

2014-01-17

Quantum Computation and Many Body Physics

D'Souza, Adam

D'Souza, A. (2014). Quantum Computation and Many Body Physics (Doctoral thesis, University of Calgary, Calgary, Canada). Retrieved from <https://prism.ucalgary.ca>. doi:10.11575/PRISM/27510
<http://hdl.handle.net/11023/1265>

Downloaded from PRISM Repository, University of Calgary

UNIVERSITY OF CALGARY

Quantum Computation and Many Body Physics

by

Adam D'Souza

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS AND ASTRONOMY

CALGARY, ALBERTA

January, 2014

© Adam D'Souza 2014

Abstract

Motivated by the underlying desire to identify novel, physically reasonable resource states for measurement-based quantum computation (MBQC), this thesis explores two seemingly unrelated topics in some detail. The first is a study of the circumstances under which multiqubit quantum states that are equivalent to cluster states of the same size under stochastic local operations and classical communication (SLOCC) are either deterministic or probabilistic resource states, with the aim of identifying new resource states that are related to, but non-trivially different from, the cluster states. The second is an analysis of the properties of the ground state of a potentially physically realisable coupled-cavity quantum electrodynamics model called the Jaynes-Cummings-Hubbard (JCH) model. The hope is that the ground state of the model can in fact serve as a universal resource for MBQC.

In the first study, I identify two classes of 1D states in the SLOCC-equivalence class of 1D cluster states that constitute resources for random-length single-qubit rotations, in one case quasi-deterministically ($N - U - N$ states) and in another probabilistically ($B - U - B$ states). In contrast to the cluster states, the $N - U - N$ states exhibit spin correlation functions that decay exponentially with distance, while the $B - U - B$ states can be arbitrarily locally pure. I also show that a two-dimensional square $N - U - N$ lattice is a universal resource for quasi-deterministic measurement-based quantum computation, and that cubic $B - U - B$ states can be locally converted to 2D universal resource states.

In the second study, the Density Matrix Renormalization Group (DMRG) algorithm is used to characterize the ground states of the 1D JCH model in the regime of low photon densities, and compare it to the 1D ground state of the Bose-Hubbard (BH) model. Numerical results indicate that a Tonks-Girardeau regime, in which the photons are strongly fermionized, appears between the Mott-insulating and superfluid phases as

a function of the intercavity coupling.

The final chapter of the thesis outlines the initial progress that I have made in determining whether the 1D JCH ground state can serve as a resource for universal single-qubit rotations in the MBQC picture, as well as directions for future investigation.

Acknowledgements

It takes a village to write a thesis, and mine was no exception. First and foremost amongst the people who made this possible was my supervisor, David Feder. Throughout the years that I have worked with him, David's sheer love of and enthusiasm for physics have been inspirational. His door and his mind have always been open, and his juggling of the roles of scientist, mentor, taskmaster, friend, psychiatrist, psychologist and chef nothing short of masterful. To him and his lovely family – Pina Colarusso, Elena and Matteo Feder – I owe my undying gratitude.

Thanks are also due to the other faculty members who have provided advice and encouragement, in particular Drs. Christoph Simon, David Hobill, Barry Sanders, Gilad Gour and Peter Høyer, as well as to Dr. Roger Melko of the University of Waterloo for trekking out to Calgary for my defence. And speaking of my defence, thanks to Joydip Ghosh and Guilherme Almeida for their special efforts in helping me to prepare. I am further indebted to the Department of Physics and Astronomy (PHAS), Natural Science and Engineering Research Council (NSERC), Alberta Innovates – Technology Futures (AITF) and the Informatics Circle of Research Excellence (iCORE) for their financial contributions to my education.

Since starting graduate school, I have had the pleasure and privilege of interacting scientifically with a number of other students that have passed through the group. Special thanks is owed to Michael Underwood, who started on the journey towards his PhD simultaneously with me, and with whom I have enjoyed four years worth of productive discussions and brainstorming sessions, not to mention friendship. Thanks also to all of David's other graduate students (and some undergraduates and postdocs) with whom I have shared intellectual spacetime: Neda Amiri, Mahdi Ebrahimi-Kahou, Tim Friesen, Mike Garrett, Adam Green, Jaewoo Joo, Farokh Mivehvar, Alexis Morris, Katie

Perrin/Underwood, Randy Squires and Kevin Van de Bogart.

A special acknowledgement is due to Jessica Jenkins and Katrine Beauregard, my thesis writing buddies. Our joint writing sessions were instrumental to the completion of this lengthy endeavour. Besides those already mentioned, a huge number of people have contributed to making my grad school experience a memorable one, for better or worse. Although it is impossible to list them all by name, I will do so anyway. Thanks to all of my wonderful grad school friends, in particular Henry Chen, Adam Cote, Julie Croskill, Janine Giles, Sukhpreet Guram, Melissa Gurney, Jyotsna Kashyap, Lisa Lambert, Itzel Lucio Martinez, Colin Macdonald, Elliot Martin, Andrew Newman, Kelly Pasolli, Ofelia Rempillo, Michelle Seguin, Joshua Slater, Val Sugrue, Menka Sull, David Torre, Mateusz Trybowski, Kim Weger and Jeremy Wotherspoon. Thanks also to Elisa Sereno-Janz, my wonderful violin teacher, and to my many friends and colleagues at the Graduate Students' Association.

It should go without saying that everything I have achieved would have been impossible without my family, but in the same spirit of gratuitous verbosity that I adopted in writing my thesis, it shall be said. To my parents, Ida and Tony: your unconditional love and support have sustained me throughout my education, and will continue to do so for the rest of my life. To my sister Olivia and brother-in-law John: well, much the same really. And to my tiny (as of the time of writing) nephew Sam: you are so cute. So very cute. Keep it up! And to the rest of my family: my grandfather, aunts, uncles, cousins, "cousins" . . . you know who you are. Thanks to all of you!

Last but most certainly not least, we come to the person to whom an adequate expression of my gratitude would be a 277-page document of its own. To my best friend and constant companion, Paul Fairie: knowing you is a privilege and a joy. Thank you for forcing me to keep working when I wanted to quit, cooking my meals and taking disproportionate care of household tasks when I was too lazy or busy to do so myself,

sharing creative ventures, holidays and hobbies, and just generally talking the most delicious nonsense I have ever heard in my life. I could go on, but it would be to the significant detriment of my hopes of ever attracting female companionship. But you know what I mean.

Table of Contents

Abstract	ii
Acknowledgements	iv
Table of Contents	vii
List of Tables	x
List of Figures	xi
1 Introduction to Measurement-Based Quantum Computing and Stochastic Local Operations and Classical Communication	1
1.1 Abstract	1
1.2 A Whirlwind Introduction to Quantum Information	1
1.2.1 Quantum computers	1
1.2.2 Qubits, qudits and single-qudit (pure) states	3
1.2.3 Multiple qudits	5
1.2.4 Single-qubit quantum logic gates	6
1.2.5 Measurements	10
1.2.6 Density matrices and mixed states	12
1.2.7 Entanglement	14
1.2.8 Schmidt decomposition and related entanglement measures	17
1.2.9 Multi-qubit gates	18
1.2.10 Universal set of gates	19
1.2.11 Application: quantum teleportation	21
1.2.12 Circuit model of quantum computing	22
1.3 Early Measurement-Based Quantum Computing	24
1.3.1 Optical Quantum Computing with Non-Linear Entangling Gates	25
1.3.2 Linear Optical Quantum Computing and the Scalability Problem	27
1.3.3 Gate teleportation I: gate-dependent resource states	29
1.3.4 Gate teleportation II: gate-independent resource states	35
1.4 Universal MBQC: the One-Way Quantum Computer	38
1.5 Resource States for MBQC	40
1.5.1 Matrix Product States	43
1.5.2 MBQC in Correlation Space via MPS Representation	47
1.5.3 Example: 1D Cluster State as a MPS	49
1.5.4 Example: Ground State of a Quantum Spin Network as a Quantum Wire	51
1.6 Projected Entangled Pair States Representation and 2D Universal Resources	53
1.7 Identifying and Characterising Resource States	54
1.7.1 Quantifying entanglement: Entanglement Monotones and LOCC	55
1.7.2 Entanglement requirements for universal MBQC: resource states are rare	57
1.7.3 Physically interesting universal resource states for MBQC	60
1.8 Stochastic Local Operations and Classical Communication	63
1.9 Preamble to Chapter 2	68

2	Strategies for measurement-based quantum computation with cluster states transformed by stochastic local operations and classical communication	69
2.1	Abstract	69
2.2	Introduction	69
2.3	Background	73
2.4	Projective measurements on SLOCC-transformed cluster states	81
2.4.1	Strategy I: Guaranteed Unitary Evolution	81
	Derivation of N-type Operators	81
	Properties of N-Transformed Cluster States	86
2.4.2	Strategy II: Guaranteed Pauli Byproduct	89
	Derivation of B-Type Operators	89
	Properties of B-Transformed Cluster States	95
2.5	Random Length Computation	98
2.5.1	Deterministic single-qubit rotations: N – U – N state	99
2.5.2	Deterministic Universal MBQC: 2D N – U – N State	102
2.5.3	Probabilistic single-qubit rotations: B – U – B state	104
2.5.4	Universal MBQC: Percolated 2D Cluster State from 3D B – U – B state	107
2.6	Discussion and Conclusions	110
2.7	Acknowledgements	114
3	Introduction to Fermionized Photons in the Ground State of One-Dimensional Coupled Cavities	115
3.1	Introduction	115
3.2	Introduction to the Hubbard and Bose-Hubbard models	116
3.3	Interlude: introduction to second quantisation	117
3.4	Derivation of Bose-Hubbard model	122
3.5	Phase transitions	127
3.5.1	Classical Phase Transitions	127
3.5.2	Quantum Phase Transitions	137
3.6	Basic low-temperature physics of Bose-Hubbard model in 3D	137
3.6.1	Ideal gas limit, $t/U \rightarrow \infty$	138
3.6.2	Bose-Einstein Condensation and Superfluidity, $t/U > (t/U)_c$	142
3.6.3	Atomic limit, $t/U = 0$	151
3.6.4	Mott insulator, $t/U < (t/U)_c$	152
3.6.5	Mott Insulator-Superfluid Transition in 3D	156
3.7	Bose-Hubbard model in 1D	157
3.7.1	BEC in 1D: Quasicondensation	159
3.7.2	Luttinger Liquids	160
3.7.3	Berezinskii-Kosterlitz-Thouless transition	162
3.7.4	Tonks-Girardeau gas	162
3.8	Jaynes-Cummings model	164
3.9	Jaynes-Cummings-Hubbard model	168
3.10	Simulating 1D Systems: Density Matrix Renormalization Group	170
3.11	Preamble to Chapter 4	175

4	Fermionization in 1D Coupled Cavities	176
4.1	Abstract	176
4.2	Introduction	176
4.3	Models and Methods	179
4.3.1	JCH model	180
4.3.2	BH Model	182
4.3.3	Tonks-Girardeau Gas	185
4.3.4	Numerical Methods	188
4.4	Results	190
4.4.1	Density phase diagrams	190
4.4.2	Correlation functions	194
	One-body density matrix	194
	Two-body correlation function	200
4.4.3	Other measures of the ground state	203
	Superfluid fraction	203
	Condensate fraction	206
	Entanglement properties	210
4.5	Conclusions	216
4.6	Acknowledgements	219
5	Outlook	220
5.1	Introduction	220
5.2	Universal MBQC with Jaynes-Cummings-Hubbard model?	221
5.3	Infinite-System Density Matrix Renormalisation Group	223
5.3.1	Finite-System DMRG	230
5.3.2	DMRG as a Variational Method over MPS	231
5.4	MPS representation for ground state of 1D JCH model	236
5.4.1	Small local and bond dimension MPS representation for MI phase	236
5.4.2	Translational invariance: site-independent matrices	238
5.4.3	Single-cavity measurements	243
5.5	Future Directions	245
	Bibliography	248

List of Tables

4.1	Numerical values (with uncertainties) of superfluid fraction f_s in thermodynamic limit, for phase space points indicated in Fig. 4.8.	206
-----	--	-----

List of Figures and Illustrations

1.1	Definitions of basic logic gates. The gates take (a) two bits and (b) one bit as input, and produce one bit as output.	8
	(a) Two input bits.	8
	(b) One input bit.	8
1.2	Milburn's Fredkin gate, figure reproduced from Ref. [1]. The qubits are encoded in a dual-rail representation, with modes a and b corresponding to the two target qubits and mode c the control qubit. A nonlinear crystal placed in the upper arm of the interferometer causes the control qubit to interact with the target qubits, with an intensity-dependent interaction strength.	26
1.3	Procedures for teleporting an arbitrary single-qubit gate $U \in C_k^{(1)}$ and the controlled-NOT gate $\text{CNOT} \in C_2^{(2)}$ according to the GC protocol [2]. The figures are taken from Ref. [2]. In each case, the states $ \Psi_n^U\rangle$ and $ \chi\rangle$ are resource states for the teleportation that do not depend on the state $ \alpha\rangle$ to be teleported through U . It can therefore be prepared offline, in advance of executing the teleportation protocol.	32
1.4	Quantum circuits for gate teleportation. Fig. (a) is the circuit for teleporting CZ; solid red boxes represent Bell state preparations, and dashed blue boxes Bell-basis measurements. Fig. (b) is a slightly modified version of Fig. (a), where (in addition to some simplifications), the CZ gate on qubits 3 and 4 from (a) has been pushed through the Pauli correction procedure and incorporated into the resource state preparation (dotted green box). Fig. (c) is the standard state teleportation circuit, which is divided into two identical stages: stage 1 (solid red area) and stage 2 (dotted blue area). Fig. (d) is stage 1 of Fig. (c), showing the output of performing just the first step of the teleportation procedure, with the measurement of the first qubit done in the eigenbasis of X . Finally, Fig. (e) shows a generalisation of (d) in which the first qubit is measured in one of a single-parameter family of measurement bases, resulting in the teleportation of one of a single-parameter family of single-qubit gates.	34
2.1	(Color online) Correlation length scale L associated with the correlation function $ C_{1,2j+1}(Z, Z) \sim \exp(-\frac{2j}{L})$ for a $N - U - N$ ring with all N -type operators identical, on a ring of 1000 qubits, as a function of the parameters of γ and θ , the parameters of N . The length scale increases as N approaches the singular limit, i.e. as θ gets close to 0 or $\pi/2$	90
2.2	$N - U - N$ state, a one-dimensional structure that can be used for deterministic random-length single-qubit rotations.	99

2.3	2D $N - U - N$ state, universal for quasi-deterministic MBQC. The qubits labeled 1, 2 and 3 can be used to implement an entangling gate, which involves measuring qubit 2 in the Y basis. Alternatively, if an entangling gate is not desired here, qubits 1 and 3 can be decoupled by measuring qubit 2 in the Z basis.	103
2.4	B – U – B state, a one-dimensional structure that can be used for probabilistic random-length single-qubit rotations.	104
2.5	3D cluster state with B-type operators and unitaries acting on alternate qubits (a) before carving and (b) after carving. Grey qubits are acted upon by B-type operators and white qubits by local unitaries. Z-basis measurements are made in the z direction in (a) to disentangle those vertical chains originating from a white qubit in the $x - y$ plane. A measurement protocol along the remaining vertical chains in (b) produces a percolated 2D cluster in the $x - y$ plane.	108
2.6	(Color online) Probability of success of the procedure for deleting a B-type operator in the plane via a random walk in the third dimension, as a function of the ratio λ of the singular values (thick blue, color online). The probabilities p_k of deleting B with exactly k even-qubit measurements are also shown for k from 1 to 10 (thin, decreasing with increasing k), and the thick blue line is the sum of these. The red dashed line is the percolation threshold.	111
3.1	Generic mean-field phase diagram of Bose-Hubbard model, reproduced from [3]. Semicircular lobes of Mott-insulating regions with integer boson occupation per site are surrounded by a superfluid region. The tips of the lobes correspond to multicritical points; the transition through these points along a line of constant mean occupation per site for the d -dimensional system is in the universality class of the $d + 1$ -dimensional XY model. Here, J corresponds to t in the text of the thesis, and V corresponds to U	158
3.2	Schematic of Jaynes-Cummings-Hubbard model on a hexagonal lattice, implemented in a photonic bandgap material (reproduced from [4].) The red spheres represent two-level atoms, and the holes in the material are the cavities. λ is the wavelength of the resonant mode of the cavity and κ characterises the intercavity photon tunnelling. η is the refractive index of the material. The symbols ϵ and ω correspond to $\hbar\omega_a$ and $\hbar\omega_c$ respectively in the text of the thesis.	169
3.3	Phase diagrams for MI-SF transition for (a) bosons in the 1D BH model and (b) polaritons in the 1D JCH model, via infinite system (a) and finite system (b) DMRG. Reproduced from Refs. [5] and [6] respectively.	174
4.1	Mean density phase diagrams for (a) polaritons in the 1D JCH model and (b) bosons in the 1D BH model, in both cases for a system size $L = 31$. Regions of constant mean density correspond to MI states.	191

4.2	(a) Mean spin and (b) photonic excitation densities in the JCH model with $L = 31$. These can vary throughout the MI phase.	193
4.3	Representative plots of the correlation function $G^{(1)}(8, s)$ for the spin excitations in the JCH model with $L = 15$. (a) $n = 1$ MI phase at point $\tilde{\mu} = -0.900$, and (b) SF phase at point $\tilde{\mu} = -0.989$. In both cases the same five values of κ/g are chosen with values given in the legends. The correlation functions are consistent with (a) exponentials and (b) power-laws.	196
4.4	Representative finite-size scaling analysis of $G^{(1)}$ in the (a) MI and (b) SF regimes of the JCH model. The data points for the MI regime are $(\tilde{\mu}, \kappa/g) = (-0.856, 0.044)$ (deep within MI lobe), $(-0.856, 0.128)$ (near generic transition), and $(-0.900, 0.156)$ (near BKT point). The SF regime points are at $(-0.989, 0.022)$ (near generic transition), $(-0.989, 0.106)$ (deep within SF phase), and $(-0.989, 0.156)$ (near BKT point).	197
4.5	Phase diagrams for ξ in the thermodynamic limit, assuming $G^{(1)}(r) \sim \exp(-r/\xi)$ in MI regime of the (a) JCH model (photons only) and the (b) BH model. The correlation length ξ approaches the system size at the edge of the Mott lobe near the BKT point.	198
4.6	Phase diagrams for α in the thermodynamic limit, assuming $G^{(1)}(r) \sim r^{-\alpha}$ in SF regime of the (a) JCH model (photons only) and the (b) BH model. These values of α indicate strong fermionization of the photons in the JCH model.	200
4.7	The spatial-dependence of the unnormalized two-body correlation function $G^{(2)}(16, s)$ for $L = 31$ is shown for representative points in the SF phase where $\alpha \approx \frac{1}{2}$. Figures (a) and (b) correspond to photons at the points $(\tilde{\mu}, \kappa/g) = (-0.989, 0.039)$ and $(-0.989, 0.106)$ in the JCH model, respectively; ‘T’ and ‘D’ in turn denote ‘theory’ and ‘data.’ Figures (c) and (d) correspond to bosons at the points $(\mu/U, t/U) = (0.019, 0.066)$ and $(0.019, 0.180)$, respectively.	201
4.8	Superfluid fraction for the (a) JCH and (b) BH model outside but near the $n = 1$ MI lobe. For the JCH model, the symbol $\tilde{\kappa} = \kappa/g$ is used for compactness. For both systems, the curve marked by red squares (color online) is beyond the critical hopping value for the BKT transition for the $n = 1$ MI lobe, while the others are not. In both cases, the superfluid fraction f_s tends to 0 in the thermodynamic limit.	205
4.9	Maximal eigenvalues of the reduced single-particle density matrix for $L = 31$. In (a), the five largest eigenvalues of the photon density matrix are shown for fixed $\tilde{\mu} = -0.678$ as a function of hopping amplitude κ/g . In (b), the largest eigenvalue of the density matrix (identified with the condensate fraction f_c in the SF phase) for spin excitations and photons in the JCH model, and for bosons in the BH model, are shown as a function of generalized hopping J/J_0 at fixed chemical potential $\tilde{\mu} = -0.678$ and $\mu/U = 0.55$, respectively. The inset shows the lower density case $\tilde{\mu} = -0.922$ and $\mu/U = 0.133$ for comparison.	207

4.10	Condensate fraction for the (a) JCH model (photons only) and the (b) BH model outside but near $n = 1$ MI lobe, in the thermodynamic limit. In both cases, the condensate fraction within the MI lobes is zeroed out. . .	210
4.11	The photon and boson entanglement entropy $S_\rho(l)$ for a contiguous block of sites of length $l = 3$ to 29 with $L = 31$, for the JCH and BH models, respectively. Two mean densities are considered, $n = 0.645$ for the (a) JCH and (c) BH models, and $n = 1.129$ for the (b) JCH and (d) BH models. The solid blue line is the exact value of S_ρ^b for an ideal lattice Bose gas, and the solid red line in (a) and (c) is the corresponding value S_ρ^f for the ideal Fermi gas.	213
5.1	Schematic of infinite-system DMRG iteration with $B-l-r-U$ superblock structure, reproduced from Ref. [7].	224
5.2	Maximum discarded weight ϵ for all iterations up to system size $L = 20$ with $\tilde{\mu} = -0.878$, for various different values of κ/g , in the $n = 1$ MI lobe. For the choices $d = 5$, $\chi = 3$, ϵ is at most of order 10^{-4} , providing evidence that important basis states are not being discarded by the block decimation procedure of the infinite-system DMRG algorithm. Some of the data points are missing because in those cases, the truncated weights were extremely close to numerical zero.	239
5.3	Energy of ground state of 1D JCH model for system size $L = 20$ with $\tilde{\mu} = -0.878$, for various different values of κ/g , in the $n = 1$ MI lobe. This provides evidence that the choices $d = 5$, $\chi = 3$ for the infinite-system DMRG algorithm yield a reasonable approximation of the true ground state at this system size.	239
5.4	Fidelity of ground state of 1D JCH model for system size $L = 20$ with respect to the ground state at $L = 18$ with $\tilde{\mu} = -0.878$, for various different values of κ/g , in the $n = 1$ MI lobe. This provides good evidence that, for all cases other than $d = 3$, $\chi = 5$, the infinite-system DMRG algorithm has reached a reasonable approximation of a fixed point at this system size.	242

Chapter 1

Introduction to Measurement-Based Quantum Computing and Stochastic Local Operations and Classical Communication

1.1 Abstract

In Ch. 2 of this thesis, I present a publication concerning Measurement-Based Quantum Computation with cluster states that have been transformed by Stochastic Local Operations and Classical Communication. The paper contains an extensive introduction providing most of the background that is necessary to understand the technical details. Nevertheless there is still much that is assumed in terms of background information, which is necessary in order for the reader to appreciate the motivation of the work and put the results in context. This chapter is therefore intended as an introduction to the different fields covered by the paper: quantum information science, Measurement-Based Quantum Computing (MBQC), Matrix Product States (MPS) and equivalence of states under Stochastic Local Operations and Classical Communication (SLOCC).

1.2 A Whirlwind Introduction to Quantum Information

1.2.1 Quantum computers

Quantum computing is the field of computing using quantum mechanics (I will make this statement more precise shortly). Computation is fundamentally a physical process; computers represent and manipulate information using electronic components (transis-

tors, resistors, capacitors, diodes, etc.) under conditions that can generally be described adequately using classical mechanics. Ordinarily, circuit component sizes, currents, energies and so on are macroscopic, and the Bohr correspondence principle [8] applies. I will refer to computing devices for which this is true as a classical computer from now on. Classical computers have been getting ever more powerful over time, with larger memories and faster processors requiring a constantly increasing number of electronic components [9]. In order for computers to remain a reasonable size, the components themselves have to become smaller and smaller over time. Eventually, integrated circuits will likely be so small that the correspondence principle will no longer apply, and they will behave quantum mechanically.

That being said, the prospect of computers using quantum devices is tremendously exciting. Such computers, which I shall henceforth call quantum computers, appear to be significantly more powerful than their classical counterparts. Algorithms have been developed for quantum computers that are better than the best-known classical methods, for wide-ranging applications including factoring [10] and other algebraic or number-theoretical problems, graph problems [11], database searching [12], efficient simulation of quantum systems [13, 14], unbreakable cryptography [15] and so on. By 'better,' I sometimes mean that the quantum algorithm requires fewer time steps than the equivalent classical algorithm (e.g. factoring), and sometimes that the task is not even known to be possible classically (e.g. unbreakable cryptography). Not even mentioned are the quantum algorithms that simply have no classical analogue. The applications and associated citations given here are by no means exhaustive, and providing a truly comprehensive list would be both impossible and outside the scope of this thesis. Fortunately, extensive references exist on these subjects; see, for example, Refs. [16, 17, 18].

If a quantum algorithm can solve a problem more efficiently (i.e. with fewer timesteps) than the best possible classical algorithm, the algorithm is often said to provide a quan-

tum speedup. In fact, although as mentioned above there are several known quantum algorithms that are more efficient in this sense than the best known classical algorithm, there is actually no proof that there fundamentally exists any quantum speedup! More formally, computer scientists divide problems into complexity classes, based on the asymptotic dependence of the number of timesteps for the solution on the size of the input. The complexity class of problems of input size n that can be solved with $\text{poly}(n)$ timesteps on a classical computer is called P, and for those that can be solved with $\text{poly}(n)$ timesteps and bounded error on a quantum computer is called BQP. It is clear that BQP contains P, because any classical algorithm can be efficiently simulated on a quantum computer with a qubit register that remains always in a computational basis state. It is not known, but rather strongly suspected and fervently hoped, that BQP is strictly larger than P [19].

1.2.2 Qubits, qudits and single-qudit (pure) states

The basic unit of information for a quantum computer is a two-level quantum system, called a qubit. A qubit can be realised in a variety of two-level systems (see Ch. 7 of Ref. [19] for an introduction). It is the quantum analogue of the classical bit, a system that has two possible states (for example, a circuit that is either open or closed), traditionally labelled ‘0’ and ‘1’. The natural basis states of the two-level system encoding the qubit are traditionally labelled $|0\rangle$ and $|1\rangle$, and the basis set $\{|0\rangle, |1\rangle\}$ is called the computational basis. Qudits are the natural generalisation of qubits to d -level systems; a qubit is the special case of a qudit for which $d = 2$, and the computational basis states for a qudit are denoted $\{|i\rangle\}_{i=0}^{d-1}$. Similarly, I will refer to classical systems with d possible states as dits.

One key difference between dits and qudits is the superposition principle. A qudit can exist not just in a computational basis states, but also in an arbitrary linear combination,

or superposition, of them. The state of a qudit in an arbitrary such superposition is written as

$$|\psi\rangle := \sum_{i=0}^{d-1} a_i |i\rangle, \quad (1.1)$$

with $\{a_i\} \in \mathbb{C}$. Actually, Eq. 1.1 is the most general representation of a particular type of quantum state called a pure state. There is another kind of state called a mixed state, which usually arises in the context of open quantum systems, or systems with multiple qudits, which I will discuss shortly. Information represented by fundamentally quantum systems, such as qudits, is called quantum information. I will defer the discussion of the physical significance of the numbers $\{a_i\}$ until I introduce measurements.

The space in which $|\psi\rangle$ is defined is a subspace of d -dimensional complex space $\mathbb{C}^d$, called a d -dimensional Hilbert space, $\mathcal{H}^d$. $\mathcal{H}^d$ is equipped with an inner product $f : \mathcal{H}^d \times \mathcal{H}^d \mapsto \mathbb{C}$. Consider two states $|\psi_1\rangle$ and $|\psi_2\rangle$ in $\mathcal{H}^d$, defined as $|\psi_1\rangle = \sum_{i=0}^{d-1} a_i |i\rangle$ and $|\psi_2\rangle = \sum_{i=0}^{d-1} b_i |i\rangle$. The inner product is then given by

$$f(|\psi_1\rangle, |\psi_2\rangle) := \langle\psi_1|\psi_2\rangle = \sum_{i=0}^{d-1} a_i^* b_i. \quad (1.2)$$

The magnitude of the inner product between two vectors can be interpreted geometrically as the length of the projection of one of these two vectors on the other. The norm of $|\psi\rangle$ is then naturally defined to be the square root of its inner product with itself, $\| |\psi\rangle \| := \sqrt{\langle\psi|\psi\rangle}$. Two vectors whose mutual inner product is 0 are called orthogonal, and a vector with norm 1 is said to be normalised (to 1). The computational basis vectors $\{|i\rangle\}$ are normalised and mutually orthogonal, so the basis is called an orthonormal basis. The technical difference between $\mathcal{H}^d$ and $\mathbb{C}^d$ is that $\mathcal{H}^d$ is a complete metric space under this norm, while $\mathbb{C}^d$ is not; the practical consequence of this distinction to quantum mechanics is that only vectors from $\mathbb{C}^d$ with finite norm are in $\mathcal{H}^d$.

According to one of the basic postulates of quantum mechanics, the vector $|\psi\rangle$ repre-

sents the same physical state as $\lambda|\psi\rangle$ for any finite, non-zero complex number λ . Thus, $|\psi\rangle$ can always be normalised by dividing the coefficients by $\sqrt{\sum_i |a_i|^2}$. Therefore, the state of a qudit can be thought of as a ray in the complex plane, or formally an element of a d -dimensional projective Hilbert space $P(\mathcal{H}^d)$. As a consequence $|\psi\rangle$ can always be regarded as normalised. In the qubit case, $d = 2$, $|\psi\rangle$ can be written in the form

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle, \quad (1.3)$$

with $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$. The angles θ and ϕ can be interpreted respectively as the polar and azimuthal angles parametrising the unit sphere, and this construction is called the Bloch sphere. The Bloch sphere is useful for visualising single-qubit states, but does not generalise to the higher-dimensional case $d > 2$.

1.2.3 Multiple qudits

The state of a system comprising $n > 1$ qudits exists in a space that is the tensor product of the Hilbert spaces of the individual qudits. Denoting the Hilbert space of qudit i by $\mathcal{H}_i^d$, the space of the total system is the d^n -dimensional Hilbert space $\mathcal{H}^{d^n} = \mathcal{H}_1^d \otimes \mathcal{H}_2^d \otimes \cdots \otimes \mathcal{H}_n^d$. A convenient basis for $\mathcal{H}^{d^n}$ is the n -fold tensor products of the computational basis states of the individual qudits:

$$\begin{aligned} |00 \dots 0\rangle &\equiv |0\rangle_1 \otimes |0\rangle_2 \otimes \cdots \otimes |0\rangle_n; \\ |00 \dots 1\rangle &\equiv |0\rangle_1 \otimes |0\rangle_2 \otimes \cdots \otimes |1\rangle_n; \\ &\vdots \\ |d-1 \ d-1 \dots d-1\rangle &\equiv |d-1\rangle_1 \otimes |d-1\rangle_2 \otimes \cdots \otimes |d-1\rangle_n. \end{aligned}$$

This basis is called the computational basis for the n -qudit system.

Considering the n -dit strings labelling the n -qudit computational basis states as binary representations of the numbers 0 to $d^n - 1$, a completely general n -qudit pure state $|\psi_n\rangle$ can be given the state vector representation

$$|\psi_n\rangle = \sum_{i=0}^{d^n-1} c_i |i\rangle, \quad \{c_i\} \in \mathbb{C}. \quad (1.4)$$

The inner product generalises in the natural way, $\langle\psi_n|\psi_n\rangle := \sqrt{\sum_{i=0}^{d^n-1} |c_i|^2}$. There are d^n computational basis states, so the dimensionality of the total Hilbert space is $\dim(\mathcal{H}) = d^n$. Now, the dimension of the logical space of n classical dits is also d^n , but crucially, only d numbers are required to specify an n dit state, since each dit is definitely in one of d states. By contrast, the superposition principle requires one to specify all d^n coefficients $\{c_i\}$ from Eq. 1.4, an exponentially larger number. Statements of this nature are usually formulated precisely using big O notation: a function $f(n)$ defined on a subset of the real numbers is said to be $O(g(n))$ if there exist constants $c > 0$ and n_0 such that $|f(n)| \leq c |g(n)|$ for all $n > n_0$. In this notation, one says that $O(d^n)$ classical dits are required to faithfully represent the state of n qudits. The precise number of dits required depends upon the precision to which the complex numbers $\{c_i\}$ are to be specified, but that is not important here; the key is that for any fixed precision, the number of dits scales exponentially in the number of qudits.

1.2.4 Single-qubit quantum logic gates

In order to process information, computers need not only to be able to represent the information, but to manipulate it as well; this is done via logic gates. A classical n -dit logic gate is a device implementing a map $G : \{0, 1, \dots, d\}^n \mapsto \{0, 1, \dots, d\}^n$, while a quantum n -qudit logic gate implements $U : \mathcal{H}^{d^n} \mapsto \mathcal{H}^{d^n}$. In this thesis, I will use the term gate to refer to both the device and the logical map it implements. Again, the

precise physical implementation of a quantum logic gate depends upon the system chosen to represent the qudits. As before, I will keep the description of logic gates completely abstract. In this subsection, I will only consider single-qubit gates, i.e. the case $n = 1$ and $d = 2$. I will discuss multi-qubit gates in the next subsection, and I will not discuss single- or multi-qudit gates for $d > 2$ at all. This is principally because the action of single-qubit gates can be nicely visualised using the Bloch sphere, but this is not the case for higher-dimensional qudits. Fortunately, this is not an important restriction; a simple counting argument shows that the state of a qudit can be faithfully represented by $O(\log_2(d))$ qubits, a constant cost. In this case, the qudit and the system of qubits have Hilbert spaces of the same dimension, allowing a one-to-one mapping between them. I will therefore assume without loss of generality that quantum information is always represented by qubits.

The state of a qubit can be transformed from one point on the surface of the Bloch sphere to another, via a rotation operator defined on elements of $\mathcal{H}^2$, corresponding to elements of the group $SU(2)$. Unsurprisingly, this is the group that is used to give the spin-1/2 representation of the quantum angular momentum operators; after all, spin-1/2 particles are physical representations of qubits. These are the most general operations that can be performed on an arbitrary single-qubit state while preserving its norm.

Time-evolution of single-qubit quantum states is generally unitary, meaning that given some initial qubit state $|\psi(t_0)\rangle$, the state of the qubit $|\psi(t)\rangle$ at some later time $t > t_0$ can be written as

$$|\psi(t)\rangle := U(t, t_0)|\psi(t_0)\rangle, \quad (1.5)$$

where $U(t, t_0) \in SU(2)$ is unitary. In quantum mechanics, $U(t, t_0)$ is called the time evolution operator; its form depends upon the physical implementation of the qubit. This picture provides the basic insight as to how a logic gate is physically implemented;

Gate	00	01	10	11
AND	0	0	0	1
OR	0	1	1	1
XOR	0	1	1	0

(a) Two input bits.

Gate	0	1
NOT	1	0

(b) One input bit.

Figure 1.1: Definitions of basic logic gates. The gates take (a) two bits and (b) one bit as input, and produce one bit as output.

a device implementing the time-evolution operator U is switched on for some time t such that the time-evolution operator $U(t, t_0)$ corresponds to the desired logical map. An important exception to unitary evolution is the effect of a measurement on the qubit, which will be described later. Any known element of $SU(2)$ acting on a qubit transforms its state in a predictable way, so in the context of quantum information science, they are called single-qubit logic gates. They are the quantum analogues of classical logic gates such as AND, OR, XOR, NOT and so on. These classical gates are standard and well-known, but for reference, their definitions are given in Fig. 1.1.

The group $SU(2)$ has three generators comprising its Lie algebra. The generators are called the single-qubit Pauli operators; they are denoted X , Y and Z respectively and obey the standard angular momentum commutation relations. Suppose the computational basis vectors are represented by column vectors:

$$|0\rangle \equiv \begin{bmatrix} 1 \\ 0 \end{bmatrix}; \quad (1.6)$$

$$|1\rangle \equiv \begin{bmatrix} 0 \\ 1 \end{bmatrix}. \quad (1.7)$$

The single-qubit Pauli operators can then be given the concrete representation

$$X \equiv \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}; \quad (1.8)$$

$$Y \equiv \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}; \quad (1.9)$$

$$Z \equiv \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \quad (1.10)$$

They are all Hermitian, and have eigenvalue spectrum $\{1, -1\}$. Each of these operators has two mutually orthogonal eigenvectors, defined by the relations $Z|j\rangle \equiv (-1)^j|j\rangle$ (where $j \in \{0, 1\}$), $X|\pm\rangle \equiv \pm|\pm\rangle$ and $Y|\pm i\rangle \equiv \pm|\pm i\rangle$. The eigenvectors of X and Y are

$$|\pm\rangle \equiv \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle); \quad (1.11)$$

$$|\pm i\rangle \equiv \frac{1}{\sqrt{2}}(|0\rangle \pm i|1\rangle), \quad (1.12)$$

while the eigenvectors of Z are by definition the computational basis vectors. Orthogonal states are antipodal on the Bloch sphere; one can thus assign x -, y - and z -axes to the Bloch sphere, passing through the eigenvectors of the X , Y and Z operators respectively. The Pauli operators can then be interpreted geometrically as rotations by an angle π about their corresponding axes. In fact, they are the generators of arbitrary rotations about these axes: a rotation operator by a real angle θ about the m -axis is given by $R_m[\theta] \equiv e^{-i\theta M/2}$, where $m = x, y, z$ and $M = X, Y, Z$.

The elements of $SU(2)$ can be parametrised by three Euler angles. For example, an arbitrary element of $SU(2)$ (neglecting the physically irrelevant global phase) can be written as $O(\xi, \theta, \zeta) \equiv R_z(\xi)R_x(\theta)R_z(\zeta)$ [19]. Any pair of orthogonal axes can be used in place of z and x . Since arbitrary elements of $SU(2)$ can be generated via products of R_x and R_z operators, these gates constitute a universal set of gates for single-qubit rotations. Of course, they are actually single-parameter families of gates since the rotation angles are arbitrary. However, there is a useful theorem that any two rotation operators

$\{R_l(\alpha), R_m(\beta)\}$ about non-parallel rotation axes l and m by angles α and β that are not rational multiples of π form a universal set [20]. For example, the gate $R_z(\pi/4)$ together with the very important Hadamard gate

$$H \equiv \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (1.13)$$

constitutes a universal set. It is not an efficient set, in the sense that a possibly infinite number of these gates may be needed to implement a particular Bloch sphere rotation. Fortunately, a result called the Solovay-Kitaev theorem shows that it is possible to find efficient finite *approximately* universal sets of gates if one is willing to tolerate some error [21].

1.2.5 Measurements

Qubits would not be useful information carriers if there were no way to read out the information they carried. Fortunately, there is a way, which is to measure the qubits. The precise technology required to implement a physical measurement depends on the physical system used to represent the qubit. However, this subject is not of concern to this thesis, so I will stick to describing measurements in a purely abstract way.

The simplest and most familiar kind of measurement is a projective measurement. Consider a system corresponding to a Hilbert space of dimension D (it could be a single qudit with $d = D$, or $\log_2(D)$ qubits, or something else). It is a probabilistic operation that maps a qudit state to a classical dit. A single-qubit projective measurement is described by a set of operators $\{M_i\}_{i=0}^{D-1}$. The operators M_i must obey the following conditions:

1. they are Hermitian: $M_i^\dagger = M_i$;

2. they are orthogonal and idempotent: $M_i^\dagger M_j = \delta_{i,j} M_j$.
3. they are complete (or undercomplete): $\sum_{i=0}^{D-1} M_i^\dagger M_i \leq I_D$.

Here, I_D refers to the D -dimensional representation of the identity operator. Given some pure state $|\psi\rangle \in \mathcal{H}^D$, the measurement M returns one of a set of possible final states $\{|\psi_i\rangle\}_{i=0}^{D-1}$, given by

$$|\psi_i\rangle = \frac{M_i |\psi\rangle}{\sqrt{\langle\psi|M_i^\dagger M_i|\psi\rangle}}, \quad (1.14)$$

with $|\psi_i\rangle$ occurring with probability $\langle\psi|M_i^\dagger M_i|\psi\rangle$.

Here, the first two conditions ensure that each operator M_i is a projector (hence the term projective measurement). This means that it can be written as $M_i \equiv |\phi_i\rangle\langle\phi_i|$ for some pure state $|\phi_i\rangle$. With this notation, the action of the M_i operators is to annihilate any state that is orthogonal to $|\phi_i\rangle$, and map any other state to $|\phi_i\rangle$ (up to a normalisation factor). The orthogonality requirement of condition 2 ensures that the possible states $\{|\phi_i\rangle\}$ are mutually orthogonal. They can also be normalised, so $\{|\phi_i\rangle\}$ constitutes an orthonormal basis for $\mathcal{H}^D$, and the measurement is said to be performed in this basis.

Every physical observable $\mathcal{O}$ is a Hermitian operator, and therefore has a complete, orthonormal eigenbasis $\{|\phi_i^\mathcal{O}\rangle\}$ with corresponding eigenvalues $\{\lambda_i^\mathcal{O}\}$. A measurement of the operator $\mathcal{O}$ is a projective measurement in the basis $\{|\phi_i^\mathcal{O}\rangle\}$. After this measurement, the system is in one of the basis states $|\phi_i^\mathcal{O}\rangle$, and the corresponding eigenvalue λ_i is called the outcome of the measurement. The third condition ensures that the sum of the probabilities of the different outcomes i corresponding to projector M_i is less than or equal to 1.

There are also more general kinds of measurements than projective ones. The only condition for projective measurements that is necessary for measurements in general is the third, the (under)completeness condition. A completely general D -dimensional measurement is defined by a set of measurement operators $\{M_i\}_{i=0}^{D'-1}$, where D' can be

less than, greater than or equal to D , satisfying just the (under)completeness condition. For example, destructive measurements (e.g. detection of a photon), or measurements in a non-orthogonal basis, are physically admissible; yet, they cannot be described directly as projective measurements.

Notice that the determination of the probability of each outcome of a projective measurement can be determined even if rather than $\{M_i\}$, only $\{F_i := M_i^\dagger M_i\}$ are given. In this case, however, the post-measurement state cannot be determined. The F_i operators are Hermitian and positive semi-definite, and give rise to a class called positive operator valued measure measurements, or POVMs for short. A POVM is specified by such a set $\{F_i\}$, and can model more general situations than projective measurements (for example, measurements in a non-orthogonal basis). General measurements (and POVMs) are perhaps harder to interpret than the more familiar projective measurements. Fortunately, an important theorem called Neumark's dilation theorem [22] guarantees that any general measurement can be modelled as a projective measurement on a larger Hilbert space.

1.2.6 Density matrices and mixed states

In fact, Eq. 1.3 is not really the most general expression for a single-qubit state. This kind of state is called a pure state, and is postulated to provide the most complete possible description of a single qubit, in the sense that every physical fact that can be known about the qubit can be extracted from $|\psi\rangle$. But what about states for which not enough information is known to construct this complete description? An artificial example of this scenario is one in which some state preparation device flips a fair coin, and then prepares a qubit in state $|0\rangle$ if the result is heads and $|1\rangle$ otherwise, but does not publish the result of the coin toss. A more physically motivated example is one in which the qubit has interacted with its environment (say, for example, a heat bath) in such a way that the microscopic details of the interaction are not known, and therefore neither

is the state of the qubit. In these scenarios, the qubit must be described as a mixed state, which requires a tool called the density matrix.

The density matrix corresponding to a D – dimensional state is an operator $\rho : \mathcal{H}^D \mapsto \mathcal{H}^D$, satisfying the following properties:

1. Hermiticity: $\rho^\dagger = \rho$;
2. Positive-semidefiniteness: $\rho \geq 0$;
3. Normalisation: $\text{Tr}(\rho) = 1$.

For a pure state $|\psi\rangle \in \mathcal{H}^D$, the corresponding density matrix is $\rho = |\psi\rangle\langle\psi|$. The defining mathematical property of a pure-state density matrix is that it is a rank-1 projector; that is, it is a projector whose eigenvalue spectrum comprises a single one and $D - 1$ zeros. However, as can be seen from the list of density matrix properties, the rank need not be 1. A density matrix defined on $\mathcal{H}^D$, having rank r satisfying $1 < r \leq D$, corresponds to a D -dimensional mixed state.

Every density matrix ρ defined on $\mathcal{H}^D$ (indeed, every Hermitian operator) can be written in terms of its spectral decomposition, $\rho = \sum_{i=0}^{D-1} p_i |v_i\rangle\langle v_i|$, where the $\{p_i\}$ and $\{|v_i\rangle\}$ are the eigenvalues and eigenvectors of ρ respectively. Without loss of generality, I will assume the eigenvalues to be sorted: $p_0 \geq p_1 \geq \dots \geq p_D \geq 0$. By the positive-semidefiniteness condition on ρ , we have $p_i \geq 0$ for all i . The normalisation condition gives $p_i \leq 1$. Therefore, the $\{p_i\}$ values can be interpreted as probabilities. For ρ a pure state, $p_0 = 1$ and the rest are zero.

A mixed state can be viewed as a classical mixture (specifically a convex combination) of the pure states $\{|v_i\rangle\}$, with each $|v_i\rangle$ weighted by the probability p_i . More generally, a density matrix ρ can be written as a convex combination of density matrices $\{\rho_i\}$ weighted by probabilities $\{p_i\}$: $\rho = \sum_{i=0}^{D-1} p_i \rho_i$, where the ρ_i are not necessarily pure states. The set $\{p_i, \rho_i\}$ is called an ensemble realising ρ ; it is not unique.

To gain some understanding into the meaning of the ensemble, consider the artificial example from the beginning of this section, of a machine that prepares a qubit in either $|0\rangle$ or $|1\rangle$, in each case with probability $1/2$, but without making known which state was prepared. Because of our ignorance, the best description we can give the state is a mixed state corresponding to an equally weighted mixture of $|0\rangle$ and $|1\rangle$

$$\rho = \frac{1}{2}|0\rangle\langle 0| + \frac{1}{2}|1\rangle\langle 1|. \quad (1.15)$$

Suppose more generally that the machine prepares one of two orthonormal states $|\psi\rangle$ and $|\psi^\perp\rangle$ with probability p and $p^\perp$ respectively. In this case, the mixed state would be

$$\rho = p|\psi\rangle\langle\psi| + p^\perp|\psi^\perp\rangle\langle\psi^\perp|.$$

As an illustration of the non-uniqueness of the ensemble realising a density matrix, consider ρ from Eq. 1.15. This state can also be viewed as an equally weighted mixture of the eigenstates of X :

$$\rho = \frac{1}{2}|+\rangle\langle +| + \frac{1}{2}|-\rangle\langle -|. \quad (1.16)$$

It can easily be verified that the density matrices from Eqs. 1.15 and 1.16 have the same effect on any single-qubit state, and therefore represent the same single-qubit density matrix.

1.2.7 Entanglement

Perhaps the most uniquely quantum characteristic of quantum systems is their ability to be entangled. Consider a two-qubit system, and suppose that qubit 1 is in some state

$|\psi\rangle$, and qubit 2 in some state $|\phi\rangle$. The state of the joint system can be written as

$$|\Psi\rangle_{12} = |\psi\rangle_1 \otimes |\phi\rangle_2. \quad (1.17)$$

The state in Eq. 1.17 is called a product state, because the individual constituent particles are each in independent pure states.

The mathematical structure of Hilbert spaces allows for the existence of states that cannot be described as product states. Consider, for example, the set of states

$$\begin{aligned} |\Phi^+\rangle &:= \frac{1}{\sqrt{2}}(|00\rangle_{12} + |11\rangle_{12}); \\ |\Phi^-\rangle &:= \frac{1}{\sqrt{2}}(|00\rangle_{12} - |11\rangle_{12}); \\ |\Psi^+\rangle &:= \frac{1}{\sqrt{2}}(|01\rangle_{12} + |10\rangle_{12}); \\ |\Psi^-\rangle &:= \frac{1}{\sqrt{2}}(|01\rangle_{12} - |10\rangle_{12}). \end{aligned}$$

Suppose one wishes to factor the state into a product $|\Phi^+\rangle = (a_0|0\rangle + a_1|1\rangle) \otimes (b_0|0\rangle + b_1|1\rangle)$. It is readily apparent that there is no solution for the coefficients a_0 , a_1 , b_0 and b_1 that gives the right result. In fact, this is true for all four states $|\Phi^\pm\rangle$ and $|\Psi^\pm\rangle$ above. These states are examples of entangled states. One can verify that these states constitute a complete, orthonormal basis for $\mathcal{H}^4$; this basis is called the Bell basis, and is very important in quantum information processing. To be more precise, a state $|\psi\rangle_{12} \in \mathcal{H}_1 \otimes \mathcal{H}_2$, where $\mathcal{H}_{1,2}$ are Hilbert spaces of dimension greater than 1, is said to be entangled with respect to the bipartition of the total system into $\mathcal{H}_1$ and $\mathcal{H}_2$ if there is no way to write $|\psi\rangle_{12}$ as a product $|\psi\rangle_1 \otimes |\phi\rangle_2$, where $|\psi\rangle_1 \in \mathcal{H}_1$ and $|\phi\rangle_2 \in \mathcal{H}_2$.

The method given in the preceding paragraph for showing that the Bell basis states are entangled is simple because of the fact that the Hilbert spaces of the individual qubits are only two-dimensional. In this case, one only needs to show that a system of four non-

linear equations in four unknowns has no solution, but in the general case, the number of unknowns is $\dim(\mathcal{H}_1) + \dim(\mathcal{H}_2)$, and the number of nonlinear equations to be satisfied is $\dim(\mathcal{H}_1) \times \dim(\mathcal{H}_2)$; for large-dimensional $\mathcal{H}_1$ and $\mathcal{H}_2$, this method quickly becomes unwieldy.

A more convenient and physically insightful description of an entangled pure state is that the states of the individual subsystems can only be described as mixed states, not pure. To make this precise requires one additional, important theoretical concept: the reduced density matrix. Consider a bipartite pure state described by a density matrix $\rho_{12} = \sum_{i,j=1}^{\dim(\mathcal{H}_1)} \sum_{k,l=1}^{\dim(\mathcal{H}_2)} \rho_{ij;kl} |i\rangle\langle j|_1 \otimes |k\rangle\langle l|_2$ defined on $\mathcal{H}_1 \otimes \mathcal{H}_2$. The partial trace of ρ with respect to subsystem 2 is defined as

$$\text{Tr}_2(\rho_{12}) = \sum_{i,j=1}^{\dim(\mathcal{H}_1)} \sum_{k,l=1}^{\dim(\mathcal{H}_2)} \rho_{ij;kl} |i\rangle\langle j|_1 \otimes \langle k|l\rangle_2, \quad (1.18)$$

which, in the case of an orthonormal basis for $\mathcal{H}_2$ simplifies to

$$\text{Tr}_2(\rho_{12}) = \sum_{i,j=1}^{\dim(\mathcal{H}_1)} \sum_{k=1}^{\dim(\mathcal{H}_2)} \rho_{ij;kk} |i\rangle\langle j|_1. \quad (1.19)$$

The result of the partial trace is an operator acting on states in $\mathcal{H}_1$, and is called the reduced density matrix $\rho_1 := \text{Tr}_2(\rho_{12})$ for subsystem 1. ρ_1 is a density matrix describing the first system alone, and if ρ_{12} is a pure state but ρ_1 is mixed, then ρ_{12} must be entangled. This method only tells you whether a state is entangled or not, but entanglement is not merely a binary concept; it comes in continuous amounts that can be quantified in various ways; this point is discussed further in Sec. 1.2.8 below.

1.2.8 Schmidt decomposition and related entanglement measures

Any quantum pure state $|\Psi\rangle$ defined on a Hilbert space $\mathcal{H}$ that can be separated into the tensor product of two subspaces $\mathcal{H}^A$ and $\mathcal{H}^B$, i.e. $\mathcal{H} = \mathcal{H}^A \otimes \mathcal{H}^B$, can be expressed in a form known as the Schmidt decomposition [19]:

$$|\Psi\rangle = \sum_{\alpha=1}^{\chi_{A,B}} \lambda_{\alpha} |\Phi_{\alpha}^{[A]}\rangle |\Phi_{\alpha}^{[B]}\rangle, \quad (1.20)$$

where $\{\lambda_{\alpha}\} > 0 \in \mathbb{R}$, $\{|\Phi_{\alpha}^{[A(B)]}\rangle\} \in \mathcal{H}^{A(B)}$, $\sum_{\alpha=1}^{\chi_{A,B}} \lambda_{\alpha}^2 = 1$ and $\langle \Phi_{\alpha}^{[A(B)]} | \Phi_{\beta}^{[A(B)]} \rangle = \delta_{\alpha,\beta}$. The number $\chi_{A,B}$ is called the Schmidt rank of $|\psi\rangle$ with respect to the bipartition defined by A and B , and satisfies $\chi_{A,B} \leq \min\{\dim(\mathcal{H}^A), \dim(\mathcal{H}^B)\}$. The quantities $\{\lambda_{\alpha}\}$ are called the Schmidt coefficients, and I will assume without loss of generality that they are sorted in descending order:

$$\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{\chi_{A,B}}. \quad (1.21)$$

The basis sets $\{|\Phi_{\alpha}^{[A]}\rangle\}$ and $\{|\Phi_{\alpha}^{[B]}\rangle\}$ are called the Schmidt bases for $\mathcal{H}^A$ and $\mathcal{H}^B$ respectively. The Schmidt rank $\chi_{A,B}$ determines whether $|\Psi\rangle$ is entangled or separable with respect to this bipartition; a pure state is separable if and only if $\chi_{A,B} = 1$, so any state satisfying $\chi_{A,B} > 1$ is entangled.

The Schmidt decomposition allows the calculation of a measure of entanglement known as the von Neumann entropy of entanglement $E_{A,B}(|\Psi\rangle)$, given by the formula

$$E_{A,B}(|\Psi\rangle) = - \sum_{i=1}^{\chi_{A,B}} \lambda_i^2 \log_2 \lambda_i^2. \quad (1.22)$$

This quantity has a minimum value of 0 when $|\Psi\rangle$ is a product state, which is true iff $\lambda_1 = 1, \lambda_{k>1} = 0$, and a maximum value of $\log_2(\chi_{A,B})$ satisfied when $\lambda_i = 1/\sqrt{\chi_{A,B}}$ for all $i = 1, \dots, \chi_{A,B}$. On the basis of this, a related entanglement measure called the Schmidt

rank width can be defined. Let $\chi := \max_{A,B} \chi_{A,B}$ be the maximum Schmidt rank of $|\Psi\rangle$ over all possible bipartitions. Then, the Schmidt rank width is given by

$$E_\chi(|\Psi\rangle) := \log_d(\chi), \quad (1.23)$$

and corresponds to the largest possible von Neumann entropy of entanglement of $|\Psi\rangle$ over all bipartitions. Clearly, the largest possible value of χ occurs when the two subspaces contain $n/2$ qudits each (or $(n-1)/2$ and $(n+1)/2$ if n is odd), so $\chi \leq \log_2(d^{\lfloor n/2 \rfloor}) = \lfloor n/2 \rfloor \log_2 d$.

1.2.9 Multi-qubit gates

Multi-qubit gates are those maps $U : \mathcal{H}^{2^n} \mapsto \mathcal{H}^{2^n}$ where $n \geq 2$. Clearly there are an uncountably infinite number of n -qubit gates (even if $n = 1$). One might reasonably demand that an n -qubit quantum computer should be able to implement an arbitrary n -qubit unitary transformation, as one would like the computer to be able to take its n -qubit register from any allowed state to any other. A set of gates that can be combined to produce an arbitrary unitary operation on any number of qubits is called a universal set of gates.

Consider first the two-qubit gates, $n = 2$. One easy way to construct them is via the tensor product of two single-qubit gates. Suppose we have two qubits, with $U^{(1)}, U^{(2)} \in \text{SU}(2)$ acting on qubits 1 and 2 respectively. One can consider the joint effect of these two gates as a single gate $U^{(1)} \otimes U^{(2)} \in \text{SU}(4)$ acting on the two-qubit system. However, this is a rather trivial example; what about gates that are fundamentally two-qubit ones, and cannot be decomposed into single-qubit gates?

A two-qubit entangling gate $U^{(12)}$ is one that maps some product state $|\psi\rangle^{(1)} \otimes |\phi\rangle^{(2)}$ of two qubits into an entangled state $|\Phi\rangle$. This gate is truly a two-qubit gate, for if it

could be decomposed into a tensor product of two single-qubit ones, $U^{(12)} = U^{(1)} \otimes U^{(2)}$, then we would have

$$\begin{aligned} U^{(12)}(|\psi\rangle^{(1)} \otimes |\phi\rangle^{(2)}) &= (U^{(1)} \otimes U^{(2)})(|\psi\rangle^{(1)} \otimes |\phi\rangle^{(2)}) \\ &= (U^{(1)}|\psi\rangle^{(1)}) \otimes (U^{(2)}|\phi\rangle^{(2)}), \end{aligned}$$

which is manifestly separable (and therefore not entangled).

1.2.10 Universal set of gates

In a classical computation with n input bits and m output bits, one would like to be able to compute an arbitrary function $f : \{0, 1\}^n \mapsto \{0, 1\}^m$. Building a unique device for every possible f on every possible n and m is cumbersome, to say the least. In practice, it would be useful to be able to compute an arbitrary f by stringing together a number of standard gates chosen from a small, finite set of gates. Such a set is called a universal set of gates. For example, assuming one has the ability to add ancilla bits in standard states to the initial input register, the set {FANOUT, CROSSOVER, AND, XOR, NOT} is universal [19]. Here, FANOUT is an operation that takes an input bit and produces two copies of it as output, and CROSSOVER exchanges the values of two bits.

Analogously in quantum computation, one would like to be able to implement an arbitrary unitary gate $U : \mathcal{H}^{2^n} \rightarrow \mathcal{H}^{2^n}$ on a system of n qubits (in general, some input qubits and some ancillae). In the quantum context, a universal set of gates is one whose elements can be composed to implement an arbitrary unitary operation on an arbitrary number of qubits. Actually, this is quite a strict definition of universality that can be relaxed in various useful ways. One way is to allow the gate set to be approximately rather than exactly universal, by simply requiring that any unitary U be implementable with arbitrary precision (which allows for universal gate sets to be finite, as per the

Solovay-Kitaev theorem [21]). Specifically this means that for any $\epsilon > 0$, it is possible to implement a unitary operation V exactly by composing gates from the set, such that $\max_{|\psi\rangle} \|(U - V)|\psi\rangle\| \leq \epsilon$. Another is to weaken the requirement that the computer must be able to accept an arbitrary quantum input state or produce an arbitrary quantum output state; I will return to this subject in Sec. 1.7.2.

It is a well-known and very useful fact that the set of all single-qubit rotations, together with any two-qubit entangling gate, is a universal set (see, for example, p. 69 of Ref. [20]). Recalling from earlier that only two single-qubit rotations are necessary to generate the entire set of single-qubit rotations (either exactly but inefficiently, or approximately but efficiently), it thus follows that one can construct a universal set of gates from as few as three elementary gates!

In light of the preceding paragraph, it is reasonable to omit any discussion of n -qubit gates for $n > 2$. I will thus end this subsection with a description of an important class of two-qubit gates. Consider a single-qubit gate $U \in \text{SU}(2)$. The two-qubit controlled- U gate, or CU_{ct} for short, is a gate that uses qubit c as a control and qubit t as a target, doing nothing if qubit c is in state $|0\rangle$, and applying U to qubit t if qubit c is in state $|1\rangle$. More explicitly, its effect on the two-qubit computational basis states is

$$\begin{aligned} CU_{ct}|00\rangle_{ct} &\equiv |00\rangle_{ct}; \\ CU_{ct}|01\rangle_{ct} &\equiv |01\rangle_{ct}; \\ CU_{ct}|10\rangle_{ct} &\equiv |1\rangle_c \otimes U|0\rangle_t; \\ CU_{ct}|11\rangle_{ct} &\equiv |1\rangle_c \otimes U|1\rangle_t. \end{aligned}$$

Two extremely important, canonical examples of two-qubit gates are the controlled versions of the Pauli X and Z gates, CX and CZ (the c and t subscripts are omitted unless absolutely necessary). The CX gate (also called $CNOT$, since X is the quantum general-

isation of the classical NOT gate), is used in the quantum teleportation procedure that I discussed in Sec. 1.2.11, and the CZ gate is central to MBQC with cluster states, the main topic of Sec. 1.3.

1.2.11 Application: quantum teleportation

Suppose Alice possesses a single-qubit quantum state $|\psi\rangle$, and wishes to send the state to Bob. One potential method is for Alice to use a classical channel (such as a telephone) to transmit the computational basis coefficients of $|\psi\rangle$ to Bob. However, since these are general complex numbers (except that their amplitude is upper-bounded by 1), she cannot accomplish this with perfect precision by communicating only a finite number of classical bits. To accomplish this task Alice and Bob must share a quantum channel (the quantum analogue of a telephone line) that allows them to transmit quantum messages by exchanging only a finite number of bits and qubits.

A protocol called quantum teleportation [23] provides just such a channel, and also serves as a fascinating demonstration of the power of entanglement as a resource for quantum information processing. This protocol is of central importance to this thesis because it is one of the basic building blocks of MBQC. The steps of quantum teleportation are as follows:

1. Alice and Bob share in advance a maximally-entangled state $|\Phi\rangle_{23} = \frac{1}{\sqrt{2}}(|0\rangle_2|0\rangle_3 + |1\rangle_2|1\rangle_3)$.
2. Alice prepares a qubit in some state $|\psi\rangle_1$, and measures qubits 1 and 2 jointly in the Bell basis. It is easy to check that the Bell basis states $|\Phi^\pm\rangle$ and $|\Psi^\pm\rangle$ can all be written as $(I \otimes X^{m_2}Z^{m_1})|\Phi^+\rangle$ with $(m_1, m_2) \in \{0, 1\}^{\times 2}$ labelling the measurement outcome.
3. Bob's qubit is now in the state $U(m_1, m_2)|\psi\rangle_3$, where $U(m_1, m_2) = X^{m_2}Z^{m_1}$

is a unitary operator determined by Alice's two bit measurement outcome.

