A_1 and A_L corresponding to the sites at the ends of the system have one virtual index instead). For a given choice of the L different local physical indices j_x, the tensors A_x are reduced to matrices (or row and column vectors at the ends of the system) $A_x^{j_x}$ which have only the virtual indices. The wavefunction amplitude corresponding to the physical state chosen is then calculated by contracting the product of a row vector, $L-2$ matrices, and a column vector as shown mathematically in the following formula:

$$\langle j_1, j_2, ...j_L|\psi\rangle = A_1^{j_1} A_2^{j_2} ... A_L^{j_L}. \tag{E.1}$$

The number of degrees of freedom within the MPS ansatz depends on the size of the virtual matrices and vectors (i.e. maximum value χ_x of the virtual indices) associated with a given state in the local Hilbert space. The virtual indices allow entanglement to be represented and therefore encode non-classical correlations. This ansatz can be made exact by using large enough χ_x, although the computational cost would still scale exponentially with the system size and offers no scaling benefit compared to the exact representation. Instead, the ansatz is usually used by choosing a lower maximum value of χ for all sites as a computationally-cheaper controlled approximation. This means the wavefunction is represented by $\leq Ld\chi^2$ complex numbers, which scales linearly with L, compared with d^L complex coefficients for the exact wavefunction.

Turning to the question of how useful MPS with a chosen value of χ are, physical states often have a more local entanglement structure than an arbitrary state in the exponentially-scaling Hilbert space because physical interactions are generally local. This implies that the physics of a local Hamiltonian only uses a small fraction of the Hilbert space. MPS target a physically-relevant corner of the Hilbert space, which satisfies area laws of the Von Neumann entanglement entropy S_{ent} (defined in equation 4.4).

Specifically, the maximum S_{ent} across a bipartition of an MPS for a given value of χ is $\ln(\chi)$ [1]. This follows the one-dimensional version of the area law, where the entanglement entropy of a contiguous bipartition of a quantum state is proportional to the area of the boundary of the subsystems rather than the volume of the smaller subsystem. This law is proven for the low-lying eigenstates of gapped, local, 1D Hamiltonians [247], rigourously justifying the description of such systems with MPS where χ does not scale with the system length. For many other systems, the area law applies with small violations [190, 248]. When applied to two dimensions, the area law suggests that the MPS ansatz requires higher χ as the system size increases, limiting its use to small or thin systems. Nevertheless, the ability to neglect a vast unphysical section of the Hilbert space by construction is an important strength of MPS.

E.2 Matrix Product Operators

Matrix product operators (MPO) allow the calculation of expectation values directly from the MPS form. Like MPS, they feature a tensor O_x at each site which has two virtual indices (or one virtual index at the ends of the system). We label the bond dimension of these virtual indices χ_{MPO}. As an operator tensor rather than

---

[1] It is rare for an MPS which acceptably approximates a physical state to approach the maximum S_{ent} for a given large χ because the singular values tend to decay gradually (rather than having χ large sigular values while all others are negligible).

a state tensor, O_x has two physical indices j_x and k_x to produce transformations between states. Similarly to MPS, for a given choice of the two sets of physical indices j_x and k_x, the operator matrix element is calculated by contracting the product of one row vector, $L-2$ matrices, and one column vector corresponding to the virtual indices of the tensors O_x using the formula

$$\langle k_1, k_2, ...k_L|\hat{O}|j_1, j_2, ...j_L\rangle = O_1^{j_1,k_1} O_2^{j_2,k_2} ... O_L^{j_L,k_L} \tag{E.2}$$

While we again do not detail the numerical implementation, the process of producing a new MPS wavefunction by operating on an existing MPS wavefunction ($|\psi'\rangle = \hat{O}\,|\psi\rangle$) is performed by contracting the physical indices of $|\psi\rangle$ with the 'dual vector' physical indices of $\hat{O}$, leaving $|\psi'\rangle$ with one set of physical indices as necessary for MPS. The new MPS wavefunction $|\psi'\rangle$ has a bond dimension equal to the product of the bond dimensions of $\hat{O}$ and $|\psi\rangle$, so it is often necessary to compress $|\psi'\rangle$ to reduce its bond dimension.

In order to perform these operations, it is necessary to express relevant physical operators as MPO. While this is straightforward for product operators (where $\chi_{MPO} = 1$), creating MPO for operators (such as Hamiltonians) which contain sums of multiple two-site terms is not trivial. Finite state machines can be used to build the MPO for the Hamiltonian [249], which allow such operators to be expressed with a relatively low bond dimension. More detail about the use of finite state machines to encode Hamiltonians can be found in Refs [223, 250].