4. Alice sends m to Bob using a classical channel, and Bob then applies $U^{-1}(m)$ to his qubit to recover Alice's original state $|\psi\rangle$.

Quantum teleportation thus allows Alice to communicate an arbitrary single-qubit pure state to Bob by sending him one qubit (his half of $|\Phi\rangle_{23}$) and two classical bits (the result of her Bell basis measurement).

1.2.12 Circuit model of quantum computing

All of the ingredients previously discussed in this chapter can be assembled together into a paradigmatic model of quantum computing called the circuit model. I will now describe how a computation on an n -qubit quantum computer is done. First, consider a classical computer executing a classical algorithm on an n -bit input string. Imagine that the n bits are all initialised in the state 0. The actual bit string constituting the input to the algorithm is then prepared by flipping some of the bits of the input register via the NOT gate. The effect of executing a classical algorithm on the input string is to compute some function $f : \{0, 1\}^n \mapsto \{0, 1\}^m$ on the input and produce an m -bit string as the output. In order to do this, the classical computer uses a register of at least $\max(n, m)$ bits, but possibly more. In practice, the function can be computed by applying a sequence of classical logic gates to this register, in such a way that m predetermined bits of the register are (either definitely or with a greater probability than for any other configuration) in states encoding the bits of the output string. The sequence is completed by the computer reading out the states of the m output bits.

The quantum circuit model is essentially the generalisation of this idea to quantum computers. The idea is to use qubits and quantum circuits to execute a quantum algorithm computing the same function as some classical algorithm would have computed,

but hopefully more efficiently in some respect (runtime, memory, etc.). I should note here that an important conjecture in computer science called the Church-Turing thesis, widely believed to be true, implies that any function that can be computed by a quantum computer can also be computed by a suitable classical computer, and vice-versa.

Now consider the function $f : \{0, 1\}^n \mapsto \{0, 1\}^m$. In the circuit model, this function is computed as follows.

- An n' -qubit register, with $n' \geq m$, is initialised in some fiducial state, for example $|0\rangle^{\otimes n}$.
- For each possible classical input string $x \in \{0, 1\}^n$ to the function f , there is a corresponding state $|\psi_{\text{in}}\rangle$ of the n' -qubit input register. A sequence of logic gates is performed on the input register to produce $|\psi_{\text{in}}\rangle$.
- A sequence of quantum logic gates (which can in general be just single- and two-qubit gates), composing a unitary operation $U \in \text{SU}(n')$, is applied to $|\psi_{\text{in}}\rangle$; this sequence is called a quantum circuit. This produces another n' -qubit state

$$|\psi_{\text{out}}\rangle = U|\psi_{\text{in}}\rangle. \quad (1.24)$$

$|\psi_{\text{out}}\rangle$ has the property that when m predetermined qubits of the register are measured, the result is an m -bit string that is (with high probability) equal to $f(x)$.

- The relevant m predetermined qubits of the register are measured in the computational basis, and the resulting m -bit string is taken to be the output of the quantum algorithm.

If the algorithm is probabilistic, it may in general have to be repeated several times in order to give an acceptably high probability that $f(x)$ has been correctly computed.

1.3 Early Measurement-Based Quantum Computing

The reason I introduced quantum computing using the quantum circuit model is because it is conceptually the most straightforward; describing its core features does not require any technical machinery beyond the concepts of qubits, quantum logic gates and measurements. Of course, there are many more details that must be addressed in order to determine whether any particular *implementation* of quantum information processing is experimentally viable, notably the DiVincenzo criteria [24], but the conceptual core is simple. But is this really the only model one can have of quantum computation?

The answer is a resounding no; there are several other models that are fundamentally different, but are nevertheless equivalent, in the sense that all of these models are able to simulate each other efficiently. For example, the topological model encodes information in systems whose excitations obey anyonic (i.e. neither bosonic nor fermionic) statistics under exchange, thereby introducing a level of native fault tolerance (see [25] for a review). Similarly, quantum walk-based methods [26] allow for universal quantum computation via particles evolving under a constant (either continuous- or discrete-time) Hamiltonian, and such approaches have led to the discovery of previously unknown quantum algorithms. The topic of the rest of this chapter, and indeed much of this thesis, is yet another distinct model: Measurement-Based Quantum Computation (MBQC). In complete contrast to the circuit model, MBQC features no unitary evolution of any physical quantum systems whatsoever. Rather, information is encoded into a many-body entangled state, and computation proceeds solely via local measurements of the quantum systems comprising this state.

The origin story of MBQC is a rather interesting one; although it is a fundamentally different model of quantum computation from the circuit model that provides a unique perspective on the role of entanglement in the process, it arose from staunchly practical

considerations. The first proposal for universal quantum computing in which measurement plays a larger role than simple readout of information was due to Knill, Laflamme and Milburn [27]. This was a linear optical scheme for quantum computing that solved the biggest problem shared by its predecessors: scalability. Linear optical quantum computing in turn arose from the technical difficulties of implementing entangling gates for photons with non-linear optical elements. The remainder of this section will lay out the story in some detail.

1.3.1 Optical Quantum Computing with Non-Linear Entangling Gates

Optical schemes for quantum computing are particularly appealing, for two principal reasons. Firstly, quantum effects are particularly easy to observe with photons, compared to other physical systems. Secondly, virtually every viable scheme for distributed quantum computing and quantum key distribution (requiring satisfaction of the sixth and seventh DiVincenzo criteria) encodes the quantum information in some degree of freedom of a photon. As a result, there were several proposals for optical quantum computing in the early days of the field. For example, in 1989 Milburn demonstrated a method for implementing a Fredkin gate, shown to be universal for reversible classical computing [28] (and therefore quantum computing), by purely optical means [1]. The Fredkin gate is also known as the controlled SWAP, or CSWAP gate. It acts on three qubits, with the first a control and the other two the targets. Its effect is to act like the identity gate if the control qubit is in state $|0\rangle$, and to swap the states of the target qubits if the control qubit is in state $|1\rangle$. The basic design of Milburn's gate is that of a Mach-Zehnder interferometer, with a nonlinear optical element, such as a crystal, inserted into each of the arms (see Fig. 1.2). The essential role of this nonlinear medium is to mediate an interaction between photons; two photons passing simultaneously through the medium experience a different phase shift from what one photon alone would. Provided that the




Figure 1.2: Milburn's Fredkin gate, figure reproduced from Ref. [1]. The qubits are encoded in a dual-rail representation, with modes a and b corresponding to the two target qubits and mode c the control qubit. A nonlinear crystal placed in the upper arm of the interferometer causes the control qubit to interact with the target qubits, with an intensity-dependent interaction strength.

input modes a_i and b_i contain at most one photon, they can each encode a qubit in the dual-rail representation, where state $|0\rangle$ ($|1\rangle$) corresponds to the absence (presence) of a photon in the mode. Thus, if there is a photon in the input control mode c_i , then a photon passing through the upper arm of the interferometer experiences a larger phase shift than one passing through the lower arm; if c_i is instead empty, then the phase shift associated with either arm is the same.

The Fredkin gate is locally equivalent to a CZ gate with a phase that is dependent upon the parameters of the non-linear medium. If there is a photon in both the control mode and the other mode passing through the non-linear medium, the net result is a phase shift that is proportional to both the intensity of the beams and the length of the medium [19]. In order to implement a perfect CSWAP gate, which is locally equivalent to a CZ gate, this phase shift must be equal to π . If the intensities of the incoming modes are very low, such as at the single-photon level, then the medium must be very

long to compensate. But in this case, there is a non-negligible chance that the photons will be absorbed or scattered before traversing the medium, in which case the gate is probabilistic. If instead a sufficiently short length is chosen such that the phase shift is π/r where r is an integer greater than 1, then r of these gates must be composed in order to effect a perfect maximal entangling gate, a scheme that is destined to fail since each of the individual gates are error-prone. In short, implementing a circuit that simulates a maximally entangling logic gate for two qubits encoded in the dual-rail representation with non-linear media is problematic.

1.3.2 Linear Optical Quantum Computing and the Scalability Problem

The difficulties intrinsic to entangling photons with non-linear optical elements motivated the search for a scheme using only linear elements, in which interactions between photons were not required. A number of such schemes were quickly developed [29, 30, 31], based on single-photon interferometry. The basic idea was to change the encoding of the qubits; rather than the dual-rail representation, which requires one photon per logical qubit (and photon-photon interactions), n qubits could be encoded in 2^n degrees of freedom of a single photon. For example, one could arrange a network of beam splitters and mirrors in such a way that there are 2^n different possible optical paths for a single photon, and encode the quantum information corresponding to n -qubits in the “which-path” degree of freedom. Unfortunately, this idea suffered from a fundamental problem of scalability: the number of linear optical elements required to produce 2^n optical paths for a single photon is itself an exponential function of n .

The importance of scalability is strikingly illustrated by one of the most industrially relevant quantum information processing applications: Shor’s factoring algorithm [10]. The factoring problem is commercially motivated by public-key cryptography, the standard method for secure transmission of messages. Suppose somebody, call her Alice,

needs to send a bit $b \in \{0, 1\}$ to a recipient, Bob, securely. This means that an eavesdropper Eve must be unable to learn the value of b even if she intercepts the message. Using the most common commercial public-key cryptosystem, RSA encryption [32], Alice can succeed under the assumption that nobody except Bob can determine the prime factors of a very large integer (the public key) chosen and published by Bob. While this assumption appears reasonable, it fails if Eve has a quantum computer that can execute Shor's algorithm and thereby factorise the public key. The size of the public key is typically 512 or 1024 bits; Eve would thus need a quantum computer having thousands of qubits. If Eve wanted her computer to be capable of error correction [33], that would increase the number of qubits needed by roughly an order of magnitude, to potentially tens of thousands. There are, of course, other problems of legitimate practical interest for quantum computation requiring a smaller number of qubits; an early example was a method for finding the eigenvalues and eigenvectors of a local Hamiltonian [34], but even here, approximately 100 unencoded qubits are needed to solve problems of practical interest, with an order of magnitude more if error correction is desired.

Thus, any proposed physical implementation of a quantum computer must be scalable to large numbers of qubits in order to be useful. The linear optical proposals based on which-path information would need approximately 2^{1000} , or roughly 10^{300} optical elements, over 200 orders of magnitude larger than the number of atoms in the observable universe. Thus, the dream of LOQC appeared to be foundering. However, a brand new idea due to Gottesman and Chuang was to come to the rescue, not only resuscitating the viability of linear optical setups for quantum computing, but also founding the field of MBQC.

1.3.3 Gate teleportation I: gate-dependent resource states

The basic idea of Gottesman and Chuang [2] (henceforth referred to as GC) was to implement logic gates via a generalisation of the quantum teleportation procedure I discussed in Sec. 1.2.11. Suppose that Alice does not communicate the result of her Bell measurement to Bob, and Bob does not apply any corrections to the state of his qubit. In this case, rather than having the state $|\psi\rangle$, Bob has the same state, but subjected to a unitary operation. Thus, it is possible to teleport quantum information in such a way that it has been transformed by a logical gate $U \neq I$ during transit; in fact, in the absence of classical communication between Bob, it is impossible for Alice to ensure she transmits the state $|\psi\rangle$ faithfully.

The gate $U(m_1, m_2)$ that would normally have to be inverted by Bob is a random element of the single-qubit Pauli group

$$C_1^{(1)} := \{I, X, Y, Z\}; \quad (1.25)$$

the process is much more transparently useful for quantum information processing if the unitary transformation is deterministic, even if the measurement outcome m is random. This is why Alice and Bob resort to classical communication of Alice's measurement result; it allows Bob to apply a correction for the unwanted random Pauli gate, thereby ensuring that the net result is to have teleported $|\psi\rangle$ deterministically, subject to the identity gate. Thus, the question is: can Alice deterministically teleport $|\psi\rangle$ to Bob, subject to $U \neq I$?

GC showed that this is indeed possible. The process, in which quantum information is teleported subject to specific unitary transformation is called gate teleportation. Their key observation was that the teleportation procedure above, implementing the identity gate, relied on Alice and Bob sharing the Bell state $|\Phi\rangle$ in advance, meaning that $|\Phi\rangle$ is

a *resource* for teleportation, which is consumed during the process. But what happens if the state that is used as a resource is not a Bell state? The major result of this work was that, for certain choices of entangled resource states, the information can be teleported through a gate $U \neq I$. In fact, they found an entire hierarchy of single- and multi-qubit gates that can be generated with the right resource state. Consider the n -qubit Pauli group $C_1^{(n)}$, comprising all the n -fold tensor products of the single-qubit Pauli operators $\{I, X, Y, Z\}$. The n -qubit Clifford group $C_2^{(n)}$ is defined as the normaliser of the n -qubit Pauli group, namely all the unitary operators that map the n -qubit Pauli group to itself under conjugation:

$$C_2^{(n)} := \left\{ U \mid UC_1^{(n)}U^\dagger \subseteq C_1^{(n)}, U \in \text{U}(2^n, \mathbb{C}) \right\}. \quad (1.26)$$

This definition can be generalised: the group $C_k^{(n)}$ for $k \geq 2, k \in \mathbb{Z}$, is defined to be the normaliser of the group $C_{k-1}^{(n)}$. Here, $C_1^{(n)}$ is the n -qubit Pauli group, and the case where $k = 2$ gives the n -qubit Clifford group. GC demonstrated that any gate in $C_k^{(n)}$ could be teleported via a suitable choice of resource state. This result is important because the union $C := \bigcup_{k,n \in \mathbb{Z}^+} C_k^{(n)}$ contains a universal set of gates! In particular, the Hadamard gate $H \in C_2^{(1)}$ is a single-qubit Clifford gate and the $\pi/4$ rotation about the z -axis $T \in C_3^{(1)}$ is in the normaliser of the single-qubit Clifford group — these two gates are a famous example of a generating set for the single-qubit unitaries $\text{SU}(2, \mathbb{C})$ (see, for example, Ref. [20]). Similarly, the maximally entangling gate $\text{CZ} \in C_2^{(2)}$ is a two-qubit Clifford gate. Therefore, these three gates are universal for quantum computing.

It turns out that the resource state $|\Psi_n^U\rangle$ necessary to teleport an n -qubit gate $U \in C_k^{(n)}$ is the $2n$ -fold tensor products of Bell states $|\Phi\rangle$, with the operation U acting collectively on the second qubit of each Bell pair. Defining $|\Phi^n\rangle$ to be the n -fold tensor product of Bell states, $|\Phi\rangle^{\otimes n}$, with the labels rearranged so that the first n regis-

ters correspond to the first qubits of each of the Bell pairs, and the last n to the second qubits, the resource state is

$$|\Psi_n^U\rangle := (I_n \otimes U) |\Phi^n\rangle. \quad (1.27)$$

The gate U can then be performed on n qubits in an arbitrary state $|\alpha\rangle$ by doing pairwise Bell-basis measurements between these qubits and the first n qubits of the resource state $|\Psi_n^U\rangle$, followed by local Pauli corrections that are determined by the outcomes of the measurements. The circuits for implementing an arbitrary single-qubit gate $U \in C_k^{(1)}$ and CNOT, taken from Ref. [2], are depicted in Fig. 1.3 as examples. Note that the scheme for teleporting U reduces exactly to the quantum teleportation protocol if the $U = I$. Of course, the argument can be made that creating $|\Psi_n^U\rangle$ is just as difficult as implementing U directly in the first place, since its constructive definition explicitly involves an application of U . But in fact there are reasonable circumstances under which this is not the case. For example, suppose that one has access to an experimental implementation of the gate U , but it is non-deterministic. If one were to incorporate this gate in a traditional quantum circuit, then failure of the gate would be catastrophic, and the entire computation would have to be restarted. However, with the gate teleportation method of GC, the probabilistic U gate need only be operated in advance, offline, in the creation of the resource state. Then, when simulating the quantum circuit, at the time when U needs to be implemented, one can use the gate teleportation method together with the pre-prepared resource state, thereby performing U deterministically.

In fact, the first scheme for (near-)deterministic, scalable LOQC incorporated the teleportation idea as a crucial ingredient. The scheme is due to Knill, Laflamme and Milburn [27], henceforth referred to as KLM, but a detailed description of their achievement is beyond the scope of this thesis. The main point of interest for the present




Figure 1.3: Procedures for teleporting an arbitrary single-qubit gate $U \in C_k^{(1)}$ and the controlled-NOT gate $\text{CNOT} \in C_2^{(2)}$ according to the GC protocol [2]. The figures are taken from Ref. [2]. In each case, the states $|\Psi_n^U\rangle$ and $|\chi\rangle$ are resource states for the teleportation that do not depend on the state $|\alpha\rangle$ to be teleported through U . It can therefore be prepared offline, in advance of executing the teleportation protocol.

discussion is that KLM showed how to implement a deterministic entangling gate in the dual-rail representation using only linear optics. They first showed how to construct a CZ gate that is non-deterministic (in fact, it succeeds only with probability $\frac{1}{16}$, but its failure is heralded). They then used their non-deterministic CZ gate to create a resource for teleporting CZ, requiring on average 8 attempts in order to produce the state successfully. Finally, they used their resource state together with the teleportation protocol to teleport CZ. The original KLM paper [27] phrased the protocol as a linear optical network, but I have reproduced their design (in turn based on the idea by GC) in the quantum circuit notation in Fig 1.4. Consider two qubits in the states $|\psi\rangle$ and $|\phi\rangle$. Rather than acting a non-deterministic CZ gate on them, one first teleports the states of these qubits to two ancillary qubits, using the standard protocol. In Fig. 1.4, the ancilla qubits are the third and fourth wires for $|\psi\rangle$ and $|\phi\rangle$ respectively. This requires creation of two Bell states, indicated by the red boxes with solid outlines, followed by two Bell-basis measurements (indicated by B) and Pauli correction operations depending on the measurement outcomes (C_1 and C_2), shown in the blue boxes with dashed outlines.

Finally, the CZ gate is performed on the ancillary qubits instead of the original ones.

Obviously, this is not helpful at all! The CZ gate is just as likely to fail on the ancillary qubits as it was on the original. However, an important trick can be played here. The CZ gate is an element of the two-qubit Clifford group $C_2^{(2)}$, which maps Pauli operators to Pauli operators. This means that it obeys a useful commutation relation with the Pauli correction operations $C_{1,2}$. The effect of conjugating the corrections by CZ (using the fact that CZ is idempotent) is

$$\text{CZ} (C_1 \otimes C_2) \text{CZ} = C'_1 \otimes C'_2 \in C_1^{(2)}.$$

Therefore,

$$\text{CZ} (C_1 \otimes C_2) = (C'_1 \otimes C'_2) \text{CZ},$$

meaning that CZ can be pushed through the Pauli correction operators in the circuit, with the only effect being that they are changed to different Pauli operators.

I have indicated the effect of pushing CZ through the correction operators in Fig. 1.4, together with some other simplifications: the notation $H|0\rangle \equiv |+\rangle$, the circuit for Bell measurements re-expressed in terms of CZ and H gates followed by computational basis measurements, and the fact that H is idempotent as well. Essentially this process looks like one where you prepare a resource state (indicated in the dotted green box), followed by (locally rotated) Bell-basis measurements and Pauli corrections, giving the same result. Provided the resource state has been prepared properly in advance, the procedure implementing CZ on the logical qubits is now deterministic (assuming the Bell-basis measurements are deterministic). The particular resource state indicated in the green dotted box is called a cluster state, an extraordinarily important resource that is at the heart of most of the discussions of this thesis.

I would like to close this section by returning briefly to the discussion of the KLM




Figure 1.4: Quantum circuits for gate teleportation. Fig. (a) is the circuit for teleporting CZ; solid red boxes represent Bell state preparations, and dashed blue boxes Bell-basis measurements. Fig. (b) is a slightly modified version of Fig. (a), where (in addition to some simplifications), the CZ gate on qubits 3 and 4 from (a) has been pushed through the Pauli correction procedure and incorporated into the resource state preparation (dotted green box). Fig. (c) is the standard state teleportation circuit, which is divided into two identical stages: stage 1 (solid red area) and stage 2 (dotted blue area). Fig. (d) is stage 1 of Fig. (c), showing the output of performing just the first step of the teleportation procedure, with the measurement of the first qubit done in the eigenbasis of X . Finally, Fig. (e) shows a generalisation of (d) in which the first qubit is measured in one of a single-parameter family of measurement bases, resulting in the teleportation of one of a single-parameter family of single-qubit gates.

CZ gate, even though it is not directly relevant to the discussion at hand. It turns out that this gate, using the teleportation method, is not deterministic either, but only succeeds with probability $\frac{1}{4}$. The reason is that the basic teleportation procedure itself is probabilistic in linear optics. There is no easy way to implement a Pauli X gate in this setting [27], so if the outcome of the measurement of the first qubit in the teleportation procedure is 1, the gate fails. Thus, the teleportation succeeds only with probability $\frac{1}{2}$, and since there are two such teleportations in the KLM implementation of CZ, the overall success probability of the gate is only $\frac{1}{4}$. Nevertheless, this is still significantly better than the original success probability of $\frac{1}{16}$. Furthermore, KLM gave a method, using a modified teleportation protocol, to boost the success probability arbitrarily close to 1.

1.3.4 Gate teleportation II: gate-independent resource states

I have now established that there exists a protocol for simulating a universal set of logic gates for quantum computing using the teleportation protocol with an entangled resource state that depends upon the gate to be teleported. But, is there a way to teleport gates using a resource state that is independent of the gate in question? For example, recall that any single-qubit gate $U \in \text{SU}(2, \mathbb{C})$ can be teleported, provided one has access to the resource state $|\psi_1^U\rangle$. An arbitrary single-qubit gate can be expressed in terms of its Euler decomposition in terms of three real parameters (neglecting a possible overall phase):

$$U \equiv U(\xi, \zeta, \eta) = R_z(\xi)R_x(\zeta)R_z(\eta), \quad (1.28)$$

where $\xi, \zeta, \eta \in \mathbb{R}$ and $R_{x(z)}(\theta)$ is a unitary that rotates a single-qubit state by an angle θ about the $x(z)$ axis of the Bloch sphere. Does there exist a resource state $|R\rangle$ that enables teleportation of an arbitrary U while itself being independent of U ? Since $|R\rangle$ is independent of U , we must find a way to introduce some degrees of freedom somewhere

else in the teleportation circuit. Why not allow some freedom in the bases in which we measure the first two qubits?

A measurement basis for a projective single-qubit measurement is completely specified by two real parameters, because the single-qubit state corresponding to measurement outcome 0 can be parametrised in terms of its two Euler angles, and then the state for outcome 1 is its unique orthogonal complement (up to an irrelevant global phase). The teleportation circuit features two projective measurements, which allows for potentially up to four degrees of freedom in the unitary that is teleported, more than the three required. Is it possible, therefore, to adapt the teleportation circuit to teleport an arbitrary U with U -independent resource state by allowing the two projective single-qubit measurements to be in arbitrary bases?

Consider the state of the system in Fig. 1.4c, after the CZ and H gates have been applied, but before the single-qubit measurements. Suppose that instead of measuring the qubits in the computational basis, we project qubit 1 onto the normalised state $|\chi\rangle := a|+\rangle + b|-\rangle$, and qubit 2 onto $|\mu\rangle := c|0\rangle + d|1\rangle$, thereby disentangling them from qubit 3. It is easy to verify, using the formalism of quantum measurements, that the resulting state on qubit 3 is $M|\psi\rangle$, where

$$M \propto \begin{bmatrix} ac & bd \\ ad & bc \end{bmatrix}. \quad (1.29)$$

Suppose we write $a = |a|e^{i\theta_a}$, $b = |b|e^{i\theta_b}$, $c = |c|e^{i\theta_c}$, and $d = |d|e^{i\theta_d}$. Imposing that $MM^\dagger = M^\dagger M = I$, i.e. that M is unitary, one finds the constraints $|a| = |b|$ and $\theta_d - \theta_c = \pm\pi/2$. So, $|\chi\rangle$ only has one remaining degree of freedom, the relative phase between the $|0\rangle$ and $|1\rangle$ components. We must thus have $|\chi\rangle = \frac{1}{\sqrt{2}}(|0\rangle + e^{i\phi}|1\rangle)$, and its orthogonal complement $|\chi^\perp\rangle = \frac{1}{\sqrt{2}}(|0\rangle - e^{i\phi}|1\rangle)$, leaving one degree of freedom in the measurement basis for qubit 1. Similarly, the phase difference must be $\pi/2$ for one basis state for the

measurement of qubit 2 and $-\pi/2$ for its orthogonal complement. The only degree of freedom is the amplitude of, say, the $|0\rangle$ component of the first state, with the other three amplitudes fixed by orthonormality of the basis. So, if we demand the teleported gate to be unitary, it can only have two real degrees of freedom, as opposed to the completely general case of three.

Consider the standard teleportation protocol for transmitting a state faithfully. I have depicted the circuit for doing this in Fig. 1.4c, but having expressed the creation of the Bell states in terms of CZ and H operators, and the measurement in the Bell basis in terms of these same operators and local computational basis measurements. Evidently, this procedure can be split into two stages, stage 1 (indicated by a solid red box in Fig. 1.4c) and stage 2 (dotted blue box), followed by the Pauli correction operations can be performed. Stage 1 and stage 2 are procedurally identical; the input is some generally unknown state in the first wire, and $|+\rangle$ in the second, the two qubits are entangled via CZ, and then the first qubit is measured in the basis $\{H|0\rangle \equiv |+\rangle, H|1\rangle \equiv |-\rangle\}$, which is the eigenbasis of the Pauli X operator. The output of stage 1 on the bottom register is used as the input for stage 2 on the top register; thus, it is worth examining the effect of the operations in stage 1 (equivalently stage 2) independently.

The circuit for stage 1 alone is presented in Fig. 1.4d. The outcome of the measurement of the top register is denoted by $m \in \{0, 1\}$. By analysing this circuit step-by-step, it is trivially shown that the resulting output state on the bottom wire is

$$|\psi'\rangle = X^m H |\psi\rangle. \quad (1.30)$$

$|\psi'\rangle$ is then used as the input state for stage 2; suppose the measurement outcome is denoted by $m' \in \{0, 1\}$. Then, the final outcome of the two-stages, before the Pauli

correction operators are applied, is

$$\begin{aligned}
 |\psi''\rangle &= X^{m'} H |\psi'\rangle \\
 &= X^{m'} H X^m H |\psi\rangle \\
 &= X^{m'} Z^m |\psi\rangle,
 \end{aligned}$$

which is the familiar result for the standard teleportation algorithm. Now, to show how to teleport a single-parameter family of gates, consider measuring the first qubit in the basis $\left\{ \frac{1}{\sqrt{2}} (|0\rangle \pm e^{i\xi} |1\rangle) \right\}$, which can be obtained by acting the operator $R_z(-\xi)H$ on the computational basis. This measurement can be implemented by implementing the transformation $HR_z(\xi)$ on the qubit, which maps the states of the desired measurement basis back to the computational basis states, and then measuring in the computational basis. The circuit implementing this procedure is shown in Fig. 1.4e. What is the output of this protocol? To see this, notice that $R_z(\xi)$ commutes with CZ, and can thus be pushed through the CZ gate to act directly on the initial state of qubit 1. The protocol is then precisely the same as the one from Fig. 1.4d with the input state $R_z(\xi)|\psi\rangle$, and so the output is clearly $X^m HR_z(\xi)|\psi\rangle$. Thus, we now have the ability to teleport an entire single-parameter family of gates with a fixed resource state, by introducing a single-parameter family of measurement bases. To clarify, when I say “fixed resource state”, I mean that it is prepared in the same way, by applying CZ to the input state $|\psi\rangle|+\rangle$, independent of the value of ξ that is desired. This is the germ of the idea behind the one-way quantum computer, to whose basic description I now turn.

1.4 Universal MBQC: the One-Way Quantum Computer

The one-way quantum computer (1WQC) of Raussendorf and Briegel [35, 36] was the first ever proposal of a scheme for universal MBQC with a uniform resource state. The

proposal takes the simple idea of gate teleportation and shows how to use this scheme to implement a universal set of logic gates with a suitable resource state. The resource states in question are called the cluster states, also known as graph states [37]. A cluster state can always be identified with a simple, undirected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, comprising a set of vertices $\mathcal{V}$ and a set of edges $\mathcal{E}$, each element of which connects a pair of vertices. The number of vertices, $|\mathcal{V}|$, is equal to the number of qubits in the cluster state. The state is constructively obtained by identifying the vertices of $\mathcal{G}$ with qubits initialised in the state $|+\rangle$, and then applying CZ gates between all pairs of qubits whose corresponding vertices are connected by edges. The gate-independent resource states that I show in the previous section are all special cases of cluster states on small one-dimensional lattices of a few qubits. Typically, $\mathcal{G}$ describes a d -dimensional rectangular lattice with $d \geq 1$, but this is not necessarily the case.

I provide a detailed description of the 1WQC in Ref. [38], which is Chapter 2 of this thesis; thus, I will not do so here. However, the key points of the 1WQC scheme are as follows. An arbitrary quantum circuit can be simulated by the 1WQC with a suitable cluster state, for example, one defined on a 2D rectangular (square) lattice. Single-qubit operations are implemented by measurements along a one-dimensional chain within the cluster state. The basis in which measurements are to be performed is the single-parameter family of measurement bases given in Sec. 1.3.4 above, and the choice of the parameter depends both upon the Euler angles of the single-qubit gate being simulated AND the results of previous measurements. Two-qubit entangling gates can be implemented via suitable local measurements on a two-dimensional subpart of the cluster state. Since universal quantum computation can be achieved with a single perfect entangling gate, it is sufficient to show how to simulate, for example, CZ; the method is closely related to the circuit for teleporting CZ shown in Fig. 1.4b.

1.5 Resource States for MBQC

In Sec. 1.4, we have seen that cluster states comprise a family of universal resource states for MBQC, via the 1WQC scheme. One question arises immediately: are there any other families of resource states that allow for universal MBQC via local measurements? There are at least two important reasons to be interested in this question. The first is a practical one: the existence of a variety of universal resource states would provide some flexibility to experimentalists, and some resources may be easier to prepare than others within particular physical implementations. However, the second is a more fundamental question: what are the requirements for quantum computing that allow for a speedup over classical computing? It is not proven that this speedup exists, but it is strongly suspected. Therefore, assuming it exists, what could possibly be the reasons?

One clear difference between classical and quantum computers is the ability for qubits to exist in superposition states, in contrast to classical bits. Superposition is clearly helpful to quantum computing; for example, Grover's algorithm [12], which allows one to find a marked item in an unordered database containing n entries, runs in $O(\sqrt{n})$ time, as compared to the provably optimal classical runtime of $O(n)$. However, this is only a quadratic speedup, and therefore does not separate BQP from P. Another important feature of quantum computers that is not available classically is the ability of multiple qubits to be entangled. Entanglement too is of obvious help; quantum teleportation [23] and quantum key distribution [15], among other communication-related tasks, would be impossible without it. But is entanglement a necessary part of the explanation for the quantum speedup? And if so, how?

There are circumstances under which a quantum computer is known to be no more powerful (in the sense of a speedup) than classical computers. By this I mean that, under certain circumstances, a quantum computer Q_C that can execute some circuit C with

$O(t_C)$ timesteps can be simulated by a classical computer with $O(\text{poly}(t_C))$ timesteps. When this is the case, I will say that Q_C can be efficiently classically simulated. An efficient classical simulation of a quantum computation on n qubits will generally involve:

- (a) finding an efficient representation, requiring $O(\text{poly}(n))$ bits, of the quantum state at each time step;
- (b) finding an efficient way of effecting the dynamics, in other words the effect of a unitary operation appearing in the quantum computation on the state.

Whenever a quantum computation can be efficiently simulated classically, it cannot be said to provide a quantum speedup.

For example, a quantum computer whose n qubits are restricted to remaining in computational basis states at all times can be trivially classically simulated. The representation of the state is via n classical bits, each encoding a qubit, whose state $0(1)$ corresponds to the qubit state $|0(1)\rangle$. Assuming without loss of generality that the quantum circuit contains gates acting on no more than two qubits, the dynamics can be determined by computing for each gate U acting on k qubits, where $k \in \{1, 2\}$, the function $f_U : \{0, 1\}^{\times k} \mapsto \{0, 1\}^{\times k}$ on the relevant classical registers. A famous result called the Gottesman-Knill theorem [39] shows that any quantum computer executing a circuit composed entirely of gates in the Clifford group $C_2^{(2)}$ can also be classically simulated; in this case, it turns out that the state and the dynamics can be efficiently represented using the stabiliser formalism [33]. This means that access to neither superposition nor entanglement can by themselves explain the quantum speedup, as both superpositions and entanglement can be produced by Clifford gates.

In the MBQC context, the underlying resource state (assuming it is defined on qubits for convenience) can always be constructed via the action of some universal set of gates, including entangling gates, on an initially uncorrelated register of qubits. What fol-

lows is nothing but local (adaptive) measurements. The entire MBQC process can thus be represented in the quantum circuit model; whenever this circuit can be efficiently classically simulated, the MBQC cannot provide a quantum speedup. Therefore, the dynamics of a MBQC on a true universal resource state must not be classically simulatable. I will explain very shortly that, for example, the dynamics of local measurements on 1D cluster states can be efficiently simulated via the Matrix Product State (MPS) representation [40, 41, 42]; 1D cluster states are therefore not universal resources. Nevertheless, they are resources for universal logical single-qubit rotations, and are still worth considering. In many circumstances, such 1D resources can be coupled together into higher-dimensional structures that are indeed universal; the MPS description then allows for a convenient description of the method by which single-qubit rotations on the 1D resources are implemented.

Since the existence of entanglement alone does not explain the quantum speedup, the precise flavour of entanglement that is present must be important to the real explanation (I will clarify what this means over the course of the rest of the chapter). A step towards providing a fuller explanation, in the form of another negative result, was provided by Vidal [40], who showed that quantum computers generating a restricted amount of entanglement (in a sense that I will make precise in Sec. 1.5.1) could also be efficiently classically simulated, using the MPS formalism. Vidal also showed how the dynamics of an algorithm running on a quantum computer whose state is always efficiently represented as a MPS can be efficiently classically simulated, thereby completing the proof that quantum computations whose entanglement is suitably restricted can be efficiently classically simulated. I will take a brief interlude to introduce the MPS representation, as it is extraordinarily useful for demonstrating that certain quantum states can be universal resource states for MBQC; I will then discuss how to simulate the dynamics of quantum circuits with MBQC on MPS.

1.5.1 Matrix Product States

The matrix product state (MPS) representation is an extraordinarily useful form in which to express states that are not very entangled, in the following sense. Consider an n -qudit state $|\Psi\rangle$, defined on a Hilbert space $\mathcal{H}$, whose qudits are labelled from 1 to n in some way. $|\Psi\rangle$ can be written as

$$|\Psi\rangle = \sum_{i_1, i_2, \dots, i_n=1}^d \psi_{i_1 i_2 \dots i_n} |i_1\rangle |i_2\rangle \dots |i_n\rangle, \quad (1.31)$$

where each $\{|i_q\rangle\}$ is a complete, orthonormal basis for the Hilbert space of qudit q , and d denotes the dimension of each such Hilbert space. This state is “not very entangled”, and can therefore be efficiently represented with a MPS representation, if there exists a labelling of the qudits from 1 to n such that the von Neumann entropy of entanglement with respect to any bipartition $1, 2, \dots, k | k+1, \dots, n$ is bounded, i.e.

$$E_{1\dots k, k+1\dots n} \leq E \quad \forall k \in \{1, \dots, n\}, \quad (1.32)$$

where E is a constant that is independent of the size of the system. This property is not satisfied by a generic quantum state; the bipartition $1 \dots \lfloor n/2 \rfloor | \lfloor n/2 \rfloor + 1 \dots n$ yields two Hilbert subspaces at least one of which has dimension $d^{\lfloor n/2 \rfloor}$, in which case $E_{1\dots k, k+1\dots n}$ can be as large as $\lfloor n/2 \rfloor \log_2 d$, a monotonically increasing linear function of the system size n . Nevertheless, it turns out that a wide assortment of interesting quantum states in nature do satisfy this property, which I will discuss further in Ch. 3. The reason why this restriction on the entanglement leads to an efficient classical description of the state will become clear after I derive the MPS representation, which I will now do following the original presentation of Vidal [40].

The matrix product state representation can be derived by applying repeated Schmidt

decompositions to $|\Psi\rangle$. Let us assume the qudits are labelled from 1 to n . First, choose the bipartition having qudit 1 by itself in one partition, and the remaining qudits in the other. This bipartition induces the Schmidt decomposition

$$|\Psi\rangle = \sum_{\alpha_1=1}^d \lambda_{\alpha_1}^{[1]} |\Phi_{\alpha_1}^{[1]}\rangle |\Phi_{\alpha_1}^{[2\dots n]}\rangle \quad (1.33)$$

$$= \sum_{\alpha_1=1}^d \lambda_{\alpha_1}^{[1]} \left(\sum_{i_1=0}^1 \Gamma_{\alpha_1}^{[1]i_1} |i_1\rangle \right) |\Phi_{\alpha_1}^{[2\dots n]}\rangle \quad (1.34)$$

$$= \sum_{\alpha_1=1}^d \sum_{i_1=1}^d \Gamma_{\alpha_1}^{[1]i_1} \lambda_{\alpha_1}^{[1]} |i_1\rangle |\Phi_{\alpha_1}^{[2\dots n]}\rangle, \quad (1.35)$$

where $\Gamma_{\alpha_1}^{[1]i_1}$ is the $d \times d$ -dimensional unitary transformation relating the local and Schmidt bases for qudit 1, i.e.

$$|\Phi_{\alpha_1}^{[1]}\rangle := \sum_{i_1=1}^d \Gamma_{\alpha_1}^{[1]i_1} |i_1\rangle. \quad (1.36)$$

So far, we have introduced $\lambda^{[1]}$ from the Schmidt decomposition in terms of the bipartition $1|2\dots n$, and then $\Gamma^{[1]}$ by relating the Schmidt basis of qudit 1 to its local basis. Now, we want to repeat this process, so that in the end we obtain an expression for $|\Psi\rangle$ containing an analogous $\lambda^{[q]}$ and a $\Gamma^{[q]}$ for each qudit q . To introduce $\lambda^{[2]}$, we use the Schmidt decomposition of $|\Psi\rangle$ with respect to the bipartition $1, 2|3\dots n$:

$$|\Psi\rangle = \sum_{\alpha_1=1}^d \lambda_{\alpha_1}^{[2]} |\Phi_{\alpha_1}^{[2]}\rangle |\Phi_{\alpha_1}^{[3\dots n]}\rangle, \quad (1.37)$$

Consider the Schmidt basis states $\{|\Phi_{\alpha_1}^{[2\dots n]}\rangle\}$ for qudits 2 to n . Rewrite each of them such that qudit 2 is expressed in its own local basis,

$$|\Phi_{\alpha_1}^{[2\dots n]}\rangle = \sum_{i_2=1}^d |i_2\rangle |\tau_{\alpha_1, i_2}^{[3\dots n]}\rangle, \quad (1.38)$$

and then write the vectors $\{|\tau_{\alpha_1, i_2}^{[3\dots n]}\rangle\}$ in terms of the Schmidt basis vectors and coefficients for qubits $3\dots n$:

$$|\tau_{\alpha_1, i_2}^{[3\dots n]}\rangle = \sum_{\alpha_2=1}^{d^2} \Gamma_{\alpha_1, \alpha_2}^{[2]i_2} |\Phi_{\alpha_2}^{[3\dots n]}\rangle. \quad (1.39)$$

Notice that the sum over α_2 in Eq. 1.39 goes from 1 to d^2 , whereas the one over α_1 in Eq. 1.35 only goes from 1 to d . This is because the smaller of the two parts of the system in the bipartition $1, 2|3\dots n$ has a Hilbert space of dimension d^2 rather than d . The next step is to substitute Eq. 1.39 into Eq. 1.38, and then the result of that back into Eq. 1.35, giving

$$|\Phi_{\alpha_1}^{[2\dots n]}\rangle = \sum_{i_2=1}^d \sum_{\alpha_2=1}^{d^2} \Gamma_{\alpha_1, \alpha_2}^{[2]i_2} \lambda_{\alpha_2}^{[2]} |i_2\rangle |\Phi_{\alpha_2}^{[3\dots n]}\rangle \quad (1.40)$$

and thus

$$|\Psi\rangle = \sum_{i_1=1}^d \sum_{i_2=1}^d \sum_{\alpha_1=1}^d \sum_{\alpha_2=1}^{d^2} \Gamma_{\alpha_1}^{[1]i_1} \lambda_{\alpha_1}^{[1]} \Gamma_{\alpha_1, \alpha_2}^{[2]i_2} \lambda_{\alpha_2}^{[2]} |i_1\rangle |i_2\rangle |\Phi_{\alpha_2}^{[3\dots n]}\rangle. \quad (1.41)$$

This process can be repeated on $|\Phi_{\alpha_2}^{[3\dots n]}\rangle$, and so on until all of the qudits have been treated, resulting in a final expression for the coefficients $\psi_{i_1 i_2 \dots i_n}$

$$\psi_{i_1 i_2 \dots i_n} = \sum_{\alpha_1, \alpha_2, \dots, \alpha_n} \Gamma_{\alpha_1}^{[1]i_1} \lambda_{\alpha_1}^{[1]} \Gamma_{\alpha_1, \alpha_2}^{[2]i_2} \lambda_{\alpha_2}^{[2]} \dots \Gamma_{\alpha_n}^{[n]i_n}, \quad (1.42)$$

where the sums over α_k for $k \in \{1, \dots, \lfloor n/2 \rfloor\}$ go from 1 to d^k , and those for $k \in \{\lfloor n/2 \rfloor + 1, \dots, n\}$ go from 1 to d^{n+1-k} . This expression can be interpreted as a multiplication of n matrices of varying dimensions; defining the d row vectors $\Gamma^{[1]}[i_1] := (\Gamma_{\alpha}^{[1]i_1})_{\alpha}$, column vectors $\Gamma^{[n]}[i_n] := (\Gamma_{\alpha}^{[n]i_n})_{\alpha}$, non-square matrices $\Gamma^{[k]}[i_k] := (\Gamma_{\alpha, \beta}^{[k]i_k})_{\alpha, \beta}$ and diagonal square matrices $\Lambda^{[k]} := (\lambda_{\alpha}^{[k]})_{\alpha, \alpha}$, we obtain the MPS representation of the

coefficients of $|\Psi\rangle$ (with the dimensions of the matrices indicated):

$$\psi_{i_1 i_2 \dots i_n} = \overbrace{\Gamma^{[1]}[i_1]}^{1 \times d} \underbrace{\Lambda^{[1]}}_{d \times d} \overbrace{\Gamma^{[2]}[i_2]}^{d \times d^2} \underbrace{\Lambda^{[2]}}_{d^2 \times d^2} \dots \overbrace{\Gamma^{[n-1]}[i_{n-1}]}^{d^2 \times d} \underbrace{\Lambda^{[n-1]}}_{d \times d} \overbrace{\Gamma^{[n]}[i_n]}^{d \times 1}. \quad (1.43)$$

Eq. 1.43 is a useful expression for the MPS representation of $|\Psi\rangle$, as its Schmidt coefficients with respect to any bipartition $1, \dots, k | k+1, \dots, n$ can be read off from the diagonal elements of $\Lambda^{[k]}$. However, for ease of notation (and another reason I shall shortly describe in the context of MBQC), the $\Gamma\Lambda$ products are often collapsed into a single matrix. Defining $A^{[k]}[i_k] := \Gamma^{[k]}[i_k]\Lambda^{[k]}$ for $k \in \{1, 2, \dots, n-1\}$ and $A^{[n]}[i_n] := \Gamma^{[n]}[i_n]$, the MPS representation for $|\Psi\rangle$ becomes

$$\psi_{i_1 i_2 \dots i_n} = \sum_{i_1, i_2, \dots, i_n=1}^d \overbrace{A^{[1]}[i_1]}^{1 \times d} \overbrace{A^{[2]}[i_2]}^{d \times d^2} \dots \overbrace{A^{[n-1]}[i_{n-1}]}^{d^2 \times d} \overbrace{A^{[n]}[i_n]}^{d \times 1}. \quad (1.44)$$

This particular form is referred to as the open boundary conditions MPS (OBC-MPS) representation, but other forms exist; see e.g. Ref. [41]. Now, specifying the form of a general n -qudit state requires d^n parameters; how many are required to give the MPS representation of the state? Since the derivation of Eq. 1.44 did not place any additional restrictions on the state, we should not expect to do significantly better (and indeed, we do not). Assume for simplicity's sake that n is even, and consider the set $\{A^{[k]}[i_k]\}$, with $k \leq \lfloor n/2 \rfloor$. There are $n/2$ such sets (labelled by k), and each set contains d matrices (labelled by i_k) with d^{2k-1} matrix elements, requiring a total of $\sum_{k=1}^{n/2} d^{2k}$ free parameters to be specified. The largest dimension of any of the MPS matrices in a MPS representation of a state is called the bond dimension of the MPS, and for this completely general qudit state, it is equal to $d^{\lfloor n/2 \rfloor}$. To describe the whole state, one needs twice as

many parameters to include the cases $n/2 < k \leq n$, for a total of

$$2 \sum_{k=1}^{n/2} d^{2k} = \frac{2}{d-1} (d^{n/2+1} - 2d + 2). \quad (1.45)$$

This is a minor improvement on the d^n parameters normally required, but is still an exponentially growing function of n . On the other hand, what if the maximum number of non-zero Schmidt coefficients (i.e. Schmidt rank) with respect to any bipartition $1 \dots k \mid k+1 \dots n$ of the system is some constant χ that is independent of the number of qudits in either partition? In that case, considering Eq. 1.43, the dimension of each matrix $\Lambda^{[k]}$ is, as before, $d^k \times d^k$ provided $d^k \leq \chi$. However, if the index k is sufficiently close to the middle index $\lfloor n/2 \rfloor$, such that $\log_d(\chi) < |k - n/2|$, then the dimension of $\Lambda^{[k]}$ becomes $\chi \times \chi$ instead. In this case, the bond dimension of the MPS is χ , a constant, and, the maximum von Neumann entropy of entanglement with respect to any bipartition is $\log_2(\chi)$. In other words, the entropy of entanglement no longer scales with the size of the subsystems. It turns out to be the case (and I will elaborate on this in Ch. 3) that the ground states of many physically realistic one-dimensional lattice Hamiltonians share this property [42], where the qudits are labelled in the natural way according to their spatial coordinates. This is a special case of what is known as an area law for the entanglement entropy. In general, the entanglement entropy of a quantum state on a z -dimensional lattice is said to obey an area law if it scales like $O(n^{z-1})$, corresponding to the surface area of the subsystems induced by the bipartitions, rather than the volumes [43].

1.5.2 MBQC in Correlation Space via MPS Representation

The MPS representation of a resource state often allows for a nice interpretation of MBQC as occurring in correlation space, the name for the vector space in which the MPS matrices are defined. To see this, consider a slight modification of Eq. 1.44: instead

of having the matrices $A^{[k]}[i_k]$ for the end points $k = 1, n$ be vectors of dimension $1 \times d$ and $d \times 1$ respectively, imagine defining the constant vectors $\langle L|$ and $|R\rangle$ such that $A^{[1]}[i_1] = \langle L|B^{[1]}[i_1]$ and $A^{[n]}[i_n] = B^{[n]}[i_n]|R\rangle$ for some matrices $\{B^{[1]}[i_i]\}$ and $\{B^{[n]}[i_n]\}$. Also define the renaming $B^{[k]}[i_k] := A^{[k]}[i_k]$ for all sites in the bulk, $2 \leq k \leq n-1$. Then, we can write the modified form of Eq. 1.44 as

$$\psi_{i_1 i_2 \dots i_n} = \langle L|B^{[1]}[i_1]B^{[2]}[i_2] \dots B^{[n]}[i_n]|R\rangle. \quad (1.46)$$

Suppose that $|\Psi\rangle$ satisfies an area law, with $\chi \leq d$. Then, the dimensions of $\langle L|$ will be $1 \times \chi$, of $\langle R|$ will be $\chi \times 1$, and of each of the $B^{[k]}[i_k]$ will be $\chi \times \chi$. In this case, even though the physical system is a chain of qudits with dimension d , we can view the effect of a MBQC in which qudits are measured sequentially along the chain as being to process a virtual qudit of dimension χ in correlation space. To see this, observe the effect of a projective local measurement on the first qudit in the chain. In particular, suppose the state of the qudit is collapsed to $|\Phi^{[1]}\rangle = \sum_{i_1=1}^d \phi_{i_1}^{[1]}|i_1\rangle$. Consider the effect of the projector $P_{\Phi^{[1]}} := |\Phi^{[1]}\rangle\langle\Phi^{[1]}|$ acting on the first qudit of $|\psi\rangle$:

$$\begin{aligned} P_{\Phi^{[1]}}|\Psi\rangle &= \sum_{i_1, i_2, \dots, i_n=1}^d \langle L|B^{[1]}[i_1]B^{[2]}[i_2] \dots B^{[n]}[i_n]|R\rangle (|\Phi^{[1]}\rangle\langle\Phi^{[1]}|i_1\rangle) |i_2 \dots i_n\rangle \\ &= \sum_{i_2, \dots, i_n=1}^d \langle L| \left(\sum_{i_1=1}^d \langle\Phi^{[1]}|i_1\rangle B^{[1]}[i_1] \right) B^{[2]}[i_2] \dots B^{[n]}[i_n]|R\rangle |\Phi^{[1]}\rangle |i_2 \dots i_n\rangle \\ &= |\Phi^{[1]}\rangle \otimes \sum_{i_2, \dots, i_n=1}^d \langle L|B^{[1]}[\Phi^{[1]}]B^{[2]}[i_2] \dots B^{[n]}[i_n]|R\rangle |i_2 \dots i_n\rangle, \end{aligned}$$

where I have defined $B^{[1]}[\Phi^{[1]}] := \sum_{i_1=1}^d \langle\Phi^{[1]}|i_1\rangle B^{[1]}[i_1]$, which is a linear combination of the MPS matrices corresponding to site 1 with coefficients determined by the result of the measurement on that site. This can be interpreted logically in correlation space as initialisation of a χ -dimensional qudit $\langle L|$, followed by a logical operation represented

by the χ -by- χ matrix $B^{[1]}[\Phi^{[1]}]$ on that qudit. Carrying on in this manner all the way along the chain and defining $\{\Phi^{[k]}\}$ and $B^{[k]}[\Phi^{[k]}]$ analogously for the remaining sites $k = 2, \dots, n$, the effect in correlation space is the logical mapping

$$\langle L | \mapsto \langle L | B^{[1]}[\Phi^{[1]}] B^{[2]}[\Phi^{[2]}] \dots B^{[n]}[\Phi^{[n]}], \quad (1.47)$$

and the probability of this particular sequence of measurement outcomes to occur is given by the overlap of this vector with $|R\rangle$.

Note that this MPS representation, with a fixed bond dimension independent of system size, is only accurate for systems obeying an entanglement area law in one dimension, and does not generalise to higher dimensions. A general state in any dimension can be represented exactly with MPS, but the bond dimension required will generally be an increasing function of the system size. For systems obeying an entanglement area law in higher dimensions, a generalisation of MPS called the Projected Entangled Pair State (PEPS) representation [44, 45] is needed. The coefficients of a state in the PEPS representation are expressed as a contraction of a tensor network whose ranks depend on the dimensionality of the system. In 1D, the tensors in the tensor networks are all rank two; they can therefore be interpreted as matrices and the contraction as a sequence of matrix multiplications – PEPS reduces to MPS in 1D.

I will now provide some examples of how to interpret local measurements on 1D resource states as processing a logical qudit in correlation space.

1.5.3 Example: 1D Cluster State as a MPS

The one-dimensional cluster state with open boundary conditions (i.e. on a chain) obtains an extraordinarily convenient representation as an OBC-MPS in the form of Eq. 1.46. Not only does the 1D cluster state obey an area law satisfying $\chi = 2$, the matrices $B^{[k]}[i_k]$

are translationally invariant, i.e. they can be chosen to be independent of the site index k . The description of the functioning of the 1D cluster state as a resource for universal single-qubit rotations I give in Ch. 2 is in terms of the quantum circuit model; however, the MPS description is also illuminating. One choice of the matrices for the 1D cluster state with open boundary conditions is given in Ref. [46]. Dropping the site label k because of the translational invariance, the MPS representation is given by

$$\langle L| := \langle +|; \quad (1.48)$$

$$|R\rangle := |0\rangle; \quad (1.49)$$

$$B[0] := |0\rangle\langle +|; \quad (1.50)$$

$$B[1] := |1\rangle\langle -|. \quad (1.51)$$

Thus, we interpret the qubit in correlation space as being initialised in the state $\langle +|$. The effect of projecting a physical qubit on to a state in the $x - y$ plane of the Bloch sphere, i.e. the state $|\Phi\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{-i\xi}|1\rangle)$ is to act in correlation space with the operator

$$B[\Phi] = \frac{1}{\sqrt{2}} (|0\rangle\langle +| + e^{-i\xi}|1\rangle\langle -|) \quad (1.52)$$

$$\equiv \frac{1}{\sqrt{2}} \left(\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 \\ e^{-i\xi} & -e^{-i\xi} \end{bmatrix} \right) \quad (1.53)$$

$$\equiv R_z[-\xi] H. \quad (1.54)$$

Similarly, projecting the physical qubit on to the orthogonal complement of $|\Phi\rangle$, given by $|\Phi'\rangle := \frac{1}{\sqrt{2}}(|0\rangle - e^{-i\xi}|1\rangle)$ gives

$$B[\Phi'] \equiv R_z[-\xi] HX. \quad (1.55)$$

Bearing in mind that these operators act on a vector to the left in correlation space, the effect in correlation space is exactly the same as the one in real space: measuring a qubit in the basis $\{|\Phi\rangle, |\Phi'\rangle\}$ teleports the virtual qubit in correlation space through the logical gate $X^m H R_z[\xi]$, where $m = 0(1)$ corresponds to the outcome $|\Phi\rangle$ ($|\Phi'\rangle$).

1.5.4 Example: Ground State of a Quantum Spin Network as a Quantum Wire

One example of a universal resource for single-qubit operations, from Refs. [46, 47], is a computational wire based on a modification of a famous 1D, nearest-neighbour, two-body lattice model on spin-1 particles, the Affleck-Kennedy-Lieb-Tasaki (AKLT) model [48].

A generalised version of the AKLT model on a chain of N spin-1 particles is given by the Hamiltonian

$$H_{\text{Haldane}} = J \sum_{j=1}^{N-1} [\mathbf{S}_j \cdot \mathbf{S}_{j+1} - \beta (\mathbf{S}_j \cdot \mathbf{S}_{j+1})^2], \quad (1.56)$$

with $J > 0$. The case where $\beta = 1/3$ corresponds to the AKLT model, whereas $\beta = 0$ is the more familiar Heisenberg model. The Hamiltonian is gapped within the entire regime $-1 < \beta < 1$ containing both the Heisenberg and AKLT points, called the Haldane phase [49]. The ground state of the AKLT model (i.e. $\beta = 1/3$) can be exactly represented as a MPS with bond dimension $\chi = 3$, making it convenient to check for the possibility of its suitability as a universal computational wire for MBQC.

The MPS representation of the AKLT ground state is available [48, 47], but for convenience of analysis the AKLT-type state of Refs. [46, 47] has a slightly different MPS representation. The logical qubits are encoded in 1D logical AKLT-like wires,

which are composed of qutrits rather than qubits. The MPS representation is

$$B[0] := H; \tag{1.57}$$

$$B[1] := 2^{-1/2}|0\rangle\langle 1|; \tag{1.58}$$

$$B[2] := 2^{-1/2}|1\rangle\langle 0| \tag{1.59}$$

(for the real AKLT ground state, the only difference is $B[0] := Z$, but the state here is also the ground state of a nearest-neighbour frustration-free gapped spin-1 Hamiltonian).

First consider the implementation of an identity gate (i.e. faithful teleportation) via measurement in the basis $\{|0\rangle, 2^{-1/2}(|1\rangle \pm |2\rangle)\}$. The three possible measurement outcomes yield the correlation space operators H , X and ZX respectively, and the measurement outcomes have equal probability. These three operators can all be considered to be unwanted byproduct operators. The operators $\{H, X, Z\}$ generate a finite group called the byproduct group $\mathcal{B}$. By stringing together a sequence of measurements in the same basis, each time one teleports the logical information through a random member of $\mathcal{B}$. The net effect after m such measurements is to have teleported the initial information through some gate $B_m \in \mathcal{B}$. After a finite expected number of measurements m' , we will obtain $B_{m'} = I$. Thus, faithful teleportation is accomplished with a random but finite number of steps; the same technique will be used in Ch. 2 when discussing the universality of $N - U - N$ states.

Now consider a universal set of single logical qubit rotations. To implement $R_z(\phi)$, one can measure in the basis $\{|0\rangle, 2^{-1/2}(|1\rangle \pm e^{i\phi/2}|2\rangle)\}$. If the outcome is $2^{-1/2}(|1\rangle + e^{i\phi/2}|2\rangle)$, the operation is successful; otherwise, some byproduct from $\mathcal{B}$ acts after $R_z(\phi)$, which can be removed by the random walk technique. Similarly, one can implement H by measuring in the computational basis $\{|0\rangle, |1\rangle, |2\rangle\}$, removing byproducts as needed. The full group $SU(2)$ is generated by $\{H, HR_z(\phi)H\}$.

In fact, it has subsequently been shown how to compute via MBQC in the degenerate ground state space of the 1D AKLT model, supplemented by two spin-1/2 particles coupled to either end of the chain [50]. The AKLT model itself has a four-fold degenerate ground state space, which is reduced to 2 by the spin-1/2 particles, and a logical qubit is encoded in this space. Later, it was shown [51] that a similar modification of the generalised AKLT model in 1D from Eq. 1.56 yields a universal quantum wire as its ground state for any value of β in the Haldane phase, even though away from the AKLT point the Hamiltonian is not frustration-free, and an analytical expression for the MPS representation of the ground state is not known. For completeness, I also mention that in Refs. [50, 51], it is discussed how to produce a true universal resource by coupling together parallel wires laid out horizontally via interactions in the vertical direction.

1.6 Projected Entangled Pair States Representation and 2D Universal Resources

So far, I have only discussed MBQC in correlation space with 1D states, which cannot be true universal resources for MBQC. A generalisation of the MPS representation allowing for efficient descriptions of higher-dimensional states exists, and is called the Projected Entangled Pair States (PEPS) representation [52, 53]. While it is not necessary for the purposes of Ch. 2 to describe this in detail, I include a brief description here for completeness.

The PEPS representation of a state is specified by giving the coefficients of the state in the computational basis as a tensor contraction, much like in the MPS case; the difference here is that the rank of the tensors appearing in the tensor network can be greater than two. For example, the 2D cluster states admit a simple PEPS representation with tensors of rank four. As in the MPS case, one specifies two tensors for each qubit,

and in general d tensors for each d -level system. It is rather cumbersome writing down the tensor contraction by listing all of the indices of each tensor and indicating which ones are summed over. Conveniently, there is a graphical representation that makes this much easier [46, 47]. In the tensor network, each tensor is represented as a box, with a number of legs emerging from the box that corresponds to the rank of the tensor. An index is contracted by joining two legs together, either from two different boxes or from the same box.

Rank-four tensors can be naturally interpreted as two-qubit operators in correlation space. A rank-four tensor is drawn as a box with four legs. One could imagine putting arrows on the legs, so that two of the legs become “inputs” to the box, and the other two become “outputs”. The operation in correlation space can then be viewed as acting on the information coming in from the input legs, which is modified by the action of the tensor, and then transmitted to the two output legs. Refs. [46, 47] provide a much more detailed description, and also use this representation to describe how to implement entangling gates in correlation space for the 2D cluster states. Furthermore, they construct a 2D version of the AKLT-like wires from above. The resource is a family of 2D states, each consisting of some number of AKLT-like chains with equal lengths, laid out side-by-side and vertically. Neighbouring qutrits in the horizontal direction are then coupled using a generalisation of the CZ gate, which applies a phase of $e^{i\pi}$ whenever both qutrits are in state $|2\rangle$.

1.7 Identifying and Characterising Resource States

Certain quantum states can act as a consumable resource for quantum information processing protocols, e.g. Bell states for quantum state teleportation, or cluster states for MBQC. As I have already discussed, the cluster states are a family of resource states

for universal MBQC. However, they are not the only possible resource state; for example, the AKLT-based resource state from the previous section [54], or indeed the AKLT state itself on spin-3/2 particles on a hexagonal lattice [55], are resources for universal MBQC as well, as indeed are the natural generalisations of cluster states from square to hexagonal, triangular or Kagome lattices [56]. However, the property of being a universal resource state appears to be extremely rare; several no-go results are known that eliminate a state or family of states from being a suitable MBQC resource. Unsurprisingly, these results generally focus on the type and quantity of entanglement present in the candidate resource state. The notions of “type” and “quantity” of entanglement need to be formalised, which I will now do.

1.7.1 Quantifying entanglement: Entanglement Monotones and LOCC

The amount of entanglement present in a multipartite quantum state $|\psi\rangle$ is often quantified via functions called type-I entanglement monotones (see Ref. [57] for an excellent review). These are functions of $|\psi\rangle$ having the property that the parties holding $|\psi\rangle$ cannot increase their values (on average) by acting locally on their part of $|\psi\rangle$, and communicating classically with each other. This set of allowed operations is called Local Operations and Classical Communication (LOCC). Suppose n parties share a single quantum state, represented by a density matrix ρ . The effect of performing LOCC operations on ρ is to map it to an ensemble $\{p_i, \rho_i\}$, where

$$\rho_i := \frac{(A_1^i \otimes \cdots \otimes A_n^i) \rho (A_1^i \otimes \cdots \otimes A_n^i)^\dagger}{p_i};$$

$$p_i := \text{Tr}((A_1^i \otimes \cdots \otimes A_n^i) \rho (A_1^i \otimes \cdots \otimes A_n^i)^\dagger).$$

Here, the $\{A_j^i\}$ are all invertible (not necessarily unitary) operators acting locally on the j th subsystem, and satisfy $(A_j^i)^\dagger A_j^i \leq I$. Furthermore, the $\{\rho_i\}$ can be chosen to

be pure states, and I will assume without loss of generality that they are. Since each party can act conditionally based on messages received from other parties, and since local measurements with random outcomes may be involved in the LOCC protocol, there may be a number of different branches the protocol can take; one can think of i as labelling the branch, ρ_i as the output of that branch and p_i as the probability of that branch occurring. If the precise branch followed is known, then the outcome is actually the pure state ρ_i ; if it is completely unknown, then the outcome is the mixed state $\rho' := \sum_i p_i \rho_i$.

The LOCC concept can be used to define measures of entanglement. A function $\mathcal{E}(\rho) : \mathbb{C}^m \mapsto \mathbb{R}^+$ from an m -dimensional quantum state to the non-negative real numbers is defined to be a type-I entanglement monotone if

1. $\mathcal{E}(\rho) = 0$ iff ρ is separable;
2. it is non-increasing on average under LOCC [58]; that is, if for every possible ensemble $\{p_i, \rho_i\}$ resulting from a LOCC protocol on ρ , $\mathcal{E}(\rho)$ satisfies

$$\mathcal{E}(\rho) \geq \sum_i p_i \mathcal{E}(\rho_i). \quad (1.60)$$

If it is possible to guarantee the output state ρ' resulting from a LOCC protocol on an input state ρ (such as, for example, if each party merely performs a local unitary), the protocol is called deterministic. In this case, ρ is said to be convertible to ρ' under deterministic LOCC, sometimes denoted as $\rho \geq_{\text{LOCC}} \rho'$. A function that is non-increasing under deterministic LOCC, i.e. $\mathcal{E}(\rho) \geq \mathcal{E}(\rho')$ whenever $\rho \geq_{\text{LOCC}} \rho'$, is called a type-II entanglement monotone.

The reader is advised that type-I entanglement monotones are more commonly referred to as simply “entanglement monotones” in the literature; my choice of nomenclature is intended to emphasise the distinction between the type-I and type-II versions, as is done in Ref. [56]. Each of these two types of measures has been used to prove

some powerful results regarding the (non-)suitability of quantum states as resources for MBQC. I will next formalise the notion of universality in MBQC, and discuss how type-I and type-II entanglement monotones rule out the vast majority of quantum states as resources for universal MBQC.

1.7.2 Entanglement requirements for universal MBQC: resource states are rare

In Ref. [56], the authors consider a form of universality they call CQ-universality, which refers to a measurement-based quantum computer that is able to accept an arbitrary classical state as input and produce an arbitrary quantum state as output. There are other forms of universality: CC-, QC- and QQ-, with the obvious meanings, but CQ-universality is the most general kind that is possible within the MBQC model. It is worth pointing out that CC-universality is generally the most that is strictly necessary, since the readout phase of a quantum algorithm ensures the final output is a classical string of bits, and that CQ-universality is a stronger requirement. In this work, universality is a property that is ascribed not to a particular quantum state but to a family of states, in order to avoid having to explicitly consider infinitely large states. The authors define a universal resource for MBQC to be a family of states $\Psi = \{|\psi_1\rangle, |\psi_2\rangle, \dots\}$ with $|\Psi| = \infty$, such that for every n -qubit quantum state $|\phi_{\text{out}}\rangle$, there exists a resource state $|\psi_i\rangle \in \Psi$ such that $|\psi_i\rangle \geq_{\text{LOCC}} |\phi_{\text{out}}\rangle$. The motivation for this definition is that the local measurements and classical feed-forward of measurement results allowed within MBQC constitute a LOCC protocol on the state (an element of the resource) acting as a substrate for the computation; therefore, if an arbitrary quantum state must be preparable as output, there must exist an element of the resource that is convertible to the desired output state under deterministic LOCC.

Type-II entanglement monotones have been used [56] to prove some powerful negative results regarding the usefulness of families of quantum states as resources for MBQC. One

of the major results of this work is that if a family of states Ψ constitutes a resource for universal MBQC, then every type-II entanglement monotone must reach its supremum over Ψ . The key observation required to prove this is that it is true for the family of $n \times n$ square cluster states on n^2 qubits. Any (CQ-)universal resource must be capable of producing any quantum state, including the cluster states, as its output. The result then follows because the MBQC protocol cannot increase the value of any type-II entanglement monotone. Therefore, if some type-II entanglement monotone does not reach its supremum on Ψ , then there must be cluster states that are not preparable from any state in Ψ .

In addition to this entanglement-based requirement for universality, the same work provides a necessary condition for a family of states to be an efficient resource. In this context, efficient means essentially that the number of particles, qubits and amount of classical side-processing required to implement some algorithm via MBQC with this resource state are all polynomial in the size of the input. In order for a family of states to be an efficient resource, certain scaling relations of various type-II entanglement monotones must be obeyed. More specifically, efficient universality of a resource Ψ (which is a family of states with different system sizes) requires every type-II entanglement monotone to grow faster than logarithmically in the system size on Ψ . Another useful quantity called the geometric measure of entanglement, defined by

$$E_g(|\psi\rangle) := \max_{|\phi\rangle \in \mathcal{H}^1 \otimes \mathcal{H}^2 \otimes \dots \otimes \mathcal{H}^n} |\langle\psi|\phi\rangle|^2 \quad (1.61)$$

is the square of the maximum overlap of an n -partite state $|\psi\rangle$ with a product state, introduced in Ref. [59] and shown to be an entanglement monotone in Ref. [60]. The closely related quantity $\pi(|\psi\rangle) := -\log_2 E_g(|\psi\rangle)$ is a type-II entanglement monotone [56], and can be used to rule out the family of W-states as constituting a universal resource for

MBQC [56]. Very similar results hold for weaker types of MBQC resources, namely those that are only required to be universal for approximate and/or stochastic MBQC [61].

These no-go theorems indicate that families of states that are insufficiently entangled (as measured by type-II entanglement monotones) are not CQ-universal resources for efficient MBQC. It is tempting to assume that the converse is also true, namely that non-universal resources are necessarily ones that are insufficiently entangled. If this were the case, it would be wonderful news, for generic multipartite quantum states are highly entangled, in several senses. For example, the entropy of entanglement of a single subsystem with respect to the rest of the state is typically close to maximal [62, 63, 64], random $2l$ -qubit states are resources for teleportation of $O(l)$ qubits [65], and have near-maximal entropy of entanglement with respect to any bipartition [66]. Sadly, this assumption turns out to be false; in fact, most quantum states are in a sense “too entangled” to be resources.

One can view measurements on the resource state as a tool to augment the power of the classical computer doing the side-processing. It turns out that generic states can in fact be replaced by an unbiased random number generator with no change in overall computational power [67, 68]. In fact, in Ref. [67], it is shown that the very same geometric measure of entanglement $\pi(|\psi\rangle)$ from the previous paragraph that must reach its supremum on any universal resource. For a family of states whose elements are labelled by the number of qubits n to be a resource, the geometric measure must thus be an increasing function of the number of qubits n and must also increase no faster than linearly in n . The fraction of n -qubit states (for $n \geq 11$) for which this is true is less than e^{-n^2} . Furthermore, while the results of Ref. [56] are valid in the case of CQ-universality, which is arguably more than is required for useful MBQC, the conclusions of Refs. [67, 68] are valid in the less stringent case of CC-universality as well.

In summary, states that are insufficiently entangled, too entangled, or whose dynamics

under LOCC can be efficiently simulated (via, for example, the stabiliser formalism or the MPS or PEPS representations) are all useless for MBQC. Clearly the property of being a universal resource is uncommon amongst quantum pure states.

1.7.3 Physically interesting universal resource states for MBQC

Despite the scarcity of universal resource states, a variety of them beyond the cluster states have been discovered [69, 54, 47, 50, 70, 71, 51, 72, 55]. The work of Gross and Eisert [54, 47], beyond developing the framework for interpreting MBQC in correlation space via the MPS representation, also provided several examples of resource states that differ materially from the cluster states in various physically relevant respects.

One example of relevance to this thesis concerns correlations present within the state. One characteristic feature of cluster states is their spin-spin correlation functions. Let the operator $S_{\alpha,i} = \vec{S}_i \cdot \hat{\alpha}$ denote the component of the spin $\vec{S}_i$ of particle i with total spin s in the direction $\hat{\alpha} \in \mathbb{R}^{2s+1}$. Then, the spin-spin correlation function

$$C_{i,j}(S_{\alpha}, S_{\beta}) = \langle S_{\alpha,i} S_{\beta,j} \rangle - \langle S_{\alpha,i} \rangle \langle S_{\beta,j} \rangle \quad (1.62)$$

is non-zero for cluster states only if qubits i and j are adjacent, irrespective of the spin projection axes $\hat{\alpha}$ and $\hat{\beta}$ chosen. This feature is actually utilised in the original proof of universality of the cluster states for MBQC; the fact that distant measurement results are uncorrelated allows the computation to be broken up into small pieces corresponding to the simulation of individual circuit model gates, with each of these pieces analysed separately. The work in Refs. [54, 47] were the first to provide an example of a resource state in which these correlation functions are not strictly nearest-neighbour, but instead are slowly decaying functions obeying a power law in $|i - j|$. The state is exactly the 2D AKLT-like state discussed in Sec. 1.6. The resulting 2D state has power-law correlations

in the horizontal direction.

Another example concerns the local entanglement properties of the individual particles. Each qubit in a 2D square cluster state shares 1 ebit of entanglement with the rest of the state; in other words, each local reduced density matrix is maximally mixed. This property too could be conjectured to be necessary for a multipartite state to be a universal resource; after all, one ebit of shared entanglement is necessary for perfect state teleportation, with any lesser amount resulting in imperfect teleportation fidelity. However, it turns out that even states where each qubit shares arbitrarily little entanglement with respect to the rest of the state can still be universal resources, as again shown in Refs. [54, 47]. In Chapter 2, I describe two different classes of resource states; the first features non-vanishing correlation functions beyond nearest-neighbour (though they decay exponentially rather than as a power law), while the second has qubits sharing less than a full ebit of entanglement with the rest of the state (though not arbitrarily little; there is a critical threshold below which the state ceases to be universal).

The AKLT-based example of Refs. [54, 47] is especially appealing as it is the unique ground state of a two-body, nearest-neighbour, gapped spin Hamiltonian (albeit one corresponding to an artificial toy model with unphysical interactions). Imposing the requirement that the resource state be the ground state of some naturally occurring system has led to the discovery of further resource states. The main desiderata usually cited for the system in question are that it should

- (I) have a ground state that is universal for MBQC;
- (II) have a unique, non-degenerate ground state;
- (III) exhibit a finite energy gap between the ground state and the first excited manifold;
- (IV) be 2-local, i.e. have only terms featuring interactions between at most 2 particles;

(V) be frustration-free, i.e. have the global ground state also be the ground state of each term in the Hamiltonian separately.

The reason for condition (I) is obvious. Conditions (II-III) are desirable to protect the state against thermal errors. Condition (IV) is motivated by the lack of intrinsic $k > 2$ -body interactions in nature. Condition (V) ensures that as the computation proceeds, the portion of the resource state that remains unmeasured is still in its ground state configuration. I should mention that conditions (II) and (V) are actually violated by the AKLT resource of Refs. [50, 51], but in those cases the information is encoded in a non-trivial global symmetry degree of freedom with topological protection from thermal errors. A natural first step is to target the cluster state on a suitable $d > 1$ -dimensional lattice as the ground state. Unfortunately, it was shown that no Hamiltonian satisfying conditions (I-V) can have the cluster state [73], or indeed any graph state [56] as its unique ground state. For the cluster states in particular, the condition that is violated is number (IV); in [73], it is shown that Hamiltonians having d -dimensional cluster states as their ground states must be at least $(2d + 1)$ -local; this corresponds to 3-local for 1D cluster states and 5-local for 2D ones. The various conditions can, of course, be relaxed; I will now discuss progress that has been made along these lines.

An early approach was to relax the requirement that the ground state be *exactly* a cluster state, and instead design a Hamiltonian that approximates the low-energy sector of the five-body parent Hamiltonian for the square cluster states with two-body operations. Approaches along these lines include schemes that utilise ancilla particles to (approximately) mediate effective $k > 5$ -body interactions [74, 75], or that have ground states approximately encoding logical n -qubit cluster states on $n' > n$ physical qubits [69]. Another possible idea is to seek Hamiltonians whose ground states encode resources other than graph/cluster states, and this has proven to be a fruitful direction of research. Additionally, there is the previously discussed AKLT-based example of Refs. [54, 47]. The

MPS representation of the 1D AKLT chain has been exploited to develop schemes for using the ground state space of the AKLT [48] model to encode computational wires with dynamical coupling between adjacent wires used to simulate two-qubit operations [50]. The proof of universality was later extended to the entire Haldane phase [76, 51], though not using the MPS representation, since it is only efficient for this model precisely at the AKLT point.

Of course, it would be more desirable to produce the 2D universal resource state directly as the ground state, and exciting progress has been made on these fronts as well. A legitimate 2D resource with a Hamiltonian satisfying conditions I-V was identified by Chen et al. [70], with the system comprising spin-5/2 particles on a honeycomb lattice. While this proposal constitutes a promising first step, it suffers from two major drawbacks. The first is that the Hamiltonian in question has a rather complicated form and would be implausible to engineer with current technology; the second is the requirement of spin-5/2 particles, which are not as easily available as lower-spin particles. Two more proposals partially address these problems; one uses the simpler AKLT model [55] and the other a slight variation [72], and both can be implemented with spin-3/2 particles. Further reductions of the spin are still desirable. It remains an open question whether there is a spin-1 Hamiltonian satisfying conditions I-V, but it has unfortunately been ruled out for spin-1/2 Hamiltonians [77].

1.8 Stochastic Local Operations and Classical Communication

As was established in Sec. 1.7.2, universal resource states for MBQC are rare. However, they do exist: besides the cluster states, Sec. 1.7.3 introduced several examples. While a few general ideas, such as the MPS and PEPS formalisms, have assisted in the identification of these states, the process has so far been ad hoc. The no-go theorems described

in Sec. 1.7.2 rule out certain families of states, but do not provide a constructive method for finding resources. Such a constructive method would be valuable.

It is therefore worth considering whether there is a wide class of pure states, defined in some systematic way, whose members are all universal resources. Furthermore, this class would be most useful if it contained every possible universal resource, and no non-resource. The identification of this class, were it to exist, would be tantamount to providing a set of necessary and sufficient conditions for a family of states to be a resource. This would then exhaustively answer the question of which physical properties a family of pure states must possess in order to be a universal resource.

The identification of such a class would be a triumph, and the accomplishment of Ch. 2 [38] is to make some modest progress in this direction. If such a class exists, then it obviously must include the family of cluster states, as well as any other previously identified universal resource states. The idea of Ch. 2 is to define a family of useful resource states by means of its relation to the cluster states, and the specific relationship chosen utilises a particular case of the LOCC concept discussed in Sec. 1.7.1, called Stochastic Local Operations and Classical Communication (SLOCC).

Two n -qubit pure states $|\psi\rangle$ and $|\phi\rangle$ are said to be LOCC-equivalent if and only if $|\psi\rangle$ can be obtained from $|\phi\rangle$, and $|\phi\rangle$ from $|\psi\rangle$ with certainty under LOCC. Now, if $|\psi\rangle$ is an element of a universal family of resource states, then it is clear that $|\phi\rangle$ is as well. Any quantum circuit $\mathcal{C}$ that can be simulated via local unitary operations and local computational basis measurements on the state $|\phi\rangle$, together with classical feed-forward of measurement results, can just as well be simulated on $|\psi\rangle$ under the same conditions. In fact, one simply needs to execute the relevant deterministic LOCC protocol to convert $|\psi\rangle$ to $|\phi\rangle$ with unit probability, and then proceed as usual. Therefore, the orbit of $|\phi\rangle$ under deterministic LOCC, also known as the LOCC-equivalence class of $|\phi\rangle$, contains only states on which $\mathcal{C}$ can be simulated perfectly via some MBQC protocol.

Suppose $|\phi\rangle$ is a cluster state on n qubits with the appropriate geometry to simulate $\mathcal{C}$. Does that mean that the LOCC-equivalence class of $|\phi\rangle$ is precisely the class containing all of the pure n -qubit resources that can simulate $\mathcal{C}$? It would be rather disappointing for this to be the case, because it turns out that the LOCC-equivalence class is precisely equal to the LU-equivalence class, those states that can be reached from $|\phi\rangle$ by applying only local unitary operations [78, 58]. For all intents and purposes, any state that is LU-equivalent to the cluster state is identical to it; no entanglement properties of the state can be modified via LU operations. Thus, we should turn our attention to a larger class of states encompassing $|\phi\rangle$; a natural one to consider is the SLOCC-equivalence class.

Two n -qubit pure states $|\psi\rangle$ and $|\phi\rangle$ are said to be SLOCC-equivalent if and only if $|\psi\rangle$ can be converted to $|\phi\rangle$, and $|\phi\rangle$ to $|\psi\rangle$, via LOCC with a non-zero probability of success [79]. The SLOCC-equivalence class first arose in the context of entanglement classification, an attempt to sort quantum states into equivalence classes on the basis of whether they can be interconverted by restricted classes of protocols, such as LU or (deterministic) LOCC. In fact, LOCC-equivalence classes do not have a very rich structure, because of the fact that the notions of LOCC- and LU-equivalence coincide [78, 58]. There are thus an infinite number of LOCC-equivalence classes even for bipartite states. On the other hand, using SLOCC-equivalence instead results in a significant coarse-graining; there are precisely two classes for two-qubit pure states, the product states in one class and all other states in the other. In other words, it is always possible to convert any two-qubit entangled pure state to any other under LOCC with a probability exceeding 0.

The result for two qubits suggests that there is only one fundamental kind of entanglement for two-qubit states. It turns out that there are two inequivalent kinds of entanglement for three-qubit pure states. In particular, there are two SLOCC-equivalence

classes for fully-entangled three-qubit states, those that cannot be factorised into a tensor product of a single-qubit pure state and a two-qubit pure state [79]. These classes have the canonical representatives

$$\begin{aligned} |\text{GHZ}\rangle &:= \frac{1}{\sqrt{2}} (|000\rangle + |111\rangle); \\ |\text{W}\rangle &:= \frac{1}{\sqrt{3}} (|100\rangle + |010\rangle + |001\rangle). \end{aligned}$$

For four or more qubits, the number of SLOCC-equivalence classes is not enumerable [79, 80, 81, 82]. Nevertheless, the important point here is that two n -qubit states with different entanglement properties, such as for example, differing entropies of entanglement with respect to any bipartition, can still be SLOCC-equivalent. Therefore, the SLOCC-equivalence class of cluster states of various sizes seems like a promising candidate for a large class of resource states differing materially from the cluster states themselves.

Of course, any state that is SLOCC-equivalent to a cluster state can be used as a stochastic, or probabilistic resource for MBQC; indeed, it can be locally converted to a cluster state with non-zero probability. The question of interest in Ch. 2 is which class of multiqubit states on n -qubits that are SLOCC-equivalent to a cluster state of a particular geometry suitable for simulating an arbitrary quantum circuit $\mathcal{C}$ of an appropriate size, can be used to simulate $\mathcal{C}$ deterministically themselves. There are two compelling (and related) reasons to consider this class of states. The first is that many multiqubit resource states have been shown to be locally equivalent to cluster states [83, 55]. Indeed, for some of these resources, the proof of universality itself is via this method. The second is that presently (to my knowledge), no qubit resource states are known that are not SLOCC-equivalent to some cluster state. For some resources, the proof of universality was via some method such as the MPS or PEPS representation, but they were later shown to be locally equivalent [83]. This is in some sense a trivial statement if one only considers

CQ-universality, as then there must exist some state from the family that can be used to prepare any particular cluster state via a series of adaptive local measurements (which clearly constitutes a LOCC protocol). However, it is not guaranteed that the cluster state will possess the same number of physical or logical qubits as the initial resource state. Furthermore, if one considers CC-universality instead, which is arguably sufficient for simulating the circuit model via MBQC, then it is not clear at all *a priori* whether it should be possible to prepare a cluster state locally from the resource. To clarify what I mean by ‘locally equivalent’: it is the same as SLOCC-equivalent if the resource state is itself on qubits. If the candidate resource is instead on d -level qudits, then it is the same as SLOCC-equivalent to a cluster state on d -level qudits as well, where only two local degrees of freedom per particle are used to encode the information.

The definition of SLOCC-equivalence I gave earlier is not very mathematically convenient, as it does not place any restrictions on the number of Kraus operators that may be necessary to describe the protocol. Fortunately, a much more convenient condition for SLOCC-equivalence was found by Dür et al. [79], namely that two n -qubit states $|\psi\rangle$ and $|\phi\rangle$ are SLOCC-equivalent if and only if they are connected by an invertible local operator, meaning that

$$|\phi\rangle = A_1 \otimes A_2 \otimes \cdots \otimes A_n |\psi\rangle, \quad (1.63)$$

where $A_i \in \text{GL}(2, \mathbb{C})$ for each i . In Ch. 2, I consider cluster states of several useful geometries, and determine certain sufficient (though not necessary) conditions on the $\{A_i\}$ such that the result of applying the invertible local operators on the cluster state is a state on which the same useful set of circuits that could be simulated by the original cluster can also be performed with the new state. In particular, I show that qubit states having long-range correlations, rather than cluster-style nearest-neighbour ones, can be produced via the action of appropriate invertible local operators on cluster states. This

provides further evidence that, perhaps, the SLOCC-equivalence class of the cluster states (or more colloquially, states having cluster-like entanglement) are the only candidates for universal MBQC resources.

1.9 Preamble to Chapter 2

Chapter 2 is a reproduction of a journal article that has been published in Physical Review A. The article is © American Physical Society (APS), 2011. The full citation to the article can be found in the Bibliography; it is Ref. [38]. The content is identical in every respect to the manuscript, except that the formatting has been changed to comply with University of Calgary thesis formatting requirements rather than those of the journal, and the bibliography has been stripped from the manuscript, with the references incorporated into the bibliography for the thesis as a whole. According to APS guidelines, it is acceptable to include an APS-published manuscript in a thesis without receiving explicit permission in writing; see <http://publish.aps.org/copyrightFAQ.html#thesis>. A preprint of the article has been placed on the arXiv (arxiv.org), with the article identifier `arXiv:1108.4909`.

I am the first author of this article, and my contribution was to perform all of the primary research whose results appear in the article, as well as write and edit a complete draft of the paper. My co-author Dr. David L. Feder, who is my supervisor, provided scientific advice and input regarding the nature of the research to be performed and the goals of the project. Additionally, Dr. Feder was heavily involved in editing the manuscript prior to submission to the journal upon receipt of a full draft from me.

Chapter 2

Strategies for measurement-based quantum computation with cluster states transformed by stochastic local operations and classical communication

2.1 Abstract

We examine cluster states transformed by stochastic local operations and classical communication, as a resource for deterministic universal computation driven strictly by projective measurements. We identify circumstances under which such states in one dimension constitute resources for random-length single-qubit rotations, in one case quasi-deterministically ($N - U - N$ states) and in another probabilistically ($B - U - B$ states). In contrast to the cluster states, the $N - U - N$ states exhibit spin correlation functions that decay exponentially with distance, while the $B - U - B$ states can be arbitrarily locally pure. A two-dimensional square $N - U - N$ lattice is a universal resource for quasi-deterministic measurement-based quantum computation. Measurements on cubic $B - U - B$ states yield two-dimensional cluster states with bond defects, whose connectivity exceeds the percolation threshold for a critical value of the local purity.

2.2 Introduction

In the Measurement-Based Quantum Computation (MBQC) model [35, 36], one starts with a highly entangled many-qubit quantum state called a resource state, and processes

logical information via single-qubit measurements on the physical qubits of the resource state. In order to compensate for the randomness of the measurement outcomes, the bases in which measurements are performed must be conditioned on the outcomes of previous measurements. Proceeding in this way, one can teleport logical quantum information situated on one part of the state to another part, but having been subjected to some desired unitary transformation. If the basic unitary transformations that can be applied via single-qubit measurements on the resource state generate a set that is dense in $SU(2)$, then the resource is said to be universal (in the terminology of Ref. [56], this is the notion of CQ-universality, and such a resource state would be called a universal state preparator).

The archetypal family of resource states known to be universal for efficient MBQC is the so-called cluster state [35, 36]. This state is special in several ways: all spin correlation functions are strictly nearest-neighbor [54, 47], the localizable entanglement between any pair of qubits is maximal [54, 47], it is the only state (up to local unitaries) on small system sizes that saturates the Tsallis and Renyi entropies of entanglement [82], it cannot be the non-degenerate ground state of a two-body spin Hamiltonian [73, 84], and so on. One might expect that one or more of these properties would be necessarily satisfied by any universal resource state. This has turned out not to be the case; several authors in recent years have identified resources that differ materially from the cluster states [69, 54, 47, 50, 70, 71, 51, 55].

The newly discovered richness in the landscape of resources notwithstanding, the property of universality is exceedingly rare; not only must a family of universal resource states saturate various measures of entanglement in the thermodynamic limit [56, 61], but the entanglement with respect to other measures must not be too high [67, 68]. Therefore, it is highly unlikely that a random pure state will be universal, so a search for new resources must be heavily constrained in order to have a reasonable chance of success.

Recently, a number of new resources [54, 47, 50, 85] have been proven to be universal by means of reduction to a known resource state [83]. The reduction strategies of interest are those composed purely of local operations, possibly augmented by classical communication. They are typically stochastic, in the sense that the known resource state is smaller than the original state. In other words, these resources all appear to be within the equivalence class of the cluster states under Stochastic Local Operations and Classical Communication (SLOCC). The SLOCC-equivalence class of an n -qubit pure state is known to be its orbit under $\text{GL}(2, \mathbb{C})^{\otimes n}$, the group of n -fold tensor products of two-by-two invertible matrices over the complex numbers. In other words, two n -qubit states $|\psi\rangle$ and $|\phi\rangle$ are equivalent under SLOCC if and only if

$$|\psi\rangle = S^{(1)} \otimes S_2 \otimes \cdots \otimes S_n |\phi\rangle \quad (2.1)$$

where the $\{S_i\} \in \text{GL}(2, \mathbb{C})$ are invertible, two-by-two complex matrices.

A natural question thus arises: are all universal MBQC resources SLOCC-equivalent to the cluster states? More precisely, is any n -qubit element of a family of universal resource states SLOCC-equivalent to an n -qubit cluster state?

In this paper, we tackle a related question, namely: what states in the SLOCC-equivalence class of the two-dimensional cluster states are universal for MBQC? It is clear, by construction, that each state in this class can be stochastically reduced to a cluster state, but what is not clear is whether it is possible to compute directly on the image of a cluster state under some invertible, local map. We show that there is a restricted subclass of invertible local transformations, strictly including the local unitaries, whose image is a set of quasi-deterministic resources for MBQC, where in general the computation is of random length and ‘repeat-until-success’ strategies must be employed (c.f. Refs. [47, 85]). In particular, we identify two types of SLOCC operators whose action can in certain

cases preserve the usefulness of the cluster state as a resource. The first type, which we call N-type operators, comprises those operators that preserve the relative norms of the computational basis states. The second type, called B-type operators, are those that preserve their orthogonality (i.e. are in a sense basis-preserving).

In particular, we show that when N-type operators act on alternating qubits in a 1D cluster state, the state remains a quasi-deterministic resource for single-qubit rotations. We refer to such 1D states as N – U – N chains. In contrast to the cluster state, the number of measurements required to implement an arbitrary single-qubit rotation with an N – U – N chain is random rather than fixed. Furthermore, the state exhibits non-zero spin-spin correlations that decay exponentially with distance. These properties are shared by other resources previously appearing in the literature [54, 50, 76, 76], notably those based on the so-called AKLT model [86, 48]. We also show how 1D N – U – N chains can be coupled together to produce a quasi-deterministic 2D resource for universal MBQC.

Next, we show that when B-type operators act on alternating qubits in a 1D cluster state, the result is in general a probabilistic resource for single-qubit rotations. We call these states B – U – B chains. We find that under a restricted subset of B-type operators, the three-dimensional analogs of B – U – B chains constitute quasi-deterministic resources for MBQC under strictly projective measurements. A similar result was exhibited in Ref. [61], in which a 2D cluster state deformed by B-type operators was shown to be reducible to a percolated 2D cluster state [87] by the action of three-element POVMs. The B – U – B states have the interesting property that each qubit can be arbitrarily locally pure, or alternatively that an individual qubit can be arbitrarily weakly entangled with the rest of the state, as measured by the von Neumann entropy of entanglement. Like the cluster states, they also exhibit vanishing long-range correlations, with no spin-spin correlation functions beyond second-nearest-neighbor surviving.

The structure of the paper is as follows: in Section 2.3, we briefly review the theory

of measurement-based quantum computation using cluster states, and introduce the various definitions and notation used in the technical part of this paper. In Section 2.4, we describe the effects of invertible local operators acting on a cluster state on the class of linear transformations that can be logically implemented via adaptive single-qubit measurements on this new state, and outline some strategies for dealing with these effects. In Section 2.5, we provide explicit examples of some structures of SLOCC-transformed cluster states that are universal for either probabilistic or deterministic single-qubit rotations or full MBQC. Finally, in Section 2.6, we discuss the relationship of our resource states with previously known quasi-deterministic resources, and outline the prospects of identifying hitherto unknown resource states by this method.

2.3 Background

An n -qubit cluster state can be defined in terms of the stabilizer formalism [33] as the unique n -qubit pure state $|\text{Cl}_n\rangle$ satisfying the n conditions

$$X_i \bigotimes_{j \in \mathcal{N}(i)} Z_j |\text{Cl}_n\rangle = |\text{Cl}_n\rangle, \quad (2.2)$$

where $i \in \{1, \dots, n\}$ labels a qubit, $\mathcal{N}(i)$ denotes the spatial neighbourhood of qubit i , and X_i and Z_i denote the standard single-qubit Pauli operators, given in the computational basis by

$$X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}; \quad (2.3)$$

$$Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \quad (2.4)$$

acting on qubit i . Alternatively, the cluster state can be identified as the result of a dynamical process in which

1. n qubits are initialized in the state $|+\rangle^{\otimes n}$, where $|+\rangle \equiv \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ is the $+1$ -eigenstate of X ;
2. CZ entangling gates, whose action on the computational basis states is given by

$$\text{CZ}_{i,j}|x,y\rangle_{i,j} = (-1)^{x \cdot y} |x,y\rangle_{i,j}, \quad (2.5)$$

are applied between each pair of neighbouring qubits (i,j) . Thus,

$$|\text{Cl}_n\rangle = \prod_{\langle i,j \rangle} \text{CZ}_{i,j} |+\rangle^{\otimes n}, \quad (2.6)$$

where $\langle i,j \rangle$ indicates that i and j label neighbouring qubits.

For notational convenience, define a global entangling operation on a lattice,

$$\mathfrak{G}_{k,l} := \prod_{j=k}^{l-1} \text{CZ}_{j,j+1}, \quad (2.7)$$

to be the tensor product of CZ gates acting between all nearest-neighbour pairs of vertices on a line with labels between k and l . Now consider a modified one-dimensional n -qubit cluster state,

$$|\text{Cl}_n^{1D}\rangle' = \mathfrak{G}_{1,n} |\psi\rangle_1 |+\rangle_{2,\dots,n}^{\otimes n-1}. \quad (2.8)$$

where the first qubit was encoded in some general pure state $|\psi\rangle$ before the global entangling operation $\mathfrak{G}_{1,n}$ was performed. The effect of projectively measuring the first qubit, the one on which $|\psi\rangle$ was initially encoded, is to teleport the quantum information corresponding to the state $|\psi\rangle$ to the next qubit, subject to some linear transformation depending upon the basis and outcome of the measurement. To see this, assume that

$|\psi\rangle = a|0\rangle + b|1\rangle$. We then find that

$$\begin{aligned}
|Cl_n^{1D}\rangle' &= \mathfrak{G}_{2,n} CZ_{1,2} (a|0+\rangle_{1,2} + b|1+\rangle_{1,2}) |+\rangle_{3,\dots,n}^{\otimes n-2} \\
&= \mathfrak{G}_{2,n} (a|0+\rangle_{1,2} + b|1-\rangle_{1,2}) |+\rangle_{3,\dots,n}^{\otimes n-2} \\
&= \mathfrak{G}_{2,n} (a|0\rangle_1 I_2 + b|1\rangle_1 Z_2) |+\rangle_{2,\dots,n}^{\otimes n-1}.
\end{aligned} \tag{2.9}$$

Projecting the first qubit via an arbitrary rank-1 projector $|m\rangle\langle m|$, the state of the system (neglecting the projected qubit and overall normalization) becomes

$$\begin{aligned}
|\Phi\rangle &= \mathfrak{G}_{2,n} (a\langle m|0\rangle I_2 + b\langle m|1\rangle Z_2) |+\rangle_{2,\dots,n}^{\otimes n-1} \\
&= \mathfrak{G}_{2,n} (a\langle m|0\rangle |+\rangle_2 + b\langle m|1\rangle |-\rangle_2) |+\rangle_{3,\dots,n}^{\otimes n-2} \\
&= \mathfrak{G}_{2,n} H_2 (a\langle m|0\rangle |0\rangle_2 + b\langle m|1\rangle |1\rangle_2) |+\rangle_{3,\dots,n}^{\otimes n-2} \\
&\propto \mathfrak{G}_{2,n} H_2 (\langle m|+\rangle I_2 + \langle m|-\rangle Z_2) |\psi\rangle_2 |+\rangle_{3,\dots,n}^{\otimes n-2},
\end{aligned}$$

where the Hadamard operator $H_i = (X_i + Z_i)/\sqrt{2}$. In other words, the quantum information has been teleported to the second qubit through the linear transformation

$$M = H (\langle m|+\rangle I + \langle m|-\rangle Z). \tag{2.10}$$

Without loss of generality, the single-qubit state acting as the projector can be written as

$$|m(\xi, \phi)\rangle = \cos \frac{\xi}{2} |+\rangle + e^{i\phi} \sin \frac{\xi}{2} |-\rangle \tag{2.11}$$

for some $0 \leq \xi < 2\pi$, $-\pi \leq \phi < \pi$. Thus, the linear transformation through which the quantum information is teleported can be written as

$$M = H \left(\cos \frac{\xi}{2} I + e^{i\phi} \sin \frac{\xi}{2} Z \right). \quad (2.12)$$

In the special case that $\phi = \pm\pi$, corresponding to $|m\rangle$ lying on the $x - y$ plane of the Bloch sphere, this transformation becomes the (familiar from cluster state MBQC) unitary transformation $HR_z[\pm\xi]$. Thus, there is an entire single-parameter family of unitary gates through which the initial state $|\psi\rangle$ can be teleported, each corresponding to a projection of the first qubit on to some state lying in the $x - y$ plane. This family is universal for single-qubit rotations: via four projections, corresponding to $\xi = 0, \xi_2, \xi_3, \xi_4$ respectively, one teleports the transformation

$$\begin{aligned} U(\xi_2, \xi_3, \xi_4) &= HR_z[\xi_4] HR_z[\xi_3] HR_z[\xi_2] HR_z[0] \\ &= R_x[\xi_4] R_z[\xi_3] R_x[\xi_2], \end{aligned} \quad (2.13)$$

which is an arbitrary single-qubit unitary decomposed in terms of Euler angles.

To this point, we have not discussed how to compensate for the randomness associated with the measurements. If one were to drive the gate teleportation described above via projective measurements, measurements must be made in an orthonormal basis containing $|m(\xi, \phi = \pi)\rangle$. For a single qubit, this basis would be $\mathcal{B}(\xi, \phi) = \{|m(\xi, \phi)\rangle, |m^\perp(\xi, \phi)\rangle\}$, where

$$|m(\xi, \phi)\rangle = \cos \frac{\xi}{2} |+\rangle + e^{i\phi} \sin \frac{\xi}{2} |-\rangle; \quad (2.14)$$

$$|m^\perp(\xi, \phi)\rangle = \sin \frac{\xi}{2} |+\rangle - e^{i\phi} \cos \frac{\xi}{2} |-\rangle. \quad (2.15)$$

From Eq. (2.10), the teleported gate associated with the application of the projector $|m^\perp\rangle\langle m^\perp|$ on the first qubit would be

$$\begin{aligned} M^\perp &= H \left[\sin \frac{\xi}{2} I - e^{i\phi} \cos \frac{\xi}{2} Z \right] \\ &\equiv XH \left[\cos \frac{\xi}{2} I - e^{-i\phi} \sin \frac{\xi}{2} Z \right], \end{aligned} \quad (2.16)$$

where in the last step we have made use of the identity $XH = HZ$, and have dropped an unimportant overall phase. Once again considering the special case that $\phi = \pm\pi$, this reduces to

$$M^\perp = XHR_z[\pm\xi] = XM.$$

The teleported gate can be summarized succinctly as follows: measuring in the basis $\{|m\rangle, |m^\perp\rangle\}$ defined by Eqs. (2.14,2.15) with $\phi = \pi$, and denoting the measurement outcome by $m = 0$ for state $|m\rangle$ and $m = 1$ for $|m^\perp\rangle$, then the teleported gate is $X^mHR_z[\xi]$. The operator X can be thought of as a byproduct operator that occurs as a result of obtaining measurement outcome 1.

If the aim is to teleport the operator $U(\xi_2, \xi_3, \xi_4)$ defined in Eq. (2.13), then apparently one runs into a problem should a measurement outcome of 1 be obtained for any of the four measurements needed to teleport this gate. In fact the X byproduct operators can be pushed through the rotations because $R_z[\xi]X = XR_z[-\xi]$. Suppose then that one performs four projective measurements on a one-dimensional cluster state, with the i th measurement being a projective measurement of qubit i in the orthonormal basis $\mathcal{B}(\theta_i, \pi)$, $\theta_1 = 0$ and the measurement outcome denoted $m_i \in \{0, 1\}$. The quantum information originally situated on qubit 1 before the global entangling operation is then teleported

through the gate

$$\begin{aligned}
M &= X^{m_4} H R_z [\theta_4] X^{m_3} H R_z [\theta_3] X^{m_2} H R_z [\theta_2] X^{m_1} H \\
&= X^{m_4} Z^{m_3} X^{m_2} Z^{m_1} R_x [(-1)^{m_3+m_1} \theta_4] \\
&\quad \times R_z [(-1)^{m_2} \theta_3] R_x [(-1)^{m_1} \theta_2].
\end{aligned} \tag{2.17}$$

Comparing Eq. (2.13) and Eq. (2.17), the choices $\theta_2 = (-1)^{m_1} \xi_2$, $\theta_3 = (-1)^{m_2} \xi_3$, and $\theta_4 = (-1)^{m_3+m_1} \xi_4$ make the implemented teleported gate

$$M = X^{m_4} Z^{m_3} X^{m_2} Z^{m_1} U(\xi_2, \xi_3, \xi_4). \tag{2.18}$$

Thus, any gate can be implemented by conditioning each of the last three measurement bases on the results of previous measurements, up to an overall Pauli byproduct operation. The byproduct is of no concern, as Z has no effect on computational basis states while X merely swaps them; this means that the effect of the byproduct can be taken into account simply by appropriate reinterpretation of the final measurement outcomes of the circuit, contingent on the intermediate measurement outcomes.

An equivalent description of this universal gate teleportation can be obtained within the Matrix-Product State (MPS) representation [54, 47] of the one-dimensional cluster state:

$$|Cl_n\rangle = \sum_{\vec{i}} A^{[1]}[i_1] A^{[2]}[i_2] \dots A^{[n]}[i_n] |i_1 i_2 \dots i_n\rangle, \tag{2.19}$$

where $\vec{i}$ is an n -bit string and the site matrices $\{A^{[j]}[i_j]\}$ are all two-by-two, except for the boundaries; the $\{A^{[1]}[i_1]\}$ are row vectors and the $\{A^{[n]}[i_n]\}$ are column vectors. The site matrices are not unique, but it is particularly convenient if they are chosen to satisfy the relation $\sum_{i_j} A^{[j]}[i_j] A^{[j]\dagger}[i_j] = I$ for each j , corresponding to the canonical form of the MPS [41]. For the left and right boundaries one obtains $A^{[1]}[0] = \frac{1}{\sqrt{2}}\langle +|$,

$A^{[1]}[1] = \frac{1}{\sqrt{2}}\langle -|$, $A^{[n]}[0] = |0\rangle$, and $A^{[n]}[1] = |1\rangle$; for the bulk sites $1 < j < n$ they are $A^{[j]}[0] = \frac{1}{\sqrt{2}}H$ and $A^{[j]}[1] = \frac{1}{\sqrt{2}}HZ = \frac{1}{\sqrt{2}}XH$. The ‘always-on’ operator H is teleported on each measurement of a qubit, and the X gate serves as the byproduct operator. In general, an MPS state is a universal resource for measurement-based single-qubit gate teleportation (a ‘computational wire’) if the bulk site matrices can be chosen to be proportional to unitaries [85]. In this case they can be written as $A^{[j]}[0] = \frac{1}{\sqrt{2}}W$ and $A^{[j]}[1] = \frac{1}{\sqrt{2}}WR_z(\phi)$ with $W \in SU(2)$ and $\phi \in \mathbb{R}$.

In the context of the calculations presented in the next two sections, it is worth pointing out that there are two special features associated with the projective single-qubit measurements on one-dimensional cluster states presented above. The first is that the linear transformation M on the quantum state $|\psi\rangle$ is guaranteed to be unitary; in practice, this means that the effect of such a measurement is not dependent on the input state $|\psi\rangle$. If the linear transformation is not unitary (for example projections of the local system outside the $x - y$ plane, as discussed below), then there is an equivalent unitary transformation resulting in the same final state vector; however, the equivalent unitary will depend on $|\psi\rangle$. The second feature is that the byproduct operator resulting from measurement outcome 1 is always X . This is beneficial as X operators can be pushed through R_z operators with an easily characterized effect, as discussed above. These properties can be summarized as follows:

- **Property IA:** the teleported gate is in general of the form $HR_z[\xi]$ where $\xi \in \mathbb{R}$, i.e. a unitary gate corresponding to a z -axis rotation by a real angle, followed by a Hadamard gate.
- **Property IIA:** the byproduct operator is always $X \equiv R_x[\pi]$.

The above two features do not hold for single-qubit projective measurements outside the $x - y$ plane. In general,

- **Property IB:** the teleported gate is in general of the form $HR_z[\xi]$ where $\xi \in \mathbb{C}$, i.e. a non-unitary gate corresponding to a z -axis rotation by a complex angle, followed by a Hadamard gate.
- **Property IIB:** the byproduct operator is in general $R_x[\eta]$ where $\eta \neq \pi$ in general, i.e. an x -axis rotation not corresponding simply to Pauli X.

Another way to view this is that in $x-y$ plane, one always teleports $HR_z[\xi + \delta_{m,1}\pi]$ with $\xi \in \mathbb{R}$, whereas in any other plane, one teleports $HR_z[\xi + \delta_{m,1}\epsilon]$ where $\xi \in \mathbb{R}$ corresponds to the angle of the z -rotation in the desired gate, and $\epsilon \in \mathbb{C}$ is a complex error that occurs on measurement outcome 1.

Single-qubit gates alone are not sufficient for universal computation; at least one multiqubit entangling gate is required as well. In the cluster state model, multiqubit gates are accomplished via measurement patterns on 2D structures. Logical qubits are processed by horizontal 1D wires, while entangling operations are mediated by vertical links between them. An entangling gate that is locally equivalent to controlled-NOT can be achieved by measuring a link qubit in the Y basis [35, 36].

As with the case of single-qubit rotations, local Pauli byproduct operators may exist as well, depending on the measurement outcomes. As before, these byproduct operators are of no concern computationally. Thus there exists for cluster states a measurement pattern that deterministically implements a two-qubit unitary entangling gate, with any byproducts that occur being of the local Pauli type. It is not immediately clear that this will be the case for states other than the cluster state; in general, the teleported two-qubit gate may be non-unitary and the byproduct may be non-local. Any candidate resource state for MBQC must be shown to be amenable to a measurement pattern implementing some suitable two-qubit entangling gate. In Section 2.4, we describe some 1D structures that are resources for single-qubit rotations and then in the examples of Section 2.5, we

demonstrate how to perform entangling gates with natural 2D or 3D extensions of the 1D structures, and how to compensate the randomness associated the measurements.

2.4 Projective measurements on SLOCC-transformed cluster states

As discussed in the previous section, the distinguishing feature of MBQC with regular cluster states is that there exists a plane of the Bloch sphere onto which successive, adaptive, single-qubit projective measurements drive an arbitrary computation that is deterministic and of fixed length. No matter which single-qubit gate is desired, it will be implemented with certainty up to an unimportant Pauli byproduct with four measurements. For an SLOCC-transformed cluster state, it is not obvious that there exists any such plane: in general it is not possible to simultaneously satisfy both Properties IA and IIA (or for the latter any another convenient Clifford gate). A natural question to ask is then: under what circumstances can either property IA or IIA be satisfied by itself? And if only one property is satisfied, does there remain a deterministic protocol for universal quantum computation? Sec. 2.4.1 and Sec. 2.4.2 discuss the circumstances under which it is possible to independently satisfy Property IA and IIA, respectively.

2.4.1 Strategy I: Guaranteed Unitary Evolution

Derivation of N-type Operators

For convenience, define

$$\mathbf{S}_{k,l} := \bigotimes_{j=k}^l S_j^{(j)} \quad (2.20)$$

where $S^{(j)} \in \text{GL}(2, \mathbb{C})$. From Eq. (2.9), it is clear that the SLOCC-transformed cluster state encoding quantum information can be written in the form

$$\mathbf{S}_{1,n} |\text{Cl}_n\rangle' = \mathbf{S}_{2,n} \mathfrak{G}_{2,n} \left(a S_1^{(1)} |0\rangle_1 I_2 + b S_1^{(1)} |1\rangle_1 Z_2 \right) |+\rangle_{2,\dots,n}^{\otimes n-1}. \quad (2.21)$$

Following the procedure discussed in Sec. 2.3, applying the projector $|m\rangle\langle m|$ to the first qubit yields the resulting state on the remaining qubits:

$$\begin{aligned} |\Phi\rangle &= \mathbf{S}_{2,n} \mathfrak{G}_{2,n} \left(a \langle m | S_1^{(1)} | 0 \rangle \mathbf{I}_2 + b \langle m | S_1^{(1)} | 1 \rangle \mathbf{Z}_2 \right) |+\rangle_{2,\dots,n}^{\otimes n-1} \\ &= \mathbf{S}_{2,n} \mathfrak{G}_{2,n} \mathbf{H}_2 \mathbf{M}_2 |\psi\rangle_2 |+\rangle_{3,\dots,n}^{\otimes n-2}, \end{aligned}$$

where

$$\mathbf{M}_2 = \mathbf{H}_2 \left[\frac{\langle m | S^{(1)} | + \rangle}{\sqrt{2}} \mathbf{I}_2 + \frac{\langle m | S^{(1)} | - \rangle}{\sqrt{2}} \mathbf{Z}_2 \right]. \quad (2.22)$$

The only way for this to correspond to a unitary gate is if

$$\begin{aligned} \frac{1}{\sqrt{2}} \langle m | S^{(1)} | + \rangle &= e^{i\alpha} \cos \frac{\xi}{2}; \\ \frac{1}{\sqrt{2}} \langle m | S^{(1)} | - \rangle &= -ie^{i\alpha} \sin \frac{\xi}{2}, \end{aligned}$$

where $0 \leq \xi < 2\pi$, and therefore

$$S^{(1)\dagger} |m\rangle = \sqrt{2} e^{-i\alpha} \left[\cos \frac{\xi}{2} |+\rangle + i \sin \frac{\xi}{2} |-\rangle \right], \quad (2.23)$$

or equivalently

$$\begin{aligned} |m\rangle &= \sqrt{2} (S^{(1)\dagger})^{-1} e^{-i\alpha} \left[\cos \frac{\xi}{2} |+\rangle + i \sin \frac{\xi}{2} |-\rangle \right] \\ &= e^{-i\alpha} (S^{(1)\dagger})^{-1} \mathbf{R}_z[\xi] |+\rangle. \end{aligned} \quad (2.24)$$

Eq. (2.24) is the condition on the state $|m\rangle$ such that measurement outcome 0 yields a unitary teleported gate. Note that there is a family of states characterized by a single parameter ξ fulfilling this condition, not including the unimportant overall phase α .

To ensure that the measurement yields a unitary teleported gate independent of the

measurement outcome, a similar condition must follow for the orthogonal complement $|m^\perp\rangle$. Orthogonality requires

$$\langle m^\perp | \propto \langle - | R_z[-\xi] S^{(1)\dagger}, \quad (2.25)$$

and therefore

$$|m^\perp\rangle = c S^{(1)} R_z[\xi] |-\rangle \quad (2.26)$$

for some constant $c \in \mathbb{C}$. Repeating the procedure that led to Eq. (2.23), but with $|m^\perp\rangle$ instead of $|m\rangle$, one obtains

$$S^{(1)\dagger} |m^\perp\rangle = \sqrt{2} e^{-i\beta} \left[\cos \frac{\xi}{2} |+\rangle + i \sin \frac{\xi}{2} |-\rangle \right]. \quad (2.27)$$

Substituting Eq. (2.26) into Eq. (2.27) yields

$$\begin{aligned} \sqrt{2} e^{-i\beta} \left[\cos \frac{\xi}{2} |+\rangle + i \sin \frac{\xi}{2} |-\rangle \right] &= c S^{(1)\dagger} S^{(1)} R_z[\xi] |-\rangle \\ &= c S^{(1)\dagger} S^{(1)} \left[\cos \left(\frac{\xi}{2} \right) |-\rangle - i \sin \left(\frac{\xi}{2} \right) |+\rangle \right]. \end{aligned} \quad (2.28)$$

Rewriting Eq. (2.28) in the computational basis results in the expression

$$|0\rangle + e^{-i\xi} |1\rangle = c' S^{(1)\dagger} S^{(1)} (|0\rangle - e^{i\xi} |1\rangle) \quad (2.29)$$

for a suitably defined constant c' . Defining $T_k := S^{(k)\dagger} S^{(k)}$ and $T_k^{i,j} := \langle i | T_k | j \rangle$, we can see from Eq. (2.29) that

$$c' (T_1^{0,0} - e^{i\xi} T_1^{0,1}) = 1; \quad (2.30)$$

$$c' (T_1^{1,0} - e^{i\xi} T_1^{1,1}) = e^{-i\xi}. \quad (2.31)$$

From here, it is easily deduced that

$$|T_1^{0,0} - e^{i\xi} T_1^{0,1}| = |T_1^{1,1} - e^{-i\xi} T_1^{1,0}|. \quad (2.32)$$

Making use of the Hermiticity of T_k , simple algebra yields

$$T_1^{0,0} = T_1^{1,1}. \quad (2.33)$$

Eq. (2.33) above has a very simple geometric interpretation: it means that $S^{(1)}$ must preserve the relative norm of the computational basis states. There is no requirement for $S^{(1)}$ to preserve their orthogonality, however, which means that $S^{(1)}$ is allowed to differ quite drastically from a unitary transformation; in fact, it can be made arbitrarily close to singular, as the relative angle between the computational basis states under the transformation by $S^{(1)}$ can be vanishingly small. We will refer to operators obeying this norm-preservation restriction as N-type operators.

Definition 2.4.1. A $\text{GL}(2, \mathbb{C})$ operator S satisfying $\langle 0|S^\dagger S|0\rangle = \langle 1|S^\dagger S|1\rangle$ is called an N-type operator.

The singular value decomposition is helpful for characterizing N-type operators. An arbitrary SLOCC operator S can be written in terms of its singular value decomposition as $S = UDV$, where U is an arbitrary two-qubit unitary, D is a positive-definite diagonal matrix

$$D = \kappa \begin{bmatrix} \cos \theta & 0 \\ 0 & \sin \theta \end{bmatrix}, \quad (2.34)$$

where $0 < \theta < \frac{\pi}{2}$, and V is an arbitrary unitary matrix parametrized via the Euler decomposition as $V = H R_z[\alpha] R_x[\beta] R_z[\gamma]$, $0 \leq \alpha, \beta, \gamma < 2\pi$, and any global phase has

been absorbed into U . It is then straightforward to determine that

$$\langle 0|S^\dagger S|0\rangle = \frac{1}{2}\kappa^2 (1 + \cos 2\theta \sin \alpha \sin \beta); \quad (2.35)$$

$$\langle 1|S^\dagger S|1\rangle = \frac{1}{2}\kappa^2 (1 - \cos 2\theta \sin \alpha \sin \beta). \quad (2.36)$$

So, if S is an N-type operator, we must have $\theta = \frac{\pi}{4}$ (in which case S is proportional to a unitary), $\alpha = 0$ or $\beta = 0$. The case where $\beta = 0$ still allows us to assume $\alpha = 0$ without loss of generality. Doing so, the $R_x[\beta]$ operator can be commuted past the H to turn into a z -rotation, and then absorbed into U . Thus we have the following characterization of N-type operators.

Lemma 2.4.2. *Every N-type operator S must either be proportional to a unitary operator, or of the form $S = UDV$, where U is an arbitrary two-by-two unitary operator, D is defined as in Eq. (2.34) and $V = HR_z[\gamma]$ with $0 \leq \gamma < 2\pi$.*

It is straightforward to obtain an expression for the byproduct angle μ' in the case of measurement outcome 1 in terms of the parameters θ , γ and ξ (recall that this is the degree of freedom in the measurement basis). Using Eq. (2.22), it can easily be checked that the teleported gate when $|m\rangle \propto (S^\dagger)^{-1} R_z[-\xi] H|0\rangle$ is $M = HR_z[\xi]$, and when $|m\rangle \propto SR_z[-\xi] H|1\rangle$ is $M^\perp = R_x[\mu'] HR_z[\xi]$, where the byproduct angle μ' obeys

$$\tan \frac{\mu'}{2} = \frac{1 - \cos 2\theta \cos(\gamma - \xi)}{\cos 2\theta \sin(\gamma - \xi)}. \quad (2.37)$$

The probabilities of the two measurement outcomes can also be easily calculated in terms of the same parameters, and are found to be

$$p(0) = \frac{1}{2} (1 + \cos 2\theta \cos 2\xi); \quad (2.38)$$

$$p(1) = \frac{1}{2} (1 - \cos 2\theta \cos 2\xi). \quad (2.39)$$

This differs from the case of gate teleportation with a perfect cluster state, where the two measurement probabilities are always exactly $\frac{1}{2}$. That said, the expected probability of obtaining a byproduct here, averaged over all ξ , is

$$\langle p(1) \rangle_\xi = \frac{1}{2}, \quad (2.40)$$

irrespective of θ . We can thus generically expect to obtain an unwanted byproduct operator that we must compensate on half of our single-qubit measurements. This point will be discussed in Sec. 2.5.

Properties of N-Transformed Cluster States

Cluster states locally transformed by N-type operators can exhibit remarkably different properties from perfect cluster states. Nevertheless, as will be shown later in Sec. 2.5.2, they can under some circumstances serve as universal resources for MBQC of random length.

Consider for the moment the Schmidt decomposition of an n -qubit cluster state on some set of qubits $\mathcal{V}$ with respect to a bipartition separating qubit k from the rest:

$$|\text{Cl}_n\rangle_{\mathcal{V}} = \frac{1}{\sqrt{2}} \left(|0\rangle_k |\text{Cl}_{n-1}\rangle_{\mathcal{V}\setminus k} + |1\rangle_k Z_{\mathcal{N}(k)} |\text{Cl}_{n-1}\rangle_{\mathcal{V}\setminus k} \right), \quad (2.41)$$

where $|\text{Cl}_{n-1}\rangle_{\mathcal{V}\setminus k}$ refers to the cluster state resulting from deleting qubit k and $Z_{\mathcal{N}(k)}$ is the tensor product of Z operators acting on all the neighbors of k . This can be checked by verifying that this state satisfies the stabilizer conditions (2.2). The Schmidt basis for the multiqubit component can be further decomposed if desired by the same technique. The equality of the Schmidt coefficients in Eq. (2.41) demonstrates that any individual qubit in a cluster state has a maximally mixed local reduced density matrix, or in other words that it is maximally entangled with the rest of the cluster with respect to the von

Neumann entanglement entropy. Likewise, exactly one ebit of entanglement is shared across any bipartition of the cluster state.

One effect of N-type operators is to change the local reduced density matrices of individual qubits within the state. In the canonical representation, the site matrices of the MPS representation [cf. Eq. (2.19)] are $A^{[j]}[0] = \frac{1}{\sqrt{2}}W = HR_z [-(\gamma^{(j)} + 2\theta^{(j)})]$ and $A^{[j]}[1] = \frac{1}{\sqrt{2}}WR_z [4\theta^{(j)}]$, which are both unitary. This immediately implies that the channel having these matrices as Kraus operators is unital, and like the ordinary cluster state one ebit of entanglement is shared across any bipartition. That said, the entanglement between any given qubit and the rest of the system need not be unity.

Consider for example the local reduced density matrix of a qubit adjacent to an endpoint of a 1D cluster state with n qubits, numbered 1 to n from left to right. Singling out first qubit 2 and then qubit 3, this state can be written as

$$|Cl_n^{1D}\rangle_{1\dots n} = \frac{1}{\sqrt{2}} (|0\rangle_2|+\rangle_1 I_3 + |1\rangle_2|-\rangle_1 Z_3) |Cl_{n-2}^{1D}\rangle_{3\dots n}. \quad (2.42)$$

Now consider the action of an N-type operator, $N^{(2)} = DHR_z[\gamma]$ on qubit 2, where the leading U operator is dropped because it can be absorbed into the measurement basis. It is easy to check that

$$N^{(2)}|Cl_n^{1D}\rangle_{1\dots n} = \frac{1}{\sqrt{2}} (\cos\theta|0\rangle_2|\Phi\rangle - i\sin\theta|1\rangle_2|\Phi^\perp\rangle), \quad (2.43)$$

where $|\Phi\rangle$ and $|\Phi^\perp\rangle$ are γ -dependent states for qubits 1, 3, $\dots$, n such that $\langle\Phi|\Phi^\perp\rangle = 0$. Thus, Eq. (2.43) remains a Schmidt decomposition. The local reduced density matrix of qubit 2 is

$$\rho_2 = \begin{bmatrix} \cos^2\theta & 0 \\ 0 & \sin^2\theta \end{bmatrix},$$

revealing that qubit 2 is no longer maximally entangled with the rest of the state. A similar calculation can be performed for qubits further from the boundary, with qualitatively similar results.

Another property of these states is the long-range behavior of two-point correlation functions, those of the form $C_{i,j}(\mathcal{A}, \mathcal{B}) := \langle \mathcal{A}_i \mathcal{B}_j \rangle - \langle \mathcal{A}_i \rangle \langle \mathcal{B}_j \rangle$ for some operators $\mathcal{A}$ and $\mathcal{B}$. Two-point correlation functions of large 1D cluster states with periodic boundary conditions can be efficiently calculated using the Matrix Product State (MPS) representation. For an n -qubit ring, calculation of the correlation functions amounts to taking traces of products of n 4×4 -dimensional matrices. Consider therefore a 1D cluster state with periodic boundary conditions (i.e. a ring), with the operation $N^{(i)}$ acting on qubit i . For this state, calculations show that all two-point Pauli correlation functions vanish except for the second-nearest-neighbour correlation function $C_{i-1,i+1}(Z, Z) = \cos 2\theta \sin \gamma$. This is in contrast to the perfect cluster state with periodic boundary conditions, for which all two-point correlation functions identically vanish.

As another example, the relevance of which will become clear in Sec. 2.5.1, consider a ring with an even number of qubits, with N acting on every alternate qubit; say, the ones with even labels. In this case, the magnitude of the same two-point correlation function between odd-numbered qubits decays exponentially:

$$|C_{1,2j+1}(Z, Z)| \sim \exp\left(-\frac{2j}{L}\right). \quad (2.44)$$

The length scale L depends on θ and γ . The same correlation function between pairs of qubits with at least one even label is zero. The numerically obtained behavior of L for a ring of 1000 qubits is shown in Fig. 2.1 as a function of γ for several values of θ between 0 and $\pi/4$. As can be seen from the figure, the length scale increases with decreasing θ over this range, i.e. as the N -type operators approach the singular limit $\theta = 0$. The

length scale is symmetric about $\theta = \pi/4$ between 0 and $\pi/2$. Viewed as a function of γ with θ held constant, the correlation function is convex and non-negative in γ over the interval from 0 to $\pi/2$ and is symmetric about $\pi/4$. For $\pi/2 \leq \gamma < \pi$, the magnitudes behave the same way as in the previous interval, but the signs alternate. The γ -behavior is periodic with period π . We note in passing that these non-zero correlation functions provides a lower bound for the localizable entanglement [88] between that pair of qubits in the state via projective measurements, with respect to the concurrence [89].

Note that a number of resources for MBQC with non-vanishing long-range correlation functions have been pointed out in the literature [47, 50, 51, 76, 55], based on the so-called spin-1 AKLT model [86, 48]. These states are quasi-deterministic resources, in the sense that measurement-based computations using these states can be made arbitrarily likely to succeed, either by reduction of the resource state to a deterministic resource or by a repeat-until-success strategy with each elementary gate requiring a random number of measurements. In Secs. 2.5.1 and 2.5.2, we describe resource states called N – U – N states that are based on cluster states transformed by N-type operators; these states share the properties of quasi-determinism and non-vanishing long-range correlations.

2.4.2 Strategy II: Guaranteed Pauli Byproduct

Derivation of B-Type Operators

Another possible strategy is to attempt to ensure that the byproduct operator is guaranteed to be Pauli-X, whether or not the teleported linear transformation is unitary. The advantage of this approach is that X has nice commutation properties through rotation operators about the z -axis, whether they be by real or complex angles, leading to the hope that the randomness inherent in the measurement process can be easily compensated.