E.3 DMRG

The most important MPS algorithm that we use in this thesis is DMRG (Density Matrix Renormalisation Group) [182, 183]. DMRG is an iterative algorithm which can be viewed as approximating the ground state with an MPS wavefunction [184]. The

MPS often has too many variational parameters to optimise them simultaneously, so the algorithm keeps the tensors on all but one site (for single-site DMRG) or two sites (for two-site DMRG) constant at one time. This fixes an effective environment ‘Hamiltonian’ for the sites being optimised, of which the lowest eigenvector is found iteratively. Over the course of a DMRG ‘sweep’, all sites are optimised in sequence. This sweeping procedure is not globally optimal but produces acceptable results at acceptable cost. Each sweep reduces the expectation value of energy and moves the state towards a local minimum of the energy. The algorithm terminates upon reaching a convergence criterion, such as the number of sweeps, the truncation error when compressing the MPS to the chosen bond dimension χ, or the energy difference between subsequent sweeps.

In this thesis, we have used two-site DMRG because it helps to avoid spurious local minima by allowing the bond dimension between two sites to vary during the local optimisation procedure (with subsequent truncation possible to compress the bond dimension). We note that other extensions to single-site DMRG can also be used for this purpose [251, 252].

While our descriptions of the algorithms focus on a chosen value of χ, we used the standard practice of running DMRG using relatively small values of χ initially to find a very approximate ground state wavefunction while benefiting from fewer energy minima and faster iterations. This wavefunction can be used as a starting point for future uses of the algorithm at higher χ, which often allows an acceptably converged ground state to be found more quickly than using the higher value of χ directly. The added information from using multiple different χ allows one to check whether the energy and other observables have converged with respect to increasing χ, which can indicate whether the chosen value of χ is sufficient to capture the ground state physics.

E.4 Two Dimensions

While MPS are not as well-suited to 2D systems as to 1D due to the necessary entanglement scaling with system size, they can still be used for small or thin 2D systems [194]. This requires the 'physical' 2D problem to be mapped to a 'numerical' 1D problem using a 2D-to-1D mapping. The 2D-to-1D mapping means that interactions between neighbouring sites on the physical lattice correspond to interactions over a large distance in the 1D chain (for common mappings, this distance scales with the smallest system length). Different choices of mapping can be particularly suited to certain entangled states, but it is usually helpful for the mapping to preserve the locality of interactions.

The ability to maintain a local interaction structure is dependent on the boundary conditions of the physical system. For open boundary conditions or cylindrical boundary conditions (open in one dimension and periodic in the other), the interaction range can be made to scale with the system width while the required χ scales exponentially with width [253]. However, if the boundary conditions are periodic in both directions, the locality structure breaks down and the computational difficulty scales exponentially with the number of sites, the same scaling behaviour as exact diagonalisation. For this reason, many MPS calculations of bulk 2D systems use cylindrical rather than fully-periodic boundary conditions even though this introduces edge effects.

The 2D-to-1D mapping means that the 'physical' 2D Hamiltonian becomes more complicated when expressed in 1D, due to both the increased range of interactions and the fact that this mapping usually destroys any translational invariance in the Hamiltonian. This makes building the MPO by hand a time-consuming and error-prone task [223]. In this thesis, we have used parsers which perform this task automatically from a 1D Hamiltonian written in a human-readable 'equation form' [250]. The necessary MPO bond dimension depends not only on the physical interactions considered, but on the system size (because the interactions become longer-ranged in 2D) and the 2D-to-1D mapping.

One problem with the introduction of long-range terms from the physical Hamiltonian or the 2D-to-1D mapping is that the energy depends explicitly on correlations between sites which are far separated in the MPS. This contrasts with two-site DMRG, which optimises two neighbouring sites at once, and can therefore struggle to recognise the energy minimisation that might arise from the long-range correlations. A useful technique to encourage long-range correlations, which we use in this thesis, is the ‘noise term’, originally developed for single-site DMRG [251]. This term adds a small perturbation based on the terms in the Hamiltonian to the state during optimisation, which can provide a ‘jump start’ to correlations, which then grow throughout subsequent sweeps if energetically favourable. This noise obviously introduces some deviation from the ground state, so it is important to reduce or remove it in the later sweeps of the DMRG algorithm.

F

Cylindrical Dipolar Bose-Hubbard Model Calculation Details

To map the 2D cylindrical lattice to a 1D matrix product state, we used a 'coil' mapping in which the MPS loops around each ring successively from $c = 1$ to $c = 8$. This means that neighbouring sites along the axis of the cylinder are separated by L_c sites in the 1D chain.

We use ≈ 80 DMRG sweeps, of which most are performed with $\chi \leq 400$. We increase χ until the maximum truncation errors in the DMRG sweep reduce to $\approx 10^{-6}$ and the energy reduction from increasing χ by over 10% is ≈ 1 part in 10^5. This required $\chi = 6400$ for the most entangled calculations, but $\chi \leq 2000$ was sufficient for regions I and II. Our calculations use the $U(1)$ particle number conservation symmetry. We use a noise term, which is reduced approximately exponentially with the number of sweeps and then removed in later sweeps, to encourage correlations between sites which are far apart in the 1D chain [251].