When projectively measuring in the orthonormal basis $\{|m\rangle, |m^\perp\rangle\}$, the two possible




Figure 2.1: (Color online) Correlation length scale L associated with the correlation function $|C_{1,2j+1}(Z, Z)| \sim \exp\left(-\frac{2j}{L}\right)$ for a $N - U - N$ ring with all N -type operators identical, on a ring of 1000 qubits, as a function of the parameters of γ and θ , the parameters of N . The length scale increases as N approaches the singular limit, i.e. as θ gets close to 0 or $\pi/2$.

operations that can occur are

$$\begin{aligned} M &= H \left[\frac{\langle m|S^{(1)}|+\rangle}{\sqrt{2}} I + \frac{\langle m|S^{(1)}|-\rangle}{\sqrt{2}} Z \right]; \\ M^\perp &= H \left[\frac{\langle m^\perp|S^{(1)}|+\rangle}{\sqrt{2}} I + \frac{\langle m^\perp|S^{(1)}|-\rangle}{\sqrt{2}} Z \right] \\ &= XH \left[\frac{\langle m^\perp|S^{(1)}|-\rangle}{\sqrt{2}} I + \frac{\langle m^\perp|S^{(1)}|+\rangle}{\sqrt{2}} Z \right]. \end{aligned} \quad (2.45)$$

Since I and Z are linearly independent, it follows that for the byproduct to be guaranteed to be proportional to Pauli- X , one must have

$$\langle m|S|+\rangle = c\langle m^\perp|S|-\rangle; \quad (2.46)$$

$$\langle m|S|-\rangle = c\langle m^\perp|S|+\rangle, \quad (2.47)$$

or equivalently,

$$\langle m|S|0\rangle = c\langle m^\perp|S|0\rangle; \quad (2.48)$$

$$\langle m|S|1\rangle = -c\langle m^\perp|S|1\rangle, \quad (2.49)$$

for some non-zero constant $c \in \mathbb{C}$. Suppose $S = UDV$ where U is an arbitrary single-qubit unitary, D is defined in Eq. (2.34), and $V = R_z[\beta] R_x[\gamma] R_z[\delta]$. Further suppose that $|m\rangle = UU'|0\rangle$ and $|m^\perp\rangle = UU'|1\rangle$, with $U' = R_z[\beta'] R_x[\gamma'] R_z[\delta']$. The reason for the appearance of U in the definitions of $|m\rangle$ and $|m^\perp\rangle$ is to compensate for the appearance of U in the singular value decomposition of S . The only effect of the $R_z[\delta']$ operation is to multiply the teleported gate by a global phase, so we can choose $\delta' = 0$ in what follows without loss of generality (it remains a free parameter for the applied

unitary U'). Having done so, Eqs. (2.48-2.49) can be rewritten as

$$\langle 0|Q|0\rangle = c\langle 1|Q|0\rangle; \quad (2.50)$$

$$\langle 0|Q|1\rangle = -c\langle 1|Q|1\rangle, \quad (2.51)$$

where we have defined

$$\begin{aligned} Q &:= (U')^\dagger DV \quad (2.52) \\ &= R_x[-\gamma'] R_z[-\beta'] D R_z[\beta] R_x[\gamma] R_z[\delta] \\ &= R_x[-\gamma'] R_z[\beta - \beta'] D R_x[\gamma] R_z[\delta] \\ &:= R_x[-\gamma'] R_z[b] D R_x[\gamma] R_z[\delta]; \end{aligned} \quad (2.53)$$

in the last line above we have defined $b := \beta - \beta'$. In the expression above, γ' and b are free parameters, while D , γ and δ are determined by the SLOCC operator S .

Return now to the constraints, Eqs. (2.50-2.51). Denoting $Q_{ij} := \langle i|Q|j\rangle$, one finds that $Q_{00}/Q_{10} = -Q_{01}/Q_{11}$. Note that neither Q_{10} nor Q_{11} can be zero; if either were zero, then the constraints would force Q and therefore S to be singular, which by assumption is not the case. This in turn means that

$$\text{Det}(Q) = 2Q_{00}Q_{11}. \quad (2.54)$$

From the definition of Q , Eq. (2.53),

$$\begin{aligned} \text{Det}(Q) &= \sin 2\theta; \\ 2Q_{00}Q_{11} &= \sin \gamma' \sin \gamma (\cos \gamma - i \cos 2\theta \sin b) \\ &\quad + \sin 2\theta (1 + \cos \gamma \cos \gamma'). \end{aligned} \quad (2.55)$$

Substituting the above expressions into Eq. (2.54) and equating real and imaginary parts gives us the two conditions

$$\sin \gamma' \sin \gamma \sin b \cos 2\theta = 0; \quad (2.56)$$

$$\cos \gamma (\sin \gamma' \sin \gamma + \cos \gamma' \sin 2\theta) = 0. \quad (2.57)$$

Recall that since S is invertible, we cannot have $\sin 2\theta = 0$ and since S is non-unitary, we cannot have $\cos 2\theta = 0$. The only ways to satisfy Eq. (2.56) are if $\sin \gamma' \sin \gamma = 0$ or $\sin b = 0$. In the first case, Eq. (2.57) immediately implies that $\cos \gamma' \cos \gamma = 0$, leaving b as a free parameter for our measurement basis. In the second case, γ' is fixed in terms of θ and γ , leaving no freedom in the measurement basis we are using. Furthermore, if we choose $\sin \gamma' = 0$, i.e. $\gamma' \in \{0, \pi\}$, then the measurement basis we are using is restricted to being the computational basis acted on by U (completely specified by S); again, no freedom. Therefore, the only solutions available to us that leave freedom in the measurement basis, and thus the teleported gate, are $\gamma \in \{0, \pi\}$ and $\gamma' \in \{\frac{\pi}{2}, \frac{3\pi}{2}\}$. Note that

$$S^\dagger S = \begin{bmatrix} 1 + \cos \gamma \cos 2\theta & -ie^{i\delta} \sin \gamma \cos 2\theta \\ ie^{-i\delta} \sin \gamma \cos 2\theta & 1 - \cos \gamma \cos 2\theta \end{bmatrix}. \quad (2.58)$$

Thus, demanding that the SLOCC operators allow for a guaranteed Pauli by-product, assuming the SLOCC operator is not unitary and thus $\cos 2\theta \neq 0$, is equivalent to demanding that $S^\dagger S$ be diagonal. Geometrically, this means that S must preserve the overlap of the computational basis states (the transformed computational basis is still orthogonal). We will refer to this kind of basis-preserving operator as a B-type operator.

Definition 2.4.3. A $\text{GL}(2, \mathbb{C})$ operator S satisfying $\langle 0|S^\dagger S|1\rangle = \langle 1|S^\dagger S|0\rangle = 0$ is called a B-type operator.

A B-type operator can therefore be written

$$B = \begin{cases} UDR_z[\beta]R_z[\delta], & \gamma = 0 \\ UDR_z[\beta]XR_z[\delta], & \gamma = \pi, \end{cases}$$

ignoring overall phases. The two possibilities above can be simplified and collapsed into one. First, note that $R_z[\beta]$ can be commuted past D and absorbed into U . Next, note that XD is itself a diagonal matrix that results from swapping the diagonal entries of D . This means that the case where $\gamma = \pi$ can be written instead as $U'D'R_z[\delta]$, where $U' = UR_z[\beta]X$ and $D' = XD$. Of course, $R_z[\delta]$ can also be absorbed into U' ; thus, a simple and completely general expression for a B-type operator is

$$B = UD. \tag{2.59}$$

The diagonal matrix D in the singular value decomposition can be expressed as

$$D \propto \text{diag}(\cos(\theta), \sin(\theta)) = \sqrt{\sin(\theta)\cos(\theta)}R_z[i\ln\cot(\theta)],$$

so that the B-type operator becomes

$$B \propto UR_z[i\ln\cot(\theta)]. \tag{2.60}$$

Because the unitary U can be absorbed directly into the measurement basis, one can interpret B-type operators as z -rotations by an imaginary angle, the value of which is related to the ratio of the singular values.

When the local operator is B-type, the single-parameter family of measurement bases satisfies $\gamma' \in \{\frac{\pi}{2}, \frac{3\pi}{2}\}$, and $\beta' \in [0, 2\pi)$ is a free parameter. When this family of bases is

used, the byproduct operator associated with measurement outcome 1 is always Z (up to a global phase). The teleported linear transformation is no longer unitary, however; it takes the form of a rotation about the z -axis of the Bloch sphere by a complex angle, followed by a Hadamard operation. The real part of the angle is completely specified by the choice of measurement basis, via the free parameter β' . The imaginary part is purely a function of the ratio of the singular values of the local $GL(2, \mathbb{C})$ operator. Denoting the measurement outcome corresponding to $\gamma' = \frac{\pi}{2}$ by $m = 0$ and that for $\gamma' = \frac{3\pi}{2}$ by $m = 1$, the teleported gate is given (up to a global phase) by

$$M = X^m H R_z [\beta' + i \ln \cot \theta]. \quad (2.61)$$

Properties of B-Transformed Cluster States

Interpreting the B-type operators as z -rotations by imaginary angles provides a simple insight into the nature of B-transformed cluster states. The R_z operator commutes with all CZ gates, so one can push it all the way through to the $|+\rangle$ states in the definition of the cluster state, Eq. (2.6). Because $R_z(\xi)|+\rangle$ is an arbitrary single-qubit state, B-transformed cluster states are equivalent to applying CZ gates between qubits in arbitrary states (not including computational basis states, which would require singular B operators).

One might assume that B-transformed cluster states are equivalent to weighted cluster states [90, 91, 92, 37, 93], but this is not in fact the case. Weighted graph states are defined as $\prod_{\langle i,j \rangle} CP(\varphi)_{i,j} |+\rangle^{\otimes n}$, where the controlled-phase entangling gate is $CP(\varphi) = \text{diag}(1, 1, 1, e^{i\varphi})$; the cluster-state edge weights are then given by $w_{ij} = \varphi_{ij}$. Consider the simplest counter-example of a three-qubit linear cluster state with the central qubit transformed by a B-type operator $B = D R_z[\gamma]$ with $D = \text{diag}(\cos \theta, \sin \theta)$. The eigenvalues of the local reduced density matrices are all $\{\frac{1}{2}(1 \pm \cos 2\theta)\}$. On the other hand, for a

three-qubit 1D weighted graph state with edge weights φ_{12} and φ_{23} , the eigenvalues of the reduced density matrix are $\frac{1}{2} (1 \pm \cos \frac{\varphi_{12}}{2})$, $\frac{1}{2} (1 \pm \frac{1}{2} \cos \varphi_{12} \cos \varphi_{23})$, and $\frac{1}{2} (1 \pm \cos \frac{\varphi_{23}}{2})$ for qubits 1 through 3, respectively. If the weighted graph and the B-transformed cluster are LU-equivalent, there must be some choice of φ_{12} and φ_{23} such that the spectra of the reduced density matrices are the same in both cases. For qubits 1 and 3 this implies $\phi_{12} = \phi_{23} = 4\theta$. For qubit 2 one obtains $\frac{1}{2} (1 \pm \frac{1}{2} \cos 4\theta^2)$. This matches the corresponding spectrum for the B-transformed cluster only when $\theta = \pm \frac{\pi}{4}, \phi = \pm \pi$, in which case both states are LU-equivalent to a perfect cluster.

Cluster states locally transformed by B-type operators also exhibit different properties from perfect cluster states. As with N-type operators, B-type operators change the local reduced density matrices of individual qubits within the state, as described in the following lemma.

Lemma 2.4.4. *Let $|\text{Cl}_n\rangle_{\mathcal{V}}$ be an n -qubit cluster state on the set of qubits $\mathcal{V}$, with some subset $\mathcal{Q} \subseteq \mathcal{V}$ acted upon by B-type operators. In particular, suppose that for each qubit $i \in \mathcal{Q}$, the B-type operator acting is given by $B^{(i)} = D^{(i)}$ with $D^{(i)} = \sqrt{2} \text{diag} (\cos \theta^{(i)}, \sin \theta^{(i)})$. Then, the local reduced density matrix for any qubit $k \in \mathcal{V}$ is given by*

$$\begin{aligned} \rho_k = & \cos^2 \theta^{(k)} |0\rangle\langle 0| + \sin^2 \theta^{(k)} |1\rangle\langle 1| \\ & + \left(\frac{1}{2} \sin 2\theta^{(k)} \prod_{j \in \mathcal{N}(k)} \cos 2\theta^{(j)} |0\rangle\langle 1| + \text{h.c.} \right), \end{aligned}$$

where $\theta^{(k)} := \frac{\pi}{4}$ if $k \notin \mathcal{Q}$.

The lemma is easily proved by taking advantage of the expression (2.41) for the Schmidt decomposition of a cluster state with one subsystem being qubit k alone, and then calculating ρ_k directly. The calculation is done by expressing the cluster state as the action of controlled-Z gates acting on the product state $|+\rangle^{\otimes n}$, and then using the

fact that the $D^{(i)}$ and controlled-Z gates are mutually commuting. A consequence of this lemma is that the reduced density matrix of a given qubit is maximally mixed if and only if the qubit itself and at least one of its neighbors are untouched by B-type operators. In general, qubits within B-transformed cluster states are not maximally entangled with the rest of the state; in fact, they can be arbitrarily weakly entangled (with respect to the von Neumann entanglement entropy).

Consider now a ring of an even number of qubits, with identical B-type operators specified by $\theta^{(2k)} = \theta$ acting on the qubits with even labels. The significance of such a state will become clear in Example 2.5.4. Two-point correlation functions can be calculated exactly for this state. The result is that the nearest-neighbor correlation functions $C_{2k,2k\pm 1}(Z, X) = \cos 2\theta$ and the next-nearest-neighbor correlation functions $C_{2k,2k\pm 2}(Z, Z) = C_{2k-1,2k+1}(X, X) = \cos^2(2\theta)$ are the only ones that are non-zero, while all the other two-point Pauli correlation functions are identically zero. Again, this differs from the perfect cluster state, where all correlation functions are zero.

Exact calculations on chains of up to 7 qubits with identical (but arbitrary) B-type operators $B_{2j} = DR_z[\gamma^{(2j)}]$ acting on even qubits $2j$ reveal that the non-zero Schmidt coefficients corresponding to any bipartition of the chain into two contiguous halves are $\{\cos \theta, \sin \theta\}$. The von Neumann entropy of entanglement is equal to the Shannon entropy of this list, and is generally less than one ebit. The MPS representation bears out this observation. In the canonical form the site matrices for the boundary qubits are

$$A^{[1]}[0] = \frac{1}{\sqrt{2}}(\cos \theta^{(2)} \langle 0| + \sin \theta^{(2)} \langle 1|); \quad (2.62)$$

$$A^{[1]}[1] = \frac{1}{\sqrt{2}}(\cos \theta^{(2)} \langle 0| - \sin \theta^{(2)} \langle 1|); \quad (2.63)$$

$$A^{[n]}[0] = |+\rangle; \quad A^{[n]}[1] = -|-\rangle, \quad (2.64)$$

those for the bulk even sites are

$$A^{[2j]}[0] = \begin{bmatrix} 0 & e^{-i\gamma^{(2j)}} \\ 0 & 0 \end{bmatrix}; \quad (2.65)$$

$$A^{[2j]}[1] = \begin{bmatrix} 0 & 0 \\ e^{i\gamma^{(2j)}} & 0 \end{bmatrix}, \quad (2.66)$$

and those for the bulk odd sites are

$$A^{[2j-1]}[0] = \begin{bmatrix} \cos \theta^{(2j)} & \sin \theta^{(2j)} \\ \cos \theta^{(2j)} & \sin \theta^{(2j)} \end{bmatrix}; \quad (2.67)$$

$$A^{[2j-1]}[1] = \begin{bmatrix} -\cos \theta^{(2j)} & \sin \theta^{(2j)} \\ \cos \theta^{(2j)} & -\sin \theta^{(2j)} \end{bmatrix}. \quad (2.68)$$

It's very easy to verify that the channels induced by the matrices on the odd sites are not unital in the sense given in Ref. [85], so a B – U – B chain is not a quantum wire.

Such a 1D state would appear not to be capable of reliably processing a single qubit. This is true, but a simple modification of the geometry from one to two dimensions yields a usable resource for random length computation. This will be elaborated upon in Example 2.5.4.

2.5 Random Length Computation

Neither Strategy I nor Strategy II discussed in the previous section directly offers a way to perform deterministic single-qubit rotations. For Strategy I, it is unclear how to compensate for a byproduct operator $R_x[\eta]$ where $\eta \neq \pi$, as such a byproduct operator does not possess convenient commutation properties with the H and R_z operations. Similarly, for Strategy II, it is unclear whether some number of non-unitary teleported gates can be combined to form a desired unitary.

Another perspective on the strategies is that a single measurement with outcome 1 teleports the gate $HR_z[\xi + \epsilon]$, where $\xi \in \mathbb{C}$ is some angle associated with the always-on operation $HR_z[\xi]$ (in the terminology of Ref. [85]) and $\epsilon \in \mathbb{C}$ is a possibly complex



Figure 2.2: $N - U - N$ state, a one-dimensional structure that can be used for deterministic random-length single-qubit rotations.

error associated with the byproduct. To correct this error in principle requires two additional measurement steps. The first measurement step should teleport the gate $HR_z[0] \equiv H$, which would cancel the previously applied Hadamard gate; a possible X byproduct operator might result depending on the measurement outcome. On the second measurement step one would attempt to teleport $HR_z[-\epsilon]$ or $HR_z[\epsilon]$ depending on the previous measurement outcome, thus cancelling the original error ϵ .

This procedure is only possible if the measurement immediately after first incurring an error cannot itself generate any further error ϵ' . One way to guarantee such a circumstance is to impose that every alternate S_i operator is in fact unitary. Thus there must exist a class of states that are a strict subset of SLOCC-transformed cluster states, which constitute resources for random-length universal gate teleportation. Likewise, a subset of SLOCC-transformed cluster states in two dimensions must be universal resources for MBQC. The remainder of this section is devoted to various explicit examples.

2.5.1 Deterministic single-qubit rotations: $N - U - N$ state

Consider a one-dimensional state of the form

$$|R\rangle = N_1^{(1)} \otimes U_2^{(2)} \otimes N_3^{(3)} \otimes U_4^{(4)} \cdots \otimes N_n^{(n)} |Cl_n\rangle, \quad (2.69)$$

where the $\{N^{(i)}\}$ are N -type operators, and the $\{U^{(j)}\}$ are local unitaries (c.f. Fig. 2.2).

The goal is to teleport the single-qubit unitary

$$U(\zeta, \eta, \xi) = R_x[\zeta] R_z[\eta] R_x[\xi]. \quad (2.70)$$

The first step is to use Strategy I to attempt a teleportation of $\text{HR}_z[0]$, analogously to the scheme with the perfect 1D cluster state. For the correct choice of basis the measurement outcome $m_1 = 0$ corresponds to success. One can then immediately measure qubit 2 in a basis that teleports $X^{m_2}\text{HR}_z[\xi]$, and then use Strategy I to attempt the teleportation of $\text{HR}_z[(-1)^{m_2}\eta]$ starting on qubit 3.

If the first measurement outcome is instead $m_1 = 1$ then one instead teleports $\text{HR}_z[\epsilon]$ with $\epsilon \in \mathbb{R}$. This error must be immediately corrected, because the next desired rotation is around an orthogonal axis. Happily, there is a local unitary $U^{(2)}$ acting on the next qubit in the chain. The Hadamard operator that effects the now-undesired transformation of the rotation axes can be eliminated by teleporting another one ($H^2 = I$). This is accomplished by measuring the next qubit in the basis $\{U^{(2)}|+\rangle, U^{(2)}|-\rangle\}$. Labelling the measurement outcome m_2 , the teleported gates are

$$X^{m_2}H\text{HR}_z[\epsilon] \equiv X^{m_2}\text{R}_z[\epsilon]. \quad (2.71)$$

The measurement basis for qubit 3 is then chosen such that measurement outcome $m_3 = 0$ results in the gate $\text{HR}_z[(-1)^{m_2+1}\epsilon]$ being teleported. In this case, the overall unitary becomes

$$\begin{aligned} \text{HR}_z[(-1)^{m_2+1}\epsilon] X^{m_2}\text{R}_z[\epsilon] &= Z^{m_2}\text{HR}_z[-\epsilon]\text{R}_z[\epsilon] \\ &= Z^{m_2}H. \end{aligned}$$

At this point one has successfully teleported a Hadamard gate and an unimportant Pauli byproduct. The next measurement on a qubit with an even label can teleport the desired $\text{HR}_z[\xi]$ gate without error. One then attempts to teleport $\text{HR}_z[\eta]$ by measuring qubit 5, using Strategy I, etc.

The procedure corresponds to the following steps:

1. Measure qubit 1 with outcome m_1 in the basis

$$\{(N^{(1)\dagger})^{-1}R_z[-\xi_1]H|0\rangle, N^{(1)}R_z[-\xi_1]H|1\rangle\}; \quad (2.72)$$

2. If $m_1 = 0$, then success;

3. If $m_1 = 1$ then one has effectively teleported the gate $R_x[\epsilon^{(1)}]HR_z[\xi_1]$, where

$$\epsilon^{(1)} = \pm 2 \arctan \frac{\cos 2\theta^{(1)} \cos \xi_1}{1 \pm \cos 2\theta^{(1)} \sin \xi_1} + \pi. \quad (2.73)$$

Note that $\epsilon^{(1)} = 0$ when $N^{(1)} = U^{(1)}$ ($\theta^{(1)} = \pi/4$), as expected. Measure qubit 2 with outcome m_2 in the basis $\{U^{(2)}X|0\rangle, U^{(2)}X|1\rangle\}$;

4. Measure qubit 3 with outcome m_3 in the basis $\{(N^{(3)\dagger})^{-1}R_z[\chi]H|0\rangle, N^{(3)}R_z[\chi]H|1\rangle\}$, where $\chi = (-1)^{m_2}\epsilon^{(1)}$;

5. Repeat steps 3 and 4 on successive qubits $2k$ and $2k+1$ until outcome 1 is achieved on an odd qubit, using $U^{(2)} \rightarrow U^{(2k)}$, $N^{(3)} \rightarrow N^{(2k+1)}$, $m_2 \rightarrow m_{2k}$, $\epsilon^{(1)} \rightarrow \epsilon^{(2k-1)}$.

The key point of this example is that as for any measurement on an odd-numbered qubit that yields the ‘correct’ outcome $m_i = 0$, one will have succeeded in implementing part of the desired single-qubit rotation. Furthermore, any errors resulting from outcomes $m_i = 1$ are correctible by making further measurements. This thus constitutes a repeat-until-success strategy, and gives rise to a quasi-deterministic random-length single-qubit rotation. The likely reason for this one-dimensional state to be capable of processing a logical qubit is that the left and right parts of the state share an ebit of entanglement with respect to any cut, as mentioned in Sec. 2.4.1.

2.5.2 Deterministic Universal MBQC: 2D $N - U - N$ State

For universal MBQC, a two-dimensional resource state is required. The precise geometry of the two-dimensional state on which MBQC occurs is determined by the specific circuit to be implemented. Ideally one would start with a state defined on a convenient and simple geometry, and then ‘carve’ the desired shape out by deleting certain qubits. For cluster-state MBQC, for example, one carves the required state out of a rectangular lattice by projectively measuring the unwanted qubits in the computational basis. The goal is to yield isolated one-dimensional wires, each of which represents a logical qubit, with links only existing between wires in places where an entangling gate between logical qubits is needed.

Consider now a regular two-dimensional lattice composed of $N - U - N$ states, as depicted in Fig. 2.3. As in the usual cluster state, logical qubits are processed by alternating horizontal wires composed of physical qubits, and entangling gates by vertical chains connecting them. Unlike the cluster case, however, the procedure for implementing single-qubit rotations with $N - U - N$ states is of random length, so it is impossible to decide in advance where the desired links between wires will occur. The computational cluster state then must be carved ‘on the fly.’

Suppose that the quantum information encoding two logical qubits resides on (yet unmeasured) N -transformed physical qubits on two different wires. If an entangling gate between logical qubits is not desired at the next step, then the link between the wires can be first severed by measuring the intervening U -transformed chain qubit in the computational basis. An example of the method to decouple qubits 1 and 3 is shown in Fig. 2.3, where qubit 2 is measured in the computational basis. This has the effect of teleporting a Z gate to each of the logical qubits if the measurement outcome is $m = 1$. Other than taking into account the possible existence of these byproduct operators, the computation subsequently proceeds as in the one-dimensional case discussed above.




Figure 2.3: 2D N – U – N state, universal for quasi-deterministic MBQC. The qubits labeled 1, 2 and 3 can be used to implement an entangling gate, which involves measuring qubit 2 in the Y basis. Alternatively, if an entangling gate is not desired here, qubits 1 and 3 can be decoupled by measuring qubit 2 in the Z basis.

The desired entangling gate is implemented as follows. At the time that an entangling gate is needed, the local part of the resource state looks like two 1D N – U – N states, each coupled via CZ operations to an ancilla initially in the state $|+\rangle$ and subsequently acted on by an arbitrary U. In Fig. 2.3, the entangling link is represented by the vertical N–U–N chain labeled by qubits 1, 2, and 3. The local part of the state is mathematically described as

$$|R\rangle = N_1^{(1)} U_2^{(2)} N_3^{(3)} CZ_{1,2} CZ_{2,3} |c + t\rangle_{123}, \quad (2.74)$$

where the states $|c\rangle$ and $|t\rangle$ could be thought of as control and target states respectively for some entangling gate. Now, qubits 1 and 3 are measured in the usual Strategy I basis (2.72) with $\xi^{(i)} = 0$, while qubit 2 is measured in the eigenbasis of the Pauli operator Y, suitably rotated by $U^{(2)}$. This procedure teleports the state initially situated on qubits 1 and 3 through an entangling gate

$$\begin{aligned} G_{1,3} &= R_x [\mu^{(1)}]_1^{m_1} X_1^{m_1+m_2} H_1 \\ &\times R_x [\mu^{(3)}]_3^{m_3} X_3^{m_2+m_3} H_3 M_{1,3} \end{aligned} \quad (2.75)$$




Figure 2.4: B – U – B state, a one-dimensional structure that can be used for probabilistic random-length single-qubit rotations.

to qubits 4 and 5, with

$$M_{1,3} = |00\rangle\langle 00|_{1,3} + i|01\rangle\langle 01|_{1,3} + i|10\rangle\langle 10|_{1,3} + |11\rangle\langle 11|_{1,3}. \quad (2.76)$$

Here, the $\{\mu^{(i)}\}$ are the standard Strategy I byproduct angles, Eq. (2.37) or (2.73). This entangling operation is related to CZ via

$$\text{CZ}_{1,3} \equiv X_1 X_3 R_z[\pi/2]_1 R_z[\pi/2]_3 M_{1,3} X_1 X_3, \quad (2.77)$$

and so $G_{i,j}$ together with single-qubit operators forms a universal set of gates.

2.5.3 Probabilistic single-qubit rotations: B – U – B state

Next consider a one-dimensional state of the form

$$|R\rangle = B_1^{(1)} \otimes U_2^{(2)} \otimes B_3^{(3)} \otimes U_4^{(4)} \cdots \otimes B_{2n+1}^{(2n+1)} |Cl_{2n+1}\rangle, \quad (2.78)$$

where the $\{B^{(2i+1)}\}$ are B-type operators, and the $\{U^{(2i)}\}$ are once again local unitaries (c.f. Fig. 2.4). This structure ensures that none of the bonds present in the structure is perfect; no particle has maximal entropy of entanglement with the rest of the state. This fact is a consequence of Lemma 2.4.4; there is no qubit unaffected by B-type operators whose neighborhood contains any unaffected qubits.

All single-qubit measurements for the odd-numbered qubits now correspond to Strategy II, in which the byproduct operator is always X if it occurs. All even-numbered qubits are measured in the $\{U^{(2i)}|+\rangle, U^{(2i)}|-\rangle\}$ basis; as in the previous example, the

only purpose of these measurements is to enable the removal of undesired contributions to the rotation angles. The main difference from the previous example is that a non-unitary gate of the form $\text{HR}_z[\xi_{2i+1}]$ is teleported, where $\xi_{2i+1} \in \mathbb{C}$. The present goal is therefore to compensate for the imaginary part of the rotation angle.

As discussed in the previous section and Eq. (2.61), the imaginary part of ξ_{2i+1} is entirely determined by the ratio of the singular values of $B^{(2i+1)}$, and can be defined as $\epsilon = i \ln(\cot \theta^{(2i+1)})$. Consider momentarily the special case where the $\{B^{(2i+1)}\}$ all have the same singular values. The gate teleported by a measurement of the B-transformed qubit $2i+1$ will then be proportional to $R_z[\epsilon(-1)^{m_{2i}+m_{2i-2}+\dots}]$, ignoring all rotations about real angles which are entirely determined by the choice of measurement basis. In short, the sign of the imaginary angle depends on the outcomes of the previous measurements on even-numbered U-transformed qubits.

The imaginary component therefore undergoes a random walk of step-length $|\epsilon|$. In particular, the walker takes its first step to the right when the first measurement outcome of an even-numbered qubit is 0, and to the left if it is 1. Subsequently, a measurement outcome of 0 on an even-numbered qubit causes the walker to take another step in the same direction as the previous step, while outcome 1 makes the walker take a step in the opposite direction.

The two possible measurement outcomes with odd qubits are always equally likely, but the probabilities with even qubits depend on the singular values of the B-type operators from the (odd) neighboring qubits, and generally speaking the walker is more likely to stray further from the origin than to step back towards it. For example, consider measuring the first two qubits of $|R\rangle$ in Eq. (2.78) in the $\{|+\rangle, |-\rangle\}$ basis. It can easily be shown, using the Schmidt decomposition (2.41) and the expression (2.59), that the

probabilities of the outcome $|\pm\rangle$ on qubit 1 are equal, and that those on qubit 2 are

$$p_{\pm,2} = \frac{1}{2} (1 \pm \cos 2\theta^{(1)} \cos 2\theta^{(3)}) . \quad (2.79)$$

Here, the random walk effectively begins at position $\ln(\tan \theta^{(1)})$ and moves to $\ln(\tan \theta^{(1)}) \pm \ln(\tan \theta^{(3)})$. For a situation to arise where the walker moves closer to the origin with probability greater than 1/2, one of the two pairs of conditions

$$|\ln(\tan \theta^{(1)}) \pm \ln(\tan \theta^{(3)})| > 2 |\ln(\tan \theta^{(1)})| ; \quad (2.80)$$

$$\pm \cos 2\theta^{(1)} \cos 2\theta^{(3)} > 0 \quad (2.81)$$

must be simultaneously satisfied, for either sign. If the $\{\theta^{(i)}\}$ are chosen uniformly at random, then the probability of this happening is only about 0.315. If this measurement procedure is continued down the chain, with qubits 3-5 relabelled 1-3 after the first two measurements and so on, the current value of $\theta^{(1)}$ tends to drift away from $\pi/4$ towards either 0 or $\pi/2$, and the range of values of $\theta^{(3)}$ for which the walker is likely to turn around and walk towards the origin progressively shrinks. Furthermore, if $\{\theta^{(1)}\}$ and $\{\theta^{(3)}\}$ are equal at any time, the walker is guaranteed to be more likely to continue in one direction than to turn around.

This procedure constitutes a probabilistic method for implementing a single-qubit rotation. Unfortunately if the walker strays too far from the origin, it becomes effectively impossible to recover and the attempted gate teleportation fails. The entire computation must then be repeated. If the singular values of the B^{2i+1} are chosen such that the imaginary components of the teleported angles are all integer multiples of each other, then the behavior of the random walk is even more deleterious. A judicious two-dimensional arrangement of B–U–B chains avoids this catastrophe, as discussed in the next example.

2.5.4 Universal MBQC: Percolated 2D Cluster State from 3D B – U – B state

In the previous example using Strategy II, a possibly infinite number of steps may be required to teleport an arbitrary single-qubit unitary. But quitting the protocol results in catastrophic failure: because the computational wire is effectively broken, the entire gate teleportation must be attempted from the beginning. A solution to these problems is to employ the 3D extension to the previous resource, corresponding to a cluster state transformed by alternating B-type operators and unitaries. This corresponds to a lattice with two interpenetrating cubic sublattices, a B-lattice and a U-lattice.

An example of this 3D resource, a cube with side length 3, is depicted in Fig. 2.5. Initially, Z-basis measurements in the z -direction (as labeled in Fig. 2.5) are used to carve out a structure in which each B-transformed qubit in the $x - y$ plane, shaded grey, is attached to a long vertical B – U – B chain. Measurements are made on the chain qubits, starting at qubit above the B-transformed qubit on the computational wire and continuing in the vertical direction until success (defined below) is achieved. The goal is to probabilistically produce perfect entanglement in the $x - y$ plane, thereby effectively eliminating the B operators in the horizontal direction. The result is a 2D cluster state in this plane with missing entanglement bonds in random locations. As long as the mean density of broken links exceeds the percolation threshold for a two-dimensional square lattice, the resource is universal for MBQC [87].

Consider the first vertical chain from the left in Fig. 2.5. Recall that one can interpret B as a z -rotation by an imaginary angle $\pm i \ln \lambda$, as shown in Eq. (2.60). For simplicity, we assume the unitary operators acting on even-numbered qubits are all equal to the identity; were they not, they could be compensated by a suitable rotation of the measurement basis for. The portion of the state corresponding to the first four qubits of the vertical chain (with the qubit that intersects the horizontal chain labeled 1), is then (ignoring




Figure 2.5: 3D cluster state with B-type operators and unitaries acting on alternate qubits (a) before carving and (b) after carving. Grey qubits are acted upon by B-type operators and white qubits by local unitaries. Z-basis measurements are made in the z direction in (a) to disentangle those vertical chains originating from a white qubit in the $x - y$ plane. A measurement protocol along the remaining vertical chains in (b) produces a percolated 2D cluster in the $x - y$ plane.

normalization)

$$\begin{aligned} |T\rangle &= B_1 B_3 C Z_{1,2} C Z_{2,3} C Z_{3,4} |++++\rangle_{1234} \\ &= R_z [i \ln \lambda]_1 R_z [i \ln \lambda]_3 |Cl_4\rangle_{1234}, \end{aligned}$$

ignoring normalization factors as usual. First, qubits 2 and 3 are measured in the $\{|+\rangle, |-\rangle\}$ basis (of course, the measurement basis for qubit 2 would need to be rotated if a local unitary $U^{(2)}$ were acting). These are commuting measurements, since they are on different qubits and not adaptive. The effect is to teleport the state $R_z [i \ln \lambda] |+\rangle$ from qubit 3 through $X^{m_2} H X^{m_1} H \equiv X^{m_2} Z^{m_1}$ to qubit 1, yielding the new state

$$\begin{aligned} |T'\rangle &= R_z [i \ln \lambda]_1 C Z_{1,4} X_1^{m_2} Z_1^{m_1} R_z [i \ln \lambda]_1 |++\rangle_{14} \\ &= C Z_{1,4} R_z [i((-1)^{m_2} + 1) \ln \lambda]_1 |++\rangle_{14} \end{aligned}$$

up to overall local unitaries on the final state.

If $m_2 = 1$, then the imaginary part of the rotation angle is completely canceled. Qubit

4 can then be measured in the computational basis (again, suitably rotated if necessary) to disentangle the rest of the vertical chain from the horizontal chain. The result is a perfect cluster along the first three qubits in the horizontal direction, and the B-type operator is effectively deleted.

If $m_2 = 0$, then the situation is similar to the original. There is still a B-type operator present in the horizontal direction, now corresponding to a z-rotation about an angle with imaginary part $2\ln\lambda = \ln\lambda^2$. In other words, the new effective B-type operator in the horizontal chain has a ratio of singular values that is the square of the original one. In order to remove the effect of the B-type operator, the chain qubits must be measured sequentially until the total number of steps towards the origin exceeds by 1 the total number of steps away.

The probabilities $p_0^{(k)}$ and $p_1^{(k)}$ of the outcomes 0 and 1 on an even qubit in the vertical chain, where $k > 0$ is the present position of the walker on the real number line, are given by

$$p_0^{(k)} = \frac{1 + \lambda^{2k+2}}{1 + \lambda^2 + \lambda^{2k} + \lambda^{2k+2}} \sim O(1); \quad (2.82)$$

$$p_1^{(k)} = \frac{\lambda^2 + \lambda^{2k}}{1 + \lambda^2 + \lambda^{2k} + \lambda^{2k+2}} \sim O(\lambda^2). \quad (2.83)$$

The total probability p_n that the effect of the B will be undone within $2n$ measurements is the sum of the probabilities of all of the possible trajectories of the walker on n or fewer steps with initial position 1, final position 0 and all intermediate positions strictly positive.

The probability p_{10} of undoing the B operator after 10 attempts (20 measurements) is shown in Fig. 2.6 as a function of the ratio of singular values λ . Calculation of the exact probability p_∞ is computationally intractable, for two reasons. First, the number of valid trajectories for the walker grows exponentially in the number of steps allowed.

Second, the probability of any particular trajectory depends on the full history of the walker, not just the number of steps. Of course, p_∞ must approach 1 as λ approaches unity (the limit that B becomes a unitary matrix). In this case, the walk reduces to the simple 1D random walk, which is known to sample the origin frequently.

If after some predetermined number of measurements along a vertical chain one has not yet succeeded in undoing B , the qubit at the root of the chain (i.e. in the computational wire) can be measured in the computational basis and thereby deleted. The result is a broken link in the 2D cluster state. The important result shown in Fig. 2.6 is that there is a critical value of λ , called λ_c , above which the probability of successfully undoing the B rises above the (bond) percolation threshold for a 2D square lattice (approximately 0.593). For this walk, the critical value obeys $\lambda_c \lesssim 0.379$. The upper bound for λ_c is read off the thick blue curve from Fig. 2.6. Thus, this procedure probabilistically yields a universal resource for MBQC provided that λ is sufficiently large. We note that a similar example was considered in [61], where the resource was a 2D cluster state with identical B -type operators acting everywhere and the percolation proceeded via two-element POVMs that either removed the B or deleted the qubit. There, the critical value of λ was found to be 0.649.

2.6 Discussion and Conclusions

Motivated by a desire to identify new resource states for measurement-based quantum computing, we have performed a (non-exhaustive) search of the equivalence class of n -qubit cluster states on a rectangular lattice, under the action of $\text{GL}(2, \mathbb{C})^{\otimes n}$. In particular, our aim was to identify which states within this class could be used as resources for MBQC, by designing explicit protocols for teleporting single-qubit gates and two-qubit entangling gates, driven by adaptive local projective measurements. We identified a




Figure 2.6: (Color online) Probability of success of the procedure for deleting a B-type operator in the plane via a random walk in the third dimension, as a function of the ratio λ of the singular values (thick blue, color online). The probabilities p_k of deleting B with exactly k even-qubit measurements are also shown for k from 1 to 10 (thin, decreasing with increasing k), and the thick blue line is the sum of these. The red dashed line is the percolation threshold.

class of one-dimensional states, the so-called $N - U - N$ states, that are deterministically universal for single-qubit rotations, although with a random number of measurements needed to teleport the desired rotation. We also identified a probabilistically universal resource for single-qubit rotations: the so-called $B - U - B$ states. We then described a three-dimensional extension of $B - U - B$ states that can yield a universal resource for deterministic MBQC beyond a percolation threshold, and a 2D $N - U - N$ state that is also universal for deterministic but random-length MBQC.

Several interesting open issues arise as the result of this work. First, it is not clear what (if any) relationship exists between the states uncovered in this work and other known resource states. For example, the probabilistic nature of the protocol with $B - U - B$ states has features in common with that of other resources for MBQC, such as photonic cluster states prepared via probabilistic entangling gates or with unreliable sources [46, 94, 95, 96]. Likewise, the quasi-deterministic $N - U - N$ states share various characteristics with the AKLT-inspired resources of Refs. [47, 50, 76, 61], in particular the exponential spin correlations and the repeat-until-success measurement-based strategies. One distinction is that only one alternating sublattice of the $N - U - N$ states exhibits non-zero correlation functions, whereas in the AKLT chain, every qubit is correlated with every other. Presumably a true identification of an AKLT-type resource with a N -transformed cluster state will require N -type operators to be present on every qubit, a case we have not handled here.

Also, it will be important to better understand the relationships of these states with the universal quantum wires of Ref. [85]. In that work, certain reasonable physical assumptions were imposed on 1D wires at the outset, for example: the possibility of producing a wire via a translationally invariant nearest-neighbour global entangling operation, the asymptotic sharing of an ebit of entanglement between the left and right halves of the chain, etc. In our work, it is not clear if a translationally invariant scheme

exists for producing $N - U - N$ or $B - U - B$ chains. Evidence from exact calculations on small chains and the explicit description of the states within the MPS representation reveals that although the left and right halves of $N - U - N$ chains share an ebit, the halves of the $B - U - B$ chains do not. This seems reasonable, as the $N - U - N$ chain is quasideterministically universal for single-qubit rotations, while the $B - U - B$ chain is only probabilistically so.

Second, it is conceivable that all states that have been hitherto identified as universal resources for MBQC are in fact SLOCC-equivalent to the family of cluster states. There is some evidence to support this conjecture. For example, the results of Ref. [83] show that many seemingly diverse resource states can be reduced to cluster states via local strategies. Similarly, the proof of the universality of the 2D AKLT state on a honeycomb lattice proceeds via local reduction to a random graph state, which can in turn be reduced to a percolated cluster state [55]. This reduction is successful despite the fact that the initial resource is defined on qutrits rather than qubits, and on a non-rectangular lattice. An even more intriguing possibility (though we believe it to be unlikely) is that all possible states for universal MBQC fall within the SLOCC-equivalence class of the cluster states. At the very least, the relative size of this class decreases exponentially with the total number of physical qubits, as expected [67, 68].

Third, while we have shown that a certain subset of the orbit of the cluster states under SLOCC are useful resources, either probabilistically or quasi-deterministically, it is not clear if the remaining SLOCC-transformed cluster states are also universal resources for MBQC. Generically, single-qubit measurements on these states teleport gates with byproduct operators that are rotations about the X axis by complex angles. One possibility is that there is a measurement protocol that can accommodate all possible byproducts that we simply haven't found. Perhaps there is another sense, besides quasi-deterministic or probabilistic, in which these states can be said to be useful for MBQC.

Another possibility is that there is some map from the full orbit to the particular subset of states considered in the work. Alternatively, the states in the full orbit may not be useful resources, though we have not attempted to prove this.

Finally, it would be useful if the resources presented in this work could be realized as the ground states of a physical Hamiltonian. A recent no-go theorem [97] shows this cannot be the case for frustration-free Hamiltonians on qubits. The question remains open for frustrated Hamiltonians, for instance the AKLT Hamiltonian in the Haldane phase [86, 48, 76]. Another possibility is that the ground states of physical Hamiltonians could be locally reduced to the resources we have found. Further research is needed to answer these and related questions.

2.7 Acknowledgements

The authors are grateful to Jens Eisert, David Gross, and Gilad Gour for stimulating discussions. This research was supported by Alberta Innovates - Technology Futures and the Natural Sciences and Engineering Research Council of Canada (NSERC).

Chapter 3

Introduction to Fermionized Photons in the Ground State of One-Dimensional Coupled Cavities

3.1 Introduction

In Chapter 4, I present a paper discussing the nature of the quantum phase transition between the Mott-insulating and superfluid phases in cavity-based quantum electrodynamics (henceforth cavity QED). The model is called the Jaynes-Cummings-Hubbard model, which has been proposed as a highly controllable simulator for the widely-known Bose-Hubbard model of interacting bosons in a lattice. The purpose of this chapter is to provide a detailed introduction to the key concepts needed as background material for Chapter 4, which should then be intelligible to a non-specialist audience.

The plan of the chapter is as follows: in Section 3.2, I introduce the Bose-Hubbard model (by way of the Hubbard model) in a qualitative fashion. Because it is a model of a many-body system of indistinguishable particles, it is most conveniently described not via the single-particle quantum mechanics of Chapters 1 and 2, but in the framework of second quantisation, which is briefly introduced in Section 3.3. Following this, I formally derive the Bose-Hubbard model as a simplification of the general interacting boson problem in Section 3.4. Because the Bose-Hubbard model exhibits a phase transition, I introduce some of the important concepts behind phase transitions in Section 3.5. Then, in Section 3.6 I discuss the basic physics of the Bose-Hubbard model in 3D, touching specifically upon Mott insulators, Bose-Einstein condensation and superfluidity. For various reasons, the physics is quite different in 1D, and I describe the 1D situation in Section 3.7. Then, I move on to cavity QED; I explain the Jaynes-Cummings model of

a single high-finesse optical cavity containing a two-level atom in Section 3.8 and then, making heavy use of the concepts already presented in the context of the Bose-Hubbard model, the situation of a network of coupled cavities — the Jaynes-Cummings-Hubbard model — in Section 3.9. Neither the Bose-Hubbard nor the Jaynes-Cummings-Hubbard model can be solved exactly in the general case, so numerical studies are often necessary. In Section 3.10, I explain the density matrix renormalisation group algorithm that I use for my results in Chapter 4, as well as its connection to the MPS formalism discussed in Chapter 1.

3.2 Introduction to the Hubbard and Bose-Hubbard models

Materials often occur (or are manufactured) as crystalline structures, comprising ions (atomic nuclei and core electrons) arranged in a regular lattice and surrounded by a sea of conduction electrons, which are not associated with any particular ion. The Hamiltonian describing a general crystalline material is complicated, and the number of ions and electrons in a macroscopic chunk of material is of the order of 10^{23} . The electronic degrees of freedom give rise to many interesting properties, but obtaining an exact solution for any static or dynamical properties of interest is intractable.

Because the underlying many-body system is so complicated, theoretical progress in the study of real materials generally relies on approximations leading to simplified models. One important example is the Hubbard model, the simplest version of which models the underlying system as fermions (electrons, for example) defined on a discrete lattice in real space, with kinetic energy described by hopping processes between neighbouring sites of the lattice, and interactions occurring only between fermions on the same site. A natural variation of the Hubbard model is the Bose-Hubbard model, first discussed in 1989 [3], which is identical to the Hubbard model except that the particles are bosons instead of

fermions.

The Hubbard model has been studied via a battery of theoretical and numerical approaches, but clean realisations of the model in real materials are hard to come by, owing to the presence of impurities, phonons and long-ranged interactions. Furthermore, the Bose-Hubbard model has no obvious natural realisation. However, in the last two decades, a technological innovation known as the optical lattice [98] has largely solved this problem. An optical lattice is a periodic electromagnetic field, created by the interference of counterpropagating monochromatic lasers, that creates a lattice potential for neutral atoms by coupling to their dipole moments. The optical lattice has no impurities, no phonon modes, and the strength of the interparticle interactions can be made arbitrarily small or large relative to the kinetic energy by varying the well depth of the lattice (or via Feshbach resonances [99] which I will not discuss here), resulting in an extremely flexible platform for simulating either the Hubbard or Bose-Hubbard model. Most notably for the present purposes, cold bosons loaded into an optical lattice have been used to demonstrate the Mott insulator to superfluid transition I will discuss shortly [100]. Since my work involves only the Bose-Hubbard model itself and not any particular physical implementation of it, there will be no further mention of optical lattices; nevertheless, they are one of the most important experimental realisations of the periodic potential necessary for implementing the model.

3.3 Interlude: introduction to second quantisation

In Chapter 1, I described n -qubit quantum states as superpositions of tensor products of single-qubit states, which can be represented as vectors; I also described quantum operators as matrices acting upon these state vectors to produce different state vectors.

In particular, an n -qubit operator $\mathcal{O}$ could be written as

$$\mathcal{O} = \sum_{i,j=1}^{2^n} \mathcal{O}_{ij} |i\rangle \langle j|, \quad (3.1)$$

and an n -qubit state $|\psi\rangle$ as

$$|\psi\rangle = \sum_{i=1}^{2^n} \psi_i |i\rangle. \quad (3.2)$$

Henceforth, I will refer to this method of describing quantum states and operators as first quantisation.

In this chapter, most of the important quantum operators act on continuous degrees of freedom like position or momentum, rather than the discrete qubit operators discussed earlier. Generalising the first quantisation representation to systems with continuous degrees of freedom (for example, position $\mathbf{r}$) is no problem; one simply replaces the summations from Eqs. (3.1) and (3.2) with integrals:

$$\mathcal{O} = \int d\mathbf{r} \int d\mathbf{r}' \mathcal{O}(\mathbf{r}, \mathbf{r}') |\mathbf{r}\rangle \langle \mathbf{r}'|, \quad (3.3)$$

and an n -qubit state $|\psi\rangle$ as

$$|\psi\rangle = \int d\mathbf{r} \psi(\mathbf{r}) |\mathbf{r}\rangle. \quad (3.4)$$

In this chapter, the quantum systems of interest are many-particle states of bosons, fermions or quasiparticles. Such states have at least two important differences from multi-qubit states that render first quantisation cumbersome:

- (i) the particles are fundamentally indistinguishable, unlike qubits;
- (ii) the number of particles is not necessarily well-defined.

Because of the fundamental indistinguishability of bosons (fermions), any multi-particle wavefunction must be appropriately symmetrised (antisymmetrised). This means

that the only effect of exchanging two particles must be to multiply the overall wavefunction by a physically undetectable global phase. For example, if $\Psi_{\pm}(\mathbf{r}_1, \dots, \mathbf{r}_n)$ is an n -particle wavefunction, where ‘+’ denotes bosons and ‘−’ fermions, it must satisfy the property

$$\Psi_{\pm}(\dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots) = \pm \Psi_{\pm}(\dots, \mathbf{r}_j, \dots, \mathbf{r}_i). \quad (3.5)$$

The (anti)symmetrisation of the n -particle wavefunction can be accomplished by the use of Slater determinants for fermions or permanents for bosons. Suppose the single-particle wavefunction of particle i has label ν_i , and is written in the position representation as $\psi_{\nu_i}(\mathbf{r}_i)$. The n -particle state is then given by

$$\Psi_{\pm}(\dots, \mathbf{r}_i, \dots, \mathbf{r}_j, \dots) = \begin{vmatrix} \psi_{\nu_1}(\mathbf{r}_1) & \psi_{\nu_1}(\mathbf{r}_2) & \dots & \psi_{\nu_1}(\mathbf{r}_n) \\ \psi_{\nu_2}(\mathbf{r}_1) & \psi_{\nu_2}(\mathbf{r}_2) & \dots & \psi_{\nu_2}(\mathbf{r}_n) \\ \vdots & \vdots & \ddots & \vdots \\ \psi_{\nu_n}(\mathbf{r}_1) & \psi_{\nu_n}(\mathbf{r}_2) & \dots & \psi_{\nu_n}(\mathbf{r}_n) \end{vmatrix}_{\pm}, \quad (3.6)$$

with the operation on the right hand side corresponding to a standard matrix determinant in the ‘−’ case, and a matrix permanent (a signless equivalent of the determinant) in the ‘+’ case. The number of terms required to properly (anti)symmetrise a fully separable n -particle configuration is $n!$, which goes like $e^{n \log n}$ for large n according to Stirling’s approximation. The notation therefore quickly becomes cumbersome as n is increased.

The possibility of a variable number of particles in the system is another problem for first quantisation. In this case, the Hilbert space for the total system takes the form of a direct sum of Hilbert spaces for each individual possible total number of particles:

$$\mathcal{H} = \bigoplus_{n=0}^{n_{\max}} \mathcal{H}_n, \quad (3.7)$$

where $n_{\max}$ is the maximum number of particles allowed, which in practice for a macroscopic condensed matter system is huge ($\sim 10^{26}$) in the canonical ensemble, and unlimited

in the grand canonical ensemble.

The problems of (anti)symmetrisation and variable particle number can both be circumvented by using an alternative formulation of quantum mechanics called second quantisation. In this picture, particles are viewed as fundamental excitations of a quantum field, rather than as fundamental objects. The energy stored in the quantum field is quantised (just like the energy of the individual particles in first quantisation). The ground state of the field is called the vacuum state, usually denoted $|0\rangle$ or $|\Phi_0\rangle$. The single-particle eigenstates from first quantisation are viewed as excitation modes of the quantum field, and the multi-particle states of second quantisation are conveniently represented in a basis that lists the occupation numbers of each mode, also known as Fock space. Second quantisation is described in a vast number of textbooks and papers; my presentation and notation in this section follow Ref. [101].

Central to the second quantisation representation are the creation and annihilation operators, which create or destroy excitations of the field. For example, consider the single-particle eigenstate $|\psi_i\rangle$. Let a_i denote the annihilation operator that destroys an excitation of the mode labelled i , and its adjoint $a_i^\dagger$ be the creation operator that creates such an excitation. The effect of these operators on a state $|N_i\rangle$ in Fock space corresponding to N_i excitations of the mode i and no other excitations, is the following:

$$a_i|N_i\rangle = \sqrt{N_i}|N_i - 1\rangle; \quad (3.8)$$

$$a_i^\dagger|N_i\rangle = \sqrt{N_i + 1}|N_i + 1\rangle. \quad (3.9)$$

From the above relations, it is easy to check that $|N_i\rangle$ is an eigenstate of the operator $a_i^\dagger a_i$ with eigenvalue N_i ; this operator is henceforth called the number operator.

The bosonic or fermionic character of the particles is accounted for by the (anti)commutation

relations of the creation and annihilation operators. These relations are

$$[a_i, a_j]_{\pm} = 0; \quad (3.10)$$

$$[a_i^{\dagger}, a_j^{\dagger}]_{\pm} = 0; \quad (3.11)$$

$$[a_i, a_j^{\dagger}]_{\pm} = \delta_{i,j}. \quad (3.12)$$

Here, the ‘+’ relations apply to bosons and the ‘−’ relations to fermions. The operator $[A, B]_{+} \equiv [A, B]$ is the commutator of the operators A and B , and $[A, B]_{-} \equiv \{A, B\}$ is the anticommutator.

States and operators in second quantisation can be conveniently represented using the creation and annihilation operators. An n -particle Fock state is denoted by an appropriate product of n creation operators acting on the vacuum state. For example, the state whose position representation is given in Eq. (3.6) can be compactly written as

$$|\Psi_{\pm}\rangle = a_{\nu_n}^{\dagger} a_{\nu_{n-1}}^{\dagger} \dots a_{\nu_1}^{\dagger} |0\rangle. \quad (3.13)$$

The (anti)symmetrisation is completely accounted for by the (anti)commutation relations obeyed by the creation operators. As for operators, consider the general first-quantised expression for a single-particle operator from Eq. (3.1). Each matrix element $\mathcal{O}_{ij}$ is associated with mapping the basis state $|j\rangle$ to $|i\rangle$. To accomplish the same effect in second-quantisation, one would associate $\mathcal{O}_{ij}$ with the annihilation of a particle in state $|j\rangle$ and the creation of one in state $|i\rangle$, leading to the representation

$$\mathcal{O} = \sum_{i,j} \mathcal{O}_{ij} a_i^{\dagger} a_j. \quad (3.14)$$

This is just a heuristic argument, but a rigorous derivation of this representation can be

found in many textbooks, including for example Ref. [101]. Completely analogously, a two-particle first-quantised operator

$$\mathcal{U} = \sum_{i,j,k,l} \mathcal{U}_{ij;kl} |ij\rangle \langle kl| \quad (3.15)$$

obtains the second-quantised representation

$$\mathcal{U} = \sum_{i,j,k,l} \mathcal{U}_{ij;kl} a_i^\dagger a_j^\dagger a_k a_l. \quad (3.16)$$

3.4 Derivation of Bose-Hubbard model

The derivation of the Bose-Hubbard model can be found in many sources, but because the model plays such an important role in Chapter 4, I will reproduce the central points here. The full Hamiltonian of the interacting N -boson problem in an external potential takes the form

$$H = \sum_{i=1}^N \left[\frac{\mathbf{p}_i^2}{2m} + V(\mathbf{r}_i) \right] + \frac{1}{2} \sum_{i,j=1}^N U(|\mathbf{r}_i - \mathbf{r}_j|), \quad (3.17)$$

where $\mathbf{p}_i$ is the momentum of particle i , $V(\mathbf{r}_i)$ is the external potential on that particle and $U(|\mathbf{r}_i - \mathbf{r}_j|)$ is the two-body interaction between particles i and j . In principle, there may be interactions with larger numbers (three or more) of participating particles, but we assume that the gas is sufficiently dilute that these interactions can be safely neglected. In deriving the Bose-Hubbard model, I assume that the potential is spatially periodic, keeping in mind the applications to crystalline or optical lattices, but make no specific reference to its physical implementation.

Because of the many-particle nature of the problem, it is helpful to express the Hamiltonian using second quantisation. The first decision is the choice of the working basis. Let us first focus on choosing an appropriate basis for the single-particle problem. In

this case the interaction term in Eq. 3.17 drops out, leaving the kinetic energy term and the periodic potential. For simplicity, I will assume the problem is one-dimensional (call the spatial coordinate x), but all of the concepts leading to the derivation of the model generalise to two or three dimensions as well. The Hamiltonian of interest is thus

$$H_0 = \frac{p^2}{2m} + V(x). \quad (3.18)$$

In the single particle case, it is immaterial whether the particle is a boson or a fermion since there is no possibility of particles exchanging positions. I said that the Bose-Hubbard model is simply the bosonic version of the Hubbard model for fermions, so sometimes I will discuss the single-particle case as if it were a single electron in a crystalline lattice. Nevertheless, I emphasise again that the actual exchange statistics of the particle, and the source of the periodic potential, are immaterial. The discussion of the tight-binding model presented here borrows heavily from Refs. [102] and [103].

The Hubbard and Bose-Hubbard models are defined on a lattice, and the locations of the potential minima will correspond to the sites of the lattice. Let L be the number of lattice sites and a be the smallest positive real number such that $V(x + a) = V(x)$, henceforth called the lattice spacing. A result known as Bloch's theorem [104] states that the single-particle wavefunctions $\{\phi_j\}$ satisfying the Schrödinger equation

$$-\frac{\hbar^2}{2m} \frac{d^2 \phi_k}{dx^2} + V(x) \phi_j = \epsilon_j \phi_j \quad (3.19)$$

must obey

$$\phi_k(x + a) = e^{ija} u_j(x), \quad (3.20)$$

where $j = 2\pi\tau/La$ with τ an arbitrary integer, and $u_j(x + a) = u_j(x)$. Notice that because of the periodicity, $e^{i(j+\nu L)a} = e^{ija}$ for any integer ν . It is convenient to change

the labelling of the wavefunctions to

$$\phi_{k,\nu}(x) = e^{ika} u_{k,\nu}(x), \quad (3.21)$$

with $k = 2\pi\tau/La$ restricted to a range of L values given by, say, $\tau \in \{0, 1, \dots, L-1\}$, and ν an arbitrary integer. The allowed range of k is called the Brillouin zone (BZ) and k is called the quasimomentum. In the crystalline setting, with $L \rightarrow \infty$, the discrete variable k becomes continuous and the allowed energies generally separate into continuous bands, labelled by ν , with the allowed values of k for each band restricted to the Brillouin zone. The label ν is called the band index, and the $\{\phi_{k,\nu}\}$ are called Bloch functions. The Bloch functions form a complete orthonormal basis and tend to be spread out over the lattice, rather than localised in space. If the temperature of the system is low and the number of particles is sufficiently small that they can all be accommodated in a single band (this is an issue for fermions because of the Pauli exclusion principle, but not for bosons), it is often reasonable to make the single-band approximation, in which one neglects all but the lowest-lying energy band. I will do so from now on, and consequently will drop the band index ν .

If the minima of the potential V are sufficiently deep that the particles are tightly confined near them most of the time, then the Bloch functions probably are not the best choice of basis for capturing the physics of the system. Instead, one defines a new basis of functions $\{w(x - x_j)\}$ called Wannier functions [105], which are related to the Bloch functions via

$$\begin{aligned} w(x - x_j) &= \frac{1}{\sqrt{L}} \sum_k e^{-ikx_j} \phi_k(x); \\ \phi_k(x) &= \frac{1}{\sqrt{L}} \sum_j e^{ikx_j} w(x - x_j). \end{aligned}$$

with $j \in \{0, 1, \dots, L-1\}$ and $\{x_j = ja\}$ the positions of the lattice sites. It can be shown that the Wannier functions also constitute a complete orthonormal set, and moreover each one is tightly localised around a particular lattice site. This, therefore, is a useful basis of choice for second-quantising the problem at hand.

In the Wannier basis, where b_i and $b_i^\dagger$ are the creation and annihilation operators corresponding to the Wannier function localised at site i , the non-interacting boson Hamiltonian in second quantisation is

$$H = \sum_{i,j} h_{ij} b_i^\dagger b_j, \quad (3.22)$$

where

$$h_{ij} = \int dx w^*(x - x_i) H_0 w(x - x_j). \quad (3.23)$$

To introduce the standard notation, let's invoke translational invariance and make the definitions $\hbar\omega_0 := h_{ii}$ for any site i and $t_{ij} := h_{ij}$. The most important assumption of the tight-binding model is that the terms t_{ij} for terms with $|i - j| > 1$ can be neglected, because for $i \neq j$, there is no place where $|w(x - x_i)|$ and $|H_0 w(x - x_j)|$ are both significant, and the already small overlap for $|i - j|$ decreases rapidly as a function of $|i - j|$. Thus, we denote the terms $t_{ij} = t_{ji} = t$ for $|i - j| = 1$ (assuming they are real), and neglect the others. The model describing this situation is called the tight-binding model, and is written in second-quantisation as

$$H_{\text{TB}} = \hbar\omega_0 \sum_i b_i^\dagger b_i + t \sum_{\langle i,j \rangle} (b_i^\dagger b_j + b_j^\dagger b_i), \quad (3.24)$$

where $\langle i, j \rangle$ indicates that i and j are nearest-neighbours. It should be noted that the tightly-localized Wannier functions described here are specialised to the tight-binding model. Other superpositions of Bloch functions can be defined in other situations, leading

to different Wannier functions; see [106] for a modern review.