We would like to acknowledge the use of the University of Oxford Advanced Research Computing (ARC) facility in carrying out this work [228].

G

Entanglement Entropy in TU State at $\theta = \frac{\pi}{2}$

This section contains details of the estimate of the central charge c_C in the TU state below the TU-PS transition. Compared to the other extracted quantities, this value is quite susceptible to finite-χ effects which we mitigate using a $\frac{1}{\chi}$ extrapolation, as shown in Figure G.1(a) and (b). In Figure G.1(c), the results of the $\chi \to \infty$ extrapolation for susbsystem lengths $L_A \leq \frac{L_z}{2}$ are shown ($L_A > \frac{L_z}{2}$ produces the same results due to symmetry). With the exception of the cylinder edges, it can be seen that S_{L_A} is well-described by a linear fit with respect to $\ln(\frac{2L_z}{\pi}\sin(\frac{\pi L_A}{L_z}))$ as expected for the relation $S_{L_A} = \frac{c_C}{6}\ln(\frac{2L_z}{\pi}\sin(\frac{\pi L_A}{L_z}))$ up to constants, which was described in the main text. For $2.5 \leq V \leq 3.4$, we find values of c_C ranging from 0.97 to 1.07. For $V \leq 3.0$, far from the lower TU-PS transition, the values of c_C are almost independent of system size for $L_z \geq 40$, which suggests that these estimates are not significantly affected by finite-size effects.

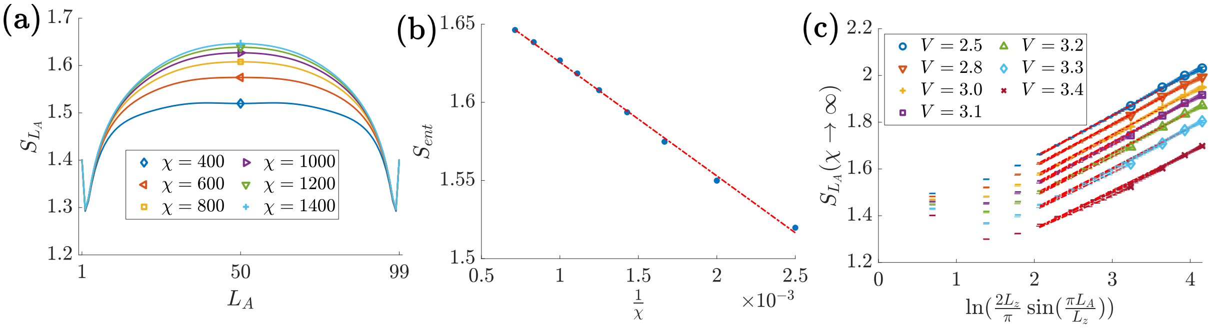

Figure G.1: Entanglement entropy as a function of subsystem length for TU state at polarisation angle $\theta = \frac{\pi}{4}$. (a) Entanglement entropy for different χ identified by markers for DDI strength $V = 3.4$ and axial length $L_z = 100$. (b) $\frac{1}{\chi}$ extrapolation (where χ is the MPS bond dimension) of entanglement entropy across central bipartition S_{ent} for $V = 3.4$ and $L_z = 100$ using $\chi = \{400, 500, 600, 700, 800, 900, 1000, 1200, 1400\}$, where the dashed red line denotes a linear fit used to extrapolate to $\chi \to \infty$. (c) Extrapolated values of S_{L_A} (entanglement entropy for a subsystem of axial length L_A) to $\chi \to \infty$ plotted against $\ln(\frac{2L_z}{\pi}\sin(\frac{\pi L_A}{L_z}))$ for $1 \leq L_A \leq \frac{L_z}{2}$, where error bars denote uncertainty in the $\chi \to \infty$ linear fit. Markers on the centremost partition identify the different values of V according to the legend. Data for $L_z = 40, 60, 80, 100$ are shown together. For each value of V, a linear fit for $4 \leq L_A \leq \frac{L_z}{2}$ is plotted in a dashed red line.

H

$L_c \to \infty$ Stripe Solid - Polar Solid

In this section, we consider the ground state for strong interaction $V >> J$. For $\theta \leq 0.63\frac{\pi}{2}$, interactions in the $\hat{\boldsymbol{z}}$ direction are attractive, which encourages stripe ground states where particles line up along the cylinder axis. For interactions truncated at next-nearest-neighbour, the energy of these states is dependent only on how many of these stripes are 'stuck together' and their azimuthal angle ϕ_c. When the number of bosons is half the number of sites and L_c is even, it is possible to avoid interactions between stripes, which means stripes will only stick together if the overall interaction between them is attractive.

Two neighbouring stripes at azimuthal angle ϕ_c feature L_z interactions in the $\hat{\boldsymbol{\phi}}$ direction and $L_z - 1$ interactions in the $\hat{\boldsymbol{\phi}} \pm \hat{\boldsymbol{z}}$ directions. Using equation 4.11, the DDI interaction energy between these stripes is

$$L_z(1 - 3\sin^2(\theta)\sin^2(\phi_c)) + \frac{L_z - 1}{\sqrt{8}}(3\sin^2(\theta)(1 - \sin^2(\phi_c)) - 1), \tag{H.1}$$

and the stripes will be encouraged to stick together if this quantity is negative. For infinite L_z, this means that stripes at $\phi_c = \frac{\pi}{2}, \frac{3\pi}{2}$ begin to stick together at $\theta \geq 0.307\frac{\pi}{2}$, which is the same condition as favouring the PS over the SS for $L_c = 8$ because the separation direction between the sites concerned is the same in both cases. As θ increases from that value, a larger range of ϕ_c feature attractive interactions between the stripes until all stripes in the range $\frac{\pi}{4} < \phi_c < \frac{3\pi}{4}$ and $\frac{5\pi}{4} < \phi_c < \frac{7\pi}{4}$ would attract their neighbours at $\theta = 0.61\frac{\pi}{2}$, forming the $L_c \to \infty$ analogue of the PS.

I

Wannier Functions for Finite Lattices

This appendix shows the derivation of the Wannier functions for finite lattices. It is based on that in the supplemental material of Ref [231] and is functionally equivalent except for a small extension for the second band.

We consider a lattice potential $A = A_0 \cos^2(\pi x)$ defined between $x = 0$ and $x = L$ with hard potential walls (after calculating the Wannier functions, we translate this to be symmetric about $x = 0$). The Wannier functions must vanish at the potential walls, so we build them using the basis functions $f_m(x) = \sqrt{2}\sin(m\pi \frac{x}{L})$ where m is a positive integer. For the lowest bands of $L = 9$, $m \leq 800$ causes negligible truncation. The kinetic Hamiltonian $\hat{H}_K$ is diagonal in these basis functions, where each basis function has kinetic energy $(\frac{m}{L})^2 E_R$. The Hamiltonian for the lattice potential H_P also couples basis states and is given by

$$\hat{H}_P = A_0 \sum_m \frac{1}{2} \left|m\right\rangle \left\langle m\right| + \frac{1}{4}(\left|m\right\rangle \left\langle m+2L\right| + \left|m+2L\right\rangle \left\langle m\right| - \left|m\right\rangle \left\langle -m+2L\right|), \tag{I.1}$$

where elements outside the chosen maximum value of m are excluded. The single-particle Hamiltonian is the sum of the kinetic and potential terms. Due to its structure, this Hamiltonian only couples the basis functions in groups defined by $(m \pm q)\text{mod}(2L) =$

0 labelled by q, where q takes integer values from 0 to L. Excluding small truncation effects, there are half the number of basis functions for $q = 0$ and $q = L$ each as for all other values of q. As in Ref [231], the lowest band is spanned by the lowest eigenstate for each value of $1 \leq q \leq L$. Meanwhile it is clear that the second-lowest eigenvalues for $1 \leq q \leq L-1$ are similar to the lowest eigenvalue for $q = 0$ and are far below the second-lowest eigenvalue for $q = L$. This motivates our identification of the lowest eigenstate for $q = 0$ and the second-lowest eigenstates for $1 \leq q \leq L - 1$ as the second band.

As in Ref [231], the Wannier functions within the space of a band are found by taking the eigenstates of the position operator

$$\hat{X} = \sum_{m_1,m_2} ((-1)^{(m_1+m_2)} - 1)\frac{4Lm_1m_2}{\pi^2(m_1^2 - m_2^2)^2} |m_1\rangle \langle m_2| \tag{I.2}$$

These can be converted into functions in real space using the definition of the basis functions. The BH model parameters can be found by calculating the matrix elements of the kinetic Hamiltonian and any additional interaction Hamiltonians in the basis defined by the Wannier functions. For the DDI and contact interaction, we numerically integrated over real space for this purpose. The tunnelling elements and on-site energies for both the lowest and second bands away from the edges of the lattice are in good quantitative agreement with the equivalent elements for infinite lattices calculated using the methods in Ref [141].

Note that this process only defines the Wannier functions up to a separate sign for each Wannier function. It is helpful for interpretation of results to fix this freedom, for example using the conventions that the lowest-band Wannier functions are positive at their potential minimum and that the second-band Wannier functions are negative(positive) for values of x immediately below(above) their potential minimum.