The next task is to consider the interaction term H_{int} . The second quantised form of the interaction $U(|x_i - x_j|)$ in the Wannier basis is

$$U = \sum_{i,j,k,l} U_{ij;kl} b_i^\dagger b_j^\dagger b_k b_l, \quad (3.25)$$

where

$$U_{ij;kl} = \int_{-\infty}^{\infty} dx_1 \int_{-\infty}^{\infty} dx_2 w^*(x_1 - x_i) w^*(x_2 - x_j) U(|x_1 - x_2|) w(x_1 - x_k) w(x_2 - x_l). \quad (3.26)$$

In the tight-binding model, because of the tightly-localised nature of the Wannier functions, one assumes that the on-site terms $U_{ii;ii}$ dominate the others, and are therefore the only ones retained. Defining $U \equiv U_{ii;ii}$, one obtains (using the bosonic commutation relations)

$$\begin{aligned} H_{\text{int}} &\approx \frac{U}{2} \sum_i b_i^\dagger b_i^\dagger b_i b_i \\ &= \frac{U}{2} \sum_i b_i^\dagger (b_i b_i^\dagger - 1) b_i \\ &= \frac{U}{2} \sum_i n_i (n_i - 1). \end{aligned}$$

Putting all of the preceding results together, one obtains the Bose-Hubbard model in the canonical ensemble:

$$H_{\text{BH}} = \sum_{i=1}^N \frac{\hbar\omega_0}{2} n_i - t \sum_{\langle i,j \rangle} (b_i^\dagger b_j + b_i b_j^\dagger) + \frac{U}{2} \sum_{i=1}^N n_i (n_i - 1). \quad (3.27)$$

Sometimes it is more convenient to work in the grand canonical ensemble, to allow for fluctuations in the particle number. In that case, introducing the chemical potential μ ,

one instead obtains

$$H_{\text{BH}} = -t \sum_{\langle i,j \rangle} (b_i^\dagger b_j + b_i b_j^\dagger) + \frac{U}{2} \sum_{i=1}^N n_i (n_i - 1) - \mu \sum_{i=1}^N n_i. \quad (3.28)$$

Since μ couples to the total particle number, it is the parameter that sets the mean total occupation of the lattice. Note that in (3.28), the factor $\hbar\omega_0/2$ from the first term in (3.27) has been absorbed into the chemical potential. In other words, the chemical potential is measured relative to the ground state energy of the M -particle noninteracting problem, which is $\hbar\omega_0 M/2$.

3.5 Phase transitions

Since Ch. 4 is principally concerned with a quantum phase transition, a general discussion of this topic is needed here. In fact, most of the formal machinery is borrowed from the theory of classical phase transitions, which are driven by thermal fluctuations of a quantity called the order parameter (to be discussed shortly) and therefore occur at finite temperature. Quantum phase transitions on the other hand occur at zero-temperature and are driven by purely quantum fluctuations of the order parameter, associated with the Heisenberg uncertainty principle. I will therefore begin this section by discussing the basics of classical phase transitions, and then provide a brief description of how quantum phase transitions are related. All of the information discussed in this section can be found in many textbooks; see for example [107] and [108].

3.5.1 Classical Phase Transitions

The nature of many physical systems in thermal equilibrium depends strongly upon the temperature at which the system is held. A familiar example is of water, which for any given value of the pressure is a solid below a temperature called the melting point (at

the standard atmospheric pressure of 101.3 kPa, this temperature is 0°C), and a fluid (usually a liquid) above it. Another familiar example is a ferromagnetic material, which has a non-zero spontaneous magnetisation even in the absence of an external magnetic field when the sample is below the so-called Curie temperature (770°C for iron), but no net magnetisation above this temperature. The transition of water from liquid to solid, or for iron from ferromagnetic to nonmagnetic, as the temperature is increased is an example of a (classical) phase transition. The temperature at which this transition occurs is called the critical temperature.

Suppose the system under consideration has some generalised set of microscopic degrees of freedom $\{x_i\}$ (for example, the direction of each magnetic moment in a ferromagnet or the positions and momenta of each atom in a monatomic fluid). In the thermodynamic limit, where the number of individual constituents (henceforth called “particles” or “bodies”) of the system and the physical size of the system go to infinity while the mean density remains constant and finite, the state of the system can usually be described by macroscopic variables (such as the volume, pressure and total energy for a fluid, or the total magnetisation and total energy for a ferromagnet). Such a description is called a macrostate; in general, many different microscopic configurations (microstates) give rise to the same macrostate, with the number of such microstates being given the symbol Ω .

The energy U of the system in a particular microscopic configuration is specified by the Hamiltonian $H(\{x_i\})$. The equilibrium state at temperature T is the one that minimises the (Helmholtz) free energy $F = U - TS$ of the system, where $S = k_B \log \Omega$ is the thermodynamic entropy of the system. At fixed temperature, the system can lower its free energy in one of two ways – either reduce its internal energy U or increase its entropy S . The microscopic configurations corresponding to small U tend to be highly ordered; Ω is small. On the other hand, S is maximised when Ω is as large as possible.

There is therefore a fundamental competition between minimising energy and maximising entropy, which is related to the temperature. At low temperatures, the contribution to F due to S is small, so minimising U is more important. Conversely, at high temperatures, maximising S is more important. The system must therefore transition from an ordered state at low T to a disordered one at high T – a phase transition. The temperature at which this occurs is called the critical temperature, denoted T_c .

For a large class of phase transitions, the phase of the system can be classified by a quantity called the order parameter, often denoted Ψ , which is zero whenever the system is in the disordered phase and non-zero in the ordered phase. The choice of order parameter depends on the system in question, and is defined as the derivative of the free energy with respect to the parameter that couples to the relevant degree of freedom in the Hamiltonian. For example, for a liquid-gas transition, the order parameter is the density, which is the derivative of the free energy with respect to the chemical potential. For another example, consider a ferromagnetic material on a lattice with local magnetic moments $\{S_i\}$ (henceforth called “spins”) in the presence of an external magnetic field H , described by the Ising model:

$$H_{\text{Ising}} = -J \sum_{\langle i,j \rangle} S_i S_j - H \sum_i S_i. \quad (3.29)$$

Here, $\langle i, j \rangle$ denotes a sum over nearest-neighbour spins only, $S_i = \pm 1$ denotes whether a particular spin is aligned (+1) or antialigned (−1) with the external field, and $J > 0$ is the strength of the interaction between neighbouring spins. Because $J > 0$, the spins have a tendency to align. In this case the order parameter for a uniform system of N spins is the total magnetisation Nm , where $m := \langle S_i \rangle$ for all i is the site-independent mean magnetisation. It can be shown that the magnetization is the derivative of the free

energy with respect to the magnetic field strength at constant temperature:

$$\Psi \equiv m := \left(\frac{\partial F}{\partial H} \right)_T. \quad (3.30)$$

Phase transitions are associated with non-analytic behaviour of the free energy at the critical temperature, usually due to a discontinuity, or the non-existence of one of its derivatives with respect to temperature or an external field strength. A phase transition in which the order parameter jumps discontinuously (as a function of T) from zero to a non-zero value at the critical temperature is called a first-order transition. The familiar solid-liquid-gas transitions are generally of this type. Sometimes, however, the order parameter itself is a continuous function of temperature but has a cusp such that one of its derivatives with respect to temperature is discontinuous. An important example is the ferromagnetic transition at the Curie temperature. Those transitions associated with continuous order parameters are called continuous phase transitions.

All important thermodynamic quantities can be obtained from the free energy (including the internal energy U). Apart from the order parameter, another example is the heat capacity C_V , which specifies the amount of heat required to raise the temperature by 1°C at constant volume. Since the first law of thermodynamics implies that the heat transfer is equal to the change in internal energy when no work is done (because the volume is constant), the heat capacity can be calculated via

$$C_V = \left(\frac{\partial U}{\partial T} \right)_V. \quad (3.31)$$

Another is the susceptibility χ , which describes the linear response of the order parameter

to an external field H and is defined by

$$\chi = \left(\frac{\partial \Psi}{\partial H} \right)_T. \quad (3.32)$$

Calculating these quantities of interest in the neighbourhood of a phase transition requires an expression for the free energy when the temperature is near the critical temperature.

An exact expression for the free energy is usually unavailable, with the exception of a small number of exactly solved models. Most of these are in 1D; exactly solved models in two or higher dimensions are exceedingly rare, one example being the 2D Ising model with no external field [109]. The reason for the difficulty is usually the fact that the individual constituents of the model (I will call them “bodies” or “particles” henceforth) are coupled via physical interactions. In the absence of interactions, a solution to the single-body problem implies a solution to the full problem: the free energy of the system is just the sum of the individual free energies. One avenue for finding an approximate solution to an interacting many-body system in arbitrary dimensions is therefore to represent it as a single-body problem in an effective external field that is due to the interactions with the other particles.

The standard method for reducing the many body problem to a single-body one is called mean field theory. The idea is to calculate the energy of a single particle by guessing the mean configuration of those with which it is interacting, and using this to approximate the free energy of the system as a whole. This guess must be made self-consistently; any quantities calculated using the mean field approximation must be consistent with the approximation itself.

For example, consider the mean-field solution of the Ising model on a hypercubic lattice of arbitrary dimension. To map it to a single-body problem, the interaction terms $-J \sum_j S_i S_j$ for a particular spin i , where the sum is over its nearest neighbours only,

must be rewritten. This is accomplished by re-expressing each spin variable S_i in terms of its average value $\langle S_i \rangle$ and the deviation $\langle S_i \rangle - S_i$ from the average:

$$H_{\text{Ising}} = -J \sum_{\langle i,j \rangle} (\langle S_i \rangle + S_i - \langle S_i \rangle) (\langle S_j \rangle + S_j - \langle S_j \rangle) - H \sum_i S_i \quad (3.33)$$

$$\begin{aligned} &= -J \sum_{\langle i,j \rangle} [\langle S_i \rangle \langle S_j \rangle + \langle S_i \rangle (S_j - \langle S_j \rangle) \\ &\quad + \langle S_j \rangle (S_i - \langle S_i \rangle) + (S_i - \langle S_i \rangle) (S_j - \langle S_j \rangle)] - H \sum_i S_i. \end{aligned} \quad (3.34)$$

One can make the ansatz that in the ground state, every spin variable has a value that is close to the mean magnetisation per spin m , or in other words that $\langle S_i \rangle = m$ for each spin i , and that the deviation $(S_i - \langle S_i \rangle)$ is a small quantity, so that the terms in Eq. 3.34 corresponding to products of two deviations can be neglected (this is an important point that I will return to shortly). In doing so, one obtains the mean-field approximation of the Hamiltonian,

$$H_{\text{Ising}} \approx H_{\text{Ising}}^{\text{MF}} = -J \sum_{\langle i,j \rangle} [(S_i + S_j) m - m^2] - H \sum_i S_i. \quad (3.35)$$

Defining the coordination number z to be the number of nearest neighbours each spin in the bulk of the lattice (i.e. away from the edges) possesses, Eq. 3.35 can be written as

$$H_{\text{Ising}}^{\text{MF}} = -(Jzm + H) \sum_i S_i + \frac{NJz}{2} m^2. \quad (3.36)$$

As advertised, Eq. 3.36 is a single-body Hamiltonian, which allows the partition function to be easily obtained, and thence the free energy. Omitting the algebraic details, the solution is

$$F^{\text{MF}} = \frac{NJzm^2}{2} - Nk_B T \ln \{2 \cosh [\beta (Jzm + H)]\}. \quad (3.37)$$

From here, the magnetisation, susceptibility and specific heat can be determined in the usual way. In particular, analysis of the magnetisation yields that it is 0 whenever $T < T_c^{\text{MF}}$ where $T_c^{\text{MF}} = Jz/k_B$, and non-zero otherwise, justifying its use as an order parameter. T_c^{MF} is therefore the critical temperature in the mean-field approximation. Furthermore, m is continuous, thereby confirming that this is a continuous phase transition.

The exact expression for the F in terms of Ψ can be rather complicated, meaning that in practice, one can often not solve exactly for Ψ or the other ground state quantities simply by minimising the free energy. However, the mechanism of Landau theory allows progress to be made. Here, one assumes that in the vicinity of a continuous phase transition, it is valid to expand F in a Taylor series about $\Psi = 0$. By doing so, one can obtain the leading behaviour of the thermodynamic quantities in the reduced temperature $t := (T_c - T)/T$. Since $\Psi = 0$ for all $t < 0$, there is no problem for negative t . One assumes that there exists a finite radius of convergence of the Taylor series $[0, \Psi_0]$ for $t > 0$, an assumption that will need to be examined further. Via this process, one can obtain the leading behaviour of the important thermodynamic quantities in t .

For the specific case of the Ising model, one finds by truncating the free energy Taylor series to fourth order in m and then minimising F , that

$$|m| = \begin{cases} 0, & t \leq 0; \\ \sqrt{3t(t+1)^3}, & t > 0. \end{cases} \quad (3.38)$$

To leading order in t , then, $m \sim t^{1/2}$ for $t > 0$. The exponent $1/2$ appearing in this relationship is called a critical exponent. The susceptibility and heat capacity near the phase transition can also be characterised by critical exponents: $\chi \sim |t|^{-1}$ and $C_V \sim \text{const.}$ These critical exponents are denoted by specific symbols in the literature:

$\Psi \sim t^\beta$ for the order parameter, $\chi \sim |t|^{-\gamma}$ for the susceptibility and $C_V \sim t^{-\alpha}$ for the heat capacity. In fact, it turns out that the mean-field description of any phase transition characterised by an order parameter produces the same critical exponents: $\alpha = 0$, $\beta = 1/2$ and $\gamma = 1$.

There is another important quantity that I have so far not discussed: the correlation length. For the Ising model, one can define the correlation function

$$g_{ij} = \langle (S_i - \langle S_i \rangle) (S_j - \langle S_j \rangle) \rangle, \quad (3.39)$$

which expresses the statistical correlation between the fluctuations of two separate spin variables about their means. In fact, the terms of the form $(S_i - \langle S_i \rangle) (S_j - \langle S_j \rangle)$ being averaged are exactly the ones that were neglected in the mean-field approximation of the Ising Hamiltonian, which implies that spin-spin correlations are completely neglected in the mean-field model. It can be shown that the contribution of these terms can only be safely neglected if

$$\frac{\sqrt{\langle S_i^2 \rangle - \langle S_i \rangle^2}}{\langle S_i \rangle} \equiv \frac{\Delta S_i}{\langle S_i \rangle} \ll 1, \quad (3.40)$$

for every choice of i .

The quantity ΔS_i , mathematically the standard deviation of S_i , is called a fluctuation, and the criterion above says that mean-field theory is only valid if the size of the spin fluctuations relative to the mean values is small; this is called the Ginzburg criterion. It turns out that the fluctuations are proportional to the susceptibility, which diverge like $|t|^{-1}$ in the vicinity of the critical temperature. This fact together with the observation that $\langle S_i \rangle \approx 0$ near the transition guarantees that there is some range of temperatures for which the Ginzburg criterion is violated. In a nutshell, correlations become very important when the system is nearly critical, and neglecting them produces completely

incorrect results. In fact, the mean-field critical exponents α , β and γ do not coincide with the values found by Onsager in his exact solution of the 2D Ising model with no external field. One must therefore be careful to ensure that the Ginzburg criterion is satisfied before taking mean-field results seriously. The importance of fluctuations tends to decrease with increasing dimensionality of the lattice, and in general there is a so-called upper critical dimension above which the mean-field results become exact. For the Ising model, this occurs for $d \geq 4$.

The behaviour of the correlation g_{ij} as a function of $|i - j|$ generally depends on the phase in which the system is to be found. As an example, in the ferromagnetic phase of the Ising model where $|m| \neq 0$, it is of the form

$$g_{ij} \propto e^{-|i-j|a/\xi}, \quad (3.41)$$

where ξ , called the correlation length, characterises the length scale over which spin-spin correlations are to be found and a is the lattice spacing. The correlation length ξ diverges as one approaches the phase transition, going like $|t|^{-\nu}$ where ν is another critical exponent. Landau theory as presented cannot be used to determine the correlation length; instead, an extension called Landau-Ginzburg theory must be used, in which the order parameter is allowed to have a spatial dependence.

Remarkably, the huge variety of quantum phase transitions in disparate systems that exists, based on very different physical mechanisms, can be partitioned into a much smaller number of equivalence classes, known as universality classes, based on the values of the critical exponents. This principle is known as universality, and it is a consequence of the fact that the enormous number of degrees of freedom of a complicated interacting system can be boiled down to a much smaller number of variables that are actually relevant to the dynamics of the system near the critical point, while certain microscopic

degrees of freedom become irrelevant. This in turn is due to the very fact that the correlation length diverges at the critical point, which means that there is no natural length scale for the system. The transitions that are correctly described by mean-field theory constitute a universality class, defined by the critical exponents $\{\alpha, \beta, \gamma, \nu\} = \{0, 1/2, 1, 1/2\}$.

Dimensionless thermodynamic properties of the system, such as for example the partition function, can be written in terms of ratios of the length scales in the system. For example, for the Ising model the length scales are the lattice spacing a , the length of the system L and the correlation length ξ . The partition function can be written in terms of the two independent ratios a/ξ and L/ξ , but when the correlation length diverges, the dependence on a/ξ in the thermodynamic limit can be removed. Thus, a microscopic parameter of the system (the lattice spacing) becomes irrelevant near criticality. The class of techniques dealing with these ideas is referred to as the renormalisation group in the literature. One important technique from this area is so-called scaling and hyperscaling relations which relate critical exponents with each other and with the dimension of the system d . Due to Rushbrooke's identity

$$\alpha + 2\beta + \gamma = 2 \tag{3.42}$$

and Widom's identity

$$2 - \alpha = \nu d, \tag{3.43}$$

(see, e.g. [107]), only two of the critical exponents $\{\alpha, \beta, \gamma, \nu\}$ are actually independent. There are more important critical exponents that I haven't discussed here, but there are also more unmentioned scaling laws. As a result, it appears that two independent critical exponents are sufficient to fully specify a universality class.

3.5.2 Quantum Phase Transitions

Quantum systems in d dimensions often bear a surprising and powerful relationship to analogous classical systems in $d + 1$ dimensions. This is useful because it allows the extensive and well-developed machinery of classical phase transitions to be applied to the quantum case. It is a remarkable fact certain quantum spin systems (such as the Ising and rotor models) on d -dimensional lattices can often be mapped exactly to the analogous classical systems in $d + 1$ dimensions, in the sense that their partition functions are identical if one maps one of the classical length dimensions to an imaginary time setting the temperature of the corresponding quantum model [108]. Naturally, when this formal and exact mapping exists, it implies that the d -dimensional quantum and $d + 1$ -dimensional classical transitions are in the same universality class.

Quantum phase transitions occur only in the thermodynamic limit, and only at temperature $T = 0$. Here, the free energy is just the ground state energy of the Hamiltonian, and there is no energy-entropy competition as in the classical case. Instead, the competition must be provided by two terms in the Hamiltonian that have competing influences (the example of the Bose-Hubbard model will be discussed in some detail shortly). As in the classical case, one identifies an order parameter, and the phase transition can be interpreted as being driven by fluctuations in the order parameter. However, in the quantum case the fluctuations are purely quantum ones associated with the uncertainty principle.

3.6 Basic low-temperature physics of Bose-Hubbard model in 3D

As first discussed in [3], the Bose-Hubbard model exhibits a quantum phase transition, arising from competition between the hopping t , which has a delocalising effect, and the onsite repulsion U that favours localisation. One can rewrite the Hamiltonian (3.28) in

units of the interaction U , so that it becomes dimensionless:

$$\frac{H_{\text{BH}}}{U} = -\frac{t}{U} \sum_{\langle i,j \rangle} (a_i^\dagger a_j + a_i a_j^\dagger) + \frac{1}{2} \sum_{i=1}^N n_i(n_i - 1) - \frac{\mu}{U} \sum_{i=1}^N n_i. \quad (3.44)$$

The prefactor t/U is a dimensionless parameter whose value can be tuned over a large range, for example in an optical lattice implementation by varying the lattice depth (see the famous experiment of Greiner et al. [100]); similarly, μ/U can in practice be tuned by varying the filling fraction of the lattice, but here it is considered to be a free parameter that controls the filling. As t/U is decreased from ∞ to 0 at fixed μ/U , the ground state changes from an ideal gas to a superfluid, and then a Mott insulator, and then finally to a system of isolated sites. The change from Mott insulator to superfluid is a continuous quantum phase transition occurring at a critical value of t/U , denoted $(t/U)_c$. A discussion of each of these regimes follows.

3.6.1 Ideal gas limit, $t/U \rightarrow \infty$

First, consider the non-interacting limit in which $|t/U| \rightarrow \infty$. In this limit, the system becomes an ideal (lattice) Bose gas and the interaction term drops out of the Hamiltonian completely, leading to the Hamiltonian

$$\frac{H_{\text{free}}}{U} = -\frac{t}{U} \sum_{\langle i,j \rangle} (a_i^\dagger a_j + a_i a_j^\dagger) - \frac{\mu}{U} \sum_{i=1}^N n_i. \quad (3.45)$$

The second term above, which involves the chemical potential, is nothing but the total number operator multiplied by a constant, so it commutes with the first term, which is manifestly number-conserving. Therefore, we can determine the eigenstates in the non-interacting limit by simply diagonalising the hopping part of the Hamiltonian; the chemical potential term will only shift the energies of these eigenstates, thereby poten-

tially changing the ordering of the states in terms of total energy.

In the case of periodic boundary conditions, we can exploit translational invariance to determine the eigenstates and spectrum, via a discrete Fourier transform of the creation and annihilation operators from position to momentum space. Suppose the lattice is a d -dimensional hypercube with lattice spacing a , with lattice sites at positions $\mathbf{x}$ defined by $\frac{\mathbf{x}}{a} = (x_1, x_2, \dots, x_d) \in \mathbb{Z}_L^d$. Let $\mathbf{k}$ be the quasimomentum corresponding to the Fourier space of the system (since $\hbar = 1$), written as $\frac{La}{2\pi}\mathbf{k} = (k_1, k_2, \dots, k_d) \in \mathbb{Z}_L^d$. Here, $\mathbf{x}$ and $\mathbf{k}$ have the usual units of length and inverse length respectively, but $\{x_i\}$ and $\{k_i\}$ are dimensionless. Then, the old and new annihilation operators can be interconverted via the expressions

$$\begin{aligned} c_{\mathbf{k}} &:= \frac{1}{L^{d/2}} \sum_{x_1, \dots, x_d=0}^{L-1} e^{i\mathbf{k} \cdot \mathbf{x}} a_{\mathbf{x}}, \\ a_{\mathbf{x}} &:= \frac{1}{L^{d/2}} \sum_{k_1, \dots, k_d=0}^{L-1} e^{-i\mathbf{k} \cdot \mathbf{x}} c_{\mathbf{k}}. \end{aligned} \quad (3.46)$$

The modes associated with the $\{c_{\mathbf{x}}\}$ operators have eigenenergies given by

$$\epsilon(\mathbf{k}) = -2t \sum_{i=1}^d \cos k_i \quad (3.47)$$

(this formula is called the dispersion relation), leading to the new hopping Hamiltonian

$$H_{\text{hop}} = \sum_{\mathbf{k}} \epsilon(\mathbf{k}) c_{\mathbf{k}}^{\dagger} c_{\mathbf{k}}. \quad (3.48)$$

In this situation, the single-particle eigenstates are d -dimensional plane waves with well-defined quasimomenta, and the ground state, corresponding to $\mathbf{k} = 0$, is perfectly uniform everywhere in real space.

By contrast, with open boundary conditions, the eigenstates for the d -dimensional

hypercubic lattice of length L are sinusoidal in nature, because of boundary effects. In 1D, the spectrum is obtained by rewriting the Hamiltonian in matrix form as

$$H_{\text{hop}} = \begin{bmatrix} a_1^\dagger & a_2^\dagger & \dots & a_{L-1}^\dagger & a_L^\dagger \end{bmatrix} \begin{bmatrix} 0 & -t & & & \\ -t & 0 & -t & & \\ & -t & \ddots & \ddots & \\ & & \dots & 0 & -t \\ & & & -t & 0 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_{L-1} \\ a_L \end{bmatrix} \\ \equiv \mathbf{a}^\dagger \cdot \mathbf{T} \cdot \mathbf{a}.$$

One then diagonalises the tridiagonal matrix $\mathbf{T}$, yielding $\mathbf{E} = \mathbf{U} \cdot \mathbf{T} \cdot \mathbf{U}^\dagger$ (where $\mathbf{E}$ is diagonal and $\mathbf{U}$ is unitary), and then writes

$$H_{\text{hop}} = (\mathbf{a}^\dagger \cdot \mathbf{U}^\dagger) (\mathbf{U} \cdot \mathbf{T} \cdot \mathbf{U}^\dagger) (\mathbf{U} \cdot \mathbf{a}) \\ \equiv \mathbf{c}^\dagger \cdot \mathbf{E} \cdot \mathbf{c} \\ = \sum_k \epsilon(k) c_k^\dagger c_k$$

where $\mathbf{c}$ and $\mathbf{c}^\dagger$ are vectors whose entries are the quasiparticle annihilation and creation operators respectively, and $\{\epsilon(k) := (\mathbf{E})_{kk}\}$ give the dispersion relation in this case. The matrix $\mathbf{T}$ is an example of a tridiagonal Toeplitz matrix, which have well-known eigenvalues and eigenvectors [110]. Performing this procedure gives the result

$$c_k = \sqrt{\frac{2}{L+1}} \sum_{x=0}^{L-1} \sin \frac{(k+1)\pi(x+1)}{L+1} a_x; \\ \epsilon(k) = 2t \left[1 - \cos \frac{(k+1)\pi}{L+1} \right], \quad (3.49)$$

where as before, $k, x \in \mathbb{Z}_L$ (I have chosen the origin in the dispersion relation to guarantee that the ground state has energy 0). Since in the general case $d > 1$ the hopping dynamics

along different dimensions are not coupled, the 1D result is easily generalised to

$$\begin{aligned}
c_{\mathbf{k}} &= \left(\frac{2}{L+1} \right)^{d/2} \sum_{x_1, \dots, x_d=0}^{L-1} \prod_{i=1}^d \sin \frac{(k_i + 1)\pi(x_i + 1)}{L+1} a_{\mathbf{x}}; \\
a_{\mathbf{x}} &= \left(\frac{2}{L+1} \right)^{d/2} \sum_{k_1, \dots, k_d=0}^{L-1} \prod_{i=1}^d \sin \frac{(k_i + 1)\pi(x_i + 1)}{L+1} c_{\mathbf{k}}; \\
\epsilon(\mathbf{k}) &= 2t(d - \sum_{i=1}^d \cos \frac{(k_i + 1)\pi}{L+1}).
\end{aligned} \tag{3.50}$$

In periodic boundary conditions, as well as for open boundary conditions in the bulk, the single-particle ground state orbital, corresponding to $\mathbf{k} = 0$, is spatially uniform. In the extreme non-interacting limit, the ground state for N bosons at temperature $T = 0$ comprises an N -fold occupation of this $\mathbf{k} = \mathbf{0}$ mode. Thus, in the non-interacting limit, the N -boson ground state of the BH model can be reasonably well represented by an ansatz corresponding to N -fold occupation of the $\mathbf{k} = 0$ state in Eq. (3.46) or (3.50). For example, for periodic boundary conditions, one obtains

$$|\psi_{\text{BEC}}\rangle \sim (c_0^\dagger)^N |0\rangle \tag{3.51}$$

$$= \left(\frac{1}{L^{d/2}} \sum_{\mathbf{x}} a_{\mathbf{x}}^\dagger \right)^N |0\rangle. \tag{3.52}$$

The non-interacting wavefunction above is featureless in real space, so it is reasonable to expect the same for the correlation functions $G^{(1)}(r)$ and $G^{(2)}(r)$, i.e. that they are r -independent constants, which is true. The reason I have called it $|\psi_{\text{BEC}}\rangle$ is because it is the state corresponding to an N -boson non-interacting Bose-Einstein condensate, which will be explained in the next section.

3.6.2 Bose-Einstein Condensation and Superfluidity, $t/U > (t/U)_c$

Now consider the case when interactions are turned on, but t/U is still greater than the critical value. In this case, the system exhibits a macroscopic quantum effect called superfluidity, in which a fluid is able to flow past obstacles without friction. The earliest example of a substance exhibiting superfluidity was ^4He , discovered independently in 1938 by Kapitza [111] and by Allen and Misener [112]. They found that when liquid ^4He was placed in a porous container and cooled below 4K, the liquid Helium actually flowed frictionlessly, right through the medium and dripped out through the bottom of container. They also found that superfluids could climb the walls of non-porous containers and spill over the brim. In describing superfluidity, it is useful to invoke concepts that are associated with another macroscopic quantum effect called Bose-Einstein condensation (BEC) [113, 114, 102]. I will therefore discuss BECs and superfluidity in parallel throughout this section.

The occupation fraction $f(\mathbf{k}) := N(\mathbf{k})/N$ for the states of a general bosonic system at finite temperature $T > 0$ with dispersion relation $\epsilon(\mathbf{k})$ is governed by the Bose-Einstein distribution [115],

$$f(\mathbf{k}) = \frac{1}{e^{(\epsilon(\mathbf{k}) - \mu)/k_B T} - 1}, \quad (3.53)$$

where k_B is the Boltzmann constant and μ is the chemical potential. Because of thermal fluctuations, the thermal population $N_{\text{th}} := \int_{\mathbf{k}=0}^{\infty} d\mathbf{k} N(\mathbf{k})$ is generally non-zero when $T > 0$; however, it can be shown in 3D both for the non-interacting Bose gas and the interacting gas (see, for example, [102]) that there is actually a finite critical temperature $T_{\text{BEC}} > 0$ such that $N_{\text{th}} = 0$, and the entire macroscopic population of the system resides in the ground state. This kind of state, actually a phase of matter, is called a Bose-Einstein condensate, or BEC. First predicted by Einstein in 1924-5 [113, 114], following up on the work of Bose [115], it was for decades believed by many to be purely of

theoretical interest, because of the low temperatures required to access the state. It was not until 1995 that real BECs were finally observed via two independent Nobel prize-winning experiments 1995, in vapours of ^{87}Rb [116] and ^{23}Na [117].

The fact that $T_{\text{BEC}} > 0$ in 3D is because of the behaviour of the density of states $\rho(E)$, which is defined by stipulating that $\rho(E)dE$ be the number of accessible quantum states with energies between E and $E + dE$. It is well-known that the density of states for the non-interacting Bose gas goes like $E^{d/2-1}$ in d dimensions [102], where the ground state is assumed to satisfy $E = 0$. This means that as $E \rightarrow 0$, the density of states vanishes like $E^{1/2}$ in 3D, but is constant in 2D and actually diverges in 1D as $E^{-1/2}$. This means that there are many more accessible states near the ground state in 2D and 1D than in 3D, which in turn means the thermal population fraction can be higher. In fact, it turns out that the critical temperature is $T_{\text{BEC}} = 0$ in 2D, and does not exist in 1D [102]. This means that no BEC phase transition is expected to occur in 2D or 1D.

In 1938, Fritz London suggested that BEC in an ideal gas was the underlying mechanism behind the superfluidity of ^4He [118]. The transition of ^4He from an ordinary fluid to a superfluid at the critical temperature $T_c = 4\text{K}$ is an example of a (classical) phase transition. Specifically, it is a continuous phase transition whose distinguishing characteristic is a discontinuity (in fact, a divergence) of the specific heat capacity at the transition temperature. Because the shape of the specific heat curve looks like the Greek letter λ , this is often called the lambda transition in the literature. London observed that a similar cusp (though not a divergence) was to be found in the specific heat capacity of an ideal Bose gas that had formed a BEC.

London's idea was followed up by the Soviet physicist Lev Landau in 1941 [119]. He imagined a Bose-condensed fluid with an obstacle flowing through it, and calculated the minimum velocity v_c required for the obstacle to be able to create an excitation. He used a Galilean transformation to calculate the energy of the ground state and the one

with a single lowest-lying excitation in the rest frame of the obstacle. He found that this lowest-lying excitation would have a phase velocity equal to the relative velocity: $v_c = \min_p \frac{\epsilon_p}{p}$, where p is the momentum of the excitation (see, for example, Sec. 10.1 of Ref. [102] for the details). The quantity v_c is called the Landau critical velocity. An important consequence of Landau's result is that, in fact, an ideal BEC is unable to support superfluidity, invalidating London's hypothesis. This is because the dispersion relation for the ideal gas of bosons with mass m is given by $\epsilon(p) = \frac{p^2}{2m}$, which means that $v_c = \min_p \frac{p}{2m} = 0$. However, the situation is rather different when interactions are allowed.

For a weakly-interacting dilute gas (more precisely, one where g and n satisfy Eq. 3.6.2), the dispersion relation is modified from the ideal case. Using the pseudopotential description of the interactions, $U(|\mathbf{r}_i - \mathbf{r}_j|) = g\delta(|\mathbf{r}_i - \mathbf{r}_j|)$, an equation for the wavefunction $\Psi(\mathbf{r}, t)$ that minimises the energy can be derived. It is a non-linear version of the Schrodinger equation, called the Gross-Pitaevskii equation [120, 121]:

$$i\hbar \frac{\partial}{\partial t} \Psi(\mathbf{r}, t) = \left(-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g |\Psi(\mathbf{r}, t)|^2 \right) \Psi(\mathbf{r}, t), \quad (3.54)$$

where $V(\mathbf{r})$ is an external potential (such as a lattice). In the dilute, weakly interacting case, it can be approximately solved by the method of Bogoliubov [122, 123]. It is a mean-field method in which the order parameter is taken to be the (presumed) macroscopically occupied BEC wave function $\Psi_0(\mathbf{r}, t) := \sqrt{\rho_c} e^{-i\mu t}$, where ρ_c is called the condensate density. Writing $\Psi = \Psi_0 + \delta\Psi$, where $\delta\Psi$ is a fluctuation, Bogoliubov derived the dispersion relation

$$\epsilon(p) = \left[\frac{gn}{m} p^2 + \left(\frac{p^2}{2m} \right)^2 \right]^{1/2}. \quad (3.55)$$

At high energies (or momenta), this spectrum reproduces the ideal gas behaviour, $\epsilon(p) \rightarrow p^2/2m$. On the other hand, for the low-lying excitations, the dispersion relation is linear:

$\epsilon(p) \rightarrow \sqrt{\frac{gn}{m}}p$. Because linear dispersion is characteristic of sound waves, the low-lying excitations are interpreted as sound waves with speed $s = \sqrt{\frac{gn}{m}}$. As in the ideal gas case, the Landau critical velocity is equal to the phase velocity of the lowest-lying excitation. Here, however, the phase velocity is

$$\frac{\epsilon(p)}{p} = s[1 + (p/2ms)^2]^{1/2}. \quad (3.56)$$

This expression is minimal when $p = 0$, leading to the conclusion $v_c = s$; the speed of sound is also the Landau critical velocity.

Without evaluating the magnitude of the fluctuations explicitly, the upper critical dimension for the MI-SF transition can be determined via some qualitative considerations. In the low-density case where the Bogoliubov approximation applies, the superfluid phase can be approximately regarded as an ideal BEC since the quantum depletion is so small and the condensate fraction so large. The ideal BEC wavefunction has a constant amplitude in space; however, an impurity or external potential such as a wall can force the amplitude of the wavefunction to zero. A property called the healing length ξ of the condensate quantifies the distance over which the amplitude of the wavefunction returns to its bulk value. The mean-field analysis is valid when the healing length is much greater than the mean interparticle spacing ℓ , i.e. $\xi/\ell \gg 1$, because this ensures that the order parameter is slowly-varying over the length scale of a single-particle wavefunction, in other words that its fluctuations are small.

The healing length can be determined by requiring that the mean interaction energy per particle is equal to the mean kinetic energy; when it is higher, the state is MI-like and when it is lower, SF-like. Define ξ to be the length scale characterising the momentum of the particles, $p := \hbar/\xi$. The kinetic energy per particle is $p^2/2m$, and the interaction energy is proportional to the total amount of overlap between the wavefunction of a

particle particle and those of all the others, and is given by the pseudopotential strength g times the density n . Equating the energies and solving for ξ give

$$\xi = \frac{\hbar}{\sqrt{2mgn}}. \quad (3.57)$$

In 3D, the interparticle spacing goes like $\ell \sim n^{-1/3}$. Thus, imposing $\xi/\ell > 1$ and using Eq. 3.57 gives

$$\frac{\hbar}{\sqrt{2mg}} n^{-1/6} > 1.$$

Thus, in 3D, the mean-field description of the transition works best at low densities.

While unlike for spin systems there is no classical analogue for the Bose-Hubbard model, it is nevertheless the case that the d -dimensional Bose-Hubbard MI-SF transition is in the universality class of the $d+1$ -dimensional classical XY-model [3]. The critical exponents and upper critical dimension are thus known. It turns out that the upper critical dimension is 3 [3], and thus the MI-SF transition critical exponents in dimensions greater than 3 have the mean field values given above. However, as described in Section 3.6.2, the mean-field limit of the transition in 3D occurs only for low density.

It should be noted that it is also possible for Fermi gases to exhibit superfluidity, in much the same way that electronic systems can exhibit superconductivity; in this case, the fermionic carriers pair up into composite bosonic quasiparticles called Cooper pairs, and it is the Cooper pairs that Bose-condense. A prominent example of superfluidity in a Fermi gas is to be found with ^3He [124, 125, 126], in which the critical temperature of 2.96mK is three orders of magnitude smaller than in ^4He (2.172K [111]).

In practice, for a BEC of N bosons, the fraction of particles $f_c = N_c/N$, where N_c is the number of Bose-condensed particles, to be found in the Bose-condensed phase is generally less than unity. At finite temperatures, thermal fluctuations lead to depopulation of the condensate, a phenomenon known as thermal depletion. Even at zero tempera-

ture, the presence of interactions mixes in components of atomic excitations ($|p| > 0$), resulting in a lower condensate population. Because it occurs at zero temperature and is purely a quantum process, this phenomenon is called quantum depletion. For the $T = 0$, dilute, interacting gas with the Bogoliubov spectrum, the occupation fraction of momentum mode p due to quantum depletion is

$$f_p^{(q)} = \frac{8}{3\sqrt{\pi}} (n |a_s|^3)^{1/2}, \quad (3.58)$$

where a_s is a quantity called the s-wave scattering length, which characterises the long-wavelength properties of the interparticle interaction. The validity criterion of the Bogoliubov approximation, Eq. 3.6.2, turns out to be equivalent to the condition that $n |a_s|^3 \ll 1$ for cold 3D gases. Thus, $f_p^{(q)}$ is small whenever the Bogoliubov approximation is valid. By contrast, liquid ^4He , has a quantum depletion is greater than 0.9, and therefore a condensate fraction not exceeding 0.1 [102], indicating that the Bogoliubov approximation breaks down in this case.

A similar concept to the condensate fraction f_c is the superfluid fraction f_s , defined as the proportion of particles that are part of the superfluid state. While superfluidity and Bose-Einstein condensation are closely related, it is not the case that f_c is always equal to f_s . In the case of cold, dilute gases as described above, the quantum depletion is small and virtually the entire system is both Bose-condensed and superfluid [102]. This suggests that in such situations, it may be reasonable to identify the superfluid component with the BEC. However, for quantum liquids such as ^4He , the assumptions of the previous sections break down, and one finds that below T_c , the superfluid fraction is close to 1, even though the condensate fraction never exceeds 0.1 [102]. Conversely, as described earlier, the ideal Bose gas at $T = 0$ has $f_c = 1$, but according to the Landau criterion does not support superfluidity, so $f_s = 0$.

It is generally necessary to determine f_c and f_s separately. Except for a few special circumstances, these quantities must be calculated numerically. A convenient way to find f_c is to use the fact, due to Penrose and Onsager [127, 128] that the single-particle density matrix $G^{(1)}$ of (3.71) is diagonal in the single-particle basis according to (3.70), and its eigenvalues are equal to the mean occupation numbers of the corresponding single-particle states. Thus, the signature of the existence of a BEC is that one of the eigenvalues of $G^{(1)}$ is macroscopic, and its value is the number of bosons N_c in the condensate; f_c is obtained from here as the ratio N_c/N .

What about the superfluid fraction? One method for determining f_s is obtained from the relationship between the superfluid velocity and the gradient of the phase of the condensate. Consider a superfluid wavefunction ψ_{SF} , which can be parametrised in terms of its magnitude and phase as

$$\psi_{\text{SF}} = \sqrt{\rho_s} e^{i\phi_s}. \quad (3.59)$$

The very fact that such a representation is possible for the wavefunction is because the phase of the superfluid is well-defined everywhere; this is not true for a generic state that is a superposition of different energy eigenstates, because the phases of each of these eigenstates evolves at a different rate. The wavefunction ψ_{SF} is the superfluid order parameter.

Now suppose the particles comprise a superfluid flowing with velocity $\vec{v}_s$. Conservation of mass is represented via the continuity equation,

$$\frac{\partial}{\partial t} n + \vec{\nabla} \cdot (n \vec{v}_s) = 0, \quad (3.60)$$

where as in first quantisation, the current density $\vec{j} = n\vec{v}_s$ is given by

$$\vec{j} = \frac{\hbar}{2mi} \left[\frac{\psi_{\text{SF}}^* \vec{\nabla} \psi_{\text{SF}} - \psi_{\text{SF}} \vec{\nabla} \psi_{\text{SF}}^*}{|\psi_{\text{SF}}|^2} \right]. \quad (3.61)$$

Using (3.59) in (3.61) straightforwardly gives

$$\vec{v}_s = \frac{\hbar}{m} \vec{\nabla} \phi. \quad (3.62)$$

Thus, the superfluid velocity is proportional to the gradient of the phase of the macroscopic wavefunction. Therefore, if a phase gradient is imposed on a general two-component Landau superfluid, the kinetic energy of the superfluid (and thus of the whole state) is boosted. On the other hand, the kinetic energy of the normal component is unchanged (roughly speaking, this is because the different momentum components of the normal fluid essentially have random phases, and the net response of the whole normal fluid to the phase gradient is nil).

Suppose that a linear phase change Θ is applied across one spatial direction of the sample of length L , so the phase gradient is Θ/L . The ground state energy $E(\Theta)$ of the system is increased from its original value $E(0)$ by an amount

$$E(\Theta) - E(0) = \frac{1}{2} N m f_s |\vec{v}_s|^2,$$

where $N m f_s$ is the total mass of the superfluid component [129], assuming that only the superfluid component responds to the phase change. Replacing $\vec{v}_s$ with the expression (3.62) and solving for f_s with the phase gradient Θ/L gives

$$f_s = \frac{2m}{\hbar^2} \frac{L^2}{N} \frac{E(\Theta) - E(0)}{\Theta^2}. \quad (3.63)$$

For the BH model, in which a lattice is present, the principle is the same, and the formula in terms of the model parameters is [130, 131, 132, 129]

$$f_s = \frac{L^2}{tN} \frac{E(\Theta) - E(0)}{\Theta^2}. \quad (3.64)$$

The actual phase gradient can be implemented in the BH model subject to periodic boundary conditions by modifying the creation and annihilation operators to include a Peierls phase factor, $a_i \mapsto a_i e^{i\Theta/L}$ [133], so that a particle that hops one site to the right picks up an incremental phase of Θ/L , thereby accumulating the correct total phase over one full traversal of the ring. In principle, if the incremental phase is large enough, it may be energetically favourable for a particle to be promoted to a higher-energy local state when it hops; in order to ensure that this does not happen, one must demand that $\Theta \ll \pi$.

Because the superfluid velocity can be written as the gradient of a phase, its curl must be zero. In other words, a superfluid is irrotational. In a uniform 2D system, there is no true BEC expected because the critical temperature is $T_{\text{BEC}} = 0$, which means that thermal fluctuations always overwhelm the BEC order. One might expect that this means there is also no superfluidity in a 2D system (or indeed a 1D system), at least in the weakly-interacting limit where the superfluid is expected to be a BEC. However, superfluid order actually can be established in 2D (at a temperature greater than $T = 0$) via a mechanism called the Berezinskii-Kosterlitz-Thouless (BKT) transition [134]. The thermal fluctuations that destroy the superfluid order at sufficiently high temperature in uniform 2D systems are of the form of vortices, which are quanta of angular momentum comprising a hole with fluid flowing around it; hence, the flow is not irrotational when vortices are present. However, at sufficiently low temperature, it can become energetically favourable for two counterrotating vortices to form a bound state of zero net angular

momentum. This vortex-antivortex pairing restores some semblance of superfluid order. However, it is not perfect; one-body correlations $G^{(1)}(r)$ do not approach a constant, but instead die off as a power law $r^{-\alpha}$; this is called quasi-off-diagonal-long-range order.

3.6.3 Atomic limit, $t/U = 0$

Having discussed the regimes in which hopping dominates the interactions, I will now discuss the regimes in which interactions dominate. In the so-called atomic limit, $t/U = 0$, the lattice sites are decoupled and the eigenstates are simply L^d -fold tensor products of the single-site ground states, which are determined by the chemical potential. In particular, the single-site spectrum is given by

$$\frac{\epsilon_n}{U} = \frac{1}{2}n(n-1) - \frac{\mu}{U}n, \quad (3.65)$$

where n labels the site occupation number, and therefore the energy difference between two consecutive eigenstates is

$$\frac{\epsilon_{n+1} - \epsilon_n}{U} = n - \frac{\mu}{U}. \quad (3.66)$$

This is a monotonically increasing function of n that is negative for $n < \mu/U$ and positive for $n > \mu/U$. Since n is a non-negative integer, this means that the single-site ground state occupation is n whenever $n-1 < \mu/U < n$. Thus, in the zero-hopping limit where the individual lattice sites are isolated, the global ground state for the entire Bose-Hubbard system on a d -dimensional hypercube with sides of length L sites and lattice spacing a is given by

$$|\psi_{\text{isol}}\rangle = \prod_{\mathbf{x}} (a_{\mathbf{x}}^\dagger)^n |0\rangle, \quad (3.67)$$

where $\frac{\mathbf{x}}{a} = (x_1, x_2, \dots, x_d) \in \mathbb{Z}_L^d$ is the Cartesian position vector of a site and $|0\rangle$ is the vacuum with respect to the $\{a_{\mathbf{x}}\}$ operators. Notice that $\mathbf{x}$ has units of length, but the quantities $\{x_i\}$ are dimensionless.

This state has exactly n bosons per lattice site, where n is a non-negative integer, and features an energy gap. In the case of repulsive interactions with $U > 0$, adding a single particle to the state incurs an energy penalty of

$$\frac{\mu_{n,p}^{(0)}}{U} = \frac{\epsilon_{n+1} - \epsilon_n}{U} = n - \frac{\mu}{U}, \quad (3.68)$$

while adding a hole (i.e. removing a particle) increases the energy by

$$\frac{\mu_{n,h}^{(0)}}{U} = \frac{\epsilon_{n-1} - \epsilon_n}{U} = \frac{\mu}{U} - (n - 1). \quad (3.69)$$

Importantly, these energy costs persist in the thermodynamic limit, in which the size of the system $L \rightarrow \infty$, while the ratio $n = N/L^d$ remains constant.

3.6.4 Mott insulator, $t/U < (t/U)_c$

The state $|\psi_{\text{isol}}\rangle$ is clearly an insulator, because the lattice sites are completely isolated from each other when $t = 0$. However, it turns out that even for small positive values of t/U (with repulsive interactions, $U > 0$), the ground state $|\psi_{\text{MI}}\rangle$ is still an insulator – specifically, a Mott Insulator (MI) [135, 136, 137]. Traditional crystalline insulators are unable to support currents because the valence band is completely filled and there is an energy gap separating the valence band from the conduction band. However, in a MI, the carriers are immobile strictly because of the repulsive interactions between them. MI states have been realised in cold atomic gases in the optical lattice setting for both bosons [100] and fermions [138]. In fact, they have been studied for much longer in fermionic systems; see Refs. [137, 139]. The MI retains the property of the atomic

limit ground state that every site is occupied by exactly n bosons, with n a non-negative integer. The essential reason the state is an insulator for sufficiently small values of t/U is simple; the kinetic energy gained by a boson hopping from one lattice site to the next is of the order t/U , but this process creates both a particle and a hole, incurring an energy penalty of approximately $\frac{\mu_{n,p}^{(0)} + \mu_{n,h}^{(0)}}{U}$ for $t \approx 0$. Such hopping processes are energetically unfavourable for very small t/U . On the other hand, as will be seen from the phase diagram, the particle and hole boundaries of the Mott lobe eventually meet, corresponding to the closing of the gap. At this point hopping is energetically favourable, leading to a non-insulating state (in fact, a superfluid).

To understand more about the physical properties of the state, it is useful to consider correlation functions. Consider first the single-particle correlation function, $G^{(1)}(i, j) := \langle a_i^\dagger a_j \rangle$. This is the position representation of the one-body density matrix (or more precisely, its i, j -th element), and is the analogue of the density matrices discussed in Ch. 1 for many-particle states. It can be used to define the correlation length ξ , discussed in Sec. 3.5. Suppose the system has L single-particle eigenstates $\{|\psi_i\rangle := \sum_j \psi_{ij}|j\rangle\}$, where the $\{|j\rangle := a_j^\dagger|\Phi\rangle\}$ are the site basis states (for example, Wannier functions). The density matrix is diagonal in the $\{|\psi_i\rangle\}$ basis, with the eigenvalues corresponding to the occupation numbers,

$$\rho := \sum_{m=1}^L N_m |\psi_m\rangle \langle \psi_m| \quad (3.70)$$

(note that here, $\text{Tr}(\rho) = N$). Now consider $G^{(1)}(i, j)$:

$$\begin{aligned} G^{(1)}(i, j) &= \sum_m N_m \langle \psi_m | a_i^\dagger a_j | \psi_m \rangle \\ &= \sum_m N_m \sum_k \langle k | \psi_{mk}^* a_i^\dagger a_j \sum_l \psi_{ml} | l \rangle \\ &= \sum_m N_m \sum_{k,l} \psi_{mk}^* \psi_{ml} \langle \Phi | a_k a_i^\dagger a_j a_l^\dagger | \Phi \rangle. \end{aligned}$$

Noting that $a_k|\Phi\rangle = \langle\Phi|a_k^\dagger = 0$ and taking advantage of the bosonic commutation relations, one obtains

$$\begin{aligned}
G^{(1)}(i, j) &= \sum_m N_m \sum_{k,l} \psi_{mk}^* \psi_{ml} \langle\Phi|(\delta_{k,i} - \cancel{a_i^\dagger a_k}) \overset{\text{annihilates vacuum}}{(\delta_{j,l} - \cancel{a_l^\dagger a_j})} |\Phi\rangle \\
&= \sum_m N_m \psi_{mi}^* \psi_{mj} \\
&= \langle i|\rho|j\rangle.
\end{aligned} \tag{3.71}$$

For the state $|\psi\rangle$ in Eq. (3.67), it is easy to show that

$$G^{(1)}(i, j) = \delta_{i,j}. \tag{3.72}$$

By construction, both $\langle\psi_{\text{isol}}|a_i^\dagger$ and $a_j|\psi_{\text{isol}}\rangle$ are Fock states, and they are identical if $i = j$, but different (and therefore orthogonal) if $i \neq j$. One would only expect the relationship (3.72) to hold exactly in the zero-hopping limit, since that is the only place where the ansatz for the ground state is exactly valid. For small but non-zero t/U , one would expect perturbative corrections to the ground state. To first order, the new ground state is

$$|\psi_{\text{MI}}\rangle \sim |\psi\rangle - \frac{t}{U} \sum_{\vec{n}} \frac{\langle\vec{n}|\sum_{\langle i,j\rangle} (a_i^\dagger a_j + \text{h.c.})|\psi\rangle}{\epsilon_{\vec{n}} - \epsilon} |\vec{n}\rangle,$$

where $\{|\vec{n}\rangle\}$ are the Fock basis states (which is the eigenbasis of the unperturbed Hamiltonian) with corresponding eigenenergies $\{\epsilon_{\vec{n}}\}$, and ϵ is the unperturbed ground state energy.

The non-zero matrix elements in the sum over $\vec{n}$ are those for which $|\vec{n}\rangle$ is the N -particle Mott insulator with an extra particle created at one site and a hole at a neighbouring site, and the corresponding energy cost of these excitations is $\epsilon_{\vec{n}} - \epsilon = \mu_p^{(t)} + \mu_h^{(t)} \in O(\Delta)$, where Δ is the gap. One can also include higher-order corrections; states in which

the particle and hole are $r = |i - j|$ sites apart appear as r^{th} -order corrections with amplitudes going like $O(\Delta^{-r})$. As a result, the true behaviour of $G^{(1)}(r)$ in the gapped MI regime is expected to be of the form

$$G^{(1)}(r) \sim \text{Exp}(-r/\xi), \quad (3.73)$$

where ξ is the correlation length, which goes like $\xi \sim \Delta^{-1}$. This echoes the behaviour of the single-body spin-spin correlation function in the gapped phases of magnetic systems like the quantum Ising or rotor models [108].

Another valuable diagnostic tool is the two-body correlation function $G^{(2)}(i, j)$, which is defined as

$$G^{(2)}(i, j) := \langle b_i^\dagger b_j^\dagger b_j b_i \rangle, \quad (3.74)$$

and its normalised equivalent, given by

$$g^{(2)}(i, j) := \frac{G^{(2)}(i, j)}{\langle b_i^\dagger b_i \rangle \langle b_j^\dagger b_j \rangle}. \quad (3.75)$$

$G^{(2)}(i, j)$ is easy to calculate for $|\psi_{\text{isol}}\rangle$. Using the bosonic commutation relations, one obtains

$$\begin{aligned} G^{(2)}(i \neq j) &= \langle \psi_{\text{isol}} | b_i^\dagger (b_i b_j^\dagger - \delta_{i,j}) b_j | \psi_{\text{isol}} \rangle \\ &= \langle \psi_{\text{isol}} | (b_i^\dagger b_i b_j^\dagger b_j - \delta_{i,j} b_i^\dagger b_j) | \psi_{\text{isol}} \rangle \\ &= n_i n_j - \delta_{i,j} n_i \\ &= n^2 - \delta_{i,j} n. \end{aligned}$$

Again, this is only exactly true in the $t/U = 0$ limit, but the generic behaviour of $G^{(2)}(i, j)$ deep in the Mott insulator regime is to be roughly constant except for a dip at $i = j$.

3.6.5 Mott Insulator-Superfluid Transition in 3D

In 3D, the ground state of the Bose-Hubbard model can be transformed from a (Mott) insulating state with a precise integer number of bosons on each site to a BEC-like superfluid state in which bosons are completely delocalised across the lattice, via a quantum phase transition. By increasing the ratio t/U from an initial value of 0 at a fixed value of μ/U , one can move continuously from the isolated-sites limit through a Mott-insulating phase and then, after passing a critical value of t/U , to a superfluid phase (and eventually the ideal Bose gas).

The MI-SF transition in the Bose-Hubbard model has been analysed analytically via at least two major but essentially equivalent approaches [103]. The first is a mean-analysis with the order parameter corresponding to the expectation value of the local annihilation operators, which for a uniform system is assumed to be site-independent: $\Psi := \langle b_i \rangle = \langle b_i^\dagger \rangle \equiv \langle b \rangle$ for all sites i . Because a MI state is a Fock state, $\Psi = 0$ everywhere in the MI regime; however, it is non-zero in the SF regime. The idea is to approximate the full Bose-Hubbard Hamiltonian by a sum of single-site ones, so that one has an effective isolated-sites problem. In the first approach, this is achieved by decoupling the hopping terms:

$$b_i^\dagger b_j \rightarrow \text{const.} + \langle b_i^\dagger \rangle b_j + b_i^\dagger \langle b_j \rangle + \dots, \quad (3.76)$$

where the ‘ $\dots$ ’ terms correspond to fluctuations in the $\{b_i\}$ that are neglected [3, 140].

This results in the single-site Hamiltonian

$$H_i^{\text{MF}} = \frac{U}{2} n_i(n_i - 1) - \mu n_i - \Psi (b_i + b_i^\dagger) + |\Psi|^2. \quad (3.77)$$

The criterion for validity is that the healing length exceed the interparticle spacing, which in 3D reduces to Eq. 3.6.2. The behaviour of the order parameter Ψ is determined by finding the ground state energy of H_i^{MF} (represented in a truncated basis with a maximum

occupancy of the site artificially imposed), and then minimising the energy with respect to Ψ .

The second approach [141, 142] is a variational method that is equivalent to the mean-field method when $T = 0$. The basic idea is to make a variational ansatz (called the Gutzwiller wavefunction for bosons) for the ground state that is a product of single-site wavefunctions, in terms of a single variational parameter $\mathbf{g}$, that is constructed from the ground state of the non-interacting problem by applying a product of single-site operators to it in such a way that the amplitudes of configurations with large potential energies are suppressed. One then finds that the expectation value of the Bose-Hubbard Hamiltonian per site is

$$\langle H_{\text{BH}} \rangle = -zt \langle b \rangle \langle b^\dagger \rangle - \mu \langle n \rangle + \frac{U}{2} \langle n(n-1) \rangle, \quad (3.78)$$

which can be variationally minimised in terms of $\mathbf{g}$ to find the ground state, which in turn allows $\langle b \rangle$ to be computed. The accuracy of the variational ansatz is a monotonically increasing function of the dimensionality d ; the method is exact in infinite dimensions and worst in one dimension.

In either case, for each value of μ/U , a critical value $(t/U)_c$ is found at which $\langle b \rangle$ becomes zero; this defines the location of the phase boundary between the MI and SF phases. The general structure of the mean-field phase diagram, depicted in Fig. 3.1 is of roughly semicircular lobes of Mott-insulator, with the rest of phase space being superfluid.

3.7 Bose-Hubbard model in 1D

The MI-SF transition in 1D is rather different from its 3D counterpart in a number of ways. The most important of them is that the mean field description of the transition




Figure 3.1: Generic mean-field phase diagram of Bose-Hubbard model, reproduced from [3]. Semicircular lobes of Mott-insulating regions with integer boson occupation per site are surrounded by a superfluid region. The tips of the lobes correspond to multicritical points; the transition through these points along a line of constant mean occupation per site for the d -dimensional system is in the universality class of the $d + 1$ -dimensional XY model. Here, J corresponds to t in the text of the thesis, and V corresponds to U .

from MI to SF works best for high densities, in direct contrast to the 3D case. This can be understood via an estimate of the ratio of the healing length to the mean interparticle spacing, as before. Eq. 3.57 is still valid, but now the mean interparticle spacing is $l = n^{-1}$, which leads to the criterion

$$\frac{\hbar}{\sqrt{2mg}} n^{1/2} > 1,$$

so mean field theory for the 1D transition works best at high densities. The nature of the low-density transition will be discussed further in the following sections.

3.7.1 BEC in 1D: Quasicondensation

It is well-known that a true BEC cannot form in 1D. The fundamental reason for this is the form of the density of states, which, in combination with the Bose-Einstein distribution, determines the thermal population of the excited states at any finite temperature. The density of states is a function $\rho(\epsilon)$ of the energy ϵ , and $\rho(\epsilon)d\epsilon$ is the number of quantum states with energy lying in the infinitesimal range between ϵ and $\epsilon + d\epsilon$. A discussion of this topic can be found in many sources (see for example [102]). The upshot is that the density of states goes to zero as $\epsilon \rightarrow 0^+$ in 3D, but it is independent of ϵ in 2D and actually diverges as $\epsilon \rightarrow 0^+$ in 1D. Therefore, the thermal population N_{th} is finite in 3D, but the relevant integral does not converge in 1D, and so there is no critical BEC temperature, and therefore no BEC.

What happens instead? The answer is a phenomenon known as condensate fragmentation [143]. Instead of all of the particles condensing into one single-particle state, the population is distributed between many different low-lying states. In fact, it is easy to see that in the $t \rightarrow 0^+$ limit, where the optical lattice sites are completely independent of each other, a lattice with L^d sites has L^d fully degenerate ground states, corresponding

to the harmonic-oscillator-like ground states of the individual lattice sites.

A detailed explanation of why fragmentation occurs in 1D but not in 3D is beyond the scope of this thesis, but a discussion can be found in [144]. The underlying reason is the nature of the interparticle interactions. When the interaction potential is expanded in the momentum basis, there are terms in which the two particles retain their initial momenta (the direct terms), and ones in which momentum is transferred from one particle to the other (the exchange terms). The exchange terms only contribute to the overall interaction energy if there is more than one single-particle state occupied (i.e. fragmentation exists). In the thermodynamic limit in 3D, this contribution is positive; however, in 1D, it is negative, which makes fragmentation energetically favourable. One consequence of this is the modification of the off-diagonal long range order associated with BECs to quasi-long-range order of the form $G^{(1)}(r) \sim r^{-\alpha}$ (although the use of the symbol α is conventional here, it is not to be confused with the critical exponent given the same symbol).

3.7.2 Luttinger Liquids

Because of the 1D geometry, the interacting bosons cannot pass each other. This means that any individual motion of a particle is quickly converted into a collective excitation, which looks like a density wave. This is in contrast to the 2D and 3D situations, in which the Fermi liquid model applies and single-particle excitations (actually quasiparticles) exist [145]. It therefore appears natural to express the Hamiltonian in a basis of collective excitations. This is conventionally accomplished via Luttinger liquid theory [146, 147], named in analogy to its 3D counterpart, Fermi liquid theory. It is an effective field theory that describes the asymptotic low-energy properties of any one-dimensional system of interacting bosons or fermions.

A detailed explication of Luttinger liquid theory can be found in [145]. For the purposes of Chapter 4, the salient points are as follows. Imagine the particles are labelled

1 to N going from left to right, which can be done unambiguously in 1D. One introduces a collective field operator $\phi_l(x)$ that takes the value $\phi_l(x_i) := 2\pi i$ at the position of the i th particle, and is chosen for convenience to be monotonically increasing. For a system with perfect crystalline order corresponding to lattice spacing a , it can be chosen to be the linearly increasing function $\phi_l^0(x) := 2\pi x/a$. One then defines the field $\phi(x)$, which expresses the displacement of the system from the perfect crystalline one, via

$$2\phi(x) = \frac{2\pi x}{a} - \phi_l(x). \quad (3.79)$$

One can then express the interacting Boson Hamiltonian as an infinite series in $\nabla\phi(x)$ and $\Pi(x)$, which is the canonically conjugate momentum to $\phi(x)$. For the low-energy properties, keeping only the lowest order terms in each operator results in the harmonic oscillator-like Hamiltonian

$$H_{\text{LL}} = \frac{\hbar}{2\pi} \int dx \left[\frac{uK}{\hbar^2} (\pi\Pi(x))^2 + \frac{u}{K} (\nabla\phi(x))^2 \right], \quad (3.80)$$

with the parameters u and K obtained from the system parameters via

$$uK = \frac{\pi\hbar}{ma}; \quad (3.81)$$

$$\frac{u}{K} = \frac{g}{\hbar\pi}. \quad (3.82)$$

Many of the universal properties of the system, such as the asymptotic decay of correlation functions, can be obtained from u and K . The most important fact for the purposes of Chapter 4 is the algebraic decay of the off-diagonal elements of the one-body density matrix, $G^{(1)}(r, s) \sim |r - s|^{-1/2K}$ [145]. This behaviour is consistent with the quasi-off-diagonal-long-range-order mentioned in Sec. 3.7.1, with $\alpha = 1/2K$.

3.7.3 Berezinskii-Kosterlitz-Thouless transition

The generic MI-SF transition features a change between commensurate and incommensurate density. However, for any Mott lobe corresponding to n particles per site, there is a line of constant density n that can be drawn from the tip of the lobe through the SF regime. Thus, a MI-SF transition corresponding to a trajectory in phase space that starts in a Mott lobe, passes through the tip and then follows this constant density contour, obviously occurs at constant density. It is thus of a different nature from the generic transition. In fact, it turns out it is in the universality class of the Berezinskii-Kosterlitz-Thouless transition [134, 3] associated with vortex-antivortex pairing that I discussed in Sec. 3.6.2. Originally this transition was found in the 2D XY-model, which is a generalisation of the spin-1/2 Ising model where the spins are allowed to point in any direction in the x - y plane, rather than just up or down. Because the superfluid-to-Mott-insulator transition in the d -dimensional Bose-Hubbard model is in the same universality class as the $d+1$ -dimensional XY model, it too exhibits a transition of the BKT type when $d = 1$.

3.7.4 Tonks-Girardeau gas

As explained in the introduction to the present section, mean field theory breaks down in 1D at low densities, and the system becomes less BEC-like. A common quantity used to characterise the transition from BEC-like behaviour to something else is the ratio of the typical interaction energy to kinetic energy of a particle. The interaction energy is determined by the amount of overlap of the individual wavefunctions. Recall that in the BEC limit, the characteristic size ℓ of these wavefunctions increases. For a 1D system of length L with N bosons, the linear density is $n = N/L$, and the mean separation between particles is its inverse $1/n$. The interaction energy is significant when there is large overlap between single-particle wavefunctions, which occurs when $\ell \gg 1/n$. By contrast, when $\ell \ll 1/n$, it is the kinetic energy of the individual particles that dominates.

The transition between BEC-like behaviour and “other” behaviour can be said to occur when $\ell \sim 1/n$.

The previous paragraph can be made more quantitative. In the BEC-limit, the mean kinetic energy E_{kin} and interaction energy E_{int} per particle can be reliably calculated via the Bogoliubov mean-field theory. The characteristic momentum of the condensate will be given by $k \sim 1/\ell$, resulting in

$$E_{\text{kin}} = \frac{\hbar^2 k^2}{2m} \sim \frac{\hbar^2}{2m\ell^2} = \frac{\hbar^2 n^2}{2m}. \quad (3.83)$$

The interaction energy for a single particle is proportional to the total overlap of its wavefunction with that of the other particles in the system, and is given by

$$E_{\text{int}} = gn. \quad (3.84)$$

The ratio γ of these two energies is therefore

$$\gamma \equiv \frac{E_{\text{int}}}{E_{\text{kin}}} \sim \frac{2mg}{\hbar^2 n}, \quad (3.85)$$

the most important feature of which is that $\gamma \rightarrow 0$ in the BEC limit, and $\gamma \gg 1$ in the “other” limit.

With reference to the phase-diagram of the Bose-Hubbard model in the (μ, t) -plane, the lowest non-zero density outside occurs for small t/U and μ/U , since increasing either t or μ generally increases the density, but outside any of the Mott-insulating lobes, because there is no MI phase with density between 0 and 1 carrier per site. This, then, is the location of the “other” limit in phase space.

The obvious question arising from this discussion is, what is the nature of this $\gamma \gg 1$ limit? Because the (repulsive) interactions are so important in this limit, the dynamics

are dominated by the particles trying to avoid each other. In other words, each particle develops an exclusion zone, or hard core. In this limit, a celebrated mapping theorem developed by Tonks [148] and Girardeau [149] allows all of the eigenstates of the hard core Bose system in 1D to be exactly mapped to an eigenstate of a 1D system of free fermions, with the exclusion hole occurring naturally as a result of the Pauli exclusion principle.

3.8 Jaynes-Cummings model

While the optical lattice setting offers a relatively clean implementation of the Bose-Hubbard model, it suffers from some limitations. Most notable among them is the very small spacing between lattice sites, typically on the order of hundreds of nanometres to microns, which makes single-site addressability difficult. Additionally, it is of intrinsic interest to be able to manufacture quantum states of strongly-interacting photons. In recent years, much progress has been made on this front using coupled arrays of optical cavities. The intercavity spacing in such systems is typically dozens of microns, allowing for single-cavity addressability with optical frequencies [150].

Optical cavities are resonators that are constructed from a pair of mirrors. Any light between the two mirrors is reflected many times, and generally remains confined within the cavity (in practice, light does leak from a cavity at a rate that is specified by a number known as the quality factor). When the intensity of the field trapped within the cavity is low (i.e. a few photons), it must be treated quantum mechanically. The nature of the modes supported by the cavity depends on the precise construction of the device, but a common scenario is for the modes to have a Gaussian transverse intensity profile, and a longitudinal profile that is a Hermite or Laguerre polynomial (see, for example, Ch. 16 of Ref. [151]). For this thesis, cavities are assumed to support only a single optical mode.

The standard way of describing the interaction of a single mode of the cavity field with the atom is via the Jaynes-Cummings model [152], but before writing down the model I will briefly describe the Hilbert space of the system.

For the cavity mode, it is most convenient to use second-quantisation, since the mode can be occupied by a variable number of indistinguishable bosons (photons). Let a and $a^\dagger$ be the creation and annihilation operators for this mode, and let $n = a^\dagger a$ be the corresponding number operator. The basis for the cavity field is then the Fock basis $\{|n\rangle\}$, which are the eigenstates of n with eigenvalues $\{n\}$. On the other hand, since the atom is a distinguishable two-level system, it is perfectly reasonable to represent its state within first quantisation. One common method is the pseudospin representation, in which the electronic states of the atom are identified with the spin states of a spin-1/2 particle, with spin raising and lowering operators $\sigma_\pm$ written in terms of the Pauli operators from Chapter 1 as

$$\sigma_\pm = \frac{1}{2} (\sigma_x \pm i\sigma_y). \quad (3.86)$$

Identifying the electronic ground state of the atom $|g\rangle$ with the spin-down state of the pseudospin, and the excited state $|e\rangle$ with spin-up, the Hilbert space of the cavity-atom system $\mathcal{H} = \mathcal{H}_{\text{field}} \otimes \mathcal{H}_{\text{atom}}$ is the Kronecker product of the Fock space for the field with the pseudospin space for the atom. The basis states for this space are $\{|n, s\rangle\}$, where $s \in \{g, e\}$ labels the electronic (pseudospin) state.

I will now sketch the derivation of the Jaynes-Cummings model using the photon and pseudospin operators. First consider the terms describing the cavity field and the qubit separately. In complete analogy to the Bose-Hubbard term $\sum_i \hbar\omega_0 n_i/2$ from Eq. (3.27), the term accounting for the energy of the field is proportional to the cavity mode occu-

pation number (i.e. the number of photons present in the cavity),

$$H_{\text{field}} = \hbar\omega_c \left(n + \frac{1}{2} \right). \quad (3.87)$$

Similarly, the term giving the energy of the atom (pseudospin), with transition frequency ω_a , is proportional to the “spin occupation number” $\sigma_+\sigma_-$,

$$\begin{aligned} H_{\text{atom}} &= \hbar\omega_a \sigma_+ \sigma_- \\ &= \frac{\hbar\omega_a}{2} (1 + \sigma_z). \end{aligned}$$

The cavity and the two-level atom interact via coupling of the cavity electric field $E = a + a^\dagger$ with the atomic polarisation vector $S = \sigma_+ + \sigma_-$:

$$\begin{aligned} H_{\text{int}} &\propto gES \\ &= g(a\sigma_+ + a^\dagger\sigma_+ + a\sigma_- + a^\dagger\sigma_-) \end{aligned} \quad (3.88)$$

(note that the simple expressions for E and S used here are only valid for a single cavity mode). The Jaynes-Cummings model is then given by

$$H = H_0 + H_{\text{int}}, \quad (3.89)$$

with

$$H_0 = H_{\text{field}} + H_{\text{atom}}. \quad (3.90)$$

In this form it is unsolved, mainly because it has no obvious symmetry that can be exploited to simplify the problem. However, the rotating wave approximation introduces just such a symmetry.

In the rotating wave approximation, one switches from the Schrödinger picture to

the interaction picture, defined by the time-evolution operator $U(t) = e^{-iH_0 t/\hbar}$. Expanding the time-evolution operator to first order in t , one obtains the interaction picture expressions for the photon and pseudospin operators:

$$\begin{aligned}
a_I(t) &= U(t)aU^\dagger(t) \\
&= e^{-i\omega_c a^\dagger a t} a e^{i\omega_c a^\dagger a t/\hbar} \\
&\approx (1 - i\omega_c a^\dagger a t) a (1 + i\omega_c a^\dagger a t) \\
&\approx a e^{i\omega_c t}; \\
a_I^\dagger(t) &\approx a^\dagger e^{-i\omega_c t}; \\
\sigma_{-,I}(t) &\approx \sigma_- e^{i\omega_a t}; \\
\sigma_{+,I}(t) &\approx \sigma_+ e^{-i\omega_a t}.
\end{aligned}$$

Putting these expressions back into Eq. (3.88), one obtains the interaction term in the interaction picture,

$$\begin{aligned}
H_{\text{int},I} &= g(a_I \sigma_{+,I} + a_I^\dagger \sigma_{+,I} + a \sigma_{-,I} + a_I^\dagger \sigma_{-,I}) \\
&= g(a_I \sigma_{+,I} + a_I^\dagger \sigma_{+,I} + a \sigma_{-,I} + a_I^\dagger \sigma_{-,I}) \\
&= g(a \sigma_+ e^{i(\omega_c - \omega_a)t} + a^\dagger \sigma_+ e^{-i(\omega_c + \omega_a)t} + a \sigma_- e^{i(\omega_c + \omega_a)t} + a^\dagger \sigma_- e^{-i(\omega_c - \omega_a)t}).
\end{aligned}$$

All of these operators are oscillatory in time. Assuming that $|\omega_c + \omega_a| \gg |\omega_c - \omega_a|$, the terms proportional to $a^\dagger \sigma_+$ and $a \sigma_-$ oscillate much more rapidly than the other two terms, and can thus be neglected. Doing this, and then returning to the Schrödinger picture, yields

$$H_{\text{int}} \approx g(a \sigma_+ + a^\dagger \sigma_-). \quad (3.91)$$

Putting all of the pieces together gives the full Jaynes-Cummings model in the rotating

wave approximation and the pseudospin representation:

$$H_{\text{JC}} = \hbar\omega_c \left(n + \frac{1}{2} \right) + \frac{\hbar\omega_a}{2} (1 + \sigma_z) + g(a\sigma_+ + a^\dagger\sigma_-). \quad (3.92)$$

Notice that because of the atom-photon interaction terms, the Hamiltonian (3.92) conserves neither the photon number n nor the excitation state of the atom, which corresponds to $(1 + \sigma_z)/2$ in the pseudospin representation. Instead, it can be exactly diagonalised in a dressed-state basis called the polariton basis. The eigenstates and eigenvalues are found in Eqs. 4.3–4.6 (note that in Chapter 4, the choice of units $\hbar = 1$ is adopted). As pointed out in Chapter 4, the structure of the polariton energy spectrum is a ladder, with energy levels $E_{N,\pm}$ labelled by the photon number N and adjacent levels separated by the energy of a single cavity-mode photon, $\hbar\omega_c$. However, the ladder is anharmonic because each of these levels is further split, due to the atom-photon interactions, into two sublevels labelled ‘+’ and ‘−,’ corresponding to two different kinds of polaritons with the energy of ‘−’ always lower. Because of this fine structure, the difference in energy (in the case of zero detuning, $\delta \equiv \omega_c - \omega_a = 0$, between the ‘−’ species in the N and $N - 1$ manifolds is equal to $\hbar\omega_c - g(\sqrt{N} - \sqrt{N - 1})$. This interaction-induced correction to the harmonic oscillator ladder of the cavity field results in an extra energy cost for multiple photon occupancy of a single cavity, a phenomenon known as photon blockade [153].

3.9 Jaynes-Cummings-Hubbard model

The Mott Insulator-superfluid transition in the BH model is a direct consequence of the competition between the kinetic energy (hopping) and the onsite repulsion between bosons (as well as density-related effects associated with the chemical potential in the grand canonical ensemble). In order to simulate this transition with cavity QED, one must mimic these effects in the new system.




Figure 3.2: Schematic of Jaynes-Cummings-Hubbard model on a hexagonal lattice, implemented in a photonic bandgap material (reproduced from [4].) The red spheres represent two-level atoms, and the holes in the material are the cavities. λ is the wavelength of the resonant mode of the cavity and κ characterises the intercavity photon tunnelling. η is the refractive index of the material. The symbols ϵ and ω correspond to $\hbar\omega_a$ and $\hbar\omega_c$ respectively in the text of the thesis.

The fundamental excitations in the Jaynes-Cummings model are polaritons, as discussed earlier. Polaritons located in different cavities can be made to interact if the photons interact, which can be arranged via the photon blockade effect discussed in the previous section. The chemical potential term is difficult to arrange physically, since photons naturally have zero chemical potential. Nevertheless, the role of the chemical potential formally is to provide a knob by which to control the boson density, which in the case of the Jaynes-Cummings model can be controlled by varying the intensity of the pump laser used to populate the cavity mode. Thus, the remaining term to be considered is the hopping term.

In the homogenous Bose-Hubbard model, the hopping is characterised by the tunnelling coefficient t , whose non-zero value arises from the overlap of the exponentially decaying tails of the single-particle wavefunctions in the finite potential wells corresponding to individual lattice sites. In the cavity QED system, the hopping mechanism can

be engineered in similar fashion. Optical cavities (in fact, all optical waveguides) have exponentially decaying tails outside their physical boundaries as well, which are known as evanescent modes. By placing two cavities sufficiently close together that their evanescent modes have non-negligible overlap, photons can be made to tunnel between the cavities, at a rate characterised by a parameter that is the equivalent of t in the Bose-Hubbard model, and is usually called κ in the literature. See, for example, Ref. [154] for details, but note that in this paper the coupling parameter is called α .

Thus, by arranging L optical cavities in a suitable geometry (for example, in the photonic bandgap implementation shown in Fig. 3.2), one can create a network of coupled cavities implementing the Hamiltonian

$$H_{\text{JCH}} := \sum_{i=1}^L H_{\text{JC},i} - \kappa \sum_{\langle i,j \rangle} \left(a_i^\dagger a_j + a_i a_j^\dagger \right) + \mu \sum_{i=1}^L a_i^\dagger a_i, \quad (3.93)$$

where $H_{\text{JC},i}$ is the Jaynes-Cummings Hamiltonian (3.92) defined on cavity i , and a_i is the annihilation operator for the single-photon mode of cavity i . The Hamiltonian (3.93) is called the Jaynes-Cummings-Hubbard (JCH) model, and has been the focus of intense study in recent years. The key results for the purposes of Ch. 4 is the prediction of a MI-SF transition of polaritons via a mean-field analysis in this model [154], supplemented by numerical evidence for the transition in 1D via analysis of the energy gap [6].

3.10 Simulating 1D Systems: Density Matrix Renormalization Group

The industry-standard method of choice for calculating static (and some dynamic) properties of the low-lying eigenstates of one-dimensional quantum models is the density-matrix renormalisation group (DMRG) algorithm [155, 156, 7, 157], first developed by Steven White in 1992. It is popular mainly because it provides spectacular precision for a wide variety of calculations in a vast array of systems, and it is immune from the

fermion sign problem that plagues other methods such as quantum Monte Carlo simulations. This Section provides a brief description of the uses and underlying ideas of DMRG, but the details of the relevant algorithms are deferred until Ch. 5. This is because for the results of Ch. 4, I used an implementation of DMRG available from the ALPS (Algorithms and Libraries for Physics Simulations) [158, 159]. Apart from the cited references, information about this project can be found on the project website at alps.comp-phys.org.

Consider a 1D lattice of L d -level quantum systems; this could, for example, be a quantum Ising chain of spin- $(d-1)/2$ particles, or the single-band Bose-Hubbard model with the maximum occupancy of a site artificially set to be $d-1$. The dimension of the Hilbert space of the system is d^L , an exponential scaling in the number of particles. Consequently, exact-diagonalisation approaches to solving the system are practically limited to a small number of particles, and do not extend directly to the thermodynamic limit. However, there are compelling arguments that most of Hilbert space is unnecessary for describing physically accessible states; as Verstraete et al. put it in the context of spin systems, “... all physical states live on a tiny submanifold of Hilbert space” (p. 146 of [42]). This is the central insight behind the success of DMRG; the algorithm provides a systematic method for truncating the Hilbert space in which the systems are numerically represented to a manageable size with minimal loss in the fidelity of the ground state as compared to the true ground state.

The idea behind DMRG is simple: to exactly solve the system for a small number of particles, and then use the solution as a model for the “environment” in which one can embed new particles being acted upon by the same Hamiltonian, slowly growing the system until the final system size is reached. In the numerical renormalisation group, the truncated basis for the environment is chosen based on the lowest-lying eigenstates of the Hamiltonian acting on appropriate subsystems of the final system, an approach

that works most famously for solving the Kondo and other impurity problems [160, 161], but fails for many other systems, usually because the ground states produced by the algorithm had unphysical cusps. DMRG, on the other hand, truncates the basis for the environment by selecting the maximal-weight eigenstates of appropriate reduced density matrices of the ground state.

Although there have been many variations developed in the past twenty years, the algorithm comes in two basic flavours. The first is the infinite-system DMRG algorithm, which is designed for finding ground state expectation values of few-body operators for infinite 1D systems (i.e. thermodynamic limit results). This is achieved by expressing all of the pertinent local operators for the system, as well as the ground state, in a truncated basis constructed by retaining only the states most relevant for accurately capturing the physics of the ground state and discarding the rest. One starts with a 1D lattice comprising L sites, each with Hilbert space dimension d , that is small enough for the ground state $\rho = |\psi\rangle\langle\psi|$ to be numerically accessible. One then cuts the lattice in half and finds the eigenbases for the reduced density matrices ρ_L and ρ_R of the left and right halves of the chain (called the Block and Universe respectively). Then, the system is grown by placing two new sites in the centre of the chain, diagonalising the Hamiltonian on the new lattice and finding the new reduced density matrices for the half-chains.

In order to avoid the dimension of the numerical representation of the Hamiltonian increasing, the Hilbert space for the joint Block-left site portion of the chain is projected onto the subspace spanned by the χ eigenvectors ρ_L with the largest eigenvalues. This process is usually called the block decimation procedure. Similarly, the Hilbert space for the right site-Universe system is projected onto the equivalent subspace of ρ_R . This makes the Hilbert space of both the Block and the Universe χ -dimensional on every iteration, and the entire system be $\chi^2 d^2$ -dimensional always. This procedure is continued until a fixed point is reached, in the sense that the physical properties of the truncated

representation of the state are close to the same as those from the previous iteration; this is why the phrase “renormalisation group” appears in the name of the algorithm.

Because the infinite-system algorithm usually does not converge to a satisfactory fixed point, another version called the finite-system algorithm was developed. In this version, one uses the infinite-system algorithm to grow the system to a predetermined finite length, and stores all of the system representations obtained at each intermediate length. One then keeps the system size fixed but grows the Block one site at a time and simultaneously shrinks the Universe by one site, until the end of the finite chain is reached, storing the representations of the system for each Block and Universe size along the way. Then one goes the other way, on each iteration incrementing the Universe size and decrementing the Block size by one. This process is called sweeping; one sweeps left and right on the finite system, over and over again until convergence of the desired physical properties is obtained. This algorithm is capable of producing remarkably accurate results for large but finite systems, and allows for thermodynamic-limit results to be obtained via finite-size scaling.

There is a fascinating connection between DMRG and the MPS formalism described in Sec. 1.5.1. In fact, both the infinite-system and finite-system algorithms have been shown to produce ground states that are MPS in character [162, 163]. Subsequently, DMRG has been explicitly rephrased as a variational method over MPS [164, 165]. Ref. [157] is an excellent and detailed review of the MPS view of DMRG, including extensions to the original algorithm that have been made possible as a result of this view. This connection will form the core of the discussion in Ch. 5.

Being such a versatile method for calculating properties of 1D systems, DMRG has found a plethora of exciting applications. A spectacular early example is from the work of White and Huse who, in 1993, applied the method to the $S = 1$ isotropic antiferromagnetic Heisenberg chain with nearest-neighbour coupling constant J , obtaining the remark-




Figure 3.3: Phase diagrams for MI-SF transition for (a) bosons in the 1D BH model and (b) polaritons in the 1D JCH model, via infinite system (a) and finite system (b) DMRG. Reproduced from Refs. [5] and [6] respectively.

ably precise values of $E_0 = -1.401484038971(4)J$ for the ground state energy (per site), and $\Delta = 0.41050(2)J$ for the gap [166]. Much more recently (owing to advances allowing the study of higher-dimensional lattices), DMRG has found applications [167, 168, 169] in attempts to explain the existence of an exotic magnetic phase called a quantum spin liquid, strongly conjectured to be the ground state of the crystalline lattice of the mineral herbertsmithite [170]. It is virtually impossible to provide an exhaustive list of successful applications of DMRG; several can be found in the reviews [7, 171, 157].

The DMRG algorithm has been used to determine the location of the phase boundary between the gapped (MI) and gapless (SF-like) phases for both the BH (infinite system in [5], finite system in [172]) and the JCH models [6] in 1D, via analysis of the closing of the energy gap. The phase diagrams thus obtained are reproduced in Fig. 3.3. The sharp cusps exhibited by the phase boundary as opposed to the smooth lobes of the generic mean-field diagram are peculiar to the 1D case.

3.11 Preamble to Chapter 4

Chapter 4 is a reproduction of a journal article that has been published in Physical Review A. The article is © American Physical Society (APS), 2013. The full citation to the article can be found in the Bibliography; it is Ref. [173]. The content is identical in every respect to the manuscript, except that the formatting has been changed to comply with University of Calgary thesis requirements rather than those of the journal, and the bibliography has been stripped from the manuscript, with the references incorporated into the bibliography for the thesis as a whole. According to APS guidelines, it is acceptable to include an APS-published manuscript in a thesis without receiving explicit permission in writing; see <http://publish.aps.org/copyrightFAQ.html#thesis>. A preprint of the article has been placed on the arXiv (arxiv.org), with the article identifier `arXiv:1311.2348`.

I am the first author of this article, and my contribution was to perform all of the primary research (comprising all of the numerical simulations and analysis of the data) whose results appear in the article, as well as write and edit a complete draft of the paper. My co-authors, Dr. Barry C. Sanders and my supervisor Dr. David L. Feder, provided scientific input regarding the direction of the research to be performed and the goals of the project. Additionally, Dr. Feder was heavily involved in editing the manuscript prior to submission to the journal upon receipt of a full draft from me, in particular with regards to interpreting the numerical results.

Chapter 4

Fermionized photons in the ground state of one-dimensional coupled cavities

4.1 Abstract

The Density Matrix Renormalization Group algorithm is used to characterize the ground states of one-dimensional coupled cavities in the regime of low photon densities. Numerical results for atom and photon densities, one- and two-body correlation functions, superfluid and condensate fractions, as well as the entanglement entropy and localizable entanglement are obtained for the Jaynes-Cummings-Hubbard (JCH) model, and are compared with those for the Bose-Hubbard (BH) model where applicable. The results indicate that a Tonks-Girardeau phase, in which the photons are strongly fermionized, appears between the Mott-insulating and superfluid phases as a function of the inter-cavity coupling. In fact, the superfluid density is found to be zero in a wide region outside the Mott-insulator phase boundary. The presence of two different species of excitation (atom and photon) in the JCH model gives rise to properties with no analog in the BH model, such as the (quasi)condensation of spin excitations and the spontaneous generation of entanglement between the atoms confined to each cavity.

4.2 Introduction

The idea of simulating a complex, many-body physical system with a simpler model system has a long and rich history. A prominent example is the Bose-Hubbard (BH) model [3], which governs the dynamics of bosons tunneling between sites of a lattice with

energy t and interacting on a given site with energy U . Originally proposed to describe the behavior of superfluid ^4He in porous Vycor glass, it is now applied to a plethora of experimental systems including Josephson junction arrays [174], ultracold atoms in optical lattices [175], photonic crystals [176], and arrays of coupled cavities [154, 4, 177, 178].

For repulsive interactions ($U > 0$), there is a competition between delocalization due to the tunneling and the tendency to localize due to the energy cost of multiple occupancy of a given site. As a consequence, the model exhibits a quantum (zero-temperature) phase transition [3]. In the strongly interacting (weak tunneling) regime $t/U \ll 1$, the ground state is predicted to be a Mott insulator (MI) [135, 136, 137], in which each site is occupied by an (identical) integer number of bosons; for $t/U \gg 1$, a superfluid (SF) state results [111, 112, 119], in which the bosons become completely delocalized. The transition between these phases was first realized in Bose-Einstein condensates confined in three-dimensional optical lattice potentials [100], where the ratio t/U was controlled by varying the well depth, and subsequently observed in many other cold-atom optical lattice experiments in one, two and three dimensions [179, 180, 181, 182, 183, 184].

There has been much recent interest in observing similar quantum phase transitions in cavity quantum electrodynamics [154, 4, 177]. The principal motivation for employing these environments is the robustness of available technology for producing, manipulating and detecting photons. Unfortunately, photons do not intrinsically interact. Various strategies have been proposed to overcome this limitation, and a vast array of interesting many-body states have been conjectured to appear as a result [185, 6, 186, 187, 188, 189, 190]. One model that has attracted particular attention is the Jaynes-Cummings-Hubbard (JCH) model [154, 191, 192], comprising a lattice of high-finesse optical cavities each containing one or more two-level atoms, with neighboring cavities coupled via the overlap of their evanescent modes. The (bosonic) photons can then be considered to

‘tunnel’ from cavity to cavity. The atom interacts with the quantized electromagnetic field present within the cavity according to the Jaynes-Cummings model [152], and photon interactions are generated via the photon-blockade mechanism [153]. Importantly, the lattice polaritons that constitute the spin-photon elementary excitations of the JCH model are also expected to undergo a phase transition from a Mott insulator to a superfluid [4, 154, 177, 6, 193, 194, 195, 196, 197, 198, 199]. While early proposals for the experimental realization of coupled cavities involved nitrogen vacancies in diamond [4], self-assembled quantum dots in photonic crystals [200], and trapped ions [201, 202], more recent proposals favor circuit QED [203, 204, 178, 205].

The close similarity between the properties of the BH and JCH models suggests that the polariton superfluid resulting from arranging cavities along a row should be rather peculiar. In one dimension, the interaction energy of free bosons is strongly enhanced relative to their kinetic energy at low densities [175], a result that will be discussed in greater detail in Sec. 4.3.3. The resulting Tonks-Girardeau gas [148, 149] is described as the hard-core limit of the (integrable) Lieb-Liniger model for bosons with delta-function interactions in one dimension [206]. These hard-core bosons can be exactly mapped to non-interacting fermions [149]; for general densities the state is well-described within the framework of Luttinger liquid theory [145]. The Tonks-Girardeau gas was first realized in ultracold atomic Bose gases confined in one-dimensional optical lattices [207, 208]. More recently, the transition between a Luttinger liquid and Mott-insulating state was observed [183]. These results suggest that a one-dimensional arrangement of coupled cavities should be sufficient to induce the mobile photons to behave entirely as fermions. In fact, this possibility was noted previously for a dissipative model [186]. The central goal of the present investigation is to show that the photons are effectively fermionized in the ground-state of the JCH model. A secondary goal is to make a careful, side-by-side comparison of the JCH and BH models in one dimension.

In this work, the ground states of the zero-temperature one-dimensional JCH and BH models in the vicinity of the MI phase boundary are obtained numerically using finite-system density matrix renormalization group (DMRG) methods. Several quantities are calculated that provide evidence for the nature of the states, including particle densities, one-particle and two-particle correlation functions, the superfluid fraction, the condensate fraction, and the entanglement entropy. Finite-size scalings are performed in order to infer the value of these quantities in the thermodynamic limit.

Two main conclusions can be drawn from our work. First, for small cavity couplings the ground state consists of a polariton MI phase as expected, though with some interesting features not shared by single-component Bose systems. Second, the results clearly reveal the strong fermionization of both the photons and spins in the JCH system throughout the so-called superfluid phase in the low-density limit. Detailed comparisons are made with the ground states of the BH model in the equivalent parameter regimes. The main conclusion that can be drawn is that superfluidity is weakly manifested in this phase, if it exists at all.

The manuscript is organized as follows. The JCH and BH models are reviewed in Sec. 4.3, and the DMRG methods used in the characterization of their ground states is described. This section also discusses the properties of the Tonks-Girardeau gas. The numerical results are presented in Sec. 4.4, and the discussion and conclusions are found in Sec. 4.5.

4.3 Models and Methods

In this work we compare the properties of the one-dimensional Jaynes-Cummings-Hubbard (JCH) and Bose-Hubbard (BH) models, using density matrix renormalization group (DMRG) methods. This section briefly provides the background to these models and

describes the numerical methods employed in the calculations.

4.3.1 JCH model

The behavior of a single two-level atom, in the pseudospin representation, confined to a single high-finesse cavity is given by the Jaynes-Cummings Hamiltonian [152], written within the rotating-wave approximation [209] as

$$H_{\text{JC}} = \omega_c \left(a^\dagger a + \frac{1}{2} \right) + \frac{1}{2} \omega_a \sigma^z + g (a^\dagger \sigma^- + a \sigma^+). \quad (4.1)$$

Here, ω_c is the natural cavity frequency ($\hbar = 1$ in this work for convenience), ω_a is the excitation frequency of the atom, a ($a^\dagger$) is the photon annihilation (creation) operator, σ^z is the spin-1/2 representation of the Pauli z operator, $\sigma^\pm$ are the spin raising and lowering operators, and g is the strength of the atom-photon coupling, proportional to the magnitude of the inner product between the dipole vector and the local field. In this work, g is assumed to be a real quantity, equivalent to assuming that the dipole and field oscillate in phase. This model describes an isolated system which ignores environmental couplings. The cavity is therefore assumed to have arbitrarily high finesse, and be in the strong-coupling limit.

In the rotating-wave approximation the total number of excitations

$$N_i = a_i^\dagger a_i + \sigma_i^+ \sigma_i^-, \quad (4.2)$$

is a conserved quantity (note that the spin number operator would normally contribute the term $\sigma_i^- \sigma_i^+$, but this corresponds to the population of atoms in the ground state; hence, it is not included in the *excitation* number operator). The eigenstates of H_{JC} are coherent superpositions of photonic and spin excitations with a definite total excitation number, known as polaritons [154]. Within a particular excitation number block N , the

eigenstates are given (c.f. Appendix A of Ref. [194]) by

$$|N-\rangle = \sin \theta_N |Ng\rangle + \cos \theta_N |(N-1)e\rangle; \quad (4.3)$$

$$|N+\rangle = \cos \theta_N |Ng\rangle - \sin \theta_N |(N-1)e\rangle, \quad (4.4)$$

with mixing angle

$$\theta_N = \frac{1}{2} \arctan \left(\frac{2g\sqrt{N}}{\delta} \right). \quad (4.5)$$

Here $\delta := \omega_c - \omega_a$ is the detuning of the cavity and atomic frequencies. The eigenenergies of Eq. (4.1) for $N \geq 1$ are given by

$$E_{N\pm} = N\omega_c + \frac{\delta}{2} \pm \left[\left(\frac{\delta}{2} \right)^2 + Ng^2 \right]^{1/2}, \quad (4.6)$$

while for $N = 0$, $E_0 = 0$. The energy levels are thus arranged in two-dimensional manifolds labeled by the polariton number N (except for the $N = 0$ sector, which is one-dimensional), separated by the energy of the single-photon cavity mode, ω_c .

The anharmonicity in the eigenenergies (4.6) of size $\sqrt{N}$ is the origin of the photon-blockade effect [153], giving rise to effective photon interactions. Consider for simplicity the zero-detuning case $\delta = 0$ giving eigenenergies $E_{N\pm} = N\omega_c \pm g\sqrt{N}$. With one photon the cavity has the lower energy eigenvalue $E_{1-} = \omega_c - g$. Naïvely, two independent photons would yield the total energy $2E_{1-} = 2\omega_c - 2g$, but in fact the two-photon energy eigenvalue is $E_{2-} = 2\omega_c - g\sqrt{2}$. The difference between these energies yields an estimate for the effective repulsive photon-photon interaction strength: $E_{2-} - 2E_{1-} = (2 - \sqrt{2})g$.

In the Jaynes-Cummings-Hubbard (JCH) model, the cavity mode leakage is no longer neglected. Instead, one imagines a regular lattice of L^d cavities in d dimensions positioned sufficiently close together that a photon emitted from one cavity can be absorbed into

an adjacent cavity with energy (rate) κ . The JCH model is written as

$$H_{\text{JCH}} = \sum_i (H_{\text{JC},i} - \mu N_i) - \kappa \sum_{\langle i,j \rangle} (a_i^\dagger a_j + a_j^\dagger a_i), \quad (4.7)$$

where $H_{\text{JC},i}$ and N_i are Eqs. (4.1) and (4.2) respectively with $\{a, a^\dagger, \sigma^z, \sigma^\pm\}$ replaced by $\{a_i, a_i^\dagger, \sigma_i^z, \sigma_i^\pm\}$, and i labels the position of a cavity. The $\langle i, j \rangle$ notation indicates that the sum is over nearest neighbors. The chemical potential μ fixes the mean polariton number, and is employed primarily to connect with the results of the BH model which is generally solved in the grand canonical ensemble. Note that in the JCH model, the atoms are fixed within the cavities, and only the photons are able to ‘tunnel’ from cavity to adjacent cavity.

For reasons that will become clearer momentarily, it is convenient to rescale the JCH energy in units of the coupling constant g ; in the limit of zero detuning $\omega_c = \omega_a$ one can rewrite the Hamiltonian (4.7) as

$$\begin{aligned} \frac{H_{\text{JCH}}}{g} &= -\frac{\kappa}{g} \sum_{\langle i,j \rangle} (a_i^\dagger a_j + a_j^\dagger a_i) \\ &+ \sum_i \left[a_i^\dagger \sigma_i^- + a_i \sigma_i^+ - \left(\frac{\mu - \omega_c}{g} \right) N_i \right], \end{aligned} \quad (4.8)$$

where unimportant additive constant terms are omitted. The first term corresponds to the photon hopping, the second to the local JC term, and the last term can be considered as a rescaled chemical potential for the total polariton density $\tilde{\mu} := (\mu - \omega_c)/g$.

4.3.2 BH Model

In the BH model, bosons tunnel between nearest neighboring sites of a lattice, and experience on-site interactions (which can be either attractive or repulsive in general).

The model is described by the Hamiltonian

$$H_{\text{BH}} = -t \sum_{\langle i,j \rangle} (b_i^\dagger b_j + b_j^\dagger b_i) + \sum_i \left[\frac{U}{2} N_i(N_i - 1) - \mu N_i \right], \quad (4.9)$$

where b_i , $b_i^\dagger$, and $N_i := b_i^\dagger b_i$ are respectively the bosonic annihilation, creation, and number operators for site i , $t > 0$ is the nearest-neighbor tunneling amplitude and U is the on-site interaction energy. The chemical potential fixes the mean boson density on the lattice. In this work only repulsive interactions $U > 0$ will be considered.

While a direct mapping between the JCH and BH models is not possible because the former has two different kinds of excitations while the latter has only one, the parameters can be chosen in such a way as to simplify comparisons. One can rescale the BH energies in terms of the interaction strength by dividing Eq. (4.9) by U . In this case the hopping amplitudes are κ/g and t/U in the JCH and BH models, respectively. The effective interaction strength between photons is $(2 - \sqrt{2})g$, which implies that the two systems should become similar for $t/U \sim \kappa/(2 - \sqrt{2})g = \left(1 + \frac{1}{\sqrt{2}}\right)(\kappa/g) \approx 1.707(\kappa/g)$. A similar connection can be obtained between the chemical potentials of the two models: $\tilde{\mu} + 1 \sim (2 - \sqrt{2})\mu/U$ or $\mu/U \sim 1.707(\tilde{\mu} + 1)$. Note that these scalings are valid only when only a single atom is confined to each cavity.

As discussed in the Introduction, in the weak tunneling (strong interactions) limit $t/U \ll 1$ the ground state is an incompressible MI characterized by localized bosons, (constant) integer occupation of a given site, and an energy gap to excitations of order U . Deep in this limit, the ground-state wavefunction can be approximated as $|\Psi\rangle \sim \prod_i b_i^\dagger |\Phi\rangle$, where $|\Phi\rangle$ is the particle vacuum state and the product is over all lattice sites. Because each site is independent of any other, the overlap of the states $b_s|\Psi\rangle$ and $b_r|\Psi\rangle$ is exactly

zero unless $r = s$. The one-body boson correlation function

$$G^{(1)}(r, s) = \langle b_r^\dagger b_s \rangle \quad (4.10)$$

and the two-body correlation function

$$G^{(2)}(r, s) = \langle b_r^\dagger b_s^\dagger b_s b_r \rangle \quad (4.11)$$

will then be zero for all $r \neq s$. In reality, for any finite t/U the gapped ground state will deviate from this simple prediction and the correlation functions should instead decrease exponentially in $|r - s|$ with a characteristic length scale $\xi \sim 1/U$ that scales as the inverse of the gap to excitations [210]. The correlation length diverges as a system becomes critical [108]. One would therefore expect ξ to increase from 0 to ∞ as the hopping goes from 0 to its critical value at the phase boundary for a fixed chemical potential.

The phase boundary in μ - t space, known as the ‘Mott lobe,’ is roughly semi-circular in profile in two and three dimensions [3]. For $\mu \notin \mathbb{Z}$, the system remains in the MI phase with increasing t until some critical value at which point the system undergoes a phase transition to SF; likewise for constant t and increasing μ . The Mott lobe becomes strongly distorted in one dimension [5], and the system displays re-entrance: at constant μ , on increasing t the ground state phase changes from MI to SF to MI and back to SF again.

In the strong tunneling (weak interactions) limit $t/U \gg 1$, the ground state of the BH model corresponds to an interacting Bose-Einstein condensate. Each boson is highly extended throughout the lattice, and the ground state can be approximated by $|\Psi\rangle \sim \left(\sum_i b_i^\dagger\right)^{N_B} |\Phi\rangle$, where N_B is the number of bosons. This compressible state is characterized by a gapless linear spectrum and long-range correlation functions (4.10)

and (4.11) that are independent of $|r - s|$. In one-dimension, however, true Bose-Einstein condensation is not possible; rather, the ground state corresponds to a quasi-condensate with only algebraic long-finite order and characterized by strong fluctuations [211]. Instead one finds $G^{(1)}(r, s) \sim 1/|r - s|^\alpha$, where the parameter α characterizes the degree of quasi-condensation ($\alpha \rightarrow 0$ for a true condensate).

4.3.3 Tonks-Girardeau Gas

At very low densities, one-dimensional repulsively interacting bosons form a Tonks-Girardeau gas, and effectively behave as non-interacting fermions [149]. In the absence of any external potential (other than the ones used for confinement), the ground state properties are governed by the kinetic and interaction potential energies T and U . The mean kinetic energy per particle scales as $T/N \sim 1/m\ell^2 \sim \langle n \rangle^2/m$, where ℓ is the mean interparticle distance which in one dimension scales as the inverse of the mean particle density $\ell \sim 1/\langle n \rangle$. (In the presence of a weak lattice, the bare boson mass m is rescaled to an effective mass $m^* \sim 1/t$). When the interaction potential can be modeled in terms of a pseudopotential (low energy, long-wavelength collisions), one can write the mean interaction potential in one dimension as $U/N \sim \langle n \rangle/m|a_{1D}|$ [212], where a_{1D} is the one-dimensional s-wave scattering length. The Tonks parameter, the ratio of the potential and kinetic energies $\gamma = U/T = 2/\langle n \rangle|a_{1D}|$ is therefore huge at low densities, in marked contrast to the situation in higher dimensions. To minimize the interaction potential, particles prefer to be as far apart from one another as possible, much like fermions.

The free fermionic wavefunction can be written in terms of a Slater determinant to guarantee the proper antisymmetrization of the wavefunction. For example, a system of

N free fermions on L sites has a wavefunction given in the position representation by

$$\Psi_F(r_1, \dots, r_L) = \det \begin{bmatrix} \phi_1(r_1) & \phi_1(r_2) & \dots & \phi_1(r_L) \\ \phi_2(r_1) & \phi_2(r_2) & \dots & \phi_2(r_L) \\ \vdots & \vdots & \ddots & \vdots \\ \phi_N(r_1) & \phi_N(r_2) & \dots & \phi_N(r_L) \end{bmatrix}, \quad (4.12)$$

where the $\{r_i\}$ indicate the positions of the lattice sites and the $\{\phi_i\}$ are single-particle wavefunctions. In the perfectly hard-core limit of the Tonks-Girardeau gas, the ground state of the fermionized bosons is simply

$$\Psi_B(r_1, \dots, r_L) = \prod_{i < j}^L \text{sgn}(r_i - r_j) \Psi_F(r_1, \dots, r_L), \quad (4.13)$$

where the factor multiplying Ψ_F ensures that all negative signs associated with the interchange of two fermions disappears.

Many properties are shared by Ψ_F and Ψ_B . For example, the local density profile of both systems in real space is the same, since $|\Psi_B(r_1, \dots, r_N)|^2 = |\Psi_F(r_1, \dots, r_N)|^2$ [213]. Similarly, all density correlation functions are the same [214]; for example, for a ring of length $L \rightarrow \infty$, the normalized two-body correlation function is

$$g^{(2)}(r, s) = \frac{\langle b_r^\dagger b_s^\dagger b_s b_r \rangle}{\langle b_r^\dagger b_r \rangle \langle b_s^\dagger b_s \rangle} = 1 - \left[\frac{\sin(\pi n |r - s|)}{\pi n |r - s|} \right]^2, \quad (4.14)$$

where n is the mean particle density. The correlation function is zero at $r = s$, reflecting the Pauli exclusion principle; this behavior is referred to as the ‘exclusion hole’. Away from this point the correlation function grows and displays Friedel oscillations [215] that decay with increasing $|r - s|$. For one-dimensional spinless fermions, the oscillations have wavelength $\lambda_F = 1/n = 2\pi/k_F$ where k_F is the Fermi wavevector. Thus, the presence

of an exclusion hole and Friedel oscillations in the two-body correlation function is a ‘smoking gun’ for the fermionization of bosons in the Tonks-Girardeau gas.

For a finite system with N free fermions on L sites with open boundary conditions, such as is considered in this work, a straightforward calculation yields

$$G^{(2)}(r, s) = \left(\frac{N}{L+1} \right)^2 [B(r, s) - A(r, s)], \quad (4.15)$$

where

$$\begin{aligned} A(r, s) &= \left[\frac{\cos \frac{\pi(N+1)(r-s)}{2(L+1)} \sin \frac{\pi N(r-s)}{2(L+1)}}{N \sin \frac{\pi(r-s)}{2(L+1)}} \right. \\ &\quad \left. - \frac{\cos \frac{\pi(N+1)(r+s)}{2(L+1)} \sin \frac{\pi N(r+s)}{2(L+1)}}{N \sin \frac{\pi(r+s)}{2(L+1)}} \right]^2; \\ B(r, s) &= \left[1 - \frac{\cos \frac{\pi(N+1)r}{(L+1)} \sin \frac{\pi N r}{(L+1)}}{N \sin \frac{\pi r}{(L+1)}} \right] \\ &\quad \times \left[1 - \frac{\cos \frac{\pi(N+1)s}{(L+1)} \sin \frac{\pi N s}{(L+1)}}{N \sin \frac{\pi s}{(L+1)}} \right]. \end{aligned} \quad (4.16)$$

It is simple to verify that $G^{(2)}(r, r) \rightarrow 0$. For large separations between particles $\left| \frac{(r-s)}{(L+1)} \right| \gg 0$, one finds that $A(r, s)$ and $B(r, s)$ oscillate in the vicinity of zero and unity, respectively, so that $G^{(2)}(r, s) \approx \left(\frac{N}{L+1} \right)^2$. Choosing the location of one particle at the center of the chain $r = \lceil L/2 \rceil$, $G^{(2)}$ far from the center oscillates about a mean value approximately equal to the square of the mean particle density n . For $\left| \frac{r-s}{L+1} \right| \gg 0$, the oscillation of $G^{(2)}$ is governed by the last term in the definition of $B(r, s)$ in Eq. (4.16). In the thermodynamic limit $N, L \rightarrow \infty$ but $n = N/L \rightarrow \text{const.}$, one obtains $B(\lceil L/2 \rceil, s) \approx 1 - \sin(2\pi n s)/2\pi n s$. The Friedel oscillation wavelength is therefore again $\lambda_F = 1/n = 2\pi/k_F$.

The single-body correlation function is not the same for the Tonks-Girardeau and free fermion gases, however: the sign function in Eq. (4.13) does not disappear when

inserted into Eq. (4.10). The calculation of this quantity is quite involved [214], but the asymptotic behavior $|r - s| \gg 0$ but $|r - s| \ll L$ is found to be

$$G^{(1)}(r, s) \sim \frac{1}{\sqrt{2n_0 L} |\sin(\pi|r - s|/L)|}. \quad (4.17)$$

For $|r - s| \ll L$ one obtains $G^{(1)}(r, s) \sim 1/\sqrt{|r - s|}$, which indicates that for the Tonks-Girardeau gas the exponent of the power law is $\alpha = 1/2$. Another ‘smoking gun’ for the Tonks-Girardeau phase is therefore the power-law behavior of the one-body density matrix with exponent $\alpha = 1/2$.

The Fourier transform of the one-particle correlation function $G^{(1)}(r)$ is the momentum distribution $n(k)$. For the Tonks-Girardeau gas, the power-law behavior at long distances translates into a power-law divergence of the momentum distribution at long wavelengths, $n(k) \sim 1/|k|^{1/2}$ for $k \rightarrow 0$. This highly peaked distribution is reminiscent of the delta-function distribution that one would expect if the bosons formed a Bose-Einstein condensate, except it is now broadened due to the finite-range phase order associated with the quasi-condensation. This distribution is dramatically different from that of a non-interacting Fermi gas, where $n(k)$ is a constant for all $k \leq k_F$ and is zero otherwise (k_F is the Fermi wavevector). The momentum distribution for a Tonks-Girardeau gas in a weak axial trapping potential has been experimentally observed [207, 208].

4.3.4 Numerical Methods

The characteristics of the BH and JCH models were obtained by means of finite-system density matrix renormalization group (DMRG) simulations. We employed the DMRG code from the Algorithms and Libraries for Physics Simulations (ALPS) project [158, 159]. Simulations were carried out for systems of size $L = 15, 19, 23, 27, 31$ with both open boundary conditions (equivalent to hard-wall boundary conditions) and periodic

boundary conditions, and a finite-size scaling analysis was performed for all quantities (unless explicitly noted) in order to infer the results for the thermodynamic limit. We use DMRG as the method because it is suitable for obtaining results that are so precise as to be considered exact [156], while being able to handle much larger finite-size systems than exact diagonalization [7, 157].

The bulk of the simulations employed open boundary conditions, in order to accelerate convergence. For the BH (JCH) model, a maximum of $N_{\max} = 5$ (6) bosons (photons) per site (cavity) were allowed, and we kept $M = 80$ (100) states. For the JCH model, this corresponds to a Hilbert dimension $D = 12$ for each cavity. For the superfluid fraction, the method chosen necessitated the use of periodic boundary conditions. Usually the number of states kept for these simulations is on the order of the square of the number chosen for open boundary conditions; however, since the method only required ground state energies and not correlation functions, the numerical requirements were not as stringent. The superfluid fraction calculation for the BH (JCH) model was performed using $N_{\max} = 7$ (6) and $M = 200$ (140). In all cases, we verified that increasing the values of M and $N_{\max}$ did not change the ground state energies or correlation functions. For the BH (JCH) system, these parameters correspond to a maximal Hilbert space dimension of 1.60×10^5 (1.44×10^6) for the simulations with open boundary conditions, and 1.96×10^6 (2.82×10^6) for periodic boundary conditions.

We used eight finite-size sweeps for all simulations, and verified that the ground state energy and correlation functions did not change by increasing the number of sweeps. The calculation times for a single run for the simulations for the BH (JCH) models were typically approximately 1 hour (24 hours) when using open boundary conditions, while the runs using periodic boundary conditions required up to approximately 8 (24) hours. In order to keep the calculation time to a minimum, we assumed that the total boson (polariton) number was a conserved quantity in every case. Note that this does not pose a

problem even in the superfluid phase because the superfluid density need not correspond to the mean boson (polariton) density.

Parameters for the hopping (t or κ for the BH or JCH models, respectively) and chemical potential μ were chosen in order to remain in the vicinity of the MI phase boundary. In 1D, the tip of the Mott $n = 1$ lobes in the $(\mu/U, t/U)$ and $(\tilde{\mu}, \kappa/g)$ planes for the BH and JCH models are found to be located at approximately $(0.09, 0.3)$ [5, 172, 216] and approximately $(-0.95, 0.2)$ [6, 202], respectively. This value of κ/g is consistent with the rescaling factor of approximately 1.7 between the BH and JCH models, discussed in Sec. 4.3.2. To span most of the Mott lobe, the range of hopping is therefore chosen to be $t/U \in [0, 0.27]$ and $\kappa/g \in [0, 0.16]$ for the two models. Likewise, the phase boundaries for $t = 0$ and $\kappa = 0$ correspond to $\mu/U = n$ [3] and $\tilde{\mu} = \sqrt{n} - \sqrt{n+1}$ [194] for the BH and JCH models in any dimension, respectively. Thus, the transition from the $n = 0$ to $n = 1$ Mott lobes at zero hopping occurs for $\mu/U = 1$ and $\tilde{\mu} = -1$ for the two models; the transition to the $n = 2$ Mott lobe occurs for $\mu/U = 2$ and $\tilde{\mu} \approx -0.41$. To capture some of the $n = 0$ lobe and approximately half of the $n = 1$ lobe, we chose chemical potentials in the range $\mu/U \in [-0.17, 0.55]$ and $\tilde{\mu} \in [-1.10, -0.68]$. Only the simplest zero-detuning case $\delta = 0$ is considered in this work. Previous work has shown that detuning can be a useful parameter, changing the effective strength of interactions and thereby the phase diagrams [194, 195, 197, 205].

4.4 Results

4.4.1 Density phase diagrams

The location of the phase boundary between the gapped MI phase and the SF phase has been previously established numerically in the thermodynamic limit with finite-size DMRG, for both the 1D Bose-Hubbard model [5, 172, 216] and the 1D JCH model [6].




Figure 4.1: Mean density phase diagrams for (a) polaritons in the 1D JCH model and (b) bosons in the 1D BH model, in both cases for a system size $L = 31$. Regions of constant mean density correspond to MI states.

The results for the mean densities of bosons and polaritons are shown in Fig. 4.1 for the parameter sets discussed in the previous section and the largest number of lattice sites $L = 31$ studied. Open boundary conditions are employed in this case; for periodic boundary conditions the density on any site would coincide with the mean density. While such density plots have not to our knowledge been previously shown in the literature, the main purpose of showing these plots here is to orient the reader to the location in phase space for which the simulations have been conducted. The ranges of the normalized hopping parameters $t/U, \kappa/g$ and effective chemical potentials $\mu/U, \tilde{\mu}$ are chosen to be equivalent, based on the scaling assumption given in Sec. 4.3.2. For both models, the goal is to explore the regions in the vicinity of the transition between the MI and SF phases. Unlike most previous studies, this work focuses particularly on the low-density region where fermionization is expected.

The mean densities of either bosons or polaritons can be used to distinguish the two phases in both models, since the mean density is pinned to an integer in the MI regime, but not in the SF regime. Consequently, the locations of the phase boundary for our finite-size systems are quite clearly visible in Fig. 4.1. The shapes of the phase boundaries

closely resemble the thermodynamic limit results of Refs. [5, 6], though their positions are shifted slightly due to the finite-size system. Depicted is the region of the phase diagram where the $n = 0$ and $n = 1$ Mott lobes meet, as well as the low-density superfluid regions near the boundaries of these lobes. Roughly, the mean density of bosons (polaritons) increases with increasing t (κ). This work is mainly concerned with the low-density superfluid regions of phase space, corresponding to low hopping and chemical potential.

The re-entrant shape of the $n = 1$ Mott lobe can be clearly seen as the constant-density region of the BH model in Fig. 4.1b, but is not as obvious in Fig. 4.1a. At certain fixed values of μ within a continuous range, monotonically increasing the hopping parameter t or κ from zero causes the system to transition from the MI to the SF, back to the MI and again back to the SF regime [6]. Viewed from within the Mott lobe, the lower part of the phase boundary (the hole boundary) is concave, while the upper part (the particle boundary) is convex. The particle and hole boundaries meet at a sharply cusped tip. The phase transition along the line of constant mean density passing through this point is in the $(d + 1)$ -dimensional XY universality class, which for $d = 1$ is of the Berezinskii-Kosterlitz-Thouless (BKT) type [3, 134] with Tomonaga-Luttinger parameter $K_b = 1/2$ [172, 216]. The MI-to-SF transition across either the particle or the hole boundary is generic (i.e. Gaussian like the condensation transition of an ideal Bose gas) and characterized by $K_b = 1$ [172]. This implies that the SF phase near the particle or hole boundaries should be characterized by one-particle correlation functions $G^{(1)}(r) \sim |r|^{-K_b/2} \sim 1/r^{1/2}$, consistent with the Tonks-Girardeau gas scaling.

In the JCH model, the mean densities of the spin and photonic species, depicted in Fig. 4.2, need not remain proportional to track each other. Consider first the mean spin excitation density, Fig. 4.2a. In the atomic limit $\kappa = 0$, the cavities are decoupled from each other and thus the overall ground state is the L -fold tensor product of the




Figure 4.2: (a) Mean spin and (b) photonic excitation densities in the JCH model with $L = 31$. These can vary throughout the MI phase.

single-cavity lower-lying polariton states

$$|\psi_{\text{ground}}\rangle = \bigotimes_{i=1}^L |1-\rangle_i. \quad (4.18)$$

From Eqs. (4.3) and (4.5) in the case of zero-detuning, the single-cavity polariton ground states $|1-\rangle$ are equal-weight superpositions of a photonic and a spin excitation. Consequently, the mean spin excitation density is equal to $\frac{1}{2}$ in this limit. Conversely, in the hopping-dominated limit $\kappa \rightarrow \infty$, the photons and spins decouple. Each atom is then in the ground state so the mean spin excitation density vanishes. At intermediate κ between these extremes, the data in Fig. 4.2 show that across the lower boundary of the Mott lobe (the hole boundary) there is a sharp drop in the mean density of spin excitations, which can be used to distinguish the two phases. By contrast, crossing the upper boundary, the mean spin excitation density varies smoothly from $\frac{1}{2}$ in the atomic limit to lower values, presumably tending towards 0 in the limit $\kappa \rightarrow \infty$.

Next consider the mean photon density. Unlike the spin excitations, the number of photons and hence the mean photon density is unbounded from above in the grand

canonical ensemble. In the large-hopping limit at fixed chemical potential, one expects the mean photon density to increase. In fact, for a lattice with coordination number z_c , an instability occurs when the hopping roughly satisfies $z_c \kappa / g > -\tilde{\mu}$; for larger values of the hopping, the ground state energy decreases without bound as a function of increasing photon number [194]. This regime is not considered in the calculations, as it corresponds to larger densities deep in the SF phase where fermionization is unlikely. The mean photon density within the $n = 1$ MI lobe tracks with the mean spin excitation density in such a way that the overall polariton density is pinned to 1 per site, as expected. Crossing the upper (particle) boundary of the Mott lobe, the mean density of photons begins to increase rapidly with increasing hopping. This indicates that in the hopping-dominated limit, the system behaves like a photon superfluid, with the effects of the spins becoming negligible. On the other hand, in the intermediate regions between the $n = 0$ and $n = 1$ lobes, the mean photon density remains low; the reason is that at low hopping in one dimension, the effective repulsive interactions between photons become strong and thus there is an energy cost associated with adding photons to the system.

4.4.2 Correlation functions

One-body density matrix

Consider now the single-particle correlation function $G^{(1)}$, defined in Eq. (4.10). This has been calculated previously via DMRG for the 1D BH model, to verify the asymptotic predictions of the Luttinger liquid theory [217], and to estimate the location of the critical value of t/U for the BKT transition [5]. Similar plots of $G^{(1)}$ for varying interaction strengths in the $n = 1$ MI lobe of the 1D BH model are shown in Ref. [218]. The normalized version $C_r(s - r) := G^{(1)}(r, s) / \sqrt{N_r N_s}$ was also considered for the 1D BH model with an additional harmonic trapping potential [219] for various different sites r , and the coexistence of the two phases was found at certain points in the phase diagram.

As discussed in Sec. 4.3.2, in the MI phase one expects this correlation function to decrease exponentially with distance, $G^{(1)}(r, s) \sim \exp(-|r - s|/\xi)$ over a correlation length ξ , while in the SF phase it should behave as a power law, $G^{(1)}(r, s) \sim 1/|r - s|^\alpha$ where α is some positive constant. The correlation functions are calculated for the JCH model (in which case $b \rightarrow a$ and σ for photons and spins, respectively) and the BH model for systems of size $L = 15, 19, 23, 27$, and 31 . Consider for concreteness the shortest length $L = 15$, for which the fits are the least reliable; the results are shown in Fig. 4.3. This size is chosen convey the worst-quality results for various system sizes. The correlations are measured with respect to the central site, in this case site number 8. The point where $r = s$ is not plotted or used for the fit, since only the asymptotic form of the correlations for $|r - s| \gg 0$ is of interest. To mitigate boundary effects, the two sites closest to the (open) boundary of the system were also not considered. The correlations for the spins and the photons track each other very closely, so the results for photons are not presented.

The results indicate that the spin excitations and photons behave in close analogy to the bosons of the BH model; that is, the single-particle correlation functions are clearly exponential within the MI lobe and follow power laws in the SF regime. The power law relationship holds up well not only for $L = 15$, but also for all larger values of L considered. The exponent associated with the power law, corresponding to the slope of the log-log fit, is increasingly precise for larger values of L , with the standard error $\sim 1/\sqrt{L}$. The correlation function decreases most rapidly deep in the MI region but increasingly slowly as the phase boundary is approached. The results shown in Fig. 4.3b correspond to the SF phase at constant μ , for values of κ/g that range from almost immediately adjacent to the MI lobe ($\kappa/g = 0.022$) almost to the edge of the phase diagram ($\kappa/g = 0.133$) in Fig. 4.2a. Unsurprisingly, the power-law fits are poor near the phase boundary (c.f. the points corresponding to $\kappa/g = 0.022$) but improve as one moves




Figure 4.3: Representative plots of the correlation function $G^{(1)}(8, s)$ for the spin excitations in the JCH model with $L = 15$. (a) $n = 1$ MI phase at point $\tilde{\mu} = -0.900$, and (b) SF phase at point $\tilde{\mu} = -0.989$. In both cases the same five values of κ/g are chosen with values given in the legends. The correlation functions are consistent with (a) exponentials and (b) power-laws.

further away.

The correlations exhibited by each type of carrier also track perfectly with each other. This result is in qualitative agreement with Ref. [198], in which the same quantity was calculated using the Variational Cluster Approximation. In fact, this feature persists for all quantities discussed below, unless mentioned explicitly. An intriguing consequence of the identical behavior for the two species of excitations is that even though one normally views the atoms as mediating photonic interactions, one could just as well think of the photons as mediating atomic interactions; though the atoms are each isolated within their own cavities, they nevertheless feel each others' presence.

The infinite-system values of the correlation length ξ and the power α are estimated using a finite-size scaling analysis, as shown in Fig. 4.4 for a few representative points in phase space. The best exponential or power-law function is fit for $G^{(1)}$ for each value of $L = 15, 19, 23, 27$, and 31 , such as is shown in Fig. 4.3 for $L = 15$. The values of ξ and




Figure 4.4: Representative finite-size scaling analysis of $G^{(1)}$ in the (a) MI and (b) SF regimes of the JCH model. The data points for the MI regime are $(\tilde{\mu}, \kappa/g) = (-0.856, 0.044)$ (deep within MI lobe), $(-0.856, 0.128)$ (near generic transition), and $(-0.900, 0.156)$ (near BKT point). The SF regime points are at $(-0.989, 0.022)$ (near generic transition), $(-0.989, 0.106)$ (deep within SF phase), and $(-0.989, 0.156)$ (near BKT point).

α are then plotted as a function of $1/L$, and the data are fit to a line whose intercept is interpreted as the corresponding value in the thermodynamic limit. The data are only weakly dependent on system size deep in the MI phase and near the generic phase boundary, but show a strong dependence near the BKT point. On the SF side the values of α are size-dependent for all three phase space points considered; this likely reflects the fact that along the $\tilde{\mu} = -0.989$ line the Mott boundary remains nearby. The same procedure is carried out for the BH model for comparison (not shown).

The phase diagram for the value of ξ in the thermodynamic limit is displayed for the JCH and BH models in Fig. 4.5. Note however that one cannot obtain ξ in this way immediately at the phase boundary. The true correlation length is expected to diverge, and a diverging correlation length cannot be accurately captured by a DMRG procedure with finite truncation M [157]. Furthermore, the precise location of the phase boundary




Figure 4.5: Phase diagrams for ξ in the thermodynamic limit, assuming $G^{(1)}(r) \sim \exp(-r/\xi)$ in MI regime of the (a) JCH model (photons only) and the (b) BH model. The correlation length ξ approaches the system size at the edge of the Mott lobe near the BKT point.

varies depending upon the size of the system, so a particular point in phase space near the boundary displayed for $L = 31$ may or may not be in the Mott-insulating lobe, depending upon the system size. In Fig. 4.5, only the phase space points that are unambiguously within the $n = 1$ lobe have been included; this accounts for the apparently smaller MI lobes than are depicted in Fig. 4.1.

The numerical results clearly indicate that the correlation length is independent of μ and is solely a function of κ or t , smoothly increasing with increasing hopping. Since everywhere within the Mott lobe the state has a well-defined number of excitations N_{tot} , the only effect of varying the chemical potential by an amount $\Delta\mu$ while fixing κ or t is to shift the entire spectrum by an amount $N_{\text{tot}}\Delta\mu$. The ground state itself at fixed hopping is therefore independent of μ within the lobe. The correlation function only increases as the BKT point is approached, not near the generic phase boundaries.

The value of α in the thermodynamic limit is calculated for all points in the superfluid phase for both the JCH and BH models, and the results are shown in Fig. 4.6. Right at the phase boundary, the exponent approaches values on the order of unity or higher. The

same caveats mentioned above for the calculation of ξ apply here as well. In addition, the spatial dependence of $G^{(1)}$ (an example of which is shown in Fig. 4.3b) and the finite-size scaling data (an example of which is shown in Fig. 4.4b) are much noisier right near the phase boundary.

Everywhere near the phase transition, however, the value of α is close to 0.5. This value is consistent with what would be expected for a Tonks-Girardeau gas of photons [c.f. Eq. (4.17)] and matches the Luttinger parameter $K_b = 1$, as discussed in Sec. 4.4.1. In fact, α has previously been used to obtain the location of the phase boundary for the BH model, using infinite system DMRG with periodic boundary conditions [5]. This indicates that the low-density regime outside of the Mott lobes is in fact not strongly superfluid in nature. Rather, the results are consistent with the complete fermionization of the BH bosons and the JCH photons in the equivalent regime. Furthermore, since $G^{(1)}$ for the spin excitations and photons in the JCH model track each other so well (not shown), the spin excitations have also been fermionized, even though the atoms are treated as spins with no particular exchange statistics. The ‘SF’ designation of this phase therefore appears to be a misnomer, but it will be kept for clarity of exposition in what follows.

The value of α decreases for increasing hopping κ or t , but the trend is slow for the parameter range studied. By the edge of the plots in Fig. 4.6, the exponent for the SF phase between the $n = 0$ and $n = 1$ lobes has dropped to almost constant (in terms of μ) values of 0.33 and 0.29 for the JCH and BH models, respectively. For larger values of μ where the density is higher the exponent drops off more rapidly, reaching a range of 0.081-0.254 for the JCH model and 0.226-0.336 for the BH model at the edge of the plots. These values are all quite different from the value $\alpha = 0$ that one would expect for an ordinary superfluid, however.




Figure 4.6: Phase diagrams for α in the thermodynamic limit, assuming $G^{(1)}(r) \sim r^{-\alpha}$ in SF regime of the (a) JCH model (photons only) and the (b) BH model. These values of α indicate strong fermionization of the photons in the JCH model.

Two-body correlation function

One of the most important signatures of fermionization within the low-density SF phase is found in the two-body correlation function $G^{(2)}(r, s)$ defined in Eq. (4.11) and its normalized variant $g^{(2)}(r, s)$ defined in Eq. (4.14), as discussed in Sec. 4.3.3. Within the $n = 1$ Mott-insulating lobe, one expects $g^{(2)}(r, s) = 1 - \delta_{r,s}$ irrespective of the model and the excitation, which is exactly what is observed. Of greater interest is the behavior of this correlation function in the SF regime for different mean densities. For a perfect superfluid the two-body correlation function is featureless, $g^{(2)}(r, s) = 1$ in the bulk, reflecting the fact that all superfluid carriers occupy the same plane wave state. On the other hand, for a system of free fermions, the two-body correlation function exhibits two important features. The first is the presence of an exclusion hole at $r = s$ reflecting the Pauli exclusion principle. The second is the characteristic Friedel oscillations appearing on either side of the exclusion hole, with a wavelength λ_F set by the mean density n or Fermi wavelength k_F , as discussed in Sec. 4.3.3.

The unnormalized two-body correlation functions $G^{(2)}(r, s)$ for the photons and spin




Figure 4.7: The spatial-dependence of the unnormalized two-body correlation function $G^{(2)}(16, s)$ for $L = 31$ is shown for representative points in the SF phase where $\alpha \approx \frac{1}{2}$. Figures (a) and (b) correspond to photons at the points $(\tilde{\mu}, \kappa/g) = (-0.989, 0.039)$ and $(-0.989, 0.106)$ in the JCH model, respectively; ‘T’ and ‘D’ in turn denote ‘theory’ and ‘data.’ Figures (c) and (d) correspond to bosons at the points $(\mu/U, t/U) = (0.019, 0.066)$ and $(0.019, 0.180)$, respectively.

excitations in the JCH model are plotted in Fig. 4.7 for representative points in the SF phase where $\alpha \approx \frac{1}{2}$; the BH results are also shown for comparison. (The momentum distribution $n(k)$ in the 1D BH model was previously considered for various interaction strengths within the $n = 1$ MI lobe [216], and for both the MI and SF regimes at unit filling [219]). The data are compared with the analytical expression for free fermions at the same mean density, Eq. (4.15). The unnormalized two-body correlation function is plotted in order to make apparent the particle densities for the particular points in phase space. The two main signatures of fermionization in the SF phase, the exclusion hole and the Friedel oscillations, are evident in all the plots. At lower densities, the match between the data and the prediction based on non-interacting fermions is excellent. This indicates that the superfluid density at this point in the SF phase is close (if not exactly equal) to zero. One would expect the systems to behave more like a superfluid as the tunneling is increased and the excitation density increases, for points in phase space that are further from the phase boundary. Indeed this is the case; the data depicted in Figs. 4.7b and 4.7d show that the exclusion hole is now slightly filled in, and the Friedel oscillations are increasingly washed out.

An interesting feature in the case of the JCH model, not present in the BH model, is that the wavelength of the oscillations for both the photons and the bosons is set by the mean density of polaritons. The amplitude of $G^{(2)}$ for each carrier is determined by the density of just that carrier, however. More precisely, to obtain a fitting curve of the form of Eq. (4.15) for $G^{(2)}(r, s)$ for the photons (spin excitations), one must use the mean number of photons (spin excitations) as the value of N in the prefactor $(N/(L+1))^2$, but the number of polaritons in the oscillatory functions appearing in $A(r, s)$ and $B(r, s)$. Hence, the spin excitations and the photons are individually fermionized, inasmuch as they have inherited this property from the polaritons.

4.4.3 Other measures of the ground state

Superfluid fraction

The most compelling evidence for the existence of a SF phase would be the presence of a non-zero superfluid order parameter. An example of such an order parameter is the superfluid fraction or superfluid stiffness f_s , the ratio of particles exhibiting superfluid flow to the total number of particles. The factor f_s can be calculated numerically by imposing periodic boundary conditions and applying a phase twist $\Theta \ll \pi$ to the boundary conditions [130, 131, 132, 129]. In practise, this can be accomplished by means of a Peierls factor applied to the bosonic creation and annihilation operators in the BH and JCH models, which has the effect of modifying the hopping terms via $a_i^\dagger a_j \mapsto a_i^\dagger a_j e^{-i\Theta/L}$ [220]. While this induces a velocity $v = (2J)\nabla\Theta$ for each quantum particle (J is the hopping coefficient corresponding to κ in the JCH model or t in the BH model) because the current density is $j = nv$, only the particles in the superfluid will respond collectively. As a result, the ground state energy will increase relative to the twist-free case solely due to the kinetic energy of the superfluid particles. From this change of the ground state energy, the superfluid fraction can be determined:

$$f_s = \frac{L}{2Jn} \left. \frac{\partial^2 E_0(\Theta)}{\partial \Theta^2} \right|_{\Theta=0} \approx \frac{L}{Jn} \frac{E_0(\Theta) - E_0}{\Theta^2}, \quad (4.19)$$

where E_0 is the ground state energy with no phase twist and $E_0(\Theta)$ is ground state energy with overall twist $\Theta \ll \pi$. The latter expression is a finite-difference approximation for the second derivative of the energy with respect to the phase twist, using the central three-point stencil. This calculation of ρ_s has previously been performed at constant density near the BKT transition of the 1D BH model across the tip of the $n = 1$ Mott lobe [217]; a significant jump in ρ_s across the transition was found in that work. This work instead examines ρ_s at representative points in the low-density SF phase far from

the transition, with considerably different results.

In the numerical calculations, we considered various points in the SF phase using five different finite-size systems $L = 15, 19, 23, 27, 31$. Periodic boundary conditions are required, which are computationally more demanding for the DMRG method than are open boundary conditions (see the discussion in Sec. 4.3.4); hence, only a set of representative points in the SF phase were considered. Recall from Sec. 4.3.4 that the $n = 1$ BKT points in the JCH and BH models are located at approximately $(\tilde{\mu}, \kappa/g) = (-0.95, 0.2)$ [6, 202] and $(\mu/U, t/U) = (0.09, 0.3)$ [5, 172, 216], respectively. We therefore considered these three points in the SF phase of the JCH model: $(-0.944, 0.133)$, $(-0.989, 0.028)$, and $(-1.00, 0.240)$. The first is just left of the BKT point, in the vicinity of the hole boundary, the second is between the $n = 0$ and $n = 1$ Mott lobes, and the third is to the right of the BKT point. In particular, the first two points correspond to mean polariton densities $n < 1$, while the third has $n > 1$. For the BH model we considered the two points: $(0, 0.08)$ and $(0, 0.5)$; again, the first is left of the BKT point in the vicinity of the hole boundary, while the second point is much to the right, in the deep SF phase well beyond the region depicted in Figs. 4.1b and 4.6b. The first of these points has mean boson density $n < 1$ and the second, $n > 1$.

The ground-state energy was obtained for these points in the SF region, for phases $\Theta/\pi \in (0.0, 1.0)$ in increments of 0.2. Note that the ground state energy is invariant under the transformation $\Theta \mapsto -\Theta$. For each value of L , the second derivative in Eq. (4.19) was then calculated using $\Theta = 0.2\pi$. The results were then verified by estimating the derivative using a central-difference approximation with five-point ($\Theta/\pi \in \{0, \pm 0.2, \pm 0.4\}$) and seven-point stencils ($\Theta/\pi \in \{0, \pm 0.2, \pm 0.4, \pm 0.6\}$); no discernible difference from the method of Eq. (4.19) was observed in the results so obtained. The values of f_s were also obtained by extracting the coefficient for the quadratic term in a polynomial fit of $E_0(\Theta)$. Again, no discernible differences from the central-difference results were found




Figure 4.8: Superfluid fraction for the (a) JCH and (b) BH model outside but near the $n = 1$ MI lobe. For the JCH model, the symbol $\tilde{\kappa} = \kappa/g$ is used for compactness. For both systems, the curve marked by red squares (color online) is beyond the critical hopping value for the BKT transition for the $n = 1$ MI lobe, while the others are not. In both cases, the superfluid fraction f_s tends to 0 in the thermodynamic limit.

using this approach. Once the value of f_s was obtained for a given system size, a finite-size scaling analysis in $1/L$ was performed to interpolate to the thermodynamic limit. As a final check on the results, the calculation in Eq. (4.19) was repeated with smaller values of the phase twist, $\Theta/\pi \in [0.0, 0.2]$ in increments of 0.02. For $\Theta/\pi < 0.06\pi$, the values of $\rho_s(L)$ could not be reliably fitted to determine the value in the thermodynamic limit (the problem of dividing one small number by another). However, for $\Theta/\pi \geq 0.06$, the results were consistent with those obtained with $\Theta/\pi = 0.2$.

The finite-size scaling results for the superfluid fraction f_s are shown in Fig. 4.8, and the resulting values and uncertainties of f_s in the thermodynamic limit are displayed in Table 4.1. These values are determined by the standard least-squares minimization procedure for linear fitting. In all cases, except for the JCH point marked by red squares in Fig. 4.8, the obtained thermodynamic limit value of the superfluid density contains 0

BH model				JCH model			
μ/U	t/U	f_s	Δf_s	$\tilde{\mu}$	κ/g	f_s	Δf_s
0	0.08	0.0007	0.0008	-0.944	0.133	-0.0020	0.0025
0	0.5	-0.0001	0.0001	-0.989	0.028	-0.0014	0.0004
				-1.00	0.240	0.0000	0.0004

Table 4.1: Numerical values (with uncertainties) of superfluid fraction f_s in thermodynamic limit, for phase space points indicated in Fig. 4.8.

within its error interval. For the one exceptional case, the thermodynamic limit f_s is an order of magnitude smaller than the finite-size values (as well as unphysically negative), indicating that the superfluid fraction should vanish in the thermodynamic limit. It is interesting to note that the finite-size scaling plot for the BH point $(\mu/U, t/U) = (0.5, 0.5)$ (not shown), which is situated above the particle boundary of the $n = 1$ lobe, is identical to the scaling of the $(0.0, 0.5)$ point shown in Fig. 4.8b. Similar behavior is found in the JCH model in the SF region to the right of the BKT point. Consider for example the point $(\tilde{\mu}, \kappa/g) = (-0.750, 0.240)$ located above the $(-1.000, 0.240)$ point in Table 4.1. The infinite-system superfluid fraction inferred from finite-size scaling is $f_s = 0.0006 \pm 0.0033$, consistent with zero. This indicates that the superfluid density is only weakly dependent on the chemical potential for a given hopping strength. The data strongly suggest that the superfluid density is zero throughout the low-density SF region studied.

Condensate fraction

The condensate fraction f_c is defined as the proportion of particles in the lowest-lying single-particle eigenstate of the system. In practice, this can be obtained by the largest eigenvalue of the single-particle density matrix $G^{(1)}(r, r')$ [127, 128]. In fact the entire spectrum, known as the entanglement spectrum [221], can be used to identify quantum phases. In the atomic limit, the single-particle ground state for a system of length L is L -fold degenerate, since there is no preferred lattice site. In the non-interacting limit, the finite-size systems have a sinusoidal single-particle ground state. At zero tempera-




Figure 4.9: Maximal eigenvalues of the reduced single-particle density matrix for $L = 31$. In (a), the five largest eigenvalues of the photon density matrix are shown for fixed $\tilde{\mu} = -0.678$ as a function of hopping amplitude κ/g . In (b), the largest eigenvalue of the density matrix (identified with the condensate fraction f_c in the SF phase) for spin excitations and photons in the JCH model, and for bosons in the BH model, are shown as a function of generalized hopping J/J_0 at fixed chemical potential $\tilde{\mu} = -0.678$ and $\mu/U = 0.55$, respectively. The inset shows the lower density case $\tilde{\mu} = -0.922$ and $\mu/U = 0.133$ for comparison.

ture, one expects macroscopic occupation of the ground state in the SF regime, tending towards unit occupation for large J/U . In the MI regime one expects the particles to be distributed evenly across many nearly-degenerate single particle states, a phenomenon known as fragmentation for large occupation of a single site [222, 143]. If the particles have fermionized, however, then no such macroscopic occupation should occur; f_c can therefore be viewed as another signature of fermionization.

The five largest eigenvalues of the single-particle photon density matrix are plotted in Fig. 4.9(a) in the JCH model for $\tilde{\mu} = -0.678$ and $L = 31$ as a function of hopping strength κ/g . This corresponds to the constant $\tilde{\mu}$ line near the very top of Fig. 4.1(a). In the zero-hopping limit $\kappa/g \rightarrow 0$, each photon is perfectly localized to each site, and the eigenvalues are precisely $1/L$ (note that the density matrix is normalized to unity rather than the

total number of particles). As κ/g increases the degeneracy is broken and the largest eigenvalues increase due to the fluctuations of the site occupations, while others decrease to preserve the normalization (not shown). At a critical hopping strength $\kappa/g \approx 0.085$ coinciding with the superfluid transition at this value of $\tilde{\mu}$, one of the eigenvalues increases precipitously relative to the others, signifying the macroscopic occupation of a single mode. This eigenvalue is associated with the (quasi)condensate fraction. For larger hopping strengths the condensate fraction tends towards unity, as expected for a non-interacting Bose gas at zero temperature.

Fig. 4.9(b) compares the values of the condensate fraction for the BH bosons with the JCH photons and spins along the same line of constant chemical potential considered above, corresponding to $\mu/U = 0.55$ in the BH model (recall that $\mu/U \sim 1.707(\tilde{\mu} + 1)$ as discussed in Sec. 4.3.2). The value of f_c for each model is plotted for fixed size $L = 31$ as a function of the hopping J , where $J = \kappa$ (t) in the JCH (BH) case, in units of the minimum hopping considered J_0 corresponding to $\kappa_0/g = 0.00556$ and $t_0/U = 1.7\kappa/g = 0.00948$. Though the onset of superfluidity occurs for smaller J/J_0 in the BH case, consistent with Fig. 4.1, the condensate fraction does not increase as quickly as that of the photons in the JCH case. This might simply reflect the fact that the mean carrier density is lower for the BH model at the top right point in the phase diagram than for the JCH model, which would discourage condensation. At lower mean particle densities on the SF side $\tilde{\mu} = -0.922$ ($\mu/U = 0.133$), shown in the inset of Fig. 4.9(b), the f_c for the BH case is larger than that for the photons of the JCH model, but neither reach 50%. The finite-size results suggest that f_c tracks the mean particle density, as is discussed further below.

Interestingly, the spin excitations in the JCH model also show strong evidence of condensation. The value of f_c on the SF side reaches approximately half that for the photons for the largest value of J/J_0 considered for this value of $\tilde{\mu}$. For lower values of $\tilde{\mu}$ in the vicinity of the $n = 0$ Mott lobe, the ratio of f_c for spin excitations to photons

approaches unity (not shown). These results are generally consistent with observations above that indicate that the spin and photon degrees of freedom follow each other closely. The condensation in the spin sector therefore appears to be driven sympathetically by the photons via the polariton excitations. The results are nevertheless somewhat surprising, suggesting that the spin excitations are delocalized.

To obtain the condensate fraction in the thermodynamic limit, the values of f_c were obtained throughout the phase diagram for each triple (μ, J, L) , where $L = 15, 19, 23, 27, 31$ using DMRG subject to open boundary conditions. The finite-size scaling analysis was performed for the SF regime only because the procedure is not robust for the points in the MI regime. Nevertheless, the finite-size results within the MI regime clearly indicate fragmentation, and since the macroscopic degeneracy in the atomic limit is true independent of system size, there is no reason to believe that the results in the thermodynamic limit would be qualitatively different.

The condensate fractions for JCH and BH models in the thermodynamic limit throughout the explored SF regime of the phase diagram are shown in Fig. 4.10(a) and (b), with the MI lobes explicitly zeroed out. Both pictures have the same color scale and are plotted with the axes corresponding to equivalent energy scales. The results are qualitatively similar for both models, but there are some quantitative differences. The value of f_c rises more rapidly with increasing hopping in the JCH model than in the BH model, attaining $f_c = 0.800$ at the highest mean densities investigated, as opposed to only $f_c = 0.500$ for the BH model. The values of f_c shown in the phase diagrams strongly resemble the mean densities of spins, photons, and bosons, such shown in Fig. 4.1. Deep in the SF regime the f_c closely follow the mean excitation densities, suggesting that the proclivity toward condensation is governed by the mean density. That said, as the Mott lobe boundary approaches the ratios of f_c to the respective mean densities is found to increase markedly even as f_c approaches zero.




Figure 4.10: Condensate fraction for the (a) JCH model (photons only) and the (b) BH model outside but near $n = 1$ MI lobe, in the thermodynamic limit. In both cases, the condensate fraction within the MI lobes is zeroed out.

While there need be no direct relationship between Bose-Einstein condensation and superfluidity, it is nevertheless somewhat surprising that the condensate fraction in the thermodynamic limit would be so large throughout the SF region where the superfluid fraction remains zero. These results are nevertheless consistent with the small values of α in the SF region, shown in Fig. 4.6 (recall that $\alpha \rightarrow 0$ for Bose-Einstein condensates), as well as with the filling in of the exclusion hole and the disappearance of the Friedel oscillations found in the spatial-dependence of the two-body correlation function (c.f. Fig. 4.7). In addition, the value of f_c is consistently small in the low-density SF regime, as expected for a fermionized gas.

Entanglement properties

Much can be learned about the ground states of physical systems by examining the properties of subsystems. A notable example is the entropy of entanglement [19]. This is obtained by partitioning the system into a block of contiguous lattice sites $b \subset \mathcal{L}$, where $\mathcal{L}$ denotes the full lattice, and its complement b' . The entropy of entanglement associated

with this bipartition is given by

$$S_{\rho_b} = -\text{Tr}(\rho_b \ln \rho_b), \quad (4.20)$$

where ρ_b is the reduced density matrix of the state over b . Analytical formulas for the entanglement entropy of non-interacting fermions and bosons on a lattice exist for the semi-infinite chain [223] but we are not aware of any for finite-size systems greater than a few sites.

The scaling of the entanglement entropy with the size of the subsystem is intimately linked with the utility of DMRG as a simulation method. For a wide array of physical systems, the entanglement entropy obeys an area law [43], meaning that the entanglement entropy S_ρ associated with ρ_b is proportional to the number of sites at the interface between b and b' (henceforth denoted l), rather than to its volume or cardinality. In one dimension, the area law corresponds to a value of S_ρ that saturates for some finite value of l :

$$S_\rho(l) \leq D, \quad (4.21)$$

where D is a constant, independent of l (of course $l = 2$ is itself constant). Only states satisfying Eq. (4.21) can be efficiently simulated by DMRG for large system sizes, because the variational ansatz used by the DMRG algorithm explicitly assumes that the area law is satisfied [40, 43]. Asymptotic scaling results [224] reveal that free fermions always logarithmically violate the area law in any dimension in both the continuum and on a lattice, whereas free bosons in 1D satisfy the area law away from criticality [225].

The entanglement entropy associated with a finite block can be used to distinguish bosonic and fermionic behavior. The entanglement entropy of the 1D non-interacting Bose gas is a smooth function, whereas that for the non-interacting Fermi gas oscillates with l . This is because at zero temperature the bosons condense into the smooth lowest-

lying eigenstate of the hopping model, while the Pauli exclusion principle forces fermions into oscillatory excited states. Generally, the entanglement entropy for bosons is larger than for fermions because the number of accessible states Ω_b is exponentially greater than Ω_f , meaning that when the system is bipartitioned, for each configuration of the left half of the chain there are an exponentially larger number of compatible configurations for the right half in the bosonic case.

The entanglement entropy $S_\rho(l)$ is plotted in Fig. 4.11 as a function of l for $3 \leq l \leq 29$ ($L = 31$) along a contour of constant density in the SF regime. The results are displayed for both the JCH and BH models at two different mean densities: $n = 0.645$ and $n = 1.129$. The corresponding entropy profiles for the free Bose and Fermi gases at the same mean density, calculated numerically, are plotted for comparison. The bosonic entanglement entropy increases approximately linearly with l for $l \ll L$. The fermionic entanglement entropy displays strong oscillations but much weaker l -dependence. At the lower mean density considered, the photon entanglement entropy profile is very similar to the ideal Fermi gas for small hopping, again providing strong evidence for fermionization. As the hopping increases, the oscillations become less pronounced and the value of the entropy increases; presumably the entropy profile would converge to that of the ideal Bose gas for very large hopping amplitudes. For the $n > 1$ case, which precludes fermionization in this single-band model, the oscillations are almost completely washed out. The magnitude of entanglement entropy remains well below the ideal Bose gas limit, however.

While the entanglement entropy can be used to help distinguish quantum phases, it does not directly quantify the possible use of this entanglement for quantum computation. Indeed, it is difficult to conceive of how one could encode quantum algorithms into the (generally delocalized) indistinguishable bosons of the BH model. The JCH model, on the other hand, has distinguishable quantum registers in the spin states of the two-level




Figure 4.11: The photon and boson entanglement entropy $S_\rho(l)$ for a contiguous block of sites of length $l = 3$ to 29 with $L = 31$, for the JCH and BH models, respectively. Two mean densities are considered, $n = 0.645$ for the (a) JCH and (c) BH models, and $n = 1.129$ for the (b) JCH and (d) BH models. The solid blue line is the exact value of S_ρ^b for an ideal lattice Bose gas, and the solid red line in (a) and (c) is the corresponding value S_ρ^f for the ideal Fermi gas.

atoms (i.e. qubits) which are each localized to a different optical cavity. If entanglement were generated between cavity atoms by the itinerant photons, then the coupled cavity QED could potentially be a natural environment for quantum computation.

It was recently proven that universal quantum computation is possible as long as the entanglement associated with arbitrary bipartitions is non-vanishing [226]. One useful measure that is closely related to the entanglement entropy is the localizable entanglement [227, 228, 88]. The localizable entanglement E_{ij}^{LE} between qubits i and j of a multiqubit system is defined as the maximal entanglement that can be concentrated between qubits i and j via local (i.e. single-spin) operations and classical communication, and is a non-negative number in the range $[0, 1]$. In a spin network, the localizable entanglement is lower-bounded by the maximal absolute value of the spin-spin correlation function $C_{ij}^{\alpha,\beta} := \langle \sigma_i^\alpha \sigma_j^\beta \rangle - \langle \sigma_i^\alpha \rangle \langle \sigma_j^\beta \rangle$ over all possible axes $\alpha, \beta \in \mathbb{R}^3$ [88, 229]. Using $\sigma_i^z = \sigma_i^+ \sigma_i^- - I/2 = n_i^{\text{spin}} - I/2$ where I is the 2×2 identity matrix, and if $\alpha = \beta = z$, the localizable entanglement is at least as large as

$$|C_{ij}^{zz}| \equiv |\langle N_i N_j \rangle - \langle N_i \rangle \langle N_j \rangle|,$$

where here $N_k = \sigma_k^+ \sigma_k^-$ is the local spin excitation number operator, satisfying $0 \leq \langle N_k \rangle \leq 1$. The localizable entanglement is therefore closely related to the (unnormalized) two-body correlation function, defined in Eq. (4.11). For $i \neq j$ (the only case of interest for entanglement), one obtains $C_{ij}^{zz} = G^{(2)}(i, j) - \langle N_i \rangle \langle N_j \rangle \approx G^{(2)}(i, j) - n^2$ (the local excitation number is approximately equal to the mean density – the total number of excitations over all sites – if the density is almost constant).

The behavior of the spin-spin correlation function in the thermodynamic limit can be estimated under the assumption that the spin excitations have completely fermionized, using Eq. (4.14) for a ring geometry or Eqs. (4.15-4.16) for open boundary conditions.

Consider the symmetric pair of sites $i_{\mp} = (L \mp a)/2$, separated by a . Using either geometry one obtains $G^{(2)}(i_{-}, i_{+}) \rightarrow n^2 [1 - \sin^2(\pi na)/(\pi na)^2]$, assuming short-ranged correlations in the bulk where a remains constant as L increases; for small n , the spin-spin correlation function approaches $-n^2$. If the separation instead scales like $a \sim L$, then $G^{(2)}(i_{-}, i_{+}) \rightarrow n^2$ and $C_{i_{-}i_{+}}^{zz} \rightarrow 0$.

The correlation function C_{ij}^{zz} was calculated throughout the phase region investigated. Restricting i and j to lie between $\lceil L/4 \rceil$ and $\lceil 3L/4 \rceil$ to avoid boundary effects, $|C_{i,j}^{zz}|$ take the largest values when $|i - j| = 1$ ($i = j$ is explicitly excluded). The values of $C_{i,i\pm 1}^{zz}$ were found to be close to zero everywhere in the MI phase, consistent with the strong density localization and small density fluctuations which are the hallmark of MI states. Likewise, the values of $C_{i,i\pm 1}^{zz}$ tend to zero for all i at high mean densities and hopping where the BEC fraction is large. This reflects the smooth and relatively constant profile of the two-body correlation function other than the remnant of the exclusion hole at short distances, as shown in Fig. 4.7.

At low density and hopping amplitude in the SF regime, the magnitudes of $C_{i,i\pm 1}^{zz}$ generally range between 0.02 to 0.05. This is a direct consequence of the large exclusion hole in the spin density-density correlation function in the vicinity of $i \sim j$, inherited from the strongly fermionized polaritons. The spin-spin correlation function reaches a maximum value in the region directly adjacent to the hole boundary of the $n = 1$ Mott lobe. At the point $(\tilde{\mu}, \kappa/g) = (-0.989, 0.0167)$ where the mean spin excitation density is $n = 0.3825$, the numerics yield $|C_{i,i\pm 1}^{zz}| = 0.075$. The fermionized theory predicts the comparable but slightly larger value of $\sin^2(\pi n)/\pi^2 \approx 0.088$. The difference between theory and computation is likely partly due to the fact that the spin excitations have not perfectly fermionized.

Because we have only considered the zz quadrature of the spin-spin correlation function, the numerical value is a lower-bound to the lower-bound of the localizable entangle-

ment. The actual value of the localizable entanglement in the Tonks-Girardeau regime could well be larger. In any case, the numerical results suggest that there is sufficient entanglement between the atoms in the ground state of coupled cavities to support universal quantum computation, if a suitable strategy to embed this environment in a quantum circuit could be found.

4.5 Conclusions

In this work, we have explored the Mott-Insulator to superfluid transition of the one-dimensional Jaynes-Cummings-Hubbard (JCH) model in the strong-coupling regime with no detuning. The purpose of the work has been two-fold. First and foremost, it has been to study the nature of the ground state in the vicinity of the phase transition, in particular to demonstrate that the photons are in fact strongly fermionized in the low-density ‘superfluid’ phase. Second, it has been to compare and contrast the properties of the ground state to that of the 1D Bose-Hubbard (BH) model, in order to highlight the unique features of the JCH model. The results were obtained by finding the ground state using the finite-system Density Matrix Renormalization Group, computing various static properties, and then performing a finite-size scaling analysis to infer the thermodynamic limit.

The main result is that in one dimension, the ground state of the JCH model in the low-density regime outside the Mott-insulating lobes is dramatically different from that of a conventional superfluid. Rather, the system in the region widely characterized as a superfluid (SF) is in fact a Tonks-Girardeau gas of strongly fermionized excitations, much as occurs in the one-dimensional BH model. This is evidenced by the power law of the single-particle density matrix in both the spin and photonic sectors of the model, by the Friedel oscillations and the fermionic exclusion hole in the two-body density matrix

for each sector, and by the strongly fermionic profile of the entanglement entropy. Thus coupled cavity QED provides a natural and accessible environment for the realization of strongly correlated photons. The photon fermionization should be readily observable in experiments using standard photon correlation spectroscopy [230].

We have calculated the superfluid fraction for both species of excitation in the JCH model, both within the low-density SF regime as well as for large tunneling beyond the expected BKT transition point, and found that the superfluid density vanishes in the thermodynamic limit in both cases. At the same time, we have found that the Bose-Einstein condensate fraction for the photons and spin excitations is non-vanishing throughout the SF region. The value of the condensate fraction is low at very low densities, consistent with fermionization, but can reach as high as 80% at high densities for the photons. This indicates the existence of a (quasi)condensate of photons and even spin excitations, despite the absence of superfluidity.

The same static properties for the 1D BH model have been calculated in an equivalent parameter regime in phase space, and the results have been compared in detail to those obtained within the JCH model. Broadly, the behavior of the two models coincide: neither exhibits a true Mott-insulator to conventional superfluid transition. However, various quantitative differences exist. The single-particle correlations in the JCH ground state within the $n = 1$ Mott lobe decay more rapidly with increasing hopping than in the BH case, while those in the intermediate region between the $n = 0$ and $n = 1$ lobes decay more slowly. Consistently with this result, the condensate fraction for the JCH photons rises more rapidly with increasing hopping than for the BH bosons. While a formal mapping between the BH and JCH models does not exist, the results indicate that the well-known manifestations of the BH model will be largely reproduced in physical systems that are well-described by the JCH model.

That said, the JCH model has some intriguing features not shared by the BH model,

owing to the presence of two species. The spin and photon degrees of freedom are inextricably linked through the fundamental polariton excitations. Thus Bose-Einstein condensation of photons implies that the spin excitations are similarly condensed. The atomic spin states are thereby effectively delocalized across the entire system in spite of the fact that each atom is confined to its respective cavity. This observation opens the intriguing possibility of inducing spin liquid-like states in cavity QED systems. In fact, the possibility of spin dimerization in the JCH model with large positive detuning $\delta \gg 0$ (which induces frustration and is not considered in the present work) has been noted very recently [231]. Likewise, atoms in different cavities are spontaneously entangled via the itinerant photons, as evidenced by the non-zero value of the localizable spin entanglement in the SF regime. It would be intriguing to systematically consider the effect of non-zero detuning on the properties of the ground states, but this is beyond the scope of the present work.

The majority of the parameter space explored in this project is in the vicinity of a phase transition, but the convergence of the DMRG algorithm is not guaranteed very close to the phase boundary. It could be fruitful to compare the results with those obtained using a method such as Multiscale Entanglement Renormalization Ansatz [232], which is specifically tailored to work well with critical systems. In a similar vein, the calculations (in particular the superfluid density) could be repeated with a method that is designed for handling infinite systems directly, such as iDMRG [233, 234]. This would avoid the need to perform finite-size scaling on small systems to make quantitative statements about the thermodynamic limit.

Cavity quantum electrodynamics is a promising candidate for quantum information processing applications, since the local atoms can be used as qubits and the photons can be used to generate entanglement between them. Our results suggest that the localizable entanglement is always finite throughout the SF region. It would be interesting

to determine if the ground state for a more complex network of coupled cavities could be a resource for measurement-based quantum computation, where universal quantum algorithms are effected solely via single-qubit measurements conditioned on previous outcomes [235]. Preliminary calculations indicate that the type of correlations between atoms in the current 1D JCH model with zero detuning are probably not suitable for gate teleportation via measurements. This will be explored more fully in future work.

4.6 Acknowledgements

The authors are grateful for research funding from the Natural Sciences and Engineering Research Council of Canada and Alberta Innovates – Technology Futures. This research has been enabled by the use of the ALPS software package, particularly the `dmrg` application, as well as by computing resources provided by WestGrid and Compute/Calcul Canada.

Chapter 5

Outlook

5.1 Introduction

The goal of this final chapter is to tie together the seemingly disparate threads of, on the one hand, novel resource states for MBQC (Chapters 1 and 2), and on the other hand, cavity QED systems (Chapters 3 and 4). The original motivation for pursuing the latter research project was the possibility that the ground state of a coupled-cavity QED system could potentially be a resource state for universal MBQC. The results and ideas presented in this Chapter are of a preliminary nature, and much more work would be needed on this front if time permitted. The general idea is to use DMRG to produce an MPS representation of the ground state of the 1D JCH model, and then check if this ground state is a universal resource for single-qubit gates. This is just a starting point, because true universal computing will require entangling operations as well; extending the study to two dimensions is an important future direction. Previous work [46, 47, 50, 85, 51, 72, 55, 38] has shown that coupling 1D states together to form a 2D resource is often possible.

In Sec. 5.2, I make some general remarks on quantum computing with cavity QED systems, and why it might be preferable to optical lattice implementations of the bare Bose-Hubbard model. Then, in Sec. 5.4 I describe two important features that a desirable MPS representation of the candidate resource state should have, and outline the work that I have done to produce such a representation for the MI phase of the 1D JCH model, using infinite-system DMRG. Finally, in Sec. 5.5, I describe the prospects for universal MBQC with this class of ground states, and outline future research directions that could

prove fruitful.

5.2 Universal MBQC with Jaynes-Cummings-Hubbard model?

In Chapters 3 and 4, I discussed at length the Jaynes-Cummings-Hubbard model, which is for various reasons considered to be a good candidate for simulating the Bose-Hubbard model. This application falls within a large area of study known as quantum simulation [13, 236, 14, 19, 237], which is one of the key potential applications of quantum computing, and was in fact the principal motivation for the original proposals of quantum computing by Feynman [13] and Deutsch [236].

The ability to simulate a specific quantum system with some other system is a weaker requirement than being able to perform universal quantum computing with the latter. Indeed, useful quantum simulations can be achieved with a modest (~ 36) number of qubits, while useful instances of, say, Shor's algorithm [10], require millions of qubits [237]. An interesting question to ask, therefore, is whether universal quantum computing is possible with many-body systems. While quantum information techniques are increasingly being used to shed light on the physics of many-body systems, the loosely-converse problem of quantum computing with such systems has received relatively little attention.

One natural model that could be considered for this purpose is the Bose-Hubbard model discussed at length in Ch. 3. Because this model is so well-known and intensely studied already, it would be fascinating to know if it also has a direct application to MBQC. One could imagine, for example, a many-body state of two-level bosons acting as qubits as the substrate for such a computation. An immediate issue is the fundamental indistinguishability of the individual bosons, which makes a straightforward identification between two-level bosons and qubits unlikely. Quantum computing with indistinguishable particles has been discussed before, most notably in the case of topological computing

with a type of quasiparticle known as non-abelian anyons [25]. Universal quantum computing with the Bose-Hubbard model [238] and with free lattice fermions [239] and other many body fermionic states [240] has also been discussed; in these cases information is encoded in the occupation number states of the model. That is not the goal of the present chapter, which is to examine the possibility of MBQC with a cavity-QED system, with the information encoded in distinguishable two-level systems.

It so happens that the optical microcavities of the Jaynes-Cummings-Hubbard model each contain a distinguishable two-level system. While there is already intense interest in implementing this model because of its promise as a quantum simulator, it would be potentially groundbreaking if one could demonstrate universal quantum computing as well. In fact, several groups have studied cavity-QED networks for the purpose of distributed quantum computing [241, 242, 243, 244, 245, 246]. In these investigations, local nodes comprising one or more atomic qubits coupled to light are interconnected via photons, with the computation proceeding via a distributed version of the circuit model. But is there a possibility of using a coupled-cavity system for universal measurement-based quantum computing as well? This question is significant because, as discussed in Ch. 1, there is intrinsic interest in proposing different kinds of resource states for MBQC. To my knowledge, there does not presently exist a proposal for MBQC with atomic qubits, whose mutual entanglement is generated via photons.

As I discussed in Chapters 1 and 2, one prominent method for demonstrating universal MBQC is to find a tensor network representation (MPS or PEPS) for a candidate resource and explicitly provide a measurement pattern that efficiently simulates (in the quantum sense, not the classical) any arbitrary quantum circuit. Since 1D resources can often be coupled together to form 2D universal resources, perhaps a good place to begin is to analyse a 1D model – find an MPS representation for it via an appropriate numerical simulation method (which will turn out to be DMRG) and see if a measurement scheme

implementing a universal set of single-qubit gates can be designed. It may seem contradictory that on the one hand I say universal resources cannot be classically simulated efficiently, but on the other hand I propose to find the MPS representation by efficiently simulating a 1D model with DMRG. The key is that the proposed 1D state is only a resource for single-qubit rotations, not for universal MBQC. However, it is nevertheless an ingredient for a true universal resource state.

The first step in the plan to design a MBQC scheme with the JCH model is to obtain a MPS representation for its ground state, via the DMRG algorithm. The ALPS implementation of DMRG does not provide user access to the internal representation of the ground state, which is needed for determining the MPS representation. I therefore implemented my own infinite-system DMRG algorithm in MATLAB, following the reduced density matrix construction of Alg. 5.3. The MPS matrices for sites situated a distance j from the centre of the chain are obtained from the matrix elements of the isometries $T_{L,R}^{(j)}$. I will now provide a detailed description of the infinite- and finite-system DMRG algorithms.

5.3 Infinite-System Density Matrix Renormalisation Group

The general philosophy behind DMRG was discussed in Sec. 3.10. In order to understand the mechanics of the algorithm, it is necessary to specify the space in which the system is defined for each iteration. One begins by considering a system of size L_0 which is small enough that it can be exactly solved, and finding the exact ground state. The size of the system increases by two sites with each iteration; let $L^{(j)} = 2(j-1) + L_0$ be the number of lattice sites at the j th iteration, where initially $j = 1$. Suppose that the Hamiltonian for $L^{(j)}$ sites is denoted by $H^{(j)}$. Now, subdivide the 1D lattice $\mathcal{L}$ into the leftmost $L^{(j)}/2 - 1$ sites B (called the Block), the rightmost $L^{(j)}/2 - 1$ sites U (the Universe)




Figure 5.1: Schematic of infinite-system DMRG iteration with $B-l-r-U$ superblock structure, reproduced from Ref. [7].

and the two sites at the centre l, r (the left and right) sites, so that $\mathcal{L} = B \cup l \cup r \cup U$. More informally, the structure of the lattice is denoted $B-l-r-U$ (this is called the superblock structure). A schematic of the infinite-system DMRG iterations with this superblock structure is depicted in Fig. 5.1. Other superblock structures are sometimes useful as well; for example, $B-l-U-r$ is usually used for systems with periodic boundary conditions, but the essence of the algorithm is unaffected by different choices (although the efficiency may be). Denote the Hilbert space for the system at the j th iteration by

$$\mathcal{H}^{(j)} := \mathcal{H}_B^{(j)} \otimes \mathcal{H}_l^{(j)} \otimes \mathcal{H}_r^{(j)} \otimes \mathcal{H}_U^{(j)}, \quad (5.1)$$

with the obvious meanings for each of the local Hilbert spaces.

For concreteness, suppose that the exact Hamiltonian is translationally invariant, is subject to open boundary conditions and consists of a sum of one-body and nearest-neighbour two-body terms only:

$$H^{(j)} = \sum_{i=1}^{L^{(j)}-1} A_i B_{i+1} + \sum_{i=1}^{L^{(j)}} C_i. \quad (5.2)$$

These assumptions are made for simplicity's sake; it is obvious how to generalise to a sum of several two-body terms acting on the same pair of sites, or incorporate k -local interactions. In terms of the superblock structure, the exact Hamiltonian can be written as

$$H^{(j)} = H_B^{(j)} + H_l^{(j)} + H_r^{(j)} + H_U^{(j)} + H_{Bl}^{(j)} + H_{lr}^{(j)} + H_{rU}^{(j)}, \quad (5.3)$$

where the “local” terms are

$$H_B^{(j)} = \sum_{i=1}^{L^{(j)}/2-2} A_i B_{i+1} + \sum_{i=1}^{L^{(j)}/2-1} C_i; \quad (5.4)$$

$$H_l^{(j)} = C_{L^{(j)}/2}; \quad (5.5)$$

$$H_r^{(j)} = C_{L^{(j)}/2+1}; \quad (5.6)$$

$$H_U^{(j)} = \sum_{i=L^{(j)}/2+2}^{L^{(j)}-1} A_i B_{i+1} + \sum_{i=L^{(j)}/2+2}^{L^{(j)}} C_i, \quad (5.7)$$

and the terms coupling different parts of the superblock structure are

$$H_{Bl}^{(j)} = A_{L^{(j)}/2-1} B_{L^{(j)}/2}; \quad (5.8)$$

$$H_{lr}^{(j)} = A_{L^{(j)}/2} B_{L^{(j)}/2+1}; \quad (5.9)$$

$$H_{rU}^{(j)} = A_{L^{(j)}/2+1} B_{L^{(j)}/2+2}. \quad (5.10)$$

The dimension of the left and right points satisfies $\dim [\mathcal{H}_l^{(j)}] = \dim [\mathcal{H}_r^{(j)}] = d$, which is constant. However, $\dim [\mathcal{H}_B^{(j)}] = \dim [\mathcal{H}_U^{(j)}] = d^{L^{(j)}/2-1}$ grows exponentially in j , so no computer will be able to store a completely faithful representation of Block or Universe operators for more than a few iterations. Furthermore, even if the amount of time required to find the ground state of a $p \times p$ sparse Hamiltonian is polynomial in p , the exponential explosion of the dimension of the exact Hilbert space in terms of

j ensures that finding the ground state of $H^{(j)}$ will be prohibitive for $j \gg 1$. In order to solve this problem, one truncates the Hilbert space of B and U at every iteration, so that the dimension is $D^{(j)} = \min\left(d_B^{N^{(j)}}, \chi\right)$, where χ is a constant. As a result, the dimension of the effective Hilbert space for the entire system is never larger than $\chi^2 d^2$, so the complexity of finding the ground state of $H^{(j)}$ is independent of j .

The idea in infinite-system DMRG is to grow the system with each iteration by absorbing l into B and r into U , and then placing two new sites (the new l and r) within an environment (the new B and U) that mimics the environment a new particle or site inserted into the real system would experience. In the $B-l-r-U$ superblock structure, the new sites are thus inserted in the centre of the chain. At each step, the Hilbert space describing the combined subsystem $B-l$ is projected onto the subspace spanned by the χ eigenvectors of the reduced density matrix of the ground state for $B-l$ (obtained by tracing out r and U), for the current system size. The operator used to perform this projection is called an isometry and the process itself is sometimes called block decimation. The Hilbert space for $r-U$ is obtained by assuming reflection symmetry and reflecting the new basis states for $B-l$ about the centre of the chain.

Information about the true system is irreversibly lost via the block decimation procedure, which is essentially a lossy compression algorithm; at the j th iteration, the algorithm produces approximate (compressed) representations $\tilde{H}_B^{(j+1)}$, $\tilde{H}_U^{(j+1)}$, $\tilde{H}_{Bl}^{(j+1)}$ and $\tilde{H}_{rU}^{(j+1)}$, which are used in the $j+1$ -th iteration to construct an approximation $\tilde{H}^{(j+1)}$ of the full Hamiltonian. The steps for the full algorithm are outlined in Algorithm 5.3, which uses Algorithms 5.1 and 5.2 to determine the isometries used to truncate the basis from the Hamiltonian.

As presented here, the algorithm provides an estimate for the ground state energy per particle, but any ground-state expectation value of any operator can be calculated; one must simply ensure that a representation of this operator is maintained in the same

truncated, “lossy” basis as the Hamiltonian and ground state. Several details that are required for an efficient implementation of DMRG are omitted; they can be found in the reviews [7, 157]. A wide variety of convergence criteria can be used; for example, that the value of the ground state energy per particle differs from its value in the previous iteration by less than some specified tolerance.

In fact, the convergence criterion for DMRG based on the ground state energy per particle is likely inadequate. I actually began the numerical simulations from Ch. 4 with my own infinite-system DMRG code in MATLAB. I found that it was relatively easy to achieve convergence in local properties such as the expectation values of the site occupation numbers. Even calculating correlation functions such as $G^{(1)}(r)$ and $G^{(2)}(r)$ met with reasonable success in the MI regime. However, the representation of correlation functions like $G^{(1)}$ and $G^{(2)}$ was very poor in the SF regime; the correlation functions I obtained were not smooth-looking in real space, were not translationally invariant for sites in the bulk, and were not even symmetric about the centre of the chain. I tried many modifications to my code to improve the situation, but ultimately concluded that the problem was the use of the infinite-system DMRG algorithm itself, rather than its finite-system variant, at which point I became aware of and decided to use the `dmrg` application from the ALPS library. Nevertheless, since we are explicitly working in the MI regime for the purposes of this chapter, which is gapped and has exponentially-decaying correlations, I do not expect the use of the infinite-system algorithm to be a terribly serious problem here.

Function LeftIsometry($\tilde{H}$)**Input:** Hamiltonian $\tilde{H}$ in $B - l - r - U$ structure, truncated dimension χ .**Output:** Isometry T_L for truncating $B - l$ basis.

/* Diagonalise Hamiltonian and obtain isometry */

 $|\psi\rangle \leftarrow$ ground state of $\tilde{H}$; $\rho \leftarrow |\psi\rangle\langle\psi|$; $\rho_{Bl} \leftarrow \text{Tr}_{r,U}\rho$; $\{\lambda_k\} \leftarrow \chi$ largest eigenvalues of ρ_{Bl} , with eigenvectors $\{|\lambda_{L,k}\rangle\}$; $T_L \leftarrow \sum_{k=1}^{\chi} |k\rangle\langle\lambda_{L,k}|$;**return** T_L **end****Algorithm 5.1:** LeftIsometry**Function RightIsometry($\tilde{H}$)****Input:** Hamiltonian $\tilde{H}$ in $B - l - r - U$ structure, truncated dimension χ .**Output:** Isometry T_R for truncating $r - U$ basis.

/* Diagonalise Hamiltonian and obtain isometry */

 $|\psi\rangle \leftarrow$ ground state of $\tilde{H}$; $\rho \leftarrow |\psi\rangle\langle\psi|$; $\rho_{rU} \leftarrow \text{Tr}_{B,l}\rho$; $\{\lambda_k\} \leftarrow \chi$ largest eigenvalues of ρ_{rU} , with eigenvectors $\{|\lambda_{R,k}\rangle\}$; $T_R \leftarrow \sum_{k=1}^{\chi} |k\rangle\langle\lambda_{R,k}|$;**return** T_R **end****Algorithm 5.2:** RightIsometry

```

Input: One-body operators  $A, B, C$  from Eqs. 5.4–5.10.
Output: approximate ground state energy per site of  $H^{(\infty)}$ .
 $\tilde{H}_B^{(1)} \leftarrow C, \tilde{H}_U^{(1)} \leftarrow C;$ 
 $\tilde{A}_B^{(1)} \leftarrow A, \tilde{B}_U^{(1)} \leftarrow B;$ 
 $j \leftarrow 1;$ 
while not converged do
    /* Diagonalise Hamiltonian and obtain isometries */
     $\tilde{H}^{(j)} \leftarrow \tilde{H}_B^{(j)} + C_l + C_r + \tilde{H}_U^{(j)} + \tilde{A}_B^{(j)} \otimes B_l + A_l \otimes B_r + A_r \otimes \tilde{B}_U^{(j)};$ 
     $T_L^{(j)} \leftarrow \text{LeftIsometry}(\tilde{H}^{(j)});$ 
     $T_R^{(j)} \leftarrow T_L^{(j)}$  with columns reflected;

    /* Block decimation procedure */
     $\tilde{H}_B^{(j+1)} \leftarrow T_L^{(j)} \left( \tilde{H}_B^{(j)} \otimes \mathbb{1}_l + \mathbb{1}_B \otimes C_l + \tilde{A}_B^{(j)} \otimes B_l \right) T_L^{(j)\dagger};$ 
     $\tilde{H}_U^{(j+1)} \leftarrow T_R^{(j)} \left( C_r \otimes \mathbb{1}_U + \mathbb{1}_r \otimes \tilde{H}_U^{(j)} + A_r \otimes \tilde{B}_U^{(j)} \right) T_R^{(j)\dagger};$ 
     $\tilde{A}_B^{(j+1)} \leftarrow T_L^{(j)} (\mathbb{1}_B \otimes A_l) T_L^{(j)\dagger};$ 
     $\tilde{B}_U^{(j+1)} \leftarrow T_R^{(j)} (B_r \otimes \mathbb{1}_U) T_R^{(j)\dagger};$ 

    /* Prepare for next iteration */
    Evaluate convergence criterion;
     $j \leftarrow j + 1;$ 
end
 $|\psi^{(j)}\rangle \leftarrow$  ground state of  $\tilde{H}^{(j)};$ 
return  $\langle \psi^{(j)} | \tilde{H}^{(j)} | \psi^{(j)} \rangle / (2j + L_0);$ 

```

Algorithm 5.3: Infinite-system DMRG

5.3.1 Finite-System Density Matrix Renormalisation Group

In most non-trivial situations, the infinite-system DMRG algorithm fails to provide satisfactory results [7]; while expectation values of local operators can be obtained, the behaviour of non-local correlation functions is not well-captured. Usually, much greater accuracy is obtained from a variant called the finite-system DMRG which, as may be expected, is used for calculating properties of the ground state of a finite-size system. In this setting, thermodynamic limit results can (and have, in many cases, like for example the 1D BH [172] and JCH models [6]) be obtained by extrapolating the finite-size results using finite-size scaling.

The basic idea of finite-system DMRG is very similar to the infinite-system version. Suppose the Hamiltonian to be considered has L sites, where for convenience I assume L is even. First, one starts with a very small system of length L_0 , which can be solved by exact diagonalisation. Then, one uses the infinite-size algorithm to grow the system to L sites (this is called the warmup phase). During this phase, the isometries $\{T_{L,R}^{(j)}\}$ and the truncated basis operators $\{\tilde{H}_B^{(j)}, \tilde{H}_U^{(j)}, \tilde{H}_{Bl}^{(j)}, \tilde{H}_{rU}^{(j)}, \tilde{A}_B^{(j)}, \tilde{B}_U^{(j)}\}$ must be saved for each value of j from 1 to $(L - L_0)/2$. Following this, a series of sweeps is performed, during which the size of the system remains constant. After the warmup phase, the Block and Universe each contain $L/2 - 1$ sites. A sweep consists of:

- (i) A left-to-right phase, where in each iteration the Block grows by one site and the Universe simultaneously shrinks by one, until the Universe contains only one site.
- (ii) A right-to-left phase, going in the opposite direction, until the Block has only one site.
- (iii) Another left-to-right phase, which ends when the Block and Universe once again have $L/2 - 1$.

During the first left-to-right phase, the stored isometries and operator representations from the warmup phase are used for the Universe (for the appropriate system size), and new isometries/representations are found and stored for each Block size from $L/2$ to $L-3$. Then, in the right-to-left phase, the Block representations obtained from the left-to-right and the warmup phases are used, while improved representations for the Universe from sizes 1 to $L/2 - 1$ and new representations from $L/2$ to $L - 3$ are found, overwriting the ones from the warmup phase as necessary. In the final left-to-right phase, improved Block representations for 1 to $L/2 - 1$ are found, overwriting the ones from the warmup phase. Multiple sweeps are performed until an appropriate convergence criterion chosen by the user is satisfied. The full procedure is found in Algorithm 5.4.

5.3.2 DMRG as a Variational Method over MPS

The reason that DMRG is so widely applicable in 1D, while generally being of more limited utility in higher dimensions, is best explained by the connection to the MPS formalism from Chapter 1. The connection was first pointed out by Östlund and Rommer in 1995 [162], who showed that the infinite-system DMRG algorithm (and, in fact, any simulation method using a block decimation scheme), produces a MPS as its ground state. Furthermore, if the DMRG has a fixed point, meaning that the isometry used for the block decimation procedure from iteration to iteration converges to some fixed operator, then the MPS matrices for the sites far away from the system boundaries are essentially site-independent.

Östlund and Rommer’s argument is ultimately very similar to Vidal’s argument [40], given in Section 1.5.1, for why states having low entanglement entropy with respect to any bipartition can be efficiently described via the MPS representation. The similarity arises from the fact that the partial trace operation (used to construct the isometries

```

Input: One-body operators  $A, B, C$  from Eqs. 5.4–5.10, length  $L$ .
Output: approximate ground state energy per site of  $H^{(L)}$ .

/* Warmup phase */
Grow system from 4 to  $L - 2$  sites via Algorithm 5.3, storing
 $\{\tilde{H}_B^{(j)}, \tilde{H}_U^{(j)}, \tilde{A}_B^{(j)}, \tilde{B}_U^{(j)}\}$  for each  $j \in \{1, \dots, L/2 - 1\}$ ;
 $s \leftarrow L/2$  /*  $s$  marks last site in Block */;
 $sweepDir \leftarrow right$ ;
/* Finite-system sweeps until convergence */
while not converged do
    while  $s \neq L/2 - 1$  OR  $sweepDir \neq right$  do /* one full sweep */
        /* Construct superblock Hamiltonian */
         $\tilde{H}^{(s)} \leftarrow \tilde{H}_B^{(s)} + C_l + C_r + \tilde{H}_U^{(s)} + \tilde{A}_B^{(s)} \otimes B_l + A_l \otimes B_r + A_r \otimes \tilde{B}_U^{(s)}$ ;

        /* Block decimation procedure */
        /* right sweep or end of left sweep */
        if  $sweepDir == right$  OR  $s == 4$  then
             $T_L^{(s)} \leftarrow \text{LeftIsometry}(\tilde{H}^{(s)})$ ;
             $\tilde{H}_B^{(s+1)} \leftarrow T_L^{(s)} (\tilde{H}_B^{(s)} \otimes \mathbb{1}_l + \mathbb{1}_B \otimes C_l + \tilde{A}_B^{(s)} \otimes B_l) T_L^{(s)\dagger}$ ;
             $\tilde{A}_B^{(s+1)} \leftarrow T_L^{(s)} (\mathbb{1}_B \otimes A_l) T_L^{(s)\dagger}$ ;
             $s \leftarrow s + 1$ ;
        end
        /* left sweep or end of right sweep */
        if  $sweepDir == left$  OR  $s == L - 4$  then
             $T_R^{(s)} \leftarrow \text{RightIsometry}(\tilde{H}^{(s)})$ ;
             $\tilde{H}_U^{(s-1)} \leftarrow T_R^{(s)} (C_r \otimes \mathbb{1}_U + \mathbb{1}_r \otimes \tilde{H}_U^{(s)} + A_r \otimes \tilde{B}_U^{(s)}) T_R^{(s)\dagger}$ ;
             $\tilde{B}_U^{(s-1)} \leftarrow T_R^{(s)} (B_r \otimes \mathbb{1}_U) T_R^{(s)\dagger}$ ;
             $s \leftarrow s - 1$ ;
        end

        /* Switch sweep direction if necessary */
        if  $s == 4$  then
             $sweepDir \leftarrow right$ ;
        else if  $s == L - 4$  then
             $sweepDir \leftarrow left$ ;
        end
        Evaluate convergence criterion;
    end
end
 $|\psi^{(s)}\rangle \leftarrow \text{ground state of } \tilde{H}^{(s)}$ ;
return  $\langle \psi^{(s)} | \tilde{H}^{(s)} | \psi^{(s)} \rangle / L$ ;

```

Algorithm 5.4: Infinite-system DMRG

$T_{L,R}$) is intimately related to the Schmidt decomposition described in Section 1.5.1. For a general bipartite system $|\Psi\rangle = \sum_{i,j} \chi_{i,j} |i\rangle_A |j\rangle_B$, suppose its Schmidt decomposition is given by

$$|\Psi\rangle = \sum_{\alpha} \lambda_{\alpha} |\Phi_{\alpha}^{[A]}\rangle |\Phi_{\alpha}^{[B]}\rangle. \quad (5.11)$$

The density matrix is

$$\begin{aligned} \rho &\equiv |\Psi\rangle\langle\Psi| \\ &= \sum_{\alpha,\beta} \lambda_{\alpha} \lambda_{\beta}^* |\Phi_{\alpha}^{[A]}\rangle |\Phi_{\alpha}^{[B]}\rangle \langle\Phi_{\beta}^{[A]}| \langle\Phi_{\beta}^{[B]}|. \end{aligned}$$

It is then a trivial matter to calculate the reduced density matrix, exploiting the orthonormality of the Schmidt bases:

$$\begin{aligned} \rho_A &\equiv \text{Tr}_B \rho \\ &= \sum_{\alpha,\beta} \lambda_{\alpha} \lambda_{\beta}^* |\Phi_{\alpha}^{[A]}\rangle \langle\Phi_{\beta}^{[A]}| \underbrace{\langle\Phi_{\alpha}^{[B]}| \Phi_{\beta}^{[B]}\rangle}_{\delta_{\alpha,\beta}} \\ &= \sum_{\alpha} |\lambda_{\alpha}|^2 |\Phi_{\alpha}^{[A]}\rangle \langle\Phi_{\alpha}^{[A]}|. \end{aligned}$$

The eigenbasis of the reduced density matrix for subsystem A is precisely the Schmidt basis for that system, and similarly for B . The eigenvalues are the squares of the Schmidt coefficients. Thus, the block decimation procedure corresponds to throwing away the subspace of the joint $B-l$ and $r-U$ subsystems spanned by the Schmidt basis vectors with the smallest Schmidt coefficients. Ultimately, therefore, the block decimation procedure produces a MPS representation of the ground state with bond dimension set by the DMRG truncation χ . Now, let's see exactly how this works.

Suppose the superblock basis at the j th iteration is labeled $\left\{ |\alpha_j \ s_j\rangle_{B-l} |s'_j \ \alpha_U^{[j]}\rangle_{r-U} \right\}$, where $1 \leq \alpha_j, \alpha_U^{[j]} \leq D$ and $1 \leq s_j, s'_j \leq d$. The isometries constructed from the reduced

density matrix for $B - l$ and $r - U$, are

$$T_L^{[j]} = \sum_{\alpha_j=1}^D |\alpha_j\rangle \langle \Phi_{\alpha_j}^{[B-l]}|;$$

$$T_R^{[j]} = \sum_{\alpha_j=1}^D |\alpha_j\rangle \langle \Phi_{\alpha_j}^{[r-U]}|,$$

where $\{|\alpha_j\rangle\}$ comprise the new basis vectors for the truncated block and universe, and $\{|\Phi_{\alpha_j}^{[B-l, r-U]}\rangle\}$ are Schmidt bases. Now, expand $|\alpha_j\rangle$ for the new Block in the composite basis for the $B - l$ subsystem:

$$|\alpha_j\rangle = \sum_{\alpha_{j-1}, s_{j-1}} \langle \alpha_{j-1} s_{j-1} | \alpha_j \rangle |\alpha_{j-1} s_{j-1}\rangle$$

$$:= \sum_{\alpha_{j-1}, s_{j-1}} \left(T_L^{[j] s_{j-1}} \right)_{\alpha_j, \alpha_{j-1}} |\alpha_{j-1} s_{j-1}\rangle.$$

Now, $|\alpha_{j-1}\rangle$ can similarly be expanded in terms of the $B - l$ basis of the previous iteration, and so on until the truncation process has been completely “reversed”. The result of doing all of this is that we recover an expansion for $|\alpha_j\rangle$ in the full product basis at the j th iteration, given by

$$|\alpha_j\rangle = \sum_{s_1, \dots, s_j=1}^d \sum_{\alpha_1, \dots, \alpha_{j-1}=1}^{\chi} \left(T_L^{[j] s_{j-1}} \right)_{\alpha_j, \alpha_{j-1}} \left(T_L^{[j-1] s_{j-2}} \right)_{\alpha_{j-1}, \alpha_{j-2}} \dots \left(T_L^{[1] s_1} \right)_{\alpha_1} |s_{j-1} s_{j-2} \dots s_1\rangle.$$

In the previous expression, the coefficient can almost be interpreted as a matrix product, except that one index α_j remains uncontracted. We can do a similar expansion for the new Universe basis states. Now, the ground state itself can be written in the

current Block-Universe basis (before inserting a new left and right site) as

$$|\psi^{[j]}\rangle := \sum_{\alpha_j, \alpha'_j} \psi_{\alpha_j, \alpha'_j} |\alpha_j\rangle_B |\alpha'_j\rangle_U. \quad (5.12)$$

Inserting the matrix product expressions for $|\alpha_j\rangle_B$ and $|\alpha'_j\rangle_U$ into the above equation, and noting that the sums over α_j and α'_j allow an interpretation of the coefficient as a completely contracted matrix product, we obtain

$$|\psi^{[j]}\rangle = \sum_{s_1, \dots, s_j=1}^d \sum_{s'_1, \dots, s'_j=1}^d T_L^{[j]s_j} \dots T_L^{[1]s_1} \Psi T_R^{[1]s'_1} \dots T_R^{[j]s'_j} |s_j \dots s_1\rangle |s'_1 \dots s'_j\rangle, \quad (5.13)$$

where we have defined $(\Psi)_{\alpha_j, \alpha'_j} := \psi_{\alpha_j, \alpha'_j}$. This is manifestly a MPS which, because of the dimensionality of the isometries, has bond dimension χ .

Subsequent to Östlund and Rommer, it was shown by Dukelsky et al. in 1998 that the finite-system DMRG not only produces MPS ground states, but that the function of the finite-size sweeping is to perform variational optimisations of the MPS matrices, leading to better and better approximations for the true ground state [163]. Subsequently, several authors (beginning with Vidal in 2004 [247]) used the variational MPS formulation of DMRG as a starting point for extending the algorithm in many ways; a comprehensive review as of 2011 can be found in [157].

It should be noted that the variational MPS view provides a clear physical reason as to why DMRG is so successful in 1D, but works less well in higher dimensions. It is because efficient MPS representations are available for states obeying an entanglement area law, which in 1D implies a constant bond dimension (independent of system size) is sufficient; this is not true in higher dimensions. Clearly, if a DMRG simulation is to be tractable, the bond dimension cannot be allowed to increase arbitrarily with system size. As mentioned earlier, it is a happy fact that the ground states of local Hamiltonians do

generically obey such area laws [43].

5.4 MPS representation for ground state of 1D JCH model

The first step in the plan to design a MBQC scheme with the JCH model is to obtain a MPS representation for its ground state, via the DMRG algorithm. For this purpose, I used my own implementation of the infinite-system DMRG algorithm, written in MATLAB. In addition to the computational overhead intrinsic to using an interpreted language like MATLAB rather than a compiled one, my implementation is also lacking in some of the optimisations that are needed for it to be truly competitive. Nevertheless, an extremely barebones implementation is likely adequate for simulations with very small superblock Hilbert space dimension. In fact, such a limited-size representation is beneficial for the present purposes. Additionally, it is desirable for the MPS representation to possess some form of translational invariance. I will elaborate on these points now.

5.4.1 Small local and bond dimension MPS representation for MI phase

For DMRG simulations of the JCH model, there are two different truncation parameters that need to be set. The first is the local truncation d (not to be confused with the number of spatial dimensions, which unfortunately is given the same symbol in Ch. 4). This is required for lattice boson models and artificially cuts off the maximum number of photons that are permitted in each cavity; without it, the photon occupation number for a single-cavity mode would be unbounded from above. The second is the truncation for the block decimation procedure χ , which is generic to DMRG irrespective of the system being simulated. The value of χ determines the Hilbert space dimension of the individual cavities, which in the MPS language corresponds to the number of matrices needed to describe each site. The value of d , on the other hand, sets the bond dimension of the MPS description, which corresponds to the dimension of each of the MPS matrices. In

short, each cavity will be described by d matrices of size $\chi \times \chi$.

There is unfortunately a tradeoff involving the magnitudes of d and χ . On the one hand the accuracy of the representation of the ground state monotonically increases in d and χ . On the other hand, the difficulty of designing a measurement pattern for a universal MBQC scheme also increases as d and χ increase. With the present methodology, keeping d and χ down to relatively small values is non-negotiable; a detailed study will therefore need to take into account that the DMRG ground state is only an approximation of the true one, and analyse the effects of the imperfect fidelity of the computed state with respect to the true ground state on the fidelity of the computation.

With respect to the bond dimension χ , the main consideration as to whether this can reasonably be kept small is whether the state obeys an entanglement area law. As long as one is not near the MI-SF phase transition, this appears to be true. An approximate assessment of this can be made by analysing the total weight of the discarded eigenvalues of the reduced density matrix in the block decimation procedure. Suppose the eigenvalues are denoted by $\{\lambda_i^2\}$, where $\{\lambda_i\}$ are the Schmidt coefficients associated with the equivalent bipartition of the system. At each iteration j , the total weight of the eigenvalues corresponding to the retained portion of the superblock Hilbert space is $w = \sum_{i=1}^{\chi} \lambda_i^2$, and the discarded weight is $\epsilon = 1 - w$. In the ideal case, when d and χ are sufficiently large for the MPS representation to be perfect, $\epsilon = 0$. This is exactly achievable whenever $\kappa = 0$, and it is plausible that it should be approximately achievable anywhere in any of the MI lobes, provided the system is not critical. I say it is plausible because the MPS description is capable of capturing exponentially-decaying correlation functions [157], which are generically present in the MI phase.

In my infinite-system DMRG implementation, if one sets the local truncation to artificially allow at most $n_{\max}$ polaritons in each cavity, one obtains a local dimension of $d = 2n_{\max} + 1$. Because the photon number fluctuations are so small in the MI regime,

setting $n_{\max}$ to an extremely modest value like 1 or 2 is sufficient to provide a rather good description of the ground state in the $n = 1$ MI lobe; the probability of ever finding more than $n_{\max}$ photons in a particular cavity is minuscule. Similarly, one would like as small a value of χ as possible, and quite a small value should provide a reasonably faithful representation of the ground state because of the existence of an energy gap. It appears that the choices $d = 5$ and $\chi = 3$ are sufficient for present purposes. Fig. 5.2 shows the maximum discarded weight ϵ over all iterations required to grow the system to size $L = 20$ for many different choices of d and χ , indicating that the maximal value of ϵ over all iterations, anywhere in the simulated region of the $n = 1$ MI lobe is of order 10^{-4} . The ground state energy of the state for $L = 20$ for the same region of parameter space is displayed in Fig. 5.3, again showing reasonably good convergence for the chosen values of d and χ . I should point out here that though it appears from Fig. 5.2 that $d = 5$, $\chi = 3$ is a good choice because the truncated weight is so small, this is almost certainly an instance when the truncated weight is not a good measure of convergence. Indeed, in Fig. 5.3, it is clear that the ground state energy is not properly captured with this choice of parameters, because the representation of the state improves monotonically as d and χ are increased, and the ground state energy curve for these parameter choices is nowhere near the curves with larger d and χ , in particular the curve for $d = 7$, $\chi = 25$ (the closest one to the exact curve).

5.4.2 Translational invariance: site-independent matrices

In Refs. [47, 54], there is a lot of flexibility in the resource states, but one common thread is that the entire state can be described with a small number of tensors. For example, with the AKLT-type resource, every spin-1 particle on the lattice has the same tensors $A[x]$, with $x \in \{0, 1, 2\}$, so only three different unique tensors appear in the tensor contraction (neglecting boundary conditions). Similarly, in the first toric-code-




Figure 5.2: Maximum discarded weight ϵ for all iterations up to system size $L = 20$ with $\tilde{\mu} = -0.878$, for various different values of κ/g , in the $n = 1$ MI lobe. For the choices $d = 5$, $\chi = 3$, ϵ is at most of order 10^{-4} , providing evidence that important basis states are not being discarded by the block decimation procedure of the infinite-system DMRG algorithm. Some of the data points are missing because in those cases, the truncated weights were extremely close to numerical zero.




Figure 5.3: Energy of ground state of 1D JCH model for system size $L = 20$ with $\tilde{\mu} = -0.878$, for various different values of κ/g , in the $n = 1$ MI lobe. This provides evidence that the choices $d = 5$, $\chi = 3$ for the infinite-system DMRG algorithm yield a reasonable approximation of the true ground state at this system size.

based example, the qubits reside on the edges of a square lattice, and correspond to tensors $K_H[s]$ or $K_V[s]$ with $s \in \{0, 1\}$ depending only on whether they are situated on horizontal or vertical edges. Thus, only four unique tensors appear in this contraction as well. It is undesirable for each particle to have its own completely unique set of tensors, as describing a systematic MBQC protocol on such a state would be more difficult.

One would therefore like a compact representation for the entire state, ideally with the number of unique tensors required being independent of the number of particles. For example, if the system being simulated is in 1D and with translational invariance (either fully with periodic boundary conditions, or in the bulk with open boundary conditions), that translational invariance should ideally be reflected in the MPS description of the state. In fact, the ideal situation would be for the site matrices (or generally, for higher dimensions, tensors) to be completely site-independent.

In principle, this is often possible. Consider the mechanism by which the infinite-system DMRG algorithm, phrased in the traditional density-matrix formalism, gives rise to a MPS in the thermodynamic limit (Sec. 5.3.2). The MPS matrices are solely dependent on the isometries T_L and T_R ; furthermore, in the case of open boundary conditions, they are often explicitly set to be mirror-symmetric to each other with respect to the centre of the chain. If the isometries reach a fixed point, in the sense that for all sufficiently late iterations $j > j_c$ one has $T_L^{(j)} = T_L^{(j-1)}$, then the MPS matrices will become site independent as well. One can then reasonably assume that these fixed matrices give the correct MPS representation for the thermodynamic limit ground state in the bulk.

In the case of the JCH model in the $n = 1$ MI lobe, I found that the isometries approximately reached a fixed point for very small system sizes, corresponding to a final size of $L \approx 20$. When such a fixed point is reached, from Eq. 5.13, we know that the infinite system DMRG algorithm produces a MPS representation for the j th iteration

ground state $|\psi^{[j]}\rangle$ whose coefficient corresponding to the basis state $|s_1 \dots s_j\rangle|s'_1 \dots s'_j\rangle$ is given by the tensor contraction

$$T_L^{[j]s_j} \dots T_L^{[1]s_1} \Psi T_R^{[1]s'_1} \dots T_R^{[j]s'_j} \quad (5.14)$$

where Ψ is a fixed matrix whose elements contain the coefficients of $|\psi^{[j]}\rangle$ in the final Block-Universe basis (note that the basis state above is the full, untruncated basis).

An easy, quantitative way to show that an approximate fixed point for the ground state has been reached is to calculate the fidelity of the state $|\psi^{[j]}\rangle$ with respect to $|\psi^{[j-1]}\rangle$; this is possible even though the states are on different physical system sizes, because the block decimation procedure ensures that the Hilbert spaces on which they are defined are of the same dimension. Fig. 5.4 shows the fidelity of the ground state for $L = 20$ with respect to that for $L = 18$ for $\tilde{\mu} = -0.878$, for various values of κ/g in the $n = 1$ MI lobe. For every case except for the one with $d = 3$, the fidelity is close to unity for each phase space point considered, indicating a good approximation of a fixed point has been reached. In fact, this fixed point is reached almost immediately, for $L = 6$ or 8 , in all cases.

I have glossed over the boundary conditions normally appearing in an MPS representation with open boundary conditions; in a complete proposal for a quantum wire using the JCH model, these need to be taken into account, as they encode the initial state of the correlation system and the measurement basis in which the final readout must be performed. In Eq. 5.14 above, the “matrices” $T_L^{[j]s_j}$ and $T_R^{[j]s'_j}$ are actually index-dependent vectors, so that the whole tensor contraction collapses to a scalar. Usually the coefficients of an MPS with open boundary conditions are written

$$\langle R | T_L^{[j]s_j} \dots T_L^{[1]s_1} \Psi T_R^{[1]s'_1} \dots T_R^{[j]s'_j} | L \rangle, \quad (5.15)$$




Figure 5.4: Fidelity of ground state of 1D JCH model for system size $L = 20$ with respect to the ground state at $L = 18$ with $\tilde{\mu} = -0.878$, for various different values of κ/g , in the $n = 1$ MI lobe. This provides good evidence that, for all cases other than $d = 3, \chi = 5$, the infinite-system DMRG algorithm has reached a reasonable approximation of a fixed point at this system size.

with $\langle R|$ and $|L\rangle$ fixed boundary conditions. For a given set of vectors $\{T_L^{[j]s_j}\}$, it is not difficult to find a vector $\langle R|$ and a set of matrices $\{\tau_L^{[j]s_j}\}$ such that $T_L^{[j]s_j} = \langle R|\tau_L^{[j]s_j}$. Doing this and inserting the result into Eq. 5.14, together with a similar rewriting $T_R^{[j]s'_j} \equiv \tau_R^{[j]s'_j}|L\rangle$, one recovers the form of Eq. 5.15. The exact boundary conditions $|L\rangle$ and $\langle R|$ are not important, as long as the intervening MPS matrices can be used to perform a universal set of single-qudit gates in correlation space.

The fixed matrix Ψ appearing in the matrix product is rather inconvenient, as there is no flexibility in what happens in correlation space when this part of the chain is reached following a measurement pattern. The simplest way to deal with this is to rewrite the matrix in terms of its singular value decomposition as $\Psi = U\Lambda V$, where Λ is a diagonal singular value matrix. If Λ is proportional to the identity matrix, which occurs when the state has maximal von Neumann entanglement entropy with respect to a cut down

the centre of the chain, then Ψ is a known unitary that will occur in correlation space at a known time, and can therefore be compensated immediately after it occurs by implementing Ψ^{-1} .

As discussed in Ch. 4, the ground state in the MI regime is independent of μ , and is completely determined by κ , assuming one is working in units of g . The precise MPS matrix elements depend on κ , but the form of the matrices anywhere in the MI regime (for the $d = 5$, $\chi = 3$ case) is found to be

$$T_L[1] = \begin{bmatrix} 0 & T_{12}^1 & 0 \\ 0 & 0 & 0 \\ T_{31}^1 & 0 & 0 \end{bmatrix}; \quad T_L[2] = \begin{bmatrix} T_{11}^2 & 0 & 0 \\ 0 & T_{22}^2 & 0 \\ 0 & 0 & T_{33}^2 \end{bmatrix}; \quad T_L[3] = \begin{bmatrix} T_{11}^3 & 0 & 0 \\ 0 & T_{22}^3 & 0 \\ 0 & 0 & T_{33}^3 \end{bmatrix}; \quad (5.16)$$

$$T_L[4] = \begin{bmatrix} 0 & 0 & T_{13}^4 \\ T_{21}^4 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}; \quad T_L[5] = \begin{bmatrix} 0 & 0 & T_{13}^5 \\ T_{21}^5 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad (5.17)$$

corresponding to the single-cavity basis state labelling

$$|1\rangle \equiv |0g\rangle; \quad |2\rangle \equiv |0e\rangle; \quad |3\rangle \equiv |1g\rangle; \quad |4\rangle \equiv |1e\rangle; \quad |5\rangle \equiv |2g\rangle, \quad (5.18)$$

with the ket labels respectively indicating the photon occupation number and the atomic state within the cavity. In each case, one finds that $T_R[s] = T_L[s]^T$, which follows directly from imposing that the basis for the Universe is the reflection of that for the Block about the centre of the chain. Because of how quickly an apparent fixed point is reached, the matrices $\{T_L[s]\}$ can be assumed to be the same for every cavity in the left half of the chain, and those in $\{T_R[s]\}$ the same throughout the right half.

5.4.3 Single-cavity measurements

For a general MBQC protocol, it will be necessary to perform single-atom measurements not just in the computational basis $\{|g\rangle, |e\rangle\}$, but also in at least a one-parameter family

of measurement bases that are superpositions of the energy levels. However, the atomic and photonic sectors of the state are in general entangled with each other. One option for dealing with this is to do a partial trace over the photon sector of the state, and try to design a measurement pattern for the resulting mixed state. Another option is to consider each photon number sector independently, and see if a unitary operator can be constructed from the MPS matrices spanning the sector. Projecting the state of the atom to $\alpha|g\rangle + \beta|e\rangle$ will implement the correlation space operation $\alpha T_L[1] + \beta T_L[2]$ in the zero-photon sector, $\alpha T_L[3] + \beta T_L[4]$ in the one-photon sector, and so on.

One could imagine first doing a destructive photon counting measurement for a particular cavity to collapse it into a particular photon number sector, and the following up with the atomic measurement. In the ideal (but unlikely) scenario, it would be possible to implement an arbitrary single-qutrit unitary gate in any photon number sector, with the atomic measurement basis being chosen on the basis of both the previous atom measurements and the photon-number-resolving measurement on the current cavity. Less restrictively, suppose it is always possible to implement the identity gate in any photon number sector, and only the N -photon number sector allows a universal set of single-qutrit gates. In this case the state can be used as a quantum wire with the number of single-atom measurements needed to implement a single-qutrit unitary in correlation space being random. Specifically, one would first resolve the photon number for the cavity, and then implement the necessary unitary only if the photon number outcome is N , or just the identity otherwise.

In fact, the correlation-space operators need not even be interpreted as qutrit operators. If the initial state of the correlation system, encoded by the boundary condition $\langle R|$, can be chosen to reside in a two-dimensional subspace, and the subsequent logical operations in correlation space keep the information in this subspace, then the correlation system can be interpreted as a qubit; there is thus more flexibility than I previously

intimated.

Unfortunately, it does not appear obviously possible to implement a one-parameter family of unitary gates in correlation space in either the full three dimensions, or in a two-dimensional subspace, in either the one- or the two-photon sector. It can readily be checked that the eigenvalues of $\alpha T_L[1] + \beta T_L[2]$ do not depend at all on α or the non-zero matrix elements of $T_L[1]$. In fact, the eigenvalues are always equal to β times the diagonal elements of $T_L[2]$. A similar story is true for the arbitrary linear combination of $T_L[3]$ and $T_L[4]$. It appears to be possible to find points in correlation space for which two of the diagonal elements of $T_L[2]$ are almost equal, but that just means that one might be able to implement the qubit identity gate in the corresponding two-dimensional subspace of correlation space; there is still no sign that a universal set of rotations is possible.

5.5 Future Directions

The results of the investigation so far are not promising, but this is not the end of the story. For one thing, I have only examined the case $d = 5$, $\chi = 3$ closely. But perhaps more importantly, the MPS representation of a quantum state, even for a specific choice of boundary conditions, is not unique. For example, the state described by a particular MPS representation is left invariant if all the matrices for some site j are right-multiplied by an invertible matrix X , and those for site $j + 1$ are left-multiplied by X^{-1} . This freedom in the choice of matrices makes it difficult to glean any physical information about the state. At any rate, an immediate avenue of exploration is whether there is a more useful form of the MPS representation of the JCH ground state than that of Eq. 5.15. This form is not ideal, principally because of the matrix Ψ that appears as a constant in the middle of the matrix product, breaking up the translational invariance of the description of the bulk.

There exist certain canonical forms for MPS representations, having various nice properties (see for example [41]). As a matter of fact, there is a canonical form for 1D quantum states (where the correlation system is a qubit) that can serve as a universal quantum wire, due to Gross and Eisert [85]. In this work, the authors show that any 1D chain of qubits that can serve as a universal quantum wire, under certain physically reasonable assumptions (including a translationally invariant preparation procedure) can be given the standard MPS representation

$$\begin{aligned} B[0] &:= 2^{-1/2}W; \\ B[1] &:= 2^{-1/2}WS(\phi). \end{aligned}$$

Here, W is a single-qubit “always-on” operation, and $S(\phi) := \text{diag}(e^{-i\phi/2}, e^{i\phi/2})$ is a rotation by ϕ about the z -axis of the Bloch sphere in correlation space. The continuous parameter ϕ is called the by-product angle and can take any value from 0 to π . In [85] it is demonstrated that in all states with MPS representation of this form, except for a set of measure zero, it is possible to use sequential measurements to simulate a universal set of single-qubit gates in correlation space (though not always deterministically). It would be useful to determine either if there is a generalisation of this MPS form for quantum wires with correlation systems of dimension greater than two, or alternatively if there is an MPS representation for the ground state of the JCH model in the MI phase containing a two-dimensional subspace in which the qubit canonical form of [85] exists.

One way of attempting to find a better MPS description is to rewrite the DMRG algorithm as an explicit variational method over MPS, and then impose useful constraints on the matrices (translational invariance, for one). The optimal constrained representations could be benchmarked against the quasiexact unconstrained DMRG solutions by comparing ground state energies, expectation values of local operators, correlation functions

and so on.

Another benefit of rewriting the DMRG algorithm in the variational MPS framework is that it allows for several modern extensions that have been developed in recent years (see [157] for an excellent review of the state of the field as of 2011). The two lines of exploration that seem most promising to me are:

1. the iDMRG algorithm due to McCulloch [233], which improves the convergence of the infinite-system DMRG algorithm and allows one to access the thermodynamic limit for certain 1D models directly;
2. simulations of 2D systems via the PEPS representation. As I mentioned earlier, 1D states cannot be universal resources, so it may be worthwhile to bypass the 1D stage entirely and go directly to 2D simulations. One path that opens up in 2D that is not available in 1D is the possibility of using local operations to transform the 2D ground state to a 2D cluster state with edge connectivity greater than the percolation threshold, like the 3D N-U-N resource of Ch. 2.

Other potential directions include computing with the mixed state resulting from tracing out the photonic sector, investigating the possibility of measurements directly in the polariton basis, encoding information in the polariton number basis rather than the atomic basis, and no doubt many others. Some or all of these subjects should be the subject of future work.

Bibliography

- [1] G. J. Milburn, “Quantum optical Fredkin gate,” *Phys. Rev. Lett.* **62**, 2124–2127 (1989).
- [2] D. Gottesman and I. L. Chuang, “Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations,” *Nature* **402**, 390–393 (1999).
- [3] M. P. A. Fisher et al., “Boson localization and the superfluid-insulator transition,” *Phys. Rev. B* **40**, 546–570 (1989).
- [4] A. D. Greentree et al., “Quantum phase transitions of light,” *Nat. Phys.* **2**, 856 – 861 (2006).
- [5] T. D. Kühner and H. Monien, “Phases of the one-dimensional Bose-Hubbard model,” *Phys. Rev. B* **58**, R14741 (1998).
- [6] D. Rossini and R. Fazio, “Mott-Insulating and Glassy Phases of Polaritons in 1D Arrays of Coupled Cavities,” *Phys. Rev. Lett.* **99**, 186401 (2007).
- [7] U. Schollwöck, “The density-matrix renormalization group,” *Rev. Mod. Phys.* **77**, 259–315 (2005).
- [8] N. Bohr, “Über die Serienspektren der Elemente,” *Zeitschrift für Physik* **2**, 423–469 (1920).
- [9] G. E. Moore, “Cramming more components onto integrated circuits,” *Electronics* **38**, 114–117 (1965).
- [10] P. W. Shor, “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,” *SICOMP* **26**, 1484 (1997).
- [11] C. Dürr et al., “Quantum Query Complexity of Some Graph Problems,” *SIAM J. Comput.* **35**, 1310–1328 (2006).
- [12] L. K. Grover, A fast quantum mechanical algorithm for database search, in proceedings of *28th Annual ACM Symposium on the Theory of Computing (STOC)*, 1996.
- [13] R. P. Feynman, “Simulating physics with computers,” *Int. J. Theor. Phys.* **21**, 467–488 (1982).
- [14] S. Lloyd, “Universal quantum simulators,” *Science* **273**, 1073–1078 (1996).
- [15] C. Bennett and G. Brassard, Quantum cryptography: Public key distribution and coin tossing, in proceedings of *Proceedings of IEEE International Conference on Computers, Systems, and Signal Processing*, edited by IEEE, 1984.

- [16] M. Mosca, “Quantum Algorithms,” , arXiv:0808.0369 [quant-ph], .
- [17] A. M. Childs and W. van Dam, “Quantum algorithms for algebraic problems,” *Rev. Mod. Phys.* **82**, 1–52 (2010).
- [18] W. van Dam and Y. Sasaki, “Quantum algorithms for problems in number theory, algebraic geometry, and group theory,” , arXiv:1206.6126 [quant-ph], .
- [19] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, 2000.
- [20] P. Kaye, R. Laflamme and M. Mosca, *An Introduction to Quantum Computing*, Oxford University Press, 2007.
- [21] A. Kitaev et al., *Classical and quantum computation*, American Mathematical Society, 2002.
- [22] A. Peres, “Neumark’s Theorem and Quantum Inseparability,” *Foundations of Physics* **20**, 1441–1453 (1990).
- [23] C. H. Bennett and S. J. Wiesner, “Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states,” *Phys. Rev. Lett.* **69**, 2881–2884 (1992).
- [24] D. P. DiVincenzo, “The Physical Implementation of Quantum Computation,” *Fortschritte der Physik* **48**, 771–783 (2000).
- [25] C. Nayak et al., “Non-Abelian anyons and topological quantum computation,” *Rev. Mod. Phys.* **80**, 1083–1159 (2008).
- [26] A. Ambainis, “Quantum walks and their algorithmic applications,” *Int. J. Quant. Inf.* **1**, 507–518 (2003).
- [27] E. Knill, R. Laflamme and G. J. Milburn, “A scheme for efficient quantum computation with linear optics,” *Nature* **409**, 46–52 (2001).
- [28] E. Fredkin and T. Toffoli, “Conservative logic,” *International Journal of Theoretical Physics* **21**, 219–253 (1982).
- [29] N. J. Cerf, C. Adami and P. G. Kwiat, “Optical simulation of quantum logic,” *Phys. Rev. A* **57**, R1477–R1480 (1998).
- [30] J. C. Howell and J. A. Yeazell, “Reducing the complexity of linear optics quantum circuits,” *Phys. Rev. A* **61**, 052303 (2000).
- [31] P. G. Kwiat et al., “Grover’s search algorithm: an optical approach,” *J. Mod. Opt.* **47**, 257–266 (2000).
- [32] R. L. Rivest, A. Shamir and L. Adleman, “A Method for Obtaining Digital Signatures and Public-Key Cryptosystems,” *Communications of the ACM* **21**, 120–126 (1978).

- [33] D. Gottesman, *Stabilizer Codes and Quantum Error Correction*, PhD thesis, California Institute of Technology, 1997.
- [34] D. S. Abrams and S. Lloyd, “Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors,” *Phys. Rev. Lett.* **83**, 5162–5165 (1999).
- [35] R. Raussendorf and H. J. Briegel, “A One-Way Quantum Computer,” *Phys. Rev. Lett.* **86**, 5188-5191 (2001).
- [36] R. Raussendorf, D. E. Browne and H. J. Briegel, “Measurement-based quantum computation on cluster states,” *Phys. Rev. A* **68**, 022312 (2003).
- [37] M. Hein et al., Entanglement in graph states and its applications, in proceedings of *International School of Physics Enrico Fermi (Varenna, Italy), Quantum computers, algorithms and chaos*, edited by P. Zoller, G. Casati, D. Shepelyansky and G. Benent, IOS Press, Amsterdam, 2006.
- [38] A. G. D’Souza and D. L. Feder, “Strategies for measurement-based quantum computation with cluster states transformed by stochastic local operations and classical communication,” *Phys. Rev. A* **84**, (2011).
- [39] D. Gottesman, “The Heisenberg Representation of Quantum Computers,” , arXiv:9807006 [quant-ph], .
- [40] G. Vidal, “Efficient Classical Simulation of Slightly Entangled Quantum Computations,” *Phys. Rev. Lett.* **91**, 147902 (2003).
- [41] D. Perez-Garcia et al., “Matrix product state representations,” *Quant. Inf. Comput.* **7**, 401-430 (2007).
- [42] F. Verstraete, V. Murg and J. I. Cirac, “Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems,” *Adv. Phys.* **57**, 143-224 (2008).
- [43] J. Eisert, M. Cramer and M. B. Plenio, “Colloquium: Area laws for the entanglement entropy,” *Rev. Mod. Phys.* **82**, 277-306 (2010).
- [44] F. Verstraete and J. I. Cirac, “Valence-bond states for quantum computation,” *Phys. Rev. A* **70**, 060302 (2004).
- [45] F. Verstraete and J. I. Cirac, “Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions,” , arXiv:0407066 [cond-mat], (2004).
- [46] D. Gross, K. Kielsing and J. Eisert, “Potential and limits to cluster-state quantum computing using probabilistic gates,” *Phys. Rev. A* **74**, 042343 (2006).
- [47] D. Gross et al., “Measurement-based quantum computation beyond the one-way model,” *Phys. Rev. A* **76**, 052315 (2007).

- [48] I. Affleck et al., “Valence Bond Ground States in Isotropic Quantum Antiferromagnets,” *Commun. Math. Phys.* **115**, 477-528 (1988).
- [49] F. Haldane, “Continuum dynamics of the 1-D Heisenberg antiferromagnet: Identification with the $O(3)$ nonlinear sigma model,” *Phys. Lett. A* **93**, 464 - 468 (1983).
- [50] G. K. Brennen and A. Miyake, “Measurement-based quantum computer in the gapped ground state of a two-body Hamiltonian,” *Phys. Rev. Lett.* **101**, 010502 (2008).
- [51] A. Miyake, “Quantum Computation on the Edge of a Symmetry-Protected Topological Order,” *Phys. Rev. Lett.* **105**, 040501 (2010).
- [52] F. Verstraete and J. I. Cirac, “Renormalization algorithms for Quantum-Many Body Systems in two and higher dimensions,” , *arXiv:0407066 [cond-mat]*, (2004).
- [53] F. Verstraete et al., “Criticality, the area law, and the computational power of projected entangled pair states,” *Phys. Rev. Lett.* **96**, 220601 (2006).
- [54] D. Gross and J. Eisert, “Novel schemes for measurement-based quantum computation,” *Phys. Rev. Lett.* **98**, 220503 (2007).
- [55] T.-C. Wei, I. Affleck and R. Raussendorf, “Affleck-Kennedy-Lieb-Tasaki State on a Honeycomb Lattice is a Universal Quantum Computational Resource,” *Phys. Rev. Lett.* **106**, 070501 (2011).
- [56] M. Van den Nest et al., “Fundamentals of universality in one-way quantum computation,” *New J. Phys.* **9**, 204 (2007).
- [57] M. B. Plenio and S. Virmani, “An introduction to entanglement measures,” *Quant. Inf. Comput.* **7**, 1-51 (2007).
- [58] G. Vidal, “Entanglement monotones,” *J. Mod. Opt.* **47 (2-3)**, 355 (2000).
- [59] A. Shimony, “Degree of Entanglement,” *Ann. N.Y. Acad. Sci.* **755**, 675–679 (1995).
- [60] H. Barnum and N. Linden, “Monotones and invariants for multi-particle quantum states,” *Journal of Physics A: Mathematical and General* **34**, 6787 (2001).
- [61] C. E. Mora et al., “Universal resources for approximate and stochastic measurement-based quantum computation,” *Phys. Rev. A* **81**, 042315 (2010).
- [62] E. Lubkin, “Entropy of an n -system from its correlation with a k -reservoir,” *Journal of Mathematical Physics* **19**, 1028-1031 (1978).
- [63] D. N. Page, “Average entropy of a subsystem,” *Phys. Rev. Lett.* **71**, 1291–1294 (1993).
- [64] S. K. Foong and S. Kanno, “Proof of Page’s conjecture on the average entropy of a subsystem,” *Phys. Rev. Lett.* **72**, 1148–1151 (1994).

- [65] P. Hayden et al., “Randomizing Quantum States: Constructions and Applications,” *Comm. Math. Phys.* **250**, 371-391 (2004).
- [66] P. Hayden, D. W. Leung and A. Winter, “Aspects of Generic Entanglement,” *Comm. Math. Phys.* **265**, 95-117 (2006).
- [67] D. Gross, S. T. Flammia and J. Eisert, “Most Quantum States Are Too Entangled To Be Useful As Computational Resources,” *Phys. Rev. Lett.* **102**, 190501 (2009).
- [68] M. J. Bremner, C. Mora and A. Winter, “Are Random Pure States Useful for Quantum Computation?,” *Phys. Rev. Lett.* **102**, 190502 (2009).
- [69] S. D. Bartlett and T. Rudolph, “A simple 2-local Hamiltonian system for which the ground state is a universal resource for quantum computation,” *Phys. Rev. A* **74**, 040302(R) (2006).
- [70] X. Chen et al., “Gapped Two-Body Hamiltonian Whose Unique Ground State Is Universal for One-Way Quantum Computation,” *Phys. Rev. Lett.* **102**, 220501 (2009).
- [71] A. C. Doherty and S. D. Bartlett, “Identifying Phases of Quantum Many-Body Systems That Are Universal for Quantum Computation,” *Phys. Rev. Lett.* **103**, 020506 (2009).
- [72] J. Cai et al., “Universal quantum computer from a quantum magnet,” *Phys. Rev. A* **82**, 052309 (2010).
- [73] M. A. Nielsen, “Cluster-state quantum computation,” *Rep. Math. Phys.* **57**, 147 - 161 (2006).
- [74] R. Oliveira and B. M. Terhal, “The complexity of quantum spin systems on a two-dimensional square lattice,” *arXiv: quant-ph/0504050* (2005).
- [75] J. Kempe, A. Kitaev and O. Regev, “The Complexity of the Local Hamiltonian Problem,” *SIAM J. Comput.* **35**, 1070-1097 (2006).
- [76] S. D. Bartlett et al., “Quantum Computational Renormalization in the Haldane Phase,” *Phys. Rev. Lett.* **105**, 110502 (2010).
- [77] Z. Ji, Z. Wei and B. Zeng, “Complete characterization of the ground-space structure of two-body frustration-free Hamiltonians for qubits,” *Phys. Rev. A* **84**, 042338 (2011).
- [78] C. H. Bennett et al., “Exact and asymptotic measures of multipartite pure-state entanglement,” *Phys. Rev. A* **63**, 012307 (2000).
- [79] W. Dür, G. Vidal and J. I. Cirac, “Three qubits can be entangled in two inequivalent ways,” *Phys. Rev. A* **62**, 062314 (2000).
- [80] F. Verstraete et al., “Four qubits can be entangled in nine different ways,” *Phys. Rev. A* **65**, 052112 (2002).

- [81] L. Lamata et al., “Inductive entanglement classification of four qubits under stochastic local operations and classical communication,” *Phys. Rev. A* **75**, 022318 (2007).
- [82] G. Gour and N. R. Wallach, “All maximally entangled four-qubit states,” *J. Math. Phys.* **51**, 112201 (2010).
- [83] X. Chen et al., “Quantum State Reduction for Universal Measurement Based Computation,” *Phys. Rev. Lett.* **105**, 020502 (2010).
- [84] M. Van den Nest et al., “Graph states as ground states of many-body spin-1/2 Hamiltonians,” *Phys. Rev. A* **77**, 012301 (2008).
- [85] D. Gross and J. Eisert, “Quantum computational webs,” *Phys. Rev. A* **82**, 040303 (2010).
- [86] I. Affleck et al., “Rigorous Results on Valence-Bond Ground-States in Antiferromagnets,” *Phys. Rev. Lett.* **59**, 799-802 (1987).
- [87] K. Kieling, T. Rudolph and J. Eisert, “Percolation, renormalization, and quantum computing with nondeterministic gates,” *Phys. Rev. Lett.* **99**, 130501 (2007).
- [88] M. Popp et al., “Localizable entanglement,” *Phys. Rev. A* **71**, 042306 (2005).
- [89] W. K. Wootters, “Entanglement of Formation and Concurrence,” *Quant. Inf. Comput.* **1**, 27-44 (2001).
- [90] W. Dür et al., “Entanglement in Spin Chains and Lattices with Long-Range Ising-Type Interactions,” *Phys. Rev. Lett.* **94**, 097203 (2005).
- [91] J. Calsamiglia et al., “Spin Gases: Quantum Entanglement Driven by Classical Kinematics,” *Phys. Rev. Lett.* **95**, 180502 (2005).
- [92] S. Anders et al., “Ground-State Approximation for Strongly Interacting Spin Systems in Arbitrary Spatial Dimension,” *Phys. Rev. Lett.* **97**, 107206 (2006).
- [93] S. Anders, H. J. Briegel and W. Dür, “A variational method based on weighted graph states,” *New J. Phys.* **9**, 361 (2007).
- [94] K. Kieling, D. Gross and J. Eisert, “Minimal resources for linear optical one-way computing,” *J. Opt. Soc. Am. B* **24**, 184-188 (2007).
- [95] K. Kieling, D. Gross and J. Eisert, “Cluster state preparation using gates operating at arbitrary success probabilities,” *New J. Phys.* **9**, 200 (2007).
- [96] K. Kieling, J. L. O’Brien and J. Eisert, “On photonic controlled phase gates,” *New J. Phys.* **12**, 013003 (2010).
- [97] J. Chen et al., “No-go theorem for one-way quantum computing on naturally occurring two-level systems,” *Phys. Rev. A* **83**, 050301 (2011).

- [98] R. Grimm, M. Weidemüller and Y. B. Ovchinnikov, “Optical Dipole Traps for Neutral Atoms.” *In* B. Bederson and H. Walther (ed.), Academic Press, pp. 95 - 170 (2000).
- [99] C. Chin et al., “Feshbach resonances in ultracold gases,” *Rev. Mod. Phys.* **82**, 1225–1286 (2010).
- [100] M. Greiner et al., “Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms,” *Nature* **415**, 39-44 (2002).
- [101] H. Bruus and K. Flensberg, *Many-Body Quantum Theory in Condensed Matter Physics*, Oxford University Press, 2004.
- [102] C. J. Pethick and H. Smith, *Bose–Einstein Condensation in Dilute Gases*, Cambridge University Press, 2008.
- [103] A. Georges and T. Giamarchi, “Strongly correlated bosons and fermions in optical lattices.” *In* C. Salomon, G. V. Shlyapnikov and L. F. Cugliandolo (ed.), *Many-Body Physics with Ultracold Gases: Lecture Notes of the Les Houches Summer School: Volume 94, July 2010*, Oxford University Press, pp. 3-70 (2012).
- [104] F. Bloch, “ber die Quantenmechanik der Elektronen in Kristallgittern,” *Z. Phys.* **52**, 555-600 (1929).
- [105] G. H. Wannier, “The Structure of Electronic Excitation Levels in Insulating Crystals,” *Phys. Rev.* **52**, 191–197 (1937).
- [106] N. Marzari et al., “Maximally localized Wannier functions: Theory and applications,” *Rev. Mod. Phys.* **84**, 1419–1475 (2012).
- [107] N. Goldenfeld, *Lectures on Phase Transitions and the Renormalization Group*, Addison-Wesley, 1992.
- [108] S. Sachdev, *Quantum Phase Transitions*, Cambridge University Press, 2000.
- [109] L. Onsager, “Crystal Statistics. I. A Two-Dimensional Model with an Order-Disorder Transition,” *Phys. Rev.* **65**, 117–149 (1944).
- [110] W. F. Trench, “On the eigenvalue problem for Toeplitz band matrices,” *Linear Algebra Appl.* **64**, 199 - 214 (1985).
- [111] P. Kapitza, “Viscosity of Liquid Helium below the λ -point,” *Nature* **141**, 74 (1938).
- [112] J. F. Allen and A. D. Misener, “Flow of Liquid Helium II,” *Nature* **141**, 75 (1938).
- [113] A. Einstein, “Quantentheorie des einatomigen idealen Gases,” *Sitzungsber. Kgl. Preuss. Akad. Wiss.* **1924**, 261 (1924).

- [114] A. Einstein, “Quantentheorie des einatomigen idealen Gases, 2,” Sitzungsber. Kgl. Preuss. Akad. Wiss. **1925**, 3 (1925).
- [115] S. N. Bose, “Plancks Gesetz und Lichtquantenhypothese,” Z. Phys. **26**, 178-181 (1924).
- [116] M. H. Anderson et al., “Observation of Bose-Einstein Condensation in a Dilute Atomic Vapor,” Science **269**, 198-201 (1995).
- [117] K. B. Davis et al., “Bose-Einstein Condensation in a Gas of Sodium Atoms,” Phys. Rev. Lett. **75**, 3969–3973 (1995).
- [118] F. London, “The λ -Phenomenon of Liquid Helium and the Bose-Einstein Degeneracy,” Nature **141**, 643 (1938).
- [119] L. D. Landau, “The theory of superfluidity of Helium II,” J. Phys. USSR **5**, 71 (1941).
- [120] E. P. Gross, “Structure of a quantized vortex in boson systems,” Nuovo Cimento **20**, 454-477 (1961).
- [121] L. P. Pitaevskii, “Vortex Lines in an imperfect Bose gas,” J. Exp. Theor. Phys. **13**, 451-454 (1961).
- [122] N. N. Bogoliubov, “On the Theory of Superfluidity,” Journal of Physics **11**, 23-32 (1947).
- [123] N. N. Bogoliubov, “On the Theory of Superfluidity,” Izv. Akademii Nauk USSR **11**, 77 (1947).
- [124] D. D. Osheroff, R. C. Richardson and D. M. Lee, “Evidence for a New Phase of Solid He^3 ,” Phys. Rev. Lett. **28**, 885–888 (1972).
- [125] D. D. Osheroff et al., “New Magnetic Phenomena in Liquid He^3 below 3 mK,” Phys. Rev. Lett. **29**, 920–923 (1972).
- [126] A. J. Leggett, “Interpretation of Recent Results on He^3 below 3 mK: A New Liquid Phase?,” Phys. Rev. Lett. **29**, 1227–1230 (1972).
- [127] O. Penrose and L. Onsager, “Bose-Einstein Condensation and Liquid Helium,” Phys. Rev. **104**, 576–584 (1956).
- [128] O. Penrose, “CXXXVI. On the quantum mechanics of helium II,” Philos. Mag. **42**, 1373-1377 (1951).
- [129] R. Roth and K. Burnett, “Phase diagram of bosonic atoms in two-component superlattices,” Phys. Rev. A **68**, 023604 (2003).

- [130] M. E. Fisher, M. N. Barber and D. Jasnow, “Helicity Modulus, Superfluidity, and Scaling in Isotropic Systems,” *Phys. Rev. A* **8**, 1111–1124 (1973).
- [131] W. Krauth, “Bethe ansatz for the one-dimensional boson Hubbard model,” *Phys. Rev. B* **44**, 9772–9775 (1991).
- [132] K. G. Singh and D. S. Rokhsar, “Disordered bosons: Condensate and excitations,” *Phys. Rev. B* **49**, 9013–9023 (1994).
- [133] R. Peierls, “Zur Theorie des Diamagnetismus von Leitungselektronen,” *Z. Phys.* **80**, 763–791 (1933).
- [134] J. M. Kosterlitz and D. J. Thouless, “Ordering, metastability and phase transitions in two-dimensional systems,” *J. Phys. C* **6**, 1181 (1973).
- [135] N. F. Mott and R. Peierls, “Discussion of the paper by de Boer and Verwey,” *Proceedings of the Physical Society* **49**, 72 (1937).
- [136] N. F. Mott, “The Basis of the Electron Theory of Metals, with Special Reference to the Transition Metals,” *Proceedings of the Physical Society. Section A* **62**, 416 (1949).
- [137] N. F. Mott, “Review Lecture: Metal-Insulator Transitions,” *Proc. R. Soc. Lond. A* **382**, 1–24 (1982).
- [138] R. Jördens et al., “A Mott insulator of fermionic atoms in an optical lattice,” *Nature* **455**, 204–207 (2008).
- [139] M. Imada, A. Fujimori and Y. Tokura, “Metal-insulator transitions,” *Rev. Mod. Phys.* **70**, 1039–1263 (1998).
- [140] K. Sheshadri et al., “Superfluid and Insulating Phases in an Interacting-Boson Model: Mean-Field Theory and the RPA,” *EPL (Europhysics Letters)* **22**, 257 (1993).
- [141] D. S. Rokhsar and B. G. Kotliar, “Gutzwiller projection for bosons,” *Phys. Rev. B* **44**, 10328–10332 (1991).
- [142] W. Krauth, M. Caffarel and J.-P. Bouchaud, “Gutzwiller wave function for a model of strongly interacting bosons,” *Phys. Rev. B* **45**, 3137–3140 (1992).
- [143] E. J. Mueller et al., “Fragmentation of Bose-Einstein condensates,” *Phys. Rev. A* **74**, 033612 (2006).
- [144] P. Nozières, “Some Comments on Bose-Einstein Condensation.” *In* A. Griffin, D. W. Snoke and S. Stringari (ed.), *Bose-Einstein Condensation*, Cambridge University Press, pp. 15 (1995).
- [145] T. Giamarchi, *Quantum Physics in One Dimension*, Oxford University Press, 2003.

- [146] F. D. M. Haldane, “Effective Harmonic-Fluid Approach to Low-Energy Properties of One-Dimensional Quantum Fluids,” *Phys. Rev. Lett.* **47**, 1840–1843 (1981).
- [147] F. D. M. Haldane, “‘Luttinger liquid theory’ of one-dimensional quantum fluids. I. Properties of the Luttinger model and their extension to the general 1D interacting spinless Fermi gas,” *Journal of Physics C: Solid State Physics* **14**, 2585 (1981).
- [148] L. Tonks, “The complete equation of state of one, two and three-dimensional gases of hard elastic spheres,” *Phys. Rev.* **50**, 955–963 (1936).
- [149] M. Girardeau, “Relationship between systems of impenetrable bosons and fermions in one dimension,” *J. Math. Phys.* **1**, 516–523 (1960).
- [150] M. Hartmann, F. Brandão and M. Plenio, “Quantum many-body phenomena in coupled cavity arrays,” *Laser & Photonics Reviews* **2**, 527–556 (2008).
- [151] A. E. Siegman, *Lasers*, University Science Books, 1986.
- [152] E. Jaynes and F. Cummings, “Comparison of quantum and semiclassical radiation theories with application to the beam maser,” **51**, 89 - 109 (1963).
- [153] K. Birnbaum et al., “Photon blockade in an optical cavity with one trapped atom,” *Nature* **436**, 87–90 (2005).
- [154] M. J. Hartmann, F. G. S. L. Brandão and M. B. Plenio, “Strongly interacting polaritons in coupled arrays of cavities,” *Nat. Phys.* **2**, 849 - 855 (2006).
- [155] S. R. White, “Density-matrix formalism for quantum renormalization groups,” *Phys. Rev. Lett.* **69**, 2863 (1992).
- [156] S. R. White, “Density-matrix algorithms for quantum renormalization groups,” *Phys. Rev. B* **48**, 10345–10356 (1993).
- [157] U. Schollwöck, “The density-matrix renormalization group in the age of matrix product states,” *Ann. Phys.* **326**, 96 - 192 (2011).
- [158] A. Albuquerque et al., “The ALPS project release 1.3: Open-source software for strongly correlated systems,” *J. Magn. Magn. Mater.* **310**, 1187 - 1193 (2007).
- [159] B. Bauer et al., “The ALPS project release 2.0: open source software for strongly correlated systems,” *J. Stat. Mech. Theor. Exp.* **2011**, P05001 (2011).
- [160] K. G. Wilson, “The renormalization group: Critical phenomena and the Kondo problem,” *Rev. Mod. Phys.* **47**, 773–840 (1975).
- [161] R. Bulla, T. A. Costi and T. Pruschke, “Numerical renormalization group method for quantum impurity systems,” *Rev. Mod. Phys.* **80**, 395–450 (2008).

- [162] S. Östlund and S. Rommer, “Thermodynamic Limit of Density Matrix Renormalization,” *Phys. Rev. Lett.* **75**, 3537–3540 (1995).
- [163] J. Dukelsky et al., “Equivalence of the variational matrix product method and the density matrix renormalization group applied to spin chains,” *EPL (Europhysics Letters)* **43**, 457 (1998).
- [164] H. Takasaki, T. Hikihara and T. Nishino, “Fixed Point of the Finite System DMRG,” *J. Phys. Soc. Jpn.* **68**, 1537-1540 (1999).
- [165] S. R. White, “Density matrix renormalization group algorithms with a single center site,” *Phys. Rev. B* **72**, 180403 (2005).
- [166] S. R. White and D. A. Huse, “Numerical renormalization-group study of low-lying eigenstates of the antiferromagnetic $S = 1$ Heisenberg chain,” *Phys. Rev. B* **48**, 3844–3852 (1993).
- [167] S. Yan, D. A. Huse and S. R. White, “Spin-Liquid Ground State of the $S = 1/2$ Kagome Heisenberg Antiferromagnet,” *Science* **332**, 1173-1176 (2011).
- [168] S. Depenbrock, I. P. McCulloch and U. Schollwöck, “Nature of the Spin-Liquid Ground State of the $S = 1/2$ Heisenberg Model on the Kagome Lattice,” *Phys. Rev. Lett.* **109**, 067201 (2012).
- [169] H.-C. Jiang, Z. Wang and L. Balents, “Identifying topological order by entanglement entropy,” *Nat Phys* **8**, 902–905 (2012).
- [170] T. Han, S. Chu and Y. S. Lee, “Refining the Spin Hamiltonian in the Spin- $\frac{1}{2}$ Kagome Lattice Antiferromagnet $\text{ZnCu}_3(\text{OH})_6\text{Cl}_2$ Using Single Crystals,” *Phys. Rev. Lett.* **108**, 157202 (2012).
- [171] K. A. Hallberg, “New Trends in Density Matrix Renormalization,” *Adv. Phys.* **55**, 477-526 (2006).
- [172] T. D. Kühner, S. R. White and H. Monien, “One-dimensional Bose-Hubbard model with nearest-neighbor interaction,” *Phys. Rev. B* **61**, 12474 (2000).
- [173] A. G. D’Souza, B. C. Sanders and D. L. Feder, “Fermionized photons in the ground state of one-dimensional coupled cavities,” *Phys. Rev. A* **88**, 063801 (2013).
- [174] H. van der Zant et al., “Field-induced superconductor to insulator transitions in Josephson-junction arrays,” *Phys. Rev. Lett.* **69**, 2971 (1992).
- [175] I. Bloch, “Ultracold quantum gases in optical lattices,” *Nat. Phys.* **1**, 23 (2005).
- [176] G. Corrielli et al., “Fractional Bloch oscillations in photonic lattices,” *Nat. Commun.* **4**, 1555 (2013).

- [177] D. G. Angelakis, M. F. Santos and S. Bose, “Photon-blockade-induced Mott transitions and XY spin models in coupled cavity arrays,” *Phys. Rev. A* **76**, 031805 (2007).
- [178] A. A. Houck, H. E. Türeci and J. Koch, “On-chip quantum simulation with superconducting circuits,” *Nat. Phys.* **8**, 292-299 (2012).
- [179] T. Stöferle et al., “Transition from a Strongly Interacting 1D Superfluid to a Mott Insulator,” *Phys. Rev. Lett.* **92**, 130403 (2004).
- [180] 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 (2007).
- [181] N. Gemelke et al., “In situ observation of incompressible Mott-insulating domains in ultracold atomic gases,” *Nature* **460**, 995-U75 (2009).
- [182] W. S. Bakr et al., “Probing the Superfluid-to-Mott Insulator Transition at the Single-Atom Level,” *Science* **329**, 547-550 (2010).
- [183] E. Haller et al., “Pinning quantum phase transition for a Luttinger liquid of strongly interacting bosons,” *Nature* **466**, 597-U1 (2010).
- [184] S. Trotzky et al., “Suppression of the critical temperature for superfluidity near the Mott transition,” *Nat. Phys.* **6**, 998-1004 (2010).
- [185] M. J. Hartmann, F. G. S. L. Brandao and M. B. Plenio, “Effective spin systems in coupled microcavities,” *Phys. Rev. Lett.* **99**, 160501 (2007).
- [186] I. Carusotto et al., “Fermionized Photons in an Array of Driven Dissipative Non-linear Cavities,” *Phys. Rev. Lett.* **103**, 033601 (2009).
- [187] M. Kiffner and M. J. Hartmann, “Dissipation-induced Tonks-Girardeau gas of polaritons,” *Phys. Rev. A* **81**, 021806 (2010).
- [188] A. Halu et al., “Phase transition of light on complex quantum networks,” *Phys. Rev. E* **87**, 022104 (2013).
- [189] A. L. C. Hayward, A. M. Martin and A. D. Greentree, “Fractional Quantum Hall Physics in Jaynes-Cummings-Hubbard Lattices,” *Phys. Rev. Lett.* **108**, 223602 (2012).
- [190] M. Schiró et al., “Phase Transition of Light in Cavity QED Lattices,” *Phys. Rev. Lett.* **109**, 053601 (2012).
- [191] M. J. Hartmann, F. G. S. L. Brandão and M. B. Plenio, “A polaritonic two-component Bose-Hubbard model,” *New J. Phys.* **10**, 033011 (2008).
- [192] M. I. Makin et al., “Quantum phase transitions in photonic cavities with two-level systems,” *Phys. Rev. A* **77**, 053819 (2008).

- [193] M. Aichhorn et al., “Quantum Fluctuations, Temperature, and Detuning Effects in Solid-Light Systems,” *Phys. Rev. Lett.* **100**, 216401 (2008).
- [194] J. Koch and K. Le Hur, “Superfluid–Mott-insulator transition of light in the Jaynes-Cummings lattice,” *Phys. Rev. A* **80**, 023811 (2009).
- [195] S. Schmidt and G. Blatter, “Strong Coupling Theory for the Jaynes-Cummings-Hubbard Model,” *Phys. Rev. Lett.* **103**, 086403 (2009).
- [196] P. Phipps, H. G. Evert and M. Hohenadler, “Excitation spectra of strongly correlated lattice bosons and polaritons,” *Phys. Rev. A* **80**, 033612 (2009).
- [197] S. Schmidt and G. Blatter, “Excitations of Strongly Correlated Lattice Polaritons,” *Phys. Rev. Lett.* **104**, 216402 (2010).
- [198] M. Knap, E. Arrigoni and W. von der Linden, “Spectral properties of coupled cavity arrays in one dimension,” *Phys. Rev. B* **81**, 104303 (2010).
- [199] M. Hohenadler et al., “Dynamical critical exponent of the Jaynes-Cummings-Hubbard model,” *Phys. Rev. A* **84**, 041608(R) (2011).
- [200] N. Na et al., “Strongly correlated polaritons in a two-dimensional array of photonic crystal microcavities,” *Phys. Rev. A* **77**, 031803(R) (2008).
- [201] P. A. Ivanov et al., “Simulation of a quantum phase transition of polaritons with trapped ions,” *Phys. Rev. A* **80**, 060301(R) (2009).
- [202] A. Mering et al., “Analytic approximations to the phase diagram of the Jaynes-Cummings-Hubbard model,” *Phys. Rev. A* **80**, 053821 (2009).
- [203] A. Nunnenkamp, J. Koch and S. M. Girvin, “Synthetic gauge fields and homodyne transmission in Jaynes-Cummings lattices,” *New J. Phys.* **13**, 095008 (2011).
- [204] C.-W. Wu et al., “Quantum phase transition of light in a one-dimensional photon-hopping-controllable resonator array,” *Phys. Rev. A* **84**, 043827 (2011).
- [205] S. Schmidt and J. Koch, “Circuit QED lattices: Towards quantum simulation with superconducting circuits,” *Ann. Phys.* **525**, 395–412 (2013).
- [206] E. H. Lieb and W. Liniger, “Exact Analysis of an Interacting Bose Gas. I. The General Solution and the Ground State,” *Phys. Rev.* **130**, 1605–1616 (1963).
- [207] B. Paredes et al., “Tonks-Girardeau gas of ultracold atoms in an optical lattice,” *Nature* **429**, 277–281 (2004).
- [208] T. Kinoshita, T. Wenger and D. S. Weiss, “Observation of a One-Dimensional Tonks-Girardeau Gas,” *Science* **305**, 1125–1128 (2004).

- [209] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics*, Cambridge University Press, 1995.
- [210] M. Hastings and T. Koma, “Spectral Gap and Exponential Decay of Correlations,” *Comm. Math. Phys.* **265**, 781-804 (2006).
- [211] V. N. Popov, *Functional integrals and collective excitations*, Cambridge University Press, 1987.
- [212] M. Olshanii, “Atomic Scattering in the Presence of an External Confinement and a Gas of Impenetrable Bosons,” *Phys. Rev. Lett.* **81**, 938–941 (1998).
- [213] V. Yukalov and M. Girardeau, “Fermi-Bose mapping for one-dimensional Bose gases,” *Laser Physics Letters* **2**, 375–382 (2005).
- [214] M. A. Cazalilla et al., “One dimensional bosons: From condensed matter systems to ultracold gases,” *Rev. Mod. Phys.* **83**, 1405 (2011).
- [215] J. Friedel, “Metallic Alloys,” *Nuovo Cimento* **7**, 287-311 (1958).
- [216] S. Ejima, H. Fehske and F. Gebhard, “Dynamic properties of the one-dimensional Bose-Hubbard model,” *EPL (Europhysics Letters)* **93**, 30002 (2011).
- [217] R. V. Pai et al., “One-Dimensional Disordered Bosonic Hubbard Model: A Density-Matrix Renormalization Group Study,” *Phys. Rev. Lett.* **76**, 2937–2940 (1996).
- [218] S. Ejima et al., “Characterization of Mott-insulating and superfluid phases in the one-dimensional Bose-Hubbard model,” *Phys. Rev. A* **85**, 053644 (2012).
- [219] C. Kollath et al., “Spatial correlations of trapped one-dimensional bosons in an optical lattice,” *Phys. Rev. A* **69**, 031601 (2004).
- [220] R. Roth and K. Burnett, “Superfluidity and interference pattern of ultracold bosons in optical lattices,” *Phys. Rev. A* **67**, 031602 (2003).
- [221] H. Li and F. D. M. Haldane, “Entanglement Spectrum as a Generalization of Entanglement Entropy: Identification of Topological Order in Non-Abelian Fractional Quantum Hall Effect States,” *Phys. Rev. Lett.* **101**, 010503 (2008).
- [222] R. W. Spekkens and J. E. Sipe, “Spatial fragmentation of a Bose-Einstein condensate in a double-well potential,” *Phys. Rev. A* **59**, 3868 (1999).
- [223] I. Peschel and V. Eisler, “Reduced density matrices and entanglement entropy in free lattice models,” *J. Phys. A: Math. Theor.* **42**, 504003 (2009).
- [224] M. M. Wolf, “Violation of the entropic area law for fermions,” *Phys. Rev. Lett.* **96**, 010404 (2006).

- [225] M. B. Plenio et al., “Entropy, Entanglement, and Area: Analytical Results for Harmonic Lattice Systems,” *Phys. Rev. Lett.* **94**, 060503 (2005).
- [226] M. Van den Nest, “Universal quantum computation with little entanglement,” *Phys. Rev. Lett.* **110**, 060504 (2013).
- [227] F. Verstraete, M. Popp and J. Cirac, “Entanglement versus Correlations in Spin Systems,” *Phys. Rev. Lett.* **92**, 027901 (2004).
- [228] F. Verstraete, M. A. Martin-Delgado and J. I. Cirac, “Diverging Entanglement Length in Gapped Quantum Spin Systems,” *Phys. Rev. Lett.* **92**, 087201 (2004).
- [229] L. Amico et al., “Entanglement in many-body systems,” *Rev. Mod. Phys.* **80**, 517 (2008).
- [230] H. J. Carmichael, P. Kochan and B. C. Sanders, “Photon Correlation Spectroscopy,” *Phys. Rev. Lett.* **77**, 631 (1996).
- [231] G. Zhu, S. Schmidt and J. Koch, “Dispersive regime of the Jaynes-Cummings and Rabi lattice,” **arXiv:1307.2505**, (2013).
- [232] G. Vidal, “A class of quantum many-body states that can be efficiently simulated,” *Phys. Rev. Lett.* **101**, 110501 (2008).
- [233] I. P. McCulloch, “Infinite size density matrix renormalization group, revisited,” **arXiv:0804.2509**, (2008).
- [234] G. M. Crosswhite, A. C. Doherty and G. Vidal, “Applying matrix product operators to model systems with long-range interactions,” *Phys. Rev. B* **78**, 035116 (2008).
- [235] H. J. Briegel et al., “Measurement-based quantum computation,” *Nat. Phys.* **5**, 19-26 (2009).
- [236] D. Deutsch, “Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer,” *Proc. R. Soc. Lond. A* **400**, 97-117 (1985).
- [237] K. L. Brown, W. J. Munro and V. M. Kendon, “Using Quantum Computers for Quantum Simulation,” *Entropy* **12**, 2268–2307 (2010).
- [238] M. S. Underwood and D. L. Feder, “Bose-Hubbard model for universal quantum-walk-based computation,” *Phys. Rev. A* **85**, 052314 (2012).
- [239] D. L. Feder, “Maximally entangled gapped ground state of lattice fermions,” *Phys. Rev. A* **85**, 012312 (2012).
- [240] Y.-J. Chiu, X. Chen and I. L. Chuang, “Fermionic measurement-based quantum computation,” *Phys. Rev. A* **87**, 012305 (2013).

- [241] J. I. Cirac et al., “Quantum State Transfer and Entanglement Distribution among Distant Nodes in a Quantum Network,” *Phys. Rev. Lett.* **78**, 3221–3224 (1997).
- [242] H.-J. Briegel et al., “Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication,” *Phys. Rev. Lett.* **81**, 5932–5935 (1998).
- [243] L.-M. Duan et al., “Long-distance quantum communication with atomic ensembles and linear optics,” *Nature* **414**, 413–418 (2001).
- [244] H. J. Kimble, “The quantum internet,” *Nature* **453**, 1023–1030 (2008).
- [245] S. Ritter et al., “An elementary quantum network of single atoms in optical cavities,” *Nature* **484**, 195–200 (2012).
- [246] R. Beals et al., “Efficient distributed quantum computing,” *Proc. R. Soc. A* **469**, 20120686 (2013).
- [247] G. Vidal, “Efficient Simulation of One-Dimensional Quantum Many-Body Systems,” *Phys. Rev. Lett.* **93**, 040502 (2004).